

Bio-inspired decision-making and control: From honeybees and neurons to network design

Vaibhav Srivastava

Naomi Ehrich Leonard

Abstract—We present nonlinear deterministic models and linear stochastic models of decision-making between alternatives that connect biological groups as diverse as honeybees and neurons. Using these models we explain how biological groups, with decentralized control and limited sensing and communication, select the highest quality alternative, flip a coin for nearly equal alternatives, optimally balance speed and accuracy, maintain robustness in the face of uncertainty, and leverage heterogeneity. Motivated by these remarkable behaviors, we present a generalizable agent-based model for the design and control of network dynamics with the advantageous features observed in the biological groups.

I. INTRODUCTION

Animals manage challenging tasks with speed, accuracy, robustness and adaptability [1]. For example, animal groups excel in tasks such as migration, foraging, and predator evasion, despite the fact that as individuals they employ decentralized strategies and face limitations on sensing, communication, and computation [2, 3]. For many applications of multi-agent system networks, ranging from transportation networks and mobile sensing networks to power networks and synthetic biological networks, features such as speed, accuracy, robustness, and adaptability are difficult to achieve yet critical for operation in complex, changing environments.

This paper focuses on a growing body of work aimed at examining animal behavior from a systems theoretic perspective and leveraging models and mechanisms for the design of high performing network dynamics and control. We present model-based investigations of the mechanisms that explain collective dynamics observed in biology in the context of decision-making among alternatives. We focus on models of honeybee swarms and networks of neurons that have been derived from empirical data, and we highlight the commonalities in their collective dynamics. To rigorously connect the mechanisms associated with decision-making in the biological groups and design methodologies for bio-inspired network dynamics and control, we describe and study a generalized multi-agent dynamic model that realizes the advantageous features observed in the biology.

Collective decision-making among alternatives is fundamental to group objectives such as navigation, migration,

foraging, and predator-avoidance. Which direction to take, when to travel, what patch to explore, whether or not to flee, are decisions that often need to be made by the group as a whole. We define a collective decision to be a group-level decision that every agent in the group ultimately takes. Typically, every agent in the group first makes an individual decision, and the collective decision follows through an external mechanism such as a quorum. Failure to come to a collective decision can lead groups to split, to lose their competitive advantage, and to become vulnerable to predation.

In this paper, we primarily consider decision-making between two alternatives. In Section II we start with collective decision-making in swarms of honeybees who choose a new nest-site among scouted out alternatives. Honeybees can choose the best quality nest-site and can also flip a coin when necessary to make a quick choice between near-equal quality alternatives. A fundamental feature of honeybee decision-making is the adaptivity and resilience that results from the interplay between the social effort of the swarm and the value of the alternatives. A model-based investigation of honeybee nest-site selection inspires control modalities for network multi-agent systems that achieve a similar adaptive and resilient performance.

In Section III, we consider collective decision-making between two alternatives in networks of neurons. Models of networks of neurons can be used to show how speed and accuracy in decision-making can be balanced in the presence of uncertainty. We present models across scales: (i) detailed biophysical models, (ii) connectionist models with abstract dynamics, and (iii) phenomenological models. We show how biophysical models can be rigorously reduced to connectionist models and how connectionist models can be reduced to phenomenological models.

We identify the similarities between the nonlinear phenomena associated with decision-making in honeybees and those associated with the biophysical models of networks of neurons. The biophysical models of neurons are nonlinear models and the phenomenological models are linear stochastic models. The nonlinear models provide valuable insights into the fundamental principles responsible for emergent collective decision-making, and the linear stochastic models provide analytic tractability essential for efficient design of systems. Connecting these models across scales enables systematic methodology to design control laws at the phenomenological level that correspond to the sophisticated mechanisms in the detailed models. This does not necessarily mean that bio-inspired control laws should be designed

This research has been supported in part by ONR grant N00014-14-1-0635, ARO grant W911NF-14-1-0431 and NSF grants ECCS-1135724 and CMMI-1635056.

V. Srivastava is with the Department of Electrical and Computer Engineering, Michigan State University, East Lansing, MI, USA, vaibhav@msu.edu

N. E. Leonard is with the Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ, USA, naomi@princeton.edu.

using linear stochastic models. In fact bio-inspired control laws have been designed using feedback of (bifurcation) parameters of the nonlinear models by introducing dynamics on these parameters at a slower time scale [4–6].

In Section IV, we present an abstract and generalizable model for design and control of bio-inspired collective decision-making in engineered multi-agent systems. We discuss various kinds of heterogeneities that can be explored in the abstract model and how parameters associated with these heterogeneities can be used as control inputs towards systematic design of decision-making networks.

We specialize the abstract model to the deterministic setting in Section V. We elucidate how the nonlinear phenomena leading to the emergence of collective decisions in models for honeybees and biophysical neuronal networks can be realized in the abstract model. We also demonstrate how value sensitive decision-making in honeybees and associated robustness explored analytically can be designed and controlled using the abstract model.

We specialize the abstract model to the linear stochastic setting in Section VI. The tractability of the linear stochastic setting allows a rigorous investigation of the influence of heterogeneities on measures of decision-making performance such as speed and accuracy. These influences can be used to design control laws to drive network dynamics as shown in [7].

Finally, in Section VII, we discuss several related research directions. These include extensions of the above models to choice among multiple alternatives, leadership as a control variable, and more sophisticated decision-making scenarios involving explore-exploit tradeoffs.

II. VALUE SENSITIVE DECISION-MAKING IN HONEYBEE NEST-SITE SELECTION

In this section we discuss a model for nest-site selection in honeybees. The presentation in this section follows from [8], [9], and the review in [10].

In the honeybee nest-site selection problem, an entire swarm must unanimously choose a good nest-site where it will live as a new colony with its queen. The choice should be made quickly, since the bees cannot survive for long without a nest, and the choice should be made accurately, since the bees cannot survive the winter in a low quality nest. The value of a candidate nest-site is determined by features such as volume, size of entrance, and height above the ground.

The process starts with a subset of scout bees that each search out a possible nest-site and ends with the swarm choosing the best of the scouted out alternatives. Each informed honeybee scout uses explicit signaling in the form of a “waggle dance” on the vertical surface of the swarm to recruit uninformed bees to commit to its discovered nest-site. Using data from an experiment in which there were two alternative sites made available to a honeybee swarm, Seeley *et al.* [9] showed that, in addition to dancing to promote their discovered site, the scouts use a cross-inhibitory stop signal to stop the dancing of the scouts recruiting for the competing site. This stop signal contributes positively

to the collective decision-making; of particular note, it was shown to facilitate breaking deadlock in the case of two near-equal value alternative sites. Because of the time pressure on the site selection process, efficient deadlock breaking can be critical for a new colony, particularly if the nearly equal sites have high value.

Seeley *et al.* [9] derived a model of the mean-field population-level dynamics of the swarm under the assumption that the total bee population size N is very large. The model describes the dynamics of $y_A = N_A/N$ and $y_B = N_B/N$, the changing fraction of bees in the population committed to nest-site A and B, respectively, and $y_U = N_U/N$, the fraction of uncommitted bees in the population. Since $N = N_A + N_B + N_U$ is constant, $y_A + y_B + y_U = 1$. Thus, the dynamics evolve on the two-dimensional unit simplex and are given by

$$\begin{aligned} dy_A &= (g_A y_U - y_A(a_A - r_A y_U + s_B y_B))dt \\ &\quad + \varrho \sqrt{y_U^2 + y_A^2 + y_U^2 y_A^2} dW_A \\ dy_B &= (g_B y_U - y_B(a_B - r_B y_U + s_A y_A))dt \\ &\quad + \varrho \sqrt{y_U^2 + y_B^2 + y_U^2 y_B^2} dW_B, \end{aligned} \quad (1)$$

where g_i is the rate of scouting discovery and commitment, a_i is the rate of abandonment of commitment, r_i is the rate of recruitment, s_i is the rate of stop signaling, ϱ is noise intensity, and dW_A and dW_B are i.i.d. Wiener process increments that are independent of each other. It is assumed that all but the stop-signal rate depend on the value v_i (quality) of the nest-site; in particular, $g_i = r_i = v_i$ and $a_i = 1/v_i$. The two stop-signal rates are assumed to be the same, i.e., $s_i = s$. Let $\Delta v = v_A - v_B$ and $\bar{v} = (v_A + v_B)/2$. When y_A or y_B crosses above a quorum threshold $\theta \in (0.5, 1]$, a collective decision is reached.

We consider the deterministic dynamics with $\varrho = 0$ studied in [8]. It was shown that for a biologically plausible parameter range, dynamics (1) admit time-scale separation and the system quickly converges to the stable manifold defined by

$$y_A y_B = \frac{2\bar{v} y_U (1 + y_A)(1 + y_B)}{s(3 - y_U)},$$

which interestingly does not depend on Δv . Thus, the dynamics (1) reduce to one-dimensional dynamics on this slow manifold. Even in presence of noise $\varrho \neq 0$, the systems quickly converges to a small neighborhood of this manifold and the one-dimensional dynamics on this manifold provide a good approximation.

For the dynamics projected onto the slow manifold and equal value nest-sites, i.e., $\Delta v = 0$, the system admits a pitchfork bifurcation at a critical stop-signaling strength $s^* = 4\bar{v}^3/(\bar{v}^2 - 1)^2$ (see Fig. 1, left panel). If $s < s^*$, then there is one stable equilibrium corresponding to deadlock ($y_A = y_B = 0.5$), i.e., half of the population committed to each alternative. If $s > s^*$, then the deadlock solution is unstable and there are two stable solutions, one corresponding to each alternative. The critical stop signal s^* decreases with the value \bar{v} of alternatives (see Fig. 2), which means that it takes higher effort to break deadlock on inferior quality nests. In

other words, for high quality choices little social effort (stop signal) is needed to convince others to reach a quorum while for low quality choices more social effort (stop signal) is needed to convince others.

When the symmetry in the value of the alternatives is broken, i.e., $\Delta v \neq 0$, the resulting bifurcation diagrams correspond to an unfolding of the pitchfork (see Fig. 1, middle panel) in which for $s \ll s^*$, there is one stable equilibrium corresponding to deadlock. For s near s^* , the only stable equilibrium point favors the superior alternative, and for $s \gg s^*$, two stable solutions emerge that correspond to each alternative.

Another important feature of the dynamics (1) is the hysteresis effect (see Fig. 1, right panel). If $\Delta v < 0$ and the dynamics have converged to alternative B , then as Δv increases, the solution does not switch to alternative A immediately as Δv becomes positive but persists for small positive values of Δv . The hysteresis effect imparts robustness to decisions in the sense that once a nest is selected and later its value is perceived to be slightly smaller than for the other nest, then the choice of nest persists.

III. COLLECTIVE DECISION-MAKING IN NEURONAL NETWORKS

In this section we review collective decision-making in neuronal networks in the context of two-alternative choice tasks. The presentation in this section follows from [11–15].

Before we discuss a biophysical model of decision-making, we define some terminology from neuroscience. Neurons are said to be excitatory or inhibitory if their activity has an excitatory or inhibitory effect on the activity of the neurons connected to them, respectively [16]. Two neurons are connected through synapses and information is transmitted across a synapse through a neurotransmitter. Several neurotransmitters in a human brain have been identified. Neurotransmitters that have a prominent effect on decision-making in two alternative choice tasks are NMDA, AMPA, and GABA [12]. For details on these neurotransmitters, we refer the reader to [17, Appendix on Neural Signaling].

A. Biophysical models

Wang [18] numerically studied a biophysically inspired model for decision-making in two alternative choice tasks. This model was analytically studied using mean-field approximations by Wong and Wang [12]. The mean-field model comprises four clusters of neurons, three of which are excitatory and one of which is inhibitory (see Fig. 3). Two of the excitatory clusters are selective to the stimulus associated with one of the two alternatives (I_1 or I_2) and the third excitatory cluster is non-selective. The excitatory connections (synapses) are modulated by NMDA and AMPA neurotransmitters while inhibitory connections are modulated by GABA neurotransmitters. A decision is made in favor of the alternative for which the neurons in the associated selective cluster first cross a fixed firing rate (decision threshold).

The postsynaptic current in cluster j is modeled as

$$I_j^{\text{type}}(t) = -g_j^{\text{type}}(\bar{V} - V_{\text{type}})N_k \omega_{kj} S_k^{\text{type}}(t), \quad (2)$$

where $k \in \{1, 2, 3, i\}$ is the presynaptic cluster, i denotes the inhibitory cluster, ω_{kj} is the coupling strength between the k -th and j -th cluster, N_k is the number of neurons in the k -th cluster, \bar{V} is the fixed average membrane potential, $\text{type} \in \{A, N, G, A_{\text{ext}}\}$ represents AMPA, NMDA, GABA, and external stimulus driven AMPA neurotransmitter, respectively, and $S_k^{\text{type}}(t)$ is the average synaptic variable, which measures the dynamics of the conductivity of the synapse.

The recurrent self-excitatory connections for selective clusters, ω_{11} and ω_{22} , play an important role in decision-making as we will show later. In the following let $\omega_{11} = \omega_{22} =: \bar{\omega}$. The steady-state firing rate of the j -th cluster can be described using a mapping $\phi_j(I_j^{\text{syn}})$, where $I_j^{\text{syn}} = \sum_{\text{type}} I_j^{\text{type}}$ is the total synaptic current entering the j -th cluster. We refer the reader to [12] for an expression for $\phi_j(\cdot)$.

Wang and Wong [12] modeled the mean-field dynamics of synaptic currents and firing rates using 11 ODEs. They then observed in simulations, for a relevant range of parameters, that the firing activity of non-selective excitatory neurons does not vary much and can be replaced by an appropriate constant. They further observed that the time constants associated with GABA and AMPA modulated synapses are much smaller than the time constants associated with NMDA modulated synapses. Thus, the currents associated with GABA and AMPA neurotransmitters evolve on faster time scales and converge to quasi-steady states quickly.

They further assumed that the firing rate dynamics associated with neuronal clusters are fast and firing rates can be approximated by their steady-state value $\phi_j(\cdot)$. Thus the model is reduced to a set of two coupled equations describing the current associated with NMDA neurotransmitters in two selective neuronal populations. From (2), it follows that the current dynamics are completely determined by

$$\frac{dS_j^N}{dt} = -\frac{S_j^N}{T_N} + \hat{\gamma}(1 - S_j^N)\phi_j(I_j^{\text{syn}}), \text{ for } j \in \{1, 2\}, \quad (3)$$

where $\hat{\gamma} = 0.641$ is computed using numerical simulations with nominal system parameters and T_N is the time constant for the NMDA modulated synapse.

The dynamics were further simplified using approximations $\phi_1 = H(y_1, y_2)$ and $\phi_2 = H(y_2, y_1)$, where H is the effective nonlinear firing rate-stimulus mapping (see [12] for details), and

$$\begin{aligned} y_1 &= J_{11}^N S_1^N - J_{21}^N S_2^N + I_1^{\text{sc}} + I_1^{\text{stim}} + I_1^{\text{noise}} \\ y_2 &= -J_{12}^N S_1^N + J_{22}^N S_2^N + I_2^{\text{sc}} + I_2^{\text{stim}} + I_2^{\text{noise}}. \end{aligned}$$

J_{ij}^N denote net self ($i = j$) and cross ($i \neq j$) connection strengths, I_j^{sc} are self-consistent currents due to GABA, AMPA and cluster 3, $I_j^{\text{stim}} = \mu(1 \pm c')$ are currents for stimulus-selective clusters, I_j^{noise} are additive Gaussian noise, and $c' \in [0, 1]$ is the coherence level of the stimuli (e.g., $c' = 1$ for perfectly clear stimuli, and $c' = 0$ for completely




Fig. 1. From [8]. Collective decision-making in the honeybee nest-site selection problem. For equal value alternatives, the decision-making is organized by a symmetric supercritical pitchfork bifurcation with the stop signal as the bifurcation parameter (left panel). For asymmetric alternatives, an unfolding of the pitchfork bifurcation favors superior alternatives for medium stop signals (middle panel). A hysteresis effect is observed as a function of Δv (right panel).

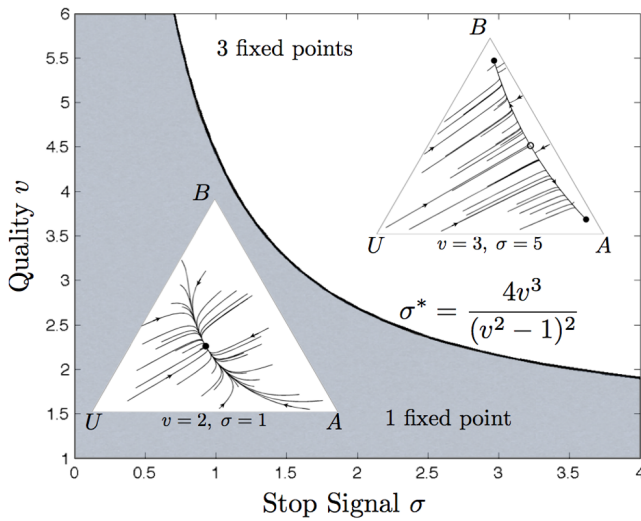


Fig. 2. From [8]. The critical value of stop signal as a function of the value of symmetric nest-sites. Smaller value of stop signal is needed for superior valued nest-sites.

random stimuli). The overall reduction procedure is summarized in Fig. 3.

For no stimulus, the system (3) admits five equilibrium points: three are stable and two are unstable. One stable node is symmetric, i.e., $S_1^N = S_2^N$, while the other two stable equilibria are asymmetric and favor one of the selective clusters. A symmetric equilibrium point corresponds to a decision-deadlock.

For stimuli with low coherence level, the system admits two stable and one unstable equilibria. The unstable equilibrium point is symmetric while the stable ones are asymmetric. As coherence is increased, the basin of attraction of the asymmetric equilibrium point favoring the correct decision increases. As coherence level is further increased the symmetric unstable node and asymmetric stable node favoring the incorrect decision are annihilated in a saddle-node bifurcation, and the system admits only one asymmetric equilibrium point associated with the correct decision. We refer the reader to [12] for detailed bifurcation diagrams.

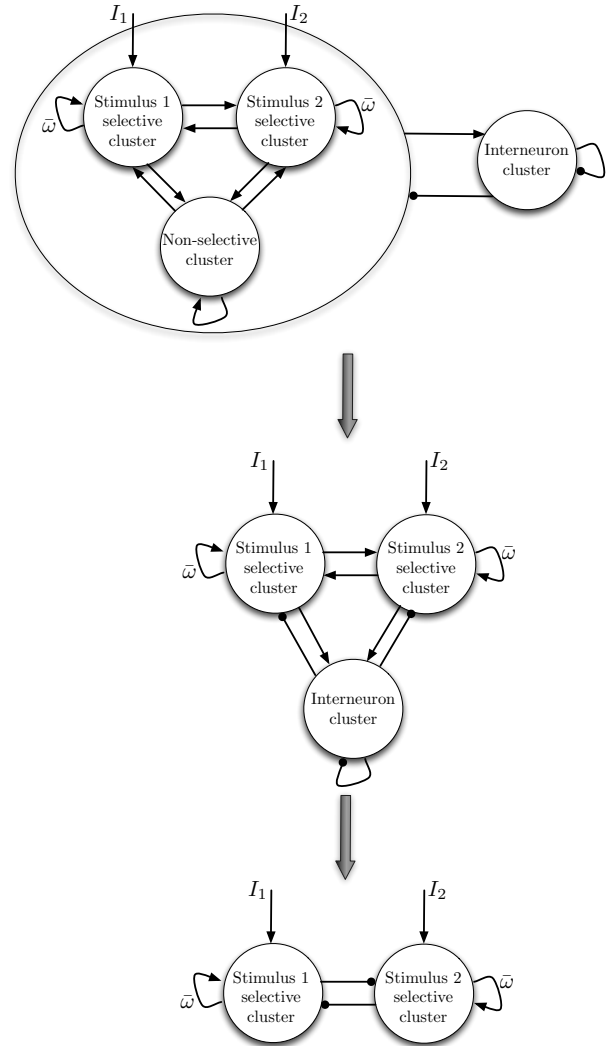


Fig. 3. Adapted from [12]. The original mean-field model that can be described using 11 equations. Approximating the activity of the non-selective cluster by a constant and using simplified frequency-current relations reduces the number of equations to 8. Using time-scale separation further reduces the number of equations to 2. Arrows represent excitatory connections and solid dots represent inhibitory connections.




Fig. 4. The short-term memory. If a stimulus of strength μ_0 is applied and it drives the system to one of the asymmetric stable equilibria associated with a decision, then even after removing the stimulus, the system stays at that asymmetric equilibrium point instead of moving to the deadlock state.

We note that the nonlinear phenomena associated with the above model have strong similarities with the honeybee nest-site selection model discussed in §II.

The bifurcation phenomenon described above exhibits rich behavior as recurrent coupling strength \bar{w} is varied. Indeed, for a given \bar{w} , as stimulus μ_0 is varied the possible set of equilibria corresponds to one of the four cases described above, but for different \bar{w} , the transition between these sets is arranged by different classes of bifurcation including supercritical pitchfork, subcritical pitchfork, and saddle-node bifurcations. We refer the reader to [12] for more details. A noteworthy behavior that is exhibited at a certain range of \bar{w} as stimulus strength μ_0 is varied is called the short-term memory shown in Fig. 4. This phenomenon is typical for a subcritical pitchfork with a fifth order stabilizing term. The presence of two saddle-node bifurcations with the subcritical pitchfork bifurcations creates a hysteresis effect. That is, if a stimulus is applied and it drives the system to one of the asymmetric stable equilibria associated with a decision, then even after removing the stimulus, the system stays at that asymmetric equilibrium point.

B. Connectionist approaches to modeling human decision-making

Unlike the biophysically detailed model described in the previous section, the connectionist model involves abstract averaged representation of the neural population. To understand decision-making at such an abstract level, firing patterns in middle temporal (MT) and lateral inter-parietal areas of the monkey brain in a moving dots experiment were studied in [19]. MT neurons are involved in motion processing and if leftward motion of the dots is more coherent, then the MT neurons sensitive to the leftward motion fire more. However, the difference between firing patterns is not substantial. In contrast, the firing pattern of LIP neurons, which are associated with eye movement control, are clearly distinguishable as time proceeds. It is believed that the LIP neurons integrate the noisy MT neuron output and this integration process is subject to inhibition

from the other population (bottom panel Fig. 3).

Let x_1 and x_2 be the mean input current to the cells in neuronal clusters selective to the two alternatives. The leaky competitive accumulator (LCA) model [15] capturing the above behavior is

$$\tau \dot{x}_1 = -x_1 - lf(x_2) + I_1 + \sigma \eta_1(t) \quad (4)$$

$$\tau \dot{x}_2 = -x_2 - lf(x_1) + I_2 + \sigma \eta_2(t), \quad (5)$$

where l is the strength of mutual inhibition, and f is a sigmoidal function that captures “current-frequency” mapping, I_j is the stimulus received by population j , η_j is additive white noise, σ is the standard deviation of the white noise, and τ is the time constant. The above set of SDEs implicitly define the integration process in the LIP neurons. Note that (4) follows the same architecture as in the bottom panel of Fig. 3.

A decision is made in favor of the alternative for which firing rate $f(x_j(t))$ first exceeds a threshold. These threshold based policies are believed to be implemented through the basal ganglia in the following way. The basal ganglia inhibits the channels for the execution of each decision in the default state. When the firing rate associated with a particular decision is sufficiently high, the inhibition of the associated channel is released leading to a winner-take-all behavior (see [13] for details).

It has been argued that neural circuits equilibrate near a point of maximum sensitivity, which corresponds to the inflection point of the sigmoidal function [11, 15]. The dynamics (4) can be linearized near the inflection point as

$$\begin{aligned} dx_1 &= (-\kappa x_1 - w x_2 + \xi_1)dt + \sigma dW_1 \\ dx_2 &= (-\kappa x_2 - w x_1 + \xi_2)dt + \sigma dW_2, \end{aligned} \quad (6)$$

where $\xi_j = I_j/\tau$, $\kappa = 1/\tau$, $w\tau/l$ is the maximum slope of f and dW_j are i.i.d. Wiener increments that are independent of each other. A decision is made in favor of alternative 1 (alternative 2) if x_1 (x_2) crosses a threshold θ .

With a change of variable $y_1 = (x_1 + x_2)/\sqrt{2}$ and $y_2 = (x_2 - x_1)/\sqrt{2}$, the above system transforms to the decoupled system

$$\begin{aligned} dy_1 &= \left(-(w + \kappa)y_1 + \frac{\xi_1 + \xi_2}{\sqrt{2}} \right)dt + \sigma d\bar{W}_1 \\ dy_2 &= \left(\ell y_2 + \frac{\xi_1 - \xi_2}{\sqrt{2}} \right)dt + \sigma d\bar{W}_2, \end{aligned} \quad (7)$$

where $\ell = w - \kappa$ and $d\bar{W}_j$ are i.i.d. Wiener increments independent of each other. The y_1 dynamics are stable and converge to a Gaussian distribution with mean $(\xi_1 + \xi_2)/\sqrt{2}(w + \kappa)$ and variance $\sigma^2/2(w + \kappa)$. Thus, if $(w + \kappa)$ is large, then y_1 very quickly converges to a neighborhood of the line $y_1 = (\xi_1 + \xi_2)/\sqrt{2}(w + \kappa)$. Along this line, the dynamics reduce to the Ornstein-Uhlenbeck (O-U) process y_2 and thresholds of x_1 and x_2 become two thresholds on y_2 defined by $\bar{\theta} = \pm 2\theta(w + \kappa) - (\xi_1 + \xi_2)/(\sqrt{2}(w + \kappa))$. When $\ell = 0$, the y_2 dynamics reduce to the so-called drift-diffusion process and the associated decision model is called the drift-diffusion model (DDM). The DDM has nice optimality properties that we describe below.

We note that the decomposition in (7) is similar to the decomposition of the fast and slow time scales in the honey-bee nest-site selection model discussed in §II. Interestingly, similar to §II, the fast dynamics are independent of $\xi_1 - \xi_2$ and depend only on $\xi_1 + \xi_2$.

C. Phenomenological model of human decision-making

We now present a phenomenological model of human decision-making. Animal and human decision-making involves several tradeoffs. It is believed that evolutionary pressure has led to behavior that is optimal in a certain sense. The fundamental tradeoff involved in two alternative choice tasks is the speed-accuracy tradeoff: making a quick decision versus making an accurate decision. This problem has been substantially studied in the engineering and mathematical statistics literature. An optimal procedure that minimizes the expected decision time for a desired level of decision error was developed by Wald [20, 21] and is called the Sequential Probability Ratio Test (SPRT).

Consider two alternatives (hypothesis) H_1 and H_2 . Suppose that the probability distribution functions of the stimulus conditional on the true alternative are $f^1(y) := f(y|H_1)$ and $f^2(y) := f(y|H_2)$. The decision-making agent may incur three types of cost: (i) a fixed cost for each observation collected, (ii) a fixed cost for a false alarm (declaring H_2 when H_1 is true), and (iii) a fixed cost for misdetection (declaring H_1 when H_2 is true). With infinitely many observations, an accurate decision can be made and consequently, the costs (ii) and (iii) can be made zero. However, this results in cost (i) becoming very large. The tradeoff between (i) versus (ii) and (iii) is called the speed-accuracy tradeoff, and the optimal solution is the SPRT described as follows:

- (i) at each time $t \in \mathbb{N}$, collect measurement y_t
- (ii) integrate evidence $\Lambda(t) := \sum_{\tau=1}^t \log \frac{f^2(y_\tau)}{f^1(y_\tau)}$
- (iii) $\begin{cases} \text{declare } H_2, & \text{if } \Lambda(t) \geq \theta, \\ \text{declare } H_1, & \text{if } \Lambda(t) \leq -\theta, \\ \text{collect another observation,} & \text{otherwise.} \end{cases}$

Here for simplicity, we have assumed that thresholds for the two hypothesis are symmetric, i.e., $\pm\theta$. In general, the magnitude of the two thresholds may be different. The choice of these thresholds dictates the speed-accuracy tradeoff: the higher the thresholds, the smaller is the probability of an erroneous decision and the larger is the expected decision time. Wald's criterion selects the two thresholds to achieve desired probabilities of misdetection and false-alarm.

The time evolution of the sufficient statistic $\Lambda(t)$ is a discrete time random walk in which the increment at each time is a random variable with mean $\beta := \mathbb{E}[\log \frac{f^2(y_\tau)}{f^1(y_\tau)} | H_k]$, and variance $\sigma^2 := \text{var}[\log \frac{f^2(y_\tau)}{f^1(y_\tau)} | H_k]$, where H_k is the true alternative. Note that $\mathbb{E}[\log \frac{f^2(y_\tau)}{f^1(y_\tau)} | H_2] = \mathcal{D}(f^2 || f^1)$ and $\mathbb{E}[\log \frac{f^1(y_\tau)}{f^2(y_\tau)} | H_1] = -\mathcal{D}(f^1 || f^2)$, where $\mathcal{D}(\cdot || \cdot)$ is the Kullback-Leibler divergence and is always non-negative.

It is well known that the continuum limit of a random walk (under certain regularity conditions) is a Wiener process with drift. Therefore, the decision-making statistic for the

SPRT converges to the following drift-diffusion process in the continuum limit:

$$dx(t) = \beta dt + \sigma dW(t), \quad x(0) = x_0, \quad (8)$$

where x_0 is the loglikelihood ratio computed using prior probabilities. For simplicity, in the following we assume an unbiased initial condition, i.e., $x_0 = 0$. As discussed above the drift rate β is positive (negative), when H_2 (H_1) is the true alternative.

Comparing (8) with (7), we note that $\beta = (\xi_1 - \xi_2)/\sqrt{2}$, i.e., the drift rate is the difference between the mean stimuli for the two alternatives. This essentially means that in a statistical sense the mean stimuli strength can be thought of as the log-likelihood of the observations $\xi_j = \mathbb{E}[f^j(y)|H_k]$. Also the optimal statistical test requires the term ℓ to be zero. In general, ℓ may not be zero, and positive values of ℓ are attributed to recency effect (old observations are given more attention) and negative values of ℓ are attributed to decay effects (newer observations are given more attention). However, fits to the behavioral data usually result in a really small value of ℓ suggesting near optimality of the decision process.

For the DDM (8) with thresholds $\pm\theta$ and drift rate $\beta \in \mathbb{R}_{>0}$, the decision time T is defined by

$$T = \inf\{t \in \mathbb{R}_{\geq 0} \mid x(t) \in \{-\theta, +\theta\}\},$$

and the probability of erroneous decision (error rate) ER is defined by $\text{ER} = \mathbb{P}(x(T) = -\theta)$. Both mean decision time and error rate can be analytically characterized for the DDM and are given by

$$\text{ET} = \frac{\theta}{\beta} \tanh\left(\frac{\beta\theta}{\sigma^2}\right) \quad \text{and} \quad \text{ER} = \frac{1}{1 + \exp\left(\frac{2\beta\theta}{\sigma^2}\right)}, \quad (9)$$

for unbiased initial condition, i.e., $x_0 = 0$. In fact, the distribution of the decision time can be analytically computed and is given by

$$f_T(t) = e^{-\frac{\beta^2 t}{2\sigma^2}} \left(e^{-\frac{2\beta\theta}{\sigma^2}} \vartheta\left(t; \frac{\theta}{\sigma_1}, \frac{2\theta}{\sigma_1}\right) + e^{\frac{2\beta\theta}{\sigma^2}} \vartheta\left(t; \frac{\theta}{\sigma}, \frac{2\theta}{\sigma}\right) \right),$$

where

$$\vartheta(t; u, v) = \sum_{k=-\infty}^{+\infty} \frac{v - u + 2kv}{\sqrt{2\pi}t^{3/2}} e^{-(v-u+2kv)^2/2t}, \quad u < v.$$

Several properties of DDM can be explicitly characterized [22] and such analysis can even be extended to cases when DDM parameters are time-varying [23]. Such analytic properties lead to efficient fitting of DDM parameters to behavioral data in two alternative choice tasks. This analytic traction and optimality make DDM a popular model to study human behavior in two alternative choice tasks.

In the context of human decision-making, the choice of threshold θ , which dictates the speed-accuracy tradeoff, is believed to be governed by reward rate maximization. The notion of reward rate maximization has its roots in the ecology literature and leads to the well-known marginal value theorem for foraging [24]. Assuming a unit reward for each

correct decision and no penalty for an incorrect decision, the reward rate (RR) for decision-making is defined by

$$\text{RR} = \frac{1 - \text{ER}}{\text{ET} + T^{\text{motor}} + D + \text{ER}D^p},$$

where T^{motor} is the motor time associated with the decision-making process, D is the response time, and D^p is the additional time that the decision-maker takes after an erroneous decision (see [11] for detailed description of the parameters). The threshold that maximizes the reward rate is the solution to the following transcendental equation:

$$e^{\frac{2\beta\theta}{\sigma^2}} - 1 = \frac{2\beta^2}{\sigma^2} \left(D + D^p + T^{\text{motor}} - \frac{\theta}{\beta} \right). \quad (10)$$

We remark that there exist regimes of parameters where biophysical models, connectionist models, and phenomenological models are mathematically equivalent, and optimality guarantees of phenomenological models extend to other models in these parameter regimes.

IV. AN ABSTRACT MODEL FOR BIO-INSPIRED NETWORK DECISION-MAKING AND CONTROL

In this section we introduce an agent-based model for design and control of bio-inspired network decision-making. The proposed model is a special case of the Hopfield network [25, 26] in which we introduce enough structure so that key behaviors can be controlled. The presentation in this section follows [6, 27–31].

We consider a set of N interconnected agents. Let $A \in \mathbb{R}^{N \times N}$ be the agent network adjacency matrix, with $a_{ij} \geq 0$ for $i \neq j$, and $a_{ii} = 0$, where $i, j \in \{1, \dots, N\}$. Let $D \in \mathbb{R}^{N \times N}$ be a diagonal matrix with $D_{ii} = d_i := \sum_{j=1}^N a_{ij}$ and let $L = D - A$ be the network Laplacian matrix. We assume that the interconnection graph is strongly connected and undirected, that is, $\text{rank}(L) = N - 1$ and $L\mathbf{1}_N = 0$, where $\mathbf{1}_N$ is the N -column-vector with all unitary entries.

Consider the following decentralized dynamics for the realization of bio-inspired collective decision-making behavior

$$dx_i = (-k_i x_i - d_i x_i + \sum_{j=1}^N a_{ij} u_j \tanh(x_j) + \beta_i) dt + \sigma d\bar{W}_i(t), \quad (11)$$

for each $i \in \{1, \dots, N\}$. Here, $\beta_i \in \{-\beta_A, 0, \beta_B\}$ is the external stimulus, $u_j \geq 0$ is a non-negative parameter representing social activity, $k_i \geq 0$ is the leadership strength, and $d\bar{W}_i$ are i.i.d. Wiener increments with $d\bar{W}_i$ and $d\bar{W}_j$ being potentially correlated. In the vector form (11) is

$$d\mathbf{x} = (-K\mathbf{x} - D\mathbf{x} + AU \tanh(\mathbf{x}) + \boldsymbol{\beta}) dt + \Sigma d\mathbf{W}(t), \quad (12)$$

where K is a diagonal matrix of k_i 's, U is a diagonal matrix of u_j 's, $\boldsymbol{\beta}$ is the vector of β_i 's, Σ is the diffusion matrix, and $d\mathbf{W}(t)$ is the i.i.d. N -dimensional Wiener increments. We interpret the state \mathbf{x} as the vector of agent opinions and treat it as a measure of ensuing decisions. We refer to (12) as the *opinion dynamics*.

The opinion dynamics model (11) may be interpreted in the following way. The term $u_j \tanh(x_j)$ in (11) is

the opinion of agent j as perceived by a generic agent i . In particular, the saturation function models an agent's assessment of the opinion of other agents: opinions with small values are assessed as they are, while the opinions with large values are assessed to have a smaller value. The parameter u_j controls this smaller value and models the social effort: higher social effort leads to a broader range of opinions being assessed correctly.

The stimulus term β_i models information or preference for agent i : $\beta_i = -\beta_A$ means information or preference for option A, $\beta_i = 0$ means no information or preference, i.e., an uninformed agent, and $\beta_i = \beta_B$ means information or preference for option B. With this interpretation, the opinion dynamics (12) are the continuous-time version of the process in which each agent at each time: (i) computes a linear combination of her opinion with the perceived opinions of her neighbors, and (ii) updates her opinion by the sum of the linear combination and the external stimulus. If $k_i = 0$, then the linear combination is a convex combination.

The opinion dynamics (12) only model the evolution of the opinion in favor of the alternatives. For a unanimous decision every agent at least must have opinion with the same sign, or the opinion of each agent should be above a positive threshold or below a negative threshold. Thus if the sum of opinions is zero or really small, then a unanimous decision cannot be reached.

Several types of common heterogeneities are captured in (11) and (12). First is the heterogeneity of leadership captured by k_i . Specifically, $k_i > 0$ for leaders and $k_i = 0$ for followers. A leader is defined as an agent that has access to some external signal and can guide the group in response to that signal. Without loss of generality we assume that the external signal is zero. Second is the heterogeneity of location in the network captured through the matrices D and A . Third is the heterogeneity of social effort captured through matrix U . The higher the social effort of an agent, the higher is its capacity to transmit its strong opinions to others. Indeed, for small values of u_j the $u_j \tanh(x_j)$ saturates at a smaller value for high values of x_j (strong opinions). Fourth is the heterogeneity of access to information captured through the β_i 's. Different values of β_i mean access to different source of information, e.g., different nest-sites in the honeybee example. $\beta_i = 0$ means no access to information. Fifth is the heterogeneity of noise captured through the diffusion matrix Σ .

In the following, we will refer to (12) with no noise, i.e., $\Sigma = 0$, $K = 0$, and $U = u\mathcal{I}_N$, where \mathcal{I}_N is the identity matrix of order N , as the *deterministic opinion dynamics*, which take the form:

$$\dot{\mathbf{x}} = -D\mathbf{x} + uA \tanh(\mathbf{x}) + \boldsymbol{\beta}. \quad (13)$$

V. DETERMINISTIC OPINION DYNAMICS AND VALUE SENSITIVE DECISION-MAKING

In this section, we discuss the deterministic opinion dynamics, and show how the model can be used to study and control for value-sensitive decision-making. The presentation follows [6, 28, 29].

A. Bifurcations in deterministic opinion dynamics

To gain intuition for the model, we consider the deterministic opinion dynamics (11) with an all-to-all network and $\beta = 0$:

$$\dot{x}_i = -(N-1)x_i + \sum_{j=1, j \neq i}^N u \tanh(x_j), \quad (14)$$

for each $i \in \{1, \dots, N\}$.

For dynamics (14), the consensus manifold is globally exponentially stable. This can be easily verified by considering a Lyapunov function $V = \sum_{i=1}^N \sum_{j=1}^N V_{ij}(\mathbf{x})$, where $V_{ij}(\mathbf{x}) = (x_i - x_j)^2/2$. It follows that

$$\begin{aligned} \dot{V}_{ij}(\mathbf{x}) &= -(N-1)(x_i - x_j)(x_i - x_j + \\ &\quad u(\tanh(x_i) - \tanh(x_j))) \\ &< -(N-1)(x_i - x_j)^2 = -2(N-1)V_{ij}, \end{aligned}$$

for all $x_i \neq x_j$. Therefore,

$$\dot{V}(\mathbf{x}) < -2(N-1)V(\mathbf{x}),$$

for all $\mathbf{x} \notin \text{span}(\mathbf{1}_N)$. Invoking LaSalle's invariance principle, it follows that the consensus manifold is globally exponentially stable.

On the consensus manifold these dynamics reduce to the scalar dynamics

$$\dot{y} = -(N-1)y + (N-1)u \tanh(y).$$

It can be verified that for the dynamics on the consensus manifold: (i) the origin is globally exponentially stable for $u \in [0, 1)$ and globally asymptotically stable for $u = 1$; and (ii) the origin is unstable and there exist two stable equilibrium points on the consensus manifold for $u > 1$. In other words, the system admits a pitchfork bifurcation at $u = 1$.

We will now discuss the deterministic opinion dynamics for $\beta \neq 0$ and for general connected undirected graphs and argue that a similar bifurcation behavior persists. The methodology that we use is that of Lyapunov-Schmidt (LS) reduction [32]. For a nonlinear set of equations $\zeta(\mathbf{x}) = 0$, a singular point is defined as a point at which the Jacobian matrix $\frac{\partial \zeta}{\partial \mathbf{x}}$ loses rank. The LS reduction decomposes the vector space into two subspaces: (i) the linear space tangent to the center manifold, i.e., the linear space spanned by eigenvectors associated with eigenvalues with zero real part, and (ii) the linear space orthogonal to linear space (i). The LS reduction then projects \mathbf{x} near the singular point on spaces (i) and (ii). Let \mathbf{x}_1 and \mathbf{x}_2 be these projections. Since by construction the projection of $\zeta(\mathbf{x})$ onto space (ii) is full rank and locally $\zeta(\mathbf{x})$ is a linear set of equations, \mathbf{x}_2 can be represented in terms of \mathbf{x}_1 . Using this representation of \mathbf{x}_2 , the projection of $\zeta(\mathbf{x})$ onto space (i) can be written explicitly in terms of \mathbf{x}_1 . Taylor series expansion of $\zeta(\mathbf{x})$ projected on subspace (i) at the singular point can be used to analyze the local nonlinear phenomena on the center manifold.

For deterministic opinion dynamics (13), it can be verified using a Lyapunov function of the form $V(\mathbf{x}) = \mathbf{x}^\top \mathbf{x}/2$ that

for $\beta = 0$, the origin is globally asymptotically stable if $0 < u \leq 1$, and is locally exponentially stable if $0 < u < 1$. At $u = 1$ the origin $\mathbf{x} = 0$ is a singular point for the fixed point equation of (13). Indeed, the linearization at the origin for $u = 1$ is $-L\mathbf{x}$, and the Laplacian matrix L has an eigenvalue at zero. Importantly, the center manifold locally corresponds to the consensus manifold, i.e., any local nonlinear phenomena on the center manifold will be the same for each agent (cf. (14) in which the consensus manifold is globally exponentially stable for all values of u and each agent's opinion bifurcates in the same way on the consensus manifold).

Let $\{0, \pm y^s\}$ be the three solution of the equation $y - uS(y) = 0$, $u > 1$. Let $P = I_N - \frac{1}{N}\mathbf{1}\mathbf{1}^\top$ the projector on $\mathbf{1}_N$ and \mathbf{d} be the vector of node degrees. Define

$$\bar{\mathbf{d}} := (L - (u-1)PA)^+ P\mathbf{d}, \quad \bar{\beta} := (L - (u-1)PA)^+ P\beta, \quad \text{and } \epsilon = (u-1)\bar{\mathbf{d}} + \mathbf{1}_N.$$

The Lyapunov-Schmidt reduction of the fixed points equation of (13) along the center manifold at $(\mathbf{x}^*, u, \beta) = (0, 1, 0)$ is

$$\begin{aligned} g(y, u, \beta) &:= \sum_{i=1}^N d_i (u \tanh(\epsilon_i y + \bar{\beta}_i) - (\epsilon_i y + \bar{\beta}_i)) + \sum_{i=1}^N \beta_i \\ &= \sum_{i=1}^N \beta_i + \sum_{i=1}^N d_i (u-1)y - \sum_{i=1}^N d_i \bar{\beta}_i y^2 - \frac{1}{3} \sum_{i=1}^n d_i y^3 \\ &\quad + O(\beta^2, (u-1)^2, y^4), \end{aligned}$$

where y is the component of \mathbf{x} along $\mathbf{1}_N$.

For $\beta = 0$, the bifurcation problem $g(y, u, 0) = \sum_{i=1}^n d_i (u-1)y - \frac{1}{3} \sum_{i=1}^n d_i y^3 + \text{h.o.t.}$ and corresponds to a symmetric pitchfork singularity. In particular, for $u > 1$ and $|u-1|$ sufficiently small, there are exactly three fixed points: the origin and $\pm y^s \mathbf{1}_N$. For $\beta \neq 0$, the bifurcation problem $g(y, u, \beta)$ is an unfolding of the symmetric pitchfork and the associated bifurcation diagrams are similar to the bifurcation diagram in the middle panel of Fig. 1.

B. Value-sensitive decision-making

As shown in §II, a noteworthy feature of collective decision making in honeybees is the value sensitivity of the social effort, i.e., the stop signal required to break the deadlock for superior nest-sites is smaller than for low quality nest-sites. We now show that such behavior can be accomplished using our agent-based model.

To this end, consider the following modification of opinion dynamics (13):

$$\dot{\mathbf{x}}(t) = -u_1 D\mathbf{x}(t) + u_2 A\mathbf{x}(t) + \mathbf{v}, \quad (15)$$

where $\mathbf{v} \in \mathbb{R}^n$ is the vector with entries equal to the value of the alternative to which the corresponding agent is exposed. For simplicity, suppose agents have to select between two alternatives both with values $v \in \mathbb{R}_{>0}$. Then, $v_i = v$ if the agent i is exposed to alternative 1, $v_i = -v$ if the agent i is exposed to alternative 2 and $v_i = 0$ if the agent i is not exposed to any alternative. By changing to time scale

$\tau = u_1 t$, we can map (15) back to the deterministic opinion dynamics (13) with $u = u_2 v$ and $\beta_i \in \{-v^2, 0, v^2\}$.

In (15), $u_1 \in \mathbb{R}_{>0}$ is the rate at which an agent commits to the alternative it is exposed to, and $u_2 \in \mathbb{R}_{>0}$ is the rate at which an agent is attracted by other agents to their alternative. We let $u_1 = \frac{1}{v}$. Such choice of u_1 means that for small values of v and a fixed value of u_2 , the opinions converge to deadlock quickly.

Assume that the number of agents exposed to each nest-site is the same. Then, for small values of u_2 , the sum of all opinions in (15) converges to zero and there is no unanimity in the group. However, at a critical value of u_2 a symmetric pitchfork bifurcation occurs and the sum of opinion converges to a positive or a negative equilibrium point that corresponds to a decision for one of the nest-sites. For all-to-all graphs, the expression u_2^* in terms of the nest-site value v can be analytically computed:

$$u_2^* = \frac{1}{v} + \frac{(1 + 3N^3)^2(N - n_3)}{9N^9} v^3 + O(v^7), \quad (16)$$

where n_3 is the number of informed agents. Thus, diminishing social effort for decision-making as a function of the value of the alternative can be captured by the abstract model (12).

VI. LINEAR STOCHASTIC OPINION DYNAMICS: INFLUENCE OF HETEROGENEITY

We now study dynamics (12) with noise. Nonlinear stochastic differential equations are very hard to analyze. Therefore, to gain tractability we linearize dynamics (12) to obtain

$$d\mathbf{x} = (\boldsymbol{\beta} - \bar{L}\mathbf{x})dt + \Sigma d\mathbf{W}(t), \quad \mathbf{x}(0) = \mathbf{x}_0, \quad (17)$$

where $\bar{L} = (K + D - AU)$. The linear stochastic model allows for explicit analysis of the influence of heterogeneity on decision-making performance, e.g., on expected decision times and error rates. The presentation in this section follows from [27, 30, 31].

For simplicity we assume that the matrix \bar{L} is diagonalizable and has all but one eigenvalue with positive real part. Further, suppose the eigenvalue with the smallest real part is purely real. Let \mathbf{e}_p and \mathbf{e}_p^\dagger be right and left eigenvectors of \bar{L} associated with eigenvalue λ_p , $p \in \{1, \dots, n\}$ such that $\bar{L} = \sum_{p=1}^n \lambda_p \mathbf{e}_p \mathbf{e}_p^\dagger$. We assume that eigenvalues are ordered in increasing order of their real parts. To analyze decision-making properties of each agent in the opinion dynamics (17), we develop a decoupled approximation to (17).

To this end, we decompose the opinion dynamics (17) into principal and residual components. We define the principal component of $\mathbf{x}(t)$ as $\mathbf{x}_{\text{prin}}(t) = x_p \mathbf{e}_1$, where $x_p = \mathbf{e}_1^\dagger \mathbf{x}(t)$. We define the residual component as $\boldsymbol{\epsilon}(t) = \mathbf{x}(t) - \mathbf{x}_{\text{prin}}(t)$. It follows that the principal component satisfies

$$\begin{aligned} dx_p(t) &= -\mathbf{e}_1^\dagger \bar{L} \mathbf{x}(t) dt + \mathbf{e}_1^\dagger \boldsymbol{\beta} dt + \mathbf{e}_1^\dagger \Sigma d\mathbf{W}_n(t) \\ &= -\lambda_1 x_p(t) dt + \beta_p dt + \sigma_p dW_1(t), \end{aligned} \quad (18)$$

where $\beta_p = \mathbf{e}_1^\dagger \boldsymbol{\beta}$, $\sigma_p = \|\mathbf{e}_1^\dagger \Sigma\|_2$, and $W_1(t)$ is the standard one-dimensional Wiener process. Thus, the principal component evolves according to an Ornstein-Uhlenbeck (O-U) process [33].

Similarly, the residual component satisfies

$$\begin{aligned} d\boldsymbol{\epsilon}(t) &= -\bar{L}\boldsymbol{\epsilon}(t)dt + \lambda_1 \mathbf{e}_1 \mathbf{e}_1^\dagger \mathbf{x}(t) + \bar{\boldsymbol{\beta}} dt + \bar{\Sigma} d\mathbf{W}_n(t) \\ &= -\bar{L}\boldsymbol{\epsilon}(t)dt + \bar{\boldsymbol{\beta}} dt + \bar{\Sigma} d\mathbf{W}_n(t), \end{aligned} \quad (19)$$

where $\bar{\boldsymbol{\beta}} = \boldsymbol{\beta} - \mathbf{e}_1^\dagger \boldsymbol{\beta} \mathbf{e}_1$ and $\bar{\Sigma} = (\mathcal{I}_N - \mathbf{e}_1 \mathbf{e}_1^\dagger) \Sigma$.

It can be shown that the expected value and variance of principal component x_p satisfy

$$\begin{aligned} \mathbb{E}[x_p(t)] &= \lim_{\lambda \rightarrow \lambda_1^+} \frac{\beta_p}{\lambda} (1 - \exp(-\lambda t)), \text{ and} \\ \text{Var}[x_p(t)] &= \lim_{\lambda \rightarrow \lambda_1^+} \frac{\sigma_p^2}{2\lambda} (1 - \exp(-2\lambda t)). \end{aligned}$$

The expected value and covariance of residual component $\boldsymbol{\epsilon}(t)$ satisfy

$$\begin{aligned} \lim_{t \rightarrow +\infty} \mathbb{E}[\boldsymbol{\epsilon}(t)] &= \sum_{p=2}^n \frac{1}{\lambda_p} \mathbf{e}_p \mathbf{e}_p^\dagger \boldsymbol{\beta}, \text{ and} \\ \lim_{t \rightarrow +\infty} \text{Cov}[\boldsymbol{\epsilon}(t)] &= \sum_{q=2}^n \sum_{r=2}^n \frac{\mathbf{e}_q^\dagger \Sigma \Sigma^\top \mathbf{e}_r^\dagger}{\lambda_q + \lambda_r} \mathbf{e}_q \mathbf{u}_r^\top, \end{aligned}$$

where \mathbf{e}_q^\dagger is the transpose of \mathbf{e}_q . Furthermore, the covariance between the principal component and the residual component satisfies

$$\lim_{t \rightarrow +\infty} \text{Cov}(x_p(t), \boldsymbol{\epsilon}(t)) = \sum_{q=2}^n \frac{\mathbf{e}_q^\dagger \Sigma \Sigma^\top \mathbf{e}_1^\dagger}{\lambda_1 + \lambda_q} \mathbf{e}_q.$$

We note that the principal component is already decoupled. In particular, the k -th principal component is $x_k^{\text{prin}}(t) = x_p(t) e_{1k}$, where e_{1k} is the k -th element of \mathbf{e}_1 , and can be modeled by the O-U process (18).

We now compute an asymptotically matching approximation to the residual component $\boldsymbol{\epsilon}(t)$. Since $\boldsymbol{\epsilon}(t)$ is the solution of the linear SDE (19) with constant coefficients and a constant input, it is a continuous Gaussian process. The asymptotic values of expected value and variance of $\boldsymbol{\epsilon}(t)$ are constants. A simple, scalar, and continuous Gaussian process with a constant asymptotic expected value and a constant asymptotic variance is the O-U process. In the following, we approximate the k -th element of the residual component $\epsilon_k(t)$ by an O-U process. Such an approximate O-U process must also capture the correlation between $x_p(t)$ and $\epsilon_k(t)$. We first introduce some notation. Let

$$\begin{aligned} \frac{1}{\mu_k} &:= \lim_{t \rightarrow +\infty} \text{Var}(\epsilon_k(t)), \quad \alpha_k := \lim_{t \rightarrow +\infty} \mathbb{E}[\epsilon_k(t)], \text{ and} \\ \gamma_k &:= \lim_{t \rightarrow +\infty} \text{Cov}(x_p(t), \epsilon_k(t)). \end{aligned}$$

We now propose the following coupled dynamics for $x_k^{\text{prin}}(t)$

and $\varepsilon_k(t)$

$$\begin{aligned} \begin{bmatrix} dx_k^{\text{prin}}(t) \\ d\varepsilon_k(t) \end{bmatrix} &= \begin{bmatrix} -\lambda_1 & 0 \\ 0 & -\frac{\mu_k}{2} \end{bmatrix} \begin{bmatrix} x_k^{\text{prin}}(t) \\ \varepsilon_k(t) \end{bmatrix} dt + \begin{bmatrix} \frac{\beta_p u_{1k}}{2} \\ \frac{\alpha_k \mu_k}{2} \end{bmatrix} dt \\ &+ \begin{bmatrix} \sigma_p u_{1k} & 0 \\ \frac{(2\lambda_1 + \mu_k)\gamma_k}{2\sigma_p} & \sqrt{1 - \frac{(2\lambda_1 + \mu_k)^2 \gamma_k^2}{4\sigma_p^2}} \end{bmatrix} \begin{bmatrix} dW_1(t) \\ dW_2(t) \end{bmatrix}. \end{aligned} \quad (20)$$

Equation (20) can be solved to show that it asymptotically matches the metrics for residual component (19), i.e.,

$$\begin{aligned} \lim_{t \rightarrow +\infty} \mathbb{E}[\varepsilon_k(t)] &= \alpha_k, \quad \lim_{t \rightarrow +\infty} \text{Var}(\varepsilon_k(t)) = \frac{1}{\mu_k}, \\ \text{and } \lim_{t \rightarrow +\infty} \text{Cov}(x_k^{\text{prin}}(t), \varepsilon_k(t)) &= u_{1k} \gamma_k. \end{aligned}$$

Notice that if \bar{L} is a symmetric matrix and Σ is an identity matrix, then the decomposition in (18) and (19) corresponds to principal components of the covariance of $\mathbf{x}(t)$. In such a case, $x_k^{\text{prin}}(t)$ and $\varepsilon_k(t)$ are uncorrelated and the diffusion matrix in (20) is a diagonal matrix. It should be noted that in contrast to standard spectral methods, which consider only the principal component or a few principal components, we consider the principal component and compute an asymptotic matching approximation to the residual components. This ensures that the original dynamics and the approximate decoupled dynamics converge to the same values.

The approximate dynamics (20) can be used to obtain the dynamics of the approximate state $y_k(t) = x_k^{\text{prin}}(t) + \varepsilon_k(t)$:

$$\begin{aligned} \begin{bmatrix} dy_k(t) \\ d\varepsilon_k(t) \end{bmatrix} &= \begin{bmatrix} -\lambda_1 & \lambda_1 - \frac{\mu_k}{2} \\ 0 & -\frac{\mu_k}{2} \end{bmatrix} \begin{bmatrix} y_k(t) \\ \varepsilon_k(t) \end{bmatrix} dt + \begin{bmatrix} \beta_p u_{1k} + \frac{\alpha_k \mu_k}{2} \\ \frac{\alpha_k \mu_k}{2} \end{bmatrix} dt \\ &+ \begin{bmatrix} \sigma_p u_{1k} + \frac{(2\lambda_1 + \mu_k)\gamma_k}{2\sigma_p} & \sqrt{1 - \frac{(2\lambda_1 + \mu_k)^2 \gamma_k^2}{4\sigma_p^2}} \\ \frac{(2\lambda_1 + \mu_k)\gamma_k}{2\sigma_p} & \sqrt{1 - \frac{(2\lambda_1 + \mu_k)^2 \gamma_k^2}{4\sigma_p^2}} \end{bmatrix} \begin{bmatrix} dW_1(t) \\ dW_2(t) \end{bmatrix}. \end{aligned} \quad (21)$$

For $K = 0$, $U = I_N$, and $\Sigma = I_N$, the dynamics (17) are called a coupled DDM and are given by

$$d\mathbf{x} = (\boldsymbol{\beta} - L\mathbf{x})dt + d\mathbf{W}(t), \quad \mathbf{x}(0) = \mathbf{x}_0. \quad (22)$$

Coupled DDM (22) is the continuum limit of distributed log-likelihood aggregation for binary sequential hypothesis testing in which at each time every agent (i) gathers an observation and computes the log likelihood ratio of the current observation, (ii) adds it to her current decision-making statistic, and (iii) computes the convex combination of her updated statistic with her neighbors. Using the updated statistic $x_k(t)$, the k -th agent decides on alternative 1 if $x_k(t)$ crosses the threshold θ_k and on alternative 2 if $x_k(t)$ crosses the threshold $-\theta_k$.

Unlike SPRT and DDM, such distributed decision-making policy is not necessarily optimal. However, such policy is asymptotically optimal, i.e., optimal when decision times are large, e.g., when thresholds are high or the drift rates are low. Here optimality is defined with respect to the speed-accuracy tradeoff, i.e., the decision-maker has the minimum expected decision time for a given error rate.

We now specialize the decoupled approximation (20) to (22). The approximate evidence at node k is $y_k(t) :=$

$\frac{1}{N} \mathbf{1}_N^\top \mathbf{x}(t) + \varepsilon_k(t)$. Here, the principal component corresponds to evidence averaged across the network, and the residual component corresponds to the deviation from this average. The evidence aggregation dynamics for the reduced model are

$$\begin{bmatrix} dy_k(t) \\ d\varepsilon_k(t) \end{bmatrix} = \begin{bmatrix} \hat{\beta} + \frac{\mu_k}{2}(\alpha_k - \varepsilon_k(t)) \\ \frac{\mu_k}{2}(\alpha_k - \varepsilon_k(t)) \end{bmatrix} dt + \begin{bmatrix} \frac{1}{\sqrt{n}} & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} dW_1(t) \\ dW_2(t) \end{bmatrix}, \quad (23)$$

with $y_k(0) = \bar{x}_0$, $\varepsilon_k(0) = 0$, and $\hat{\beta} = \frac{1}{N} \mathbf{1}_N^\top \boldsymbol{\beta}$.

We can further assume that the only heterogeneity among agents is the heterogeneity of location in the network, i.e., $\boldsymbol{\beta} = \beta \mathbf{1}_N$. Also, for simplicity let $\bar{x}_0 = 0$. Then, the reduced model (23) further reduces to

$$\begin{bmatrix} dy_k(t) \\ d\varepsilon_k(t) \end{bmatrix} = \begin{bmatrix} \beta - \frac{\mu_k \varepsilon_k(t)}{2} \\ -\frac{\mu_k \varepsilon_k(t)}{2} \end{bmatrix} dt + \begin{bmatrix} \frac{1}{\sqrt{n}} & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} dW_1(t) \\ dW_2(t) \end{bmatrix}. \quad (24)$$

Note that as we made the noise i.i.d. across agents, i.e., set $\Sigma = I_N$ in (17), γ_k goes to 0. Similarly, making the drift rates identical, i.e., $\boldsymbol{\beta} = \beta \mathbf{1}_N$, α_k goes to 0. In this sense, the parameters μ_k , α_k , and γ_k primarily capture the influence of heterogeneity due to location in the network, heterogeneity due to access to different quality of information, and heterogeneity due to coherence of noise, respectively.

For the coupled DDM (22), the e_k 's are the eigenvectors of the Laplacian matrix, and μ_k is called the node certainty index defined by

$$\frac{1}{\mu_k} = \sum_{p=2}^n \frac{1}{2\lambda_p} e_k^{(p)2}.$$

It has been shown in [27] that the node certainty index is equivalent to the *information centrality* [34] which is defined as the inverse of the harmonic mean of the resistance distance [35] of the given node to every other node in the network graph. In particular,

$$\frac{1}{\mu_k} = \frac{\sigma^2}{2} \left(\frac{1}{\kappa_{\text{info}}(k)} - \frac{K_f}{n^2} \right),$$

where $\kappa_{\text{info}}(k)$ is the information centrality of the k -th node and K_f is the Kirchhoff index of the network graph.

It should be noted that for the coupled DDM (22), the evidence $y_k(t) = x_k^{\text{prin}}(t) + \varepsilon_k(t) = \frac{1}{N} \mathbf{1}_N^\top \mathbf{x}(t) + \varepsilon_k(t)$, and the principal component is the same for each agent. So any difference in performance among agents is attributed to $\varepsilon_k(t)$. Consequently, if the only heterogeneity is the heterogeneity of location in the network, then the difference in performance of agents is completely determined by μ_k and hence, by the information centrality of that agent.

Without loss of generality we assume that $\beta > 0$. Then the decision time for the k -th agent is $T_k = \inf\{t \geq 0 \mid y_k(t) \in \{-\theta_k, \theta_k\}\}$. The error rate, i.e., the probability of making an incorrect decision is $\text{ER}_k = \mathbb{P}(y_k(T_k) = -\theta_k)$. It was observed in [30] that the effect of location heterogeneity can be accurately captured by keeping the opinion evolution of each


agent the same and equal to $y_k(t) = x_k^{\text{prin}}(t) = \frac{1}{N} \mathbf{1}_N^\top \mathbf{x}(t)$ and delegating the effect of $\epsilon_k(t)$ to the threshold $\pm\theta_k$. In particular, it was observed that for large thresholds, the decision-making performance of the k -th agent is equivalent to the performance of the DDM

$$dx_{\text{cen}}(t) = \beta dt + 1/\sqrt{N} dt, \quad x_{\text{cen}}(0) = 0, \quad (25)$$


with thresholds $\pm\theta_k \mp \bar{K}(\beta)/\sqrt{\mu_k}$. The function $\bar{K} : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ was empirically computed and is defined by

$$\bar{K}(\beta) = \frac{e^{-\frac{1}{4\sqrt{\beta}}}}{\sqrt{\beta}(1 + \beta/3)}.$$

Therefore if each agent has the same thresholds, then the effective thresholds for (22) are higher for a more centralized agent as defined by information centrality. Consequently, a more central agent has a larger expected decision time and a smaller error rate. This effect can be seen in Fig. 5.



(a) Interaction graph for decision-makers



(b) Log-likelihood ratio of no error (c) Expected decision times

Fig. 5. From [30]. The decision-making performance of agents interacting according to (22) and (25) with corrected thresholds. The blue \times , the red $+$, and the green triangles represent the performance of the coupled DDM (22) for decision-makers 1, 2, and 6, respectively. The blue dashed lines, the red solid lines with dots, and the green solid lines represent the performance of the threshold corrected DDM (25) for decision-makers 1, 2, and 6, respectively.

Consider the scenario when each agent may choose thresholds to maximize their reward rates defined in §III-C. Let $\pm\theta^*$ be the optimal thresholds that maximize the reward rate for DDM (25). Then, the k -th agent can choose her thresholds as $\pm\theta^* \pm \bar{K}(\beta)/\sqrt{\mu_k}$ to achieve the optimal threshold. Thus, in this setting the more central agent has lower thresholds. Moreover, each agent has the same error rate and expected decision time.

VII. DISCUSSION

We presented models for collective decision-making in fundamental decision-making tasks, namely, the two alternative choice tasks. In this section, we summarize the ideas

presented in the paper and point to several interesting and related research directions that we did not cover.

We presented a model for collective decision-making in a swarm of honeybees that must decide between two potential nest-sites and discussed how a decision-deadlock is broken using social effort (stop signal). This model was originally studied in [8, 9]. We showed that the fundamental nonlinear phenomenon associated with the decision-making in nest-site selection by honeybees is the pitchfork bifurcation and its unfolding. This is not the only scenario in which the collective decision-making in animal groups is realized through a pitchfork bifurcation. In the context of foraging decisions in a school of fish [36–38], the pitchfork bifurcation plays a pivotal role. In particular, it was shown in [38] that a democratic collective decision emerges through the unfolding of a pitchfork bifurcation when the number of uninformed individuals is treated as a bifurcation parameter. Another context in which decision-making emerges through a pitchfork bifurcation is in the replicator-mutator equations from evolutionary dynamics [39]. These equations serve as a model for the evolution of language and behavioral dynamics in social networks.

We discussed collective decision-making in neuronal networks at several levels of abstraction. At the biophysical level, we discussed the pivotal role the pitchfork bifurcation plays in decision-making using the model from [12, 18]. For different combinations of the stimulus strength and recurrent connection strength in the neuronal clusters, there may exist only one stable symmetric equilibrium point, two asymmetric equilibrium points, or three stable equilibrium point of which two are asymmetric and one is symmetric. Two stable equilibrium points are realized through a supercritical pitchfork bifurcation, while three stable equilibrium points are realized through a subcritical pitchfork bifurcation.

We discussed a connectionist model proposed in [15] and studied in detail in [40]. The connectionist models rely on similar architectures as in the biophysical models but employ a simpler abstract dynamics. In particular, they rely on evidence integration and mutual-inhibition. The neural data suggesting such dynamics was observed in [19, 41, 42]. We discussed how, using time-scale separation and linearization, these models can be reduced to the OU model of decision-making of which DDM is a special case.

We discussed phenomenological models of collective decision-making. The underlying philosophy of such models is that the evolutionary forces have led to optimal behavior when faced with decision-making tradeoffs. We discussed the SPRT which is the optimal decision-making algorithm in the context of the speed-accuracy tradeoff. The SPRT in the continuum limit converges to the DDM. The use of SPRT and DDM to study human decision-making goes back to Laming [43]. Ratcliff [44] made significant advances to DDM-like models of human decision-making. For a review of such decision models see [11, 45]. Extensions of DDM to tasks in which contextual information needs to be learned are discussed in [46]. Luce [47] discusses how decision-making performance captured through reaction times and error rates

can inform about mental organization.

We discussed an abstract model for realization of bio-inspired collective decision-making in engineered networks. These models were inspired by Hopfield networks [26]. We presented different kind of heterogeneities and individual differences captured by the abstract model. For the deterministic version of the model, we elucidated the nonlinear phenomena associated with the model using Lyapunov-Schmidt reduction and singularity theory. Specifically, the abstract model realizes the pitchfork bifurcation and its unfolding along the consensus manifold, i.e., the abstract model achieves a unanimous decision using nonlinear phenomena similar to those in animal groups and neuronal networks. We also showed that the abstract model can realize value-sensitive decision-making as in the honeybee nest-site selection problem.

We presented the linearized stochastic version of the abstract model. The linearized stochastic model is amenable to explicit analysis of the influence of heterogeneities on the decision-making performance of individuals. We showed that if the only heterogeneity is due to agent location in the network, then the performance of an agent is determined by the information centrality of the agent in the network graph. Specifically, if each agent has the same thresholds, then a more (information) central agent is slower and more accurate. The linearized stochastic model is the continuum limit of distributed sequential hypothesis test and is asymptotically optimal. Such distributed hypothesis testing problems have been extensively studied in the engineering literature [48–51]. The bio-inspired abstract model provides insights into how social efforts can be incorporated in these distributed models to control for better performance.

We primarily focused on heterogeneities due to location in the network, due to access to different information, and due to coherence of noise across agents. An important class of heterogeneity is leadership, which has received a lot of attention recently. A significant focus is on the leader selection problem, i.e., which agents should be selected as leaders such that the overall system uncertainty is minimized [52–55]. It can be argued that minimum overall uncertainty may lead to smaller overall decision-making error. The study of emergence of leadership from an evolutionary perspective [56–58] is another area that has received attention.

We focused on two alternative choice tasks in this paper. However, the kind of analysis we presented is not limited to two alternative choice tasks. The bifurcation model can be extended to multiple choices. However, the associated nonlinear phenomenon become more complex. For example, in the context of the replicator-mutator equations of evolutionary dynamics, the nonlinear phenomenon associated with the three alternative choice task studied in [59] is shown in Fig. 6. Indeed, the bifurcation diagram is more complicated than a simple pitchfork bifurcation. Nevertheless, it still provides valuable insights into the remarkable adaptive and robust properties of collective decision-making.

The connectionist models for collective decision-making in neuronal networks also extend to multiple alternatives [15]. The multiple alternative equivalent of DDM is




Fig. 6. Bifurcations associated with three symmetric choices. For large values of u we observe a decision-deadlock among three alternatives, as u is decreased six new equilibrium points appear in three symmetric saddle node bifurcations. The new stable nodes correspond to the choice of each alternative, and the three saddle nodes correspond to deadlock between each pair. Finally as u is further increased, three saddle nodes cross the stable node corresponding to deadlock among all three alternatives in an S_3 symmetric transcritical bifurcation and three choices emerge as the only stable equilibrium points.

called the race model [60] and is the continuum limit of the multiple sequential probability ratio test (MSPRT) [61]. In MSPRT, the farther a hypothesis is from the true hypothesis in the sense of KL divergence between the underlying distributions, the sooner that hypothesis is rejected. After some time, the MSPRT reduces to a test of hypothesis between the true alternative and the second best alternative. This is remarkably similar to the bifurcation diagram in Fig. 6: as three new stable equilibrium points appear, the farthest hypothesis is rejected. We are not aware of biophysical models that capture decision-making behavior in multiple alternative choice tasks. However, we believe that the nonlinear phenomena in such models for three-alternative choices tasks should be similar to Fig. 6.

Another decision-making tradeoff that is significant both in natural systems and engineered networks is explore-vs-exploit, i.e., choosing between the most informative and the most rewarding actions. A multiarmed bandit problem [62] is a canonical example that captures this tradeoff. Consequently, these problems have received significant attention in both mathematical biology and the engineering literature. These problems have been used to study foraging decisions in animal groups [63–65]. In the context of neuronal networks, these problems have been studied at the phenomenological level [66, 67]. Explore-exploit problems using the linear stochastic abstract framework discussed here have been studied in [68, 69]. The importance of studying the biophysics and neuroscience of sensory systems with ecological and evolutionary processes in the context of explore-exploit tension has been argued [70] and is an active research direction.

In conclusion, there is a lot of commonality in decision-making problems that exist in seemingly disparate fields of ecology, neuroscience, and engineering. A synergistic approach that leverages the salient features of approaches in these fields towards a unified framework for adaptive, re-

silient, and tractable design and control of network systems is needed. We summarized some of our efforts in this direction, but a lot still needs to be explored for the realization of such a framework.

REFERENCES

- [1] J. K. Parrish and L. Edelman-Keshet. Complexity, pattern, and evolutionary trade-offs in animal aggregation. *Science*, 284(5411):99–101, 1999.
- [2] J. Krause and G. D. Ruxton. *Living in Groups*. Oxford University Press, 2002.
- [3] D. J. T. Sumpter. *Collective Animal Behavior*. Princeton University Press, 2010.
- [4] K. Özcimder, B. Dey, R. J. Lazier, D. Trueman, and N. E. Leonard. Investigating group behavior in dance: an evolutionary dynamics approach. In *American Control Conference*, pages 6465–6470, 2016.
- [5] A. J. Krener, W. Kang, and D. E. Chang. Control bifurcations. *IEEE Transactions on Automatic Control*, 49(8):1231–1246, 2004.
- [6] R. Gray, A. Franci, V. Srivastava, and N. E. Leonard. Honeybee-inspired dynamics for multi-agent decision-making. *arXiv preprint arXiv:1503.08526*, February 2017.
- [7] G. Punzo, G. F. Young, M. Macdonald, and N. E. Leonard. Using network dynamical influence to drive consensus. *Scientific Reports*, 6(26318), 2016.
- [8] D. Pais, P. M. Hogan, T. Schlegel, N. R. Franks, N. E. Leonard, and J. A. R. Marshall. A mechanism for value-sensitive decision-making. *PLoS ONE*, 8(9):e73216, 2013.
- [9] T. D. Seeley, P. K. Visscher, T. Schlegel, P. M. Hogan, N. R. Franks, and J. A. R. Marshall. Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, 335(6064):108–111, 2012.
- [10] N. E. Leonard. Multi-agent system dynamics: Bifurcation and behavior of animal groups. *IFAC Annual Reviews in Control*, 38:171–183, 2014.
- [11] R. Bogacz, E. Brown, J. Moehlis, P. Holmes, and J. D. Cohen. The physics of optimal decision making: A formal analysis of performance in two-alternative forced choice tasks. *Psychological Review*, 113(4):700–765, 2006.
- [12] K-F Wong and X-J Wang. A recurrent network mechanism of time integration in perceptual decisions. *The Journal of Neuroscience*, 26(4):1314–1328, 2006.
- [13] R. Bogacz. Optimal decision-making theories: linking neurobiology with behaviour. *Trends in Cognitive Sciences*, 11(3):118–125, 2007.
- [14] P. Eckhoff and P. Holmes. Mathematical Neuroscience: MAT323/APC323 Course Notes, Princeton University. with contributions by K. Sadeghi, M. Schwemmer and K. Wong-Lin, Spring 2014.
- [15] M. Usher and J. L. McClelland. The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3):550, 2001.
- [16] P. Dayan and L. F. Abbott. *Theoretical Neuroscience*. Cambridge, MA: MIT Press, 2001.
- [17] D. Purves, R. Cabeza, S. A. Huettel, K. S. LaBar, M. L. Platt, M. G. Woldorff, and E. M. Brannon. *Cognitive Neuroscience*. Sunderland: Sinauer Associates, Inc, 2008.
- [18] X-J Wang. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36(5):955–968, 2002.
- [19] M. N. Shadlen and W. T. Newsome. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86(4):1916–1936, 2001.
- [20] A. Wald. Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, 16(2):117–186, 1945.
- [21] A. Wald and J. Wolfowitz. Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, 19(3):326–339, 1948.
- [22] V. Srivastava, P. Holmes, and P. Simen. Explicit moments of decision times for single- and double-threshold drift-diffusion processes. *Journal of Mathematical Psychology*, 75:96 – 109, 2016. Special Issue in Honor of R. Duncan Luce.
- [23] V. Srivastava, S. Feng, J. D. Cohen, N. E. Leonard, and A. Shenhav. A martingale analysis of first passage times of time-dependent Wiener diffusion models. *Journal of Mathematical Psychology*, 2017. In press.
- [24] E. L. Charnov. Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, 9(2):129–136, 1976.
- [25] J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982.
- [26] J. J. Hopfield. Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences*, 81(10):3088–3092, 1984.
- [27] I. Poulakakis, G. F. Young, L. Scardovi, and N. E. Leonard. Information centrality and ordering of nodes for accuracy in noisy decision-making networks. *IEEE Transactions on Automatic Control*, 61(4):1040–1045, 2016.
- [28] A. Franci, V. Srivastava, and N. E. Leonard. A realization theory for bio-inspired collective decision making. *arXiv preprint arXiv:1503.08526*, March 2015.
- [29] R. Gray, A. Franci, V. Srivastava, and N. E. Leonard. An agent-based framework for bio-inspired value-sensitive decision-making. In *IFAC World Congress*, Toulouse, France, July 2017.
- [30] V. Srivastava and N. E. Leonard. Collective decision-making in ideal networks: The speed-accuracy tradeoff. *IEEE Transactions on Control of Network Systems*, 1(1):121–132, 2014.
- [31] V. Srivastava and N. E. Leonard. On first passage time problems in collective decision making with heterogeneous agents. In *American Control Conference*, pages 2113–2118, Chicago, IL, June 2015.
- [32] M. Golubitsky and D. G. Schaeffer. *Singularities and Groups in Bifurcation Theory*, volume 51 of *Applied Mathematical Sciences*. Springer-Verlag, 1985.
- [33] C. Gardiner. *Stochastic Methods: A Handbook for the Natural and Social Sciences*. Springer, fourth edition, 2009.
- [34] K. Stephenson and M. Zelen. Rethinking centrality: Methods and examples. *Social Networks*, 11(1):1–37, 1989.
- [35] D. J. Klein and M. Randić. Resistance distance. *Journal of Mathematical Chemistry*, 12(1):81–95, 1993.
- [36] I. D. Couzin, J. Krause, N. R. Franks, and S. A. Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516, 2005.
- [37] N. E. Leonard, T. Shen, B. Nabet, L. Scardovi, I. D. Couzin, and S. A. Levin. Decision versus compromise for animal groups in motion. *Proceedings of the National Academy of Sciences*, 109(1):227–232, 2012.
- [38] I. D. Couzin, C. C. Ioannou, G. Demirel, T. Gross, C. J. Torney, A. Hartnett, L. Conradt, S. A. Levin, and N. E. Leonard. Uninformed individuals promote democratic consensus in animal groups. *Science*, 334(6062):1578–1580, 2011.
- [39] N. L. Komarova and S. A. Levin. Eavesdropping and language dynamics. *Journal of Theoretical Biology*, 264(1):104–118, 2010.
- [40] E. Brown, J. Gao, P. Holmes, R. Bogacz, M. Gilzenrat, and J. D. Cohen. Simple neural networks that optimize decisions. *International Journal of Bifurcation and Chaos*, 15(3):803–826, 2005.
- [41] K. H. Britten, M. N. Shadlen, W. T. Newsome, and A. J. Movshon. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *The Journal of Neuroscience*, 12(12):4745–4765, 1992.
- [42] J. D. Schall. Neural basis of deciding, choosing and acting. *Nature Reviews Neuroscience*, 2(1):33–42, 2001.
- [43] D. R. J. Laming. *Information theory of choice-reaction times*. Academic Press, 1968.
- [44] R. Ratcliff. A theory of memory retrieval. *Psychological Review*, 85(2):59–108, 1978.
- [45] R. Ratcliff and J. N. Rouder. Modeling response times for two-choice decisions. *Psychological Science*, 9(5):347–356, 1998.
- [46] M. Shvartsman, V. Srivastava, and J. D. Cohen. A theory of decision making under dynamic context. In *Advances in Neural Information Processing Systems 28*, pages 2476–2484, Montréal, Canada, December 2015.
- [47] R. Duncan Luce. *Response times: Their role in inferring elementary mental organization*. Number 8. Oxford University Press, 1986.
- [48] R. S. Blum, S. A. Kassam, and H. V. Poor. Distributed detection with multiple sensors II. Advanced topics. *Proceedings of the IEEE*, 85(1):64–79, 1997.
- [49] P. Braca, S. Marano, V. Matta, and P. Willett. Asymptotic optimality of running consensus in testing binary hypotheses. *IEEE Transactions on Signal Processing*, 58(2):814–825, 2010.
- [50] R. Olfati-Saber, E. Franco, E. Frazzoli, and J. S. Shamma. Belief consensus and distributed hypothesis testing in sensor networks. In P. J. Antsaklis and P. Tabuada, editors, *Network Embedded Sensing and Control. (Proceedings of NESC’05 Workshop)*, Lecture Notes in

- Control and Information Sciences, pages 169–182. Springer, 2006.
- [51] D. Bajovic, D. Jakovetic, J. Xavier, B. Sinopoli, and J. M. F. Moura. Distributed detection via Gaussian running consensus: Large deviations asymptotic analysis. *IEEE Transactions on Signal Processing*, 59(9):4381–4396, 2011.
 - [52] K. Fitch and N. E. Leonard. Joint centrality distinguishes optimal leaders in noisy networks. *IEEE Transactions on Control of Network Systems*, 2015. conditionally accepted.
 - [53] F. Lin, M. Fardad, and M. R. Jovanović. Algorithms for leader selection in stochastically forced consensus networks. *IEEE Transactions on Automatic Control*, 59(7):1789–1802, 2014.
 - [54] S. Patterson and B. Bamieh. Leader selection for optimal network coherence. In *IEEE Conf. on Decision and Control*, pages 2692–2697, Atlanta, GA, December 2010.
 - [55] A. Clark, L. Bushnell, and R. Poovendran. A supermodular optimization framework for leader selection under link noise in linear multi-agent systems. *IEEE Transactions on Automatic Control*, 59(2):283–296, 2014.
 - [56] D. Pais and N. E. Leonard. Adaptive network dynamics and evolution of leadership in collective migration. *Physica D*, 267(1):81–93, 2014.
 - [57] C. J. Torney, S. A. Levin, and I. D. Couzin. Specialization and evolutionary branching within migratory populations. *Proceedings of the National Academy of Sciences*, 107(47):20394–20399, 2010.
 - [58] V. Guttal and I. D. Couzin. Social interactions, information use, and the evolution of collective migration. *Proceedings of the National Academy of Sciences*, 107(37):16172–16177, 2010.
 - [59] D. Pais, C. H. Caicedo-Núñez, and N. E. Leonard. Hopf bifurcations and limit cycles in evolutionary network dynamics. *SIAM Journal on Applied Dynamical Systems*, 11(4):1754–1784, 2012.
 - [60] T. McMillen and P. Holmes. The dynamics of choice among multiple alternatives. *Journal of Mathematical Psychology*, 50(1):30–57, 2006.
 - [61] V. P. Dragalin, A. G. Tartakovsky, and V. V. Veeravalli. Multihypothesis sequential probability ratio tests: I. Asymptotic optimality. *IEEE Transactions on Information Theory*, 45(7):2448–2461, 1999.
 - [62] S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Machine Learning*, 5(1):1–122, 2012.
 - [63] C. R. Reid, H. MacDonald, R. P. Mann, J. A. R. Marshall, T. Latty, and S. Garnier. Decision-making without a brain: How an amoeboid organism solves the two-armed bandit. *Journal of The Royal Society Interface*, 13(119):20160030, 2016.
 - [64] J. R. Krebs, A. Kacelnik, and P. Taylor. Test of optimal sampling by foraging great tits. *Nature*, 275(5675):27–31, 1978.
 - [65] V. Srivastava, P. Reverdy, and N. E. Leonard. Optimal foraging and multi-armed bandits. In *Allerton Conf. on Communications, Control and Computing*, pages 494–499, Monticello, IL, USA, October 2013.
 - [66] P. Reverdy, V. Srivastava, and N. E. Leonard. Modeling human decision making in generalized Gaussian multiarmed bandits. *Proceedings of the IEEE*, 102(4):544–571, 2014.
 - [67] R. C. Wilson, A. Geana, J. M. White, E. A. Ludvig, and J. D. Cohen. Humans use directed and random exploration to solve the explore-exploit dilemma. *Journal of Experimental Psychology: General*, 143(6):2074–2081, 2014.
 - [68] P. Landgren, V. Srivastava, and N. E. Leonard. On distributed cooperative decision-making in multiarmed bandits. In *European Control Conference*, pages 243 – 248, Aalborg, Denmark, June 2016.
 - [69] P. Landgren, V. Srivastava, and N. E. Leonard. Distributed cooperative decision-making in multiarmed bandits: Frequentist and Bayesian algorithms. In *IEEE Conf. on Decision and Control*, pages 167–172, Las Vegas, NV, December 2016.
 - [70] A. M. Hein, F. Carrara, D. R. Brumley, R. Stocker, and S. A. Levin. Natural search algorithms as a bridge between organisms, evolution, and ecology. *Proceedings of the National Academy of Sciences*, 113(34):9413–9420, 2016.