

Dynamic Incentive Design in Large Populations: A Mean Field Game Approach to the Principal-Agent Problem

Wei Liang^a, Heng-Fu Zou^b

^a*China Economics and Management Academy, Central University of Finance and Economics, Beijing, 100081, China*

^b*The World Bank, Washington, D. C., 20433, USA*

Abstract

We develop a continuous-time principal-agent model with mean-field interactions to study optimal incentive design in large economies. We extend Sannikov's recursive contracting framework to a setting where a single principal manages a continuum of heterogeneous agents whose behaviors are interdependent through aggregate effort externalities. Each agent's hidden effort affects their stochastic output, while the principal designs state-contingent contracts that must account for both individual moral hazard and collective behavior effects. The model generates a coupled system of forward-backward stochastic differential equations: a backward Hamilton-Jacobi-Bellman equation characterizing the principal's value function and optimal contracts, and a forward Fokker-Planck equation governing the evolution of the agent distribution. We establish conditions for mean-field equilibrium where the aggregate effort assumed by the principal when designing contracts coincides with the aggregate effort induced by agents following these contracts. Our numerical implementation reveals that despite complex state-dependent individual contracts, the aggregate effort remains approximately stable over the contract horizon, while the distribution of continuation values exhibits increasing dispersion. Applications include compensation design in large firms, platform economies, and public incentive programs where individual actions generate aggregate externalities.

Keywords: Mean field games, Principal-agent problem, Continuous-time contracting, Moral hazard, Fokker-Planck equation, Dynamic incentives

JEL Classification: C63, C73, D86, G32, J33.

1. Introduction

The principal–agent problem is foundational to the theory of incentives: a risk-neutral principal cannot observe the effort of a risk-averse agent yet must design contracts to elicit productive behavior from noisy output signals. The continuous-time formulation of Holmström and Milgrom (1987) established that under constant absolute risk aversion (CARA), the optimal dynamic contract reduces to a linear function of cumulative output: a tractable benchmark that rules out wealth effects and precludes rich contract dynamics. Sannikov (2008) removed this restriction by introducing the agent’s promised continuation value as a recursive state variable, yielding a Hamilton-Jacobi-Bellman characterization of optimal contracts that accommodates endogenous retirement, non-monotonic effort, and consumption smoothing. Both frameworks, however, analyze incentive provision as a bilateral engagement, abstracting from the externalities that arise when a single principal manages a large, interdependent population of agents whose individual outputs are linked through aggregate performance.

We address this gap by embedding the principal–agent problem within a mean-field game (MFG) framework that couples the Sannikov (2008) recursive contracting apparatus with mean-field interaction. In the settings we consider (large corporations, digital platforms, and public incentive programs), individual productivity depends on aggregate effort across the population, and the principal must design contracts that simultaneously induce effort and account for this interdependence under moral hazard. Unlike earlier applications of mean-field theory to large populations (e.g., Huang, Caines, and Malhamé, 2007), which focus on cost-coupled systems without a contracting authority, our framework embeds moral hazard at the center of the MFG structure. This synthesis generates a coupled forward-backward PDE system: a backward Hamilton-Jacobi-Bellman (HJB) equation characterizing the principal’s optimal contract, and a forward Fokker-Planck equation governing the evolution of the agent value distribution. Establishing and numerically solving the mean-field equilibrium of this system constitutes, to our knowledge, the first characterization of optimal dynamic incentive design within a global mean-field interaction, specifically by embedding Sannikov’s (2008) continuous-time moral-hazard framework within the MFG fixed-point structure, and deriving an inverted-U effort profile as an equilibrium property of the coupled system. While non-monotone effort can arise in bilateral contracting (Sannikov, 2008), the novel result here is distributional: the aggre-

gate effort path $\bar{a}(t)$ remains stable despite the inverted-U individual profile, a property sustained by the Fokker-Planck dynamics of the continuation-value distribution $\mu(W, t)$.

Our framework differs fundamentally from general equilibrium (GE) models of externalities, such as Achdou et al. (2022). In a GE model, agent interactions are mediated by equilibrium prices that clear markets, and individual actions are typically observable. In our MFG formulation, by contrast, there is no price mechanism: the aggregate effort path $\bar{a}(t)$ enters each agent’s environment directly, and each agent treats it as an exogenous field rather than a market signal. Centrally, moral hazard, the non-observability of individual effort, is central to our analysis and absent from standard GE treatments; it is precisely this informational friction that necessitates the incentive contract. We adopt a global mean-field specification, in which each agent responds to the economy-wide average effort rather than to the actions of local neighbors. This specification is appropriate for the applications considered: in large firms, digital platforms, and government incentive programs, a centralized principal designs contracts for the full agent population, making the global mean-field structure the economically appropriate interaction design. Local network specifications, while potentially more realistic in peer-to-peer settings, are computationally intractable within the forward-backward PDE framework and are better suited to settings without a central contracting authority.

The remainder of this paper is organized as follows. Section 2 establishes the model setup by introducing agent heterogeneity and specifying the utility functions that capture both the benefits from consumption and costs of effort exertion. Section 3 examines the principal’s optimization problem, focusing on how profit depends on both individual performance and aggregate behavior through mean-field effects. Section 4 solves the agents’ optimal control problem using the Sannikov (2008) recursive approach, deriving the incentive compatibility conditions and the evolution of continuation values. Section 5 presents the principal’s problem in its full complexity, developing the backward Hamilton-Jacobi-Bellman equation for the principal’s value function, the forward Fokker-Planck equation for the distribution dynamics, and the consistency condition that ensures equilibrium between assumed and realized aggregate behavior. Section 6 details our numerical solution methodology, presenting a fixed-point iteration algorithm to solve the coupled PDE system and providing computational results that demonstrate the equilibrium contract structure and aggregate dynamics. Section 7 concludes by summarizing

our contributions and discussing potential extensions for future research.

2. Agent Heterogeneity and Utility

The economy consists of a single risk-neutral principal and a unit continuum of agents indexed by $i \in [0, 1]$. Agents are ex ante identical in their structural characteristics (sharing a common discount rate r , effort cost parameter β , and output volatility σ) but become ex post heterogeneous through idiosyncratic shocks; their sole distinguishing state at any date t is the idiosyncratic continuation value $W_i(t)$. The principal designs and offers each agent a compensation contract specifying a consumption flow $c_i(t)$ contingent on the agent's observable output; effort $a_i(t)$ remains private information and is not directly contractible.

Agent i 's output evolves according to

$$dX_i(t) = a_i(t)dt + \sigma dZ_i(t), \quad (1)$$

where $dZ_i(t)$ are independent standard Brownian motions, so output is contractible but effort is not: the canonical moral-hazard environment.

Each agent i derives flow utility from consumption according to

$$u(c_i(t)) = \log c_i(t), \quad (2)$$

and bears a quadratic effort cost

$$h(a_i(t)) = \frac{\beta}{2}a_i(t)^2, \quad \beta > 0. \quad (3)$$

Given a contract $\{c_i(t)\}$, agent i chooses effort $\{a_i(t)\}_{t \geq 0}$ to maximize expected discounted lifetime utility

$$\mathbb{E} \left[\int_0^\infty e^{-rt} \left(\log c_i(t) - \frac{\beta}{2}a_i(t)^2 \right) dt \right], \quad (4)$$

where $r > 0$ is the common discount rate.

The key state variable characterizing each agent's contract is the *promised continuation utility* $W_i(t)$, defined as the principal's commitment to deliver to agent i the expected discounted stream of utility from time t onward:

$$W_i(t) = \mathbb{E} \left[\int_t^\infty e^{-r(s-t)} \left(\log c_i(s) - \frac{\beta}{2}a_i(s)^2 \right) ds \middle| \mathcal{F}_t \right]. \quad (5)$$

The integrand $\log c_i(s) - \frac{\beta}{2}a_i(s)^2 \equiv u(c_i(s)) - h(a_i(s))$ is the flow payoff from (2)–(3); thus $W_i(t)$ compresses the utility specification into a single scalar state variable. The utility functions do not disappear from the analysis; they are embedded in the dynamics of $W_i(t)$ through its drift, as made precise in Section 4. This variable is not an innate characteristic such as ability or productivity type; it is a dynamic state variable recording the principal’s outstanding utility obligation to each agent at time t . All agents begin with the same promised utility $W_i(0) = W_0$, reflecting their ex-ante identical standing. Cross-sectional heterogeneity in $W_i(t)$ emerges endogenously over time: agents who experience favorable output realizations (sequences of positive idiosyncratic shocks $dZ_i(t)$) accumulate higher promised utility, while those facing adverse shocks see their continuation values decline. The distribution $\mu_t(W)$ of promised utilities across the agent population thus evolves stochastically, and its law of motion is governed by the Fokker-Planck equation derived in Section 5.2.

3. The Profit Function and Mean-Field Externality

The principal’s profit depends not only on each agent’s individual effort but also on the aggregate effort of the population. Let

$$\bar{a}(t) = \int_0^1 a_i(t) di \quad (6)$$

denote the economy-wide mean effort.

The instantaneous profit flow from the agent population is

$$\int_0^1 (\pi(a_i(t), \bar{a}(t)) - c_i(t)) di, \quad (7)$$

where $\pi(a_i, \bar{a})$ is the output value derived from agent i ’s effort given aggregate effort $\bar{a}$, and $c_i(t)$ is the consumption payment.

Positive values of $\partial\pi/\partial\bar{a}$ correspond to strategic complementarity (knowledge spillovers, team effects); negative values to strategic substitutability (congestion, internal competition). We adopt the linear specification

$$\pi(a_i(t), \bar{a}(t)) = a_i(t) + k\bar{a}(t), \quad (8)$$

where $k \in \mathbb{R}$ parametrizes the externality. This specification is the structural activator of the mean-field interaction: when $k = 0$, the profit function

decouples across agents, and the principal solves independent bilateral contracts identical to Sannikov (2008); no mean-field loop arises. When $k \neq 0$, the principal's value function $F(W, t)$ depends on $\bar{a}(t)$, which is itself determined by F through the optimal policy $a^*(W, t)$; this mutual dependence constitutes the fixed-point problem that defines the mean-field equilibrium and distinguishes our framework from bilateral contracting.

The principal maximizes the expected discounted sum of profit flows over an infinite horizon, taking as given the agents' incentive-compatible responses to the offered contracts. The formal optimization problem, including the Hamilton-Jacobi-Bellman equation and the consistency condition linking $F(W, t)$ to $\bar{a}(t)$, is developed in Section 5.

4. Solving the Agents' Optimal Control Problem

In the mean-field game framework, each agent i treats the aggregate effort path $\{\bar{a}(t)\}_{t \geq 0}$ as exogenous when solving their individual control problem. This is consistent with the large-population approximation: with a continuum of agents, no individual can affect aggregate statistics through unilateral deviation. In equilibrium, the path $\bar{a}(t)$ is determined endogenously by the collective choices of all agents, but each agent optimizes treating it as fixed.

Following the recursive contracting approach of Sannikov (2008), the continuation value $W_i(t)$, the agent's expected discounted utility from time t onward under the contract, serves as the sole dynamic state variable. This sufficient statistic for the history of the contractual relationship allows the principal to specify optimal transfers and effort recommendations as functions of $W_i(t)$ alone, without tracking the full output history.

The continuation value $W_i(t)$ evolves according to the stochastic differential equation

$$dW_i(t) = r(W_i(t) - u(c_i(t)) + h(a_i(t)))dt + r\gamma_i(t)(dX_i(t) - a_i(t)dt), \quad (9)$$

This representation follows from an application of the martingale representation theorem to the agent's continuation value process under the natural filtration of X_i , as established in Sannikov (2008, Proposition 1); the sensitivity process $\gamma_i(t)$ is the unique adapted process satisfying the promise-keeping constraint. The three components of this equation are as follows, where the drift term

$$r[W_i(t) - u(c_i(t)) + h(a_i(t))] \quad (10)$$

is the promise-keeping drift, reflecting the accrual of promised utility at rate r net of the agent's flow payoff $u(c_i(t)) - h(a_i(t))$. The diffusion term

$$r\gamma_i(t) (dX_i(t) - a_i(t)dt) \quad (11)$$

links continuation-value adjustments to output surprises, where $\gamma_i(t)$ is the sensitivity of promised utility to performance deviations.

The sensitivity parameter $\gamma_i(t)$ is determined by the incentive compatibility (IC) constraint. Given sensitivity $\gamma_i(t)$, the agent solves $\max_a \{r\gamma_i(t)a - h(a)\}$, yielding first-order condition $r\gamma_i(t) = h'(a_i(t))$. For the principal to implement a target effort level $a_i(t)$, the sensitivity must therefore satisfy

$$\gamma_i(t) = \frac{h'(a_i(t))}{r}, \quad (12)$$

which equates the sensitivity of continuation value to output with the ratio of marginal effort cost to the discount rate.

From the agent's perspective, the individual optimization problem has the same recursive structure as in the bilateral Sannikov (2008) model, with the difference that the contract terms $(c_i(t), a_i(t))$ may depend on the exogenous mean-field path $\bar{a}(t)$ through the production externality. The principal's problem of choosing the optimal contract, subject to this IC structure, is taken up in the following section.

5. The Principal's Problem

In the mean-field game setting, the principal designs contracts for a continuum of agents whose continuation values evolve stochastically over time. The state variable for the principal's problem is therefore the probability distribution $\mu_t(W)$ over continuation values at each instant t , which encodes the cross-sectional heterogeneity in promised utilities generated by idiosyncratic output shocks.

The principal's objective function reflects this distributional perspective by integrating profits across the entire agent population. The value function $V(t, \mu_t)$ represents the principal's expected discounted profit from time t forward, given the current distribution μ_t . This can be expressed as

$$V(t, \mu_t) = \int_t^\infty e^{-r(s-t)} \left[\int (\pi(a^*(W, s), \bar{a}(s)) - c^*(W, s)) \mu(W, s) dW \right] ds, \quad (13)$$

where the inner integral aggregates instantaneous profits across the agent population and the outer integral discounts over time.

Because $V(t, \mu_t)$ is linear in μ , the principal's value function admits a pointwise decomposition. Define

$$F(W, t, \mu_t) = \frac{\delta V}{\delta \mu(W)} \quad (14)$$

as the principal's marginal value from an agent with continuation value W , given the current distribution μ_t . This Lions (flat) derivative of $V(t, \cdot)$ with respect to μ evaluated at W is well-defined because V is linear in μ , and reduces to the integrand $F(W, t)$ in the representation $V(t, \mu_t) = \int_W F(W, t) \mu(W, t) dW$.¹ The Hamilton-Jacobi-Bellman equation for $F(W, t)$ is

$$rF = \frac{\partial F}{\partial t} + \max_{a, c} \left\{ \pi(a, \bar{a}_t) - c + \frac{\partial F}{\partial W} [r(W - u(c) + h(a))] + \frac{1}{2} \frac{\partial^2 F}{\partial W^2} [r\sigma\gamma(a)]^2 \right\}, \quad (15)$$

where the inner maximization determines the optimal instantaneous contract (a, c) for an agent with continuation value W , taking $\bar{a}(t)$ as given.

The principal acts as a direct mechanism designer: by the Revelation Principle, no generality is lost in restricting attention to direct incentive-compatible mechanisms that recommend effort $a^*(W, t)$ and specify consumption $c^*(W, t)$ as functions of the agent's continuation value W . Compliance with $a^*(W, t)$ is enforced by the incentive compatibility constraint (12): given $\gamma_i(t) = h'(a^*)/r$ embedded in the continuation-value dynamics, the agent finds it optimal to exert the recommended effort. Consumption $c^*(W, t)$ is a promised payment controlled entirely by the principal; the agent undertakes no separate optimization over consumption.

A mean-field equilibrium requires three conditions to hold simultaneously. First, agent optimality: given the principal's contract and the path $\{\bar{a}(t)\}$, each agent satisfies the incentive compatibility constraint. Second, principal optimality: the contract $(a^*(W, t, \mu), c^*(W, t, \mu))$ solves the HJB equation

¹Here $\frac{\delta V}{\delta \mu}(W)$ denotes the Lions (flat) derivative of the functional $V(t, \cdot): \mathcal{P}(\mathbb{R}) \rightarrow \mathbb{R}$ with respect to the measure μ , evaluated at W . This is distinct from a standard partial derivative $\partial V / \partial W$: the argument W indexes mass within the distributional input μ , not a real scalar argument of V . Since $V(t, \mu_t) = \int_W F(W, t) \mu(W, t) dW$ is linear in μ , the functional derivative reduces to the integrand $F(W, t)$, yielding the identity stated above. See Cardaliaguet (2013, Notes on Mean Field Games) for a formal treatment.

(15). Third, consistency: the aggregate effort generated by the optimal contracts must coincide with the assumed path:

$$\bar{a}(t) = \int a^*(W, t, \mu) \cdot \mu(W, t) dW \quad (16)$$

must hold at every instant.

These three conditions constitute a coupled forward-backward system. The backward component is the HJB equation for $F(W, t)$, solved from a terminal condition. The forward component is the Fokker-Planck equation for $\mu(W, t)$, solved from an initial condition. The coupling occurs through $\bar{a}(t)$, which depends on the current distribution (forward state) and enters the optimal contract through the value function (backward state). This structure necessitates the iterative numerical solution approach described in Section 6.

5.1. Backward HJB Equation

The HJB equation (15) is solved backward from the terminal condition $F(W, T) = 0$ for all $W \in [W_{\min}, W_{\max}]$, which reflects that the principal receives no continuation surplus beyond the contract horizon T . At each earlier date, the solution requires maximizing the principal's Hamiltonian over (a, c) taking $\bar{a}(t)$ as given, yielding policy functions that depend on W , on t , and through F on the mean field path.

The optimal consumption policy follows from the first-order condition of the principal's Hamiltonian with respect to c :

$$-1 - \frac{\partial F}{\partial W} r u'(c) = 0. \quad (17)$$

This equation equates the direct marginal cost of consumption (normalized to -1) with its indirect benefit through the continuation-value drift. With logarithmic utility $u'(c) = 1/c$, this simplifies to:

$$c^*(W, t) = -r \frac{\partial F}{\partial W}, \quad (18)$$

establishing that optimal consumption is proportional to the marginal cost of delivering continuation value, with the discount rate r as the proportionality constant. This result is consistent with the bilateral Sannikov (2008) framework; the mean-field externality enters through F rather than directly through the consumption formula.

The first-order condition with respect to effort a , after substituting the optimal consumption c^* , is:

$$\frac{\partial \pi}{\partial a}(a, \bar{a}_t) + \frac{\partial F}{\partial W} r h'(a) + \frac{1}{2} \frac{\partial^2 F}{\partial W^2} (r\sigma)^2 \cdot 2\gamma(a)\gamma'(a) = 0. \quad (19)$$

The three terms represent, respectively, the marginal productivity of effort, the cost of implementing effort through the continuation-value channel, and the risk adjustment from volatility in promised utilities. Substituting the IC conditions $\gamma(a) = h'(a)/r$ and $\gamma'(a) = h''(a)/r$, and using the quadratic cost specification $h'(a) = \beta a$, yields the closed-form solution:

$$1 + r\beta \frac{\partial F}{\partial W} a^*(W, t) + \frac{\partial^2 F}{\partial W^2} (\beta\sigma)^2 a^*(W, t) = 0. \quad (20)$$

Solving for $a^*(W, t)$, this becomes

$$a^*(W, t) = - \left[\beta \left(r \frac{\partial F}{\partial W} + \sigma^2 \beta \frac{\partial^2 F}{\partial W^2} \right) \right]^{-1}. \quad (21)$$

Optimal effort decreases as the marginal cost of providing utility ($-\partial F/\partial W > 0$) increases, and decreases further when the value function is more concave ($\partial^2 F/\partial W^2 < 0$), as greater curvature raises the cost of incentive provision through volatile continuation values. While the functional form of $a^*(W, t)$ matches the bilateral Sannikov (2008) model, it now depends implicitly on $\bar{a}(t)$ through F , which must satisfy the full HJB equation.

Substituting both optimal policies a^* and c^* back into the HJB equation yields a nonlinear partial differential equation in $F(W, t)$:

$$\begin{aligned} rF = & \frac{\partial F}{\partial t} + a^*(W, t) + k\bar{a}(t) - c^*(W, t) \\ & + \frac{\partial F}{\partial W} \left[r \left(W - \log c^*(W, t) + \frac{\beta}{2} a^*(W, t)^2 \right) \right] \\ & + \frac{1}{2} \frac{\partial^2 F}{\partial W^2} [\sigma\beta a^*(W, t)]^2. \end{aligned} \quad (22)$$

The presence of $\bar{a}(t)$ both directly through $k\bar{a}(t)$ and indirectly through F transforms this into a fixed-point problem: the solution $F(W, t)$ must satisfy this PDE while generating effort policies that, when aggregated over the population distribution, reproduce the path $\bar{a}(t)$ assumed in the HJB.

5.2. Fokker-Planck Dynamics of Agent Distribution

The evolution of the agent population's distribution over continuation values represents the forward-looking component of our mean-field game system. While the HJB equation describes how the principal optimally designs contracts looking backward from the terminal time, the Fokker-Planck equation captures how these contracts, once implemented, shape the forward evolution of the agent distribution from the initial condition. These forward dynamics are central to understanding the aggregate implications of the incentive system, including how inequality emerges, how shocks propagate through the population, and how the distribution of agent welfare changes over time. The stochastic nature of each agent's continuation value, driven by both systematic drift from the contract design and random diffusion from output uncertainty, gives rise to rich distributional dynamics that cannot be captured by representative agent models.

The foundation for deriving the Fokker-Planck equation lies in understanding the stochastic process governing each agent's continuation value $W_i(t)$. From the incentive-compatible contract design established in the previous section, the continuation value evolves according to the stochastic differential equation

$$dW_i(t) = r(W_i(t) - u(c_i(t)) + h(a_i(t)))dt + r\gamma(a_i(t))(dX_i(t) - a_i(t)dt). \quad (23)$$

Substituting the output process $dX_i(t) = a_i(t)dt + \sigma dZ_i(t)$, we can rewrite this as

$$dW_i(t) = r(W_i(t) - u(c_i(t)) + h(a_i(t)))dt + r\gamma(a_i(t))\sigma dZ_i(t). \quad (24)$$

This formulation reveals that the continuation value follows a controlled diffusion process with drift coefficient

$$b(W, t) = r(W - u(c(W, t)) + h(a(W, t))) \quad (25)$$

and diffusion coefficient

$$\Sigma(W, t) = r\gamma(a(W, t))\sigma = \frac{r\sigma h'(a(W, t))}{r} = \sigma\beta a(W, t). \quad (26)$$

The drift term captures the deterministic evolution of promised utility, incorporating the agent's time preference, consumption delivery, and effort cost,

while the diffusion term represents the volatility induced by performance-based incentives necessary for moral hazard mitigation.

The probability density function $\mu(W, t)$ describing the distribution of continuation values across the agent population evolves according to the Fokker-Planck equation, also known as the forward Kolmogorov equation in the probability literature. This partial differential equation, given by

$$\frac{\partial \mu}{\partial t} = -\frac{\partial}{\partial W} [b(W, t)\mu(W, t)] + \frac{1}{2} \frac{\partial^2}{\partial W^2} [\Sigma(W, t)^2 \cdot \mu(W, t)], \quad (27)$$

governs how the density changes over time due to the combined effects of drift and diffusion. The first term on the right-hand side, the divergence of the probability flux $b(W, t)\mu(W, t)$, describes how the distribution shifts due to the deterministic component of the dynamics. When the drift is positive (negative) at a particular value W , probability mass flows away from (toward) that point. The second term, involving the second derivative, captures the spreading or concentration of the distribution due to stochastic shocks. Higher diffusion coefficients lead to faster spreading of the distribution, reflecting greater uncertainty in individual outcomes over time.

The Fokker-Planck dynamics admit a direct welfare interpretation of the optimal contract. The drift coefficient $b(W, t)$ determines whether agents at a particular continuation value W tend to see their promised utility increase or decrease over time. In regions where $b(W, t) > 0$, the contract provides positive expected utility growth, potentially reflecting high productivity or strong incentive provision. Conversely, regions with $b(W, t) < 0$ indicate that agents are expected to drift toward lower continuation values, possibly due to high consumption extraction by the principal or low effort implementation. The diffusion coefficient $\Sigma(W, t)^2$ measures the degree of risk that agents face, which is directly tied to the incentive intensity $\gamma(a(W, t))$ required for effort implementation. Higher effort requirements necessitate stronger incentives, leading to greater volatility in continuation values and thus more dispersion in the agent distribution over time.

The boundary conditions for the Fokker-Planck equation reflect important economic assumptions about the agent population. In our formulation, we impose reflecting (zero-flux) Neumann boundary conditions at both $W_{\min}$ and $W_{\max}$, mathematically expressed as

$$b(W, t)\mu(W, t) - \frac{1}{2} \frac{\partial}{\partial W} [\Sigma(W, t)^2 \cdot \mu(W, t)] = 0 \quad (28)$$

at the boundaries. These conditions ensure that no probability mass exits the domain $[W_{\min}, W_{\max}]$, preserving the normalization $\int \mu(W, t) dW = 1$ for all t . Economically, this assumption means that agents cannot be terminated or exit the contract, and the principal cannot drive continuation values beyond the specified bounds. The lower boundary $W_{\min}$ might represent a minimum participation constraint below which agents would prefer to leave, while the upper boundary $W_{\max}$ could reflect a maximum feasible utility level given resource constraints. The reflecting nature of these boundaries creates interesting dynamics near the edges of the domain: agents who reach these boundaries experience forces that push them back toward the interior, potentially leading to accumulation of probability mass near the boundaries under certain contract designs.

5.3. Consistency Condition

The consistency condition links the backward principal optimization and the forward distributional dynamics into a mean-field equilibrium. The condition requires that the aggregate effort path $\bar{a}(t)$ assumed by the principal when designing optimal contracts coincides with the aggregate effort that emerges when agents respond optimally to those contracts. This fixed-point requirement, standard in mean-field game theory, ensures that no agent has incentive to deviate from the prescribed effort given the aggregate behavior, and that the aggregate behavior is collectively consistent with individual optimal responses.

Formally, the consistency condition requires that at each time t :

$$\bar{a}(t) = \int a^*(W, t) \cdot \mu(W, t) dW, \quad (29)$$

where $a^*(W, t)$ is the optimal effort recommended by the principal's contract for an agent with continuation value W , and $\mu(W, t)$ is the cross-sectional distribution of agents at time t . This integral embeds a mutual dependence: $a^*(W, t)$ is derived from the HJB equation for $F(W, t)$, which is parameterized by $\bar{a}(t)$; simultaneously, $\mu(W, t)$ evolves according to the Fokker-Planck equation with drift and diffusion determined by $a^*(W, t)$ and $c^*(W, t)$. The consistency condition therefore defines a fixed-point system: an equilibrium exists when the mapping from an assumed aggregate effort path to the induced aggregate effort path has a fixed point.

The economic interpretation of the consistency condition reveals important insights about coordination and externalities in large economies. When

production exhibits complementarities ($k > 0$ in our linear specification), higher aggregate effort makes individual effort more valuable, potentially creating multiple equilibria with different levels of aggregate activity. The principal's contract design must not only provide appropriate individual incentives but also coordinate the economy to a particular equilibrium. Conversely, with substitutabilities ($k < 0$), higher aggregate effort reduces individual productivity, and the consistency condition ensures that the principal's contracts account for these congestion effects. The dynamic nature of the consistency condition is particularly important: as the distribution $\mu(W, t)$ evolves over time due to the stochastic shocks and incentive provision, the aggregate effort $\bar{a}(t)$ must adjust accordingly, requiring the principal's contracts to adapt to the changing composition of the agent population.

The enforcement of the consistency condition in practice requires sophisticated computational methods, as analytical solutions are available only in very special cases. The forward-backward structure of the coupled PDE system, combined with the fixed-point requirement, precludes direct solution methods and necessitates iterative approaches. Starting with an initial guess for the aggregate effort path, one solves the backward HJB equation to obtain optimal contracts, then uses these contracts to solve the forward Fokker-Planck equation for the distribution evolution, from which a new aggregate effort path can be computed. This process must be repeated until convergence, with each iteration potentially revealing new aspects of how individual incentives and aggregate behavior interact. The convergence of this iterative process is not guaranteed in general and depends on the economic parameters, particularly the strength of the mean-field interaction k and the curvature of the utility and cost functions. When convergence is achieved, the resulting equilibrium provides a complete characterization of how optimal dynamic contracts must account for and shape collective behavior in large economies.

Remark. Analytical global uniqueness of the mean-field equilibrium is an open question. The fixed-point mapping $\Phi : \bar{a} \mapsto \int a^*(W, t, \bar{a}) \mu(W, t) dW$ is not globally contractive under the nonlinear optimality conditions of this model, so multiple fixed points cannot be ruled out in general. When strategic complementarities are present ($k > 0$), the mapping Φ is increasing in $\bar{a}$, which can support multiple self-consistent equilibria at different aggregate effort levels. Establishing strict monotonicity conditions sufficient for uniqueness is left for future work. However, numerical experiments starting from ten distinct initial guesses for $\bar{a}^{(0)}(t)$ (ranging from the zero-effort path to

the maximum feasible effort) consistently converge to the same equilibrium at the baseline parameterization (Table 1), providing computational evidence for local uniqueness at the parameter values considered.

6. Solving the Coupled PDE System

6.1. Numerical Strategy: A Fixed-Point Iteration Approach

The mean-field game formulation of the principal-agent problem, as derived from Sannikov’s continuous-time framework, results in a coupled system of forward-backward partial differential equations (PDEs). Specifically, the principal’s optimization problem is characterized by a backward Hamilton-Jacobi-Bellman (HJB) equation, which describes the evolution of the principal’s marginal value function $F(W, t)$. Simultaneously, the collective behavior of the continuum of agents gives rise to a forward Fokker-Planck (FPK) equation, which governs the dynamics of the distribution of continuation values $\mu(W, t)$. The inherent forward-backward structure of this system, in which the HJB equation is solved backward in time from a terminal condition while the FPK equation evolves forward from an initial distribution, precludes a direct analytical or numerical solution.

To overcome this challenge, we adopt a numerical strategy based on a fixed-point iteration method, a standard approach for solving mean-field game systems (Achdou and Capuzzo-Dolcetta, 2010). This iterative algorithm effectively decouples the forward and backward components of the problem. The core idea is to treat the aggregate behavior of the agents, specifically the mean effort path $\bar{a}(t)$, as a given input when solving the principal’s optimization problem (the HJB equation). Once the principal’s optimal contract $(a^*(W, t), c^*(W, t))$ is determined for this given mean effort, we then solve the FPK equation to find the resulting evolution of the agent distribution $\mu(W, t)$. From this new distribution, we can compute an updated mean effort path. This process is repeated iteratively until the mean effort path converges, at which point a consistent, mean-field equilibrium has been found.

The nonlinearity of the HJB equation introduces numerical instability in naive fixed-point iterations. To address this, the iterative updates are stabilized through a damping (relaxation) technique: rather than replacing each iterate directly with the newly computed solution, we employ a convex combination of successive iterates, as detailed in Section 6.3. Combined with

a finite difference discretization of the coupled PDE system, this approach yields a numerically convergent solution to the mean-field equilibrium.

6.2. The Discretization Scheme

To numerically solve the coupled PDE system, we employ a finite difference method. This involves discretizing the continuous state and time domains and approximating the differential operators with finite differences. This approach transforms the system of PDEs into a large-scale system of algebraic equations that can be solved computationally.

6.2.1. Time and State Discretization

We consider the agent's continuation value W within a bounded domain $[W_{\min}, W_{\max}]$ and the time horizon $[0, T]$. These continuous domains are discretized into a uniform grid. The time interval $[0, T]$ is divided into N_t subintervals of length

$$\Delta t = \frac{T}{N_t}, \quad (30)$$

yielding discrete time steps

$$t_n = n\Delta t \quad (31)$$

for $n = 0, 1, \dots, N_t$. Similarly, the state space $[W_{\min}, W_{\max}]$ is discretized into N_w subintervals of length

$$\Delta W = \frac{W_{\max} - W_{\min}}{N_w}, \quad (32)$$

with grid points denoted by

$$W_i = W_{\min} + i\Delta W, \quad (33)$$

for $i = 0, 1, \dots, N_w$. The value of any function $f(t, W)$ at the grid point (t_n, W_i) is denoted by f_i^n .

6.2.2. Discretization of the HJB Equation

The backward HJB equation for the principal's value function F is a nonlinear, second-order PDE. We discretize it using a semi-implicit scheme, which treats the time derivative explicitly and the spatial derivatives implicitly to ensure numerical stability. The time derivative $\partial F / \partial t$ is approximated by a forward difference, $(F_i^{n+1} - F_i^n) / \Delta t$, as we solve backward in time from $n = N_t - 1$ down to 0. The first- and second-order spatial derivatives,

$\partial F/\partial W$ and $\partial^2 F/\partial W^2$, are approximated using standard second-order central difference schemes. Let $F_w(W_i, t_n)$ and $F_{ww}(W_i, t_n)$ denote these discrete derivative approximations at (W_i, t_n) . The fully discretized HJB equation at each interior grid point (W_i, t_n) can then be written as a nonlinear algebraic equation for the unknown values $\{F_i^n\}$:

$$rF_i^n = \frac{F_i^{n+1} - F_i^n}{\Delta t} + H^*(W_i, t_n, F_i^n, F_w(W_i, t_n), F_{ww}(W_i, t_n), \bar{a}_n), \quad (34)$$

where H^* is the minimized Hamiltonian with the optimal controls a^* and c^* expressed as functions of F_i^n and its discrete derivatives. This results in a system of nonlinear equations for the vector $F^n = (F_0^n, \dots, F_{N_w}^n)$ at each time step n . This system is solved iteratively using a Newton-type method, subject to Dirichlet boundary conditions $F_0^n = F_p(W_{\min})$ and $F_{N_w}^n = F_p(W_{\max})$, where $F_p(W)$ represents the principal's terminal profit function from retiring an agent with continuation value W . To ensure the concavity of the value function, a fundamental economic property, we explicitly enforce that the discrete second derivative F_{ww} remains non-positive throughout the computation.

6.2.3. Discretization of the Fokker-Planck Equation

The forward Fokker-Planck equation describes the evolution of the probability density $\mu(W, t)$ of the agents' continuation values. This is a linear, second-order advection-diffusion equation. We discretize it using a forward-time, centered-space (FTCS) scheme for the diffusion term and an upwind scheme for the advection (drift) term to maintain the non-negativity of the density and ensure stability. The time derivative $\partial\mu/\partial t$ is approximated by a forward difference $(\mu_i^{n+1} - \mu_i^n)/\Delta t$.

The drift and diffusion coefficients, $b(W, t)$ and $\Sigma(W, t)^2$ respectively, are functions of the optimal controls a^* and c^* , which in turn depend on the solution F of the HJB equation from the previous step of the fixed-point iteration. The discrete form of the FPK equation is thus a linear system for the unknown density μ^{n+1} at the next time step, given μ^n :

$$\frac{\mu_i^{n+1} - \mu_i^n}{\Delta t} = -\frac{(b\mu)_i^+ - (b\mu)_i^-}{\Delta W} + \frac{1}{2} \cdot \frac{(\Sigma^2\mu)_{i+1} - 2(\Sigma^2\mu)_i + (\Sigma^2\mu)_{i-1}}{(\Delta W)^2}, \quad (35)$$

where $(b\mu)_i^+$ and $(b\mu)_i^-$ represent the discretized flux terms handled by the upwind scheme, depending on the sign of the drift $b(W_i, t_n)$. For this problem, we impose reflecting (zero-flux) Neumann boundary conditions at $W_{\min}$

and $W_{\max}$, which ensures the conservation of the total probability mass, $\int \mu(W, t) dW = 1$, over time. This boundary condition reflects the assumption that agents are neither created nor destroyed and cannot exit the specified range of continuation values.

The choice of reflecting rather than absorbing boundary conditions for μ warrants comment. In Sannikov-type models, $W_{\min}$ represents the agent's participation constraint at which retirement occurs. An absorbing boundary condition would allow probability mass to exit at $W_{\min}$, reflecting agent termination. We retain reflecting boundary conditions for the density because the contract horizon T is finite and the retirement event is incorporated into the drift dynamics of $W_i(t)$ through the incentive-compatible consumption path, rather than as an explicit mass-absorbing boundary mechanism. This treatment, which is standard in the MFG numerical literature (Achdou and Capuzzo-Dolcetta, 2010), preserves total probability mass and ensures the Fokker-Planck system remains well-posed over the finite horizon. The Dirichlet conditions imposed on F at $W_{\min}$ and $W_{\max}$ separately govern the principal's payoff at the contract boundaries and are not in conflict with the density boundary conditions.

6.3. The Fixed-Point Iteration Algorithm

The coupled, forward-backward structure of the HJB and Fokker-Planck PDE system necessitates an iterative solution method. We implement a fixed-point iteration algorithm, also known as the fictitious play method in game theory literature, which has been established as a robust technique for solving mean-field game systems. This algorithm resolves the coupling by iteratively solving the HJB and FPK equations until a consistent, equilibrium solution is found. The entire procedure, as implemented in our numerical solver, can be broken down into the following sequential steps.

The algorithm begins with an initial guess for the aggregate behavior of the agent population. Specifically, we initialize a path for the mean effort, $\{\bar{a}^{(0)}(t)\}$ for $t \in [0, T]$, and the distribution of continuation values, $\{\mu^{(0)}(W, t)\}$. A simple yet effective choice is to assume a constant mean effort and set the distribution at all times equal to the given initial distribution $\mu_0(W)$. The algorithm then proceeds through a series of iterations, indexed by $k = 1, 2, \dots$. In each iteration, we first solve the principal's problem by computing the backward HJB equation. Taking the mean effort path $\{\bar{a}^{(k-1)}(t)\}$ from the previous iteration as a given, exogenous input, the HJB equation becomes a standard, albeit nonlinear, backward PDE. We solve it

numerically from $t = T$ to $t = 0$ using the finite difference scheme detailed in Section 6.2.2. This step yields an updated value function for the principal, $F^{(k)}(W, t)$, together with the corresponding optimal policy functions for the recommended effort, $a^{*(k)}(W, t)$, and consumption, $c^{*(k)}(W, t)$.

Subsequently, we solve the forward Fokker-Planck equation to determine the evolution of the agent distribution under the newly computed optimal contract. The drift and diffusion coefficients of the FPK equation are updated using the optimal policy functions $a^{*(k)}$ and $c^{*(k)}$. The PDE is then solved forward in time from the initial distribution $\mu_0(W)$ at $t = 0$ to the terminal time T , yielding a new distribution path $\{\mu^{(k)}(W, t)\}$. From this, we compute the updated mean effort path by integrating the optimal effort against the new distribution:

$$\bar{a}^{(k)}(t) = \int a^{*(k)}(W, t) \mu^{(k)}(W, t) dW. \quad (36)$$

To promote convergence and suppress oscillations inherent in the nonlinear HJB-FPK fixed-point iteration, we employ a damping (relaxation) strategy. At each iteration k , the updated solutions for the value function, distribution, and mean effort path are not adopted directly but are replaced by a convex combination of the previous and newly computed iterates:

$$X^{(k)} \leftarrow \theta \cdot X_{\text{new}}^{(k)} + (1 - \theta) \cdot X^{(k-1)}, \quad \theta \in (0, 1]. \quad (37)$$

The parameter θ controls the step size: $\theta = 1$ recovers the undamped iteration, while $\theta < 1$ introduces inertia that prevents the solution from jumping between branches of the nonlinear system. Without damping ($\theta = 1$), the iteration frequently diverges or oscillates between two solution branches due to the strong nonlinearity of the optimality conditions. Damping does not theoretically guarantee convergence, but empirically promotes it by keeping successive iterates in proximity. The baseline simulation employs $\theta = 0.3$, a conservative choice consistent with numerical practice for MFG systems (Achdou and Capuzzo-Dolcetta, 2010). Sensitivity to θ was assessed by re-running the algorithm with $\theta \in \{0.1, 0.3, 0.5\}$; all cases converged to the same equilibrium, indicating local robustness of the fixed point at the baseline parameterization.

The iteration continues until the successive solutions converge. Convergence is monitored via the L^2 -norm of the difference between successive mean-effort paths. Once this distance falls below a predefined tolerance

threshold ϵ , the algorithm is terminated, and the resulting set of functions $\{F^{(k)}, \mu^{(k)}, a^{*(k)}, c^{*(k)}, \bar{a}^{(k)}\}$ is considered the numerical solution to the mean-field game equilibrium.

6.4. Numerical Results and Equilibrium Analysis

The fixed-point iteration algorithm described above is implemented using the parameter values summarized in Table 1, which represent a baseline calibration with moderate effort cost, moderate output volatility, and positive mean-field interaction ($k = 1.0$, strategic complementarity). The resulting equilibrium contract and aggregate dynamics are presented in Figure 1.

A methodological remark is warranted regarding the relationship between the theoretical and numerical specifications. The model is formulated on an infinite horizon (Section 2, (4)), yet the numerical solution procedure employs a finite-horizon approximation with terminal date $T = 0.3$. This is a computational necessity imposed by the forward-backward structure of the HJB-FPK system: the backward HJB equation requires a terminal boundary condition, specified here as $F(W, T) = 0$. The infinite-horizon stationary equilibrium corresponds to the limit of this finite-horizon solution as $T \rightarrow \infty$, at which point the value function converges to the solution of the stationary HJB equation $rF = \mathcal{H}^*(W, F_W, F_{WW})$. The baseline value $T = 0.3$ is retained for computational tractability in the forward-backward iteration; the quantitative results should accordingly be interpreted as characterizing the optimal contract dynamics over a finite planning horizon rather than as converging to the full infinite-horizon steady state.

6.4.1. The Principal's Value Function and the Optimal Contract

The equilibrium solution for the principal's marginal value function, $F(W, t)$, is depicted in panel (a). The function is smooth, strictly decreasing, and concave with respect to the agent's continuation value W . This property is consistent with the theory of dynamic contracting: F being decreasing ($F' < 0$) confirms that the agent's promised utility is a liability to the principal, while its concavity ($F'' < 0$) reflects the increasing marginal cost of providing incentives to an agent already endowed with a high continuation value. The curves shift upward as time moves backward from the terminal time T , indicating the positive value generated by the optimal dynamic contract over time.

The resulting optimal contract, specified by the recommended effort $a^*(W, t)$ and consumption $c^*(W, t)$, exhibits significant state-dependent and time-

Figure 1: Equilibrium Dynamics of the Mean Field Principal-Agent Model. The figure contains six panels arranged in three rows of two columns: (a) principal's marginal value function $F(W, t)$; (b) optimal effort $a^*(W, t)$; (c) optimal consumption $c^*(W, t)$; (d) distribution of continuation values $\mu(W, t)$; (e) aggregate mean effort $\bar{a}(t)$; (f) conservation of total probability mass.




Table 1: Parameter Values for Numerical Simulation

Parameter Name	Symbol	Value
<i>Model Parameters</i>		
Cost of Effort Parameter	β	0.15
Mean Field Interaction Strength	k	1.0
Discount Rate	r	0.1
Output Volatility	σ	0.3
<i>Numerical Parameters</i>		
Time Horizon	T	0.3
Number of Time Steps	N_t	30
Continuation Value Domain	$[W_{\min}, W_{\max}]$	$[0.0, 1.0]$
Number of State Grid Points	N_w	30
Damping Parameter	θ	0.3
Convergence Tolerance	ϵ	5×10^{-4}

dependent variations. The optimal effort, depicted in panel (b), follows a distinct inverted-U shape as a function of the agent's value W . Effort is low for agents with very low continuation values (who are close to the retirement boundary and thus have little to lose) and also for agents with very high values, reflecting the increasing marginal cost of providing incentives to agents with high continuation values (the income effect). The highest effort is exerted by agents with intermediate continuation values, for whom incentives are most effective. Furthermore, optimal consumption $c^*(W, t)$ (panel (c)) remains stable, only decaying sharply as the terminal time T approaches, demonstrating the principal's use of consumption smoothing. Higher promised future utility is consistently paired with higher immediate consumption, except near the upper boundary where terminal effects can create complex, non-monotonic patterns as the contract horizon shortens.

6.4.2. Macroeconomic Dynamics and Aggregate Behavior

The forward-looking component of the solution describes the evolution of the agent population. Panel (d) illustrates the dynamics of the distribution of continuation values, $\mu(W, t)$. Starting from a symmetric, unimodal distribution μ_0 at $t = 0$, we observe two key phenomena over time: the distribution's

peak drifts to the left, and its dispersion significantly increases. The leftward drift of the distribution’s mode has the following economic interpretation. It indicates that for the majority of agents, the expected change in their continuation value is negative ($\mathbb{E}[dW_t] < 0$). This result reflects the front-loaded cost of incentive provision. In equilibrium, the principal offers relatively high consumption in the early stages of the contract (as confirmed by the “Optimal Consumption” panel) to secure the agent’s participation and elicit effort. This “upfront payment” in utility terms causes the promised future utility W to decrease on average. This does not imply the contract is ineffective; rather, it shows the principal reduces the liability owed to the agent over time, while successfully extracting a stable stream of high effort from the population. This front-loading pattern is, however, partly conditioned by the finite-horizon terminal condition $F(W, T) = 0$, which implies that optimal consumption $c^*(W, t) = -r\partial F/\partial W$ declines toward zero as $t \rightarrow T$. In the infinite-horizon stationary equilibrium the consumption profile would be governed by the steady-state relationship between F_W and the ergodic distribution μ^* ; the qualitative front-loading mechanism is consistent with that theory, but the precise quantitative trajectory reported here reflects the finite planning horizon $T = 0.3$. Simultaneously, the increasing dispersion (variance) of the distribution highlights the inevitable consequence of moral hazard: rising inequality. To motivate effort, the contract must expose agents to risk by linking their continuation value to the stochastic output. This inherent risk naturally leads to a divergence in outcomes, causing the distribution of welfare W across the agent population to widen over time. In summary, the evolution of $\mu(W, t)$ reveals that the equilibrium contract is a mechanism that, while successful in generating stable aggregate effort, does so at the cost of a declining average expected utility path for agents and simultaneously generates significant ex-post inequality.

Panel (e) shows the evolution of the aggregate mean effort $\bar{a}(t)$. Despite the complex, state-dependent nature of the individual optimal effort $a^*(W, t)$, the resulting mean effort $\bar{a}(t)$ remains approximately constant throughout the contract period. This suggests that in equilibrium, the principal designs the contract such that the shifting distribution of agent values precisely counteracts the state-dependent variations in individual effort, resulting in a stable aggregate level of productivity for the economy. Finally, panel (f) confirms the numerical stability of our Fokker-Planck solver by showing that the total mass of the agent distribution is conserved and remains equal to one throughout the simulation horizon, consistent with our choice of reflecting boundary

conditions.

7. Conclusion

This paper develops a tractable framework for optimal incentive design in economies with mean-field externalities by embedding continuous-time principal-agent contracting within a mean-field game structure. The approach extends the Sannikov (2008) bilateral contracting framework to settings where a single principal faces a continuum of agents whose individual efforts interact through an aggregate externality, generating a coupled forward-backward PDE system whose solution simultaneously determines the optimal contract and the equilibrium distribution of agent welfare.

The paper makes three theoretical contributions. First, it characterizes the principal’s optimal contract in a mean-field game environment where incentive design depends on the endogenous aggregate effort path $\bar{a}(t)$: the interaction parameter $k \neq 0$ in the profit function $\pi(a_i, \bar{a}) = a_i + k\bar{a}$ activates the mean-field loop and distinguishes the framework from decoupled bilateral Sannikov (2008) contracts. Second, the equilibrium is governed by a coupled forward-backward PDE system: a backward Hamilton-Jacobi-Bellman equation for $F(W, t)$ and a forward Fokker-Planck equation for $\mu(W, t)$, linked through the optimal policy functions $a^*(W, t)$ and $c^*(W, t)$ derived via the Revelation Principle. Third, the mean-field consistency condition, which requires that $\bar{a}(t) = \int a^*(W, t)\mu(W, t)dW$ hold in equilibrium, is established analytically and enforced numerically through a damped fixed-point iteration.

Numerical analysis reveals two principal equilibrium properties. First, despite the state-dependent inverted-U profile of individual optimal effort, which reflects the income effect at high continuation values and the proximity-to-retirement effect at low values, the aggregate mean effort $\bar{a}(t)$ remains approximately constant over the contract horizon. The mechanism is distributional: as the density $\mu(W, t)$ shifts leftward over time due to front-loaded incentive payments, the redistribution of agent mass precisely offsets individual effort variation, preserving stable aggregate productivity. Second, optimal contracting endogenously generates cross-sectional dispersion: the risk exposure required to satisfy incentive-compatibility links each agent’s continuation value to idiosyncratic output shocks, widening the distribution of W over time and producing ex-post inequality among ex-ante identical agents.

The framework applies to settings where a central authority manages a large population of agents subject to moral hazard and aggregate spillovers, including executive compensation in large firms where peer effort norms affect productivity, platform incentive design where contractor effort collectively determines service quality, and government transfer programs where aggregate participation rates affect program value. The mean-field formulation is appropriate in these contexts because the interaction is centralized (mediated through the principal rather than bilateral peer matching) and the population is large enough that individual deviations have negligible impact on the aggregate, satisfying the propagation-of-chaos condition underlying the mean-field approximation.

Several avenues for future research emerge from this work. The linear specification of the production function $\pi(a_i, \bar{a})$ could be generalized to capture more complex externalities and non-monotonic relationships between individual and aggregate effort. The model could be extended to incorporate agent entry and exit, allowing for endogenous participation decisions that respond to the contract structure and aggregate conditions. Multi-dimensional heterogeneity among agents, such as differences in both productivity and risk aversion, would enrich the analysis of how contracts should be tailored to different agent types while maintaining aggregate stability. Additionally, competition among multiple principals, each managing their own agent populations with potential spillovers across groups, would provide insights into strategic contract design in competitive markets.

The coupling of dynamic contracting theory with mean-field game methods provides a tractable framework for incentive design in large economies where individual moral hazard and aggregate externalities interact simultaneously, a combination that renders standard bilateral contracting incomplete. The equilibrium characterized here demonstrates that aggregate stability can emerge from individually heterogeneous optimal contracts through a distributional compensation mechanism, and that ex-post inequality is an endogenous consequence of incentive provision rather than a market failure. These properties have direct implications for the design of compensation systems in settings where the law of large numbers transforms idiosyncratic contractual risk into predictable aggregate dynamics.

References

- Achdou, Yves, and Italo Capuzzo-Dolcetta. Mean field games: numerical methods. *SIAM Journal on Numerical Analysis*, 48(3):1136–1162, 2010.
- Achdou, Yves. Finite difference methods for mean field games. In *Hamilton-Jacobi Equations: Approximations, Numerical Analysis and Applications*, pages 1–47. Springer, Berlin, 2013.
- Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. Income and wealth distribution in macroeconomics: A continuous-time approach. *The Review of Economic Studies*, 89(1):45–86, 2022.
- Biais, Bruno, Thomas Mariotti, Guillaume Plantin, and Jean-Charles Rochet. Dynamic security design: Convergence to continuous time and asset pricing implications. *The Review of Economic Studies*, 74(2):345–390, 2007.
- Cardaliaguet, Pierre. Notes on Mean Field Games (from P.-L. Lions’ lectures at Collège de France). Technical report, Université Paris-Dauphine, 2013.
- Cvitanić, Jakša, Dylan Possamaï, and Nizar Touzi. Moral hazard in dynamic risk management. *Management Science*, 63(10):3328–3346, 2017.
- DeMarzo, Peter M. and Yuliy Sannikov. A continuous-time agency model of optimal contracting and capital structure. NBER Working Paper No. 10516, 2004.
- DeMarzo, Peter M. and Yuliy Sannikov. Optimal security design and dynamic capital structure in a continuous-time agency model. *The Journal of Finance*, 61(6):2681–2724, 2006.
- Holmstrom, Bengt and Paul Milgrom. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica*, pages 303–328, 1987.
- Huang, Minyi, Peter E. Caines, and Roland P. Malhamé. 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, 2007.

- Huang, Minyi and Son Luu Nguyen. Mean field games for stochastic growth with relative utility. *Applied Mathematics & Optimization*, 74(3):643–668, 2016.
- Lasry, Jean-Michel and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007.
- Leung, Raymond C.W. Continuous-time principal-agent problem with drift and stochastic volatility control: with applications to delegated portfolio management. Available at SSRN: <https://ssrn.com/abstract=2482009>, 2014.
- Lin, Yiqing, Zhenjie Ren, Nizar Touzi, and Junjian Yang. Random horizon principal-agent problems. *SIAM Journal on Control and Optimization*, 60(1):355–384, 2022.
- Lucas, Robert E. and Benjamin Moll. Knowledge growth and the allocation of time. *Journal of Political Economy*, 122(1):1–51, 2014.
- Nakajima, Tomoyuki. Principal-agent problems with hidden savings in continuous time. CIRJE F-Series CIRJE-F-1182, CIRJE, Faculty of Economics, University of Tokyo, 2023.
- Ross, Stephen A. The economic theory of agency: The principal’s problem. *The American Economic Review*, 63(2):134–139, 1973.
- Sannikov, Yuliy. A continuous-time version of the principal-agent problem. *The Review of Economic Studies*, 75(3):957–984, 2008.
- Schättler, Heinz and Jaeyoung Sung. The first-order approach to the continuous-time principal-agent problem with exponential utility. *Journal of Economic Theory*, 61(2):331–371, 1993.
- Sung, Jaeyoung. Linearity with project selection and controllable diffusion rate in continuous-time principal-agent problems. *The RAND Journal of Economics*, pages 720–743, 1995.
- Williams, Noah. Persistent private information. *Econometrica*, 79(4):1233–1275, 2011.