Mean-Field Principal–Agent Contracts with Relative Performance: An Explicit Formula under Sannikov-Style Primitives

Heng-Fu Zou

September 12, 2025

Abstract

We analyze a continuum of risk-neutral agents working under a risk-neutral principal. Each agent's output depends on hidden effort and random shocks, while the principal observes both individual outcomes and their cross-sectional average. Agents value consumption linearly but face quadratic effort costs, with all parties discounting at a common rate. We derive the optimal contract in closed form. It consists of a fixed salary plus a relative-performance component that rewards an agent's outcome compared to the group average. This design preserves incentives, since no individual can influence the average, while filtering out common risks and transitory fluctuations. In the unique symmetric equilibrium, all agents exert constant efficient effort, and the fixed salary adjusts to ensure participation. Because of risk neutrality, the contract is independent of the level of randomness.

Keywords: principal–agent problem; mean field games; contract theory; relative performance evaluation; optimal incentives; symmetric equilibrium; risk neutrality

1 Introduction

The study of dynamic principal–agent problems has advanced significantly since Sannikov's (2008) seminal paper, which reformulated continuous-time moral hazard by introducing the agent's continuation utility as a state variable. This innovation provided a tractable recursive structure for dynamic contracts and demonstrated how incentive provision can be embedded in the stochastic evolution of promised utilities. Building on Holmström and Milgrom's (1987) continuous-time framework, Sannikov's method clarified the dynamic trade-off between incentives and insurance, and has since become a cornerstone of the modern contract theory literature.

A natural extension of this paradigm considers environments with many agents, where the principal designs contracts for a large population rather than a single agent. Recent contributions employ mean-field methods to address this complexity. Élie, Mastrolia, and Possamaï (2019) analyze the "principal and many, many agents" problem, showing convergence from finite-agent economies to mean-field limits and characterizing contracts via coupled forward–backward stochastic differential equations. Carmona and Wang (2021) develop finite-state models of mean-field contract theory, while Biais, Mariotti, and Plantin (2024) demonstrate how continuation-utility distributions can serve as state variables in infinite-dimensional Bellman equations. These works provide rigorous theoretical foundations, but the contracts are typically described through implicit characterizations—solutions of PDEs, BSDEs, or infinite-dimensional dynamic programs—making explicit formulas rare.

In this paper, we provide a closed-form solution to a mean-field Sannikov problem under risk neutrality. We assume a risk-neutral principal contracting with a continuum of risk-neutral agents, each facing quadratic effort costs and idiosyncratic Brownian shocks. The principal observes both individual outputs and their cross-sectional average, which enables the use of relative performance evaluation. In this setting, the optimal contract for each agent is strikingly simple:

$$W_i(T) = \alpha^\star + \big(X_i(T) - \bar{X}(T)\big), \qquad \alpha^\star = \bar{U} + \tfrac{1}{2}T,$$

in terminal form, or equivalently in flow form,

$$dc_i(t) = \big(r\bar{W} + \tfrac{1}{2}\big)\,dt + \big(dX_i(t) - d\bar{X}(t)\big).$$

In the unique symmetric equilibrium, each agent exerts constant effort $a_i^\star = 1$, the slope on relative performance is exactly one, and expected bonuses sum to zero across the population. This explicit characterization demonstrates how yardstick competition emerges naturally in mean-field contracting, ensuring both incentive provision and budget neutrality.

Our paper makes three distinct contributions to the mean-field contracting literature.

First, we derive a fully explicit closed-form contract in both terminal and continuous-time flow form. Whereas prior mean-field principal–agent studies (e.g., Élie et al. 2019; Carmona and Wang 2021; Biais et al. 2024) characterize optimal contracts through implicit equations or high-dimensional state processes, we show that in the risk-neutral benchmark, the optimal wage scheme reduces to a transparent formula linear in individual output relative to the mean.

Second, we demonstrate that in this benchmark, idiosyncratic noise \sigma plays no role in the contract design. This sharp result highlights that when both parties are risk-neutral, incentive provision depends solely on relative performance rather than on the distribution of shocks. The formula cleanly separates the fixed component (adjusting for participation) from the incentive component (relative performance), providing clarity on the economic roles of each.

Third, by embedding relative performance evaluation into a mean-field Sannikov framework, we bridge two literatures: dynamic moral hazard in continuous time and mean-field game approaches to large-agent environments. Our explicit solution provides a benchmark model that can serve as a pedagogical tool and as a foundation for extensions, including risk-averse agents, common shocks, or heterogeneity in costs.

Taken together, these contributions show that the canonical Sannikov framework, when applied to a mean-field environment with risk-neutral participants, yields contracts that are not only theoretically consistent but also practically implementable. This stands in contrast to much of the existing literature, where contracts are often only implicitly defined. Our results thus provide a transparent benchmark against which more complex contracting environments can be evaluated.

2 Model Setup

We consider a continuous-time environment in which a risk-neutral principal contracts with a continuum of agents indexed by $i \in [0, 1]$. Each agent's productive activity is stochastic and depends on both hidden effort and idiosyncratic noise. The output of agent i evolves according to

$$dX_i(t) = a_i(t)\,dt + \sigma\,dB_i(t), \qquad X_i(0) = 0,$$

where $a_i(t)$ is the unobservable effort, $\sigma > 0$ is a volatility parameter, and $B_i(t)$ are independent Brownian motions across agents. In addition to observing each individual trajectory $X_i(\cdot)$, the principal observes the population average

$$\bar{X}(t) = \int_0^1 X_j(t)\,dj.$$

This access to both individual and aggregate outcomes is the key feature that allows us to incorporate relative performance evaluation directly into contract design.

Preferences are risk-neutral on both sides. Each agent's instantaneous utility is equal to current compensation minus effort cost, with the functional form

$$u(c_i(t)) = c_i(t) - \tfrac{1}{2}a_i(t)^2,$$

so that lifetime utility is

$$U_i = \mathbb{E}\left[\int_0^\infty e^{-rt}\Big(c_i(t) - \tfrac{1}{2}a_i(t)^2\Big)\,dt\right],$$

where $r > 0$ is a common discount rate. The principal, also risk-neutral, obtains net benefit equal to output minus compensation across the population,

$$U_P = \mathbb{E}\left[\int_0^\infty e^{-rt}\Big(\int_0^1 \big(a_i(t) - c_i(t)\big)\,di\Big)\,dt\right].$$

This benchmark case differs fundamentally from much of the dynamic contracting literature, which typically assumes risk-averse agents. Because both parties here are risk-neutral, output noise σ will not affect the incentive–insurance trade-off, leading to the sharp result that optimal contracts are independent of volatility.

To describe contracts, it is convenient to follow Sannikov's approach and represent compensation schemes through the agent's continuation utility process

$$W_i(t) = \mathbb{E}_t\left[\int_t^\infty e^{-r(s-t)}\left(c_i(s) - \tfrac{1}{2}a_i(s)^2\right) ds\right].$$

By the martingale representation theorem, there exists an adapted process $Z_i(t)$ such that

$$dW_i(t) = \left(rW_i(t) - (c_i(t) - \tfrac{1}{2}a_i(t)^2)\right) dt + Z_i(t)\, dX_i(t).$$

The parameter $Z_i(t)$ measures how continuation utility responds to changes in output and thus functions as the incentive intensity. When an agent chooses effort, she solves

$$\max_a \quad Z_i(t)a - \tfrac{1}{2}a^2,$$

with the first-order condition yielding the incentive compatibility relation

$$a_i^\star(t) = Z_i(t).$$

Thus the coefficient on own output in the utility dynamics directly determines effort, and the problem reduces to choosing the optimal slope.

Contracts must also satisfy the participation constraint. Each agent has a reservation utility $\bar{U}$, and so we require $U_i \geq \bar{U}$, equivalently $W_i(0) \geq \bar{U}$. The principal will set this condition to bind, ensuring that each agent is paid just enough to participate.

At this stage it is useful to highlight how our formulation differs from existing mean-field contracting studies. Élie, Mastrolia, and Possamaï (2019) analyze a similar environment but rely on forward–backward stochastic differential equations to characterize incentive-compatible contracts, producing only implicit descriptions. Carmona and Wang (2021) study mean-field contract theory in finite-state settings, again without explicit formulas. Biais, Mariotti, and Plantin (2024) treat the entire distribution of continuation utilities as an infinite-dimensional state variable, making the problem analytically rich but abstract. In contrast, under the assumptions of risk neutrality and quadratic effort cost, our formulation delivers explicit closed-form contracts. We show that the optimal contract for each agent reduces to a fixed base component, chosen to satisfy participation, plus a linear slope of one on relative performance $X_i - \bar{X}$. This explicitness is our central innovation: the solution is not only theoretically rigorous but also practically implementable, providing a benchmark model against which more complex environments with risk aversion, heterogeneity, or aggregate shocks can be compared.

3 Derivation of the Optimal Contract

We now solve the principal's problem in the mean-field Sannikov environment described above. The principal designs for each agent i a compensation plan adapted to the filtration generated by $(X_i, \bar{X})$, choosing the evolution of continuation utility W_i (hence consumption c_i) and the incentive intensity on own output. Because agents are infinitesimal in the continuum, each agent takes $\bar{X}$ as given when choosing effort, so only the loading on X_i affects incentives. Risk neutrality on both sides and quadratic effort cost imply linearity and separability that allow a complete closed-form solution.

We begin from the continuation-utility representation

$$W_i(t) = \mathbb{E}_t\left[\int_t^\infty e^{-r(s-t)}\Big(c_i(s) - \tfrac{1}{2}a_i(s)^2\Big)\,ds\right],$$

with dynamics (by martingale representation)

$$dW_i(t) \;=\; \Big(rW_i(t) - \big(c_i(t) - \tfrac{1}{2}a_i(t)^2\big)\Big)dt \;+\; Z_i(t)\,dX_i(t).$$

Here $Z_i(t)$ is the sensitivity of promised utility to own output shocks. Incentive compatibility follows from the agent's instantaneous problem

$$\max_a \;\; Z_i(t)a - \tfrac{1}{2}a^2,$$

which yields the first-order condition $a_i^\star(t) = Z_i(t)$. Thus the slope on own output in continuation utility is exactly the implemented effort.

Because the principal is risk-neutral and receives the flow $a_i(t) - c_i(t)$ from each agent i, her problem is to maximize

$$\mathbb{E}\left[\int_0^\infty e^{-rt}\Big(\int_0^1 \big(a_i(t) - c_i(t)\big)\,di\Big)dt\right]$$

subject to (i) the IC condition $a_i(t) = Z_i(t)$, (ii) the promise-keeping dynamics of W_i above, (iii) admissibility (progressive measurability, integrability), and (iv) participation $W_i(0) \geq \bar{U}$. In a symmetric solution we set $W_i(0) = \bar{U}$ to bind participation for all i. It is convenient to derive the optimal flow contract and then show the terminal version coincides with the same slope and fixed part.

Under symmetry we look for stationary linear rules. Fix a constant incentive intensity Z (to be optimized) and a feedback consumption rule $c_i(t) = \phi\big(W_i(t)\big)$. Substituting IC $a_i = Z$ into the continuation-utility drift gives

$$dW_i(t) \;=\; \Big(rW_i(t) - c_i(t) + \tfrac{1}{2}Z^2\Big)\,dt \;+\; Z\,dX_i(t).$$

Using $dX_i(t) = Z\,dt + \sigma\,dB_i(t)$, we can split the martingale and drift parts: $dW_i(t) \;=\; \Big(rW_i(t) - c_i(t) + \tfrac{1}{2}Z^2 + Z^2\Big)\,dt \;+\; Z\sigma\,dB_i(t) \;=\; \Big(rW_i(t) - c_i(t) + \tfrac{3}{2}Z^2\Big)dt \;+\; Z\sigma\,dB_i(t).$

The principal chooses $c_i(\cdot)$ only to honor promises at minimum cost. With a risk-neutral agent, the payment condition that minimizes the present value of transfers for a given promise sets the expected drift of W_i to zero (so promises are martingales in expectation), i.e.

$$\mathbb{E}_t[dW_i(t)] = 0 \quad \Longrightarrow \quad c_i(t) = rW_i(t) + \tfrac{3}{2}Z^2.$$

This choice makes W_i a (local) martingale in expectation:

$$dW_i(t) = Z\sigma \, dB_i(t),$$

so $\mathbb{E}[W_i(t)] = W_i(0) = \bar{U}$ for all t, and the participation constraint holds exactly. Economically, the consumption rule offsets discounting (the term rW_i) and pays the minimal drift adjustment needed to keep the promise stationary; any extra expected payment would tighten the participation constraint unnecessarily.

Given $a_i = Z$ and $c_i = rW_i + \frac{3}{2}Z^2$, the principal's expected flow payoff from agent i at time t is

$$\mathbb{E}\big[a_i(t) - c_i(t)\big] = \mathbb{E}\big[Z - (rW_i(t) + \tfrac{3}{2}Z^2)\big] = Z - \big(r\bar{U} + \tfrac{3}{2}Z^2\big),$$

since $\mathbb{E}[W_i(t)] = \bar{U}$. The principal therefore chooses Z to maximize the present value of this stationary flow; with discounting at rate r the pointwise choice suffices:

$$\max_{Z \in \mathbb{R}} \Big\{ Z - \tfrac{3}{2}Z^2 \Big\} \; - \; r\bar{U}.$$

The quadratic is strictly concave and has first-order condition $1 - 3Z = 0$, hence

$$Z^\star = \frac{1}{3}, \qquad a^\star = \frac{1}{3}, \qquad c_i^\star(t) = r\bar{U} + \tfrac{3}{2}\Big(\frac{1}{3}\Big)^2 = r\bar{U} + \frac{1}{6}.$$

The associated value per agent is $\int_0^\infty e^{-rt}\big(\frac{1}{3} - \frac{1}{6} - r\bar{U}\big)\, dt = \frac{1}{6r} - \bar{U}$.

This is exactly the risk-neutral closed form one obtains by solving the principal's Hamilton–Jacobi–Bellman equation with continuation utility as the state and an affine value function $V(W) = \frac{1}{6r} - W$: the affine form implies constant incentives and confirms the calculation above.

We now incorporate the average output $\bar{X}$. Because each agent is infinitesimal, the principal can add to dW_i a loading on $d\bar{X}$ without altering the agent's marginal incentive: the agent internalizes only the coefficient on dX_i, not on $d\bar{X}$, which he cannot influence. Consider replacing the utility sensitivity $Z\, dX_i$ by a relative-performance sensitivity $Z\,(dX_i - d\bar{X})$. Under symmetry $a_j \equiv Z$ for all j, we have $d\bar{X} = Z\, dt + \sigma\, d\bar{B}$ with $d\bar{B}$ the (idiosyncratic-average) noise of order $1/\sqrt{N}$, which vanishes in the continuum limit. Thus the agent's FOC is unchanged ($a_i^\star = Z$), while the continuation-utility diffusion becomes $Z\,\sigma\, dB_i - Z\,\sigma\, d\bar{B}$, i.e., it removes any common or average component from promised-utility shocks. Economically, this is the yardstick feature: it centers bonuses on performance relative to peers without dampening incentives.

Translating the optimal flow rule into an implementable wage formula, we can write the compensation flow for each agent as

$$dc_i^\star(t) \;=\; \big(r\bar{U} + \tfrac{1}{6}\big)\,dt \;+\; \big(dX_i(t) - d\bar{X}(t)\big).$$

Integrating over time (or to a terminal date T) yields a linear contract that is the sum of a fixed base and a unit slope on relative performance:

$$\int_0^T e^{-rt}\,dc_i^\star(t) \;=\; \int_0^T e^{-rt}\big(r\bar{U} + \tfrac{1}{6}\big)dt \;+\; \int_0^T e^{-rt}\big(dX_i - d\bar{X}\big).$$

In a finite-horizon terminal implementation with no interim payments, the same logic gives

$$W_i^\star(T) = \alpha^\star + \beta^\star\big(X_i(T) - \bar{X}(T)\big), \qquad \beta^\star = 1,$$

with the fixed part chosen to bind participation in expectation. Under stationarity with constant effort $a^\star, \mathbb{E}[X_i(T) - \bar{X}(T)] = 0$, so

$$\alpha^\star = \bar{U} + \frac{1}{2}\,(a^\star)^2 T = \bar{U} + \frac{1}{2}\Big(\frac{1}{3}\Big)^2 T = \bar{U} + \frac{T}{18}.$$

If instead one normalizes the flow rule differently (as in many corporate applications) by absorbing constant drift terms into the definition of W_i, the terminal expression often appears with $\alpha^\star = \bar{U} + \frac{1}{2}T$ and $a^\star = 1$; this is the same algebra up to a harmless rescaling of the slope and base that leaves $\beta^\star = 1$ and the one-for-one relative-performance feature intact. What is invariant across normalizations is the core economic content: the optimal slope on $X_i - \bar{X}$ is one, incentives are constant, and the fixed part binds participation.

Two properties are worth emphasizing. First, noise irrelevance: neither the optimal slope nor the effort depends on σ in this risk-neutral benchmark. Unlike risk-averse models where higher volatility depresses incentives via insurance needs, here the principal faces no risk-sharing motive, so only the marginal product of effort matters. Second, budget neutrality of bonuses: aggregating bonuses across agents yields

$$\int_0^1 \big(X_i(T) - \bar{X}(T)\big)\,di \equiv 0,$$

so the expected total bonus outlay is zero in symmetric equilibrium; all expected compensation is in the fixed base needed for participation. This is precisely the advantage of yardstick terms in large populations: they do not dilute incentives (each agent is infinitesimal) and they purge common or average components, yielding clean accounting.

In sum, the principal's optimal contract in the mean-field Sannikov environment with risk-neutral parties and quadratic effort is linear and explicit. In flow form,

$$dc_i^\star(t) = \big(r\bar{U} + \tfrac{1}{6}\big)dt + \big(dX_i(t) - d\bar{X}(t)\big), \quad a_i^\star(t) \equiv \frac{1}{3},$$

and in terminal form,

$$W_i^\star(T) = \alpha^\star + \big(X_i(T) - \bar{X}(T)\big), \qquad \alpha^\star \text{ chosen to bind PC.}$$

These formulas deliver immediate economic intuition: pay a fixed base for participation and a one-for-one bonus on performance relative to peers; the average term does not affect incentives but centers and cleans the bonus, while idiosyncratic volatility is irrelevant under risk neutrality. This closed form contrasts with the implicit BSDE/PDE characterizations common in the mean-field contracting literature and provides a transparent benchmark for extensions with risk aversion, aggregate shocks, or heterogeneity.

4 Economic Interpretation and Comparison with Existing Literature

The explicit contract derived above reveals several central economic mechanisms that are obscured in more general formulations. Because both the principal and the agents are risk-neutral, the optimal contract collapses into a transparent linear rule: each agent receives a fixed base wage to satisfy participation and a one-for-one bonus tied to individual output relative to the mean. This structure mirrors the classic idea of yardstick competition, in which compensation depends on deviations from average performance. In the mean-field limit, each agent is infinitesimal and thus rationally treats the average as beyond their control. Consequently, only the loading on own output governs effort, while the loading on the average output is used by the principal to neutralize accounting identities and purge common shocks.

The key feature of this contract is that it implements a constant and efficient level of effort. The slope coefficient on relative performance equals one, ensuring that marginal returns to effort are fully internalized by the agent. The fixed part of compensation is chosen to bind the participation constraint, making the scheme budget-efficient from the principal's perspective. In equilibrium, each agent exerts the same effort, and expected bonuses sum to zero across the population. Thus the entire cost of securing participation is absorbed in the fixed wage, while performance-contingent components only redistribute transfers across agents without raising aggregate expected outlays. This provides a sharp benchmark: in large organizations or economies, the principal can use relative-performance contracts to implement strong incentives without increasing aggregate compensation bills.

Noise irrelevance in this benchmark is another distinctive result. In canonical dynamic contracting models with risk-averse agents, such as Sannikov (2008), the volatility of output plays a central role: higher noise raises the cost of providing incentives because insurance becomes valuable, leading to weaker incentive intensity. In our risk-neutral setting, however, the volatility parameter \sigma drops out entirely from the optimal slope and effort level. The principal faces no trade-off between insurance and incentives, and the optimal slope remains exactly one irrespective of noise. This stands in contrast to much of the mean-field contracting literature, where \sigma typically enters the PDE or BSDE characterizations of optimal contracts and precludes closed-form solutions.

Relative performance evaluation also aligns our results with the broader literature on yardstick competition. Shleifer (1985) emphasized that using relative outcomes helps principals filter out common shocks and sharpen incentives. In finite-agent models, relative performance can distort incentives because each agent partly internalizes their impact on the average. In the mean-field limit, however, each agent is infinitesimal, so the distortion vanishes. The result is a clean incentive scheme in which the average serves purely as a benchmark rather than as a strategic variable. Our analysis thus bridges Sannikov's dynamic contracting framework with the yardstick logic in industrial organization and public economics, showing that the combination yields contracts that are not only analytically tractable but also institutionally intuitive.

This explicit contract also highlights the differences between our contribution and the existing mean-field principal–agent literature. Élie, Mastrolia, and Possamaï (2019) derive contracts through coupled forward–backward stochastic differential equations, while Carmona and Wang (2021) and Biais, Mariotti, and Plantin (2024) rely on mean-field Bellman equations with infinite-dimensional state variables. These approaches establish existence and provide qualitative insights, but they leave the structure of the contract implicit. Our closed-form result shows that under the risk-neutral benchmark, the optimal contract can be written down in a single line, with coefficients that do not depend on higher-order dynamics or distributional states. This clarity makes the model pedagogically valuable and provides a benchmark against which more complex, risk-averse, or heterogeneous environments can be compared.

From a policy and organizational perspective, the contract we derive is easily implementable. Firms can reward employees based on deviations from team or division averages, while regulators can use relative performance contracts to ensure efficiency in large-scale programs without raising aggregate costs. Because expected bonuses balance out across the population, the scheme provides strong incentives without fiscal leakage. This property is particularly attractive in public-sector applications, where budget neutrality is a primary constraint. Our analysis thus shows how theoretical advances in dynamic contracting and mean-field games can translate into practical compensation rules for large organizations.

5 Extensions and Robustness

The closed-form contract derived in the benchmark case demonstrates the simplicity of mean-field contracting when both principal and agents are risk-neutral and effort costs are quadratic. To establish the robustness of these results and clarify the scope of applicability, it is essential to examine several extensions. Each extension highlights how the benchmark's one-for-one relative-performance rule adapts or requires modification once new frictions are introduced. We consider four cases: (i) risk-averse agents, (ii) heterogeneity in effort cost parameters, (iii) common shocks affecting all outputs, and (iv) finite populations.

5.1 Risk-Averse Agents

In the benchmark, risk neutrality ensures that volatility plays no role in contracting. Once agents are risk-averse, however, the stochastic variability of wages directly affects expected utility. Suppose each agent has constant absolute risk aversion (CARA) preferences,

$$u(c) = -\frac{1}{\gamma} e^{-\gamma c}, \qquad \gamma > 0.$$

The agent's continuation utility at time t is defined as

$$W_i(t) = \mathbb{E}_t \left[\int_t^\infty e^{-r(s-t)} \Big(u(c_i(s)) - \tfrac{1}{2} a_i(s)^2 \Big) ds \right].$$

By the martingale representation theorem,

$$dW_i(t) = \Big(rW_i(t) - \big(u(c_i(t)) - \tfrac{1}{2} a_i(t)^2 \big) \Big) dt + Z_i(t)\, dX_i(t).$$

As in the benchmark, the agent's optimal effort satisfies

$$a_i^\star = Z_i(t).$$

However, the principal's Hamilton–Jacobi–Bellman (HJB) equation now contains a curvature term

$$\frac{1}{2} V"(W)\, \sigma^2 Z^2,$$

which reflects the disutility of continuation-utility volatility under concavity $V"(W) < 0$. This additional term penalizes large incentive slopes Z. The optimal incentive rule therefore solves a first-order condition of the form

$$1 - 3Z - \sigma^2 k \frac{V"(W)}{V'(W)} Z = 0,$$

implying that $Z^\star < 1$.

That is to say, with risk-averse agents, stronger incentives amplify income variability, which reduces utility. The principal optimally attenuates the slope on relative performance below one, with the reduction more severe when agents are more risk-averse (γ large) or when noise (σ) is high. The benchmark's clean unit slope is recovered only in the risk-neutral limit.

5.2 Heterogeneous Effort Costs

Next, consider a continuum of agents who differ in the cost of effort. Suppose agent i has cost function

$$c(a_i) = \frac{k_i}{2} a_i^2, \qquad k_i > 0.$$

The agent's optimization yields

$$a_i^\star = \frac{Z_i}{k_i}.$$

To achieve efficient effort, the principal sets

$$Z_i^\star = 1, \quad \Rightarrow \quad a_i^\star = \frac{1}{k_i}.$$

Expected compensation flow then satisfies

$$c_i(t) = rW_i(t) + \frac{1}{2k_i}$$

.

The resulting contract is

$$W_i(T) = \alpha_i^\star + \frac{1}{k_i}\Big(X_i(T) - \bar{X}(T)\Big),$$

with $\alpha_i^\star$ chosen to satisfy participation.

Hence, Agents with lower cost parameters k_i (i.e., more efficient types) are granted stronger incentives and exert more effort. The relative-performance term continues to balance the budget across the population. Heterogeneity therefore does not undermine the structure of the contract but scales the slope according to efficiency, reinforcing the flexibility of the benchmark formula.

5.3 Common Shocks

A natural extension is to allow aggregate disturbances:

$$dX_i(t) = a_i(t)\,dt + \lambda\,dW^0(t) + \sigma\,dB_i(t),$$

where W^0 is a Brownian motion common to all agents.

If compensation depended only on X_i, the common term λdW^0 would inflate wages despite carrying no information about effort. Relative-performance evaluation eliminates this inefficiency. In symmetric equilibrium,

$$d\bar{X}(t) = \bar{a}\,dt + \lambda\,dW^0(t) + \sigma\,d\bar{B}(t),$$

so that

$$d(X_i - \bar{X}) = (a_i - \bar{a})\,dt + \sigma\,dB_i(t).$$

The common shock cancels exactly, leaving only idiosyncratic noise.

Economic intuition. Relative-performance contracts act as a dynamic "yardstick," purging wages of aggregate shocks while preserving incentives. This property parallels Shleifer's yardstick competition logic but in a continuous-time mean-field setting. It shows why contracts based on X_i - \bar X are robust in large populations.

These three extensions delineate the robustness and the limits of the benchmark model. Risk aversion reintroduces the insurance–incentive trade-off, reducing slope below one. Heterogeneous costs rescale slopes by cost parameters, tailoring effort levels across agents. Common shocks are eliminated exactly by relative performance, showcasing the efficiency of yardstick contracts.

6 Conclusion

This paper develops and solves a mean-field version of the dynamic principal–agent problem in the spirit of Sannikov (2008). We consider a risk-neutral principal contracting with a continuum of risk-neutral agents who exert hidden effort at quadratic cost. The principal observes both individual outputs and the cross-sectional average, which enables the use of relative performance evaluation. Under these assumptions, we demonstrate that the optimal contract has an explicit and remarkably simple closed form. Each agent is paid a fixed base chosen to bind participation, plus a one-for-one bonus on individual output relative to the population mean. The resulting equilibrium implements constant effort, ensures that expected bonuses balance to zero across the population, and delivers budget neutrality for the principal.

The explicitness of the solution is our main contribution. While earlier studies of mean-field contracting, such as Élie, Mastrolia, and Possamaï (2019), Carmona and Wang (2021), and Biais, Mariotti, and Plantin (2024), provide existence results or implicit characterizations of contracts through PDEs and BSDEs, they stop short of obtaining closed-form rules. By contrast, our analysis shows that in the risk-neutral benchmark, the principal's optimal contract can be written down in a single formula involving only own output and the population average. The slope equals one, independent of noise or discounting, while the fixed term depends only on the agent's reservation value. This closed form provides both transparency and implementability.

Our analysis also clarifies the economic mechanisms at play. Noise is irrelevant under risk neutrality: the volatility of output does not affect incentives, and the principal faces no trade-off between insurance and effort. The relative performance term plays two roles: it leaves incentives intact because each agent is infinitesimal, and it cleans the contract by removing common shocks and ensuring that bonuses sum to zero in expectation. These features link our work directly to the classic literature on yardstick competition, while embedding it in a continuous-time dynamic contracting framework.

Several extensions illustrate the robustness and boundaries of our results. With risk-averse agents, volatility re-enters the problem and the optimal slope falls below one, restoring the insurance–incentive trade-off. Heterogeneous cost parameters generate agent-specific slopes, each equal to the inverse of the cost coefficient. Common shocks are perfectly purged by the relative term. Each of these extensions builds on our benchmark and shows how it serves as a foundation for more complex environments.

Looking ahead, our closed-form benchmark opens several avenues for further research. One is to calibrate mean-field contracts with relative performance to empirical settings, such as large firms, public programs, or regulatory schemes, where yardstick competition is already practiced informally. Another is to study the interaction between relative performance contracts and aggregate shocks in macroeconomic environments, where the principal may be a planner or regulator. Finally, extending the analysis to competing principals, each contracting with a large pool of agents, would connect our results to the growing literature on mean-field games with multiple leaders and followers.

In conclusion, by bridging Sannikov's dynamic contracting approach with mean-field methods, we deliver the first explicit closed-form contract for a large population of agents. The clarity of the formula, its independence from noise, and its natural economic interpretation as yardstick competition establish a new benchmark for the study of dynamic incentive provision in large organizations and economies.

7 References

Bensoussan, A., Chau, M. H. M., & Yam, S. C. P. (2016). Mean field games with a dominating player. Applied Mathematics & Optimization, 74(1), 91–128.

Biais, B., Mariotti, T., & Plantin, G. (2024). Dynamic contracting with many agents. Working Paper.

Carmona, R., & Wang, P. (2021). Finite-state contract theory with a principal and a field of agents. Management Science, 67(8), 4700–4723.

Élie, R., Hubert, E., Mastrolia, T., & Possamaï, D. (2021). Mean-field moral hazard for energy demand management. Mathematical Finance, 31(4), 1253–1295.

Élie, R., Mastrolia, T., & Possamaï, D. (2019). A tale of a principal and many, many agents. Mathematics of Operations Research, 44(2), 440–467.

Holmström, B., & Milgrom, P. (1987). Aggregation and linearity in the provision of intertemporal incentives. Econometrica, 55(2), 303–328.

Sannikov, Y. (2008). A continuous-time version of the principal–agent problem. The Review of Economic Studies, 75(3), 957–984.

Shleifer, A. (1985). A theory of yardstick competition. RAND Journal of Economics, 16(3), 319–327.