A theory of collaborative interaction: From interactive to coordinated decision-making in conditional games

Luca Tummolini

Institute of Cognitive Sciences and Technologies
Italian National Research Council

Wynn Stirling

Department of Electrical and Computer Engineering
Brigham Young University

Abstract

From prehistoric collaborative foraging to contemporary international research networks, a collaborative group is one in which individuals are prepared to make a concerted effort to achieve an outcome. How should we understand such willingness to engage in coordinated action? We argue that collaborative interaction relies on agents who are responsive to motivation of others and who also aim at coordinated group behavior when deciding how to act. While game-theoretic models of reciprocity focus on the former kind ("reciprocal" players), models of team-reasoning aim at those who are motivated by group-level interests ("team" players). Neither approach by itself, however, provides an adequate theory of "collaborative" players. Using the framework of conditional game theory, we present a dynamic model of joint decision-making between collaborative players that integrates essential features of reciprocity and of team reasoning approaches. The theory is used to analyze the interaction between collaborative players in the Snowdrift game, a social dilemma situation in which each party directly benefits from her own efforts but there is also a temptation to free-ride on those of others. By comparing the classical analysis to that afforded by conditional game theory, we identify the conditions that support the joint decision to act in coordination with others, i.e. when collaboration is successful.

1 Introduction

It is common experience that before undertaking a collaborative project, all the involved parties are able to identify the mutual benefits that might ensue from it and can be sincerely motivated to opt in.

As a familiar example, consider research collaborations. It has been shown that research that emerges from collaboration networks attracts more attention and is cited more frequently than that undertaken by solo authors (Wuchty et al., 2007). Since international research collaborations are thriving (Adams, 2013; Grayson and Pincock, 2015), when scientists decide to collaborate they expect to achieve more than they can do by solitary inquiry. Still, even if everyone can see the benefits, collaboration often involves additional costs in terms of money, time and additional effort to integrate each contribution in a coherent whole. In fact, the costs of collaborating with others might be such that a team might settle for less than what was originally envisioned or, most commonly, that one or just a few members of the team end up doing all the work while the rest enjoys the free ride. Other times, however, the team members engage in their joint project in such a way that they are able to overcome its costs and successfully complete it. What is the difference between these two possible trajectories?

2 The Snowdrift game and the challenge of collaboration

To provide a more precise picture of the challenge of collaboration, let $\{X_1, \ldots, X_n\}$ denote a collective such that each X_i possesses its own finite action set $\mathcal{A}_i$ from which it is empowered to choose one element for implementation. The structure of interaction is such that the benefit to each individual is a function of the simultaneous actions of all participants. The collaboration problem is then for the individuals to come to an agreement regarding which set of simultaneous actions should be implemented in order to achieve the desired task.

When framed in this way, the collaboration problem fits the general structure of a strategic game, which is perhaps the most well known mathematical structure within which to couch such collective decision problems. The fundamental structure of the game theoretic model is that each participant comes to the social engagement with ex ante preferences expressed as payoffs that accrue to each participant as a result of their interactive decisions. These payoffs are then juxtaposed into a payoff array and each participant chooses an action according to whatever criterion it deems to be important. The most well-known such concept is that each individual chooses an action that maximizes its individual benefit under the constraint that all others are behaving similarly—a Nash equilibrium.

Consider, in particular, the following scenario. Two individuals with different backgrounds and expertise desire to accomplish a task that requires both skill sets. This task involves a division of labor where each individual performs some portion of a task. The success of the collaboration, however, depends on the effectiveness of the integration of the different skill sets. If each works independently and concentrates only on its portion of the task, the integration may be poor and the result may be of marginal value. But integration requires that the participants invest in each other's area of expertise, which involves additional cost. If the cost is too high, the collaboration may fail, even though both participants are willing to collaborate. Thus, there is both a cooperative utility (the benefit of achieving the task) and a conflictive disutility (the cost of integration).

Each participant thus has two possible actions: cooperate (C) or defect (D). One cooperates by expending the additional effort required to integrate, and one defects by concentrating only on one's own portion of the task. The individual preference ordering over the four possible outcomes are as follows.

- The best individual payoff obtains when the other collaborator does all of the integration work.
- The next-best individual payoff obtains when both participate in the integration.
- The next-worst individual payoff obtains when one does all of the integration alone.
- The worst individual payoff obtains when neither integrates.

From the perspective of classical game theory, this collaboration maybe viewed as a Snowdrift game (Sugden, 1986), with ordinal payoff matrix as given in Table 1.

Table 1: Payoff matrix for the Snowdrift game. X_r and X_c are the row a column players, respectively

	X_c	
X_r	C	D
C	(3, 3)	(2, 4)
D	(4, 2)	(1, 1)

Ordinal ranking: 4 = best; 3 = next-best; 2 = next-worst; 1 = worst

The storyline for this game involves two neighbors whose road between them is blocked by snow. Each player may either take the initiative to shovel the snow (cooperate) or wait for the other to shovel (defect). The Snowdrift has the same underlying structure of a Chicken game and it captures a situation with an opportunity for mutual benefit in which the agents are in dispute over something that they both want to avoid, e.g. the additional effort to integrate.

In its one-shot version, Snowdrift has two pure strategy Nash equilibria in which one agent does all the effort while the other free rides. Alternatively, if the game is re-cast as a mixed strategy scenario, the participants choose their respective strategies according to their chosen probability distributions. The resulting game has a unique mixed-strategy Nash equilibrium, where the frequency of mutual cooperation is given by the joint probability of their independent decisions to choose cooperation. Since their independent choices do not aim at coordination over the cooperative outcome, such coordination, if it happens at all, is just accidental. Collaboration in this perspective is doomed from the start.

In contrast with this traditional picture, here we suggest that a collaborative interaction is possible for individuals who are prepared to make a concerted effort to achieve an outcome, who, in other words, engage in coordinated action purposefully. How should such willingness be understood?

2.1 Related Work

Game theory provides a powerful, yet simple, framework within which to analyze social behavior. Thus, much effort has gone into attempts to adapt the theory to more socially sophisticated contexts. Many attempts have been made to modify the content of the payoffs to account for social interests. *Behavioral game theory* (Rabin, 1993; Fehr and Schmidt, 1999; Camerer, 2003; Camerer et al., 2004b,a; Henrich et al., 2005; Bolton and Ockenfels, 2005) has developed in response to the desire to introduce psychological realism into game theory by incorporating social concepts such as fairness and reciprocity into the utilities in addition to considerations of material benefit. The closely related field of *psychological game theory* (Geanakoplos et al., 1989; Dufwenberg and Kirchsteiger, 2004; Battigalli and Dufwenberg, 2009) also employs utilities that account for beliefs as well as actions and takes into consideration belief-dependent motivations such as guilt aversion, reciprocity, regret, and shame. Others have attempted to account for social context by modifying the solution concept with the introduction of models drawn from biological and social evolutionary processes (Axelrod, 1984; Bicchieri, 1993; Sartorius, 2003; Fefferman and Ng, 2007; Bossert et al., 2012). Such approaches provide important models of the emergence of social relationships in repeated-play environments where individuals' fitness for long-term survival is taken into consideration in addition to their short-term material payoffs, and the propensity to behave in ways that extend beyond narrow self-interest is viewed as the end product of natural or social evolution.

Another approach to the collaboration problem is offered by Bacharach (2006) and Sugden (2000, 2003), who develop a theory of *team reasoning*. The central thesis of this approach is that individuals may frame the decision scenario in multiple ways—one in terms of their individual payoffs, and one in terms of the payoffs for the team with which they associate. The result is for the grop to cast the issue as "What shall we do?" rather than "What shall I do?" Bacharach argues that the solution concept of team reasoning will emerge endogenously when circumstances exist that tend to make people identify primarily as a group rather than as individuals.

2.2 Our Contribution

The common element of the game-theoretic models mentioned above is that they all rely, at the end of the day, on the assumption that each individual comes to the social engagement with a fully defined preference ordering over all profiles of action. Although much effort and creativity has gone into the *content* of the payoffs and on the solution concepts, the *structure* of the payoffs has remained unchanged. Once defined, the payoffs, or utilities, are fixed and immutable—they are categorical. Furthermore, they are independent of context. An individual endowed with categorical preferences can only react to a pre-assumed social situation; it cannot respond to any particular social situation. This assumption greatly limits the ability of individuals to respond to the interests of others in an attempt to reach a collaborative solution. In addition, whether the agents will in fact coordinate their choices or not entirely depends on the pre-existing alignment between their categorical preferences. In the Snowdrift game, for instance, such alignment is only partial and, as a consequence, coordination on the cooperative outcome happens only in an equilibrium in mixed-strategies that is, by definition, accidental.

The goal of this paper is to present a model of collaborative interaction, couched in the game-theoretic context, that overcomes the limitation imposed by relying on categorical utilities. Our approach is to endow individuals with *conditional utilities* that enable them to modulate their preferences as a function of the social influence that others exert on them. We follow Stirling (2012), who developed a theory of conditional games that provides the mathematical structure and theoretical development to enable individuals to respond to social influence. That development, however, is restricted to acyclic influence, that is, two-way influence is not permitted. In this paper, however, we extend conditional game theory to account for influence cycles, which is here considered an essential attribute of a collaborative interaction.

2.3 The Catch-and-Toss model of Collaboration

In contrast with traditional assumptions, collaboration is not really an activity where participants come to the engagement with fixed utilities, juxtapose them in a payoff matrix, and render a final decision that governs the collaboration. Collaboration is more realistically modeled as a process involving repeated interchanges between the participants in a sort of "catch and toss" dynamics. Think, for instance, as the process of drafting a paper in which an author first "tosses" an idea to another author who "catches" it, revises or extends it, and then "tosses" it back (Huebner et al., forthcoming). Beside the material progress, such catch and toss dynamics is also a way in which information about one another's preferences circulates in the group and information about others' preferences can influence one's own preferences as the literature on peer effects have largely documented (Thoni and Gachter, 2015) .

Thus, a more natural way to model this collaboration is to view it as a time sequence of back-and-forth influence relationships—a cyclic game of the form

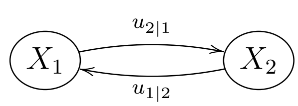

(1)

where $u_{i|j}$represent the social influence that X_i exerts on X_j. The existence of the cycle, however, introduces a dynamic component into the decision process, where X_i influences X_j, who in turn influences X_i, and the cycle is repeated infinitely many times. Thus, the preferences of each member of the cycle will be subject to continual change. At first glance, this situation may appear to be a manifestation of the concern expressed by Bacharach (2006) that attempting to derive joint intentions from individual intentions can lead to an infinite regress with no resolution. This is a legitimate concern that must be addressed. The issue devolves around the question of convergence: Does the cycle of influence propagation oscillate endlessly with each individual repeatedly changing its mind, or does it converge in the sense that each individual ultimately possesses an individual *steady-state* utility? Answering this question requires the development of additional theory regarding conditional games.

3 Summary of Conditional Game Theory

A classical *normal form game* comprises set of agents (players) $\{X_1, \ldots, X_n\}$, each of whom possesses a finite action set $\mathcal{A}_i = \{z_{i1} \ldots z_{iN_i}\}$ and a categorical utility u_i defined over the product set $\boldsymbol{\mathcal{A}} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$, the set of *action profiles.* A *payoff array* is an $N_1 \times \cdots \times N_n$ dimensional structure such that the $(k_1, \cdots, k_n)$th entry is the sub-array $(u_1(z_{1k_1}), \ldots, u_n(z_{nk_n}))$, where $u_i(z_{ik_i})$ is the payoff to X_i receives if the profile $(z_{ik_1}, \ldots, z_{nk_n})$ is actualized. When $n = 2$, the payoff array is termed a *payoff matrix.* Game theory is a powerful prescriptive model for multi-agent decision making. One of the reasons for its success is the virtually perfect match between the mathematical structure of the game and the notion of individual rationality used to render a decision. If one is motivated by narrow self-interest, the natural mechanism by which to express those interests is with a linear ordering over the set of all outcomes. Conversely, if one possesses a linear ordering over all outcomes, the natural approach is to choose the outcome that maximizes individual benefit. Thus, the Nash equilibrium solution provides an optimal solution, subject to the mutual constraints imposed by all players.

3.1 Conditionalization

As a consequence of the social influence exerted by others, a player must take that influence into consideration when forming its preferences and, therefore, must extend its concept of individual interest to account for social, as well as individual interests. A natural way to incorporate social considerations into a collective is to view it as a network whose members are linked together by some means of communication or control that enables them to exert social influence on each other.

Conditional game theory, as introduced by Stirling (2012), is an extension to classical noncooperative game theory, and serves as a model for behavior for the behavior of a collective of individuals who exert social influence on each other. To introduce this concept, consider the simple network scenario involving two agents, X_1 and X_2, as illustrated by the acyclic graph

$$X_1 \longrightarrow X_2, \qquad (2)$$

where the direction of the arrow indicates that X_1 directly influences X_2 but X_2 does not directly influence X_1. We assume the following conditions.

Directionality: Although social influence is directed from the one who influences to the one who is influenced, it is the receiver of the influence who activates the relationship. Thus, the influencee can modulate its preferences in response according to its own volition. The influencer does not control, dictate, or otherwise force preferences or behavior on the influencee.

Conditionality: The influencee does not require knowledge of the preferences of the influencer in order to establish the influence linkage.

These assumptions may appear to be problematic. How can the influencee respond without knowledge of the influencer's preferences? The answer to this question lies in the logical structure of the influence mechanism, namely, the logic of *conditionalization*—a key concept of Bayesian epistemology. Conditionalization takes the form of a hypothetical proposition "If ... then ____," where ... is the *antecedent* and ____ is the *consequent*. In an epistemological context, one incorporates evidence into an assessment of the probability that an event B is realized, conditioned on knowledge that event A is realized, by the conditional probability $P(B|A)$. The antecedent is the hypothesis that A is realized, and the consequent is the resulting probability that B is realized. The great strength and wide applicability of probability theory is the facility to deal with hypothetical propositions.

We employ this same conditionalization logic to model preferences. In this case, the antecedent is a hypothesis that a particular profile is under consideration by or in behalf of the influencer as the one that is intended to be actualized. Conditioned on that hypothesis, the consequent is the influenee's ordering over the set of profiles. It is important to emphasize that the influencee does not need to know if or why the influencer might have intentions regarding that profile, since it is merely a hypothesis. All the influencee needs to consider is what its response should be given the antecedent.

3.2 Graphical Network Models

Modeling the influence mechanism according to the logical structure of conditionalization enables agents to anticipate all possible scenarios regarding the influencer's preferences, thereby providing the influencee maximum flexibility in responding, especially if the influencer is itself an influencee of a third agent. In such cases, both influencees will be uncertain regarding their response until the preferences of the influencers are determined. Uncertainty, in this context, is *not* regarding beliefs; rather it is uncertainty regarding preferences. The conventional application of the probability syntax as means of expressing epistemological uncertainty, but this same logical structure may be used to expressing *behavioral* uncertainty. In other words, one is epistemologically uncertain if one does not have complete knowledge that a proposition is realized, and one is behaviorally uncertain if one is not completely decisive that an alternative should be actualized. To comply with this logical structure, we focus on preference structures that admit conditionalization. To proceed, it is convenient to employ the structure and syntax of graph theory.

Definition 1. *The* graph *of a network consists of a set of* vertices *comprising the individuals X_i, $i = 1, \ldots, n$, and a set of* edges*, also termed* linkages*, that serve as the medium by which influence is propagated between individuals. An edge is* directed *(denoted with the arrow symbol "$\rightarrow$") if the propagation is unidirectional: $X_j \rightarrow X_i$ means that X_j directly influences X_i. A* path *from X_j to X_i is a sequence of directed edges from X_j to X_i, denoted $X_j \mapsto X_i$. A path is a* cycle*, or* closed path*, if $X_j \mapsto X_j$. A graph is said to be a* directed acyclic graph *if all edges are directed and there are no cycles.*

The directionality structure of a directed graph is such that the child (the influencee) responds to the parent (the influencer) by forming a hypothesis for each of the possible action profiles that may be viewed by the parent as he one that should or will be implemented. Such a hypothesis is termed a conjecture. We set the following notation.

Definition 2. *A* conjecture *for X_i is an action profile, denoted $X_i \models \mathbf{a}_i = (a_{i1}, \ldots, a_{in})$, where the jth entry of $\mathbf{a}_i$ is the element $z_{jk} \in \mathcal{A}_j$ such that $a_{ij} = z_{jk}$, $j = 1, \ldots, n$, that* *is hypothetically asserted by or in behalf of* X_i *as the one that will or should be actualized. The* n^2 *dimensional array* $(\mathbf{a}_1, \dots, \mathbf{a}_n) \in \boldsymbol{\mathcal{A}} \times \cdots \times \boldsymbol{\mathcal{A}}$ *is a* joint conjecture*, where* $X_i \models \mathbf{a}_i$*,* $i = 1, \dots, n$*.*

Since a conjecture is a hypothetical assertion, why or by whom it is asserted is irrelevant; thus, we simply say that $\mathbf{a}_i$ is a conjecture *for* X_i unless the context clearly identifies the source of the assertion. There is no need to assume that a conjectured profile is preferred by the influencer (although that is one possible motivation).

Definition 3. *The* parent *set for* X_i*, denoted* $\mathrm{pa}(X_i) = \{X_j \colon X_j \to X_i\}$*, is the subset of individuals that directly influence* X_i*. If* X_i *has* $p_i > 0$ *parents, then* $\mathrm{pa}(X_i) = \{X_{i_1}, \dots, X_{i_{p_i}}\}$*, where* $X_{i_k} \to X_i$*,* $k = 1, \dots, p_i$*. For notational convenience, let* $\mathrm{pa}(i) = \{i_1, \dots, i_{p_i}\}$ *denote the indices corresponding to the elements of* $\mathrm{pa}(X_i)$*. Also, let* $\boldsymbol{\alpha}_i = (\mathbf{a}_{i_1}, \dots, \mathbf{a}_{i_{p_i}})$*, where* $X_{i_k} \models \mathbf{a}_{i_k}, k = 1, \dots, p_i$*, denote a* conditioning conjecture *for the parents of* X_i*.*

Definition 4. *Given a parent set* $\mathrm{pa}(X_i)$ *and a conditioning conjecture* $\boldsymbol{\alpha}_i$*, A* conditional utility *is a function* $u_{i|\mathrm{pa}(i)}(\cdot|\boldsymbol{\alpha}_i) \colon \boldsymbol{\mathcal{A}} \to \mathbb{R}$ *such that, for each conditioning conjecture* $\boldsymbol{\alpha}_i \in \boldsymbol{\mathcal{A}}^{p_i}$ *for* $\mathrm{pa}(X_i)$*,* $u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_i)$ *defines the utility of* $X_i \models \mathbf{a}_i$*, given that* $\mathrm{pa}(X_i) \models \boldsymbol{\alpha}_i = (\mathbf{a}_{i_1}, \dots, \mathbf{a}_{i_{p_i}})$*. The conditioning symbol "|" separates the conditioned entity on the left from the conditioning entity on the right. If* $p_i = 0$*, then* $u_{i|\mathrm{pa}(i)} = u_i$*, a categorical utility for* X_i*.*

Definition 5. *A* conditional network game *comprises a set of individuals* $\{X_1, \dots, X_n\}$*, a set of conjecture profiles* $\boldsymbol{\mathcal{A}} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ *and a family of sets of conditional utilities of the form*

$$\{u_{i|\mathrm{pa}(i)}(\cdot|\boldsymbol{\alpha}_i)\, \forall \boldsymbol{\alpha}_i \in \boldsymbol{\mathcal{A}}^{p_i}, 1, \dots, n\} . \tag{3}$$

Conditional utilities provide the mechanism by which individuals may incorporate the interests of others into their own interests. In this way, any influence that $\mathrm{pa}(X_i)$ exerts on X_i can be encoded into X_i's family of conditional utilities. As a simple example of a network game, consider the system illustrated by (4), where X_1 possesses a categorial utility u_1 and X_2 possesses a conditional utility $u_{2|1}$, which we indicate by the following expression:

$$\textcircled{X_1} \xrightarrow[u_{2|1}]{} \textcircled{X_2} . \tag{4}$$

Migrating the notion of conditionalization from the epistemological domain to the behavioral domain also invites the application of the probability syntax to utilities. So doing, however, imposes a constraint on the structure of the utilities: they must be *normalized*, meaning that they must be nonnegative and must sum to unity; that is,

$$\begin{aligned} u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_i) &\geq 0 \ \forall \boldsymbol{\alpha}_i \in \boldsymbol{\mathcal{A}}^{p_i} \\ \sum_{\mathbf{a}_i} u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_i) &= 1 \ \forall \boldsymbol{\alpha}_i \in \boldsymbol{\mathcal{A}}^{p_i} . \end{aligned} \tag{5}$$

Since positive affine transformations of utility preserve preference orderings uniquely, we may assume normalization without loss of generality. We shall refer to normalized utilities as *utility mass functions*.

The application of the probability syntax to the behavioral domain is a departure from the classical applications of probability theory which predominantly fall into the epistemological domain. As the following lemma establishes, however, the two domains are connected by an order isomorphism.

Lemma 1. *An order isomorphism exists between ordering the strength of belief regarding propositions and ordering the strength of preference regarding alternatives.*[1] *This isomorphism applies to both categorical and conditional orderings.*

This lemma, a proof of which can be found in the Appendix, establishes an isomorphic relationship between probability mass functions and utility mass functions. Accordingly, all of the operations associated with probability, such as marginalization, independence, Bayes rule, and the chain rule are well defined mathematical procedures. However, to be meaningful, we must supply behavioral interpretations.

3.3 Coordination

When dealing with multiagent decision making in contexts where some notion of systematic group-level behavior is a factor, the notion of coordination becomes relevant. Coordination comes from the Latin: *co* (together) + *ordinare* (to regulate). The Oxford English Dictionary defines *coordinate* as "to place or arrange (things) in proper position relative to each other and to the system of which they form parts; to bring into proper combined order as parts of a whole" (Murray et al., 1991). Applying this definition to a collective of autonomous decision-making individuals, coordination requires that they (the parts) must combine their behavior in a systematic way to form a properly constructed society (the whole). For a collective to coordinate, the individuals must fit into some nontrivial social structure, and coordination occurs if the group possesses some notion of systematic behavior.

To motivate our concept of coordination, let us first consider an analogous probabilistic situation. Let Y_1 and Y_2 be discrete random variables defined over a finite set Ω Let $p_1\colon \Omega \to [0,1[$ denote a marginal probability mass function for Y_1, and let $p_{2|1}\colon \Omega \to [0,1]$ for each $\omega_1 \in \Omega$ denote a conditional probability mass function for Y_2. We may then synthesize a joint probability mass function $p_{12}\colon \Omega \times \Omega \to [0,1]$ as

$$p_{12}(\omega_1, \omega_2) = p_1(\omega_1) p_{2|1}(\omega_2|\omega_1)\,. \tag{6}$$

Now applying the isomorphism between probability theory and utility theory, let X_1 possess a categorical utility mass function $u_1\colon \boldsymbol{\mathcal{A}} \to [0,1]$ and let X_2 possess a conditional utility mass function $u_{2|1}\colon \boldsymbol{\mathcal{A}} \to [0,1]$ for each $\mathbf{a} \in \boldsymbol{\mathcal{A}}$. Then we may synthesize a joint utility mass function $u_{12}\colon \boldsymbol{\mathcal{A}} \times \boldsymbol{\mathcal{A}} \to [0,1]$ as

$$u_{12}(\mathbf{a}_1, \mathbf{a}_2) = u_1(\mathbf{a}_1) u_{2|1}(\mathbf{a}_2|\mathbf{a}_1)\,, \tag{7}$$

The synthesis of a joint utility via (7) provides a mechanism by which a notion of strict individual rationality, as expressed by the categorical utility u_1, and an expanded notion of

[1] Two sets are *order isomorphic* if one of the orderings can be obtained from the other by renaming the members of the set (Itô, 1987).

individual rationality whereby one takes into account the influence of others, as expressed by the conditional utility $u_{2|1}$, are combined to form a notion of social rationality, as expressed by u_{12}. We note, however, that joint utility is not an ordering over $\boldsymbol{\mathcal{A}}$, the set of conjeture profiles. Rather, it is an ordering over $\boldsymbol{\mathcal{A}} \times \boldsymbol{\mathcal{A}}$, the set of joint conjectures. This is an important difference from the classical notion of a social ordering as an ordering over $\boldsymbol{\mathcal{A}}$. Since u_{12} is a function of the conjectures of both agents, it enables us to evaluate the connection between the two agents as a function of their individual conjectures. In particular, it provides a vehicle with which to investigate coordination. To develop this concept, it is convenient to return to the isomorphism.

Definition 6. *A* Bayesian network *is a directed acyclic graph that satisfies the following conditions.*

- *The ith vertex corresponds to a discrete random variable Y_i taking values in a finite set Ω.*
- *The conditional probability that $Y_i = \omega_i \in \Omega$, given that its parents* $\mathrm{pa}\,(Y_i) = \{Y_{i_1}, \ldots, Y_{i_{q_i}}\}$ *assume the values $\boldsymbol{\omega}_i = (\omega_{i_1}, \ldots, \omega_{i_{q_i}}) \in \Omega^{p_i}$, is specified as $p_{i|\mathrm{pa}(i)}(y_i|\boldsymbol{\omega}_i)$. If* $\mathrm{pa}\,(Y_i) = \emptyset$, *then Y_i is a root vertex and $p_{i|\mathrm{pa}(i)} = p_i$, the unconditional marginal distribution of Y_i.*

A Bayesian network provides a powerful framework within which to analyze the behavior of either artificial or naturally existing collectives of interacting elements. It is natural to start by considering the way the local behavior of specific elements is influenced by other elements that are in close proximity either spatially, temporally, or functionally. From such local models of behavior one can build a global model by piecing together the local components in appropriate ways. This approach, pioneered by Pearl (1988), provides a powerful synthesis tool for the analysis of human networks and for the design and synthesis of artificial networks. For additional discussions of Bayesian networks, see Cowell et al. (1999), Lauritzen (1996), and Jensen (2001).

An acyclic conditional network game as expressed by Definition 5 is isomorphic to a Bayesian network. A fundamental property of Bayesian networks is that the joint probability mass function of the vertices, denoted $p_{1:n}$, is the product of the conditional mass functions of all non-root vertices and the marginal mass functions of the root vertices, or

$$p_{1:n}(\omega_1, \ldots, \omega_n) = \prod_{i=1}^{n} p_{i|\mathrm{pa}(i)}(\omega_i|\boldsymbol{\omega}_i)\,. \tag{8}$$

Thus, by the isomorphism, the *joint utility* of a social network is

$$u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_n) = \prod_{i=1}^{n} u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_i)\,, \tag{9}$$

where, if $\mathrm{pa}\,(X_i) = \varnothing$, then $u_{i|\mathrm{pa}(i)} = u_i$, a categorical utility. The joint utility of a social network is isomorphic to the joint distribution of a Bayesian network whose vertices are random vectors and whose edges are conditional utilities with respect to the random vectors.

In the epistemological context, the operational concept associated with the joint probability $p_{1:n}$ is that it orders the set of joint events $(\omega_1, \ldots, \omega_n) \in \Omega^n$ in terms of their probability of being jointly realized. In the behavioral context, the joint utility $u_{1:n}$ orders the set of joint conjectures $(\mathbf{a}_1, \ldots, \mathbf{a}_n) \in \boldsymbol{\mathcal{A}}^n$, but we have yet to establish an operational definition of what such an ordering means. It is intuitively appealing to say that, if the joint utility of $(\mathbf{a}_1, \ldots, \mathbf{a}_n)$ is greater than the joint utility of $(\mathbf{a}_1', \ldots, \mathbf{a}_n')$, then the former joint conjecture is more coordinated than the latter. To illustrate, suppose two individuals are attempting to pass through a narrow doorway that just wide enough for the two to squeeze past each other. The action set for each is $\mathcal{A}_1 = \mathcal{A}_2 = \{r, \ell\}$; that is, to move right (r) or left (ℓ), where the sense of direction is from the perspective of the individual. There are sixteen possible joint conjectures of the form

$$[(a_{11}, a_{12}), (a_{21}, a_{22})] \in (\mathcal{A}_1 \times \mathcal{A}_2) \times (\mathcal{A}_1 \times \mathcal{A}_2) \,. \tag{10}$$

If we assume that the emergent systematic behavior of the group is cooperative (which would be an emergent result of their interrelationships), then a natural preference ordering of the joint conjectures would be

$$\begin{aligned} u_{12}[(r,r),(r,r)] = u_{12}[(\ell,\ell),(\ell,\ell)] > u_{12}[(r,r),(r,\ell)] = u_{12}[(\ell,\ell),(\ell,r)] \\ > \cdots > u_{12}[(r,\ell),(\ell,r)] = u_{12}[(\ell,r),(r,\ell)] \,, \end{aligned} \tag{11}$$

On the other hand, if the group preference is conflictive, then the natural ordering would be reversed. In either case, however, such expressions admit an interpretation as an ordering of the coordination characteristics of the set of joint conjectures. This example illustrates that coordination is a neutral concept. The first scenario illustrates cooperative coordination, with latter illustrating conflictive coordination.

Definition 7. *Given a conditional network acyclic game, the joint utility as defined by (9), is termed the* coordination function *for the network.*

The coordination function captures all of the social relationships that emerge as the conditional preferences propagate through the network. Since it is a function of n^2 independent variables, the coordination function does not lead directly to a group-level utility. We may, however, construct a group-level utility as the marginal utility by considering the utility that obtains when each agent conjectures only its component of a profile.

Definition 8. *Given the coordination function $u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_n)$, let a_{ij} denote the jth element of $\mathbf{a}_i$, that is, $\mathbf{a}_i = (a_{i1}, \ldots, a_{in})$ is a conjecture for X_i, and form the profile $(a_{11}, \ldots, a_{nn})$ by taking the ith element of each X_i's conjecture profile, $i = 1, \ldots, n$. Now sum the coordination function over all elements of each $\mathbf{a}_i$ except the iith elements to form the* group utility $w_{1:n}$ *for $\{X_1, \ldots, X_n\}$, yielding*

$$w_{1:n}(a_{11}, \ldots, a_{nn}) = \sum_{\sim a_{11}} \cdots \sum_{\sim a_{nn}} u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_n) \,, \tag{12}$$

where the notation $\sum_{\sim a_{ii}}$ means the sum is taken over all arguments of $u_{1:n}$ except a_{ii}.

The individual decision function *w_i of X_i is the ith marginal of $w_{1:n}$, that is,*

$$w_i(a_{ii}) = \sum_{\sim a_{ii}} w_{1:n}(a_{11}, \ldots, a_{1n}) = \sum_{\sim a_{ii}} u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_n) \,. \tag{13}$$

The group utility takes into account all of the social relationships and provides an ordering over $\boldsymbol{\mathcal{A}}$ that defines the emergent preference ordering for the group. The individual decision functions provide each member of the collective a decision rule that also takes into account all of the social influence that is exerted on it. In terms of coordination, the individual decision functions constitute the parts, and the group utility constitutes the whole. Together, they establish the coordination characteristics of the collective. To the extent that the group utility and the individual decision functions are compatible, the collective engages in truly coordinated behavior.

4 Extension to Cyclic Networks

The results presented in Section 3 apply to acyclic networks, and therefore cannot be directly applied to a the cyclic network given by (1). However, we may view the cyclic structure as a sequence of acyclic networks of the form $X_i \to X_j \to X_i \to \cdots$, where a time interval lapses between each network. In other words, the cyclic network can be viewed as a dynamic network comprising the concatenation of an infinite sequence of acyclic networks. To avoid an infinite regress, we must establish necessary and sufficient conditions for this dynamic network to converge in the sense that, in the limit, each individual possesses a constant utility that can be used to render a joint decision. Essentially, this dynamic process is a mathematical representation of the iterative process involved in a collaboration.

4.1 Matrix Form Dynamics Model

In this section we develop the theory for the convergence of a cyclic network.

Definition 9. *A* path *from X_j to X_i is a sequence of directed edges from X_j to X_i, denoted $X_j \mapsto X_i$. A path is* closed *if $X_j \mapsto X_j$. A k-member closed path is called a* simple k-cycle *if every vertex has exactly one incoming edge and one outgoing edge.*

Consider the k-cycle illustrated in Figure 1. We may view this cyclic network as a sequence of influence operations that evolve in time, Let $t = 0, \delta, 2\delta, \cdots$, and consider the equivalent dynamic network displayed in Figure 2.

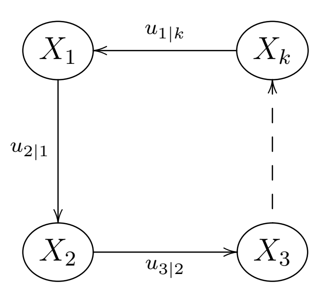

Figure 1: A simple k cycle.

Let us first consider the segment 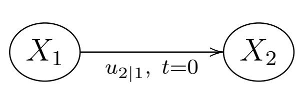

.

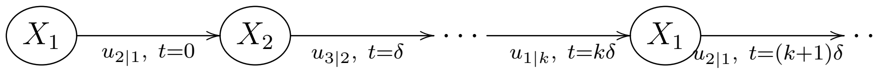

Figure 2: An equivalent dynamic network.

At $t = 0$, X_1's marginal utility mass function is $v_1(\mathbf{a}_1, 0)$. At $t = \delta$, the coordination function is

$$\mathbf{u}_{12}(\mathbf{a}_1, \mathbf{a}_2, \delta) = u_{2|1}(\mathbf{a}_2|\mathbf{a}_1)v_1(\mathbf{a}_1, 0)\,, \tag{14}$$

from which we may compute X_2's marginal:

$$v_2(\mathbf{a}_2, \delta) = \sum_{\mathbf{a}_1} \mathbf{u}_{12}(\mathbf{a}_1, \mathbf{a}_2, \delta)\,. \tag{15}$$

Now consider the next segment 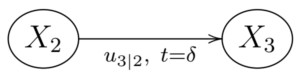

. At $t = 2\delta$ the coordination function is

$$\mathbf{u}_{23}(\mathbf{a}_2, \mathbf{a}_3, 2\delta) = u_{3j2}(\mathbf{a}_3|\mathbf{a}_2)v_2(\mathbf{a}_2, \delta)\,, \tag{16}$$

and X_3's marginal is

$$v_3(\mathbf{a}_3, 2\delta) = \sum_{\mathbf{a}_2} \mathbf{u}_{23}(\mathbf{a}_2, \mathbf{a}_3, 2\delta)\,. \tag{17}$$

We may continue this process for $t = 3\delta$, $t = 4\delta$, etc. To do so, however, it is convenient to introduce matrix notation. Let us denote the elements of $\boldsymbol{\mathcal{A}}$ as follows

$$\boldsymbol{\mathcal{A}} = \{\mathbf{z}_1, \ldots, \mathbf{z}_N\} \tag{18}$$

and define the *utility mass vector*

$$\mathbf{v}_i(t) = \begin{bmatrix} v_i(\mathbf{z}_1, t) \\ \vdots \\ v_i(\mathbf{z}_N, t) \end{bmatrix}. \tag{19}$$

We next define the *state-to-state transition matrix*

$$T_{i+1|i} = \begin{bmatrix} u_{i+1|i}(\mathbf{z}_1|\mathbf{z}_1) & \cdots & u_{i+1|i}(\mathbf{z}_1|\mathbf{z}_N) \\ \vdots & \vdots & \vdots \\ u_{i+1|i}(\mathbf{z}_N|\mathbf{z}_1) & \cdots & u_{i+1|i}(\mathbf{z}_N|\mathbf{z}_N) \end{bmatrix}. \tag{20}$$

With this notation, we may combine the operations defined by (14) and (15) with the single expression

$$\mathbf{v}_2(\delta) = T_{1|2}\mathbf{v}_1(0)\,, \tag{21}$$

and replace (16) and (17) with

$$\mathbf{v}_3(2\delta) = T_{3|2}\mathbf{v}_1(\delta)\,. \tag{22}$$

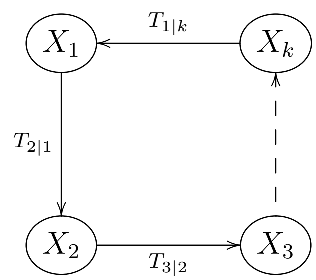

Figure 3: A k-cycle

Figure 3 displays the k-cycle with the linkages represented by the transition matrices. As we trace the path from X_i around the cycle back to X_i, the marginal mass vector v_i is updated k times, where the indices are incremented mod k:

$$\begin{aligned}
\mathbf{v}_{i+1}(\delta) &= T_{i+1|i}\mathbf{v}_i(0) \\
\mathbf{v}_{i+2}(2\delta) &= T_{i+2|i+1}T_{i+1|i}\mathbf{v}_i(0) \\
&\vdots \\
\mathbf{v}_{i+k-1}((k-1)\delta) &= T_{i+k-1|i+k-2}\cdots T_{i+2|i+1}T_{i+1|i}\mathbf{v}_i(0)\,.
\end{aligned}$$

The loop is closed with the final update of the cycle, yielding

$$\mathbf{v}_{i+k}(k\delta) = T_{i+k|i+k-1}T_{i+k-1|i+k-2}\cdots T_{i+2|i+1}T_{i+1|i}\mathbf{v}_i(0) \tag{23}$$

or, since all indices are incremented mod k,

$$\mathbf{v}_i((k\delta) = T_{i|i+k-1}T_{i+k-1|i+k-2}\cdots T_{i+2|i+1}T_{i+1|i}\mathbf{v}_i(0)\,. \tag{24}$$

Now let us define the *closed-loop transition matrix*

$$T_i = T_{i|i+k-1}T_{i+k-1|i+k-2}\cdots T_{i+2|i+1}T_{i+1|i}\,. \tag{25}$$

Also, it is convenient to express time in in units equal to the interval $k\delta$. Thus, we may write (24) as

$$\mathbf{v}_i(1) = T_i\mathbf{v}_i(0) \tag{26}$$

for $i = 1, \ldots, k$. The closed-loop transition matrices for the cycle are as follows.

$$\begin{aligned}
T_1 &= T_{1|k}T_{k|k-1}\cdots T_{3|2}T_{2|1} \\
T_2 &= T_{2|1}T_{1|k}\cdots T_{4|3}T_{3|2} \\
&\vdots \\
T_k &= T_{k|k-1}T_{k-1|k-2}\cdots T_{3|2}T_{2|1}
\end{aligned}$$

After t cycles, we have

$$\begin{aligned}\mathbf{v}_i(t) &= T_i\mathbf{v}_i(t-1)\\ &= T_iT_i\mathbf{v}_i(t-2)\\ &\vdots\\ &= T_i\cdots T_i\mathbf{v}_i(0)\\ &= T_i^t\mathbf{v}_i(0)\,.\end{aligned}$$

The key issue devolves around the behavior of T_i^t as $t \to \infty$. To address this issue, we must explore the convergence properties of this matrix.

4.2 Convergence of Closed-Loop Transition Matrices

Definition 10. *Let $T = [t_{ij}]$ be a square matrix. T is nonnegative, denoted $T \not< 0$, if $t_{ij} \not< 0\,\forall\, i,j$. T is positive, denoted $T \geq 0$, if $t_{ij} \not< 0\,\forall\, i,j$ and $t_{ij} > 0$ for at least one element. T is strictly positive, denoted $T > 0$, if $t_{ij} > 0\,\forall\, i,j$.*

The key theoretical results underlying this approach are the following theorems:

Theorem 1 (Frobeneius-Peron)**.** *If a square matrix T^k is strictly positive for some finite integer k, then T has a unique largest eigenvalue with positive eigenvector.*

Definition 11. *A square matrix T is a* regular transition matrix *if T^k is strictly positive for some finite integer k and each column sums to unity.*

Applying the Frobeneius-Peron theorem to regular transition matrices yields the following result.

Theorem 2 (Markov Convergence)**.** *If T is a regular transition matrix, there exists a unique mass vector $\overline{\mathbf{v}}$ such that*

- $T\mathbf{x} = \overline{\mathbf{v}}$
- $\overline{T} = \lim_{t\to\infty} T^t = \begin{bmatrix}\overline{\mathbf{v}} & \cdots & \overline{\mathbf{v}}\end{bmatrix}$
- $\overline{\mathbf{v}} = \overline{T}\mathbf{v}(0)$ *for every initial mass vector* $\mathbf{v}(0)$

4.3 Closed-Loop Coordination Function

In this discussion, we restrict attention to networks composed of a simple n-cycle. In this case, we may compute the coordination function starting with any member of the cycle and propagate the influence through each member of the cycle, terminating with the cycle that returns to the starting element. As we propagate from X_j in state $\overline{\mathbf{v}}_j$ to X_{j+1}, the state will transit to $\overline{\mathbf{v}}_{j+1}$ according to $\overline{\mathbf{v}}_{j+1} = T_{j+1|j}\overline{\mathbf{v}}_j$, $j = 1,\ldots,n-1$. Starting with X_1, the coordination function is of the form

$$\overline{u}_{1:n}(\mathbf{a}_1,\ldots,\mathbf{a}_n) = u_{n|n-1}(\mathbf{a}_n|\mathbf{a}_{n-1})\cdots u_{2|1}(\mathbf{a}_2|\mathbf{a}_1)\overline{v}_1(\mathbf{a}_1)\,, \tag{27}$$

from which we may compute the corresponding group utility and individual decision functions according to (12) and (13).

5 Bilateral Interdisciplinary Collaboration

We now apply this theory to the collaboration model introduced Section 2. As a specific example, we propose the following scenario. X_r and X_c are collaborating on a paper that is intended to appeal to both computer science (CS) and social science (SS) audiences. X_r is in charge of writing the introduction, and X_c is in charge of writing the discussion. Thus the action sets are

$$\begin{aligned} X_r\colon \quad & \begin{cases} e & \text{write introduction for both } CS \text{ and } SS \text{ audiences} \\ c & \text{write introduction primarily for } CS \text{ audience} \end{cases} \\ X_c\colon \quad & \begin{cases} b & \text{write discussion for both } CS \text{ and } SS \text{ audiences} \\ s & \text{write discussion primarily for } SS \text{ audience} \end{cases} \end{aligned} \tag{28}$$

5.1 Snowdrift Game Formulation

It will be convenient to ascribe numerical parameters to the Snowdrift rankings. Let $\alpha, \pi \in (0,1)$ be such that $\pi > \alpha > 1/2$. Then the payoff matrix displayed in Table 2 conforms to the ordinal rankings given in Table 1. Figure 4 displays the (α, π) region, bounded below by the line $\pi = \alpha$ and to the left by the line $\alpha = 1/2$, of the $[0,1] \times [0,1]$ rectangle that complies with the conditional Snowdrift game. Thus, a one-shot application of the Snowdrift model for any (α, π) values in this region will result in a mixed decision with two conflicting Nash equilibria.

Table 2: Parameterized payoff matrix for the Snowdrift game.

	X_c	
X_r	b	s
e	(α, α)	$(1-\alpha, \pi)$
c	$\pi, 1-\alpha)$	$(1-\pi, 1-\pi)$

$\pi > \alpha > 1-\alpha > 1-\pi$

5.2 Cyclic Conditional Game Formulation

By formulating collaboration as a conditional game, we enable the players to deliberate and seek a truly collaborative solution. Our goal is to define conditional utilities for two players that provide conditional responses for each agent given conjectures for the other that accounts for both the benefits and costs of collaboration. To define these preferences we must consider each agent's response to each of the four possible conjectures for the other. Let us first consider the conjecture (e, b) for X_c. Under this hypothesis, X_c views mutual cooperation as the outcome that should be actualized. Given such a commitment by X_c, if X_r is serious about collaboration, then X_r's highest priority should also be to collaborate, since stubbornly demanding that X_c operate against it's best interest (i.e., ascribing highest priority to (c, s))in

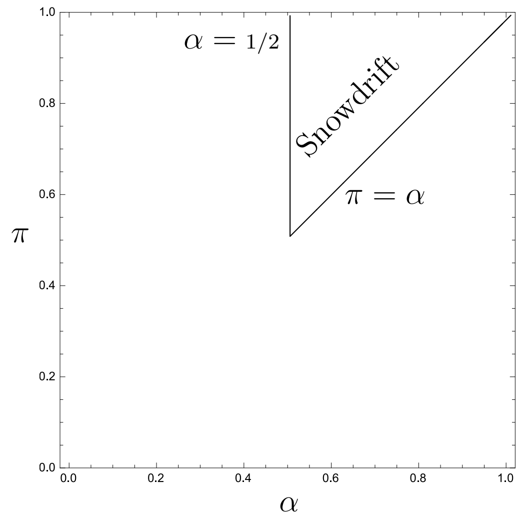

Figure 4: (α, π) regions for the Snowdrift game.

the face of evidence to the contrary would be manifestly counterproductive. But X_r is justified in ascribing the next-best ranking to $(c, b$, the next-worst ranking to (e, s), and the worst ranking to (c, s), all in full accordance with the categorical Snowdrift ordering defined by Table 1. A similar logic applies to X_c given the conjecture (e, b) for X_f. Conjectures (e, s) and (c, b) for X_c are assertions that X_r should either exploit or be exploited by the other. In such cases, X_r's preferences should be aligned with the categorical orderings, with a similar logic applied to X_c's conditional utilities in response to the same conjectures for X_r. Finally, let us consider the conjecture (c, s) for X_c. Offering this conjecture would essentially mean that X_c views the collaboration as a failure and should be terminated. Given such evidence, the rational response by X_r should be to concur since, otherwise, pursuing a doomed collaboration would be counterproductive. Thus, X_r should ascribe its highest ranking also to (c, s), its worst ranking to (e, b), and the intermediate rankings according to the categorical rankings defined by Table 1. By similarity, X_c's conditinal utility conditioned on (c, s) should be constructed by identical logic. Tables 3(a) and 3(b) display, in ordinal form, the conditional utilities for (a) X_r given conjectures for X_c and (b) X_c given conjectures for X_r.

The entries in Tables 3 may be used to parametrically construct the transition matrices $T_{r|c}$ and $T_{c|r}$ between the row and column players in terms of (α, π). The resulting transition matrices are as follows (notice that the entries in these matrices are divided by $1/2$ to ensure

Table 3: Ordinal form conditional utilities for the cyclic conditional collaboration game.

$u_{r|c}(a_{r1}, a_{r2}|a_{c1}, a_{c2})$

X_r	X_c (a_{c1}, a_{c2})			
a_{r1}, a_{r2}	e, b	e, s	c, b	c, s
e, b	4	3	3	1
e, s	3	2	2	3
c, b	2	4	4	2
c, s	1	1	1	4

$u_{c|r}(a_{c1}, a_{c2}|a_{r1}, a_{r2})$

X_c	X_r (a_{r1}, a_{r2})			
a_{c1}, a_{c2}	e, b	e, s	c, b	c, s
e, b	4	3	3	1
e, s	2	4	4	2
c, b	3	2	2	3
c, s	1	1	1	4

Ordinal ranking: 4 = best; 3 = next-best; 2 = next-worst; 1 = worst

that the columns sum to unity):

$$\begin{aligned} T_{r|c} &= {}^1\!/_2 \begin{bmatrix} \pi & \alpha & \alpha & 1-\pi \\ \alpha & 1-\alpha & 1-\alpha & \alpha \\ 1-\alpha & \pi & \pi & 1-\alpha \\ 1-\pi & 1-\pi & 1-\pi & \pi \end{bmatrix} \\ T_{c|r} &= {}^1\!/_2 \begin{bmatrix} \pi & \alpha & \alpha & 1-\pi \\ 1-\alpha & \pi & \pi & 1-\alpha \\ \alpha & 1-\alpha & 1-\alpha & \alpha \\ 1-\pi & 1-\pi & 1-\pi & \pi \end{bmatrix} \end{aligned} \tag{29}$$

where we have assumed that both individuals are governed by the same parameter values. This assumption, although not required by the theory, is made in the interest of simplicity and clarity. Restricting α and π to the unit interval creates no loss of generality, since utilities are invariant with respect to scale. Furthermore, as we shall see, the behavior of this network depends on, and only on, the ratio α/π.

The closed-loop transition matrix T_r is

$$T_r = T_{r|c} T_{c|r} \,. \tag{30}$$

By inspection, we see that these transition matrices are related by the permutation matrix

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} , \tag{31}$$

that is,

$$\begin{aligned} T_{r|c} &= P T_{c|r} P \\ T_{c|r} &= P T_{r|c} P \end{aligned} \tag{32}$$

and, therefore,

$$T_c = T_{c|r} T_{r|c} = P T_{r|c} P P T_{c|r} P = P T_{r|c} T_{c|r} P = P T_r P \,, \tag{33}$$

since $PP = I$, the identity matrix. Furthermore, since $\overline{\mathbf{v}}_r$ is the unique eigenvector of T_r corresponding to the unit eigenvalue, we have

$$\overline{\mathbf{v}}_c = T_c\overline{\mathbf{v}}_c = PT_rP\overline{\mathbf{v}}_c$$

and, pre-multiplying all terms by P yields

$$\begin{aligned} P\overline{\mathbf{v}}_c = PT_c\overline{\mathbf{v}}_c = PPT_rP\overline{\mathbf{v}}_c \\ T_rP\overline{\mathbf{v}}_c \end{aligned} \tag{34}$$

and, since $\overline{\mathbf{v}}_r$ is the unique eigenvector of T_r corresponding to the unit eigenvalues, it follows that $\overline{\mathbf{v}}_r = P\overline{\mathbf{v}}_c$. Similarly, $\overline{\mathbf{v}}_c = P\overline{\mathbf{v}}_r$. Solving for $\overline{\mathbf{v}}_r$, we obtain (via Mathematica)

$$\overline{\mathbf{v}}_r = \begin{bmatrix} v_r(e,b) \\ v_r(e,s) \\ v_r(c,b) \\ v_r(c,s) \end{bmatrix} = \begin{bmatrix} \frac{(\alpha(2-\pi)-(-1+\pi)^2)}{(2+\alpha-\pi)(3-2\pi)} \\ \frac{1+\alpha^2-\alpha\pi}{4+2\alpha-2\pi} \\ \frac{1-\alpha^2+\alpha\pi}{4+2\alpha-2\pi} \\ \frac{1-\pi}{3-2\pi} \end{bmatrix} \tag{35}$$

and, consequently,

$$\overline{\mathbf{v}}_c = \begin{bmatrix} v_c(e,b) \\ v_c(e,s) \\ v_c(c,b) \\ v_c(c,s) \end{bmatrix} = \begin{bmatrix} \frac{(\alpha(2-\pi)-(-1+\pi)^2)}{(2+\alpha-\pi)(3-2\pi)} \\ \frac{1-\alpha^2+\alpha\pi}{4+2\alpha-2\pi} \\ \frac{(1+\alpha^2-\alpha\pi)}{4+2\alpha-2\pi} \\ \frac{1-\pi}{3-2\pi} \end{bmatrix}. \tag{36}$$

The utility vectors $\overline{\mathbf{v}}_r$ and $\overline{\mathbf{v}}_c$ correspond to the steady-state marginal utilities for X_r and X_c, respectively. Using these utilities, we may form the *ex post* payoff matrix displayed in Table 4. This payoff will correspond to a Snowdrift game if

$$\begin{aligned} \overline{v}_r(c,b) > \overline{v}_r(e,b) > \overline{v}_r(e,s) > \overline{v}_r(c,s) \\ \overline{v}_c(e,s) > \overline{v}_c(e,b) > \overline{v}_c(c,b) > \overline{v}_c(c,s)\,. \end{aligned} \tag{37}$$

Table 4: The *ex post* payoff matrix.

	X_c	
X_r	b	s
e	$(\overline{v}_r(e,b), \overline{v}_c(e,b))$	$(\overline{v}_r(e,s), \overline{v}_c(e,s))$
c	$(\overline{v}_r(c,b), \overline{v}_c(c,b))$	$(\overline{v}_r(c,s), \overline{v}_c(c,s))$

Figure 5 illustrates the (α, π) values in $(1/2, 1) \times (1/2, 1)$ where the *ex post* payoff matrix conforms to the Snowdrift scenario—the region to the left of the line of demarcation. To the right of this line, the game becomes a Concord game with ordinal-form payoff matrix given by Table 5, and has a dominant equilibrium (e, b).

Table 5: Payoff matrix for the Concord game.

	X_c	
X_r	b	s
e	(4, 4)	(2, 3)
c	(3, 2)	(1, 1)

Ordinal ranking: 4 = best; 3 = next-best; 2 = next-worst; 1 = worst

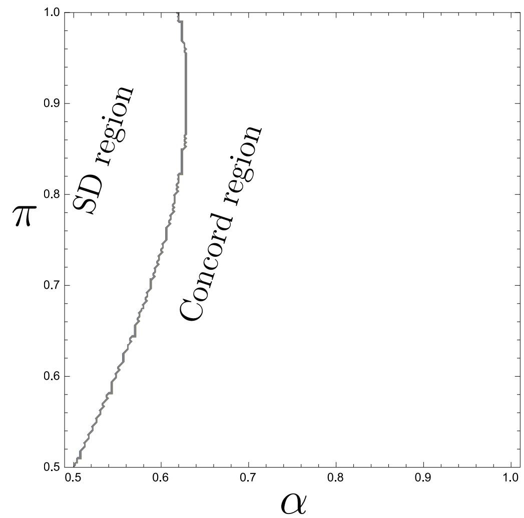

Figure 5: Region for transition from Snowdrift form to Concord form.

5.3 The Coordinated Solution

Although the *ex post* payoff matrix takes into account all of the social relationships that exist between the individuals, applying solution concepts such as Nash equilibria are still fundamentally individually based, and do not take into consideration the creation of a group utility. To gain a more complete assessment of the coordination characteristics of this scenario, it is necessary to examine the group utility and individual decision functions. To do so, we first compute the coordination function as defined in Section 4.3. The steady-state coordination function, via (27), is

$$\overline{u}_{rc}[(a_{r1}, a_{r2}), (a_{c1}, a_{c2})] = u_{c|r}(a_{c1}, a_{c2}|a_{r1}, a_{r2})\overline{v}_r(a_{r1}, a_{r2}) \,, \tag{38}$$

from which we may compute the group utility

$$\overline{w}_{rc}(a_{r1}, a_{c2}) = \sum_{a_{r2}, a_{c1}} \overline{u}_{rc}[(a_{r1}, a_{r2}), (a_{c1}, a_{c2})] \,, \tag{39}$$

yielding

$$\begin{aligned}
\overline{w}(e,b) &= \frac{-3+4\pi^2-2\pi^3+\alpha^2(-7+4\pi)+\alpha(-2+3\pi-2\pi^2)}{4(2+\alpha-\pi)(-3+2\pi)} \\
\overline{w}(e,s) &= \frac{-7+\alpha^2+12\pi-8\pi^2+2\pi^3+\alpha(-6+7\pi-2\pi^2)}{4(2+\alpha-\pi)(-3+2\pi)} \\
\overline{w}(c,b) &= \frac{-7+\alpha^2+12\pi-8\pi^2+2\pi^3+\alpha(-6+7\pi-2\pi^2)}{4(2+\alpha-\pi)(-3+2\pi)} \\
\overline{w}(c,s) &= \frac{-7+\alpha^2(5-4\pi)+4\pi+4\pi^2-2\pi^3+\alpha(2-9\pi+6\pi^2)}{4(2+\alpha-\pi)(-3+2\pi)}\,.
\end{aligned} \tag{40}$$

Because of the symmetric structure of the conditional utilities, we have $\overline{w}(e,s)=\overline{w}(c,b)$; that is, the outcomes where one cooperates and the other defects have equal group utility. The individual decision functions

$$\begin{aligned}
\overline{w}_r(a_{r1}) &= \sum_{a_{c2}} \overline{w}_{rc}(a_{r1},a_{c2}) \\
\overline{w}_c(a_{c2}) &= \sum_{a_{r1}} \overline{w}_{rc}(a_{r1},a_{c2})
\end{aligned} \tag{41}$$

become

$$\begin{aligned}
\overline{w}_r(e) = \overline{w}_c(b) &= \frac{-5+6\pi-2\pi^2+\alpha^2(-3+2\pi)+\alpha(-4+5\pi-2\pi^2)}{2(2+\alpha-\pi)(-3+2\pi)} \\
\overline{w}_r(c) = \overline{w}_c(s) &= \frac{-7+\alpha^2(3-2\pi)+8\pi-2\pi^2+\alpha(-2-\pi+2\pi^2)}{2(2+\alpha-\pi)(-3+2\pi)}\,.
\end{aligned} \tag{42}$$

Straightforward calculations establish that

$$(e,b) = \arg\max_{a_{r1},a_{c2}} \overline{w}(a_{r1},a_{c2}) \tag{43}$$

and

$$\begin{aligned}
\overline{w}_r(e) &> \overline{w}_r(c) \\
\overline{w}_c(b) &> \overline{w}_r(s)
\end{aligned} \tag{44}$$

for all $1/2 < \alpha < \pi < 1$. Thus, even though the game is formulated initially as a Snowdrift game and may even remain so after deliberation, when examined in terms of the emergent group-level and individual-level preferences, the individuals (the parts) combine in an orderly way to form a group (the whole) that is highly coordinated.

Example 1. *Let $(\alpha,\pi)=(0.6,0.9)$. The* ex post *payoff matrix is given by Table 6, where we see that the steady-state game is a Snowdrift. However, the group utility is*

$$\begin{aligned}
\overline{w}(e,b) &= 0.37 \\
\overline{w}(e,s) &= 0.20 \\
\overline{w}(c,b) &= 0.20 \\
\overline{w}(c,s) &= 0.23\,,
\end{aligned} \tag{45}$$

which establishes that the emergent group ordering over all possible outcomes is such that mutual cooperation is the favored outcome for the collective. Also, the individual decision functions are

$$\begin{aligned} \overline{w}_r(e) &= \overline{w}_c(b) = 0.57 \\ \overline{w}_r(c) &= \overline{w}_c(s) = 0.43\,, \end{aligned} \tag{46}$$

which establish that the individually best outcome for each is to cooperate.

Table 6: The *ex post* payoff matrix for $(\alpha, \pi) = (0.6, 0.9)$.

	X_c	
X_r	b	s
e	$(0.33, 0.33)$	$(0.24, 0.35)$
c	$(0.35, 0.24)$	$(0.08, 0.08)$

Example 2. *Let* $(\alpha, \pi) = (0.6, 0.65)$*. The* ex post *payoff matrix is given by Table 7, where we see that the steady-state game is a Concord game. The group utility is*

$$\begin{aligned} \overline{w}(e, b) &= 0.30 \\ \overline{w}(e, s) &= 0.23 \\ \overline{w}(c, b) &= 0.23 \\ \overline{w}(c, s) &= 0.24 \end{aligned} \tag{47}$$

and the individual decision functions are

$$\begin{aligned} \overline{w}_r(e) &= \overline{w}_c(b) = 0.53 \\ \overline{w}_r(c) &= \overline{w}_c(s) = 0.47\,. \end{aligned} \tag{48}$$

For this example, we see that the ex post *Nash equilibrium solution also corresponds to the coordinated decision.*

Table 7: The *ex post* payoff matrix for $(\alpha, \pi) = (0.6, 0.65)$.

	X_c	
X_r	b	s
e	$(0.28, 0.28)$	$(0.25, 0.26)$
c	$(0.26, 0.25)$	$(0.21, 0.21)$

6 Conclusion

Be it prehistoric collective hunting (Gavrilets, 2015) or contemporary international research networks, collaboration requires the ability and willingness to align goals and efforts with one another to be successfully carried out. When individuals are able to collaborate, the desired outcome can be obtained with less individual effort and more efficiently. In this paper we have used the Snowdrift game as a stylized strategic encounter to capture the challenge faced by individuals who take part in collaborative endeavours. Rational egoists in this game are victims of the so-called "tragedy of the commune" (Doebeli et al., 2004): in a cooperative system where each individual contributes to a common good and benefit from its own efforts, some individuals may end up investing a lot while others do not invest anything at all. Though this inefficient outcome is a fact of life, it marks also the absence of any collaborative ability. In contrast, if we relax the standard assumption of self-interest and allow for context-dependent social influence relationships between the agents, we are free to identify the social mechanism that is behind more efficient solutions. Here, we have presented a dynamic model of "catch-and-toss" collaboration in which two agents are able to reciprocally adjust their preferences and reach a steady-state that is conducive to the cooperative outcome. More complex structures and dynamics are matters of further research.

A Proof of Isomorphism Lemma

Proof. Without loss of generality, we restrict attention to a two-agent collective $\{X_1, X_2\}$ defined over the product set $\boldsymbol{\mathcal{A}} = \mathcal{A}_1 \times \mathcal{A}_2$, with X_1 possessing a categorical utility $u_1\colon \boldsymbol{\mathcal{A}} \to \mathbb{R}$ and X_2 possessing a family of conditional utilities $\{u_{2|1}(\cdot|\mathbf{a}_1)\colon \boldsymbol{\mathcal{A}} \to \mathbb{R}\ \forall\ \mathbf{a}_1 \in \boldsymbol{\mathcal{A}}\}$. Let Ω_1 and Ω_2 be arbitrary sets of propositions of distinct elements with cardinalities equal to the cardinalities of $\mathcal{A}_1$ and $\mathcal{A}_2$, respectively, and let $\boldsymbol{\Omega} = \Omega_1 \times \Omega_2$. Let $g\colon \boldsymbol{\mathcal{A}} \to \boldsymbol{\Omega}$ be a bijective mapping $g\colon \mathbf{a} \mapsto \boldsymbol{\omega}$, and define functions $b_1\colon \boldsymbol{\Omega} \to \mathbb{R}$ and $\{b_{2|1}(\cdot|\boldsymbol{\omega}_1)\colon \boldsymbol{\Omega} \to \mathbb{R}\ \forall \boldsymbol{\omega}_1 \in \boldsymbol{\Omega}\}$ such that

$$b_1(\boldsymbol{\omega}) = u_1[g^{-1}(\boldsymbol{\omega})] = u_1(\mathbf{a}) \tag{49}$$

and

$$b_{2|1}(\boldsymbol{\omega}_2|\boldsymbol{\omega}_1) = u_{2|1}[g^{-1}(\boldsymbol{\omega}_2)|g^{-1}(\boldsymbol{\omega}_1)] = u_{2|1}(\mathbf{a}_2|\mathbf{a}_1)\,. \tag{50}$$

This construction defines a marginal belief function b_1 over $\boldsymbol{\Omega}$ such that, for $\boldsymbol{\omega}, \boldsymbol{\omega}' \in \boldsymbol{\Omega}$, $b_1(\boldsymbol{\omega}) \geq b_1(\boldsymbol{\omega}')$ means that the belief that $\boldsymbol{\omega}$ will be realized is at least as great as the belief that $\boldsymbol{\omega}'$ will be realized. It also defines a family of conditional belief functions $\{b_{2|1}(\cdot|\boldsymbol{\omega}_1)\colon \boldsymbol{\Omega} \to \mathbb{R}\ \forall \boldsymbol{\omega}_1 \in \boldsymbol{\Omega}\}$ such that $b_{2|1}(\boldsymbol{\omega}_2|\boldsymbol{\omega}_1) \geq b_{2|1}(\boldsymbol{\omega}_2'|\boldsymbol{\omega}_1)$ means that the belief that $\boldsymbol{\omega}_2$ is realized is at least a great as the belief that $\boldsymbol{\omega}_2'$ is realized, given that $\boldsymbol{\omega}_1$ is realized. This mapping establishes the structural equivalence of the preference criterion regarding $\boldsymbol{\mathcal{A}}$ and the belief criterion regarding $\boldsymbol{\Omega}$, thereby establishing the order isomorphism. □

References

J. Adams. Collaborations: The fourth age of research. *Nature*, 497:557–560, 2013.

R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.

M. Bacharach. *Beyond Individual Choice: Teams and Frames in Game Theory*. Princeton University Press, Princeton, NJ, 2006.

P. Battigalli and M. Dufwenberg. Dynamic psychological games. *Journal of Economic Theory*, 144:1–35, 2009.

C. Bicchieri. *Rationality and Coordination*. Cambridge University Press, Cambridge, UK, 1993.

G. E. Bolton and A. Ockenfels. A stress test of fairness measures in models of social utility. *Economic Theory*, 24(4), 2005.

W. Bossert, C. X. Qi, and J. A. Weymark. Measuring group fitness in a biological hierarchy: an axiomatic social choice approach. *Economics and Philosophy*, 29:301–323, 2012.

C. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton Univ. Press, Princeton, NJ, 2003.

C. Camerer, G. Lowenstein, and M. Rabin, editors. *Advances in Behavorial Economics*. Princeton Univ. Press, Princeton, NJ, 2004a.

C. Camerer et al. *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-scale Societies*. Oxford University Press, Oxford, UK, 2004b.

R. G. Cowell, A. P. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter. *Probabilistic Networks and Expert Systems*. Springer Verlag, New York, NY, 1999.

M. Doebeli, C. Hauert, and T. Killingback. The evolutionary origin of cooperators and defectors. *Science*, 306:859Äì862, 2004.

M. Dufwenberg and G. Kirchsteiger. A theory of sequential reciprocity. *Games and Economic Behavior*, 47:268–298, 2004.

N. H. Fefferman and K. L. Ng. The role of individual choice in the evolution of social complexity. *Annales Zoologici Fennici*, 44:58–69, 2007.

E. Fehr and K. Schmidt. A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114:817–868, 1999.

S. Gavrilets. Collective action and the collaborative brain. *Journal of The Royal Society Interface*, 102(12), 2015.

J. Geanakoplos, D. Pearce, and E. Stacchetti. Psychological games and sequential rationality. *Games and Economic Behavior*, 1:60–79, 1989.

M. Grayson and S. Pincock. Nature index 2015 collaborations. *Nature*, 527:S49, 2015.

J. Henrich et al. "Economic Man" in Cross-culturative Perspective: Behavorial Experiments in 15 Small-scale Societies. *Behavioral and Brain Sciences*, 28(6):795–855, 2005.

B. Huebner, R. Kukla, and E. Winsberg. Making an author. In L.T. Boyer, C. Mayo-Wilson, and M. Weisberg, editors, *Scientific collaboration and collective knowledge*. Oxford University Press, Oxford, forthcoming.

K. Itô, editor. *Encyclopedic Dictionary of Mathematics*, volume II.12 (311 E). MIT Press, Cambridge, MA, 1987.

F. V. Jensen. *Bayesian Networks and Decision Graphs*. Springer Verlag, New York, NY, 2001.

S. L. Lauritzen. *Graphical Models*. Springer Verlag, New York, NY, 1996.

J. A. H. Murray, H. Bradley, W. A. Craigie, and C. T. Onions, editors. *The Compact Oxford English Dictionary*, Oxford, UK, 1991. The Oxford Univ. Press.

J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, CA, 1988.

M. Rabin. Incorporating fairness into game theory and economics. *American Economic Review*, 83:1281–1302, 1993.

C. Sartorius. *An Evolutionary Approach to Social Welfare*. Routledge, London, UK, 2003.

W. C. Stirling. *Theory of Conditional Games*. Cambridge University Press, Cambridge, UK, 2012.

R. Sugden. *The Economics of Rights, Co-operation, and Welfare*. Basil Blackwell, Oxford, 1986.

R. Sugden. Team preferences. *Economics and Philosophy*, 16:175–204, 2000.

R. Sugden. The logic of team reasoning. *Philosophical Explorations*, 6:165–181, 2003.

C. Thoni and S. Gachter. Peer effects and social preferences in voluntary cooperation: A theoretical and experimental analysis. *Journal of Economic Psychology*, 48:72–88, 2015.

S Wuchty, B F Jones, and B Uzzi. The increasing dominance of teams in production of knowledge. *Science*, 316(5827):1036–1039, May 2007.