

SLS-BRD: A system-level approach to seeking generalised feedback Nash equilibria

Otacílio B. L. Neto, Michela Mulas, and Francesco Corona

Abstract—This work proposes a policy learning algorithm for seeking generalised feedback Nash equilibria in N_P -players non-cooperative dynamic games. We consider linear-quadratic games with stochastic dynamics and design a best-response dynamics in which players update and communicate a parametrisation of their state-feedback policies. Our approach leverages the System Level Synthesis framework to formulate each player’s update rule as the solution to a tractable robust optimisation problem. Under certain conditions, rates of convergence to a feedback Nash equilibrium can be established. The algorithm is showcased in exemplary problems ranging from the decentralised control of unstable systems to competition in oligopolistic markets.

Index Terms—Non-cooperative games, Feedback Nash equilibrium, Best-response dynamics, System level synthesis

I. INTRODUCTION

MODERN cyber-physical systems are often comprised of interacting subsystems operated locally by selfish decision-making agents. Ideally, agents operate these systems according to feedback policies that optimise local objectives, while satisfying global requirements and being robust to their rivals’ interference. However, the large-scale and decentralised nature of most applications, alongside the usual lack of coordination between agents, hinders most traditional approaches to policy design. Dynamic game theory provides an alternative framework based on the concept of competitive equilibria (e.g., the Nash equilibrium [1, 2]), which describe locally optimal, yet strategically stable, operating conditions for each non-cooperative agent. A methodology for policy design based on Nash equilibrium seeking thus presents a promising venue for enabling a decentralised control of cyber-physical systems.

In general, solving a non-cooperative game can concern distinct goals: *i*) For the unbiased observer, to examine the behaviour of rational agents from a local to a global scale; *ii*) for the game designer, to actuate the self-interested players towards system-wide objectives; and *iii*) for the player, to determine locally optimal and robust policies in competitive environments. Regardless, computing a Nash equilibrium (NE) solution is a notoriously difficult task [3]. Algorithmic game theory thus emerges as the field concerned with designing procedures to bridge this computational gap [4]. An important class of algorithms are those which place to the players the task of converging towards a non-cooperative equilibrium: These include best-response dynamics (BRD, [5]–[10]) and

no-regret learning (NRL, [11]–[15]). Recently, online feedback optimisation has also been employed for equilibrium seeking problems [16]. In general, these are fixed-point iteration methods centred on players updating their strategies, simultaneously and independently, using only the information available to them. In particular, best-response dynamics stands out as a simple, yet fundamental, model of policy learning for uncoordinated but communicating players. These algorithms have become important tools in economics and engineering, with applications including the design and control of networked systems [17]–[19], robotics [20, 21], and resource management [22]–[24].

For dynamic games, any solution concept depends qualitatively on a description of the data available to each player; a notion termed *information pattern* [1]. Under *open-loop* information patterns, when players only access the initial state of the game, the analysis and solution of NE problems are well-understood. Conversely, games under *closed-loop* (or *feedback*) information patterns, when players monitor the state of the game, are still under active research. A feedback Nash equilibrium (FNE) solution is often more attractive than its open-loop counterpart: It renders the players robust to disturbances and decision errors. However, designing routines for computing FNE is demanding for all but simple cases. A notably challenging class of games consists of generalised feedback Nash equilibrium (GFNE, [2]) problems: When players’ policies are required to satisfy coupled constraints (e.g., restrictions on the state of the game). To the best of our knowledge, methods for FNE seeking exist only for linear-quadratic games under restrictive assumptions [25]–[28], unconstrained nonlinear and control-affine games [29]–[32], and application-specific problems [33, 34]. Recently, [35] expanded these results and presented a systematic (approximate) solution for GFNE problems. Despite remarkable, these efforts are restricted to finite-duration games and thus exclude decision-making processes in which agents are continuously operating the underlying system. Finally, a systematic solution is also lacking for GFNE games described by stochastic dynamics.

In this work, we investigate algorithms for GFNE seeking in (stationary) difference games with linear-stochastic dynamics. Leveraging the System Level Synthesis (SLS, [36]) framework, we propose a best-response approach based on a parametrisation of all stabilising policies for the game. The algorithm consists of players iteratively, and simultaneously, updating the parameters of their individual policies, then announcing these changes through some communication network (Figure 1). Under our approach, each player’s update rule is formulated as the solution to a tractable robust optimisation problem. Moreover, these problems allow for constraints to be imposed directly on the policies’ structure, thus having the ability to encode information

Otacílio B. L. Neto and Francesco Corona are with the Department of Chemical and Metallurgical Engineering, Aalto University, Finland (e-mails: {otacilio.neto, francesco.corona}@aalto.fi).

Michela Mulas is with the Department of Teleinformatics Engineering, Federal University of Ceará, Brazil (e-mail: michela.mulas@ufc.br).

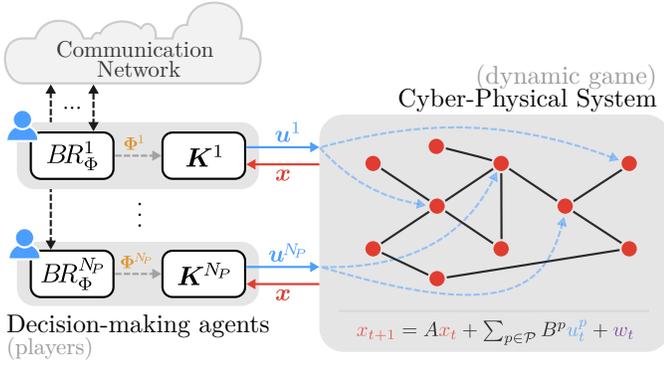


Figure 1. SLS-BRD: Control architecture and the learning dynamics.

patterns at the synthesis level. Finally, this proposed learning dynamics do not depend on the state and actions applied to the system, and thus can be performed simultaneously with the game's execution. In summary, our contributions are:

- (i) A system-level best-response dynamics (SLS-BRD) algorithm for GFNE seeking in linear-stochastic dynamic games. Specifically, we design a class of best-response maps which compute an optimal system level parametrisation of a player's policy in response to its rivals' choices.
- (ii) A realisation of the best-response mappings as robust finite-dimensional programs amenable to numerical solutions. Considering these best-responses, the SLS-BRD algorithm is formally an ε -GFNE seeking algorithm for an equilibrium gap $\varepsilon > 0$ which can be explicitly established.
- (iii) For a specific but important class of dynamic games, an analysis of the conditions and rates of convergence of the SLS-BRD to feedback Nash equilibria.

This policy learning algorithm is showcased in simulated experiments on the decentralised control of an unstable network and price management in a competitive oligopolistic market.

The paper is organised as follows: Section II overviews the classes of (generalised) static and dynamic games, and the best-response dynamics algorithm. In Section III, we provide a system level parametrisation for linear-stochastic games, then design a best-response dynamics for GFNE seeking. Finally, Section IV illustrates this learning dynamics in simulated examples, and Section V provides some concluding remarks. Towards a concise presentation, only some essential theorems are proved in the main text: The remaining are in the Appendix.

A. Notation

We use Latin letters to denote vectors (lowercase) and mappings (uppercase), and use boldface to distinguish signals, operators, and their respective spaces. Sets are in calligraphic font; exceptions are the usual \mathbb{R} and \mathbb{N} , and the sets of $N \times N$ symmetric (\mathbb{S}^N), positive semidefinite (\mathbb{S}_+^N), and positive definite matrices (\mathbb{S}_{++}^N). In particular, sequences are denoted $\mathbf{x} = (x_t)_{t \in \mathcal{I}}$ for a countable set $\mathcal{I} \subseteq \mathbb{N}$, or $\mathbf{x} = (x_t)_{t=0}^T$ if $\mathcal{I} = \{0, \dots, T\}$. For $p \in (0, \infty)$, we use the space of N_x -dimensional vector-sequences $\ell_p^{N_x}(\mathcal{I}) = \{\mathbf{x} : \|\mathbf{x}\|_{\ell_p} = (\sum_{t \in \mathcal{I}} \|x_t\|^p)^{1/p} < \infty\}$,

with $\ell_\infty^{N_x}$ and $\ell_{\infty, e}^{N_x}$ the space of all bounded sequences and all sequences, respectively. $\mathcal{L}(\mathcal{X}, \mathcal{Y})$ is the set of all bounded linear operators $\mathbf{A} : \mathcal{X} \rightarrow \mathcal{Y}$ and we sometimes denote transformed signals by $\mathbf{A}\mathbf{x} = (Ax_t)_{t \in \mathcal{I}}$. Whenever $(\mathcal{X}, \mathcal{Y})$ are finite-dimensional, we let M_A refer to a matrix representation of \mathbf{A} . We use the standard definition of Hardy spaces \mathcal{H}_∞ and \mathcal{RH}_∞ , and write $\frac{1}{z}\mathcal{RH}_\infty$ to refer to the set of real-rational strictly proper transfer functions. Finally, some standard signals and operators used in this paper are: The impulse signal $\delta = (\delta_t)_{t \in \mathcal{I}}$, the identity operator I and the identity matrix I_{N_x} , and the shift operator $\mathbf{S}_+ : (x_0, x_1, \dots) \mapsto (0, x_0, \dots)$.

We distinguish set-valued mappings from ordinary functions using the notation $F : \mathcal{X} \rightrightarrows \mathcal{Y}$. For any tuple $s = (s^p)_{p \in \mathcal{P}} \in \mathcal{S}$ we frequently write $s = (s^p, s^{-p})$ to highlight the element s^p ; this should not be interpreted as a re-ordering. Similarly, if $\mathcal{S} = \prod_{p \in \mathcal{P}} \mathcal{S}^p$, we define the product $\mathcal{S}^{-p} = \prod_{\bar{p} \in \mathcal{P} \setminus \{p\}} \mathcal{S}^{\bar{p}}$.

II. NON-COOPERATIVE GAMES AND BEST-RESPONSE DYNAMICS

A (static) N_P -player game, denoted by a tuple

$$\mathcal{G} := (\mathcal{P}, \{\mathcal{S}^p\}_{p \in \mathcal{P}}, \{L^p\}_{p \in \mathcal{P}}), \quad (1)$$

defines the problem in which *players* $p \in \mathcal{P} = \{1, \dots, N_P\}$ each decides on a strategy $s^p \in \mathcal{S}^p(s^{-p}) \subseteq \mathcal{S}^p$ to minimise an objective function $L^p : \mathcal{S}^1 \times \dots \times \mathcal{S}^{N_P} \rightarrow \mathbb{R}$. The strategy spaces \mathcal{S}^p ($\forall p \in \mathcal{P}$) determine the actions available to the players, with the mappings $\mathcal{S}^p : \mathcal{S}^{-p} \rightrightarrows \mathcal{S}^p$ restricting this choice based on the actions from their opponents. As such, both the players' objectives and feasible strategies depend explicitly on their rivals' strategies. Finally, the players are assumed to be rational, non-cooperative, and acting simultaneously.

A solution to the game \mathcal{G} is understood as a strategy profile $s = (s^1, \dots, s^{N_P}) \in \mathcal{S}$, $\mathcal{S} = \mathcal{S}^1 \times \dots \times \mathcal{S}^{N_P}$, having some specified property that makes it agreeable to all players if they act rationally. In non-cooperative settings, a widely accepted solution concept is that of a generalized Nash equilibrium: The game is *solved* when no player can improve its objective by unilaterally deviating from the agreed strategy profile. Formally,

Definition 1. A strategy profile $s^* = (s^{1*}, \dots, s^{N_P*}) \in \mathcal{S}$ is a generalized Nash equilibrium (GNE) for the game \mathcal{G} if

$$L^p(s^{p*}, s^{-p*}) \leq \min_{s^p \in \mathcal{S}^p(s^{-p*})} L^p(s^p, s^{-p*}). \quad (2)$$

holds for every player $p \in \mathcal{P}$.

In general, the set of GNEs that solve a game \mathcal{G} ,

$$\Omega_{\mathcal{G}} := \{s^* \in \mathcal{S} : s^* \text{ satisfies Eq. (2)}\},$$

is not a singleton. As such, players might favour different solutions based on their individual objectives. We thus seek to characterise which solutions are "acceptable", in the sense that no other equilibrium can improve the objective of all players simultaneously. Consider the partial ordering of $\Omega_{\mathcal{G}}$ defined by

$$s^* \prec \tilde{s}^* \iff L^p(s^{p*}) \leq L^p(\tilde{s}^{p*}), \quad \forall p \in \mathcal{P}, \quad (3)$$

with at least one inequality being strict, for any pair $s^*, \tilde{s}^* \in \Omega_{\mathcal{G}}$. A solution $s^* \in \Omega_{\mathcal{G}}$ is then characterised as an *admissible* GNE

if there are no other equilibrium $\tilde{s}^* \in \Omega_G$ such that $\tilde{s}^* \prec s^*$. Note that an admissible GNE can still be favourable to only a subset of players (that is, $\tilde{s}^* \not\prec s^* \not\prec \tilde{s}^*$), and thus it is possible that \mathcal{G} has no solution that is “fair” to all players. The game might also not admit any GNE (i.e., $\Omega_G = \emptyset$), then being characterised as unsolvable. Hereafter, we ensure that the problems being discussed are well-posed by considering the following conditions on their primitives:

Assumption 1. For each player $p \in \mathcal{P}$,

a) the objective $L^p : \mathcal{S}^1 \times \dots \times \mathcal{S}^{N_P} \rightarrow \mathbb{R}$ is jointly continuous in all of its arguments and convex in the p -th argument, $s^p \in S^p(s^{-p})$, for every $s^{-p} \in \mathcal{S}^{-p}$.

b) the mapping $S^p : \mathcal{S}^p \rightrightarrows \mathcal{S}^{-p}$ takes the form

$$S^p(s^{-p}) := \{s^p \in \mathcal{S}^p : (s^p, s^{-p}) \in \mathcal{S}_G\},$$

where \mathcal{S}_G is some global constraint set shared by all players. Moreover, S^p and \mathcal{S}_G are both compact convex sets and they satisfy $\mathcal{S}_G \cap (\mathcal{S}^1 \times \dots \times \mathcal{S}^{N_P}) \neq \emptyset$.

Under Assumption 1, a generalisation of the Kakutani fixed-point theorem ensures that \mathcal{G} has a GNE, that is, $\Omega_G \neq \emptyset$ [37]. In practice, these conditions consider each objective to have a unique optimal value, while imposing the feasible set of strategies to be nonempty and coupled only through a common constraint. Although restrictive, these assumptions still cover a broad class of problems of practical relevance.

We investigate algorithms for solving \mathcal{G} . A direct computation of a GNE is equivalent to solving N_P optimisation problems simultaneously, as implied by Definition 1. Such an approach would require players to be coordinated and their objectives to be public. Conversely, we consider adaptive procedures in which the players learn their GNE strategies independently. In this direction, consider that \mathcal{G} admits episodic repetitions, and let $s_k := (s_k^1, \dots, s_k^{N_P})$ be the strategy profile taken by players \mathcal{P} at the k -th episode. A prototypical learning procedure for equilibrium seeking is outlined in Algorithm 1, where

- $T^p : \mathcal{S}^p \times \mathcal{S}^{-p} \rightrightarrows \mathcal{S}^p$ describes how the p -th player updates its strategy, based on its opponents’ predicted next actions and given its individual objective;
- $R^p : \mathcal{S} \rightrightarrows \mathcal{S}^{-p}$ describes how the p -th player’s predicts its opponents’ strategies for the next episode, based on the strategy profile currently being played.

Algorithm 1: Prototypical learning dynamics

Input: Game $\mathcal{G} := (\mathcal{P}, \{S^p\}_{p \in \mathcal{P}}, \{L^p\}_{p \in \mathcal{P}})$
Output: GNE $s^* = (s^{1*}, \dots, s^{N_P*})$

- 1 Initialize $s_0 = (s_0^1, \dots, s_0^{N_P})$ and $k = 0$;
 - 2 **for** $k = 1, 2, \dots$ **do**
 - 3 **if** $s_k \in T(s_k)$ **then return** s_k ;
 - 4 **for** $p \in \mathcal{P}$ **do**
 - 5 Update $s_{k+1}^p \in T^p(s_k^p, R^p(s_k) \mid L^p)$;
-

The procedure in Algorithm 1 belongs to the class of fixed-point methods: Its termination implies that s^* is a fixed-point of both

T^p and R^p , that is, $s^* \in T^p(s^{p*}, R^p(s^*)) \subseteq T^p(s^{p*}, s^{-p*})$. In its general form, the conditions (and convergence rates) for these learning dynamics to approach an equilibrium are difficult to establish. In this work, we build upon a fundamental instance from this class of algorithms: The best-response dynamics (BRD). This routine is overviewed in the following.

Best-response dynamics: Let the mapping $BR^p : \mathcal{S}^{-p} \rightrightarrows \mathcal{S}^p$,

$$BR^p(s^{-p}) := \arg \min_{s^p \in S^p(s^{-p})} L^p(s^p, s^{-p}) \quad (4)$$

denote the best-response of $p \in \mathcal{P}$ to other players’ strategies. Collectively, $BR(s) := BR^1(s^{-1}) \times \dots \times BR^{N_P}(s^{-N_P}) \subseteq \mathcal{S}$ is the joint best-response to any given profile $s \in \mathcal{S}$. Then,

Theorem 1. A strategy profile $s^* = (s^{1*}, \dots, s^{N_P*}) \in \mathcal{S}$ is a GNE for \mathcal{G} if and only if $s^* \in BR(s^*)$ or, equivalently,

$$s^{p*} \in BR^p(s^{-p*}), \quad \forall p \in \mathcal{P}. \quad (5)$$

The task of computing a Nash equilibrium can thus be translated into searching for a fixed-point of the set-valued mapping $BR : \mathcal{S} \rightrightarrows \mathcal{S}$. The set of GNE solutions for \mathcal{G} is the set of all such fixed-points, $\Omega_G := \{s^* \in \mathcal{S} : s^* \in BR(s^*)\}$. A natural procedure for GNE seeking thus consists of players adapting their strategies towards best-responses to their rivals’ strategies, which they assume will remain constant. Formally,

$$T^p(s_k^p, R^p(s_k)) := (1-\eta)s_k^p + \eta BR^p(R^p(s_k)), \quad (6)$$

given $R^p(s_k) = s_k^{-p}$ and a learning rate factor of $\eta \in (0, 1)$. This learning dynamics, summarised in Algorithm 2, is known as best-response dynamics or fictitious play.

Algorithm 2: Best-Response Dynamics (BRD)

Input: Game $\mathcal{G} := (\mathcal{P}, \{S^p\}_{p \in \mathcal{P}}, \{L^p\}_{p \in \mathcal{P}})$
Output: GNE $s^* = (s^{1*}, \dots, s^{N_P*})$

- 1 Initialize $s_0 = (s_0^1, \dots, s_0^{N_P})$ and $k = 0$;
 - 2 **for** $k = 0, 1, 2, \dots$ **do**
 - 3 **if** $s_k \in BR(s_k)$ **then return** s_k ;
 - 4 **for** $p \in \mathcal{P}$ **do**
 - 5 Update $s_{k+1}^p \in (1-\eta)s_k^p + \eta BR^p(s_k^{-p})$;
-

After each episode, the strategy profile gets updated as

$$s_{k+1} = T(s_k) = (1-\eta)s_k + \eta BR(s_k), \quad (7)$$

given the global update rule $T = (1-\eta)I + \eta BR$. Notably, the mappings T and BR share the same set of fixed-points: The GNEs Ω_G . We can then establish the following result.

Theorem 2. Let $BR : \mathcal{S} \rightrightarrows \mathcal{S}$ be a non-expansive mapping. Then, the best-response dynamics $s_{k+1} = T(s_k)$ converge monotonically to a GNE solution $s^* \in \Omega_G$, that is,

$$\lim_{k \rightarrow \infty} \inf_{s^* \in \Omega_G} \|T(s_k) - s^*\| = 0 \quad (8)$$

given any appropriate norm $\|\cdot\|$ for \mathcal{S} .

This convergence result stems from fixed-point theory, where the BRD algorithm is interpreted as belonging to the class of

averaged (or Mann-Krasnosel'skii) iteration methods [38, 39]. For non-generalised quadratic problems, we show in Section III-B that the best-response mapping BR is non-expansive under specific conditions on L^1, \dots, L^{N_P} . Importantly, if the best-response mapping is a contraction, then the BRD converge geometrically to a GNE solution $s^* \in \Omega_G$, which is unique.

Theorem 3. *Let $BR : \mathcal{S} \rightrightarrows \mathcal{S}$ be L_{BR} -Lipschitz, $L_{BR} < 1$. Then, from any feasible $s_0 \in \mathcal{S}$, the best-response dynamics $s_{k+1} = T(s_k)$ converge to the unique GNE $s^* \in \Omega_G$ with rate*

$$\frac{\|s_k - s^*\|}{\|s_0 - s^*\|} \leq ((1-\eta) + \eta L_{BR})^k \quad (9)$$

given any appropriate norm $\|\cdot\|$ for \mathcal{S} .

Importantly, these results might not hold in practice whenever the best-response maps, $\{BR^p\}_{p \in \mathcal{P}}$, are only approximated (e.g., by solving Eq. (4) numerically). However, such *inexact* averaged operators are still known to converge under specific assumptions on the accuracy of this approximation [40]. The learning rate η also plays a central role in the numerical stability of the BRD algorithm: A careful choice is required to ensure that strategy updates do not escape the feasible set, that is, to ensure that $T(s_k) \in \mathcal{S}_G \cap \mathcal{S}$ for all $s_k \in \mathcal{S}_G \cap \mathcal{S}$. Otherwise, the learning dynamics become undefined and the aforementioned convergence results do not apply. For non-generalised games, T trivially satisfy the constraints for any $\eta \in (0, 1)$ and thus a careful design of the learning rate might not be necessary. In such cases, $\eta \rightarrow 1$ is the optimal choice if BR is a contraction and the convergence rate in Eq. (9) simplifies to L_{BR}^k .

Finally, we note that the stopping criteria in Algorithm 2 can be modified to allow for earlier termination. In this case, interrupting the best-response dynamics at some episode $k_f > 0$ will produce a strategy profile $s_{k_f} \in \mathcal{S}$ for which

$$L^p(s_{k_f}^p, s_{k_f}^{-p}) \leq \min_{s^p \in \mathcal{S}^p(s_{k_f}^{-p})} L^p(s^p, s_{k_f}^{-p}) + \varepsilon \quad (10)$$

holds for every player $p \in \mathcal{P}$ with an ‘‘equilibrium gap’’ $\varepsilon > 0$. This profile characterises an ε -GNE: No player can improve its cost more than ε by unilaterally changing his strategy. The set of all ε -GNEs is denoted $\Omega_\varepsilon^G = \{s^\varepsilon \in \mathcal{S} : s^\varepsilon \text{ satisfies Eq. (10)}\}$.

A. Infinite-horizon dynamic games

An infinite-horizon dynamic N_P -player game, denoted by tuple

$$\mathcal{G}_\infty := (\mathcal{P}, \mathcal{X}, \{\mathcal{U}^p\}_{p \in \mathcal{P}}, \mathcal{W}, \{J^p\}_{p \in \mathcal{P}}), \quad (11)$$

is defined by the stochastic linear dynamics

$$x_{t+1} = Ax_t + \sum_{p \in \mathcal{P}} B^p u_t^p + w_t, \quad x_0 \text{ given}, \quad (12)$$

describing how the state of the game, $\mathbf{x} = (x_t)_{t \in \mathbb{N}} \in \mathcal{X}$, evolves in response to the players’ actions $\mathbf{u}^p = (u_t^p)_{t \in \mathbb{N}} \in \mathcal{U}^p$ ($\forall p \in \mathcal{P}$) and the additive random noise $\mathbf{w} = (w_t)_{t \in \mathbb{N}} \in \mathcal{W}$. For each realisation $\mathbf{w} \in \mathcal{W}$ and initial x_0 , the state is explicitly expressed as $\mathbf{x} = \mathbf{F}_w \mathbf{u}$ via the causal affine operator

$$\mathbf{F}_w : \mathbf{u} \mapsto (I - \mathbf{S}_+ \mathbf{A})^{-1} \left(\sum_{p \in \mathcal{P}} \mathbf{S}_+ B^p \mathbf{u}^p + \mathbf{S}_+ \mathbf{w} + \delta x_0 \right), \quad (13)$$

with $\mathbf{A} : \mathbf{x} \mapsto (Ax_t)_{t \in \mathbb{N}}$ and $\mathbf{B}^p : \mathbf{u}^p \mapsto (B^p u_t^p)_{t \in \mathbb{N}}$. Because known, the dependency on x_0 is omitted to simplify notation. Moreover, we assume $E w_t = 0$ and $E(w_t w_{t'}^T) = \delta_{t-t'} \Sigma_w$, given a covariance matrix $\Sigma_w \in \mathbb{S}_{++}^{N_x}$, for every $t, t' \in \mathbb{N}$. Finally, the sets $\mathcal{X}, \mathcal{U}^p$ ($\forall p \in \mathcal{P}$), and \mathcal{W} define all permissible state, action, and noise sequences; they take the form

$$\begin{aligned} \mathcal{X} &:= \{\mathbf{x} \in \ell_\infty^{N_x}(\mathbb{N}) : x_t \in \mathcal{X}, t \in \mathbb{N}\}; \\ \mathcal{U}^p &:= \{\mathbf{u}^p \in \ell_\infty^{N_u^p}(\mathbb{N}) : u_t^p \in \mathcal{U}^p, t \in \mathbb{N}\}; \\ \mathcal{W} &:= \{\mathbf{w} \in \ell_\infty^{N_x}(\mathbb{N}) : w_t \in \mathcal{W}, t \in \mathbb{N}\}, \end{aligned}$$

given sets $\mathcal{X} \subseteq \mathbb{R}^{N_x}, \mathcal{U}^p \subseteq \mathbb{R}^{N_u^p}$ ($\forall p \in \mathcal{P}$), and $\mathcal{W} \subseteq \mathbb{R}^{N_x}$.

In dynamic stochastic games, each player decides on a plan of action $\mathbf{u}^p \in U^p(\mathbf{u}^{-p})$ to minimize its objective functional

$$J^p(\mathbf{u}^p, \mathbf{u}^{-p}) := E \left[\sum_{t=0}^{\infty} L^p(x_t, u_t^p, \mathbf{u}^{-p}) \right], \quad (14)$$

defined by cost function $L^p : \mathcal{X} \times \mathcal{U}^1 \times \dots \times \mathcal{U}^{N_P} \rightarrow \mathbb{R}$. The mappings $U^p : \mathcal{U}^{-p} \rightrightarrows \mathcal{U}^p$ restrict the permissible strategies for each player based on their rivals’ strategies. Under this formulation, the dynamic game \mathcal{G}_∞ is stationary and can be interpreted as a static game defined on the appropriate functional spaces. A plan of action $\mathbf{u} = (\mathbf{u}^1, \dots, \mathbf{u}^{N_P}) \in \mathcal{U}$ can then be characterised as a generalised Nash Equilibrium solution to \mathcal{G}_∞ when no player can improve its objective by unilaterally deviating from this agreed profile. Formally,

Definition 2. A strategy profile $\mathbf{u}^* = (\mathbf{u}^{1*}, \dots, \mathbf{u}^{N_P*}) \in \mathcal{U}$ is a generalized Nash equilibrium (GNE) for the game \mathcal{G}_∞ if

$$J^p(\mathbf{u}^{p*}, \mathbf{u}^{-p*}) \leq \min_{\mathbf{u}^p \in U^p(\mathbf{u}^{-p*})} J^p(\mathbf{u}^p, \mathbf{u}^{-p*}) \quad (15)$$

holds for every player $p \in \mathcal{P}$.

As before, the set of GNEs that solve \mathcal{G}_∞ ,

$$\Omega_{\mathcal{G}_\infty} := \{\mathbf{u}^* \in \mathcal{U} : \mathbf{u}^* \text{ satisfies Eq. (15)}\},$$

is not necessarily a singleton and can include non-admissible equilibria. The game is considered unsolvable if $\Omega_{\mathcal{G}_\infty} = \emptyset$. The following assumptions are taken for this class of games:

Assumption 2. For each player $p \in \mathcal{P}$ and noise $\mathbf{w} \in \mathcal{W}$,

- the cost functional $J^p : \mathcal{U}^1 \times \dots \times \mathcal{U}^{N_P} \rightarrow \mathbb{R}$ is jointly continuous in all of its arguments and convex in the p -th argument, $\mathbf{u}^p \in U^p(\mathbf{u}^{-p})$, for every sequence $\mathbf{u}^{-p} \in \mathcal{U}^{-p}$.
- the mapping $U^p : \mathcal{U}^{-p} \rightrightarrows \mathcal{U}^p$ takes the form

$$U^p(\mathbf{u}^{-p}) := \{\mathbf{u}^p \in \mathcal{U}^p : (\mathbf{u}^p, \mathbf{u}^{-p}) \in \mathbf{F}_w^\dagger(\mathcal{X}) \cap \mathcal{U}_G\},$$

where \mathbf{F}_w^\dagger is a left-inverse of the operator in Eq. (13) and

$$\mathcal{U}_G = \{\mathbf{u} \in \prod_{p \in \mathcal{P}} \ell_\infty^{N_u^p}(\mathbb{N}) : (u_t^p, \mathbf{u}^{-p}) \in \mathcal{U}_G, t \in \mathbb{N}\}$$

is a global constraint set shared by all players. The sets U^p, \mathcal{U}_G , and \mathcal{X} are all non-empty, compact, and convex. Finally, we have that $\mathbf{F}_w^\dagger(\mathcal{X}) \cap \mathcal{U}_G \cap (\mathcal{U}^1 \times \dots \times \mathcal{U}^{N_P}) \neq \emptyset$.

These conditions are analogous to those of Assumption 1: They ensure the existence of GNE solutions to \mathcal{G}_∞ , that is, $\Omega_{\mathcal{G}_\infty} \neq \emptyset$. In this case, the players not only share a global constraint \mathcal{U}_G ,

but are also required to ensure that state trajectories lie in a feasible set, $x \in \mathcal{X}$, against all possible noise realisations. These constraints often describe operational desiderata and/or limitations concerning the game \mathcal{G}_∞ . Moreover, note that Assumption 2 implies that the operator F_w is injective, that is, that distinct control signals cannot produce the same state trajectory. A sufficient condition for this property consists on the matrix $B = [B^1 \ B^2 \ \dots \ B^{N_P}]$ being full-rank.

Any equilibrium $\mathbf{u}^* \in \Omega_{\mathcal{G}_\infty}$ incorporates an open-loop information pattern: Actions $u_t^* = (u_t^{1*}, \dots, u_t^{N_P*})$ depend explicitly only on initial $x_0 \in \mathcal{X}$ and stage $t \in \mathbb{N}$. A plan of action with such representation is undesirable, as players become sensible to noise disturbances and decision errors. Conversely, a state-feedback decision policy, $\mathbf{u}^* = K(x)$ for some $K : \mathcal{X} \rightarrow \mathcal{U}$, often carry desirable robustness properties. In this sense, a feedback (respectively, open-loop) representation of $\mathbf{u}^* \in \Omega_{\mathcal{G}_\infty}$ is said to be strongly (weakly) time consistent [1]. In this work, we investigate state-feedback solutions to the game \mathcal{G}_∞ .

We consider a closed-loop information pattern and assume that each p -th player's actions are represented as

$$\mathbf{u}^p := \mathbf{K}^p \mathbf{x}, \quad \mathbf{K}^p : \mathbf{x} \mapsto \Phi^p * \mathbf{x}, \quad (16)$$

given a causal operator $\mathbf{K}^p \in \mathcal{C}^p \subseteq \mathcal{L}(\mathcal{X}, \mathcal{U}^p)$ defined by its convolution kernel $\Phi^p = (\Phi_n^p)_{n \in \mathbb{N}} \in \ell_1(\mathbb{N})$. The sets $\{\mathcal{C}^p\}_{p \in \mathcal{P}}$ describe the operators that satisfy some \mathcal{G}_∞ -related restrictions (e.g., information patterns incurred by communication, actuation, and sensing delays). In this setting, players do not devise their actions explicitly but rather by designing a state-feedback policy profile $\mathbf{K} := (\mathbf{K}^1, \dots, \mathbf{K}^{N_P}) \in \mathcal{C}$, with $\mathcal{C} = \mathcal{C}^1 \times \dots \times \mathcal{C}^{N_P}$. The solution concept that naturally arises is that of a generalized feedback Nash equilibrium.

Definition 3. A policy profile $\mathbf{K}^* = (\mathbf{K}^{1*}, \dots, \mathbf{K}^{N_P*})$ is a generalized feedback Nash equilibrium (GFNE) for \mathcal{G}_∞ if

$$J^p(\mathbf{u}^{p*}, \mathbf{u}^{-p*}) \leq \min_{\mathbf{u}^p \in \mathcal{U}^p(\mathbf{u}^{-p*})} J^p(\mathbf{u}^p, \mathbf{u}^{-p*}), \quad (17)$$

where $\mathbf{u}^* \in \text{Ker}(I - \mathbf{K}^* \mathbf{F}_w)$, holds for every $p \in \mathcal{P}$.

The set of GFNE that solve \mathcal{G}_∞ is defined as

$$\Omega_{\mathcal{G}_\infty}^{\mathbf{K}} := \{\mathbf{K}^* \in \mathcal{C} : \mathbf{u}^* = \mathbf{K}^* \mathbf{x}^* \text{ satisfies Eq. (17)}\}.$$

We consider $\mathbf{K}^* \in \Omega_{\mathcal{G}_\infty}^{\mathbf{K}}$ to be admissible only if it renders the game stable, that is, if the closed-loop evolution

$$\mathbf{x}^* = (I - \mathbf{S}_+(A - \sum_{p \in \mathcal{P}} B^p \mathbf{K}^{p*}))^{-1} (\mathbf{S}_+ \mathbf{w} + \delta \mathbf{x}_0)$$

is bounded ($\mathbf{x}^* \in \ell_\infty^{N_x}$) for all bounded noise ($\mathbf{w} \in \ell_\infty^{N_w}$). A policy satisfying this requirement is said to be *stabilising*. From Assumption 2, we have $\Omega_{\mathcal{G}_\infty}^{\mathbf{K}} \neq \emptyset$ when $\mathcal{C} = \mathcal{L}(\mathcal{X}, \mathcal{U})$. For a more specific $\mathcal{C} \subseteq \mathcal{L}(\mathcal{X}, \mathcal{U})$, establishing the existence (and, especially, uniqueness) of a solution is demanding [41]. In practice, the set $\Omega_{\mathcal{G}_\infty}^{\mathbf{K}}$ can be constructed from open-loop equilibria $\mathbf{u}^* \in \Omega_{\mathcal{G}_\infty}$ by *i*) parametrising the set of all possible trajectories $\{\mathbf{x}_w^* = \mathbf{F}_w \mathbf{u}^*\}_{w \in \mathcal{W}}$, then *ii*) identifying policies $(\mathbf{K}^{1*}, \dots, \mathbf{K}^{N_P*})$ that satisfy $\{\mathbf{u}^{p*} = \mathbf{K}^{p*} \mathbf{x}_w^*\}_{w \in \mathcal{W}, p \in \mathcal{P}}$. Highlighting this equivalence, we refer to such \mathbf{u}^* as the *open-loop realisation* of the *closed-loop policy* \mathbf{K}^* , and vice-versa.

Best-response dynamics for GFNE seeking: The mapping

$$BR^p(\mathbf{u}^{-p}) := \arg \min_{\mathbf{u}^p \in \mathcal{U}^p(\mathbf{u}^{-p})} J^p(\mathbf{u}^p, \mathbf{u}^{-p}) \quad (18)$$

is the best-response of $p \in \mathcal{P}$ to other players' plan of action. Under its feedback representation, $\mathbf{u}^p = \mathbf{K}^p \mathbf{x} \in BR^p(\mathbf{u}^{-p})$ is a solution to the infinite-horizon control problem

$$\text{minimize}_{\mathbf{u}^p := \mathbf{K}^p \mathbf{x}} \quad \mathbb{E} \left[\sum_{t=0}^{\infty} L^p(x_t, u_t^p, u_t^{-p}) \right] \quad (19a)$$

$$\text{subject to} \quad x_{t+1} = Ax_t + \sum_{\bar{p} \in \mathcal{P}} B^{\bar{p}} u_t^{\bar{p}} + w_t, \quad (19b)$$

$$x_t \in \mathcal{X}, \quad u_t^p \in \mathcal{U}^p, \quad (u_t^p, u_t^{-p}) \in \mathcal{U}_{\mathcal{G}}, \quad (19c)$$

$$\mathbf{K}^p \in \mathcal{C}^p, \quad (19d)$$

$$(x_0 \text{ given}). \quad (19e)$$

While posed in terms of action signals ($\mathbf{u}^p, p \in \mathcal{P}$), Problem (19) should be interpreted as the direct search for a best-response policy \mathbf{K}^p against the (fixed) plan of action from other players, $\mathbf{u}^{-p} := (\mathbf{u}^{\bar{p}})_{\bar{p} \in \mathcal{P} \setminus \{p\}}$. We slightly abuse notation and let $BR^p(\mathbf{K}^{-p})$ be its solutions, when the problem is parametrised by $\mathbf{u}^{-p} := \mathbf{K}^{-p} \mathbf{x} = (\mathbf{K}^{\bar{p}} \mathbf{x})_{\bar{p} \in \mathcal{P} \setminus \{p\}}$. The mapping $BR : \mathcal{C} \rightrightarrows \mathcal{C}$, defined by $BR(\mathbf{K}) = BR^1(\mathbf{K}^{-1}) \times \dots \times BR^{N_P}(\mathbf{K}^{-N_P})$, is the joint best-response to a strategy profile \mathbf{K} . The GFNEs of \mathcal{G}_∞ are thus the fixed-points of this mapping: That is, $\Omega_{\mathcal{G}_\infty}^{\mathbf{K}} = \{\mathbf{K}^* \in \mathcal{C} : \mathbf{K}^* \in BR(\mathbf{K}^*)\}$. Due to constraints $(\mathcal{X}, \mathcal{U}^p, \mathcal{U}_{\mathcal{G}})$ and \mathcal{C}^p , an analytical solution to Problem (19) does not exist. Moreover, because infinite-dimensional, its numerical approximation cannot be obtained.

The BRD method for GFNE seeking is given in Algorithm 3. As \mathcal{G}_∞ is dynamic and stationary, the procedure needs not episodic repetitions of the game. Instead, the learning dynamics occurs simultaneously with the game's execution: Players learn and announce their new policies at stages $t \in \{(k+1)\Delta T\}_{k \in \mathbb{N}}$. $\mathbf{K}_k := (\mathbf{K}_k^1, \dots, \mathbf{K}_k^{N_P})$ denotes the strategy profile after $k \in \mathbb{N}$ updates. The period $\Delta T \geq 1$ defines the rate at which policies are updated, reflecting some communication structure (e.g., the time needed for each $p \in \mathcal{P}$ to collect $\{\mathbf{K}_k^{\bar{p}}\}_{\bar{p} \in \mathcal{P} \setminus \{p\}}$).

Algorithm 3: BRD for GFNE seeking (BRD-GFNE)

Input: Game $\mathcal{G}_\infty := (\mathcal{P}, \mathcal{X}, \{\mathcal{U}^p\}_{p \in \mathcal{P}}, \mathcal{W}, \{J^p\}_{p \in \mathcal{P}})$

Output: GFNE $\mathbf{K}^* = (\mathbf{K}^{1*}, \dots, \mathbf{K}^{N_P*})$

- 1 Initialize $\mathbf{K}_0 = (\mathbf{K}_0^1, \dots, \mathbf{K}_0^{N_P})$ and $k = 0$;
 - 2 **for** $t = 0, 1, 2, \dots$ **do**
 - 3 */** Players apply actions $\{u_{k,t}^p = \mathbf{K}_k^p x_{k,t}\}_{p \in \mathcal{P}}$ **/*
 - 4 **if** $\mathbf{K}_k \in BR(\mathbf{K}_k)$ **then return** \mathbf{K}_k ;
 - 5 **if** $t = (k+1)\Delta T$ **then**
 - 6 **for** $p \in \mathcal{P}$ **do**
 - 7 Update $\mathbf{K}_{k+1}^p \in (1-\eta)\mathbf{K}_k^p + \eta BR^p(\mathbf{K}_k^{-p})$;
 - 7 $k := k + 1$;
-

If we verbally execute Algorithm 3, we have the following:

- The players $p \in \mathcal{P}$ act on \mathcal{G}_∞ according to the policies

$$\mathbf{u}_k^p = \mathbf{K}_k^p \mathbf{x}_k, \quad k \in \mathbb{N},$$

where $\mathbf{u}_k^p = (u_t^p)_{t \in \mathcal{T}_k}$ and $\mathbf{x}_k = (x_t)_{t \in \mathcal{T}_k}$ are the signals restricted to the interval $\mathcal{T}_k = [k\Delta T, (k+1)\Delta T)$.

- At $t = (k+1)\Delta T$, every p -th player updates its policy,

$$\mathbf{K}_{k+1}^p \in (1-\eta)\mathbf{K}_k^p + \eta BR^p(\mathbf{K}_k^{-p}),$$

which is then announced to the other players.

The BRD-GFNE induces an operator $T = (1-\eta)I + \eta BR$ which is equivalent to the update rule of its static counterpart. Thus, it possesses the same properties: The learning dynamics converge if BR is non-expansive and the convergence rate is geometric if BR is also a contraction (Theorems 2–3). These properties can also be stated in terms of stage indices $t \in \mathbb{N}$ by substituting $k = \lfloor t/\Delta T \rfloor$. As in the static case, a careful choice of the learning rate $\eta \in (0, 1)$ is required to ensure that this fixed-point iteration is well-defined. Moreover, the BRD-GFNE can be interrupted at any episode $k_f > 1$ (producing an ϵ -GFNE, $\epsilon > 0$) and, if necessary, a central coordinator assigned to project the resulting strategy profile \mathbf{K}_{k_f} onto the feasible set $\mathbf{F}_w^+(\mathcal{X}) \cap \mathcal{U}_G \cap (\mathcal{U}^1 \times \dots \times \mathcal{U}^{N_P})$. Finally, we note that such considerations can be disregarded when the game is not generalised, that is, when $\mathcal{X} = \mathbb{R}^{N_x}$ and $\mathcal{U}_G = \prod_{p \in \mathcal{P}} \mathbb{R}^{N_u^p}$.

III. BEST-RESPONSE DYNAMICS VIA SYSTEM LEVEL SYNTHESIS

In this section, we present an approach for GFNE seeking in (stationary) stochastic dynamic games. Firstly, we introduce the system level parametrisation of the players' feedback policies (\mathbf{K}^p , $p \in \mathcal{P}$) and reformulate their best-response mappings (BR^p , $p \in \mathcal{P}$) through finite-dimensional robust optimisation problems. Then, a modified BRD-GFNE procedure is proposed and its convergence properties are investigated.

We focus on N_P -players linear-quadratic stochastic games $\mathcal{G}_{\infty}^{\text{LQ}} = (\mathcal{P}, \mathcal{X}, \{\mathcal{U}^p\}_{p \in \mathcal{P}}, \mathcal{W}, \{J^p\}_{p \in \mathcal{P}})$ with dynamics

$$x_{t+1} = Ax_t + \sum_{p \in \mathcal{P}} B^p u_t^p + w_t, \quad x_0 \text{ given}, \quad (20)$$

and objective functionals

$$J^p(\mathbf{u}^p, \mathbf{u}^{-p}) = \mathbb{E} \left[\sum_{t=0}^{\infty} \left(\|C^p x_t\|_2^2 + \|\sum_{\tilde{p} \in \mathcal{P}} D^{p\tilde{p}} u_t^{\tilde{p}}\|_2^2 \right) \right], \quad (21)$$

defined by matrices $C^p \in \mathbb{R}^{N_z \times N_x}$ and $D^{p\tilde{p}} \in \mathbb{R}^{N_z \times N_u^{\tilde{p}}}$ with dimension $N_z \geq N_x + N_u$. The following assumptions ensure that GFNE of $\mathcal{G}_{\infty}^{\text{LQ}}$ are admissible (that is, stabilising):

Assumption 3. For each player $p \in \mathcal{P}$,

- The pair (A, B^p) is stabilisable;
- The pair (C^p, A) is detectable;
- The matrix D^{pp} is full column rank, i.e., $D^{pp\top} D^{pp} \in \mathbb{S}_{++}^{N_u^p}$. Moreover, $D^{p\tilde{p}\top} C^p = 0 = C^{p\top} D^{p\tilde{p}}$ for all $\tilde{p} \in \mathcal{P}$.

Finally, the constraint sets \mathcal{X} and \mathcal{U}^p ($\forall p$) are convex polyhedra and contain their respective origins: $0 \in \mathcal{X}$ and $0 \in \mathcal{U}^p$.

The class $\mathcal{G}_{\infty}^{\text{LQ}}$ describe problems in which N_P non-cooperative agents have to agree on stationary policies that jointly stabilise

a global system, robustly to the noise process, while penalising state- and input-deviations differently. While representative of many practically relevant problems, this choice is not restrictive. Our derivations should follow similarly for any collection of cost functions $\{L^p\}_{p \in \mathcal{P}}$ satisfying Assumption 2.

A. System-level best-response mappings

System level synthesis (SLS, [36]) is a novel methodology for controller design which lifts the search for optimal controllers into the search of optimal closed-loop responses to disturbances. Under this approach, the synthesis of stationary control policies can be equivalently posed as the solution to scalable numerical problems, even when operational constraints on the state and control signals need to be enforced. Additionally, the method allows for constraints to be enforced directly on the structure of the control policies. In this section, we present a system-level parametrisation for the best-response mappings in $\mathcal{G}_{\infty}^{\text{LQ}}$.

We start by assuming a stabilising profile $(\mathbf{K}^1, \dots, \mathbf{K}^{N_P})$, guaranteed by Assumption 3. Each policy is associated with a transfer matrix $\hat{\mathbf{K}}^p \in \mathcal{RH}_{\infty}$, $\hat{\mathbf{K}}^p = \sum_{n=0}^{\infty} \frac{1}{z^n} \Phi_n^p$, which defines the state-feedback $\hat{\mathbf{u}}^p = \hat{\mathbf{K}}^p \hat{\mathbf{x}}$ in the frequency domain. Considering the linear dynamics Eq. (20),

$$z\hat{\mathbf{x}} = A\hat{\mathbf{x}} + \sum_{p \in \mathcal{P}} B^p \hat{\mathbf{u}}^p + \hat{\mathbf{w}}; \quad (22a)$$

$$\hat{\mathbf{u}}^p = \hat{\mathbf{K}}^p \hat{\mathbf{x}}, \quad (\forall p \in \mathcal{P}), \quad (22b)$$

the signals $(\hat{\mathbf{x}}, \hat{\mathbf{u}}^1, \dots, \hat{\mathbf{u}}^{N_P})$ can be expressed in terms of $\hat{\mathbf{w}}$,

$$\begin{bmatrix} \hat{\mathbf{x}} \\ \hat{\mathbf{u}}^1 \\ \vdots \\ \hat{\mathbf{u}}^{N_P} \end{bmatrix} = \begin{bmatrix} (zI - A - \sum_{p \in \mathcal{P}} B^p \hat{\mathbf{K}}^p)^{-1} \\ \hat{\mathbf{K}}^1 (zI - A - \sum_{p \in \mathcal{P}} B^p \hat{\mathbf{K}}^p)^{-1} \\ \vdots \\ \hat{\mathbf{K}}^{N_P} (zI - A - \sum_{p \in \mathcal{P}} B^p \hat{\mathbf{K}}^p)^{-1} \end{bmatrix} \hat{\mathbf{w}}; \quad (23)$$

$$= \begin{bmatrix} \hat{\Phi}_x \\ \hat{\Phi}_u^1 \\ \vdots \\ \hat{\Phi}_u^{N_P} \end{bmatrix} \hat{\mathbf{w}},$$

where the introduced transfer matrices $(\hat{\Phi}_x, \hat{\Phi}_u^1, \dots, \hat{\Phi}_u^{N_P})$ are denoted as *system level responses* or *closed-loop maps*. Under this representation, the following result holds.

Theorem 4 (System level parametrisation, [36]). *Consider the dynamics Eq. (22) under state-feedback $\hat{\mathbf{u}}^p = \hat{\mathbf{K}}^p \hat{\mathbf{x}}$ ($\forall p \in \mathcal{P}$). The following statements are true:*

- The affine space

$$\begin{bmatrix} zI - A & -B^1 & \dots & -B^{N_P} \end{bmatrix} \begin{bmatrix} \hat{\Phi}_x \\ \hat{\Phi}_u^1 \\ \vdots \\ \hat{\Phi}_u^{N_P} \end{bmatrix} = I, \quad (24)$$

with $\hat{\Phi}_x, \hat{\Phi}_u^1, \dots, \hat{\Phi}_u^{N_P} \in \frac{1}{z} \mathcal{RH}_{\infty}$, parametrizes all system responses from $\hat{\mathbf{w}}$ to $(\hat{\mathbf{x}}, \hat{\mathbf{u}}^1, \dots, \hat{\mathbf{u}}^{N_P})$ achievable by internally stabilising policies $(\hat{\mathbf{K}}^1, \dots, \hat{\mathbf{K}}^{N_P})$.

b) Any response $(\hat{\Phi}_x, \hat{\Phi}_u^1, \dots, \hat{\Phi}_u^{N_P})$ satisfying Eq. (24) is achieved by the policies $\hat{K}^p = \hat{\Phi}_u^p \hat{\Phi}_x^{-1}$ ($\forall p \in \mathcal{P}$), which are internally stabilising and can be implemented as

$$z\hat{\xi} = \tilde{\Phi}_x \hat{\xi} + \hat{x}; \quad (25a)$$

$$\hat{u}^p = \tilde{\Phi}_u^p \hat{\xi}, \quad (25b)$$

with $\tilde{\Phi}_x = z(I - \hat{\Phi}_x)$ and $\tilde{\Phi}_u^p = z\hat{\Phi}_u^p$.

Proof. To simplify the notation, we momentarily define the matrices $B := [B^1 \ B^2 \ \dots \ B^{N_P}]$, $\hat{\Phi}_u := [\hat{\Phi}_u^1 \ \hat{\Phi}_u^2 \ \dots \ \hat{\Phi}_u^{N_P}]^T$, and $\hat{K} := [\hat{K}^1 \ \hat{K}^2 \ \dots \ \hat{K}^{N_P}]^T$. For the first statement, consider internally stabilising policies $(\hat{K}^1, \dots, \hat{K}^{N_P})$ such that $\hat{u}^p = \hat{K}^p \hat{x}$ ($\forall p \in \mathcal{P}$). Then, from Eq. (23), we have that

$$\begin{aligned} & [zI - A \quad -B] \begin{bmatrix} (zI - A - B\hat{K})^{-1} \\ \hat{K}(zI - A - B\hat{K})^{-1} \end{bmatrix} \\ &= (zI - A)(zI - A - B\hat{K})^{-1} - B\hat{K}(zI - A - B\hat{K})^{-1} \\ &= (zI - A - B\hat{K})(zI - A - B\hat{K})^{-1} \\ &= I. \end{aligned}$$

For the second statement, we first show that \hat{K} achieves the desired response, and then that these policies are internally stabilising. Note that Eq. (24) implies that $\hat{\Phi}_x$ has the leading spectral component $\Phi_{x,1} = I_{N_x}$, which is invertible, and thus $\hat{\Phi}_x^{-1}$ exists. Then, $\hat{K} = \hat{\Phi}_u \hat{\Phi}_x^{-1}$ is well-defined, and we have

$$\begin{aligned} \hat{x} &= (zI - A - B\hat{\Phi}_u \hat{\Phi}_x^{-1})^{-1} \hat{w} \\ &= \hat{\Phi}_x \left((zI - A) \hat{\Phi}_x - B\hat{\Phi}_u \right)^{-1} \hat{w} \\ &= \hat{\Phi}_x \hat{w}, \end{aligned}$$

due to Eq. (24), and $\hat{u} = \hat{K} \hat{x} = \hat{\Phi}_u \hat{\Phi}_x^{-1} \hat{\Phi}_x \hat{w} = \hat{\Phi}_u \hat{w}$. Thus, \hat{K} achieves the response $(\hat{\Phi}_x, \hat{\Phi}_u)$, or, equivalently, $(\hat{K}^1, \hat{K}^2, \dots, \hat{K}^{N_P})$ achieve $(\hat{\Phi}_x, \hat{\Phi}_u^1, \dots, \hat{\Phi}_u^{N_P})$. To show that this policy is internally stabilising, consider its equivalent representation $\tilde{K} = \tilde{\Phi}_u (zI - \tilde{\Phi}_x)^{-1} \tilde{\Phi}_y$ with $\tilde{\Phi}_x = z(I - \hat{\Phi}_x)$, $\tilde{\Phi}_u = z\hat{\Phi}_u$, and $\tilde{\Phi}_y = I$, as depicted in Figure 2.

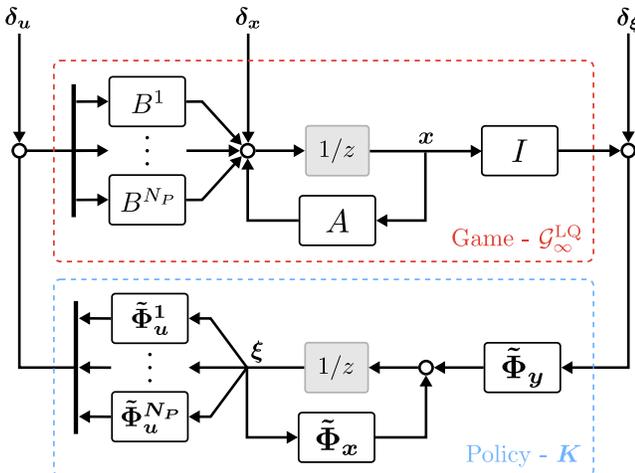


Figure 2. Feedback structure for the policy $\tilde{K} = \tilde{\Phi}_u (zI - \tilde{\Phi}_x)^{-1} \tilde{\Phi}_y = \hat{\Phi}_u^p \hat{\Phi}_x^{-1}$, equivalent to the internal representation in Eq. (25).

Introducing external perturbations $\{\delta_x, \delta_u, \delta_\xi\} \subseteq \ell_\infty$ into the game, it suffices to verify that the transfer matrices from $(\hat{\delta}_x, \hat{\delta}_u, \hat{\delta}_\xi)$ to $(\hat{x}, \hat{u}, \hat{\xi})$ are stable. In this case, we have

$$\begin{bmatrix} \hat{x} \\ \hat{u} \\ \hat{\xi} \end{bmatrix} = \begin{bmatrix} \hat{\Phi}_x & \hat{\Phi}_x B & \hat{\Phi}_x (zI - A) \\ \hat{\Phi}_u & I + \hat{\Phi}_u B & \hat{\Phi}_u (zI - A) \\ \frac{1}{z} I & \frac{1}{z} B & \frac{1}{z} (zI - A) \end{bmatrix} \begin{bmatrix} \hat{\delta}_x \\ \hat{\delta}_u \\ \hat{\delta}_\xi \end{bmatrix}. \quad (26)$$

Since $\hat{\Phi}_x, \hat{\Phi}_u \in \frac{1}{z} \mathcal{RH}_\infty$, the transfer matrices in Eq. (26) are all stable. We thus conclude that the policy \tilde{K} is internally stabilising, and so are the policies $(\tilde{K}^1, \tilde{K}^2, \dots, \tilde{K}^{N_P})$. \square

We refer to the system level responses through their kernels, $\Phi_x = (\Phi_{x,n})_{n \in \mathbb{N}} \in \ell_2(\mathbb{N})$ and $\Phi_u = (\Phi_{u,n})_{n \in \mathbb{N}} \in \ell_2(\mathbb{N})$, for all $p \in \mathcal{P}$. Due to strict causality, $\Phi_{x,0} = 0$ and $\Phi_{u,0} = 0$. From Theorem 4, the operators $\{K^p \in \mathcal{C}^p\}_{p \in \mathcal{P}}$ and the transfer matrices $\{\hat{K}^p \in \mathcal{RH}_\infty\}_{p \in \mathcal{P}}$ are equivalent representations of the feedback policies. Hence, provided there is no confusion, we use exclusively the first notation. In particular, $K^p = \Phi_u^p \Phi_x^{-1}$ ($p \in \mathcal{P}$) denotes the policy parametrised by $(\hat{\Phi}_x, \hat{\Phi}_u^p)$ and $K = (\Phi_u^1, \dots, \Phi_u^{N_P}) \Phi_x^{-1}$ denotes the corresponding profile. A time-domain characterisation of K is given in the following.

Corollary 4.1. A policy $K^p = \Phi_u^p \Phi_x^{-1}$ ($\forall p \in \mathcal{P}$) is defined by the kernel $\Phi^p = \Phi_u^p * \Phi_x^{-1}$, and implemented as

$$\xi_t = -\sum_{\tau=1}^t \Phi_{x,\tau+1} \xi_{t-\tau} + x_t; \quad (27a)$$

$$u_t^p = \sum_{\tau=0}^t \Phi_{u,\tau+1}^p \xi_{t-\tau}, \quad (27b)$$

using an auxiliary internal state $\xi = (\xi_n)_{n \in \mathbb{N}}$ with $\xi_0 = x_0$.

The system level parametrisation enables a methodology for policy synthesis consisting of searching the space of stabilising policies (in)directly through (Φ_x, Φ_u^p) , $p \in \mathcal{P}$. In particular, this parametrisation can be leveraged to reformulate the best-response dynamics maps in $\mathcal{G}_\infty^{\text{LQ}}$ as tractable numerical programs. In this direction, consider that players design stabilising policies $K = (K^1, \dots, K^{N_P})$ by choosing their desired system level responses $\Phi_u = (\Phi_u^1, \dots, \Phi_u^{N_P})$ simultaneously. From the affine space Eq. (24), the signal Φ_x , common to all players, satisfies the (deterministic) linear dynamics

$$\Phi_{x,n+1} = A\Phi_{x,n} + \sum_{p \in \mathcal{P}} B^p \Phi_{u,n}^p, \quad \Phi_{x,1} = I_{N_x}, \quad (28)$$

or $\Phi_x = F_\Phi \Phi_u$ given the causal affine operator

$$F_\Phi : \Phi_u \mapsto (I - S_+ A)^{-1} \left(\sum_{p \in \mathcal{P}} S_+ B^p \Phi_u^p + \delta I_{N_x} \right), \quad (29)$$

Using Parseval's Theorem [42], we substitute Eq. (23) into Eq. (21) to redefine the objective functionals of $\mathcal{G}_\infty^{\text{LQ}}$ as

$$\begin{aligned} & J^p(\Phi_u^p, \Phi_u^{-p}) \\ &= \mathbb{E} \left[\sum_{n=1}^{\infty} \left(\|C^p \Phi_{x,n} w_n\|_F^2 + \|\sum_{\tilde{p} \in \mathcal{P}} D^{p\tilde{p}} \Phi_{u,n}^{\tilde{p}} w_n\|_F^2 \right) \right] \\ &= \sum_{n=1}^{\infty} \left(\|C^p \Phi_{x,n} \Sigma_w^{\frac{1}{2}}\|_F^2 + \|\sum_{\tilde{p} \in \mathcal{P}} D^{p\tilde{p}} \Phi_{u,n}^{\tilde{p}} \Sigma_w^{\frac{1}{2}}\|_F^2 \right). \end{aligned}$$

The game $\mathcal{G}_\infty^{\text{LQ}}$ thus induces a *system-level* dynamic game,

$$\mathcal{G}_\infty^\Phi := (\mathcal{P}, \mathcal{C}_x, \{\mathcal{C}_u^p\}_{p \in \mathcal{P}}, \mathcal{W}, \{J^p\}_{p \in \mathcal{P}}), \quad (30)$$

defining the problem in which players $p \in \mathcal{P}$ each plans a closed-loop response $\Phi_u^p \in U_{\Phi}^p(\Phi_u^{-p}) \subseteq \mathcal{C}_u^p$ to minimise its individual cost functional $J^p : \mathcal{C}_u^1 \times \dots \times \mathcal{C}_u^{N_P} \rightarrow \mathbb{R}$. Here, the set-valued mappings $U_{\Phi}^p : \mathcal{C}_u^{-p} \rightrightarrows \mathcal{C}_u^p$ are defined as

$$U_{\Phi}^p(\Phi_u^{-p}) := \{\Phi_u^p \in \mathcal{C}_u^p : (\Phi_u^p, \Phi_u^{-p}) \in F_{\Phi}^{\dagger}(\mathcal{C}_x), \\ \Phi_u^p * w \in U^p(\Phi_u^{-p} * w)\},$$

which incorporate the constraints U^p from the original $\mathcal{G}_{\infty}^{\text{LQ}}$. The sets $(\mathcal{C}_x, \mathcal{C}_u^p)$ are designed to enforce the policy constraints $K^p \in \mathcal{C}^p$ directly through the kernels (Φ_x, Φ_u^p) : They are related as $\mathcal{C}_u^p = \{K^p \mathcal{C}_x : K^p \in \mathcal{C}^p\}$. We refer to a joint response $\Phi_u = (\Phi_u^1, \dots, \Phi_u^{N_P}) \in \mathcal{C}_u$, $\mathcal{C}_u = \mathcal{C}_u^1 \times \dots \times \mathcal{C}_u^{N_P}$, as a system-level strategy profile. Finally, the set of (open-loop) system-level GNEs for this game is denoted as $\Omega_{\mathcal{G}_{\infty}^{\Phi}}$.

The best-response mappings for $\mathcal{G}_{\infty}^{\Phi}$ take the form

$$BR_{\Phi}^p(\Phi_u^{-p}) := \arg \min_{\Phi_u^p \in U_{\Phi}^p(\Phi_u^{-p})} J^p(\Phi_u^p, \Phi_u^{-p}),$$

consisting of the set of closed-loop maps Φ_u^p which are best-responses to the maps of other players, $\Phi_u^{-p} = (\Phi_u^{\bar{p}})_{\bar{p} \in \mathcal{P} \setminus \{p\}}$. They are solutions to the system level synthesis problem

$$\text{minimize}_{\Phi_u^p} \sum_{n=1}^{\infty} \left(\|C^p \Phi_{x,n} \Sigma_w^{\frac{1}{2}}\|_F^2 + \left\| \sum_{\bar{p} \in \mathcal{P}} D^{p\bar{p}} \Phi_{u,n}^{\bar{p}} \Sigma_w^{\frac{1}{2}} \right\|_F^2 \right) \quad (31a)$$

$$\text{subject to } \forall n \in \mathbb{N}^+ \quad \Phi_{x,n+1} = A\Phi_{x,n} + \sum_{\bar{p} \in \mathcal{P}} B^{\bar{p}} \Phi_{u,n}^{\bar{p}}, \quad (31b)$$

$$(\Phi_x * w)_n \in \mathcal{X}, \quad (\Phi_u^p * w)_n \in \mathcal{U}^p, \\ ((\Phi_u^p * w)_n, (\Phi_u^{-p} * w)_n) \in \mathcal{U}_{\mathcal{G}}, \quad (31c)$$

$$\Phi_x \in \mathcal{C}_x, \quad \Phi_u^p \in \mathcal{C}_u^p, \quad (31d)$$

$$\Phi_{x,1} = I_{N_x}. \quad (31e)$$

Collectively, the mapping $BR_{\Phi} : \mathcal{C}_u \rightrightarrows \mathcal{C}_u$, defined by $BR_{\Phi}(\Phi_u) = BR_{\Phi}^1(\Phi_u^{-1}) \times \dots \times BR_{\Phi}^{N_P}(\Phi_u^{-N_P})$, is the joint best-response to a system-level strategy profile Φ_u . As before, the GNEs of $\mathcal{G}_{\infty}^{\Phi}$ are equivalent to the fixed-points of this map, $\Omega_{\mathcal{G}_{\infty}^{\Phi}} = \{\Phi_u^* \in \mathcal{C}_u : \Phi_u^* \in BR_{\Phi}(\Phi_u^*)\}$. Considering how $\mathcal{G}_{\infty}^{\text{LQ}}$ induces $\mathcal{G}_{\infty}^{\Phi}$, a relationship can be established between the best-responses BR and BR_{Φ} and, consequently, between their fixed-points, $\Omega_{\mathcal{G}_{\infty}^{\text{LQ}}}^K$ and $\Omega_{\mathcal{G}_{\infty}^{\Phi}}$. In Section III-B, we formalise this relationship and propose a learning dynamics for GFNE seeking based on the system-level best-response mappings.

Remark 1. Problem (31) is independent of the matrix $\Sigma_w^{\frac{1}{2}}$ due to its linear-quadratic structure and the fact that $\Sigma_w \in \mathbb{S}^{N_x}$. Hence, we remove it from J^p ($\forall p$) to simplify notation.

The best-responses $\{BR_{\Phi}^p\}_{p \in \mathcal{P}}$ are still intractable: *i*) They are defined by infinite-dimensional problems with no general solution and *ii*) that require full knowledge of the noise process $(w_n)_{n \in \mathbb{N}}$ to formulate the constraints U_{Φ}^p . In the following, we tackle both issues and provide a class of finite-dimensional robust optimisation programs that approximate Problem (31). We conclude the section by presenting a class of system level constraints which enforce a richer feedback information pattern.

Finite-horizon approximation: The programs in $\{BR_{\Phi}^p\}_{p \in \mathcal{P}}$ can be made finite-dimensional by restricting the closed-loop maps to the set of finite-impulse responses (FIR),

$$\mathcal{C}_x = \{\Phi_x \in \ell_2[0, N] : \Phi_{x,n} \in \mathcal{C}_{x,n}, n \in [0, N], \Phi_{x,N} = 0\}; \\ \mathcal{C}_u^p = \{\Phi_u^p \in \ell_2[0, N] : \Phi_{u,n}^p \in \mathcal{C}_{u,n}^p, n \in [0, N]\},$$

given horizon $N \in [2, \infty)$. We enforce $\Phi_x \in \mathcal{C}_x$ and $\Phi_u^p \in \mathcal{C}_u^p$ in Problem (31) by adding the terminal constraint $\Phi_{x,N} = 0$, and restricting the spectral factors to satisfy $(\Phi_{x,n} \in \mathcal{C}_{x,n})_{n=1}^N$ and $(\Phi_{u,n}^p \in \mathcal{C}_{u,n}^p)_{n=1}^{N-1}$. The constraint sets $\mathcal{C}_{x,n} \subseteq \mathbb{R}^{N_x \times N_x}$ and $\mathcal{C}_{u,n}^p \subseteq \mathbb{R}^{N_u \times N_x}$ ($\forall p$) have no restriction other than being compact convex sets. The resulting problems are of the form

$$\text{minimize}_{\Phi_u^p = (\Phi_{u,n}^p)_{n=1}^{N-1}} \sum_{n=1}^N \left(\|C^p \Phi_{x,n}\|_F^2 + \left\| \sum_{\bar{p} \in \mathcal{P}} D^{p\bar{p}} \Phi_{u,n}^{\bar{p}} \right\|_F^2 \right) \quad (32a)$$

$$\text{subject to } \forall n \in [1, N) \quad \Phi_{x,n+1} = A\Phi_{x,n} + \sum_{\bar{p} \in \mathcal{P}} B^{\bar{p}} \Phi_{u,n}^{\bar{p}}, \quad (32b)$$

$$(\Phi_x * w)_n \in \mathcal{X}, \quad (\Phi_u^p * w)_n \in \mathcal{U}^p, \\ ((\Phi_u^p * w)_n, (\Phi_u^{-p} * w)_n) \in \mathcal{U}_{\mathcal{G}}, \quad (32c)$$

$$\Phi_{x,n} \in \mathcal{C}_{x,n}, \quad \Phi_{u,n}^p \in \mathcal{C}_{u,n}^p, \quad (32d)$$

$$\Phi_{x,1} = I_{N_x}, \quad \Phi_{x,N} = 0, \quad (32e)$$

which define finite-dimensional convex programs that can be solved numerically. Since $(\mathcal{C}_x, \mathcal{C}_u^p)$ are finite-dimensional, the operators (F_{Φ}, K) and the convolutions in Eq. (32c) can be represented by matrix multiplications: Consider any kernel $\Phi \in \{\Phi_x\} \cup \{\Phi^p, \Phi_u^p\}_{p \in \mathcal{P}}$ and some signal $z \in \ell_{\infty}$. Then,

$$z^{\text{out}} = \Phi * z \iff z^{\text{out}} = M_{\Phi} z,$$

given block-Toeplitz matrix $M_{\Phi} = [\Phi_{i-j}]_{i \in [1, 2N+1], j \in [1, N+1]}$ and (z, z^{out}) being a vector representation of these signals, that is, $z = \text{col}(z_0, \dots, z_N)$ and $z^{\text{out}} = \text{col}(z_0^{\text{out}}, \dots, z_{2N}^{\text{out}})$. Moreover, $z_n^{\text{out}} = (\Phi * z)_n = [M_{\Phi}]_n z$, where $[M_{\Phi}]_n$ is the n -th block-row of M_{Φ} . Thus, Problem (32) can be reformulated into a program requiring only matrix algebra. Finally, we remark that policies $(K^p, p \in \mathcal{P})$ are implemented either by using Corollary 4.1 or by directly computing kernels $\Phi^p = ([M_{\Phi_u^p}]_n \Phi_x^{-1})_{n \in \mathbb{N}}$.

Although realising Problem (31) into a tractable program, the Problem (32) is only feasible when the pair (A, B^p) is full-state controllable. This is a difficult requirement in multi-agent settings, as often $N_u^p \ll N_x$ for all $p \in \mathcal{P}$, leading to Problem (32) becoming overdetermined. Furthermore, enforcing FIR constraints is known to result in *deadbeat* policies: Control actions are excessively large in magnitude for small $N < \infty$. Alternatively, we restrict the system level responses to the sets

$$\mathcal{C}_x = \{\Phi_x \in \ell_2[0, N] : \Phi_{x,n} \in \mathcal{C}_{x,n}, n \in [0, N], \|\Phi_{x,N}\|_F^2 \leq \gamma\}; \\ \mathcal{C}_u^p = \{\Phi_u^p \in \ell_2[0, N] : \Phi_{u,n}^p \in \mathcal{C}_{u,n}^p, n \in [0, N]\},$$

with $\|\Phi_{x,N}\|_F^2 = \sum_i \sigma_i(\Phi_{x,N})^2 \leq \gamma$ for some factor $\gamma \in (0, 1)$ and $\sigma_i(\cdot)$ denoting the i -th largest singular value of a matrix.

These are referred to as the set of approximately FIR for $N > 0$. The p -th player best-response map thus takes the form

$$\begin{aligned} \underset{\Phi_u^p = (\Phi_{u,n}^p)_{n=1}^{N-1}}{\text{minimize}} \quad & \sum_{n=1}^{N-1} \left(\|C^p \Phi_{x,n}\|_F^2 + \left\| \sum_{\bar{p} \in \mathcal{P}} D^{p\bar{p}} \Phi_{u,n}^{\bar{p}} \right\|_F^2 \right) \\ & + \|C^p \Phi_{x,N}\|_F^2 \end{aligned} \quad (33a)$$

$$\text{subject to } \Phi_{x,n+1} = A\Phi_{x,n} + \sum_{\bar{p} \in \mathcal{P}} B^{\bar{p}} \Phi_{u,n}^{\bar{p}}, \quad (33b)$$

$$\begin{aligned} (\Phi_x * w)_n \in \mathcal{X}, \quad (\Phi_u^p * w)_n \in \mathcal{U}^p, \\ ((\Phi_u^p * w)_n, (\Phi_u^{\bar{p}} * w)_n) \in \mathcal{U}_G, \end{aligned} \quad (33c)$$

$$\Phi_{x,n} \in \mathcal{C}_{x,n}, \quad \Phi_{u,n}^p \in \mathcal{C}_{u,n}^p, \quad (33d)$$

$$\Phi_{x,1} = I_{N_x}, \quad \|\Phi_{x,N}\|_F^2 \leq \gamma. \quad (33e)$$

The solutions to Problem (33) approximate those of the infinite-horizon Problem (31): With respect to N , the performance of the former converges to that achieved by the latter [36]. In this case, feasibility only requires (A, B^p) stabilisable and a sufficiently large horizon N to ensure that $\|\Phi_{x,N}\|_F^2 \leq \gamma$ is achievable for some $\Phi_u^p \in U_\Phi^p(\Phi_u^{-p})$. Computationally, this is still a finite-dimensional convex problem which can be solved numerically. Here, we let $\widehat{BR}_\Phi^p : \mathcal{C}_u^{-p} \rightrightarrows \mathcal{C}_u^p$ be the solutions of Problem (33) parametrised by Φ_u^{-p} . The map $\widehat{BR}_\Phi : \mathcal{C}_u \rightrightarrows \mathcal{C}_u$, $\widehat{BR}_\Phi(\Phi_u) = \widehat{BR}_\Phi^1(\Phi_u^{-1}) \times \cdots \times \widehat{BR}_\Phi^{N_p}(\Phi_u^{-N_p})$, is the joint (approximately) best-response to the system-level profile Φ_u .

Conversely to BR_Φ , the fixed-points of \widehat{BR}_Φ do not coincide with the set of GNEs $\Omega_{\mathcal{G}_\Phi}$, but are rather contained in the set of ε -GNEs $\Omega_{\mathcal{G}_\Phi}^\varepsilon$ for some equilibrium gap $\varepsilon > 0$. This is clear from the fact that the original Problem (31) and the approximation Problem (33) have different optimal values. Under certain conditions, this fact can be shown explicitly.

Theorem 5. Consider a fixed-point $\Phi_u^\varepsilon \in \widehat{BR}_\Phi(\Phi_u^\varepsilon)$ and assume that $\|\Phi_{x,N}^*\|_F^2 \leq \gamma$ for $\Phi_x^* = F_\Phi \Phi_u^*$ obtained from the original best-response $\Phi_u^* \in BR_\Phi(\Phi_u^\varepsilon)$. Then, the profile $\Phi_u^\varepsilon = (\Phi_u^{1^\varepsilon}, \dots, \Phi_u^{N_p^\varepsilon})$ is an ε -GNE of \mathcal{G}_∞^Φ satisfying

$$J^p(\Phi_u^{p^\varepsilon}, \Phi_u^{-p^\varepsilon}) \leq \min_{\Phi_u^p \in U_\Phi^p(\Phi_u^{-p^\varepsilon})} J^p(\Phi_u^p, \Phi_u^{-p^\varepsilon}) + \varepsilon \quad (34)$$

with $\varepsilon = \max_{p \in \mathcal{P}} \gamma J^p(\Phi_u^{p^\varepsilon}, \Phi_u^{-p^\varepsilon})$ for every player $p \in \mathcal{P}$.

The equilibrium gap associated with $\Phi_u^\varepsilon \in \widehat{BR}_\Phi(\Phi_u^\varepsilon)$ is thus proportional to the parameter γ , assuming that the terminal constraint Eq. (33e) also holds for solutions to the original Problem (31). Hereafter, $U_\Phi^p : \mathcal{C}_u^{-p} \rightrightarrows \mathcal{C}_u^p$ ($\forall p$) are assumed to incorporate the pair $(\mathcal{C}_x, \mathcal{C}_u^p)$ defining the set of (soft-constrained) FIR approximations defined above for a $N > 0$. We write $U_{\Phi,N}^p$ whenever this needs to be made explicit (e.g., to distinguish this choice from a more general \mathcal{C}_x and \mathcal{C}_u^p).

Robust operational constraints: From Assumption 3, the sets \mathcal{X} , \mathcal{U}^p , and \mathcal{U}_G , can be expressed by linear inequalities,

$$\mathcal{X} = \{x_t \in \mathbb{R}^{N_x} : G_x x_t \leq g_x\};$$

$$\mathcal{U}^p = \{u_t^p \in \mathbb{R}^{N_u^p} : G_u^p u_t^p \leq g_u^p\};$$

$$\mathcal{U}_G = \{u_t \in \prod_{\bar{p} \in \mathcal{P}} \mathbb{R}^{N_u^{\bar{p}}} : \sum_{\bar{p} \in \mathcal{P}} G_G^{\bar{p}} u_t^{\bar{p}} \leq g_G\},$$

given some matrices $G_x \in \mathbb{R}^{N_x \times N_x}$, $G_u^p \in \mathbb{R}^{N_u^p \times N_u^p}$ and $G_G^{\bar{p}} \in \mathbb{R}^{N_{u_G} \times N_{u_G}^{\bar{p}}}$ ($\forall \bar{p} \in \mathcal{P}$), and vectors $g_x \in \mathbb{R}_{\geq 0}^{N_x}$, $g_u^p \in \mathbb{R}_{\geq 0}^{N_u^p}$,

and $g_G \in \mathbb{R}_{\geq 0}^{N_{u_G}}$. The map $U^p(u^{-p})$ is then equivalent to the actions $u^p = \Phi_u^p * w$ whose associated response Φ_u^p satisfy

$$[G_x]_i (\Phi_x * w)_n \leq [g_x]_i, \quad i = 1, \dots, N_x; \quad (35)$$

$$[G_u^p]_j (\Phi_u^p * w)_n \leq [g_u^p]_j, \quad j = 1, \dots, N_{u^p}; \quad (36)$$

$$\sum_{\bar{p} \in \mathcal{P}} [G_G^{\bar{p}}]_l (\Phi_u^{\bar{p}} * w)_n \leq [g_G]_l, \quad l = 1, \dots, N_{u_G}, \quad (37)$$

with $([G_x]_i, [g_x]_i)$, $([G_u^p]_j, [g_u^p]_j)$, and $([G_G^{\bar{p}}]_l, [g_G]_l)$ denoting, respectively, the i -th, j -th, and l -th rows of the corresponding matrices and vectors. In this work, players are assumed to synthesise policies that satisfy these constraints for any $w \in \mathcal{W}$. At the expense of conservativeness, we cast Problem (33) as a robust optimization problem by considering the worst-case realisation of the noise. Specifically, we reformulate Eq. (35),

$$\sum_{n'=0}^N \sup_{w \in \mathcal{W}} [G_x]_i \Phi_{x,n'} w \leq [g_x]_i, \quad i = 1, \dots, N_x \quad (38)$$

and similarly for Eqs. (36)–(37), then exploit our knowledge of \mathcal{W} to obtain an analytical solution for the supremum.

A common instance of $\mathcal{G}_\infty^{\text{LQ}}$ considers the problem in which $(w_n)_{n \in \mathbb{N}}$ is uniformly bounded. In these problems, the inequality Eq. (38) can be expressed in terms of a dual norm on the appropriate vector space. We highlight two important cases.

- \mathcal{W} is an ellipsoid centred at zero. Formally, the set is

$$\mathcal{W} = \{w_t \in \mathbb{R}^{N_x} : \|P w_t\|_2 \leq 1\},$$

given a matrix $P \in \mathbb{S}_{++}^{N_x}$. This corresponds to noise processes with uniformly bounded energy, as encoded by $\|\cdot\|_2$. For such cases, Eqs. (35)–(37) can be enforced by the second-order conic (SOC) constraints

$$\begin{aligned} \sum_{n=0}^N \left\| ([G_x]_i \Phi_{x,n} P^{-1})^\top \right\|_2 &\leq [g_x]_i, \quad (\forall i); \\ \sum_{n=0}^{N-1} \left\| ([G_u^p]_j \Phi_{u,n}^p P^{-1})^\top \right\|_2 &\leq [g_u^p]_j, \quad (\forall j); \\ \sum_{n=0}^{N-1} \left\| \sum_{\bar{p} \in \mathcal{P}} ([G_G^{\bar{p}}]_l \Phi_{u,n}^{\bar{p}} P^{-1})^\top \right\|_2 &\leq [g_G]_l, \quad (\forall l). \end{aligned}$$

- \mathcal{W} is a polyhedron, symmetric around zero. Formally,

$$\mathcal{W} = \{w_t \in \mathbb{R}^{N_x} : \|P w_t\|_\infty \leq 1\},$$

given a full-rank matrix $P \in \mathbb{R}^{N_{\mathcal{W}} \times N_x}$. This corresponds to noise processes with uniformly bounded intensity, as encoded by $\|\cdot\|_\infty$. For such cases, Eqs. (35)–(37) can be enforced by the first-order conic constraints

$$\begin{aligned} \sum_{n=0}^N \left\| ([G_x]_i \Phi_{x,n} P^\dagger)^\top \right\|_1 &\leq [g_x]_i, \quad (\forall i); \\ \sum_{n=0}^{N-1} \left\| ([G_u^p]_j \Phi_{u,n}^p P^\dagger)^\top \right\|_1 &\leq [g_u^p]_j, \quad (\forall j); \\ \sum_{n=0}^{N-1} \left\| \sum_{\bar{p} \in \mathcal{P}} ([G_G^{\bar{p}}]_l \Phi_{u,n}^{\bar{p}} P^\dagger)^\top \right\|_1 &\leq [g_G]_l, \quad (\forall l). \end{aligned}$$

Remark 2. If $\mathcal{X} = \mathbb{R}^{N_x}$, $\mathcal{U}^p = \mathbb{R}^{N_u^p}$, and $\mathcal{U}_G = \prod_{p \in \mathcal{P}} \mathbb{R}^{N_u^p}$, these constraints are trivially satisfied for any $w \in \mathcal{W}$, and thus can be removed from Problem (31). Conversely, if $\mathcal{W} = \mathbb{R}^{N_x}$ when either \mathcal{X} , \mathcal{U}^p , or \mathcal{U}_G is bounded, then no stabilising policy can enforce those constraints for all possible w .

Structural constraints: The sets $(\mathcal{C}_{x,n}, \mathcal{C}_{u,n}^p)_{n \in \mathbb{N}^+}$ ($\forall p \in \mathcal{P}$) are designed to impose some structure directly on the policy \mathbf{K}^p (Corollary 4.1), often in the form of sparsity constraints. An important class of such structural constraints encodes information patterns incurred by the presence of actuation and sensing delays: Let $\mathbf{K}^p \in \mathcal{C}^p$ ($\forall p$) satisfy the restrictions

$$\mathcal{C}^p = \{ \mathbf{K}^p \in \mathcal{L}(\mathcal{X}, \mathcal{U}^p) : \text{Actions } [B^p(K^p x)_t]_i \text{ respectively} \\ \text{affect and feedback the component} \\ [x_t]_i \text{ with the delays } d_a, d_s > 0 \}, \quad (39)$$

and consider the operators $\mathbf{S}_x : \Phi_x \mapsto (S_{x,n} \odot \Phi_{x,n})_{n \in \mathbb{N}^+}$ and $\mathbf{S}_{u^p}^p : \Phi_{u^p}^p \mapsto (S_{u,n}^p \odot \Phi_{u,n}^p)_{n \in \mathbb{N}^+}$ ($\forall p$), given the matrices

$$(S_{x,n})_{n \in \mathbb{N}^+} = \left(\text{Sp}(A^{\max(0, \lfloor \frac{n-d_a}{d_s} \rfloor)}) \right)_{n \in \mathbb{N}^+}, \\ (S_{u,n}^p)_{n \in \mathbb{N}^+} = \left(\text{Sp}(B^{p\top} A^{\max(0, \lfloor \frac{n-d_a}{d_s} \rfloor)}) \right)_{n \in \mathbb{N}^+},$$

with $\text{Sp}(\cdot)$ denoting the sparsity pattern of a matrix¹. It can be shown that $\mathbf{K}^p = \Phi_{u^p}^p \Phi_x^{-1} \in \mathcal{C}^p$ ($\forall p$) if its parametrisation satisfy $(\Phi_{x,n} \in \mathcal{C}_{x,n})_{n \in \mathbb{N}^+}$ and $(\Phi_{u,n}^p \in \mathcal{C}_{u,n}^p)_{n \in \mathbb{N}^+}$ with

$$\mathcal{C}_{x,n} = \{ \Phi_{x,n} \in \mathbb{R}^{N_x \times N_x} : \Phi_{x,n} = S_{x,n} \odot \Phi_{x,n} \}; \quad (40a)$$

$$\mathcal{C}_{u,n}^p = \{ \Phi_{u,n}^p \in \mathbb{R}^{N_u^p \times N_x} : \Phi_{u,n}^p = S_{u,n}^p \odot \Phi_{u,n}^p \}. \quad (40b)$$

The constraints Eq. (40) enforce that the closed-loop response to the noise obeys an information pattern induced by the dynamics of the game $\mathcal{G}_{\infty}^{\text{LQ}}$. Specifically, $[S_{x,n}]_{i,j} = 0$ (respectively, $[S_{u,n}^p]_{i,j} = 0$) indicates that any disturbance to the j -th component of the state, $[x_t]_j$, does not affect $[x_{t+n}]_i$ (does not feedback into action $[u_{t+n}^p]_i$). Finally, note that $(\mathcal{C}_{x,n}, \mathcal{C}_{u,n}^p)_{n \in \mathbb{N}^+}$ are compact convex sets, as they correspond to the kernels of the linear operators $(I - \mathbf{S}_x)$ and $(I - \mathbf{S}_{u^p}^p)$.

The ability to impose a desired policy structure using convex constraints is a central feature of the SLS framework. In the context of dynamic games, it allows for describing and, most importantly, solving problems where players have asymmetric information patterns; a major challenge for feedback Nash equilibrium problems [41, 43]. For $(\mathcal{C}_{x,n}, \mathcal{C}_{u,n}^p)_{n \in \mathbb{N}^+}$ from Eq. (40), conditions for the existence of GFNE (i.e., $\Omega_{\mathcal{G}_{\infty}^{\text{LQ}}}^{\mathbf{K}} \neq \emptyset$), can be stated in terms of the delay parameters d_a and d_s . Unless stated otherwise, the mappings $\{BR_{\Phi}^p\}_{p \in \mathcal{P}}$ (and $\{\bar{B}R_{\Phi}^p\}_{p \in \mathcal{P}}$) are assumed to include these constraints with $d_a = d_s = 1$.

B. System-level best-response dynamics

A learning dynamics based on the system-level best-responses $\{BR_{\Phi}^p\}_{p \in \mathcal{P}}$ relies on the following central result.

Theorem 6. A policy profile $\mathbf{K}^* = (\Phi_{u^1}^{1*}, \dots, \Phi_{u^{N_P}}^{N_P*}) \Phi_x^{-1} \in \mathcal{C}$, is a GFNE of $\mathcal{G}_{\infty}^{\text{LQ}}$ if $\Phi_{u^p}^p \in BR_{\Phi}^p(\Phi_{u^*}^p)$, or, equivalently,

$$\Phi_{u^p}^{p*} \in BR_{\Phi}^p(\Phi_{u^*}^{-p*}), \quad \forall p \in \mathcal{P}. \quad (41)$$

Proof. Consider an arbitrary fixed-point $\Phi_{u^*}^p \in BR_{\Phi}^p(\Phi_{u^*}^p)$. From Theorem 4, we have $\Phi_x^* = \mathbf{F}_{\Phi} \Phi_{u^*}^p$. Now, consider

¹The operator $\text{Sp}(A)$ produces a binary matrix such that $[\text{Sp}(A)]_{i,j} = 1$ if $[A]_{i,j} \neq 0$ and $[\text{Sp}(A)]_{i,j} = 0$ otherwise.

policies $\mathbf{K}^{p*} = \Phi_{u^p}^{p*} (\Phi_x^*)^{-1}$, $p \in \mathcal{P}$. Clearly, $\Phi_{u^p}^{p*} = \mathbf{K}^{p*} \Phi_x^*$. As a consequence, for any $w \in \mathcal{W}$,

$$\Phi_{u^p}^{p*} w = \mathbf{K}^{p*} \Phi_x^* w \iff u^{p*} = \mathbf{K}^{p*} x^*.$$

and, by definition, $\Phi_{u^p}^{p*} \in U_{\Phi}^p(\Phi_{u^*}^{-p*})$ imply $u^{p*} \in U^p(u^{-p*})$. Thus, u^* is the open-loop realisation of the policy \mathbf{K}^* . Finally, since $J^p(u^{p*}, u^{-p*}) \cong J^p(\Phi_{u^p}^{p*}, \Phi_{u^*}^{-p*})$ and $\Phi_{u^p}^{p*} \in \Omega_{\mathcal{G}_{\infty}^{\Phi}}$, we conclude that no player can obtain an admissible policy that unilaterally improves its cost, that is, $\mathbf{K}^* \in \Omega_{\mathcal{G}_{\infty}^{\text{LQ}}}^{\mathbf{K}}$. \square

The relationship between BR and BR_{Φ} implies that a GFNE of $\mathcal{G}_{\infty}^{\text{LQ}}$ can be obtained analytically from a GNE of $\mathcal{G}_{\infty}^{\Phi}$. This fact allows us to adapt the BRD-GFNE procedure (Algorithm 3) and propose a procedure for GFNE seeking in constrained infinite-horizon dynamic games based on the mappings $\{BR_{\Phi}^p\}_{p \in \mathcal{P}}$. This system-level best-response dynamics (SLS-BRD) approach is given in Algorithm 4. We remark on some technical aspects:

- The pair $(\Phi_{x,k}^p, \Phi_{u,k}^p)$ defines the p -th player's parametrisation after $k \in \mathbb{N}$ updates, that is, $\Phi_{x,k}^p = (\Phi_{x,k,t}^p)_{t \in \mathbb{N}}$ and $\Phi_{u,k}^p = (\Phi_{u,k,t}^p)_{t \in \mathbb{N}}$. The update index $k \in \mathbb{N}$ should not be mistaken for the stage index $t \in \mathbb{N}$.
- Updating the policy \mathbf{K}_{k+1}^p consists of either employing Corollary 4.1 with maps $(\Phi_{x,k+1}^p, \Phi_{u,k+1}^p)$, or directly computing its kernel $\Phi_{k+1}^p = ([M_{\Phi_{u,k+1}^p}^p]_n \Phi_{x,k+1}^{-1})_{n \in \mathbb{N}}$.
- Responses $\{\Phi_{x,k}^p\}_{p \in \mathcal{P}}$ are most likely distinct at $k < \infty$, that is, $\Phi_{x,k}^p \neq \Phi_{x,k}^{\tilde{p}}$ for $p \neq \tilde{p}$. Consequently, the system level parametrisation (Eq. 24) might not hold for the profile $\mathbf{K}_k = (\Phi_{u,k}^1 (\Phi_{x,k}^1)^{-1}, \dots, \Phi_{u,k}^{N_P} (\Phi_{x,k}^{N_P})^{-1})$ and this could lead to stability issues. However, this policy still satisfies a robust variant of Theorem 4 when the distances $\|\Phi_{x,k}^p - \mathbf{F}_{\Phi} \Phi_{u,k}^p\|$ ($\forall p \in \mathcal{P}$) are sufficiently small [36].

The SLS-BRD induces an operator $T_{\Phi} = (1 - \eta)I + \eta BR_{\Phi}$, which defines the global update $\Phi_{u,k+1}$ from $\Phi_{u,k}$. As before, T_{Φ} and BR_{Φ} share the same fixed-points: The GNEs $\Omega_{\mathcal{G}_{\infty}^{\Phi}}$. From Theorem 6, convergence to a response $\Phi_{u^*}^p \in T_{\Phi}(\Phi_{u^*}^p)$ then implies convergence to a policy $\mathbf{K}^* \in \Omega_{\mathcal{G}_{\infty}^{\text{LQ}}}^{\mathbf{K}}$. Hence, this learning dynamics is a formal procedure for GFNE seeking.

Algorithm 4: System-level BRD (SLS-BRD)

Input: Game $\mathcal{G}_{\infty}^{\text{LQ}} := (\mathcal{P}, \mathcal{X}, \{\mathcal{U}^p\}_{p \in \mathcal{P}}, \mathcal{W}, \{J^p\}_{p \in \mathcal{P}})$
Output: GFNE $\mathbf{K}^* = (\mathbf{K}^{1*}, \dots, \mathbf{K}^{N_P*})$

- 1 Initialize $\mathbf{K}_0 = (\mathbf{K}_0^1, \dots, \mathbf{K}_0^{N_P})$ and $k = 0$;
 - 2 **for** $t = 0, 1, 2, \dots$ **do**
 - 3 /* Players apply actions $\{u_{k,t}^p = K_k^p x_{k,t}\}_{p \in \mathcal{P}} \quad */$
 - 4 **if** $\Phi_{u,k} \in BR_{\Phi}(\Phi_{u,k})$ **then return** \mathbf{K}_k ;
 - 5 **if** $t = (k+1)\Delta T$ **then**
 - 6 **for** $p \in \mathcal{P}$ **do**
 - Update \mathbf{K}_{k+1}^p by computing the kernels
 - $\Phi_{u,k+1}^p := (1-\eta)\Phi_{u,k}^p + \eta BR_{\Phi}^p(\Phi_{u,k}^{-p})$,
 - $\Phi_{x,k+1}^p := \mathbf{F}_{\Phi}(\Phi_{u,k+1}^p, \Phi_{u,k}^{-p})$
 - 7 $k := k + 1$;
-

In this general form, Algorithm 4 is still unpractical due to $\{BR_{\Phi}^p\}_{p \in \mathcal{P}}$ being intractable. The SLS-BRD can be adapted to consider instead $\{\widehat{BR}_{\Phi}^p\}_{p \in \mathcal{P}}$: The players' updates become

$$\Phi_{u,k+1}^p := (1-\eta)\Phi_{u,k}^p + \eta\widehat{BR}_{\Phi}^p(\Phi_{u,k}^{-p}).$$

The global update rule induced by this modified algorithm is $\widehat{T}_{\Phi} = (1-\eta)I + \eta\widehat{BR}_{\Phi}$. In this case, the fixed-points of \widehat{T}_{Φ} coincide with those of \widehat{BR}_{Φ} . As such, this (approximately)best-response dynamics is a procedure for ϵ -GFNE seeking, $\epsilon > 0$.

Convergence of the SLS-BRD: The convergence of Algorithm 4 depends on BR_{Φ} (or \widehat{BR}_{Φ} , for its tractable version) being at least non-expansive (Theorems 2–3). Formally,

Theorem 7. *Let the map $\widehat{BR}_{\Phi} : \mathcal{C}_u \rightrightarrows \mathcal{C}_u$ be $L_{\widehat{BR}_{\Phi}}$ -Lipschitz, with $L_{\widehat{BR}_{\Phi}} < 1$. Then, the SLS-BRD $\Phi_{u,k+1} = \widehat{T}_{\Phi}(\Phi_{u,k})$ converge to the unique ϵ -GNE $\Phi_u^* \in \Omega_{\mathcal{G}_{\Phi}}^{\epsilon}$ with rate*

$$\frac{\|\Phi_{u,k} - \Phi_u^*\|_{\ell_2}}{\|\Phi_{u,0} - \Phi_u^*\|_{\ell_2}} \leq ((1-\eta) - \eta L_{\widehat{BR}_{\Phi}})^k \quad (42)$$

from any feasible initial $\Phi_{u,0}$.

In general, determining a Lipschitz constant for such mappings is challenging. However, for linear-quadratic games $\mathcal{G}_{\infty}^{LQ}$ where \mathcal{W} is a polyhedron, best-responses are piecewise-affine operators and their Lipschitz properties are straightforward. In the following, we use this fact to establish some conditions for convergence of the SLS-BRD for a specific class of games.

Consider the (approximately)best-response maps $\{\widehat{BR}_{\Phi}^p\}_{p \in \mathcal{P}}$ and assume N sufficiently large to ensure that $\|\Phi_{x,N}\|_F^2 < \gamma$ is strictly satisfied. For notational convenience, let us introduce the operators $F_{\Phi}^p = (I - S_+ A)^{-1} S_+ B^p$ ($\forall p \in \mathcal{P}$) and signal $F_{\Phi}^0 = (I - S_+ A)^{-1} \delta I_{N_x}$, and the objective-related mappings $C^p : \Phi_x \mapsto (C^p \Phi_{x,n})_{n \in \mathbb{N}}$ and $D^{p\tilde{p}} : \Phi_u^{\tilde{p}} \mapsto (D^{p\tilde{p}} \Phi_{u,n}^{\tilde{p}})_{n \in \mathbb{N}}$, with $D^{p0} = 0$. Moreover, for all $p, \tilde{p} \in \mathcal{P}$, define the operators

$$H^{p\tilde{p}} = (C^p F_{\Phi}^p + D^{pp})^T (C^{\tilde{p}} F_{\Phi}^{\tilde{p}} + D^{p\tilde{p}}), \quad (43)$$

and $H^{p,-p} = (H^{p,\tilde{p}})_{\tilde{p} \in \mathcal{P}}$. Because $\{\Phi_u^p\}_{p \in \mathcal{P}}$ are FIR, all the elements defined above have equivalent matrix representations. Using this notation, Problem (33) can be reformulated as

$$\begin{aligned} \text{minimize} \quad & \text{Tr}[\Phi_u^T H^{pp} \Phi_u^p \\ & + 2(\sum_{\tilde{p} \in \mathcal{P} \setminus \{p\}} H^{p\tilde{p}} \Phi_u^{\tilde{p}} + H^{p0})^T \Phi_u^p] \end{aligned} \quad (44)$$

The maps $BR^p(\Phi_u^{-p})$, $p \in \mathcal{P}$, thus correspond to the solution of quadratic programs with convex constraints $U_{\Phi}^p(\Phi_u^{-p})$, for $\Phi_u^{-p} \in \mathcal{C}_u^{-p}$. In the case of $\mathcal{W} = \{w_t \in \mathbb{R}^{N_x} : \|w_t\|_{\infty} \leq 1\}$, and without structural constraints, that is, $S_x = I$ and $S_u^p = I$,

$$\begin{aligned} U_{\Phi}^p(\Phi_u^{-p}) &= \{\Phi_u^p \in \ell_2[0, N] : \\ & \Phi_{x,n+1} = A\Phi_{x,n} + \sum_{\tilde{p} \in \mathcal{P}} B^{\tilde{p}} \Phi_{u,n}^{\tilde{p}}, \quad \Phi_{x,1} = I_{N_x}, \\ & \sum_{n=0}^N \|([G_x]_i \Phi_{x,n})^T\|_1 \leq [g_x]_i, \\ & \sum_{n=0}^{N-1} \|([G_u^p]_j \Phi_{u,n}^p)^T\|_1 \leq [g_u^p]_j, \\ & \sum_{n=0}^{N-1} \|\sum_{\tilde{p} \in \mathcal{P}} ([G_{\tilde{p}}^p]_l \Phi_{u,n}^{\tilde{p}})^T\|_1 \leq [g_{\tilde{p}}^p]_l\}. \end{aligned} \quad (45)$$

Using standard techniques from optimisation, the constraints Eq. (45) can be incorporated into Problem (44) as linear inequalities. The solutions for these problems, that is, the best-responses $\{\widehat{BR}_{\Phi}^p\}_{p \in \mathcal{P}}$, are thus piecewise affine in $\Phi_u^{-p} \in \mathcal{C}_u^{-p}$ [44]. Consequently, also \widehat{BR}_{Φ} must be piecewise affine. A local Lipschitz constant can then be derived for each region of \mathcal{C}_u^{-p} that leads to a subset of the operational constraints being active. For non-generalised games (i.e., when $\mathcal{X} = \mathbb{R}^{N_x}$ and $\mathcal{U}_G = \prod_{p \in \mathcal{P}} \mathbb{R}^{N_u^p}$), the structure of this affine mapping can be exploited to derive a global Lipschitz constant for \widehat{BR}_{Φ} .

Theorem 8. *Consider $\mathcal{X} = \mathbb{R}^{N_x}$ and $\mathcal{U}_G = \prod_{p \in \mathcal{P}} \mathbb{R}^{N_u^p}$. Then, The map \widehat{BR}_{Φ} is $L_{\widehat{BR}_{\Phi}}$ -Lipschitz with*

$$L_{\widehat{BR}_{\Phi}} = \sqrt{\sum_{p \in \mathcal{P}} (L_{\widehat{BR}_{\Phi}^p}^p)^2}, \quad (46)$$

given the player-specific constants

$$L_{\widehat{BR}_{\Phi}^p}^p = (1 + \kappa(H^{pp})) \| (H^{pp})^{\dagger} H^{p,-p} \|_{2 \rightarrow 2}, \quad (47)$$

with condition number $\kappa(H^{pp}) = \| (H^{pp})^{\dagger} \|_{2 \rightarrow 2} \| H^{pp} \|_{2 \rightarrow 2}$.

The proof of Theorem 8 is extensive: The reader is referred to Appendix C for the full details. Importantly, this Lipschitz constant is not tight and thus $L_{\widehat{BR}_{\Phi}} < 1$ using Eq. (46) is only a sufficient (but not necessary) condition for Theorem 7 to hold in practice. Regardless, it highlights some intuitive, but non-trivial, facts about the convergence of the SLS-BRD:

- The condition $L_{\widehat{BR}_{\Phi}} < 1$ implies that the block-operator $H = [H^{p\tilde{p}}]_{p, \tilde{p} \in \mathcal{P}}$ is diagonally dominant. This highlights the relationship between the SLS-BRD and the class of Jacobi iterative methods, where diagonal dominance of the linear system being solved is a strict requirement.
- The convergence rate of the SLS-BRD is inversely proportional to $\|D^{pp}\|_2^2$ ($\forall p$) through H^{pp} . As such, faster convergence is expected for games in which players apply strong penalties to their own actions.
- The convergence rate of the SLS-BRD depends on the number of players: Let $L_{\widehat{BR}_{\Phi}^p}^p = \sqrt{\alpha + \beta^p}$ for all $p \in \mathcal{P}$ given some $\alpha, \beta^p > 0$. Then, $L_{\widehat{BR}_{\Phi}}^2 = (\alpha N_p)^2 \sum_p \beta^p$. In large-scale games, players might need to become more conservative in order to ensure that \widehat{BR} is a contraction.
- The convergence rate of the SLS-BRD is dominated by the slowest player: Whenever there exists a $p \in \mathcal{P}$ such that $L_{\widehat{BR}_{\Phi}^p}^p \gg L_{\widehat{BR}_{\Phi}^{\tilde{p}}}^{\tilde{p}}$ for all $\tilde{p} \in \mathcal{P}$, then $L_{\widehat{BR}_{\Phi}} \approx L_{\widehat{BR}_{\Phi}^p}^p$.
- The norms $\|H^{pp}\|_{2 \rightarrow 2} = \|C^p F_{\Phi}^p + D^{pp}\|_{2 \rightarrow 2}$ increase with the FIR horizon N , specially for unstable games. In some cases, players might need to proportionally decrease the state-penalties C^p to ensure (or improve) convergence.

At the cost of interpretability, similar Lipschitz constants as in Eqs. (46)–(47) can be obtained also when $(\mathcal{X}, \mathcal{U}_G, \mathcal{U}^p)$ and P are general, and when structural constraints are present. Moreover, we remark that Theorem 8 is obtained under the assumption that $\|\Phi_{x,N}\|_F^2 < \gamma$ holds strictly. The best-response maps $\{\widehat{BR}_{\Phi}^p\}_{p \in \mathcal{P}}$ when this constraint is active, or when \mathcal{W} is an ellipsoid, become quadratically-constrained quadratic programs and their Lipschitz properties are less intuitive.

IV. EXAMPLES

In the following, we demonstrate the SLS-BRD algorithm in two exemplary problems: Decentralised control of an unstable network and management of a competitive market. In both problems, we set the initial profile $\mathbf{K}_0 = (\Phi_{u,0}^1, \dots, \Phi_{u,0}^{N_p}) \Phi_{x,0}^{-1}$ by projecting the zero-response $\hat{\Phi}_{u,0} = \mathbf{0}$ into the feasible set.² A SLS-BRD routine is then simulated by having players update their policies through best-responses to the parametrisation of their rivals' policies, assumed to be available. Due to numerical limitations, we interrupt the updates whenever the condition $\|\Phi_{u,k}^p - \Phi_{u,k-1}^p\|_{\ell_2} / \|\Phi_{u,k}^p\|_{\ell_2} \leq 10^{-8}$ ($\forall p \in \mathcal{P}$) is satisfied.

A. Stabilisation of a bidirectional chain network

Consider a game $\mathcal{G}_{\infty}^{\text{LQ}} = \{\mathcal{P}, \mathcal{X}, \{\mathbf{U}^p\}_{p \in \mathcal{P}}, \mathcal{W}, \{J^p\}_{p \in \mathcal{P}}\}$ with players $\mathcal{P} = \{1, 2, 3\}$ operating a chain network of $N_x = 14$ single-state nodes whose dynamics are described by

$$A = \begin{bmatrix} 1 & 0.2 & & & \\ -0.2 & \ddots & \ddots & & \\ & \ddots & \ddots & \ddots & \\ & & & & 0.2 \\ & & & -0.2 & 1 \end{bmatrix}; \left\{ B^p = \begin{bmatrix} \mathbf{0}_{6(p-1) \times 2} \\ I_2 \\ \mathbf{0}_{6(3-p) \times 2} \end{bmatrix} \right\}_{p \in \mathcal{P}}, \quad (48)$$

where $A \in \mathbb{R}^{N_x \times N_x}$ and $B^p \in \mathbb{R}^{N_x \times N_u^p}$ with $N_u^p = 2$ ($\forall p$). This game has unstable dynamics, since $\rho(A) = 1.073 > 1$, however it is stabilisable for each (A, B^p) . Moreover, the game is subjected to a noise process described by $w_t \sim \text{Uniform}(\mathcal{W})$, $t \in \mathbb{N}$, defined over $\mathcal{W} = \{w_t \in \mathbb{R}^{N_x} : \|w_t\|_{\infty} \leq 1\}$. In this problem, players are interested in stabilising the game $\mathcal{G}_{\infty}^{\text{LQ}}$, while minimising their individual objective functionals,

$$J^p(u^p, u^{-p}) = \mathbb{E} \left[\sum_{t=0}^{\infty} \left(\|x_t\|_2^2 + \beta^p \|u_t^p\|_2^2 \right) \right],$$

equivalent to Eq. (21) after setting $C^p = [I_{N_x} \ \mathbf{0}_{N_x \times N_u}]^T$, $D^{pp} = [\mathbf{0}_{N_u \times N_x} \ \sqrt{\beta^p} I_{N_u}]^T$, and $D^{p\tilde{p}} = 0$ for all $\tilde{p} \in \mathcal{P} \setminus \{p\}$. The players' actions are subjected to operational constraints, $\mathbf{u}^p \in U^p(\mathbf{u}^{-p})$, defined by the constraint sets

$$\begin{aligned} \mathcal{X} &= \mathbb{R}^{N_x}; \\ U^p &= \{u_t^p \in \mathbb{R}^{N_u^p} : \frac{1}{10} [\mathbf{1}_{N_u^p} \ -\mathbf{1}_{N_u^p}]^T u_t^p \leq [1 \ 1]^T\}; \\ \mathcal{U}_G &= \prod_{p \in \mathcal{P}} \mathbb{R}^{N_u^p}, \end{aligned}$$

enforcing $-10 \leq [u_t^p]_1 + [u_t^p]_2 \leq 10$ for every $t \in \mathbb{N}$ and $p \in \mathcal{P}$. We assume that players design their state-feedback policies, $\mathbf{K}^p = \Phi_{u,0}^p \Phi_{x,0}^{-1} \in \mathcal{C}^p$, considering a FIR horizon of $N = 50$ and the constraints $(\Phi_{x,n} \in \mathcal{C}_{x,n})_{n=1}^N$ and $(\Phi_{u,n}^p \in \mathcal{C}_{u,n}^p)_{n=1}^N$,

$$\begin{aligned} \mathcal{C}_{x,n} &= \{\Phi_{x,n} \in \mathbb{R}^{N_x \times N_x} : \Phi_{x,n} = \text{Sp}(A^{n-1}) \odot \Phi_{x,n}\}; \\ \mathcal{C}_{u,n}^p &= \{\Phi_{u,n}^p \in \mathbb{R}^{N_u^p \times N_x} : \Phi_{u,n}^p = \text{Sp}(B^{pT} A^{n-1}) \odot \Phi_{u,n}^p\}. \end{aligned}$$

Under this setup, $\mathcal{G}_{\infty}^{\text{LQ}}$ belongs to the class of dynamic potential games (DPG, [45]) and the (unique) GFNE can be obtained in advance by solving a centralised optimisation problem.

²That is, $\Phi_{u,0} = \arg \min_{\Phi_{u,0} \in \mathcal{U}_{\Phi}} \|\Phi_{u,0} - \hat{\Phi}_{u,0}\|_{\ell_2}^2$ given the feasible set $\mathcal{U}_{\Phi} = \{\Phi_{u,0} \in \ell_2[0, N) : \Phi_{u,0}^p \in U_{\Phi}^p(\Phi_{u,0}^{-p}), p \in \mathcal{P}\}$.

In this experiment, we simulate a GFNE seeking procedure for each different value $\beta \in \{(10, 40, 10), (2, 8, 2), (0.4, 1.6, 0.4)\}$. The game $\mathcal{G}_{\infty}^{\text{LQ}}$ is executed alongside the updating of players' policies according to some learning dynamics. The players seek ϵ -GFNE policies by adhering to the SLS-BRD routine (Algorithm 4) using their (approximately)best-response maps, $\{\widehat{BR}_{\Phi}^p\}_{p \in \mathcal{P}}$, with $\gamma = 0.95$. The policies are updated simultaneously every $\Delta T = 1$ stage with a learning rate of $\eta = 1/2$.

The convergence of the SLS-BRD routine to the fixed-point $\mathbf{K}^* = \Phi_{u,0}^* \Phi_{x,0}^{-1} = (\Phi_{u,0}^{1*}, \dots, \Phi_{u,0}^{N_p*}) \Phi_{x,0}^{-1}$ is shown in Figure 3. The results demonstrate that the players' policies are sufficiently close to the ϵ -GFNE profile after 260, 420, and 510 iterations for each respective weighting configuration. In each case, the (soft) FIR constraints are satisfied strictly after the initial profile (that is, $\|\Phi_{x,k,N}^p\|_F^2 < 0.95$, $k > 2$, $p \in \mathcal{P}$). The Lipschitz properties of the global update rule T_{Φ} seem to follow the intuition provided in Section III-B: The convergence rate improves when players apply stronger penalties to their actions ($\beta = (10, 40, 10)$) when compared to that obtained by weaker control penalties ($\beta = (0.4, 1.6, 0.4)$). However, we stress that the Lipschitz constants from Theorem 8 do not consider structural constraints (\mathcal{S}_x and \mathcal{S}_u^p , $\forall p \in \mathcal{P}$) and cannot be applied to this experiment. Regardless, the convergence rates are shown to be geometric, indicating that the best-response maps \widehat{BR}_{Φ} in these scenarios are contractive. If not interrupted, and disregarding numerical limitations, the SLS-BRD should continue to approach the fixed-point \mathbf{K}^* at this rate.

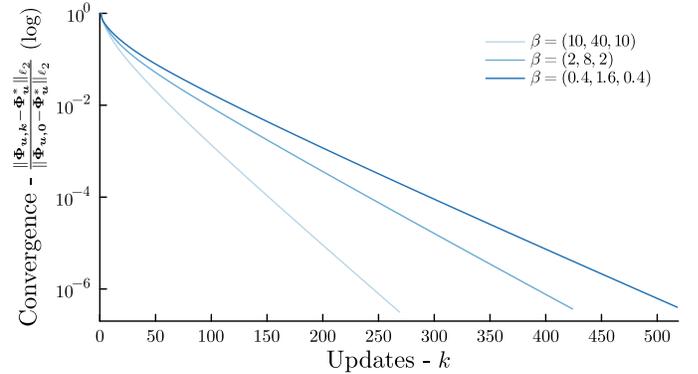


Figure 3. N_p -Chain game: Convergence of the SLS-BRD routine.

In Figure 4 (top), we show the relative distance between the individual updates $\{\Phi_{u,k}^p, \Phi_{u,k-1}^p\}_{p \in \mathcal{P}}$ from each player. The updates are of similar magnitude for all players $p \in \mathcal{P}$, in each scenario, except for player $p = 3$ which shows a slightly faster convergence. In general, these local changes become numerically negligible at a faster rate than the global convergence in Figure 3. Specifically, when the relative distance between updates approaches the aforementioned threshold of 10^{-8} , the corresponding policy profile has become closer to the ϵ -GFNE by a factor of 10^{-6} . Finally, we consider the relative distances $\|\Phi_{x,k}^p - \Phi_{x,k-1}^p\|_{\ell_2} / \|\Phi_{x,k}^p\|_{\ell_2}$, $k \in \mathbb{N}_+$, between the responses $\{\Phi_{x,k}^p\}_{p \in \mathcal{P}}$ and $\Phi_{x,k} = \mathbf{F}_{\Phi} \Phi_{u,k}$ obtained from the system-level parametrisation associated with $\Phi_{u,k}$ (Theorem 4). As shown in Figure 4 (bottom), these distances decrease at a similar rate and are relatively small since the initial stages of the game.

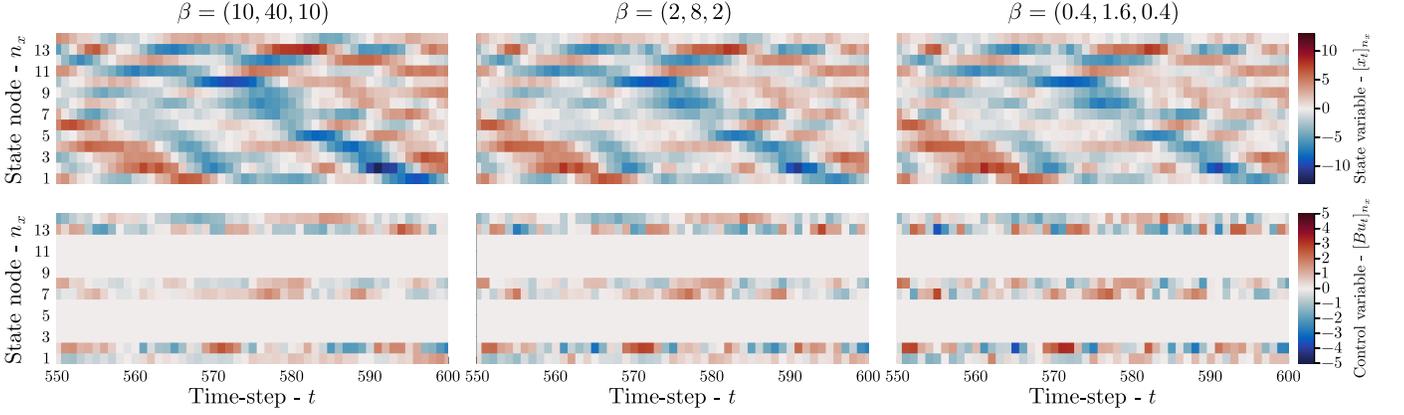


Figure 5. N_P -Chain game, $t \in (550, 600]$: State \mathbf{x} (top panels) and applied control $\mathbf{B}\mathbf{u}$ (bottom panels) trajectories for each execution of $\mathcal{G}_\infty^{\text{LQ}}$ with $\beta \in \{(10, 40, 10), (2, 8, 2), (0.4, 1.6, 0.4)\}$. The vertical axis represents each node in the chain-networked system.

The policies $\mathbf{K}_k = (\Phi_{\mathbf{u},k}^1 (\Phi_{\mathbf{x},k}^1)^{-1}, \dots, \Phi_{\mathbf{u},k}^{N_P} (\Phi_{\mathbf{x},k}^{N_P})^{-1})$ thus approximate $(\Phi_{\mathbf{u},k}^1, \dots, \Phi_{\mathbf{u},k}^{N_P}) \Phi_{\mathbf{x},k}^{-1}$ as $k \rightarrow \infty$ and are expected to stabilise the system during this learning dynamics.

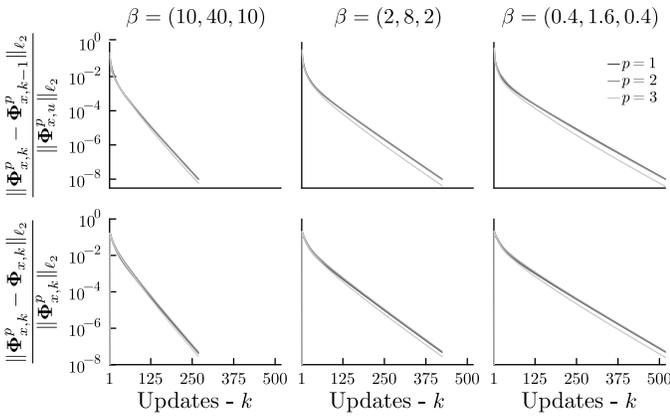


Figure 4. N_P -Chain game: Relative distance (in log-scale) between updates $(\Phi_{\mathbf{u},k}^p, \Phi_{\mathbf{u},k+1}^p)$, top panels, and responses $(\Phi_{\mathbf{x},k}^p, \Phi_{\mathbf{x},k+1}^p)$, bottom panels.

The evolution of the game given the actions by each player is displayed in Figure 5 for the period $t \in (550, 600]$, when the policy updates have been interrupted. The results demonstrate that the policies obtained by the SLS-BRD routine are capable of jointly stabilising the networked system, robustly against the random noise. These policies are also shown to satisfy the operational constraints: Each player's actions become (roughly) symmetrical and satisfy $|[u_t]_1| + |[u_t]_2| \leq 5$, $t \in \mathbb{N}$, in all scenarios. The enforcement of this strategy highlights the conservativeness resulting from the robust operational constraints. Finally, note that the state- and control-trajectories from operating the system through these policies are similar for all β configurations; however, the policies with $\beta = (2, 8, 2)$ and $\beta = (0.4, 1.6, 0.4)$ achieve a slightly better noise rejection at the expense of more aggressive control actions. In this experiment, the choice $\beta = (10, 40, 10)$ seems to be preferable as it converges faster while still achieving satisfactory performance.

B. Price competition in oligopolistic markets

Consider a game $\mathcal{G}_\infty^{\text{LQ}} = \{\mathcal{P}, \mathcal{X}, \{\mathbf{U}^p\}_{p \in \mathcal{P}}, \mathcal{W}, \{J^p\}_{p \in \mathcal{P}}\}$ consisting of a set of companies $\mathcal{P} = \{1, 2, 3, 4\}$ participating in a single-product market. Assume that these companies have equivalent production capacities and are able to satisfy the demand for their products. The product offered by company $p \in \mathcal{P}$ has a daily local demand $\mathbf{d}^p = (d^p(t))_{t \in \mathbb{R}_{\geq 0}}$ which evolves according to the continuous-time dynamics

$$\tau \dot{d}^p(t) = \underbrace{d_{\text{base}}^p(t) - \sum_{\tilde{p} \in \mathcal{P}} \tilde{B}^{p\tilde{p}} u^{\tilde{p}}(t)}_{\text{linear price-demand curve}} - d^p(t),$$

where $\mathbf{u}^{\tilde{p}} = (u^{\tilde{p}}(t))_{t \in \mathbb{R}_{\geq 0}}$ are price changes around the value at which $\tilde{p} \in \mathcal{P}$ sell its products and $\mathbf{d}_{\text{base}}^p = (d_{\text{base}}^p(t))_{t \in \mathbb{R}_{\geq 0}}$ is some fluctuating baseline demand. Specifically, the baseline demands are of the form $d_{\text{base}}^p(t) = \bar{d}_{\text{base}}^p + v^p(t)$, for all $p \in \mathcal{P}$, given fixed $\bar{d}_{\text{base}}^p = 10$ and noise process $\mathbf{v}^p = (v^p(t))_{t \in \mathbb{R}_{\geq 0}}$. The market parameters $\tau \in \mathbb{R}_{\geq 0}$ and $\tilde{B}^{p\tilde{p}} \in \mathbb{R}_{\geq 0}$ ($\forall p, \tilde{p} \in \mathcal{P}$) describe how local demands respond to price changes: We set $\tau = 1.2$ and sample $B^{p\tilde{p}} \sim \text{Uniform}(0.5, 1.5)$ for all $p, \tilde{p} \in \mathcal{P}$. In this problem, companies aim at devising pricing policies to stabilise their demands around \bar{d}_{base}^p , which provides a stable profit margin, while satisfying a price-cap regulation enforcing

$$-\bar{u}_{\text{avg}} \leq \frac{1}{N_P} \sum_{p \in \mathcal{P}} u^p(t) \leq \bar{u}_{\text{avg}}, \quad \bar{u}_{\text{avg}} = 0.5.$$

Define $\mathbf{x} = (\mathbf{x}^1, \dots, \mathbf{x}^{N_P})$, with $\mathbf{x}^p = (d^p(t) - \bar{d}_{\text{base}}^p)_{t \in \mathbb{R}_{\geq 0}}$, and $\tilde{B}^p = [\tilde{B}^{p1} \dots \tilde{B}^{pN_P}]$, for all $p \in \mathcal{P}$. Considering a zero-order hold of inputs with period $\Delta t = 1/4$ [days], the game $\mathcal{G}_\infty^{\text{LQ}}$ can be described by the discrete-time dynamics³

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \sum_{p \in \mathcal{P}} \mathbf{B}^p u_t^p + \mathbf{w}_t$$

with $\mathbf{A} = \exp(-\tau \Delta t) \mathbf{I}_{N_P}$ and $\{\mathbf{B}^p = -\frac{1}{\tau} (\mathbf{A} - \mathbf{I}_{N_P}) \tilde{B}^p\}_{p \in \mathcal{P}}$, and the noise process $\mathbf{w} = (-\frac{1}{\tau} (\mathbf{A} - \mathbf{I}_{N_P}) \mathbf{v}_t)_{t \in \mathbb{N}}$ [46]. The companies assume that the baseline demand fluctuations satisfy $\mathbf{w}_t \in \mathcal{W} = \{\mathbf{w}_t \in \mathbb{R}^{N_P} : \|\mathbf{w}_t\|_\infty \leq 1\}$ for all $t \in \mathbb{N}$. Under

³With a slight abuse of notation, we use t for both the continuous-time $(x(t), t \in \mathbb{R}_{\geq 0})$ and discrete-time $(x_t, t \in \mathbb{N})$ signals.

this representation, each player's objective is formulated as solving for a policy which minimises the functional

$$J^p(u^p, u^{-p}) = \mathbb{E} \left[\sum_{t=0}^{\infty} \left(\alpha^p \|x_t\|_2^2 + \beta^p \|u_t^p\|_2^2 \right) \right],$$

given weights $\alpha^p, \beta^p \in \mathbb{R}_{\geq 0}$. We sample $\alpha^p \sim \text{Uniform}(5, 15)$ and $\beta^p \sim \text{Uniform}(0.3, 0.6)$ for each $p \in \mathcal{P}$. The operational constraints, $u^p \in U^p(u^{-p})$, are defined by the constraint sets

$$\begin{aligned} \mathcal{X} &= \mathbb{R}^{N_x}; \\ \mathcal{U}^p &= \mathbb{R}^{N_u^p}; \\ \mathcal{U}_G &= \{u_t \in \prod_{p \in \mathcal{P}} \mathbb{R}^{N_u^p} : \frac{1}{N_P \cdot \bar{u}_{\text{avg}}} [\mathbf{1}_{N_u} \ -\mathbf{1}_{N_u}]^T u_t \leq [1 \ 1]^T\}, \end{aligned}$$

We consider that players design their state-feedback policies, $K^p = \Phi_u^p \Phi_x^{-1} \in \mathcal{C}^p$, with a FIR horizon of $N = 16$ and no structural constraints, that is, $S_x = I$ and $S_u^p = I (\forall p \in \mathcal{P})$. In practice, this implies that companies have perfect information of any demand x^p and price u^p changes in the market.

As in Section IV-A, we simulate an instance of the game $\mathcal{G}_{\infty}^{\text{LQ}}$ alongside the SLS-BRD routine (Algorithm 4) with players using (approximately) best-response maps, $\{\widehat{BR}_{\Phi}^p\}_{p \in \mathcal{P}}$, given $\gamma = 0.95$. Policies are updated simultaneously every $\Delta T = 1$ stage with learning rate $\eta = 1/4$. Under the above setup, the game $\mathcal{G}_{\infty}^{\text{LQ}}$ is not a potential game and thus a GFNE cannot be easily computed in advance. Finally, we remark that solutions to this problem are neither unique nor necessarily admissible: Different initial profiles result in convergence to different fixed-points, which might be unfair to a subset of players.

The convergence of the SLS-BRD routine to a fixed-point $K^* = \Phi_u^* \Phi_x^{-1} = (\Phi_u^1, \dots, \Phi_u^{N_P}) \Phi_x^{-1}$ is shown in Figure 6. Since the exact fixed-point to which the routine will converge is not known for this problem, we let $\Phi_u^* \approx \Phi_{u, k_f}$ for $k_f = 828$ (when the policy updates are interrupted) and analyse the convergence with respect to this point. In this case, the iterates approach the fixed-point mostly at a geometric rate with the policies requiring roughly 800 updates (or 200 days in-game) before changes become numerically negligible. The rate of convergence is also observed to slow around $k = 15$, when the generalised constraints (\mathcal{U}_G) become active. We stress that the best-responses $\{\widehat{BR}_{\Phi}^p\}_{p \in \mathcal{P}}$ cannot be (globally) contractive in this case, as the set of fixed points $\Omega_{\mathcal{G}_{\infty}^{\text{LQ}}}^{\varepsilon}$ is not a singleton.

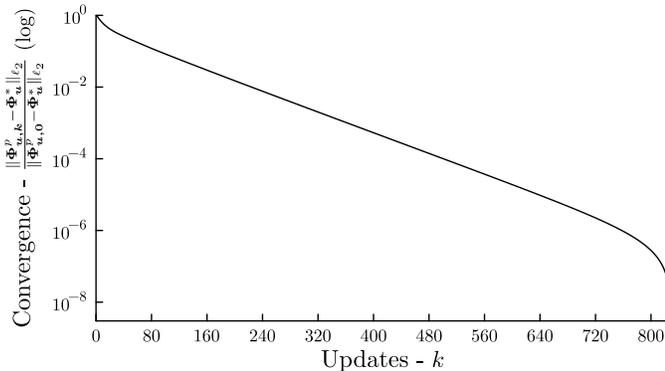


Figure 6. Market game: Convergence of the SLS-BRD routine.

In Figure 7 (top), the relative distances between the individual updates $\{\Phi_{u,k}^p, \Phi_{u,k-1}^p\}_{p \in \mathcal{P}}$ are displayed. As in the previous example, the updates are of similar magnitude for all players $p \in \mathcal{P}$ and they become numerically negligible at a faster rate than the global convergence in Figure 6. The convergence of the distance between updates is shown to slow considerably, especially for $p \in \{1, 4\}$, around $k = 50$. Thereafter, these relative distances continue to decrease at a steady rate. In Figure 7 (bottom), the relative distances $\|\Phi_{x,k}^p - \Phi_{x,k}\|_{\ell_2} / \|\Phi_{x,k}^p\|_{\ell_2}$, $k \in \mathbb{N}_+$, between responses $\{\Phi_{x,k}^p\}_{p \in \mathcal{P}}$ and $\Phi_{x,k} = F_{\Phi} \Phi_{u,k}$ are shown. As before, these distances decrease at a similar rate and are relatively small since the initial stages of the game. The policies $K_k = (\Phi_{u,k}^1 (\Phi_{x,k}^1)^{-1}, \dots, \Phi_{u,k}^{N_P} (\Phi_{x,k}^{N_P})^{-1})$ thus approximate the stabilising profile $(\Phi_{u,k}^1, \dots, \Phi_{u,k}^{N_P}) \Phi_{x,k}^{-1}$ as $k \rightarrow \infty$ and are expected to be stabilising for all iterations.

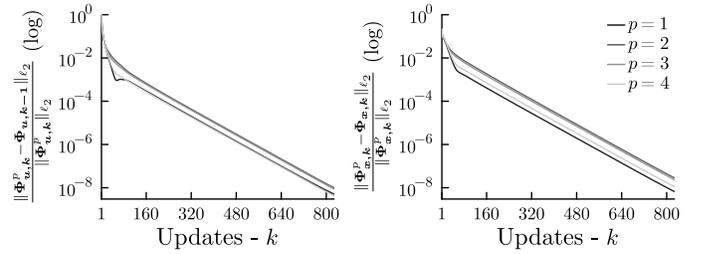


Figure 7. Market game: Relative distance between updates $(\Phi_{u,k}^p, \Phi_{u,k+1}^p)$, top panels, and responses $(\Phi_{x,k}^p, \Phi_{x,k})$, bottom panels, for all $p \in \mathcal{P}$.

The evolution of the game given each player's actions is displayed Figure 5 for the *in-game* period $t \in (260, 400]$ days, after the policy updates have been interrupted. We compare the performance of the policy profile $K^* = \Phi_u^* \Phi_x^{-1}$ with the evolution obtained through the open-loop operation $u(t) = 0$. During this period, we simulate a worst-case fluctuation on the baseline demand for the companies' product by defining

$$w(t) = \begin{cases} (1, 1, 1, 1) & \text{for } t \in [265, 279]; \\ (1, -1, 1, -1) & \text{for } t \in [293, 307]; \\ (-1, 1, -1, 1) & \text{for } t \in [321, 335]; \\ (1, 1, -1, -1) & \text{for } t \in [349, 363]; \\ (-1, -1, -1, -1) & \text{for } t \in [377, 391], \end{cases}$$

and $w(t) = 0$ otherwise. The results show that the policy profile obtained by the SLS-BRD routine allows the companies to efficiently respond to changes in their local demands. In general, the players coordinate price changes to alleviate the deviations from the baseline demand caused by the disturbances. Moreover, the player's actions satisfy the price-cap constraint $\frac{1}{N_P} |\sum_{p \in \mathcal{P}} u^p(t)| \leq 0.5$ in all disturbance scenarios. The best performance is observed for companies $p \in \{3, 4\}$, whereas $p = 1$ behaves noticeably worse than all players (especially for $w(t) = (1, -1, 1, -1)$ and $w(t) = (-1, 1, -1, 1)$, when the open-loop operation seems to attain better results). This highlights the fact that fixed-points obtained through the SLS-BRD routine are not necessarily *admissible* GFNE, and thus might be unfavourable for a subset of players. Finally, we note that these policies are still not able to completely reject the effect of the extreme disturbances: Achieving zero-offset would require incorporating integral action into the player's policies.

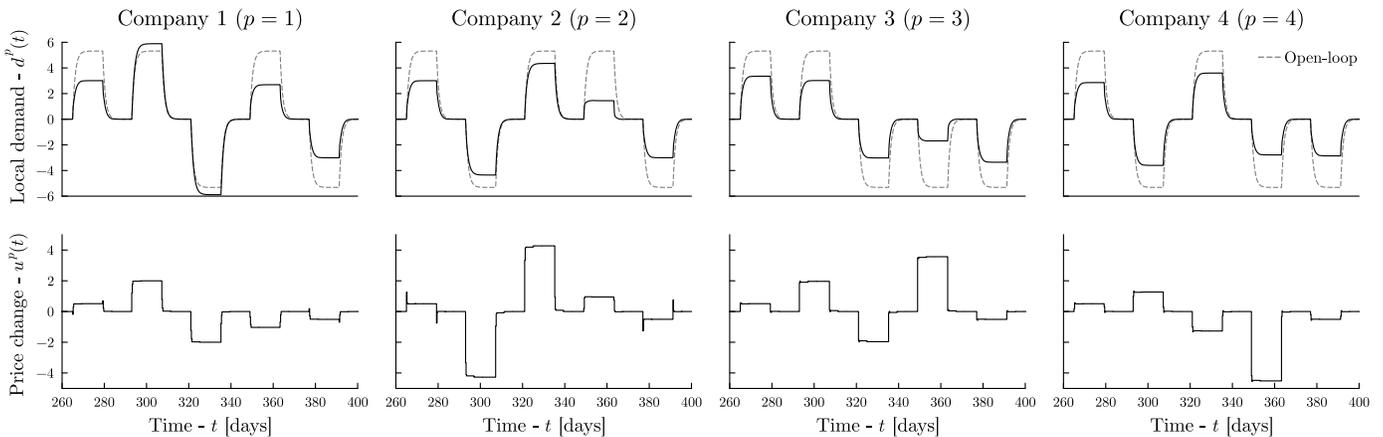


Figure 8. Market game, $t \in (260, 400]$ days: State \mathbf{x} (top panels) and applied control \mathbf{Bu} (bottom panels) trajectories for an execution of the game $\mathcal{G}_{\infty}^{\text{LQ}}$. The dashed lines refer to the state-trajectories resulting from an open-loop operation of the market with $u(t) = 0$ for all $t \in \mathbb{R}_{\geq 0}$

V. CONCLUDING REMARKS

This work presents the SLS-BRD, an algorithm for generalised feedback Nash equilibrium seeking in N_P -players non-cooperative games. The method is based on the class of best-response dynamics algorithms for decentralised learning and consists of players updating and announcing a parametrisation of their policies until converging to an equilibrium. Because not updating control actions explicitly, this learning dynamics can be performed alongside the execution of the game. Using results from operator theory, we propose that convergence to a Nash equilibrium is geometric if certain conditions are met: Namely, that the best-response maps are contractive operators. Our approach leverages the System Level Synthesis framework to formulate each player's best-response map as the solution to robust finite-horizon optimal control problems. This framework benefits the SLS-BRD algorithm also by allowing richer information patterns to be enforced directly at the synthesis level. After the main theoretical and practical aspects are discussed, the algorithm is demonstrated on an exemplary problems on the decentralised control of multi-agent systems.

REFERENCES

- [1] T. Başar and G. J. Olsder, *Dynamic noncooperative game theory*. SIAM, 1998.
- [2] F. Facchinei and C. Kanzow, "Generalized Nash equilibrium problems," *Annals of Operations Res.*, vol. 175, no. 1, pp. 177–211, 2010.
- [3] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou, "The complexity of computing a Nash equilibrium," *Communications of the ACM*, vol. 52, no. 2, pp. 89–97, 2009.
- [4] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. Cambridge University Press, 2011.
- [5] J. S. Shamma and G. Arslan, "Dynamic fictitious play, dynamic gradient play, and distributed convergence to Nash equilibria," *IEEE Transactions on Automatic Control*, vol. 50, no. 3, pp. 312–327, 2005.
- [6] J. Hofbauer and S. Sorin, "Best response dynamics for continuous zero-sum games," *Discrete and Continuous Dynamical Systems*, vol. 6, no. 1, pp. 215–224, 2006.
- [7] J. R. Marden, G. Arslan, and J. S. Shamma, "Joint strategy fictitious play with inertia for potential games," *IEEE Transactions on Automatic Control*, vol. 54, no. 2, pp. 208–220, 2009.
- [8] A. Cortés and S. Martínez, "Self-triggered best-response dynamics for continuous games," *IEEE Transactions on Automatic Control*, vol. 60, no. 4, pp. 1115–1120, 2014.
- [9] S. Grammatico, "Dynamic control of agents playing aggregative games with coupling constraints," *IEEE Transactions on Automatic Control*, vol. 62, no. 9, pp. 4537–4548, 2017.
- [10] B. Swenson, R. Murray, and S. Kar, "On best-response dynamics in potential games," *SIAM Journal on Control and Optimization*, vol. 56, no. 4, pp. 2734–2767, 2018.
- [11] A. Jafari, A. Greenwald, D. Gondek, and G. Ercal, "On no-regret learning, fictitious play, and Nash equilibrium," in *ICML*, vol. 1, pp. 226–233, 2001.
- [12] N. Cesa-Bianchi and G. Lugosi, *Prediction, learning, and games*. Cambridge university press, 2006.
- [13] G. J. Gordon, A. Greenwald, and C. Marks, "No-regret learning in convex games," in *Proceedings of the 25th international conference on Machine learning*, pp. 360–367, 2008.
- [14] J. Hartline, V. Syrgkanis, and E. Tardos, "No-regret learning in bayesian games," *Advances in Neural Information Processing Systems*, vol. 28, 2015.
- [15] A. Celli, A. Marchesi, G. Farina, and N. Gatti, "No-regret learning dynamics for extensive-form correlated equilibrium," *Advances in Neural Information Processing Systems*, vol. 33, pp. 7722–7732, 2020.
- [16] G. Belgioioso, D. Liao-McPherson, M. H. de Badyn, S. Bolognani, R. S. Smith, J. Lygeros, and F. Dörfler, "Online feedback equilibrium seeking," *arXiv preprint arXiv:2210.12088*, 2022.
- [17] G. Scutari, S. Barbarossa, and D. Palomar, "Potential games: A framework for vector power control problems with coupled constraints," in *2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*, vol. 4, pp. IV–IV, 2006.
- [18] W. Saad, Z. Han, H. V. Poor, and T. Basar, "Game-theoretic methods for the smart grid: An overview of microgrid systems, demand-side management, and smart grid communications," *IEEE signal processing magazine*, vol. 29, no. 5, pp. 86–105, 2012.
- [19] S. Maghsudi and S. Stańczak, "Channel selection for network-assisted d2d communication via no-regret bandit learning with calibrated forecasting," *IEEE Transactions on Wireless Communications*, vol. 14, no. 3, pp. 1309–1322, 2015.
- [20] J. F. Fisac, E. Bronstein, E. Stefanosson, D. Sadigh, S. S. Sastry, and A. D. Dragan, "Hierarchical game-theoretic planning for autonomous vehicles," in *2019 International conference on robotics and automation (ICRA)*, pp. 9590–9596, IEEE, 2019.

- [21] R. Spica, E. Cristofalo, Z. Wang, E. Montijano, and M. Schwager, “A real-time game theoretic planner for autonomous two-player drone racing,” *IEEE Transactions on Robotics*, vol. 36, no. 5, pp. 1389–1403, 2020.
- [22] Y. Chen and J. W. Vaughan, “A new understanding of prediction markets via no-regret learning,” in *Proceedings of the 11th ACM conference on Electronic commerce*, pp. 189–198, 2010.
- [23] M. Braverman, J. Mao, J. Schneider, and M. Weinberg, “Selling to a no-regret buyer,” in *Proceedings of the 2018 ACM Conference on Economics and Computation*, pp. 523–538, 2018.
- [24] D. Aussel, L. Brotcorne, S. Lepaul, and L. von Niederhäusern, “A trilevel model for best response in energy demand-side management,” *European Journal of Operational Research*, vol. 281, no. 2, pp. 299–315, 2020.
- [25] J. Engwerda, W. Van Den Broek, and J. M. Schumacher, “Feedback Nash equilibria in uncertain infinite time horizon differential games,” in *Proceedings of the 14th International Symposium of Mathematical Theory of Networks and Systems, MTNS 2000*, pp. 1–6, Unknown Publisher, 2000.
- [26] D. Vrabie and F. Lewis, “Integral reinforcement learning for online computation of feedback Nash strategies of nonzero-sum differential games,” in *49th IEEE Conference on Decision and Control (CDC)*, pp. 3066–3071, IEEE, 2010.
- [27] P. V. Reddy and G. Zaccour, “Feedback Nash equilibria in linear-quadratic difference games with constraints,” *IEEE Transactions on Automatic Control*, vol. 62, no. 2, pp. 590–604, 2017.
- [28] P. V. Reddy and G. Zaccour, “Open-loop and feedback Nash equilibria in constrained linear-quadratic dynamic games played over event trees,” *Automatica*, vol. 107, pp. 162–174, 2019.
- [29] G. Kossioris, M. Plexousakis, A. Xepapadeas, A. de Zeeuw, and K.-G. Mäler, “Feedback Nash equilibria for non-linear differential games in pollution control,” *Journal of Economic Dynamics and Control*, vol. 32, no. 4, pp. 1312–1331, 2008.
- [30] R. Kamalapurkar, J. R. Klotz, and W. E. Dixon, “Concurrent learning-based approximate feedback-Nash equilibrium solution of n-player nonzero-sum differential games,” *IEEE/CAA journal of Automatica Sinica*, vol. 1, no. 3, pp. 239–247, 2014.
- [31] D. Fridovich-Keil, E. Ratner, L. Peters, A. D. Dragan, and C. J. Tomlin, “Efficient iterative linear-quadratic approximations for nonlinear multi-player general-sum differential games,” in *2020 IEEE international conference on robotics and automation (ICRA)*, pp. 1475–1481, IEEE, 2020.
- [32] D. Fridovich-Keil, V. Rubies-Royo, and C. J. Tomlin, “An iterative quadratic method for general-sum differential games with feedback linearizable dynamics,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2216–2222, IEEE, 2020.
- [33] A. Tanwani and Q. Zhu, “Feedback Nash equilibrium for randomly switching differential-algebraic games,” *IEEE Transactions on Automatic Control*, vol. 65, no. 8, pp. 3286–3301, 2019.
- [34] Z. Wang, R. Spica, and M. Schwager, “Game theoretic motion planning for multi-robot racing,” in *Distributed Autonomous Robotic Systems: The 14th International Symposium*, pp. 225–238, Springer, 2019.
- [35] F. Laine, D. Fridovich-Keil, C.-Y. Chiu, and C. Tomlin, “The computation of approximate generalized feedback Nash equilibria,” *SIAM Journal on Optimization*, vol. 33, no. 1, pp. 294–318, 2023.
- [36] J. Anderson, J. C. Doyle, S. H. Low, and N. Matni, “System level synthesis,” *Annual Reviews in Control*, vol. 47, pp. 364–393, 2019.
- [37] I. L. Glicksberg, “A further generalization of the kakutani fixed theorem, with application to Nash equilibrium points,” *Proceedings of the American Mathematical Society*, vol. 3, no. 1, pp. 170–174, 1952.
- [38] E. K. Ryu and S. Boyd, “Primer on monotone operator methods,” *Appl. comput. math.*, vol. 15, no. 1, pp. 3–43, 2016.
- [39] E. K. Ryu and W. Yin, *Large-scale convex optimization: algorithms & analyses via monotone operators*. Cambridge University Press, 2022.
- [40] J. Liang, J. Fadili, and G. Peyré, “Convergence rates with inexact non-expansive operators,” *Mathematical Programming*, vol. 159, pp. 403–434, 2016.
- [41] T. Başar, “Stochastic differential games and intricacy of information structures,” in *Dynamic Games in Economics*, pp. 23–49, Springer, 2014.
- [42] K. Zhou and J. C. Doyle, *Essentials of robust control*, vol. 104. Prentice hall Upper Saddle River, NJ, 1998.
- [43] A. Nayyar and T. Başar, “Dynamic stochastic games with asymmetric information,” in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pp. 7145–7150, IEEE, 2012.
- [44] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*. Cambridge University Press, 2017.
- [45] S. Zazo, S. Valcarcel Macua, M. Sánchez-Fernández, and J. Zazo, “Dynamic potential games with constraints: Fundamentals and applications in communications,” *IEEE Transactions on Signal Processing*, vol. 64, no. 14, pp. 3806–3821, 2016.
- [46] K. J. Åström and B. Wittenmark, *Computer-controlled systems: theory and design*. Dover Publications, 2013.
- [47] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.

APPENDIX

A. Proof of Corollary 4.1

Corollary. A policy $\mathbf{K}^p = \Phi_u^p \Phi_x^{-1}$ ($\forall p \in \mathcal{P}$) is defined by the kernel $\Phi^p = \Phi_u^p * \Phi_x^{-1}$, and implemented as

$$\xi_t = -\sum_{\tau=1}^t \Phi_{x,\tau+1} \xi_{t-\tau} + x_t; \quad (49a)$$

$$u_t^p = \sum_{\tau=0}^t \Phi_{u,\tau+1}^p \xi_{t-\tau}, \quad (49b)$$

using an auxiliary internal state $\xi = (\xi_n)_{n \in \mathbb{N}}$ with $\xi_0 = x_0$.

Proof. The statement $\Phi^p = \Phi_u^p * \Phi_x^{-1}$ follows directly from the inverse Z-Transform of $\hat{\mathbf{K}}^p = \hat{\Phi}_u^p \hat{\Phi}_x^{-1}$ and $\Phi^p = (\Phi_n^p)_{n \in \mathbb{N}}$ being the convolution kernel of \mathbf{K}^p . The operations Eq. (49) are obtained as the inverse Z-Transform of Eq. (25) and the fact that $\mathcal{Z}^{-1}[z(I - \hat{\Phi}_x)\hat{\xi}] = \xi_{t+1} - \xi_t - \sum_{\tau=2}^{t+1} \Phi_{x,\tau} \xi_{t+1-\tau}$. \square

B. Proof of Theorem 5

Theorem. Consider a fixed-point $\Phi_u^\varepsilon \in \widehat{BR}_\Phi(\Phi_u^\varepsilon)$ and assume that $\|\Phi_{x,N}^*\|_F \leq \gamma$ for $\Phi_x^* = \mathbf{F}_\Phi \Phi_x^*$ obtained from the original best-response $\Phi_x^* \in \widehat{BR}_\Phi(\Phi_u^\varepsilon)$. Then, the profile $\Phi_u^\varepsilon = (\Phi_u^{1\varepsilon}, \dots, \Phi_u^{N\varepsilon})$ is an ε -GNE of \mathcal{G}_∞^Φ satisfying

$$J^p(\Phi_u^{p\varepsilon}, \Phi_u^{-p\varepsilon}) \leq \min_{\Phi_u^p \in U_\Phi^p(\Phi_u^{-p\varepsilon})} J^p(\Phi_u^p, \Phi_u^{-p\varepsilon}) + \varepsilon \quad (50)$$

with $\varepsilon = \max_{p \in \mathcal{P}} \gamma J^p(\Phi_u^{p\varepsilon}, \Phi_u^{-p\varepsilon})$ for every player $p \in \mathcal{P}$.

Proof. Let $\Phi_u^* \in \widehat{BR}_\Phi(\Phi_u^\varepsilon)$. We construct a candidate fixed-point by defining $\tilde{\Phi}_u^\varepsilon = (\tilde{\Phi}_{u,n}^\varepsilon)_{n=1}^{N-1}$. Clearly, $\tilde{\Phi}_u^\varepsilon$ satisfies all the constraints in Problem (33) by construction and also $(\tilde{\Phi}_{x,n}^*)_{n=1}^N = \tilde{\Phi}_x^\varepsilon = \mathbf{F}_\Phi \tilde{\Phi}_u^\varepsilon$ satisfies $\|\Phi_{x,N}^*\|_F \leq \gamma$ by our assumption. Now, consider $\Phi_u^\varepsilon \in \widehat{BR}_\Phi(\Phi_u^\varepsilon)$. From optimality,

$$J^p(\Phi_u^{p\varepsilon}, \Phi_u^{-p\varepsilon}) \leq J^p(\tilde{\Phi}_u^{p\varepsilon}, \Phi_u^{-p\varepsilon}) \leq \frac{1}{1-\gamma} J^p(\Phi_u^*, \Phi_u^{-p\varepsilon}), \quad (51)$$

where the second inequality derives from the fact that the quadratic objective functional is larger for the infinite-impulse response Φ_u^* and that $\frac{1}{1-\gamma} > 1$. Finally, by definition,

$$\Phi_u^{p*} = \arg \min_{\Phi_u^p \in U_\Phi^p(\Phi_u^{-p\varepsilon})} J^p(\Phi_u^p, \Phi_u^{-p\varepsilon})$$

and thus the inequality Eq. (50) follows directly from Eq. (51) by defining $\varepsilon = \max_{p \in \mathcal{P}} \gamma J^p(\Phi_u^{p\varepsilon}, \Phi_u^{-p\varepsilon})$. \square

C. Proof of Theorem 8

In order to prove this theorem, we start by introducing some useful lemmas. Firstly, let $\mathcal{W} = \{w_t \in \mathbb{R}^{N_x} : \|w_t\|_\infty \leq 1\}$ and assume $\|\Phi_{x,N}\|_F^2 \leq \gamma$ is strictly satisfied, and define the operators $\{\mathbf{H}^{p\tilde{p}}\}_{p,\tilde{p} \in \mathcal{P}}$ as in Section III-B. In the following, we slightly abuse notation and apply matrix algebra directly to signals and operators: These should be understood as applied to their equivalent matrix representations.

Lemma 9. Define the signal $\tilde{\Phi}_u^p = (\text{vec } \Phi_{u,n}^p)_{n=1}^{N-1}$ for each $p \in \mathcal{P}$. The (approximately) best-response mappings satisfy $\tilde{\Phi}_u^{p*} \in \widehat{BR}_\Phi^p(\tilde{\Phi}_u^{-p*})$, where $\tilde{\Phi}_u^{p*}$ is obtained from a solution of

$$\underset{t, \tilde{\Phi}_u^p}{\text{minimize}} \quad \tilde{\Phi}_u^{p\top} \tilde{\mathbf{H}}^{pp} \tilde{\Phi}_u^p + 2(\tilde{\mathbf{H}}^{p,-p} \tilde{\Phi}_u^{-p} + \tilde{\mathbf{H}}^{p0})^\top \tilde{\Phi}_u^p$$

$$\text{subject to} \quad -\tilde{t} \leq \tilde{\mathbf{G}}_u^p \tilde{\Phi}_u^p \leq \tilde{t},$$

$$\mathbf{1}^\top t_j = [g_u^p]_j, \quad j = 1, \dots, N_{U^p}$$

given the operators

$$\tilde{\mathbf{H}}^{p\tilde{p}} = \text{blkdiag}(\mathbf{H}^{p\tilde{p}}, \dots, \mathbf{H}^{p\tilde{p}}) = I_{N_x} \otimes \mathbf{H}^{p\tilde{p}}$$

$$\tilde{\mathbf{G}}_u^p = \text{blkdiag}(G_u^p, \dots, G_u^p) = I_{(N-1)N_x} \otimes G_u^p,$$

for $\tilde{p} \in \mathcal{P}$. Here, t is an auxiliary block-vector of appropriate dimensions and $\tilde{t} = U^\top t$ for some unitary transformation U .

Proof. Consider Problem (33) with the aforementioned assumptions. Firstly, we convert the linear dynamics,

$$\text{vec}(\Phi_x) = \text{vec}(\mathbf{S}_+ \mathbf{A} \Phi_x + \sum_p \mathbf{S}_+ \mathbf{B}^p \Phi_u^p + \delta I_{N_x}),$$

$$\Rightarrow \tilde{\Phi}_x = (I_{N_x} \otimes \mathbf{S}_+ \mathbf{A}) \tilde{\Phi}_x + \sum_p (I_{N_x} \otimes \mathbf{S}_+ \mathbf{B}^p) \tilde{\Phi}_u^p + \tilde{\delta} \tilde{\Phi}_{x,1}$$

$$\Rightarrow \tilde{\Phi}_x = I_{N_x} \otimes (I - \mathbf{S}_+ \mathbf{A})^{-1} (\sum_p \mathbf{S}_+ \mathbf{B}^p \tilde{\Phi}_u^p + \tilde{\delta} \tilde{\Phi}_{x,1})$$

$$\Rightarrow \tilde{\Phi}_x = I_{N_x} \otimes (\sum_p \mathbf{F}_\Phi^p \tilde{\Phi}_u^p + \mathbf{F}_\Phi^0),$$

where $\tilde{\Phi}_{x,1} = \text{vec } I_{N_x}$. After some standard manipulations, the objective functional becomes

$$\begin{aligned} J^p(\cdot) &= \tilde{\Phi}_u^{p\top} (I_{N_x} \otimes \mathbf{H}^{pp}) \tilde{\Phi}_u^p \\ &\quad + 2((I_{N_x} \otimes \mathbf{H}^{p,-p}) \tilde{\Phi}_u^{-p} + (I_{N_x} \otimes \mathbf{H}^{p0}))^\top \tilde{\Phi}_u^p \\ &\quad + (\text{affine terms}). \end{aligned}$$

Consider now the vectorised constraints,

$$\sum_{n=0}^{N-1} \|\text{vec}([G_{u,j}^p] \Phi_{u,n}^p)\|_1 \leq [g_u^p]_j$$

$$\Rightarrow \sum_{n=0}^{N-1} \|(I_{N_x} \otimes [G_{u,j}^p]) \tilde{\Phi}_{u,n}^p\|_1 \leq [g_u^p]_j,$$

$$\Rightarrow -t_{j,n} \leq (I_{N_x} \otimes [G_{u,j}^p]) \tilde{\Phi}_{u,n}^p \leq t_{j,n}, \quad n = 1, \dots, N-1$$

where we introduced the auxiliary vector $t_{j,n} \in \mathbb{R}^{N_x}$ then reformulated the sum of 1-norms into its epigraph form by adding the constraint $\mathbf{1}^\top t_j = [g_u^p]_j$ with $t_j = \text{c}\circ\mathbf{1}(t_{j,1}, \dots, t_{j,N-1})$. Equivalently, these inequalities can be written as

$$-t_j \leq (I_{(N-1)N_x} \otimes [G_{u,j}^p]) \tilde{\Phi}_u^p \leq t_j, \quad \mathbf{1}^\top t_j = [g_u^p]_j.$$

Since $I_{(N-1)N_x} \otimes G_u^p = U \text{c}\circ\mathbf{1}(I_{(N-1)N_x} \otimes [G_{u,j}^p])_{j=1}^{N_{U^p}}$ for some unitary matrix U , we concatenate all inequalities into

$$-U^\top t \leq (I_{(N-1)N_x} \otimes G_u^p) \tilde{\Phi}_u^p \leq U^\top t,$$

with $t = \text{c}\circ\mathbf{1}(t_1, \dots, t_{N_{U^p}})$. Finally, we define $\{\tilde{\mathbf{H}}^{p\tilde{p}}\}_{\tilde{p} \in \mathcal{P}}$ and $\tilde{\mathbf{G}}_u^p$ as in the lemma's statement and reformulate the best-response mapping accordingly. \square

The above result shows that $\{\widehat{BR}_\Phi^p\}_{p \in \mathcal{P}}$ are the solutions of parametric quadratic programs with linear constraints: They consist of piecewise affine operators. The following lemma relates the Lipschitz constants of a collection of affine operators with that of the operator defined by their concatenation.

Lemma 10. Let $T = (T^1, \dots, T^{N_P}) : \mathcal{X} \rightarrow \mathcal{Y}$ be a mapping constructed from the affine operators $T^p x = A^p x + b^p$, $p \in \mathcal{P}$. Then, $L_T = \|(A^1, \dots, A^{N_P})\|_{2 \rightarrow 2}$ is the (tightest) Lipschitz constant for T , which can be relaxed by the upper-bound

$$L_T^2 \leq \sum_{p \in \mathcal{P}} \|A^p\|_{2 \rightarrow 2}^2 \quad (53)$$

Proof. Firstly, $Tx = (A^1, \dots, A^{N_P})x + (b^1, \dots, b^{N_P})$, that is, T is an affine operator with $A = (A^1, \dots, A^{N_P}) \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$. From definition, the operator norm $\|(A^1, \dots, A^{N_P})\|_{2 \rightarrow 2} = L_T$ is smallest Lipschitz constant for this mapping. Finally,

$$\begin{aligned} L_T^2 &= \|(A^1, \dots, A^{N_P})\|_{2 \rightarrow 2}^2 \\ &= \|\sum_{p \in \mathcal{P}} A^p A^p\|_{2 \rightarrow 2} \\ &\leq \sum_{p \in \mathcal{P}} \|A^p A^p\|_{2 \rightarrow 2} \\ &= \sum_{p \in \mathcal{P}} \|A^p\|_{2 \rightarrow 2}^2, \end{aligned}$$

since $(A^1, \dots, A^{N_P})^*(A^1, \dots, A^{N_P}) = \sum_{p \in \mathcal{P}} A^p A^p$. \square

Finally, we can proceed to prove the Theorem 7.

Theorem. Consider $\mathcal{X} = \mathbb{R}^{N_x}$ and $\mathcal{U}_G = \prod_{p \in \mathcal{P}} \mathbb{R}^{N_u^p}$. Then, The map \widehat{BR}_Φ is $L_{\widehat{BR}_\Phi}$ -Lipschitz with

$$L_{\widehat{BR}_\Phi} = \sqrt{\sum_{p \in \mathcal{P}} (L_{\widehat{BR}_\Phi}^p)^2}, \quad (54)$$

given the player-specific constants

$$L_{\widehat{BR}_\Phi}^p = (1 + \kappa(\mathbf{H}^{pp})) \|(\mathbf{H}^{pp})^\dagger \mathbf{H}^{p,-p}\|_{2 \rightarrow 2}, \quad (55)$$

with condition number $\kappa(\mathbf{H}^{pp}) = \|(\mathbf{H}^{pp})^\dagger\|_{2 \rightarrow 2} \|\mathbf{H}^{pp}\|_{2 \rightarrow 2}$.

Proof. Consider the reformulation of the (approximately)best-response maps \widehat{BR}_Φ introduced in Lemma 9. Further, assume that the active constraints are known for each $\Phi_u^{-p} \in \mathcal{C}_u^{-p}$: We let $\tilde{\mathbf{G}}_{u,\mathcal{A}}^p = e_{\mathcal{A}} \tilde{\mathbf{G}}_u^p$ denote the rows associated with the active inequalities indexed by \mathcal{A} , where the sign has been absorbed into the selector matrix $e_{\mathcal{A}}$. Using the KKT conditions [47], a solution of this problem is obtained from its Lagrangian as

$$\tilde{\Phi}_u^* = -(\tilde{\mathbf{H}}^{pp})^\dagger (\tilde{\mathbf{H}}^{p,-p} \tilde{\Phi}_u^{-p} + \tilde{\mathbf{H}}^{p0} + \frac{1}{2} \tilde{\mathbf{G}}_u^{\mathcal{A}^\top} \lambda^p), \quad (56)$$

where λ^p is an optimal solution to the associated dual problem. From complementary slackness, $\lambda_i^p = 0$ for all $i \notin \mathcal{A}$, and the non-zero lagrangian multiplier are

$$\begin{aligned} \lambda_{\mathcal{A}}^p &= \tilde{\mathbf{G}}_{u,\mathcal{A}}^p \tilde{\Phi}_u^* - \tilde{t}_{\mathcal{A}} \\ &= -2 [\tilde{\mathbf{G}}_{u,\mathcal{A}}^p (\tilde{\mathbf{H}}^{pp})^\dagger \tilde{\mathbf{G}}_{u,\mathcal{A}}^{\mathcal{A}^\top}]^\dagger [\tilde{\mathbf{G}}_{u,\mathcal{A}}^p (\tilde{\mathbf{H}}^{pp})^\dagger \tilde{\mathbf{H}}^{p,-p}] \tilde{\Phi}_u^{-p} \\ &\quad + (\text{affine terms}). \end{aligned} \quad (57)$$

Now, introduce the auxiliary operator $\tilde{\mathbf{V}}^p = (\tilde{\mathbf{G}}_{u,\mathcal{A}}^p)^\dagger (\tilde{\mathbf{G}}_{u,\mathcal{A}}^p)$. After some algebra, we can combine Eqs. (56)–(57) into

$$\begin{aligned} \tilde{\Phi}_u^* &= -\left\{ (\tilde{\mathbf{H}}^{pp})^\dagger \tilde{\mathbf{H}}^{p,-p} \right. \\ &\quad \left. + (\tilde{\mathbf{H}}^{pp})^\dagger \tilde{\mathbf{V}}^{p^\top} \tilde{\mathbf{H}}^{pp} \tilde{\mathbf{V}}^p (\tilde{\mathbf{H}}^{pp})^\dagger \tilde{\mathbf{H}}^{p,-p} \right\} \tilde{\Phi}_u^{-p} \\ &\quad + (\text{affine terms}). \end{aligned} \quad (58)$$

Thus, the optimal solution takes the form of an affine operator $\tilde{\Phi}_u^* = (0, \mathbf{A}^{-p}) \tilde{\Phi}_u + \mathbf{b}^p$ from \mathcal{C}_u to \mathcal{C}_u^p , with $(\mathbf{A}^{-p}, \mathbf{b}^p)$

obtained from Eq. (58). The collective best-response is the concatenation of these solutions,

$$\tilde{\Phi}_u^* = ((0, \mathbf{A}^{-1}), \dots, (0, \mathbf{A}^{-N_P})) \tilde{\Phi}_u + (\mathbf{b}^1, \dots, \mathbf{b}^{N_P}).$$

Finally, note that converting this formula in terms of the original matrix-valued signal requires only a unitary transformation that preserves its Lipschitz properties. Thus, using Lemma 10,

$$L_{\widehat{BR}_\Phi}^2 = \sum_{p \in \mathcal{P}} \|(0, \mathbf{A}^{-p})\|_{2 \rightarrow 2}^2.$$

Now, consider that

$$\begin{aligned} \|(0, \mathbf{A}^{-p})\|_{2 \rightarrow 2} &\leq (1 + \kappa(\tilde{\mathbf{H}}^{pp}) \|\tilde{\mathbf{V}}^p\|_{2 \rightarrow 2}^2) \|(\tilde{\mathbf{H}}^{pp})^\dagger \tilde{\mathbf{H}}^{p,-p}\|_{2 \rightarrow 2} \\ &= (1 + \kappa(\tilde{\mathbf{H}}^{pp})) \|(\tilde{\mathbf{H}}^{pp})^\dagger \tilde{\mathbf{H}}^{p,-p}\|_{2 \rightarrow 2} \\ &= (1 + \kappa(\mathbf{H}^{pp})) \|(\mathbf{H}^{pp})^\dagger \mathbf{H}^{p,-p}\|_{2 \rightarrow 2} \\ &= L_{\widehat{BR}_\Phi}^p \end{aligned}$$

where for the first inequality we applied to Eq. (58) the triangle inequality and submultiplicative properties of operators norms, on the first equality we used the fact that $\|\tilde{\mathbf{V}}^p\| = 1$, and for the last equality we used the fact that $\|I_m \otimes \mathbf{Z}\| = \|\mathbf{Z}\|$ for any operator norm, $m > 0$, and operator \mathbf{Z} . Finally, since

$$L_{\widehat{BR}_\Phi}^2 \leq \sum_{p \in \mathcal{P}} (L_{\widehat{BR}_\Phi}^p)^2,$$

we can relax this Lipschitz constant for \widehat{BR}_Φ by considering this expression to hold as an equality. \square