Theory of Coordinated Agency

Wynn C. Stirling
Department of Electrical and Computer Engineering
Brigham Young University

Workshop on Conditional Games and Strategic Risk
hosted by
The Center for the Economic Analysis of Risk
Villino Volterra, Rome, Italy
19-20 September 2016

Abstract

Part I: Theory of Coordinated Agency — Coordinated decision making has long been the topic of philosophical debate. In addition to standard game-theoretic models of coordination based on individual rationality, explicitly group-level approaches, such as joint intentionality and team reasoning, have been recently advanced as rational concepts of coordinated group behavior. Whether the approaches are based on either individual or social notions of rationality, however, coordination as discussed in the extant literature is implemented via solution concepts based on psychological, sociological, and economic considerations that are exogenously attributed to the individuals after individual preferences are established and the game's mathematical structure has been defined. Rather than focusing on *ex post* solution concepts that overlay the preference model, this paper takes the position that social considerations should be incorporated *ex ante* into the individual preference models, thereby enabling an endogenous concept of coordinated behavior to emerge as the game is engaged. Conditional game theory provides a framework within which to explore this alternative perspective. The key feature of this approach is to apply the concepts of Bayesian conditionalization to enable agents to modulate their individual preferences by conditioning them on the intentions of others, thereby actively responding to social influence. This paper establishes a coordination mechanism that generates an emergent social model as conditional preferences propagate through the group. The result is an operational definition of coordinated agency and the creation of coordinated decision rules. The theory is expanded to account for cyclic influence propagation by the use of the Markov convergence theorem to establish the convergence of reciprocal preference propagation. Additionally, the concept of mutual information, as developed via Shannon information theory, is applied to define a quantitative measure of the innate ability of a group to coordinate as a result of social influence. Finally, the theory is extended to account for stochastic entities that can influence the behavior of deterministic agents and vice versa.
Part II: Application — The theory is applied to Bacharach's three puzzles: Hi-Lo, Matching Pennies, and the Prisoner's Dilemma.

Contents

Part I: Theory of Coordinated Agency

1 Introduction

> Order is not pressure which is imposed on society from without, but an equilibrium which is set up from within.
>
> José Ortega y Gasset, *Mirabeau: An Essay on the Nature of Statesmanship*

Coordinated multiagent decision making has long been an important subtopic of game theory. Seminal works by Schelling (1960), Lewis (1969), Bicchieri (1993), and many others provide in-depth analyses of coordination scenarios. Schelling introduces the notions of "tacit coordination," where "the player's objective is to make contact with the other player through some imaginative process of introspection, of searching for shared clues" (Schelling, 1960, p. 96). Lewis (1969) argues that social convention is the basis for much of coordinated behavior. Bicchieri (1993) proposes that coordination is learned as the end result of social evolution. Cooper (1999) expresses coordination in terms of technological complementarity. Recently, Sugden (1993, 2000, 2003, 2014) and Bacharach (1999, 2006) have proposed concepts of team reasoning, where individuals identify as members of a team and modify their choices to conform with team aspirations.

Team reasoning is part of a larger conversation on collective intentionality, the study of group action in terms of the mental states of the individuals, and involves notions of "we-intention" and "joint commitments" (Searle, 1990, 1995; Gilbert, 1989) that, it is argued, cannot be constructed from individual-level states. Bratman (1993, 1999, 2014) introduces a concept of "augmented individualism," and asserts that there is no discontinuity between individual and joint intentionality Rather, shared intentions consist primarily of interrelated attitudes of the individuals. Ross (2014) is skeptical of building joint behavioral models that rely on internal mental phenomena, arguing that, since mental states require context, they do not provide an objective account of the world. Although intentional states may be needed to describe a scientific worldview, they are best understood as descriptions of coupled informational and behavioral patterns.

Bratman's and Ross's observations provide a natural point of departure for this paper. We assume that all relevant mental states have been incorporated into each individual's preference model, and that the key issue is to understand how these preferences generate coordinated behavior. 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). To comply with this definition, coordination must be more than cooperation. Even purely selfish individuals may cooperate if their interests happen to coincide, and no notion of group behavior need be relevant. Coordination, however, is a more complex concept that involves the relationship between individual and group behavior. For coordination to occur, the individuals (the parts) must combine their behavior to form a properly constructed group (the whole).

This paper develops a concept of *coordinated agency* that is distinct from approaches in the extant literature. On the one hand, it differs from Bacharach's (2006) team-reasoning in that it does not require individuals to undergo agent and utility transformations from individual to group, and it also differs from Sugden's (2014) concept of intentional cooperation, where individuals reason as members of a group with the intention of playing their parts in the interest of mutual benefit. On the other hand, it is also distinct from conventional game-theoretic notions of coordination which retain the individual game-theoretic payoff structure and focus on coordinated solution concepts. We propose a middle ground that retains the game-theoretic assumption of individual preferences, but extends the concept of individual rationality to include the incorporation of the interests of others into one's own interest, thereby enabling the emergent creation of mutual interest as a consequence of social relationships. We establish an explicit operational definition of coordination by introducing specific properties of the concept in terms of its existence and quantity. We a) present a mathematical framework within which intentional attributes may be incorporated into individual preferences, b) introduce a social interaction mechanism that produces endogenously coordinated multiagent decisions, and c) provide an explicit metric to express the degree to which a group possesses the intrinsic ability to coordinate as a function of its social structure.

1.1 Coordination Vis-à-vis Performance

We focus on the coordination of groups whose members possess the ability to respond to the social influence exerted on them by each other. Examples include cooperative groups, such as teams and business entities, mixed organizations such as families, which can encompass both cooperative and conflictive influence, and adversarial groups such as athletic contests and military engagements that express conflictive coordination. Team members coordinate by cooperating in the pursuit of a common goal, business partners coordinate by dividing the labor, family members coordinate by respecting (or not) each other's opinions and priorities, and military opponents coordinate by opposing each other in some systematic way.

Coordination is a principle of behavior on a parallel with, but different from, performance. *Individuals perform; groups coordinate.* Performance deals with operational measures of efficiency and effectiveness of individual behavior. Coordination, however, is an attribute of organizational structure regarding how the members of a group function together. In terms of overall functionality, it is often the case that the propensity of a group to coordinate is more relevant than the propensity of the individuals to optimize. It is more relevant for a team to win the game than for each player to maximize the number of points he or she scores. It is more relevant for a business entity to settle on a productive division of labor than for each partner to maximize individual control. It is more relevant for a family to function in a civil and equitable way than for the members to focus exclusively on what is individually best for themselves. It is more relevant to the conducting of a war for each opponent to seek victory rather than simply to destroy as many enemy resources as possible.

Focusing on performance without considering coordination is an incomplete characterization of group behavior. Similarly, focusing on coordination without considering performance is an incomplete characterization of individual behavior. A football team may possess the organizational structure required to win the game, but that structure is useless if the players do not attempt to maximize the number of goals scored. A business firm may be well organized in terms of individual responsibilities, but unless the partners exert control, the entity will not prosper. A family may possess fair and equitable rules of conduct but will still be dysfunctional if the members do not pursue their individual goals within that context. An army may have a strategic plan to win the war, but that still requires the solders to fight effectively. Coordination without performance is unproductive, and performance without coordination is equivocal. A full understanding of the functionality of a group requires the assessment of both attributes. Coordination occurs when individual contributions appropriately fit together to form a coherent organizational structure.

The synergistic relationship between coordination and performance paves the way for the reconciliation of group and individual rationality. Specifically, we view group rationality in terms of coordination, and we view individual rationality in terms of performance. Developing this relationship, however, requires an operational definition of coordination that meshes with the operational definition of performance. A truly operational notion of coordination must a) establish a mathematical representation of material interests and social relationships; b) define a mechanism to combine the individual interests (the parts) to create a coordinated society (the whole); and c) define solution concepts through which the group achieves coordinated behavior and the individuals achieve individually rational performance within the group context.

1.2 The Role of Game Theory

Game theory occupies a dominant position as a framework within which to model coordinated decision making (cf. Schelling (1960); Lewis (1969); Bicchieri (1993); Cooper (1999); Goyal (2007); Bacharach (1999, 2006); Sugden (1993, 2000, 2003, 2014); Jackson (2008); Shoham and Leyton-Brown (2009); Easley and Kleinberg (2010); Gintis (2016)). A complex social problem is defined and factors that are deemed to be relevant are encoded into mathematical expressions, while those factors considered to be irrelevant are ignored. The classical game-theoretic approach is to make minimal assumptions about individual preferences and investigate how they interact. The standard approach is to endow each agent (player) with an individual action set and a linear (i.e., reflexive, antisymmetric, transitive, and complete) preference ordering over the set of joint actions (profiles) that produce its payoffs. Such preference orderings are assumed to account for all of the interests that can affect the individual's welfare. Once specified, they are *categorical:* unconditional, fixed, and immutable. The payoffs are then juxtaposed into a payoff array, and each agent invokes a strategy according to its solution concept.

A key attribute of the classical approach is that it divides the labor between the specification of preferences and the specification of solution concepts. It complies with the observation by Friedman (1962, p. 13) that "economic theory proceeds largely to take wants as fixed. This is primarily a division of labor. The economist has little to say about the formation of wants; this is the province of the psychologist. The economist's task is to trace the consequences of any given set of wants." Friedman's model is designed for scenarios where individual behavior is governed by narrowly construed self-interest and payoffs are expressed in terms of material performance. The critical feature of such an environment is that individuals behave *reactively.* They form their strategies according to pre-calculated preferences as a reaction to the possible strategies that others can invoke.

Since it is based on categorical individual preferences, classical game theory is expressly designed to study individual performance. It is *not* designed to account for coordination. In fact, attempting to deduce any notion of group rationality from individual rationality is problematic. Shubik (1982) cautions against the "anthropomorphic trap" of building on "the shaky analogy between individual and group psychology," and argues that "It may be meaningful, in a given setting, to say that group 'chooses' or 'decides' something. It is rather less likely to be meaningful to say that the group 'wants' or 'prefers' something .... Game theory makes a special point of *not* requiring 'society' to be a generalized person, capable of making choices and judgments among actions or outcomes on the basis of some sort of welfare function" (Shubik, 1982, p. 123-124). Luce and Raiffa (1957) also argue that "the notion of group rationality is neither a postulate of the model nor does it appear to follow as a logical consequence of individual rationality" (p. 193), and conclude that "it may be too much to ask that any sociology be derived from the single assumption of individual rationality" (p. 196).

Simply put, relying on categorical preferences to model coordinated behavior limits the ability to account for the social relationships that are necessary for the proper combining of the parts to form a whole, since there no explicit mechanism to account for interrelationships. Although the members of a team may have individual preferences that are consistent with cooperative behavior, if there is no social linkage, they may each choose behaviors that, at least ostensibly, are consistent with coordination but, in reality, are governed by individual goals and may result in poor results. A family that does not coordinate is likely to be dysfunctional. If opposing armies have individual preferences that are consistent with fighting but do not engage in reconnaissance or some other form of coordination, their behaviors may not be consistent with the group-level behavior of a war.

It is tempting to encode social interests into categorical individual preferences, but doing so runs the risk of conflating social issues with material issues. For example, including parameters to account for "self-centered inequity aversion" (fairness), as introduced by Fehr and Schmidt (1999), is an attempt to embed a fundamentally social concern into the evaluation of individual material payoff. Essentially, it puts a price on fairness. Social interests, however, cannot be easily expressed as material payoffs or other tangible or easily quantifiable rewards. Rather, they are intentions that connect the individuals to each other in ways that enable them to achieve appropriate material payoffs in a social context. To pursue such payoffs, individuals must possess the ability to *respond* to social influence. To illustrate, suppose an athletic team member can either attempt to score a goal directly or pass to a team member who has a better shot. Both options are consistent with the material performance of scoring a goal, but the latter option requires a connection to the other player in pursuit of the social interest of winning the game.

The question, therefore, is, where should such a connecting mechanism reside? Relegating the mechanism to the solution concept that *ex post* overlays the payoff model perpetuates the division of labor between modeling preferences and modeling rational behavior. Notions of coordination that arise through the solution concept, rather than explicitly encoded into the payoff structure, are *extrinsic*. This concept of coordination is imposed exogenously on the network — it is not an innate property of an explicitly defined social structure.

Noncooperative game theory is based on the assumption (at least for single play) that the players are unable to communicate during play, for if they did, they could renegotiate their preferences and define a different game. Instead, game theory expects each player to have taken all relevant issues into consideration at the moment of truth when the game is engaged. But if individuals are responsive to social influence, then, as they engage, social relationships will develop dynamically, thereby enabling an endogenous notion of *intrinsically* coordinated behavior to emerge as the parts (the individuals) combine in a systematic way to form a whole (the society). A key issue in this regard is how to express such emergent social behavior with a mechanism that does not violate the communication assumption.

Conditional game theory, as developed by Stirling (2012), provides a framework that accommodates the modeling of direct social influence and, consequently, enables the creation of an operational definition of coordination that comports with the etymological and dictionary interpretations of the concept. The key features of conditional game theory are as follows:

- Societies are modeled as networks — directed graphs with individuals as the vertices and social influence linkages as the edges.
- Categorical individual preferences are replaced by conditional individual preferences that enable players to incorporate the interests of others into their own rationality.
- Social relationships are formed as social influence propagates through the network, thereby creating an emergent comprehensive coordination model.

- An operational definition of coordination is defined in terms of a group-level coordination metric, from which coordinated utilities may be derived.
- The coordination model may be used to generate a *coordination index* as an explicit measure of *coordinatability*, or the theoretical degree to which the social structure of the group enables coordination.
- Conditions are established for individuals to deliberate together to converge to a coordinated solution.
- Stochastic entities that influence, and can be influenced by, decision-making agents may be incorporated into the network.

A classical noncooperative game becomes a special case of conditional game for a network with no edges, that is, a group where the preferences of all of the players are categorical. In this case, there is no explicit social influence. This does not mean, of course, that the group cannot coordinate extrinsically. But if they do, coordination is because of the coincidental alignment of their individual categorical preferences and an appropriately applied solution concept, rather than because of direct social influence.

Part I is organized as follows. In Section 2 we set notation, extend the concept of individual rationality to incorporate the influence of others into one's own rationality, and merge game theory and network theory to define a conditional network game. Section 3 builds on this structure and reviews the basic concepts of conditional game theory. Section 4 invokes Markov Convergence theorem to extend conditional game theory to account for networks that contain influence cycles, and Section 5 applies Shannon information theory to the coordination problem and introduces the coordination index as a measure of the intrinsic coordinatability of a network. Section 6 then extends the theory further to account for stochastic agents. A discussion of results is offered in Section 7.

2 Game Theory Models

> Complexity is no argument against a theoretical approach if the complexity arises not out of the theory itself but out of the material which any theory ought to handle.
>
> Frank R. Palmer, *Grammar*

2.1 Classical Game Theory Models

Definition 1. *A* 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* $(k_1, \ldots, k_n)$*th entry is the sub-array* $[u_1(z_{ik_1}, \ldots, z_{nk_n}), \ldots, u_n(z_{ik_1}, \ldots, z_{nk_n})]$*, where* $u_i(z_{ik_1}, \ldots, z_{nk_n})$ *is the payoff that* 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 multiagent decision making. One of the reasons for its success is the virtually perfect match between the mathematical structure of the payoffs and the notion of individual rationality used to render a decision. If one is motivated by narrow self-interest, a 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, a natural solution concept is to choose the outcome that maximizes individual performance. But this straightforward model of rational behavior can oversimplify and may even distort the decision problem. Arrow has identified its limitations.

> Rationality in application is not merely a property of the individual. Its useful and powerful implications derive from the conjunction of individual rationality and other basic concepts of neoclassical theory — equilibrium, competition, and completeness of markets. . . . When these assumptions fail, the very concept of rationality becomes threatened, because perceptions of others and, in particular, their rationality become part of one's own rationality (Arrow, 1986, p. 203).

Despite Arrow's implicit warning, reliance on categorical utilities as the vehicle with which to express preferences has essentially remained unchanged, although it has not remained unchallenged. Sen notably threw down the gauntlet long ago:

> A person is given *one* preference ordering, and as and when the need arises this is supposed to reflect his interests, represent his welfare, summarize his idea of what should be done, and describe his actual choices and behavior. Can one preference ordering do all these things? A person thus described may be "rational" in the limited sense of revealing no inconsistencies in his choice behavior, but if he has no use for these distinctions between quite different concepts, he must be a bit of a fool. The *purely* economic man is indeed close to being a social moron. Economic theory has been much preoccupied with this rational fool decked in the glory of his *one* all-purpose ordering. To make room for the different concepts related to his behavior **we need a more elaborate structure** [italic emphasis in original, bold emphasis added] (Sen, 1977, pp. 335-336).

A narrow interpretation of individual rationality may be appropriate when individuals function in an economic environment governed by the *price system* (Hayek, 1945; Friedman, 1962). Under this system, prices guide both users and providers of products as they make decisions regarding the various transactions they undertake. The price system frees individuals to focus their attention on, and only on, their own interests, since the social effects of their behavior are automatically regulated. If an individual changes the price of some product in the interest of its own welfare, that signal will propagate through the society and others will respond by adjusting their demand for the product in the interest of their individual welfare. This group-level automatic regulation mechanism makes it possible to justify pursuing one's own interests without concern for the welfare of others. As Arrow (1974, p. 21) observes, "It makes a virtue out of selfishness."

Perhaps few would dispute that the price system is a valuable characterization of behavior that fits economic scenarios where markets clear and competition dominates behavior. But, as Arrow (1974, p. 22) observes, "it cannot be made the complete arbiter of social life." In particular, the assumptions imposed by the price system are not applicable to the formation of teams, where it is more natural to consider choices in the light of group-level behavior. Replacing solution concepts based on narrowly construed self-interest with socially amenable solution concepts address a significant shortcoming of classical game theory as a model of human behavior in social environments that are not based on competition and the clearance of markets. Nevertheless, doing so continues to rely on the mathematical model structure used by classical game theory — each individual comes to the game with completely defined *ex ante* categorical preferences. Although these preferences may be transformed into group-level preferences by an *ex post* psychological mechanism, the underlying mathematical structure of the game remains unchanged. It is not sufficient, however, simply to assert that individuals behave in a certain way because the are motivated by psychological considerations. This is the nub of the issue. As Elster argues, an action cannot be considered rational unless it can be explained and justified.

> Once we have constructed a normative theory of rational choice, we may go on to employ it for explanatory purposes. We may ask, that is, whether a given action was performed *because* it was rational. To show that it was, it is not sufficient to show that the action was rational, since people are sometimes lead by accident or coincidence to do what is in fact best for them. We must show, in addition, that the action arose in the proper way, through a proper kind of connection to desires, beliefs, and evidence [emphasis in original] (Elster, 1986, p. 2).

One might rightly ask, in the light of Elster's assertion: Is there evidence in the mathematical structure of the game that the players are disposed to act in accordance with any particular notion of rational behavior? And if the model provides no such evidence, then one might ask for the source of the connections between the preference model and rational behavior.

To establish a concept of socially amenable rational choice in a way that complies with Elster's injunction, the connections must be directly built into the preference models *ex ante*, rather than applied *ex post* as an overlay of a structure that is designed from the perspective of narrowly construed self-interest. In fact, such an incorporation is essential to the synthesis of artificial multiagent systems which must be designed to function in accord with socially sophisticated human behavior. It is one thing, from a third-party analysis perspective, to attribute observed behavior to psychological motives, but it is quite another thing to imbue the members of a group of artificial agents with "personalities" that exhibit the desired social attributes. This is the distinction between non-causal and causal models. As an analysis tool, a model is non-causal — it may explain or predict behavior, but it does not dictate behavior. But as a synthesis tool, a model is causal — it generates the behavior of the artificial entities. In its role as an analysis tool, a model is used to reduce reality to an abstraction, but in its role as a synthesis tool, it is used to create a (man-made, or artificial) reality from an abstraction.

2.2 Network Theory Models

A natural way to incorporate social considerations into a group 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. It must be emphasized that influence is inherently neutral. It can be positive, in the sense of representing cooperative intentions, it can be negative, in the sense of representing conflictive intentions, or it can be mixed, where propensities for both cooperation and conflict co-exist.

A response for Sen's (1977) call for a "more elaborate structure" is to expand beyond categorical preferences by incorporating the influence linkages into the preference model. To introduce such a structure, consider the simple network scenario involving two agents, X_1 and X_2, as illustrated by the directed graph

$$\textcircled{X_1} \longrightarrow \textcircled{X_2}\,, \tag{1}$$

where the direction of the arrow indicates that X_1 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 onto the influencee.

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

Directionality and conditionality 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 a Bayesian 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 event that B is realized as governed by the conditional probability. The great strength and wide applicability of probability theory is the facility to deal with hypothetical propositions. Indeed, one might say that this facility is its *raison d'être*. As succinctly expressed by Glenn Shafer (cited in Pearl (1988, p. 15)): "probability is not really about numbers; it is about the structure of reasoning."

We employ this same conditionalization logic to model preferences. In this case, the antecedent is a hypothesis that the outcome corresponding to a profile $\mathbf{a} \in \boldsymbol{\mathcal{A}}$ is intended by the influencer. Conditioned on that hypothesis, the consequent is that the outcome corresponding to a profile $\mathbf{a}' \in \boldsymbol{\mathcal{A}}$ is intended by the influencee, and the influencee responds by defining its conditional preference ordering according to its psychological disposition. Once the antecedent is specified, the resulting preference ordering is a conventional linear ordering over the set of outcomes. This is the mechanism by which influencees may express their intentions, such as social propensities for cooperation, fairness, and altruism, or overtly antisocial propensities such as conflict, avarice, and malevolence, as conditional responses to each of the hypothesized intentions of the influencer.

Modeling the influence mechanism according to the logical structure of conditionalization is compatible with the concept of virtual bargaining, where individuals posit "agreements that the social participants anticipate they would make, were they to engage in explicit bargaining ... [and] operates within the framework of rational-choice theory ... [by] extend[ing] the scope of rational-choice models of interaction" Misyak et al. (2014, p. 512). Conditionalization thus serves as a mathematical vehicle with which to conduct though experiments such as virtual bargaining. An individual may conduct a thought experiment and determine its response for each of the possible intentions of those who influence it, thereby providing maximum flexibility in responding, especially if the influencer is itself an influencee of a third agent. In such a case, both influencees will be uncertain regarding their preferences until the preferences of the influencers are determined. Uncertainty, in this context, is *not* regarding beliefs; rather it is uncertainty regarding preference. The conventional application of the probability syntax is as a means of expressing epistemological uncertainty regarding belief, but this same logical structure may be used to expressing behavioral uncertainty regarding preference. 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 action should be taken. To comply with this logical structure, we focus on preference structures that admit conditionalization. It is convenient to employ the structure and syntax of graph theory.

Definition 2. *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 *of length* k *from vertex* X_{i_1} *to vertex* X_{i_k} *is a sequence* $\{X_{i_1}, \ldots, X_{i_k}\}$ *of distinct vertices such that an edge exists between* X_{i_j} *and* $X_{i_{j+1}}$*,* $j = 1 \ldots, k-1$*. A path never crosses itself, and movement along the path never violates the directed-edge condition. If all of the edges along a path are directed, then the path is a* directed path*. We write* $X_i \mapsto X_j$ *if there is a directed path from* X_i *to* X_j*. A path is a* cycle*, or* closed path*, if* $X_j \mapsto X_j$ *for any* X_j*. A graph is said to be a* directed acyclic graph *if all edges are directed and there are no cycles.*

Definition 3. *The* parent *set for* X_i*, denoted* $\mathrm{pa}\,(X_i) = \{X_j\colon X_j \rightarrow X_i\}$*, is the subset of individuals that directly influence* X_i*. If* X_i *has* $q_i > 0$ *parents, then* $\mathrm{pa}\,(X_i) = \{X_{i_1}, \ldots, X_{i_{q_i}}\}$*, where* $X_{i_k} \rightarrow X_i$*,* $k = 1, \ldots, q_i$*. For notational convenience, let* $\mathrm{pa}\,(i) = \{i_1, \ldots, i_{q_i}\}$ *denote the indices corresponding to the elements of* $\mathrm{pa}\,(X_i)$*.*

Definition 4. *A* conjecture *by* X_i*, denoted* $a_{ii} \in \mathcal{A}_i$*, is a hypothetically intended action by* X_i*, and we write* $X_i \models a_{ii}$*. A conjecture for* X_j *by* X_i*, denoted* a_{ij}*, is an action that* X_i *hypothesizes as intended by* X_j*. A* conjecture profile *by* X_j*, denoted* $\mathbf{a}_i = (a_{i1}, \ldots, a_{in}) \in \boldsymbol{\mathcal{A}}$*, is a profile of hypothetically intended actions for the group* $\{X_1, \ldots, X_n\}$*, where* $a_{ij} \in \mathcal{A}_j$*. If* $\mathbf{a}_i$ *is a conjecture profile by* X_i*, we write* $X_i \models \mathbf{a}_i$*. A* joint conjecture set *for* $\{X_1, \ldots, X_n\}$ *is a set of conjecture profiles* $(\mathbf{a}_1, \ldots, \mathbf{a}_n) \in \boldsymbol{\mathcal{A}}^n$*, where* $X_j \models \mathbf{a}_j$*,* $j = 1, \ldots, n$*. A* conditioning conjecture set *for* $\mathrm{pa}\,(X_i) = \{X_{i_1}, \ldots, X_{i_{q_i}}\}$ *is a set of conjecture profiles* $\boldsymbol{\alpha}_{\mathrm{pa}(i)} = (\mathbf{a}_{i_1}, \ldots, \mathbf{a}_{i_{q_i}}) \in \boldsymbol{\mathcal{A}}^{q_i}$*, where* $X_{i_k} \models \mathbf{a}_{i_k}$*, and we write* $\mathrm{pa}\,(X_i) \models \boldsymbol{\alpha}_{\mathrm{pa}(i)}$*.*

We introduce the notation

$$\mathcal{H}_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_{\mathrm{pa}(i)})\colon\ \mathrm{pa}\,(X_i) \models \boldsymbol{\alpha}_{\mathrm{pa}(i)} \Rightarrow X_i \models \mathbf{a}_i\,, \tag{2}$$

to express the hypothetical proposition that, if $\boldsymbol{\alpha}_{\mathrm{pa}(i)}$ is a conditioning conjecture set for $\mathrm{pa}\,(X_i)$, then X_i will conjecture $\mathbf{a}_i$. The conditioning symbol "|" separates the conditioned entity (the consequent) on the left from the conditioning entity (the antecedent) on the right.

Definition 5. *Given a parent set* $\mathrm{pa}\,(X_i) = \{X_{i_1}, \ldots, X_{i_{q_i}}\}$ *for* X_i *and a conditioning conjecture set* $\boldsymbol{\alpha}_{\mathrm{pa}(i)} = (\mathbf{a}_{i_1}, \ldots, \mathbf{a}_{i_{q_i}})$ *for* $\mathrm{pa}\,(X_i)$*, the* conditional utility *of the hypothetical proposition* $\mathcal{H}_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_{\mathrm{pa}(i)})$ *is* $u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_{\mathrm{pa}(i)})\colon \boldsymbol{\mathcal{A}} \rightarrow \mathbb{R}$*. If* $\mathrm{pa}\,(i) = \varnothing$*, then* $u_{i|\mathrm{pa}(i)} = u_i$*, a categorical utility for* X_i*.*

A conditional utility involves a special logical structure as the consequent of a hypothetical proposition whose antecedent is an assertion attributed to the conditioning entity. Thus, there is a significant operational difference between a categorical utility and a conditional utility.[1] The former provides sufficient information for the individual to take action, whereas the latter is used to define situational, or context dependent, relationships, and taking action requires the appropriate context to be actualized. Thus, whereas the statement $u_i(\mathbf{a}_i) > u_i(\mathbf{a}'_i)$ means that X_i prefers $\mathbf{a}_i$ to $\mathbf{a}'_i$ under all circumstances, the statement "$u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_{\mathrm{pa}(i)}) > u_{i|\mathrm{pa}(i)}(\mathbf{a}'_i|\boldsymbol{\alpha}_{\mathrm{pa}(i)})$ means that X_i

[1] A concept of a conditional utility that is syntactically similar to this approach is the notion of attribute dominance introduced by Abbas and Howard (2005) and Abbas (2009), who create a conditional utility from a joint and marginal utility, whereas we create a joint utility from conditional and marginal utilities. Although they share the same syntax, the two usages are inverses of each other.

prefers $\mathbf{a}_i$ to $\mathbf{a}_i'$, given that $\boldsymbol{\alpha}_{\mathrm{pa}(i)}$ is a conjecture for $\mathrm{pa}(X_i)$.[2] We emphasize that a conjecture is not a strategy. A strategy is an action that conforms to a rule that defines what action *should* be taken, but a conjecture is an action that *might* be intended, given appropriate circumstances.

Conditional utilities provide the mechanism by which individuals may incorporate the behavioral influence of others into their own preferences. Given two agents X_1 and X_2 forming preferences over a set $\boldsymbol{\mathcal{A}}$ such that $\mathrm{pa}(X_1) = \varnothing$ and $\mathrm{pa}(X_2) = \{X_1\}$, the conditional $u_{2|1}(\mathbf{a}_2|\mathbf{a}_1)$ characterizes the degree to which the hypothesis $X_1 \models \mathbf{a}_1$ influences X_2's preference for $X_2 \models \mathbf{a}_2$. Notice that this is exactly the same syntax employed by probability theory to incorporate the statistical influence that other random phenomena have on a random phenomenon. Given two random variables Y_1 and Y_2 taking values in $\mathcal{Y}$, the conditional probability mass function $p_{2|1}(y_2|y_1)$ characterizes the degree to which the hypothesis $Y_1 = y_1$ influences the belief that $Y_2 = y_2$.

Furthermore, since positive affine transformations of utility preserve preference orderings uniquely, we may assume without loss of generality that the utilities are *normalized*, meaning that they are nonnegative and sum to unity; that is,

$$\begin{aligned} u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_{\mathrm{pa}(i)}) &\geq 0 \ \forall \ \boldsymbol{\alpha}_{\mathrm{pa}(i)} \in \boldsymbol{\mathcal{A}}^{q_i} \\ \sum_{\mathbf{a}_i} u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_{\mathrm{pa}(i)}) &= 1 \ \forall \ \boldsymbol{\alpha}_{\mathrm{pa}(i)} \in \boldsymbol{\mathcal{A}}^{q_i} \,. \end{aligned} \tag{3}$$

We refer to such 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.[3]

Lemma 1. *An order isomorphism exists between ordering the strength of belief regarding propositions and ordering the strength of preference regarding alternatives.*[4] *This isomorphism applies to both categorical and conditional orderings.*

Proof. Without loss of generality, we restrict attention to a two-agent group $\{X_1, X_2\}$ 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 $\mathcal{Y}_1$ and $\mathcal{Y}_2$ be arbitrary sets of random propositions of distinct elements with cardinalities equal to the cardinalities of $\mathcal{A}_1$ and

[2] The notion of conditional preferences used here employs syntax similar to the well-known concept of state-dependent preferences (cf. Karni and Schmeidler (1981); Karni et al. (1983); Karni (1985); Drèze (1987)), where the decision maker's preferences are modulated by the state of nature. The two notions, however, have different semantics. State dependence yields a preference ordering corresponding to a particular state of nature, where a state, as defined by Arrow (1971, p. 45), is "a description of the world so complete that, if true and known, the consequences of every action would be known." The state of nature is an assumption imposed on the decision maker that constrains the preferences to a particular environment. The concept of conditional preferences used herein, however, deals with an individual's ability to modulate its preferences to account for the varying preferences of other individuals who are participants in the multiagent decision dynamics. Rather than constraining the model to a particular environment, the intent of such preferences is to extend the model to a more complex social environment. Conventional state-dependent preferences are used to account for the presence of uncertainty, while the conditional preferences are used herein to extend beyond narrow self-interest and accommodate a complex social structure.

[3] The isomorphic relationship simply means that beliefs and preferences can be expressed with the same mathematical syntax. It does not imply any causal relationships between them, and doing so would be irrational: One who interprets preferences as beliefs is a wishful thinker (if it is best, it must be true), and one who interprets beliefs as preferences is a fatalist (if it is true, it must be best).

[4] 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).

$\mathcal{A}_2$, respectively, and let $\boldsymbol{\mathcal{Y}} = \mathcal{Y}_1 \times \mathcal{Y}_2$. Let $g\colon \boldsymbol{\mathcal{A}} \to \boldsymbol{\mathcal{Y}}$ be a bijective mapping $g\colon \mathbf{a} \mapsto \mathbf{y}$, and define functions $b_1\colon \boldsymbol{\mathcal{A}} \to \mathbb{R}$ and $\{b_{2|1}(\cdot|\mathbf{y}_1)\colon \boldsymbol{\mathcal{Y}} \to \mathbb{R}\ \forall \mathbf{y}_1 \in \boldsymbol{\mathcal{Y}}\}$ such that

$$b_1(\mathbf{y}) = u_1[g^{-1}(\mathbf{y})] = u_1(\mathbf{a}) \tag{4}$$

and

$$b_{2|1}(\mathbf{y}_2|\mathbf{y}_1) = u_{2|1}[g^{-1}(\mathbf{y}_2)|g^{-1}(\mathbf{y}_1)] = u_{2|1}(\mathbf{a}_2|\mathbf{a}_1)\,. \tag{5}$$

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

Definition 6. *A* conditional network game *comprises a set of individuals $\{X_1, \ldots, X_n\}$, a set of conjecture profiles $\boldsymbol{\mathcal{A}} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$, and a family of sets of conditional utility mass functions of the form*

$$\{u_{i|\mathrm{pa}(i)}(\cdot|\boldsymbol{\alpha}_{\mathrm{pa}(i)})\,\forall \boldsymbol{\alpha}_{\mathrm{pa}(i)} \in \boldsymbol{\mathcal{A}}^{q_i}, i = 1, \ldots, n\}\,. \tag{6}$$

2.3 Conventional Solutions to Network Games

An obvious extension to classical noncooperative game theory is to define Nash equilibria for network games involving conditional utilities.

Definition 7. *Let $\mathbf{a} = (a_{11}, \ldots, a_{nn})$ be a fixed conjecture profile. For each X_i, let $\boldsymbol{\alpha}_{\mathrm{pa}(i)} = (\mathbf{a}, \ldots, \mathbf{a}) \in \boldsymbol{\mathcal{A}}^{p_i}$ denote the conditioning conjecture set for $\mathrm{pa}\,(X_i) = \{X_{i_1}, \ldots, X_{i_{p_i}}\}$ (notice that all parents conjecture the same profile). Let $\mathbf{a}_i = (a_{11}, \ldots, a_{ii}', \ldots, a_{nn})$ where $a_{ii}' \neq a_{ii}$; that is, $\mathbf{a}_i$ differs from $\mathbf{a}$ in the ith position. The profile $\mathbf{a}$ is a* conditioned Nash equilibrium *if*

$$u_{i|\mathrm{pa}(i)}(\mathbf{a}|\boldsymbol{\alpha}_{\mathrm{pa}(i)}) \geq u_{i|\mathrm{pa}(i)}(\mathbf{a}_i'|\boldsymbol{\alpha}_{\mathrm{pa}(i)}) \tag{7}$$

for all $a_i' \neq a_i$ and for all $i = 1, \ldots, n$. If all utilities are categorical, the conditioned Nash equilibrium becomes a classical Nash equilibrium.

Example 1. *The Battle of the Sexes game is often viewed as a coordination scenario. A man (M – the row player) and a woman (W – the column player) can each attend either the dog race (D) or the ballet (B). Thus, $\mathcal{A}_M = \mathcal{A}_W = \{D, B\}$, yielding $\boldsymbol{\mathcal{A}} = \mathcal{A}_M \times \mathcal{A}_W$. First, assume that M and W possess categorical preferences defined by the payoff matrix displayed in Table 1. There are two Nash equilibria, (D, D) and (B, B) but, unfortunately, game theory does not provide a definitive solution, and instead invites a mixed strategy as a function of probabilities.*

Problems such as this have received considerable attention in the literature. One well-known approach introduced by Schelling (1960) is the concept of a focal point *that would cause the players to concentrate their attention on one of the equilibria. Suppose the players reside in a culture where the man defers to the woman when deciding which social event to attend. The focal point would then be (B, B), since that outcome is most preferred by the woman.*

Relying on focal points forces the players to invoke social criteria that are not encoded into the utilities. But if social criteria are relevant, then perhaps they should be explicitly incorporated into the utilities. We may do so by establishing the influence link

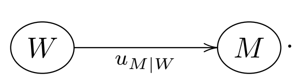

(8)

Table 1: The payoff matrix in ordinal form for the Battle of the Sexes game.

	W	
M	D	B
D	4, 3	2, 2
B	1, 1	3, 4

Key: 4 = best; 3 = next-best; 2 = next-worst; 1 = worst

We may convert W's ordinal preferences into a parameterized utility mass function as

$$u_W(D,D) = \alpha, \quad u_W(D,B) = 0, \quad u_W(B,D) = 0, \quad u_W(B,B) = 1 - \alpha\,, \tag{9}$$

where $0 \leq \alpha < 1/2$.[5] *We form the conditional propositions for M as*

$$\mathcal{H}_{M|W}(\mathbf{a}_M|\mathbf{a}_W)\text{: } X_W \vDash \mathbf{a}_W \Rightarrow X_M \vDash \mathbf{a}_M \;\forall\, (\mathbf{a}_M, \mathbf{a}_W) \in \boldsymbol{\mathcal{A}} \times \boldsymbol{\mathcal{A}}. \tag{10}$$

The corresponding conditional utility mass function $u_{M|W}$ *is specified as follows. If the conjecture profile for W were* (D,D), (D,B), *or* (B,D), *then M would concentrate his entire utility on* (D,D), *since that is his most preferred outcome. But if W's conjecture were* (B,B), *then M would apportion* $\beta \in [0, 1/2)$ *to* (D,D) *and* $1-\beta$ *to* (B,B). *The resulting conditional utilities are displayed in Table 2.*[6]

Table 2: M's conditional utility $u_{M|W}$ for the conditioned Battle of the Sexes game.

$u_{M\|W}(a_{MM}, a_{MW}\|a_{WM}, a_{WW})$				
	a_{WM}, a_{WW}			
a_{MM}, a_{MW}	D,D	D,B	B,D	B,B
D,D	1	1	1	β
D,B	0	0	0	0
B,D	0	0	0	0
B,B	0	0	0	$1-\beta$

The conditioned payoff matrices are defined by juxtaposing W's utility given by (9) and each column of Table 2, and are displayed in 3. Conditioning with respect to (D,D) , (D,B) *and* (B,D) *yields the payoffs displayed in Table 3(a), where we see that* (D,D) *and* (B,B) *are both Nash equilibria but neither is dominant. Thus, as with the conventional case, conditioning on these outcomes does not resolve the issue. Table 3(b) displays the payoff matrix when conditioning on* (B,B), *which also reveals* (D,D) *and* (B,B) *as equilibria, but* (B,B) *is dominant, since* $\alpha, \beta < 1/2$. *Thus, directly incorporating the social context into the utility structure apparently generates the appropriate outcome. But a closer analysis reveals that applying conditional Nash equilibrium logic merely begs the question — focal point analysis is ultimately used to prefer the equilibrium conditioned on* (B,B) *over the other conditioned equilibria. Thus, at the end of the day, we are no closer to a resolution that we were with the original formulation.*

[5] When the individuals have specific names, it is convenient to replace the numerical indices with literations to represent the individuals. Thus we write X_M for X_1, X_W for X_2, a_{MM} for a_{11}, a_{MW} for a_{12}, etc.

[6] In Section 3.2 we show how this formulation can easily be simplified.

Table 3: The Conditioned payoff matrices for the Battle of the Sexes game: (a) corresponds to conditioning on (D, D) (D, B), and (B, D), and (b) corresponds to conditioning on (B, B).

	W	
M	D	B
D	$1, \alpha$	$0, 0$
B	$0, 0$	$0, 1-\alpha$

(a)

	W	
M	D	B
D	β, α	$0, 0$
B	$0, 0$	$1-\beta,\ 1-\alpha$

(b)

3 Coherent Coordination

> A mathematical formalism may be operated in ever new, uncovenanted ways, and force on our hesitant minds the expression of a novel conception.
>
> Michael Polanyi, *Personal Knowledge*

A conditioned Nash equilibrium takes into consideration the social influence that exists in the network and, therefore, constitutes a useful extension of classical noncooperative game theory. If that were the end of the story, then the introduction of conditional utilities would be a nice, but somewhat modest, extension to game theory. But there is more to be said. The introduction of explicitly modeled social influence into a network offers the possibility of defining entirely new solution concepts that are not possible with the classical game-theoretic model. To identify such concepts, we continue to exploit the mathematical machinery of probability theory.

The application of the conditionalization syntax as a means of modeling social influence, together with the order isomorphism between beliefs and preferences, suggests that there are natural connections between a social network in the behavioral domain and a Bayesian network in the epistemological domain.

Definition 8. *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 $\mathcal{Y}$.*
- *The incoming edges to Y_i constitute the conditional probability, denoted $p_{i|\mathrm{pa}(i)}(y_i|\Upsilon_{\mathrm{pa}\,(i)})$, that $Y_i = y_i \in \mathcal{Y}$, given that its parents $\mathrm{pa}\,(Y_i) = \{Y_{i_1}, \ldots, Y_{i_{q_i}}\}$ assume the values $\Upsilon_{\mathrm{pa}\,(i)} = (y_{i_1}, \ldots, y_{i_{q_i}}) \in \mathcal{Y}^{q_i}$. If $\mathrm{pa}\,(Y_i) = \emptyset$, then Y_i is a root vertex and $p_{i|\mathrm{pa}(i)} = p_i$, the unconditional marginal probability mass function for Y_i.*

A Bayesian network provides a powerful framework within which to analyze the behavior of either artificial or naturally existing groups of interacting elements. It is natural to start by considering the way the behavior of a given element 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 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).

Bayesian network theory establishes that a unique joint probability mass function, denoted $p_{1:n}$, can be constructed as the product of the conditional probability mass functions of all non-root vertices and the marginal probability mass functions of the root vertices, yielding

$$p_{1:n}(y_1, \ldots, y_n) = \prod_{i=1}^{n} p_{i|\text{pa}(i)}(y_i|\Upsilon_{\text{pa}\,(i)})\,. \tag{11}$$

In particular, consider the two-agent Bayesian network given by

$$\boxed{Y_1} \xrightarrow[p_{2|1}]{} \boxed{Y_2}\,. \tag{12}$$

The joint probability mass function is then computed as

$$p_{12}(y_1, y_2) = p_1(y_1)p_{2|1}(y_2|y_1)\,. \tag{13}$$

The synthesis of a joint probability mass function is the mechanism by which a notion of strict individual belief, as expressed by the marginal probability p_1, and an expanded notion of conditional individual belief, whereby one takes into account the influence that the beliefs of others have on one's own beliefs as expressed by the conditional probability $p_{2|1}$, are combined to form a notion of belief connectivity as expressed by the joint probability p_{12}. This connection does *not* generate a group belief; belief is an innately individual concept. If a group were to believe something, then it must function as a single organism.

It is immediate that the social network given by (1) is isomorphic to (12) and, accordingly, we may combine the categorical and conditional utilities to form a new function defined over the $\boldsymbol{\mathcal{A}} \times \boldsymbol{\mathcal{A}}$ of the form

$$u_{12}(\mathbf{a}_1, \mathbf{a}_2) = u_1(\mathbf{a}_1)u_{2|1}(\mathbf{a}_2|\mathbf{a}_1)\,. \tag{14}$$

The synthesis of such a function is the mechanism by which a notion of strict individual preference, as expressed by the categorical utility u_1, and an expanded notion of conditional individual preference, $u_{2|1}$, whereby one takes into account the influence that the preferences of others have on one's own preferences, are combined to form a notion of preference connectivity as expressed by the joint utility u_{12}. Preference connectivity, however, is *not* group preference in the sense of the group possessing wants or desires as a single entity. Preference, as expressed via the individual utilities, is a measure of individual material benefit. Furthermore, no concept of group-level material benefit is provided in the problem formulation. Thus, u_{12} *cannot* be a measure of material benefit for the group. To be meaningful, it must be a measure of a phenomenon that is intrinsically social. To motivate the identification of such a phenomenon, let us revisit the Battle of the Sexes game.

Example 2. *Returning to the Battle of the Sexes game, let W's categorical utility be given by (9) and M's conditional utility be given by Table 2. Applying (14) yields*

$$u_{MW}[(a_{MM}, a_{MW}), (a_{WM}, a_{WW})] = u_{M|W}(a_{MM}, a_{MW}|a_{WM}, a_{WW})u_W(a_{WM}, a_{WW})\,, \tag{15}$$

the results of which are displayed in Table 4. There are only three nonzero joint conjecture sets: $[(D,D),(D,D)]$, $[(D,D),(B,B)]$, *and* $[(B,B),(B,B)]$. *By straightforward calculations, these joint conjecture sets are ordered as follows:*

$$\begin{aligned} &u_{MW}[(D,D),(D,D)] > u_{MW}[(B,B),(B,B)] > u_{MW}[(D,D),(B,B)] \text{ for } \frac{1-2\alpha}{1-\alpha} < \beta < 1/2 \\ &u_{MW}[(B,B),(B,B)] > u_{MW}[(D,D),(B,B)] > u_{MW}[(D,D),(D,D)] \text{ for } \frac{\alpha}{1-\alpha} < \beta < 1/2 \\ &u_{MW}[(B,B),(B,B)] > u_{MW}[(D,D),(D,D)] > u_{MW}[(D,D),(B,B)] \text{ for } \beta < \min\left\{\frac{\alpha}{1-\alpha}, \frac{1-2\alpha}{1-\alpha}\right\}. \end{aligned} \tag{16}$$

Table 4: The coordination function $u_{MW}[(a_{MM}, a_{MW}), (a_{WM}, a_{WW})]$ for the conditional Battle of the Sexes game.

	a_{WM}, a_{WW}			
a_{MM}, a_{MW}	D, D	D, B	B, D	B, B
D, D	α	0	0	$\beta - \alpha\beta$
D, B	0	0	0	0
B, D	0	0	0	0
B, B	0	0	0	$1 - \alpha - \beta + \alpha\beta$

To analyze this situation, let us focus on the ordering

$$u_{MW}[(B, B), (B, B)] > u_{MW}[(D, D), (B, B)] > u_{MW}[(D, D), (D, D)] , \tag{17}$$

which obtains when $\frac{\alpha}{1-\alpha} < \beta < 1/2$. *This string of inequalities generates a profound question. On the one hand, the relation* $u_{MW}[(B, B), (B, B)] > u_{MW}[(D, D), (B, B)]$ *makes intuitive sense: It seems reasonable for the group to be more coordinated if both players conjecture the same outcome than for them to conjecture different outcomes. On the other hand, the relation* $u_{MW}[(D, D), (B, B)] > u_{MW}[(D, D), (D, D)]$ *indicates that the group is more coordinated if the players conjecture different outcomes rather than conjecturing the same outcome — a not so intuitive relation. To understand these results, we note that the* (α, β) *region where this seeming conundrum occurs is where* $\beta > \alpha$, *that is,* M*'s conditional utility for* (D, D), *given that* W *conjectures* (B, B), *is greater than* W*'s utility of* (D, D). *In other words,* M*'s stubborn insistance on* (D, D) *conflicts with* W*'s relatively weak preference for* (D, D). *In the face of this conflict, it is more coordinated for them to express their differences than for* W *to cave in to pressure. Clearly, the most coordinated relationship is for both to conjecture* (B, B).

Example 2 illustrates the kinds of complexity that can arise for even a rather simple social relationship as a result of social influence. The function $u_{MW}[(a_{MM}, a_{MW}), (a_{WM}, a_{WW})]$ provides an ordering of all joint conjecture profiles with respect to their compatibility, thereby providing an assessment of the seriousness of disputes and the possibilities for compromise. A natural way to interpret this function is as a measure of the degree to which the joint conjecture sets correspond to some emergent notion of systematic behavior — in other words, to coordinate.

Definition 9. *The* coordination function *of a conditional network game comprising* $\{X_1, \ldots, X_n\}$, $\boldsymbol{\mathcal{A}}$, *and* $\{u_{i|\text{pa}(i)}(\cdot|\boldsymbol{\alpha}_{\text{pa}(i)}) \,\forall \boldsymbol{\alpha}_{\text{pa}(i)} \in \boldsymbol{\mathcal{A}}^{q_i}, i = 1, \ldots, n\}$, *is given by*

$$u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_n) = \prod_{i=1}^{n} u_{i|\text{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_{\text{pa}(i)}) , \tag{18}$$

where, if $\text{pa}(X_i) = \varnothing$, *then* $u_{i|\text{pa}(i)} = u_i$, *a categorical utility.*[7]

[7]To be precise, the coordination function is analogous to the joint probability mass function of a family of random vectors $\{\mathbf{y}_i, 1, \ldots, n\}$ with $\mathbf{y}_i = \{y_{i1}, \ldots, y_{in}\}$, where the jth y_{ij} a scalar random variable. This additional complexity does not affect the validity of the isomorphism, since a Bayesian network can easily be extended such that each vertex is a random vector comprising n random variables with marginal and conditional probability mass functions defined over n variables.

The coordination function provides an ordering over all possible combinations of joint conjecture sets $(\mathbf{a}_1, \ldots, \mathbf{a}_n) \in \boldsymbol{\mathcal{A}} \times \cdots \times \boldsymbol{\mathcal{A}}$. In this sense it captures all of the social relationships that develop as the players interact. Since it is a function of n^2 independent variables, it does not directly establish a decision criterion. Rather, it serves as a foundation for the identification of coordinated solution concepts.

3.1 Solution Concepts

3.1.1 Nash Equilibria

Building on classical game-theoretic solution concepts, the most obvious approach is exploit the syntax of probability theory. In the epistemological context, once given the joint probability mass function $p_{12}(y_1, y_2)$, one may compute the marginal probability mass functions

$$p_1(y_1) = \sum_{y_2} p_{12}(y_1, y_2) \text{ and } p_2(y_2) = \sum_{y_1} p_{12}(y_1, y_2)\,. \tag{19}$$

The marginal probability mass function expresses the degree of belief that $Y_i = y_i$ after taking into account the statistical relationships that exist between Y_1 and Y_2.

By the isomorphism, we may compute the *ex post* marginal utilities for each X_i, yielding

$$v_i(\mathbf{a}_i) = \sum_{\sim \mathbf{a}_i} u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_n)\; i = 1, \ldots, n\,, \tag{20}$$

where the notation $\sum_{\sim \mathbf{a}_i}$ means the sum is taken over all arguments of $u_{1:n}$ except $\mathbf{a}_i$. We may use these marginals to form an *ex post* payoff array. For a 2×2 (two-peson, two-move) game, we endow X_1, the row player, with $\mathcal{A}_1 = \{z_{11}, z_{12}\}$, and X_2, the column player, with $\mathcal{A}_2 = \{z_{21}, z_{22}\}$. The *ex post* payoff matrix is of the form displayed in Table 5, to which we may apply classical solution concepts.

Definition 10. *Let* $\mathbf{a}^* = (a_{11}^*, \ldots, a_{ii}^*, \ldots, a_{nn}^*)$ *and let* $\mathbf{a}_i = (a_{11}^*, \ldots, a_{ii}', \ldots, a_{nn}^*)$*. Then* $\mathbf{a}^*$ *is an* ex post Nash equilibrium *if* $v_i(\mathbf{a}^*) \geq v_i(\mathbf{a}_i)$ *for all* $a_{ii}' \neq a_{ii}^*$ *and for all* $i = 1, \ldots, n$*, where* v_i *is defined by (20).*

Table 5: The *ex post* payoff matrix for a 2×2 game.

	X_2	
X_1	z_{21}	z_{22}
z_{11}	$v_1(z_{11}, z_{21}),\ v_2(z_{11}, z_{21})$	$v_1(z_{11}, z_{22}),\ v_2(z_{11}, z_{22})$
z_{12}	$v_1(z_{12}, z_{21}),\ v_2(z_{12}, z_{21})$	$v_1(z_{12}, z_{22}),\ v_2(z_{12}, z_{22})$

Example 3. *To compute the* ex post *Nash equilibrium for the Battle of the Sexes game, we compute the marginal utility for* M *using Table 4, yielding*

$$\begin{aligned} v_M(D,D) = \alpha + \beta - \alpha\beta \quad & v_M(D,B) = 0 \\ v_M(B,D) = 0 \quad & v_M(B,B) = 1 - \alpha - \beta + \alpha\beta\,. \end{aligned} \tag{21}$$

Since W*'s* ex ante *utility is categorical, her* ex post *marginal is the same as her* ex ante *categorical utility, thus*

$$v_W(D,D) = \alpha, \quad v_W(D,B) = 0, \quad v_W(B,D) = 0, \quad v_W(B,B) = 1 - \alpha\,. \tag{22}$$

The payoff matrix for the ex post *Battle of the Sexes game is displayed in Table 6. Although there are the same two Nash equilibria as before,* (B, B) *is the dominant equilibrium for all* $(\alpha, \beta) \in (0, 1/2) \times (0, 1/2)$*. Notice that the* ex post *payoff matrix for the Battle of the Sexes actually turns into a Hi-Lo game — one of Bacharach's puzzles. Thus, a definitive solution remains unattained. (The Hi-Lo game will be analyzed in detail in Section 1.1 of Part II.)*

Table 6: The payoff matrix for the *ex post* Battle of the Sexes game.

	W	
M	D	B
D	$\alpha + \beta - \alpha\beta,\ \alpha$	0, 0
B	0, 0	$1 - \alpha - \beta + \alpha\beta,\ 1 - \alpha$

3.1.2 Coordination, Individual Performance, and Team Reasoning

Although the *ex post* payoffs incorporate social relationships into the individual unconditional preferences, that does not change the fact that Nash equilibria solutions are based on individual rationality and therefore do not establish a notion of socially rational behavior. This result represents a true generalization of classical game theory, but if we were to stop the analysis here, then the coordination analysis offered by conditional game theory would be little more than an exercise to refine the specification of categorical utilities in preparation for the ultimate application of conventional theory.

But there is still more to be said. Recall that each conjecture profile is of the form $\mathbf{a}_i = (a_{i1}, \ldots, a_{in})$, where $a_{ii} \in \mathcal{A}_i$ is a conjecture by X_i and $a_{ij} \in \mathcal{A}_j$ is a conjecture for X_j by X_i. Each X_i, however, has control only over a_{ii}, its own component of its conjecture profile. In terms of coordinated behavior, what is most relevant is how the individual conjectures a_{ii}, $i = 1, \ldots, n$, fit together to generate a coordinated outcome. To make this assessment, we compute the marginal coordination function with respect to the conjectures a_{ii} for each X_i.

Definition 11. *Given a joint conjecture set* $(\mathbf{a}_1, \ldots, \mathbf{a}_n)$*, form the profile* $(a_{11}, \ldots, a_{nn})$ *by taking the* i*th element of each* X_i*'s conjecture profile,* $i = 1, \ldots, n$*, and summing the coordination function over all elements of each* $\mathbf{a}_i$ *except the* ii*th elements to form the* coordination utility $w_{1:n}$ *for* $\{X_1, \ldots, X_n\}$*, yielding*

$$w_{1:n}(a_{11}, \ldots, a_{nn}) = \sum_{\sim a_{11}} \cdots \sum_{\sim a_{nn}} u_{1:n}[(a_{11}, \ldots, a_{1n}), \ldots, (a_{n1} \ldots, a_{nn})] \,. \tag{23}$$

The function $w_{1:n}$: $\boldsymbol{\mathcal{A}} \to [0, 1]$ is a group-level function but, as mentioned earlier, it is not a measure of group-level material benefit. We reserve the term "utility" to refer to individual expressions of material benefit (either conditional or unconditional). Instead, $w_{1:n}$ is a measure of the degree to which the actions (the parts) of profile $(a_{11}, \ldots, a_{nn}) \in \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ fit together to form systematic social behavior (the whole). If the whole is cooperative, then coordination will be high when players perform as a team, and if the whole is conflictive (e.g, a military engagement), then coordination will be high when players are in opposition. The coordination utility is the mechanism by which individual decisions can be defined.

Definition 12. *The* coordinated individual rationality decision function *for X_i is the ith marginal of $w_{1:n}$, that is,*

$$w_i(a_{ii}) = \sum_{\sim a_{ii}} w_{1:n}(a_{11}, \ldots, a_{1n}) \,. \tag{24}$$

The coordinated individual decision *for X_i is defined as*

$$a_{ii}^* = \arg\max_{a_{ii} \in \mathcal{A}_i} w_i(a_{ii}) \,. \tag{25}$$

Also, from the perspective of team reasoning, an alternative solution concept is for each individual to select its component of the profile that maximizes the coordination utility.

Definition 13. *Given a coordination utility $w_{1:n}$, the* generalized team-reasoning*, or GTR, solution is for each X_i to choose $a_{ii}^\dagger$, its component of $\mathbf{a}^\dagger = (a_{11}^\dagger, \ldots, a_{nn}^\dagger)$, where*

$$\mathbf{a}^\dagger = \arg\max_{\mathbf{a} \in \boldsymbol{\mathcal{A}}} w_{1:n}(\mathbf{a}) \,. \tag{26}$$

If $a_{ii}^\dagger = a_{ii}^*$ for $i = 1, \ldots, n$, then $\mathbf{a}^*$ is a consensus *solution.*

Unlike the Nash equilibrium solution concept that is imposed exogenously, the coordinated individually rational and *GTR* solution concepts emerge endogenously as a consequence of the social structure. Thus, this approach replaces the division of labor concept of separately defining the preferences and the solution concept. Each agent's decision automatically incorporates the social influence that others exert on it, thereby integrating preference specification and the solution concept. The coordination utility and the individually rational decision functions provide the network with the ability to simultaneously evaluate the outcomes in terms of coordination (a social concept) and performance (an individual concept).

Example 4. *Returning again to the Battle of the Sexes conditional game and applying (23) and Table 4, the coordination utility is*

$$\begin{aligned}
w_{MW}(D,D) &= u_{MW}[(D,D),(D,D)] + u_{MW}[(D,D),(B,D)] + u_{MW}[(D,B),(D,D)] + u_{MW}[(D,B),(B,D)]\\
w_{MW}(D,B) &= u_{MW}[(D,D),(D,B)] + u_{MW}[(D,D),(B,B)] + u_{MW}[(D,B),(D,B)] + u_{MW}[(D,B),(B,B)]\\
w_{MW}(B,D) &= u_{MW}[(B,D),(D,D)] + u_{MW}[(B,D),(B,D)] + u_{MW}[(B,B),(D,D)] + u_{MW}[(B,B),(B,D)]\\
w_{MW}(B,B) &= u_{MW}[(B,D),(D,B)] + u_{MW}[(B,D),(B,B)] + u_{MW}[(B,B),(D,B)] + u_{MW}[(B,B),(B,B)\,,
\end{aligned} \tag{27}$$

yielding

$$\begin{aligned}
w_{MW}(D,D) &= \alpha\\
w_{MW}(D,B) &= \beta - \alpha\beta\\
w_{MW}(B,D) &= 0\\
w_{MW}(B,B) &= 1 - \alpha - \beta + \alpha\beta
\end{aligned} \,. \tag{28}$$

The coordinated utilities are

$$\begin{aligned}
w_M(D) &= w_{MW}(D,D) + w_{MW}(D,B) = \alpha + \beta - \alpha\beta\\
w_M(B) &= w_{MW}(B,D) + w_{MW}(B,B) = 1 - \alpha - \beta + \alpha\beta
\end{aligned} \tag{29}$$

and

$$\begin{aligned}
w_W(D) &= w_{MW}(D,D) + w_{MW}(B,D) = \alpha\\
w_W(B) &= w_{MW}(D,B) + w_{MW}(B,B) = 1 - \alpha
\end{aligned} \,. \tag{30}$$

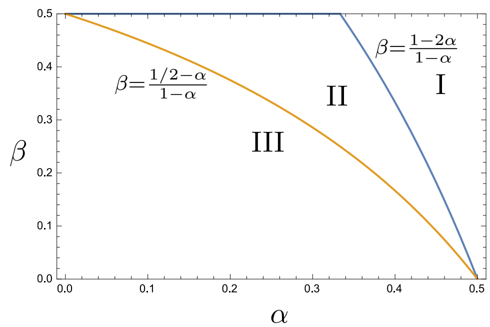

Figure 1: Plot of coordination regions for the Battle of the Sexes conditional game

The coordination utility is maximized at (B, B) *for all* (α, β) *such that* $\beta < \min\{1/2, \frac{1-2\alpha}{1-\alpha}\}$ *(Regions I and II), and is maximized at* (D, D) *for* $\frac{1-2\alpha}{1-\alpha} < \beta < 1/2$ *(Region III). The coordinated utility for* W *satisfies* $w_W(B) > w_W(D)$ *for* $0 < \alpha < 1/2$*, but the coordinated utility for* M *satisfies*

$$\begin{aligned} w_M(B) &> w_M(D) \text{ for } 0 < \beta < \frac{1/2 - \alpha}{1 - \alpha} \\ w_M(D) &> w_M(B) \text{ for } \frac{1/2 - \alpha}{1 - \alpha} < \beta < 1/2 \,. \end{aligned} \tag{31}$$

Figure 1 displays the (α, β) *regions for the coordination utility and for the coordinated decisions, and Table 7 interprets the three* (α, β) *regions.*

Table 7: The coordination utility and the coordinated individually rational decision functions for the conditional Battle of the Sexes game.

(α, β) Region	Coordination utility ordering	Individual M ordering	Individual W ordering	*GTR* decision	Consensus choice?
I	$D, D \succ_{MW} B, B$	$D \succ_M B$	$B \succ_W D$	D	no
II	$B, B \succ_{MW} D, D$	$D \succ_M B$	$B \succ_W D$	B	no
III	$B, B \succ_{MW} D, D$	$B \succ_M D$	$B \succ_W D$	B	yes

3.2 Sociation

The assumption underlying game theory is that each individual's payoff is a function the joint actions of all agents, hence the requirement that categorical utilities are of the form $u_i(\mathbf{a}_i)$; that is, the payoff is assumed to be a function of all elements of $\mathbf{a}_i \in \boldsymbol{\mathcal{A}}$. The development presented in Section applies to conditional utilities of the form $u_{i|\text{pa}(i)}\mathbf{a}_i|\boldsymbol{\alpha}_{\text{pa}(i)})$, where $\mathbf{a}_i \in \boldsymbol{\mathcal{A}}$ and $\boldsymbol{\alpha}_{\text{pa}(i)} \in \boldsymbol{\mathcal{A}}^{q_i}$. Although this structure is more complex than the utility structure of classical game theory, it also offers important opportunities for simplifications that the classical theory does not offer. Many interesting multiagent decision problems are such that influence is with respect only to the individual conjectures of the parents, rather than the entire conjecture profile. Suppose that X_j is influenced by X_i's component of $\mathbf{a}_j$; that is, X_j preferences are conditioned on, but only on, X_i's individual action. To the extent that the influence of others depends only on the conjectured actions of others, rather than on the entire conjecture profile of others, the influence is said to be *conjecture dissociated*, and is *completely conjecture dissociated* if the influence depends on, and only on, the conjectured actions of those who influence it.

Furthermore, to the extent that an individual's conditional utility is a function of a subset of its own conjectured profile, then the individual is said to be *utility dissociated.* If its utility is a function of, and only of, its own conjectured actions, then it is *completely utility dissociated.*

If an individual is both completely conjecture dissociated and completely utility dissociated, then it is *completely dissociated.* It should be noted that this simplification is not generally applicable to games with categorical utilities.

Definition 14. *A network is* completely dissociated *if all of its members are completely dissociated, in which case, the coordination function is of the form*

$$u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\mathbf{a}_{i_1},\ldots,\mathbf{a}_{i_{q_i},i_{q_i}}) = u_{i|\mathrm{pa}(i)}(a_{ii}|a_{i_1 i_1},\ldots,a_{i_{q_i} i_{q_i}}) \tag{32}$$

for $=1,\ldots,n$. *In such a case,* X_i*'s utility depends on, and only on, its own action, conditioned on the actions of, and only of, all who influence it. For a completely dissociated network, the coordination utility coincides with the coordination function, yielding*

$$w_{1:n}(a_{11},\ldots,a_{nn}) = u_{1:n}(a_{11},\ldots,a_{nn}) = \prod_{i=1}^{n} u_{i|\mathrm{pa}(i)}(a_{ii}|a_{i_1 i_1},\ldots,a_{i_{q_i} i_{q_i}})) \tag{33}$$

The *ex post* marginal utility mass function for the members of a completely dissociated network coincides with the individual decision function, yielding

$$w_i(a_{ii}) = v_i(a_{ii}) = \sum_{\sim a_i} w_{1:n}(a_{11},\ldots,a_{nn}) = \sum_{\sim a_i} u_{1:n}(a_{11},\ldots,a_{nn})\,. \tag{34}$$

Example 5. *We once again revisit the Battle of the Sexes game, and assume that both* M *and* W *are completely dissociated; that is,* W*'s categorical utility mass function is defined over only over her actions, yielding*

$$u_W(D) = \alpha, \quad \tilde{u}_W(B) = 1-\alpha\,. \tag{35}$$

Also, suppose M *has defined a conditional utility mass function over only his his actions given* W*'s actions, yielding*

$$\begin{array}{ll} u_{M|W}(D|D) = 1 & u_{M|W}(B|D) = 0 \\ u_{M|W}(D|B) = \beta & u_{M|W}(B|B) = 1-\beta \end{array}. \tag{36}$$

Since both M *and* W *are completely dissociated, the coordination function collapses to the coordination utility, hence*

$$w_{MW}(a_M,a_W) = u_{M|W}(a_M|a_W)u_W(a_W)\,, \tag{37}$$

yielding

$$\begin{array}{rl} w_{MW}(D,D) = \alpha & w_{MW}(B,D) = 0 \\ w_{MW}(D,B) = \beta-\alpha\beta, & w_{MW}(B,B) = 1-\alpha-\beta+\alpha\beta\,, \end{array} \tag{38}$$

and the coordinated individually rational decision becomes

$$u_M(D) = \alpha + \beta - \alpha\beta, \quad u_M(B) = 1 - \alpha - \beta + \alpha\beta\,, \tag{39}$$

which is identical with the results obtained with the fully sociated model previously developed.

It is instructive to compare the four versions of the Battle of the Sexes game. With the classical formulation, M and W are modeled as selfish individuals with narrow-defined concepts of self-interest. This model is based on minimal assumptions and it is difficult for one to view this as a realistic social scenario, since no social relationships are included in the model. Furthermore, the game resolves nothing. To settle the issue, the game must be overlaid with a focal point assumption. This situation is a manifestation of the limited ability to express social interests with a mechanism that is designed to express material interest.

If a social assumption is required to settle the issue, then why not build it directly into the model in the first place? Neither conditioned Nash equilibrium nor *ex post* Nash equilibrium resolve the issue. The fully sociated coordinated approach solves the problem, but requires M to consider all of the possible joint conjectures for W when defining his conditional utilities. This entails considerable work, and seems, at least in this case, to be a bit over overkill for such a simple problem. The dissociated conditional game scenario, however, is much simpler than either the traditional game formulation or the fully sociated conditional game formulation, yet is captures all of the relevant social structure. Given that there is an actual social relationship between the two individuals and W has preferences over the venue choices only, then all that is required of M is that his preferences over the venue choices are governed by her preferences. This formulation of the issue is no more complicated than it needs to be. What could be simpler and more natural?

3.3 Coherence

Thus far in the development, we have invoked the logic of conditionalization to motivate the migration of the syntax of probability theory from the epistemological domain to the behavioral domain. We have not, however, justified the wholesale adoption of the probability syntax as a meaningful way to express and combine agent preferences, nor have we discussed any delimitations of such a practice. Thus, the next step is to provide additional rationale for this approach.

Let $\mu_{1:n}$ denote a function defined over $\boldsymbol{\mathcal{A}}$ (not necessarily generated by (23)), and let μ_i denote individual utilities defined over $\mathcal{A}_i$,$i = 1, \ldots, n$ (not necessarily generated by (24)), and suppose there were to exist an agent X_i such that, if

$$\mu_i(a_{ii}^*) > \mu_i(a_{ii}) \tag{40}$$

for all $a_{ii} \neq a_{ii}^*$, then

$$\mu_{1:n}(a_{11}, \ldots, a_{ii}^*, \ldots, a_{nn}) < \mu_{1:n}(a_{11}, \ldots, a_{ii}, \ldots, a_{nn}) \tag{41}$$

for all $a_{ii} \neq a_{ii}^*$ and for all $(a_{11}, \ldots, a_{i-1\,i-1}, a_{i+1\,i+1}, \ldots, a_{nn}) \in \mathcal{A}_1 \times \mathcal{A}_{i-1} \times \mathcal{A}_{i+1} \times \mathcal{A}_n$. Such an agent would be *subjugated*. No matter which of its possible actions it considered best, the coordination of any profile that contained that action would be worse (in terms of coordination) for the group than any other profile. A subjugated agent would be the victim of a particularly harsh and irrational form of discrimination: The mere fact that such an individual would prefer any action, no matter what it would be, essentially renders the group dysfunctional.

Subjugation is a generalization of the social choice notion of *suppression* discussed by Fishburn (1973, p. 211). An individual is suppressed if, whenever it prefers alternative a to a', then society chooses a' over a.[8] It is also important to observe that the opposite notion of subjugation is *subversion*. An individual is a subverter if the inequality in (41) is reversed. Such an individual would possess the power to enforce its will on the group — an absolute dictator.

The ability to subjugate an individual is an extremely powerful collective attribute, but having that power does not mean that a group will use it. Also, the ability to subvert the group is a powerful attribute, but having that power does not mean that an individual will use it. Nevertheless, these attributes are extremely dangerous and have the potential to render a group unstable if not dysfunctional. Denying such power is perhaps the weakest concept of behavior that could qualify as democratic. It would ensure that every member of the group has a "seat at the table" in that its most preferred action, no matter what it might be, is a possible coordination maximizing action.

Perhaps surprisingly, stipulating that it must be impossible to for any individual to subvert the group or for any individual to be subjugated provides a powerful constraint on the operational concept of coordination. It is perhaps equally unsurprising, however, given the isomorphism, that complying with this constraint leads directly to a connection to probability theory. We proceed by establishing that subjugation is isomorphic to the gambling notion of a sure loss — a gamble such that, no matter what the outcome, the gambler's payout is less than the entry fee.

Lemma 2. The Isomorphism Lemma *Subjugation is isomorphic to sure loss.*

Proof. Without loss of generality, we restrict attention to a two-agent completely dissociated group. Let $\mathcal{Y}_1$ and $\mathcal{Y}_2$ be two sets of distinct propositions with cardinalities equal to the cardinalities of $\mathcal{A}_1$ and $\mathcal{A}_2$, respectively. Let u_1: $\mathcal{A}_1 \to \mathbb{R}$ be a categorical utility and let $u_{2|1}(\cdot|a_1)$: $\mathcal{A}_2 \to \mathbb{R}$ for each $a_1 \in \mathcal{A}_1$be conditional utilities defined over $\mathcal{A}_2$. Let g_i: $\mathcal{A}_i \to \mathcal{Y}_i$, $i = 1, 2$, be bijective mappings and define the belief functions b_1: $\mathcal{Y}_1 \to \mathbb{R}$ and b_{12}:$\mathcal{Y}_1 \times \mathcal{Y}_2 \to \mathbb{R}$ such that

$$\begin{aligned} b_1(y_1) &= u_1[g_1^{-1}(y_1)] \\ b_{12}(y_1, y_2) &= u_{12}[g_1^{-1}(y_1), g_2^{-1}(y_2)]\,. \end{aligned} \tag{42}$$

Let $y_1^* \in \mathcal{Y}_1$ be such that

$$b_1(y_1^*) > b_1(y_1)\ \forall y_1 \in \mathcal{Y}_1 \backslash \{y_1^*\}\,, \tag{43}$$

but

$$b_{12}(y_1^*, y_2) < b_{12}(y_1, y_2) \tag{44}$$

for all $y_1 \in \mathcal{Y}_1 \backslash \{y_1^*\}$ and for all $y_2 \in \mathcal{Y}_2$. Thus, even though y_1^* is the most strongly believed proposition in $\mathcal{Y}_1$, the belief regarding the realization of any joint proposition for which y_1^* is X_1's component is weaker than the belief regarding the realization of the corresponding joint proposition with any other y_1 as X_1's component.

If, on the basis of (43) one were to enter a lottery to earn \$1 if y_1^* is realized, a fair entry fee would be $q_1 > 1/2$. On the other hand, if, on the basis of (44), one were to earn \$1 if y_1^* is not realized, then a fair entry fee for that lottery would be $q_2 > 1/2$. By combining these two lotteries into one with an entry fee of $q_1 + q_2 > 1$ with the (false) hope of winning \$2, one would win exactly \$1 regardless of the outcome — a sure loss. It is immediate by the order isomorphism that the relationships given by (40) and (43) and by (41) and (44) are identical. □

It follows immediately by reversing the relevant inequality (i.e., exchanging rolls of the gambler and the bookie) that subversion is isomorphic to a scenario where the gambler's payout is guaranteed to be greater than the entry fee — a sure win. A gambling scenario where the gambler suffers a sure loss is called a Dutch book. The key result in this regard is the famous Dutch Book Theorem.

[8] As Fishburn put it, one who is suppressed is "a dictator turned upside down" (Fishburn, 1973, p. 211).

Theorem 1. The Dutch Book Theorem. *Suppose a gambler places a bet to win a payout of S. A fair entry fee for this gamble is pS, where p is the gambler's degree of belief of winning. (Necessity) If the degree of belief p violates the probability axioms, then it is possible to construct a lottery such that the payout is less (more) than the entry fee — a sure loss (win). (Sufficiency) If p conforms to the probability axioms, then it is not possible to construct a lottery such that the gambler sustains a sure loss (win).*

Necessity (the original theorem) was independently established by de Finetti (1937) (who introduced the terminology that a belief system that avoids sure loss is *coherent*) and Ramsey (1950), and sufficiency (the converse theorem) was independently established by Kemeny (1955) and Lehman (1955). Combining the Isomorphism Lemma and the Dutch Book Theorem results immediately in the following theorem.

Theorem 2. *For it to be impossible for either subjugation or subversion to occur, the coordination utility and the individual decision functions must comply with the axioms and syntax of probability theory.*

Following de Finetti's terminology, we offer the following definition.

Definition 15. *A network is* socially coherent *if the individual utilities (both categorical and conditional) are expressed and combined according to the syntax of probability theory.*

Since the coordination utility is the marginal of the coordination function, it follows that coherence requires that the conditional utilities must be utility mass functions, and the coordination function must be synthesized as given by (18).

The lack of social coherence does not mean that if the utilities are not structured according to the probability theorem then subjugation is sure to occur. Rather, it means that if the utilities are so structured and combined, then subjugation is impossible.

The restriction to acyclicity limits the generality of this approach. Nevertheless, this model still represents an important generalization from mutual social independence, where all agents possess categorial utilities and thus all social influence is trivially acyclical. In Section 4 we extend the theory to cyclical networks.

3.4 Invariance

The assumptions of directionality, conditionality, and coherence make the appropriation of the probability syntax an attractive and potentially useful framework within which to model social influence. Before simply adopting this syntax, however, we must ensure that this formulation complies with one additional condition that is assumed, often implicitly, with probability theory.

Invariance: The coordination function must be independent of the way the coordination function is synthesized, as long as exactly the same information is used for its construction.

To address invariance, we again resort the probability analogy. One of the key properties of probability theory is that a joint probability mass function is invariant to the way it is framed, that is,

$$p_{12}(y_1, y_2) = p_1(y_1)p_{2|1}(y_2|y_1) = p_{1|2}(y_1|y_2)p_2(y_2) = p_{21}(y_2, y_1)\,. \tag{45}$$

Invariance is easily verified when probability is used with frequentist scenarios such as predicting outcomes of coin flips and dice rolls. The situation is more problematic, however, under subjective interpretations involving beliefs. Let y_1 be the event that an athletic team wins the game, with probability $p(y_1)$, and let y_2 be the event that the star player does not play, with probability $p(y_2)$. Now let $p_{1|2}(y_1|y_2)$ denote the conditional probability that the team wins, given that the star player does not participate, and let $p_{2|1}(y_2|y_1)$ denote the conditional probability that star player does not play, given that the team wins the game. Invariance requires that the joint probability satisfies (45). There is no objective mechanism, however, that relates these conditional events in the way that the number of faces that turn up when dice are rolled are conditionally related. These latter relationships can be empirically computed, but the conditional relationships between the team winning and the star player not playing are based on subjective beliefs.

The vulnerability of subjective probability to invariance violations is often ignored in practice because the analyst would most likely use only one framing and would not be inclined to confirm consistency with alternate framings. Nevertheless, invariance is implicitly invoked in the synthesis of probability models used to characterize the way people organize information. Pearl argues that, in practice, multivariate distributions are rarely determined by specifying all of the entries in a joint-distribution table. "Probabilistic judgments on a small number of propositions are issued swiftly and reliably, while judging the likelihood of a conjunction of propositions entails much difficulty and hesitancy. This suggests that the elementary building blocks of human knowledge are not entries of a joint-distribution table. Rather, they are low-order marginal and conditional probabilities defined over small clusters of propositions" (Pearl, 1988, p. 78). Given this hypothesis, a natural way to construct a joint distribution is to synthesize it from conditional and marginal distributions. But that is *not* how the classical development of probability theory actually works. Under the conventional development of probability theory, conditional and marginal probability mass functions are derived from a joint probability mass function as the primitive component. Nevertheless, as Pearl observes, it is common and extremely useful to view the conditional and marginal probability mass functions as primitives with which to construct the joint mass function.

Since the goal is to synthesize a coordination function, invariance is also relevant in the behavioral domain. The assumption behind invariance is that, although there are many possible ways to organize a network, if they all are based on the same information set and if the information is used in a consistent and logical way without distorting or discarding any data, then they should all generate the same coordination function. It is not presumed that invariance will apply to all social situations, but it is a reasonable condition that will apply to many human decision making scenarios. Its application to the design of artificial systems, however, is less controversial. Artificial agents must operate according to the model that is used to design them. If they are designed to use all of the available information in a consistent way, then it is not unreasonable to require invariance. In fact, it would be highly desirable as a fundamental regularity property that would reduce or eliminate inconsistent or contradictory behavior.

4 Reciprocity

> One is not, in tacit coordination, trying to guess what another will do in an objective situation; one is trying to guess what the other will guess one' self to guess the other to guess, and so on ad infinitum.
>
> Thomas C. Schelling, *The Strategy of Conflit*

Thus far, our considerations of social influence have been unidirectional — from parent to child. With this model, the child is able to modulate its preferences according to the hypothesized intentions of the parent, but there is no reciprocal way for the child to influence the parent. In other words there is no opportunity for dialogue — a dynamic process whereby individuals may work together with the hope of converging to a result that is rational, fair, and legitimate. Such communication can be accomplished either by explicit bargaining or, if that is not possible, by virtual bargaining as discussed by Misyak et al. (2014). Conditionalization provides a natural mechanism, especially in the case of virtual bargaining, for individuals to engage in discourse. In this section we extend conditional game theory to accommodate networks that contain cycles; that is, where influence propagates bidirectional between agents.

In the probabilistic context, acyclicity is imposed because conditional relationships are tightly coupled. Given $P(A|B)$, one can compute the inverse relationship with the formula $P(B|A) = P(A|B)P(B)/P(A)$. Probability theory does not permit independent specifications of the two conditional probabilities. Many useful social networks, however, do involve the independent specification of conditional utilities. Consider the two-agent scenario

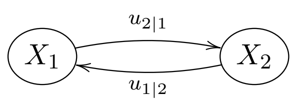

, (46)

where X_1 influences X_2 who in turn influences X_1 and so forth, and for which there is no constraint binding $u_{1|2}$ and $u_{2|1}$. Such a cyclic influence relationship could result in an infinite regression with no ultimate resolution. Indeed, Bacharach (2006) argues that such phenomena make it impossible to generate group-level intentions from individual intentions.

But the situation need not be so bleak. In fact, there is a natural and powerful way to resolve this problem. The key observation is to recognize that a cyclic network can be expressed dynamically as a time sequence of acyclic networks. From this perspective, we may view (46) as

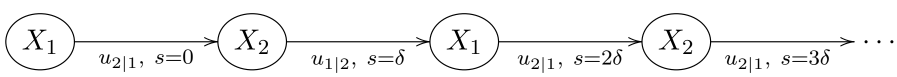

, (47)

where s denotes time and δ is the time interval between transmission of information between the agents.

Given this dynamic structure, the issue devolves around whether this time sequence of updates converges. To develop this theory, it is convenient to introduce notation that is amenable to the dynamic nature of this representation.

4.1 Matrix Form Dynamics Model

Definition 16. *Recall, from Definition 2, that a path X_j to X_j is a cycle. A k-member cycle is called a* simple k-cycle *if every vertex has exactly one incoming edge and one outgoing edge.*

Consider the simple k-cycle illustrated in Figure 2. We may view this cyclic network as a sequence of influence operations that evolve in time, Let $s = 0, \delta, 2\delta, \cdots$, and consider the equivalent dynamic network displayed in Figure 3.

Let us first consider the segment 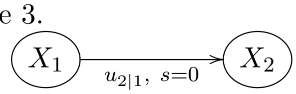

.

At $s = 0$, X_1's marginal utility mass function is $v_1(\mathbf{a}_1, 0)$. At $t = \delta$, the coordination function is, following (14),

$$u_{12}(\mathbf{a}_1, \mathbf{a}_2, \delta) = u_{2|1}(\mathbf{a}_2|\mathbf{a}_1)v_1(\mathbf{a}_1, 0)\,, \tag{48}$$

from which we may compute X_2's marginal using (20):

$$v_2(\mathbf{a}_2, \delta) = \sum_{\mathbf{a}_1} u_{12}(\mathbf{a}_1, \mathbf{a}_2, \delta)\,. \tag{49}$$

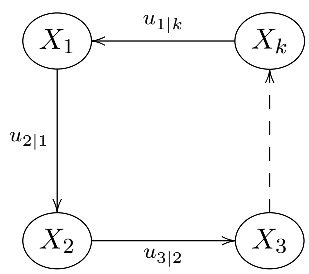

Figure 2: A simple k cycle.

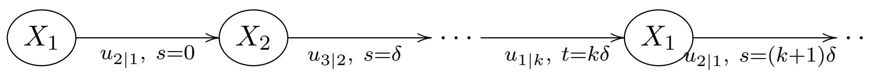

Figure 3: An equivalent dynamic network.

Now consider the next segment 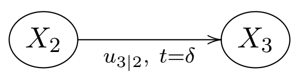

. At $s = 2\delta$ the coordination function is

$$u_{23}(\mathbf{a}_2, \mathbf{a}_3, 2\delta) = u_{3j2}(\mathbf{a}_3|\mathbf{a}_2)v_2(\mathbf{a}_2, \delta)\,, \tag{50}$$

and X_3's marginal is

$$v_3(\mathbf{a}_3, 2\delta) = \sum_{\mathbf{a}_2} u_{23}(\mathbf{a}_2, \mathbf{a}_3, 2\delta)\,. \tag{51}$$

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

$$\boldsymbol{\mathcal{A}} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n = \{\mathbf{z}_1, \ldots, \mathbf{z}_N\}\,, \tag{52}$$

where each profile $\mathbf{z}_k$ is of the form

$$\mathbf{z}_k = (z_{1k_1}, \ldots, z_{nk_n}) \tag{53}$$

where z_{ik_i} is the k_ith element of $\mathcal{A}_i$, and define the *utility mass vector*

$$\mathbf{v}_i(s) = \begin{bmatrix} v_i(\mathbf{z}_1, s) \\ \vdots \\ v_i(\mathbf{z}_N, s) \end{bmatrix}\,. \tag{54}$$

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

$$T_{i+1|i} = \begin{bmatrix} u_{i+1|i}(\mathbf{z}_1|\mathbf{z}_1) & \cdots & u_{i+1|i}(\mathbf{z}_1|\mathbf{z}_N) \\ \vdots & \vdots & \vdots \\ u_{i+1|i}(\mathbf{z}_N|\mathbf{z}_1) & \cdots & u_{i+1|i}(\mathbf{z}_N|\mathbf{z}_N) \end{bmatrix}\,. \tag{55}$$

With this notation, we may combine the operations defined by (48) and (49) with the single expression

$$\mathbf{v}_2(\delta) = T_{1|2}\mathbf{v}_1(0)\,, \tag{56}$$

and replace (50) and (51) with

$$\mathbf{v}_3(2\delta) = T_{3|2}\mathbf{v}_2(\delta)\,. \tag{57}$$

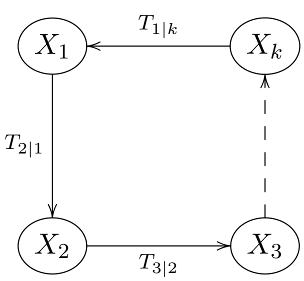

Figure 4: A k-cycle

Figure 4 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

$$\mathbf{v}_{i+k}(k\delta) = T_{i+k|i+k-1}T_{i+k-1|i+k-2}\cdots T_{i+2|i+1}T_{i+1|i}\mathbf{v}_i(0) \tag{58}$$

or, since all indices are incremented mod k,

$$\mathbf{v}_i(k\delta) = T_{i|i+k-1}T_{i+k-1|i+k-2}\cdots T_{i+2|i+1}T_{i+1|i}\mathbf{v}_i(0)\,. \tag{59}$$

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

$$T_i = T_{i|i+k-1}T_{i+k-1|i+k-2}\cdots T_{i+2|i+1}T_{i+1|i}\,. \tag{60}$$

Also, it is convenient to express time in units equal to the interval $k\delta$. Thus, we may write (59) as

$$\mathbf{v}_i(1) = T_i\mathbf{v}_i(0) \tag{61}$$

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 17. *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 3 (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 18. *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 4 (Markov Convergence)**.** *If T is a regular transition matrix, there exists a unique mass vector $\overline{\mathbf{v}}$ such that*

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

Thus, a network whose dynamic behavior is governed by regular transition matrices will, in the limit, converge to a network where the agents possess constant marginal utility mass functions, that is,

$$\lim_{t\to\infty} \mathbf{v}_i(t) = \overline{\mathbf{v}}_i = \begin{bmatrix}\overline{v}_i(z_{i1})\\ \vdots \\ \overline{v}(z_{iN})\end{bmatrix}. \tag{62}$$

These utilities are termed *steady-state utilities,* and correspond to the marginal utilities defined for acyclic networks by (20). As been established by Doob (1953), the convergence is exponentially fast. Thus, as a practical matter, convergence will effectively be reached in only a few cycles.

Throughout this development, we have assumed a condition of *stationarity*: the transition matrix does not change as time progresses. This constraint, however, is not a necessary condition. It is also possible for a non-stationary process to converge. Let $T(k)$ denote the transition matrix at time k. Then the study of convergence requires convergence of the product $T(k)T(k-1)\cdots T(0)$ rather than the much simpler product T^k. This analysis rather complex and will not be detailed in this paper, but we briefly sketch the essential concept. If the elements of the transition matrices are of bounded variation, then, subject to technical constraints, the columns of $T(k)T(k-1)\cdots T(0)$ all converge to the same vector as $k \to \infty$. For details, see Anily and Federgruen (1987). Essentially, this means that if the variations in the transition matrices diminish sufficiently over time, then a steady-state solution exists.

4.3 Coordination Function for Networks with Cycles

When a network contains cycles, we must first let each cycle achieve steady-state and then compute the coordination function with the utilities that correspond to the members of each cycle replaced by their steady-state unconditional utilities. In this paper we focus on simple n-cycle networks. In the interest of brevity and without loss of generality, we further restrict attention to 2×2 networks defined by (46) with steady-state marginal utilities $\overline{v}_1$ and $\overline{v}_2$ as defined by (62). The steady-state coordination function may be synthesized as

$$\overline{u}_{12}[(a_{11}, a_{12}), (a_{21}, a_{22})] = u_{1|2}(a_{11}, a_{12}|a_{21}a_{22})\overline{v}_2(a_{21}, a_{22}) \tag{63}$$

or, by invoking invariance,

$$\overline{u}_{12}[(a_{11}, a_{12}), (a_{21}, a_{22})] = u_{2|1}(a_{21}, a_{22}|a_{11}a_{12})\overline{v}_2(a_{11}, a_{12})\,. \tag{64}$$

We may then extract the coordination utility by computing the marginal

$$\overline{w}_{12}(a_{11}, a_{22}) = \sum_{a_{12}, a_{21}} \overline{u}_{12}[(a_{11}, a_{12}), (a_{21}, a_{22})] \tag{65}$$

and the steady-state coordinated individually rational decision functions become

$$\begin{aligned} \overline{w}_1(a_{11}) &= \sum_{a_{22}} \overline{w}_{12}(a_{11}, a_{22}) \\ \overline{w}_2(a_{22}) &= \sum_{a_{11}} \overline{w}_{12}(a_{11}, a_{22})\,. \end{aligned} \tag{66}$$

Example 6. *We recast the dissociated Battle of the Sexes as a cyclic game as follows:*

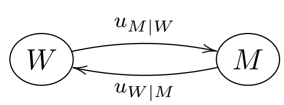

(67)

with conditional utilities given by

$$\begin{aligned} u_{W|M}(D|D) &= 1-\alpha & u_{M|W}(D|D) &= 1 \\ u_{W|M}(B|D) &= \alpha & u_{M|W}(B|D) &= 0 \\ u_{W|M}(D|B) &= 0 & u_{M|W}(D|B) &= \beta \\ u_{W|M}(B|B) &= 1 & u_{M|W}(B|B) &= 1-\beta\,, \end{aligned} \tag{68}$$

with corresponding state-to-state transition matrices

$$T_{W|M} = \begin{bmatrix} 1-\alpha & 0 \\ \alpha & 1 \end{bmatrix} \qquad T_{M|W} = \begin{bmatrix} 1 & \beta \\ 0 & 1-\beta \end{bmatrix}. \tag{69}$$

The closed-loop transition matrices are

$$T_W = T_{W|M}T_{M|W} = \begin{bmatrix} 1-\alpha & \beta-\alpha\beta \\ \alpha & 1-\beta+\alpha\beta \end{bmatrix} \tag{70}$$

and

$$T_M = T_{M|W}T_{W|M} = \begin{bmatrix} 1-\alpha+\alpha\beta & \beta \\ \alpha-\alpha\beta & 1-\beta \end{bmatrix}, \tag{71}$$

resulting in the steady-state marginal utilities (which, due to dissociation, become the individual decision functions)

$$\overline{\mathbf{w}}_W = \begin{bmatrix} \overline{w}_W(D) \\ \overline{w}_W(B) \end{bmatrix} = \frac{1}{\alpha + \beta - \alpha\beta} \begin{bmatrix} \beta - \alpha\beta \\ \alpha \end{bmatrix} \tag{72}$$

and

$$\overline{\mathbf{w}}_M = \begin{bmatrix} \overline{w}_M(D) \\ \overline{w}_M(B) \end{bmatrix} = \frac{1}{\alpha + \beta - \alpha\beta} \begin{bmatrix} \beta \\ \alpha - \alpha\beta \end{bmatrix} . \tag{73}$$

It is immediate that M's coordinated decision rule is to choose D when $\beta < \frac{\alpha}{1-\alpha}$, and W will choose D when $\beta < \frac{\alpha}{1+\alpha}$. Furthermore, we may compute the steady-state coordination utility as

$$\overline{w}_{MW}(a_{MM}, a_{WW}) = u_{M|W}(a_{MM}|a_{WW})\overline{w}_W(a_{WW}) , \tag{74}$$

yielding

$$\begin{aligned} \overline{w}_{MW}(D, D) &= \frac{\beta - \alpha\beta}{\alpha + \beta - \alpha\beta} \\ \overline{w}_{MW}(D, B) &= \frac{\alpha\beta}{\alpha + \beta - \alpha\beta} \\ \overline{w}_{MW}(B, D) &= 0 \\ \overline{w}_{MW}(B, B) &= \frac{\alpha - \alpha\beta}{\alpha + \beta - \alpha\beta} . \end{aligned} \tag{75}$$

Figure 5 displays the steady-state regions as a function of (α, β), and Table 8 interprets the four regions.

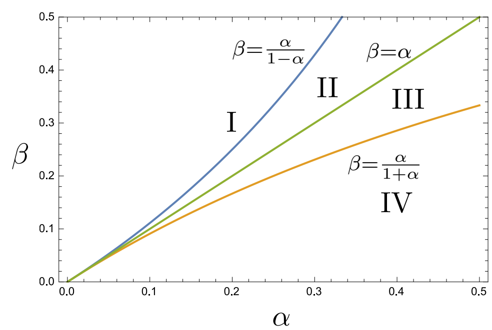

Figure 5: Plot of coordination regions for the cyclic Battle of the Sexes conditional game

5 Coordinatability

> Any good mathematical commodity is worth generalizing.
>
> Michael Spivak, *Calculus on Manifolds*

Table 8: The coordination utility and the coordinated individual decision functions for the cyclic Battle of the Sexes game.

(α, β) Region	Coordination utility ordering	i M ordering	ndividual W ordering	GTR decision	Consensus choice?
I	$D, D \succ_{MW} B, B$	$D \succ_M B$	$D \succ_W B$	D	yes
II	$D, D \succ_{MW} B, B$	$B \succ_M D$	$D \succ_W B$	D	no
III	$B, B \succ_{MW} D, B$	$B \succ_M D$	$D \succ_W B$	B	no
IV	$B, B \succ_{MW} D, B$	$B \succ_M D$	$B \succ_W D$	B	yes

The coordination utility and the coordinated individual decision functions can be used to form an operational definition of coordination. To introduce this development, we continue to explore the isomorphism and focus on the two-dimensional network defined by (1) and (12), its epistemological dual. Using (13), we may form the marginal probabilities via . If Y_1 and Y_2 are statistically independent, then conditioning on Y_1 has no effect on Y_2, and $p_{2|1} \equiv p_2$, in which case the joint probability mass function defined by (13) assumes the form $p_{12}(y_1, y_2) = p_1(y_1)p_2(y_2)$. But if Y_1 and Y_2 are not statistically independent, then it is not possible to synthesize the joint probability mass function as the product of the marginals. Thus, marginalization, as defined by (5), and synthesis, as defined by (13), are not reversible operations. It is of great theoretical interest, therefore, to investigate the effect of assuming that Y_1 and Y_2 are statistically independent when, in fact, they are not. In other words, if we are provided only the marginals and use them to synthesize the joint mass function under the assumption of independence, how different are $p_{12}(y_1, y_2)$ and $p_1(y_1)p_2(y_2)$? One way to address this question is to compute the *mutual information* between Y_1 and Y_2, as defined by the expression

$$I(Y_1, Y_2) = \sum_{y_1, y_2} p_{12}(y_1, y_2) \log_2 \frac{p_{12}(y_1, y_2)}{p_1(y_1)p_2(y_2)} \,. \tag{76}$$

It can be shown (e.g., (Cover and Thomas, 1991)) that mutual information is nonnegative and is zero if, and only if, Y_1 and Y_2 are statistically independent.[9] Mutual information is a measure of the degree of statistical interdependence between Y_1 and Y_2 and serves as an operational definition of statistical dependence.[10] The larger the mutual information, the more significant the dependence. The notion of mutual information generalizes to multivariate case comprising n random variables.

Now let us now apply the concept of mutual information to the behavioral domain by computing the mutual information with respect to the coordination utility (23) and the individual decisions functions (24), yielding

$$I(X_1, \ldots, X_n) = \sum_{a_{11}, \ldots, a_{nn}} w_{1:n}(a_{11}, \ldots, a_{nn}) \log_2 \frac{w_{1:n}(a_{11}, \ldots, a_{nn})}{w_1(a_{11}) \cdots w_n(a_{nn})} \,. \tag{77}$$

[9] Mutual information is a special case of a more general concept, the *Kullback-Leibler divergence* (Kullback and Leibler, 1951) between two probability mass functions p and q, defined as $D(p||q) = \sum_y p(y) \log \frac{p(y)}{q(y)}$. The Kullback-Liebler divergence is not a true metric (it is not symmetric and does not satisfy the triangle inequality), but it is nonnegative and is zero if and only if $p \equiv q$. Mutual information is a key concept of information theory, as developed by Shannon (1948) to characterize the theoretical performance of communication systems. Base 2 logarithms are used because digital communication systems typically encode messages with *bi*nary digi*ts*, or *bits*.

[10] It is interesting that introductory treatments of probability theory provide an operational definition of *statistical independence*, and treat a condition of statistical dependence simply as *not* statistical independence.

Operationally, *coordination* is the behavioral analogue to the statistical concept of *dependence*.[11] Thus, mutual information, viewed in the behavioral domain, is a measure of the degree of coordination among $X_1, \ldots, X_n$. If $I(X_1, \ldots, X_n) = 0$, then all individuals possess categorical utilities and there is no social linkage among the agents — they are *socially uncoordinated*. This is true even if the individual interests of all agents are identical. *Coincidence of interest is not the same as coordinated interest*. With the former, the agents just happen to be in the fortuitous situation where they can each take advantage of the fact that they all desire the same output without having any actual social contact. An example is the Hi-Lo game, where two individuals independently choose between a high and a low reward and receive their chosen rewards if they agree, otherwise both receive nothing. Under the classical formulation, the payoffs are categorical, and the mutual information (i.e., the strength of the social linkage between them) is zero. But if there is a social relationship between them such that $I(X_1, X_2) > 0$, then a degree of shared interest exists and there is an opportunity to make a coordinated decision. In Section 5 we will discuss this concept more thoroughly.

5.1 Coordination Index for Two Agents

Mutual information for a 2×2 network $\{X_1, X_2\}$ is

$$I(X_2, X_2) = \sum_{a_{11}, a_{22}} w_{12}(a_{11}, a_{22}) \log_2 \frac{w_{12}(a_{11}, a_{22})}{w_1(a_{11}) w_2(a_{22})} \,. \tag{78}$$

To develop this concept more thoroughly, we introduce the notion of *entropy*, which we approach first from the epistemological perspective.

Definition 19. *Let $\{Y_1, Y_2\}$ be discrete random variables defined over finite sets $\mathcal{Y}_1$ and $\mathcal{Y}_2$, let p_1, and p_2 be marginal probability mass functions for Y_1 and Y_2, respectively, and let p_{12} be the joint probability mass function for $\{Y_1, Y_2\}$. The* entropy *of Y_i is*

$$H(Y_i) = -\sum_{y_i} p_i(y_i) \log_2 p_i(y_i) \;\; i = 1, 2 \,, \tag{79}$$

and the joint entropy *of $\{Y_1, Y_2\}$ is*

$$H(Y_1, Y_2) = -\sum_{y_1, y_2} p_{12}(y_1, y_2) \log_2 p_{12}(y_1, y_2) \,. \tag{80}$$

For a detailed treatment of entropy, see Cover and Thomas (1991). Entropy is a central concept in the development of information theory, as introduced by Shannon (1948), and is used extensively as an analysis tool to study the behavior of communication systems. In this context, entropy is viewed as a numerical measure of the average epistemic uncertainty associated with a random phenomenon. It is straightforward to see that entropy is zero if all of the probability mass is concentrated on one outcome, and it is maximized if the probability mass is equally distributed between the two outcomes. It is also straightforward to show that if Y_1 and Y_2 are independent random variables, then

$$H(Y_1, Y_2) = H(Y_1) + H(Y_2) \,. \tag{81}$$

[11] Arrow (1974) formed the analogy of passing data between economic agents with the communications activity of passing a message from a transmitter to a receiver. Although he broached the idea of using Shannon information theory for the analysis of communication within an organization, he focused mainly on the qualitative insight that the analogy offers, rather than the quantitative possibilities.

In other words, if Y_1 and Y_2 are independent, then the average uncertainty of their joint realization is the sum of the average uncertainties of each individual realization.

It is straightforward to show that there is a close relationship between mutual information, as defined by (76) and entropy, namely,

$$I(Y_1, Y_2) = H(Y_1) + H(Y_2) - H(Y_1, Y_2)\,. \tag{82}$$

We interpret this expression as follows. If we assume (incorrectly) that Y_1 and Y_2 are independent, then the joint entropy is the sum of the two individual entropies. The difference between this incorrect expression for joint entropy and the actual joint entropy is a measure of the seriousness of the incorrect assumption. Another way to think of that error is that it is the amount that knowledge is increased (or, equivalently, uncertainty is decreased) regarding the outcome of one random variable, given the outcome of the other random variable — hence the interpretation as mutual information, or information that is shared between the two random variables.

The isomorphism between probability theory and utility theory allows us to apply entropy in the behavioral domain. For a two-agent network $\{X_1, X_2\}$ and a profile set $\mathcal{A}_1 \times \mathcal{A}_2$, let w_1 and w_2 be coordinated utilities for X_1 and X_2, respectively, as defined by (24), and let w_{12} be the coordination utility for $\{X_1, X_2\}$ as defined by (23). The *entropy* of X_i is

$$H(X_i) = -\sum_{a_{ii}} w_i(a_{ii}) \log_2 w_i(a_{ii}) \;\; i = 1, 2\,, \tag{83}$$

the *joint entropy* of $\{X_1, X_2\}$ is

$$H(X_1, X_2) = -\sum_{a_{11}, a_{22}} w_{12}(a_{11}, a_{22}) \log_2 w_{12}(a_{11}, a_{22})\,, \tag{84}$$

and the *mutual information* of X_1 and X_2 is

$$I(X_1, X_2) = H(X_1) + H(X_2) - H(X_1, X_2)\,. \tag{85}$$

In the behavioral context, entropy is a measure of the average behavioral uncertainty regarding preference (i.e.,the lack of decisiveness) when choosing from among alternatives. In other words, it is a measure of the average opportunity cost (i.e., the utility of the alternatives not chosen) involved in making a choice. Entropy is zero if all of the utility mass is concentrated on one outcome, meaning that there is no opportunity cost when a choice is decisive, but if the utilities are information to the unit interval, thereby providing an automatic calibration of the degree of coordination.

Definition 20. *For a two-agent network $\{X_1, X_2\}$, the function*

$$d(X_1, X_2) = H(X_1, X_2) - I(X_1, X_2) = 2H(X_1, X_2) - H(X_1) - H(X_2) \tag{86}$$

is the dispersion function*.*

Theorem 5. *Let $\{X_1, X_2, X_3\}$ be three agents. The dispersion function satisfies the following conditions.*

$$d(X_1, X_2) = d(X_2, X_1) \;\; (\textit{symmetry}) \tag{87}$$

$$d(X_1, X_2) \geq 0 (\textit{non-negativity}) \tag{88}$$

$$d(X_1, X_2) = 0 \textit{ if and only if } X_1 \Leftrightarrow X_2 (\textit{strict positivity}) \tag{89}$$

$$d(X_1, X_2) \leq d(X_1, X_3) + d(X_3, X_2) \;\; (\textit{triangle inequality})\,, \tag{90}$$

where the notation $X_1 \Leftrightarrow X_2$ means that there exist permutations $\pi_{1|2}, \pi_{2|1} \colon \boldsymbol{\mathcal{A}} \to \boldsymbol{\mathcal{A}}$ such that

$$u_{j|i}(\mathbf{a}_j|\mathbf{a}_i) = \begin{cases} 0 & \text{if } \mathbf{a}_j \neq \pi_{j|i}(\mathbf{a}_i) \\ 1 & \text{if } \mathbf{a}_j = \pi_{j|i}(\mathbf{a}_i) \end{cases}, i, j \in \{1,2\}, i \neq j\,, \tag{91}$$

in which case X_j is slaved *to X_i. Thus $d(X_1, X_2)$is a true metric.*

A proof of this result may be found in Kraskov et al. (2003) and Li et al. (2001).

The dispersion function (86) expresses the "distance" between the interests of X_1 and X_2 as a function of their social relationship. This distance is minimized when they are slaved together and social influence is maximized. The distance is maximized if X_1 and X_2 are socially uncoordinated and thus exert no social influence on each other. We may refine this concept by normalizing the distance by the joint entropy, which gives rise to an explicit quantification of the degree of coordination as a direct result of social influence.

Definition 21. *The* relative dispersion $\mathcal{D}(X_i, X_j)$ *is defined by*

$$\mathcal{D}(X_i, X_j) = \frac{d(X_i, X_j)}{H(X_i, X_j)} \tag{92}$$

and the coordination index *is given by*

$$\mathcal{C}(X_i, X_j) = 1 - \mathcal{D}(X_i, X_j) = \frac{H(X_i) + H(X_j) - H(X_i, X_j)}{H(X_i, X_j)}\,. \tag{93}$$

Example 7. *Let us compute the coordination index for the cyclic Battle of the Sexes. We compute the coordination index as the product of the conditional utilities (75) and the steady-state utilities (68), yielding*

$$w_{MW}(a_{MM}|a_{WW})\overline{w}_W(a_{WW})], . \tag{94}$$

The coordination index is given by

$$\mathcal{C}(M, W) = \frac{H(M) + H(W) - H(M, W)}{H(M, W)}. \tag{95}$$

Figure 6 illustrates a contour plot of the coordination index for the cyclic Battle of the Sexes game over the region $(\alpha, \beta) \in [0, 1/2] \times [0, 1/2]$.

5.2 Coordination Index for $n > 2$ Agents

When $n > 2$, we may expand the definition of dispersion to become

$$d(X_1, \ldots, X_n) = nH(X_1, \ldots, X_n) - \sum_{i=1}^{n} H(X_i) \tag{96}$$

The relative dispersion function for $n > 2$ becomes

$$\mathcal{D}(X_1, \ldots, X_n) = \frac{1}{n-1} \frac{d(X_1, \ldots, X_n)}{H(X_1, \ldots, X_n)} \tag{97}$$

This function is symmetric and non-negative, and is zero if and only if there exist permutations $\pi_{i|j} \colon \boldsymbol{\mathcal{A}} \to \boldsymbol{\mathcal{A}}$ such that

$$u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\mathbf{a}_{i_1}, \ldots, \mathbf{a}_{i_{q_i}}) = \begin{cases} 0 & \text{if } \mathbf{a}_i \neq \pi_{i|i_k}(\mathbf{a}_i), k = 1, \ldots, q_i \\ 1 & \text{if } \mathbf{a}_i = \pi_{i|i_k}(\mathbf{a}_i), k = 1, \ldots, q_i \end{cases},\ i = 1, \ldots, n\,. \tag{98}$$

$\mathcal{D}(X_1, \ldots, X_n)$ is a a measure of how much the utilities of the group are in conflict, and achieves its maximum when all of the X_i's are mutually uncoordinated, in which case $\mathcal{D}(X_1, \ldots, X_n) = 1$.

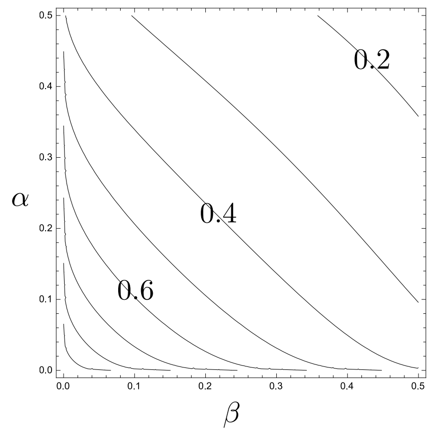

Figure 6: The coordination index for the cyclic Battle of the Sexes game.

Definition 22. *The* coordination index *of the network is*

$$\mathcal{C}(X_1, \ldots, X_n) = 1 - \mathcal{D}(X_1, \ldots, X_n) = \frac{\sum_{i=1}^{n} H(X_i) - H(X_1, \cdots, X_n)}{(n-1)H(X_1, \cdots, X_n)} . \tag{99}$$

The coordination index is a measure of the degree to which the members of a network are socially connected, and serves as a measure of the intrinsic ability of the individuals to align their interests *as a result of direct social influence.* Alignment can be positive, in which case the individuals have compatible shared interests (e.g., teams), or it can be negative, in which case they have conflicting shared interests (e.g, athletic contests). When all of the utilities are categorical, as is the case with classical noncooperative game theory, the individuals are mutually socially uncoordinated and the relative dispersion is maximum. Consequently, the coordination index is zero — the members have no shared social interests (although they may have shared material interests). This is true even if $u_1 = \cdots = u_n$. Common material interests do not imply shared social interests. This does not mean, of course, that mutually socially uncoordinated individuals cannot be aligned, nor does a lack of an intrinsic ability to align their interests mean that the agents will not function harmoniously. Rather, it means that if they do, it is simply by coincidence, rather than by social design.

The coordination index is the theoretical upper bound on the intrinsic ability of a network to align its interests. Basically, it tells how theoretically amenable a given organizational or network structure is for socially coordinated behavior. But it cannot tell whether that structure supports the kind of coordination that one desires.

One way to think of the coordination index is as a measure of the ecological fitness of a given network to function appropriately in its environment. It can also be used as a design tool to evaluate the ecological fitness of a proposed network structure. A low coordination index may prompt structural changes such as inserting additional linkages or modifying existing links.

6 Extension to Stochastic Agents

Conditional game theory has appropriated the syntax of probability theory in several ways: the concept of conditioning to form the social influence linkages between individuals; the concept of coherence (avoiding subjugation) to motivate the need to express preferences in terms of conditional mass functions that are manipulated according to the probability syntax; Bayesian network theory to develop an explicit social model as a coordination ordering over conjectures; Shannon information theory to define an operational concept of coordination; and Markov convergence theory has been applied to deal with cyclic networks. What we have not yet done, however, is apply probability theory in its traditional usage as a means to model epistemological uncertainty.

One way for uncertainty to arise is if some of the members of the network are random variables whose realizations can influence the interests of the deterministic members of the network. For example, agents may be influenced by phenomena beyond their control, such as environmental factors, and ignoring such factors could distort or invalidate the social model. Consider the network illustrated in Figure 7, where Y is a discrete random variable defined over some finite set $\mathcal{Y}$ that influences X, where X is whose a deterministic agent defined over an alternative set $\mathcal{A}$, and Z is a discrete random variable defined over a finite set $\mathcal{Z}$.

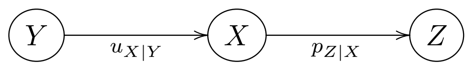

Figure 7: Transitions between stochastic and deterministic individuals.

The function $u_{X|Y}$ has the syntax of a conditional utility mass function whose conditioning entity is a stochastic phenomenon, rather than another individual, and $p_{Z|X}$ has the syntax of a conditional probability mass function whose conditioning entity is a deterministic individual, rather than another stochastic phenomenon. This structure involves the mixing of epistemological phenomena with behavioral phenomena. By invoking the isomorphism, coherence, and invariance, the syntax of probability theory and the syntax of utility theory are compatible. Referring to Figure 7, the function $u_{X|Y}$ is a conditional utility mass function whose conditioning entity is a random phenomenon, and $p_{Z|X}$ is a probability mass function whose conditioning entity is a behavioral phenomenon. Thus, epistemological entities can be seamlessly absorbed into a behavioral network and treated without distinction as far as constructing the coordination function is concerned.

Definition 23. *Let $\{X_1, \ldots, X_n\}$ denote a network of n deterministic agents, and let $\{Y_1, \ldots, Y_m\}$ denote a set of discrete stochastic agents (random variables) where Y_j is defined over a finite set $\mathcal{B}_j$, $j = 1, \ldots, m$. The combined set of individuals $\{X_1, \ldots, X_n, Y_1, \ldots, Y_m\}$ is defined over the product set $\mathcal{A}^n \times \mathcal{B}_1 \times \cdots \times \mathcal{B}_m$. comprise a* stochastic network.

A stochastic conjecture profile *is an array $(\mathbf{a}_1, \ldots, \mathbf{a}_n, b_1, \ldots, b_m) \in \mathcal{A}^n \times \mathcal{B}_1 \times \cdots \times \mathcal{B}_m$ such that $(\mathbf{a}_1, \ldots, \mathbf{a}_n)$ is a conjecture profile for $\{X_1, \ldots, X_n\}$ and $(b_1, \ldots, b_m)$ is a* stochastic profile *for $\{Y_1, \ldots, Y_m\}$.*

A stochastic coordination function *is a function $u_{1:n\,1:m}$: $\boldsymbol{\mathcal{A}}^n \times \mathcal{B}_1 \times \cdots \times \mathcal{B}_m \to [0, 1]$ that characterizes all of the social relationships that exist among the deterministic agents, the stochastic dependency relationships that exist among the stochastic agents, the stochastic relationships that the deterministic agents exert on the stochastic agents, and the social relationships that the stochastic agents exert on the deterministic agents.*

Suppose X_i has p_i deterministic parents and q_i stochastic parents, and Y_j has r_j deterministic parents and s_j stochastic parents, that is,

$$\begin{aligned} \mathrm{pa}\,(X_i) &= \{X_{i_1}, \ldots, X_{i_{p_i}}, Y_{k_1}, \ldots, Y_{k_{q_i}}\} \\ \mathrm{pa}\,(Y_j) &= \{X_{j_1}, \ldots, X_{j_{r_j}}, Y_{l_1}, \ldots, Y_{l_{s_j}}\} \end{aligned} . \tag{100}$$

For any stochastic conjecture profile $(\mathbf{a}_1, \dots, \mathbf{a}_n, b_1, \dots, b_m)$, let

$$\begin{aligned} \boldsymbol{\alpha}_{\mathrm{pa}(i)} &= (a_{i_1}, \dots, a_{i_{p_i}}, b_{k_1}, \dots, b_{k_{q_i}}) \\ \mathbf{b}_{\mathrm{pa}\,(j)} &= (a_{j_1}, \dots, a_{j_{r_j}}, b_{l_1}, \dots, b_{l_{s_j}}) \end{aligned} \tag{101}$$

denote the conditioning conjectures of $\mathrm{pa}\,(X_i)$ and $\mathrm{pa}\,(Y_j)$, respectively.

The *stochastic coordination function* is of the form

$$\hat{u}_{1:n\,1:m}(\mathbf{a}_1, \dots, \mathbf{a}_n, b_1, \dots, b_m) = \prod_{i=1}^{n} \prod_{j=1}^{m} u_{i|\mathrm{pa}(i)}(\mathbf{a}_i | \boldsymbol{\alpha}_{\mathrm{pa}(i)}) p_{j|\,\mathrm{pa}\,(j)}(b_j | \mathbf{b}_{\mathrm{pa}\,(j)})\,. \tag{102}$$

Although the stochastic agents influence the preferences of the deterministic agents, they are not able to make choices. Thus, to make a choice, we must compute the *expected coordination function* obtained by summing over the stochastic states, yielding

$$\hat{u}_{1:n}(\mathbf{a}_1, \dots, \mathbf{a}_n) = \sum_{b_1, \dots, b_m} u_{1:n\,1:m}(\mathbf{a}_1, \dots, \mathbf{a}_n, b_1, \dots, b_m)\,. \tag{103}$$

The *expected coordination utility* and the *expected individual decision functions* are thus

$$\hat{w}_{1:n}(a_{11}, \dots, a_{nn}) = \sum_{\sim a_{11}} \cdots \sum_{\sim a_{nn}} \hat{u}_{1:n}[(a_{1n}, \dots, a_{1n}), \dots, (a_{n1} \dots, a_{nn})]\,. \tag{104}$$

and

$$\hat{w}_i(a_{ii}) = \sum_{\sim a_{ii}} \hat{w}_{1:n}(a_{11}, \dots, a_{1n})\,. \tag{105}$$

Example 8. *Returning to the Battle of the Sexes, let A denote the availability of tickets to the Ballet, and let S denote a sold-out venue. Suppose that W's preferences are influenced by the availability of tickets to the Ballet. Let Y be a discrete random variable defined over the sample space $\mathcal{Y} = \{A, S\}$, with probability mass function*

$$\begin{aligned} p_Y(A) &= \gamma \\ p_Y(S) &= 1 - \gamma\,. \end{aligned} \tag{106}$$

Building on the completely dissociated model given in Example 5, suppose, under the hypothetical proposition that tickets are available, that W's preferences are as defined previously, but under the hypothetical proposition that the venue is sold out, then W will abandon her desire to attend the Ballet and place her entire conditional mass on D. The resulting conditional utility for W is

$$\begin{array}{ll} u_{W|Y}(D|A) = \alpha & u_{W|Y}(B|A) = 1 - \alpha \\ u_{W|Y}(D|S) = 1 & u_{W|Y}(B|S) = 0 \end{array}\,. \tag{107}$$

We continue to suppose that M's conditional preferences are as given (36). The stochastic coordination utility thus becomes

$$w_{MWY}(a_M, a_W, y) = u_{M|W}(a_M | a_W) u_{W|Y}(a_W | y) p_Y(y)\,, \tag{108}$$

yielding

$$
\begin{aligned}
w_{MWY}(D,D,A) &= \alpha\gamma \\
w_{MWY}(D,D,S) &= 1-\gamma \\
w_{MWY}(D,B,A) &= (1-\alpha)\beta\gamma \\
w_{MWY}(D,B,S) &= 0 \\
w_{MWY}(B,D,A) &= 0 \\
w_{MWY}(B,D,S) &= 0 \\
w_{MWY}(B,B,A) &= (1-\alpha)(1-\beta)\gamma \\
w_{MWY}(B,B,S) &= 0\,.
\end{aligned} \tag{109}
$$

The expected coordination utility is

$$\hat{w}_{MW}(a_M,a_W) = \sum_y w_{MWY}(a_M,a_W,y)\,, \tag{110}$$

yielding

$$
\begin{aligned}
\hat{w}_{MW}(D,D) = 1-\gamma+\alpha\gamma \quad & \hat{w}_{MW}(B,D) = 0 \\
\hat{w}_{MW}(D,B) = (1-\alpha)\beta\gamma \quad & \hat{w}_{MW}(B,B) = (1-\alpha)(1-\beta)\gamma\,,
\end{aligned}\,. \tag{111}
$$

and the expected individual decision functions are

$$\hat{w}_M(D) = 1-(1-\alpha)(1-\beta)\gamma \quad \hat{w}_M(B) = (1-\alpha)(1-\beta)\gamma \tag{112}$$

and

$$\hat{w}_W(D) = 1-\gamma+\alpha\gamma \quad \hat{w}_W(B) = \gamma-\alpha\gamma\,. \tag{113}$$

Notice that when $\gamma = 1$ (i.e., the probability of ticket availability is unity), then the result agrees with Example 5.

Example 9. *As a final example involving the Battle of the Sexes, we consideer the model*

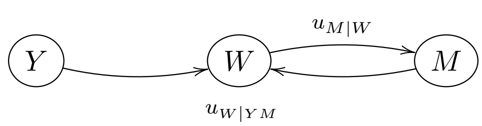

(114)

W's preferences are conditioned on Y and M thus become

$$
\begin{aligned}
u_{W|YM}(D|A,D) &= 1-\alpha \\
u_{W|YM}(B|A,D) &= \alpha \\
u_{W|YM}(D|A,B) &= 0 \\
u_{W|YM}(B|A,B) &= 1 \\
u_{W|YM}(D|S,D) &= 1 \\
u_{W|YM}(B|S,D) &= 0 \\
u_{W|YM}(D|S,B) &= 1 \\
u_{W|YM}(B|S,B) &= 0\,,
\end{aligned} \tag{115}
$$

yielding

$$T_{W|YM} = \begin{bmatrix} 1-\alpha & 0 & 1 & 1 \\ \alpha & 1 & 0 & 0 \end{bmatrix}\,. \tag{116}$$

To compute the conditional utility of Y and M, we note that Y and M are conditionally independent, given W, thus $u_{YM|W} = p_{Y|W}u_{M|W} = p_Y u_{M|W}$, where the last equality holds since Y is not influenced by W. There resulting conditional utility for $\{Y, M\}$ is

$$\begin{aligned}
u_{YM|W}(A, D|D) &= p_Y(A)u_{M|W}(D|D) = \gamma \\
u_{YM|W}(A, B|D) &= p_Y(A)u_{M|W}(B|D) = 0 \\
u_{YM|W}(S, D|D) &= p_Y(S)u_{M|W}(D|D) = (1-\gamma) \\
u_{YM|W}(S, B|D) &= p_Y(S)u_{M|W}(B|D) = 0 \\
u_{YM|W}(A, D|B) &= p_Y(A)u_{M|W}(D|B) = \gamma\beta \\
u_{YM|W}(A, B|B) &= p_Y(A)u_{M|W}(B|B) = \gamma(1-\beta) \\
u_{YM|W}(S, D|B) &= p_Y(S)u_{M|W}(D|B) = (1-\gamma)\beta \\
u_{YM|W}(S, B|B) &= p_Y(S)u_{M|W}(B|B) = (1-\gamma)(1-\beta),
\end{aligned} \tag{117}$$

yielding the transition matrix

$$T_{YM|W} = \begin{bmatrix} \gamma & \gamma\beta \\ 0 & \gamma(1-\beta) \\ 1-\gamma & (1-\gamma)(1-\beta) \\ 0 & (1-\gamma)(1-\beta)\,, \end{bmatrix} \tag{118}$$

from which the closed-loop transition matrix is computed as

$$T_W = T_{W|YM}T_{YM|W} = \begin{bmatrix} 1-\alpha\gamma & 1-\gamma+\beta\gamma-\alpha\beta\gamma \\ \alpha\gamma & \gamma-\beta\gamma+\alpha\beta\gamma\,. \end{bmatrix} \tag{119}$$

The expected steady-state coordinated individual decision vectors are

$$\begin{bmatrix} \overline{w}_W(D) \\ \overline{w}_W(B) \end{bmatrix} = \begin{bmatrix} \frac{1-\gamma+\beta\gamma-\alpha\beta\gamma}{1-\gamma+\beta\gamma+\alpha\gamma-\alpha\beta\gamma} \\ \frac{\alpha\gamma}{1-\gamma+\beta\gamma+\alpha\gamma-\alpha\beta\gamma} \end{bmatrix} \tag{120}$$

$$\begin{bmatrix} \overline{w}_M(D) \\ \overline{w}_M(B) \end{bmatrix} \begin{bmatrix} \frac{1-\gamma+\beta\gamma}{1-\gamma+\beta\gamma+\alpha\gamma-\alpha\beta\gamma} \\ \frac{\alpha\gamma-\alpha\beta\gamma}{1-\gamma+\beta\gamma+\alpha\gamma-\alpha\beta\gamma} \end{bmatrix}. \tag{121}$$

Notice that when $\gamma = 1$, these results become identical with the results given in Example 6.

7 Discussion

Classical game theory provides a model of how individuals might behave in group settings; it does not provide a model of how a group might behave. Except for what Shubik terms "inessential games," games are not really "solved" in the sense that an incontrovertible solution exists.

> The general n-person game postulates a separate "free will" for each of the contending parties and is therefore fundamentally indeterminate. To be sure, there are limiting cases, which game-theorists call "inessential games," in which indeterminacy can be resolved satisfactorily by applying the familiar principle of self-seeking utility maximization or individual rationality. But there is no principle of societal rationality, of comparable resolving power, that can cope with the "essential" game, and none is in sight. Instead, deep-seated paradoxes, challenging our intuitive ideas of what kinds of behavior should be called "rational," crop up on all sides as soon as we cross the line from "inessential" to "essential" (Shubik, 1982, p. 2).

This paper does not claim to introduce a notion of societal rationality in the sense of establishing concepts of group preference, group performance, or group choice. It certainly does not deny the existence of such phenomena, but does suggest that classical game theory is not an adequate mathematical tool with which to model them. Instead, we circumvent such on-going psychological/sociological/philosophical discussions and argue that coordinated agency is the operative concept when discussing group behavior. We do not conflict with the classical game-theoretic premise that preferences are innately individual. We do assert, however, that coordination is innately a group-level concept, and that a complete model of a group requires the integration of both aspects of behavior.

The logical mechanism that integrates individual performance and group coordination is *conditionalization*. This mechanism enables individuals to modulate their preferences in response to the social context. This logic is identical to the conditionalization logic of Bayesian epistemology, where conditional probability serves as the vehicle by which individual belief regarding the realization of some random phenomenon is modulated by belief regarding the realization of another random phenomenon. Bayesian conditionalization enables these beliefs to be connected in a coherent and systematic way. This connection does *not* result in a concept of group belief; rather, it establishes statistical dependence — a state where the beliefs regarding different phenomena influence each other to some extent but nevertheless remain individual.

Since beliefs and preferences are order isomorphic, the conditionalization syntax serves as the vehicle by which individual preferences regarding the actualization by some agent is modulated by preferences regarding the actualization by another agent. This approach enables the preferences to be connected in a coherent and systematic way. This connection does *not* result in a concept of group preference; rather, it establishes social coordination — a state where the preferences of different individuals influence each other but nevertheless remain individual.

This research has appropriated several well-known mathematical operations associated with probability theory: conditionalization, Bayesian networks, coherence (via the Dutch Book theorem), Shannon information theory, and Markov convergence theory. Although these operations are traditionally associated with probability theory as applied to epistemological uncertainty, they are context-neutral and apply equally well as a framework within which to model behavioral uncertainty.

Table 9 summarizes the contrast between classical noncooperative game theory and conditional game theory.

Table 9: Contrasts between classical and conditional game theory.

Issue	Classical Game Theory	Conditional Game Theory
Game structure	Divison of labor: preference model and solution concept defined separately	Social influence integrated into conditional preference model
Preference model	Categorical — fixed and immutable	Conditional — modeled after Bayesian conditionalization
Rationality concept	Individual rationality with regard to material performance	Individual rationality with regard to material benefit as modulated by social influence
Solution concept	Individual: constrained optimization (Nash equilibrium) Group: individuals possess a concept of shared intentions (team reasoning)	Individual: maximize the coordinated individual rationality utility Group: maximize coordination utility (generalized team reasoning)
Coordination concept	Extrinsic: exogenously defined social solution concept applied to payoffs	Intrinsic: endogenously emerges as social influence propagates through the network

Part II: The Three Bacharach Puzzles

In his posthumously published book *Beyond Individual Choice*, Bacharach (2006) discusses what he termed “the three puzzles of game theory,” comprising three games for which uniquely best solutions are seemingly obvious, but game theory fails to identify them: Hi-Lo, Matching Pennies, and Prisoner’s Dilemma. Hi-Lo is game involving two players who may choose between alternatives A and B, with A the one of greater value. They both receive the reward if, and only if, they agree; otherwise, neither receives anything. There are two pure-strategy Nash equilibria, (A, A) and (B, B), but game theory does not provide a formal rational mechanism for choosing the one with greater value. Bacharach argues that such paradoxes arise because solution concepts based exclusively on narrow self-interest provide an incomplete model of rational behavior: “Ultimately, the reason why (B, B) is a solution is that it is consistent with all the facts about rationality that game theory can muster. (B, B) is a solution because game theory has mustered no fact about rationality that excludes B” (Bacharach, 2006, p. 47). The paradox he associates with Matching Pennies is that, although people are much more likely to choose heads over tails, game theory’s best answer is a mixed strategy with equal probabilities (which is Parato-dominated by the two pure strategy equilibria). Bacharach’s problem with the Prisoner’s Dilemma is that, although mutual defection is the unique pure-strategy Nash equilibrium, it has been empirically established that humans often do not behave accordingly. We now analyze each of these examples by recasting them as games of bilateral social coordination, where the members of a two-agent network seek a coordinated solution. Notationally, we designate the row player as X_r and the column player as X_c.

1 Hi-Lo and Matching Pennies

Hi-Lo is a coordination game with payoff matrix displayed in Table 10(a). Without communicating, each player must choose between A and B. Bacharach argues that, for situations such as this, where the intuitive choice is incontrovertible, a theory of rational choice ought to single out (A, A) as the only rational outcome. But game theory identifies (B, B) as rational as well, since it is also a Nash equilibrium. The paradox exists because game theory cannot definitively predict that (A, A) will be chosen. Thus, the only way to resolve the paradox is to overlay the payoff matrix with psychological assumptions that are not part of the game definition. This suggests that something is missing in the way the game is defined.

The team reasoning approach does indeed impose an additional psychological component to the model; namely, that each agent undergoes an agent transformation from the status of a completely autonomous individual to a member of a team, and also undergoes a utility transformation from an individual utility to a group utility wherewith the group acts as an individual entity. The payoff matrix for the team reasoning formulation is given in Table 10(b) and clearly resolves the paradox in favor of (A, A) as the uniquely best team response.

The Matching Pennies scenario is as follows: Two players who are not allowed to communicate are shown a coin with one side “heads” (H) and the other side “tails” (T), and are given the following instruction: “Complete the sentence: A penny was tossed. It came down ____.” If both complete the sentence in the same way, they both receive a penny, but if they differ, then neither receives anything. The payoff matrix for this game is displayed in Table 11. This game has two Nash equilibria, both with the same payoffs, and classical game theory provides no way to distinguish between them. Experimental evidence establishes, however, that in such scenarios a substantial

Table 10: The payoff matrix for the Hi-Lo game: (a) The original payoff matrix, (b) the payoff array for the team $\{X_r, X_c\}$.

	X_c	
X_r	A	B
A	$2,2$	$0,0$
B	$0,0$	$1,1$

(a)

Outcome	Team Payoff
A, A	$2,2$
B, B	$1,1$
A, B	$0,0$
B, A	$0,0$

(b)

majority of players both choose "heads," although, on the face of it, there is no compelling reason to favor "heads" over "tails."

Table 11: Payoff matrix for Matching Pennies with H = heads, T= tails.

	X_c	
X_r	H	T
H	$1,1$	$0,0$
T	$0,0$	$1,1$

Under the assumptions of both Hi-Lo and Matching Pennies, the players do not communicate. However, as argued by Misyak et al. (2014), the players may certainly engage in *virtual* communication by independently reasoning how the other might be reasoning. Such scenarios are examples of Schelling's concept of *focal points*, also termed *salience*. "People *can* often concert their intentions or expectations with others if each knows that the other is trying to do the same. Many situations — perhaps every situation for people who are practiced at this kind of game — provide some clue for coordinating behavior, some focal point for each person's expectations of what the other expect him to expect to be expected to do" (Schelling, 1960, p. 57). This kind of conditional reasoning is of the form described by Bacharach (2006, p. 138): "I will if and only if you will; but I know that you will if and only if you know that I will; hence I will if and only if I know you know that I will; hence, similarly, I will if and only if I know you know I know you will; and so on." Bacharach argues that this leads to an endless sequence of implications that does not result in either joint or individual categorical intentions.

We assert, however, that such a pessimistic conclusion is premature, and that it is indeed possible for such iterated reasoning to converge to unconditional *coordinated* intentions. The only difference between the two scenarios is that, with Hi-Lo, the issue is choosing definitively between two equilibra with different material preference, and with Matching Pennies, the issue is choosing definitively between two equilibria with the same material preference, but with different social preference.

1.1 Hi-Lo

With the Hi-Lo scenario, the action set is $\mathcal{A}_r = \mathcal{A}_c = \{A, B\}$ and the hypothetical propositions are of the form

$$\begin{aligned} &\mathcal{H}_{i|j}(A|A)\colon X_j \models A \Rightarrow X_i \models A\\ &\mathcal{H}_{i|j}(B|A)\colon X_j \models A \Rightarrow X_i \models B\\ &\mathcal{H}_{i|j}(A|B)\colon X_j \models B \Rightarrow X_i \models A\\ &\mathcal{H}_{i|j}(B|B)\colon X_j \models B \Rightarrow X_i \models B \end{aligned} \tag{122}$$

for $i, j, \in \{r, c\}$, $i \neq j$. The task is to define conditional utilities that comport with these hypothetical propositions. Consider the hypothetical scenarios $\mathcal{H}_{i|j}(A|A)$ and $\mathcal{H}_{i|j}(B|B)$. Under $\mathcal{H}_{i|j}(A|A)$, X_i would rationally assume that X_j prefers A to B, which agrees with the fact that both know that $A > B$. Under $\mathcal{H}_{i|j}(B|B)$, X_i would also rationally assume that X_j prefers B to A, which would indicate to X_i that X_j has made a mistake. Let h_i denote X_i's probability that X_j **has not** made a mistake by conjecturing A, and let ℓ_i denote X_i's probability that X_j **has** made a mistake by conjecturing B. Then $1 - h_i$ is X_i's probability that X_j has made a mistake by conjecturing A, and $1 - \ell_i$ is X_i's probability that X_j has not made a mistake by conjecturing B.

Suppose that X_i assumes that X_j intends to be rational but is also succecptable to error. Then X_i would constrain h_i and ℓ_i to be greater than one half. If we also suppose that X_i's probability of X_j correctly conjecturing A is greater than X_i's probability that X_j will mistakenly conjecture B, then $h_i > \ell_i > 1/2$. In the interest of brevity and in keeping with the symmetry of the original problem formulation, we assume that $h_r = h_c = h$ and $\ell_r = \ell_c = \ell$, thereby yielding the following symmetric conditional utilities.

$$\begin{aligned} u_{j|i}(A|A) &= u_{i|j}(A|A) = h\\ u_{j|i}(B|A) &= u_{i|j}(B|A) = 1 - h\\ u_{j|i}(A|B) &= u_{i|j}(A|B) = 1 - \ell\\ u_{j|i}(B|B) &= u_{i|j}(B|B) = \ell\,, \end{aligned} \tag{123}$$

where $h > \ell > 1 - \ell > 1 - h$. This formulation of Hi-Lo is completely dissociated (see Section 3.2), and results in a cyclic network of the form

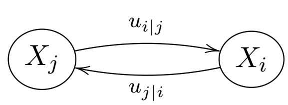

. (124)

Following (55), the state-to-state transition matrices are

$$T_{i|j} = T_{j|i} = \begin{bmatrix} h & 1-\ell \\ 1-h & \ell \end{bmatrix} \tag{125}$$

and, following (60), the closed-loop transition matrices are

$$T_i = T_j = T_{j|i}T_{i|j} = \begin{bmatrix} h^2 + (1-h)(1-\ell) & (h+\ell)(1-\ell) \\ (h+\ell)(1-h) & (1-h)(1-\ell) + \ell^2 \end{bmatrix} \tag{126}$$

for $i, j \in \{r, c\}$. The steady coordinated decision function is the eigenvector corresponding to the unit eigenvalue of (126), which is

$$\overline{\mathbf{v}}_r = \overline{\mathbf{v}}_c = \overline{\mathbf{v}} = \begin{bmatrix} \overline{v}(A) \\ \overline{v}(B) \end{bmatrix} = \begin{bmatrix} \frac{1-\ell}{2-h-\ell} \\ \frac{1-h}{2-h-\ell} \end{bmatrix}\,. \tag{127}$$

Following the development in Section 4.3 and applying (123) and (127), the steady-state coordination utility is

$$\overline{\mathbf{w}}_{rc} = \begin{bmatrix} \overline{w}_{rc}(A,A) \\ \overline{w}_{rc}(A,B) \\ \overline{w}_{rc}(B,A) \\ \overline{w}_{rc}(B,B) \end{bmatrix} = \begin{bmatrix} u_{r|c}(A|A)\overline{v}(A) \\ u_{r|c}(A|B)\overline{v}(B) \\ u_{r|c}(B|A)\overline{v}(B) \\ u_{r|c}(B|B)\overline{v}(B) \end{bmatrix} = \begin{bmatrix} \frac{h(1-\ell)}{2-h-\ell} \\ \frac{(1-\ell)(1-h)}{2-h-\ell} \\ \frac{(1-h)(1-\ell)}{2-h-\ell} \\ \frac{\ell(1-h)}{2-h-\ell} \end{bmatrix} . \tag{128}$$

The coordinated utilities are

$$\begin{aligned} \overline{\mathbf{w}}_r &= \begin{bmatrix} \overline{w}_r(A) \\ \overline{w}_r(B) \end{bmatrix} = \begin{bmatrix} \overline{w}_{rc}(A,A) + \overline{w}_{rc}(A,B) \\ \overline{w}_{rc}(B,A) + \overline{w}_{rc}(B,B) \end{bmatrix} = \begin{bmatrix} \frac{1-\ell}{2-h-\ell} \\ \frac{1-h}{2-h-\ell} \end{bmatrix} \\ \overline{\mathbf{w}}_c &= \begin{bmatrix} \overline{w}_c(A) \\ \overline{w}_c(B) \end{bmatrix} = \begin{bmatrix} \overline{w}_{rc}(A,A) + \overline{w}_{rc}(B,A) \\ \overline{w}_{rc}(A,B) + \overline{w}_{rc}(B,B) \end{bmatrix} = \begin{bmatrix} \frac{1-\ell}{2-h-\ell} \\ \frac{1-h}{2-h-\ell} \end{bmatrix} , \end{aligned} \tag{129}$$

which are consistent with (34) and (127). It immediately follows that the coordination utility is maximized at (A, A) and the coordinated utilities are identical and are maximized at A. Thus, the coordinated decision is for each to choose A.

We next investigate the coordination properties of this network. It will be instructive first to consider extrinsic coordination, that is coordination in terms of the solution concept. From both the coordinated utilities and the payoff matrix of the steady-state game, there is no conflict, and both agents would, even from a position of narrow self-interest, choose A, since $h > \ell$. Thus, intuitively, the network is highly extrinsically coordinated.

Intrinsic coordination, however, is not a function of the solution concept. Rather, it is a function of the influence relationships, and is determined by the coordination index, which is a function of the values of h and ℓ. From (129), we have

$$H(X_i) = -\overline{w}_i(A)\log_2 \overline{w}_i(A) - \overline{w}_i(B)\log_2 \overline{w}_i(B),\ \ i = r, c\,, \tag{130}$$

and, using (128),

$$H(X_r, X_c) = \sum_{a_{rr}}\sum_{a_{cc}} \overline{w}_{rc}(a_{rr}, a_{cc}) \log_2 \overline{w}_{rc}(a_{rr}, a_{cc})\,, \tag{131}$$

yielding the coordination index

$$\mathcal{C}(X_r, X_c) = \frac{H(X_r) + H(X_c) - H(X_r, X_c)}{H(X_r, X_c)}\,. \tag{132}$$

Figure 8 displays the coordination index over the (h, ℓ) range $(1/2, 1) \times (1/2, 1)$. We see that coordination increases as $\ell \to h \to 1$. Since, by constraint, $\ell < h$, we may set $\ell = h - \epsilon$, where $\epsilon > 0$, and take the limit as $h \to 1$, yielding

$$\lim_{h\to 1} \begin{bmatrix} \overline{w}(A) \\ \overline{w}(B) \end{bmatrix} = \lim_{h\to 1} \begin{bmatrix} \frac{1-h-\epsilon}{2-2h-\epsilon} \\ \frac{1-h}{2-2h-\epsilon} \end{bmatrix} = \begin{bmatrix} \frac{\epsilon}{\epsilon} \\ \frac{0}{\epsilon} \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} . \tag{133}$$

Since this limit holds for all ϵ, it follows that $\overline{w}(A) \to 1$ and $\overline{w}(B) \to 0$ as $h \to 1$ and $\ell \to h$. Thus, the network is maximally coordinated when the probability of no error under conjecture A and the probability of error under conjecture B are both unity, in which case both agents choose A with probability one.

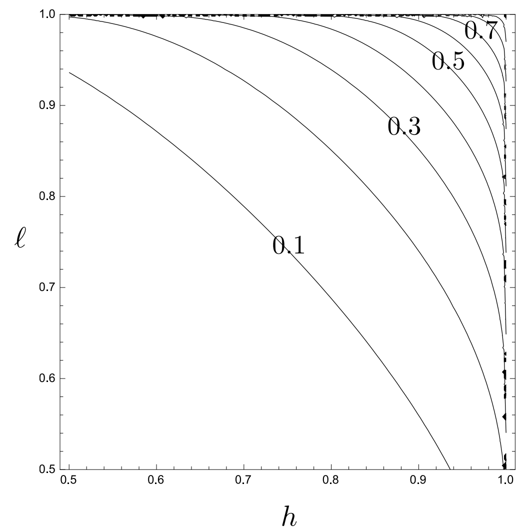

Figure 8: Coordination index for the Hi-Lo game.

This example illustrates the seamlessness of the isomorphism between conditional utility theory and Bayesian epistemology: The conditional utilities are not only isomorphic to conditional probabilities, they actually are conditional probabilities. In fact, had the Hi-Lo scenario been originally expressed via probability theory rather than game theory, this scenario would never have been viewed as a puzzle. Clearly, probability theory (and conditional utility theory via the isomorphism) is a more powerful mathematical vehicle than game theory with which to model this scenario, since it provides a natural mechanism with which to incorporate *ex ante* contextual information into the preference model.

1.2 Matching Pennies

The action set for the Matching Pennies game is $\mathcal{A}_r = \mathcal{A}_c = \{H, T\}$, and the hypothetical propositions are

$$\begin{aligned} \mathcal{H}_{i|j}(H|H)&\colon X_j \models H \Rightarrow X_i \models H \\ \mathcal{H}_{i|j}(T|H)&\colon X_j \models H \Rightarrow X_i \models T \\ \mathcal{H}_{i|j}(H|T)&\colon X_j \models T \Rightarrow X_i \models H \\ \mathcal{H}_{i|j}(T|T)&\colon X_j \models T \Rightarrow X_i \models T\,. \end{aligned} \tag{134}$$

As with Hi-Lo, we examine the relationship between $\mathcal{H}_{i|j}(H|H)$ and $\mathcal{H}_{i|j}(T|T)$. Given that H is a focal point, we may assume, by arguments similar to the development of the Hi-Lo conditional game, that the probability that X_j will correctly prefer H to T, given the conjecture H, is greater than the probability that it will mistakenly prefer T to H, given the conjecture T, and that the probability of preferring T to H, given the conjecture T, is greater than the probability of preferring H to T, given the conjecture T. Assuming symmetry, the analysis of this game is identical with that of Hi-Lo by replacing $\{A, B\}$ with $\{H, T\}$ and defining the conditional utilities

$$\begin{aligned} u_{j|i}(A|A) &= u_{i|j}(A|A) = h \\ u_{j|i}(B|A) &= u_{i|j}(B|A) = 1 - h \\ u_{j|i}(A|B) &= u_{i|j}(A|B) = 1 - t \\ u_{j|i}(B|B) &= u_{i|j}(B|B) = t\,, \end{aligned} \tag{135}$$

where $h > t > 1-t > 1-h$, where h is the conditional probability that X_j does not make a mistake under conjecture H, and t is the conditional probability that X_j makes a mistake under conjecture T.

2 Prisoner's Dilemma

2.1 Model Definition

Whereas Hi-Lo and Matching Pennies are obvious coordination scenarios where there is no overt conflict between the players, the Prisoner's Dilemma provides a more complex and challenging problem. With this game, the players, denoted $\{X_r, X_c\}$, each take actions in the sets $\mathcal{A}_r = \mathcal{A}_c = \{C, D\}$ (where C corresponds to cooperation and D corresponds to defection). The joint action set is

$$\boldsymbol{\mathcal{A}} = \{(C,C), (C,D), (D,C)), (D,D)\}\,. \tag{136}$$

The standard normal form payoff matrix for the Prisoner's Dilemma for X_r and X_c is given in Table 12. To qualify as a Prisoner's Dilemma, the payoff values must comply with the *Axelrod conditions* (Axelrod, 1984): $T > R > P > S$ and $R > (T+S)/2$. The Prisoner's Dilemma serves as a model of

Table 12: The payoff matrix for Prisoner's Dilemma: R = reward for mutual coopertion; S = sucker's payoff; T = temptation to defect; and P = punishment for mutual defection.

	X_c	
X_r	C	D
C	R, R	S, T
D	T, S	P, P

scenarios where there is an incentive to cooperate, but doing so leaves one vulnerable to exploitation. The puzzle arises because the unique individually rational choice is the Nash equilibrium (D, D) (the next worst outcome for both players), but the Pareto efficient choice (C, C) (the next best outcome for both) is the outcome that many allegedly rational people play when subjected to psychological experimentation. The problem for Bacharach is that game theory does not provide a definitive argument for either case: On the one hand, defensive play is prudent and understandable given the risk of exploitation while, on the other hand, cooperation is intuitively appealing and yields a greater payoff. This latter sentiment is motivation for a team reasoning approach to this game, which, according to Bacharach's thesis, invokes the group-identity hypothesis: "I come now to the hypothesis that perceived 'interdependence' prompts group identification .... The overwhelmingly most frequent example of a scenario in which a sense of interdependence is said to promote group identification is certainly a case of strong interdependence. It is the Prisoner's Dilemma" (Bacharach, 2006, p. 84).

We may test Bacharach's hypothesis by expressing the Prisoner's Dilemma as a cyclic conditional game of the form

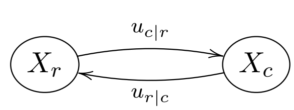

(137)

to see if deliberation alters the logic. Unlike Hi-Lo and Matching Pennies, for which a completely dissociated conditional game model is adequate, the Prisoner's Dilemma requires a fully sociated model involving the hypothetical propositions

$$H_{i|j}(a_{ir}, a_{ic}|a_{jr}, a_{jc})\colon X_j \models (a_{jr}, a_{jc}) \Rightarrow X_i \models (a_{ir}, a_{ic})\,, \tag{138}$$

where $(a_{ir}, a_{ic}), (a_{jr}, a_{jc}) \in \boldsymbol{\mathcal{A}}$ for $i, j \in \{r, c\}$, $i \neq j$. The corresponding conditional utilities are of the form $u_{i|j}(a_{ir}, a_{ic}|a_{jr}, a_{jc})$. We are guided in the valuation of these utilities by two assumptions. First, we assume the obvious: given that X_j conjectures mutual defection (D, D), then X_i will ascribe highest preference also to (D, D), rather than stubbornly insisting on a higher, but clearly unattainable, outcome. This assumption is entirely consistent with the Nash equilibrium solution. The second assumption is perhaps less obvious, but still consistent with the Prisoners Dilemma scenario: Given that X_j conjectures mutual cooperation (C, C), then X_i should also favor (C, C), the Pareto efficient outcome, even though exploiting the other perhaps within the realm of possibility, given that X_j is at least disposed toward cooperation. In other words, if the other were committed to cooperation, then it would be in one's best interest to agree. Failure to do so would eliminate any possibility of achieving a coordinated result. Invoking these assumptions, we structure the remaining entries in the conditional utility according to the categorial ordering defined by Table 12, yielding the ordinal conditional payoffs as defined in Table 13.

Table 13: The ordinal conditional preference orderings for the Prisoner's Dilemma game: (a) corresponds to $u_{r|c}$ (row player conditioned on column player) and (b) corresponds to $u_{c|r}$ (column player conditioned on row player).

$u_{r\|c}(a_{rr}, a_{rc}\|a_{cr}, a_{cc})$				
	a_{cr}, a_{cc}			
a_{rr}, a_{rc}	C, C	C, D	D, C	D, D
C, C	4	3	3	2
C, D	1	1	1	1
D, C	3	4	4	3
D, D	2	2	2	4

(a)

$u_{c\|r}(a_{cr}, a_{cc}\|a_{rr}, a_{rc})$				
	a_{rr}, a_{rc}			
a_{cr}, a_{cc}	C, C	C, D	D, C	D, D
C, C	4	3	3	2
C, D	3	4	4	3
D, C	1	1	1	1
D, D	2	2	2	4

(b)

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

To express these preferences in terms of the probability syntax, we must ascribe numerical values to each ranking. Let $\alpha, \pi \in (0, 1)$ be ordered such that

$$\pi > \alpha > 1 - \alpha > 1 - \pi\,. \tag{139}$$

Notice that this ordering corresponds to the Axelrod conditions by setting $\pi = T$, $\alpha = R$, $1-\alpha = P$, and $1 - \pi = S$, and enables defining the payoff matrix in terms of these parameters, as displayed in Table 14.

Restricting α and π to the unit interval creates no loss of generality, since utilities are invariant with respect to positive affine transformations. The paramaters (α, π) can be used to define the psychological state of the agents which, in the interest of clarity, simplicity, and consistency with the symmetric structure of the classical categorical formulation, we assume are the same for both (although not required by the theory). The psychological trait that might be associated with α is *avarice*. One with a high avarice index is focused exclusively on material gain without regard for the effect doing so has on others. We associate the psychological trait of *pragmatism* with π. A

Table 14: The payoff matrix for Prisoner's Dilemma expressed in terms of the parameters (α, π) with $\pi > \alpha > 1-\alpha > 1-\pi$.

	X_c	
X_r	C	D
C	$\alpha,\ \alpha$	$1-\pi,\ \pi$
D	$\pi,\ 1-\pi$	$1-\alpha,\ 1-\alpha$

high pragmatism index suggests that one behaves logically and deliberately and is not swayed by emotions. Certainly, other psychological traits may be attributed to these parameters, but these interpretations adequate for this development. This model assumes that pragmatism dominates avarice, which would motivate the individuals to engage in compromise to achieve at least its security level, rather than stubbornly holding out for a better outcome that entails risk.

The transition matrices $T_{r|c}$ and $T_{c|r}$ are constructed according to the orderings given in Table 13 with the modification that each entry must be divided by 2 to ensure that each column sums to unity. The resulting transition matrices are as follows:

$$\begin{aligned} T_{r|c} &= 1/2 \begin{bmatrix} \pi & \alpha & \alpha & 1-\alpha \\ 1-\pi & 1-\pi & 1-\pi & 1-\pi \\ \alpha & \pi & \pi & \alpha \\ 1-\alpha & 1-\alpha & 1-\alpha & \pi \end{bmatrix} \\ T_{c|r} &= 1/2 \begin{bmatrix} \pi & \alpha & \alpha & 1-\alpha \\ \alpha & \pi & \pi & \alpha \\ 1-\pi & 1-\pi & 1-\pi & 1-\pi \\ 1-\alpha & 1-\alpha & 1-\alpha & \pi \end{bmatrix}. \end{aligned} \tag{140}$$

Clearly, the closed-loop transition matrices $T_i = T_{i|j}T_{j|i}$ are regular for all pairs $(\alpha, \pi) \in (0,1) \times (0,1)$, $i, j \in \{r, c\}$, $i \neq j$, thus ensuring that the network will converge to steady-state unconditional utilities. According to the Markov convergence theorem, the steady-state utilities for X_i correspond to the eigenvector associated with the unique unit eigenvalue of T_i, $i = r, c$. To compute these eigenvectors, first note that T_r and T_c are related by a permutation transform defined by the matrix

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{141}$$

It is straightforward to see that

$$T_{i|j} = PT_{j|i}P\,, i, j \in \{1, 2\}\,, i \neq j \tag{142}$$

and it follows that

$$T_i = T_{i|j}T_{j|i} = PT_{j|i}PPT_{i|j}P = PT_{j|i}T_{i|j}P = PT_jP\,, \tag{143}$$

since $PP = I$. Now let $\overline{\mathbf{v}}_i$ and $\overline{\mathbf{v}}_j$ denote the eigenvectors of T_i and T_j that correspond to the unit eigenvalues of T_i and T_j, respectively. Then

$$T_i\overline{\mathbf{v}}_i = PT_jP\overline{\mathbf{v}}_i = \overline{\mathbf{v}}_i \tag{144}$$

and, multiplying all terms by P yields

$$\begin{aligned} PT_i\overline{\mathbf{v}}_i &= PPT_jP\overline{\mathbf{v}}_i = P\overline{\mathbf{v}}_i \\ &= T_jP\overline{\mathbf{v}}_i = P\overline{\mathbf{v}}_i\,. \end{aligned} \tag{145}$$

Thus, $P\overline{\mathbf{v}}_i = \overline{\mathbf{v}}_j$ and the steady-state utilities are related by the permutation matrix P. Consequently, both eigenvectors can be obtained from any column of either T_r or T_c. The first column of T_r is (with the help of *Mathematica*)

$$\overline{\mathbf{v}}_r = \begin{bmatrix} \overline{v}_r(C,C) \\ \overline{v}_r(C,D) \\ \overline{v}_r(D,C) \\ \overline{v}_r(D,D) \end{bmatrix} = \begin{bmatrix} \frac{1+\alpha^2-\alpha\pi}{6+\alpha-\alpha^2-5\pi+\pi^2} \\ \frac{1-\pi}{2} \\ \frac{\alpha+\pi+\alpha\pi-\pi^2}{4+2\alpha-2\pi} \\ \frac{1-\alpha}{3-\alpha-\pi} \end{bmatrix} \tag{146}$$

and, since $\overline{\mathbf{v}}_c = P\overline{\mathbf{v}}_r$, it follows that

$$\overline{\mathbf{v}}_c = \begin{bmatrix} \overline{v}_c(C,C) \\ \overline{v}_c(C,D) \\ \overline{v}_c(D,C) \\ \overline{v}_c(D,D) \end{bmatrix} = P\overline{\mathbf{v}}_r = \begin{bmatrix} \frac{1+\alpha^2-\alpha\pi}{6+\alpha-\alpha^2-5\pi+\pi^2} \\ \frac{\alpha+\pi+\alpha\pi-\pi^2}{4+2\alpha-2\pi} \\ \frac{1-\pi}{2} \\ \frac{1-\alpha}{3-\alpha-\pi} \end{bmatrix} \tag{147}$$

Thus, the two steady-state utilities differ by reversing the order of the outcomes (C,D) and (D,C). We now investigate several solution concepts.

2.2 Socially Coordinated Solution

Given the steady-state utilities $\overline{v}_r$ and $\overline{v}_c$ for X_r and X_c, respectively, we may compute the *ex post* payoff array as defined by Table 5, yielding the payoff matrix displayed in Table 15. The

Table 15: The payoff matrix for the steady-state Prisoner's Dilemma game.

	X_c	
X_r	C	D
C	$\overline{v}_r(C,C),\ \overline{v}_c(C,C)$	$\overline{v}_r(C,D),\ \overline{v}_c(C,D)$
D	$\overline{v}_r(D,C),\ \overline{v}_c(D,C)$	$\overline{v}_r(D,D),\ \overline{v}_c(D,D)$

classical solution concept would be to base decisions on *ex post* Nash equilibria. Although an *ex post* Nash equilibrium takes the social relationships into consideration, it is a solution concept based on individual rationality, and does not constitute a truly coordinated decision. To compute a socially coordinated solution, we must compute the coordination utility and the coordinated individual decision functions. Following (65) and (66), the steady-state coordination utility is

$$\overline{\mathbf{w}}_{rc} = \begin{bmatrix} \overline{w}_{rc}(C,C) \\ \overline{w}_{rc}(C,D) \\ \overline{w}_{rc}(D,C) \\ \overline{w}_{rc}(D,D) \end{bmatrix} = \begin{bmatrix} \frac{(-8-\alpha^3(-1+\pi)+17\pi-17\pi^2+7\pi^3-\pi^4+\alpha^2(-2-\pi+\pi^2)+\alpha(-7+15\pi-7\pi^2+\pi^3)}{4(2+\alpha-\pi)(-3+\alpha+\pi)} \\ \frac{(-8+\alpha^3(-1+\pi)+5\pi+5\pi^2-5\pi^3+\pi^4-\alpha^2\pi(1+\pi)-\alpha(-5+9\pi-7\pi^2+\pi^3)}{4(2+\alpha-\pi)(-3+\alpha+\pi)} \\ \frac{(-8+\alpha^3(-1+\pi)+5\pi+5\pi^2-5\pi^3+\pi^4-\alpha^2\pi(1+\pi)-\alpha(-5+9\pi-7\pi^2+\pi^3)}{4(2+\alpha-\pi)(-3+\alpha+\pi)} \\ \frac{(-\alpha^3(-1+\pi)+\alpha^2(6+3\pi+\pi^2)+\alpha(-7+3\pi-7\pi^2+\pi^3)-\pi(7-3\pi-3\pi^2+\pi^3)}{4(2+\alpha-\pi)(-3+\alpha+\pi)} \end{bmatrix}, \tag{148}$$

and the coordinated utilities are

$$\overline{\mathbf{w}}_r = \overline{\mathbf{w}}_c = \overline{\mathbf{w}} = \begin{bmatrix} \overline{w}(C) \\ \overline{w}(D) \end{bmatrix} = \begin{bmatrix} \frac{-8+11\pi-6\pi^2+\pi^3-\alpha^2(1+\pi)+\alpha(-1+3\pi)}{2(2+\alpha-\pi)(-3+\alpha+\pi)} \\ -\frac{4+\alpha+\pi+3\alpha\pi-4\pi^2+\pi^3-\alpha^2(3+\pi)}{2(2+\alpha-\pi)(-3+\alpha+\pi)} \end{bmatrix} . \tag{149}$$

Notice that each agent will use exactly the same coordinated decision function to make a coordinated decision.

It is instructive to examine the behavior of the network as a function of the psychological parameters (α, π). To comply with the Axelrod conditions as expressed via the payoff matrix defined by Table 14 in terms of psychological parameters, we require $1/2 < \alpha < \pi$. For the *ex post* payoff matrix also to comply with the Axelrod conditions, we require

$$\begin{aligned} \overline{v}_r(D,C) > \overline{v}_r(C,C) > \overline{v}_r(D,D) > \overline{v}_r(C,D) \\ \overline{v}_c(C,D > \overline{v}_c(C,C) > \overline{v}_c(D,D) > \overline{v}_c(D,C) \end{aligned} \tag{150}$$

and

$$\begin{aligned} \overline{v}_r(C,C) > \frac{\overline{v}_r(D,C) + \overline{v}_r(D,C)}{2} \\ \overline{v}_c(C,C) > \frac{\overline{v}_c(C,D) + \overline{v}_c(C,D)}{2} . \end{aligned} \tag{151}$$

Straightforward calculations reveal that the *ex post* payoff matrix does indeed satisfy the Axelrod conditions for all $1/2 < \alpha < \pi < 1$. In addition, the coordination utility $\overline{w}_{rc}$ is maximized with $(a_{rr}, a_{cc}) = (D, D)$, and the coordinated decision rule $\overline{w}$ is maximized at $a_{rr} = a_{cc} = D$. Thus, the coordinated decision agrees with the Nash solution.

The coordination index is computed with respect to the coordination utility and the coordinated utilities, yielding

$$\mathcal{C}(X_r, X_c) = \frac{(H(X_r) + H(X_c) - H(X_r, X_c)}{H(X_r, X_c)} , \tag{152}$$

where

$$H(X_r) = H(X_c) = -\overline{w}(C) \log_2 \overline{w}(C) - \overline{w}(D) \log_2 \overline{w}(D) \tag{153}$$

and

$$\begin{aligned} H(X_r, X_c) = -\overline{w}_{rc}(C,C) \log_2 \overline{w}_{rc}(C,C) - \overline{w}_{rc}(C,D) \log_2 \overline{w}_{rc}(C,D) \\ - \overline{w}_{rc}(D,C) \log_2 \overline{w}_{rc}(D,C) - \overline{w}_{rc}(D,D) \log_2 \overline{w}_{rc}(D,D) . \end{aligned} \tag{154}$$

Computer simulations establish that the coordination index $\mathcal{C}(X_r, X_c) \approx 0$ for all $1/2 < \alpha < \pi < 1$. This is a significant result that tells us a lot about the Prisoner's Dilemma. There is essentially no intrinsic ability for the players to form a meaningful team in order to cooperate, and there is also essentially no intrinsic ability for the players to form adversarial relationships to compete in any meaningful way. Contrary to Bacharach's argument that the Prisoner's Dilemma is a model of situations where a sense of interdependence promotes group identification, we argue instead that the Prisoner's Dilemma is a model of situations that are virtually impervious to team formation. The fact that literally thousands of papers and books have been devoted to this game over the years without resolving the issue is itself a testament to the profoundness of the social conundrum invoked by the game. Thus, the claim that the Prisoner's Dilemma is a puzzle because it defies intuitive concepts of rational behavior is really no puzzle at all. In fact, it would be puzzling if, under any meaningful notion of individually rational behavior, individuals who are really playing according to preferences defined by the payoff matrix were to deviate.

But this is not the end of the story, nor does it explain why people often do deviate from the "rational" choice. We have expressed the model in terms of the psychological parameters (α, π), where we have ascribed the attribute of avarice to α and pragmatism to π. If pragmatism really does dominate avarice (i.e., $\pi > \alpha$) for both players, then, cold, hard logic suggests that the risk of attempting to cooperate with someone who is very willing to exploit one's cooperative aspirations is unacceptably high. But what if, in reality, $\alpha > \pi$? What if both players discount the risk of exploitation and focus more heavily on increasing their payoff? To address this question, let us examine the behavior of this system in more detail.

If $1/2 < \pi < \alpha < 1$, then the payoff matrix structure displayed in Table 14 no longer meets the Axelrod conditions. In such a situation, even though the actual payoffs are defined by the payoff matrix displayed in Table 12, if $\alpha > \pi$, the players should behave, at least according to their psychological attributes, as if they are really playing Concord, a no-conflict game whose ordinal preferences are given in Table 16, for which (C, C) is the unique dominant equilibrium.[12]

Table 16: Payoff matrices for the Concord game.

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

If we simply assume that both players really intend to play a Concord game in terms of preferences but with actual payoffs determined by the Prisoner's Dilemma, then they would immediately agree to (C, C) and the game would be over. But to leap immediately to that explanation would ignore Bacharach's demand to address the puzzle and would also obviate Elster's demand for evidence of connections between rationality and behavior. Suppose the players begin their deliberations with $\pi > \alpha$ but, as time progresses, the condition $\alpha > \pi$ eventually obtains. (In Section 4.2 we briefly discuss the convergence of non-stationary Markov chains.) Incorporating this eventual preference reversal into the transition matrices given by (140) results in ordinal rankings of the conditional utilities as displayed in Table 17. This ordering generates preferences for which meaningfully coordinated behavior appears to be even less likely than with the original Axelrod-based ordering. On the one hand, if the conditioning conjecture were either (C, C) or (D, D), one's highest ranked outcome would be to exploit the other. On the other hand, if the conditioning conjecture were either to exploit or be a sucker, then one's highest-ranked outcome would be mutual cooperation. In other words, if the other were to favor any agreement, then one would favor conflict, but if the other were to favor any conflict, then one would favor agreement. These conditional preferences would appear to lead to behavior that is bizzarre, if not irrational. This certainly does not look like a Concord game. Nevertheless, it is instructive to examine the steady-state behavior as the end result of deliberation.

Figure 9 illustrates the behavior of the steady-state game for the (α, π) rectangle $[1/2, 1 \times [1/2, 1]$. The triangular section above the dashed line is the region where $\pi > \alpha > 1/2$, the parameter values that correspond to the Axelrod conditions as expressed via the psychological payoff matrix given

[12] Using the terminology provided by Bruns (2015), Concord is a 2×2 no-conflict game where moves following dominant strategies raise the payoffs by one rank and Harmony is a no-conflict game where moves following dominant strategies raise the payoff by two ranks.

Table 17: The ordinal conditional preference orderings for the Prisoner's Dilemma game with reversed parametric ordering ($\alpha > \pi$): (a) corresponds to $u_{r|c}$ and (b) corresponds to $u_{c|r}$ (compare with Table 13).

(a) $u_{r|c}(a_{rr}, a_{rc}|a_{cr}, a_{cc})$

	a_{cr}, a_{cc}			
a_{rr}, a_{rc}	C, C	C, D	D, C	D, D
C, C	3	4	4	1
C, D	2	2	2	2
D, C	4	3	3	4
D, D	1	1	1	3

(b) $u_{c|r}(a_{cr}, a_{cc}|a_{rr}, a_{rc})$

	a_{rr}, a_{rc}			
a_{cr}, a_{cc}	C, C	C, D	D, C	D, D
C, C	3	4	4	1
C, D	4	3	3	4
D, C	2	2	2	2
D, D	1	1	1	3

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

in Table 14. The (α, π) values between the dashed line and the curve labeled "PD/SD" denotes the region where the steady-state game satisfies the Axelrod conditions, even though the preference reversal has occurred. This artifact is an indication of the robustness of the Prisoner's Dilemma scenario.

The region between the curves labeled "PD/SD" and "SD/Concord" corresponds the parameter values where the *ex post* payoff matrix is such that

$$\begin{aligned}&\overline{v}_r(C,C) > \overline{v}_r(D,C) > \overline{v}_r(C,D) > \overline{v}_r(D,D)\\&\overline{v}_c(C,C > \overline{v}_c(C,D) > \overline{v}_c(D,C) > \overline{v}_c(D,D)\,.\end{aligned} \tag{155}$$

Interestingly, these conditions correspond to a Snowdrift game (also called Chicken), whose ordinal-form payoff matrix given is in Table 18. The difference between Snowdrift and Prisoner's Dilemma

Table 18: Payoff matrix for the Snowdrift game.

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

is that with Snowdrift, mutual defection is the worst outcome for both, rather than next-worst. Snowdrift has two pure-strategy Nash equilibria — the conflictive outcomes (C, D) and (D, C) and, hence, does not provide a definitive solution (although a mixed-strategy solution can be defined in terms of probabilities, but that approach is also not very satisfying).

The curve labeled "$\overline{w}(C) = \overline{w}(D)$" denotes the locus where $\overline{w}(C) = \overline{w}(D)$. Above that line, $\overline{w}(D) > \overline{w}(C)$ and the coordination utility is maximized at (D, D), and below that line, $\overline{w}(C) > \overline{w}(D)$ and the coordination utility is maximized at (C, C). Thus, above this locus, the coordinated decision is mutual defection, and is mutual cooperation below the line. For parameter values below the "SD/Concord" line, the steady-state game becomes a Concord scenario where the coordination utility is maximized at (C, C) and $\overline{w}(C) > \overline{w}(CD)$.

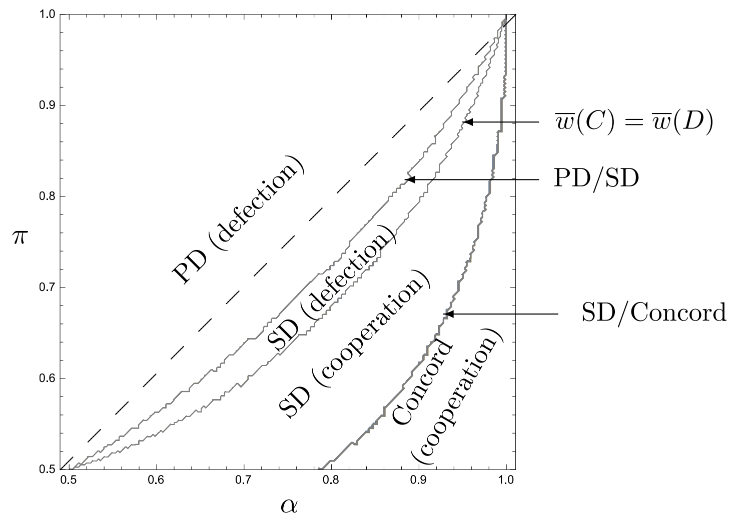

Figure 9: Parameter regions corresponding to *ex post* payoff structures for the cyclic Prisoner's Dilemma.

Thus, as the ratio α/π increases, the steady-state game starts as a Prisoner's Dilemma with mutual defection as the coordinated outcome, then transforms into a Snowdrift with mutual defection remaining the coordinated outcome, then transforms into a Snowdrift for which mutual cooperation becomes the coordinated outcome, and finally transforms into a Concord game with mutual cooperation as the obvious coordinated outcome. Thus, for α sufficiently greater than π, the players converge after deliberation to the game that their psychological parameters dictate that they should be playing.

Table 19 displays the *ex post* payoff matrices for four scenarios: (a) a Prisoner's Dilemma scenario with $(\alpha, \pi) = (0.7, 0.75)$ (mutual defection); (b) a Snowdrift scenario with $(\alpha, \pi) = (0.85, 0.75)$ (mutual defection); (c) a Snowdrift scenario with $(\alpha, \pi) = (0.9, 0.75)$ (mutual cooperation); and (d) a Concord scenario with $(\alpha, \pi) = (0.98, 0.75)$ (mutual cooperation). The corresponding coordination utilities, coordinated utilities, and coordination indices are displayed in Table 20.

Table 19: *Ex post* payoffs for the Prisoner's Dilemma game for (a) $(\alpha, \pi) = (0.70, 0.75)$, (b) $(\alpha, \pi) = (0.85, 0.75)$, (c) $(\alpha, \pi) = (0.90, 0.75)$, and (d) $(\alpha, \pi) = (0.99, 0, 75)$.

	X_c	
X_r	C	D
C	$0.32, 0.32$	$0.13, 0.36$
D	$0.36, 0.13$	$0.19, 0.19$

(a)

	X_c	
X_r	C	D
C	$(0.37, 0.37)$	$(0.13, 0.40)$
D	$(0.40, 0.13)$	$(0.11, 0.11)$

(b)

	X_c	
X_r	C	D
C	$0.27, 0.27$	$0.13, 0.41$
D	$0.41, 0.13$	$0.07, 0.07$

(c)

	X_c	
X_r	C	D
C	$0.44, 0.44$	$0.13, 0.43$
D	$0.43, 0.13$	$0.01, 0.01$

(d)

Table 20: The coordination utilities, coordinated individual decision functions, and coordination index for the Prisoner's Dilemma game for $(\alpha, \pi) = (0.70, 0.75)$, $(\alpha, \pi) = (0.85, 0.75)$, $(\alpha, \pi) = (0.90, 0.75)$, and $(\alpha, \pi) = (0.99, 0.75)$.

	(α, π)			
	(0.70, 0.75)	(0.85, 0.75)	(0.90, 0.75)	(0.99, 0.75)
$\overline{w}_{rc}(C, C)$	0.22	0.25	0.27	0.30
$\overline{w}_{rc}(C, D)$	0.23	0.24	0.25	0.27
$\overline{w}_{rc}(D, C)$	0.23	0.24	0.25	0.27
$\overline{w}_{rc}(D, D)$	0.33	0.27	0.24	0.17
$\overline{w}(C)$	0.44	0.49	0.52	0.56
$\overline{w}(D)$	0.56	0.51	0.48	0.44
$\mathcal{C}(X_r, X_c)$	0.003	0.0005	0.0006	0.002

References

A. E. Abbas. From bayes' nets to utility nets. *Proceedings of the 29th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, pages 3–12, 2009.

A. E. Abbas and R. A. Howard. Attribute dominance utility. *Decision Analysis*, 2(4):185–206, 2005.

S. Anily and A. Federgruen. Ergodicity in parametric nonstationary markov chains: An application to simulated annealing methods. *Operations Research*, 35(6):867–874, 1987.

K. J. Arrow. *Essays in the Theory of Risk Bearing.* Markham Publishing Co., Chicago, Il, 1971.

K. J. Arrow. *The Limits of Organization.* W. W. Norton & Company, New York, NY, 1974.

K. J. Arrow. Rationality of self and others in an economic system. In R. M. Hogarth and M. W. Reder, editors, *Rational Choice*. University of Chicago Press, Chicago, 1986.

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

M. Bacharach. Interactive team reasoning: A contribution to the theory of cooperation. *Research in Economics*, 23:117–147, 1999.

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

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

M. Bratman. *Faces of Intention*. Cambridge University Press, New York, 1999.

M. Bratman. *Shared Agency*. Oxford University Press, Oxford, UK, 2014.

M.E. Bratman. Shared intention. *Ethics*, 104:97–113, 1993.

B. R. Bruns. Names for games: Locating 2×2 games. *Games*, 6:496–520, 2015.

R. W. Cooper. *Coordination Games*. Cambridge University Press, Cambridge, UK, 1999.

T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley, New York, 1991.

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

B. de Finetti. La prévision: ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincaré*, 7:1–68, 1937. translated as 'Forsight. Its Logical Laws, Its Subjective Sources', in *Studies in Subjective Probability*, H. E. Kyburg Jr. and H. E. Smokler (eds.), Wiley, New York, NY, 1964, pages 93–158.

J. L. Doob. *Stochastic Processes*. John Wiley & Sons, New York, NY, 1953.

J. H. Drèze. *Essays on Economic Decisions under Uncertainty*. Cambridge University Press, Cambridge, UK, 1987.

D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge University Press, Cambridge, 2010.

J. Elster, editor. *Rational Choice*. Basil Blackwell, Oxford, UK, 1986.

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

P. C. Fishburn. *The Theory of Social Choice*. Princeton University Press, Princeton, NJ, 1973.

M. Friedman. *Price theory*. Aldine Press, Chicago, Il, 1962.

M. Gilbert. *On Social Facts*. Routledge, New York, NY, 1989.

H. Gintis. Homo ludens: Social rationality and political behavior. *Journal of Economic Behavior & Organization*, 126(PB):95–109, 2016. URL `http://EconPapers.repec.org/RePEc:eee:jeborg:v:126:y:2016:i:pb:p:95-109`.

S. Goyal. *Connections*. Princeton University Press, Princeton, NJ, 2007.

F. A. Hayek. The use of knowledge in society. *American Economic Review*, 35:519–530, 1945.

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

M.O. Jackson. *Social and Economic Networks*. Princeton University Press, Princeton, NJ, 2008.

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

E. Karni. *Decision Making Under Uncertainty*. Harvard University Press, 1985.

E. Karni and D. Schmeidler. An expected utility theory for state-dependent preferences. Working paper 48-80, the Foerder Institute of Economic Research, Tel Aviv University, Tel Aviv, 1981.

E. Karni, D. Schmeidler, and K. Vind. On state-dependent preferences and subjective probabilities. *Econometrica*, 51:1021–1032, 1983.

J. Kemeny. Fair bets and inductive probabilities. *Journal of Symbolic Logic*, 20(1):263–273, 1955.

A. Kraskov, H. Stöbauer, R. G. Andrzejak, and P. Grassberger. Hierarchical clustering based on mutual information. *ArXiv q-bio/0311039 (http://arxiv.org/abs/q-bio/0311039)*, 2003.

S. Kullback and R. A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.

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

R. S. Lehman. On conformation and rational betting. *Journal of Symbolic Logic*, 20(1):263–273, 1955.

D. K. Lewis. *Convention*. Harvard University Press, Cambridge, MA, 1969.

M. Li, J. H. Badger, X. Chen, S. Kwong, P. Kearney, and H. Zang. An information-based sequence distance and its application to whole mitochondrial genome phylogeny. *Bioinformatics*, 17(2): 149–154, 2001.

R. D. Luce and H. Raiffa. *Games and Decisions*. John Wiley, New York, 1957.

J. B. Misyak, T. A. Melkonyan, H. Zeitoun, and N. Chater. Unwritten rules: virtual bargaining underpins social interaction, culture, and society. *Trends in Cognitive Sciences*, 18:512–519, 2014.

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.

José Ortega y Gasset. Mirabeau: An essay on the nature of statesmanship, 1975. Historical Conservation Society, Manila.

F. R. Palmer. *Grammar*. Harmondsworth, Penguin, Harmondsworth, Middlesex, 1971.

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

M. Polanyi. *Personal Knowledge*. University Chicago Press, Chicago, 1962.

F. P. Ramsey. Truth and probability. In R. B. Braithwaite, editor, *The Foundations of Mathematics and Other Logical Essays*. The Humanities Press, New York, NY, 1950.

D. Ross. *Philosophy of Economics*. Palgrave Macmillan, Houndmills, Basingstoke, UK, 2014.

T. C. Schelling. *The Strategy of Conflict*. Harvard University Press, Cambridge, MA, 1960.

J. Searle. Collective intentions and actions. In P. Cohen, J. organ, and M. E. Pollack, editors, *Intentions in Communication*. Bradford Books, MIT Press, Cambridge, MA, 1990.

J. Searle. *The Construction of Social Reality*. Free Press, New York, NY, 1995.

A. K. Sen. Rational fools: A critique of the behavioral foundations of economic theory. *Philosophy & Public Affairs*, 6(4):317–344, 1977.

C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27: 379–423,623–656, 1948. This paper was printed in two parts.

Y. Shoham and K. Leyton-Brown. *Multiagent Systems*. Cambridge University Press, Cambridge, UK, 2009.

M. Shubik. *Game Theory in the Social Sciences*. MIT Press, Cambridge, MA, 1982.

M. Spivak. *Calculus on Manifolds*. W. A. Benjamin, Inc., 1965.

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

R. Sugden. Thinking as a team: Towards an explanation of nonselfish behavior. *Social Philosophy and Policy*, 10:69–89, 1993.

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

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

R. Sugden. Team reasoning and intentional cooperation for mutual benefit. *Journal of Social Ontology*, 1(1):143–166, 2014. ISSN (Online) 2196-9963, ISSN (Print) 2196-9655, DOI: 10.1515/jso-2014-0006, November 2014.