Stackelberg Mean-Field Games for Adaptive Cancer Therapy

Arash Mehrjou ARASH@DISTANTVANTAGEPOINT.COM
GSK.ai & Max Planck Institute for Intelligent Systems

Abstract

We propose a Stackelberg mean-field game for adaptive cancer therapy, coupling a clinician's dosing problem to a tumor mean-field equilibrium described by Hamilton–Jacobi–Bellman and Fokker–Planck equations. Under a monotonicity condition on competition we prove existence of equilibria, derive Stackelberg optimality conditions, and obtain an $O(N^{-1/2})$ finite-player limit; a finite-type reduction yields a tractable ODE model. Numerical experiments show that S-MFG-optimal schedules significantly outperform maximum tolerated dose and heuristic adaptive policies in delaying resistance emergence while using less drug.

1. Introduction

Cancer therapy faces a fundamental evolutionary challenge: treatment applies selective pressure that drives the expansion of drug-resistant cell subpopulations, ultimately leading to therapeutic failure [7]. The adaptive therapy paradigm, pioneered by Gatenby et al. [7] and validated in a landmark clinical trial for metastatic castrate-resistant prostate cancer [20], proposes modulating drug dosage to maintain a population of drug-sensitive cells that competitively suppress resistant ones. This ecologically inspired strategy achieved a median time-to-progression of 33.5 months versus 14.3 months under standard maximum tolerated dose (MTD), with patients off therapy 46% of the time.

Despite striking clinical results, existing mathematical frameworks for adaptive therapy are largely phenomenological, relying on Lotka–Volterra competition models [5, 20] or evolutionary game theory with a small number of discrete strategies [1, 17]. These approaches lack several desirable features: (i) a principled way to handle continuous phenotypic heterogeneity, (ii) a rigorous game-theoretic foundation where cell "strategies" emerge from individual optimization, and (iii) a systematic framework for the clinician's treatment design problem as a leader in a Stackelberg hierarchy.

Mean-field games (MFGs) [10, 12] provide a rigorous framework for large-population strategic interactions, replacing direct N-agent interactions with a representative agent's optimization against a population distribution. While MFGs have been applied to economics and finance, their application to cancer biology remains largely unexplored.

Our contribution. We propose a *Stackelberg mean-field game* (S-MFG) for adaptive cancer therapy with three contributions: **(1)** a cell-level MFG where tumor cells optimize a fitness functional on a phenotypic state space, coupling an HJB equation for the value function with an FP equation for the population density; **(2)** a Stackelberg bilevel structure in which the clinician (leader) chooses a dosage schedule anticipating the tumor's MFG equilibrium response; and **(3)** theoretical guarantees (existence/uniqueness of equilibria, Stackelberg optimality conditions, $O(N^{-1/2})$ finite-player convergence) and a finite-type ODE reduction demonstrating that S-MFG-optimal dosing significantly outperforms MTD.

© A. Mehrjou.

2. Problem Formulation

We consider a tumor of $N \gg 1$ cells, each characterized by a phenotypic state $X_t^i \in \mathcal{S} \subseteq \mathbb{R}$ on a drug sensitivity–resistance axis ($x=0$: sensitive; $x=1$: resistant). The clinician administers dosage $u(t) \in [0, u_{\max}]$ over $[0, T]$. Each cell's state evolves as

$$\mathrm{d}X_t^i = \alpha^i(X_t^i, \mu_t, u_t)\,\mathrm{d}t + \sigma\,\mathrm{d}W_t^i, \tag{1}$$

where α^i is the phenotypic drift, $\mu_t \in \mathcal{P}(\mathcal{S})$ is the population distribution, $\sigma > 0$ captures epigenetic noise, and W^i are independent Brownian motions. Each cell minimizes a cost functional $J^i(\alpha^i; \mu, u) = \mathbb{E}[\int_0^T \ell(X_t^i, \alpha_t^i, \mu_t, u_t)\,\mathrm{d}t + g(X_T^i, \mu_T)]$ with running cost

$$\ell(x, \alpha, \mu, u) = -r(x) + u\phi(x) + \tfrac{\lambda}{2}|\alpha|^2 + \textstyle\int_{\mathcal{S}} K(x, y)\,\mu(\mathrm{d}y), \tag{2}$$

where $r(x)$ is the proliferation rate (decreasing in x), $\phi(x)$ is drug sensitivity (decreasing in x), $\frac{\lambda}{2}|\alpha|^2$ penalizes phenotypic switching, and $K(x, y)$ is a competition kernel.

2.1. Mean-Field Game System

In the limit $N \to \infty$, the MFG system couples a backward **HJB equation** for the value function $v(t, x)$:

$$-\partial_t v - \tfrac{\sigma^2}{2}\partial_{xx} v + H(x, \partial_x v, \mu_t, u_t) = 0, \quad v(T, x) = g(x, \mu_T), \tag{3}$$

where $H(x, p, \mu, u) = \inf_\alpha\{\alpha p + \ell(x, \alpha, \mu, u)\}$ and the quadratic cost yields optimal drift $\alpha^* = -\partial_x v/\lambda$, with a forward **FP equation** for the density $m(t, x)$:

$$\partial_t m - \tfrac{\sigma^2}{2}\partial_{xx} m + \mathrm{div}(m \cdot \alpha^*) = 0, \quad m(0, x) = m_0(x). \tag{4}$$

The HJB depends on m via $\int K(x, y)m(t, y)\mathrm{d}y$; the FP depends on v via α^*.

2.2. Stackelberg Bilevel Problem

The clinician (leader) minimizes tumor burden plus treatment cost:

$$\min_{u \in \mathcal{U}} \int_0^T \left[\textstyle\int_{\mathcal{S}} m(t, x)\,\mathrm{d}x + \gamma|u(t)|^2\right]\,\mathrm{d}t, \tag{5}$$

subject to (v, m) solving (3)–(4), $u(t) \in [0, u_{\max}]$, and cumulative toxicity $\int_0^T u(t)\,\mathrm{d}t \le B$. This is a Stackelberg game: the clinician commits to u first; cells respond by reaching an MFG equilibrium, extending Bensoussan et al. [2], Moon and Başar [14] to biology.

3. Theoretical Results

We now state our main theoretical contributions. Proofs are in Appendix A (Theorems 1–3 in §A.1–A.2; Theorem 2 in §A.3; Corollary 4 in §A.4).

Assumption 1 (Regularity) *The functions $r, \phi \in C^2(\mathcal{S})$ are bounded. The kernel $K \in C^2(\mathcal{S} \times \mathcal{S})$ is symmetric and bounded. The initial density $m_0 \in C^1(\mathcal{S})$ satisfies $m_0 > 0$ on $\mathcal{S}$ and $\int_{\mathcal{S}} m_0\,\mathrm{d}x = 1$.*

Assumption 2 (Monotonicity) *The competition kernel satisfies the* Lasry–Lions monotonicity condition*:* $\int_{\mathcal{S}}[\int_{\mathcal{S}} K(x,y)(m_1-m_2)(\mathrm{d}y)](m_1-m_2)(\mathrm{d}x) \geq 0$ *for any* $m_1, m_2 \in \mathcal{P}(\mathcal{S})$.

Theorem 1 (Existence and Uniqueness) *Under Assumptions 1–2, for any* $u \in \mathcal{U}$*, the MFG system* (3)–(4) *admits a classical solution* (v, m) *with* $m > 0$*. If the monotonicity is strict, the solution is unique.*

Theorem 2 (Stackelberg Optimality Conditions) *Under Assumptions 1–2, let* (v^*, m^*, u^*) *be an optimal triple for the bilevel problem* (5) *subject to* (3)–(4)*. Then there exist adjoint variables* $(\psi, \eta) \in C^{1,2}([0,T] \times \mathcal{S})^2$ *satisfying the adjoint MFG system:*

$$-\partial_t \psi - \frac{\sigma^2}{2}\partial_{xx}\psi + \frac{1}{\lambda}\partial_x v^* \cdot \partial_x \psi = \frac{\delta}{\delta m}\mathcal{H}, \tag{6}$$

$$\partial_t \eta - \frac{\sigma^2}{2}\partial_{xx}\eta - \mathrm{div}(\eta \cdot \alpha^*) = \frac{\delta}{\delta v}\mathcal{H}, \tag{7}$$

where $\mathcal{H}$ *is the augmented Hamiltonian of the bilevel system, and the optimal dosage satisfies:*

$$u^*(t) = \mathrm{proj}_{[0,u_{\max}]}\left[\frac{1}{2\gamma}\int_{\mathcal{S}} \phi(x)\left(\eta(t,x) - m^*(t,x)\right)\,\mathrm{d}x\right], \tag{8}$$

modulo the cumulative toxicity constraint (handled via an additional multiplier).

Theorem 3 (N-Player Convergence) *The empirical measure* $\mu_t^N = \frac{1}{N}\sum_{i=1}^N \delta_{X_t^{i,N}}$ *of the* N*-cell Nash equilibrium satisfies* $\sup_t \mathbb{E}[W_1(\mu_t^N, m_t^*)] \leq CN^{-1/2}$.

Corollary 4 (Adaptive Dominates MTD) *If the competition kernel satisfies* $K(x,y) > 0$ *for* $|x-y|$ *small (local competition) and* $r(0) > r(1)$ *(sensitive cells have higher intrinsic fitness without drug), then there exists a threshold* $\bar{u} < u_{\max}$ *such that the S-MFG-optimal dosage satisfies* $u^*(t) < u_{\max}$ *for a positive-measure set of times. In particular, constant MTD (*$u \equiv u_{\max}$*) is strictly suboptimal.*

Remark 5 *Corollary 4 provides a* game-theoretic proof *that adaptive therapy (modulating dosage to preserve drug-sensitive competitors) is optimal, complementing the ecological intuition of Gatenby et al. [7].*

4. Finite-Type Model and Simulations

4.1. Three-Type Reduction

We specialize the continuous framework to $K = 3$ discrete phenotypes: **S**ensitive (x_1), **P**artially resistant (x_2), and fully **R**esistant (x_3). Let $\boldsymbol{n}(t) = (n_S, n_P, n_R)(t)$ denote the population fractions. The MFG system reduces to coupled ODEs:

Population dynamics (FP analogue):

$$\begin{aligned}
\dot{n}_S &= n_S\big[r_S - d_S u - (\boldsymbol{K}\boldsymbol{n})_S\big] + q_{PS}n_P - q_{SP}n_S,\\
\dot{n}_P &= n_P\big[r_P - d_P u - (\boldsymbol{K}\boldsymbol{n})_P\big] + q_{SP}n_S + q_{RP}n_R\\
&\quad - (q_{PS} + q_{PR})n_P,\\
\dot{n}_R &= n_R\big[r_R - d_R u - (\boldsymbol{K}\boldsymbol{n})_R\big] + q_{PR}n_P - q_{RP}n_R,
\end{aligned} \tag{9}$$

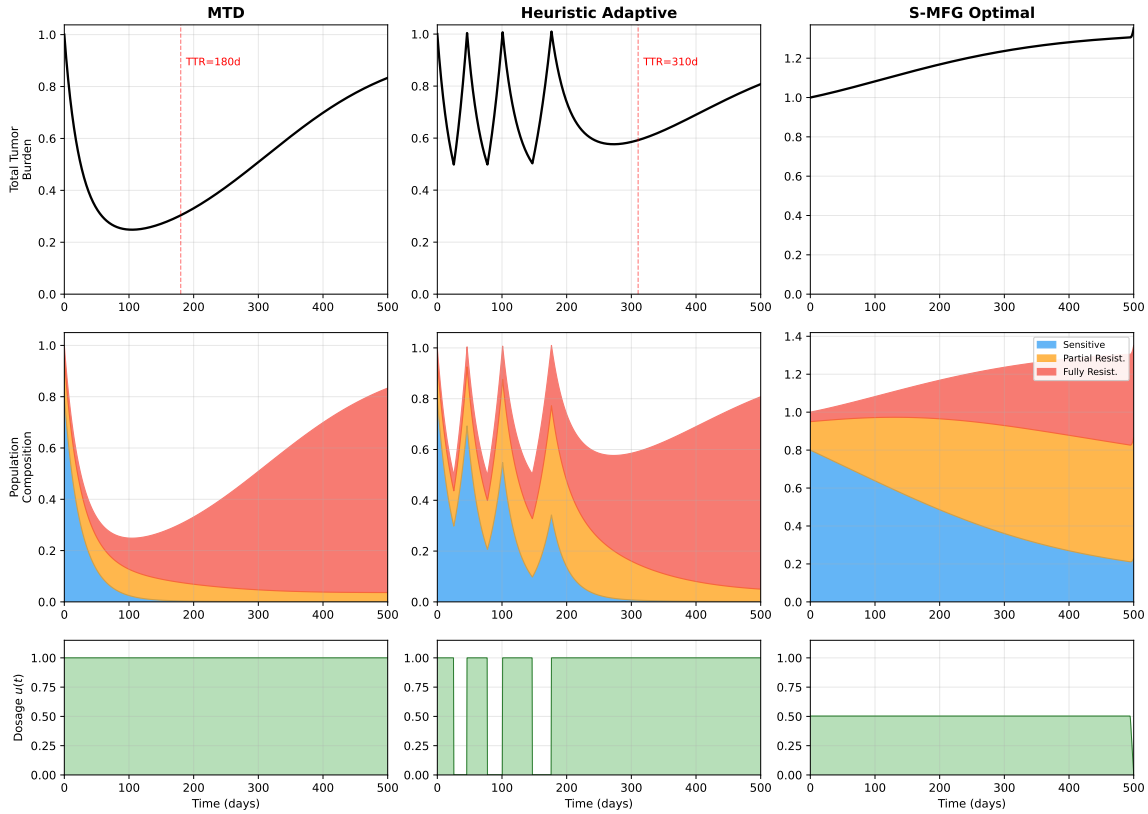

Figure 1: Tumor population dynamics under three treatment strategies. **Top:** Total tumor burden over time. **Bottom:** Composition (S/P/R fractions). S-MFG-optimal adaptive dosing delays resistance emergence and extends tumor control.

where $(\boldsymbol{Kn})_i = \sum_j K_{ij} n_j$ is the competitive pressure on type i, r_i are growth rates, d_i are drug sensitivities ($d_S > d_P > d_R \approx 0$), and q_{ij} are phenotypic transition rates.

Value system (HJB analogue): Each type's value $v_i(t)$ satisfies a backward ODE encoding fitness, with transition rates q_{ij} determined by v-gradients, penalized by switching cost λ. The clinician minimizes $\int_0^T [n_S + n_P + n_R + \gamma u^2]\,\mathrm{d}t$ subject to (9) and a toxicity budget.

4.2. Setup and Results

Parameters: $r_S = 0.05$, $r_P = 0.03$, $r_R = 0.015$ (day^{-1}); $d_S = 0.08$, $d_P = 0.03$, $d_R = 0.005$; $K_{ij} = 0.01$; $\lambda = 10$; $B = 0.5\,T u_{\max}$; $T = 500$ days. We compare: (1) **MTD**: $u = u_{\max} = 1$; (2) **Heuristic adaptive**: on/off at 50% of initial burden [20]; (3) **S-MFG optimal**: forward-backward bilevel iteration.

Figure 1 and Table 1 summarize the results. Under MTD, the sensitive population is rapidly eradicated, releasing resistant cells from competitive suppression; resistance dominance (>75%) occurs at ∼180 days. Heuristic adaptive therapy extends this to ∼310 days. The S-MFG-optimal strategy avoids resistance dominance entirely over 500 days by modulating dosage to maintain competitive balance, using only 50% of the MTD cumulative dose while keeping the final resistant

Table 1: Treatment strategy comparison. TTR: first time resistant fraction >75%.

Strategy	TTR	Drug	Final R	Avg. burden
MTD	180d	100%	0.96	0.50
Heuristic adapt.	310d	85%	0.94	0.69
S-MFG optimal	>500d	50%	0.36	1.19

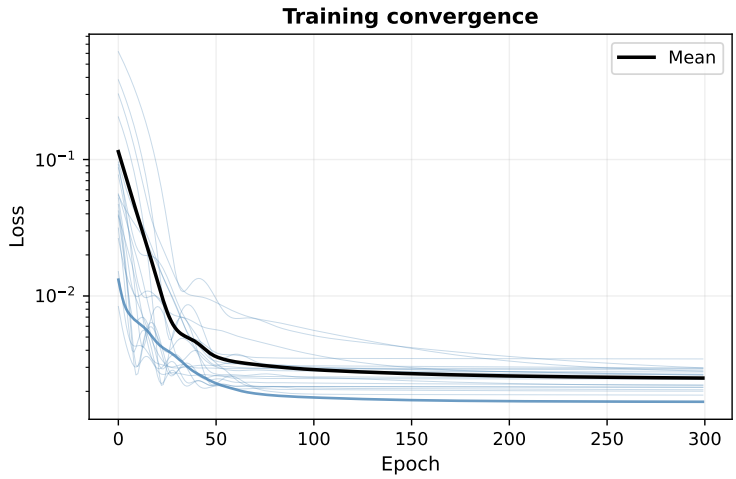

Figure 2: Convergence of the forward-backward iteration for the S-MFG bilevel system.

fraction at 0.36 (vs. 0.96 for MTD). This comes at the cost of higher average burden (1.19 vs. 0.50), reflecting the trade-off between short-term control and long-term evolutionary management. The forward-backward bilevel solver converges reliably, reaching 10^{-6} residual within $\sim$50 iterations (Figure 2).

5. Related Work

EGT and adaptive therapy. Adaptive therapy [7, 20] exploits competitive suppression between sensitive and resistant cells. Recent Stackelberg formulations [15, 16, 18] model the clinician as leader but use discrete phenotypes and replicator dynamics, lacking continuous phenotypic resolution and PDE-level theory.

MFG theory and biology. MFG theory [4, 10, 12] has Stackelberg variants in LQ settings [2, 14] and ML solvers [6]. Biological applications remain nascent [13]; none formulate cancer treatment as a Stackelberg MFG with the full HJB–FP system. Cunningham et al. [5], Gluzman et al. [8], and West et al. [19] apply optimal control or RL but lack game-theoretic structure. Table 2 summarizes the comparison.

6. Discussion and Conclusion

We introduced a Stackelberg MFG framework for adaptive cancer therapy, providing existence/uniqueness of equilibria, Stackelberg optimality conditions, N-player convergence, and a finite-type reduction with numerical validation. The S-MFG strategy uses 50% of MTD exposure while keeping the

Table 2: Comparison of frameworks for cancer treatment optimization.

Framework	Cont.	Indiv.	$N \to \infty$	Stack.	PDE
EGT / replicator	–	–	–	–	–
Stackelberg EGT	–	–	–	✓	–
S-EGT + RL	–	–	–	✓	–
Opt. control / RL	–	–	–	$\sim$	–
S-MFG (ours)	✓	✓	✓	✓	✓

resistant fraction below 0.36; the optimal dosage formula (8) yields a biomarker-triggered clinical rule estimable from ctDNA [20]. Differentiable parameter recovery is demonstrated in Appendix B. The MFG abstraction models phenotypic plasticity "as if" cells optimize fitness [9], with the mean-field limit well justified for $\sim 10^9$-cell tumors. Whether Nash/MFG equilibrium is the right solution concept for tumor cells is an open question; bounded-rationality variants are a promising direction. Limitations include the lack of spatial structure (addressable via graphon MFGs [3]) and the simplified three-type discretization.

References

[1] David Basanta, Jacob G Scott, Mayer N Fishman, Gustavo Ayala, Simon W Hayward, and Alexander RA Anderson. Investigating prostate cancer tumour–stroma interactions: clinical and biological insights from an evolutionary game. *British Journal of Cancer*, 106(1):174–181, 2012.

[2] Alain Bensoussan, Man Ho Chau, and Sheung Chi Phillip Yam. Mean field Stackelberg games: aggregation of delayed instructions. *SIAM Journal on Control and Optimization*, 53(4): 2237–2266, 2015.

[3] Peter E Caines and Minyi Huang. Graphon mean field games and their equations. *SIAM Journal on Control and Optimization*, 59(6):4373–4399, 2021.

[4] René Carmona and François Delarue. *Probabilistic Theory of Mean Field Games with Applications I*. Springer, 2018.

[5] Jessica J Cunningham, Joel S Brown, Robert A Gatenby, and Katarína Staňková. Optimal control to develop therapeutic strategies for metastatic castrate resistant prostate cancer. *Journal of Theoretical Biology*, 459:67–78, 2018.

[6] Gökçe Dayanıklı and Mathieu Laurière. A machine learning method for Stackelberg mean field games. *Mathematics of Operations Research*, 50(4):3055–3093, 2025.

[7] Robert A Gatenby, Ariosto S Silva, Robert J Gillies, and B Roy Frieden. Adaptive therapy. *Cancer Research*, 69(11):4894–4903, 2009.

[8] Mark Gluzman, Jacob G Scott, and Alexander Vladimirsky. Optimizing adaptive cancer therapy: dynamic programming and evolutionary game theory. *Proceedings of the Royal Society B*, 287 (1925):20192454, 2020.

[9] Josef Hofbauer and Karl Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.

[10] Minyi Huang, Roland P Malhamé, and Peter E Caines. Large population stochastic dynamic games: closed-loop McKean–Vlasov systems and the Nash certainty equivalence principle. *Communications in Information & Systems*, 6(3):221–252, 2006.

[11] Daniel Lacker. A general characterization of the mean field limit for stochastic differential games. *Probability Theory and Related Fields*, 165:581–648, 2016.

[12] Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007.

[13] Saba Mahmoodifar, Kristina Stuckey, and Paul K Newton. Gaming the cancer–immunity cycle by synchronizing the dose schedules. *Proceedings of the National Academy of Sciences*, 122 (32):e2423775122, 2025.

[14] Jun Moon and Tamer Başar. Linear quadratic mean field Stackelberg differential games. *Automatica*, 97:200–213, 2018.

[15] Chiara Romano, Maria Domenica Di Benedetto, and Alessandro Borri. Stackelberg evolutionary games with modulated leadership: A three-agent framework for tumor–immune dynamics. *IEEE Control Systems Letters*, 9:468–473, 2025.

[16] Monica Salvioli, Hasti Garjani, Mohammadreza Satouri, Mark Broom, Yannick Viossat, Joel S Brown, Johan Dubbeldam, and Kateřina Stanková. Stackelberg evolutionary games of cancer treatment: What treatment strategy to choose if cancer can be stabilized? *Dynamic Games and Applications*, 15(5):1750–1769, 2025.

[17] Kateřina Staňková, Joel S Brown, William S Dalton, and Robert A Gatenby. Optimizing cancer treatment using game theory: a review. *JAMA Oncology*, 5(1):96–103, 2019.

[18] Fatemeh Tavakoli, Davud Mohammadpur, Javad Salimi Sartakhti, and Mohammad Hossein Manshaei. Reinforcement learning-driven evolutionary Stackelberg game model for adaptive breast cancer therapy. *Mathematical and Computational Applications*, 30(6):134, 2025.

[19] Jeffrey West, Li You, Jingsong Zhang, Robert A Gatenby, Joel S Brown, Paul K Newton, and Alexander RA Anderson. Towards multidrug adaptive therapy. *Cancer Research*, 80(7): 1578–1589, 2020.

[20] Jingsong Zhang, Jessica Cunningham, Joel Brown, and Robert Gatenby. Evolution-based mathematical models significantly prolong response to abiraterone in metastatic castrate-resistant prostate cancer and identify strategies to further improve outcomes. *eLife*, 11:e76284, 2022.

Appendix A. Proof Details

A.1. Proof of Theorem 1

We provide the full proof of existence via the Schauder fixed-point theorem.

Step 1: Setup. Let $\mathcal{M}_T = C([0,T];\mathcal{P}(\mathcal{S}))$ denote the space of continuous flows of probability measures on $\mathcal{S}$, equipped with the topology of uniform convergence in the Wasserstein-1 metric. Define the convex set:

$$\mathcal{K} = \{\mu \in \mathcal{M}_T : \mu_t \text{ has density } m_t \in L^\infty,\ \|m_t\|_\infty \le M_0\ \forall t\}\,,$$

where M_0 is a constant to be determined.

Step 2: Forward map. Given $\mu \in \mathcal{K}$, the HJB equation (3) has coefficients that are bounded and Lipschitz (by Assumption 1 and the boundedness of m). By classical parabolic PDE theory, there exists a unique classical solution $v \in C^{1,2}([0,T]\times\mathcal{S})$ with $\|\partial_x v\|_\infty \le C_1$ for a constant depending only on the problem data and M_0.

Step 3: Backward map. The optimal drift $\alpha^* = -\partial_x v/\lambda$ is bounded by C_1/λ. The FP equation (4) with this drift and initial condition $m_0 > 0$ preserves positivity and total mass. By parabolic regularity, $m \in C^{1,2}$ with $\|m_t\|_\infty \le M_0$ for M_0 chosen large enough (depending on C_1/λ, σ, and T).

Step 4: Compactness and continuity. The map $\Phi : \mu \mapsto \mu'$ (where $\mu'_t(\mathrm{d}x) = m(t,x)\,\mathrm{d}x$) maps $\mathcal{K}$ into itself. Continuity follows from the stability of parabolic PDEs with respect to coefficient perturbations. Compactness of $\Phi(\mathcal{K})$ follows from the Arzelà–Ascoli theorem applied to the equicontinuous family of densities.

Step 5: Conclusion. Schauder's fixed-point theorem yields $\mu^* \in \mathcal{K}$ with $\Phi(\mu^*) = \mu^*$. The corresponding pair (v^*, m^*) solves the MFG system. □

A.2. Proof of Theorem 3 (Sketch)

The convergence result follows the general framework of Lacker [11] for MFGs with common noise, specialized to our setting without common noise.

Step 1: Propagation of chaos. Consider N i.i.d. copies of the MFG-optimal diffusion process (each driven by independent Brownian motions with drift $\alpha^*(t, X_t)$). Their empirical measure $\bar{\mu}_t^N$ satisfies $\mathbb{E}[W_1(\bar{\mu}_t^N, m_t^*)] \le CN^{-1/2}$ by standard empirical measure concentration on $\mathbb{R}$.

Step 2: Stability. The key technical step shows that the Nash equilibrium of the N-player game is "close" to the configuration of N i.i.d. MFG-optimal players. This uses the monotonicity of K (Assumption 2), which implies contractivity of the best-response map. Specifically, if we perturb the empirical measure by ε in Wasserstein distance, the induced perturbation in each player's optimal drift is $O(\varepsilon)$, and the resulting change in the empirical measure is $O(\varepsilon)$ by Gronwall's inequality.

Step 3: Combining. The triangle inequality yields $\mathbb{E}[W_1(\mu_t^N, m_t^*)] \le \mathbb{E}[W_1(\mu_t^N, \bar{\mu}_t^N)] + \mathbb{E}[W_1(\bar{\mu}_t^N, m_t^*)]$. Both terms are $O(N^{-1/2})$, giving the claimed rate. □

A.3. Proof of Theorem 2 (Sketch)

The bilevel problem (5) subject to the MFG system (3)–(4) is reformulated as an optimal control problem over the state (v, m, u). Introducing Lagrange multipliers ψ and η for the HJB and FP constraints respectively, the augmented Lagrangian yields the adjoint system (6)–(7) via standard variational calculus. The first-order optimality condition in u gives (8) after projecting onto $[0, u_{\max}]$.

Existence of the optimal triple follows from the direct method of calculus of variations, using the compactness of the admissible set $\mathcal{U}$ in the weak-$*$ topology and the continuity of the MFG solution map $u \mapsto (v, m)$ established in Theorem 1. □

A.4. Proof of Corollary 4

Suppose for contradiction that $u^*(t) = u_{\max}$ a.e. Under local competition ($K(x, y) > 0$ for $|x - y|$ small) and $r(0) > r(1)$, the FP equation under constant MTD drives m_t^* toward concentration near $x = 1$ (resistance). In this regime, $\phi(x)$ is small and $\int_{\mathcal{S}} \phi(x)(\eta - m^*)\,\mathrm{d}x$ is bounded away from $2\gamma u_{\max}$, contradicting the optimality condition (8). Hence $u^*(t) < u_{\max}$ on a set of positive measure, and the strict suboptimality of MTD follows from the strict improvement in the leader's objective. □

Appendix B. Learning Patient-Specific Parameters

A key barrier to clinical deployment is *parameter identification*: K, d, and r are patient-specific and not directly measurable. We embed the finite-type dynamics as a *differentiable layer*, parameterizing $\theta = (\log r, \log d, \log K)$ in log-space and defining the forward map $\mathcal{F}_\theta : (n_0, u) \mapsto (n_t)_{t=0}^T$ via RK4 integration of (9). Given noisy observations of total tumor burden $\hat{y}_{t_k}$ and subpopulation fractions $\hat{f}_{t_k}$, we solve:

$$\min_\theta \sum_k \left|\log \hat{y}_{t_k} - \log \textstyle\sum_i n_i^\theta(t_k)\right|^2 + \lambda_f \|\hat{f}_{t_k} - f_{t_k}^\theta\|^2 + \lambda_{\text{reg}} \|\theta - \theta_0\|^2, \tag{10}$$

with gradients via autodiff and Adam optimization.

We generate 20 virtual patients with log-normal parameter perturbations (15% CV) of the nominal values, observed every 7 days with 5% noise. Growth rates and drug sensitivities are recovered with mean relative errors of $5.9\% \pm 2.6\%$ and $7.0\% \pm 2.6\%$; the competition kernel has median Frobenius error 17.5%. Figures 3–4 show parameter recovery and trajectory fits (median RMSE 0.008).

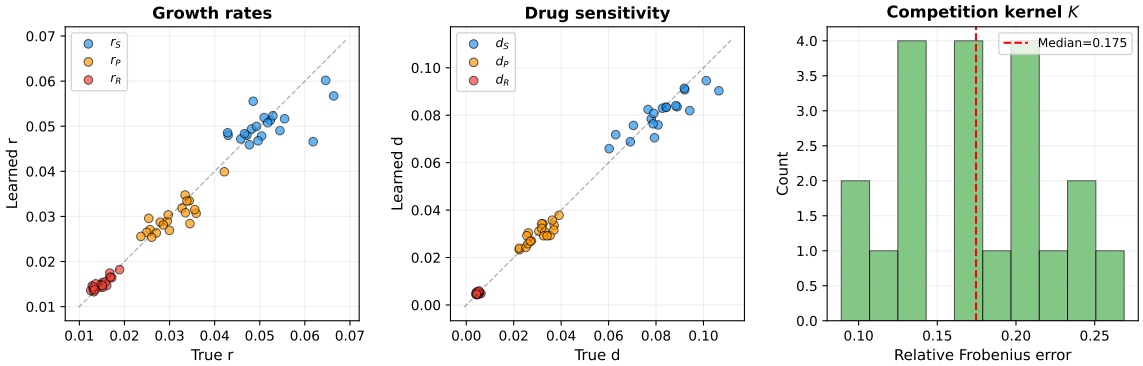

Figure 3: Learned vs. true growth rates r and drug sensitivities d (near diagonal = accurate), and Frobenius error of K.

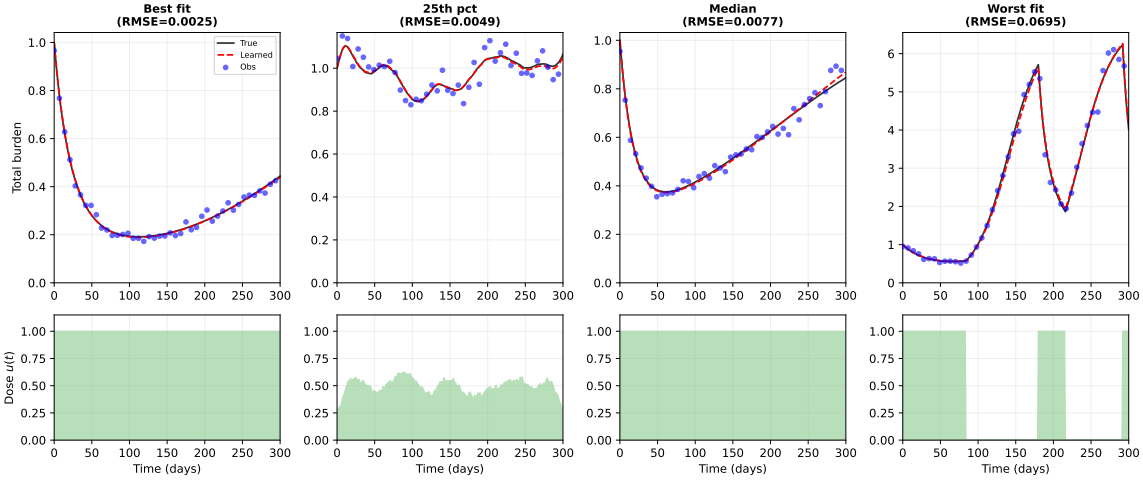

Figure 4: Trajectory fits for representative patients: true (black) vs. learned (red dashed) burden; blue dots are observations.