Conditional Coordination Games and Team Reasoning

Wynn C. Stirling

Abstract: This paper presents a coordination model to characterize how individuals may combine their individual interests to form a coordinated decision. Using the concept of team reasoning as a point of departure, conditional game theory is presented as a vehicle for individuals to incorporate the rationality of others into their own rationality. A model is developed to enable individuals to deliberate via the cyclic propagation of social influence expressed in the form of conditional utilities. Necessary and sufficient conditions are established for this process to converge to unconditional individual utilities that incorporate all social influence, thereby enabling a coordinated choice. The theory is illustrated via analysis of Bacharach's "three puzzles of game theory."

1 Introduction

Game theory has been at the forefront of multiagent decision theory for decades and has provided a powerful, yet simple, framework. 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 way to construct a game-theoretic model is to make only minimal assumptions about the behavior of the individuals and then investigate what can be deduced about the behavior of the collective. Any such model must focus, therefore, on only selected aspects of the notion of rationality that governs the mutual interaction of the individuals. Perhaps the least complex model of rational behavior is that individuals are motivated by narrow self-interest: one does best by maximizing individual welfare, regardless of the effect doing so has on others. Game theory comports with this model by requiring that preferences must be defined at the individual level and must be expressed in terms of utility functions that, once defined, are fixed and immutable—the are *categorical.* In fact, game theory makes a point of not requiring society to make decisions as a single entity in order to avoid what Shubik (1982, pp. 123-124) calls the "anthropomorphic trap" of building on "the shaky analogy between individual and group psychology." 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).

A consequence of the categorical model structure, however, is that it leaves unanswered, or at least not completely answered, questions that, even to one whose values are ostensibly compatible with self-interest, might be legitimately considered to be obvious when considered in a larger social context. In his posthumously published book entitled *Beyond*

Individual Choice, Bacharach argues that such paradoxes arise because solution concepts based exclusively on narrow self-interest provide an incomplete model of behavior. His thesis is that, in many social scenarios, *group identification* is more fundamental than narrow self-interest. When individuals identify with a group, the association motivates team reasoning, from which joint, rather than individual, intentions are created, leading directly to a group utility. Bacharach thereby either escapes from or falls into Shubik's anthropomorphic trap, depending on one's point of view.

This paper presents an alternative that, in a sense to be made precise, wants to have it both ways: On the one hand, rationality lies, ultimately, with individuals, and on the other, group-level coordination can emerge through the interaction of individuals who exert social influence on each other. This approach is based on the premise that, if social influence exists among the players, then it should be explicitly and directly incorporated into the mathematical expressions that define the model, rather than relegated to solution concepts that overlay the mathematical model. In other words, rather than assuming a division of labor between the way preferences are formed and the way a solution is reached, the labor should be shared. This approach is in contrast to the assertion 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." Game theory, per se, is not concerned with how one manages to express all of one's social, economic, psychological, and ethical features in to a single linear ordering over all outcomes. It simply assumes that such an exercise can be successfully accomplished and that it is adequate to motivate an individual's behavior for the social context to which the game applies.

This blanket assertion has been challenged by many, notably by Sen, who calls for a replacement. "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 [emphasis in original]" (Sen, 1977, pp. 335-336). This paper offers a way to respond to Sen's invitation by providing a mechanism for the explicit incorporation of social, as well as material, benefit into the utility structure, thereby enabling utilities to be sensitive and responsive to social relationships. As these utilities are combined, they generate an endogenous social model that leads to the emergence of a coordinated choice. Frankly put, categorical utilities are simply not up to the job. They must be replaced with a more complex utility structure that enables individuals to respond to the social context. The classical game theory model then becomes a special case of this more complex structure.

To establish the context for this development, it is important first to summarize the team reasoning approach, which represents a marked departure from the classical assumption that preferences are formed at the individual level. Next, it is important to examine current game-theoretic research regarding the modeling of complex social structures under the assumption that interests are indeed vested in the individual.

1.1 Team Reasoning

In his book, Bacharach addresses what he termed "the three puzzles of game theory." He argues that game theory does not provide a definitive answer to problems for which a uniquely best solution is obvious. His first example of such a game is Hi-Lo, where two players may choose between 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, but game theory does not provide a formal rational mechanism for choosing the one with greater value. Bacharach argues thusly: "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 other two members of Bacharach's trio of puzzles are the Matching Pennies (MP) and Prisoner's Dilemma (PD) games. The paradox he associates with MP is that, although people ar 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 PD is that, although mutual defection is the unique pure-strategy Nash equilibrium, it has been empirically established that humans often do not behave accordingly.

The approach taken by Bacharach (2006) and Sugden (2000, 2003), is to develop a theory of *team reasoning*. The central thesis of this approach is that individuals may frame the decision scenario in multiple ways—one in terms of their individual payoffs, and one in terms of the payoffs for the team with which they associate. Bacharach argues that this transformation involves a process of entification: "if and when the entification is from within, that is, the individual sees herself as part of this entified humanity, she will be led through transformations of utility and agency to seek to play her part in the profile of actions that maximizes the expected value of total happiness" (Bacharach, 2006, p. 137). This transformation thus involves framing the issues as 'What shall we do?' rather than 'What shall I do?' Bacharach argues that this concept of team reasoning will emerge endogenously when there are circumstances that tend to make people group-identify.

Essentially, Bacharach's position is that when game scenarios include social relationships, then the classical assumption that each player is categorically intent on maximizing her own benefit, regardless of the effect doing so on others, is not an adequate characterization. Rather, in a social environment, behavior is not confined to the consideration of individual intentions, but is more appropriately viewed as a manifestation of joint intentions. The issue, then, is to understand the mechanism that underlies the creation of joint intentions. Bacharach rejects the premise, held by Bratman (1993) and others, that joint intentions can be derived from individual intentions, arguing that attempting to do so leads inevitably to an infinite regress. Instead, he argues that team reasoning produces a joint intention based on common knowledge of conditional intentions (C1) and (C2) of the form "(C1) P1 intends to do x if and only if she believes that P2 will do y" (Bacharach, 2006, p. 139), with (C2) the counterpart condition.

To illustrate with the Hi-Lo scenario, Bacharach argues that, since interests of the players perfectly coincide and there are no countervailing features, an obvious group identity exists, and the group payoff is the common individual payoff, thereby yielding A, the higher value.

Bacharach succinctly summarizes this process: "The heart of this explanation is endogenous group identification in the player set, and the stimulation by group identity of team reasoning" (Bacharach, 2006, p. 144). Team reasoning then leads to joint intentions, resulting in a group utility, which can then be unequivocally maximized.

1.2 Related Work

Social psychologists and mathematicians have studied social influence network theory since the 1950s, with much of the research focusing on the organizational structure of so-called *small groups*, defined as as loosely coupled collectives of mutually interacting autonomous individuals Weick (1995). Specifically, much of the emphasis has been placed on the structure of such organizations (cf. French (1956); DeGroot (1974); Friedkin (1986); Friedkin and Johnson (1990); Friedkin (1990); Friedkin and Johnson (1997); Friedkin (1998); Friedkin and Johnson (2011)). The basic model is that an individual's socially adjusted utility is a convex combination of its own categorical utility and a weighted sum of the categorical utilities of those agents who influence it. Karni and Schmeidler (1981) introduce the concept of state-dependent preferences where the decision maker's preferences are modulated by the state of nature. Hu and Shapley (2003a,b) apply a command structure to model player interactions by simple games. Grabisch and Rusinowska (2010, 2011, 2013) and Förster et al. (2013), build on this structure to study social influence in the context of simple "yes" or "no" voting games by introducing influence functions that allow agents to modify their opinions as a result of the inclinations of others.

To account for social relationships that exist among the members of a collective, several innovations have been applied to classical game theory. Behavioral game theory (Bolton and Ockenfels, 2005; Fehr and Schmidt, 1999; Camerer, 2003; Camerer et al., 2004b,a; Henrich et al., 2005) is a response to the desire to introduce psychological realism and social influence into game theory by incorporating notions such as fairness and reciprocity into the utilities, in addition to considerations of material benefit. The closely related field of psychological game theory (Geanakoplos et al., 1989; Dufwenberg and Kirchsteiger, 2004; Colman, 2003; Battigalli and Dufwenberg, 2009; Gilboa and Schmeidler, 1988) also employs utilities that account for beliefs as well as actions and takes into consideration belief-dependent motivations such as guilt aversion, reciprocity, regret, and shame. Regardless of the issues used to define the preferences, however, these approaches to game theory differ from the approach taken in this paper in that they use unconditional linear preference orderings (i.e., categorical utilities) over the outcomes and, therefore, the solution concepts used by classical game theory continue to apply.

2 Conditional Game Theory

In one sense, conditional game theory is similar to Bacharach's, in that it also employs a form of conditioning. However, whereas conditioning *à la* Bacharach is with respect to intentions, conditioning in the conditional game-theoretic sense is with respect to the preferences that players assign to consequences. Rather than expressing preference with categorical utilities that are fixed and immutable once the game is engaged, preferences are expressed with conditional utilities that enable players to modulate their preferences as a function of specific social situations. This approach is motivated by Arrow's recognition that a narrow view of individually rational behavior is limited to a specific set of conditions.

> 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).

If the concept of self-interest expands such that the rationality of others becomes part of one's own rationality, then the utilities used to characterize one's preferences should also be expanded in order to incorporate the preferences of others into one's own preferences. For a collective to function as a team, each member ought to have an expanded concept of preference, but it should still be an individual preference. An alternative concept is for the group to have a goal, and for each member of the group to do its part in order to achieve the goal. However, there is no guarantee that the players all have the same concept of how to achieve the goal. Suppose a manager assigns two employees with complementary skills to perform some task that requires their complete cooperation with the expectation that both will do their parts as they collectively pursue the goal. But even if both have the same group-level goal, that does not mean that they automatically agree regarding how each individual should behave. Suppose one member unilaterally chooses a particular approach without taking into consideration how that will affect the productivity of the other team member, with the likely result that the goal will not be achieved, or at least will not be achieved as well as it could be if they were to coordinate their actions in such a way that achieving the task emerges as a consequence of their coordinated behavior. In other words, coordination requires both individuals to incorporate the rationality of the other into their individual concepts of rational behavior.

This paper presents an extension of classical noncooperative game theory that is based on three premises. First, true coordination is an emergent phenomenon that arises, if it does, as a result the social influence that the players exert on each other. Even the existence of a strong exogenously imposed group identity does not guarantee that the players will function effectively in accordance with the group-level objective. Second, social influence is not easily expressed in terms of material benefit. A more natural way to express social influence is for the players to evaluate their material benefit in response to social influence. Third, the extension should not be a departure from the basic game-theoretic philosophy that preferences should be defined at the individual level. Thus, classical noncooperative game theory should be a special case of this extended theory.

2.1 Network Games

Conditional game theory, as introduced by Stirling (2012), is an extension to classical noncooperative game theory, and serves as a model for an *influence network:* a finite collective of individuals who exert social influence on each other. It models such a collective as a graph whose vertices are the individuals, denoted $\{X_1, \ldots, X_n\}$, and whose edges represent the influence linkages between individuals. An edge is *directed* (denoted with the arrow symbol "$\rightarrow$") if the linkage is unidirectional: $X_j \rightarrow X_i$ means that X_j directly influences X_i. 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. A *path* from X_j to X_i is a sequence of directed edges from X_j to X_i, denoted $X_j \mapsto X_i$. A path is a *cycle*, or *closed path*, if $X_j \mapsto X_j$. A graph is said to be a *directed acyclic graph* if all edges are directed and there are no cycles.

Following the standard game-theoretic setup, each X_i possesses an action set $\mathcal{A}_i$, $i = 1, \ldots, n$. The Cartesian product set $\boldsymbol{\mathcal{A}} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ comprises the set of *profiles* of the form $\mathbf{a} = (a_1, \ldots, a_n)$, where $a_i \in \mathcal{A}_i$. A profile $\mathbf{a}_i \in \boldsymbol{\mathcal{A}}$ is a *conjecture* for X_i if it is under consideration for X_i as the profile that will or should be activated. Symbolically, the statement that X_i conjectures $\mathbf{a}_i$ is written $X_i \models \mathbf{a}_i$.

The effect of influence is modeled by a utility that governs X_i's preferences over conjectures for it given conjectures for those who influence it. For $\mathrm{pa}\,(X_i) = \{X_{i_1}, \ldots, X_{i_{p_i}}\}$, let $\mathrm{pa}\,(i) = \{i_1, \ldots, i_{p_i}\}$ denote the indices corresponding to the elements of $\mathrm{pa}\,(X_i)$. Also, let $\mathrm{cpa}\,(X_i) = \{\mathbf{a}_{i_1}, \ldots, \mathbf{a}_{i_{p_i}} \colon X_{i_k} \models \mathbf{a}_{i_k}, k = 1, \ldots, p_i\}$ denote the set of *conditioning conjectures* for the parents of X_i. Let $u_{i|\mathrm{pa}(i)}[\cdot|\,\mathrm{cpa}\,(X_i)] \colon \boldsymbol{\mathcal{A}} \rightarrow \mathbb{R}$ denote a *conditional utility*. The expression $u_{i|\mathrm{pa}(i)}[\mathbf{a}_i|\,\mathrm{cpa}\,(X_i)]$ is the consequent for $X_i \models \mathbf{a}_i$ of the hypothetical proposition whose antecedent is the joint conjecture $\{X_{i_1}, \ldots, X_{i_{p_i}}\} \models \{\mathbf{a}_{i_1} \ldots, \mathbf{a}_{i_{p_i}}\}$. This expanded utility structure results in an extension to classical noncooperative game theory, defined as follows.

Definition 1 *A normal form* network game *comprises a) a collective* $\{X_1, \ldots, X_n\}$*; b) a finite action set* $\mathcal{A}_i$ *for each* X_i*; and c) a family of conditional utilities*

$$\{u_{i|\mathrm{pa}(i)}[\cdot|\,\mathrm{cpa}\,(X_i)] \colon \boldsymbol{\mathcal{A}} \rightarrow \mathbb{R}, \forall\, \mathrm{cpa}\,(X_i) \in \boldsymbol{\mathcal{A}}^{p_i}, i = 1, \ldots, n\}\,.$$

If an individual is not influenced by any others, then it possesses a categorical utility $u_i \colon \boldsymbol{\mathcal{A}} \rightarrow \mathbb{R}$. Whereas a conditional utility is designed to respond to context—the social environment of those who influence it—a categorial utility is context independent. In general, it will be the case that a network will comprise a mixture of individuals with conditional and categorical utilities. Thus, if $\mathrm{pa}\,(i) = \varnothing$, then $u_{i|\mathrm{pa}(i)} = u_i$.

2.2 Coordination

As influence propagates through the network, social relationships are created as the individuals interact, resulting in an emergent social structure that governs the organizational behavior of the collective. One of the issues associated with this emergent structure is the innate ability of the individuals to coordinate in the sense of serving as components of a societal whole. Thus, in contrast to classical noncooperative game theory, which makes no assumptions regarding group-level preferences or choices, the social model that emerges as a result of social influence may indeed possess some, perhaps abstract, systematic notion of emergent group-level behavior.

A natural way to account for the social relationships that emerge as a result of influence propagation is to define a *coordination function* that characterizes all of the social relationships the arise as the members of a society interact.[1] Indeed, according to the Oxford English Dictionary, "[To coordinate is] 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). Whereas cooperation and conflict involve the compatibility of agents with respect to their individual objectives and do not require even the existence of group-level objectives (e.g., the classical PD model), coordination involves the compatibility of individual objectives (the parts) with respect to group-level objectives (the whole).

There are many possible ways to define a coordination function. For such a function to be meaningful, however, it should possess the following properties. First, it should be *endogenous*, in that it emerges as a result of the social interrelationships that exist among the members of the collective—it should not be exogenously imposed. Second, it should be *comprehensive*, in that it provides a means of assessing the degree or coordination among the individuals, accounts for the seriousness of disputes, and enables the evaluation of the possibilites for compromise. Both endogeny and comprehensiveness conditions are met by requiring the coordination function to be of the form

$$u_{1:n}(\mathbf{a}_1; \cdots ; \mathbf{a}_n) = f\big[u_{i|\mathrm{pa}(i)}[\mathbf{a}_i |\, \mathrm{cpa}\,(X_i)],\, i = 1, \ldots, n\big]\,, \tag{1}$$

where $\mathrm{cpa}\,(X_i) = \{\mathbf{a}_{i_1}, \ldots, \mathbf{a}_{i_{p_i}}\} \subset \{\mathbf{a}_1, \ldots, \mathbf{a}_n\}$ and f is to be determined. This structure guarantees that coordination arises from, and only from, the social relationships that emerge as a result of social interrelationships. By requiring that coordination be a function of the conjectures of all individuals, it enables the evaluation of coordination when the individuals conjecture different profiles.

The third criterion is that coordination should be *socially coherent* in that it does not permit pathological or manifestly undemocratic behavior. Specifically, no individual should have absolute dictatorial power over the behavior of the society, and no individual should be the victim of systematic and categorical discrimination by the rest of society (i.e., "a dictator turned upside-down" —Fishburn (1973, p. 211)). These two pathological behaviors are termed *tyranny* and *subjugation*, respectively.

Suppose there exists $\mathbf{a} \in \boldsymbol{\mathcal{A}}$ such that, if

$$u_i(\mathbf{a}) > u_i(\mathbf{a}')\ \forall \mathbf{a}' \in \boldsymbol{\mathcal{A}} \backslash \{\mathbf{a}\}\,, \tag{2}$$

then

$$u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_{i-1}, \mathbf{a}, \mathbf{a}_{i+1}, \ldots, \mathbf{a}_n) > u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_{i-1}, \mathbf{a}', \mathbf{a}_{i+1}, \ldots, \mathbf{a}_n) \tag{3}$$

for all $\mathbf{a}' \in \boldsymbol{\mathcal{A}} \backslash \{\mathbf{a}\}$ and for all $\{\mathbf{a}_1, \ldots, \mathbf{a}_{i-1}, \mathbf{a}_{i+1}, \ldots, \mathbf{a}_n\} \in \boldsymbol{\mathcal{A}}^{n-1}$. Then X_i is a *tyrant.* If X_i is a tyrant, then X_i's most preferred profile maximizes coordination, regardless of the profiles preferred by others. Although X_i cannot force others to change their preferences, it overrides them to enforce whatever organizational behavior it favors.

[1] In Stirling (2012), the coordination function is called the *concordance function*. Concordance, however, typically conotes "a state in which things agree and do not conflict with each other" (Merriam-Webster, 2015). Coordination, by contrast, is a neutral concept, and captures the notions of both cooperation and conflict (e.g., members of a team would coordinate cooperatively, but opponents in an athletic contest or military engagement would coordinate conflictively).

Alternatively, suppose there exists $\mathbf{a} \in \boldsymbol{\mathcal{A}}$ such that, if

$$u_i(\mathbf{a}) > u_i(\mathbf{a}') \ \forall \mathbf{a}' \in \boldsymbol{\mathcal{A}} \backslash \{\mathbf{a}\}\,, \tag{4}$$

then

$$u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_{i-1}, \mathbf{a}, \mathbf{a}_{i+1}, \ldots, \mathbf{a}_n) < u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_{i-1}, \mathbf{a}', \mathbf{a}_{i+1}, \ldots, \mathbf{a}_n) \tag{5}$$

for all $\mathbf{a}' \in \boldsymbol{\mathcal{A}} \backslash \{\mathbf{a}\}$ and for all $\{\mathbf{a}_1, \ldots, \mathbf{a}_{i-1}, \mathbf{a}_{i+1}, \ldots, \mathbf{a}_n\} \in \boldsymbol{\mathcal{A}}^{n-1}$. X_i is then *subjugated*. If X_i is subjugated, then any joint conjecture that includes X_i's most preferred profile will be less coordinated than all profiles that do not include X_i's most preferred profile.

Definition 2 *A coordination function is* socially coherent *if neither tyranny nor subjugation is possible.*

The key result of conditional game theory is that there is a unique mathematical structure that complies with the endogeny, comprehensiveness, and coherence criteria. Results for acyclic networks are considered first, and Section 3 then generalizes to the cyclic case.

Theorem 1 (Coordination) *The coordination function for an acyclic network game is socially coherent if, and only if, a) the individual utilities are expressed as* utility mass functions, *that is,*

$$u_{i|\mathrm{pa}(i)}[\mathbf{a}_i | \,\mathrm{cpa}\,(X_i)] \geq 0\,,$$

and

$$\sum_{\mathbf{a} \in \boldsymbol{\mathcal{A}}} u_{i|\mathrm{pa}(i)}[\mathbf{a}_i | \,\mathrm{cpa}\,(X_i)] = 1$$

for all $\mathrm{cpa}\,(X_i) \in \boldsymbol{\mathcal{A}}^{p_i}$, $i = 1, \ldots, n$; *and b) the coordination function is of the form*

$$u_{1:n}(\mathbf{a}_1; \cdots; \mathbf{a}_n) = \prod_{i=1}^{n} u_{i|\mathrm{pa}(i)}[\mathbf{a}_i | \,\mathrm{cpa}\,(X_i)]\,, \tag{6}$$

where $\mathrm{cpa}\,(X_i) = \{\mathbf{a}_{i_1}, \ldots, \mathbf{a}_{i_{p_i}}\} \subset \{\mathbf{a}_1, \ldots, \mathbf{a}_n\}$.

Proof This result is established by first appealing to the Dutch Book theorem (de Finetti, 1937; Ramsey, 1950) and its converse (Lehman, 1955; Kemeny, 1955), which together establish that a sure loss (win) is impossible if, and only if, the gambler's beliefs conform to the probability axioms.[2] The Dutch Book theorem is connected to the present problem by the following lemma, which is proven in Appendix A.

Lemma 1 *Subjugation is isomorphic to sure loss and tyranny is isomorphic to sure win.*

By applying the isomorphism to the Dutch Book theorem, it follows that neither subjugation nor tyranny can occur if and only if all utility functions possess the mathematical structure of probability mass functions and are combined and manipulated according to the mathematical syntax of probability theory. The resulting acyclic network is then isomorphic to a Bayesian

[2] A Dutch Book a gambling scenario where the gambler's outcome is always less (greater) than the entry fee—a sure loss (win). This is an alternative to the proof provided in Stirling (2012), which relies on the associativity equation (Cox, 1946; Jaynes, 2003).

network (Pearl, 1988; Cowell et al., 1999; Jensen, 2001). A fundamental property of Bayesian networks is the *Markov condition:* nondescendent nonparents of a vertex have no influence on the vertex, given the state of its parent vertices (Cozman, 2000). Consequently, the joint probability distribution of a collective of random vectors is the product of the probability distributions of each of its members as conditioned on the realizations of its parents. By the isomorphism, the coordination function of a collective of agents is the product of the utilities of the agents as conditioned by the preferences of its parents. In other words, the coordination function is defined by the chain rule of probability. □

Conditional utilities may be interpreted analogously to conditional probability mass functions and combined according to probability syntax. For example, given two discrete random vectors Y_1 and Y_2 taking values $\mathbf{y}_1, \mathbf{y}_2 \in \mathbb{R}^n$, $p_{2|1}(\mathbf{y}_2|\mathbf{y}_1)$ is the conditional probability of $Y_2 = \mathbf{y}_2$, given that $Y_1 = \mathbf{y}_2$, and the joint probability mass function is the product: $p_{12}(\mathbf{y}_1, \mathbf{y}_2) = p_1(\mathbf{y}_1)p_{2|1}(\mathbf{y}_2|\mathbf{y}_1)$. The analogous game-theoretic interpretation is that, given two individuals possessing action profiles $\mathbf{a}_1, \mathbf{a}_2 \in \boldsymbol{\mathcal{A}}$, $u_{2|1}(\mathbf{a}_2|\mathbf{a}_1)$ is the conditional utility of $X_2 \models \mathbf{a}_2$ given that $X_1 \models \mathbf{a}_2$, and the coordination function is the product: $u_{12}(\mathbf{a}_1, \mathbf{a}_2) = u_1(\mathbf{a}_1)u_{2|1}(\mathbf{a}_2|\mathbf{a}_1)$. Thus, the coordination function is syntactically equivalent to a joint probability mass function for a set of discrete random vectors. Just as a joint probability mass function captures all of the innate statistical dependencies among a set of random vectors, the coordination function captures all of the innate social relationships among a set of networked individuals.

Since the coordination function is analogous to the probability mass function of a collective of random vectors, all of the probabilistic syntax may be applied. In particular, the *ex post* utility for each X_i may be computed by marginalization, yielding

$$v_i(\mathbf{a}_i) = \sum_{\sim \mathbf{a}_i} u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_n)\,, \tag{7}$$

where the notation $\sum_{\sim \mathbf{a}_i}$ means the sum is taken over all arguments of $u_{1:n}$ except $\mathbf{a}_i$. These *ex post* utilities are now unconditional, and standard game-theoretic solutions, such as Nash equilibrium, may be applied. However, such solution concepts are not necessarily coordinated, since they respond to, and only to, the narrow self-interest of the players.

Thus, to obtain a truly coordinated decision, more must be done. Fortunately, there is much more information contained in the coordination function since it is a function of all joint conjectures $(\mathbf{a}_1, \ldots, \mathbf{a}_n\}$—a function of n^2 variables. It is therefore possible to define a group-level utility as a function of n variables as follows. Let a_{ii} denote the ith element of $\mathbf{a}_i = (a_{i1}, \ldots, a_{in})$, where $a_{ij} \in \mathcal{A}_j$; that is, a_{ij} is the jth component of X_i's conjecture profile. The group utility is then the group-level marginal of the coordination function, given by

$$w_{1:n}(a_{11}, \ldots, a_{nn}) = \sum_{\sim a_{11}} \cdots \sum_{\sim a_{nn}} u_{1:n}(a_{11}, \ldots, a_{1n}; \cdots ; a_{n1}, \ldots, a_{nn})\,. \tag{8}$$

The *maximally coordinated decision* is the profile that maximizes $w_{1:n}$, yielding

$$\mathbf{a}^* = \arg\max_{\mathbf{a} \in \boldsymbol{\mathcal{A}}} w_{1:n}(\mathbf{a})\,. \tag{9}$$

Summarizing, the coordination process involves two steps. First, the *ex ante* conditional utilities are combined via (6) to form a coordination function, and then a group-level utility is derived by computing the group-level marginal utility, accounting for all of the social relationships that emerge as the *ax ante* conditional utilities propagate through the network.

3 Network Games with Influence Cycles

A necessary condition for Theorem 1 to apply is that the network is acyclic; that is, if $X_i \mapsto X_j$, then $X_j \not\mapsto X_i$. This condition restricts the application of the theorem to hierarchical networks. Many important networks, however, involve cycles. In particular, coordination between individuals is difficult if they cannot deliberate and reach a joint resolution, or at least find a compromise. Consider the two-agent cyclic network

$$\textcircled{X_1} \rightleftarrows \textcircled{X_2} . \tag{10}$$

X_1 influences X_2, who influences X_1, and so on, *ad infinitum.* At first glance, this situation may appear to be a manifestation of the concern expressed by Bacharach (2006) that attempting to derive joint intentions from individual intentions can lead to an infinite regress with no resolution. This is a legitimate concern that must be addressed. The issue devolves around the question of convergence: Does the cycle of influence propagation oscillate endlessly with each individual repeatedly changing its mind, or does it converge in the sense that each individual ultimately possesses an individual *steady-state* utility? Thus, to extend conditional game theory to cyclic networks, it is necessary to establish conditions for convergence. If the utilities do converge, their steady-state values represent the end point of bilateral deliberations with each individual taking the preferences of the other into consideration.

To study the convergence properties of a cyclic network, first recall that a network that satisfies the coherence condition is isomorphic to a collective of random vectors. Thus, the probability syntax may be used to study cyclic networks. The key observation in this regard is that the network defined by (10) is a Markov chain, and the convergence of such probabilistic entities is well studied, which is now reviewed.

3.1 Matrix Form Dynamics Model

A k-member closed path is called a *simple k-cycle* if every vertex in the path has exactly one incoming edge and one outgoing edge.[3] Consider the simple k-cycle illustrated in Figure 1. Insight can be gained by viewing this cyclic network as a sequence of influence operations that evolve in time. Let δ denote the time interval required for information to transit between individuals. At time $t = 0$, when the cycle is initiated, each X_i will be in some utility state $v_i(\mathbf{a}_i, 0)$. Without loss of generality, analysis my begin with X_1. For $t = 0, \delta, 2\delta, \cdots$, consider the equivalent dynamic network displayed in Figure 2.

Let us first consider the segment $\textcircled{X_1} \xrightarrow[t=0]{u_{2|1}} \textcircled{X_2}$. At $t = \delta$, X_1's conditional preferences are received by X_2, and the resulting coordination function for $\{X_1, X_2\}$ is, following

[3]Although this discussion is restricted to the analysis of simple cycles, the theory easily generates to more complicated cycles.

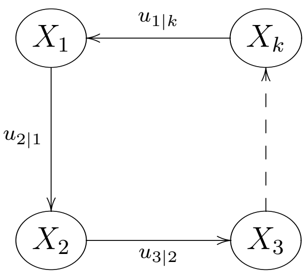

Figure 1: A simple k cycle.

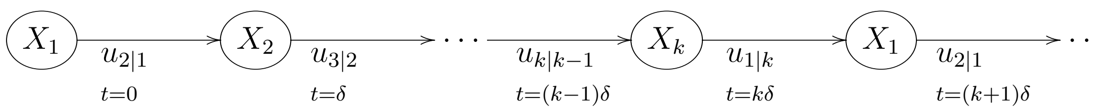

Figure 2: An equivalent dynamic network.

(6),

$$u_{12}(\mathbf{a}_1, \mathbf{a}_2, \delta) = u_{2|1}(\mathbf{a}_2|\mathbf{a}_1)v_1(\mathbf{a}_1, 0)\,, \tag{11}$$

from which X_2's *ex post* marginal may be computed via (7):

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

Now consider the next segment 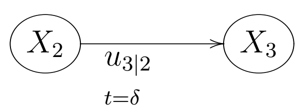

. At $t = 2\delta$ the coordination function for $\{X_2, X_3\}$ is

$$u_{23}(\mathbf{a}_2, \mathbf{a}_3, 2\delta) = u_{3|2}(\mathbf{a}_3|\mathbf{a}_2)v_2(\mathbf{a}_2, \delta)\,, \tag{13}$$

and X_3's *ex post* marginal is

$$v_3(\mathbf{a}_3, 2\delta) = \sum_{\mathbf{a}_2} u_{23}(\mathbf{a}_2, \mathbf{a}_3, 2\delta)\,. \tag{14}$$

Continuing this process for $t = 3\delta$, $t = 4\delta$, $\cdots$, may be facilitated by reformulating the issue in terms of matrix notation. Let us denote the elements of $\boldsymbol{\mathcal{A}}$ as

$$\boldsymbol{\mathcal{A}} = \{\mathbf{z}_1, \ldots, \mathbf{z}_N\} \tag{15}$$

and define the *utility mass vector* at time t by

$$\mathbf{v}_i(t) = \begin{bmatrix} v_i(\mathbf{z}_1, t) \\ \vdots \\ v_i(\mathbf{z}_N, t) \end{bmatrix}\,. \tag{16}$$

Next, define the *state-to-state transition matrix*

$$T_{i+1|i} = \begin{bmatrix} u_{i+1|i}(\mathbf{z}_1|\mathbf{z}_1) & \cdots & u_{i+1|i}(\mathbf{z}_1|\mathbf{z}_N) \\ \vdots & & \vdots \\ u_{i+1|i}(\mathbf{z}_N|\mathbf{z}_1) & \cdots & u_{i+1|i}(\mathbf{z}_N|\mathbf{z}_N) \end{bmatrix}\,. \tag{17}$$

With this notation, the operations defined by (11) and (12) can be combined in to the single expression

$$\mathbf{v}_2(\delta) = T_{1|2}\mathbf{v}_1(0)\,, \tag{18}$$

and (13) and (14) can be combined to yield

$$\mathbf{v}_3(2\delta) = T_{3|2}\mathbf{v}_1(\delta)\,. \tag{19}$$

Figure 3 displays the k-cycle with the linkages represented by the transition matrices. Tracing 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:

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

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

$$\mathbf{v}_{i+k}[(k+1)\delta] = T_{i+k|i+k-1}T_{i+k-1|i+k-2}\cdots T_{i+2|i+1}T_{i+1|i}\mathbf{v}_i(0) \tag{20}$$

or, since all indices are incremented mod k.

$$\mathbf{v}_i[(k+1)\delta)] = T_{i|i+k-1}T_{i+k-1|i+k-2}\cdots T_{i+2|i+1}T_{i+1|i}\mathbf{v}_i(0)\,. \tag{21}$$

Now define the *closed-loop transition matrix*

$$T_i = T_{i|i+k-1}T_{i+k-1|i+k-2}\cdots T_{i+2|i+1}T_{i+1|i}\,. \tag{22}$$

Also, it is convenient to express time in units equal to the interval $k\delta$, resulting in expressing (21) as

$$\mathbf{v}_i(1) = T_i\mathbf{v}_i(0) \tag{23}$$

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,

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

The key issue devolves around the behavior of T_i^t as $t \to \infty$, which motivates a study of the convergence properties of this matrix.

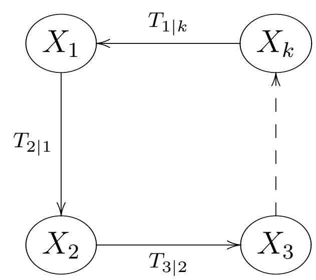

Figure 3: A k-cycle expressed in terms of transition linkages.

3.2 Convergence of Closed-Loop Transition Matrices

Definition 3 *Let $T_i = [t_{jk}]$ be a square matrix with t_{jk} denoting the entry in the jth row and kth column. T_i is nonnegative, denoted $T \not< 0$, if $t_{jk} \not< 0 \,\forall\, j,k$. T_i is positive, denoted $T_i \geq 0$, if $t_{jk} \not< 0 \,\forall\, j,k$ and $t_{jk} > 0$ for at least one element. T_i is strictly positive, denoted $T_i > 0$, if $t_{jk} > 0 \,\forall\, j,k$.*

Definition 4 *A square matrix T_i is a* regular transition matrix *if T_i^k is strictly positive for some finite integer k and each column sums to unity.*

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

Theorem 2 (Frobeneius-Peron) *If a square matrix T_i^k is strictly positive for some finite integer k, then T_i has a unique largest eigenvalue with positive eigenvector. In particular, the unique largest eigenvalue of a regular transition matrix is unity.*

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

Theorem 3 (Markov Convergence) *If T_i is a regular transition matrix, there exists a unique mass vector $\mathbf{u}_i$ such that*[4]

- $T_i \mathbf{u}_i = \mathbf{u}_i$,
- $\overline{T}_i = \lim_{t\to\infty} T_i^t = \begin{bmatrix} \mathbf{u}_i & \cdots & \mathbf{u}_i \end{bmatrix}$, *and*
- $\mathbf{u}_i = \overline{T}_i \mathbf{v}_i$ *for every mass vector* $\mathbf{v}_i$.

For proofs of these theorems, see Gantmacher (1959); Doob (1953).

The content of the convergence theorem is that the cyclic network illustrated in Figure 1 with linkages expressed by *ex ante* conditional utilities can be replaced by a network with no edges, as displayed in Figure 4, with each vertex possessing a constant unconditional utility of the form

$$\overline{\mathbf{u}}_i = \begin{bmatrix} \overline{u}_i(\mathbf{z}_1) \\ \vdots \\ \overline{u}_i(\mathbf{z}_N) \end{bmatrix} = \lim_{t\to\infty} \mathbf{v}_i(t)\,. \tag{24}$$

[4]A mass vector has all nonnegative entries that sum to unity.

Mathematically, these utilities are structured the same as *ex ante* categorical utilities, but they are now *ex post* utilities as the end result of social influence. In other words, the *ex ante* conditional utilities $u_{i|i-1}(\mathbf{a}_i|\mathbf{a}_{i-1})$ are replaced by steady-state unconditional utilities $\overline{u}_i(\mathbf{a}_i)$. The coordination function (6) thus becomes

$$\overline{u}_{1:n}(\mathbf{a}_1; \cdots ; \mathbf{a}_n) = \prod_{i=1}^{n} \overline{u}_i(\mathbf{a}_i) \tag{25}$$

and, from (8), the coordination utility is

$$\overline{w}_{1:n}(a_{11}, \ldots, a_{nn}) = \sum_{\sim a_{11}} \cdots \sum_{\sim a_{nn}} \prod_{i=1}^{n} \overline{u}_i(\mathbf{a}_i)\,. \tag{26}$$

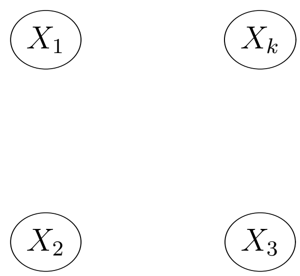

Figure 4: A converged cyclic network.

4 A Bilateral Coordination Model

Bilateral coordination occurs when the members of a two-agent network seek a solution that meets the objectives of the network as well as the individuals. Developing coordination models for Bacharach's puzzles, however, first requires investigating the possible mechanisms by which influence is propagated. Second, understanding the significance of the time dynamics that underly the convergence process is critical.

One possible mechanism is the existence of an explicit communication linkage between the players, whereby they may transmit information back and forth. Such a situation is tantamount to a collaboration problem where the individuals iteratively seek to achieve a common goal, such as two employees who are tasked by their manager to accomplish some task. Let A and B denote two distinct approaches and suppose, for example, that X_1 initially prefers approach A, and X_2 initially prefers approach B and the group is therefore likely, at the outset, to settle for an inferior outcome. The conditional utilities, however, are defined independently of the initial preferences of either. In fact, it may turn out that both X_1 and X_2 initially prefer B. But as they begin to take into consideration the conditional preferences of the other, the eventually (and perhaps rapidly) converge to a consensus.

Suppose, however, that there is no direct communication between the two individuals. The lack of an explicit means of transmitting information, however, does not imply the non-existence of a social connection. It is possible, by methods similar to those involved in Bacharach's concept of group identification, for each player to construct a model of how it would respond to the preferences of the other. This type of reasoning is consistent with Schelling's concept of focal points: "Most situations ... provide some clue for coordinating behavior, some focal point for each person's expectation of what the other expects him to expect to be expected to do" (Schelling, 1960, p. 57): The mechanism advanced by Schelling involves an imaginative process of introspection: "In the pure-coordination game, 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). Thus, each player could engage in a *virtual deliberation*, or thought experiment, whereby they define the model and establish convergence.

Regardless of whether influence is propagated through direct communication or by some imaginative process of virtual deliberation, the Markov convergence theorem involves a time sequence of influence propagation. In a sequential context where information is transmitted, it cannot be done instantaneously. The time increment δ between transmission and reception must be non-zero, and it need not be constant. All that matters is that enough communication transpires to obtain a reliable approximation to the steady-state values. In the case of implicit influence propagation, each participant performs its own convergence exercise, which is simply a matter of deliberation, and the cyclic model, along with the convergence theory, is really nothing more than a mathematical characterization of such a deliberative process.

4.1 Sociation

In the interest of complete generality, conditional game theory has been developed under the assumption that conditioning will be with respect to $\mathrm{cpa}\,(X_i) = \{\mathbf{a}_{i_1}, \ldots, \mathbf{a}_{i_{p_i}}\}$, the entire profiles of all members of $\mathrm{pa}\,(X_i) = \{X_{i_1}, \ldots, X_{i_{p_i}}\}$. It may turn out, however, that the conditioning will not be with respect to all components of each $\mathbf{a}_{i_k}$. To the extent that the influence of others depends only on the conjectured actions others, rather on on the entire conjecture profile of others, the conditioning conjectures are *conjecture dissociated*. The conditioning conjectures are *completely conjecture dissociated* if the influence depends on, and only on, the conjectured actions of those who influence it.

It may also turn out that an individual's utility may not be a function of every component of $\mathbf{a}_i$. 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*, in which case the coordination function is of the form

$$u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\mathbf{a}_{i_1}, \ldots, \mathbf{a}_{i_{p_i}}) = u_{i|\mathrm{pa}(i)}(a_i|a_{i_1}, \ldots, a_{i_{p_i}})\,, \tag{27}$$

that is, X_i's utility depends on, and only on, its own action, conditioned on the actions of, and only of, all who influence it.

A network is *uniformly dissociated* if all of its members are completely dissociated. For a uniformly dissociated network, the coordination function (6) becomes

$$u_{1:n}(a_1, \ldots, a_n) = \prod_{i=1}^{n} u_{i|\mathrm{pa}(i)}(a_i|a_{i_1}, \ldots, a_{i_{p_i}}) \,. \tag{28}$$

Furthermore, in the uniformly dissociated case, the group utility coincides with the coordination function, since the latter is now a function of only n variables. Thus,

$$w_{1:n}(a_1, \ldots, a_n) = \prod_{i=1}^{n} u_{i|\mathrm{pa}(i)}(a_i|a_{i_1}, \ldots, a_{i_{p_i}}) \,. \tag{29}$$

Also, the individual marginal utilities (7) become, in the uniformly dissociate case,

$$v_i(a_i) = \sum_{\sim a_i} w_{1:n}(a_1, \ldots, a_n) \, i = 1, \ldots, n \,. \tag{30}$$

Although uniformly dissociated networks constitute a small subset of general networks, they still serve as useful behavioral models for many interesting networks. If a network is uniformly dissociated, then each player's preferences are with respect to its own choices only, rather than the entire profile, and the conditioning is with respect to the conjectured choices only of its parents. Under the special case where all players have categorical utilities, a uniformly dissociated model would have preferences over its own action set only, and the game would reduce to the trivial case, comprising a set of n completely decoupled agents whose behavior does not affect each other. But when influence linkages influence the players, a uniformly dissociated game is far from trivial. It will be convenient to use this reduced structure to study the Hi-Lo and MP games, but it will be necessary to revert to the completely sociated model when discussing the PD game.

4.2 Hi-Lo

The payoff matrix for the Hi-Lo game is displayed in Table 1. To reiterate, the paradox defined by Friedman (1962, p. 13)Bacharach (2006) is that, although (A, A) is clearly the unique best outcome, classical game theory does not provide a formal theoretical reason for eliminating (B, B) since, as a Nash equilibrium, that outcome enjoys equal status with (A, A) as a rational group choice. Under the team reasoning approach, both players agents

Table 1: The Hi-Lo game.

	X_2	
X_1	A	B
A	$2, 2$	$0, 0$
B	$0, 0$	$1, 1$

undergo an agent transformation from the status of a completely autonomous individual to a member of a team, and also undergo a utility transformation from an individual utility to the group utility. The payoffs, however, remain unchanged. The essential difference between this formulation and the classical non-cooperative game formulation is that the utilities are no longer correspond to individual payoffs. The payoff matrix for the team reasoning formulation is given in Table 2.

Table 2: Payoff matrix for the team reasoning Hi-Lo game.

	X_1, X_2
A, A	$2, 2$
B, B	$1, 1$

Consider modeling this scenario as a conditional game, where each player's preferences are modulated by the preferences of the other. Since this game is symmetric, the utilities of both players are the same. Before proceeding, observe that the team reasoning approach required a utility transformation, but that did not change the game, since the players are still subject to the original payoffs as defined by Table 1. Similarly, this approach also requires that the utilities be such that, although not defined as the payoffs, they will reflect the values associated with the payoffs depending on the social circumstances. Whatever the decision, however, the payoffs will be distributed according to the entries in Table 1.

The goal is to define conditional utilities of the form $u_{i|j}$, for $i, j \in \{1, 2\}$, with $i \neq j$. Under the uniformly dissociated assumption, each player's preferences over its own action set are conditioned on the other player's preferences over its action set. If the completely sociated model were employed, then each player would need to define preferences over the joint action set conditioned by the preferences of the other, also over the joint action set (which would require the specification of sixteen conditional utility functions with much of the information being redundant). Because of the simple structure of this game, however, the uniformly dissociated formulation retains the essence of the completely sociated case, and results in a simpler and more intuitive solution (requiring the specification of only four utility functions).

Given that X_i conjectures A, X_j would place its utility mass on A according to the strength of belief that X_i conjectures A. X_j would then place the remaining utility mass on B. Thus, the conditional utilities are of the form

$$u_{j|i}(A|A) = u_{i|j}(A|A) = h \tag{31}$$
$$u_{j|i}(B|A) = u_{i|j}(B|A) = 1 - h \tag{32}$$
$$u_{j|i}(A|B) = u_{i|j}(A|B) = 1 - \ell \tag{33}$$
$$u_{j|i}(B|B) = u_{i|j}(B|B) = \ell\,, \tag{34}$$

where $0 < \ell < h < 1$, which is consistent with the rational assumption that X_j more strongly believes that X_i will conjecture A rather than B. The resulting cyclic network is

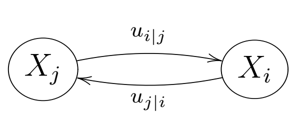

. (35)

To study the convergence properties of the Hi-Lo game, it is necessary first express the player's utilities using matrix notation and define the transition matrix. Following (16), the utility mass vectors are given by

$$\mathbf{v}_i(t) = \begin{bmatrix} v_i(A,t) \\ v_i(B,t) \end{bmatrix} \tag{36}$$

$$\mathbf{v}_j(t) = \begin{bmatrix} v_j(A,t) \\ v_j(B,t) \end{bmatrix} \tag{37}$$

Also, following (17), the state-to-state transition matrices are

$$T_{i|j} = T_{j|i} = \begin{bmatrix} h & 1-\ell \\ 1-h & \ell \end{bmatrix} \tag{38}$$

and, following (22), the closed-loop transition matrix is

$$T_i = T_j = T_{i|j}T_{j|i} = \begin{bmatrix} h^2 + (1-h)(1-\ell) & (h+\ell)(1-\ell) \\ (h+\ell)(1-h) & (1-h)(1-\ell) + \ell^2 \end{bmatrix}. \tag{39}$$

The closed-loop transition matrix is clearly regular for $h, \ell \in (0,1)$, thus a steady state exists and the cyclic network defined by (10) converges to the steady state network

$$\textcircled{X_1} \qquad\qquad \textcircled{X_2}, \tag{40}$$

with unconditional utilities given by the steady-state utility vectors $\overline{\mathbf{u}}_i$ and $\overline{\mathbf{u}}_j$, which, by the convergence theorem, are equal to the eigenvector of T_i corresponding to the unit eigenvalue. Thus,

$$\overline{\mathbf{u}}_i = \overline{\mathbf{u}}_j = \begin{bmatrix} \overline{u}_j(h) \\ \overline{u}_j(\ell) \end{bmatrix} = \begin{bmatrix} \frac{1-\ell}{2-h-\ell} \\ \frac{1-h}{2-h-\ell} \end{bmatrix}. \tag{41}$$

By the structure of the game, $0 < \ell < h < 1$ and, since the two components of the eigenvector have the same denominator, it is clear that $\overline{u}_j(h) > \overline{u}_j(\ell)$, This analysis definitively establishes that A is the unique best solution, thereby resolving the paradox that game theory does not provide a definitive formal mechanism for establishing that (A, A) is preferred to (B, B). In fact, it is not necessary that the game be symmetric. The result holds if X_1 and X_2 have different values of h and ℓ, so long as they satisfy the constraint that $0 < \ell_i < h_i$, $i \in \{1, 2\}$.

The limiting closed-loop game is equivalent to a classical noncooperative game with payoff matrix given by Table 3, for which $(1-\ell, 1-\ell)$ is the unique pure-strategy Nash equilibrium. It must be stressed, however, that this limiting game is not a Hi-Lo game, since the payoffs for non-agreement are nonzero. However, the actual payoffs received by the players are in accord with the values displayed in Table 1.

4.3 Matching Pennies

The payoff matrix for MP is displayed in Table 4. The only difference between this game and Hi-Lo is that the rewards are the same for both of the pure-strategy Nash equilibria.

Table 3: The payoff matrix for the limiting closed-loop Hi-Lo game.

	X_2	
X_1	A	B
A	$1-\ell, 1-\ell$	$1-\ell, 1-h$
B	$1-h, 1-\ell$	$1-h, 1-h$

Bacharach's puzzle arises because game theory does not identify heads as the salient outcome. A closed-loop formulation of this game, therefore, will have exactly the same structure as the Hi-Lo game, except that the conditional utilities will express preferences that correspond to salience. Thus, heads will be associated with the Hi-Lo outcome A, and tails with B, yielding the limiting closed-loop solution of both players calling heads.

Table 4: Payoff matrix for Matching Pennies with H = heads, T= tails.

	X_2	
X_1	H	T
H	$1, 1$	$0, 0$
T	$0, 0$	$1, 1$

4.4 Prisoner's Dilemma

The payoff matrix for PD for X_1 (the row player) and X_2 (the column player), is given in Table 5, with C corresponding to cooperation and D to defection. To qualify as a PD, the payoff values must comply with what may be termed the *Axelrod conditions* (Axelrod, 1984): $T > R > P > S$ and $R > (T + S)/2$.

Table 5: 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_2	
X_1	C	D
C	R, R	S, T
D	T, S	P, P

A team reasoning approach to this game would invoke the group-identity hypothesis: "I come now to the hypothesis that perceived 'interdependence' prompts group identification" (Bacharach, 2006, p. 84). He then argues that "The overwhelmingly most frequent example of a scenario in which a sense of interdependence is said to prompt group identification is currently a case of strong interdependence. It is the Prisoner's Dilemma" (Bacharach, 2006, p. 84).

As an alternative to relying on a notion of group identification, PD can be addressed using individual, albeit conditional, utilities. Recasting it as a network game results in the following structure.

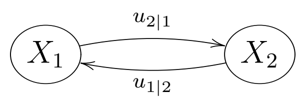

(42)

The goal is to identify appropriate structures for $u_{1|2}$ and $u_{2|1}$. Unlike Hi-Lo, however, the non-agreement components of PD are antisymmetric, thus this game is considerably more complex in structure than Hi-Lo. To capture this complexity in a conditional game it is necessary to use the completely sociated model.

The traditional game theory model expresses the preferences of the individuals in terms of, and only of, material benefit. This model, however, does not account for the social dispositions of the players, who may possess traits that influence their behavior in addition to the desire for material benefit. There are many possible ways to introduce such dispositions, but for the purposes of this development, attention is focused on two traits that are consistent with the usual PD scenario. Let $\alpha_i \in [0,1]$ denote an *assertiveness* index: if $\alpha_i \approx 1$, then X_i is confident or decisive but if $\alpha_i \approx 0$, then X_i is diffident or reserved. Also, let π_i denote an *pragmatism* index: if $\pi_i \approx 1$, then X_i is practical or efficient. If $\pi_i \approx 0$, then X_i unrealistic or impractical.

Of course, other behavioral traits could be used to characterize a social situation such as a PD, but the assertiveness and pragmatism descriptors are deontologically neutral and uncorrelated (knowing the assertiveness of an individual tells little about her pragmatism).[5] However, even if one objects to these descriptors, they at least serve the purpose of illustrating how social parameters can be inserted into a game.

In the PD context, pragmatism is the more dominant behavioral trait, followed by assertiveness, and both of them should be high, resulting in the following ordering:

$$\pi_i > \alpha_i > 1 - \alpha_i > 1 - \pi_i\,. \tag{43}$$

To define the conditional utilities in terms of these social parameters, it is necessary to examine the consequent ordering for X_i, given the antecedent by X_j, $i, j \in \{1,2\}$, $i \neq j$. Consider X_1's responses to conjectures for X_2. Clearly, the worst (least pragmatic) response in all cases is (C, D). Now let us consider X_1's response given that X_2 conjectures (C, C). X_1's pragmatic response is also to cooperate, yielding (C, C). The choice of which is the next-best and next-worst out of (D, C) and (D, D) is somewhat arbitrary, given the antecedent. However, since X_1's payoff is greater for (D, C) than for (D, D) the more assertive response is (D, C). If X_2 were to conjecture (C, D), X_1 would be conflicted with regard to what is its best response. However, since the individual payoff is greater for (D, C) than for (C, C), the pragmatic response is (D, C), since it has a higher individual payoff, with (C, C) next-best. Next, if X_2 were to conjecture (D, C), the best (most pragmatic) response for X_1 is obviously (D, C), which would maximize X_1's individual welfare. In keeping with prior

[5]In Stirling (2012), an acyclic version of PD is developed. In that formulation, the behavioral traits are specified in terms of *cooperation* and *exploitation* indices.

reasoning, the next-best response would be (C, C). Finally, if X_2 were to conjecture (D, D), X_1's pragmatic response would clearly be (D, D), and next-best response would be (D, C). By a similar analysis, the conditional utilities for X_2 given X_1's conjectures may be defined. The resulting conditional preference orderings for X_1 and X_2 are displayed in Table 6(a) for $u_{1|2}$ and in Table 6(b) for $u_{2|1}$. The columns correspond to X_j's conjecture, (a_{j1}, a_{j2}), and the corresponding row entries define X_i's resulting conditional utility for its conjecture (a_{i1}, a_{i2}). The transition matrices $T_{1|2}$ and $T_{2|1}$ comprise the entries in these tables with the modification that each entry must be divided by 2 to ensure that each column sums to unity.

Table 6: The conditional preference orderings for the Prisoner's Dilemma game: (a) corresponds to $u_{1|2}$ (row player) and (b) corresponds to $u_{2|1}$ (column player).

$u_{1|2}(a_{11}, a_{12}|a_{21}, a_{22})$

X_1	X_2			
	a_{21}, a_{22}			
a_{11}, a_{12}	C, C	C, D	D, C	D, D
C, C	π_1	α_1	α_1	$1-\alpha_1$
C, D	$1-\pi_1$	$1-\pi_1$	$1-\pi_1$	$1-\pi_1$
D, C	α_1	π_1	π_1	α_1
D, D	$1-\alpha_1$	$1-\alpha_1$	$1-\alpha_1$	π_1

(a)

$u_{2|1}(a_{21}, a_{22}|a_{11}, a_{12})$

X_2	X_1			
	a_{11}, a_{12}			
a_{21}, a_{22}	C, C	C, D	D, C	D, D
C, C	π_2	α_2	α_2	$1-\alpha_2$
C, D	α_2	π_2	π_2	α_2
D, C	$1-\pi_2$	$1-\pi_2$	$1-\pi_2$	$1-\pi_2$
D, D	$1-\alpha_2$	$1-\alpha_2$	$1-\alpha_2$	π_2

(a)

Clearly, the closed-loop transition matrices $T_i = T_{i|j}T_{j|i}$ are regular for all pairs $(\alpha_i, \pi_i) \in (0, 1) \times (0, 1)$, $i, j \in \{1, 2\}$, $i \neq j$, thus ensuring that the network will converge to steady-state unconditional utilities. In the interest of maintaining symmetry between the two players, attention is restricted to $\alpha_1 = \alpha_2 = \alpha$ and $\pi_1 = \pi_2 = \pi$.

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, $i = 1, 2$. To compute these eigenvectors, first note that the structure of T_1 and T_2 are related by a permutation transform defined by the permutation matrix

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} . \tag{44}$$

It is straightforward to see that

$$T_{i|j} = PT_{j|i}P\,, i, j \in \{1, 2\}\,, i \neq j \tag{45}$$

and it follows that

$$T_i = T_{i|j}T_{j|i} = PT_{j|i}PPT_{i|j}P = PT_{j|i}T_{i|j}P = PT_jP\,, \tag{46}$$

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

$$T_i\mathbf{u}_i = PT_jP\mathbf{u}_i = \mathbf{u}_i \tag{47}$$

and, multiplying all terms by P yields

$$\begin{aligned} PT_i\mathbf{u}_i &= PPT_jP\mathbf{u}_i = P\mathbf{u}_i \\ &= T_jP\mathbf{u}_i = P\mathbf{u}_i \,. \end{aligned} . \tag{48}$$

Thus, $P\mathbf{u}_i = \mathbf{u}_j$ and the steady-state utilities are related by the permutation matrix P. Consequently, both eigenvectors can be obtained from any column of either T_1 or T_2. The first column of T_1 is

$$\overline{\mathbf{u}}_1 = \begin{bmatrix} \overline{u}_1(C,C) \\ \overline{u}_1(C,D) \\ \overline{u}_1(D,C) \\ \overline{u}_1(D,D) \end{bmatrix} . \tag{49}$$

Since $\mathbf{u}_2 = P\mathbf{u}_1$, it follows that

$$\overline{\mathbf{u}}_2 = \begin{bmatrix} \overline{u}_2(C,C) \\ \overline{u}_2(C,D) \\ \overline{u}_2(D,C) \\ \overline{u}_2(D,D) \end{bmatrix} = P\overline{\mathbf{u}}_1 = \begin{bmatrix} \overline{u}_1(C,C) \\ \overline{u}_1(D,C) \\ \overline{u}_1(C,D) \\ \overline{u}_1(D,D) \end{bmatrix} . \tag{50}$$

The limiting coordination function (26) is

$$\overline{w}_{12}(a_{11}, a_{22}) = \sum_{a_{12}} \sum_{a_{21}} \overline{u}_1(a_{11}, a_{12})\overline{u}_2(a_{21}, a_{22}) \,, \tag{51}$$

and the maximally coordinated decision is

$$(a_{11}^*, a_{22}^*) = \max_{a_{11}, a_{22}} \overline{w}_{12}(a_{11}, a_{22}) \,. \tag{52}$$

Computer simulations confirm that

$$\max\{w_{12}(C,C), w_{12}(D,D)\} \geq \max\{w_{12}(C,D), w_{12}(C,D)\} \tag{53}$$

for all values of $(\alpha, \pi) \in (0,1) \times (0,1)$. Thus, the maximally coordinated decision is either (C,C) or (D,D). Figure 5 displays a contour plot of the ratio $w_{12}(C,C)/w_{12}(D,D)$. Points above the line of demarcation where the $ratio = 1$ correspond to (D,D), mutual defection, as the maximally coordinated decision, with points below the line corresponding to (C,C), mutual cooperation, as the maximally coordinated choice. This result ensures complete coordination: The individuals will either mutually cooperate or mutually defect. Thus, mixed outcomes are impossible when both players have the same assertiveness and pragmatism indices.

Expressing preferences in terms of behavioral traits provides a resolution to Bacharach's complaint with PD. Whereas the categorical utilities provided in Table 5 are with respect to material benefit without regard to social context, the conditional preferences provided in Table 6 ascribe material benefit in a social context. Of course, once the decisions are made, the players receive the rewards given in Table 5.

Notice that it is also possible to recast the PD game in terms of the social parameters by setting $T = \pi$, $R = \alpha$, $P = 1 - \alpha$, and $S = 1 - \pi$, as displayed in Table 7. The (α, π) pairs

Table 7: The Prisoner's Dilemma payoff matrix expressed in terms of the assertiveness and pragmatism indices, where $\pi > \alpha > 1/2$.

	X_2	
X_1	C	D
C	α, α	$1-\pi, \pi$
D	$\pi, 1-\pi$	$1-\alpha, 1-\alpha$

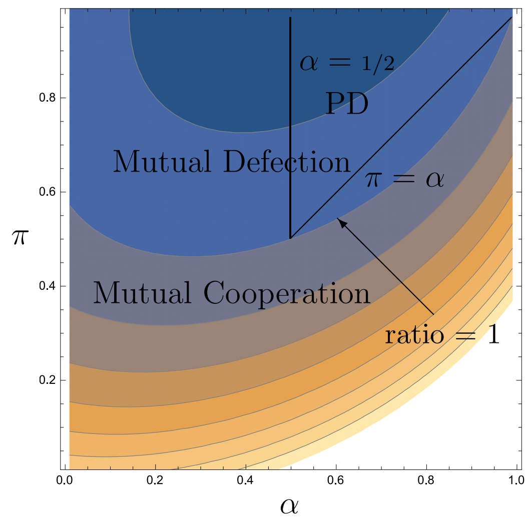

Figure 5: Contour plots of the ratio $w_{12}(C,C)/w_{12}(D,D)$ for Prisoner's Dilemma.

that meet the Axelrod criterion (43) are illustrated in Figure 5, and lie in the region bounded by the lines $\alpha = 1/2$ and $\pi = \alpha$, marked PD. Thus, the coordination regions defined by the line of demarcation include PD as a special case.

This theory extends, of course, to the asymmetric case, where the individuals have different assertiveness and pragmatism indices, but the behavior illustrated in Figure 5 does not apply, since there are four parameters in play, rather than just two. In such cases, all four outcomes are possible, depending on the alignment of interests among the players.

5 Conclusion

The coordination game model serves as an alternative to the notion of team reasoning advanced by Bacharach (2006) and Sugden (2000, 2003) as a way to account for social context. Rather than relying on concepts of group identity and manifesting such phenomena through variable frames and utility transformations, social relationship are incorporated by modeling the collective as a network with individuals as vertices and influence linkages as edges in the form of conditional preferences. The preference model employed by standard game theory then becomes a special case—a network with no edges whose socially isolated vertices are characterized by categorical utilities. This approach appears to be philosophically compatible to the notion of *team-centered choice* suggested by Ross (2014), who argues that any such theory should be technically integrated with traditional game theory.

A fundamental distinction between this approach and team reasoning is that, whereas team reasoning defines intensions at the group level, this approach retains the fundamentally individualistic foundations of game theory, but has enlarged the notion of individual rationality to accommodate the interests of others into one's own rationality. Group preferences, if they exist, then emerge endogenously as a consequence of social interaction as social influence propagates throughout a network of individuals who are connected by explicit social linkages.

Conditional game theory, as originally introduced by Stirling (2012), was restricted to acyclical influence structures. This paper extends the theory to account for influence cycles. This extension enables the development of conditional coordination games, where the players exert bilateral social influence on each other in order to achieve a coordinated decision. To the extent that the *ex post* interests are socially aligned, the resulting joint decision corresponds to emergent systematic group behavior. It is important to appreciate that such group-level behavior is not exogenously imposed. It emerges as a result of, and only of, the interest that is shared by the members of the network. Furthermore, the sense in which interest is shared is neutral—it applies to cooperative agents for whom the emergent systematic group-level behavior corresponds to teamwork, and it applies equally well to conflictive agents for whom the emergent systematic group-level behavior is to work in opposition, as would be the case for athletic contests and military engagements.

Formulating a decision problem such as the Prisoner's Dilemma results in a structure that is more complex than the original payoff matrix and reliance on Nash equilibria as the solution concept. But it is not more complex than it needs to be in order to capture the social aspects of the issue, which are completely lacking with the classical formulation. As expressed Palmer (1971, p. 184), "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."

A Proof of the Isomorphism Lemma

Lemma 1 *Subjugation is isomorphic to sure loss and tyranny is isomorphic to sure win.*

Proof Establishing this result requires first proving that the categorical and conditional utilities are order isomorphic to marginal and conditional probability mass functions. Without loss of generality, attention is restricted to a two-agent collective $\{X_1, X_2\}$ defined over the product set $\boldsymbol{\mathcal{A}} = \mathcal{A}_1 \times \mathcal{A}_2$, with X_1 possessing a categorical utility $u_1\colon \boldsymbol{\mathcal{A}} \to \mathbb{R}$ and X_2 possessing a family of conditional utilities $\{u_{2|1}(\cdot|\mathbf{a}_1)\colon \boldsymbol{\mathcal{A}} \to \mathbb{R} \ \forall\ \mathbf{a}_1 \in \boldsymbol{\mathcal{A}}\}$. Let Ω_1 and Ω_2 be arbitrary sets of propositions of distinct elements with cardinalities equal to the cardinalities of $\mathcal{A}_1$ and $\mathcal{A}_2$, respectively, and let $\boldsymbol{\Omega} = \Omega_1 \times \Omega_2$. Let $b_1\colon \boldsymbol{\Omega} \to \mathbb{R}$ be a belief function such that, for $\boldsymbol{\omega}, \boldsymbol{\omega}' \in \boldsymbol{\Omega}$, $b_1(\boldsymbol{\omega}) \geq b_1(\boldsymbol{\omega}')$ means that the belief that $\boldsymbol{\omega}$ will be realized is at least as great as the belief that $\boldsymbol{\omega}'$ will be realized. Also, let $\{b_{2|1}(\cdot|\boldsymbol{\omega}_1)\colon \boldsymbol{\Omega} \to \mathbb{R}\ \forall \boldsymbol{\omega}_1 \in \boldsymbol{\Omega}\}$ be a family of conditional belief functions over $\boldsymbol{\Omega}$ such that $b_{2|1}(\boldsymbol{\omega}_2|\boldsymbol{\omega}_1) \geq b_{2|1}(\boldsymbol{\omega}_2'|\boldsymbol{\omega}_1)$ means that the belief that $\boldsymbol{\omega}_2$ is realized is at least a great as the belief that $\boldsymbol{\omega}_2'$ is realized, given that

$\boldsymbol{\omega}_1$ is realized. Now let g: $\boldsymbol{\mathcal{A}} \to \boldsymbol{\Omega}$ be a bijective mapping of the form g: $\mathbf{a} \mapsto \boldsymbol{\omega}$ such that

$$u_1(\mathbf{a}) = b_1[g(\mathbf{a})] = b_1(\boldsymbol{\omega}) \tag{54}$$

and

$$u_{2|1}(\mathbf{a}_2|\mathbf{a}_1) = b_{2|1}[g(\mathbf{a}_2)|g(\mathbf{a}_1)] = b_{2|1}(\boldsymbol{\omega}_2|\boldsymbol{\omega}_1) \tag{55}$$

for all $\mathbf{a}_1 \in \boldsymbol{\mathcal{A}}$. It is immediate that this mapping establishes the structural equivalence of the benefit criterion regarding $\boldsymbol{\mathcal{A}}$ and the belief criterion regarding $\boldsymbol{\Omega}$. Furthermore, the conjecture $\mathbf{a}_1$ and the realization assertion $\boldsymbol{\omega}_1$ are both antecedents of hypothetical propositions whose consequents are $u_{2|1}(\mathbf{a}_2|\mathbf{a}_1)$ and $b_{2|1}[g(\mathbf{a}_2)|g(\mathbf{a}_1)]$, respectively. This establishes the order isomorphism.

To establish the isomorphism between subjugation and sure loss, Let b_{12}: $\boldsymbol{\Omega}^2 \to \mathbb{R}$ be a belief function such that $b_{12}(\boldsymbol{\omega}_1, \boldsymbol{\omega}_2) \geq b_{12}(\boldsymbol{\omega}_1', \boldsymbol{\omega}_2')$ means that the belief that $(\boldsymbol{\omega}_1, \boldsymbol{\omega}_2)$ is realized is at least as great as the belief that $(\boldsymbol{\omega}_1', \boldsymbol{\omega}_2')$ is realized.

Suppose there exists $\boldsymbol{\omega} \in \boldsymbol{\Omega}$ such that

$$b_1(\boldsymbol{\omega}) > b_1(\boldsymbol{\omega}') \;\forall \boldsymbol{\omega}' \in \boldsymbol{\Omega} \backslash \{\boldsymbol{\omega}\}\,, \tag{56}$$

but

$$b_{12}(\boldsymbol{\omega}, \boldsymbol{\omega}_2) < b_{12}(\boldsymbol{\omega}', \boldsymbol{\omega}_2) \tag{57}$$

for all $\boldsymbol{\omega}' \in \boldsymbol{\Omega} \backslash \{\boldsymbol{\omega}\}$ and for all $\boldsymbol{\omega}_2 \in \boldsymbol{\Omega}$. Thus, even though $\boldsymbol{\omega}$ is the most strongly believed event, the belief regarding the realization of any joint event for which $\boldsymbol{\omega}$ is the realization is weaker than the belief regarding the realization of the corresponding joint event with any other $\boldsymbol{\omega}'$ the realization.

If, on the basis of (56) one were to enter a lottery to earn \$1 if $\boldsymbol{\omega}$ is realized, a fair entry fee would be $q_1 > \frac{1}{2}$. On the other hand, if, on the basis of (57), one were to earn \$1 if $\boldsymbol{\omega}$ is not realized, then a fair entry fee would be $q_2 > \frac{1}{2}$. By combining these two bets 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 (4) and (56) and by (5) and (57) are identical.

A sure win can occur as follows. On the basis of (56), a fair entry fee for a bet that $\boldsymbol{\omega}$ is not realized would be $q_1 < \frac{1}{2}$. If the positons of $\boldsymbol{\omega}$ and $\boldsymbol{\omega}'$ in (57) are reversed, then a fair entry fee for a bet that $\boldsymbol{\omega}$ sill be realized is $q_2 < \frac{1}{2}$. Thus, one would win exactly \$1 with an entry fee of $q_1 + q_2 < 1$—a sure win. Thus, a sure loss is isomorphic to subjugation and a sure win is isomorphic to tyranny. □

References

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

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

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

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

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

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

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

A. M. Colman. Cooperation, psychological game theory, and limitations of rationality in social interaction. *Behavioral and Brain Sciences*, 26:139–198, 2003.

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

R. T. Cox. Probability, frequency, and reasonable expectation. *American Journal of Physics*, 14:1–13, 1946.

F. G. Cozman. Credal networks. *Artificial Intelligence*, 120:199–233, 2000.

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.

M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69: 118–121, 1974.

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

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

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. Förster, M. Grabisch, and A. Rusinowska. Anonymous social influence. *Games and Economic Behavior*, 82:621–635, 2013.

J. R. P. French. A formal theory of social power. *The Psychology Review*, 63:181–194, 1956.

N. E. Friedkin. A formal theory of social power. *Journal of Mathematical Sociology*, 12: 103–136, 1986.

N. E. Friedkin. Social networks in structural equation models. *Social Psychology Quarterly*, 53:316–328, 1990.

N. E. Friedkin. *A Structural Theory of Social Influence*. Cambridge University Press, Cambridge, UK, 1998.

N. E. Friedkin and E. C. Johnson. Social influence and opinions. *Journal of Mathematical Sociology*, 15:193–205, 1990.

N. E. Friedkin and E. C. Johnson. Social positions in influence networks. *Social networks*, 19:209–222, 1997.

N. E. Friedkin and E. C. Johnson. *Social Influence Network Theory*. Cambridge University Press, Cambridge, UK, 2011.

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

F. Gantmacher. *Applications of the Theory of Matrices*. Chelsea Publishing, 1959.

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

I. Gilboa and D. Schmeidler. Information dependent games. *Economics Letters*, 27:215–221, 1988.

M. Grabisch and A. Rusinowska. A model of influence in social networks. *Theory and Decision*, 69(1):69–96, 2010.

M. Grabisch and A. Rusinowska. Influence functions, followers and command games. *Games and Economic Behavior*, 72:123–138, 2011.

M. Grabisch and A. Rusinowska. A model of influence based on aggregation. *Mathematical Social Sciences*, 66:316–330, 2013.

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

X. Hu and L. S. Shapley. On authority distributions in organizations: Controls. *Games and Economic Behavior*, 45:153–170, 2003a.

X. Hu and L. S. Shapley. On authority distributions in organizations: Equilibrium. *Games and Economic Behavior*, 45:132–152, 2003b.

E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge, UK, 2003.

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

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.

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

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

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

Merriam-Webster. Online dictionary. `http://www.merriam-webster.com/dictionary`, 2015.

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.

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

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

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.

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

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

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

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

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

K. E. Weick. *Sensemaking in organizations*. Sage Publications, Thousand Oaks, CA, 1995.