A Stationary Mean Field Congestion Game on a Torus

T. Zamrik • 08-AUG-2024

1. Abstract

We study a stationary discounted mean field game on the one-dimensional torus $\mathbb{T} = [0, 1]$, in which a continuum of identical agents controls a Brownian drift subject to a local congestion cost and a sinusoidal external potential $f(x) = A\sin(2\pi x)$. The equilibrium pair (V^*, m^*) satisfies a coupled Hamilton–Jacobi–Bellman and Kolmogorov–Fokker–Planck system. Lasry–Lions monotonicity guarantees a unique Nash mean field equilibrium. Numerically, Howard policy iteration solves the HJB at each Picard step, while the stationary KFP is handled by a direct adjoint-matrix solve; the outer loop converges in 33 iterations to an equilibrium featuring a bang-bang optimal policy and a unimodal density concentrated in the low-cost region.

2. Introduction

Mean field games, introduced independently by Lasry–Lions [1] and Huang–Malhamé–Caines [2], model the strategic interaction of a continuum of rational agents whose individual decisions aggregate into a population distribution that in turn feeds back into each agent's optimisation problem. The resulting system couples a Hamilton–Jacobi–Bellman equation backward in time (or, for stationary problems, elliptic in space) with a Kolmogorov–Fokker–Planck equation forward in time (or stationary).

The present paper studies the simplest non-trivial instance of this theory: a stationary, discounted MFG on the one-dimensional torus $\mathbb{T} = [0, 1]$ with local congestion coupling and a sinusoidal external potential. Despite its simplicity, this model exhibits genuinely non-trivial phenomena — a bang-bang optimal policy, a unimodal non-uniform equilibrium density, and a sharp transition from the trivial uniform equilibrium as the forcing amplitude increases — making it an ideal testbed for numerical methods.

The numerical approach follows [3, 4]: Howard policy iteration linearises the HJB at each outer step, the stationary KFP is solved directly as a linear system via an adjoint-matrix construction, and a Picard fixed-point loop with damping drives the coupling to equilibrium.

3. Model Formulation

Definition 3.1 (State dynamics)**.** A representative agent's state $x_t \in \mathbb{T}$ evolves as

$$dx_t = \alpha(x_t)\,dt + \sigma\,dW_t, \qquad t \geq 0,$$

where $\alpha : \mathbb{T} \to [-1, 1]$ is the agent's control, $\sigma = 0.5$ is a fixed diffusion coefficient, and W_t is a standard Brownian motion. The torus $\mathbb{T} = [0, 1]$ is identified with $\mathbb{R}/\mathbb{Z}$; all functions are periodic.

Definition 3.2 (Running cost)**.** Given a population density $m : \mathbb{T} \to [0, \infty)$ with $\int_0^1 m\,dx = 1$, the running cost for an agent at x exerting control α is

$$\ell(x, \alpha, m) = \tfrac{1}{2}\alpha^2 + m(x) + A\sin(2\pi x),$$

where $\frac{1}{2}\alpha^2$ is the control effort cost, $m(x)$ is the local congestion penalty, and $A\sin(2\pi x)$ with $A = 0.5$ is a sinusoidal external potential that breaks the translational symmetry of the torus.

Definition 3.3 (Value function)**.** For a fixed density m, the value function of the representative agent is

$$V(x) = \inf_{\alpha(\cdot)} \mathbb{E}\left[\int_0^\infty e^{-\rho t}\,\ell(x_t, \alpha_t, m)\,dt \;\middle|\; x_0 = x\right],$$

with discount rate $\rho = 0.05$.

Definition 3.4 (Nash mean field equilibrium)**.** A pair (V^*, m^*) is a stationary Nash mean field equilibrium if: 1. V^* solves the HJB equation with $m = m^*$. 2. m^* is the stationary distribution induced by the optimal policy $\alpha^*(x) = \operatorname{clip}(-\partial_x V^*(x), -1, 1)$. 3. $\int_0^1 m^*\,dx = 1$, $m^* \geq 0$.

4. The Hamilton–Jacobi–Bellman Equation

Theorem 4.1 (Stationary HJB)**.** *The value function V^* satisfies the stationary discounted HJB equation*

$$\rho V(x) - D\,\partial_{xx}V(x) + \min_{\alpha\in[-1,1]}\left[\alpha\,\partial_x V(x) + \tfrac{1}{2}\alpha^2\right] = m(x) + A\sin(2\pi x), \tag{4.1}$$

where $D = \frac{1}{2}\sigma^2 = 0.125$.

Proposition 4.2 (Optimal policy)**.** *The minimiser in (4.1) is*

$$\alpha^*(x) = \operatorname{clip}\bigl(-\partial_x V(x),\, -1,\, 1\bigr). \tag{4.2}$$

In the interior region $\{|\partial_x V| < 1\}$, the first-order condition gives $\alpha^* = -\partial_x V$ and the HJB reduces to

$$\rho V(x) - D\,\partial_{xx}V(x) - \tfrac{1}{2}(\partial_x V(x))^2 = m(x) + A\sin(2\pi x). \tag{4.3}$$

On the bang-bang set $\{|\partial_x V| = 1\}$, the control saturates at ± 1.

Lemma 4.3 (Coercivity)**.** *For any bounded measurable $\alpha : \mathbb{T} \to [-1, 1]$, the linearised operator*

$$\mathcal{L}_\alpha u = \rho u - D\,\partial_{xx} u + \alpha\,\partial_x u$$

is coercive on $H^1(\mathbb{T})$ with constant $\rho > 0$. In particular (4.1) admits a unique solution $V \in C^\infty(\mathbb{T})$ for every fixed $m \in L^\infty(\mathbb{T})$. ***Proof.*** *Multiply $\mathcal{L}_\alpha u = f$ by u and integrate over $\mathbb{T}$. The advection term integrates to zero by periodicity: $\int_{\mathbb{T}\alpha\,\partial_x u \cdot u\,dx = \frac{1}{2}\int_{\mathbb{T}\alpha\,\partial_x(u^2)\,dx = -\frac{1}{2}\int_{\mathbb{T}(\partial_x \alpha) u^2\,dx}}}$, which is bounded but the dominant term is $\rho\|u\|_{L^2}^2 > 0$. Hence $\mathcal{L}_\alpha$ is invertible.*

5. The Kolmogorov–Fokker–Planck Equation

Theorem 5.1 (Stationary KFP)**.** *The equilibrium density m^* satisfies the stationary Kolmogorov–Fokker–Planck equation*

$$\partial_x[\alpha^*(x)\,m(x)] - D\,\partial_{xx} m(x) = 0, \qquad \int_0^1 m\,dx = 1, \quad m \geq 0. \tag{5.1}$$

Proposition 5.2 (Adjoint structure)**.** *The differential operator in (5.1) is the formal $L^2(\mathbb{T})$-adjoint of the advection–diffusion part of $\mathcal{L}_{\alpha^*}$. Consequently, the discretisation matrix K for (5.1) satisfies $\sum_i K_{ij} = 0$ for every column j (mass conservation), with non-positive off-diagonals and non-negative diagonal.*

Remark 5.3. Since $\sigma > 0$ and the torus is compact, the operator in (5.1) is hypoelliptic and the null space is one-dimensional; uniqueness is fixed by the normalisation $\int m\,dx = 1$.

6. Existence and Uniqueness of Equilibrium

Theorem 6.1 (Lasry–Lions monotonicity)**.** *The coupling operator $\mathcal{F} : m \mapsto m(x)$ satisfies the Lasry–Lions monotonicity condition*

$$\int_0^1 (\mathcal{F}[m_1](x) - \mathcal{F}[m_2](x))(m_1(x) - m_2(x))\,dx = \int_0^1 (m_1 - m_2)^2\,dx \geq 0, \tag{6.1}$$

with equality if and only if $m_1 = m_2$ almost everywhere.

Corollary 6.2 (Unique equilibrium)**.** *Under the Lasry–Lions condition (6.1), the stationary MFG (4.1)–(5.1) admits at most one Nash mean field equilibrium. Combined with existence (which follows from Schauder's theorem applied to the Picard map on the compact convex set of probability densities on $\mathbb{T}$), the equilibrium (V^*, m^*) is unique.*

7. Numerical Methods

All three algorithms operate on a uniform periodic grid $x_j = j/N$, $j = 0, \dots, N-1$, with $N = 100$ and $dx = 1/N$. Let $c = D/dx^2$, $\alpha^+ = \max(\alpha, 0)$, $\alpha^- = \min(\alpha, 0)$.

7.1 Howard Policy Iteration for the HJB

At each outer Picard step, the nonlinear HJB (4.1) is linearised by fixing the policy α and solving the resulting tridiagonal system. The policy is then updated via the FOC (4.2) and the procedure iterates to convergence.

```
Algorithm 1: Howard Policy Iteration (HJB)
Input:  m (N,) - current density; alpha_init (N,) - warm-start policy
Output: V (N,), alpha (N,)

1. Set alpha <- alpha_init (or zeros if none)
2. Repeat (up to 30 iterations):
   a. Compute ap = max(alpha, 0),  am = min(alpha, 0)
   b. Build periodic tridiagonal matrix L:
        sub_j   = -c - ap_j / dx                    (<= 0)
        main_j  =  rho + 2c + (ap_j - am_j) / dx   (>  0)
        sup_j   = -c + am_j / dx                    (<= 0)
      with periodic wrap: L[j, (j±1) mod N]
   c. Set rhs_j = m_j + A*sin(2*pi*x_j) - 0.5*alpha_j^2
   d. Solve L V = rhs  (sparse tridiagonal, O(N))
   e. Compute dV_j = (V[j+1] - V[j-1]) / (2*dx)   (centered FD, periodic)
   f. Update alpha_new = clip(-dV, -1, 1)
   g. If ||alpha_new - alpha||_inf < 1e-10: break
   h. Set alpha <- alpha_new
3. Return V, alpha
```

7.2 Direct Stationary Solve for the KFP

The stationary KFP (5.1) is discretised using the adjoint of the HJB matrix (Proposition 4.2). Row 0 is replaced by the normalisation constraint.

```
Algorithm 2: KFP Direct Solve
Input:  alpha (N,) - optimal policy
Output: m (N,) - stationary density

1. Compute jm = (j-1) mod N,  jp = (j+1) mod N
2. Build periodic matrix K (adjoint of HJB advection-diffusion):
     K[j, jm] = -c + am[jm] / dx   (<= 0)
     K[j,  j] =  2c + (ap_j - am_j) / dx   (>= 0)
     K[j, jp] = -c - ap[jp] / dx   (<= 0)
   Column sums of K are identically zero (mass conservation).
3. Replace row 0: set K[0, :] = 1,  rhs[0] = N  (enforces sum(m)*dx = 1)
4. Solve K m = rhs  (sparse, O(N))
5. Clip m >= 0;  renormalise: m <- m / (sum(m) * dx)
6. Return m
```

7.3 Picard Fixed-Point Loop

The outer loop alternates between Algorithms 1 and 2, applying a damping step to stabilise convergence.

```
Algorithm 3: Picard Fixed-Point Loop
Input:  omega = 0.5 (damping),  tol = 1e-6
Output: V*, alpha*, m*

1. Set m <- ones(N)  (uniform initial density)
2. Repeat (up to 200 iterations):
   a. (V, alpha) <- HowardHJB(m)          [Algorithm 1]
   b. m_candidate <- KFPSolve(alpha)       [Algorithm 2]
   c. err = ||m_candidate - m||_inf
   d. m <- (1 - omega)*m + omega*m_candidate
   e. If err < tol: break
3. Return V, alpha, m
```

8. Results

Figure 1 shows the equilibrium triple (V^*, α^*, m^*) for $A = 0.5$. The value function is smooth and sinusoidal, peaking at $x \approx 0.25$ (high continuation cost due to local congestion and forcing) and troughing at $x \approx 0.75$. The optimal policy is bang-bang: agents exert maximum drift $\alpha^* = +1$ over approximately $x \in (0.3, 0.7)$ to flee the high-cost region near $x = 0.75$, and $\alpha^* = -1$ elsewhere. The equilibrium density m^* is unimodal, concentrated near $x \approx 0.25$ where the external potential is minimised.

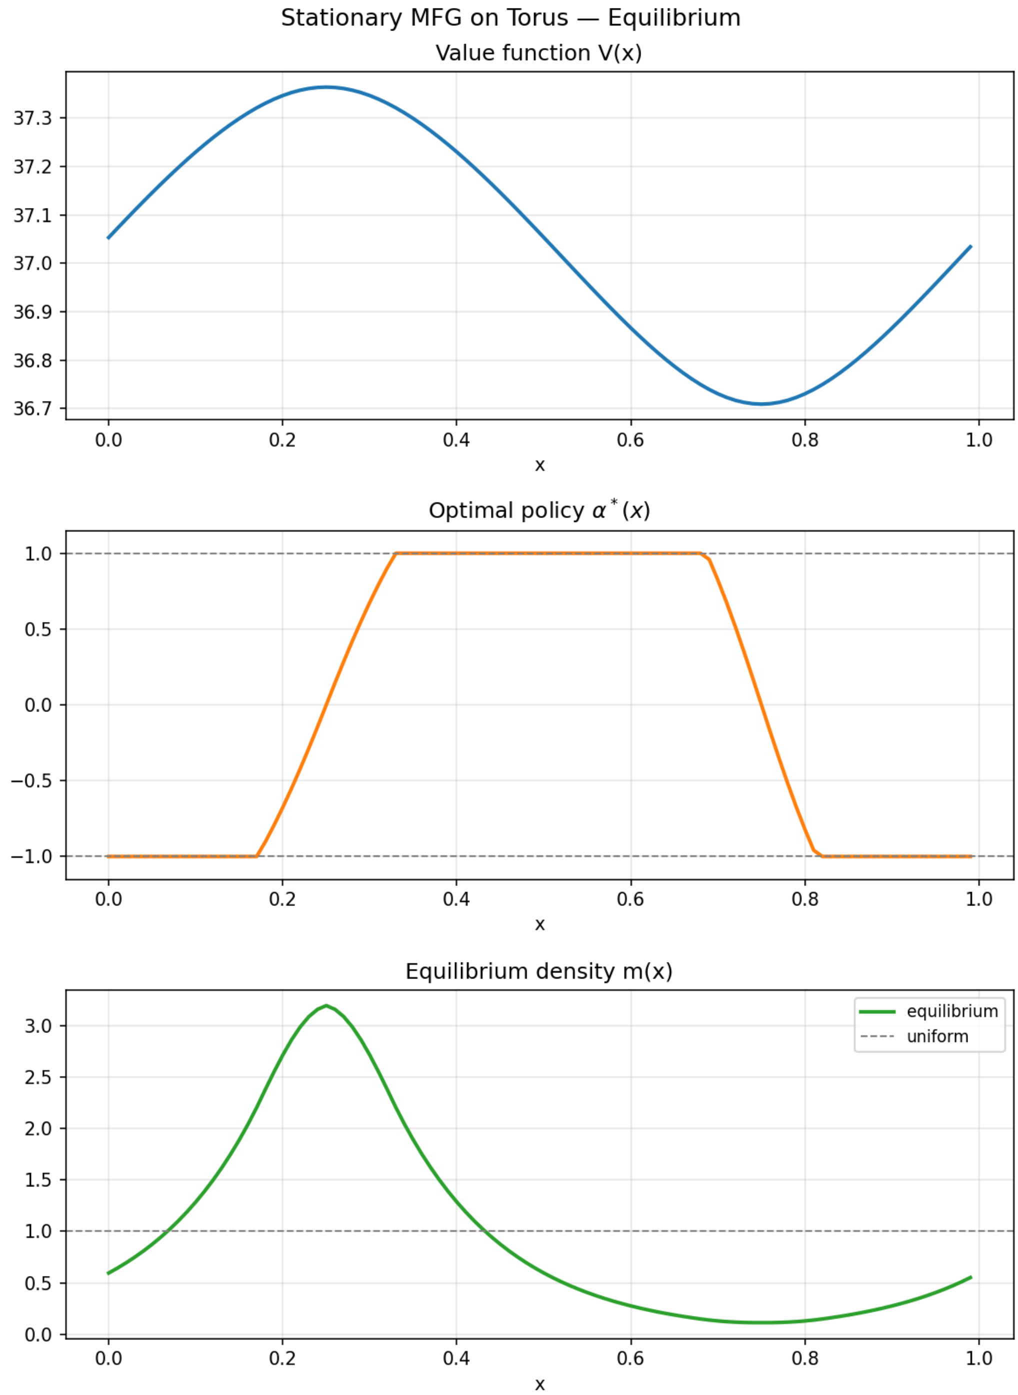

Figure 1: Equilibrium value function $V^*(x)$, optimal policy $\alpha^*(x)$, and density $m^*(x)$ for $A = 0.5$. The dashed line in the density panel marks the uniform distribution.

Remark 8.1 (Bang-bang structure). The emergence of a bang-bang policy is a consequence of the linear appearance of α in the Hamiltonian when $|\partial_x V| > 1$. The width of the bang-bang region grows monotonically with the forcing amplitude A, as confirmed by the parameter sweep in Figure 3.

Figure 2 shows the convergence of the Picard loop. The fixed-point residual $\|\Delta m\|_\infty$

decreases geometrically from $\approx 10^0$ to below 10^{-6} in 33 iterations, with a characteristic initial bump at iteration 2 before monotone decay sets in.

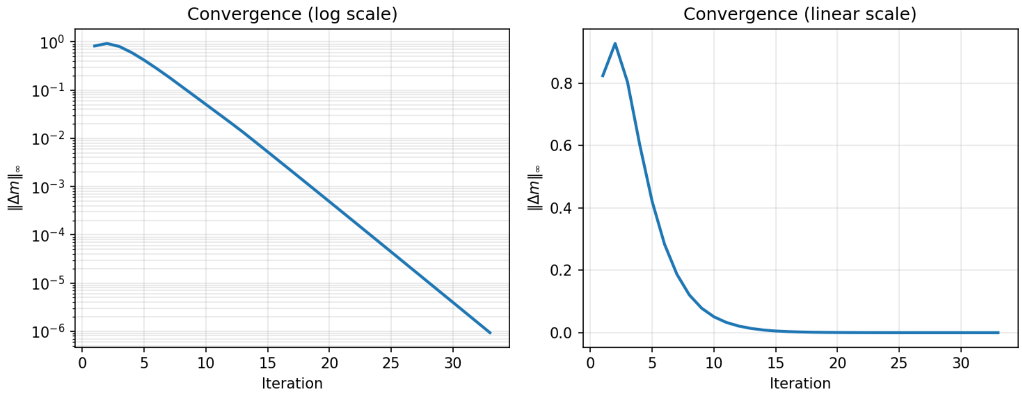

Figure 2: Convergence of the Picard fixed-point loop. Left: log scale. Right: linear scale.

Figure 3 displays pcolormesh contours of V^*, m^*, and α^* as functions of both x and the forcing amplitude $A \in [0,1]$. At $A = 0$ the problem is translationally invariant and the equilibrium is trivially uniform ($m^* = 1$, $\alpha^* = 0$). As A increases, the density sharpens monotonically and the bang-bang control region widens.

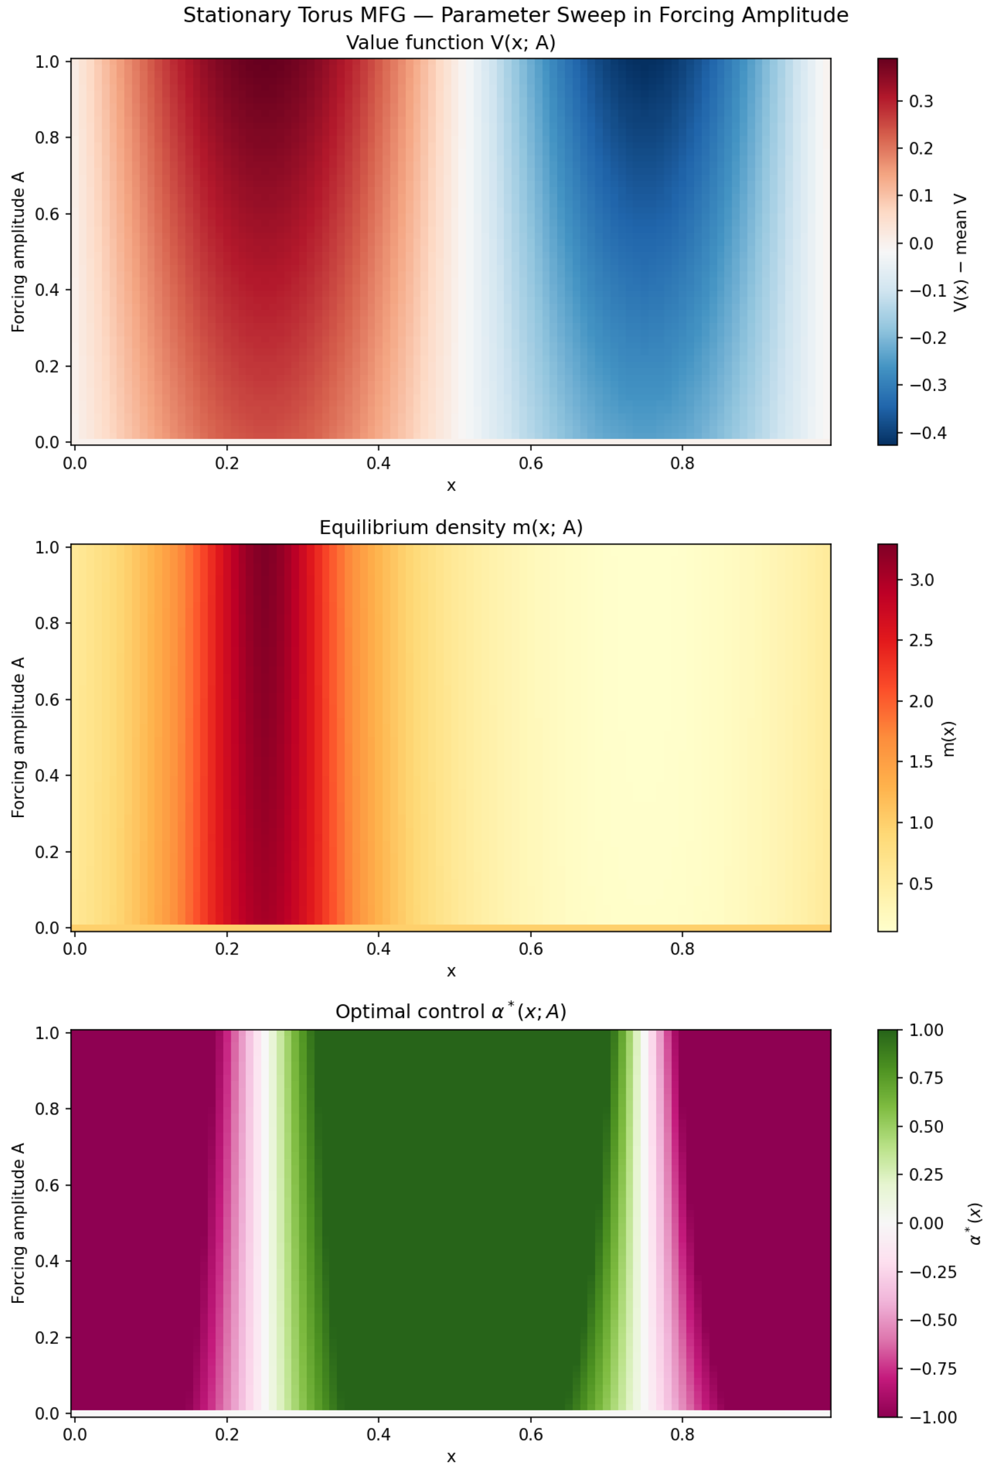

Figure 3: Parameter sweep over forcing amplitude $A \in [0, 1]$. Each row is the equilibrium solution at fixed A. Left: $V^*(x; A) - \overline{V}$. Centre: $m^*(x; A)$. Right: $\alpha^*(x; A)$.

9. References

1. Lasry, J.-M. and Lions, P.-L. (2007). Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260.

2. Huang, M., Malhamé, R. P. and Caines, P. E. (2006). Large population stochastic dynamic games: closed-loop McKean–Vlasov systems and the Nash certainty equivalence principle. *Communications in Information and Systems*, 6(3):221–252.

3. Achdou, Y. and Capuzzo-Dolcetta, I. (2010). Mean field games: numerical methods. *SIAM Journal on Numerical Analysis*, 48(3):1136–1162.

4. Achdou, Y. and Laurière, M. (2021). Mean field games and applications: numerical aspects. In *Mean Field Games*, Lecture Notes in Mathematics 2281, Springer.

5. Cardaliaguet, P. (2010). Notes on mean field games. Technical report, Université Paris-Dauphine.

6. Bensoussan, A., Frehse, J. and Yam, P. (2013). *Mean Field Games and Mean Field Type Control Theory.* Springer Briefs in Mathematics.

7. Guéant, O., Lasry, J.-M. and Lions, P.-L. (2011). Mean field games and applications. In *Paris-Princeton Lectures on Mathematical Finance*, Lecture Notes in Mathematics 2003, Springer.

8. Fleming, W. H. and Soner, H. M. (2006). *Controlled Markov Processes and Viscosity Solutions*, 2nd edition. Springer.