A Huggett-Moll Mean Field Game with Mean-Reverting Kou Jump-Diffusion Productivity

PIDE Reduction to PDE via Exponential Jump Structure

T. Zamrik

16-AUG-2025

1. Abstract

We formulate a Huggett-Moll mean field game in which agents' labor productivity follows a mean-reverting Ornstein-Uhlenbeck process with Kou double-exponential jumps, capturing the empirically documented persistence of earnings together with asymmetric income shocks — sudden collapses and rare windfalls. Working in log-productivity space $x = \log z$, the multiplicative jump structure becomes additive, and both the Hamilton-Jacobi-Bellman and Fokker-Planck partial integro-differential equations admit a complete reduction to coupled PDE systems via auxiliary first-order equations derived from the exponential jump kernel. General equilibrium is closed by a Cobb-Douglas production sector, with aggregate capital $K[m] = \iint a\, m(a,x)\, da\, dx$ determining wages and returns endogenously. We state the MFG fixed-point problem and derive the auxiliary PDE systems explicitly. The numerical implementation assembles a sparse generator matrix whose block-diagonal Kou component is the discrete realization of the auxiliary integral operators, solving the HJB by implicit time stepping with direct FOC policy recovery and the FP as a stationary linear system.

2. Introduction

Heterogeneous-agent models in the tradition of Huggett [1] and Aiyagari [2] have become the workhorse framework for studying wealth and income distribution in macroeconomics. Their continuous-time formulation, developed by Achdou, Han, Lasry, Lions, and Moll [3], casts the problem as a mean field game (MFG): a Hamilton-Jacobi-Bellman (HJB) equation governing optimal savings, coupled to a Fokker-Planck (FP) equation governing the joint distribution of wealth and income.

The income process in these models is typically a diffusion. Empirically, individual earnings exhibit two key features: strong persistence, with autocorrelations close to but strictly less than one [10], and pronounced tail asymmetry, with large negative shocks (job losses, pay cuts) more frequent than positive windfalls. Neither feature is captured by a pure random walk. The Ornstein-Uhlenbeck (OU) process restores mean-reversion, ensuring a well-defined stationary distribution; the Kou double-exponential jump-diffusion [4] superimposes asymmetric jump shocks with separate intensities and tail parameters for positive and negative events.

The natural extension of Huggett-Moll to jump-diffusion income introduces partial integro-differential equations (PIDEs) in both the HJB and FP systems, substantially complicating both analysis and computation. The key contribution of this paper is to show that the exponential kernel of the Kou model — when the PIDE is written in log-productivity space — allows complete elimination of the integral operators. Each jump direction introduces one auxiliary first-order PDE, derived by differentiating the convolution integral. The result is a coupled PDE system with no integral terms in either the HJB or the FP equation, making the jump-diffusion MFG numerically equivalent in structure to the pure-diffusion benchmark. The mean-reverting OU drift enters only as a linear term $-\kappa x\,\partial_x$ in the HJB and its adjoint $\partial_x(\kappa x\cdot)$ in the FP; neither term affects the PIDE reduction.

Mean field games were introduced by Lasry and Lions [5] and independently by Huang, Malhamé, and Caines [6]. Their application to heterogeneous-agent macroeconomics is developed in [3] and extended to richer income processes in subsequent work [7, 8]. Jump-diffusion processes in optimal control are treated in [9]; the Kou auxiliary ODE reduction has been applied to option pricing [4] and optimal stopping, but not, to our knowledge, to mean field game systems.

3. Model

Definition 3.1 (State Space)**.** The state of each agent is the pair $(a, x) \in [\underline{a}, \infty) \times \mathbb{R}$, where a denotes financial wealth and $x = \log z$ denotes log-labor-productivity.

A unit mass of households is indexed by (a, x). Each household chooses a consumption path $c(t) \geq 0$ to maximize

$$\mathbb{E}_0\left[\int_0^\infty e^{-\rho t} u(c_t)\,dt\right], \qquad u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \tag{1}$$

where $\rho > 0$ is the discount rate and $\gamma > 0$ is the coefficient of relative risk aversion.

Definition 3.2 (OU-Kou Log-Productivity Dynamics)**.** Log-productivity $x_t = \log z_t$ follows a mean-reverting Ornstein-Uhlenbeck process with Kou double-exponential jumps:

$$dx_t = -\kappa x_t\,dt + \sigma\,dW_t + dJ_t^+ - dJ_t^-, \tag{2}$$

where $\kappa > 0$ is the mean-reversion speed, $\sigma > 0$ is the diffusion coefficient, W_t is a standard Brownian motion, J^+ is a compound Poisson process with intensity $\lambda^+ > 0$ and jump sizes $Y^+ \sim \mathrm{Exp}(\eta^+)$, and J^- is an independent compound Poisson process with intensity $\lambda^- > 0$ and jump sizes $Y^- \sim \mathrm{Exp}(\eta^-)$.

Working directly in log-productivity space has two advantages: the OU drift is linear in x, entering the generator as $-\kappa x\,\partial_x$, and the multiplicative Kou jump structure in z becomes additive in x, enabling the PIDE reduction of Section 4.

Remark 3.3. In the absence of jumps, x_t has stationary distribution $\mathcal{N}(0, \sigma^2/2\kappa)$, so $z_t =$

e^{x_t} is log-normally distributed with finite mean and variance. The jump component adds asymmetric tails without disrupting the stationarity guaranteed by mean reversion.

Definition 3.4 (Budget Constraint)**.** Wealth evolves as

$$\dot{a}_t = we^{x_t} + ra_t - c_t, \qquad a_t \geq \underline{a}, \tag{3}$$

where r is the return on wealth, w is the wage, and $\underline{a}$ is the borrowing limit.

Definition 3.5 (Production Sector)**.** A representative firm produces with Cobb-Douglas technology $Y = K^\alpha L^{1-\alpha}$, $\alpha \in (0, 1)$, with inelastic labor $L = 1$. Profit maximization yields:

$$r = \alpha K^{\alpha-1} - \delta, \qquad w = (1-\alpha)K^\alpha, \tag{4}$$

where $\delta > 0$ is the depreciation rate and $K = K[m]$ is aggregate capital.

4. The HJB–Fokker-Planck System

The MFG system consists of two coupled equations with complementary economic interpretations.

The **HJB equation** models the decision problem of a single atomistic agent. Taking wages w and the interest rate r as fixed — the agent has no market power — she chooses a consumption path $c(a, x)$ at every state (a, x) to maximize expected lifetime utility. The value function $V(a, x)$ records the maximum attainable utility from state (a, x) onward. The HJB is solved **backward**: it asks what is the best action today, given the value of being in each future state.

The **Fokker-Planck equation** models the aggregate population. Given that every agent follows the optimal policy $c^*(a, x)$ derived from the HJB, the FP describes how the cross-sectional distribution $m(a, x)$ of agents evolves across the state space. At stationarity, m is the invariant measure: the fraction of agents at each (a, x) is constant over time because inflows and outflows balance. The FP is solved **forward**: mass is pushed by the drift $(we^x + ra - c^*)$ in wealth, the OU-Kou dynamics in productivity, and the Kou jumps. The mean-field coupling closes when m is used to compute aggregate capital $K[m] = \iint a\, m\, da\, dx$, which determines the prices (r, w) that the HJB took as given.

4.1 Hamilton-Jacobi-Bellman Equation

Definition 4.1 (Jump Operators)**.** For a function $V(a, x)$, define the jump expectation operators:

$$\mathcal{J}^+V(a, x) = \eta^+ \int_0^\infty V(a,\, x+y)\, e^{-\eta^+ y}\, dy,$$

$$\mathcal{J}^-V(a, x) = \eta^- \int_0^\infty V(a,\, x-y)\, e^{-\eta^- y}\, dy.$$

$\mathcal{J}^+$ integrates V to the right of x (where the agent lands after a positive jump); $\mathcal{J}^-$ integrates to the left (after a negative jump).

Proposition 4.2 (HJB PIDE)**.** *The stationary value function $V(a,x)$ solves:*

$$\rho V = \max_{c\geq 0}\Big\{u(c) + \partial_a V\cdot(we^x + ra - c) - \kappa x\,\partial_x V + \tfrac{\sigma^2}{2}\partial_{xx}V + \lambda^+[\mathcal{J}^+V - V] + \lambda^-[\mathcal{J}^-V - V]\Big\}. \tag{5}$$

The drift $-\kappa x\,\partial_x V$ reflects mean reversion of log-productivity toward zero.

Corollary 4.3 (Optimal Consumption)**.** *The first-order condition gives $c^*(a,x) = (\partial_a V)^{-1/\gamma}$.*

4.2 Fokker-Planck Equation

Definition 4.4 (Adjoint Jump Operators)**.** The L^2-adjoints of $\mathcal{J}^\pm$ are:

$$(\mathcal{J}^+)^* m(a,x) = \eta^+\int_0^\infty m(a,\,x-y)\,e^{-\eta^+ y}\,dy,$$

$$(\mathcal{J}^-)^* m(a,x) = \eta^-\int_0^\infty m(a,\,x+y)\,e^{-\eta^- y}\,dy.$$

The adjoint operators integrate m from the direction of the jump: $(\mathcal{J}^+)^*$ looks left (mass that jumped right to reach x); $(\mathcal{J}^-)^*$ looks right (mass that jumped left to reach x).

Proposition 4.5 (FP PIDE)**.** *The stationary joint density $m(a,x)$ solves:*

$$0 = -\partial_a[(we^x + ra - c^*)m] + \partial_x[\kappa x\, m] + \tfrac{\sigma^2}{2}\partial_{xx}m + \lambda^+[(\mathcal{J}^+)^*m - m] + \lambda^-[(\mathcal{J}^-)^*m - m]. \tag{6}$$

The term $\partial_x(\kappa x\, m)$ is the Fokker-Planck adjoint of the OU drift $-\kappa x\,\partial_x$; it pushes mass toward $x = 0$ (productivity $z = 1$).

5. PIDE Reduction to Coupled PDE Systems

5.1 Auxiliary Variables for the HJB

Definition 5.1. Define auxiliary functions:

$$I^+(a,x) = \mathcal{J}^+V(a,x) = \eta^+\int_0^\infty V(a,\,x+y)\,e^{-\eta^+ y}\,dy,$$

$$I^-(a,x) = \mathcal{J}^-V(a,x) = \eta^-\int_0^\infty V(a,\,x-y)\,e^{-\eta^- y}\,dy.$$

Theorem 5.2 (HJB Auxiliary PDEs)**.** *The functions I^+ and I^- satisfy the first-order PDEs:*

$$\partial_x I^+ - \eta^+ I^+ = -\eta^+ V, \qquad I^+(a,x)\to 0 \text{ as } x\to+\infty, \tag{7}$$

$$\partial_x I^- + \eta^- I^- = \eta^- V, \qquad I^-(a,x)\to 0 \text{ as } x\to-\infty. \tag{8}$$

Proof. *Write $I^+(a,x) = \eta^+ e^{\eta^+ x}\int_x^\infty V(a,u)e^{-\eta^+ u}\,du$ and differentiate: $\partial_x I^+ = \eta^+ I^+ - \eta^+ V$, which rearranges to (7). The derivation for I^- is analogous.*

Corollary 5.3 (HJB Coupled PDE System)**.** *Substituting $\mathcal{J}^\pm V = I^\pm$, the HJB PIDE* (5) *is equivalent to the PDE system:*

$$\rho V = \max_c \Big\{ u(c) + \partial_a V \cdot (we^x + ra - c) - \kappa x\, \partial_x V + \tfrac{\sigma^2}{2} \partial_{xx} V + \lambda^+ (I^+ - V) + \lambda^- (I^- - V) \Big\},$$

$$\partial_x I^+ - \eta^+ I^+ = -\eta^+ V, \qquad I^+ \to 0 \text{ as } x \to +\infty,$$

$$\partial_x I^- + \eta^- I^- = \eta^- V, \qquad I^- \to 0 \text{ as } x \to -\infty.$$

No integral operators remain. The OU drift $-\kappa x\, \partial_x V$ *passes through the reduction unchanged.*

5.2 Auxiliary Variables for the FP

Definition 5.4. Define adjoint auxiliary functions:

$$J^+(a,x) = (\mathcal{J}^+)^* m(a,x) = \eta^+ \int_0^\infty m(a,\, x-y)\, e^{-\eta^+ y}\, dy,$$

$$J^-(a,x) = (\mathcal{J}^-)^* m(a,x) = \eta^- \int_0^\infty m(a,\, x+y)\, e^{-\eta^- y}\, dy.$$

Note the reversed integration directions relative to $I^\pm$: J^+ integrates m to the left of x; J^- integrates to the right.

Theorem 5.5 (FP Auxiliary PDEs)**.** *The functions* J^+ *and* J^- *satisfy:*

$$\partial_x J^+ + \eta^+ J^+ = \eta^+ m, \qquad J^+(a,x) \to 0 \text{ as } x \to -\infty, \tag{9}$$

$$\partial_x J^- - \eta^- J^- = -\eta^- m, \qquad J^-(a,x) \to 0 \text{ as } x \to +\infty. \tag{10}$$

The boundary conditions are on opposite sides to those of (7)–(8)*, reflecting the adjoint symmetry.* ***Proof.*** *Write* $J^+(a,x) = \eta^+ e^{-\eta^+ x} \int_{-\infty}^x m(a,u) e^{\eta^+ u}\, du$*. Differentiating:* $\partial_x J^+ = -\eta^+ J^+ + \eta^+ m$*, giving* (9)*. The derivation for* J^- *is analogous.*

Corollary 5.6 (FP Coupled PDE System)**.** *The FP PIDE* (6) *is equivalent to:*

$$0 = -\partial_a[(we^x + ra - c^*)m] + \partial_x[\kappa x\, m] + \tfrac{\sigma^2}{2} \partial_{xx} m + \lambda^+ (J^+ - m) + \lambda^- (J^- - m),$$

$$\partial_x J^+ + \eta^+ J^+ = \eta^+ m, \qquad J^+ \to 0 \text{ as } x \to -\infty,$$

$$\partial_x J^- - \eta^- J^- = -\eta^- m, \qquad J^- \to 0 \text{ as } x \to +\infty.$$

No integral operators remain.

Remark 5.7 (Independence from drift and connection to numerics)*.* The PIDE reduction depends only on the exponential kernel of the Kou distribution, not on the form of the drift. The OU drift $-\kappa x$ enters the HJB as $-\kappa x\, \partial_x V$ and its adjoint $\partial_x(\kappa x\, m)$ in the FP; neither term involves the integral operators $\mathcal{J}^\pm$ or their adjoints, and both pass through the reduction unchanged. The reduction is therefore valid for any drift — GBM, OU, or otherwise — as long as the jump kernel is exponential. In the numerical implementation, the integral operators $\mathcal{J}^\pm$ and $(\mathcal{J}^\pm)^*$ are not assembled via the auxiliary sweep but are instead discretized directly as a block-diagonal sparse matrix J_{sp} added to the generator T.

This is the discrete realization of the same reduction: each block encodes the exponential weighting $\eta^{\pm}e^{-\eta^{\pm}|x_k-x_j|}$ between grid nodes, reproducing the action of $I^{\pm}$ and $J^{\pm}$ on the discrete grid. The two approaches — auxiliary ODE sweep and block-diagonal matrix — are equivalent; the matrix form is chosen here for its compatibility with sparse direct solvers.

6. General Equilibrium

Definition 6.1 (Aggregate Capital)**.**

$$K[m] = \int_{\underline{a}}^{\infty} \int_{-\infty}^{\infty} a\, m(a,x)\, dx\, da. \tag{11}$$

Definition 6.2 (MFG Equilibrium)**.** A stationary MFG equilibrium is a tuple (V, m, c^*, r, w, K) such that: 1. **Optimality**: V solves the HJB coupled PDE system with $c^* = (\partial_a V)^{-1/\gamma}$. 2. **Consistency**: m solves the FP coupled PDE system under c^*. 3. **Market clearing**: $K = \iint a\, m\, da\, dx$ and $r = \alpha K^{\alpha-1} - \delta$, $w = (1-\alpha)K^{\alpha}$. 4. **Normalization**: $\iint m\, da\, dx = 1$. 5. **Borrowing constraint**: no-flux condition at $a = \underline{a}$.

Remark 6.3. The mean-reversion parameter $\kappa > 0$ is essential for existence of a stationary distribution m. Without it, log-productivity drifts without bound and no normalizable stationary density exists.

7. Numerical Method

The HJB equation is solved by **implicit time stepping with direct first-order condition (FOC) policy update**, following the finite difference framework of AHLLM [3]. A sparse generator matrix T encodes the upwind wealth drift, centred diffusion in x, and the block-diagonal Kou jump operator. At each iteration:

$$\left(\tfrac{1}{\Delta} + \rho\right) V^{k+1} - T[c^k]V^{k+1} = u(c^k) + \tfrac{1}{\Delta}V^k, \tag{12}$$

with step size Δ large enough to ensure stability. The policy is recovered at each step directly from the upwind first-order condition:

$$c^* = \left(\partial_a V\right)^{-1/\gamma}, \tag{7.1}$$

selecting forward, backward, or income-consumption finite differences according to the sign of the savings drift $we^x + ra - c$. This is a direct FOC update: the policy is not iterated separately but extracted immediately from the current V.

Remark 7.1 (Howard Acceleration). A more efficient alternative, which we leave for future work, is **Howard policy iteration**: freeze c^k, solve $(\rho I - T(c^k))V = u(c^k)$ exactly by a single sparse direct solve, then update c^{k+1} via FOC. This exact policy evaluation step guarantees convergence in 3–5 iterations rather than hundreds, with identical limiting solution. Implementation of Howard acceleration for the OU-Kou MFG is planned for the next research paper in this series.

The **stationary FP equation** $T^\top m = 0$ is solved as a linear system by replacing one row with the normalization constraint $\iint m\, da\, dx = 1$. The OU drift $-\kappa x$ enters the generator as a spatially varying upwind coefficient, handled by per-node sign detection.

Remark 7.2 (MFG Coupling and Fictitious Play). The present scheme couples the HJB and FP equations through prices alone: the full HJB is solved to convergence, then the FP is solved once under the resulting policy c^*, and aggregate capital $K[m]$ updates the prices (r, w) via bisection. A stronger form of coupling — **simultaneous fictitious-play iteration** — would instead interleave the two solves: solve HJB with the current density m^k, solve FP with the resulting c^{k+1}, update prices from m^{k+1}, and repeat without waiting for either equation to fully converge. This is a distinct algorithm with different convergence properties, not employed by AHLLM [3] and not implemented here, but relevant for non-stationary or mean-field coupling beyond the price channel.

7.1 Algorithm 1: Outer Equilibrium Loop

```
Input: initial guess K_0, tolerance eps_K
Repeat:
  r = alpha * K^{alpha-1} - delta
  w = (1 - alpha) * K^alpha
  Solve HJB via implicit time stepping (AHLLM) -> V(a,x), c*(a,x)
  Solve FP via T^T m = 0 with normalization row -> m(a,x)
  K_new = integral of a * m(a,x) da dx
  if |K_new - K| < eps_K: break
  K = bisection_update(K, K_new)
Output: V, m, c*, r, w, K
```

7.2 Algorithm 2: Auxiliary Sweep

```
Input: V (or m) on grid (a_i, x_j), parameters eta_p, eta_m
HJB sweep for I+:
  right-to-left in x (BC: I+ -> 0 as x -> +inf)
  I+[i,j] = (eta_p * V[i,j] + I+[i,j+1] / dx) / (eta_p + 1/dx)
HJB sweep for I-:
  left-to-right in x (BC: I- -> 0 as x -> -inf)
  I-[i,j] = (eta_m * V[i,j] + I-[i,j-1] / dx) / (eta_m + 1/dx)
FP sweep for J+:
  left-to-right in x (BC: J+ -> 0 as x -> -inf)
  J+[i,j] = (eta_p * m[i,j] + J+[i,j-1] / dx) / (eta_p + 1/dx)
FP sweep for J-:
  right-to-left in x (BC: J- -> 0 as x -> +inf)
  J-[i,j] = (eta_m * m[i,j] + J-[i,j+1] / dx) / (eta_m + 1/dx)
Output: I+, I-, J+, J-
```

7.3 Algorithm 3: Implicit HJB Step

```
Input: V^k, c^k (frozen policy), Delta, rho
Build generator T(c^k):
```

```
  - upwind in a: sign(w*exp(x) + r*a - c*) selects forward/backward
  - upwind in x: sign(-kappa * x) selects forward/backward (OU drift)
  - centred diffusion in x: sigma^2/2 * d^2/dx^2
  - block-diagonal Kou jump operator J_sp
Solve linear system:
  M = (1/Delta + rho) * I - T
  rhs = u(c^k) + V^k / Delta
  V^{k+1} = M \ rhs
Update policy: c^{k+1} = (upwind Va)^{-1/gamma}
```

8. Figures

All figures are computed at the stationary MFG equilibrium with fixed parameters $\rho = 0.09$, $\kappa = 0.6$, $\sigma = 0.35$, $\lambda^+ = \lambda^- = 0.4$, $\eta^+ = \eta^- = 6$, $\alpha = 0.36$, $\delta = 0.08$, $K_0 = 4$. Each figure shows two panels side by side: left at $\gamma = 1.5$ (low risk aversion) and right at $\gamma = 4$ (high risk aversion). The comparison isolates the effect of risk aversion on equilibrium outcomes while holding all other parameters constant.

Value function (Figure 1). In both panels $V(a, z)$ is strictly increasing in wealth a and productivity z, and strictly concave in a — standard properties of CRRA preferences. The gradient $\partial_a V$ is steeper at low a under high risk aversion: an agent near the borrowing constraint faces a large marginal utility of wealth because the prospect of being unable to smooth consumption through a bad productivity shock is more painful when $\gamma = 4$. The iso-value contours tilt: under low γ, productivity z matters relatively more (agents care less about wealth insurance); under high γ, wealth dominates (buffer stock is more valuable than productivity upside).

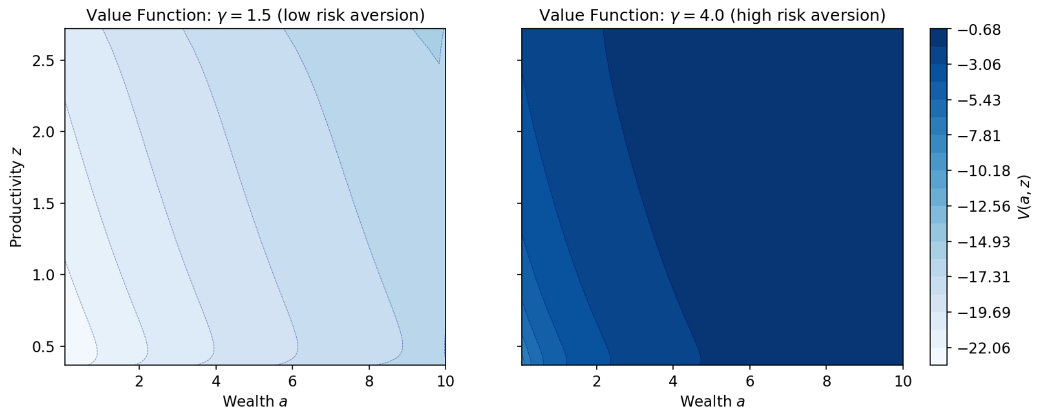

Figure 1: Value function $V(a, z)$ — left: $\gamma = 1.5$; right: $\gamma = 4$. Higher risk aversion steepens the gradient near the borrowing constraint and compresses the value range at high wealth.

Mean-field density (Figure 2). This figure tells the central economic story. Under low risk aversion ($\gamma = 1.5$), agents are relatively tolerant of consumption fluctuations and behave near hand-to-mouth: the equilibrium density peaks sharply at low wealth close to the borrowing constraint $\underline{a}$, with mass decaying rapidly to the right. Agents do not find it worthwhile to accumulate a large buffer stock. Under high risk aversion ($\gamma = 4$), the precautionary saving motive is much stronger: agents actively build wealth to self-insure against bad productivity draws. The density peak shifts substantially rightward, the distribution spreads across a wider range of a, and the borrowing constraint is less binding. The implied aggregate capital is correspondingly higher under $\gamma = 4$ — a direct general equilibrium consequence of precautionary demand.

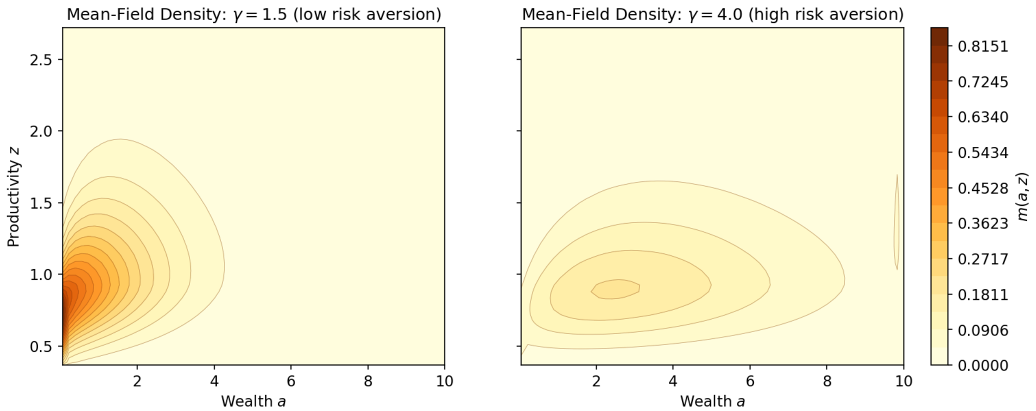

Figure 2: Mean-field density $m(a, z)$ — left: $\gamma = 1.5$, mass concentrated near the borrowing constraint (hand-to-mouth); right: $\gamma = 4$, mass shifted rightward with wider spread (buffer-stock saving). The contrast illustrates how risk aversion shapes the equilibrium wealth distribution.

Optimal consumption policy (Figure 3). The policy $c^*(a, z)$ is monotone increasing in both a and z in both panels, as expected from the envelope condition and the fact that higher productivity raises permanent income. Under $\gamma = 1.5$, agents consume more aggressively out of current wealth and income: the marginal propensity to consume is higher and the contours are more steeply sloped in a. Under $\gamma = 4$, consumption is more cautious — agents suppress current consumption to protect the buffer stock, particularly at low a where the precautionary motive is strongest. The gap between the two panels is largest near the borrowing constraint and narrows at high wealth, where the constraint is slack and both types of agent behave more similarly.

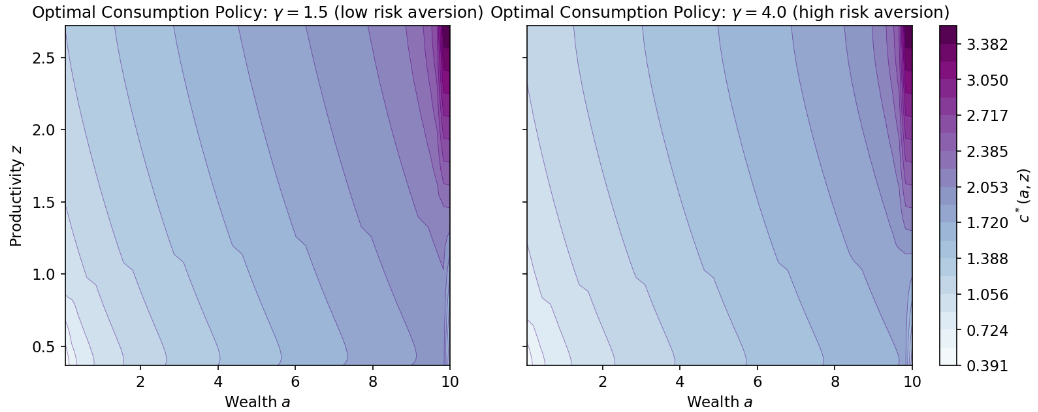

Figure 3: Optimal consumption policy $c^*(a, z)$ — left: $\gamma = 1.5$, more aggressive; right: $\gamma = 4$, more cautious. The largest difference appears near the borrowing constraint, where precautionary motives are strongest.

9. Conclusion and Future Work

This paper establishes that the Huggett-Moll mean field game with OU-Kou productivity admits a complete PIDE-to-PDE reduction, yielding a system of four auxiliary first-order equations that eliminates all integral operators. The reduction is clean, exact, and independent of the form of the drift — it exploits only the exponential kernel of the Kou distribution. The numerical implementation solves the resulting PDE system on a 2D grid via implicit time stepping and direct FOC policy recovery, with the FP equation solved as a stationary linear system. The equilibrium figures confirm that risk aversion is a quantitatively important determinant of the wealth distribution, consistent with the buffer-stock savings literature.

Several directions for future work present themselves, ordered roughly by ambition.

More realistic income dynamics. The symmetric OU-Kou model used here imposes equal intensities $\lambda^+ = \lambda^-$ and tail parameters $\eta^+ = \eta^-$. Empirically, negative income shocks are larger and more frequent than positive ones [10]. Asymmetric Kou calibrated to matched-moments from longitudinal earnings data (PSID or administrative records) would sharpen the quantitative predictions. Beyond Kou, the Variance Gamma process and the Normal Inverse Gaussian distribution offer additional flexibility in fitting the full empirical income distribution, at the cost of losing the exponential-kernel reduction — requiring either numerical PIDE solvers or a generalized auxiliary ODE approach.

Heterogeneous agents with multiple state variables. The two-dimensional state space (a, x) of this paper could be extended to include health status, human capital accumulation, or age in an overlapping-generations structure. The PIDE reduction generalizes naturally to any number of jump-driven state variables with exponential kernels, producing one auxiliary equation per jump direction per state. However, the grid grows exponentially in dimension, motivating sparse grid or deep learning approximations for value functions in four or more dimensions.

Richer market structure. The Huggett-Moll model closes markets through a single asset earning return r. A more realistic extension would include housing (an illiquid asset with transaction costs), equity (a risky asset with correlated shocks), and a borrowing constraint that tightens endogenously with productivity — capturing the empirical finding that credit access is procyclical. The MFG system would then involve a higher-dimensional control problem with portfolio choice alongside consumption.

Howard policy iteration. As noted in Section 6, the present solver recovers c^* by direct FOC at each implicit step. True Howard policy iteration — exact policy evaluation by solving $(\rho I - T(c^k))V = u(c^k)$ followed by FOC improvement — converges in 3–5 iterations rather than hundreds and is the method used in AHLLM [3]. Implementing Howard for the OU-Kou generator (which includes a spatially varying upwind coefficient $-\kappa x$) is the immediate next computational step.

Simultaneous HJB-FP coupling via fictitious play. The present outer loop couples HJB and FP only through prices (r, w): the HJB is solved to full convergence, then the FP is solved once, and capital $K[m]$ drives a bisection update. A more aggressive coupling — **fictitious play** — would interleave the two solves: update c^* from a partially converged HJB, immediately solve FP under the new policy, update prices, and repeat. This can converge faster for non-stationary problems or when the mean-field coupling operates beyond the price channel (e.g., when the HJB depends on m directly through externalities or congestion effects). Establishing convergence rates for fictitious play under OU-Kou dynamics is an open theoretical question.

Deep learning and neural solvers. For high-dimensional extensions, physics-informed neural networks (PINNs) and deep Galerkin methods [11] offer mesh-free approximations of both V and m. The auxiliary PDE structure derived here translates directly into additional loss terms, replacing the four integral constraints with four first-order PDE residuals — a natural fit for automatic differentiation frameworks. Whether neural solvers can match the accuracy of sparse direct linear solvers in two dimensions remains to be established; the value of the deep approach lies in dimensions four and above.

10. References

1. Huggett M (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control* 17(5–6): 953–969.

2. Aiyagari SR (1994). Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics* 109(3): 659–684.

3. 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.

4. Kou SG (2002). A jump-diffusion model for option pricing. *Management Science* 48(8): 1086–1101.

5. Lasry J-M, Lions P-L (2007). Mean field games. *Japanese Journal of Mathematics* 2(1): 229–260.

6. Huang M, Malhamé RP, Caines PE (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.

7. Achdou Y, Buera FJ, Lasry J-M, Lions P-L, Moll B (2014). Partial differential equation models in macroeconomics. *Philosophical Transactions of the Royal Society A* 372: 20130397.

8. Nuño G, Moll B (2018). Social optima in economies with heterogeneous agents. *Review of Economic Dynamics* 28: 150–180.

9. Øksendal B, Sulem A (2007). *Applied Stochastic Control of Jump Diffusions*. Springer, Berlin.

10. Meghir C, Pistaferri L (2004). Income variance dynamics and heterogeneity. *Econometrica* 72(1): 1–32.

11. Sirignano J, Spiliopoulos K (2018). DGM: A deep learning algorithm for solving partial differential equations. *Journal of Computational Physics* 375: 1339–1364.