

Dynamic Mean Field Contracting: Relative Performance Evaluation, Tournaments, and Incentives under Competition

Heng-Fu Zou*

September 19, 2025

Abstract

This paper develops a dynamic theory of contracting with one principal and many agents, extending the promised-utility framework of Sannikov (2008) to environments with competition, correlated outputs, and relative performance evaluation (RPE). We show that the first-order approach remains valid when the principal filters performance through generalized least squares (GLS) residuals of peer outputs, which minimize variance while preserving incentive slopes. The principal's Hamilton–Jacobi–Bellman equation links optimal effort to three forces: marginal revenue, signal informativeness, and intertemporal insurance. In a linear–quadratic environment, equilibrium effort is uniquely determined by a contraction mapping, resolving the multiplicity of equilibria found in static models such as Mookherjee (1984). We further unify Lazear–Rosen (1981) tournaments and linear RPE contracts as diffusion-driven special cases—Gaussian, CIR, and lognormal—within a mean field game system. The framework rationalizes empirical puzzles in executive pay, sales contests, and financial benchmarking, and generates new testable predictions about incentives and organizational scale.

Keywords: Dynamic contracts; Principal–agent problem; Multi-agent incentives; Relative performance evaluation; Tournaments; Mean field games; Hamilton–Jacobi–Bellman equation; Fokker–Planck dynamics; Incentives and insurance; Executive compensation; Organizational design.

1 Introduction

The design of incentives in organizations with many agents has long been a central concern in economics. The classic agency formulation of Ross (1973) introduced the principal's problem of providing incentives under moral hazard. Holmström (1979) established the conditions under which incentive compatibility constraints can be reduced to first-order conditions, laying the analytical

*IAS, Wuhan University, Wuhan, China 430079. Email: Henryzoucema@gmail.com

foundation for tractable contracting models. Subsequent work extended the analysis to environments with multiple tasks (Holmström and Milgrom 1991), correlated outputs (Holmström 1982), and rank-order incentives (Lazear and Rosen 1981; Green and Stokey 1983; Rosen 1986; Nalebuff and Stiglitz 1983). A central theme in this literature is the trade-off between risk-sharing and incentives, with different contractual forms—linear relative performance evaluation (RPE) contracts versus tournament schemes—emphasizing different margins.

In static settings, Mookherjee (1984) formalized the design of optimal incentive schemes with many agents and demonstrated how relative performance evaluation can eliminate common shocks, thereby improving the informativeness of contracts without distorting marginal incentives. Yet the static framework is characterized by multiplicity of equilibria and cannot capture the intertemporal evolution of incentives and utilities. Dynamic models, pioneered by Sannikov (2008) and extended by Williams (2011), resolved this difficulty for the single-agent case, showing how promised utilities can serve as sufficient statistics in continuous time and how the principal’s problem can be represented as a Hamilton–Jacobi–Bellman (HJB) equation. These contributions established tractability in dynamic principal–agent models but left open the problem of multi-agent interactions, competition, and correlated outputs.

This paper develops a dynamic contracting framework with one principal and a continuum of agents, embedding competition and relative performance evaluation in a continuous-time promised-utility model. Each agent exerts hidden effort that affects output subject to both idiosyncratic and aggregate shocks. The principal cannot observe actions but can design contracts based on filtered signals of performance, which include linear benchmarking against peers. These filters generalize Holmström’s (1982) correlation principle and Mookherjee’s (1984) multi-agent results: in equilibrium, the contract strips away common shocks by benchmarking against others, leaving only the residual variation most closely tied to individual effort. Incentives then take a tractable form: the strength of the link between effort and compensation is stronger when effort is more productive, weaker when effort is more costly, and more powerful when residual performance signals are less noisy. By adjusting the exposure of continuation utilities to these signals, the principal fine-tunes how strongly each agent responds to competition.

We show that the Hamilton–Jacobi–Bellman representation of the principal’s problem naturally extends to the many-agent case and that the joint HJB–Fokker–Planck system characterizes the evolution of continuation utilities and the aggregate distribution across agents. In the large-population limit, equilibrium coincides with the fixed point of a mean field game in the sense of Lasry and Lions (2007, 2010): given the aggregate effort path, each individual’s best response matches the aggregate in equilibrium. This mean field game structure guarantees existence, uniqueness, and stability of equilibrium, while embedding both relative performance evaluation and tournament-like incentives in a unified dynamic model.

The analysis also unifies the tournament and relative performance literatures within a single dynamic framework. Lazear and Rosen’s (1981) tour-

nament logic, where effort is motivated by rank-order prize spreads, emerges endogenously when contract exposures scale with promised utilities, generating lognormal dispersion and fat-tailed outcomes. Mookherjee’s (1984) linear RPE model is recovered when exposures are independent of promised utility, yielding Gaussian benchmarking with symmetric dispersion. Intermediate cases, such as Cox–Ingersoll–Ross–type exposures (Cox, Ingersoll, and Ross 1985), produce skewed Gamma-type distributions. This taxonomy highlights how the curvature of the principal’s value function and the structure of exposures determine whether equilibrium resembles smooth benchmarking, skewed contests, or tournament-like winner-take-all outcomes.

Empirically, the framework provides a unified explanation for several puzzles in organizational economics and finance. The limited reliance on relative performance evaluation in CEO contracts (Gibbons and Murphy 1990; Aggarwal and Samwick 1999) is consistent with dynamic insurance motives rather than inefficient governance. The muted prize spreads observed in sales contests, relative to static tournament predictions, follow from the attenuation imposed by intertemporal smoothing. Benchmark-based pay in financial firms and asset managers corresponds exactly to GLS residualization of common shocks, in line with the theoretical filter. These results underscore the applied relevance of the model and suggest new testable hypotheses about how organizational scale, market structure, and variance reduction shape observed compensation practices.

In this way, the paper unifies several major strands of economic theory. It combines Ross’s (1973) formulation of agency, Holmström’s (1979, 1982) insights on observability and correlation, Lazear and Rosen’s (1981) tournament contracts, Mookherjee’s (1984) multi-agent optimal schemes, Sannikov’s (2008) dynamic promised-utility approach, Williams’s (2011) continuous-time extensions, and the mean field game methodology of Lasry and Lions (2007, 2010). The contributions are both theoretical and empirical: establishing FOA validity in multi-agent continuous time, characterizing the HJB–FPK system with competition, deriving closed-form equilibria in the linear–quadratic environment, unifying tournament and RPE outcomes within explicit diffusion families, and rationalizing empirical regularities in executive pay, sales contests, and financial benchmarking. To our knowledge, this is the first paper to unify tournaments, RPE, and mean field dynamic contracting within a promised-utility framework.

The remainder of the paper is organized as follows. Section 2 introduces the model and defines filtered signals. Section 3 derives the principal’s Hamilton–Jacobi–Bellman equation and optimal incentive intensity. Section 4 characterizes the optimal relative performance evaluation filter as the generalized least squares projection. Section 5 specializes to a linear–quadratic environment, deriving closed-form equilibria and proving existence and uniqueness in the mean-field limit. Section 6 analyzes competition and outcomes across Ornstein–Uhlenbeck, Cox–Ingersoll–Ross, and lognormal cases. Section 7 formally compares tournaments and linear RPE within the dynamic mean field framework. Section 8 discusses empirical and applied implications, and Section 9 concludes.

2 Model and Signals

We study a continuous-time economy with one principal and $N \geq 2$ agents. Uncertainty is generated by idiosyncratic Brownian motions $(W_i)_{i=1}^N$ and a common Brownian motion B . Agent i chooses an effort process $a_i(t) \geq 0$, adapted and square-integrable. Effort is hidden, but it affects the agent's observable output

$$dY_i(t) = (\mu a_i(t) + \phi \bar{a}_{-i}(t))dt + \sigma dW_i(t) + \sigma_0 dB(t), \quad (2.1)$$

where $\mu > 0$ is own productivity, $\phi \in \mathbb{R}$ captures technological interaction with peers (substitutes if $\phi < 0$, complements if $\phi > 0$), $\sigma > 0$ is idiosyncratic volatility, $\sigma_0 > 0$ is common volatility, and $\bar{a}_{-i}(t) = \frac{1}{N-1} \sum_{j \neq i} a_j(t)$. Collect outputs as $Y = (Y_1, \dots, Y_N)^\top$. The instantaneous covariance of increments is

$$\Sigma := \text{Var}(dY/\sqrt{dt}) = \sigma^2 I_N + \sigma_0^2 \mathbf{1}\mathbf{1}^\top. \quad (2.2)$$

Agents have CARA utility $u(c) = -\frac{1}{\gamma} e^{-\gamma c}$ with risk aversion $\gamma > 0$ and quadratic effort cost $\frac{1}{2} k a_i^2$ ($k > 0$). Given discount rate $r > 0$, the promised (continuation) utility state for agent i is

$$W_i(t) = \mathbb{E}_t \left[\int_t^\infty e^{-r(s-t)} (u(c_i(s)) - \frac{1}{2} k a_i(s)^2) ds \right]. \quad (2.3)$$

Promised utility is a sufficient state: once $W_i(t)$ is fixed, all future contractual obligations to agent i are summarized by W_i .

The principal observes outputs but not efforts, and can therefore condition transfers only on $\mathcal{F}_t^Y = \sigma(Y(s) : s \leq t)$. Following Holmström's informativeness logic and its multi-agent version, the principal improves signal quality by benchmarking an agent's output against peers. For agent i , consider the linear filter

$$dZ_i^\omega(t) = dY_i(t) - \sum_{j \neq i} \omega_{ij}(t) dY_j(t). \quad (2.4)$$

Write $dZ_i^\omega = e_i^\top dY$ with $e_i = (0, \dots, 1, \dots, 0)^\top - \sum_{j \neq i} \omega_{ij} e_j^{\text{std}}$. The drift and variance of the filter are

$$\mathbb{E} \left[\frac{dZ_i^\omega}{dt} \right] = \mu a_i + \phi \bar{a}_{-i} - \sum_{j \neq i} \omega_{ij} (\mu a_j + \phi \bar{a}_{-j}), \quad v_i(\omega_i) = e_i^\top \Sigma e_i. \quad (2.5)$$

The filter leaves the own-action sensitivity equal to μ (the term $\sum_{j \neq i} \omega_{ij}$ does not depend on a_i), so there is no arbitrary scale in the signal's sensitivity to a_i . The innovation (standardized residual) is

$$d\tilde{B}_i(t) := \frac{dZ_i^\omega(t) - \mathbb{E}[dZ_i^\omega(t) | \mathcal{F}_t^Y]}{\sqrt{v_i(\omega_i)}}, \quad (2.6)$$

which is a standard Brownian increment (possibly correlated across agents).

The contract for agent i consists of consumption $c_i(t)$ and an exposure $\theta_i(t)$ to the innovation. By the martingale representation theorem, any $\mathcal{F}_t^Y$ -adapted contract can be written in “transfer + exposure” form. Under any feasible contract $(c_i, \theta_i, \omega_i)$, promised utility evolves as

$$dW_i(t) = \left(rW_i(t) - u(c_i(t)) + \frac{1}{2}k a_i(t)^2 \right) dt + \theta_i(t) d\tilde{B}_i(t). \quad (2.7)$$

Because the filtered signal is linear-Gaussian and its variance does not depend on effort, the first-order approach (FOA) is valid. A local perturbation argument (Girsanov) yields the Gateaux derivative of the agent’s continuation payoff with respect to effort at time t equal to $\theta_i(t)\mu/\sqrt{v_i(\omega_i)} - k a_i(t)$. Hence the best response is linear:

$$a_i(t) = \lambda_i(\omega_i) \theta_i(t), \quad \lambda_i(\omega_i) := \frac{\mu}{k\sqrt{v_i(\omega_i)}}. \quad (2.8)$$

This implementability slope increases with productivity μ , decreases with cost k , and increases with informativeness (smaller residual variance $v_i(\omega_i)$).

To make the information design explicit, we now compute the variance-minimizing filter. Partition Σ conformably as

$$\Sigma = \begin{pmatrix} \Sigma_{ii} & \Sigma_{i,-i} \\ \Sigma_{-i,i} & \Sigma_{-i,-i} \end{pmatrix}, \quad e_i = \begin{pmatrix} 1 \\ -\omega_i \end{pmatrix}. \quad (2.9)$$

Then

$$v_i(\omega_i) = e_i^\top \Sigma e_i = \Sigma_{ii} - 2\omega_i^\top \Sigma_{-i,i} + \omega_i^\top \Sigma_{-i,-i} \omega_i. \quad (2.10)$$

The unique minimizer (GLS projection) and minimized variance (Schur complement) are

$$\omega_i^* = \Sigma_{-i,-i}^{-1} \Sigma_{-i,i}, \quad v_i^* = \Sigma_{ii} - \Sigma_{i,-i} \Sigma_{-i,-i}^{-1} \Sigma_{-i,i}. \quad (2.11)$$

Intuitively, the principal subtracts from Y_i its best linear predictor based on peers Y_{-i} ; the residual variance is exactly the mean-squared prediction error. This is the continuous-time GLS analogue of Holmström’s correlation principle and the multi-agent filtering in Mookherjee (1984).

In the symmetric covariance case $\Sigma = \sigma^2 I_N + \sigma_0^2 \mathbf{1}\mathbf{1}^\top$, symmetry implies equal peer weights $\omega_{ij} = w$ for $j \neq i$. Write $Z_i = Y_i - w \sum_{j \neq i} Y_j$. Using $\text{Var}(Y_i) = \sigma^2 + \sigma_0^2$, $\text{Var}(\sum_{j \neq i} Y_j) = (N-1)\sigma^2 + (N-1)^2\sigma_0^2$, and $\text{Cov}(Y_i, \sum_{j \neq i} Y_j) = (N-1)\sigma_0^2$, we obtain the quadratic

$$v_i(w) = \text{Var}(Z_i) = (\sigma^2 + \sigma_0^2) + w^2((N-1)\sigma^2 + (N-1)^2\sigma_0^2) - 2w(N-1)\sigma_0^2. \quad (2.12)$$

The first-order condition $dv_i/dw = 0$ yields the unique minimizer

$$w^* = \frac{\sigma_0^2}{\sigma^2 + (N-1)\sigma_0^2}, \quad (2.13)$$

and substituting back gives a closed-form residual variance

$$v_i^* = \sigma^2 + \sigma_0^2 - \frac{(N-1)\sigma_0^4}{\sigma^2 + (N-1)\sigma_0^2} = \sigma^2 \frac{\sigma^2 + N\sigma_0^2}{\sigma^2 + (N-1)\sigma_0^2}. \quad (2.14)$$

Two useful limits follow. First, as the peer set widens,

$$v_i^* \downarrow \sigma^2 \quad \text{as } N \rightarrow \infty, \quad (2.15)$$

so optimal benchmarking purges the common shock asymptotically and drives the implementability slope to its idiosyncratic frontier,

$$\lambda_i(\omega_i^*) \uparrow \frac{\mu}{k\sigma} \quad \text{as } N \rightarrow \infty. \quad (2.16)$$

Second, the cross-correlations of standardized innovations $d\tilde{B}_i$ induced by Z_i vanish with N : after substituting w^* , $\text{Cov}(Z_i, Z_j) = O(1/N)$ for $i \neq j$ under $\Sigma = \sigma^2 I + \sigma_0^2 \mathbf{1}\mathbf{1}^\top$, so the covariance matrix of $(d\tilde{B}_i)$ converges to the identity. This justifies treating a representative agent facing an idiosyncratic innovation in the mean-field limit.

Combining filtering and effort choice, we can rewrite the promised-utility dynamics under an optimal filter as

$$dW_i(t) = \left(rW_i(t) - u(c_i(t)) + \frac{1}{2}k a_i(t)^2 \right) dt + \theta_i(t) d\tilde{B}_i(t), \quad a_i(t) = \lambda^* \theta_i(t), \quad \lambda^* := \frac{\mu}{k\sqrt{v_i^*}}, \quad (2.17)$$

where v_i^* is given by (2.14). The economic content is transparent. The principal sets exposure θ_i to translate filtered performance innovations into movements of promised utility; the agent's best response scales linearly with θ_i at slope λ^* . Better benchmarking (larger N or larger σ_0^2 relative to σ^2) lowers v_i^* , steepens λ^* , and strengthens incentives for a given exposure; but in the dynamic problem the principal balances this against the insurance motive embedded in the curvature of her value function. These forces—marginal revenue from effort, informativeness via GLS, and intertemporal insurance—are assembled in Section 3's HJB and, in the mean-field limit where $v_i^* \rightarrow \sigma^2$ and innovation correlations vanish, deliver a tractable mean-field game of dynamic multi-agent contracting.

3 Principal's HJB and Incentive Intensity

The principal is risk-neutral and discounts at rate $r > 0$. Given a contract $(c_i, \theta_i, \omega_i)_{i=1}^N$ and the filtered signals from Section 2, the promised utility of agent i evolves as

$$dW_i(t) = \left(rW_i(t) - u(c_i(t)) + \frac{1}{2}k a_i(t)^2 \right) dt + \theta_i(t) d\tilde{B}_i(t), \quad (3.1)$$

where $d\tilde{B}_i$ is the innovation associated with the variance-minimizing filter ω_i^* . Incentive compatibility can be handled by the first-order approach. To see this,

consider a small, adapted perturbation $a_i^\varepsilon = a_i + \varepsilon h$ supported on $[t, t + \Delta]$. Under the Girsanov change of measure implied by (optimal) linear-Gaussian filtering, the likelihood ratio for the innovation changes by a drift $\Delta m = \varepsilon (\mu / \sqrt{v_i(\omega_i)}) \int_t^{t+\Delta} h_s ds$, so the Radon–Nikodym derivative equals $1 + \varepsilon (\mu / \sqrt{v_i}) \int_t^{t+\Delta} h_s d\tilde{B}_i(s) - \frac{1}{2} O(\varepsilon^2)$. Taking expectations of (3.1) under the perturbed measure and dividing by $\varepsilon \Delta$ yields the Gateaux derivative $\theta_i(t) \mu / \sqrt{v_i(\omega_i)} - k a_i(t)$. Hence the agent's unique best reply at time t satisfies

$$a_i(t) = \lambda_i(\omega_i) \theta_i(t), \quad \lambda_i(\omega_i) := \frac{\mu}{k \sqrt{v_i(\omega_i)}}, \quad (3.2)$$

that is, effort is linear in exposure with slope increasing in productivity, decreasing in cost, and decreasing in the residual noise of the filtered signal. Here we define the action vector

$$\mathbf{a} := (a_1, \dots, a_N), \quad a_i = \lambda_i(\omega_i) \theta_i.$$

Turning to the principal's problem, let $V(W, t)$ denote her value when the state is the vector $W = (W_1, \dots, W_N)$. Applying Itô's formula to V along the diffusion (3.1) for each agent and allowing for possibly correlated innovations $(d\tilde{B}_i)$ with instantaneous covariance matrix $C = (C_{ij})$ defined by $d\langle \tilde{B}_i, \tilde{B}_j \rangle_t = C_{i,j} dt$, we obtain

$$\begin{aligned} dV(W, t) &= V_t dt + \sum_{i=1}^N V_{W_i} dW_i + \frac{1}{2} \sum_{i,j=1}^N V_{W_i W_j} dW_i dW_j \\ &= \left[V_t + \sum_i V_{W_i} \left(rW_i - u(c_i) + \frac{1}{2} k a_i^2 \right) + \frac{1}{2} \sum_{i,j} V_{W_i W_j} \theta_i \theta_j C_{ij} \right] dt + \sum_i V_{W_i} \theta_i d\tilde{B}_i. \end{aligned} \quad (3.3)$$

Taking conditional expectations and noting that the last term is a martingale with zero mean, the dynamic programming principle implies the Hamilton–Jacobi–Bellman equation

$$rV(W, t) = \sup_{c, \theta, \omega} \left\{ \mathbb{E}[R(a)] - \sum_i c_i + \sum_i V_{W_i} \left(rW_i - u(c_i) + \frac{1}{2} k a_i^2 \right) + \frac{1}{2} \sum_{i,j} V_{W_i W_j} \theta_i \theta_j C_{ij} + V_t \right\}, \quad (3.4)$$

where $a_i = \lambda_i(\omega_i) \theta_i$ and $a = (a_1, \dots, a_N)$. The four groups of terms have standard meanings: expected revenue, the resource cost of transfers, promise-keeping in continuation utility units, and the Itô variance correction reflecting the risk borne by exposing promised utilities to innovations.

The first-order condition with respect to consumption is immediate. Differentiating the HJB objective in c_i gives $-1 + V_{W_i} u'(c_i) = 0$, so optimal transfers satisfy

$$V_{W_i}(W, t) u'(c_i^*(W, t)) = 1. \quad (3.5)$$

Thus the shadow value of promised utility scales marginal utility of consumption one-for-one.

The first-order condition with respect to exposure θ_k requires more care because θ_k affects both drift and diffusion and (potentially) enters cross-variance terms through C . Differentiating the objective in (3.4) with respect to θ_k and using $\partial a_i / \partial \theta_k = \lambda_i(\omega_i) \mathbf{1}\{i = k\}$ yields

$$\frac{\partial}{\partial \theta_k} [\mathbb{E}[R(a)]] = \text{MR}_k(a) \lambda_k(\omega_k), \quad \frac{\partial}{\partial \theta_k} \left[\sum_i V_{W_i} \frac{1}{2} k a_i^2 \right] = V_{W_k} k \lambda_k(\omega_k)^2 \theta_k, \quad (3.6)$$

and

$$\frac{\partial}{\partial \theta_k} \left[\frac{1}{2} \sum_{i,j} V_{W_i W_j} \theta_i \theta_j C_{ij} \right] = \sum_{j=1}^N V_{W_k W_j} C_{kj} \theta_j, \quad (3.7)$$

where $\text{MR}_k(a) := \partial \mathbb{E}[R(a)] / \partial a_k$. Setting the derivative of the HJB objective to zero gives the linear system

$$\text{MR}_k(a) \lambda_k(\omega_k) + V_{W_k} k \lambda_k(\omega_k)^2 \theta_k + \sum_{j=1}^N V_{W_k W_j} C_{kj} \theta_j = 0, \quad k = 1, \dots, N. \quad (3.8)$$

In compact matrix notation, let $\Lambda := \text{diag}(\lambda_1, \dots, \lambda_N)$, $D := \text{diag}(V_{W_1} k \lambda_1^2, \dots, V_{W_N} k \lambda_N^2)$, and $H := (V_{W_i W_j})$ the Hessian of V . Then the system reads

$$\Lambda \text{MR}(a) + (D + H \circ C) \theta = 0, \quad (3.9)$$

where $\circ$ denotes the Hadamard (entrywise) product. Provided the matrix $D + H \circ C$ is negative definite—which is guaranteed under the concavity of V and standard integrability conditions—the unique optimal exposure vector is

$$\theta^* = -(D + H \circ C)^{-1} \Lambda \text{MR}(a), \quad a^* = \Lambda \theta^*. \quad (3.10)$$

This is the general multi-agent condition allowing for correlated innovations. It makes transparent that exposure solves a linear system balancing the marginal revenue gain against two costs: the promise-keeping cost $V_W k \lambda^2 \theta$ that scales with the effort cost delivered in continuation utility units, and the diffusion cost $(H \circ C) \theta$ induced by the curvature of the principal's value and the covariance of innovations.

In many environments of interest—most notably after generalized least-squares filtering with many peers—the innovation covariance becomes approximately diagonal: cross-agent correlations vanish as the common shock is purged and the peer set grows. In that case $C \approx I_N$, $H \circ C$ reduces to the diagonal of the Hessian, and (3.9) decouples into scalar first-order conditions. Using the standard promised-utility envelope $V_{W_i} k = 1 - V_{W_i W_i}$ (obtained by differentiating the HJB with respect to W_i and using the dynamics (3.1); intuitively, a

marginal relaxation of the promise enters one-for-one in drift and is offset by curvature in diffusion), each component reduces to

$$\text{MR}_i(a) \lambda_i(\omega_i) = (1 - V_{W_i W_i}) \theta_i, \quad \theta_i^* = \frac{\lambda_i(\omega_i) \text{MR}_i(a)}{1 - V_{W_i W_i}}, \quad a_i^* = \frac{\lambda_i(\omega_i)^2 \text{MR}_i(a)}{1 - V_{W_i W_i}}. \quad (3.11)$$

These scalar expressions make the economics transparent. The left-hand side is the marginal benefit of stronger incentives: the marginal revenue from effort multiplied by the implementability slope λ_i , which is higher when signals are more informative after GLS filtering. The right-hand side is the marginal insurance cost: exposure raises the variance of promised utility and, through promise-keeping, the required delivery of effort costs; both scale with the curvature of the principal's value in promised utility. Effort therefore increases with productivity and informativeness and decreases with competition (via MR_i) and with the stringency of intertemporal insurance (via $V_{W_i W_i}$).

Aggregation closes the mean-field game. Let m_t denote the cross-sectional law of promised utilities generated by (3.1) under the optimal policy, and let $a^*(W, t)$ be the induced best response from (3.10) (or from (3.11) in the approximately diagonal case). The aggregate effort at time t is

$$A_t = \int a^*(W, t) m_t(dW), \quad (3.12)$$

and the marginal revenue term $\text{MR}_i(a)$ appearing in the first-order condition is evaluated at this same aggregate. In the linear-quadratic revenue environment used later (Cournot-type demand), $\text{MR}_i(A) = \bar{m} - \bar{\beta}A$, so the mapping $A \mapsto \int a^*(W, t; A) m_t(dW)$ is affine with negative slope. Under the curvature and informativeness restrictions delivered by the GLS filter and the concavity of V , this mapping is a contraction, guaranteeing existence, uniqueness, and global stability of the symmetric mean-field equilibrium. In short, the HJB and the Fokker-Planck dynamics jointly determine individual exposures and the evolution of promised utilities, while the mean-field fixed point pins down the aggregate that feeds back into marginal revenue.

4 Optimal Relative Performance Evaluation Filter

The central determinant of incentive strength in this dynamic contracting framework is the informativeness of the filtered performance signals. As established in the previous sections, the agent's induced effort is proportional to the implementability slope, which in turn is inversely proportional to the square root of the variance of the filtered signal. The principal therefore faces a canonical information design problem: how to construct filters that minimize the variance of the innovation associated with an individual agent's effort while preserving sensitivity to that agent's action.

Formally, recall that for agent i the filtered signal takes the linear form

$$dZ_i^\omega(t) = dY_i(t) - \sum_{j \neq i} \omega_{ij}(t) dY_j(t), \quad (4.1)$$

where ω_{ij} are contractible weights. Writing $dZ_i^\omega = e_i^\top dY$ with $e_i = (1, -\omega_i^\top)^\top$, the variance of the innovation is

$$v_i(\omega_i) = e_i^\top \Sigma e_i, \quad (4.2)$$

where $\Sigma = \text{Var}(dY/\sqrt{dt})$ is the covariance matrix of output increments. The drift sensitivity of the filter to a_i is normalized to μ , since subtracting peers does not alter the coefficient on own effort. This normalization ensures that the signal has correct scale: the mapping from effort to drift remains unchanged, so that minimization of variance does not artificially inflate or deflate incentives.

The principal's problem is therefore equivalent to

$$\min_{\omega_i} v_i(\omega_i) \quad \text{subject to normalization of own effort sensitivity.} \quad (4.3)$$

Partition Σ conformably as

$$\Sigma = \begin{pmatrix} \Sigma_{ii} & \Sigma_{i,-i} \\ \Sigma_{-i,i} & \Sigma_{-i,-i} \end{pmatrix}, \quad e_i = \begin{pmatrix} 1 \\ -\omega_i \end{pmatrix}. \quad (4.4)$$

Then

$$v_i(\omega_i) = \Sigma_{ii} - 2\omega_i^\top \Sigma_{-i,i} + \omega_i^\top \Sigma_{-i,-i} \omega_i. \quad (4.5)$$

This is a convex quadratic function in ω_i . The unique minimizer is obtained by solving the first-order condition

$$\frac{\partial v_i}{\partial \omega_i} = -2\Sigma_{-i,i} + 2\Sigma_{-i,-i}\omega_i = 0, \quad (4.6)$$

so that

$$\omega_i^* = \Sigma_{-i,-i}^{-1} \Sigma_{-i,i}. \quad (4.7)$$

Substituting back yields the minimized residual variance

$$v_i^* = \Sigma_{ii} - \Sigma_{i,-i} \Sigma_{-i,-i}^{-1} \Sigma_{-i,i}. \quad (4.8)$$

Equation (4.7) shows that the optimal filter is the generalized least squares (GLS) projection of Y_i on the outputs of all other agents, and equation (4.8) shows that the minimized variance is the Schur complement of $\Sigma_{-i,-i}$ in Σ . This regression interpretation is central: the optimal contract extracts precisely the unpredictable component of an agent's output after controlling linearly for peer outputs.

To illustrate, consider the symmetric covariance structure from Section 2, $\Sigma = \sigma^2 I_N + \sigma_0^2 \mathbf{1}\mathbf{1}^\top$. Symmetry implies that all peer weights are identical, $\omega_{ij} = w$ for $j \neq i$. Writing

$$Z_i = Y_i - w \sum_{j \neq i} Y_j, \quad (4.9)$$

the variance expands as

$$v_i(w) = (\sigma^2 + \sigma_0^2) + w^2((N-1)\sigma^2 + (N-1)^2\sigma_0^2) - 2w(N-1)\sigma_0^2. \quad (4.10)$$

The first-order condition $dv_i/dw = 0$ gives the unique minimizer

$$w^* = \frac{\sigma_0^2}{\sigma^2 + (N-1)\sigma_0^2}. \quad (4.11)$$

Substituting yields the residual variance

$$v_i^* = \sigma^2 + \sigma_0^2 - \frac{(N-1)\sigma_0^4}{\sigma^2 + (N-1)\sigma_0^2} = \sigma^2 \frac{\sigma^2 + N\sigma_0^2}{\sigma^2 + (N-1)\sigma_0^2}. \quad (4.12)$$

This expression is strictly decreasing in the number of agents N . As $N \rightarrow \infty$,

$$v_i^* \downarrow \sigma^2, \quad (4.13)$$

which is the idiosyncratic variance. Thus in large organizations the common shock is nearly eliminated from performance signals through benchmarking, leaving only irreducible individual noise.

With CARA–Normal preferences and quadratic costs, the implementability slope is

$$\lambda_i^* = \frac{\mu}{k\sqrt{v_i^*}}. \quad (4.14)$$

Equations (4.12) and (4.14) show explicitly how organizational scale affects incentives. A larger number of agents drives v_i^* down toward σ^2 , which raises λ_i^* and steepens the mapping from exposure to effort. Benchmarking is therefore intrinsically more powerful in larger organizations, because a larger pool of peers provides a better predictor of common shocks.

The regression characterization also clarifies the empirical practice of benchmarking executive pay against industry indices or peer groups. In effect, these benchmarks are nothing more than the residual from regressing an individual firm's outcome on industry performance. The GLS solution derived above formalizes why such practice is optimal: it strips out industry-wide shocks, leaving a residual that isolates the manager's individual contribution.

The economic interpretation is thus twofold. On one hand, relative performance evaluation reduces variance and strengthens incentives, a point emphasized by Holmström's (1982) correlation principle and empirically investigated

by Gibbons and Murphy (1990). On the other hand, the variance-reducing effect is bounded by the intertemporal insurance motive discussed in Section 3, since even highly informative signals will not be fully loaded if the principal values smooth continuation utilities. The model therefore reconciles the informativeness principle with the observed moderate reliance on benchmarking in practice.

In summary, the optimal relative performance evaluation filter is a GLS projection of each agent's output onto the outputs of others, producing a residual with minimized variance. The residual variance decreases with organizational size and converges to the idiosyncratic component as the number of agents grows, thereby explaining why large organizations can design highly informative performance measures. Yet in dynamic contracting, these gains are always tempered by the principal's insurance motive, a balance that is made explicit in the subsequent equilibrium analysis.

5 Existence and Uniqueness in a Linear–Quadratic Case

To study equilibrium behavior in a tractable setting, it is useful to specialize the contracting environment to a linear–quadratic (LQ) structure. This specification preserves the essential trade-offs of incentives, competition, information, and intertemporal insurance, while permitting closed-form solutions for induced effort and explicit equilibrium conditions. Moreover, the linear–quadratic form connects our dynamic contracting framework to the classical industrial organization literature on Cournot competition and to the recent mean field game approach to large-agent economies (Lasry and Lions 2007, 2010).

We begin with the principal's revenue. Let total output be $Q(t) = \sum_{i=1}^N Y_i(t)$. Suppose that the inverse demand function is linear, $P(Q) = \alpha - \beta Q$, where $\alpha > 0$ is the intercept and $\beta > 0$ the slope. Then revenue is

$$R(t) = P(Q)Q = (\alpha - \beta Q)Q. \quad (5.1)$$

Differentiating with respect to agent i 's action yields the marginal revenue of effort, after substituting the law of motion for outputs,

$$MR(A) = \bar{m} - \bar{\beta}A, \quad (5.2)$$

where $A(t) = \frac{1}{N} \sum a_j(t)$ denotes mean effort. The constants $\bar{m}$ and $\bar{\beta}$ are proportional to $\mu\alpha$ and $\mu\beta N$, respectively. Equation (5.2) captures the Cournot logic: as aggregate effort rises, the marginal revenue from any single agent's action declines, reflecting competition among agents.

Turning to incentives, recall from Section 3 that under the first-order approach the induced effort satisfies

$$a_i(t) = \frac{\lambda^2}{1 - V_{W_i W_i}(t)} MR(A), \quad (5.3)$$

where $\lambda = \mu/(k\sqrt{v^*})$ is the implementability slope determined by technology, cost, and the minimized residual variance v^* from Section 4. The denominator $1 - V_{W_i W_i}(t)$ reflects the curvature of the principal's value function in promised utility: greater concavity reduces exposure and attenuates incentives.

By symmetry, in equilibrium all agents exert the same effort, so $a_i(t) = A(t)$ for all i . Substituting into (5.3) and suppressing the time index yields the fixed-point equation

$$A = \frac{\lambda^2}{1 - \underline{V}_{WW}} (\bar{m} - \bar{\beta}A), \quad (5.4)$$

where $\underline{V}_{WW}$ denotes the common curvature across agents. Rearranging gives the explicit closed-form equilibrium effort

$$A^* = \frac{\bar{m}}{\bar{\beta} + \frac{1 - \underline{V}_{WW}}{\lambda^2}}. \quad (5.5)$$

Expression (5.5) makes transparent the determinants of effort. Higher baseline productivity $\bar{m}$ raises effort, stronger competition $\bar{\beta}$ reduces it, more informative signals (smaller residual variance v^*) raise λ and therefore increase effort, while greater concavity $\underline{V}_{WW}$ reduces the denominator $1 - \underline{V}_{WW}$, thereby lowering effort. The balance of these four forces—technology, competition, information, and insurance—completely characterizes incentives in equilibrium.

To establish existence and uniqueness, define the best-response mapping

$$\Phi(A) = \frac{\lambda^2}{1 - \underline{V}_{WW}} (\bar{m} - \bar{\beta}A). \quad (5.6)$$

The slope of this mapping is $-\kappa$, where

$$\kappa = \frac{\bar{\beta}\lambda^2}{1 - \underline{V}_{WW}}. \quad (5.7)$$

If $0 < \kappa < 1$, the mapping is a contraction on $\mathbb{R}_+$. By Banach's fixed-point theorem, there exists a unique fixed point A^* , and from any initial guess of aggregate effort the best-response iteration converges globally and geometrically to this equilibrium. Thus, unlike in static multi-agent contracting models such as Mookherjee (1984), where multiple equilibria can arise, the dynamic promised-utility formulation delivers a unique and stable symmetric equilibrium under mild conditions.

Further insight comes from examining the mean field limit as $N \rightarrow \infty$. In this limit, the variance of the optimal GLS-filtered signal simplifies because the common shock is diversified away by benchmarking against an arbitrarily large peer group. Specifically, as shown in Section 4,

$$v_i^* \downarrow \sigma^2, \quad (5.8)$$

the idiosyncratic variance, and the implementability slope converges to

$$\lambda = \frac{\mu}{k\sigma}. \quad (5.9)$$

The equilibrium effort in the mean field economy is therefore

$$A^* = \frac{\bar{m}}{\bar{\beta} + \frac{1 - V_{WW}}{\mu^2/(k^2\sigma^2)}}. \quad (5.10)$$

This expression shows that in very large organizations, incentive provision approaches the idiosyncratic frontier: all common noise is eliminated through benchmarking, leaving only individual shocks. Yet effort is still attenuated by the intertemporal insurance motive, reflected in V_{WW} .

The linear-quadratic specialization thus achieves three objectives. First, it provides explicit closed-form solutions that clarify the role of productivity, competition, information, and insurance in determining equilibrium effort. Second, it establishes existence, uniqueness, and global stability of equilibrium under the contraction property, resolving multiplicity problems of earlier static models. Third, it connects the dynamic contracting framework to the mean field game literature: in the large- N limit, the equilibrium condition (5.10) is precisely a mean field fixed point, ensuring consistency between individual best responses and aggregate behavior.

6 Competition and Outcomes in the Mean Field Game Formulation with Explicit Solvable Cases

Having established the general structure of dynamic contracting and the linear-quadratic benchmark, we now extend the analysis to the mean field game formulation, in which the number of agents is large and each individual is infinitesimal relative to the aggregate. In this setting, competition is naturally introduced through relative performance evaluation: each agent's output is benchmarked against the average of all others, and in symmetric equilibrium the aggregate effort is simply the mean of all optimal responses. The mean field framework allows us to characterize the law of motion of the cross-sectional distribution of promised utilities, and in certain cases to obtain explicit closed-form solutions.

Formally, let W denote an individual agent's promised utility, and let m_t be the distribution of W across the unit interval of agents at time t . The principal's problem is described by a backward-forward system: the Hamilton-Jacobi-Bellman (HJB) equation for the value function $V(W, t)$, and the forward Fokker-Planck (FPK) equation for the distribution m_t . Specifically, the HJB takes the form

$$rV(W, t) = \sup_{c, \theta, \omega} \left\{ \mathbb{E}[R(A_t, m_t)] - c + V_W \left(rW - u(c) + \frac{1}{2}ka^2 \right) + \frac{1}{2}V_{WW}\theta^2 \right\}, \quad (6.1)$$

subject to incentive compatibility $a = \lambda(\omega)\theta$, where the slope $\lambda(\omega) = \mu/(k\sqrt{v(\omega)})$ depends on the informativeness of the chosen filter. The associated continuation utility evolves as

$$dW_t = \left(rW_t - u(c^*(W, t)) + \frac{1}{2}k(a^*(W, t))^2 \right) dt + \theta^*(W, t) d\tilde{B}_t, \quad (6.2)$$

and the law of motion of the distribution is given by

$$\frac{\partial m_t(W)}{\partial t} = -\frac{\partial}{\partial W} \left((rW - u(c^*(W, t)) + \frac{1}{2}k(a^*(W, t))^2) m_t(W) \right) + \frac{1}{2} \frac{\partial^2}{\partial W^2} \left((\theta^*(W, t))^2 m_t(W) \right). \quad (6.3)$$

A mean field equilibrium is a pair (V, m) such that the HJB yields optimal controls, the FPK evolves consistently with these controls, and the fixed-point condition holds:

$$A_t = \int a^*(W, t) m_t(dW). \quad (6.4)$$

Competition is encoded in the marginal revenue schedule. In the linear demand framework of Section 5, each agent faces

$$\text{MR}_i(A) = \bar{m} - \bar{\beta}A, \quad (6.5)$$

so that effort depends negatively on the aggregate A . In equilibrium, aggregate effort is precisely the mean of optimal responses, so competition manifests as a fixed point linking individual best responses to the collective outcome.

Although the general HJB–FPK system rarely admits explicit solutions, three diffusion families for the continuation utility W are analytically tractable: the Ornstein–Uhlenbeck (OU) case with Gaussian distribution, the Cox–Ingersoll–Ross (CIR) case with Gamma distribution, and the geometric Brownian motion (GBM) case with lognormal distribution. Each corresponds to a different curvature of the principal’s value function and hence to a different exposure schedule, producing distinct equilibrium implications.

In the OU case, the curvature of the value function is constant, so exposures and induced efforts are independent of promised utility. Substituting into the promise-keeping condition yields linear dynamics of the form

$$dW_t = (\alpha - \kappa W_t) dt + \sigma d\tilde{B}_t, \quad (6.6)$$

an Ornstein–Uhlenbeck process. The stationary distribution is normal with mean α/κ and variance $\sigma^2/(2\kappa)$. The mean-field fixed point is

$$A_{OU}^* = \frac{\lambda^2 \bar{m}}{\bar{\beta} \lambda^2 + 1 - \alpha_2}, \quad (6.7)$$

where α_2 is the constant curvature of V . This case represents Gaussian competition: agents exert the same effort, and continuation utilities are symmetrically dispersed.

In the CIR case, curvature is chosen so that exposure scales with the square root of promised utility, $\theta^*(W) = \zeta\sqrt{W}$. Effort is then proportional to $\sqrt{W}$, and continuation utility follows a CIR process:

$$dW_t = \kappa(\Theta - W_t)dt + \sigma\sqrt{W_t}d\tilde{B}t. \quad (6.8)$$

The stationary distribution is Gamma with mean Θ and variance $\Theta\sigma^2/(2\kappa)$. The mean-field fixed point is

$$A_{CIR}^* = \frac{\lambda c_0 \mathbb{E}[\sqrt{W}] \bar{m}}{1 + \lambda c_0 \mathbb{E}[\sqrt{W}] \bar{\beta}}, \quad (6.9)$$

with $\mathbb{E}[\sqrt{W}]$ evaluated under the Gamma law. This case yields skewed outcomes: continuation utilities are dispersed asymmetrically, resembling tournament environments with winners and losers.

In the GBM case, curvature is such that exposure scales linearly with promised utility, $\theta^*(W) = \eta W$. Effort is proportional to W , and continuation utility follows geometric Brownian motion:

$$dW_t = \tilde{\alpha}W_t dt + \tilde{\sigma}W_t d\tilde{B}t. \quad (6.10)$$

The cross-sectional distribution is lognormal, with mean $\mathbb{E}[W_t] = W_0 e^{\tilde{\alpha}t}$. The mean-field fixed point is

$$A_{GBM}^* = \frac{\lambda d_0 \mathbb{E}[W] \bar{m}}{1 + \lambda d_0 \mathbb{E}[W] \bar{\beta}}. \quad (6.11)$$

This case produces fat-tailed outcomes: a small fraction of agents attain disproportionately high continuation utilities, mimicking the winner-takes-all character of Lazear–Rosen tournaments.

Despite their differences, these three families share a common structure. In each case equilibrium effort satisfies

$$A^* = \frac{\Xi \mathcal{M} \bar{m}}{1 + \Xi \mathcal{M} \bar{\beta}}, \quad (6.12)$$

where Ξ is a contract coefficient depending on technology, information, and curvature, while $\mathcal{M}$ is a moment of the stationary distribution of promised utilities: $\mathcal{M} = 1$ in the OU case, $\mathcal{M} = \mathbb{E}[\sqrt{W}]$ in the CIR case, and $\mathcal{M} = \mathbb{E}[W]$ in the GBM case. The slope of the best-response map is $-\Xi \mathcal{M} \bar{\beta}$; if this lies in $(-1, 0)$, the map is a contraction and the equilibrium is unique and globally stable.

The economic implications are clear. Gaussian competition yields mild, symmetric dispersion in utilities and relatively smooth incentives. CIR competition introduces skewness, producing outcomes reminiscent of moderate tournaments with visible differentiation. GBM competition generates extreme inequality, with lognormal tails and strong winner-takes-all features. By choosing the curvature of the value function—and hence the exposure schedule—the principal effectively selects the competitive regime that governs organizational outcomes.

7 Formal Comparison of Tournaments, Relative Performance Evaluation, and Unification via Mean Field Games

The multi-agent incentive literature has developed along two largely distinct paths. The first, initiated by Lazear and Rosen (1981), views tournaments as efficient incentive schemes: compensation depends on rank rather than absolute performance, and common shocks that affect all agents cancel out in relative comparisons. The second, pioneered by Mookherjee (1984) and following Holmström’s correlation principle (1982), develops linear relative performance evaluation (RPE) contracts: compensation depends linearly on own output and on peer outputs, with coefficients chosen to minimize residual variance. Both approaches emphasize the role of correlation and competition, but they differ in the functional form of incentives, the distribution of residual risk, and the equilibrium outcomes. In this section we formally compare the two frameworks and show how our mean field contracting model unifies them as special cases of a dynamic HJB–FPK system.

The Lazear–Rosen tournament model begins with two risk-neutral agents who exert costly effort $a_i \geq 0$ to produce outputs

$$y_i = a_i + \varepsilon_i, \quad (7.1)$$

where the shocks ε_i are independent and identically distributed with density f and cumulative distribution F . The contract specifies two prizes: a high wage W_H for the winner (the agent with higher output) and a low wage W_L for the loser. Writing the spread as $\Delta = W_H - W_L$, the expected utility of agent i is

$$U_i(a_i, a_j) = W_L + \Delta \cdot \Pr(y_i > y_j) - c(a_i), \quad (7.2)$$

where $c(a_i)$ is the convex cost of effort. In symmetric equilibrium, the first-order condition is

$$c'(a^*) = \Delta \cdot f_D(0), \quad (7.3)$$

where $D = \varepsilon_j - \varepsilon_i$ is the difference in shocks and f_D its density, with $f_D(0) = \int f(x)^2 dx$. Equation (7.3) captures the essence of tournament incentives: marginal cost of effort equals the marginal benefit of winning, which is the increase in winning probability evaluated at the zero gap, scaled by the prize spread. The condition highlights two properties. First, the power of incentives is directly proportional to the prize spread. Second, common shocks vanish in ranks, so agents are fully insured against aggregate noise.

By contrast, the linear RPE model assumes agents are risk-averse with CARA utility and normally distributed correlated outputs. Suppose there are N agents with outputs

$$y = a + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \Sigma), \quad (7.4)$$

where $a = (a_1, \dots, a_N)^\top$ is the vector of efforts and Σ the covariance matrix. Contracts are linear:

$$w_i = \alpha_i + \beta_i y_i + \sum_{j \neq i} \delta_{ij} y_j. \quad (7.5)$$

Expected utility is

$$U_i = \alpha_i + \beta_i a_i + \sum_{j \neq i} \delta_{ij} a_j - c(a_i) - \frac{\gamma}{2} \text{Var}(w_i). \quad (7.6)$$

Here marginal incentives are governed entirely by β_i , since $\partial U_i / \partial a_i = \beta_i - c'(a_i)$. The role of the cross-weights δ_{ij} is not to provide incentives but to reduce variance. The variance-minimizing choice of coefficients is the generalized least squares (GLS) residualization:

$$\delta_i^* = -\Sigma_{-i,-i}^{-1} \Sigma_{-i,i}, \quad (7.7)$$

which makes the filtered signal

$$Z_i = y_i - \sum_{j \neq i} \omega_{ij} y_j, \quad (7.8)$$

have minimized variance equal to the Schur complement,

$$v_i^* = \Sigma_{ii} - \Sigma_{i,-i} \Sigma_{-i,-i}^{-1} \Sigma_{-i,i}. \quad (7.9)$$

The induced effort is therefore

$$a_i^* = \frac{\beta_i}{k} = \lambda_i(\omega_i^*) \theta_i, \quad \lambda_i(\omega_i^*) = \frac{\mu}{k \sqrt{v_i^*}}, \quad (7.10)$$

with θ_i the contract's exposure to the filtered innovation. Thus, unlike in tournaments, incentives are continuous and directly linked to the variance of residual signals, while common shocks are purged through regression against peer outputs.

The technical distinction between tournaments and linear RPE is crucial. In tournaments, the distribution of shocks matters only through the density at zero of their differences, $f_D(0)$, which pins down the slope of winning probabilities. In linear RPE, the distribution matters only through the covariance matrix Σ , which determines the variance of filtered signals. In both cases, common shocks cancel out, but the residual risk is allocated differently: all-or-nothing in tournaments, smoothly in linear contracts.

Our mean field contracting model unifies these two approaches by embedding them in a continuous-time HJB–FPK system. The key lies in the curvature of the principal's value function, which determines the exposure schedule. If curvature is constant, exposures are independent of promised utility, continuation utilities follow an Ornstein–Uhlenbeck process, and equilibrium resembles linear Gaussian RPE with symmetric dispersion. If curvature is state-dependent such

that exposures scale with $\sqrt{W}$, continuation utilities follow a CIR process with Gamma distribution, producing skewed dispersion. If curvature is chosen such that exposures scale linearly with W , continuation utilities follow a lognormal law, generating fat-tailed distributions characteristic of tournaments.

Formally, all three cases collapse to the unified fixed-point form

$$A^* = \frac{\Xi \mathcal{M} \bar{m}}{1 + \Xi \mathcal{M} \bar{\beta}}, \quad (7.11)$$

where Ξ is a contract coefficient reflecting technology and information, and $\mathcal{M}$ is a moment of the distribution of promised utilities: $\mathcal{M} = 1$ in the Gaussian case, $\mathcal{M} = \mathbb{E}[\sqrt{W}]$ in the CIR case, and $\mathcal{M} = \mathbb{E}[W]$ in the lognormal case. Uniqueness and global stability obtain whenever $\Xi \mathcal{M} \bar{\beta} \in (0, 1)$, so that the best-response map is a contraction.

The economic interpretation is that both tournaments and linear RPE are specializations of a more general dynamic contracting framework. The Lazear–Rosen tournament emerges as the dynamic analogue of GBM exposure schedules: effort affects the drift of log-utilities, winning probabilities depend on the density at zero of log differences, and equilibrium outcomes are highly unequal. The Mookherjee linear RPE contract emerges as the OU specialization: efforts are proportional to constant exposures, continuation utilities are symmetrically distributed, and incentives are smoothed by variance-minimizing regression. The CIR case provides an intermediate structure, capturing situations where incentives are state-dependent but less extreme than in full tournaments.

Thus, the mean field game formulation reveals a unifying principle: the curvature of the principal’s value function selects the diffusion family governing continuation utilities, which in turn determines whether competition among agents manifests as smooth RPE-type incentives, skewed intermediate regimes, or extreme tournament-like outcomes. By embedding these possibilities in a continuous-time model, the analysis connects and extends the literatures on tournaments, RPE, and dynamic contracting, and provides a single framework in which their empirical predictions can be compared and tested.

8 Empirical and Applied Implications

The theoretical framework developed here provides a unifying dynamic account of incentives under competition, relative performance evaluation, and tournaments, and therefore has direct empirical implications for organizational economics, executive compensation, and corporate finance. By embedding the interaction of variance-minimizing filters, exposure schedules, and mean field fixed points into a continuous-time contracting environment, the analysis clarifies several puzzles in the empirical literature and generates new testable hypotheses.

A first domain concerns executive compensation. Holmström’s (1982) informativeness principle suggests that contracts should employ any signal that reduces the variance of performance measures, implying widespread use of relative performance evaluation (RPE). Yet empirical studies such as Gibbons

and Murphy (1990) and Aggarwal and Samwick (1999) find only modest use of benchmarking in chief executive officer pay. Our model explains this apparent discrepancy by introducing dynamic insurance. Even when peer-relative signals dramatically reduce variance, the principal optimally attenuates exposures in order to smooth continuation utilities for risk-averse agents. In other words, the weak reliance on RPE observed empirically need not indicate inefficient contracting; it may instead reflect the intertemporal insurance constraint highlighted by the curvature of the principal’s value function in our model.

A second application concerns tournaments and sales contests. Lazear and Rosen (1981) showed that tournaments provide high-powered incentives because rank order filters out common shocks, and empirical work confirms that sales organizations frequently use contests to motivate effort. Yet prize spreads in practice are often smaller than the static tournament model predicts (Rosen 1986; Kräkel and Schöttner 2010). Our dynamic framework reconciles this puzzle: while common shocks can indeed be purged in large organizations, thereby preserving the informativeness of rank comparisons, the principal must also respect the promise-keeping and insurance constraints implied by continuation utilities. The result is that equilibrium prize spreads are endogenously muted relative to the static first-best, producing exactly the attenuated tournament incentives observed in practice.

A third domain arises in financial firms, asset management, and investment banking. Benchmark-relative compensation schemes—such as rewarding fund managers for outperforming the S&P 500 or peer group averages—are pervasive. In our model, these contracts correspond directly to generalized least squares (GLS) filters: the optimal weights in the linear benchmarking rule are precisely those that purge common market shocks while preserving residual variation tied to skill. This prediction aligns with practices documented in asset management, where larger peer groups and industry indices are routinely used to evaluate performance. Our mean field formulation further predicts that the power of such benchmarking increases with organizational scale: as the number of comparable peers rises, common shocks are diversified away, and the residual variance converges to the idiosyncratic component, thereby steepening the implementability slope of incentives.

The role of organizational scale thus emerges as a central empirical implication. In large organizations, relative performance evaluation becomes increasingly effective as variance reductions from benchmarking accumulate. This generates the prediction that larger firms or sectors with many comparable units should make heavier use of benchmarking in compensation contracts. At the same time, our model shows that individual incentives do not scale without bound: the intertemporal insurance motive ensures that exposures remain finite, even with perfect filtering. Hence reliance on benchmarking should increase with organizational size, but the effective strength of incentives relative to variance will eventually plateau.

Finally, the model sheds light on comparative institutional outcomes. The Ornstein-Uhlenbeck specialization implies smooth Gaussian dispersion of utilities and corresponds to modest, continuous incentives, consistent with industries

where benchmarking is formalized but not extreme. The CIR case produces skewed outcomes with moderate inequality, resembling settings such as internal career ladders where performance differentials accumulate gradually. The GBM case generates fat-tailed, lognormal distributions of continuation utilities, corresponding to winner-takes-all environments such as tournaments for top executive positions, financial “star” traders, or superstar markets more generally. Observing the shape of cross-sectional compensation distributions therefore provides indirect evidence of which curvature regime the principal has implemented, offering a bridge between contract theory and empirically measurable inequality patterns.

Taken together, these implications demonstrate that a dynamic mean field contracting framework is both theoretically illuminating and empirically grounded. It rationalizes long-standing puzzles such as the limited use of RPE in CEO contracts and the muted prize spreads in tournaments, while also predicting observable links between organizational scale, variance reduction, and incentive intensity. By unifying the insights of Holmström (1979, 1982), Lazear and Rosen (1981), Mookherjee (1984), and Sannikov (2008) within a single continuous-time model, the analysis offers a rigorous foundation for understanding how competition and benchmarking are deployed in real-world organizations.

9 Conclusion

This paper has developed a unified dynamic framework for contracting with one principal and many agents, embedding competition, relative performance evaluation, and tournament incentives within a continuous-time promised-utility formulation. By allowing the principal to construct filtered performance signals from agents’ outputs, we showed that contract design reduces to the choice of consumption transfers, exposures to innovations, and benchmarking weights across peers. The first-order approach remains valid in this environment, with induced effort proportional to contractual exposure and scaled by the informativeness of the filtered signal. This tractability provides a rigorous basis for analyzing multi-agent moral hazard in continuous time, while avoiding the multiplicity and indeterminacy that have long challenged static models.

The Hamilton–Jacobi–Bellman representation of the principal’s problem revealed that dynamic incentives depend on three interacting forces: the marginal revenue of effort, the statistical precision of relative performance signals, and the curvature of the principal’s value function that captures intertemporal insurance needs. Marginal revenue reflects the competitive or cooperative structure of the environment, the filtering of signals generalizes Holmström’s (1982) correlation principle and Mookherjee’s (1984) GLS characterization, and the curvature term introduces a dynamic attenuation mechanism absent in static models. Together, these elements unify the Lazear–Rosen (1981) tournament framework with linear RPE contracting in a single dynamic system.

A central technical contribution was to show that the optimal filter corresponds to the generalized least squares residual of an agent’s output on those

of their peers, with minimized variance given by the Schur complement of the covariance matrix. This provides a transparent regression-based interpretation of relative performance evaluation and demonstrates that in large organizations, where common shocks can be diversified away, performance measures approach the idiosyncratic frontier. A further contribution was to establish, in a linear-quadratic specialization, that equilibrium effort takes a closed form, with existence, uniqueness, and stability guaranteed by a contraction mapping. In the mean field limit, equilibrium is globally stable, overcoming the multiplicity of equilibria in static models such as Mookherjee (1984).

The framework also connected theoretical predictions to empirical patterns. The limited reliance on relative performance evaluation in executive compensation, the moderate size of tournament prize spreads relative to static predictions, and the pervasive use of benchmark-relative contracts in finance all emerge as natural consequences of dynamic insurance considerations. Organizational scale was shown to play a dual role: larger peer groups improve informativeness and increase the power of RPE, yet dynamic smoothing prevents incentive intensity from rising without bound. The diffusion-based characterization of continuation utilities further explained why Gaussian, Gamma, or lognormal distributions of compensation arise in different organizational environments, ranging from smooth benchmarking to winner-takes-all tournaments.

By integrating tournaments, linear RPE, and dynamic promised-utility models into a mean field contracting framework, this paper extends the continuous-time agency literature initiated by Sannikov (2008) and Williams (2011) beyond its single-agent focus. It demonstrates that the structure of competition among agents can be captured by the curvature of the principal’s value function and the induced diffusion of promised utilities. In doing so, the analysis resolves long-standing issues of tractability and equilibrium selection in multi-agent contracting and provides a coherent explanation of diverse empirical practices in executive pay, sales contests, and financial benchmarking.

The contributions are therefore both theoretical and applied. On the theoretical side, the paper advances dynamic contract theory by unifying previously distinct literatures and proving existence, uniqueness, and stability of equilibrium in a many-agent setting. On the applied side, it provides a rigorous basis for interpreting compensation practices and organizational design choices observed in real-world data. By situating tournaments, RPE, and dynamic insurance in a common mean field game framework, the analysis offers a comprehensive and empirically relevant theory of incentives under competition.

10 References

- Aggarwal, R. K., & Samwick, A. A. (1999). Executive compensation, strategic competition, and relative performance evaluation: Theory and evidence. *Journal of Finance*, 54(6), 1999–2043.
- Cox, J. C., Ingersoll, J. E., & Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2), 385–407.

- Gibbons, R., & Murphy, K. J. (1990). Relative performance evaluation for chief executive officers. *Industrial and Labor Relations Review*, 43(3), 30–51.
- Green, J., & Stokey, N. L. (1983). A comparison of tournaments and contracts. *Journal of Political Economy*, 91(3), 349–364.
- Holmström, B. (1979). Moral hazard and observability. *Bell Journal of Economics*, 10(1), 74–91.
- Holmström, B. (1982). Moral hazard in teams. *Bell Journal of Economics*, 13(2), 324–340.
- Holmström, B., & Milgrom, P. (1991). Multitask principal–agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, and Organization*, 7, 24–52.
- Karatzas, I., & Shreve, S. E. (1991). *Brownian Motion and Stochastic Calculus*. Springer.
- Kräkel, M., & Schöttner, A. (2010). Prizes versus wages: An agency perspective. *Journal of Labor Economics*, 28(4), 867–879.
- Lasry, J.-M., & Lions, P.-L. (2007). Mean field games. *Japanese Journal of Mathematics*, 2(1), 229–260.
- Lasry, J.-M., & Lions, P.-L. (2010). Mean field games with a common noise. *Annals of Probability*, 38(4), 1444–1491.
- Lazear, E. P., & Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89(5), 841–864.
- Mookherjee, D. (1984). Optimal incentive schemes with many agents. *Review of Economic Studies*, 51(3), 433–446.
- Nalebuff, B. J., & Scharfstein, D. S. (1987). Testing for asymmetric information using quantity variations. *RAND Journal of Economics*, 18(2), 291–308.
- Nalebuff, B. J., & Stiglitz, J. E. (1983). Prizes and incentives: Toward a general theory of compensation and competition. *Bell Journal of Economics*, 14(1), 21–43.
- Rosen, S. (1986). Prizes and incentives in elimination tournaments. *American Economic Review*, 76(4), 701–715.
- Ross, S. A. (1973). The economic theory of agency: The principal’s problem. *American Economic Review*, 63(2), 134–139.
- Sannikov, Y. (2008). A continuous-time version of the principal–agent problem. *Review of Economic Studies*, 75(3), 957–984.
- Williams, N. (2011). On dynamic principal–agent problems in continuous time. *Journal of Economic Theory*, 146(2), 720–737.