

Dynamic Multi-Agent Contracting with Relative Performance Evaluation and Competition

Heng-Fu Zou*

October 22, 2025

Abstract

This paper develops a continuous-time model of contracting between a principal and many agents under moral hazard, linking dynamic contract theory to the mean field game (MFG) framework. Outputs depend on effort, competition, and both idiosyncratic and common shocks. The principal designs contracts using filtered signals that benchmark agents against peers. We show that the first-order approach remains valid: effort is linear in exposure with slope determined by the informativeness of filtered outputs. The Hamilton–Jacobi–Bellman equation implies that optimal incentives depend jointly on marginal revenue, signal variance, and intertemporal insurance. The optimal filter is the generalized least squares residual, which asymptotically eliminates common shocks as group size grows. In a linear–quadratic specialization, equilibrium exists, is unique, and is globally stable, with the uniqueness result coinciding with the Lasry–Lions MFG fixed-point condition. The framework unifies tournaments, RPE, and dynamic contracts, and explains empirical puzzles in executive pay, sales teams, and finance.

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

Incentive design in environments where a single principal interacts with multiple agents has long stood at the center of organizational and contract theory. The canonical contributions of Ross (1973) and Holmström (1979) provided the first rigorous treatments of the principal–agent problem under moral hazard, showing how contracts must balance the provision of incentives with the sharing of risk. Yet those models are built around a single agent. Real organizations, by contrast, are populated by many agents whose actions are interdependent,

*IAS, Wuhan University, Wuhan, China 430079

whose outputs are correlated, and whose performance is often evaluated relative to that of peers. These features are not peripheral but central: they shape the way firms design pay-for-performance schemes, governments allocate incentives across bureaucracies, and financial institutions structure compensation for managers and traders.

The problem of competition and correlation among agents introduces two layers of difficulty absent in the single-agent model. First, when outputs are subject to both idiosyncratic and common shocks, the principal must determine how to filter signals to isolate individual effort. Second, when the value of an agent's action depends on the actions of others, marginal revenues become endogenous to the equilibrium profile of efforts. In such settings, the principal's challenge is to design contracts that induce effort through competition or cooperation while simultaneously providing insurance and filtering noise. The theoretical difficulty is that standard tools often fail: static models can exhibit multiple Nash equilibria in agents' actions, and dynamic models with many agents quickly lose tractability.

Three strands of literature have tackled aspects of this problem but in isolation. The first is the theory of tournaments, pioneered by Lazear and Rosen (1981), who showed that rank-order contracts can provide strong incentives even when individual outputs are noisy, because relative rank eliminates common shocks. This literature was further developed by Rosen (1986), Green and Stokey (1983), and Nalebuff and Stiglitz (1983), and has been applied extensively to sales contests, sports, and executive pay. The second is the literature on relative performance evaluation (RPE), where Holmström (1982) formulated the correlation principle: any signal that improves the informativeness of performance evaluation should enter the contract. Mookherjee (1984) extended this logic to environments with many agents, analyzing the trade-offs between independent contracts and tournaments when shocks are correlated. A large empirical literature has since debated why RPE appears weakly implemented in practice, with contributions by Gibbons and Murphy (1990), Hermalin (1992), and Aggarwal and Samwick (1999). The third is the theory of dynamic contracting in continuous time, where Sannikov (2008) and Williams (2011) developed the promised-utility approach for single-agent models, enabling a tractable recursive formulation of incentive constraints and providing a foundation for modern dynamic agency theory.

Despite these advances, there has been no unified treatment of dynamic contracting with many agents that simultaneously incorporates tournaments, RPE, and competition. Tournaments explain rank-based rewards but neglect dynamic insurance and intertemporal smoothing. RPE highlights the value of correlation structures but is largely static and struggles with multiplicity of equilibria in effort choices. Dynamic models capture intertemporal insurance and continuous renegotiation-proofness but almost always assume a single agent. In practice, however, principals face the combined problem: motivating many agents whose outputs are correlated, whose actions are technologically linked, and whose compensation must be designed dynamically to balance incentives and insurance.

This paper develops a model that resolves these gaps. We study a principal contracting with many agents in continuous time, where outputs follow stochastic processes driven by individual efforts, technological interactions, and both idiosyncratic and aggregate shocks. Contracts are written on filtered signals constructed as linear combinations of outputs, so that the principal can benchmark each agent against their peers. Consumption transfers and incentive exposures to innovations in these signals define the agents' continuation utilities. The analysis shows that the first-order approach remains valid: agent efforts are linear in incentive exposures, with slopes proportional to the signal-to-noise ratios of the filtered signals. This ensures tractability and rules out global deviations, extending Holmström's correlation principle and Mookherjee's multi-agent logic into a dynamic setting.

The principal's problem is then expressed in a Hamilton–Jacobi–Bellman equation. The optimal exposure for each agent depends on the marginal revenue of their effort, the informativeness of the filtered signal, and the curvature of the principal's value function in continuation utilities. The first factor reflects how competition or cooperation shapes the benefit of individual actions, the second reflects how relative performance filters improve identification, and the third reflects the cost of intertemporal insurance that weakens incentives compared to static benchmarks. The optimal filter weights are shown to be precisely the generalized least squares projection of an agent's output on those of their peers, with minimized variance equal to the Schur complement of the covariance matrix. Under common shocks, the filter asymptotically eliminates aggregate risk as the number of agents grows, making incentives progressively more powerful in large organizations.

In contrast to the existing literature, which has largely treated competition, correlation, and dynamic insurance in isolation, our paper unifies these dimensions in a single continuous-time framework. The tractability of the model rests on a new result: the first-order approach remains valid in multi-agent environments with correlated signals, ensuring that incentive compatibility can be reduced to a linear proportionality between effort and exposure. This extends the promised-utility method of Sannikov (2008) and Williams (2011), which were restricted to single-agent problems, to settings with multiple agents interacting under competition and common shocks. Whereas Holmström and Milgrom (1991) studied multitask contracts within a single agent, our focus is on multiple agents whose outputs are technologically linked and stochastically correlated, making the design of relative performance evaluation indispensable.

The paper makes four key contributions. First, it establishes the validity of the first-order approach in dynamic multi-agent contracting with relative performance evaluation, thereby extending the tractable tools of dynamic contract theory to settings that had been regarded as analytically intractable. Second, it derives the principal's Hamilton–Jacobi–Bellman equation in the presence of competition, showing that optimal incentive intensity depends simultaneously on marginal revenue, signal variance after filtering, and intertemporal insurance. Third, it proves that the optimal relative-performance filter is precisely the generalized least squares residual of an agent's output on those of peers, providing

a direct bridge between theory and empirical benchmarking practices. Fourth, in a linear–quadratic specialization, it establishes existence, uniqueness, and global stability of symmetric equilibrium in the mean-field limit, resolving the multiplicity concerns familiar from static multi-agent models such as Mookherjee (1984). These contributions place the paper at the intersection of three classic literatures—tournaments, relative performance evaluation, and dynamic contracting—while moving beyond each to provide a unified and empirically relevant framework.

2 Model and Signals

We consider a continuous-time setting in which a single principal contracts with a set of $N \geq 2$ agents. Each agent i chooses an effort process $a_i(t) \geq 0$, which is not directly observable by the principal but influences a stochastic output process $Y_i(t)$. The dynamics of each output combine the agent’s own action, technological interactions with other agents’ actions, and both idiosyncratic and aggregate random shocks. Formally, the process for $Y_i(t)$ is specified as

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

where $\mu > 0$ is the productivity of an agent’s own effort, $\phi \in \mathbb{R}$, measures technological spillovers or competitive effects from peers, and $\bar{a}_{-i}(t) = \frac{1}{N-1} \sum_{j \neq i} a_j(t)$ is the average effort of the other agents. The shocks consist of an idiosyncratic Brownian component $W_i(t)$ with variance coefficient σ^2 , and a common Brownian component $B(t)$ with variance coefficient σ_0^2 . This specification embeds two important features of multi-agent environments. First, the parameter ϕ allows for both competition and complementarity: when $\phi < 0$, rivals’ efforts reduce the marginal return to agent i ’s action, as in Cournot competition; when $\phi > 0$, the environment approximates teamwork or joint production, where others’ actions enhance an agent’s productivity. Second, the presence of both idiosyncratic and common shocks makes it essential for the principal to design compensation schemes that separate effort-driven performance from random noise.

Each agent derives utility from a stream of consumption $\{c_i(t)\}$ provided by the principal and suffers disutility from exerting costly effort. Preferences are modeled by constant absolute risk aversion with utility $u(c) = -\frac{1}{\gamma}e^{-\gamma c}$, where $\gamma > 0$ is the coefficient of risk aversion. The instantaneous cost of effort is quadratic, $\frac{1}{2}ka_i^2$ with $k > 0$. Given a discount rate $r > 0$, the continuation utility of agent i at time t is

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

This promised utility serves as the sufficient state variable summarizing the agent’s contractual position going forward. The dynamic contracting problem therefore reduces to choosing control processes $(c_i(t), a_i(t))$ in such a way that promise-keeping is satisfied while incentives and insurance are balanced.

Because effort is hidden, the contract cannot condition directly on $a_i(t)$. Instead, the principal must design transfers based on observable output processes. To increase informativeness, the principal is free to construct filtered signals that benchmark one agent's output against those of others. For agent i , a general filter takes the linear form

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

where the weights $\omega_{ij}(t)$ are contractible and can be chosen by the principal. If all weights are zero, the signal corresponds to absolute performance evaluation. If the weights are positive, the contract effectively subtracts the performance of peers, creating relative evaluation. In general, the filter can capture any linear benchmarking scheme, from pure tournaments to correlation-based RPE rules.

To study the informativeness of these signals, let $\Sigma = \text{Var}(dY/\sqrt{dt})$ denote the covariance matrix of output increments. Writing $dZ_i^\omega = e_i^\top dY$, where e_i is a vector with one in the i th position and entries $-\omega_{ij}$ in the j th positions for $j \neq i$, the instantaneous variance of the filtered signal is

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

while its expected drift is

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

The design of ω_i therefore determines the balance between informativeness and bias: a well-chosen filter reduces the variance of noise but may also reduce sensitivity to effort if rivals' contributions are strongly correlated with an agent's own. The principal's task is to minimize $v_i(\omega_i)$ while maintaining adequate sensitivity to a_i .

The innovation associated with each filtered signal is

$$\frac{dZ_i^\omega - \mathbb{E}[dZ_i^\omega | \mathcal{F}_t]}{\sqrt{v_i(\omega_i)}},$$

which is a standardized Brownian increment relative to the available information. By the martingale representation theorem, any admissible contract can be written as a process involving consumption transfers plus linear loadings $\theta_i(t)$ on these standardized innovations. These exposures represent the power of incentives in the contract: larger θ_i increases the sensitivity of the agent's continuation utility to innovations in the filtered signal. Because the distribution of dZ_i^ω is Gaussian with mean linear in effort and variance independent of effort, the first-order approach remains valid, and the agent's optimization problem can be reduced to a quadratic trade-off between effort costs and incentive benefits.

The economic intuition is clear. If outputs are driven by a large common shock, absolute performance evaluation is uninformative: the contract would

reward or punish agents for fluctuations they cannot influence. Relative evaluation improves informativeness by subtracting out common noise, allowing the principal to focus on residual variation more closely tied to individual effort. When efforts are substitutes, so that $\phi < 0$, the performance of peers carries negative information about one's own contribution, making relative signals particularly valuable. When efforts are complements, so that $\phi > 0$, peer outputs carry positive information, and aggressive filtering may inadvertently remove useful signals. In either case, the law of large numbers plays an important role: as the number of agents grows, common shocks can be increasingly well diversified away by benchmarking against many peers, and the variance of the filtered signal converges to its idiosyncratic minimum.

Thus the model delivers a compact but general representation of incentives under competition and correlation. Each agent's contract consists of consumption transfers, exposures to filtered innovations, and filters defined by peer weights. These instruments determine the evolution of promised utilities and the strength of incentives, while also capturing the statistical trade-off between informativeness and insurance. In subsequent sections, we analyze how the principal chooses exposures and filters optimally, derive the Hamilton–Jacobi–Bellman representation of the contracting problem, solve for the variance-minimizing filters, and establish equilibrium existence and uniqueness in the mean-field limit.

3 Principal's HJB and Optimal Incentive Intensity

The foundation of the principal's problem lies in describing how agents' promised utilities evolve under a contract and how incentive compatibility links exposures to effort. Once signals have been filtered as in Section 2, the continuation utility of agent i evolves according to a stochastic differential equation that reflects both the deterministic effect of transfers and the stochastic effect of exposure to innovations. Specifically, using the martingale representation theorem, the dynamics of promised utility satisfy

$$dW_i(t) = \left(rW_i(t) - u(c_i(t)) + \frac{1}{2}ka_i(t)^2 \right) dt + \theta_i(t) \frac{dZ_i^\omega(t) - \mathbb{E}[dZ_i^\omega(t)|\mathcal{F}_t]}{\sqrt{v_i(\omega_i)}},$$

where $c_i(t)$ denotes consumption, $\theta_i(t)$ is the contractual exposure to the standardized innovation of the filtered signal, and $a_i(t)$ is the effort chosen by the agent. The term rW_i arises from discounting, the subtraction of $u(c_i)$ ensures promise-keeping, and the quadratic cost term $\frac{1}{2}ka_i^2$ captures the disutility of effort. The stochastic term reflects the risk that the agent bears through contract exposure, with variance proportional to θ_i^2 .

Incentive compatibility is preserved by the first-order approach. Given a linear exposure θ_i to the filtered innovation, the induced effort is

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

where $\lambda_i(\omega_i)$ is the implementability slope. This parameter measures how strongly effort responds to exposure: it is proportional to productivity μ and inversely related to both the cost parameter k and the standard deviation of the filtered signal. More informative signals (smaller $v_i(\omega_i)$) make λ_i larger, meaning that the same contractual exposure induces higher effort. The validity of the FOA therefore reduces incentive compatibility to a simple proportionality condition between effort and exposure.

Let the principal's value function be denoted by $V(\mathbf{W}, t)$, where $\mathbf{W} = (W_1, \dots, W_N)$ is the vector of agents' promised utilities. Applying Itô's lemma to this function and substituting the continuation utility dynamics yields the Hamilton–Jacobi–Bellman equation for the principal's optimization problem:

$$rV(\mathbf{W}, t) = \sup_{\{c_i, \theta_i, \omega_i\}} \left\{ \mathbb{E}[R(\mathbf{a})] - \sum_{i=1}^N c_i + \sum_{i=1}^N V_{W_i} \left(rW_i - u(c_i) + \frac{1}{2} k a_i^2 \right) + \frac{1}{2} \sum_{i=1}^N V_{W_i W_i} \theta_i^2 \right\}.$$

The first term represents expected revenue, which depends on the effort profile $\mathbf{a} = (a_1, \dots, a_N)$. The second term subtracts the cost of transfers to agents. The third term enforces the promise-keeping constraints, ensuring that continuation utilities evolve consistently with the contract. The final term arises from the variance of continuation utility due to contract exposure, scaled by the curvature of the principal's value function in promised utilities. Because $V_{W_i W_i} \leq 0$ under concavity, this term reflects the risk-sharing cost of exposing the agent to uncertainty.

The first-order conditions clarify the structure of optimal contracts. Differentiation with respect to consumption gives

$$V_{W_i} u'(c_i) = 1,$$

ensuring that the marginal cost of providing consumption equals the shadow value of promised utility. Differentiation with respect to exposure θ_i , using the incentive-compatibility condition $a_i = \lambda_i(\omega_i) \theta_i$, yields

$$\frac{\partial \mathbb{E}[R]}{\partial a_i} \cdot \frac{\partial a_i}{\partial \theta_i} = \text{MR}_i(\mathbf{a}) \lambda_i(\omega_i) = (1 - V_{W_i W_i}) \theta_i,$$

where $\text{MR}_i(\mathbf{a}) = \frac{\partial \mathbb{E}[R]}{\partial a_i}$ denotes the marginal revenue of agent i 's effort. Solving for the optimal exposure gives

$$\theta_i^* = \frac{\lambda_i(\omega_i) \text{MR}_i(\mathbf{a})}{1 - V_{W_i W_i}},$$

and substituting back into the effort condition gives

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

The economic interpretation of this condition is straightforward but powerful. Optimal effort is proportional to the squared implementability slope and the marginal revenue of the agent’s action, and inversely related to the curvature of the principal’s value function. Three distinct forces are therefore at work. First, the marginal revenue term reflects the competitive or technological environment: when outputs are substitutes, marginal revenue falls with the effort of others, weakening incentives; when outputs are complements, marginal revenue rises with others’ effort, strengthening them. Second, the implementability slope reflects signal informativeness: better relative performance evaluation reduces variance, raises λ_i , and therefore increases induced effort. Third, the curvature term reflects the principal’s valuation of insurance: greater concavity of V implies a stronger demand to smooth continuation utilities, which leads to weaker exposure and attenuated incentives.

This decomposition unifies the insights of several literatures. The marginal revenue term generalizes the tournament logic of Lazear and Rosen (1981), where competition for prizes determines the slope of effort incentives. The signal-informativeness term generalizes Holmström’s (1982) correlation principle, where filtering improves the identification of effort. The curvature term introduces a dynamic attenuation force absent in static models: because continuation utilities must evolve smoothly across time, principals optimally dampen incentives relative to the static first-best. This explains why empirical evidence, such as the limited reliance on relative performance evaluation in CEO compensation, may be consistent with optimal dynamic contracting rather than evidence of inefficiency.

The HJB formulation therefore makes explicit the central trade-off faced by the principal in designing incentives for many agents under uncertainty. Incentive intensity is shaped jointly by the competitive structure of revenues, the statistical properties of signals, and the intertemporal need for insurance. The next step of the analysis is to characterize the optimal choice of relative-performance filters themselves. As we show in Section 4, the variance-minimizing filter corresponds to the generalized least squares projection of one agent’s output onto those of their peers, providing a closed-form and economically transparent solution to the filtering problem.

4 Optimal RPE Filter

The effectiveness of incentives in this dynamic contracting environment depends on the variance of the filtered signal that links effort to compensation. As established in the previous section, induced effort satisfies

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

with the implementability slope given by $\lambda_i(\omega_i) = \mu/(k\sqrt{v_i(\omega_i)})$. The smaller the residual variance $v_i(\omega_i)$, the steeper the slope, and the stronger the incentives for a given exposure. The central design problem for the principal is therefore to

choose filter weights ω_{ij} that minimize the variance of the filtered signal while preserving sensitivity to the effort of agent i .

Formally, recall that the filter for agent i can be written as

$$dZ_i^\omega = e_i^\top dY,$$

where e_i is a weight vector with entry 1 in the i th position, $-\omega_{ij}$ in the positions corresponding to other agents $j \neq i$, and zeros elsewhere. Let $\Sigma = \text{Var}(dY/\sqrt{dt})$ denote the instantaneous covariance matrix of outputs. The variance of the filtered signal is

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

and the drift sensitivity to agent i 's own effort must remain normalized to one in order to ensure that the scale of the signal correctly reflects the action of interest. This normalization is essential: otherwise, the principal could simply scale both exposures and filters to induce arbitrary levels of effort, which would not meaningfully change incentives but would distort the analysis. With the normalization in place, the design of ω_i reduces to a variance minimization problem.

This problem has a natural interpretation as generalized least squares regression. Partition the covariance matrix as

$$\Sigma = \begin{pmatrix} \Sigma_{ii} & \Sigma_{i,-i} \\ \Sigma_{-i,i} & \Sigma_{-i,-i} \end{pmatrix},$$

where the block $\Sigma_{-i,-i}$ is the covariance matrix of the outputs of all agents other than i , the block $\Sigma_{i,-i}$ records covariances between agent i and the others, and the scalar Σ_{ii} is the variance of agent i 's own output. The variance-minimizing choice of weights is

$$\omega_i^* = \Sigma_{-i,-i}^{-1} \Sigma_{-i,i},$$

and the minimized variance is given by the Schur complement,

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

In words, the optimal filter is the residual from regressing agent i 's output on the outputs of all others, and its variance is precisely the variance of the regression residual. This characterization provides an explicit solution to the principal's filtering problem: optimal relative performance evaluation is nothing more than generalized least squares projection.

The intuition is straightforward. When outputs share common components, simply rewarding agents for their own outcomes is inefficient: they would be paid for fluctuations that are unrelated to their own actions. By projecting agent i 's output onto the outputs of others and subtracting the predicted component, the principal strips away noise that is common across agents. The remaining residual variation is more closely tied to agent i 's own effort. This filtering procedure is

the dynamic analogue of Holmström's (1982) correlation principle, which stated that all informative signals should enter the contract. Here, informativeness is defined in terms of the regression residual after controlling for peers' outputs, and the principal implements it continuously in real time.

To illustrate the structure more concretely, consider the common case in which outputs consist of idiosyncratic shocks with variance σ^2 and a shared common shock with variance σ_0^2 . In this case the covariance matrix takes the form $\Sigma = \sigma^2 I_N + \sigma_0^2 \mathbf{1}\mathbf{1}^\top$, where $\mathbf{1}$ denotes the vector of ones. Because of symmetry, the optimal weights are identical across peers, and the solution becomes

$$\omega_{ij}^* = \frac{\sigma_0^2}{\sigma^2 + (N-2)\sigma_0^2}, \quad j \neq i.$$

The minimized variance of the filtered signal is then

$$v_i(\omega^*) = \sigma^2 \cdot \frac{\sigma^2 + (N-1)\sigma_0^2}{\sigma^2 + (N-2)\sigma_0^2}.$$

This expression demonstrates clearly how relative performance evaluation improves informativeness. When the number of agents is small, common shocks cannot be perfectly eliminated, and residual variance remains inflated above the idiosyncratic level. But as the number of agents increases, the denominator grows, and the variance converges to σ^2 . In the limit of large N , the common shock is almost entirely filtered away by benchmarking against peers, and the contract isolates the idiosyncratic risk as the only remaining source of noise.

The economic implications are significant. First, the analysis shows that the value of relative performance evaluation depends fundamentally on the covariance structure of outputs. When correlation is positive because of common shocks, subtracting peer outcomes is beneficial; when correlation arises from genuine complementarities in effort, excessive filtering may remove valuable information and weaken incentives. Second, organizational scale matters: larger groups make RPE more effective, because more peers provide a better predictor of common shocks, reducing variance and strengthening incentives. Third, the regression characterization clarifies the link to empirical practice: benchmarking pay against industry indices or peer groups, as is common in executive compensation and finance, is exactly the implementation of GLS filtering.

Thus the design of the optimal RPE filter provides the crucial link between information, incentives, and insurance in the dynamic principal-agent problem. By minimizing variance through regression-based subtraction of common shocks, the principal raises the implementability slope and strengthens incentives, subject only to the intertemporal insurance trade-off reflected in the curvature of the value function. In the next section, we turn to equilibrium analysis in a linear-quadratic setting, where explicit solutions can be derived and existence, uniqueness, and stability of equilibrium can be established in the mean-field limit.

5 Existence and Uniqueness in a Linear–Quadratic Case

To make the dynamic contracting framework tractable and to study equilibrium behavior in detail, we now specialize to a linear–quadratic environment. This structure allows closed-form characterizations of effort, provides transparent comparative statics, and makes it possible to establish existence and uniqueness of equilibrium in the mean-field limit as the number of agents becomes large. The linear–quadratic specification captures the essential trade-offs between incentives, competition, signal informativeness, and insurance, while remaining mathematically manageable.

Consider an industry where the principal’s revenue depends on aggregate output through a linear demand function. Total output is $Q(t) = \sum_{i=1}^N Y_i(t)$, and price is $P(Q) = \alpha - \beta Q$ with $\alpha > 0$ and $\beta > 0$. Revenue is therefore $R(t) = P(Q)Q = (\alpha - \beta Q)Q$. Differentiating with respect to the effort of agent i , and applying the law of motion for output given in Section 2, yields a marginal revenue of the form

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

where $A(t) = \frac{1}{N} \sum a_j(t)$ is mean effort, $\bar{m}$ is a constant proportional to $\mu\alpha$, and $\bar{\beta}$ is proportional to $\mu\beta N$, capturing the depressing effect of competition on marginal revenue. The linear–quadratic form thus embodies the Cournot logic: as rivals exert more effort, the marginal revenue of any individual’s action falls, reducing the incentive to exert effort.

From Section 3, induced effort under optimal exposure and filtering is

$$a_i(t) = \frac{\lambda^2 v^*}{1 - V_{W_i W_i}(t)} (\bar{m} - \bar{\beta}A(t)),$$

where $\lambda = \mu/(k\sqrt{v^*})$ is the implementability slope and v^* is the minimized residual variance of the optimal relative-performance filter derived in Section 4. Symmetry implies that in equilibrium all agents choose the same effort, so $a_i(t) = A(t)$ for all i . Substituting into the fixed-point condition gives

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

where $\underline{V}_{WW}$ denotes the common curvature of the principal’s value function with respect to agents’ promised utilities. Rearranging yields the closed-form equilibrium effort

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

This expression illustrates how equilibrium effort is determined by the interaction of technology, information, competition, and insurance. Effort increases

with the marginal revenue intercept $\bar{m}$, which represents technological productivity and baseline demand. Effort falls with the competition parameter $\bar{\beta}$, since greater competition depresses marginal returns. Effort increases with the informativeness of signals, as a smaller residual variance v^* raises the implementability slope λ and therefore strengthens incentives. Finally, effort decreases with more negative curvature $\underline{V}_{WW}$, since stronger concavity reflects greater value of smoothing continuation utilities, which attenuates exposures and weakens incentives.

The fixed-point mapping provides conditions for existence and uniqueness. Define $\Phi(A) = \frac{\lambda^2 v^*}{1 - \underline{V}_{WW}} (\bar{m} - \bar{\beta} A)$. The slope of this mapping is $-\kappa$, where $\kappa = \frac{\bar{\beta} \lambda^2 v^*}{1 - \underline{V}_{WW}}$. If $0 < \kappa < 1$, the mapping is a contraction, and the Banach fixed-point theorem guarantees the existence of a unique symmetric equilibrium. Moreover, convergence to the equilibrium is global and geometric: starting from any initial mean effort, iterative best responses converge to A^* at rate κ . The contraction property therefore ensures both uniqueness and stability, resolving the multiplicity problem that arises in static models with many agents, such as Mookherjee's (1984) framework, where multiple Nash equilibria may coexist.

The mean-field limit provides additional insight. As $N \rightarrow \infty$, the covariance structure simplifies: the variance of the optimal filtered signal converges to its idiosyncratic component σ^2 , since the common shock is diversified away by benchmarking against a large pool of agents. In this limit, the implementability slope converges to $\lambda = \mu/(k\sigma)$, and the equilibrium effort expression reduces to

$$A^* = \frac{\bar{m}}{\bar{\beta} + \frac{1 - \underline{V}_{WW}}{\mu^2/k^2\sigma^2}}.$$

This uniqueness and stability result can also be interpreted through the lens of mean field games. In the large- N limit, each agent becomes infinitesimal relative to the aggregate, so the equilibrium condition reduces to a fixed point: given the mean effort trajectory $A(t)$, each agent optimally chooses an effort $a^*(W, t; A(\cdot))$, and in equilibrium this coincides with the aggregate, $a^*(W, t) = A(t)$. The contraction mapping established above is precisely the mean field fixed-point condition, ensuring that the aggregate trajectory is uniquely determined and dynamically stable. Far more details will be presented in section 6 with various examples.

Thus in large organizations, incentives approach their idiosyncratic frontier: common shocks are asymptotically eliminated, and effort is constrained only by the trade-off between marginal revenue and intertemporal insurance. This highlights the dual role of scale. Larger organizations allow the principal to filter away more noise, strengthening incentives, but insurance constraints remain binding, preventing effort from reaching the static first-best.

The linear-quadratic specialization therefore demonstrates clearly how dynamic contracting with many agents generates robust equilibrium predictions. Incentives are weaker under stronger competition, stronger under better relative performance evaluation, and always attenuated by the need to smooth continuation utilities over time. The equilibrium is unique and globally stable in

the mean-field case, overcoming the multiplicity of equilibria that characterizes static multi-agent models. These results provide both theoretical clarity and empirical relevance, laying the groundwork for applications to executive pay, team incentives, and organizational design.

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

In this section we analyze how the principal induces competition among agents in the mean field contracting environment and derive the equilibrium outcomes across the three diffusion families for which the forward Fokker–Planck equation admits explicit solutions. The analysis demonstrates how the contract structure—through the choice of relative performance filters, incentive exposures, and state-dependent curvature of the principal’s value function—determines both the aggregate effort and the cross-sectional distribution of continuation utilities.

We first present the general Hamilton–Jacobi–Bellman (HJB) and Fokker–Planck (FPK) system for the continuum of agents. We then specialize to three diffusion families—Ornstein–Uhlenbeck (OU/Gaussian), Cox–Ingersoll–Ross (CIR), and geometric/lognormal—that admit closed-form transition densities, thereby permitting explicit derivation of the principal’s contracts in the mean field setting.

Let W denote an agent’s promised utility, and m_t the cross-sectional distribution of W at time t . The principal’s value function $V(W, t)$ satisfies the backward HJB equation

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

where effort is determined by the first-order approach as $a = \lambda(\omega)\theta$ with slope $\lambda(\omega) = \mu/(k\sqrt{v(\omega)})$. Given optimal contracts $(c^*, \theta^*, \omega^*)$, the continuation utility evolves as

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

The law of motion of the distribution is then governed by the forward Fokker–Planck equation

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

A mean field equilibrium is a pair (V, m) such that the HJB equation yields optimal policies and the FPK equation evolves consistently with those policies, with the additional fixed-point condition that the aggregate effort satisfies $A_t = \int a^*(W, t) m_t(dW)$.

6.1 Competition in the mean field game.

Competition is introduced through the principal's use of relative performance evaluation (RPE). Each agent $i \in [0, 1]$ is evaluated on the filtered signal

$$dZ_i = dY_i - \int_{[0,1] \setminus \{i\}} \omega_i^*(j) dY_j,$$

where ω_i^* is chosen to minimize the conditional variance of the innovation. In symmetric equilibrium this reduces to benchmarking each agent against the average performance of all others. The resulting marginal revenue of effort takes the form

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

where $A = \int_0^1 a_j dj$ is aggregate effort. Each agent therefore optimizes taking A as given, while in equilibrium A must coincide with the average of optimal responses. This fixed-point condition captures the essence of competition: every agent's payoff depends on their relative standing vis-à-vis the mean, which itself is determined endogenously in equilibrium.

6.2 Ornstein–Uhlenbeck (Gaussian) case.

Under CARA utility, Gaussian shocks, and quadratic value functions, the curvature V_{WW} is constant, yielding exposures and efforts independent of the agent's current promise:

$$\theta^*(t) = \frac{\lambda(\bar{m} - \bar{\beta}A_t)}{1 - \alpha_2(t)}, \quad a^*(t) = \frac{\lambda^2(\bar{m} - \bar{\beta}A_t)}{1 - \alpha_2(t)}.$$

The promise-keeping condition gives the consumption transfer

$$c^*(W, t) = \frac{1}{\gamma} \ln \left(\frac{1}{\alpha_1(t) + \alpha_2(t)W} \right).$$

Substituting these rules yields the continuation utility dynamics

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

an Ornstein–Uhlenbeck process with Gaussian distribution

$$m_t(W) = \frac{1}{\sqrt{2\pi v(t)}} \exp \left(-\frac{(W - \bar{W}(t))^2}{2v(t)} \right),$$

with mean and variance solving $\dot{\bar{W}} = \alpha - \kappa \bar{W}$ and $\dot{v} = \sigma^2 - 2\kappa v$. The stationary distribution is normal with mean α/κ and variance $\sigma^2/(2\kappa)$. The mean-field fixed point is

$$A_{\text{OU}}^* = \frac{\lambda^2 \bar{m}}{\bar{\beta} \lambda^2 + 1 - \alpha_2}.$$

Competition in this case is mild: all agents exert the same effort in equilibrium, while continuation utilities are symmetrically dispersed around the mean in a Gaussian distribution.

6.3 Cox–Ingersoll–Ross (CIR) case.

If the curvature $V_{WW}(W)$ is chosen so that exposure scales as $\theta^*(W, t) = \zeta(t)\sqrt{W}$, then effort becomes

$$a^*(W, t) = \lambda \zeta(t) \sqrt{W}.$$

The continuation utility evolves as a Cox–Ingersoll–Ross process,

$$dW_t = \kappa(\Theta - W_t) dt + \sigma \sqrt{W_t} d\tilde{B}_t,$$

whose stationary distribution is Gamma with mean Θ and variance $\Theta\sigma^2/(2\kappa)$. The equilibrium effort is determined by

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

where the expectation is taken with respect to the Gamma law. At stationarity,

$$\mathbb{E}[\sqrt{W}] = \sqrt{\frac{\sigma^2}{2\kappa}} \cdot \frac{\Gamma(\nu + \frac{1}{2})}{\Gamma(\nu)}, \quad \nu = \frac{2\kappa\Theta}{\sigma^2}.$$

Competition in this case generates greater heterogeneity: effort levels remain symmetric across agents, but continuation utilities are dispersed according to a skewed, right-tailed distribution, resembling tournament outcomes with clear winners and losers.

6.4 Geometric Brownian (lognormal) case.

If the curvature is chosen such that $\theta^*(W, t) = \eta(t)W$, then effort scales linearly with promised utility:

$$a^*(W, t) = \lambda \eta(t) W.$$

Continuation utility then follows a geometric Brownian motion,

$$dW_t = \tilde{\alpha} W_t dt + \tilde{\sigma} W_t d\tilde{B}_t,$$

implying a lognormal cross-sectional distribution,

$$m_t(W) = \frac{1}{W \tilde{\sigma} \sqrt{2\pi t}} \exp\left(-\frac{(\ln W - \ln W_0 - (\tilde{\alpha} - \frac{1}{2}\tilde{\sigma}^2)t)^2}{2\tilde{\sigma}^2 t}\right).$$

The mean-field fixed point is

$$A_{\text{Log}}^* = \frac{\lambda d_0 \mathbb{E}[W_t] \bar{m}}{1 + \lambda d_0 \mathbb{E}[W_t] \bar{\beta}},$$

with $\mathbb{E}[W_t] = W_0 e^{\tilde{\alpha}t}$. Here competition produces extreme inequality: while all agents choose the same proportional effort, multiplicative shocks create a fat-tailed lognormal distribution, with a small set of agents emerging as large “winners.”

6.5 Comparative Implications.

Across all three families, the mean-field fixed point takes the common form

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

where Ξ is a contract coefficient depending on curvature and information parameters, and $\mathcal{M}$ is a moment of the promised-utility distribution ($\mathcal{M} = 1$ in the OU case, $\mathcal{M} = \mathbb{E}[\sqrt{W}]$ under CIR, and $\mathcal{M} = \mathbb{E}[W]$ under lognormal). The structure of competition is thus unified, but the distributional outcomes differ sharply: Gaussian competition yields mild symmetric dispersion, CIR competition produces skewed outcomes with moderate inequality, and lognormal competition generates strongly tournament-like outcomes with extreme inequality. By selecting the curvature of the value function—and thus the exposure schedule—the principal can tune the intensity and distributional consequences of competition among agents. We explore these issues much deeper using the mean field game framework in next section.

7 Formal Comparison: Tournaments vs. Linear RPE, and Unification via Mean Field Games

A central branch of the multi-agent incentive literature contrasts tournament contracts, introduced by Lazear and Rosen (1981), with linear relative performance evaluation (RPE) contracts as characterized by Mookherjee (1984). Both approaches emphasize the role of correlated outputs and the possibility of filtering common shocks, but they differ markedly in the technical structure of incentive provision. Our framework embeds both strands as special cases within a continuous-time, mean field contracting model.

The tournament model assumes that the principal can commit to a rank-order payoff scheme: the worker who produces the highest observed output receives a high prize W_H , while others receive a lower prize W_L . Letting $\Delta = W_H - W_L$ denote the prize spread, an agent’s probability of winning is $\Pr(\varepsilon_j - \varepsilon_i < a_i - a_j)$. With two risk-neutral agents of symmetric ability and i.i.d. shocks, the symmetric Nash equilibrium effort a^* satisfies

$$c'(a^*) = \Delta f(0),$$

where $f(0)$ is the density of the noise difference at zero. This condition follows because the marginal benefit of an increase in effort is the increase in the probability of winning, multiplied by the prize spread. Hence effort is proportional to

Δ . The principal, who observes only the winner, can replicate first-best effort by setting

$$\Delta^* = \frac{c'(a^*)}{f(0)}.$$

Thus tournaments are efficient, in the sense that the optimal spread can induce first-best effort even when effort is unobservable. The crucial feature is that only relative ranks matter; common shocks that shift both agents' outputs equally drop out, thereby providing full insurance against aggregate noise.

The linear RPE framework of Mookherjee (1984) starts instead with risk-averse agents with CARA utility and normally distributed output shocks. Outputs are correlated, with covariance matrix Σ . Each agent's contract is assumed linear:

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

The agent's expected utility is

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

The incentive condition is simply $c'(a_i) = \beta_i$, so marginal incentives depend only on the slope with respect to own output. The role of the cross-terms δ_{ij} is to reduce variance without altering the marginal incentive: the principal sets them equal to the best linear predictor coefficients from regressing y_i on $\{y_j : j \neq i\}$. This is precisely the generalized least squares (GLS) filter, which minimizes the variance of the innovation. The result is that incentives β_i are pinned down by technology and cost, while relative performance terms optimally insure against common shocks.

A technical distinction arises here. In tournaments, the distribution of shocks enters via the density $f(0)$, which governs how sensitive the probability of winning is to small deviations in effort. The condition $c'(a^*) = \Delta f(0)$ ties optimal spread to the noise distribution. In linear RPE contracts, the distribution of shocks enters instead through the covariance matrix Σ , which determines the GLS filter and hence the residual variance v_i^* . The incentive slope then becomes

$$a_i^* = \frac{\mu}{k \sqrt{v_i^*}} \theta_i,$$

where θ_i is the contract's exposure parameter. Thus, whereas in tournaments incentives are mediated by the density of the noise difference, in linear contracts they are mediated by the residual variance after filtering.

Economically, tournaments produce discrete, high-powered incentives because only ranks matter; the marginal return to effort is the probability of winning. This structure naturally produces "winner-takes-all" outcomes, large inequality, and the possibility of sabotage or collusion. Linear RPE contracts,

in contrast, produce continuous, lower-powered incentives with smoother risk-sharing, since each agent is rewarded on a continuous scale against the filtered outputs of others. Both forms economize on risk by eliminating common noise, but they distribute residual idiosyncratic risk differently: tournaments allocate it in a discontinuous all-or-nothing fashion, while linear RPE contracts smooth it through variance reduction.

Our contribution situates both approaches within a continuous-time mean field game. By modeling continuation utilities explicitly, we show that different specifications of the curvature of the value function $V(W)$ yield different families of promised-utility dynamics. The Ornstein–Uhlenbeck case corresponds to the Gaussian linear RPE framework, where incentives are smooth and utilities are symmetrically dispersed. The lognormal case reproduces a tournament-like structure: exposures scale with promised utility, inducing multiplicative dispersion and fat-tailed outcomes, much like the winner-takes-all outcomes of Lazear–Rosen. The CIR case provides an intermediate benchmark: exposures scale with $\sqrt{W}$, generating skewed outcomes with more moderate inequality.

From a theoretical standpoint, our mean field game formulation clarifies that both tournaments and linear RPE contracts are limiting cases of dynamic incentive provision under different diffusion specifications. The technical connection is that both rely on FOA validity: in tournaments, FOA appears as the agent’s probability of winning being differentiable in effort, yielding a closed-form FOC; in linear RPE, FOA appears as a linear mapping between incentives and effort via the GLS-filtered variance. By embedding these within dynamic HJB–FPK systems, we unify the two literatures in a single framework that permits both comparative statics and explicit analytical solutions. We present formal proofs in the following three propositions and one corollary.

7.1 Proposition 1 (Lazear–Rosen tournament effort condition):

Let two agents $i \in \{1, 2\}$ exert efforts $a_i \geq 0$ at cost $c(a_i)$ with $c' > 0, c'' > 0$. Outputs are

$$y_i = a_i + \varepsilon_i,$$

where $\varepsilon_1, \varepsilon_2$ are i.i.d. with density f and cdf F . A tournament contract pays W_H to the winner ($y_i > y_j$) and W_L otherwise; write the prize spread $\Delta := W_H - W_L$. In a symmetric equilibrium $a_1 = a_2 = a^*$, the first-order condition for each agent is

$$\boxed{c'(a^*) = \Delta \cdot f_D(0)} \quad \text{with } D := \varepsilon_j - \varepsilon_i,$$

where f_D is the density of D (for i.i.d. shocks with density f , f_D is the self-convolution density, symmetric with $f_D(0) = \int f(x)^2 dx$). Hence equilibrium effort is proportional to the prize spread.

Proof.

Fix $a_j = a$. Agent i 's probability of winning is

$$p(a_i; a) = \Pr(y_i > y_j) = \Pr(a_i - a + \varepsilon_i - \varepsilon_j > 0) = \Pr(D < a_i - a) = F_D(a_i - a).$$

Her expected wage is $W_L + \Delta p(a_i; a)$. Expected utility is

$$U(a_i; a) = W_L + \Delta F_D(a_i - a) - c(a_i).$$

FOC at the symmetric profile $a_i = a$ gives

$$0 = \frac{\partial U}{\partial a_i} \Big|_{a_i=a} = \Delta f_D(0) - c'(a),$$

which rearranges to the stated condition. The SOC holds since $c'' > 0$ and f_D is bounded.

The economic intuition is quite simple: The marginal benefit of effort equals the marginal increase in win probability evaluated at zero gap (because the opponent exerts the same effort), multiplied by the spread Δ . Common shocks cancel in ranks, so incentives are pinned down by idiosyncratic dispersion $f_D(0)$. Larger Δ or tighter noise (larger $f_D(0)$) raise effort.

7.2 Proposition 2 (Linear RPE with GLS filtering à la Mookherjee):

Consider $N \geq 2$ agents with outputs

$$y = a + \varepsilon, \quad a = (a_1, \dots, a_N)^\top, \quad \varepsilon \sim \mathcal{N}(0, \Sigma),$$

and each agent i has CARA utility with risk aversion $\gamma > 0$ and quadratic effort cost $\frac{1}{2}ka_i^2$. Let wages be linear in all outputs:

$$w_i = \alpha_i + \beta_i y_i + \sum_{j \neq i} \delta_{ij} y_j = \alpha_i + b_i^\top y, \quad b_i := (\beta_i, \{\delta_{ij}\}_{j \neq i}).$$

Then, for a given incentive slope β_i , the variance-minimizing choice of cross-coefficients is the GLS projection

$$\boxed{\delta_i^* = -\Sigma_{-i}, -i^{-1} \Sigma_{-i,i}} \quad (\text{so } b_i^* = e_i - [0; \delta_i^*]),$$

and the induced wage variance is

$$\text{Var}(w_i) = b_i^{*\top} \Sigma b_i^* = \Sigma_{ii} - \Sigma_{i,-i} \Sigma_{-i,-i}^{-1} \Sigma_{-i,i} =: v_i^*.$$

Moreover the agent's FOC is $ka_i = \beta_i$, so at an interior optimum

$$\boxed{a_i^* = \frac{\beta_i}{k}}.$$

Equivalently, writing the contract in filtered-signal form

$$dZ_i^\omega = dY_i - \sum_{j \neq i} \omega_{ij} dY_j, \quad \omega_i^* = \Sigma_{-i, -i}^{-1} \Sigma_{-i, i},$$

the residual variance is v_i^* (Schur complement) and the FOA mapping in our dynamic notation is

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

Proof.

With CARA–Normal, agent i 's certainty equivalent is linear in mean minus $\frac{\gamma}{2}$ times variance. Since $\mathbb{E}[w_i] = \alpha_i + \beta_i a_i + \sum_{j \neq i} \delta_{ij} a_j$, the marginal incentive is $\partial \mathbb{E}[w_i] / \partial a_i = \beta_i$. Thus $k a_i = \beta_i$. For risk, $\text{Var}(w_i) = b_i^\top \Sigma b_i$. Holding β_i fixed, minimizing variance over δ_i is the standard GLS residualization of y_i on y_{-i} : the minimizing coefficients are $\delta_i^* = -\Sigma_{-i, -i}^{-1} \Sigma_{-i, i}$, delivering residual variance v_i^* equal to the Schur complement. Rewriting in filtered form simply renames b_i^* as the coefficients of Z_i^ω . The FOA mapping follows since the signal-to-noise slope equals $\mu / (k \sqrt{v_i^*})$ under the linear-Gaussian technology.

Hence, in linear RPE the cross-weights serve solely to insure (reduce variance) for a given marginal incentive β_i . Optimal cross-weights implement the “benchmark against peers” rule: subtract the best linear predictor of peers to purge common shocks. Compared to tournaments, incentives are continuous, and dispersion in pay comes from Gaussian residual risk rather than winner-take-all jumps.

7.3 Proposition 3 (Unification in the dynamic mean-field game: a single fixed-point form):

In the continuous-time model with the first-order approach, suppose the principal selects the curvature of the value function V_{WW} , and thereby the exposure schedule θ^* , such that the continuation utility W follows one of three diffusion families. In the Ornstein–Uhlenbeck (Gaussian) case, the exposure θ^* is independent of W . In the CIR case, the exposure scales with the square root of promised utility, so that $\theta^*(W) \propto \sqrt{W}$. In the GBM case, the exposure scales linearly with promised utility, so that $\theta^*(W) \propto W$. Let the Cournot-style marginal revenue be $\text{MR}(A) = \bar{m} - \beta A$. Then in all three families the symmetric mean-field equilibrium effort A^* satisfies the unified closed form

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

where the contract coefficient $\Xi > 0$ and the statistical moment $\mathcal{M}$ depend on the diffusion family. Specifically, in the Gaussian case $\Xi = \lambda^2 / (1 - \alpha_2)$ and $\mathcal{M} = 1$; in the CIR case $\Xi = \lambda c_0$ and $\mathcal{M} = \mathbb{E}[\sqrt{W}]$; and in the GBM case $\Xi = \lambda d_0$ and $\mathcal{M} = \mathbb{E}[W]$. Here $\lambda = \mu / (k \sqrt{v^*})$ denotes the implementability slope associated with GLS filtering of signals, α_2 is the constant curvature of the quadratic value

function in the Gaussian case, and $c_0, d_0 > 0$ are constants determined by the mapping $1 - V_{WW}(\cdot)$ that delivers square-root or linear exposure schedules. In each family, global uniqueness holds provided that $\Xi\mathcal{M}\bar{\beta} \in (0, 1)$, since in this case the best-response map $A \mapsto \Xi\mathcal{M}(\bar{m} - \bar{\beta}A)$ is a contraction.

Proof. The first-order approach combined with the principal's HJB delivers the generic mapping between effort and exposure,

$$a^*(\cdot) = \frac{\lambda \text{MR}(A)}{1 - V_{WW}(\cdot)} \lambda = \Xi(\cdot)(\bar{m} - \bar{\beta}A),$$

with $\Xi(\cdot) = \lambda^2/(1 - V_{WW}(\cdot))$. The mean effort is therefore $A = \mathbb{E}[a^*(\cdot)]$. In the Gaussian case, the curvature is constant, $V_{WW} = \alpha_2$, so $\Xi(\cdot) \equiv \lambda^2/(1 - \alpha_2) =: \Xi$ and the expectation reduces to $\mathcal{M} = 1$. This yields $A = \Xi(\bar{m} - \bar{\beta}A)$, which rearranges to the closed form stated. In the CIR case, the mapping satisfies $1 - V_{WW}(W) \propto 1/\sqrt{W}$, so $\Xi(W) = \lambda^2/(1 - V_{WW}(W)) = \lambda c_0 \sqrt{W}$. Taking expectations gives $A = \lambda c_0 \mathbb{E}[\sqrt{W}](\bar{m} - \bar{\beta}A)$, with $\mathcal{M} = \mathbb{E}[\sqrt{W}]$. In the GBM case, $1 - V_{WW}(W) \propto 1/W$, so $\Xi(W) = \lambda d_0 W$, which implies $A = \lambda d_0 \mathbb{E}[W](\bar{m} - \bar{\beta}A)$. In every case, rearranging produces

$$A = \frac{\Xi\mathcal{M}\bar{m}}{1 + \Xi\mathcal{M}\bar{\beta}}.$$

The slope of the best-response map is $-\Xi\mathcal{M}\bar{\beta}$. If this lies in $(-1, 0)$, then the map is a contraction, ensuring existence, uniqueness, and global stability by Banach's fixed-point theorem.

This unification demonstrates that the form of competition is governed by the single sufficient statistic $\Xi\mathcal{M}$, which combines the curvature of the contract (insurance versus incentive strength) with the informational moment of the continuation-utility distribution. When $\mathcal{M}$ is large and highly dispersed, as in the GBM case where $\mathcal{M} = \mathbb{E}[W]$, the resulting distribution of continuation utilities is fat-tailed, producing tournament-like outcomes in which a small subset of agents pull far ahead. When $\mathcal{M} = 1$, as in the Gaussian linear RPE environment, dispersion is mild and symmetric, corresponding to continuous benchmarking with smoother risk-sharing. The CIR case occupies an intermediate position: $\mathcal{M} = \mathbb{E}[\sqrt{W}]$ generates skewed distributions with moderate inequality. In all cases stronger competition (a higher $\bar{\beta}$) depresses equilibrium effort, while improved information (a smaller residual variance v^* and thus larger λ) or reduced demand for insurance (a lower curvature V_{WW}) raise effort.

7.4 Corollary (GBM exposures are the dynamic analogue of Lazear–Rosen tournaments):

Consider the GBM specification of Proposition 3 in which the principal chooses curvature so that the optimal exposure satisfies $\theta^*(W, t) = \eta(t)W$ and the agent's best response is $a^*(W, t) = \lambda \eta(t)W$. Let promised utility evolve as

$$dW_t = \tilde{\alpha}(t)W_t dt + \tilde{\sigma}(t)W_t dB_t, \quad \ln W_t = \ln W_0 + \int_0^t (\tilde{\alpha}(s) - \frac{1}{2}\tilde{\sigma}(s)^2)ds + \int_0^t \tilde{\sigma}(s)dB_s.$$

Fix a horizon $T > 0$ and suppose the terminal prize depends on relative rank in W_T (equivalently in $\ln W_T$). Then, for small local deviations in effort paths around a symmetric mean-field equilibrium, the agent's first-order condition for optimal effort coincides with the Lazear–Rosen tournament condition:

$$(\text{marginal cost of effort}) = \Delta_T \cdot f_{\Delta \ln W_T}(0),$$

where Δ_T is an effective prize spread induced by the contract and $f_{\Delta \ln W_T}(0)$ is the cross-sectional density at zero of the log-utility difference between two agents at time T . Hence GBM exposure schedules implement a dynamic, continuous-time tournament: the probability of finishing ahead depends on mean shifts of $\ln W_T$ caused by effort, and the effective prize spread scales with the state W , delivering winner-take-all features.

Proof (sketch). Work under symmetry and consider two agents i, j with identical primitives and exposure $\theta^* = \eta W$. Let $\xi_t := \ln W_t$. From Itô's formula,

$$d\xi_t = (\tilde{\alpha}(t) - \tfrac{1}{2}\tilde{\sigma}(t)^2)dt + \tilde{\sigma}(t)dB_t.$$

A marginal deviation δa in effort raises θ by $\delta\theta = \delta a/\lambda$. Because $\theta = \eta W$ and $a = \lambda\eta W$, a small increase δa at time s raises the mean drift of ξ by a constant increment $\delta\mu_s := \frac{\partial}{\partial a}\mathbb{E}[d\xi_s|\cdot] \cdot \delta a$. Aggregating from 0 to T , the induced shift in the mean of ξ_T relative to a symmetric opponent is

$$\Delta\mu_T = \int_0^T \Gamma(s) \delta a_s ds,$$

for some deterministic kernel $\Gamma(s) > 0$ pinned down by the HJB shadow price of promised utility (formally, $\Gamma = \partial\{\text{drift of } \xi\}/\partial a$ evaluated at the symmetric equilibrium control). Since ξ_T is Gaussian with variance $\Sigma_T := \int_0^T \tilde{\sigma}(s)^2 ds$, the difference $\Delta\xi_T := \xi_T^i - \xi_T^j$ is Gaussian with mean $\Delta\mu_T$ and variance $2\Sigma_T$. The probability that agent i finishes ahead equals

$$p_T(\Delta\mu_T) = \Pr(\Delta\xi_T > 0) = \Phi\left(\frac{\Delta\mu_T}{\sqrt{2\Sigma_T}}\right).$$

Differentiating at the symmetric point $\Delta\mu_T = 0$ gives

$$\left.\frac{d}{d\varepsilon} p_T(\varepsilon)\right|_{\varepsilon=0} = \frac{1}{\sqrt{2\pi} \sqrt{2\Sigma_T}} = f\Delta\xi_T(0),$$

the density at zero of the log-difference. If terminal pay implements a rank-dependent prize with effective spread Δ_T (possibly the shadow value of promised utility times a discrete bonus, or a steep slope of continuation value around the terminal ranking), the agent's expected terminal payoff changes at the margin by

$$\Delta_T \cdot \left.\frac{d}{d\varepsilon} p_T(\varepsilon)\right|_{\varepsilon=0} = \Delta_T f\Delta\xi_T(0).$$

Equating this marginal benefit with the marginal cost $\int_0^T c'(a_s) \delta a_s ds$ and normalizing by the shape of the deviation δa yields the tournament-type FOC at the symmetric equilibrium:

$$c'(a_s^*) = \Delta_T f_{\Delta\xi_T}(0) \Gamma(s) \quad \text{for a.e. } s \in [0, T].$$

Integrating over s (or considering stationary controls) collapses to the usual Lazear–Rosen structure “marginal cost equals prize spread times density at zero,” with the effective spread $\Delta_T \Gamma(\cdot)$ determined by the GBM exposure and the HJB shadow price. Because $\theta^*(W) = \eta W$, high- W agents induce larger mean shifts in ξ per unit of effort, so their effective spread scales with W . This scaling, combined with the lognormal cross-sectional law, generates the same qualitative selection and dispersion as rank-order tournaments.

Therefore, the GBM specification converts the continuation utility into a multiplicative “score.” Effort raises the drift of $\ln W$, so the chance of beating the field is a smooth function of how much effort shifts the mean of $\ln W_T$ relative to its Gaussian dispersion. The first-order change in winning probability is exactly the density-at-zero term characteristic of tournaments. The exposure rule $\theta^*(W) = \eta W$ implies that the effective prize spread is state-contingent and increasing in W : high- W agents face steeper stakes, push further ahead, and generate the heavy right tail observed under GBM. Thus GBM exposures replicate the core mechanics of Lazear–Rosen tournaments within a continuous-time, mean-field contract: competition is about relative rank, common shocks vanish in ranks, and incentives are governed by a prize-spread-times-density-at-zero formula—now endogenously produced by the contract and the HJB shadow prices rather than imposed as a primitive.

8 Empirical and Applied Implications

The theoretical framework developed in this paper not only advances the theory of dynamic contracting but also provides insights into a wide range of empirical puzzles in organizational economics and corporate finance. By embedding relative performance evaluation, tournaments, and competition within a continuous-time model, the analysis produces testable predictions about the use of benchmarking, the role of organizational size, and the observed strength of incentives in practice.

A first domain of application is executive compensation. Holmström’s informativeness principle suggests that contracts should exploit any available signal that reduces variance in the measurement of effort, which would imply widespread use of relative performance evaluation against industry peers or indices. Yet empirical studies such as Gibbons and Murphy (1990) and Aggarwal and Samwick (1999) have found only weak evidence of such benchmarking in chief executive officer contracts. Our model offers a dynamic explanation. Because the principal must deliver smooth continuation utilities to risk-averse agents, the curvature of the principal’s value function attenuates exposures, reducing the optimal use of relative performance signals even when they are highly

informative in a static sense. Thus the apparent underuse of RPE in executive pay does not necessarily indicate inefficiency or poor governance, but rather reflects an intertemporal insurance motive that static models omit.

A second application lies in sales teams and tournament-style incentives. Lazear and Rosen (1981) argued that tournaments generate strong incentives because relative rank automatically filters out common shocks. Empirical evidence from sales contests and internal competitions supports this prediction, yet observed prize spreads are often smaller than static models would suggest. The dynamic perspective clarifies this discrepancy. While tournaments indeed remain effective in large organizations—because common shocks can be filtered almost completely when many peers are available—the intertemporal smoothing of continuation utilities reduces the exposure of agents to the variance of relative signals. The result is that prizes and penalties are optimally muted relative to the static benchmark. The model therefore reconciles the strong presence of tournaments with the moderate intensity of incentives observed in practice.

A third area concerns financial firms and asset management. Compensation schemes in hedge funds and investment banks often involve bonuses based on benchmarks such as outperforming the S&P 500 or beating a peer group of funds. These benchmarks are precisely the optimal filters derived in Section 4: they correspond to generalized least squares projections that purge common market shocks while preserving residual variation attributable to individual effort or skill. The model further predicts that such benchmarking is more powerful in larger organizations or peer groups, where the law of large numbers allows common shocks to be nearly eliminated. This insight explains why small funds exhibit more variable pay-for-performance sensitivity, while larger asset managers rely heavily on index-relative performance contracts.

The framework also has implications for the role of organizational scale more broadly. As the number of agents increases, relative performance evaluation becomes increasingly effective, because the residual variance of filtered signals converges to the idiosyncratic level. Large organizations therefore have a natural informational advantage in designing contracts that motivate effort while limiting exposure to aggregate risk. At the same time, the dynamic insurance channel prevents incentives from scaling without bound: even with perfect filtering, the principal chooses to limit exposures to smooth continuation utilities. This generates a non-monotonic empirical prediction. Reliance on RPE should increase with organizational size, as variance reductions become more significant, but the effective strength of individual incentives relative to variance may plateau once insurance constraints dominate.

Taken together, these empirical implications demonstrate the value of embedding tournaments and relative performance evaluation in a dynamic contracting model. The framework rationalizes puzzling features of compensation practices that have long challenged static theories: the weak reliance on RPE in executive contracts, the smaller-than-predicted prize spreads in tournaments, and the widespread use of benchmarks in finance. It also generates new testable hypotheses: larger organizations should exhibit stronger reliance on RPE, sales contests should show muted prize spreads relative to static predictions, and fi-

nancial firms should employ benchmarks in ways that correspond to GLS residuals of common shocks. By linking dynamic theory to observed practices, the model not only unifies the literatures on moral hazard, tournaments, and relative evaluation but also provides a platform for empirical analysis in diverse organizational settings.

9 Conclusion

This paper has developed a dynamic theory of contracting with one principal and many agents, extending the promised-utility framework of Sannikov (2008) to incorporate competition, relative performance evaluation, and the presence of both idiosyncratic and common shocks. By allowing the principal to construct filtered signals through linear benchmarking of agents' outputs, we showed that incentive design reduces to the selection of consumption transfers, exposures to innovations, and peer-weighting schemes. The analysis demonstrated that the first-order approach remains valid in this multi-agent environment: effort is proportional to exposure with a slope determined by the informativeness of the filtered signal. This tractability provides a foundation for studying multi-agent moral hazard in continuous time without the multiplicity problems that arise in static models.

The Hamilton–Jacobi–Bellman representation of the principal's problem revealed the fundamental structure of dynamic incentives. Optimal effort depends on three forces: the marginal revenue of an agent's action, the signal-to-noise ratio of the filtered performance measure, and the curvature of the principal's value function in continuation utilities. The first reflects the economic environment, including competition or complementarities between agents. The second reflects information design, with relative performance evaluation reducing variance and strengthening incentives. The third reflects the demand for insurance over time, which attenuates incentives relative to the static first-best. Together, these forces unify the insights of Lazear and Rosen's tournament model, Holmström's correlation principle, and Mookherjee's analysis of multi-agent moral hazard within a dynamic framework.

A key technical result established that the optimal filter is the generalized least squares projection of one agent's output onto those of their peers, with the minimized variance equal to the Schur complement of the covariance matrix. This regression-based characterization makes relative performance evaluation transparent: the principal isolates residual variation by subtracting the best linear predictor of peer outcomes. In environments with common shocks, residual variance converges to the idiosyncratic level as the number of agents grows, demonstrating how large organizations can nearly eliminate aggregate noise from performance signals.

In a linear–quadratic specialization, the model delivered closed-form solutions and strong equilibrium results. Effort levels were shown to depend positively on productivity, negatively on competition, positively on informativeness of signals, and negatively on the strength of intertemporal insurance needs.

The equilibrium mapping was shown to be a contraction, guaranteeing existence, uniqueness, and global stability of symmetric equilibrium in the mean-field limit. This result resolves the long-standing problem of multiple equilibria in static multi-agent contracting and highlights the robustness of the dynamic promised-utility approach.

The model also provides a natural explanation for several empirical puzzles. The limited use of relative performance evaluation in executive pay, despite its informativeness, can be rationalized by intertemporal insurance demands. The prevalence of tournaments in sales teams, alongside muted prize spreads relative to static predictions, follows from the same dynamic insurance channel. The reliance of financial firms on benchmark-relative contracts corresponds exactly to the GLS residualization of common shocks implied by the theory. More broadly, the framework predicts that larger organizations should rely more heavily on relative performance evaluation as variance reductions become significant, though the strength of individual incentives is ultimately bounded by insurance considerations.

Our analysis therefore extends dynamic contract theory beyond its traditional single-agent focus. The continuous-time promised-utility models of Sannikov (2008) and Williams (2011) established tractability in individual moral hazard environments, but did not address competition, correlation, or relative performance evaluation. Holmström and Milgrom (1991) studied multitask contracts within a single agent, but not interactions across many agents. By proving the validity of the first-order approach in a multi-agent setting, deriving the principal’s Hamilton–Jacobi–Bellman equation with competition, and showing that optimal filters correspond to generalized least squares residuals, our framework generalizes these dynamic studies to organizational environments with many interacting agents. The linear–quadratic specialization further demonstrates existence, uniqueness, and global stability of symmetric equilibria, overcoming the multiplicity that plagues static multi-agent models such as Mookherjee (1984).

In this way, the paper unifies the tournament perspective of Lazear and Rosen (1981), the correlation principle of Holmström (1982), and the multi-agent analysis of Mookherjee (1984) with the dynamic tools of Sannikov (2008) and Williams (2011). By combining these strands, it establishes a comprehensive and empirically relevant theory of dynamic multi-agent contracting. The contributions are both theoretical, in resolving long-standing questions of tractability and equilibrium selection, and applied, in providing a rigorous foundation for understanding observed practices in executive compensation, sales contests, and financial benchmarking.

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.
- Che, Y.-K., & Yoo, S.-W. (2001). Optimal incentives for teams. *American*

Economic Review, 91(3), 525–541.

- 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.
- Hermalin, B. E. (1992). The effects of competition on executive behavior. *RAND Journal of Economics*, 23(3), 350–365.
- 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. (1995). *Personnel economics*. MIT Press.
- 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.