

Individual Price of Strategic Myopia in Mean-Field-Type Games

Hamidou Tembine

Julian Barreiro-Gomez

Abstract—We consider a class of discrete-time mean-field-type games with a finite number of players. We quantify the exact strategic inefficiency that arises in these games when players derive their feedback strategies using classical Bellman system of equations that treat the state and/or action probability law as exogenous, termed to as “oblivious” strategies, rather than solving the true mean-field-type game. Working in a tractable deterministic setting with non-quadratic power costs, we (i) derive closed-form expressions for the feedback equilibria in both the oblivious (mean-taking) and mean-field-type (mean-accounting) formulations, (ii) compute the associated value functions, and (iii) establish a closed-form, nonnegative expression for the *Individual Price of Oblivion* or strategic myopia. Our analysis reveals a structural failure of standard equilibrium concepts in strongly coupled regimes: the cost of obliviousness is not merely an approximation error but a strategic liability that grows unbounded as the mean-field type coupling strength increases. Furthermore, we identify a “multiplier of ignorance” effect, showing that strategic competition exacerbates the performance degradation, relative inefficiency is strictly higher in competitive games than in single-agent mean-field-type control. These results prove the non-validity of standard “frozen-law” oblivious solutions for mean-field-type games and highlight the necessity of formulating equilibria using proper Master Adjoint Systems or Bellman systems on the space of measures.

I. INTRODUCTION

Earlier foundations of what is now called “mean-field games” in discrete time go back at least to 1982 and were laid out by Jovanovic and Rosenthal (1988, [1]), who considered a continuum of agents each controlling a private Markov chain, with the aggregate distribution thereby reducing a large-population game to a representative single-agent optimization against the mean field. Their 1988 model includes not only the state-mean-field but also the action-mean-field and the joint state-action mean-field. Their work has already Kolmogorov equation in it and also Bellman equation on it. These equations combined together are the key driving forces of such systems. Continuous time mean-field games came much later (see [8]–[10]). The analysis of systems subject to mean-field interactions has evolved along distinct theoretical trajectories, and it is called *Mean-Field-Type Games* [2], [3]. As we will see in this paper, MFTG has a huge difference with the freezing approach

(widely used in the literature) when the payoffs is non-linear in the probability measure of own state-action. It includes MFTG driven by McKean-Vlasov dynamics (which has the entire probability law of state/action that enters in the state/cost coefficients), MFTG driven by McKean-Vlasov dynamics of Nemytskii-type (which has Radon-Nikodym derivative (density) that enters in the state/payoff coefficients, see [13]). The case of one single agent corresponds to the centralized optimization of cost functionals depending nonlinearly on the state distribution [4]. McKean-Vlasov Forward-Backward Stochastic Differential Equations [6] and the analytic framework of Bellman equations on the Wasserstein space of probability measures [7]. See also [5], [12] for dependence on trends, delays, quantiles etc. Although the asymptotic convergence of interacting particle systems to their mean-field limits is classical [11], and the structural gap between the game-theoretic (Nash) solution and the centralized (social) optimum has been extensively studied as the *Price of Anarchy* or *Inefficiency gap* (see [12] and the references therein). The specific performance degradation incurred by incorrectly applying the standard, state-space HJB/Bellman system of equations to a true MFTG problem remains underexplored in deterministic settings.

This work isolates that precise error, extending the concept of the *Price of Oblivion* implicitly touched upon in discussions of time-inconsistency and restricted information structures [2] to provide closed-form quantification of the suboptimality inherent in ignoring the measure-derivative terms of the true dynamic programming principle in the space of measures.

For more than half a century, the Hamilton-Jacobi-Bellman (HJB) (for continuous-time problems) and Bellman (for discrete-time problems) system of equations has stood as the central analytical pillar of dynamic game theory. However, *the classical HJB/Bellman system is not universally valid for generating equilibria in mean-field-type games*. It relies, often implicitly, on a structural assumption that a player’s cost depends only on the state and the actions of specific adversaries. As soon as the cost functional depends on the *distribution*, density, quantile, variance, expectile, extremile, or *aggregate* whether through a local squared mean, a variance, or a nonlinear functional of own-state law, the classical *best response* machinery ceases to encode the correct strategic incentives. Despite the explosion of interest in Mean-Field-Type Games (MFTG), this limitation remains largely underappreciated, particularly among *communities accustomed to treating aggregate trends as exogenous signals rather than endogenous consequences of collective play/behavior*.

Why does this misunderstanding persist? Because in deter-

H. Tembine is with Department of Electrical Engineering and Computer Science, School of Engineering at Université du Québec à Trois-Rivières UQTR, Quebec, Canada, and with Timadie, Guinaga, Grabal, AI Mali, TF, WETE, MFTG, LnG Lab, CI4SI; (e-mail: tembine@ieee.org).

J. Barreiro-Gomez is with KU Center for Autonomous Robotic Systems, Department of Computer and Information Engineering, Khalifa University, Abu Dhabi 127788, UAE. (julian.barreirogomez@ku.ac.ae).

This research was funded by Khalifa University of Science and Technology through the KU Start-Up Grant Program under Project ID: KU-INT-FSU-2025-8474000774. Julian Barreiro-Gomez is profoundly grateful to The Almighty for the blessings of health and life.

ministic game settings, the “law” of the population coincides with the state trajectory itself. This creates a powerful temptation to freeze or externalize the distributional term (treating the mean field as a background weather pattern rather than a variable subject to control) and proceed with standard differential game theory. This shortcut, *what we call oblivious equilibrium* (or “mean-taking” behavior) appears innocuous. However, it breaks the game-theoretic structure in a subtle but consequential way. The resulting equilibrium is not merely inefficient; it is the solution to a fundamentally different, incorrectly simplified game.

The purpose of this paper is to make this strategic limitation manifest, quantitatively, and in the starkest possible terms. Working deliberately in a minimal deterministic game setting so that no stochastic technicalities can obscure the core argument, we:

- derive the exact oblivious solution that agents would reach under the “mean-taking” assumption that the population law is fixed and exogenous;
- derive the true Mean-Field-Type (MFT) Nash Equilibrium (“mean-accounting”) obtained when players correctly internalize their influence on the aggregate via the Master Adjoint Systems or HJB/Bellman system on the space of measures;
- compute the exact closed-form value functions associated with each equilibrium; and
- establish a fully explicit, nonnegative, closed-form formula for the strategic cost gap: the precise efficiency loss incurred by the group when they compete obliviously.

The resulting formula is remarkably transparent and straightforward. It shows that the missing sensitivity term, the derivative with respect to the measure, is not a subtle second-order correction; it is a dominant driver of strategic behavior. Even in the simplest linear-quadratic game, the gap grows with the number of players and the initial condition, and it disappears *only* in the degenerate case where the cost functional contains no mean-field dependence. *In other words, the standard oblivious Nash solution is valid in MFTG if and only if the game is not Mean-Field Type.*

This work therefore provides something that has long been missing from the literature: an exact, pedagogically minimal demonstration that classical differential game theory is structurally invalid for MFTG, even in settings that appear perfectly compatible with traditional methods. By isolating the cost penalty explicitly, we offer a diagnostic tool for game theorists: it tells precisely *when* the classical oblivious solution can be trusted (rarely), *when* it cannot (generically), and *how much* social welfare is lost when the wrong solution concept is applied.

II. DISCRETE-TIME MFTG WITH POWER COST

Consider a differential game on discrete-time infinite horizon with $k \in [0, \infty)$ and player set $\mathcal{I} = \{1, \dots, I\}$, $I \geq 2$.

The state dynamics are assumed to be linear as follows:

$$X_{k+1} = X_k + \eta \alpha_k^i + \eta \sum_{j=1, j \neq i}^I \alpha_k^j = X_k(1 - \eta S), \quad (1)$$

where η is a parameter in the dynamics, $\alpha_k^i = -\kappa_k^i X_k$, for all $i \in \mathcal{I}$, and $S = \sum_{j=1}^I \kappa_k^j$. The set of admissible control laws of $i \in \mathcal{I}$ is denoted by $\mathcal{A}^i$. We restrict our study to feedback controls of the form $\alpha_k^i = \phi^i(X_k)$. Let $\alpha := (\alpha^1, \dots, \alpha^I)$, $\alpha^{-i} := (\alpha^1, \dots, \alpha^{i-1}, \alpha^{i+1}, \dots, \alpha^I)$, and the instantaneous cost for player $i \in \mathcal{I}$ is

$$L^i(X_k, m_k, \alpha_k) = |X_k|^{2p} + \lambda^i |m_k|^{2p} + \frac{1}{2p} |\alpha_k^i|^{2p},$$

where m_k is known as the mean-field term, and $1 \leq p \in \mathbb{Z}$. Note that the system is deterministic and homogeneous, so the mean state is $m_k = X_k$. Each player $i \in \mathcal{I}$ has heterogeneous sensitivity to the mean-field term m_k given by the parameter $\lambda^i \geq 0$. Each player $i \in \mathcal{I}$ selects a feedback strategy to minimize the total cost functional

$$J^i(\alpha) = \sum_{k=0}^{\infty} L^i(X_k, m_k, \alpha_k) dt.$$

A profile of feedback strategies $\alpha^* := (\alpha^{i,*})_{i \in \mathcal{I}}$ is a Nash equilibrium if for every $i \in \mathcal{I}$,

$$J^i(\alpha^*) = \min_{\alpha^i \in \mathcal{A}^i} J^i(\alpha^i, \alpha^{-i,*})$$

We compare two concepts that differ in how players treat the mean-field term. In the oblivious setting, the i -th player treats the trajectory m_k as *exogenous*. Thus, the term $\lambda^i m_k^{2p}$ becomes an *irrelevant time-dependent constant* and the effective optimization problem reduces to considering $\lambda^i = 0$. On the contrary, when considering the mean-field term as *endogenous*, the parameters $\lambda^i > 0$. In general, player i minimizes the $2p$ -power cost functional:

$$J^i(\alpha) = \eta \sum_{k=0}^{\infty} \left((1 + \lambda^i) |X_k|^{2p} + \frac{1}{2p} |\alpha_k^i|^{2p} \right).$$

We focus our attention to the domain where the aggregated feedback satisfies the stability condition $0 < \eta S < 2$, ensuring that $|1 - \eta S| < 1$. We search for value functions of the form $V^i(X) = P^i |X|^{2p}$ with $P^i > 0$. The stationary Bellman equation for player i is as follows:

$$P^i |X|^{2p} = \min_{\alpha^i} \left\{ \eta (1 + \lambda^i) |X|^{2p} + \frac{\eta}{2p} |\alpha^i|^{2p} + P^i |(1 - \eta(S - \kappa^i))X + \eta \alpha^i|^{2p} \right\}.$$

The first-order necessary and sufficient condition for the optimal control α^i (assuming $X > 0$ –without loss of generality– due to symmetry) yields the structural relationship:

$$\begin{aligned} (\alpha_k^i)^{2p-1} + 2p P^i (X_{k+1})^{2p-1} &= 0, \\ \implies \alpha_k^i &= -(2p P^i)^{\frac{1}{2p-1}} X_{k+1}. \end{aligned} \quad (2)$$

Imposing the linear feedback form $\alpha_k^i = -\kappa^i X_k$ and noting $X_{k+1} = (1 - \eta S)X_k$, we derive the linkage between the feedback gain κ^i and the value coefficient P^i :

$$\kappa^i = (1 - \eta S)(2pP^i)^{\frac{1}{2p-1}} \iff P^i = \frac{1}{2p} \left(\frac{\kappa^i}{1 - \eta S} \right)^{2p-1}. \quad (3)$$

With this problem statement in hand and the emerging Bellman system, we introduce in the following the concept of oblivious strategy and cost, i.e., treating the state and/or action probability law as exogenous.

A. Oblivious Strategy and Oblivious Cost

The oblivious strategy corresponds to the baseline case where players ignore the specific mean-field penalty (i.e., setting $\lambda^i = 0$) and assume a symmetric environment.

Lemma 1 (Oblivious Strategy and Cost): Let $I \geq 1$ and $p \geq 1$. The discrete-time oblivious strategy is given by the feedback gain κ_{obl} which is the unique positive solution to the scalar fixed-point equation:

$$\left(\frac{\kappa}{1 - \eta I \kappa} \right)^{2p-1} \left[1 - (1 - \eta I \kappa)^{2p} \right] = 2p\eta + \eta \kappa^{2p}, \quad (4)$$

subject to the stability constraint $0 < \eta I \kappa < 1$. Under this profile

$$J_{\text{obl}}^i(x_0) = x_0^{2p} \cdot \frac{\eta \left(1 + \lambda^i + \frac{1}{2p} \kappa_{\text{obl}}^{2p} \right)}{1 - (1 - \eta I \kappa_{\text{obl}})^{2p}}. \quad (5)$$

is the oblivious cost for a player with sensitivity λ^i .

Proof: The oblivious player solves the Bellman equation assuming $\lambda^i = 0$ and a symmetric aggregate $S = I\kappa$. Substituting the linkage into the Bellman value identity:

$$P \left(1 - (1 - \eta S)^{2p} \right) = \eta(1 + 0) + \frac{\eta}{2p} \kappa^{2p}.$$

Substitute $P = \frac{1}{2p} \left(\frac{\kappa}{1 - \eta S} \right)^{2p-1}$ and $S = I\kappa$:

$$\frac{1}{2p} \left(\frac{\kappa}{1 - \eta I \kappa} \right)^{2p-1} \left(1 - (1 - \eta I \kappa)^{2p} \right) = \eta + \frac{\eta}{2p} \kappa^{2p}.$$

Multiplying by $2p$ yields (4). The LHS is strictly increasing in κ (singular at $1/\eta I$) and the RHS is strictly increasing. Since the LHS starts at 0 and goes to $+\infty$ as $\eta I \kappa \rightarrow 1$, while the RHS is finite, a unique solution κ_{obl} exists in the stable region.

Given the feedback κ_{obl} , the closed-loop decay rate is $r_{\text{obl}} = 1 - \eta I \kappa_{\text{obl}}$. The state evolves as $X_k = x_0(r_{\text{obl}})^k$. The cost functional evaluates to:

$$\begin{aligned} J_{\text{obl}}^i &= \eta \sum_{k=0}^{\infty} \left((1 + \lambda^i) + \frac{1}{2p} \kappa_{\text{obl}}^{2p} \right) (x_0 r_{\text{obl}}^k)^{2p}, \\ &= \eta \left(1 + \lambda^i + \frac{1}{2p} \kappa_{\text{obl}}^{2p} \right) x_0^{2p} \sum_{k=0}^{\infty} (r_{\text{obl}}^{2p})^k. \end{aligned}$$

The geometric series sum is $1/(1 - r_{\text{obl}}^{2p})$. Substituting this yields (5). ■

Now, we repeat the same procedure and compute the strategy and cost for the mean-field-type equilibrium. Thus, we could compare both solutions.

B. Equilibrium Strategy and Equilibrium Cost

The Mean-Field-Type (MFT) strategy accounts for the heterogeneous sensitivity λ^i and the equilibrium coupling.

Lemma 2 (MFT Equilibrium Strategy and Cost): For a parameter vector $\lambda = (\lambda^1, \dots, \lambda^I)$, the Mean-Field-Type Nash equilibrium is defined by the unique vector of gains $\kappa^* = (\kappa^{1,*}, \dots, \kappa^{I,*})$ and aggregate $S^* = \sum \kappa^{j,*}$ satisfying the coupled system for all $i \in \mathcal{I}$:

$$\left(\frac{\kappa^i}{1 - \eta S^*} \right)^{2p-1} \left[1 - (1 - \eta S^*)^{2p} \right] = 2p\eta(1 + \lambda^i) + \eta(\kappa^i)^{2p}, \quad (6)$$

and

$$J_{\text{MFT}}^i(x_0) = x_0^{2p} \cdot \frac{\eta \left(1 + \lambda^i + \frac{1}{2p} (\kappa^{i,*})^{2p} \right)}{1 - (1 - \eta S^*)^{2p}}. \quad (7)$$

is the equilibrium cost for player i given explicitly.

Proof: We utilize the Bellman equation for player i with heterogeneous λ^i . The aggregate S^* is fixed at the equilibrium value. Substituting the optimal P^i from (3) into the Bellman identity:

$$P^i \left[1 - (1 - \eta S^*)^{2p} \right] = \eta(1 + \lambda^i) + \frac{\eta}{2p} (\kappa^i)^{2p}.$$

Replacing P^i with $\frac{1}{2p} \left(\frac{\kappa^i}{1 - \eta S^*} \right)^{2p-1}$ and multiplying by $2p$:

$$\left(\frac{\kappa^i}{1 - \eta S^*} \right)^{2p-1} \left[1 - (1 - \eta S^*)^{2p} \right] = 2p\eta(1 + \lambda^i) + \eta(\kappa^i)^{2p}.$$

This defines $\kappa^{i,*}$ as an implicit function of S^* and λ^i . Summing these roots yields the fixed point equation for S^* . Existence and uniqueness follow from the monotonicity of the LHS and RHS with respect to κ and S .

Under the equilibrium strategy, the system decays at rate $r^* = 1 - \eta S^*$. The cost calculation follows the same geometric series summation as in the oblivious case, but using the specific Nash gains $\kappa^{i,*}$ and the Nash decay rate r^* .

$$J_{\text{MFT}}^i = \eta \left(1 + \lambda^i + \frac{1}{2p} (\kappa^{i,*})^{2p} \right) x_0^{2p} \frac{1}{1 - (r^*)^{2p}}.$$

This matches the stated formula (7). ■

III. PRICE OF ANARCHY AND PRICE OF SIMPLICITY

We quantify the inefficiency introduced by the lack of mean-field awareness in a strategic I -player game with $2p$ -power costs and player-specific mean-field sensitivities $\lambda^i \geq 0$. In this context, the Oblivious solution ignores the mean-field term, while the Mean-Field-Type Nash Equilibrium accounts for each λ^i individually. We retain the structural constant $\gamma_{I,p} := \frac{1}{2pI-1} \in (0, 1]$. Then, from above

$$\begin{aligned} J_{\text{MFT}}^i(x_0) &= \eta |x_0|^{2p} \frac{1 + \lambda^i + \frac{1}{2p} (\kappa^{i,*})^{2p}}{1 - (1 - \eta S^*)^{2p}}, \\ J_{\text{obl}}^i(x_0) &= \eta |x_0|^{2p} \frac{1 + \lambda^i + \frac{1}{2p} \kappa_{\text{obl}}^{2p}}{1 - (1 - \eta I \kappa_{\text{obl}})^{2p}}, \end{aligned}$$

where $\kappa^{i,*}$ and $S^* = \sum_{j=1}^I \kappa^{j,*}$ are the MFT equilibrium gains solving the coupled system and κ_{obl} is the symmetric oblivious gain.

Definition 1 (Game Price of Anarchy): The Price of Anarchy for player i is

$$\text{PoA}_{I,2p}^i := \frac{J_{\text{obl}}^i(x_0)}{J_{\text{MFT}}^i(x_0)},$$

if $J_{\text{MFT}}^i(x_0) \neq 0$ (i.e. $x_0 \neq 0$).

Definition 2 (Game Price of Simplicity): The Price of Simplicity for player i is

$$\text{PoS}_{I,2p}^i := J_{\text{obl}}^i(x_0) - J_{\text{MFT}}^i(x_0)$$

if $J_{\text{MFT}}^i(x_0) \neq 0$ (i.e. $x_0 \neq 0$).

Lemma 3 (Exact PoA and PoS): For each $i \in \mathcal{I}$, the exact expressions are

$$\begin{aligned} \text{PoA}_{I,2p}^i &= \frac{1 - (1 - \eta I \kappa_{\text{obl}})^{2p}}{1 - (1 - \eta S^*)^{2p}} \cdot \frac{1 + \lambda^i + \frac{1}{2p} \kappa_{\text{obl}}^{2p}}{1 + \lambda^i + \frac{1}{2p} (\kappa^{i,*})^{2p}}, \\ \text{PoS}_{I,2p}^i &= \eta |x_0|^{2p} \cdot \left[\frac{1 + \lambda^i + \frac{1}{2p} \kappa_{\text{obl}}^{2p}}{1 - (1 - \eta I \kappa_{\text{obl}})^{2p}} - \frac{1 + \lambda^i + \frac{1}{2p} (\kappa^{i,*})^{2p}}{1 - (1 - \eta S^*)^{2p}} \right], \end{aligned}$$

for the Price of Anarchy and the Price of Simplicity.

Lemma 4 (Asymmetric λ^i : Cost Penalty Growth): Fix $I \geq 1$, $p \geq 1$, and let at least one $\lambda^i > 0$. Then, for each player i :

$$\text{PoA}_{I,2p}^i > 1, \quad \text{PoS}_{I,2p}^i > 0.$$

Moreover, as $\lambda^i \rightarrow \infty$, the penalties grow unbounded:

$$\text{PoA}_{I,2p}^i \sim \frac{1}{1 - (1 - \eta S^*)^{2p}} (\lambda^i)^{1/(2p)}, \quad (8a)$$

$$\text{PoS}_{I,2p}^i \sim \frac{\eta |x_0|^{2p}}{1 - (1 - \eta S^*)^{2p}} \lambda^i, \quad (8b)$$

for the Price of Anarchy and the Price of Simplicity.

Proof: The mapping $\lambda^i \mapsto \kappa^{i,*}$ is strictly increasing by the monotonicity of the Bellman FOC. Increasing λ^i increases the numerator in (8) and decreases the effective denominator $(1 - \eta S^*)^{2p}$ since S^* is increasing. Therefore, both PoA and PoS are strictly increasing in λ^i . In the asymptotic regime $\lambda^i \gg 1$, the MFTG gain $\kappa^{i,*} \sim (\lambda^i)^{1/(2p)}$, leading to the indicated scaling in (8). ■

Lemma 5 (Divergence of the Price of Simplicity): Let $I \rightarrow \infty$. Assume the sensitivity parameters are uniform and scale with I , i.e. $\lambda^i(I) \equiv \Lambda_I$. If

$$\Lambda_I = \Omega\left(I^{1-\frac{1}{2p}+\varepsilon}\right), \text{ for some } \varepsilon > 0,$$

then the exact cost gap, given by

$$\Delta^i(x_0, \{\Lambda_I\}) := J_{\text{obl}}^i(x_0) - J_{\text{MFT}}^i(x_0),$$

diverges:

$$\lim_{I \rightarrow \infty} \Delta^i(x_0, \{\Lambda_I\}) = +\infty.$$

Moreover, when Λ_I dominates the aggregate dynamics,

$$\Delta^i \approx \frac{x_0^{2p}}{2p} \cdot \frac{\min_i \lambda^i(I)}{I^{1-\frac{1}{2p}}}. \quad (9)$$

Proof: Oblivious gain satisfies $\kappa_{\text{obl}} \sim I^{-1/(2p)}$, hence

$$I \kappa_{\text{obl}} \sim I^{1-\frac{1}{2p}}.$$

Thus

$$J_{\text{obl}}^i = \frac{x_0^{2p}}{2p} \frac{1 + \Lambda_I + O(I^{-1})}{I \kappa_{\text{obl}}} \approx \frac{x_0^{2p}}{2p} \frac{\Lambda_I}{I^{1-\frac{1}{2p}}}.$$

For the mean-accounting cost, with uniform Λ_I , the fixed point relation implies

$$S^* \sim I^{1-\frac{1}{2p}} (\Lambda_I)^{\frac{1}{2p}}.$$

It yields

$$J_{\text{MFT}}^i = \frac{x_0^{2p}}{2p} \frac{1 + \Lambda_I + O((\kappa^*)^{2p})}{S^*} \sim \frac{x_0^{2p}}{2p} \frac{(\Lambda_I)^{1-\frac{1}{2p}}}{I^{1-\frac{1}{2p}}}.$$

Since Λ_I grows faster than $I^{1-\frac{1}{2p}}$, the leading term in the gap is J_{obl}^i , giving the approximation (9) and the divergence conclusion. ■

We introduce a structural interpretation that isolates the manner, in which an oblivious player's failure to update beliefs about the mean field, induces a multiplicative penalty that intensifies with population size. This effect is formalized through the *Ignorance Ratio*, which measures the discrepancy between the true system coupling and the coupling presumed when computing the oblivious strategy.

Let Λ_I satisfy the super-linear growth condition $\Lambda_I \gg I^{1-\frac{1}{2p}}$. Then the exact cost gap $\Delta^i(x_0, \{\Lambda_I\})$ admits the asymptotic decomposition

$$\Delta^i(x_0, \{\Lambda_I\}) \approx J_{\text{base}}(x_0) \mathcal{M}_{\text{ign}}(I, \Lambda_I),$$

where:

- 1) $J_{\text{base}}(x_0) = \frac{x_0^{2p}}{2p}$ is the uncoupled state-cost contribution;
- 2) the multiplier

$$\mathcal{M}_{\text{ign}}(I, \Lambda_I) \approx \frac{\Lambda_I}{\text{Aggregate Effort}_{\text{obl}}} = \frac{\Lambda_I}{I^{1-\frac{1}{2p}}}$$

quantifies the amplification induced by ignoring the true mean-field sensitivity.

The *Multiplier of Ignorance* arises because the oblivious strategy breaks the structural link between the penalty parameter λ and the stabilizing control gain. In contrast, the Nash equilibrium maintains this link: higher penalty induces higher control, suppressing the state variance. By ignoring this dependence, the oblivious player exposes themselves to a multiplicative amplification of cost proportional to the discrepancy between actual coupling and assumed coupling. The following section introduces the Price of Oblivion within the game-theoretic framework. We define this metric to quantify the specific “tax” imposed on a strategic ecosystem when players collectively ignore the mean-field coupling.

IV. THE PRICE OF OBLIVION IN COMPETITIVE ENVIRONMENTS

While the Price of Anarchy measures cost ratios, the Price of Oblivion (PoO) focuses on the relative *excess burden* incurred by individual players. In the heterogeneous case, each player i has a distinct mean-field sensitivity $\lambda^i \geq 0$, and the PoO quantifies the relative inefficiency due to the collective use of oblivious strategies compared to the fully aware MFT Nash Equilibrium.

Definition 3 (Game Price of Oblivion): Let $\mathcal{G}$ be the I -player $2p$ -power game with individual couplings $\lambda^i \geq 0$. Let α^{obl} and α^{MFT} denote the joint strategy profiles at the Oblivious solution and MFT equilibria, respectively. The *Price of Oblivion* for player i is

$$\text{PoO}_{I,2p}^i(\lambda^i) := \frac{J^i(\alpha^{\text{obl}}; x_0)}{J^i(\alpha^{\text{MFT}}; x_0)} - 1 = \text{PoA}_{I,2p}^i(\lambda^i) - 1,$$

assuming $J^i(\alpha^{\text{MFT}}; x_0) \neq 0$.

The term PoO can be seen as the ‘‘Cognitive Tax.’’

- $\text{PoO}_{I,2p}^i(\lambda^i) = 0$ indicates player i suffers no additional penalty by being oblivious.
- As $\lambda^i \rightarrow \infty$, $\text{PoO}_{I,2p}^i(\lambda^i) \rightarrow \infty$, reflecting an unbounded tax on model simplicity for high mean-field sensitivity.
- Players with larger λ^i pay disproportionately more in the presence of competition, illustrating heterogeneity in strategic cost exposure.

The heterogeneous PoO formalizes the individual burden of model simplification in strategic systems:

- PoA^i gives the multiplicative inefficiency,
- PoS^i gives the absolute cost penalty,
- PoO^i gives the relative ‘‘tax’’ for using oblivious strategies.

Together, they provide a complete quantitative picture of the cost of ignoring mean-field interactions in competitive games with asymmetric player sensitivities.

V. NUMERICAL INVESTIGATION

We present a numerical study quantifying the inefficiency induced by ignoring the mean-field coupling in I -player $2p$ -power games. We focus on a range of cost exponents $p \in \{1, 2, 3, 10\}$ and player counts $I \in \{1, 2, 5, 10\}$, evaluating the divergence between the Oblivious Nash Equilibrium and the Mean-Field-Type Nash Equilibrium for varying individual coupling strengths λ^i . Fig. 1 depicts the absolute cost difference, $\text{PoS}_{I,2p}^i = J_{\text{obl}}^i - J_{\text{MFT}}^i$, across the same parameter ranges.

- The PoS grows approximately linearly with λ^i , in accordance with the asymptotic scaling predicted by the analytic formulas for heterogeneous λ^i .
- The slope of the PoS curve increases with the number of players I , illustrating that the absolute penalty for using oblivious strategies is amplified in more competitive environments.

- Nonlinearity in p subtly shapes the curvature of the PoS for intermediate λ^i , but the linear asymptotic growth dominates for strong coupling.

We can simplify the feedback gains as: Oblivious Gain $k_{\text{obl}}^{2p} = 2p\gamma_{I,p}$, and Nash Gain given by $(k^*)^{2p} = 2p\gamma_{I,p}(1 + \lambda)$. We substitute these into the two cost components: the Oblivious Cost (J_{obl}^i), and the Mean-Field-Type Nash Cost (J_{MFT}^i). For the Oblivious Cost (J_{obl}^i), the oblivious player uses fixed gain k_{obl} and the term $\frac{1}{2p}k_{\text{obl}}^{2p}$ simplifies to $\gamma_{I,p}$. Then,

$$J_{\text{obl}}^i = \frac{x_0^{2p}}{2p} \left[\frac{1 + \lambda + \gamma_{I,p}}{Ik_{\text{obl}}} \right]$$

For the Mean-Field-Type Nash Cost (J_{MFT}^i), The MFT Nash player adapts $k^* = k_{\text{obl}}(1 + \lambda)^{\frac{1}{2p}}$. The numerator term $\frac{1}{2p}(k^*)^{2p}$ simplifies to $\gamma_{I,p}(1 + \lambda)$.

$$\text{Num} = 1 + \lambda + \gamma_{I,p}(1 + \lambda) = (1 + \lambda)(1 + \gamma_{I,p}),$$

$$\begin{aligned} J_{\text{MFT}}^i &= \frac{x_0^{2p}}{2p} \left[\frac{(1 + \lambda)(1 + \gamma_{I,p})}{Ik_{\text{obl}}(1 + \lambda)^{\frac{1}{2p}}} \right], \\ &= \frac{x_0^{2p}}{2p} \left[\frac{(1 + \gamma_{I,p})}{Ik_{\text{obl}}} (1 + \lambda)^{1 - \frac{1}{2p}} \right]. \end{aligned}$$

Subtracting J_{MFT}^i from J_{obl}^i yields the explicit form:

$$\begin{aligned} \text{PoS}_{I,2p}^i(\lambda) &= \frac{\overbrace{x_0^{2p}}^{\text{System scaling factor}}}{2pI \left(\frac{2p}{2pI-1} \right)^{\frac{1}{2p}}} \\ &\quad \left[\underbrace{\left(1 + \lambda + \frac{1}{2pI-1} \right)}_{\text{Linear Oblivious Cost}} - \underbrace{\left(1 + \frac{1}{2pI-1} \right) (1 + \lambda)^{1 - \frac{1}{2p}}}_{\text{Sub-linear Nash Cost}} \right]. \end{aligned}$$

We have that the *interaction strength that grows with the number of players*. Let the interaction parameter scale linearly with population size, $\lambda^i(I) = I$. Then, for any exponent $p \geq 1$, the individual Price of Simplicity diverges, i.e., $\lim_{I \rightarrow \infty} \text{PoS}_{I,2p}^i(I) = \infty$. A convenient representation is

$$\begin{aligned} \text{PoS}_{I,2p}^i(I) &= \frac{x_0^{2p}}{2p} \left[\frac{1 + I + \gamma_{I,p}}{I(2p\gamma_{I,p})^{\frac{1}{2p}}} - \frac{(1 + \gamma_{I,p})(1 + I)^{1 - \frac{1}{2p}}}{I(2p\gamma_{I,p})^{\frac{1}{2p}}} \right]. \end{aligned}$$

Taken together, the numerical results in Fig. 1 validate and extend the theoretical findings:

- 1) In strongly coupled games, the cost of ignoring the endogenous mean field is both multiplicatively and absolutely unbounded as $\lambda^i \rightarrow \infty$.
- 2) Competition among players ($I > 1$) exacerbates the inefficiency, effectively multiplying the cost of obliviousness relative to single-agent scenarios.
- 3) The nonlinearity parameter p affects intermediate regimes, but the qualitative divergence of PoA and linear growth of PoS are robust across all p considered.
- 4) From a control or learning perspective, oblivious strategies are fundamentally inadequate in high-interaction




Fig. 1. (a) The shape of the price of simplicity vs interaction strength, for different values of I and for different power exponent p . (b) The shape of the price of simplicity of all players vs interaction strength. (c) The shape of the price of simplicity when $\lambda^i(I)$ scales to infinity with I .

regimes: the combination of strong coupling and multi-agent competition leads to catastrophic deviations from the MFT benchmark.

These findings quantitatively formalize the “cost of freezing the mean field”, i.e., neglecting the strategic feedback in high-interaction environments is not only theoretically sub-optimal but practically intolerable for system performance.

VI. CONCLUSION AND FUTURE WORK

In this paper, we have established a definitive quantitative boundary between classical differential game theory and MFTG. By deriving exact, closed-form solutions in a minimal deterministic setting, we demonstrated that the “oblivious” derivation of Nash equilibria where players treat the population law or own-state-action law as an exogenous signal, is not merely a heuristic approximation, but a fundamental strategic error. Our analysis of the *Individual Price of Oblivion* proves that this inefficiency is unbounded. The performance penalty incurred by the oblivious coalition grows arbitrarily large as the mean-field coupling λ increases. Furthermore, we identified a “multiplier of ignorance”: the relative cost of being oblivious is strictly higher in competitive environments ($I > 1$) than in single-agent settings. This result serves as a critical diagnostic for game theorists: while the classical “frozen-law” Nash concept may suffice for weakly coupled systems ($\lambda \approx 0$), it fails catastrophically in regimes of strong interaction, where the correct equilibrium concept relying on the Master Adjoint Systems or HJB/Bellman systems on the space of measures is strictly required. Future research should extend these explicit quantifications to stochastic regimes, particularly those involving common noise and mixed atomic/non-atomic populations, where the governing equation becomes a stochastic Master Adjoint System. There is significant value in analyzing the *Price of Oblivion* in Strategic Learning contexts, where independent learners that fail to track the evolving population distribution effectively converge to sub-optimal oblivious equilibria. Further work is needed to map these theoretical gaps onto scenarios involving heterogeneous agents and partial information, where the strategic tax of model simplicity may be even more pronounced.

REFERENCES

- [1] B. Jovanovic and R. W. Rosenthal, “Anonymous sequential games,” *Journal of Mathematical Economics*, vol. 17, no. 1, pp. 77–87, 1988.
- [2] T. Başar, B. Djehiche, H. Tembine: *Mean-Field-Type Game Theory I-II*, Birkhäuser Cham, 2026.
- [3] J. Barreiro-Gomez and H. Tembine: *Mean-Field-Type Games for Engineers*, CRC Press Taylor & Francis Group, 2021.
- [4] D. Andersson and B. Djehiche, “A maximum principle for SDEs of mean-field type,” *Applied Mathematics and Optimization*, vol. 63, no. 3, pp. 341–356, 2011.
- [5] A. Bensoussan, J. Frehse, and P. Yam, *Mean Field Games and Mean Field Type Control Theory*, Springer, New York, 2013. ISBN: 978-1-4614-8507-0 (print), 978-1-4614-8508-7 (eBook).
- [6] R. Carmona and F. Delarue, *Probabilistic Theory of Mean Field Games with Applications I-II*, Springer, Cham, 2018. (Volumes I & II). ISBN (Vol I): 978-3-319-58920-6; ISBN (Vol II): 978-3-319-56435-7.
- [7] W. Gangbo and A. Święch, “Existence of a solution to an equation arising from the theory of Mean Field Games,” *Journal of Differential Equations*, vol. 259, no. 11, pp. 6573–6643, 2015.
- [8] M. Huang, P. E. Caines, and R. P. Malhamé, “Large population stochastic dynamic games: closed-loop McKean–Vlasov systems and the Nash certainty equivalence principle,” *Communications in Information & Systems*, vol. 6, no. 3, pp. 221–252, 2006.
- [9] J. M. Lasry and P. L. Lions, “Mean field games,” *Japanese Journal of Mathematics*, vol. 2, no. 1, pp. 229–260, 2007.
- [10] P. L. Lions, “Cours au Collège de France: Théorie des jeux à champ moyen,” audio lectures, 2007–2011.
- [11] A. S. Sznitman, “Topics in propagation of chaos,” in *École d’Été de Probabilités de Saint-Flour XIX — 1989*, Lecture Notes in Mathematics, vol. 1464, Springer, Berlin/Heidelberg, 1991, pp. 165–251.
- [12] B. Djehiche, A. Tcheukam, and H. Tembine, “Mean-Field-Type Games in Engineering,” *AIMS Electronics and Electrical Engineering*, vol. 1, no. 1, pp. 18–73, 2017.
- [13] S. Grube, “Strong solutions to McKean–Vlasov SDEs with coefficients of Nemytskii type: the time-dependent case,” *Journal of Evolution Equations*, vol. 24, article no. 37, 2024.
- [14] J. F. Nash, *Equilibrium Points in N-Person Games*. Proceedings of the National Academy of Sciences, 36(1), 48–49, 1950.
- [15] J. F. Nash, *Non-Cooperative Games*. Annals of Mathematics, 54(2), 286–295, 1951.
- [16] A. Rapoport and A. M. Chammah. *Prisoner’s Dilemma: A Study in Conflict and Cooperation*. Ann Arbor: University of Michigan Press, 1965.
- [17] A. C. Pigou. *The Economics of Welfare*. London: Macmillan, 1920.
- [18] J. Maynard Smith and G. R. Price. *The Logic of Animal Conflict*. Nature, 246, 15–18, 1973.
- [19] D. Fudenberg and E. Maskin. *The Folk Theorem in Repeated Games with Discounting*. Econometrica, 54(3), 533–554, 1986.
- [20] P. Dubey. *Inefficiency of Nash Equilibria*. IIASA Working Paper WP-83-074, 1983.