

Information sharing for a coordination game in fluctuating environments

Keith Paarporn^a, Ceyhun Eksin^b, Joshua S. Weitz^{a,c,d}

^a*School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332*

^b*Industrial & Systems Engineering Department, Texas A&M University, College Station, TX 77843*

^c*School of Biological Sciences, Georgia Institute of Technology, Atlanta, GA 30332*

^d*School of Physics, Georgia Institute of Technology, Atlanta, GA 30332*

Abstract

Collective action dilemmas pervade the social and biological sciences - from human decision-making to bacterial quorum sensing. In these scenarios, individuals sense cues from the environment to adopt a suitable phenotype or change in behavior. However, when cues include signals from other individuals, then the appropriate behavior of each individual is linked. Here, we develop a framework to quantify the influence of information sharing on individual behavior in the context of two player coordination games. In this framework, the environment stochastically switches between two states, and the state determines which one of two actions players must coordinate on. Given a stochastically switching environment, we then consider two versions of the game that differ in the way players acquire information. In the first model, players independently sense private environmental cues, but do not communicate with each other. We find there are two types of strategies that emerge as Nash equilibria and fitness maximizers - players prefer to commit to one particular action when private information is poor, or prefer to employ phenotypic plasticity when it is good. The second model adds an additional layer of communication, where players share social cues as well. When the quality of social information is high, we find the socially optimal strategy is a novel “majority logic” strategy that bases decision-making on social cues. Our game-theoretic approach offers a principled way of investigating the role of communication in group decision-making under uncertain conditions.

Keywords: Information sharing, Coordination, Uncertain environments, Game theory

1. Introduction

In stochastically fluctuating environments, organisms adapt their behavior to increase their chances of survival. For instance, bacteria process information from extracellular cues to reduce their uncertainty about the environmental state and employ stochastic phenotype switching in proportion to the remaining uncertainty (Lopez et al., 2009; Perkins and Swain, 2009). However in many circumstances, coordination of behavior with others directly influences survival. In bacteria, some cooperative behaviors are facilitated by quorum sensing, a cell-to-cell communication system where individual cells secrete and sense autoinducer molecules to obtain information about the environment and to gauge local cell density (Miller et al., 2001; Henke and Bassler, 2004). Quorum sensing can induce formation of biofilms for protection against a host’s immune system, secretion of virulence factors to consolidate colonization of the host, motility, and many other behaviors (West et al., 2007; Nadell et al., 2008; de Kievit and Iglewski, 2000; Atkinson et al., 2006). In honeybee and ant colonies, individuals process and share information to collectively reach an informed collective decision about

Email addresses: kpaarporn@gatech.edu (Keith Paarporn), ceyhuneksin@gmail.com (Ceyhun Eksin), jsweitz@gatech.edu (Joshua S. Weitz)

the best nesting site (Franks et al., 2002; Pratt et al., 2002). The individual-level mechanisms that produce collective behaviors in animal groups is an area of ongoing research (Couzin, 2009; Sasaki and Pratt, 2018). Inspired by these examples, the aim of this paper is to develop a game-theoretic framework in which to study the individual-level decision-making processes which produce collective behavior under environmental uncertainty and noisy communication.

There is an extensive body of work on individual (as opposed to group) decision-making in fluctuating environments (Perkins and Swain, 2009; Lopez et al., 2009). In many scenarios, an individual must match its phenotype or behavior to changing conditions by using sensory cues from the environment, signals from other individuals in the population, or both (Lachmann et al., 2000; Kussell and Leibler, 2005; Rivoire and Leibler, 2011; Donaldson-Matasci et al., 2010, 2013). An individual employing the optimal bet-hedging strategy diversifies behaviors at frequencies that mirror the posterior knowledge of environmental fluctuations. These results quantify an information-theoretic connection between the optimal population-level growth rate and the amount of information about the environment available to the individual. However, the resulting individual fitness is independent of the actions of others in the population. Hence, these models do not address the interplay between information and collective decision-making in fluctuating environments. In bacterial quorum sensing, the autoinducer signaling molecules that individual cells send serve as social cues that indicate local cell density (Miller et al., 2001; Nadell et al., 2008; Popat et al., 2015). Substantial experimental research has been done in recent years to unravel the mechanisms of this complex communication system. Nonetheless, questions remain regarding why such communication systems are utilized, particularly from an evolutionary standpoint (Whiteley et al., 2017).

Game theory offers a framework to explain rational behaviors when an individual's well-being depends on the actions of others. To understand the role of communication in collective decision-making under uncertainty, we recognize there are two components to the decision-making process: 1) a *communication system*, or the way individuals acquire information and 2) *strategies*, or the way individuals use acquired information to make a decision. In this vein, two recent works have attempted to understand the role of communication in group coordination under uncertain fluctuating environments from a game-theoretic perspective. Pacheco et al. (2015) studies the evolutionary outcomes of communication systems in an N -person volunteer's dilemma game, where all players adopt a majority rules strategy. Burgos and Polani (2016) consider how the choice of communication systems affect levels of cooperation between two populations of microbes in an information exchange game. In that model, each individual is assumed to follow a bet-hedging strategy. Both works assume the players' strategies are fixed while the communication systems are evolvable, and do not investigate how players may adapt their strategies to a given communication system.

In this paper, we study a two-player, two-action game where the environment stochastically switches between two possible states. The state determines game payoffs, where the players must coordinate on the correct action corresponding to the environment. Here, we study two versions of this game in order to highlight the value of information sharing. In the first, players only receive a private cue from the environment. In the second, players receive a shared social cue in addition to private cues. By considering social interactions and their effects on a single stage payoff, our game-theoretic model differs from bet-hedging models that focus on the link between variation in individual strategies and long-term growth rates of the population. Additionally, complementary to Pacheco et al. (2015) and Burgos and Polani (2016), we seek to identify the optimal strategies that promote coordination given the limitations of a fixed, noisy communication system of variable fidelity. Most notable is a "majority logic" type strategy that is socially optimal when environmental sensing has intermediate reliability, and information sharing is very reliable. As we show, this allows an individual to act upon their inference of the environmental state when it is validated by social information.

2. Model and methods

2.1. A coordination game in fluctuating environments

We consider a two action, two player game that is played repeatedly in stages $t = 1, 2, \dots$. We denote the set of players $\mathcal{N} \equiv \{1, 2\}$ with generic member i , and the set of actions $\mathcal{X} \equiv \{A, B\}$. At each stage t , the environment $E(t)$ takes one of two possible states, e_A or e_B . If $E(t) = e_A$ then $E(t+1) = e_B$ with probability v_{AB} , and if $E(t) = e_B$ then $E(t+1) = e_A$ with probability v_{BA} . We assume that $E(1)$ is arbitrarily determined and the switching probabilities $v_{AB} > 0$ and $v_{BA} > 0$ are fixed. Hence, the environmental state evolves according to a two-state Markov chain that spends a fraction $v_A \equiv \frac{v_{BA}}{v_{BA}+v_{AB}}$ and $v_B \equiv \frac{v_{AB}}{v_{BA}+v_{AB}}$ of time in states e_A and e_B , respectively. Once $E(t)$ is determined, players select their actions $x_i(t) \in \mathcal{X}$. They do not know $E(t)$ with certainty, and the state $E(t)$ governs the players' utility functions at each stage. We are interested in the average payoff of the players given the stochastically switching environment.

We define the game that is played at each stage as follows. Both must coordinate on action A (B) if $E(t) = e_A$ (e_B) to receive a payoff $b_A > 0$ ($b_B > 0$). If they mis-coordinate, or if they coordinate on the incorrect action, the payoff is zero. The environment-dependent payoff matrices are illustrated in Figure 1. Note that the payoff value to either player is the same. Thus, a generic player's utility function is defined as

$$U(x_1, x_2, E(t)) \equiv \begin{cases} b_A & \text{if } x_1 = x_2 = A \text{ and } E(t) = e_A \\ b_B & \text{if } x_1 = x_2 = B \text{ and } E(t) = e_B \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Given the realization of the environment $E(t)$, a normal form game is played at stage t between the set of players \mathcal{N} equipped with action set \mathcal{X} , and with utility function $U(\cdot, \cdot, E(t))$. We denote this normal form game by the triple $\mathcal{G} \equiv (\mathcal{N}, \mathcal{X}, U(\cdot, \cdot, E(t)))$, following standard game-theoretic notation. In the following, we introduce two models where players receive noisy information about $E(t)$, and present our definition of fitness, which we refer to as *time-averaged payoffs*.

2.2. The game \mathcal{G}_p with environmental sensing

At each stage t , suppose each player ($i = 1, 2$) independently senses the environment $E(t)$ by receiving a private cue $\alpha_i(t) \in \{e_A, e_B\}$. The cue $\alpha_i(t)$ matches the true environmental state with probability $p \in [1/2, 1]$, and mismatches with probability $1 - p$.

$$P(\alpha_i(t) = X | E(t) = Y) = \begin{cases} p & \text{if } X = Y \\ 1 - p & \text{if } X \neq Y \end{cases}, \quad X, Y \in \{e_A, e_B\} \quad (2)$$

Thus, $\alpha_i(t)$ is the output of a binary noisy channel of fidelity p whose input is $E(t)$. We will refer to the parameter p as the sensing fidelity.

A *strategy* is a mapping from the set of cues $\{e_A, e_B\}$ to the set of actions \mathcal{X} . In other words, a strategy is a contingency plan or policy that a player adopts that assigns the action to take upon receiving a particular cue. Hence, it is a description of how a player makes informed decisions. Since there are two possible cues and two actions, each player can choose among $2^2 = 4$ strategies. We assume player i chooses only one strategy s_i , which is fixed for all stages $t = 1, 2, \dots$. For notational convenience, we will denote a strategy by an ordered two-vector whose entries are either A or B . For instance, $s_i = [A, B]$ denotes the strategy where i plays action A when $\alpha_i(t) = e_A$ and action B when $\alpha_i(t) = e_B$. We write $s_i(e_A) = A$ and $s_i(e_B) = B$. We denote the set of all four strategies \mathcal{S}^4 . When the context is clear, we will also omit the notation (t) indicating variables realized at stage t .

The list of all four strategies is given in Table 1. They are read "Only A" (OA), "Follow Cue" (FC), "FC bar" ($\overline{\text{FC}}$), and "Only B" (OB). When $p = 1/2$, s_{FC} and $s_{\overline{\text{FC}}}$ reduce to the strategy that uniformly randomizes between A and B .

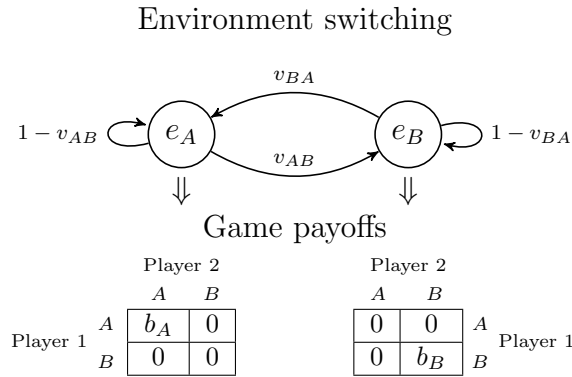


Figure 1: After each stage, the environment stochastically switches between e_A and e_B and determines the payoff matrix. In e_A (e_B), players need to coordinate on the A (B) action to receive a positive payoff b_A (b_B).

Remark 1. These strategies, with the exception of \overline{FC} , can be classified as a max likelihood estimate of the environmental state, depending on the environmental switching probabilities v_A, v_B and private cue fidelity p . Specifically, the strategy

$$s_{MLE}(\alpha_i) = \arg \max_{X=A,B} P(E = e_X | \alpha_i) \quad (3)$$

is precisely OA ($v_B < p < v_A$), OB ($v_A < p < v_B$), or FC ($v_A < p, v_B < p$). In this paper, we omit analysis of the s_{MLE} strategy because it simply corresponds to one strategy in Table 1 for any given set of parameters.

α_i / s_i	s_{OA}	s_{FC}	$s_{\overline{FC}}$	s_{OB}
e_A	A	A	B	B
e_B	A	B	A	B

Table 1: List of all strategies in the game with no information sharing.

A measure of fitness is the fraction of time the players coordinate on the correct action, weighted by the benefits b_A and b_B accordingly. We calculate this measure as follows. Consider the environment $E(t)$, which evolves according to the two-state Markov chain with stationary distribution (v_A, v_B) . Additionally, the players' cues are drawn independently of each other, but conditionally on the state $E(t)$. Since neither the cues nor player actions affect the environment switching probabilities, the fraction of time spent in the aggregate state (α_1, α_2, E) is given by the following stationary distribution

$$\pi_p(\alpha_1, \alpha_2, E) = \begin{cases} v_A P(\alpha_1 | e_A) P(\alpha_2 | e_A), & \text{if } E = e_A \\ v_B P(\alpha_1 | e_B) P(\alpha_2 | e_B), & \text{if } E = e_B \end{cases} \quad (4)$$

For instance, the value of the stationary distribution π_p at the entry (e_A, e_A, e_B) is $v_B(1-p)^2$. We define the *time-averaged payoff* as the expected utility with respect to the stationary distribution,

$$f_p(s_1, s_2) \equiv \sum_{\alpha_1, \alpha_2, E} \pi_p(\alpha_1, \alpha_2, E) U(s_1(\alpha_i), s_2(\alpha_j), E). \quad (5)$$

Via (5), we define the normal form game $\mathcal{G}_p \equiv (\mathcal{N}, \mathcal{S}^4, f_p)$ played between players \mathcal{N} equipped with strategy space \mathcal{S}^4 , and with utility function f_p for a given sensing fidelity $p \in [1/2, 1]$. Figure 2

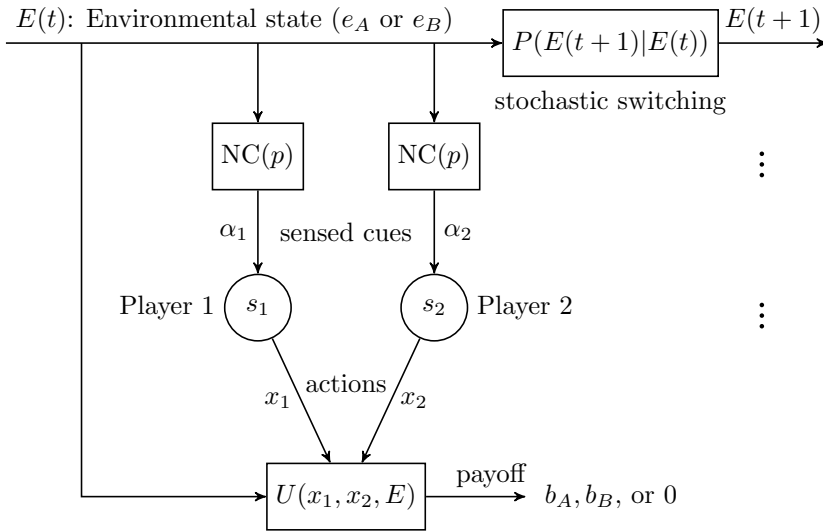


Figure 2: Diagram of the game with no information sharing \mathcal{G}_p . At stage t , each player independently senses a cue α_i directly from the environment $E(t)$ through a binary noisy channel with fidelity p , denoted here as $\text{NC}(p)$. That is, $\alpha_i = E(t)$ with probability p , and $\alpha_i \neq E(t)$ with probability $1 - p$. Player i 's strategy $s_i(\alpha_i)$ determines its action $x_i \in \{A, B\}$. The payoff $U(x_1, x_2, E)$ (eq. (1)) to both players is either b_A , b_B , or 0, which is determined by the current actions of both players and the current environment. The environmental state stochastically switches to $E(t + 1)$ at stage $t + 1$.

shows a diagram of the stage game and information system that underlies \mathcal{G}_p . We can represent \mathcal{G}_p with the following 4×4 matrix of time-averaged payoffs $f_p(s_1, s_2)$,

$$\begin{array}{c}
 \begin{array}{cccc}
 & s_{OA} & s_{FC} & s_{\overline{FC}} & s_{OB} \\
 s_{OA} & c_A & c_A p & c_A \bar{p} & 0 \\
 s_{FC} & c_A p & (c_A + c_B) p^2 & (c_A + c_B) p \bar{p} & c_B p \\
 s_{\overline{FC}} & c_A \bar{p} & (c_A + c_B) p \bar{p} & (c_A + c_B) \bar{p}^2 & c_B \bar{p} \\
 s_{OB} & 0 & c_B p & c_B \bar{p} & c_B
 \end{array} \\
 \end{array} \quad (6)$$

where $\bar{p} \equiv 1 - p$ and

$$c_A \equiv b_A v_A \text{ and } c_B \equiv b_B v_B \quad (7)$$

are the relative benefits of each environmental state. Note the payoff matrix is symmetric, which gives 10 unique strategy profiles. Therefore, the identity of the players do not matter. An expression to calculate each entry of (6) is given in Appendix A of the SI. We note that a “normalized” fitness value for all the entries above can be attained by dividing by c_B . This captures all relative payoffs in terms of a single environment-related parameter,

$$\kappa \equiv \frac{c_A}{c_B} \quad (8)$$

We call κ the *ratio of relative benefits*, and it is positive and nonzero. We will later see in Section 3 that κ is useful for parameterizing Nash equilibrium and fitness maximizer regions.

Remark 2. The average payoff (5) can also be viewed as the *ex-ante* expected utility in the one-shot Bayesian game \mathcal{B}_p consisting of the set of players \mathcal{N} , action set A , external states $\{e_A, e_B\}$, type space $\{e_A, e_B\}$, utility U (eq. (1)), beliefs $P(\alpha_1, \alpha_2, E)$, and a common prior on E , $P(e_A) = v_A$ and $P(e_B) = v_B$. The ex-ante expected utility is defined as the expectation of U with respect to

the belief $P(\alpha_1, \alpha_2, E)$, which is the induced probability distribution over the aggregate state of the world and coincides with (4). In an ex-ante setup, players evaluate their utilities before receiving their signal α_i , and therefore must reason about the possible environmental states, the signal of the other player, and its own signal. The (normal form) Nash equilibrium solution concept applied to the game \mathcal{G}_p coincides with the definition of ex-ante Bayesian Nash equilibrium (BNE) of \mathcal{B}_p . A Bayesian Nash equilibrium describes a state of rationality where no player can profitably deviate by changing its strategy given its belief of the world. We focus our attention on the normal-form formulation \mathcal{G}_p in this paper, keeping in mind that a Nash equilibrium in \mathcal{G}_p can also be interpreted as a Bayesian Nash equilibrium of \mathcal{B}_p . A treatment of Bayesian games can be found in Ch. 6 of Vega-Redondo (2003).

2.3. The game \mathcal{G}_{pq} with information sharing

The game \mathcal{G}_p is extended by allowing players to share their private cues with each other before deciding on an action. At stage t , after players sense $\alpha_i(t)$, player i sends a social cue $\beta_j(t)$ to player j ($j \neq i$), which matches i 's private cue $\alpha_i(t)$ with probability $q \in [1/2, 1]$.

$$P(\beta_i(t) = X | \alpha_j(t) = Y) = \begin{cases} q & \text{if } X = Y \\ 1 - q & \text{if } X \neq Y \end{cases}, \quad X, Y \in \{e_A, e_B\} \quad (9)$$

Thus, $\beta_i(t)$ is the output of a binary noisy channel of fidelity q whose input is $\alpha_j(t)$. We will refer to parameter q as the *sharing fidelity*. We assume here that players signal honestly, i.e. each player attempts to share their true private cue. This is a reasonable assumption because the players' interests are aligned - they must attempt to coordinate. In different contexts where player interests conflict, dishonest signalling becomes a rational alternative. For example, male fiddler crabs with inferior claws can bluff fighting ability to ward off other males (Backwell et al., 2000).

Player i 's information is now composed of the pair $y_i(t) \equiv (\alpha_i(t), \beta_i(t)) \in \{e_A, e_B\}^2$. Similarly for this model, a strategy is a mapping from the set of information pairs $\{e_A, e_B\}^2$ to actions, so player i can now choose among $2^4 = 16$ strategies. We represent strategies as ordered four-vectors whose entries are A or B . For instance, $s_i = [A, B, B, A]$ is the strategy where i plays action A when $y_i = (e_A, e_A)$, B when $y_i = (e_A, e_B)$, B when $y_i = (e_B, e_A)$, and A when $y_i = (e_B, e_B)$. We note that the four strategies available in \mathcal{G}_p are also strategies in this game. They are now represented by the vectors $s_{OA} = [A, A, A, A]$, $s_{FC} = [A, A, B, B]$, $s_{\overline{FC}} = [B, B, A, A]$, $s_{OB} = [B, B, B, B]$. These four strategies base decision-making either on no information at all (OA and OB), or only on the private cue α_i (FC and \overline{FC}). The 12 new strategies base decisions on both private and shared signals. The set of all 16 strategies is denoted \mathcal{S}^{16} .

Similarly to (4), we calculate the time-averaged payoff by considering the aggregate state (y_1, y_2, E) , which occurs a fraction of the time according to the following stationary distribution

$$\pi_{pq}(y_1, y_2, E) = \begin{cases} v_A P(y_1, y_2 | e_A), & \text{if } E = e_A \\ v_B P(y_1, y_2 | e_B), & \text{if } E = e_B \end{cases} \quad (10)$$

One can expand $P(y_1, y_2 | E) = P(\alpha_1 | E)P(\alpha_2 | E)P(\beta_1 | \alpha_2)P(\beta_2 | \alpha_1)$. For instance, the value of the stationary distribution at the entry $((e_A, e_A), (e_A, e_B), e_A)$ is $v_A p^2 q(1 - q)$. The *time-averaged payoff* at $(p, q) \in [1/2, 1]^2$ is defined as

$$f_{pq}(s_1, s_2) \equiv \sum_{y_1, y_2, E} \pi_{pq}(y_1, y_2, E) U(s_1(y_1), s_2(y_2), E). \quad (11)$$

Via (11) and following the arguments from Remark 2, we define the normal form game $\mathcal{G}_{pq} \equiv (\mathcal{N}, \mathcal{S}^{16}, f_{pq})$ played between the set of players \mathcal{N} equipped with strategy space \mathcal{S}^{16} , and with utility function f_{pq} for a given pair of sensing and sharing fidelities $(p, q) \in [1/2, 1]^2$. Figure 3 shows

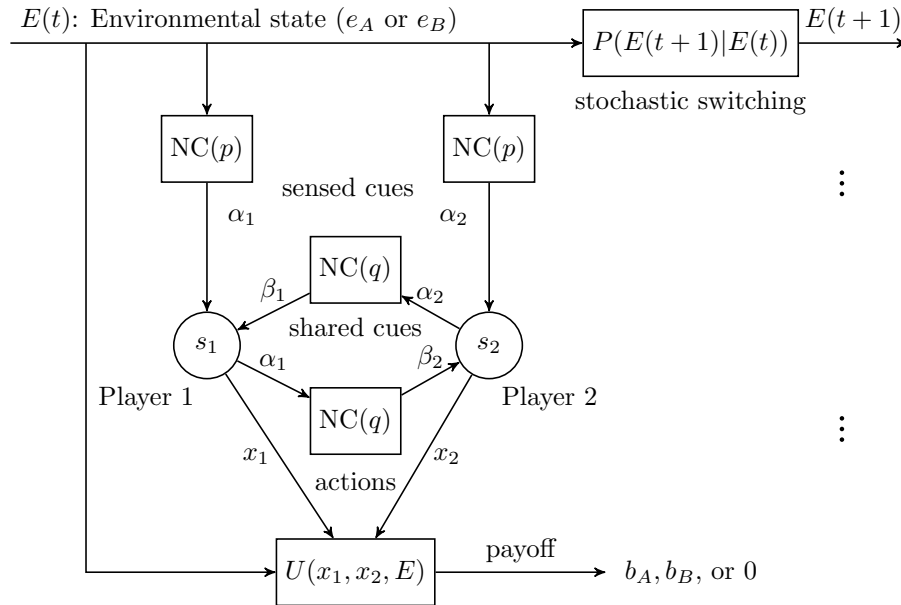


Figure 3: Diagram of the game with information sharing \mathcal{G}_{pq} . At stage t , each player independently senses a private cue α_i from the environment $E(t)$ through a binary noisy channel with fidelity p , denoted here as $\text{NC}(p)$. Then, each player receives a signal β_j from the other's private cue through a separate noisy channel of fidelity q . Given player i 's information (α_i, β_i) , its action is determined by its strategy, $s_i(\alpha_i, \beta_i) = x_i$. The payoff $U(x_1, x_2, E)$ (eq. (1)) to both players is either b_A , b_B , or 0, which is determined by the actions of both players and the current environment. The environmental state stochastically switches to $E(t+1)$ at stage $t+1$.

a diagram of the stage game and communication system that underlies \mathcal{G}_{pq} . An expression to calculate each entry of the resulting 16×16 symmetric payoff matrix is given in Appendix A of the SI. Given the symmetry of the payoff matrix, there are 136 unique strategy profiles. A summary of all relevant parameters that define the games \mathcal{G}_p and \mathcal{G}_{pq} is listed in Table 2.

Parameter	Description	Notes
e_X	environmental state, $X = \{A, B\}$	
p	environment sensing fidelity	$\in [1/2, 1]$
q	sharing fidelity	$\in [1/2, 1]$
α_i	private environment cue for player i	$\in \{e_A, e_B\}, i = 1, 2$
β_i	social cue player i receives from player j	$\in \{e_A, e_B\}, j \neq i$
b_X	payoff of coordinating in $e_X, X = \{A, B\}$	> 0
v_{AB}, v_{BA}	prob. of switching to e_B from e_A , vice versa	both $\in [1/2, 1]$
v_X	fraction of time spent in $e_X, X = \{A, B\}$	$v_A + v_B = 1$
c_X	relative payoff of coordinating in e_X	$= v_X b_X$
κ	ratio of relative payoffs	$= c_A/c_B$

Table 2: List of parameters that define the games \mathcal{G}_p and \mathcal{G}_{pq} .

Remark 3. We note that the payoff matrices of \mathcal{G}_p and \mathcal{G}_{pq} are symmetric. Hence, the identity of the player (player 1 or 2) does not matter. Such games fall into the class of potential games (Monderer and Shapley, 1996; Sandholm, 2001). A game is a potential game if the incentives of the players align with a global *potential* function. Here, the change in payoff from a unilateral deviation of a single player is equivalent to the change in global potential, given all other players remain the

same. In our formulation, the potential function is the average player fitness, or welfare.

3. Results

We consider the Nash equilibrium solution concept and fitness maximizing strategies as notions of rational and optimal behavior, respectively. The comparison of the fitness maximizing strategies between the two models (\mathcal{G}_p and \mathcal{G}_{pq}) presented offers a principled way to investigate the role of sensing and communication in collective decision-making in systems such as quorum sensing.

Specifically, we are interested in how the quality of the communication system, defined by the parameter p for \mathcal{G}_p and (p, q) for \mathcal{G}_{pq} , dictates which strategy profiles are Nash equilibria and fitness maximizers. A Nash equilibrium is the classical solution concept in game theory which describes a rational steady-state configuration where no player has an incentive to deviate from its strategy. A Nash equilibrium in \mathcal{G}_p is a strategy profile (s_1^*, s_2^*) that satisfies

$$\begin{aligned} f_p(s_1^*, s_2^*) &\geq f_p(s_1, s_2^*) \\ f_p(s_1^*, s_2^*) &\geq f_p(s_1^*, s_2) \end{aligned} \quad (12)$$

for all $s_1 \in \mathcal{S}^4$, $s_1 \neq s_1^*$, and for all $s_2 \in \mathcal{S}^4$, $s_2 \neq s_2^*$. A strict Nash equilibrium is a Nash equilibrium satisfying (12) with strict inequality. We say a strategy profile (\hat{s}_1, \hat{s}_2) is a *fitness maximizer* at p if it satisfies

$$f_p(\hat{s}_1, \hat{s}_2) = \max_{s_1, s_2 \in \mathcal{S}^4} f_p(s_1, s_2). \quad (13)$$

The same definitions above apply for \mathcal{G}_{pq} , where f_p is replaced by f_{pq} , and \mathcal{S}^4 by \mathcal{S}^{16} . Because the payoff matrices of \mathcal{G}_p and \mathcal{G}_{pq} are symmetric, the following fact holds: in \mathcal{G}_p and \mathcal{G}_{pq} , the fitness maximizer(s) is necessarily a Nash equilibrium. However, the converse is not true. We note that in our model, there is no fitness cost to having higher sensing and sharing fidelities p and q . Our aim is not to investigate such evolutionary tradeoffs, but to identify the types of strategies that ensure coordination in fluctuating environments given that the players utilize a communication system of quality $(p, q) \in [1/2, 1]^2$.

3.1. Nash equilibria and fitness maximizers in \mathcal{G}_p

The base game \mathcal{G} (see Eq. (1)) admits a single strict Nash equilibrium at the correct coordinated action. We might expect the game \mathcal{G}_p , represented by (6), also has a coordination structure. That is, the players are better off if they coordinate on the same strategy in \mathcal{S}^4 . In other words, they should play the same action given they receive the same signal. We can show that this intuition is indeed correct and state the following fact: For $p \in [1/2, 1]$, all Nash equilibria and fitness maximizers of \mathcal{G}_p are necessarily symmetric strategy profiles, i.e. $s_1^* = s_2^*$. The proof of this fact, given in Appendix B of the SI, relies on showing that for any asymmetric strategy profile, there is one player that can switch to the other player's strategy to improve the fitness f_p .

Due to this result, we can limit our analysis to the four symmetric strategy profiles. To refer to symmetric strategy profiles (s, s) for the four strategies we simply write "OA" to denote (s_{OA}, s_{OA}) , and similarly for the other three. Their resulting fitnesses correspond to the diagonal entries of (6).

We find that OA and OB are always Nash equilibria regardless of the value for sensing fidelity p and the relative benefits c_A, c_B . We find that $\overline{\text{FC}}$ is not a Nash equilibrium for any set of parameters. Figure 4 (left) illustrates the Nash equilibrium region of FC. We can characterize the equilibrium region of FC in the game \mathcal{G}_p with respect to just two parameters, $(p, \kappa) \in [1/2, 1] \times (0, \infty)$, where we recall $\kappa = c_A/c_B$ is the ratio of relative benefits. This is possible by normalizing the fitness, f_p/c_B , and applying (12) to solve for the region's conditions on p and κ .

Figure 4 (right) illustrates the regions where OA, OB, and FC are fitness maximizers. While OA and OB are Nash equilibria everywhere, there is a unique fitness maximizer (OA, OB, or FC),

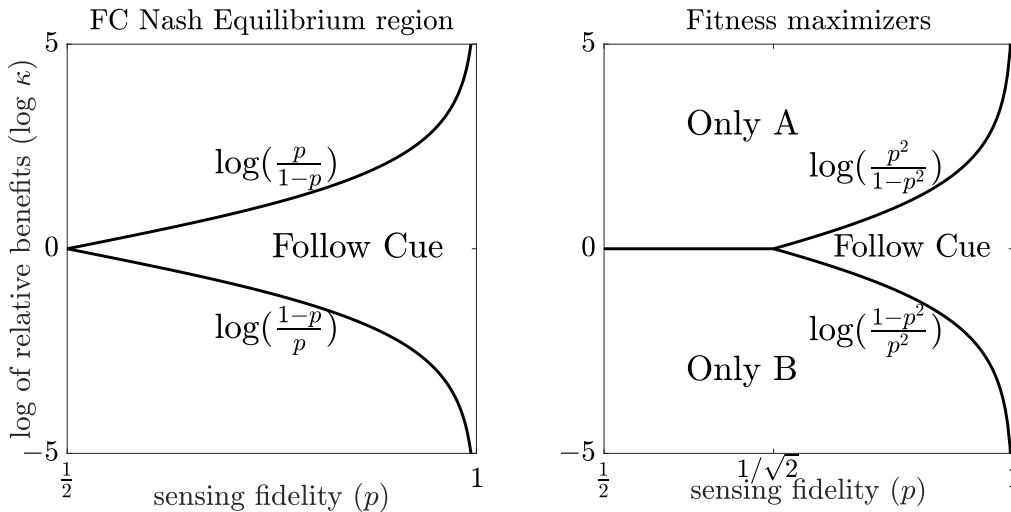


Figure 4: Characterization of strategies in the game \mathcal{G}_p with respect to the parameter space (p, κ) with $\kappa \equiv \frac{c_A}{c_B}$ ranging in log scale from -5 to 5. Log scale is used to show symmetry between the ranges $0 < \kappa < 1$ and $1 < \kappa < \infty$. The quantity κ is the ratio of relative benefits between environment e_A and e_B . (Left) The region where FC is a strict Nash equilibrium : $\frac{1-p}{p} < \kappa < \frac{p}{1-p}$. The strategies OA and OB are Nash equilibria everywhere. (Right) Disjoint regions where the strategies OA, OB, and FC are fitness maximizers. OA: $\kappa > \max\left\{1, \frac{p^2}{1-p^2}\right\}$. OB: $\kappa < \min\left\{1, \frac{1-p^2}{p^2}\right\}$. FC: $\frac{1-p^2}{p^2} < \kappa < \frac{p^2}{1-p^2}$

except for on the boundaries dividing each region, for any given (p, κ) value in $[1/2, 1] \times (0, \infty)$. On the boundary lines, the fitnesses of the strategy profiles that are separated coincide. We note here that the region where FC maximizes fitness is a subset of its Nash equilibrium region. This reveals regions in the parameter space where FC is a suboptimal Nash equilibrium to the OA/OB strategies.

The fitness for FC increases quadratically in p : $f_p(\text{FC}) = (c_A + c_B)p^2$. Hence, with higher sensing fidelity, players are better able to infer the correct environmental state and select the correct action. The p^2 term appears because each player is independently sensing E , and coordination depends on both players independently receiving the correct cue which occurs with probability p^2 . The value of employing FC diminishes as $\kappa \rightarrow \infty$ or 0. In either extreme, one environment is favored over the other. Either one occurs far more frequently over the long run, or its coordination benefit outweighs that of the other state. Hence, players do best by adopting either OA or OB (corresponding to the environment that is favored) instead of FC in these situations. FC is most desirable when $\kappa \approx 1$, where either the environment fluctuates frequently and $b_A \approx b_B$, or a rare environment offers an enormous coordination benefit compared to the other. These are situations where acting on knowledge of the environment is most crucial. By playing OA or OB, players miss out on half of the fitness benefit opportunity whereas players employing FC are able to adapt to changing conditions.

The s_{FC} strategy possesses similarities to the optimal individual bet-hedging strategy that maximizes long-term population growth rate in fluctuating environments (Kussell and Leibler, 2005; Donaldson-Matasci et al., 2008, 2010) because it adapts actions based on the player's inference on the environmental state. However in our model, the social context of coordination renders it suboptimal when private information is unreliable (low p), though it is optimal for sufficiently high values of p .

3.2. Nash equilibria and fitness maximizers in \mathcal{G}_{pq}

We now turn to the game $\mathcal{G}_{pq} = (\mathcal{N}, \mathcal{S}^{16}, f_{pq})$. The Nash condition remains the same as in (12), except the time-averaged payoff is now given by f_{pq} (11). In comparison to \mathcal{G}_p , the parameter space of the information sharing game \mathcal{G}_{pq} has another dimension, the social fidelity q . We search for the Nash equilibrium strategy profiles and fitness maximizers of \mathcal{G}_{pq} over the parameter space $(p, q, \kappa) \in [1/2, 1]^2 \times (0, \infty)$. First, we observe that unlike \mathcal{G}_p , the coordination structure in the new set of strategies \mathcal{S}^{16} is not preserved in \mathcal{G}_{pq} . Indeed, we find that the asymmetric strategy profile (s_1, s_2) with $s_1 = s_{\text{FC}} = [A, A, B, B]$ and $s_2 = [A, B, A, B]$ is a Nash equilibrium in a region of low sensing fidelity p and high sharing q (see SI Appendix C for derivation). Hence, not all Nash equilibria are necessarily symmetric strategy profiles for all $(p, q) \in [1/2, 1]^2$, and we cannot restrict our search for Nash equilibria and fitness maximizers to symmetric strategy profiles. In essence, players may now select new strategies that utilize the social cues as a means to coordinate.

There are 136 unique strategy profiles (by symmetry). In Figure 5, we display the multiplicity of Nash equilibria across the range of fidelity parameters p and q . Here, we do not count OA and OB because they are always Nash equilibria. We find the maximum number of Nash equilibria in the upper left region of the parameter space $(p, q) \in [1/2, 1]^2$, where private information is unreliable but information sharing has high fidelity. This suggests reliable social cues are utilized as a means to coordinate (see Section 4 for further discussion).

The strategy profiles that maximize fitness are shown in Figure 6 in the parameter space $(p, q) \in [1/2, 1]^2$ and for cross-sections of κ . We note that the set of fitness maximizers consist only of symmetric strategy profiles. We find that OA, OB, and FC are fitness maximizers, along with a new type of strategy we term ‘‘Majority Logic’’ (ML). When environment A is favored ($\kappa > 1$), ML_A appears as a fitness maximizer, and similarly ML_B when $\kappa < 1$. These strategies are written

$$\begin{aligned} s_{\text{ML}_A} &\equiv [A, A, A, B] \\ s_{\text{ML}_B} &\equiv [A, B, B, B]. \end{aligned} \tag{14}$$

Recall that in our strategy vector notation, $[A, A, A, B]$ is the strategy where the player chooses action A when $y_i = (e_A, e_A)$, action A when $y_i = (e_A, e_B)$, action A when $y_i = (e_B, e_A)$, and action B when $y_i = (e_B, e_B)$. The majority logic strategies ML_A and ML_B were not previously available in \mathcal{G}_p . Their fitnesses increase quadratically in q , and their functional forms are given in Appendix A.

The notation ML_A, ML_B differentiates which action is assigned, A or B , to the cues $y_i = (e_A, e_B)$ and (e_B, e_A) , i.e. the middle two entries in the strategy vector. We will write just ‘‘ML’’ when generally speaking of the majority logic strategies (e.g. when $\kappa = 1$).

The region where the ML strategies maximize fitness is characterized by intermediate sensing fidelity p and high sharing fidelity q . The shared cues β_i are highly accurate so each player will have reliable knowledge of the other’s private cue. Our interpretation of why ML thrives in this regime is as follows. With reliable information sharing, players can detect when their private cues agree and when they disagree. As an example, consider when both players employ strategy $s_{\text{ML}_A} = [A, A, A, B]$ and suppose environment A is favored over B ($\kappa > 1$). When both of their private cues coincide, i.e. $y_1 = y_2 = (e_A, e_A)$ or (e_B, e_B) , they choose the correct action. This occurs with overall probability p^2q^2 . When they believe their private cues disagree with each other (e.g. $y_1 = (e_A, e_B)$ and $y_2 = (e_B, e_A)$), they will both play action A , and obtain a positive payoff with overall probability $v_{AP}(1-p)q^2$. Moreover, the above situations will occur far more frequently than players obtaining polarized beliefs (e.g. $y_1 = (e_A, e_A)$ and $y_2 = (e_B, e_B)$), where they mis-coordinate actions with probability $p(1-p)(1-q)^2$. The ML_A or ML_B strategy profile appears as the fitness maximizer depending on which environment is favored (Figure 6 right).

The s_{ML_A} strategy vector $[A, A, A, B]$ ($[A, B, B, B]$) differs from $s_{\text{FC}} = [A, A, B, B]$ and $s_{\text{OA}} = [A, A, A, A]$ by only one action, whereas s_{FC} differs from s_{OA} by two actions (similarly for s_{ML_B} and s_{OB}). Thus, we interpret s_{ML} to be a hybrid of s_{FC} and s_{OA} . It acts as an estimator of




Figure 5: The number of Nash equilibria not including OA or OB, that exist across fidelity values $(p, q) \in [1/2, 1]^2$, for $\kappa = 1$ (Left) and $\kappa = 5$ (Right) in the game \mathcal{G}_{pq} . The data is numerically calculated by sweeping through $(p, q) \in [1/2, 1]^2$ in a uniform grid of spacing .001, and exhaustively verifying whether each 136 strategy profile satisfies its Nash condition, given in (12). We verify whether a strategy profile is a Nash equilibrium or not by using a simplified expression of the Nash condition in (12) that is amenable to numerical evaluation, given in Appendix C of the SI.

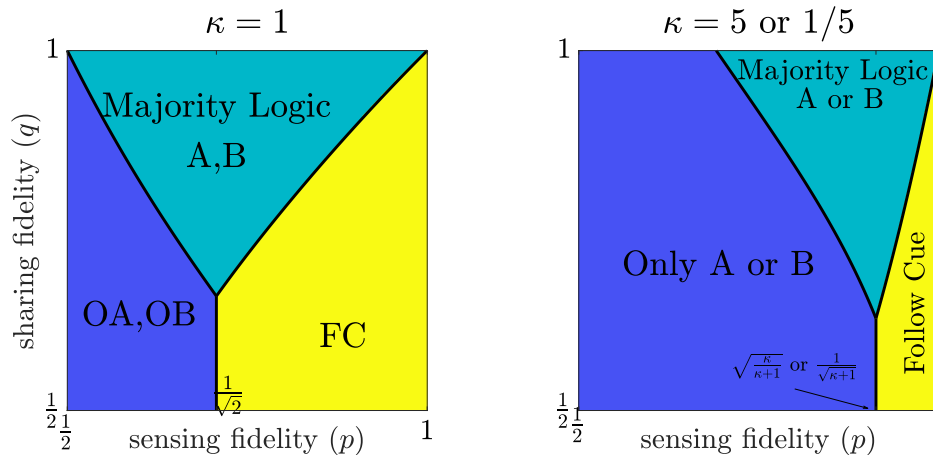


Figure 6: Three classes of strategies emerge as the fitness maximizers in \mathcal{G}_{pq} : 1) the pure strategies OA and OB, 2) follow cue (FC), and 3) majority logic (ML_A and ML_B). The regions are drawn by first sweeping $(p, q) \in [1/2, 1]^2$ in a uniform grid with spacing .001, and exhaustively searching all 136 strategy profiles for the fitness maximizer. Each grid cell is then filled with the unique color corresponding to the fitness-maximizing strategy profile. Upon observing the emergence of the three strategy classes, the boundaries are analytically solved by equating their fitnesses. Hence, results are accurate within a spacing tolerance of .001. (Left) When $\kappa = 1$, players do not prefer any environment over the other. Hence, OA and OB give the same fitness value and we indicate this by OA/OB (similarly for ML_A/ML_B). (Right) When $\kappa \neq 1$, the region boundaries are the same for both 5 and $\frac{1}{5}$ but the fitness maximizer is either OA or OB, depending on whether κ is greater or less than 1, respectively (similarly, either ML_A or ML_B).

the environment when the posterior belief on the environmental state is very high. This is the reasoning for the term ‘‘Majority Logic’’ - players act upon their inference of the environmental state only when their private information is validated by social information, i.e. when $y_i = (e_A, e_A)$ or (e_B, e_B) . Otherwise, when $y_i = (e_A, e_B)$ or (e_B, e_A) , the posterior is not as high and the player disregards its information altogether, playing the default action A .

3.3. The fitness value of information sharing

Majority logic stands alone among all possible strategies in \mathcal{G}_{pq} that can outperform the fitness maximizers of \mathcal{G}_p (OA, OB, and FC). However, recall from Figure 6 (Left, for example) that it requires the sharing fidelity q to be sufficiently high for a fixed sensing fidelity p . This critical threshold value, which we denote with $q_c(p)$, is the value of q above which the majority logic strategy is the fitness maximizer. Hence, the values of $q_c(p)$ are parameterized by the boundary lines that separate the ML strategies from OA/OB and FC in Figure 6. Full parameterizations of $q_c(p)$ are given in Appendix D of the SI.

The critical thresholds $q_c(p)$ when the ML strategies outperform the optimal strategies of \mathcal{G}_p suggests there is no value for players to share signals unless sharing fidelity is sufficiently high, $q > q_c(p)$. Higher sharing fidelity is needed for extreme values of p : $q_c(p)$ increases up to 1 as p decreases towards $1/2$, as well as when p increases towards 1. When p is near 1, players prefer FC because they are able to independently detect the correct environment with very high probability and act accordingly. When p is near $1/2$, private cues are effectively random because they contain no information about E . Consequently, the shared signals are also effectively random. Players then receive any of the four signals y_i with equal probability. By employing ML in this regime, players will miscoordinate more often than they would if they committed to OA or OB.

We are also interested in quantifying the fitness benefit of sharing signals over no sharing. Let

$$F_p(\kappa) \equiv \max_{X=OA,OB,FC} f_p(X) \quad (15)$$

be the maximum fitness attainable in the game \mathcal{G}_p at (p, κ) , and let

$$F_{pq}(\kappa) \equiv \max_{X=OA,OB,FC,ML_A,ML_B} f_{pq}(X) \quad (16)$$

be the maximum fitness attainable in the game \mathcal{G}_{pq} at (p, q, κ) . Then we define the *fitness value of information sharing* as

$$V(p, q, \kappa) \equiv \frac{F_{pq}(\kappa)}{F_p(\kappa)}, \quad (17)$$

which gives the ratio of the maximum fitness in \mathcal{G}_{pq} to the maximum fitness in \mathcal{G}_p at the parameters (p, q, κ) . By definition, $V(p, q, \kappa) = 1$ for $q \leq q_c(p)$. In other words, there is no fitness benefit to sharing signals when q does not exceed the threshold $q_c(p)$. A contour map of V is shown in Figure 7. When $\kappa = 1$ (Figure 7, Left), V is maximized at $p = 1/\sqrt{2}$ and $q = 1$, giving $V_{\max} = \sqrt{2}$ (using Eq. (A.12) and divide by c_B). Thus, information sharing can improve fitness by approximately 41% when the environment fluctuates frequently. For $\kappa > 1$, it is maximized at $q = 1$, $p = \sqrt{\kappa}/(1 + \kappa)$ (Figure 7, Right). The general form of $V_{\max}(\kappa)$ is the following piecewise continuous function,

$$V_{\max}(\kappa) = \begin{cases} \sqrt{\frac{\kappa}{\kappa+1}} \left(2 - \sqrt{\frac{\kappa}{\kappa+1}} (1 - \kappa^{-1}) \right) & \text{if } \kappa > 1 \\ \sqrt{2} & \text{if } \kappa = 1 \\ \frac{1}{\sqrt{\kappa+1}} \left(2 - \frac{1}{\sqrt{\kappa+1}} (1 - \kappa) \right) & \text{if } 0 < \kappa < 1 \end{cases} \quad (18)$$

and is plotted in Figure 8. As $\kappa \rightarrow 0$ or ∞ , the improvement ratio V_{\max} degrades as one environment becomes favored over the other. In these extreme scenarios, either OA or OB become optimal for increasingly larger regions.




Figure 7: The fitness value of information sharing $V(p, q, \kappa)$. (Left) When $\kappa = 1$, the maximum value of V is attained at $(p = 1/\sqrt{2}, q = 1, \kappa = 1)$, where information sharing leads to approximately a 41% increase in fitness over no sharing. (Right) When $\kappa = 5$, information sharing leads to approximately 15% increase in fitness at best ($p = \sqrt{\frac{\kappa}{1+\kappa}}, q = 1$).

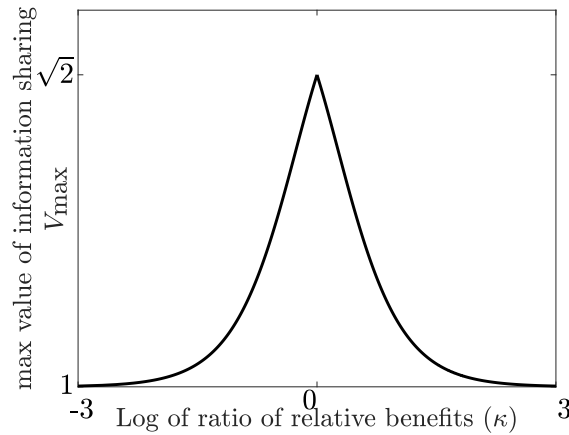


Figure 8: The maximum value of information sharing $V_{\max}(\kappa)$ (Eq. (18)), which is the value of $V(p, q, \kappa)$ attained at $q = 1$ and $p = 1/\sqrt{1+\kappa}$ when $\kappa < 1$ and $p = \sqrt{\kappa/(1+\kappa)}$ when $\kappa \geq 1$. The peak of fitness improvement occurs when no environment is favored over the other ($\kappa = 1$), where there is an improvement ratio of $\sqrt{2}$. The value for V_{\max} is $f_{pq}(\text{ML}_B)/c_B$ when $\kappa < 1$, and $f_{pq}(\text{ML}_A)/c_A$ when $\kappa > 1$ (they coincide when $\kappa = 1$). The fitness improvement ratio degrades as $\kappa \rightarrow 0$ or ∞ , i.e. when one environment is favored over the other.

We also note that $V(p, q, \kappa)$ increases quadratically in q for $q > q_c(p)$. Hence, the rate at which V increases with respect to q , $\frac{dV}{dq}$, increases linearly in q . This is due to a majority logic strategy being the fitness maximizer in this region. We also note that $V(p, q, \kappa)$ decreases as p deviates away from the intermediate value $\sqrt{\kappa/(1+\kappa)}$. At the extremes $p = 1/2$ and $p = 1$, the strategies *OA* and *FC* begin to dominate, respectively, as the line $q_c(p)$ tapers to 1.

4. Discussion

We have presented a two-player, two action coordination game in a fluctuating environment where players independently sense private cues from the environment and share their cues with each other. In our analysis, we found the optimal strategies that promote coordination across all levels of sensing and sharing fidelities (Figures 4 and 6). When individual sensing is very reliable, there is no need to share signals because players can accurately infer the environmental state independently. When sensing is unreliable, players prefer to ignore their information altogether and always commit to a single action. The “majority logic” strategy is optimal when private information has intermediate fidelity and social information has high fidelity. This strategy highlights the importance of information sharing because they are the only optimal strategies that utilize the shared social cues. Because these optimal strategies are derived from a first-principles approach, their appearance in our model offers insight into the mechanisms that maintain coordinated group behaviors.

The Majority Logic strategies strike a balance between when to use information as a predictor of the environment and when to use information as a means to coordinate. Essentially, it allows players to coordinate when their individual inferences (private information) about the environment conflict, and to choose the correct action when they agree. The interplay of private and social information in our model draws similarities to Condorcet’s jury theorem, where in King and Cowlshaw (2007), states that good reliability of private information is a requirement to effectively make group decisions based on social information in the large population limit. In their model, individuals disregard social information when private information is poor. When it is more accurate, individuals rely more on social information to make a decision. In our model, these two situations are akin to the fitness maximizing regions of *OA* and *OB* for poor sensing fidelity, and the region where *ML* thrives, respectively. These results also corroborate with controlled experimental lab work conducted on nine-spined sticklebacks (van Bergen et al., 2004). When private experience about foraging sites was 100% reliable, sticklebacks based foraging decisions only on private information. When it was less reliable, they followed social cues instead. Without perfect sensing capabilities, organisms need to rely on social information to survive (van Bergen et al., 2004; King and Cowlshaw, 2007; Pérez-Escudero and de Polavieja, 2011; Arganda et al., 2012; Miller et al., 2013).

In the region of unreliable private information and reliable social information (Fig 5, right), there is an abundance of Nash equilibria (ten) that are suboptimal to the *OA* and *OB* strategies. This suggests that in the context of coordination games, social cues serve as a coordination device rather than as an additional source of information about the environment. However, the strategies included in this region, which includes majority logic and its variants, will often coordinate on the wrong action. This is because the social cues carry no information about the environment, as the private cues are themselves uninformative. The best an individual can do to infer the environment is simply to guess. However, if both guess independently, they will only coordinate 1/4 of the time. If they play *OA*, they are guaranteed to coordinate for the fraction of time the environment spends in state e_A . Hence, committing to a single action corresponding to the most frequent environment (*OA* or *OB*) is the best the group can do when information about the environment is poor. Therefore in principle, information sharing is useful only when the information that is being shared is itself reliable.

Our game-theoretic model portrays situations where communicating individuals must coordinate behaviors in uncertain fluctuating environments. These situations pervade collective behaviors in

groups of organisms across the animal kingdom from social insects to bacterial colonies. The information flow in our model is particularly inspired by quorum sensing in bacterial populations. Such a communication system enables bacteria to display complex social behaviors (West et al., 2007). Therefore, game theory is a natural framework in which to study microbial decision-making to consolidate experimental understanding of the phenomena. Advancing such knowledge may also have therapeutic applications. For example, a deeper understanding of how bacteria communicate to form harmful biofilms presents an opportunity to develop methods to inhibit the fidelity of quorum sensing systems (Rutherford and Bassler, 2012; Popat et al., 2015).

Several questions remain in our study. In animal communication, better quality signalling entails increasing fitness costs (Bergstrom and Lachmann, 1997; Brown and Johnstone, 2001; Skyrms, 2010; Meacham et al., 2013; Huttegger et al., 2014). Our work has not yet investigated such evolutionary tradeoffs. Instead, we have presented a systematic, centralized analysis of the optimal strategies given a fixed, costless communication system. Consequently, we have yet to address whether the optimal strategies identified are stable in an evolutionary sense, with or without costly signalling.

Moving forward, several evolutionary dynamics can be applied to our model in the context of population games (Sandholm, 2010). Population games are a framework to describe the interactions between a well-mixed, continuous mass of agents that select from the same set of strategies. In the population, interactions between agents are probabilistic and pair-wise, which allows two-player normal form games to be represented as population games. Such a framework works well with our aim to describe a population of organisms that can adopt a variety of communication-based strategies, *e.g.*, quorum sensing bacteria. Furthermore, due to the potential structure of our games \mathcal{G}_p and \mathcal{G}_{pq} (see Remark 3), they admit evolutionary dynamics that have certain stability guarantees (Sandholm, 2001, 2010). Embedding our model into population games will be necessary to identify which local maximizers of average fitness are likely to be reached. In doing so, we hope that our model encourages the integration of social interactions and communication into efforts to understand coordination, cooperation, and conflict in complex environments.

Acknowledgements

This work is supported by ARO grant #W911NF-14-1-0402. We thank William H. Sandholm, Georgy Loginov, and William C. Ratcliff for their comments on the manuscript. We thank Marvin Whiteley and Dan Cornforth for insightful discussions. We thank two anonymous reviewers for their constructive feedback. K.P. and C.E. also thank Matthieu Bloch for insightful discussions.

Code availability

The code for our simulations is available online at <http://dx.doi.org/10.5281/zenodo.1179068>.

References

- S. Arganda, A. Pérez-Escudero, and G. G. de Polavieja. A common rule for decision making in animal collectives across species. *Proceedings of the National Academy of Sciences*, 109(50):20508–20513, 2012.
- S. Atkinson, C-Y Chang, R. E. Sockett, M. Cámara, and P. Williams. Quorum sensing in yersinia enterocolitica controls swimming and swarming motility. *Journal of Bacteriology*, 188(4):1451–1461, 2006.
- P. R. Y. Backwell, J. H. Christy, S. R. Telford, M. D. Jennions, and N. I. Passmore. Dishonest signalling in a fiddler crab. *Proceedings: Biological Sciences*, 267(1444):719–724, 2000.

- Carl T. Bergstrom and Michael Lachmann. Signalling among relatives. i. is costly signalling too costly? *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 352(1353):609–617, 1997. ISSN 0962-8436.
- S. P. Brown and R. A. Johnstone. Cooperation in the dark: Signalling and collective action in quorum-sensing bacteria. *Proceedings of the Royal Society B: Biological Sciences*, 268(1470):961–965, 2001.
- A. C. Burgos and D. Polani. Cooperation and antagonism in information exchange in a growth scenario with two species. *Journal of Theoretical Biology*, 399:117 – 133, 2016. ISSN 0022-5193.
- I. D. Couzin. Collective cognition in animal groups. *Trends in Cognitive Sciences*, 13(1):36 – 43, 2009. ISSN 1364-6613.
- T. R. de Kievit and B. H. Iglewski. Bacterial quorum sensing in pathogenic relationships. *Infection and Immunity*, 68(9):4839–4849, 2000.
- M. C. Donaldson-Matasci, M. Lachmann, and C. T. Bergstrom. Phenotypic diversity as an adaptation to environmental uncertainty. *Evolutionary Ecology Research*, 10:493–515, 2008.
- M. C. Donaldson-Matasci, C. T. Bergstrom, and M. Lachmann. The fitness value of information. *Oikos*, 119(2):219–230, 2010.
- M. C. Donaldson-Matasci, C. T. Bergstrom, and M. Lachmann. When unreliable cues are good enough. *The American Naturalist*, 182(3):313–327, Sept 2013.
- N. R. Franks, S. C. Pratt, E. B. Mallon, N. F. Britton, and D. J. T. Sumpter. Information flow, opinion polling and collective intelligence in house-hunting social insects. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 357(1427):1567–1583, 2002. ISSN 0962-8436. doi: 10.1098/rstb.2002.1066.
- J. M. Henke and B. L. Bassler. Bacterial social engagements. *Trends in Cell Biology*, 14(11):648–656, 2004.
- S. Huttegger, B. Skyrms, P. Tarrés, and E. Wagner. Some dynamics of signaling games. *Proceedings of the National Academy of Sciences*, 111(Supplement 3):10873–10880, 2014. doi: 10.1073/pnas.1400838111.
- A. J. King and G. Cowlshaw. When to use social information: the advantage of large group size in individual decision making. *Biology Letters*, 3(2):137–139, 2007. ISSN 1744-9561. doi: 10.1098/rsbl.2007.0017.
- E. Kussell and S. Leibler. Phenotypic diversity, population growth, and information in fluctuating environments. *Science*, 309(5743):2075–2078, 2005. ISSN 0036-8075.
- M. Lachmann, G. Sella, and E. Jablonka. On the advantages of information sharing. *Proceedings of the Royal Society of London B: Biological Sciences*, 267(1450):1287–1293, 2000. ISSN 0962-8452.
- D. Lopez, H. Vlamakis, and R. Kolter. Generation of multiple cell types in bacillus subtilis. *FEMS Microbiology Reviews*, 33(1):152–163, 2009. ISSN 1574-6976.
- F. Meacham, A. Perlmutter, and C. T. Bergstrom. Honest signalling with costly gambles. *Journal of The Royal Society Interface*, 10(87), 2013.
- M. B. Miller, , and B. L. Bassler. Quorum sensing in bacteria. *Annual Review of Microbiology*, 55(1):165–199, 2001. PMID: 11544353.

- N. Miller, S. Garnier, A. T. Hartnett, and I. D. Couzin. Both information and social cohesion determine collective decisions in animal groups. *Proceedings of the National Academy of Sciences*, 110(13):5263–5268, 2013.
- D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124 – 143, 1996.
- C. D Nadell, J. B Xavier, S. A Levin, and K. R Foster. The evolution of quorum sensing in bacterial biofilms. *PLOS Biology*, 6(1):1–9, 01 2008.
- J. M. Pacheco, V. V. Vasconcelos, F. C. Santos, and B. Skyrms. Co-evolutionary dynamics of collective action with signaling for a quorum. *PLOS Computational Biology*, 11(2):1–12, 02 2015.
- A. Pérez-Escudero and G. G. de Polavieja. Collective animal behavior from bayesian estimation and probability matching. *PLOS Computational Biology*, 7(11):1–14, 11 2011.
- T. J Perkins and P. S Swain. Strategies for cellular decision-making. *Molecular Systems Biology*, 5(1), 2009. ISSN 1744-4292.
- R. Popat, D. M. Cornforth, L. McNally, and S. P. Brown. Collective sensing and collective responses in quorum-sensing bacteria. *Journal of The Royal Society Interface*, 12(103), 2015. ISSN 1742-5689.
- S. C. Pratt, E. B. Mallon, D. J. Sumpter, and N. R. Franks. Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant *leptothorax albipennis*. *Behavioral Ecology and Sociobiology*, 52(2):117–127, Jul 2002. ISSN 1432-0762. doi: 10.1007/s00265-002-0487-x.
- O. Rivoire and S. Leibler. The value of information for populations in varying environments. *Journal of Statistical Physics*, 142(6):1124–1166, 2011. ISSN 1572-9613.
- S. T. Rutherford and B. L. Bassler. Bacterial quorum sensing: Its role in virulence and possibilities for its control. *Cold Spring Harbor Perspectives in Medicine*, 2(11):a012427, 11 2012.
- W. H. Sandholm. Potential games with continuous player sets. *Journal of Economic Theory*, 97(1): 81 – 108, 2001. ISSN 0022-0531.
- W. H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, 2010. ISBN 978-0262195874.
- T. Sasaki and S. C. Pratt. The psychology of superorganisms: Collective decision making by insect societies. *Annual Review of Entomology*, 63(1):259–275, 2018. doi: 10.1146/annurev-ento-020117-043249.
- B. Skyrms. *Signals: Evolution, Learning, and Information*. Oxford University Press, 2010.
- Y. van Bergen, I. Coolen, and K. N. Laland. Nine-spined sticklebacks exploit the most reliable source when public and private information conflict. *Proceedings of the Royal Society B: Biological Sciences*, 271(1542):957–962, 2004.
- F. Vega-Redondo. *Economics and the Theory of Games*. Cambridge University Press, 2003.
- S. A. West, S. P. Diggle, A. Buckling, A. Gardner, and A. S. Griffin. The social lives of microbes. *Annual Review of Ecology, Evolution, and Systematics*, 38(1):53–77, 2007.
- M. Whiteley, S. P. Diggle, and E. P. Greenberg. Progress in and promise of bacterial quorum sensing research. *Nature*, 551:313–320, 2017.