Neuromorphic Realization of Best Response in Finite-Action Games

Himani Sinhmar[a], Vaibhav Srivastava[b], Naomi Ehrich Leonard[a]

***Abstract*— We develop a mechanistic dynamical-systems formulation of best response in finite-action games with relational structure on the action set. The proposed neuromorphic decision dynamics realize best response as the stable outcome of an internal state-space process, rather than as an externally imposed choice rule. This provides a deterministic account of commitment formation, symmetry resolution through basins of attraction, and hysteresis and decision persistence under perturbations. For action spaces with circulant coupling, we prove using Lyapunov-Schmidt reduction that the action-coupling operator determines which components of evidence govern decision formation. We further show that the dynamics implicitly compute a geometry-aware utility, converge exponentially to the corresponding best response with rate independent of the number of actions, and switch only when evidence is sufficiently strong. In contrast, supplying the same geometry-aware utility directly to logit dynamics does not recover these properties, showing that relational structure must be embedded in the decision mechanism itself. We illustrate the framework in a repeated coverage game, prove that the induced game is an exact potential game, and show that its Nash equilibria are reached by the neuromorphic dynamics.**

I. Introduction

Decision-making problems often require an agent to select an action from a finite set and to do so repeatedly as conditions evolve over time [1]. When there are multiple agents and factors such as congestion limitations, shared resources, or coordination constraints, actions are not isolated alternatives but rather carry relational structure. For example, when an agent selects a motion direction, nearby directions may yield similar motion outcomes. When choosing a sensing sector, adjacent sectors may overlap in coverage. When allocating effort across candidate tasks, tasks requiring overlapping capabilities may be partial substitutes. When committing to one among several resource-seeking actions, actions close in heading or location may compete for the same resource. Actions sharing downstream paths or bottlenecks may have coupled payoffs. When such relations matter, decision making should account not only for the utility of each action in isolation but also for how utility is distributed across related actions.

Research was sponsored by the Army Research Office and was accomplished under Grant Number W911NF2410126. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

[a] Department of Mechanical and Aerospace Engineering, Princeton University; {himani.sinhmar,naomi}@princeton.edu

[b] Department of Electrical and Compute Engineering, Michigan State University vaibhav@msu.edu

Classic best response selects the action with highest utility under the current utility profile [2]. Smooth variants such as logit replace this deterministic choice by a probabilistic response map [3], [4]. These models prescribe an action-selection rule, but they do not leverage relational structure among actions beyond including action coupling in utility construction, which is not robust. They also fall short in how they deal with time-varying, corrupted, noisy, and/or partial evidence [5]. Importantly, classic approaches do not provide a mechanism for how best response is realized, how nearties are resolved, or how commitment persists under transient perturbations.

To address these shortcomings, we propose and study a mechanistic dynamical system formulation of best response as an internal decision process. Our formulation uses neuromorphic decision dynamics as a state-space realization of best response in finite-action games with relational structure on the action set. Each agent maintains an internal decision state over actions. External evidence enters as an input and a mixed-sign coupling operator on the action set shapes how evidence is integrated during decision formation. Action relations may admit positive coupling between similar or overlapping actions and negative coupling between competitive or incompatible actions. The selected action is realized by the stable attractor reached by the internal dynamics instead of an externally imposed $\arg\max$ or by stochastic sampling [6], as in logit-type dynamics that converge more slowly to a distribution over actions corresponding to a quantal equilibrium [6]. Our decision dynamics go beyond classic best response by giving rich deterministic behavior including fast commitment through attractor structure, basin-dependent symmetry resolution, and, in the subcritical regime, hysteresis and persistence.

Our approach connects three lines of work. First, congestion and potential games provide the benchmark setting for repeated finite-action coordination and equilibrium convergence [7], [8]. Second, smooth and stochastic response models study equilibrium selection through perturbed or probabilistic choice [3], [4]. Third, recent nonlinear opinion-dynamics models show that saturating recurrent interactions can generate multistability, tunable sensitivity, fast decision formation, and game-theoretic behaviors [9]–[13]. *Our contribution is to connect these lines through a mechanistic dynamical systems perspective, showing how neuromorphic decision dynamics realize best response as the stable outcome of an internal state-space process and provide significant advantages to performance.*

We specialize in this paper to symmetric circulant mixed-sign coupling among actions. This provides a canonical

model of geometry-aware action spaces and admits an explicit spectral and bifurcation analysis. The dominant eigenspace determines which components of evidence govern equilibrium selection near onset, while noncritical components are attenuated through stable modes. The result is an action-structure-aware best-response mechanism and, in the subcritical regime, a hysteretic one.

Our **contributions** are as follows:

i) We present a mechanistic neuromorphic formulation of best response decision dynamics in finite-action games with relational action structure. Our method realizes action-selection as the stable outcome of an internal state-space process.
ii) We prove using Lyapunov–Schmidt reduction that a coupled action space governs the components of evidence that drive decision formation through the underlying coupling operator.
iii) We demonstrate that the neuromorphic dynamics formulation implicitly computes geometry-aware utility, exponentially converges to the best response, is independent of the number of actions, and in the subcritical regime, committed decisions persist under perturbations and switch only with sufficient evidence.
iv) We show that providing geometry-aware utility directly to logit dynamics is insufficient, emphasizing that the decision mechanism must encode relational structure underlying the action-set.
v) We instantiate our formulation in a repeated coverage game and prove that it is an exact potential game whose Nash equilibria are reached by neuromorphic dynamics.

Section II formulates the repeated game and the neuromorphic decision dynamics. Section III develops the bifurcation analysis and proves the realization of geometry-aware best response. Section IV studies the coverage game instantiation and proves convergence, optimality, and hysteresis. Section V presents conclusions.

II. Problem Formulation

A. Problem Setup

We consider a repeated multi-agent decision problem with agent set $\mathcal{N} = \{1, \dots, N\}$ and common finite action set $\mathcal{A} = \{1, \dots, K\}$. At each epoch, agent i selects $a_i \in \mathcal{A}$; the joint profile is $a \in \mathcal{A}^N$. The environment provides each agent i with an evidence vector $\boldsymbol{b}_i(a) \in \mathbb{R}^K$, where $b_{ik}(a)$ is the raw signal to agent i for action k, reflecting local observations about costs, rewards, or event densities.

We equip $\mathcal{A}$ with a symmetric *action-coupling operator* $A \in \mathbb{R}^{K\times K}$, where A_{kj} encodes geometry or relational affinity between actions k and j (e.g., spatial proximity, functional substitutability, shared resources). We define *utility* of agent i as the projection of evidence onto the dominant eigenspace of A: $U_i(a) = \Psi_A(\boldsymbol{b}_i(a))$, where $\Psi_A : \mathbb{R}^K \to \mathbb{R}^K$ aggregates evidence across actions weighted by coupling strength. $U_i \in \mathbb{R}^K$ is the effective utility in the game-theoretic sense [1]: it is the quantity whose maximization defines rational action selection when the action space carries relational structure. The *geometry-aware best response* for agent i is then

$$\mathrm{BR}_i^A(a_{-i}) = \arg\max_{k\in\mathcal{A}} \left[\Psi^A\big(\boldsymbol{b}_i(k, a_{-i})\big)\right]_k, \tag{1}$$

where $a_{-i} \in \mathcal{A}^{N-1}$ is the action profile of all agents except i. A profile $a^\star$ is a projected Nash equilibrium if $a_i^\star \in \mathrm{BR}_i^A(a_{-i}^\star)$ for all $i \in \mathcal{N}$. In the following, we propose neuromorphic dynamics that realize this best response.

B. Neuromorphic Decision Dynamics

We propose neuromorphic decision dynamics (NDD) for a cyclic action set $\mathcal{A}$ equipped with a circulant, symmetric action-coupling operator A. The action set is viewed as a discrete ring, and A defines a shift-invariant interaction pattern analogous to ring-attractor networks [12] whose spectral structure we exploit in Section III. Each agent i maintains an internal decision $\boldsymbol{z}_i(t) \in \mathbb{R}^K$, where each component z_i^k represents preference for action $k \in \mathcal{A}$ evolving according to

$$\tau_z \dot{\boldsymbol{z}}_i(t) = -d_z \boldsymbol{z}_i(t) + S\big(\alpha A \boldsymbol{z}_i(t) + \boldsymbol{b}_i\big), \tag{2}$$

where $\tau_z > 0$, $d_z > 0$, $\alpha > 0$, and $\boldsymbol{b}_i \in \mathbb{R}^K$ are the decision time constant, leakage rate, coupling gain, and evidence bias (Section II), respectively. The nonlinearity $S : \mathbb{R} \to \mathbb{R}$ acts componentwise and is odd, $\mathcal{C}^3$, sigmoidal, with $S'(0) = 1$ and $s_3 \triangleq S'''(0) \neq 0$. Action selection corresponds to a stable equilibrium of (2), i.e., a solution of

$$F(\boldsymbol{z}, \alpha, U) \triangleq -d_z \boldsymbol{z} + S(\alpha A \boldsymbol{z} + \boldsymbol{b}) = 0. \tag{3}$$

The gain α acts as the primary bifurcation parameter: as α increases, the trivial equilibrium $\boldsymbol{z} = 0$ loses stability and committed decisions emerge as stable nontrivial attractors [9], [14].

Circulant Action Coupling: We index $\mathcal{A}$ by equally spaced points on a discrete ring with sector angles $\theta_j = 2\pi j/K$, and take A to be real, symmetric, and circulant: $A_{jk} = \varphi(j - k \bmod K)$ for a symmetric profile φ. The canonical choice is a Mexican-hat profile with local excitation and surround, for example, $\texttt{toeplitz}([+, +, 0, -, \dots, -, 0, +])$. By standard results [15], A is diagonalized by the real discrete Fourier basis: $\boldsymbol{\phi}_k \triangleq \{\cos(k\theta_j)\}_{j=0}^{K-1}$, $\boldsymbol{\psi}_k \triangleq \{\sin(k\theta_j)\}_{j=0}^{K-1}$, $k = 0, \dots, \lfloor K/2 \rfloor$, with $A\boldsymbol{\phi}_k = \lambda_k \boldsymbol{\phi}_k$, $A\boldsymbol{\psi}_k = \lambda_k \boldsymbol{\psi}_k$, real eigenvalues satisfying $\lambda_k = \lambda_{K-k}$, and orthonormality $\langle \boldsymbol{\phi}_k, \boldsymbol{\phi}_k \rangle = \langle \boldsymbol{\psi}_k, \boldsymbol{\psi}_k \rangle = \frac{1}{2}$, $\langle \boldsymbol{\phi}_k, \boldsymbol{\psi}_k \rangle = 0$ under $\langle \boldsymbol{u}, \boldsymbol{v} \rangle \triangleq \frac{1}{K}\sum_j u_j v_j$. The linearization of (3) at $\boldsymbol{z} = 0$ is $L(\alpha) = -d_z I + \alpha A$, sharing eigenvectors with A and scalar gain $-d_z + \alpha\lambda_k$ on eigenmode k. Thus the trivial equilibrium loses stability eigenmode-by-eigenmode as α increases, first along the eigenmode with the largest eigenvalue.

Assumption 1 (Unique dominant mode): There exists a unique index $k_\star$ such that $\lambda_{k_\star} > \lambda_k$ for all $k \neq k_\star$.

Under Assumption 1, the critical gain is $\alpha_c = d_z/\lambda_{k_\star}$, and instability first occurs along the two-dimensional subspace $E_c = \mathrm{span}\{\boldsymbol{\phi}_{k^\star}, \boldsymbol{\psi}_{k^\star}\}$ where $\mu = -d_z + \alpha\lambda_{k_\star}$ measures distance from bifurcation.

Assumption 2 (Local near-bifurcation regime): We focus on a neighborhood of the critical gain $\alpha_c = \frac{d_z}{\lambda_{k^\star}}$. assuming that $|\mu|$, $\|\boldsymbol{b}\|$, and $\|\boldsymbol{z}\|$ are sufficiently small, equivalently, that $|\alpha - \alpha_c|$ is small and the equilibrium under consideration lies in a sufficiently small neighborhood of $\boldsymbol{z} = 0$.

III. Analysis of Neuromorphic Decision Dynamics

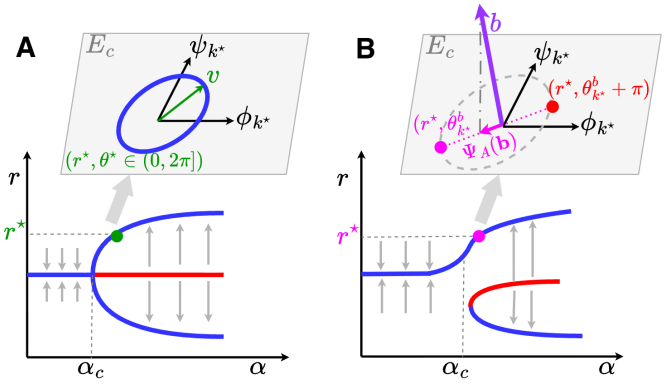

Fig. 1. **Bifurcation geometry and action-selection in the critical eigenspace** $E_c = \mathrm{span}\{\boldsymbol{\phi}_{k^\star}, \boldsymbol{\psi}_{k^\star}\}$**. (A)** At the onset $\alpha = \alpha_c$, a supercritical pitchfork generates a continuum of equilibria parameterized by $(r^\star, \theta^\star)$, forming a ring in E_c (blue) of phase-indeterminate committed states along the stable branch $r = r^\star$. **(B)** In the presence of an input $\boldsymbol{b}$ (purple), its projection $\Psi_A(\boldsymbol{b})$ onto E_c (magenta) selects a unique equilibrium on the ring. The intersection determines the phase $\theta^b_{k^\star}$ and amplitude $r^\star$, yielding the committed state $(r^\star, \theta^b_{k^\star})$.

As α increases through α_c, $\boldsymbol{z} = 0$ loses stability, and a family of committed decision states emerges (Figure 1). Under Assumptions 1– 2, we formalize this transition by reducing the full K-dimensional equilibrium equation to a two-dimensional problem on E_c using Lyapunov–Schmidt reduction [16].

A. Low-dimensional reduction

At $\alpha = \alpha_c$, the linearization $L(\alpha_c)$ is singular: its kernel is E_c and, since $L(\alpha_c)$ is self-adjoint with respect to $\langle\cdot,\cdot\rangle$, its range is $E_s = E_c^\perp$. Standard implicit function theorem arguments therefore fail at this point, and a reduction is required to characterize equilibria near the onset of decision formation. Write $\boldsymbol{z} = \boldsymbol{v} + \boldsymbol{w}$ with $\boldsymbol{v} \in E_c$ and $\boldsymbol{w} \in E_s$. Projecting (3) onto E_s and E_c yields the coupled system $\mathcal{G}(\boldsymbol{v}, \boldsymbol{w}, \alpha, \boldsymbol{b}) := \Psi_A F(\boldsymbol{v} + \boldsymbol{w}, \alpha, \boldsymbol{b})$ and $\mathcal{H}(\boldsymbol{v}, \boldsymbol{w}, \alpha, \boldsymbol{b}) := \Psi_A^\perp F(\boldsymbol{v} + \boldsymbol{w}, \alpha)$. Since all noncritical eigenvalues of $L(\alpha_c)$ are strictly negative, $D_w \mathcal{H}(0, 0, \alpha_c, 0) = \Psi_A^\perp L(\alpha_c)$ is invertible. By the implicit function theorem, there exists a unique $\mathcal{C}^3$ reduction map $W : E_c \times \mathbb{R} \times \mathbb{R}^K \to E_s$ solving $\mathcal{H} = 0$ identically, with $W(0, \alpha_c, 0) = 0$ and $D_v W(0, \alpha_c, 0) = 0$. The latter follows because $L(\alpha_c)\delta\boldsymbol{v} \in E_c$ for any $\delta\boldsymbol{v} \in E_c$, which is annihilated by $\Psi_A^\perp$, so variations in $\boldsymbol{v}$ do not linearly force the range equation. The odd symmetry of S eliminates quadratic terms, giving $W(\boldsymbol{v}, \alpha, \boldsymbol{b}) = \mathcal{O}(\|\boldsymbol{b}\|) + \mathcal{O}(\|\boldsymbol{v}\|^3)$: the stable modes are uniquely determined (locally) to the cubic and linear-in-evidence corrections of the critical mode. Substituting $\boldsymbol{w} = W(\boldsymbol{v}, \alpha, \boldsymbol{b})$ into the critical equation yields the reduced condition $\mathcal{G}(\boldsymbol{v}, \alpha, \boldsymbol{b}) = 0$ posed entirely on E_c.

Lemma 1 (Reduced critical equation): Under Assumptions 1 and 2

$$\mathcal{G}(\boldsymbol{v}, \mu, \boldsymbol{b}) = \mu\boldsymbol{v} + \Psi_A(\boldsymbol{b}) + \frac{s_3}{6}\alpha_c^3\,\Psi_A\big((A\boldsymbol{v})^{\circ 3}\big) + \frac{s_3}{2}\alpha_c^2\,\Psi_A\big((A\boldsymbol{v})^{\circ 2} \circ \boldsymbol{b}\big) + \mathcal{R}, \tag{4}$$

where $\mathcal{R} = \mathcal{O}(\|\boldsymbol{v}\|^5) + \mathcal{O}(\|\boldsymbol{b}\|\|\boldsymbol{v}\|^2) + \mathcal{O}(\|\boldsymbol{b}\|^2\|\boldsymbol{v}\|) + \mathcal{O}(\|\boldsymbol{b}\|^3)$. The linear term $\mu\boldsymbol{v}$ controls onset of instability, the cubic self-interaction saturates decision amplitude, and the mixed cubic term captures how evidence $\boldsymbol{b}$ biases the emerging decision through the coupling structure A.

Proof: Expand F using $S(u) = u + \frac{s_3}{6}u^{\circ 3} + \mathcal{O}(\|u\|^5)$ and project onto E_c: $\mathcal{G} = \Psi_A L(\alpha)(\boldsymbol{v} + \boldsymbol{w}) + \Psi_A(\boldsymbol{b}) + \frac{s_3}{6}\Psi_A\big(\alpha A(\boldsymbol{v} + \boldsymbol{w}) + \boldsymbol{b}\big)^{\circ 3} + \mathcal{O}_{\geq 5}$ For the linear terms: $\boldsymbol{v} \in E_c$ is an eigenvector of A with eigenvalue $\lambda_{k_\star}$, so $\Psi_A L(\alpha)\boldsymbol{v} = \mu\boldsymbol{v}$. Since E_s is invariant under A for any symmetric circulant matrix [15] (each real Fourier plane is an eigenspace), $A\boldsymbol{w} \in E_s$ and thus $\Psi_A L(\alpha)\boldsymbol{w} = 0$. The linear contribution is exactly $\mu\boldsymbol{v} + \Psi_A(\boldsymbol{b})$. For the cubic term, expand at $(\alpha_c, 0)$ and retain the displayed leading orders: $\big(\alpha_c A(\boldsymbol{v} + \boldsymbol{w}) + \boldsymbol{b}\big)^{\circ 3} = (\alpha_c A\boldsymbol{v})^{\circ 3} + 3(\alpha_c A\boldsymbol{v})^{\circ 2} \circ \boldsymbol{b}$. Using $\|\boldsymbol{w}\| = \mathcal{O}(\|\boldsymbol{v}\|^3) + \mathcal{O}(\|\boldsymbol{b}\|)$, the remainder terms satisfy $(\alpha_c A\boldsymbol{v})^{\circ 2} \circ A\boldsymbol{w} = \mathcal{O}(\|\boldsymbol{v}\|^5) + \mathcal{O}(\|\boldsymbol{b}\|\|\boldsymbol{v}\|^2)$, $(\alpha_c A\boldsymbol{v}) \circ (\alpha_c A\boldsymbol{w} + \boldsymbol{b})^{\circ 2} = \mathcal{O}(\|\boldsymbol{b}\|^2\|\boldsymbol{v}\|) + \text{h.o.t.}$, and $(\alpha_c A\boldsymbol{w} + \boldsymbol{b})^{\circ 3} = \mathcal{O}(\|\boldsymbol{b}\|^3) + \text{h.o.t.}$ The projected cubic expansion contains ths displayed terms $\Psi_A((A\boldsymbol{v})^{\circ 3})$ and $\Psi_A((A\boldsymbol{v})^{\circ 2} \circ \boldsymbol{b})$, while the remaining contributions, including those involving $\boldsymbol{w}$, are absorbed into $\mathcal{R}$. Collecting terms yields (4). ■

B. Polar Reduced Equations

Any $\boldsymbol{v} \in E_c$ admits a polar representation $\boldsymbol{v} = re_\theta$ with $r \geq 0$ and $e_\theta = \sqrt{2}(\cos\theta\,\boldsymbol{\phi}_{k^\star} + \sin\theta\,\boldsymbol{\psi}_{k^\star})$. The orthonormal frame $\{e_\theta, e_{\theta+\pi/2}\}$ on E_c yields two scalar equations by projecting (4) onto the radial and angular directions. The evidence $\boldsymbol{b} \in \mathbb{R}^K$ decomposes over the Fourier basis as $b_j = a_0 + \sum_m (a_m \cos(m\theta_j) + \beta_m \sin(m\theta_j))$, with amplitude $A^b_m = \sqrt{a_m^2 + \beta_m^2}$ and phase $\theta^b_m = \mathrm{atan2}(\beta_m, a_m)$ at mode m, for $(a_m, \beta_m) \neq (0, 0)$.

Proposition 1 (Local polar reduced equations): With $\Gamma = \frac{s_3}{4}\alpha_c^3\lambda_{k_\star}^3$ and $\eta = \frac{s_3}{4}\alpha_c^2\lambda_{k_\star}^2$, the reduced equilibrium conditions $g_r = 0$ and $g_\theta = 0$ are, up to cubic order,

$$g_r = \mu r + \Gamma r^3 + \frac{A^b_{k_\star}}{\sqrt{2}}\cos(\theta - \theta^b_{k_\star}) + \frac{\eta r^2}{\sqrt{2}}\Big[3A^b_{k_\star}\cos(\theta - \theta^b_{k_\star}) + A^b_{3k_\star}\cos(3\theta - \theta^b_{3k_\star})\Big] \tag{5}$$

$$g_\theta = \mathbf{1}(b \neq 0)\Big(\frac{A^b_{k_\star}}{\sqrt{2}}\sin(\theta^b_{k_\star} - \theta) + \frac{\eta r^2}{\sqrt{2}}\Big[3A^b_{k_\star}\sin(\theta^b_{k_\star} - \theta) + A^b_{3k_\star}\sin(\theta^b_{3k_\star} - 3\theta)\Big]\Big). \tag{6}$$

Proof: Substitute $v = re_\theta$ into the reduced equation (4) from Lemma 1. Since $\boldsymbol{e}_\theta$ is an eigenvector of A with eigenvalue $\lambda_{k_\star}$, we have $A\boldsymbol{v} = \lambda_{k_\star} r\boldsymbol{e}_\theta$, and thus $(A\boldsymbol{v})^{\circ 3} = \lambda_{k_\star}^3 r^3\,\boldsymbol{e}_\theta^{\circ 3}$ and $(A\boldsymbol{v})^{\circ 2} \circ \boldsymbol{b} = \lambda_{k_\star}^2 r^2\,\boldsymbol{e}_\theta^{\circ 2} \circ \boldsymbol{b}$. For the cubic self-interaction, the identity $\Psi_A(\boldsymbol{e}_\theta^{\circ 3}) = \frac{3}{2\sqrt{2}}\boldsymbol{e}_\theta$ holds for circulant A [15], yielding the coefficient $\Gamma = \frac{s_3}{6}\alpha_c^3\lambda_{k_\star}^3 \cdot \frac{3}{2} =$

$\frac{s_3}{4}\alpha_c^3\lambda_{k_\star}^3$. For the mixed cubic term, decompose $\boldsymbol{e}_\theta^{\circ 2} \circ \boldsymbol{b}$ in the Fourier basis. Using $\boldsymbol{e}_\theta^{\circ 2} = \frac{1}{\sqrt{2}}(\Phi_0 + \cos 2\theta\, \Phi_{2k_\star} + \sin 2\theta\, \Psi_{2k_\star})$, the convolution with $\boldsymbol{b}$ projects onto modes $k_\star$ and $3k_\star$ under Ψ_A: $\Psi_A(\boldsymbol{e}_\theta^{\circ 2} \circ \boldsymbol{b}) = \frac{1}{\sqrt{2}}\Big[A^b_{k_\star}\cos(\theta - \theta^b_{k_\star})\,\boldsymbol{e}_\theta + A^b_{3k_\star}\cos(3\theta - \theta^b_{3k_\star})\,\boldsymbol{e}_\theta + \text{angular terms}\Big]$, with coefficient $\eta = \frac{s_3}{4}\alpha_c^2\lambda_{k_\star}^2$. Similarly, projecting $\mathcal{G} = 0$ onto $\boldsymbol{e}_\theta$ (radial) and $\boldsymbol{e}_{\theta+\pi/2}$ (angular) yields (5)–(6) directly, while the remaining contributions, including those induced through the reduction map $\boldsymbol{w}$ are absorbed into the stated remainder terms in (4). ∎

C. Supercritical Bifurcation and Decision Formation

Unforced case ($\boldsymbol{b} = 0$): The polar equations reduce to $g_r = \mu r + \Gamma r^3$ and $g_\theta \equiv 0$. For $\Gamma < 0$, a local supercritical branch $r^\star = \sqrt{-\mu/\Gamma}$ emerges for $\mu > 0$ within the regime of Assumption 2. The equilibrium direction $\theta^\star$ is indeterminate: the entire ring $\{r^\star \boldsymbol{e}_\theta : \theta \in [0, 2\pi)\}$ consists of neutrally stable equilibria, and the system commits to whichever $\theta^\star$ is selected by initial conditions (Figure 1). At $(r^\star, \theta^\star)$, the Jacobian of the reduced equilibrium map satisfies $\partial_r g_r(r^\star) = \mu + 3\Gamma(r^\star)^2 = -2\mu < 0$, and $\partial_\theta g_\theta(r^\star, \theta^\star) = 0$. Thus the branch is locally stable in amplitude, while the 0 angular eigenvalue reflects neutral phase degeneracy due to rotational symmetry.

Lemma 2 (Suppression of noncritical evidence): Let $\boldsymbol{z} = \boldsymbol{v} + \boldsymbol{w}$ with $\boldsymbol{v} \in E_c$ and $\boldsymbol{w} \in E_s$ be the Lyapunov–Schmidt decomposition. For each noncritical Fourier mode $m \neq k_\star$, the corresponding coefficient of $\boldsymbol{w}$ satisfies $w_m = -\frac{b_m}{\sigma_m} + \text{h.o.t.}$, where $\sigma_m := -d_z + \alpha_c\lambda_m < 0$ is the stable eigenvalue of $L(\alpha_c)$ at mode m. Hence, if

$$|\sigma_m| \gg A^b_m, \qquad \forall\, m \neq k_\star, \tag{7}$$

noncritical harmonics enter the reduced critical equation only through higher-order corrections in $\boldsymbol{w}$ and do not affect phase selection at leading order.

Proof: The range equation is $\mathcal{H}(\boldsymbol{v}, \boldsymbol{w}, \alpha, \boldsymbol{b}) = \Psi_A^\perp F(\boldsymbol{v}+\boldsymbol{w}, \alpha, \boldsymbol{b}) = 0$. Since $D_w\mathcal{H}(0, 0, \alpha_c, 0) = \Psi_A^\perp L(\alpha_c)$ is invertible on E_s, the implicit function theorem gives $\boldsymbol{w} = -\big(L(\alpha_c)\big)^{-1}\big(\Psi_A^\perp \boldsymbol{b} + \text{h.o.t.}\big)$. In the Fourier basis of the circulant operator A, this yields $w_m = -b_m/\sigma_m + \text{h.o.t.}$ for each $m \neq k_\star$. Therefore, under (7), these modes are attenuated by the stable spectrum and contribute only higher-order corrections to the reduced angular equation. ∎

Forced case ($\boldsymbol{b} \neq 0$): Evidence breaks rotational symmetry and by Proposition 1, $g_\theta = \frac{A^b_{k_\star}}{\sqrt{2}}\sin(\theta^b_{k_\star} - \theta) + \frac{nr^2}{\sqrt{2}}\Big[3A^b_{k_\star}\sin(\theta^b_{k_\star} - \theta) + A^b_{3k_\star}\sin(\theta^b_{3k_\star} - 3\theta)\Big]$. By Lemma 2, harmonics $m \neq k_\star$ are filtered through the stable component w and do not enter the phase equation at leading order. The only leading correction on E_c is the $3k_\star$ harmonic generated by cubic mixing. If $3A^b_{k_\star} > A^b_{3k_\star}$, the critical harmonic dominates near onset, so $g_\theta = 0$ has, to leading order, the two solutions $\theta^\star = \theta^b_{k_\star}$ and $\theta^\star = \theta^b_{k_\star} + \pi$. At these two phases, the Jacobian of the reduced equilibrium map satisfies $\partial_\theta g_\theta(\theta^b_{k_\star}) = -\frac{A^b_{k_\star}}{\sqrt{2}} + \mathcal{O}(r^2) < 0$, and $\partial_\theta g_\theta(\theta^b_{k_\star} + \pi) = \frac{A^b_{k_\star}}{\sqrt{2}} + \mathcal{O}(r^2) > 0$, for sufficiently small r. Thus the aligned phase is locally stable, the anti-aligned phase is unstable, and forcing pins the selected phase to the phase of the projected evidence $\Psi_A(\boldsymbol{b})$ as shown in Figure 1.

D. Subcritical Bifurcation and Hysteretic Best Response

When $\Gamma > 0$, the cubic term destabilizes the bifurcating branch. To recover a bounded subcritical bifurcation, we introduce a state-dependent gain $\alpha = \alpha_0 + \kappa\|z\|^2$, giving $\mu = \mu_0 + \kappa\lambda_{k_\star} r^2$, where $\mu_0 = -d_z + \alpha_0\lambda_{k_\star}$ and $\|z\|^2 = r^2 + \mathcal{O}(r^6)$ Substituting into the radial equation gives $g_r = \mu_0 r + \Psi_A(\boldsymbol{b}) + \tilde{\Gamma} r^3 + \Delta r^5 = 0$, where $\tilde{\Gamma} = \Gamma + \kappa\lambda_{k_\star}$ is the effective cubic coefficient and $\Delta < 0$ bounds the subcritical branch via the next-order Taylor expansion of S. Subcriticality requires $\tilde{\Gamma} > 0$, or equivalently $\kappa > -\frac{\Gamma}{\lambda_{k_\star}}$. The nontrivial branch then bends backward, stabilized by the negative quintic term, producing a saddle-node bifurcation at $\mu_{\text{SN}} < \mu_0$. For $\mu_0 < \mu_{\text{SN}}$, only $r = 0$ is stable. For $\mu_0 \in (\mu_{\text{SN}}, 0)$, the system is bistable: the trivial and a large-amplitude committed state are both stable, separated by an unstable branch. At $\mu_0 = 0$, the large-amplitude branch persists, so committed decisions survive below the critical gain. This bistable region produces *hysteretic best response*. As μ_0 increases through zero, the decision state jumps to the large-amplitude branch, and as μ_0 decreases, commitment persists until μ_{SN}. The system therefore resists oscillatory switching — once committed, it does not reverse unless projected evidence shifts enough to cross the unstable branch. Since $g_\theta = 0$ continues to select $\theta^\star = \theta^b_{k_\star}$ when $\Psi_A(\boldsymbol{b}) \neq 0$, the subcritical structure modifies only amplitude persistence, not the identity of the selected action.

E. Realization of Best Response

The bifurcation analysis reveals that only the projection of $\boldsymbol{b}$ onto E_c enters the reduced equations at leading order, $\Psi_A(\boldsymbol{b}) = A^b_{k_\star}\big(\cos\theta^b_{k_\star}\,\boldsymbol{\phi}_{k_\star} + \sin\theta^b_{k_\star}\,\boldsymbol{\psi}_{k_\star}\big)$, which is the utility $U_i = \Psi_A(\boldsymbol{b}_i)$ of Section II: a weighted sum of evidence along the directions $\{\boldsymbol{\phi}_{k_\star}, \boldsymbol{\psi}_{k_\star}\}$, fully characterizing how evidence biases the emergent decision. Since $\theta^\star = \theta^b_{k_\star}$, the selected action $a^\star = \arg\max_k U_{ik}$ is BR_i^A of (1). Hence, under Assumption 2, and provided $3A^b_{k^\star} > A^b_{3k^\star}$ together with (7), the action selected by the locally stable equilibrium of (6) coincides, near onset, with the projected best response BR_i^A (1). For $k_\star = 1$, $\Psi_A(\boldsymbol{b}_i)$ is unimodal, so the induced choice is winner-take-all.

This perspective also suggests an inverse design problem: given a desired projection of evidence onto a low-dimensional subspace, synthesize an action-coupling operator A whose dominant eigenstructure realizes that projection. With graph constraints, this becomes a structured eigenstructure assignment problem. In symmetric settings (ring), the admissible eigenvectors are fixed by symmetry, and design acts through eigenvalue placement among Fourier modes. Finally, while the evidence lies in $\mathbb{R}^K$, the effective feature dimension is determined by the dominant eigenspace of A. In this sense, the proposed dynamics combine linear function

approximation with a mechanistic dynamical-systems realization of best response.

IV. A Coverage Game Analysis

In the proposed framework, the projected utility $U_{ik} = [\Psi_A(\boldsymbol{b}_i(t))]_k$ is a linear projection of the evidence field, with the projection determined by the dominant eigenspace of the action-coupling operator A. Thus, A specifies which evidence components become decision-relevant in realizing the best response. The coverage game is a natural example because the action utility is a linear functional of the underlying evidence, but the same structure also arises in other finite-action problems with relational action structure, including correlated task allocation, discretized heading selection, and structured hypothesis selection.

We consider a discrete ring of K sectors $\mathcal{A} = \{1, \dots, K\}$ with event density $V : \mathcal{A} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, $\sum_k V(k,t) = 1$, varying on a slow time scale $\tau_b \gg \tau_z$ relative to (2). Each agent observes evidence $b_{ik}(t) = V(k,t) - \rho \sum_{j \neq i} \mathbf{1}[a_j(t) = k]$, $\rho > 0$, balancing event density against congestion [17]. For any symmetric circulant A satisfying Assumption 1, the local bifurcation analysis of Section III implies, under Assumption 2, that the dynamics (2) implicitly computes the effective utility $U_{ik}(a_{-i}, t) = [\Psi_A(\boldsymbol{b}_i(t))]_k = \sum_{s \in \mathcal{A}} C_{ks}\, b_{is}(t)$, where $C_{ks} := \cos\big(k^\star(\theta_s - \theta_k)\big)$ are the entries of the symmetric interaction kernel C induced by the dominant mode, and the locally stable equilibrium of NDD selects the corresponding projected best response near bifurcation onset. These utilities are the marginal increments of the welfare function (8), so the game is an exact potential game (Lemma 3) whose maximizers are Nash equilibria reached exponentially fast at rate $\min(1, 2|\mu|)$, independent of K (Proposition 2). Under subcritical bifurcation, amplitude and phase hysteresis ensures robustness to perturbations below $\delta^\star$ (Proposition 3).

Lemma 3 (Exact potential for the projected game): Fix t. Define $\bar{V}_k(t) := \sum_{s \in \mathcal{A}} C_{ks} V(s,t)$. Then, the game with the utilities U_{ik} is an exact potential game with potential

$$W(a,t) := \sum_{i=1}^{N} \bar{V}_{a_i}(t) - \frac{\rho}{2} \sum_{i \neq j} C_{a_i, a_j} \tag{8}$$

Consequently, $a^\star$ is a projected Nash equilibrium if and only if it is a local maximizer of $W(\cdot, t)$.

Proof: By definition, $U_{ik}(a_{-i}, t) = \sum_{s \in \mathcal{A}} C_{ks} V(s,t) - \rho \sum_{j \neq i} \sum_{s \in \mathcal{A}} C_{ks} \mathbf{1}[a_j = s] = \bar{V}_k(t) - \rho \sum_{j \neq i} C_{k, a_j}$. Consider a unilateral deviation of agent i from a_i to a_i'. Since C is symmetric, $\Delta W := W(a_i', a_{-i}, t) - W(a_i, a_{-i}, t) = \bar{V}_{a_i'}(t) - \bar{V}_{a_i}(t) - \rho \sum_{j \neq i} \big(C_{a_i', a_j} - C_{a_i, a_j}\big) = U_i(a_i', a_{-i}, t) - U_i(a_i, a_{-i}, t)$. Therefore W is an exact potential for the projected game. The equivalence between projected Nash equilibria and local maximizers of W is the standard characterization for finite exact potential games [8]. ∎

Proposition 2 (Local convergence to projected BR): Let $\Gamma < 0$, $A^b_{k_\star} \neq 0$, and and suppose (7) together with Assumptions 1–2 hold. Then, for almost any initial condition, the dynamics (2) converge locally exponentially at rate $\min(1, 2|\mu|)$ to a geometry-aware best response $\mathrm{BR}^A(t)$ of the coverage game at each t.

Proof: Since $\tau_b \gg \tau_z$, $\boldsymbol{b}_i(t)$ is quasi-static and (2) tracks the moving attractor adiabatically. At each t, we establish convergence to a best response in amplitude and angle separately, with overall rate $\min(1, 2|\mu|)$: $\alpha > \alpha_c$ and $\Gamma < 0$ give stable branch $r^\star = \sqrt{-\mu/\Gamma}$ with $\partial_r g_r(r^\star) = -2|\mu| < 0$ (Section III-C), so the radial component converges exponentially at rate $2|\mu|$. $A^b_{k_\star} \neq 0$ gives a unique stable angular solution $\theta^\star = \theta^b_{k_\star}(t)$ for any $\theta(0)$, since $\partial_\theta g_\theta(\theta^b_{k_\star}) = -1$ globally at leading order, yielding exponential convergence at rate 1. By Section III-E and Lemma 2, the attractor $(r^\star, \theta^\star)$ selects $a_i^\star = \mathrm{BR}_i^A(t)$, the geometry-aware best response of agent i at each t. The overall convergence rate is $\min(1, 2|\mu|)$, determined by the slower of the two modes. ∎

Proposition 3 (Hysteresis in commitment and phase): Assume the state-dependent-gain extension of Section III-D and suppose the effective cubic coefficient satisfies $\tilde{\Gamma} > 0$, so that the reduced radial dynamics possess a local bistable region for $\mu_0 \in (\mu_{SN}, 0)$. Under Assumptions 1 and 2, there exists a threshold $\delta^\star > 0$ such that (i) *(Radial hysteresis)* If $\|\Psi_A(\delta \boldsymbol{b}_i)\| < \delta^\star$, transient fluctuations $\delta \boldsymbol{b}_i(t)$ do not destroy commitment, i.e., $r(t)$ remains near $r^\star$. (ii) *(Phase hysteresis)* If $A^b_{k_\star} |\cos(\Delta\theta)| < \delta^\star$, where $\Delta\theta$ is the angular shift in $\theta^b_{k_\star}(t)$, the committed phase $\theta^\star$ does not follow the new $\theta^b_{k_\star}$ and the selected sector is unchanged. Above $\delta^\star$, the agent switches deterministically to BR_i^A at the new $\theta^b_{k_\star}$.

Proof: By Section III-D, when $\mu_0 \in (\mu_{\mathrm{SN}}, 0)$ the reduced radial dynamics have an unstable branch at $r_u = \sqrt{-\mu_0/\tilde{\Gamma}}$, which locally bounds the basin of attraction of the committed equilibrium on E_c. A perturbation $\delta \boldsymbol{b}_i$ changes the projected forcing along the committed direction by an amount proportional to $A^b_{k_\star} |\cos(\Delta\theta)|$. If this change is too small to cross the unstable branch, the trajectory stays in the same basin and returns to the committed equilibrium. If it is large enough to cross the boundary, the trajectory leaves that basin and relaxes to the branch selected by the perturbed phase, using the local phase stability established in Section III-C. ∎

A. Numerical Experiment

We simulate a repeated coverage game with $N = 10$ agents on a ring of $K = 18$ sectors over $T = 700$ steps. The coupling matrix A is sampled from a symmetric circulant Mexican-hat kernel (Figure 2(B)), with dominant eigenmode $k_\star = 1$ (Assumption 1). This encodes the goal of positioning each agent at the event density center of mass of its neighborhood. The event density $V(k,t)$ has three peaks of unequal amplitude concentrated near sectors $\{3,\ 8\text{–}10,\ 14\text{–}16\}$ (Figure 2(A)), with transient ambiguity at $t \approx 280$ and a directional shift at $t \approx 430$, on timescale $\tau_b \gg \tau_z$. The global objective is to maximize weighted coverage $W(a,t)$ (Lemma 3) whose local maximizers are under Nash equilibria. Figure 3 plots the agent trajectories

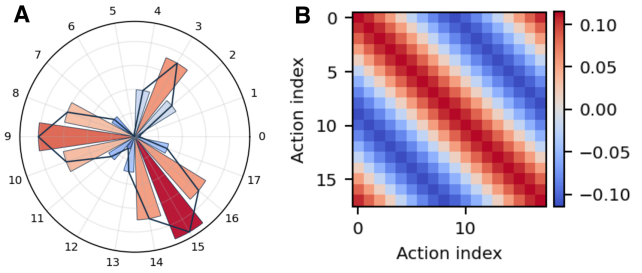

Fig. 2. (A) Raw event density $V(k, 0)$ on the action ring. (B) Mexican-hat action-coupling matrix A with dominant eigenmode $k_\star = 1$.

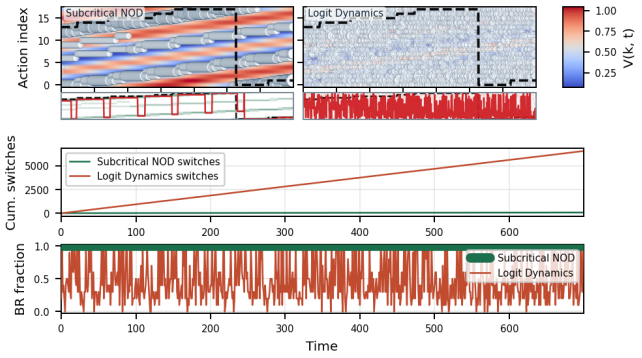

Fig. 3. Subcritical NDD vs. logit dynamics on a coverage task. *Top:* Heatmap shows event density $V(k, t)$; white circles are agents scaled by population fraction at each action. The strip below each heatmap shows the modal action (red, largest population fraction) and effective utility (black). Neuromorphic agents commit to dominant peaks and track them stably; logit agents scatter without committing. *Middle:* Cumulative switches over 700 steps: $\mathcal{O}(10)$ for NDD vs. $\mathcal{O}(5000)$ for logit. *Bottom:* NDD maintains BR fraction = 1 throughout; logit fluctuates erratically.

(white circles), the cumulative switches, and the BR fraction (fraction of agents simultaneously playing $\mathrm{BR}_i^A(t)$ at a given step) for both NDD and logit dynamics. Subcritical NDD ($\alpha = 0.96 < \alpha_c = 1$) is compared against logit dynamics which is supplied $\Psi_A(\boldsymbol{b}_i)$ directly. NDD recovers $\Psi_A(\boldsymbol{b}_i)$ from raw evidence $\boldsymbol{b}_i$ alone. Figure 3 shows that NDD outperforms logit across all four metrics. NDD agents spontaneously spread into three stable clusters around the dominant peaks of $V(\cdot, t)$, maximizing $W(a, t)$ bidirectionally, while logit agents disperse incoherently across all 18 sectors. NDD maintains BR fraction $= 1.0$ for all 700 steps (Proposition 2), whereas logit fluctuates between 0 and 1 and never stabilizes. Subcritical hysteresis (Proposition 3) renders NDD immune to the ambiguity window near $t \approx 280$ which makes the logit dynamics BR fraction collapse repeatedly. At $t \approx 430$, NDD executes a single deterministic reallocation when the directional shift exceeds $\delta^\star$ and immediately recommits, accumulating $O(10)$ total switches over 700 steps versus $O(5000)$ for logit.

V. Conclusion

We developed a neuromorphic realization of best response for finite-action games with relational action structure. Unlike classical best response rules, which specify only the selected action, or logit models, which specify a response distribution, our approach models decision making as a nonlinear dynamical process whose stable attractors realize committed actions. This mechanistic viewpoint explains how commitment forms, how near-symmetric alternatives are resolved through basins of attraction, and why decisions persist under transient perturbations. For symmetric circulant coupling, the coupling operator selects the evidence components governing decision formation, yielding a geometry-aware best response with exponential convergence, perturbation robustness, and hysteretic commitment in the subcritical regime. The coverage game instantiation showed that logit dynamics is prone to oscillatory switching even with geometry-aware utility, whereas NOD avoids oscillations and recovers Nash equilibria. These properties are especially relevant for resource-constrained robotic systems, where coordination must be achieved with limited computation, noisy observations, and little tolerance for indecisive switching.

VI. Acknowledgments

We sincerely thank Giovanna Amorim, Ian Xul Belaustegui, and Anastasia Bizyaeva for helpful discussions during the development of this research.

References

[1] D. Paccagnan, R. Chandan, and J. R. Marden, "Utility and mechanism design in multi-agent systems: An overview," *Annual Reviews in Control*, vol. 53, pp. 315–328, 2022.
[2] 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.
[3] R. D. McKelvey and T. R. Palfrey, "Quantal response equilibria for normal form games," *Games and Economic Behavior*, vol. 10, no. 1, pp. 6–38, 1995.
[4] L. E. Blume, "The statistical mechanics of strategic interaction," *Games and Economic Behavior*, vol. 5, no. 3, pp. 387–424, 1993.
[5] Z. Wang, Y. Shen, M. M. Zavlanos, and K. H. Johansson, "Convergence analysis of the best response algorithm for time-varying games," in *IEEE Conference on Decision and Control*, 2023, pp. 1144–1149.
[6] V. Auletta, D. Ferraioli, F. Pasquale, P. Penna, and G. Persiano, "Convergence to equilibrium of logit dynamics for strategic games," in *ACM Symposium on Parallelism in Algorithms and Architectures*, 2011, pp. 197–206.
[7] R. W. Rosenthal, "A class of games possessing pure-strategy nash equilibria," *International Journal of Game Theory*, vol. 2, no. 1, pp. 65–67, 1973.
[8] D. Monderer and L. S. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, no. 1, pp. 124–143, 1996.
[9] A. Bizyaeva, A. Franci, and N. E. Leonard, "Nonlinear opinion dynamics with tunable sensitivity," *IEEE Transactions on Automatic Control*, vol. 68, no. 3, pp. 1415–1430, 2023.
[10] S. Park, A. Bizyaeva, M. Kawakatsu, A. Franci, and N. E. Leonard, "Tuning cooperative behavior in games with nonlinear opinion dynamics," *IEEE Control Systems Letters*, vol. 6, pp. 2030–2035, 2022.
[11] R. Moreno-Morton, A. Bizyaeva, N. E. Leonard, and A. Franci, "Fast-and-flexible decision-making with modulatory interactions," *IEEE Control Systems Letters*, vol. 8, pp. 3333–3338, 2024.
[12] G. Amorim, A. Bizyaeva, A. Franci, and N. E. Leonard, "Spatially-invariant opinion dynamics on the circle," *IEEE Control Systems Letters*, vol. 8, pp. 3231–3236, 2024.
[13] J. D. Seelig and V. Jayaraman, "Neural dynamics for landmark orientation and angular path integration," *Nature*, vol. 521, no. 7551, pp. 186–191, 2015.
[14] N. E. Leonard, A. Bizyaeva, and A. Franci, "Fast and flexible multiagent decision-making," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 7, 2024.
[15] R. M. Gray, "Toeplitz and circulant matrices: A review," 2006.
[16] M. Golubitsky, I. Stewart, and D. G. Schaeffer, *Singularities and groups in bifurcation theory*. Springer Science & Business Media, 2012, vol. 2.
[17] J. R. Marden, G. Arslan, and J. S. Shamma, "Cooperative control and potential games," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 39, no. 6, pp. 1393–1407, 2009.