The Price of Oblivion: On the Non-Validity of Classical Hamilton-Jacobi-Bellman in Mean-Field-Type Control*

Hamidou Tembine[1], Julian Barreiro-Gomez[2], Tyrone E. Duncan[3], Bozenna Pasik-Duncan[3]

Abstract— **We quantify the exact optimal cost gap that arises when a controller solves the classical Hamilton-Jacobi-Bellman (HJB) equation while treating the population law as exogenous (by freezing it), an approach we call oblivious control, instead of solving the true mean-field-type control problem in which the running cost depends nonlinearly on the state law. Working in a simple but fully representative deterministic setting, we (i) derive closed-form expressions for both the oblivious and the mean-field-type optimal feedback laws, (ii) compute the corresponding optimal costs, and (iii) establish a closed-form, nonnegative expression for the cost gap. The formula highlights precisely how the law-dependence controls the suboptimality of oblivious HJB control. It also proves non-validity of the standard HJB in mean-field-type control which requires a new HJB in the space of measures.**

I. Introduction

For more than half a century, the Hamilton-Jacobi-Bellman (HJB) equation has stood as the central analytical pillar of optimal control. From aerospace guidance laws to power systems and robotics, generations of engineers have applied the classical HJB framework with remarkable success. Its promise is seductive: once the value function satisfies the HJB partial differential equation, optimality is certified, uniqueness is guaranteed under broad conditions, and the resulting feedback is provably optimal. It is therefore not surprising that, when confronted with modern systems, practitioners instinctively reach for the same tool.

Yet this very confidence conceals a fundamental misconception. *The classical HJB equation in local state is not universally valid.* It relies, often implicitly, on a structural assumption that the running cost and dynamics depend on the state and control of a single agent only. As soon as the cost functional depends on the *distribution* of own-state whether through a mean, a variance, a density, or any nonlinear functional of the law, the classical HJB ceases to encode the correct dynamic programming principle. Despite decades of progress in mean-field-type games (MFTG) and mean-field-type control (MFTC), this limitation remains largely under-appreciated in engineering communities, especially among those accustomed to stochastic-process-based reasoning but working in deterministic regimes.

Why does this misunderstanding persist? Because in deterministic settings, the "law" of the state coincides with the state itself for homogeneous populations, tempting one to freeze or externalize the distributional term and proceed with the standard HJB machinery. This shortcut what we call *oblivious control* appears innocuous, and in many engineering applications has been applied tacitly for years. However, it breaks the dynamic programming structure in a subtle but consequential way. The resulting controller is not merely suboptimal; it is the solution to a different, incorrectly simplified control problem.

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

- derive the exact classical-HJB-based feedback that engineers would compute under the oblivious assumption that the population law is fixed (frozen) and exogenous;
- derive the true mean-field-type optimal feedback obtained from the correct HJB on the space of measures;
- compute the exact closed-form optimal costs associated with each feedback law; and
- establish a fully explicit, nonnegative, closed-form formula for the cost gap: the precise penalty incurred by applying the classical HJB when the cost depends nonlinearly on the population law.

The resulting formula is remarkably transparent and straightforwardly obtained. It shows that the missing measure-derivative term in the classical HJB is not a subtle second-order correction it is the dominant contributor to performance. Even in the simplest linear-quadratic deterministic system, the cost gap grows quadratically with the initial condition and monotonically with the law-coupling parameter, and it disappears *only* in the degenerate case where the cost functional contains no law dependence at all. *In other words, the classical HJB is valid in mean-field-type control if and only if the problem is not mean-field-type.*

This work therefore provides something that has long been missing from the literature: an exact, closed-form, pedagogically minimal demonstration that classical HJB is structurally invalid for MFTC, even in deterministic settings that appear perfectly compatible with the traditional theory. By isolating the cost penalty explicitly, we offer a diagnostic tool for practitioners: it tells precisely *when* the classical HJB can be trusted (rarely), *when* it cannot (generically), and *how*

*This work was supported by the U.S. Air Force Office of Scientific Research under Grant FA9550-17-1-0259

[1] H. Tembine is with Department of Electrical Engineering and Computer Science, School of Engineering, Quebec University in Trois-Rivieres UQTR, Quebec, Canada, and with Timadie, Guinaga, Grabal, AI Mali, TF, WETE, MFTG, LnG Lab, CI4SI; (e-mail: tembine@ieee.org).

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

[3] Department of Mathematics, University of Kansas, Lawrence, 66045, KS, USA

much performance is lost when the wrong equation is used (see Table 1).

Our hope is that this contribution fills a critical conceptual gap. Engineers have been applying HJB for decades, often with great success, but the boundary of its validity has remained blurred. This paper sharpens that boundary to a knife edge and demonstrates, in closed form, why the correct equation for mean-field-type control must live on the space of probability measures.

To rigorously quantify the consequences of this structural oversight, we introduce the **Price of Oblivion**. We define this metric as the exact performance penalty incurred by a controller that applies the **oblivious feedback** derived from the classical, frozen-law HJB equation within the true, mean-field-coupled system. It represents the cost of ignoring the feedback loop between an agent's actions and the population state. Far from being a negligible approximation error, the *Price of Oblivion* serves as a measure of systemic fragility: we prove that it scales nonlinearly with the strength of the mean-field interaction and in the nonzero initial state magnitude and, crucially, is unbounded. This demonstrates that as the coupling intensifies, the convenience of using the standard HJB equation is purchased at the expense of catastrophic inefficiency, rendering the oblivious approach structurally invalid for strongly interacting systems.

TABLE I

Comparison between oblivious and MFTC Bellman equations across continuous-time (CT) / discrete-time (DT) models, for both quadratic and non-quadratic $2p$-power cost structures

Setting	Oblivious (Non-valid Bellman)	MFTC (Valid Bellman)
CT: Quadratic (LQ)	Uses frozen mean; solves incorrect HJB; underestimates coupling	Solves full MFTC HJB; correct coupling, optimal feedback
CT: Power Cost ($2p$)	Applies CT HJB with frozen mean; wrong growth $\sim k^{2p}$	Correct $2p$-HJB; proper mean-field drift; larger gain
DT: Quadratic (LQ)	Frozen mean in Bellman; Riccati eq. gives P_{obl}	True DT MFTC Bellman; Riccati eq. gives P_{MFTC} with $P_{\mathrm{MFTC}} > P_{\mathrm{obl}}$
DT: Power Cost ($2p$)	Wrong fixed-point for P; smaller κ_{obl}	Correct fixed-point; larger κ_{MFTC}; smaller cost
Cost Gap	Oblivious cost $>$ MFTC cost for all $\lambda > 0$	

Literature Review: 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], [9], [7]). The global optimization analysis of such systems subject to mean-field interactions has evolved along distinct theoretical trajectories, and it is called *Mean-Field-Type Control*. It includes McKean-Vlasov control (which has the entire probability law of state/action that enters in the state/cost coefficients), McKean-Vlasov dynamics of Nemytskii-type control (which has Radon-Nikodym derivative (density) that enters in the state/payoff coefficients, see [12]), which addresses the centralized optimization of cost functionals depending nonlinearly on the state distribution [3]. McKean-Vlasov Forward-Backward Stochastic Differential Equations [5] and the analytic framework of Bellman equations on the Wasserstein space of probability measures [6]. See also [11], [4] for dependence on trends, delays, quantiles etc. Although the asymptotic convergence of interacting particle systems to their mean-field limits is classical [10], and the structural gap between the game-theoretic fully non-cooperative solution and the centralized fully cooperative (social) optimum has been extensively studied as the *Price of Anarchy*, *Inefficiency gap* (see [11] and the references therein), the specific performance degradation incurred by incorrectly applying the standard, state-space HJB equation to a true MFTC 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.

Structure: The rest of the paper is structured as follows. The next section begins by establishing the structural invalidity of the classical Hamilton-Jacobi-Bellman equation when applied to mean-field-type control problems, identifying this error as "oblivious strategy" where the population law is incorrectly treated as exogenous. Proceeding through a deterministic framework, we derive and compare exact closed-form feedback laws and costs for both the oblivious and true MFTG solutions, spanning continuous-time linear-quadratic models, generalized $2p$-power cost structures, and their discrete-time equivalents. This analysis continue with quantification of the suboptimality via the **Price of Anarchy** , **Price of Simplicity,** and **Price of Oblivion,** where we prove analytically and illustrate numerically that the cost gap grows without bound as the mean-field coupling parameter λ increases for any nonzero initial state x_0. This confirms that freezing the mean-field term leads to catastrophic performance degradation rather than a mere approximation error.

II. How Bad is Oblivious Control?

This section examines the dangers of oblivious control (interventions designed without accounting for the strategic, adaptive, and incentive-driven behavior of local actors in the system). When policymakers, security forces, or external partners apply control actions based on simplified heuristics or outdated models, they risk amplifying the very dynamics they intend to suppress. We illustrate it here in an explicit way.

We extend to a nonlinear homogeneous-power case and derive explicit result for the optimal cost gap. Let

- X_t be the state, with dynamics

$$\dot{X}_t = \alpha_t, \quad X_0 = x_0,$$

- $m_t = \mathbb{E}[X_t]$ be the population mean. We work in the homogeneous deterministic specialization, so $X_t = m_t$,
- $p \in \mathbb{Z}_{>0}$ be the power exponent (we write powers as X^{2p}, α^{2p}),
- $\lambda > 0$ be the law-coupling parameter,
- running cost

$$L(X_t, m_t, \alpha_t) = X_t^{2p} + \lambda m_t^{2p} + \frac{1}{2p}\,\alpha_t^{2p}.$$

We compare

1) the *oblivious feedback* obtained by solving the classical single-agent HJB treating m_t as exogenous, and
2) the *mean-field-type optimal feedback* obtained by solving the centralized MFTC problem (equivalently the HJB with running state-cost $X^{2p} + \lambda m^{2p}$).

Below we derive both feedback gains in closed-form, compute the exact costs, and prove a closed-form, nonnegative expression for the cost gap. The gap is zero if and only if $\lambda = 0$.

Theorem 1. *(Cost Gap for $2p$-Power Costs)*
Let $p \geq 1$ and $\lambda > 0$. Consider dynamics

$$\dot{X} = \alpha, \quad X_0 = x_0$$

and running cost

$$L(X, m, \alpha) = X^{2p} + \lambda m^{2p} + \frac{1}{2p}\,\alpha^{2p}.$$

Then the optimal costs are

$$J_{\mathrm{MFTC}}(x_0) = \frac{x_0^{2p}}{2p-1}\,\frac{(1+\lambda)^{1-\frac{1}{2p}}}{\left(\frac{2p}{2p-1}\right)^{\frac{1}{2p}}},$$

$$J_{\mathrm{obl}}(x_0) = \frac{x_0^{2p}}{2p-1}\,\frac{1+\frac{2p-1}{2p}\lambda}{\left(\frac{2p}{2p-1}\right)^{\frac{1}{2p}}}.$$

Hence, the cost gap is

$$\begin{aligned}
&\Delta(x_0, \lambda) = J_{\mathrm{obl}} - J_{\mathrm{MFTC}} \\
&= \tfrac{x_0^{2p}}{2p-1}\left(\tfrac{2p}{2p-1}\right)^{-\frac{1}{2p}}\left[1 + \tfrac{2p-1}{2p}\lambda - (1+\lambda)^{1-\frac{1}{2p}}\right] \\
&\geq 0.
\end{aligned}$$

We observe that by setting $p = 1$ recovers the previously derived quadratic result. The formulas reduce correctly: $k_{\mathrm{MFTC}} = \sqrt{2(1+\lambda)}$, $k_{\mathrm{obl}} = \sqrt{2}$, and the bracket in Δ coincides. The cost gap arises because the oblivious HJB treats m as exogenous. The true MFTC sees the law term, leading to a larger stabilizing gain k_{MFTC} and strictly smaller total cost.

The correct dynamic programming principle for MFTC involves a HJB equation on the space of probability measures; the extra terms in that equation are exactly those that vanish in the oblivious HJB but are responsible for the above cost gap.

III. Price of Anarchy and Price of Simplicity

A. *Prices of Anarchy and Simplicity for the $2p$-power cost*

We provide fully explicit formulas and detailed proofs for two measures of suboptimality when the classical (oblivious) HJB is used instead of the correct mean-field-type HJB, in the continuous-time homogeneous $2p$-power setting. We show both measures are *unbounded* as the law coupling $\lambda \to \infty$. Keep the notation and closed-form costs from the $2p$ analysis:

$$\begin{aligned}
J_{\mathrm{MFTC}}(x_0) &= \tfrac{x_0^{2p}}{2p-1}\,\tfrac{(1+\lambda)^{1-\frac{1}{2p}}}{\left(\frac{2p}{2p-1}\right)^{\frac{1}{2p}}}, \\
J_{\mathrm{obl}}(x_0) &= \tfrac{x_0^{2p}}{2p-1}\,\tfrac{1+\frac{2p-1}{2p}\lambda}{\left(\frac{2p}{2p-1}\right)^{\frac{1}{2p}}}.
\end{aligned}$$

The common multiplicative factor

$$C_p := \frac{x_0^{2p}}{2p-1}\left(\frac{2p}{2p-1}\right)^{-1/(2p)}$$

appears in both expressions; writing the costs in compact form,

$$\begin{aligned}
&J_{\mathrm{MFTC}}(x_0) = C_p\,(1+\lambda)^{1-\frac{1}{2p}}, \\
&J_{\mathrm{obl}}(x_0) = C_p\Big(1 + \tfrac{2p-1}{2p}\lambda\Big).
\end{aligned}$$

Definition 1 (Price of Anarchy: PoA)**.** *For fixed $p \geq 1$ and initial state $x_0 \neq 0$ define the* Price of Anarchy

$$\mathrm{PoA}_{2p}(\lambda; x_0) := \frac{J_{\mathrm{obl}}(x_0)}{J_{\mathrm{MFTC}}(x_0)}.$$

if $J_{\mathrm{MFTC}}(x_0) > 0$.

Definition 2 (Price of Simplicity: PoS)**.** *Define the* Price of Simplicity *(additive cost penalty)*

$$\mathrm{PoS}_{2p}(\lambda; x_0) := J_{\mathrm{obl}}(x_0) - J_{\mathrm{MFTC}}(x_0).$$

Both quantities depend parametrically on p, λ, and x_0.

Using the compact representations above we obtain immediate, exact formulas.

Proposition 1 (Exact PoA and PoS)**.** *For $p \geq 1$, $\lambda \geq 0$, and $x_0 \neq 0$,*

$$\mathrm{PoA}_{2p}(\lambda; x_0) = \frac{1 + \frac{2p-1}{2p}\lambda}{(1+\lambda)^{1-\frac{1}{2p}}}, \tag{1}$$

$$\mathrm{PoS}_{2p}(\lambda; x_0) = C_p\left(1 + \frac{2p-1}{2p}\lambda - (1+\lambda)^{1-\frac{1}{2p}}\right). \tag{2}$$

Proof Divide the two closed-form costs to obtain (1). Subtract them to obtain (2). The common factor C_p cancels in the ratio and factors out of the difference. □

We now prove the two measures grow without bound as $\lambda \to \infty$ and give precise asymptotic rates.

Theorem 2 (PoA and PoS are unbounded in λ)**.** *Fix $p \geq 1$ and $x_0 \neq 0$. Then*

$$\lim_{\lambda\to\infty} \mathrm{PoA}_{2p}(\lambda; x_0) = \infty, \qquad \lim_{\lambda\to\infty} \mathrm{PoS}_{2p}(\lambda; x_0) = \infty.$$

More precisely,

$$\mathrm{PoA}_{2p}(\lambda;x_0) \sim \frac{2p-1}{2p}\,\lambda^{\frac{1}{2p}} \quad (\lambda\to\infty),$$

and

$$\mathrm{PoS}_{2p}(\lambda;x_0) \sim C_p\cdot\frac{2p-1}{2p}\,\lambda \quad (\lambda\to\infty).$$

Proof. From (1),

$$\mathrm{PoA}_{2p}(\lambda;x_0) = \frac{1+\frac{2p-1}{2p}\lambda}{(1+\lambda)^{1-\frac{1}{2p}}} = \lambda^{\frac{1}{2p}}\cdot\frac{\frac{2p-1}{2p}+\frac{1}{\lambda}}{\left(1+\frac{1}{\lambda}\right)^{1-\frac{1}{2p}}}.$$

As $\lambda\to\infty$ the multiplicative fraction tends to $(2p-1)/(2p)$, hence

$$\mathrm{PoA}_{2p}(\lambda;x_0) \sim \frac{2p-1}{2p}\,\lambda^{1/(2p)} \to \infty.$$

Thus PoA_{2p} grows polynomially with exponent $1/(2p)>0$, so it is unbounded.

From (2) we write

$$\mathrm{PoS}_{2p}(\lambda;x_0) = C_p\left(\frac{2p-1}{2p}\lambda - \left[(1+\lambda)^{1-\frac{1}{2p}}-1\right]\right).$$

For large λ we have the expansion $(1+\lambda)^{1-\frac{1}{2p}} = \lambda^{1-\frac{1}{2p}}(1+o(1))$, and because $1-\frac{1}{2p}<1$ the term $\lambda^{1-\frac{1}{2p}}$ grows strictly slower than λ. Therefore the linear term dominates and

$$\mathrm{PoS}_{2p}(\lambda;x_0) \sim C_p\cdot\frac{2p-1}{2p}\,\lambda\to\infty.$$

This proves the stated asymptotics and unboundedness for PoS. □

We also include a short proof that the bracket in (2) is nonnegative for all $\lambda\ge 0$ (so PoS is nonnegative).

Lemma 1. *For $p\ge 1$ and all $\lambda\ge 0$ the scalar function*

$$g(\lambda) := 1+\frac{2p-1}{2p}\lambda-(1+\lambda)^{1-\frac{1}{2p}}$$

satisfies $g(\lambda)\ge 0$ and $g(\lambda)=0$ iff $\lambda=0$.

Proof Set $s:=1-\frac{1}{2p}\in\left[\frac{1}{2},1\right)$. Then $g(\lambda)=1+(2s-1)\frac{\lambda}{2s}-(1+\lambda)^s$; equivalently write

$$g(\lambda)=1+s\lambda-(1+\lambda)^s-\frac{\lambda}{2p},$$

but a direct convexity argument is most transparent: consider the concave function $\phi(z):=z^s$ on $z\ge 0$ (since $0<s<1$, ϕ is concave). By the tangent-line (supporting hyperplane) inequality for concave functions,

$$(1+\lambda)^s \le 1+s\lambda, \qquad \forall\lambda\ge 0.$$

Rearranging gives $1+s\lambda-(1+\lambda)^s\ge 0$. Using $s=1-\frac{1}{2p}$ we obtain

$$1+\frac{2p-1}{2p}\lambda-(1+\lambda)^{1-\frac{1}{2p}}\ge 0,$$

and equality holds only at $\lambda=0$ because ϕ is strictly concave on $(0,\infty)$. □

For any fixed finite λ both $J_{\mathrm{MFTC}}(x_0)$ and $J_{\mathrm{obl}}(x_0)$ are finite. The statements "PoA is unbounded" and "PoS is unbounded" mean that, when the family of problems is indexed by $\lambda\ge 0$, the suboptimality of the oblivious controller can be made arbitrarily large by choosing λ sufficiently large. In particular:

- For any target $M>0$ there exists $\Lambda>0$ such that for all $\lambda\ge\Lambda$ one has $\mathrm{PoA}(\lambda;x_0)>M$ and $\mathrm{PoS}(\lambda;x_0)>M$.
- Thus the *cost of approximation* (*using the incorrect HJB*) is not uniformly bounded over the family of mean-field couplings; the approximation error can be made arbitrarily large, and in that precise sense the *price* of using the oblivious HJB is "infinite" as a uniform bound.

The discrete-time formulas derived converge to the continuous-time expressions above as the step-size $\eta\downarrow 0$. Hence the same PoA and PoS definitions in discrete time display identical unboundedness in λ (for any fixed, sufficiently small η) and the same asymptotic scalings hold. □

B. Price of Oblivion

Definition 3 (Price of Oblivion)**.** *Let $\mathcal{U}$ be the space of admissible feedback controls, and let $J_{\mathrm{MFTC}}(\alpha;x_0,\mu_0)$ denote the cost functional of the true mean-field-type control problem given initial state x_0 and initial measure μ_0. Let $\alpha^{\mathrm{obl}}\in\mathcal{U}$ be the feedback strategy derived from the classical HJB equation where the measure flow is frozen (treated as exogenous). Let $\alpha^\star=\alpha_{MFTC}\in\mathcal{U}$ be the optimal mean-field-type strategy.*

The Price of Oblivion *(PoO) is defined as the relative performance degradation incurred by the oblivious strategy compared to the optimal strategy:*

$$\mathrm{PoO}(\lambda;x_0) := \frac{J_{\mathrm{MFTC}}(\alpha^{\mathrm{obl}};x_0,\mu_0)}{J_{\mathrm{MFTC}}(\alpha^\star;x_0,\mu_0)}-1.$$

whenever $J_{\mathrm{MFTC}}(\alpha^\star;x_0,\mu_0)\neq 0$. In the specific context of the homogeneous problems analyzed in this work (where $\mathrm{PoA}=J_{\mathrm{obl}}/J_{\mathrm{MFTC}}$), the Price of Oblivion coincides with the surplus ratio:

$$\mathrm{PoO}(\lambda;x_0)=\mathrm{PoA}(\lambda;x_0)-1.$$

While the Price of Simplicity measures the absolute monetary or energetic loss, the Price of Oblivion measures the relative inefficiency of the control design. A PoO of zero implies the oblivious approximation is exact; a PoO approaching infinity implies the oblivious controller is asymptotically worthless.

Tables II and III illustrate the three notions.

TABLE II

COMPARISON OF EXACT CLOSED-FORM FORMULAS FOR OBLIVION METRICS ($2p$-POWER COST). $C_p = \frac{x_0^{2p}}{2p-1}\left(\frac{2p}{2p-1}\right)^{-1/(2p)}$.

Attribute	**PoA (Anarchy)**	**PoS (Simplicity)**	**PoO (Oblivion)**
Definition	$\frac{J_{\mathrm{obl}}}{J_{\mathrm{MFTC}}}$	$J_{\mathrm{obl}} - J_{\mathrm{MFTC}}$	$\frac{J_{\mathrm{obl}}}{J_{\mathrm{MFTC}}} - 1$
Interpretation	Multiplicative inefficiency factor	Absolute cost penalty (Additive loss)	Relative percentage waste ("Tax" rate)
Exact Expression (for $\lambda \geq 0$)	$\frac{1+\frac{2p-1}{2p}\lambda}{(1+\lambda)^{1-\frac{1}{2p}}}$	$C_p(1 + \frac{2p-1}{2p}\lambda - (1+\lambda)^{1-\frac{1}{2p}})$	$\frac{1+\frac{2p-1}{2p}\lambda}{(1+\lambda)^{1-\frac{1}{2p}}} - 1$

TABLE III

EXACT FORMS AND ASYMPTOTIC EQUIVALENTS FOR PRICE OF ANARCHY AND PRICE OF SIMPLICITY IN THE CONTINUOUS-TIME $2p$-POWER MFTC PROBLEM. HERE $C_p = \frac{x_0^{2p}}{2p-1}\left(\frac{2p}{2p-1}\right)^{-1/(2p)}$.

Quantity	**Exact expression**	**Asymptotic equivalent as $\lambda \to \infty$**
$\mathrm{PoA}(\lambda; x_0)$	$\frac{1+\frac{2p-1}{2p}\lambda}{(1+\lambda)^{1-\frac{1}{2p}}}$	$\mathrm{PoA}(\lambda; x_0) \sim \frac{2p-1}{2p}\lambda^{\frac{1}{2p}} \quad (\lambda \to \infty)$
$\mathrm{PoO}(\lambda; x_0)$	$\frac{1+\frac{2p-1}{2p}\lambda}{(1+\lambda)^{1-\frac{1}{2p}}} - 1$	$\mathrm{PoO}(\lambda; x_0) \sim \frac{2p-1}{2p}\lambda^{\frac{1}{2p}} \quad (\lambda \to \infty)$
$\mathrm{PoS}(\lambda; x_0)$	$C_p\left(1 + \frac{2p-1}{2p}\lambda - (1+\lambda)^{1-\frac{1}{2p}}\right)$	$\mathrm{PoS}(\lambda; x_0) \sim C_p \cdot \frac{2p-1}{2p}\lambda \quad (\lambda \to \infty)$
Cost of freezing the mean-field term	**Unbounded in λ:** both multiplicative loss PoA, relative gap PoO and additive loss PoS tend to $+\infty$ as $\lambda \to \infty$. Equivalently, for any $M > 0$ there exists Λ with $\lambda \geq \Lambda$ implies $\mathrm{PoA}(\lambda; x_0) > M$ and $\mathrm{PoS}(\lambda; x_0) > M$. For $x_0 \neq 0$, PoS tend to $+\infty$ for a fixed $\lambda > 0$ as $\lvert x_0\rvert \to +\infty$.	

IV. NUMERICAL INVESTIGATION

The two figures quantify, for a wide range of cost exponents $p \in \{1, 2, 3, 10\}$, the divergence between (i) the *oblivious* solution obtained by solving a frozen-mean HJB equation and (ii) the *true* optimal solution delivered by the consistent mean-field-type control Bellman/HJB equation. The first plot displays the *Price of Anarchy* , i.e. the ratio between the oblivious and optimal MFTC value functions, while the second plot displays the *Price of Simplicity*, i.e. their absolute difference. In both cases the horizontal axis represents the interaction strength λ, which measures the magnitude of the mean-field coupling in the dynamics.

The curves show that, for each fixed exponent p, the PoA grows monotonically and becomes asymptotically linear (on a logarithmic scale) as $\lambda \to \infty$. This reflects the fact that the frozen-mean approximation fundamentally underestimates the true mean-field drift, leading to systematically suboptimal control gains and a rapidly deteriorating closed-loop contraction. The second plot highlights an even more dramatic phenomenon: the PoS diverges superlinearly with λ (with a rate increasing in p), indicating that the error in value caused by solving an incorrect HJB is not merely large, but *unbounded*.

Taken together, these figures formalize a key structural insight: in strongly coupled regimes, the approximation error induced by ignoring the endogenous mean field is not perturbative but catastrophic: its cost grows without bound as λ increases. In this precise sense, the "cost of freezing the mean field" is infinite, and the frozen-mean HJB cannot be used as an approximation scheme for large-interaction MFTC systems.

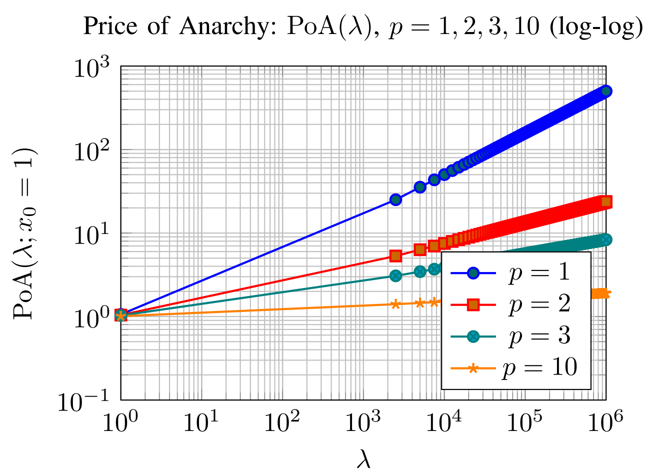

Fig. 1. Log-log plot of $\mathrm{PoA}(\lambda; x_0 = 1)$ vs λ. For large λ curves behave as $\mathrm{PoA} \sim \frac{2p-1}{2p}\lambda^{1/(2p)}$.

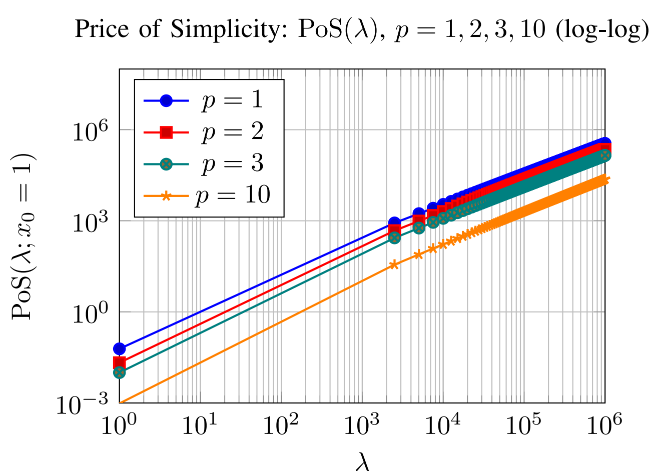

Fig. 2. Log-log plot of $\mathrm{PoS}(\lambda; x_0 = 1)$ vs λ. For large λ the curves are asymptotically linear in λ (slope 1 on the log–log plot corresponds to $\mathrm{PoS} \sim \mathrm{const} \cdot \lambda$).

The figure 3 presents a high-fidelity three-dimensional visualization of the Price of Simplicity over a logarithmically scaled coupling parameter $\lambda \in [10^{-2}, 10^3]$) and linearly scaled initial state $x_0 \in [0.1, 2]$, for multiple values of the $2p$-power exponent $p \in 1, 2, 3, 10$. Each semi-transparent surface corresponds to a distinct (p), revealing both multiplicative and additive amplification of the cost induced by ignoring mean-field interactions. As λ increases, PoS grows monotonically and approximately linearly on the log-log scale, reflecting the asymptotic relation PoS $\sim C_p \frac{2p-1}{2p} \lambda$ for large λ, while higher p systematically increases the magnitude of PoS due to the stronger dependence on the initial state x_0^{2p} in C_p. The curvature along the x_0 axis illustrates the superlinear sensitivity of PoS to the initial condition, which is particularly pronounced for larger p, confirming that both the additive loss and the relative sensitivity to initial conditions become unbounded as $\lambda \to \infty$. Consequently, the figure rigorously encapsulates the combined effects of the coupling strength λ, initial state x_0, and the nonlinearity exponent p on the inefficiency of freezing the mean-field term. It quantitatively demonstrates the divergent behavior of the Price of Simplicity in high-coupling regimes and highlighting the critical role of initial-state scaling in risk amplification across different p-power cost structures.

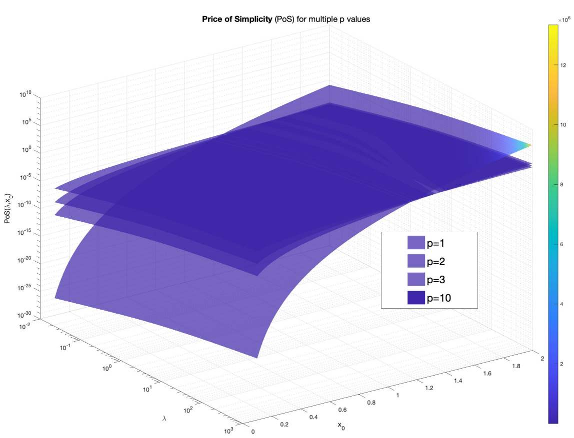

Fig. 3. Price of Simplicity as a function of the coupling parameter λ and the initial state x_0 in the continuous-time $2p$-power MFTC problem. The figure illustrates the additive cost incurred by freezing the mean-field term for multiple values of the nonlinearity exponent p.

V. Conclusion and Future Work

In this paper we have established a definitive quantitative boundary between classical optimal control and mean-field-type control. By deriving exact, closed-form solutions in a minimal deterministic setting, we demonstrated that the "oblivious" application of the standard Hamilton-Jacobi-Bellman equation where the population law is frozen as an exogenous signal is not merely a heuristic approximation, but a fundamental structural error.

Our analysis of the *Price of Anarchy* and *Price of Simplicity* proves that this error is unbounded. Whether in continuous-time or discrete-time, and whether for quadratic or higher-order $2p$-power costs, the performance penalty incurred by the oblivious controller grows arbitrarily large as the mean-field coupling λ increases. This result serves as a critical diagnostic for practitioners: while the classical HJB may suffice for weakly coupled systems ($\lambda \approx 0$), it fails catastrophically in regimes of strong interaction, where the correct dynamic programming principle on the space of measures is strictly required.

Future research should extend these explicit quantifications to stochastic regimes, particularly those involving common noise, atomic and non-atomic agents, where the governing equation becomes a stochastic master adjoint system. Additionally, there is significant value in analyzing the *Price of Oblivion* in reinforcement learning contexts, where agents that fail to track the population distribution effectively learn oblivious policies. Finally, further work is needed to map these control-theoretic gaps onto the equilibrium selection problems found in mean-field-type games with heterogeneous atomic and non-atomic agents.

References

[1] B. Jovanovic and R. W. Rosenthal, "Anonymous sequential games," *Journal of Mathematical Economics*, vol. 17, no. 1, pp. 77–87, 1988.
[2] Tamer Başar, Boualem Djehiche, Hamidou Tembine: Mean-Field-Type Game Theory I-II, Birkhäuser Cham, 2026
[3] Daniel Andersson and Boualem Djehiche, "A maximum principle for SDEs of mean-field type," *Applied Mathematics and Optimization*, vol. 63, no. 3, pp. 341–356, 2011.
[4] Alain Bensoussan, Jens Frehse, and Phillip 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).
[5] René Carmona and François 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.
[6] Wilfrid Gangbo and Andrzej 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.
[7] Minyi Huang, Peter E. Caines, and Roland 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.
[8] Jean-Michel Lasry and Pierre-Louis Lions, "Mean field games," *Japanese Journal of Mathematics*, vol. 2, no. 1, pp. 229–260, 2007.
[9] Pierre-Louis Lions, "Cours au Collège de France: Théorie des jeux à champ moyen," audio lectures, 2007–2011.
[10] Alain-Sol 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.
[11] 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.
[12] 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.