Les Cahiers du GERAD ISSN: 0711–2440

Generalized Nash equilibria for GMFG systems with control mean fields

Z. Huang, P. E. Caines

G–2026–88

October 2026

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

Citation suggérée : Z. Huang, P. E. Caines (Octobre 2026). Generalized Nash equilibria for GMFG systems with control mean fields, Rapport technique, Les Cahiers du GERAD G– 2026–88, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (https://www.gerad.ca/fr/papers/G-2026-88) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

Suggested citation: Z. Huang, P. E. Caines (October 2026). Generalized Nash equilibria for GMFG systems with control mean fields, Technical report, Les Cahiers du GERAD G–2026–88, GERAD, HEC Montréal, Canada.

Before citing this technical report, please visit our website (https://www.gerad.ca/en/papers/G-2026-88) to update your reference data, if it has been published in a scientific journal.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2026
– Bibliothèque et Archives Canada, 2026

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2026
– Library and Archives Canada, 2026

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Generalized Nash equilibria for GMFG systems with control mean fields

Ziqi Huang [a, b]

Peter E. Caines [a, b]

[a] *GERAD, Montréal (Qc), Canada, H3T 1J4*

[b] *Department of Electrical and Computer Engineering, McGill University, Montréal (Qc), Canada, H3A 0E9*

ziqi.huang2@mail.mcgill.ca
peter.caines@mcgill.ca

October 2026
Les Cahiers du GERAD
G–2026–88

Copyright © 2026 Huang, Caines

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract : A generalized Nash equilibrium (GNE) framework is introduced for a graphon mean field game of controls (GMFGC). Each agent's drift, cost, and hard control constraints depend on local and graphon-weighted moments of the joint state–control laws. Uniform linear independence constraint qualification and strong convexity yield stable feasible sets and a Lipschitz constrained best response. The associated GMFGC Hamilton–Jacobi–Bellman equation has a classical solution and the GMFGC Fokker–Planck–Kolmogorov equation a generalized energy solution; under weak coupling, the resulting GNE is unique. Furthermore, the ϵ-Nash property is established for finite-network, finite-population constrained GMFGC systems.

Keywords: Graphon mean field games; mean field games of controls; mean-field-dependent control constraints; generalized Nash equilibrium; finite-population approximation

Mathematics Subject Classification: 49N80, 35Q89, 91A16, 90C31, 93E20

Acknowledgements: This work was supported in part by AFOSR Grant FA9550–23–1–0015 and NSERC Grant RGPIN–2019–05336.

1 Introduction

Mean field game (MFG) theory studies Nash equilibria of large populations of asymptotically negligible agents through an infinite-population limit. Its foundational ideas appeared independently in the work of Huang, Caines, and Malhamé [31, 32] and Lasry and Lions [39]. For standard treatments, see [3, 16]. In controlled diffusion models, equilibrium is characterized by a backward Hamilton–Jacobi–Bellman (HJB) equation for the representative agent's value function, coupled with a forward Fokker–Planck–Kolmogorov (FPK) equation for the population law.

Mean field games of controls (MFGC) generalize MFG theory to the frameworks where agents interact through the joint state–control law, or distribution, rather than through the state law alone. Potential applications include trade crowding [13], relative consumption [33], congestion with endogenous prices [4], and residential energy trading [2]. These models have been analyzed using probabilistic methods [17, 40], PDE methods [13, 26], HJB-based explicit solution [33], and Lagrangian methods [4].

State and control constraints weaken MFG HJB–FPK regularity. Constrained MFGs are studied by Lagrangian methods based on probability measures over optimal constrained trajectories [1, 4, 12]. In finite-agent generalized Nash equilibria (GNEs), shared constraints make each agent's feasible set depend on the others' strategies [21]. Dynamic GNEs have been formulated as differential variational inequalities [18]; mean-field limits of constrained finite-agent differential GNEs are analyzed in [28].

Graphon limits [41] provide a tractable framework for asymptotically dense heterogeneous networks. Static graphon games model strategic interactions through graphon-weighted aggregates [15, 43]. For graphon mean field games (GMFGs) in stochastic dynamic settings, the GMFG equations and finite-network ϵ-Nash approximation theory were developed in [7, 8, 9], with subsequent work addressing linear-quadratic-Gaussian models [23] and classical solutions [30].

Graphon limits cannot retain the local structure of bounded-degree sparse networks since they are void; however, stochastic games on specific sparse networks were studied in [37]. Concerning sparse networks in general, graphings describing bounded-degree network limits are presented in [41]. In [34] there is an analysis of the control of populations of agents on local–global convergent sequences of networks, and these necessarily converge to graphings, see [27]. Graphexes model sparse exchangeable random graphs [46] including those with the small world property and permitting the formulation of graphex MFG equilibria [20]. The vertexon–graphexon framework covers dense and sparse embedded network limits enabling the definition of the graphexon MFG equations [6, 10, 11], while related measure-based approaches establish Vlasov mean field limits for processes on digraph measures [35], co-evolutionary signed networks [25] and directed hypergraph measures [36].

The presenting model combines graphon population interactions with endogenous control constraints. It is motivated by macroscopic resource constraints, under which aggregate control use must respect network capacity. Related formulations include density caps enforced by pressure prices in MFGs [14], population-level safety constraints [29], and constructive budget-constrained incentives in finite-agent games [48].

For parameterized constrained optimization, metric regularity, i.e., the local control of the distance to the feasible set by the constraint violation under local perturbations, is a standard stability property [5, 45]. In the present setting, Robinson's Constraint Qualification [5] yields such an error bound, through which Hausdorff stability of the feasible sets with respect to the mean-field parameter is established.

The Karush–Kuhn–Tucker (KKT) system can be formulated as a generalized equation. Strong regularity of its solution [5, 44], i.e., locally unique and Lipschitz continuous solvability of its perturbed linearization, implies that in the present setting the optimizer and associated Lagrange multipliers depend Lipschitz continuously on the state, costate, and mean-field parameter. These properties provide the Hamiltonian and feedback regularity needed for the subsequent HJB–FPK well-posedness analysis.

1.1 Main contributions

To the best of the authors' knowledge, this is the first work to combine graphon interactions, dependence on the joint state–control law, and endogenous shared control constraints. We refer to the resulting model as a constrained graphon mean field game of controls (GMFGC). In the present study, each agent's feasible control set depends on the population's endogenous joint state–control distribution, making the equilibrium generalized Nash. This coupling captures network-wide resource constraints and creates the key difficulty of moving feasible sets.

Hausdorff–Lipschitz stability of the moving feasible sets and KKT strong regularity jointly yield Lipschitz dependence of the best response on the mean-field flow. Weak coupling then gives existence and uniqueness of the GNE.

In finite games, empirical moments can make the limiting feedback infeasible. Feasibility transport and a finite-feedback fixed point yield an exactly feasible ϵ-GNE, with an explicit error bound vanishing as network and cluster populations grow.

1.2 Organization and proof strategy

The MFG approach treats the mean field as a fixed deterministic environment and solves the stochastic control problem of a representative agent. Consistency then requires that the population law induced by the optimal control reproduce the prescribed mean field.

Given a candidate mean-field distributional environment, represented by a flow of joint state–control laws, its local and graphon-weighted moments determine the representative-agent problem. With this environment fixed, the inner fixed point enforces moment-parameter consistency, i.e., the path-law ensemble generated by the optimal feedback reproduces the moment parameters inserted into its closed-loop SDE; the conditioning mean-field flow is not updated in this step. The outer update replaces this candidate flow by the joint state–control marginals generated by that path law; under weak coupling, this update is a contraction, yielding the unique fixed point solution to the overall constrained GMFGC problem.

Section 2 gives the model and assumptions. Section 3 establishes feasible-set and Hamiltonian regularity, HJB well-posedness, and feedback estimates. Section 4 proves FPK well-posedness and equilibrium. Section 5 constructs feasible finite-game ϵ-GNE profiles with an explicit error bound.

2 Preliminaries

2.1 Notation

Let $\alpha \in [0,1]$ index graphon nodes and $[0,T]$ the horizon. At node α, the state and control take values in $\mathbb{R}^n$ and $\mathbb{R}^m$, respectively. All filtered probability spaces are complete and satisfy the usual conditions; all Brownian motions are $\mathbb{R}^n$-valued. $|\cdot|$ is used for Euclidean and induced operator norms. For $R > 0$, set $B_R^d = \{z \in \mathbb{R}^d : |z| \le R\}$, omitting d when clear; vector inequalities are componentwise.

Let (E, d_E) be a complete separable metric space, and let $\mathcal{P}(E)$ denote the Borel probability measures on E. For a fixed $x_0 \in E$, define

$$\mathcal{P}_1(E) \coloneqq \left\{ \mu \in \mathcal{P}(E) : \int_E d_E(x, x_0)\, \mu(dx) < \infty \right\}.$$

This definition is independent of the choice of x_0. For $\mu, \nu \in \mathcal{P}_1(E)$, set

$$W_1^E(\mu, \nu) \coloneqq \inf_{\pi \in \Pi(\mu,\nu)} \int_{E \times E} d_E(x, y) \pi(dx, dy),$$

where $\Pi(\mu,\nu)$ is the set of probability measures on $E\times E$ with first marginal μ and second marginal ν (the couplings). The superscript E is suppressed when the underlying space is clear. For a Borel map F, $F_{\#}\mu$ denotes the pushforward of μ.

Let $\mathcal{O}$ be a subset of $\mathbb{R}^d$. For a Lipschitz map $h:\mathcal{O}\to\mathbb{R}^q$, write $\mathrm{Lip}(h) \coloneqq \sup_{x\neq y\in\mathcal{O}} |h(x)-h(y)|/|x-y|$. For $j\in\mathbb{N}_0$ and $\vartheta\in(0,1)$, set

$$[h]_{\vartheta;\mathcal{O}} \coloneqq \sup_{\substack{x,y\in\mathcal{O}\\ x\neq y}} \frac{|h(x)-h(y)|}{|x-y|^{\vartheta}}, \qquad \|h\|_{C^{j+\vartheta}(\mathcal{O})} \coloneqq \sum_{|a|\le j} \|D^a h\|_{\infty;\mathcal{O}} + \sum_{|a|=j} [D^a h]_{\vartheta;\mathcal{O}}.$$

Let $C_b^{j+\vartheta}(\mathcal{O})$ denote the space of functions with finite $C^{j+\vartheta}(\mathcal{O})$ norm.

For an interval $I\subset\mathbb{R}$, let $Q = I\times\mathcal{O}$. The parabolic seminorm and norm are

$$\begin{aligned}
[h]_{\vartheta/2,\vartheta;Q} &\coloneqq \sup_{\substack{(t,x),(s,y)\in Q\\ (t,x)\neq(s,y)}} \frac{|h(t,x)-h(s,y)|}{(|t-s|^{1/2}+|x-y|)^{\vartheta}},\\
\|h\|_{C^{1+\vartheta/2,2+\vartheta}(Q)} &\coloneqq \sum_{2r+|a|\le 2} \|\partial_t^r D_x^a h\|_{\infty;Q} + [\partial_t h]_{\vartheta/2,\vartheta;Q} + \sum_{|a|=2} [D_x^a h]_{\vartheta/2,\vartheta;Q}\\
&\quad + \sum_{|a|=1} \sup_{\substack{x:(t,x),(s,x)\in Q\\ t,s\in I,\ t\neq s}} \frac{|D_x^a h(t,x)-D_x^a h(s,x)|}{|t-s|^{(1+\vartheta)/2}}.
\end{aligned}$$

Let $C^{1+\vartheta/2,2+\vartheta}_{\mathrm{loc}}(Q)$ denote the space of functions with finite $C^{1+\vartheta/2,2+\vartheta}(Q')$ norm for every compact cylinder $Q'\Subset Q$. For $Q_T \coloneqq (0,T)\times\mathbb{R}^n$, set

$$V_2(Q_T) \coloneqq L^\infty(0,T;L^2(\mathbb{R}^n)) \cap L^2(0,T;H^1(\mathbb{R}^n)). \tag{1}$$

Here $H^1(\mathbb{R}^n)$ denotes the first-order Sobolev space, and $H^{-1}(\mathbb{R}^n)$ denotes its dual.

For nonempty compact sets $A,B\subset\mathbb{R}^d$, the Pompeiu–Hausdorff distance is

$$d_H(A,B) \coloneqq \max\bigl\{\sup_{a\in A}\operatorname{dist}(a,B),\ \sup_{b\in B}\operatorname{dist}(b,A)\bigr\}.$$

For a closed convex set $C\subset\mathbb{R}^d$ and $z\in C$, its convex-analysis normal cone is

$$N_C(z) \coloneqq \{w\in\mathbb{R}^d : \langle w, y-z\rangle \le 0 \text{ for every } y\in C\} = \bigcap_{y\in C}\left\{w\in\mathbb{R}^d : \langle w,y-z\rangle\le 0\right\}.$$

For each $(\alpha,t)\in[0,1]\times[0,T]$, write $\mu_{\alpha,t}\in\mathcal{P}_1(\mathbb{R}^n)$ for the state law and $\nu_{\alpha,t}\in\mathcal{P}_1(\mathbb{R}^n\times\mathbb{R}^m)$ for the joint state–control law. The corresponding measure flows are $\mu_\alpha = \{\mu_{\alpha,t}\}_{t\in[0,T]}$ and $\nu_\alpha = \{\nu_{\alpha,t}\}_{t\in[0,T]}$. Their graphon-indexed ensembles are

$$\mu_G = \{\mu_\alpha\}_{\alpha\in[0,1]}, \qquad \nu_G = \{\nu_\alpha\}_{\alpha\in[0,1]}. \tag{2}$$

The time-t sections are denoted by $\mu_G(t)$ and $\nu_G(t)$. Condition (C1) requires $\nu_G(t)$, and hence its state marginal $\mu_G(t)$, to be W_1-continuous in α.

For each α, fix a filtered probability space $(\Omega_\alpha,\mathcal{F}_\alpha,\mathbb{F}_\alpha,\mathbb{P}_\alpha)$ satisfying the usual conditions, where $\mathbb{F}_\alpha = \{\mathcal{F}_{\alpha,t}\}_{t\in[0,T]}$. It carries an $\mathcal{F}_{\alpha,0}$-measurable initial random variable ξ_α and an $\mathbb{F}_\alpha$-Brownian motion B_α. Write $\mathbb{E}_\alpha$ for expectation under $\mathbb{P}_\alpha$ and $\mathcal{L}_\alpha(Y) \coloneqq Y_{\#}\mathbb{P}_\alpha$ for the law of a random variable Y on this space. These nodewise stochastic bases are used for the continuum model. The finite game is instead realized on the common stochastic basis introduced in Section 2.2.1.

2.2 Problem setting

2.2.1 Finite-network, finite-agent constrained game

Let $\mathcal{W} : [0,1]^2 \to \mathbb{R}$ be the graphon kernel. Write $z = (y, v) \in \mathbb{R}^{n+m}$ for a generic joint state–control variable. Let $f_0, f : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^{n+m} \to \mathbb{R}^n$ denote the local and network drift kernels, and let l_0, l be the corresponding running-cost kernels. Let $K \subset \mathbb{R}^m$ be the fixed compact convex control set. For $\lambda = (\lambda^{loc}, \lambda^{net}) \in \mathbb{R}^{2(n+m)}$, representing the local and network joint moments, or their empirical counterparts in the finite game, let $\gamma_\alpha(u, \lambda)$ encode the moment-dependent control constraints.

The finite game motivating the continuum model is defined first. At level k, let G^k be a finite clustered network with N_k agents partitioned into M_k clusters $C_1^k, \dots, C_{M_k}^k$. Let $W^k = (W_{\ell r}^k)_{\ell,r=1}^{M_k}$ be the cluster-level weight matrix and set $n_\ell^k = |C_\ell^k|, n_{\min}^k = \min_{1\le\ell\le M_k} n_\ell^k, N_k = \sum_{\ell=1}^{M_k} n_\ell^k$. Assume $n_{\min}^k \ge 2$. The cells $I_\ell^k = [(\ell-1)/M_k, \ell/M_k)$ partition $[0,1]$, with the last cell closed at 1, and α_ℓ^k is the midpoint of I_ℓ^k. Cluster C_ℓ^k represents I_ℓ^k. If $i \in C_\ell^k$, write $c(i) = \ell$ and $\alpha_i^k = \alpha_\ell^k$.

Let $\mathcal{I}_k := \{1, \dots, N_k\}$ be the agent set. For state and control vectors write $\mathbf{x} = (x_i)_{i\in\mathcal{I}_k} \in (\mathbb{R}^n)^{N_k}, \mathbf{u} = (u_i)_{i\in\mathcal{I}_k} \in K^{N_k}$, and use the product sup norm $\|\mathbf{x}\|_\infty = \max_i |x_i|$, with the analogous notation for control vectors. A strategy profile is denoted by $\mathbf{U} = (U_i)_{i\in\mathcal{I}_k}$, and $\mathbf{U}_{-i}$ denotes the profile with its ith component removed.

Let $q(x,u) = (x,u) \in \mathbb{R}^{n+m}$. For $i \in C_\ell^k$, define the full cluster mean and its agent-i leave-one-out version by

$$\widehat{z}_\ell^{N,k}(t) = \frac{1}{n_\ell^k} \sum_{j\in C_\ell^k} q(X_j^{N,k}(t), U_j^{N,k}(t)),$$

$$\widehat{z}_r^{N,k,-i}(t) = \begin{cases} \dfrac{1}{n_\ell^k - 1} \displaystyle\sum_{j\in C_\ell^k, j\neq i} q(X_j^{N,k}(t), U_j^{N,k}(t)), & r = \ell, \\ \widehat{z}_r^{N,k}(t), & r \neq \ell. \end{cases} \tag{3}$$

The leave-one-out network statistic and complete finite parameter seen by agent i are

$$\widehat{n}_i^{N,k,-i}(t) = \frac{1}{M_k} \sum_{r=1}^{M_k} W_{\ell r}^k \widehat{z}_r^{N,k,-i}(t), \tag{4}$$

$$\widehat{\lambda}_i^{N,k,-i}(t) = \big(\widehat{z}_{c(i)}^{N,k,-i}(t), \widehat{n}_i^{N,k,-i}(t)\big) \in \mathbb{R}^{2(n+m)}. \tag{5}$$

For each k, realize the finite game on a common filtered probability space $(\Omega^{N,k}, \mathcal{F}^{N,k}, \mathbb{F}^{N,k}, \mathbb{P}^{N,k})$, where $\mathbb{F}^{N,k} = \{\mathcal{F}_t^{N,k}\}_{t\in[0,T]}$ is the usual augmentation of the filtration generated by $\{\xi_j, B_j(s) : j \in \mathcal{I}_k, 0 \le s \le t\}$. The variables $\{\xi_i\}_{i\in\mathcal{I}_k}$ are i.i.d. with density p_0, the processes $\{B_i\}_{i\in\mathcal{I}_k}$ are independent Brownian motions, and the two families are mutually independent. Let $\mathbb{E}^{N,k}$ denote expectation under $\mathbb{P}^{N,k}$. The finite dynamics and costs are

$$dX_i^{N,k}(t) = \widehat{f}(X_i^{N,k}(t), U_i^{N,k}(t), \widehat{\lambda}_i^{N,k,-i}(t))dt + \sigma dB_i(t), \qquad X_i^{N,k}(0) = \xi_i, \tag{6}$$

$$J_i^{N,k}(\mathbf{U}) = \mathbb{E}^{N,k}\Big[\int_0^T \widehat{l}(X_i^{N,k}(t), U_i^{N,k}(t), \widehat{\lambda}_i^{N,k,-i}(t))dt + g(X_i^{N,k}(T))\Big]. \tag{7}$$

A finite strategy is a K-valued, $\mathbb{F}^{N,k}$-progressively measurable process. It is admissible for agent i if

$$U_i^{N,k}(t) \in K_{\alpha_i^k}(\widehat{\lambda}_i^{N,k,-i}(t)) \quad \text{for } (dt \otimes \mathbb{P}^{N,k})\text{-a.e. } (t,\omega). \tag{8}$$

Agent i minimizes $J_i^{N,k}(\mathbf{U})$ over its admissible controls.

Assuming the limit exists in the appropriate sense as $M_k \to \infty$, the cluster populations diverge, W^k approximates $\mathcal{W}$, and the empirical laws converge to $\{\nu_{\beta,t}\}_{\beta\in[0,1]}$; the finite sums then converge to graphon integrals, while the finite parameters $\widehat{\lambda}_i^{N,k,-i}$ converge to $\lambda_\alpha[\nu_G]$. Section 5 imposes the approximation hypotheses and proves the vanishing ϵ-GNE bound.

2.2.2 Infinite-network, infinite-population constrained GMFGC

At the infinite population infinite network limit, the representative-agent problem underlying the GMFGC system is specified as follows.

Given a deterministic flow ν_G of joint state–control laws, the representative agent at node α chooses an $\mathbb{F}_\alpha$-progressively measurable control u_α satisfying $u_\alpha(t) \in K_\alpha(\lambda_\alpha[\nu_G](t))$ almost surely for a.e. t. The flow specifies distributions, not the other agents' states or Brownian motions. The controlled dynamics and cost are given by

$$dx_\alpha(t) = \widetilde{f}_\alpha\big(x_\alpha(t), u_\alpha(t), \nu_G(t)\big)dt + \sigma dB_\alpha(t), \qquad x_\alpha(0) = \xi_\alpha, \tag{9}$$

$$J_\alpha^{\nu_G}(u_\alpha) = \mathbb{E}_\alpha\Big[\int_0^T \widetilde{l}_\alpha(x_\alpha(t), u_\alpha(t), \nu_G(t))dt + g(x_\alpha(T))\Big]. \tag{10}$$

Here $\mathcal{L}_\alpha(\xi_\alpha) = p_0(x)\,dx$, with p_0 common across α, and ξ_α is independent of B_α. At consistency, (9) is the representative-agent equation of the graphon McKean–Vlasov limit associated with the finite dynamics (6). Section 5 quantifies this connection through exactly feasible profiles and a vanishing ϵ-GNE error.

For the limiting joint flow, the law-dependent aggregate drift and running costs are

$$\widetilde{f}_\alpha(x,u,\nu_G(t)) = \int_{\mathbb{R}^{n+m}} f_0(x,u,z)\nu_{\alpha,t}(dz) + \int_0^1 \mathcal{W}(\alpha,\beta)\int_{\mathbb{R}^{n+m}} f(x,u,z)\nu_{\beta,t}(dz)d\beta,$$

$$\widetilde{l}_\alpha(x,u,\nu_G(t)) = \int_{\mathbb{R}^{n+m}} l_0(x,u,z)\nu_{\alpha,t}(dz) + \int_0^1 \mathcal{W}(\alpha,\beta)\int_{\mathbb{R}^{n+m}} l(x,u,z)\nu_{\beta,t}(dz)d\beta.$$

The moment vector $\lambda_\alpha[\nu_G](t) \in \mathbb{R}^{2(n+m)}$ has components

$$\lambda_\alpha^{loc}(t) \coloneqq \int_{\mathbb{R}^{n+m}} z\,\nu_{\alpha,t}(dz), \qquad \lambda_\alpha^{net}(t) \coloneqq \int_0^1 \mathcal{W}(\alpha,\beta)\lambda_\beta^{loc}(t)\,d\beta. \tag{11}$$

For $\lambda = (\lambda^{loc}, \lambda^{net}) \in \mathbb{R}^{2(n+m)}$, define the network-dependent feasible set by

$$K_\alpha(\lambda) = \{u \in K : \gamma_\alpha(u,\lambda) \le 0\}. \tag{12}$$

For a prescribed flow ν_G, define the admissible control class at node α by

$$\mathcal{A}_\alpha(\nu_G) \coloneqq \big\{u_\alpha : [0,T]\times\Omega_\alpha \to K : u_\alpha \text{ is } \mathbb{F}_\alpha\text{-progressively measurable}, \\ u_\alpha(t,\omega) \in K_\alpha(\lambda_\alpha[\nu_G](t)) \quad (dt\otimes\mathbb{P}_\alpha)\text{-a.e.}\big\}. \tag{13}$$

For a prescribed flow ν_G, the representative-agent problem at node α, denoted by $(\mathrm{P}_\alpha[\nu_G])$, is to minimize $J_\alpha^{\nu_G}(u_\alpha)$ over $u_\alpha \in \mathcal{A}_\alpha(\nu_G)$, where the state in the cost functional is generated by (9).

Thus (9), (10), and (12) constitute the graphon mean-field limit of the finite-network, finite-population game. The existence and uniqueness of solutions to this class of constrained GMFGC problems shall now be analyzed in terms of the associated formulation of HJB–FPK equations.

2.2.3 Dynamic programming HJB–FPK system

For fixed ν_G, define the control-dependent Hamiltonian integrand at node α and its minimum over the feasible set by

$$\widetilde{H}_\alpha(t,x,u,p,\nu_G) \coloneqq \widetilde{f}_\alpha(x,u,\nu_G(t))\cdot p + \widetilde{l}_\alpha(x,u,\nu_G(t)), \tag{14}$$

$$\mathcal{H}_\alpha(t,x,p,\nu_G) \coloneqq \min_{u\in K_\alpha(\lambda_\alpha[\nu_G](t))} \widetilde{H}_\alpha(t,x,u,p,\nu_G). \tag{15}$$

The minimum is attained under Assumptions (H1) and (H2) by compactness and continuity. The variable $p \in \mathbb{R}^n$ is the HJB costate. In the HJB equation, it is evaluated at $p = \nabla V^\alpha(t,x)$, where $V^\alpha : [0,T] \times \mathbb{R}^n \to \mathbb{R}$ is the value function.

For a candidate equilibrium flow ν_G^*, set $\lambda_\alpha^*(t) \coloneqq \lambda_\alpha[\nu_G^*](t)$ and $\widetilde{H}_\alpha^*(t,x,u,p) \coloneqq \widetilde{H}_\alpha(t,x,u,p,\nu_G^*)$. Let $\rho_\alpha : [0,T] \times \mathbb{R}^n \to [0,\infty)$ denote the state density and $\phi_\alpha^* : [0,T] \times \mathbb{R}^n \to K$ the optimal feedback. The terminal cost $g : \mathbb{R}^n \to \mathbb{R}$ and initial density $p_0 : \mathbb{R}^n \to [0,\infty)$ are common across α and satisfy Assumptions (H7) and (H5), respectively. The associated HJB–FPK characterization seeks ν_G^*, V^α, ρ_α, and ϕ_α^* satisfying

$$\begin{aligned}
&[\mathrm{HJB}](\alpha) \;\; -\partial_t V^\alpha - \frac{\sigma^2}{2}\Delta V^\alpha = \mathcal{H}_\alpha(t,x,\nabla V^\alpha(t,x),\nu_G^*), V^\alpha(T,x) = g(x),\\
&[\mathrm{FPK}](\alpha)\; \partial_t \rho_\alpha + \nabla\cdot\big(\widetilde{f}_\alpha(x,\phi_\alpha^*(t,x),\nu_G^*(t))\rho_\alpha\big) = \frac{\sigma^2}{2}\Delta\rho_\alpha, \rho_\alpha(0,x) = p_0(x),\\
&[\mathrm{BR}](\alpha)\; \phi_\alpha^*(t,x) = \arg\min_{u\in K_\alpha(\lambda_\alpha^*(t))} \widetilde{H}_\alpha^*(t,x,u,\nabla V^\alpha(t,x)),\\
&[\mathrm{CONS}](\alpha)\; \nu_{\alpha,t}^* = (\mathrm{id},\phi_\alpha^*(t,\cdot))_\#\big(\rho_\alpha(t,x)\,dx\big).
\end{aligned} \tag{16}$$

Here $\mathrm{id} : \mathbb{R}^n \to \mathbb{R}^n$ is the identity.

Definition 2.1 (Generalized Nash Equilibrium for a Constrained GMFGC)**.** Subject to conditions yielding existence and uniqueness of solutions to (9)–(12) and (14)–(16), let (ν_G^*, ϕ^*) be the solution pair to the GMFGC equations (16), where, for every graphon node α, (i) the feedback ϕ_α^* minimizes $\widetilde{H}_\alpha^*$ over all controls in the endogenous feasible set $K_\alpha(\lambda_\alpha[\nu_G^*])$, and (ii) the resulting closed-loop joint state–control distribution is ν_α^*. Then, (ν_G^*, ϕ^*) is called a generalized Nash equilibrium of the constrained GMFGC problem determined by (9), (10) and (12).

The justification of this definition is as follows. For each $\alpha \in [0,1]$, let x_α^* be the associated closed-loop state and set $u_\alpha^*(t) \coloneqq \phi_\alpha^*(t, x_\alpha^*(t))$. The two equilibrium properties are consequences of the HJB–FPK system: (i) $[\mathrm{BR}](\alpha)$ selects an admissible feedback from the feasible set. A verification argument then yields optimality over admissible control trajectories. Specifically, under the assumptions of Lemma 3.5, the identity (35) shows that $u_\alpha^* \in \mathcal{A}_\alpha(\nu_G^*)$ solves $(\mathrm{P}_\alpha[\nu_G^*])$; (ii) $[\mathrm{FPK}](\alpha)$ together with $[\mathrm{CONS}](\alpha)$ gives $\nu_{\alpha,t}^* = \mathcal{L}_\alpha((x_\alpha^*(t), u_\alpha^*(t)))$ for $t \in [0,T]$.

The equilibrium is termed generalized Nash, rather than Nash, for two reasons. First, the best-response property holds for an agent at node α when all other agents' actions are those prescribed by the GMFGC solution control laws. Second, the feasible-control correspondence itself depends on the equilibrium law and hence on the strategies of the other agents.

The HJB equation is interpreted classically. However, in general, active-set changes may render the feedback nondifferentiable in x; accordingly, the FPK equation is interpreted in the energy space $V_2(Q_T)$; see Lemmas 3.5 and 4.5.

2.3 Assumptions

Constants are uniform in $\alpha \in [0,1]$, and $C > 0$ is generic.

(H1) [Feasible sets] For every (α,λ), $K_\alpha(\lambda)$ is a nonempty closed convex subset of a fixed compact convex set $K \subset \mathbb{R}^m$.

(H2) [Moment-reduced coefficients] There are maps $\widehat{f} : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^{2(n+m)} \to \mathbb{R}^n$, and $\widehat{l} : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^{2(n+m)} \to \mathbb{R}$ such that, for every admissible joint flow ν_G,

$$\widetilde{f}_\alpha(x,u,\nu_G(t)) = \widehat{f}(x,u,\lambda_\alpha[\nu_G](t)), \quad \widetilde{l}_\alpha(x,u,\nu_G(t)) = \widehat{l}(x,u,\lambda_\alpha[\nu_G](t)).$$

The map $\widehat{f}$ is globally Lipschitz on $\mathbb{R}^n \times K \times \mathbb{R}^{2(n+m)}$ and satisfies $|\widehat{f}(x,u,\lambda)| \le C_f(1+|\lambda|)$. For every $R > 0$, $\widehat{l}$ is bounded and Lipschitz on $\mathbb{R}^n \times K \times B_R^{2(n+m)}$. Both maps are C^2 in u.

For every $R > 0$, their first and second u-derivatives are jointly continuous and bounded on $\mathbb{R}^n \times K \times B_R^{2(n+m)}$, and their first u-derivatives are Lipschitz there. These local constants may depend on R.

(H3) [Strong convexity] For every $P, R > 0$, there is $\kappa_{P,R}^{\mathrm{sc}} > 0$ such that, uniformly for $x \in \mathbb{R}^n$, $p \in B_P^n$, and $\lambda \in B_R^{2(n+m)}$, the map $u \mapsto \widehat{f}(x,u,\lambda)\cdot p + \widehat{l}(x,u,\lambda)$ is strongly convex on K with modulus $\kappa_{P,R}^{\mathrm{sc}}$.

(H4) [Graphon regularity] The kernel $\mathcal{W}$ satisfies:
- (i) $\sup_{\alpha\in[0,1]} \int_0^1 |\mathcal{W}(\alpha,\beta)|\, d\beta \le M_{\mathcal{W}}$ for some $M_{\mathcal{W}} < \infty$;
- (ii) for every $\zeta \in L^\infty(0,1)$, the map $\alpha \mapsto \int_0^1 \mathcal{W}(\alpha,\beta)\zeta(\beta)\, d\beta$ is continuous.

(H5) [Initial law] The nonnegative density $p_0 \in L^2(\mathbb{R}^n)$ has unit mass and satisfies $\int_{\mathbb{R}^n} |x|^2 p_0(x)\, dx < \infty$.

(H6) [Constraint regularity] The constraint system satisfies:
- (i) $K = \{u : k(u) \le 0\}$ for a componentwise convex C^2 map $k : \mathbb{R}^m \to \mathbb{R}^{n_K}$. For each α, $\gamma_\alpha : \mathbb{R}^m \times \mathbb{R}^{2(n+m)} \to \mathbb{R}^{n_c}$ is C^2 in u and C^1 in λ. The mappings $D_u^r\gamma_\alpha$, $r = 0,1,2$, $D_\lambda\gamma_\alpha$, and $D_\lambda D_u\gamma_\alpha$ are jointly continuous in (u,λ). Each component is convex in u.
- (ii) For every $R > 0$, $D_u^r\gamma_\alpha$, $r = 0,1,2$, $D_\lambda\gamma_\alpha$, and $D_\lambda D_u\gamma_\alpha$ are bounded on $K \times B_R^{2(n+m)}$ uniformly in α. Moreover, $D_u^r\gamma_\alpha$ is uniformly continuous in α there, with common modulus $\omega_{\gamma,R}$, for $r = 0,1,2$.
- (iii) Set $G_\alpha(u,\lambda) = (\gamma_\alpha(u,\lambda), k(u)) \in \mathbb{R}^{n_c+n_K}$, so that $K_\alpha(\lambda) = \{u : G_\alpha(u,\lambda) \le 0\}$, and define the active index set as $I_\alpha(u,\lambda) = \{i : G_{\alpha,i}(u,\lambda) = 0\}$; then the uniform linear independence constraint qualification (LICQ), as in [42, Def. 12.4], holds: for every $R > 0$, there is $\kappa_{\mathrm{LICQ},R} > 0$ such that the active gradient Hessian $D_u G_{\alpha,I_\alpha(u,\lambda)}(u,\lambda)$ constitute linearly independent rows and $\sigma_{\min}\big(D_u G_{\alpha,I_\alpha(u,\lambda)}(u,\lambda)\big) \ge \kappa_{\mathrm{LICQ},R}$ whenever $\lambda \in B_R^{2(n+m)}$, $u \in K_\alpha(\lambda)$, and $I_\alpha(u,\lambda) \neq \varnothing$.

(H7) [Parabolic regularity] Suppose $\sigma > 0$ and $g \in C_b^{2+\theta}(\mathbb{R}^n)$ for some $\theta \in (0,1)$. Set $\vartheta = \theta/2$.

(H8) [Contractive weak coupling] The constants C_ϕ and $C_{\mathbf{V}}$ in Lemmas 4.3 and 4.10 can be chosen so that $C_\phi C_{\mathbf{V}} < 1$.

2.4 Admissible measure-flow class

To formulate the global fixed-point map, fix $M_{\mathcal{V}}, C_h > 0$ and consider candidate joint flows ν_G satisfying:

(C1) [Graphon-label continuity] For each $t \in [0,T]$, $\nu_G(t) \in C([0,1]; \mathcal{P}_1(\mathbb{R}^n \times K))$, where $\mathcal{P}_1$ is endowed with the metric W_1.

(C2) [Uniform first-moment bound] $\sup_{\alpha,t} \int_{\mathbb{R}^n\times K} |(x,u)|\, \nu_{\alpha,t}(dx,du) \le M_{\mathcal{V}}$.

(C3) [Temporal Hölder regularity] For all $\alpha \in [0,1]$ and $s,t \in [0,T]$, $W_1(\nu_{\alpha,t}, \nu_{\alpha,s}) \le C_h|t-s|^\vartheta$.

(C1)–(C3) are assumed to hold for a candidate environment flows; consistency with the joint marginal flow generated by the controlled SDE under the best-response feedback will necessarily hold at the solved fixed point. By Condition (C2) and Assumption (H4)(i), $\sup_{\alpha,t} |\lambda_\alpha[\nu_G](t)| \le (1 + M_{\mathcal{W}})M_{\mathcal{V}}$. Fix $R_{\mathcal{V}} > (1 + M_{\mathcal{W}})M_{\mathcal{V}}$; then $\lambda_\alpha[\nu_G](t) \in B_{R_{\mathcal{V}}}^{2(n+m)}$ for all $\nu_G \in \mathcal{V}_{[0,T]}$, $\alpha \in [0,1]$, and $t \in [0,T]$.

Definition 2.2 (Admissible Measure-Flow Class and Metric). $\mathcal{V}_{[0,T]}$ shall denote the class of joint measure flows satisfying Conditions (C1)–(C3), with constants $M_{\mathcal{V}}$ and C_h, equipped with the metric

$$d_{\mathcal{V}}(\nu_G, \widetilde{\nu}_G) := \sup_{\alpha\in[0,1],\, t\in[0,T]} W_1(\nu_{\alpha,t}, \widetilde{\nu}_{\alpha,t}). \tag{17}$$

The space $(\mathcal{V}_{[0,T]}, d_{\mathcal{V}})$ is complete. Indeed, since $\mathbb{R}^n \times K$ is complete and separable, $(\mathcal{P}_1(\mathbb{R}^n \times K), W_1)$ is complete [47, Thm. 6.18]. Hence every $d_{\mathcal{V}}$-Cauchy sequence $(\nu_G^k)_k$ admits a family $\nu_G = (\nu_{\alpha,t})_{\alpha,t}$ such that $d_{\mathcal{V}}(\nu_G^k, \nu_G) \to 0$. It remains to verify that $\nu_G \in \mathcal{V}_{[0,T]}$. For fixed t, the triangle inequality

gives

$$\begin{aligned} W_1(\nu_{\beta,t},\nu_{\alpha,t}) &\le W_1(\nu_{\beta,t},\nu^k_{\beta,t}) + W_1(\nu^k_{\beta,t},\nu^k_{\alpha,t}) + W_1(\nu^k_{\alpha,t},\nu_{\alpha,t}) \\ &\le 2d_{\mathcal{V}}(\nu^k_G,\nu_G) + W_1(\nu^k_{\beta,t},\nu^k_{\alpha,t}). \end{aligned} \tag{18}$$

Since ν^k_G satisfies (C1), taking $\limsup_{\beta\to\alpha}$ and then $k\to\infty$ proves (C1) for ν_G. Since $z\mapsto|z|$ is 1-Lipschitz, (C2) passes to the limit. Similarly, the decomposition in (18) gives $W_1(\nu_{\alpha,t},\nu_{\alpha,s}) \le 2d_{\mathcal{V}}(\nu^k_G,\nu_G) + C_h|t-s|^{\vartheta}$. Letting $k\to\infty$ proves (C3) with the same constant C_h.

Definition 2.3 (Path and Path-Measure Spaces). Let $C_T = C([0,T];\mathbb{R}^n\times K)$, with canonical coordinate process (X_t,U_t) defined by $(X_t(\omega),U_t(\omega)) := \omega(t)$ for $\omega\in C_T$. Write

$$\|\omega\|_{C_T} := \sup_{0\le s\le T}\big(|X_s(\omega)| + |U_s(\omega)|\big).$$

Thus $P\in\mathcal{P}_1(C_T)$ means, equivalently, $\int_{C_T}\|\omega\|_{C_T}\,P(d\omega)<\infty$; changing the origin in this norm gives the same first-moment class. For $\omega,\omega'\in C_T$ and $t\in[0,T]$, define

$$d_t(\omega,\omega') := \sup_{0\le s\le t}\big(|X_s(\omega)-X_s(\omega')| + |U_s(\omega)-U_s(\omega')|\big). \tag{19}$$

Let D_t be the 1-Wasserstein pseudometric induced by d_t; d_T and D_T are metrics. Set $\mathbf{V}^G_T := C([0,1];\mathcal{P}_1(C_T))$. For $\mathbf{P}_G=(P_\alpha)_{\alpha\in[0,1]}$, $\mathbf{Q}_G=(Q_\alpha)_{\alpha\in[0,1]}\in\mathbf{V}^G_T$, define

$$D_t(\mathbf{P}_G,\mathbf{Q}_G) := \sup_{\alpha\in[0,1]} D_t(P_\alpha,Q_\alpha). \tag{20}$$

For $P_\alpha\in\mathcal{P}_1(C_T)$ and $t\in[0,T]$, set $\mathrm{Marg}_t(P_\alpha) := (X_t,U_t)_{\#}P_\alpha$. For $\mathbf{P}_G=(P_\alpha)_{\alpha\in[0,1]}$, set $\mathrm{Marg}(\mathbf{P}_G) := (\mathrm{Marg}_t(P_\alpha))_{(\alpha,t)\in[0,1]\times[0,T]}$.

For each $\mathbf{P}_G\in\mathbf{V}^G_T$, compactness of $[0,1]$ and D_T-continuity yield a first-moment bound uniform in α. Since K is compact, (C_T,d_T) is complete and separable, so $(\mathcal{P}_1(C_T),D_T)$ is complete [47, Thm. 6.18]. Consequently, $(\mathbf{V}^G_T,D_T)$ is complete.

Proof architecture. The nested fixed-point construction is

$$\overbrace{\mathbf{P}_G \xrightarrow{\mathrm{Marg}} \underbrace{\nu_G \xrightarrow{\lambda,\mathrm{HJB}} \phi[\nu_G] \xrightarrow{\mathrm{SDE}} \mathbf{P}^{\nu}_G}_{\substack{\text{moment-parameter consistency;}\\ \nu_G \text{ fixed}}} \xrightarrow{\mathrm{Marg}} \nu^{\mathrm{new}}_G}^{\text{mass–environment update}} \xRightarrow{\text{weak coupling}} \underbrace{\substack{\mathbf{P}^{\nu}_G=\mathbf{P}_G,\\ \nu^{\mathrm{new}}_G=\nu_G}}_{\text{unique global fixed point}}.$$

For $\nu_G=\mathrm{Marg}(\mathbf{P}_G)\in\mathcal{V}_{[0,T]}$, the condition (57) enforces the inner moment-parameter consistency described above, with all constraints retained. The outer condition $\mathbf{P}^{\nu}_G=\mathbf{P}_G$ yields $\nu^{\mathrm{new}}_G=\nu_G$, aligning feasibility, optimality, and the generated law. Lemma 4.9 proves inner consistency; Theorem 4.11 closes the outer loop, whose weak-coupling contraction gives uniqueness.

3 Analysis for a prescribed joint measure flow

This section establishes the best-response side of the construction for a fixed candidate flow. Lemma 3.1 controls the moving feasible sets; Lemmas 3.2–3.5 establish HJB well-posedness and uniform estimates; and Corollary 3.6 yields a regular optimal feedback.

3.1 Feasible-set stability

Lemma 3.1 (Feasible Set Stability). Under Assumptions (H1), (H4), and (H6), for every $R>0$ there exist a constant $C_{K,R}>0$ and a modulus of continuity $\omega_{K,R}$ such that, for all $\alpha,\beta\in[0,1]$ and $\lambda,\eta\in B^{2(n+m)}_R$,

$$d_H\big(K_\alpha(\lambda),K_\beta(\eta)\big) \le C_{K,R}|\lambda-\eta| + \omega_{K,R}(|\alpha-\beta|). \tag{21}$$

Consequently, there exists $C_{\mathcal{V}} > 0$ such that, for all $\nu_G \in \mathcal{V}_{[0,T]}$, $\alpha \in [0,1]$, $s,t \in [0,T]$,

$$d_H\big(K_\alpha(\lambda_\alpha(t)), K_\alpha(\lambda_\alpha(s))\big) \le C_{\mathcal{V}}|t-s|^{\vartheta}. \tag{22}$$

Proof. Fix $R > 0$, and let $\alpha_0 \in [0,1]$, $\lambda_0 \in B_R^{2(n+m)}$, and $u_0 \in K_{\alpha_0}(\lambda_0)$. If the active set $I_0 = I_{\alpha_0}(u_0,\lambda_0)$ is nonempty, put $A_0 = D_u G_{\alpha_0,I_0}(u_0,\lambda_0)$. By the uniform LICQ in Assumption (H6), A_0 has full row rank, so $d_0 = -A_0^\top(A_0A_0^\top)^{-1}\mathbf{1}$ satisfies $A_0 d_0 = -\mathbf{1} < 0$. Thus the Mangasarian–Fromovitz constraint qualification (MFCQ) holds for the finite C^1 inequality system $G_{\alpha_0}(\cdot,\lambda_0) \le 0$ [5, Sec. 2.3.4, Eq. (2.191)]. Here, MFCQ is equivalent to Robinson's Constraint Qualification [5, Def. 2.86, Cor. 2.101].

By the mean value theorem and Assumption (H6), there is $L_{DG,R} > 0$ such that, uniformly for $\alpha,\beta \in [0,1]$, $v,w \in K$, and $\lambda,\eta \in B_R^{2(n+m)}$,

$$\|D_v G_\alpha(v,\lambda) - D_v G_\beta(w,\eta)\| \le L_{DG,R}\big(|v-w| + |\lambda-\eta|\big) + \omega_{\gamma,R}(|\alpha-\beta|). \tag{23}$$

Hence $D_v G_\alpha(v,\lambda)$ is jointly continuous in (α,λ,v). An analogous estimate shows that $G_\alpha(v,\lambda)$ is jointly continuous in the same variables. Consequently, the Stability Theorem for Robinson Constraint Qualification [5, Thm. 2.87], applied with parameter (α,λ) and range $\mathbb{R}_-^{n_c+n_K}$, gives a constant $c > 0$ such that $\operatorname{dist}(v, K_\alpha(\lambda)) \le c|[G_\alpha(v,\lambda)]_+|$ for all (α,λ,v) sufficiently close to (α_0,λ_0,u_0), where $[\cdot]_+$ denotes the componentwise positive part.

The feasible graph $\mathcal{S}_R = \{(\alpha,\lambda,v) \in [0,1] \times B_R^{2(n+m)} \times K : G_\alpha(v,\lambda) \le 0\}$ is compact. A finite cover yields $C'_R > 0$ such that the preceding error bound holds on a relative neighborhood of $\mathcal{S}_R$. On its compact complement, if nonempty, the continuous residual has a minimum $\delta_R > 0$, while Assumption (H1) gives $\operatorname{dist}(v, K_\alpha(\lambda)) \le \operatorname{diam}(K)$. Hence, with $C_R = \max\{C'_R, \operatorname{diam}(K)/\delta_R\}$,

$$\operatorname{dist}(v, K_\alpha(\lambda)) \le C_R|[G_\alpha(v,\lambda)]_+|, \qquad (\alpha,\lambda,v) \in [0,1] \times B_R^{2(n+m)} \times K. \tag{24}$$

Let $\alpha,\beta \in [0,1]$, $\lambda,\eta \in B_R^{2(n+m)}$, and $u \in K_\alpha(\lambda)$. Feasibility and Assumption (H6) give $|[G_\beta(u,\eta)]_+| \le L_{\gamma,R}|\eta-\lambda| + \omega_{\gamma,R}(|\beta-\alpha|)$. Applying (24) at (β,η), taking the supremum over $u \in K_\alpha(\lambda)$, and reversing the parameter pairs proves (21), with $C_{K,R} = C_R L_{\gamma,R}$ and $\omega_{K,R} = C_R\omega_{\gamma,R}$.

By the definition of $R_{\mathcal{V}}$, the parameters $\lambda_\alpha[\nu_G](t)$ belong to $B_{R_{\mathcal{V}}}^{2(n+m)}$. For $e \in \mathbb{R}^{n+m}$ with $|e| \le 1$, set $q_e(z) = e \cdot z$. Since $\sup_{|e|\le 1}|e \cdot a| = |a|$, Kantorovich–Rubinstein (KR) duality [47, Thm. 5.10, Rem. 6.5], Condition (C3), and Assumption (H4)(i) give

$$\begin{aligned}
|\lambda_\alpha[\nu_G](t) - \lambda_\alpha[\nu_G](s)| &\le |\lambda_\alpha^{loc}(t) - \lambda_\alpha^{loc}(s)| + |\lambda_\alpha^{net}(t) - \lambda_\alpha^{net}(s)| \\
&\le (1+M_{\mathcal{W}}) \sup_{\substack{\beta\in[0,1]\\ |e|\le 1}} \left| \int_{\mathbb{R}^{n+m}} q_e(z)(\nu_{\beta,t} - \nu_{\beta,s})(dz) \right| \\
&\le (1+M_{\mathcal{W}}) \sup_{\beta\in[0,1]} W_1(\nu_{\beta,t},\nu_{\beta,s}) \le (1+M_{\mathcal{W}})C_h|t-s|^{\vartheta}.
\end{aligned} \tag{25}$$

Taking $R = R_{\mathcal{V}}$ and $\beta = \alpha$ in (21), and using (25), proves (22), with $C_{\mathcal{V}} = C_{K,R_{\mathcal{V}}}(1+M_{\mathcal{W}})C_h$. □

3.2 Well-posedness of the parameterized constrained HJB equation

Fix an admissible joint flow $\nu_G \in \mathcal{V}_{[0,T]}$ and write

$$\lambda_\alpha(t) = \lambda_\alpha[\nu_G](t), \quad \widehat{f}_\alpha^{\nu}(t,x,u) := \widehat{f}(x,u,\lambda_\alpha(t)), \quad \widehat{l}_\alpha^{\nu}(t,x,u) := \widehat{l}(x,u,\lambda_\alpha(t)). \tag{26}$$

The constrained HJB equation is

$$\begin{aligned}
-\partial_t V^\alpha(t,x) - \frac{\sigma^2}{2}\Delta V^\alpha(t,x) &= \mathcal{H}_\alpha\big(t,x,\nabla V^\alpha(t,x),\nu_G\big), \qquad V^\alpha(T,x) = g(x), \\
\mathcal{H}_\alpha(t,x,p,\nu_G) &:= \min_{u\in K_\alpha(\lambda_\alpha(t))} \left\{ \widehat{f}_\alpha^{\nu}(t,x,u)\cdot p + \widehat{l}_\alpha^{\nu}(t,x,u) \right\}.
\end{aligned} \tag{27}$$

Lemma 3.2 (Regularity of the Reduced Coefficients)**.** Under Assumptions (H1), (H2), and (H4), fix $\nu_G \in \mathcal{V}_{[0,T]}$ and $\widehat{h} \in \{\widehat{f}, \widehat{l}\}$, define $h_\alpha(t,x,u) := \widehat{h}\big(x,u,\lambda_\alpha[\nu_G](t)\big)$. Then, for all $\alpha \in [0,1]$, $s,t \in [0,T]$, $x,y \in \mathbb{R}^n$, $u,v \in K$, and $\lambda,\eta \in B_{R_\mathcal{V}}^{2(n+m)}$,

$$|h_\alpha(t,x,u) - h_\alpha(t,y,v)| \le C\big(|x-y| + |u-v|\big), \tag{28a}$$

$$|h_\alpha(t,x,u) - h_\alpha(s,x,u)| \le C|t-s|^\vartheta, \tag{28b}$$

$$|\widehat{h}(x,u,\lambda) - \widehat{h}(x,u,\eta)| \le C|\lambda - \eta|. \tag{28c}$$

The constant C is uniform over $\alpha \in [0,1]$ and $\nu_G \in \mathcal{V}_{[0,T]}$.

Proof. Estimates (28a) and (28c) are the (x,u)- and λ-Lipschitz bounds in Assumption (H2), respectively; combining the latter with (25) gives (28b). □

Lemma 3.3 (Regularity of the Minimized Hamiltonian)**.** Suppose Assumptions (H1), (H2), (H4), and (H6) hold. There is a constant $C > 0$, independent of C_h, such that, for all $\nu_G \in \mathcal{V}_{[0,T]}$, $\alpha \in [0,1]$, $t \in [0,T]$, $x \in \mathbb{R}^n$, and $p \in \mathbb{R}^n$,

$$|\mathcal{H}_\alpha(t,x,p,\nu_G)| \le C(1+|p|). \tag{29}$$

Moreover, for every $R > 0$, $s,t \in [0,T]$, $x,y \in \mathbb{R}^n$, and $|p|,|q| \le R$,

$$\begin{aligned} &\big|\mathcal{H}_\alpha(t,x,p,\nu_G) - \mathcal{H}_\alpha(s,y,q,\nu_G)\big| \\ &\quad \le C\big[|p-q| + (1+R)\big(|x-y| + (1+C_h)|t-s|^\vartheta\big)\big]. \end{aligned} \tag{30}$$

Proof. Fix $\alpha \in [0,1]$ and $\nu_G \in \mathcal{V}_{[0,T]}$. Set $A_t = K_\alpha(\lambda_\alpha(t))$ and $F_t(x,p,u) = \widehat{f}_\alpha^\nu(t,x,u)\cdot p + \widehat{l}_\alpha^\nu(t,x,u)$. Write $H_t(x,p) := \mathcal{H}_\alpha(t,x,p,\nu_G) = \min_{u\in A_t} F_t(x,p,u)$. The constant $C > 0$ below is generic and uniform in α and ν_G.

Since $|\lambda_\alpha(t)| < R_\mathcal{V}$, Assumption (H2) gives $|\widehat{f}_\alpha^\nu(t,x,u)| + |\widehat{l}_\alpha^\nu(t,x,u)| \le C$. Thus $|F_t(x,p,u)| \le C(1+|p|)$. Taking the minimum gives (29).

For functions F, G on the same compact set A, $|\min_A F - \min_A G| \le \sup_A |F-G|$. The preceding coefficient bound, (28a), and the minimum inequality give $|H_t(x,p) - H_t(x,q)| \le C|p-q|$ and, for $|p| \le R$, $|H_t(x,p) - H_t(y,p)| \le C(1+R)|x-y|$.

For the time estimate, let $u_s \in A_s$ minimize $F_s(x,p,\cdot)$. Lemma 3.1 gives $u_t \in A_t$ such that $|u_t - u_s| \le C(1+C_h)|t-s|^\vartheta$. By minimality, (28a), and (28b), $|H_t(x,p) - H_s(x,p)| \le F_t(x,p,u_t) - F_s(x,p,u_s) \le C(1+R)(1+C_h)|t-s|^\vartheta$. Interchanging s and t gives the reverse inequality. □

Lemma 3.4 (A Priori Estimates for the Value Function)**.** Suppose Assumptions (H1), (H2), (H4), and (H7) hold. There exists $C_1 > 0$ such that, for every $\nu_G \in \mathcal{V}_{[0,T]}$, the corresponding value function of (10) satisfies

$$\sup_{t,x,\alpha} |V^\alpha(t,x)| + \sup_{t,\alpha} \operatorname{Lip}_x V^\alpha(t,\cdot) \le C_1. \tag{31}$$

Proof. For fixed α and ν_G, the admissible control set is independent of the initial state. Using the same admissible control and Brownian motion for initial states x and y, the coefficient Lipschitz bounds, Gronwall's inequality, and taking infima give

$$|V^\alpha(t,x) - V^\alpha(t,y)| \le (TL_l + L_g)e^{L_f T}|x-y|,$$

uniformly in α and ν_G. Bounded running and terminal costs give the L^∞ estimate. □

Lemma 3.5 (Solvability and Second-Order Regularity of the HJB)**.** Let Assumptions (H1)– (H7) hold. Fix $\nu_G \in \mathcal{V}_{[0,T]}$ and set $\vartheta = \theta/2$. For every $\alpha \in [0,1]$, the parabolic HJB problem (27) has a unique bounded classical solution

$$V^\alpha \in C_b([0,T]\times\mathbb{R}^n) \cap C^{1+\vartheta/2,2+\vartheta}_{\mathrm{loc}}([0,T]\times\mathbb{R}^n), \tag{32}$$

with terminal condition $V^\alpha(T,x) = g(x)$ specified in (27). For every $R > 0$,

$$\|V^\alpha\|_{C^{1+\vartheta/2,2+\vartheta}([0,T]\times B_R)} \le C_R. \tag{33}$$

The constants C_R may depend on R and C_h, but are independent of α. Moreover, there is a constant C_0, independent of α, ν_G, and C_h, such that

$$\sup_{\alpha,t,x}\big(|V^\alpha(t,x)| + |V^\alpha_x(t,x)| + \|D^2_x V^\alpha(t,x)\| + |V^\alpha_t(t,x)|\big) \le C_0. \tag{34}$$

In particular, V^α_x is globally Lipschitz in x, uniformly in t and α.

Proof. Fix α and ν_G, set $a = \sigma^2/2$, and abbreviate $H(t,x,p) = \mathcal{H}_\alpha(t,x,p,\nu_G)$. By Lemma 3.3, H is globally Lipschitz in p, locally parabolic ϑ-Hölder in (t,x) on bounded p-sets, and $|H(t,x,0)| \le C$. Thus the semilinear-parabolic argument of [9, Lem. 3.3] applies once the moving constraint has been absorbed into H by Lemma 3.3.

For completeness, set $r = T - t$, $W(r,x) = V^\alpha(T-r,x)$, and let P_r be the heat semigroup for $a\Delta$. The mild map is $P_r g + \int_0^r P_{r-s} H(T-s,\cdot,Dw(s))ds$. The standard estimates $\|P_r h\|_\infty \le \|h\|_\infty$ and $\|DP_r h\|_\infty \le Cr^{-1/2}\|h\|_\infty$ [38, Ch. IV, Secs. 1–2] make it a contraction on bounded uniformly continuous (w, Dw) for small time. Iteration over finitely many intervals gives a unique bounded mild solution on $[0,T]$.

For any $\beta \in (\vartheta, 1)$, the heat-semigroup Hölder estimates [38, Ch. IV, Sec. 2, (2.1)–(2.3)] give, through the mild formula, $DW \in C^{\beta/2,\beta}_{\mathrm{loc}}$. Thus $H(T-r,x,DW(r,x))$ is locally parabolic ϑ-Hölder; the Schauder estimate [38, Ch. IV, Sec. 5, Thm. 5.1] gives, for every $R > 0$, $\|W\|_{C^{1+\vartheta/2,2+\vartheta}([0,T]\times B_R)} \le C_R$, with C_R possibly depending on C_h. Every solution in (32) satisfies the same mild formula, hence is unique.

The barriers $\pm(\|g\|_\infty + Cr)$ and the maximum principle give $\|V^\alpha\|_\infty \le \|g\|_\infty + CT$. Lemma 3.1 and uniqueness of the Hamiltonian minimizer imply, by compactness, that ϕ_α is continuous and hence Borel. Its closed-loop drift is bounded, so Girsanov's theorem [22, Ch. V, Sec. 10, Thm. 10.2] gives a weak solution. For any admissible u with state X^u, localized Itô verification gives

$$J_{t,x}(u) - V^\alpha(t,x) = \mathbb{E}\int_t^T \Big[\widehat{f}^\nu_\alpha(s,X^u_s,u_s)\cdot V^\alpha_x(s,X^u_s) + \widehat{l}^\nu_\alpha(s,X^u_s,u_s) - \mathcal{H}_\alpha(s,X^u_s,V^\alpha_x(s,X^u_s),\nu_G)\Big]ds \ge 0. \tag{35}$$

The integrand vanishes for $u_s = \phi_\alpha(s,X^u_s \mid \nu_G)$, so V^α is the value function. Lemma 3.4 gives $\sup_{\alpha,t,x}|V^\alpha_x(t,x)| \le C_1$, where C_1 is independent of ν_G and C_h.

For the global Hessian bound, smoothly approximate H, g while preserving their spatial Lipschitz and terminal Hessian bounds. Local Schauder compactness [24, Lemma 6.33] and uniqueness give local $C^{1,2}$ convergence of the solutions W^k to W. Put $M_k(r) = \|D^2_x W^k(r,\cdot)\|_\infty$. The cancellation $\int D^2 p_r = 0$ yields $\|D^2 P_r h\|_\infty \le Cr^{-1/2}\,\mathrm{Lip}(h)$ [38, Ch. IV, Sec. 2, (2.5)]. Since $\mathrm{Lip}_x H^k(T-s,\cdot,DW^k(s,\cdot)) \le C(1+M_k(s))$, the mild formula gives

$$M_k(r) \le \|D^2 g^k\|_\infty + C\int_0^r (r-s)^{-1/2}\big(1 + M_k(s)\big)\,ds. \tag{36}$$

Iterating once and using $\int_s^r (r-q)^{-1/2}(q-s)^{-1/2}\,dq \le C$ reduces this to Gronwall, so $\sup_{k,r} M_k(r) \le C$ independently of C_h. Passing to the limit gives the same Hessian bound for W. The equation then bounds W_r, proving (34); the Hessian bound makes V^α_x globally Lipschitz. □

Corollary 3.6 (Regularity of the Optimal Feedback Law)**.** Let Assumptions (H1)–(H7) hold. For every $\nu_G \in \mathcal{V}_{[0,T]}$, the feedback

$$\phi_\alpha(t,x|\nu_G) = \arg\min_{u\in K_\alpha(\lambda_\alpha(t))} \{\widehat{f}^{\nu}_\alpha(t,x,u)\cdot V^\alpha_x(t,x) + \widehat{l}^{\nu}_\alpha(t,x,u)\} \tag{37}$$

is well defined, bounded, Borel measurable, and globally Lipschitz in x: there is $L_\phi > 0$ such that, for all $\alpha \in [0,1]$, $\nu_G \in \mathcal{V}_{[0,T]}$, $t \in [0,T]$, and $x, y \in \mathbb{R}^n$,

$$|\phi_\alpha(t,x|\nu_G) - \phi_\alpha(t,y|\nu_G)| \le L_\phi |x-y|. \tag{38}$$

Proof. Compactness and strong convexity give a unique minimizer; the continuity argument in the proof of Lemma 3.5 shows that the feedback is continuous, hence Borel. For fixed t, let u_x, u_y denote the minimizers at x, y. The feasible set is common to x, y, so we may test the two variational inequalities with u_y, u_x, respectively. With $d = u_x - u_y$, their sum and strong convexity with modulus $\kappa_{\mathrm{sc}} > 0$ give

$$\kappa_{\mathrm{sc}}|d|^2 \le C\big(|x-y| + |V^\alpha_x(t,x) - V^\alpha_x(t,y)|\big)|d|.$$

For $d \neq 0$, divide by $|d|$ and use the global Lipschitz bound for V^α_x from Lemma 3.5 to obtain $|u_x - u_y| \le C(1 + \operatorname{Lip} V^\alpha_x)|x-y|/\kappa_{\mathrm{sc}}$. All constants are uniform over the admissible flow class, which proves the asserted uniformity. □

4 Existence of the consistent joint measure flow

A perturbation of ν_G affects the optimization through the feasible set $K_\alpha(\lambda_\alpha)$, the Hamiltonian coefficients, and the HJB costate V^α_x. The analysis in this section controls these three effects and closes the fixed-point argument. The first subsection establishes KKT and costate sensitivity; the second proves FPK and path-law stability; the final subsection applies the Contraction Principle to the global consistency map.

4.1 Stability of the feasible sets and the minimized Hamiltonian

Lemma 4.1 (KKT Stability for the Constrained Hamiltonian)**.** Under Assumptions (H1)–(H6), the pointwise problem defining (37) has a unique minimizer. Let $\mu_\alpha = (\mu^\gamma_\alpha, \mu^K_\alpha) \in \mathbb{R}^{n_c+n_K}_+$ denote the associated multipliers for the moving constraints γ_α and the fixed constraints k, respectively. Write $\xi = (x,p,\lambda)$. For every $P, R > 0$, there is $L_{P,R} > 0$, independent of α, such that

$$|(u_\alpha,\mu_\alpha)(\xi) - (u_\alpha,\mu_\alpha)(\xi')| \le L_{P,R}|\xi - \xi'|, \quad \forall \xi,\xi' \in \mathbb{R}^n \times B^n_P \times B^{2(n+m)}_R.$$

Proof. Put $\xi = (x,p,\lambda)$, $\ell_\xi(u) := \widehat{f}(x,u,\lambda)\cdot p + \widehat{l}(x,u,\lambda)$, and $A_\xi := \{u : G_\alpha(u,\lambda) \le 0\}$. The set A_ξ is nonempty, compact, and convex; strong convexity of ℓ_ξ therefore gives a unique minimizer. Full-constraint LICQ gives a unique multiplier $\mu = (\mu^\gamma, \mu^K) \ge 0$. Because either the moving or fixed constraints may be active, LICQ is applied to the active rows of $G_\alpha = (\gamma_\alpha, k)$. The KKT conditions are

$$D_u\ell_\xi(u) + D_uG_\alpha(u,\lambda)^\top\mu = 0,\ G_\alpha(u,\lambda) \le 0,\ \mu_iG_{\alpha,i}(u,\lambda) = 0,\ \forall i = 1,\dots,n_c+n_K. \tag{39}$$

If I is the active set, then $D_uG^\top_{\alpha,I}\mu_I = -D_u\ell_\xi$. The uniform LICQ modulus and the bound on $D_u\ell_\xi$ for $|p| \le P$ hence give a uniform multiplier bound.

With $z := (u,\mu)$, the KKT system is the generalized equation

$$0 \in F_\alpha(z,\xi) + N_{\mathbb{R}^m\times\mathbb{R}^{n_c+n_K}_+}(z), \qquad F_\alpha(z,\xi) = \begin{pmatrix} D_u\ell_\xi(u) + D_uG_\alpha(u,\lambda)^\top\mu \\ -G_\alpha(u,\lambda)\end{pmatrix}. \tag{40}$$

Let $\mathcal{L}_\xi(u,\mu) := \ell_\xi(u)+\langle \mu, G_\alpha(u,\lambda)\rangle$. Convexity of the constraints and $\mu \geq 0$ imply $\langle D^2_{uu}\mathcal{L}_\xi(u,\mu)h, h\rangle \geq \kappa^{\mathrm{sc}}_{P,R}|h|^2$ for every $h \in \mathbb{R}^m$. Since this estimate holds in all directions, it implies the second-order condition [5, eq. (5.83)]. Combined with LICQ, [5, Prop. 5.38] then gives strong regularity [5, Def. 5.12] of (40).

Fix $P, R > 0$. The preceding bound and the uniform convexity, LICQ, and derivative bounds make the strong-regularity radii and moduli independent of x and α. On these neighborhoods, $|F_\alpha(z,\xi') - F_\alpha(z,\xi)| \leq C_{P,R}|\xi' - \xi|$. Robinson's implicit-function theorem [44, Thm. 2.1 and Cor. 2.2] therefore gives a locally single-valued map $z_\alpha(\xi) = (u_\alpha(\xi), \mu_\alpha(\xi))$ and constants $r_{P,R}, L_{P,R} > 0$, independent of x and α, such that $|z_\alpha(\xi') - z_\alpha(\xi)| \leq L_{P,R}|\xi' - \xi|$ whenever $|\xi' - \xi| < r_{P,R}$. Subdividing the segment between two parameters into pieces shorter than $r_{P,R}$ and summing these local estimates proves the asserted bound. Uniqueness makes the local branches agree. □

Lemma 4.2 (Costate Stability). Let Assumptions (H1)–(H7) hold. Let V^α and $\widetilde{V}^\alpha$ be the HJB solutions generated by flows $\nu_G, \widetilde{\nu}_G \in \mathcal{V}_{[0,T]}$. Then there exists $C_p > 0$, independent of α, the two flows, and C_h, such that

$$\sup_{\alpha\in[0,1]} \sup_{t\in[0,T]} \sup_{x\in\mathbb{R}^n} |V^\alpha_x(t,x) - \widetilde{V}^\alpha_x(t,x)| \leq C_p d_{\mathcal{V}}(\nu_G, \widetilde{\nu}_G). \tag{41}$$

Proof. Let $\delta = d_{\mathcal{V}}(\nu_G, \widetilde{\nu}_G)$ and set $a := \sigma^2/2$. Write $H^\nu_\alpha(t,x,p) = \mathcal{H}_\alpha(t,x,p,\nu_G)$ and define $H^{\widetilde{\nu}}_\alpha$ analogously. By (11) and Assumption (H4)(i), $\sup_{\alpha,t}|\lambda_\alpha[\nu_G](t) - \lambda_\alpha[\widetilde{\nu}_G](t)| \leq (1 + M_{\mathcal{W}})\delta$. For a minimizer under ν_G, Lemma 3.1 selects a feasible comparison point under $\widetilde{\nu}_G$ at distance $O(\delta)$; applying this construction in both directions and using Assumption (H2) gives, on the common bounded costate range,

$$\sup_{\alpha,t,x,|p|\leq P_0} |H^\nu_\alpha(t,x,p) - H^{\widetilde{\nu}}_\alpha(t,x,p)| \leq C\delta. \tag{42}$$

Set $Z = V^\alpha - \widetilde{V}^\alpha$, $q(t,x) = \nabla V^\alpha(t,x)$, and $\widetilde{q}(t,x) = \nabla \widetilde{V}^\alpha(t,x)$. Define

$$b(t,x) := \begin{cases} \dfrac{H^\nu_\alpha(t,x,q(t,x)) - H^\nu_\alpha(t,x,\widetilde{q}(t,x))}{|q(t,x) - \widetilde{q}(t,x)|^2}\big(q(t,x) - \widetilde{q}(t,x)\big), & q(t,x) \neq \widetilde{q}(t,x), \\ 0, & q(t,x) = \widetilde{q}(t,x), \end{cases}$$

$$S(t,x) := H^\nu_\alpha(t,x,\widetilde{q}(t,x)) - H^{\widetilde{\nu}}_\alpha(t,x,\widetilde{q}(t,x)).$$

Subtracting the two HJB equations and using these definitions gives

$$-Z_t - a\Delta Z - b(t,x)\cdot\nabla Z = S(t,x), \qquad Z(T,\cdot) = 0. \tag{43}$$

Boundedness of $\widehat{f}$ on the relevant parameter range implies that H^ν_α is globally Lipschitz in p, and hence $\|b\|_\infty \leq C$. Moreover, (42) gives $\|S\|_\infty \leq C\delta$. Apply the Duhamel and heat-semigroup argument from the proof of Lemma 3.5 to (43). With $G(t) = \|\nabla Z(t)\|_\infty$ and $(P_r)_{r\geq 0}$ denoting the heat semigroup generated by $a\Delta$, Duhamel's formula gives

$$Z(t) = \int_t^T P_{s-t}\big(b(s,\cdot)\cdot\nabla Z(s) + S(s,\cdot)\big)\,ds,$$

The heat-semigroup bound used above, namely $\|\nabla P_r h\|_\infty \leq Cr^{-1/2}\|h\|_\infty$, yields $G(t) \leq C\int_t^T (s-t)^{-1/2}(G(s)+\delta)\,ds$. Iterating once, as in the argument after (36) in Lemma 3.5, and using $\int_t^r (s-t)^{-1/2}(r-s)^{-1/2}ds = \pi$, gives $G(t) \leq C\delta + C\int_t^T G(r)\,dr$. Gronwall's inequality therefore gives $\sup_t G(t) \leq C_p\delta$, where C_p is independent of α, the two flows, and C_h. □

Lemma 4.3 (Sensitivity of the Best Response). Let Assumptions (H1)–(H7) hold, and let $\nu_G, \widetilde{\nu}_G \in \mathcal{V}_{[0,T]}$. Let $\phi_\alpha(\cdot,\cdot|\nu_G)$ and $\phi_\alpha(\cdot,\cdot|\widetilde{\nu}_G)$ be the feedback laws generated by the corresponding HJB equations. Then there is $C_\phi > 0$ such that

$$\sup_{\alpha\in[0,1]} \sup_{t\in[0,T]} \sup_{x\in\mathbb{R}^n} |\phi_\alpha(t,x|\nu_G) - \phi_\alpha(t,x|\widetilde{\nu}_G)| \leq C_\phi d_{\mathcal{V}}(\nu_G, \widetilde{\nu}_G). \tag{44}$$

Proof. By Lemma 4.1, there are constants L_p^ϕ, L_λ^ϕ with

$$|\phi_\alpha(t,x|\nu_G) - \phi_\alpha(t,x|\widetilde{\nu}_G)| \le L_p^\phi |V_x^\alpha(t,x) - \widetilde{V}_x^\alpha(t,x)| + L_\lambda^\phi |\lambda_\alpha[\nu_G](t) - \lambda_\alpha[\widetilde{\nu}_G](t)|. \tag{45}$$

By Lemma 4.2 and (4.1), the conclusion holds with $C_\phi := L_p^\phi C_p + (1 + M_{\mathcal{W}})L_\lambda^\phi$. In particular, C_ϕ is independent of C_h. □

Lemma 4.4 (Temporal Hölder Stability of the Feedback Law)**.** Under Assumptions (H1)–(H7), there exists a constant $A_\phi > 0$, independent of α, ν_G, and C_h, such that

$$|\phi_\alpha(t,x|\nu_G) - \phi_\alpha(s,x|\nu_G)| \le (A_\phi + C_\phi C_h)|t-s|^\vartheta. \tag{46}$$

The feedback is also bounded and globally Lipschitz in x, uniformly in α.

Proof. The spatial properties follow from Corollary 3.6. Put $p_\alpha = V_x^\alpha$ and $\lambda_\alpha = \lambda_\alpha[\nu_G]$. The difference-quotient argument and (34) give $|p_\alpha(t,x) - p_\alpha(s,x)| \le A_p|t-s|^\vartheta$, where A_p is independent of C_h. Conditions (C3) and (H4)(i) give $|\lambda_\alpha(t) - \lambda_\alpha(s)| \le (1 + M_{\mathcal{W}})C_h|t-s|^\vartheta$. Applying the optimizer estimate used in (45) and the definition of C_ϕ proves (46) with $A_\phi = L_p^\phi A_p$, since

$$|\phi_\alpha(t,x|\nu_G) - \phi_\alpha(s,x|\nu_G)| \le \big(L_p^\phi A_p + C_\phi C_h\big)|t-s|^\vartheta. \qquad \square$$

Lemma 4.5 (Generalized FPK Well-Posedness)**.** Under Assumptions (H1)– (H7), fix $\nu_G \in \mathcal{V}_{[0,T]}$ and set

$$b_\alpha(t,x) = \widehat{f}\big(x, \phi_\alpha(t,x|\nu_G), \lambda_\alpha[\nu_G](t)\big).$$

With initial density p_0, the FPK equation

$$\partial_t \rho_\alpha - \frac{\sigma^2}{2}\Delta\rho_\alpha + \nabla\cdot(b_\alpha\rho_\alpha) = 0 \quad \text{in } Q_T \tag{47}$$

has a unique generalized solution $\rho_\alpha \in V_2(Q_T)$. Its $C([0,T];L^2(\mathbb{R}^n))$ representative is nonnegative, has unit mass at every time, and satisfies energy bounds uniform in α.

Proof. The drift b_α is bounded and Borel, uniformly in α. In the energy formulation of [38, Ch. III, Sec. 5], the solution obeys, for every $\psi \in C_c^\infty([0,T)\times\mathbb{R}^n)$,

$$\int_{Q_T}\left(-\rho_\alpha\partial_t\psi + \frac{\sigma^2}{2}\nabla\rho_\alpha\cdot\nabla\psi - b_\alpha\rho_\alpha\cdot\nabla\psi\right)dxdt = \int_{\mathbb{R}^n} p_0(x)\psi(0,x)dx. \tag{48}$$

Uniform ellipticity and boundedness of the lower-order coefficient give existence and uniqueness by [38, Ch. III, Sec. 5, Thm. 5.2]. For completeness, take smooth approximations b_α^r, p_0^r preserving the uniform drift bound, nonnegativity, unit mass, and the L^2 and second-moment bounds. Testing the smooth equation with ρ_α^r and using Young's inequality give

$$\frac{1}{2}\frac{d}{dt}\|\rho_\alpha^r\|_2^2 + \frac{\sigma^2}{4}\|\nabla\rho_\alpha^r\|_2^2 \le C\|\rho_\alpha^r\|_2^2, \qquad \sup_{t\le T}\|\rho_\alpha^r(t)\|_2^2 + \int_0^T \|\nabla\rho_\alpha^r\|_2^2 dt \le C. \tag{49}$$

The equation also bounds $\partial_t\rho_\alpha^r$ in $L^2(0,T;H^{-1})$. Weak and local compactness permit passage to (48), while the Lions–Magenes time-continuity theorem, namely the continuous Gelfand-triple embedding $L^2(0,T;H^1)\cap H^1(0,T;H^{-1}) \hookrightarrow C([0,T];L^2)$, gives the continuous representative [19, Sec. 5.9.2, Thm. 3, p. 287]; the difference estimate proves uniqueness.

Testing with the negative part gives $\rho_\alpha^r \ge 0$. Testing with truncated versions of $|x|^2$ and applying the Gronwall's inequality yield

$$\sup_{r,\alpha,t}\int_{\mathbb{R}^n}|x|^2\rho_\alpha^r(t,x)dx \le C\left(1 + \int_{\mathbb{R}^n}|x|^2 p_0(x)dx\right). \tag{50}$$

Thus the limit is nonnegative. Finally test against cutoffs $\chi_R \uparrow 1$ with $|\nabla\chi_R| \le C/R$ and $|\Delta\chi_R| \le C/R^2$, first passing $r\to\infty$ and then $R\to\infty$. The moment bound eliminates escape of mass, while the drift and diffusion errors vanish. Thus the representative has unit mass at every time. □

4.2 Generalized McKean–Vlasov map

For a fixed candidate mean-field distributional environment ν_G and its constrained feedback, $\Phi_{\mathbf{V}}$ maps a path-law ensemble to the closed-loop law generated by its induced moments. Its fixed points satisfy inner moment-parameter consistency. Such a fixed point is a GMFGC equilibrium exactly when the outer condition also holds: $\nu_G^{\mathrm{new}} := \mathrm{Marg}(\mathbf{P}_G^{\nu}) = \nu_G$.

Definition 4.6 (Generalized McKean–Vlasov Mapping on Path Space)**.** Fix $\nu_G \in \mathcal{V}_{[0,T]}$ and its feedback $\phi_\alpha(t,x|\nu_G)$ from (37). For $\mathbf{P}_G = (P_\alpha)_\alpha \in \mathbf{V}_T^G$, set, for every (α,t),

$$\mu_{\alpha,t}^{\mathbf{P}} = \mathrm{Marg}_t^x(P_\alpha), \qquad \overline{\nu}_{\alpha,t}^{\mathbf{P}} = \big(\mathrm{id}, \phi_\alpha(t,\cdot|\nu_G)\big)_{\#}\mu_{\alpha,t}^{\mathbf{P}}, \qquad \overline{\lambda}_\alpha^{\mathbf{P}}(t) = \lambda_\alpha[\overline{\nu}_G^{\mathbf{P}}](t). \tag{51}$$

For each $\alpha \in [0,1]$, let $X_\alpha : [0,T]\times\Omega_\alpha \to \mathbb{R}^n$ be the unique $\mathbb{F}_\alpha$-adapted strong solution, with $X_\alpha(0) = \xi_\alpha$, of

$$dX_\alpha(t) = \widehat{f}\big(X_\alpha(t), \phi_\alpha(t, X_\alpha(t)|\nu_G), \overline{\lambda}_\alpha^{\mathbf{P}}(t)\big)dt + \sigma dB_\alpha(t), \qquad t\in[0,T],$$

then the associated control process is defined by

$$U_\alpha(t) = \phi_\alpha(t, X_\alpha(t)|\nu_G), \qquad t\in[0,T], \tag{52}$$

and $\Phi_{\mathbf{V}}[\mathbf{P}_G \mid \nu_G] := \big(\mathcal{L}(X_\alpha(\cdot), U_\alpha(\cdot))\big)_{\alpha\in[0,1]}$, with $\Phi_{\mathbf{V}}[\,\cdot \mid \nu_G] : \mathbf{V}_T^G \to \big(\mathcal{P}_1(C_T)\big)^{[0,1]}$. □

Definition 4.6 uses strong solutions on the prescribed nodewise stochastic bases; the globally Lipschitz closed-loop drift gives the stated unique adapted solutions. Lemma 4.9 proves that $\Phi_{\mathbf{V}}[\mathbf{V}_T^G \mid \nu_G] \subseteq \mathbf{V}_T^G$, so the map is a self-map on $\mathbf{V}_T^G$, but equality of the two sets is not asserted.

Lemma 4.7 (Continuity in the Node Label)**.** Under Assumptions (H1)–(H7), fix $\nu_G \in \mathcal{V}_{[0,T]}$. For every $R>0$, as $\beta\to\alpha$,

$$\sup_{t\le T} |\lambda_\beta[\nu_G](t) - \lambda_\alpha[\nu_G](t)| \longrightarrow 0, \tag{53a}$$

$$\sup_{\substack{t\le T\\ x\in B_R}} \big(|V^\beta(t,x) - V^\alpha(t,x)| + |V_x^\beta(t,x) - V_x^\alpha(t,x)|\big) \longrightarrow 0, \tag{53b}$$

$$\sup_{\substack{t\le T\\ x\in B_R}} |\phi_\beta(t,x|\nu_G) - \phi_\alpha(t,x|\nu_G)| \longrightarrow 0. \tag{53c}$$

Proof. Put $q_\gamma(t) = \int_{\mathbb{R}^{n+m}} z\,\nu_{\gamma,t}(dz)$. The uniform moment bound makes $q_\gamma(t)$ bounded, and Condition (C1) makes $\gamma\mapsto q_\gamma(t)$ continuous. At each grid point t_j, the local moment term converges by Condition (C1), while the graphon term converges by Assumption (H4)(ii). Consequently $|\lambda_\beta(t_j)-\lambda_\alpha(t_j)|\to 0$. On a time grid of mesh h, the common temporal modulus then implies

$$\sup_{t\le T}|\lambda_\beta(t)-\lambda_\alpha(t)| \le \max_j |\lambda_\beta(t_j)-\lambda_\alpha(t_j)| + 2Ch^\vartheta.$$

First let $\beta\to\alpha$ and then $h\downarrow 0$. The reduced coefficients and feasible sets consequently converge locally uniformly. For every $R>0$, the uniform Schauder estimate of Lemma 3.5 bounds $\{V^\beta\}_\beta$ in $C^{1+\vartheta/2,2+\vartheta}([0,T]\times B_{R+1})$. Hence, for every $0<\vartheta'<\vartheta$, the compact embedding of parabolic Hölder spaces yields a subsequence converging in $C^{1+\vartheta'/2,2+\vartheta'}([0,T]\times B_R)$. Passing to the limit in the HJB equation, using the local-uniform convergence of the reduced coefficients and feasible sets, shows that the subsequential limit solves the HJB equation at node α. Uniqueness identifies this limit with V^α and removes the need to extract a subsequence. Consequently, $V^\beta\to V^\alpha, V_x^\beta\to V_x^\alpha$ locally uniformly as $\beta\to\alpha$. Finally, the KKT stability estimate of Lemma 4.1 gives, on $[0,T]\times B_R$,

$$|\phi_\beta-\phi_\alpha| \le C_R\big(|\lambda_\beta-\lambda_\alpha| + |V_x^\beta - V_x^\alpha| + \omega_{\gamma,R}(|\beta-\alpha|)\big).$$

Assumption (H6) and uniform strong regularity imply the third limit in (53). □

Because the next lemma concerns only path laws and their marginals, a weak realization of the closed-loop SDE is sufficient for its estimates.

Lemma 4.8 (Self-Mapping Regularity of the Consistent Marginal Flow)**.** Let Assumptions (H1)–(H7) hold, and fix $\nu_G \in \mathcal{V}_{[0,T]}$. Suppose $\mathbf{P}_G = (P_\alpha)_{\alpha\in[0,1]} \in \mathbf{V}_T^G$ is a consistent closed-loop path ensemble for this fixed environment, in the sense that, with $\nu_G^{new} := \mathrm{Marg}(\mathbf{P}_G)$, the canonical process (X_α, U_α) on C_T under P_α is a weak solution, i.e., admits a realization on a filtered probability space carrying an $\mathbb{R}^n$-valued Brownian motion B_α, of

$$\begin{aligned} dX_\alpha(t) &= \widehat{f}\big(X_\alpha(t), \phi_\alpha(t, X_\alpha(t)|\nu_G), \lambda_\alpha[\nu_G^{new}](t)\big)dt + \sigma dB_\alpha(t), \\ X_\alpha(0) &\sim p_0, \qquad U_\alpha(t) = \phi_\alpha(t, X_\alpha(t)|\nu_G). \end{aligned} \tag{54}$$

Then there are constants $M_0, A_0 > 0$, independent of C_h, such that

$$\sup_{\alpha,t} \int_{\mathbb{R}^n\times K} |(x,u)|\, \nu_{\alpha,t}^{new}(dx,du) \le M_0, \quad W_1(\nu_{\alpha,t}^{new}, \nu_{\alpha,s}^{new}) \le \big(A_0 + C_\phi C_h\big)|t-s|^\vartheta. \tag{55}$$

And then, under the condition $C_\phi < 1$, $\nu_G^{new} \in \mathcal{V}_{[0,T]}$ whenever $M_\mathcal{V} \ge M_0$ and $C_h \ge A_0/(1-C_\phi)$.

Proof. For each α, fix a weak realization of (54). The following estimates depend only on its law.

Conditions (C2) and (C3). We first verify the moment and temporal conditions. For $t \le T$, set $Q(t) = \sup_\alpha \mathbb{E}\sup_{r\le t}|X_\alpha(r)|^2$. The moment definition, Assumption (H4)(i), and compactness of K give the bound $C(1+Q(t))$ for the squared moment parameter; hence the global drift growth in Assumption (H2) and the Burkholder–Davis–Gundy inequality give $Q(t) \le C(1+\int_0^t Q(s)ds)$. Gronwall's lemma and Assumption (H5) yield, independently of C_h,

$$\sup_\alpha \mathbb{E}\sup_{t\le T}|X_\alpha(t)|^2 + \sup_{\alpha,t}\mathbb{E}|U_\alpha(t)|^2 \le C.$$

This gives the first estimate in (55) for some M_0 independent of C_h. Moreover, $\mathbb{E}|X_\alpha(t) - X_\alpha(s)| \le C_T|t-s|^\vartheta$ since $\vartheta \le 1/2$. By spatial Lipschitz continuity and (46),

$$\mathbb{E}|U_\alpha(t) - U_\alpha(s)| \le L_\phi \mathbb{E}|X_\alpha(t) - X_\alpha(s)| + (A_\phi + C_\phi C_h)|t-s|^\vartheta. \tag{56}$$

Coupling the two joint marginals along the same path therefore gives

$$W_1(\nu_{\alpha,t}^{new}, \nu_{\alpha,s}^{new}) \le (A_0 + C_\phi C_h)|t-s|^\vartheta,$$

where A_0 is independent of C_h. If $C_\phi < 1$, the choices stated in the lemma make this bound at most $C_h|t-s|^\vartheta$.

Condition (C1). This follows directly from $\mathbf{P}_G \in \mathbf{V}_T^G$. Indeed, the time-$t$ evaluation map from (C_T, d_T) to $\mathbb{R}^n \times K$ is 1-Lipschitz. Hence, for every $t \in [0,T]$,

$$W_1(\nu_{\beta,t}^{new}, \nu_{\alpha,t}^{new}) \le D_T(P_\beta, P_\alpha) \longrightarrow 0 \qquad \text{as } \beta \to \alpha,$$

which is Condition (C1). □

Lemma 4.9 (Existence and Uniqueness of the Consistent Path Law)**.** Under Assumptions (H1)–(H7), for any fixed $\nu_G \in \mathcal{V}_{[0,T]}$, the map $\Phi_\mathbf{V}[\cdot|\nu_G]$ admits a unique fixed point $\mathbf{P}_G^\nu \in \mathbf{V}_T^G$; that is,

$$\Phi_\mathbf{V}[\mathbf{P}_G^\nu \mid \nu_G] = \mathbf{P}_G^\nu \quad \text{in } (\mathbf{V}_T^G, D_T). \tag{57}$$

Proof. *Self-map.* Let $\mathbf{P}_G \in \mathbf{V}_T^G$. The D_T-continuous image of $[0,1]$ is W_1-compact and hence has uniformly integrable path norms. Together with compactness of K and Lemma 4.7, a standard truncation argument gives $\sup_{t\le T} W_1(\overline{\nu}_{\beta,t}^\mathbf{P}, \overline{\nu}_{\alpha,t}^\mathbf{P}) \to 0$ as $\beta \to \alpha$. For $q_\gamma^\mathbf{P}(t) = \int z\, \overline{\nu}_{\gamma,t}^\mathbf{P}(dz)$, Assumption (H4)(ii) and dominated convergence then give $\int_0^T |\overline{\lambda}_\beta^\mathbf{P}(t) - \overline{\lambda}_\alpha^\mathbf{P}(t)|dt \to 0$. Couple the output SDEs with the

same initial state and Brownian motion. The preceding parameter convergence, the Lipschitz estimates, truncation, and Gronwall give convergence of both state and control paths in L^1 of the supremum norm. The same moment estimate gives a uniform first moment; hence $\Phi_{\mathbf{V}}[\mathbf{P}_G|\nu_G] \in \mathbf{V}_T^G$.

Contractive iterate. For $\mathbf{P}_G, \mathbf{Q}_G \in \mathbf{V}_T^G$, the map $x \mapsto (x, \phi_\alpha(t,x|\nu_G))$ is $(1+L_\phi)$-Lipschitz. Hence the induced marginal and parameter differences are bounded by $CD_t(\mathbf{P}_G, \mathbf{Q}_G)$. A synchronous coupling and the Lipschitz drift give

$$D_t\big(\Phi_{\mathbf{V}}[\mathbf{P}_G|\nu_G], \Phi_{\mathbf{V}}[\mathbf{Q}_G|\nu_G]\big) \le C\int_0^t D_s(\mathbf{P}_G, \mathbf{Q}_G)ds. \tag{58}$$

Its r-fold iterate has Lipschitz constant $(CT)^r/r!$ by induction. Choose r so this is below one. Banach's theorem on the complete space $\mathbf{V}_T^G$ gives a fixed point of the iterate, also fixed by $\Phi_{\mathbf{V}}$. Equation (58) and Gronwall give uniqueness. □

Lemma 4.10 (Path-Law Sensitivity)**.** Under Assumptions (H1)–(H7), let $\mathbf{P}_G^\nu$ and $\mathbf{P}_G^{\widetilde{\nu}}$ be the fixed-point path ensembles corresponding to $\nu_G, \widetilde{\nu}_G \in \mathcal{V}_{[0,T]}$. There is $C_{\mathbf{V}} \ge 1$, independent of C_h, such that

$$D_T(\mathbf{P}_G^\nu, \mathbf{P}_G^{\widetilde{\nu}}) \le C_{\mathbf{V}} \sup_{\alpha,t,x} |\phi_\alpha(t,x|\nu_G) - \phi_\alpha(t,x|\widetilde{\nu}_G)|. \tag{59}$$

Proof. Set $\Delta_\phi = \sup_{\alpha,t,x} |\phi_\alpha(t,x|\nu_G) - \phi_\alpha(t,x|\widetilde{\nu}_G)|$ and synchronously couple the two fixed-point systems. The common spatial Lipschitz bound gives

$$\begin{aligned} W_1(\overline{\nu}_{\alpha,t}^\nu, \overline{\nu}_{\alpha,t}^{\widetilde{\nu}}) &\le (1+L_\phi)D_t(P_\alpha^\nu, P_\alpha^{\widetilde{\nu}}) + \Delta_\phi, \\ |\overline{\lambda}_\alpha^\nu(t) - \overline{\lambda}_\alpha^{\widetilde{\nu}}(t)| &\le C\big(D_t(\mathbf{P}_G^\nu, \mathbf{P}_G^{\widetilde{\nu}}) + \Delta_\phi\big). \end{aligned} \tag{60}$$

Writing $Y_\alpha(t) = \mathbb{E}\sup_{r\le t} |X_\alpha^\nu(r) - X_\alpha^{\widetilde{\nu}}(r)|$, the Lipschitz drift and feedback estimates yield

$$\begin{aligned} Y_\alpha(t) &\le C\int_0^t \big(Y_\alpha(s) + D_s(\mathbf{P}_G^\nu, \mathbf{P}_G^{\widetilde{\nu}}) + \Delta_\phi\big)ds, \\ \mathbb{E}\sup_{r\le t} |U_\alpha^\nu(r) - U_\alpha^{\widetilde{\nu}}(r)| &\le L_\phi Y_\alpha(t) + \Delta_\phi. \end{aligned}$$

Taking the supremum in α and applying Gronwall gives $D_T(\mathbf{P}_G^\nu, \mathbf{P}_G^{\widetilde{\nu}}) \le C\Delta_\phi$. By (34), L_ϕ, and hence C, is independent of C_h; take $C_{\mathbf{V}} = \max\{1, C\}$. □

4.3 Contraction and equilibrium existence

Let

$$\mathrm{K}_T^G := \{\mathbf{P}_G \in \mathbf{V}_T^G : \mathrm{Marg}(\mathbf{P}_G) \in \mathcal{V}_{[0,T]}\}. \tag{61}$$

For $\mathbf{P}_G \in \mathrm{K}_T^G$, define $\Psi(\mathbf{P}_G)$ as follows. Put $\nu_G = \mathrm{Marg}(\mathbf{P}_G)$ and let $\Psi(\mathbf{P}_G)$ be the unique consistent path law generated by $\Phi_{\mathbf{V}}[\cdot|\nu_G]$ in Lemma 4.9. Thus Ψ first solves the moment-parameter consistency problem with the mean-field distributional environment ν_G frozen; its fixed-point condition then enforces mass–environment consistency.

Theorem 4.11 (Well-Posedness and Existence of a Generalized Nash Equilibrium)**.** Under Assumptions (H1)–(H8), there exist admissible-class constants $M_{\mathcal{V}}, C_h > 0$ such that:

(i) (K_T^G, D_T) is a nonempty complete metric space, and $\Psi : \mathrm{K}_T^G \to \mathrm{K}_T^G$ is a strict contraction.

(ii) There exists a unique fixed point $\mathbf{P}_G^* \in \mathrm{K}_T^G$. With $\nu_G^* = \mathrm{Marg}(\mathbf{P}_G^*)$ and $\phi_\alpha^*(t,x) = \phi_\alpha(t, x \mid \nu_G^*)$, the pair (ν_G^*, ϕ^*) is a generalized Nash equilibrium: ϕ_α^* is optimal for the representative problem with environment ν_G^*, and the law of $(X_\alpha^*, \phi_\alpha^*(\cdot, X_\alpha^*))$ has time-t marginal $\nu_{\alpha,t}^*$.

(iii) For every $\alpha \in [0,1]$, the equilibrium closed-loop FPK equation admits a unique generalized solution $\rho^*_\alpha \in V_2(Q_T)$.

Proof. *(i)* By Lemma 4.10, $C_{\mathbf{V}} \geq 1$. Set $q := C_{\mathbf{V}} C_\phi$. Assumption (H8) gives $q < 1$, hence $C_\phi < 1$. Lemma 4.8 therefore shows that choosing $M_{\mathcal{V}} \geq M_0$ and $C_h \geq A_0/(1-q)$ ensures that every output marginal flow belongs to $\mathcal{V}_{[0,T]}$. The class K^G_T is nonempty: for any $u_0 \in K$, the Dirac ensemble concentrated on the constant path $(X_t, U_t) = (0, u_0)$ has a constant marginal flow satisfying (C1)–(C3). It is complete because the marginal projection is 1-Lipschitz and the defining continuity and Hölder bounds of $\mathcal{V}_{[0,T]}$ are closed under uniform W_1 convergence. Lemmas 4.9, 4.10, and 4.3 show that Ψ is a self-map and, for all $\mathbf{P}_G, \widetilde{\mathbf{P}}_G \in \mathrm{K}^G_T$,

$$D_T(\Psi(\mathbf{P}_G), \Psi(\widetilde{\mathbf{P}}_G)) \leq q\, d_{\mathcal{V}}(\mathrm{Marg}(\mathbf{P}_G), \mathrm{Marg}(\widetilde{\mathbf{P}}_G)) \leq q\, D_T(\mathbf{P}_G, \widetilde{\mathbf{P}}_G). \tag{62}$$

Thus Ψ is a strict contraction with contraction modulus $q < 1$, proving (i).

(ii) The Banach fixed-point theorem applied to Ψ on (K^G_T, D_T) gives a unique fixed point $\mathbf{P}^*_G$. Its fixed-point identity makes the closed-loop law consistent, and the verification argument proves that ϕ^* is optimal in the induced environment ν^*_G. Therefore (ν^*_G, ϕ^*) is the stated generalized Nash equilibrium.

(iii) For each $\alpha \in [0,1]$, Lemma 4.5 gives existence and uniqueness of the generalized solution $\rho^*_\alpha \in V_2(Q_T)$ of the equilibrium closed-loop FPK equation. □

5 Finite-network and finite-population approximation

The limiting graphon MFG equilibrium is approximated by games with finite networks and agent populations, using the same local and graphon-weighted joint state–control statistics. The key convention is a *leave-one-out* empirical distributional environment: agent i's parameter excludes its own state–control pair, avoiding a self-referential feasible set at the single-agent error quantified below.

5.1 Finite approximation assumptions

The finite game is defined in Section 2.2.1; here we add the approximation and feasibility-transport hypotheses.

For $u \in K_\alpha(\lambda)$ and $\overline{\lambda} \in \mathbb{R}^{2(n+m)}$, define the feasibility-transport map

$$\mathcal{R}_\alpha(u, \lambda, \overline{\lambda}) := \arg \min_{v \in K_\alpha(\overline{\lambda})} \frac{1}{2}|v-u|^2, \qquad \forall\, u \in K_\alpha(\lambda). \tag{63}$$

Here $\mathrm{Gr}(K_\cdot) = \{(\alpha, \lambda, u) : \alpha \in [0,1], \lambda \in \mathbb{R}^{2(n+m)}, u \in K_\alpha(\lambda)\}$ is the feasible graph, λ is the source parameter, and $\overline{\lambda}$ is the target moment parameter. The following are the finite approximation hypotheses.

(F1) [Growing clustered networks and populations] $M_k \to \infty$, $n^k_{\min} \geq 2$, and $n^k_{\min} \to \infty$.

(F2) [Finite-network graphon approximation] The sectional error satisfies

$$e^k_W = \max_{1 \leq \ell \leq M_k} \sum_{r=1}^{M_k} \left| \frac{W^k_{\ell r}}{M_k} - \int_{I^k_r} \mathcal{W}(\alpha^k_\ell, \beta) d\beta \right| \longrightarrow 0. \tag{64}$$

Consequently, by Assumption (H4)(i), $\overline{W} := \sup_{k,\ell} \frac{1}{M_k} \sum_{r=1}^{M_k} |W^k_{\ell r}| < \infty$.

(F3) [Global finite-system coefficients] The finite dynamics, costs, and constraints use the same $\widehat{f}, \widehat{l}, g$, and $K_\alpha(\lambda)$ as the limiting model. The maps $\widehat{f}$ and $\widehat{l}$ are globally Lipschitz in (x, u, λ), with constants independent of k.

(F4) [Global uniform constraint regularity] The bounds in Assumption (H6)(ii) can be chosen independently of R, and the maps $D_u G_\alpha$ and $D_\lambda G_\alpha$ are globally Lipschitz in (u,λ), uniformly in α. Moreover, there is a constant $\kappa_{\mathrm{fin}} > 0$ such that $\sigma_{\min}\big(D_u G_{\alpha, I_\alpha(u,\lambda)}(u,\lambda)\big) \geq \kappa_{\mathrm{fin}}$ whenever $\lambda \in \mathbb{R}^{2(n+m)}$, $u \in K_\alpha(\lambda)$, and the active set is nonempty.

(F5) [Finite-feedback contraction] With $L_{\Lambda,u} := 1 + \overline{W}$ and $L_{\mathcal{R},\lambda}$ denoting the transport constant furnished by Lemma 5.1 below, assume $q_c := L_{\mathcal{R},\lambda} L_{\Lambda,u} < 1$.

Because $\nu_G^* \in \mathcal{V}_{[0,T]}$, conditions (C1) and (C3), a finite time-grid argument, and compactness of $[0,1]$ imply the uniform spatial modulus

$$\omega_\nu(h) := \sup_{\substack{|\alpha-\beta|\leq h \\ t\in[0,T]}} W_1(\nu^*_{\alpha,t}, \nu^*_{\beta,t}) \longrightarrow 0. \tag{65}$$

Empirical moments can make the limiting feedback infeasible in a finite game, motivating the feasibility transport and feedback construction that follow.

5.2 Feasibility transport and finite feedback

Lemma 5.1 (Feasibility Transport)**.** Under Assumptions (H1), (H6), and (F4), $\mathcal{R}_\alpha$ has a jointly measurable, single-valued version. More precisely, there is a constant $L_{\mathcal{R},\lambda} < \infty$ such that for every $\alpha \in [0,1]$, $\lambda, \eta, \overline{\lambda}, \overline{\eta} \in \mathbb{R}^{2(n+m)}$, $u \in K_\alpha(\lambda)$, and $v \in K_\alpha(\eta)$,

$$|\mathcal{R}_\alpha(u,\lambda,\overline{\lambda}) - u| \leq L_{\mathcal{R},\lambda}|\overline{\lambda} - \lambda|, \tag{66}$$

$$|\mathcal{R}_\alpha(u,\lambda,\overline{\lambda}) - \mathcal{R}_\alpha(v,\eta,\overline{\eta})| \leq |u - v| + L_{\mathcal{R},\lambda}|\overline{\lambda} - \overline{\eta}|. \tag{67}$$

Proof. Fix α and write $P_{\alpha,\theta} z$ for the unique minimizer of $\frac{1}{2}|w - z|^2$ over $G_\alpha(w,\theta) \leq 0$. Its KKT Lagrangian has w-Hessian at least I, since the constraints are convex. Uniform LICQ and the bounds in (F4) therefore give uniform strong regularity [5, Prop. 5.38]. Robinson's implicit-function theorem [44, Thm. 2.1 and Cor. 2.2] then yields $|P_{\alpha,\theta} z - P_{\alpha,\vartheta} z| \leq L_{\mathcal{R},\lambda}|\theta - \vartheta|$.

Projection onto a fixed closed convex set is nonexpansive; hence $|P_{\alpha,\theta} z - P_{\alpha,\vartheta} z'| \leq |z - z'| + L_{\mathcal{R},\lambda}|\theta - \vartheta|$. Since $P_{\alpha,\lambda} u = u$ for $u \in K_\alpha(\lambda)$, these two estimates give (66) and (67). Joint continuity makes the feasible graph Borel; measurable projection and uniqueness give the jointly measurable version. □

Recall from (11) and Theorem 4.11 that $\lambda^*_\alpha(t) = \lambda_\alpha[\nu^*_G](t)$, $\phi^*_\alpha(t,x) = \phi_\alpha(t, x \mid \nu^*_G)$ are, respectively, the limiting equilibrium local-network parameter and feedback. Although $\phi^*_\alpha(t,x) \in K_\alpha(\lambda^*_\alpha(t))$, it need not be feasible for an agent's empirical parameter. Define the finite constrained implementation implicitly by

$$\widehat{U}^{N,k}_i(t) = \mathcal{R}_{\alpha^k_i}\Big(\phi^*_{\alpha^k_i}(t, \widehat{X}^{N,k}_i(t)), \lambda^*_{\alpha^k_i}(t), \widehat{\lambda}^{N,k,-i}_i(t)\Big). \tag{68}$$

Here $\widehat{\lambda}^{N,k,-i}_i(t)$ uses the current finite profile, including other agents' controls. Thus (68) is a fixed point in the finite control vector.

For state and control vectors, direct averaging in (3)–(5) gives

$$\max_i \big|\widehat{\lambda}^{N,k,-i}_i(t;\mathbf{x},\mathbf{u}) - \widehat{\lambda}^{N,k,-i}_i(t;\mathbf{y},\mathbf{v})\big| \leq L_{\Lambda,x}\|\mathbf{x} - \mathbf{y}\|_\infty + L_{\Lambda,u}\|\mathbf{u} - \mathbf{v}\|_\infty, \tag{69}$$

$$L_{\Lambda,x} = L_{\Lambda,u} = 1 + \overline{W}. \tag{70}$$

Lemma 5.2. Under Assumptions (F2), (F4), and (F5), for every t and state vector $\mathbf{x}$, (68) has a unique feasible solution $\widehat{\mathbf{U}}^{N,k} = \Phi^{N,k}(t,\mathbf{x})$. Moreover, $\Phi^{N,k}$ is jointly Borel and globally Lipschitz in $\mathbf{x}$ uniformly in k, its closed-loop control is $\mathbb{F}^{N,k}$-progressively measurable, and

$$\max_i \big|\widehat{U}^{N,k}_i(t,\mathbf{x}) - \phi^*_{\alpha^k_i}(t,x_i)\big| \leq \frac{L_{\mathcal{R},\lambda}}{1 - q_c} \max_i \big|\widehat{\lambda}^{N,k,-i}_i(t;\mathbf{x},\phi^*(t,\mathbf{x})) - \lambda^*_{\alpha^k_i}(t)\big|, \tag{71}$$

where $\phi^*(t,\mathbf{x}) := (\phi^*_{\alpha^k_1}(t,x_1), \ldots, \phi^*_{\alpha^k_{N_k}}(t,x_{N_k}))$.

Proof. For fixed $(t,\mathbf{x})$ define the map on finite control vectors by

$$(\mathcal{T}_{\mathbf{x}}(\mathbf{u}))_i = \mathcal{R}_{\alpha_i^k}\Big(\phi^*_{\alpha_i^k}(t,x_i), \lambda^*_{\alpha_i^k}(t), \widehat{\lambda}_i^{N,k,-i}(t;\mathbf{x},\mathbf{u})\Big). \tag{72}$$

Lemma 5.1 and (69) give $\|\mathcal{T}_{\mathbf{x}}(\mathbf{u})-\mathcal{T}_{\mathbf{x}}(\mathbf{v})\|_\infty \le q_c\|\mathbf{u}-\mathbf{v}\|_\infty$. Thus Banach's theorem gives a unique fixed point, which is feasible by (63). Let $D = \|\widehat{\mathbf{U}}^{N,k} - \phi^*(t,\mathbf{x})\|_\infty$ and let $B = \max_i |\widehat{\lambda}_i^{N,k,-i}(t;\mathbf{x},\phi^*(t,\mathbf{x})) - \lambda^*_{\alpha_i^k}(t)|$. By (66) and (69) with $\mathbf{x}$ fixed, the corresponding parameter difference is at most $L_{\Lambda,u}D$. Hence $D \le L_{\mathcal{R},\lambda}(L_{\Lambda,u}D + B) = q_c D + L_{\mathcal{R},\lambda}B$, proving (71).

Picard iteration from a Borel initial vector proves Borel measurability. For two state vectors, the same estimates give $\|\mathcal{T}_{\mathbf{x}}(\mathbf{u}) - \mathcal{T}_{\mathbf{y}}(\mathbf{v})\|_\infty \le C\|\mathbf{x}-\mathbf{y}\|_\infty + q_c\|\mathbf{u}-\mathbf{v}\|_\infty$. Absorbing q_c gives the uniform state Lipschitz bound. The resulting closed-loop SDE has a unique strong solution, and its Borel feedback is progressively measurable. □

5.3 Approximation error and ϵ-generalized Nash equilibrium

The ϵ-generalized Nash proof compares the finite feasible profile (68), the finite system after an admissible unilateral deviation, and independent limiting representatives

$$dY_i^\infty(t) = \widehat{f}(Y_i^\infty(t), V_i^\infty(t), \lambda^*_{\alpha_i^k}(t))dt + \sigma dB_i(t), \qquad V_i^\infty(t) = \phi^*_{\alpha_i^k}(t, Y_i^\infty(t)). \tag{73}$$

Define $\eta_k = 1/\sqrt{n^k_{\min}} + e^k_W + \omega_\nu(1/M_k)$. The three terms are the cluster sampling, graphon approximation, and cell discretization errors. The influence of one agent is measured by $a_k := 1/(n^k_{\min}-1) + \max_{\ell,r}\big(|W^k_{\ell r}|/(M_k(n^k_r-1))\big)$. The row-sum bound (F2) gives $a_k \le C/(n^k_{\min}-1) \to 0$.

For each α, let $\mathcal{A}^*_\alpha$ contain the progressively measurable K-valued controls on any filtered basis carrying an initial state with density p_0 and a Brownian motion that remains Brownian in that filtration, with $U(t) \in K_\alpha(\lambda^*_\alpha(t))$ for $(dt\otimes\mathbb{P})$-almost every (t,ω). On this basis, $J^*_\alpha(U)$ is defined by the right-hand side of (10) with $\nu_G = \nu^*_G$ and expectation under $\mathbb{P}$. HJB verification on every such basis includes auxiliary randomization and the transported deviations below; finite deviations range over (8).

Lemma 5.3 (Finite GMFGC Consistency)**.** Under Assumptions (H1)–(H8) and (F1)–(F5), let $(\widehat{X}_i^{N,k}, \widehat{U}_i^{N,k})$ and (Y_i^∞, V_i^∞) be the finite-implementation and limiting representative state–control pairs at node α_i^k, respectively. Then some $C>0$, independent of i,k, satisfies

$$\sup_i \mathbb{E}^{N,k}\Big[\sup_{0\le t\le T}|\widehat{X}_i^{N,k}(t) - Y_i^\infty(t)|\Big] + \sup_i \mathbb{E}^{N,k}\int_0^T |\widehat{U}_i^{N,k}(t) - V_i^\infty(t)|dt \le C\eta_k,$$
$$\sup_i \big|J_i^{N,k}(\widehat{\mathbf{U}}^{N,k}) - J^*_{\alpha_i^k}(\phi^*_{\alpha_i^k})\big| \le C\eta_k.$$

Proof. Couple the finite system and independent limiting representatives using common initial states and Brownian motions. Set $d_X(t) = \sup_i \mathbb{E}^{N,k}|\widehat{X}_i^{N,k}(t) - Y_i^\infty(t)|$ and define $d_U(t)$ analogously. Also set $D_X(t) = \sup_i \mathbb{E}^{N,k}\sup_{s\le t}|\widehat{X}_i^{N,k}(s) - Y_i^\infty(s)|$ and $D_U(t) = \sup_i \mathbb{E}^{N,k}\int_0^t |\widehat{U}_i^{N,k}(s) - V_i^\infty(s)|ds$.

For $Z_j^\infty = (Y_j^\infty, V_j^\infty)$, let $\overline{\lambda}_i^{N,k,-i}$ be defined by (3)–(5) with $q(X_j,U_j)$ replaced by Z_j^∞. Put $m_\alpha = \int z\,\nu^*_{\alpha,t}(dz)$ and, for $i \in C^k_\ell$, define $\lambda_i^{W,k} = (m_{\alpha_i^k}, M_k^{-1}\sum_r W^k_{\ell r} m_{\alpha_r^k})$ and $\lambda_i^{C,k} = (m_{\alpha_i^k}, \sum_r \int_{I_r^k}\mathcal{W}(\alpha_i^k,\beta)m_{\alpha_r^k}d\beta)$, suppressing t. Then

$$\overline{\lambda}_i^{N,k,-i} - \lambda^*_{\alpha_i^k} = (\overline{\lambda}_i^{N,k,-i} - \lambda_i^{W,k}) + (\lambda_i^{W,k} - \lambda_i^{C,k}) + (\lambda_i^{C,k} - \lambda^*_{\alpha_i^k}). \tag{74}$$

The terms are, respectively, the sampling, graphon, and cell errors. Uniform second moments, (64)–(65) and the row-sum bound give

$$\sup_{i,t}\mathbb{E}^{N,k}|\overline{\lambda}_i^{N,k,-i}(t) - \lambda^*_{\alpha_i^k}(t)| \le C\left(\frac{1}{\sqrt{n^k_{\min}}} + e^k_W + \omega_\nu(1/M_k)\right) = C\eta_k. \tag{75}$$

Coupling the finite and limiting pairs then bounds the actual parameter error by $C(d_X(t)+d_U(t)+\eta_k)$.

The Lipschitz properties of $\mathcal{R}_\alpha$ and ϕ^*_α give $d_U \le Cd_X + q_c d_U + C\eta_k$. Absorbing $q_c < 1$, synchronous state comparison yields $D_X(t)+D_U(t) \le C\int_0^t D_X(s)ds + C\eta_k$. Gronwall gives $D_X(T)+D_U(T) \le C\eta_k$, proving the first claim. The Lipschitz bounds for $\widehat{l}$ and g then give the cost estimate. □

Lemma 5.4 (Lower Bound for Finite Deviations). Under Assumptions (H1)–(H8) and (F1)–(F5), fix an agent ι. Let $U_\iota^{N,k}$ be any finite-admissible deviation, and suppose all agents $j \neq \iota$ use the finite-feasible GMFGC implementation. Then there is $C > 0$, independent of k, ι, and $U_\iota^{N,k}$, such that

$$J_\iota^{N,k}(U_\iota^{N,k}, \widehat{\mathbf{U}}_{-\iota}^{N,k}) \ge \inf_{U \in \mathcal{A}^*_{\alpha_\iota^k}} J^*_{\alpha_\iota^k}(U) - C(\eta_k + a_k). \tag{76}$$

The infimum equals $J^*_{\alpha_\iota^k}(\phi^*_{\alpha_\iota^k})$ by optimality.

Proof. *Population Response.* Denote the baseline finite system by A and the system with agent ι's deviation by B. Thus $U_j^A = \Phi_j^{N,k}(t, \mathbf{X}^A)$, while $U_\iota^B = U_\iota^{N,k}$ and $U_j^B = \Phi_j^{N,k}(t, \mathbf{X}^B)$ for $j \neq \iota$. Couple the nondeviators synchronously and set $\Delta_X(t) = \max_{j\neq\iota} \mathbb{E}^{N,k} \sup_{s\le t} |X_j^A(s) - X_j^B(s)|$ and $\Delta_U(t) = \max_{j\neq\iota} \mathbb{E}^{N,k}|U_j^A(t) - U_j^B(t)|$. Linear growth and compactness give uniform first moments. Since the deviator carries only weights $(n_r^k - 1)^{-1}$ and $|W_{\ell r}^k|/[M_k(n_r^k - 1)]$, the definition of a_k and the row-sum bound give a field difference at most $C(\Delta_X(t) + \Delta_U(t) + a_k)$. Transport and $q_c < 1$ imply $\Delta_U \le C(\Delta_X + a_k)$; synchronous state comparison and Gronwall yield $\Delta_X(T) + \int_0^T \Delta_U(t)dt \le Ca_k$.

Transported Deviation. Let $\widehat{\lambda}_\iota^B$ be the deviator's leave-one-out parameter. The preceding response bound and Lemma 5.3 give $\mathbb{E}^{N,k}\int_0^T |\widehat{\lambda}_\iota^B - \lambda^*_{\alpha_\iota^k}|dt \le C(\eta_k + a_k)$. Set $\widetilde{U}_\iota(t) = \mathcal{R}_{\alpha_\iota^k}(U_\iota^{N,k}(t), \widehat{\lambda}_\iota^B(t), \lambda^*_{\alpha_\iota^k}(t))$. This control is progressively measurable and belongs to $\mathcal{A}^*_{\alpha_\iota^k}$. Couple its state $\widetilde{X}_\iota$ with X_ι^B using the same Brownian motion. Transport, linear growth, and Gronwall yield $\mathbb{E}^{N,k}\int_0^T |U_\iota^{N,k} - \widetilde{U}_\iota|dt + \mathbb{E}^{N,k}\sup_{t\le T}|X_\iota^B - \widetilde{X}_\iota| \le C(\eta_k + a_k)$.

Cost Comparison. Global Lipschitz continuity of $\widehat{l}$ and g and the above estimates give $|J_\iota^{N,k}(U_\iota^{N,k}, \widehat{\mathbf{U}}_{-\iota}^{N,k}) - J^*_{\alpha_\iota^k}(\widetilde{U}_\iota)| \le C(\eta_k + a_k)$. Since $\widetilde{U}_\iota \in \mathcal{A}^*_{\alpha_\iota^k}$, optimality proves the claim. □

Theorem 5.5 (Finite ϵ-Generalized Nash Equilibrium). Assume (H1)–(H8) and (F1)–(F5). Let $\widehat{\mathbf{U}}^{N,k}$ be the exactly feasible strategy profile from (68). For each k, there is an error $\epsilon = \epsilon(k) \ge 0$, with $\epsilon(k) \to 0$, such that, for every agent i and every finite-admissible deviation $U_i^{N,k}$,

$$J_i^{N,k}(\widehat{\mathbf{U}}^{N,k}) \le J_i^{N,k}(U_i^{N,k}, \widehat{\mathbf{U}}_{-i}^{N,k}) + \epsilon, \qquad \epsilon \le C(\eta_k + a_k) \longrightarrow 0. \tag{77}$$

Thus $\widehat{\mathbf{U}}^{N,k}$ is a finite-network, finite-population ϵ-generalized Nash equilibrium. Moreover, (F2) gives $a_k \le C/(n_{\min}^k - 1)$.

Proof. For any agent i and finite-admissible $U_i^{N,k}$, Lemmas 5.3 and 5.4 give

$$J_i^{N,k}(\widehat{\mathbf{U}}^{N,k}) \le J^*_{\alpha_i^k}(\phi^*_{\alpha_i^k}) + C\eta_k \le J_i^{N,k}(U_i^{N,k}, \widehat{\mathbf{U}}_{-i}^{N,k}) + C(\eta_k + a_k). \tag{78}$$

As $i, U_i^{N,k}$ are arbitrary and $\eta_k + a_k \to 0$ by (F1)– (F2) and (65), (77) follows. □

6 Conclusion

This paper develops a constrained GMFGC model with law-coupled admissibility. Under the stated regularity and contraction conditions, its GNE is unique and its generalized FPK flow is well posed. Feasibility transport yields finite-network, finite-population ϵ-generalized Nash profiles satisfying the shared constraints. Equilibrium existence and uniqueness under weaker conditions than contraction remain open, as do measure-dependent constrained problems.

References

[1] Y. Achdou, P. Mannucci, C. Marchi, and N. Tchou, Deterministic mean field games with control on the acceleration and state constraints, SIAM J. Math. Anal., 54 (2022), pp. 3757–3788.

[2] M. Aziz, H. Suprême, H. Dagdougui, and A. Lesage-Landry, Optimal energy trading in residential prosumer clusters via graphon mean field games, Electr. Power Syst. Res., 263 (2027), p. 113745, https://doi.org/10.1016/j.epsr.2026.113745.

[3] A. Bensoussan, J. Frehse, and P. Yam, Mean Field Games and Mean Field Type Control Theory, Springer, 2013.

[4] J. F. Bonnans, J. Gianatti, and L. Pfeiffer, A Lagrangian approach for aggregative mean field games of controls with mixed and final constraints, SIAM J. Control Optim., 61 (2023), pp. 105–134.

[5] J. F. Bonnans and A. Shapiro, Perturbation analysis of optimization problems, Springer, 2013.

[6] P. E. Caines, Embedded vertexon-graphons and embedded GMFG systems, in Proc. 2022 IEEE 61st Conf. Decis. Control (CDC), IEEE, 2022, pp. 5550–5557.

[7] P. E. Caines and M. Huang, Graphon mean field games and the GMFG equations, in Proc. 2018 IEEE Conf. Decis. Control (CDC), IEEE, 2018, pp. 4129–4134, https://doi.org/10.1109/CDC.2018.8619367.

[8] P. E. Caines and M. Huang, Graphon mean field games and the GMFG equations: ϵ-Nash equilibria, in Proc. 2019 IEEE 58th Conf. Decis. Control (CDC), IEEE, 2019, pp. 286–292, https://doi.org/10.1109/CDC40024.2019.9029871.

[9] P. E. Caines and M. Huang, Graphon mean field games and their equations, SIAM J. Control Optim., 59 (2021), pp. 4373–4399.

[10] P. E. Caines and M. Huang, Mean field games on dense and sparse networks: the graphexon MFG equations, in Proc. 2024 Amer. Control Conf. (ACC), IEEE, 2024, pp. 4230–4235.

[11] P. E. Caines and M. Huang, Sparse network mean field games: Ring structures and related topologies, in Proc. 2024 IEEE 63rd Conf. Decis. Control (CDC), IEEE, 2024, pp. 2584–2590, https://doi.org/10.1109/CDC56724.2024.10886458.

[12] P. Cannarsa and R. Capuani, Existence and uniqueness for mean field games with state constraints, in PDE models for multi-agent phenomena, Springer, 2018, pp. 49–71.

[13] P. Cardaliaguet and C.-A. Lehalle, Mean field game of controls and an application to trade crowding, Math. Financ. Econ., 12 (2018), pp. 335–363.

[14] P. Cardaliaguet, A. R. Mészáros, and F. Santambrogio, First order mean field games with density constraints: pressure equals price, SIAM J. Control Optim., 54 (2016), pp. 2672–2709.

[15] R. Carmona, D. B. Cooney, C. V. Graves, and M. Laurière, Stochastic graphon games: I. The static case, Math. Oper. Res., 47 (2022), pp. 750–778, https://doi.org/10.1287/moor.2021.1148.

[16] R. Carmona, F. Delarue, et al., Probabilistic theory of mean field games with applications I-II, vol. 3, Springer, 2018.

[17] R. Carmona and D. Lacker, A probabilistic weak formulation of mean field games and applications, Ann. Appl. Probab., 25 (2015), pp. 1189–1231.

[18] X. Chen and Z. Wang, Differential variational inequality approach to dynamic games with shared constraints, Math. Program., 146 (2014), pp. 379–408.

[19] L. C. Evans, Partial Differential Equations, vol. 19 of Grad. Stud. Math., Amer. Math. Soc., Providence, RI, 1998.

[20] C. Fabian, K. Cui, and H. Koeppl, Learning mean field games on sparse graphs: A hybrid graphex approach, in Proc. 12th Int. Conf. Learn. Represent. (ICLR), 2024.

[21] F. Facchinei and C. Kanzow, Generalized Nash equilibrium problems, Ann. Oper. Res., 175 (2010), pp. 177–211.

[22] W. H. Fleming and R. W. Rishel, Deterministic and stochastic optimal control, Springer, 2012.

[23] S. Gao, P. E. Caines, and M. Huang, LQG graphon mean field games: Analysis via graphon-invariant subspaces, IEEE Trans. Automat. Control, 68 (2023), pp. 7482–7497.

[24] D. Gilbarg and N. S. Trudinger, Elliptic Partial Differential Equations of Second Order, Classics Math., Springer, 2 ed., 2001, https://doi.org/10.1007/978-3-642-61798-0.

[25] M. A. Gkogkas, C. Kuehn, and C. Xu, Mean field limits of co-evolutionary signed heterogeneous networks, Eur. J. Appl. Math., 37 (2026), pp. 643–686, https://doi.org/10.1017/S0956792524000858.

[26] D. A. Gomes and V. K. Voskanyan, Extended deterministic mean-field games, SIAM J. Control Optim., 54 (2016), pp. 1030–1055, https://doi.org/10.1137/130944503.

[27] H. Hatami, L. Lovász, and B. Szegedy, Limits of locally–globally convergent graph sequences, Geom. Funct. Anal., 24 (2014), pp. 269–296.

[28] M. Hintermüller, T. M. Surowiec, and M. Theiss, On a differential generalized Nash equilibrium problem with mean field interaction, SIAM J. Optim., 34 (2024), pp. 2821–2855.

[29] A. Hu and Z. Lyu, Mean-field games with constraints, arXiv:2510.11843, (2025).

[30] M. Huang and P. Caines, Classical solutions to graphon MFG equations with affine control: Lipschitz mappings on Hölder spaces, Les Cahiers du GERAD ISSN, 711 (2024), p. 2440.

[31] M. Huang, P. E. Caines, and R. P. Malhamé, Individual and mass behaviour in large population stochastic wireless power control problems: centralized and Nash equilibrium solutions, in Proc. 42nd IEEE Conf. Decis. Control (CDC), vol. 1, IEEE, 2003, pp. 98–103.

[32] M. Huang, R. P. Malhamé, and P. E. Caines, Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle, Commun. Inf. Syst., 6 (2006), pp. 221–252.

[33] M. Huang and S. L. Nguyen, Mean field games for stochastic growth with relative utility, Appl. Math. Optim., 74 (2016), pp. 643–668, https://doi.org/10.1007/s00245-016-9395-8.

[34] Y. Jiang, E. Neuman, and S. Tuschmann, Local-global convergence and optimal control on graphings. SSRN preprint, 2026, https://doi.org/10.2139/ssrn.7476398.

[35] C. Kuehn and C. Xu, Vlasov equations on digraph measures, J. Differential Equations, 339 (2022), pp. 261–349, https://doi.org/10.1016/j.jde.2022.08.023.

[36] C. Kuehn and C. Xu, Vlasov equations on directed hypergraph measures, Partial Differ. Equ. Appl., 6 (2025), p. 9, https://doi.org/10.1007/s42985-025-00313-6.

[37] D. Lacker and A. Soret, A case study on stochastic games on large graphs in mean field and sparse regimes, Math. Oper. Res., 47 (2022), pp. 1530–1565, https://doi.org/10.1287/moor.2021.1179.

[38] O. A. Ladyzhenskaia, V. A. Solonnikov, and N. N. Ural'tseva, Linear and quasi-linear equations of parabolic type, vol. 23, Amer. Math. Soc., 1968.

[39] J.-M. Lasry and P.-L. Lions, Mean field games, Jpn. J. Math., 2 (2007), pp. 229–260.

[40] M. Li, C. Mou, Z. Wu, and C. Zhou, Linear-quadratic mean field games of controls with non-monotone data, Trans. Amer. Math. Soc., 376 (2023), pp. 4105–4143, https://doi.org/10.1090/tran/8868.

[41] L. Lovász, Large networks and graph limits, vol. 60, Amer. Math. Soc., 2012.

[42] J. Nocedal and S. J. Wright, Numerical optimization, Springer, 2006.

[43] F. Parise and A. Ozdaglar, Graphon games: A statistical framework for network games and interventions, Econometrica, 91 (2023), pp. 191–225, https://doi.org/10.3982/ECTA17564.

[44] S. M. Robinson, Strongly regular generalized equations, Math. Oper. Res., 5 (1980), pp. 43–62.

[45] R. T. Rockafellar and R. J. Wets, Variational analysis, Springer, 1998.

[46] V. Veitch and D. M. Roy, The class of random graphs arising from exchangeable random measures. arXiv:1512.03099, 2015, https://doi.org/10.48550/arXiv.1512.03099.

[47] C. Villani, Optimal Transport: Old and New, vol. 338, Springer, 2009.

[48] Y. Yan and T. Hayakawa, Pareto-improving incentive mechanism for noncooperative dynamical systems under sustainable budget constraint, IEEE Trans. Automat. Control, 69 (2023), pp. 4291–4306.