Conditional Game Theory as a Model for Coordinated Decision Making

Wynn C. Stirling[1] and Luca Tummolini[2]
[1]*Department of Electrical and Computing Engineering, Brigham Young University, Provo, Utah, U.S.A.*
[2]*Institute of Cognitive Sciences and Technologies, Italian National Research Council, Rome, Italy*

Keywords: Game Theory, Coordination, Social Influence, Social Utility, Network Theory.

Abstract: Standard game theory is founded on the premise that choices in interactive decision situations are strategically rational—best reactions to the expected actions of others. However, when studying groups whose members are responsive to one another's interests, a relevant notion of behavior is for them to coordinate in the pursuit of coherent group behavior. Conditional game theory provides a framework that facilitates the study of coordinated rational behavior of human social networks and the synthesis of artificial social influence networks. This framework comprises three elements: a socialization model to characterize the way individual preferences are defined in a social context; a diffusion model to define the way individual preferences propagate through the network to create an emergent social structure; and a deduction model that establishes the structure of coordinated individual choices.

1 INTRODUCTION

Coordinated decision making is one of the fundamental attributes of intelligent behavior. Indeed, the word *intelligent* comes from the Latin roots *inter* (between) + *legĕre* (to choose). Accordingly, much effort has been devoted to defining what it means for a choice to be "rational." And appending the modifier "coordinated" adds a level of complexity that moves beyond the hypothesis that each of the individual decision makers should behave as if it were solving a constrained maximization problem without overt regard for the welfare of others.

Coordination, as used in this paper, has a precise meaning, as expressed by 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).

Coordination is a principle of behavior on a parallel with, but different from, performance. Individuals *perform*; the group *coordinates*. Performance deals with operational measures of efficiency and effectiveness of individual behavior in terms of individual payoffs. Coordination, however, is an attribute of organizational structure regarding how the members of a group function together.

An important class of collectives comprises entities that possess the ability to respond to the social influence that they exert on one another. 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 tennis players who exert conflictive influence on each other. 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 tennis players (an anti-team?) coordinate by opposing each other in some systematic way.

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

Stirling, W. and Tummolini, L.
Conditional Game Theory as a Model for Coordinated Decision Making.
DOI: 10.5220/0006956702950302
In *Proceedings of the 10th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K 2018) - Volume 2: KEOD*, pages 295-302
ISBN: 978-989-758-330-8

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. Tennis players may collectively understand the rules and best practices of the game, but unless each is able to execute those practices, playing the game will be unrewarding. 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.

Coordination requires individuals to possess some notion of social connectivity in addition to concerns for their own material welfare. There are two extreme methodologies for incorporating coordination into a multilateral decision scenario. One way is for the participants to come to a social engagement with a global view of the way the group is intended to behave. Under this view, coordination is built-in: Each participant performs its ex ante assigned part. Another way is for participants to come to the engagement with local views of how they will behave as they interact with others. Under this view, coordination is emergent: It occurs (or not) as each participant responds to the social influence exerted by others. We argue that the latter approach is the appropriate way to design a collective of autonomous decision makers (agents), and present a general framework for the analysis of human social networks and the design and synthesis of artificially intelligent networks. For coordination to be designed into such a network, however, the social relationships must be defined *operationally*—they must be characterized via mathematical expressions that explicitly model social influence.

A social influence network comprises a group of agents whose choices can depend on the attitudes and opinions of others as well as their own welfare. More precisely, it is a collective of agents who are empowered to make individual choices under the following conditions: a) the combination of the choices of all generates an outcome that affects the welfare of each, and b) the preferences over outcomes for each can be influenced by the preferences of others. The first condition is the usual scenario for standard game theory, but the second condition introduces a social component that is not explicitly modeled by the standard theory. With a social influence network, there can be a difference between what constitutes rational behavior when viewing the anticipated behavior of others as a constraint on the pursuit of narrow self-interest (e.g., material benefit) and what constitutes rational behavior from a socially oriented perspective of viewing oneself as a part of a coordinated whole—a society. Thus, the ability of the individuals to make their choices in a way that responds to social influence, while at the same time retaining their individuality and concern for their own welfare, is of prime importance.

This position paper argues that conditional game theory, introduced by (Stirling, 2012) and (Stirling and Felin, 2013), provides a framework within which to model social influence networks. Conditional game theory comprises three components: a *socialization* mechanism by which indviduals may incorporate the interests of others into their own self-interest without compromising their individuality; a *diffusion* mechanism by which the preferences resulting from an expanded view of self-interest can be conglomerated to produce a comprehensive social model that accounts for all social interrelationships; and a *deduction* mechanism by which coordinated individual decisions may be deduced from the social model. A critical feature of this theory is that it is consistent with the fundamental assumptions of game theory; in fact, conventional noncooperative game theory is a special case of this extended theory.

2 SOCIALIZATION

With conventional game theory, preferences are *categorical*—fixed, immutable, and unconditional. The mathematical mechanism used to express categorical preferences is a *payoff* function. Given a collective of agents $\{X_1, \ldots, X_n\}$ for $n \geq 2$, let $\mathcal{A}_i$ denote a finte set of actions for X_i, and let the Cartesian product set $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ denote the *outcome set*. The function $u_i \colon \mathcal{A} \to \mathbb{R}$ quantifies the payoff to X_i as a function of the combined actions of the collective. Of course, X_i is free to define its preferences in whatever way it chooses, be it egocentric, altruistic, or other. Once defined, u_i is the formal expression of X_i's notion of self-interest. The payoff function establishes a *global ordering* of the outcome set. The most well-known solution concept associated with such payoffs is to juxtapose them into a payoff array and identify Nash equilibria—the set of outcomes such that, if any agent were to make a unilateral change, it's payoff would either decrease or remain unchanged.

The innovation provided by conditional game theory is to allow agents to possess *conditional payoffs*. Establishing this concept requires the application of graph theory. A *network graph* $G(\mathbf{X}, E)$ comprises a set of *vertices* $\mathbf{X} = \{X_1, \ldots, X_n\}$ (the set of agents) and a set $E \subset \mathbf{X} \times \mathbf{X}$ of pairs of vertices such that there is an explicit connection between them that serves as the medium by which influence is propagated between X_i and X_j. Specifically, the expression $X_i \longrightarrow X_j$ means that the influence propagates in only one direction—a *directed edge* from X_i to X_j. 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* if $X_j \mapsto X_j$. A graph is said to be a *directed acyclic graph*, or DAG, if all edges are directed and there are no cycles. For each X_i, its *parent* set is $\mathrm{pa}(X_i) = \{X_{i_1}, \ldots, X_{i_{q_i}}\}$, where $X_{i_k} \longrightarrow X_i$, $k = 1, \ldots, q_i$. If $\mathrm{pa}(X_i) = \varnothing$ then X_i is a *root vertex*.

A *conjecture profile* $\mathbf{a}_i = (a_{i1}, \ldots, a_{in}) \in \mathcal{A}$ is a profile hypothesized by X_i as the outcome to be actualized. The expression $X_i \models \mathbf{a}_i$ means that X_i conjectures $\mathbf{a}_i$. The element a_{ii} is X_i's *self-conjecture*, denoted $X_i \models a_{ii}$, and a_{ij}, $j \neq i$, is an *other-conjecture* by X_i for X_j, denoted $X_i \models a_{ij}$. The array $(\mathbf{a}_1, \ldots, \mathbf{a}_n)$ is termed a *joint conjecture set*.

A *conditioning conjecture profile* by X_i for X_{i_k}, denoted $\mathbf{a}_{i_k} = (a_{i_k1}, \ldots, a_{i_kn})$, is a profile that X_i hypothesizes that $X_{i_k} \models \mathbf{a}_{i_k}$, $k = 1, \ldots, q_k$. A *conditioning conjecture set* $\boldsymbol{\alpha}_{\mathrm{pa}(i)} = (\mathbf{a}_{i_1}, \ldots, \mathbf{a}_{i_{q_i}}) \in \mathcal{A}^{q_i}$ by X_i for $\mathrm{pa}(X_i)$ is the set of conditioning conjecture profiles by X_i for its parents, denoted $\mathrm{pa}(X_i) \models \boldsymbol{\alpha}_{\mathrm{pa}(i)}$.

A *conditional payoff* given $\boldsymbol{\alpha}_{\mathrm{pa}(i)}$, denoted $u_{i|\mathrm{pa}(i)}(\cdot|\boldsymbol{\alpha}_{\mathrm{pa}(i)})$: $\mathcal{A} \to \mathbb{R}$, is an ordering function such that, given the antecedent $\mathrm{pa}(X_i) \models \boldsymbol{\alpha}_{\mathrm{pa}(i)}$, then

$$u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_{\mathrm{pa}(i)}) \geq u_{i|\mathrm{pa}(i)}(\mathbf{a}'_i|\boldsymbol{\alpha}_{\mathrm{pa}(i)}) \tag{1}$$

if X_i prefers the conjecture profile $\mathbf{a}_i$ to $\mathbf{a}'_i$ or is indifferent, given that $\mathrm{pa}(X_i) \models \boldsymbol{\alpha}_{\mathrm{pa}(i)}$. If $\mathrm{pa}(X_i) = \varnothing$, then $u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_{\mathrm{pa}(i)}) = u_i(\mathbf{a}_i)$, a categorical payoff.

A *conditional network game* is a triple $\{\mathbf{X}, \mathcal{A}, \mathcal{U}\}$, where $\mathbf{X} = \{X_1, \ldots, X_n\}$, $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$, and

$$\mathcal{U} = \{u_{i|\mathrm{pa}(i)}(\cdot|\boldsymbol{\alpha}_{\mathrm{pa}(i)})\ \forall \boldsymbol{\alpha}_{\mathrm{pa}(i)} \in \mathcal{A}^{q_i},\ i = 1, \ldots, n\}. \tag{2}$$

A conditional network game degenerates to a standard noncooperative normal-form game when $\mathcal{U} = \{u_i,\ i = 1, \ldots, n\}$. Thus, conditional game theory is an extension of standard noncooperative game theory.

The conditional structure of the preferences enables agents to extend their spheres of interest beyond strategic self-interest without surrendering individuality, and therefore differs fundamentally from the categorical preference structure of a standard noncooperative game. The players of a standard game react to the fixed categorical preferences; there is no opportunity for adaptation as the game is played—the preferences are static. A conditional game enables players to adapt to the social environment, since they are able to respond to the preferences of others *as they interact*—the preferences are dynamic. The representation of a conditional network game as a graph with agents as vertices and linkages as conditional payoffs fully integrates the individual preferences into an organizational structure that enables the synthesis of a comprehensive model of the way individual preferences interact. An example of a three-agent social influence network is

$$\begin{tikzcd} & X_1 \arrow[dl, "u_{2|1}"'] \arrow[dr, "u_{3|12}"] & \\ X_2 \arrow[rr] & & X_3 \end{tikzcd} \tag{3}$$

where X_1 is a root vertex and thus must have a categorical payoff $u_1(\mathbf{a}_1)$, $\mathrm{pa}(X_2) = \{X_1\}$ with conditional payoff $u_{2|1}(\mathbf{a}_2|\mathbf{a}_1)$, and $\mathrm{pa}(X_3) = \{X_1, X_2\}$ with conditional payoff $u_{3|12}(\mathbf{a}_3|\mathbf{a}_1, \mathbf{a}_2)$.

3 DIFFUSION

There is a distinct operational difference between categorial and conditional preferences. Given that X_i categorically prefers $\mathbf{a}_i$ to $\mathbf{a}'_i$, X_i has sufficient information to choose between the two conjectures. But if X_i only conditionally prefers $\mathbf{a}_i$ to $\mathbf{a}'_i$, X_i does not have sufficient information to choose between them without entering into the conditioning social relationships. As the conditional preferences propagate through the group, emergent social interrelationships are established between its members. This process, termed diffusion, involves conglomerating the individual conditional payoffs to form a social model that provides a comprehensive expression of the emergent social structure. Conglomeration is superficially related to the concept of aggregation as employed by social choice theory, but serves a different purpose. With social choice theory, the votes of the individuals are aggregated to form a group-level decision. Conglomeration, by contrast, is a process of combining a collective of parts to form a whole while remaining distinct entities.

Given a conditional game $\{\mathbf{X}, \mathcal{A}, \mathcal{U}\}$, a *coordination functional* is a mapping $F\colon \mathcal{U} \to [0,1]$ that generates a *social model* $u_{1:n}\colon \mathcal{A}^n \to [0,1]$ of the form

$$u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_n) = F[u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\mathbf{a}_{i_1}, \ldots, \mathbf{a}_{i_{q_i}}),\ i = 1, \ldots, n]. \tag{4}$$

The intended role of the coordination functional is to provide a measure of the degree of compatibility of

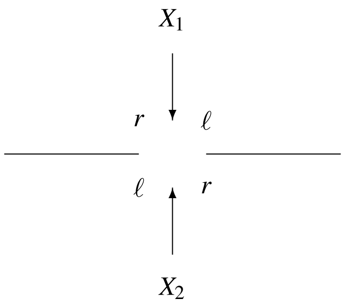

Figure 1: The doorway game.

the agents as they conjecture their various outcomes. To illustrate the manifestation of coordination, consider a scenario involving $\{X_1, X_2\}$, who approach a doorway from opposite directions, as illustrated in Figure 1. Suppose the doorway is just wide enough for two agents to pass simultaneously if they both move either to their respective right (r) or left (ℓ) sides of the doorway. Let $\mathcal{A}_1 = \mathcal{A}_2 = \{r, \ell\}$, yielding the four-element outcome set

$$\mathcal{A} = \mathcal{A}_i \times \mathcal{A}_j = \{(\ell,\ell),(\ell,r),(r.\ell),(r,r)\}, \quad (5)$$

with influence relations

$$\begin{array}{c} X_1 \xrightarrow[u_{2|1}]{} X_2 \\ X_1 \xleftarrow{u_{1|2}} X_2 \end{array} \quad (6)$$

for $i, j, \in \{1,2\}$, $i \neq j$. Suppose X_i possesses a conditional payoff $u_{i|j}(a_{i1}, a_{i2}|a_{j1}, a_{j2})$ defined over $\mathcal{A} \times \mathcal{A}$, and views X_j as a root vertex possessing a categorical payoff $u_j(a_{j1}, a_{j2})$ defined over $\mathcal{A}$.

Let $u_{ij}(\mathbf{a}_i, \mathbf{a}_j)$ be a social model as defined by (4) and consider the joint conjecture sets $[\mathbf{a}_i, \mathbf{a}_j] = [(\ell,\ell),(r,r)]$ and $[\mathbf{a}'_i, \mathbf{a}'_j] = [(\ell,\ell),(\ell,r)]$. The former joint conjecture set corresponds to a scenario where, although the players do not agree regarding which way they should turn, they do agree that they should cooperate, whereas the latter joint conjecture set corresponds to a scenario where X_i conjectures cooperation and X_j conjectures conflict. Assuming that $u_{i|j}(\ell,\ell|r,r) > u_{i|j}(\ell,\ell|\ell,r)$, it would be reasonable that

$$u_{ij}[(\ell,\ell),(r,r)] > u_{ij}[(\ell,\ell),(\ell,r)], \quad (7)$$

meaning that, the joint conjecture set $[(\ell,\ell),(r,r)]$ is more coordinated than $[(\ell,\ell),(\ell,r)]]$. However, both of these joint conjecture sets are less coordinated than $[(\ell,\ell),(\ell,\ell)]$.

The choice of a suitable coordination functional is a critical component of conditional game theory. To motivate such a choice, it is instructive to recognize the analogical relationship between conditional payoffs and conditional probabilities. Indeed, syntax of conditional payoffs and conditional probabilities are in the form of hypothetical propositions of the form "If p then q", where p is the antecedent and q is the consequent. Furthermore, the topology of the network illustrated by (3) is similar to that of a Bayesian network. Thus, the structure and syntax of a social influence network will be identical to that of a Bayesian network if the conditional payoffs are expressed using the mathematical structure of probability theory, namely, that the conditional payoffs are conditional mass functions, that is,

$$\begin{aligned} u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_{\mathrm{pa}(i)}) &\geq 0 \text{ for all } \mathbf{a}_i \in \mathcal{A} \\ \sum_{\mathbf{a}_i} u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_{\mathrm{pa}(i)}) &= 1 \text{ for all } \boldsymbol{\alpha}_{\mathrm{pa}(i)} \in \mathcal{A}^{q_i} \end{aligned} \quad (8)$$

Furthermore, the analogy with a Bayesian network can be made exact by defining the coordination functional according to the fundamental theorem of Bayesian networks, namely,

$$\begin{aligned} u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_n) &= F[u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_{\mathrm{pa}(i)}),\ i = 1, \ldots, n] \\ &= \prod_{i=1}^{n} u_{i|\mathrm{pa}(i)}(\mathbf{a}_i|\boldsymbol{\alpha}_{\mathrm{pa}(i)}). \end{aligned} \quad (9)$$

This structure is attractive for three key reasons. First, it takes advantage of one of the great strengths of probabilistic reasoning, which has long been recognized as an important model of human reasoning. Indeed, as Glenn Shafer has noted, "Probability is not really about numbers; it is about the structure of reasoning" (Pearl, 1988, quoted by). Second, the social model is analogous to the joint distribution of a set of random variables. Analogous to the way a joint probability mass function that captures all of the statistical relationships that exist among a collective of random variables, the social model captures all of the social influence relationships that exist among a collective of agents. Third, adopting (9) as the diffusion functional ensures that no agent can be categorically subjugated by the group in that whatever it chooses as its most preferred outcome is socially unacceptable to the collective. To explain this concept, suppose X_i possesses a categorical payoff u_i, and let $u_{1:n}$ be a social model defined by (4). X_i is *subjugated* if, for every fixed $\mathbf{a}_i \in \mathcal{A}$,

$$u_i(\mathbf{a}_i) > u_i(\mathbf{a}'_i) \text{ for all } \mathbf{a}'_i \neq \mathbf{a}_i \quad (10)$$

holds, then

$$\begin{aligned} &u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_{i-1}, \mathbf{a}_i, \mathbf{a}_{i+1}, \ldots, \mathbf{a}_n) \\ &\quad < u_{1:n}(\mathbf{a}_1, \ldots, \mathbf{a}_{i-1}, \mathbf{a}'_i, \mathbf{a}_{i+1}, \ldots, \mathbf{a}_n) \end{aligned} \quad (11)$$

for all joint conjecture sets $(\mathbf{a}_1, \ldots, \mathbf{a}_{i-1}, \mathbf{a}'_i, \mathbf{a}_{i+1}, \ldots, \mathbf{a}_n) \in \mathcal{A}^n$ with $\mathbf{a}'_i \neq \mathbf{a}_i$. If X_i is subjugated, then, no matter which outcome it most prefers, all joint conjecture sets with X_i conjecturing its most preferred outcome have lower coordination than all joint conjecture sets with X_i not conjecturing its most preferred outcome. In other words, X_i's participation in the group is so toxic that the very fact that it even has a preference destroys the functionality of the group. Thus, avoiding even the potential for any agent to be subjugated is an important consideration for the design of a coordination functional.

The notion of subjugation is mathematically equivalent to the notion of a sure loss gambling scenario; that is, a Dutch book, where the gambler loses more than the entry fee regardless of the outcome. [1] The Dutch book theorem establishes that a sure loss is impossible if, and only if, the gambler's beliefs and actions conform to the axioms of probability theory. Thus, subjugation is impossible if, and only if, the preferences and actions of the agents also conform to the probability axioms.

4 DEDUCTION

The ordering provided by the social model is with respect to joint conjecture sets $\boldsymbol{\alpha}_{1:n} = (\mathbf{a}_1, \ldots, \mathbf{a}_n)$, with each conjecture profile of the form $\mathbf{a}_i = (a_{i1}, \ldots, a_{in})$, where $a_{ii} \in \mathcal{A}_i$ is a self-conjecture by X_i and $a_{ij} \in \mathcal{A}_j$ is an other-conjecture for X_j by X_i. This model is comprehensive in the sense that it captures all of the social relationships that exist among the individuals. It contains all of the information necessary for each agent to deduce the actions that are consistent with its need for individual performance as well as the social influence exerted by others. The deduction process comprises two phases: First, the extraction of an explicit measure of the degree of coordination associated with each conjecture outcome $\mathbf{a} \in \mathcal{A}$,and, second, an ordering over its own self-conjecture $a_{ii} \in \mathcal{A}_i$.

Vilfredo Pareto understood the distinction between individual preference and group sociality. Individual behavior is expressed in terms of the way one makes choices according to one's preferences over alternatives, and group behavior is expressed in terms of the way its members interact as a consequence of their preferences. He employs the notion of "social utility" as a characterization of the degree of satisfaction associated with an alternative. For individuals, the utility of an alternative can be expressed economically with operational measures such as payoffs or other manifestations of individual benefit. According to Pareto, however, the utility of a group should be analyzed sociologically, and may not coincide with the economic payoffs of its individual members.

> In pure economics a community cannot be regarded as a person. In sociology it can be considered, if not as a person, at least as a unit. There is no such thing as the ophelimity of a community; but a community utility can roughly be assumed. So in pure economics there is no danger of mistaking the maximum of ophelimity *for* a community for a non-existent maximum of ophelimity *of* a community. In sociology, instead, we must stand watchfully on guard against confusing the maximum of utility *for* a community with the maximum utility *of* a community, since they both are there [emphasis in original] (Pareto, 1935, pp. 1471, par. 2133).

Coser elaborates on Pareto's distinction between economic utility and social utility.

> By making his distinction between the utility *for* and the utility *of* a community, Pareto moved from classical liberal economics, where it was assumed that total benefits for a community simply involved a sum total of the benefits derived by each individual ("the greatest happiness of the greatest number"), to a sociological point of view in which society is treated as a total unit and sub-groups or individuals are considered from the viewpoint of their contribution to the overall system as well as in terms of their peculiar wants and desires. System needs and individual or sub-group needs are distinguished [emphasis in original]. (Coser, 1971, p. 401)

Although the social model provides a ranking of the sociality of the network with respect to the joint conjecture sets of the network, each X_i has direct control over only a_{ii}, its own self-conjecture. Thus, what is most relevant with respect to coordination is a ranking of how individual self-conjectures a_{ii}, $i = 1, \ldots, n$ combine to form a notion of coordination.

Given a joint conjecture set $\boldsymbol{\alpha}_{1:n} = (\mathbf{a}_1, \ldots, \mathbf{a}_n)$, form the *coordination profile* $\mathbf{a} := (a_{11}, \ldots, a_{nn})$ comprising the set of self-conjectures, and compute the marginal of the social model $u_{1:n}[(a_{11}, \ldots, a_{1n}), \ldots, (a_{n1}, \ldots, a_{nn})]$ with respect to the coordination profile by summing the social model over all elements of each $\mathbf{a}_i$ except the self-conjectures to form the *social*

[1] To establish this equivalence, suppose one were to place a \$1 bet on the event (10), with an fair entry fee of $p > 1/2$ and, simultaneously, to place a bet on the event (11) with a fair entry fee of $q > 1/2$. Regardless of the outcome, the gambler wins \$1 but pays $p + q > 1$—a sure loss.

utility $w_{1:n}$ for $\{X_1,\ldots,X_n\}$, yielding

$$w_{1:n}(a_{11},\ldots,a_{nn}) = \sum_{\neg a_{11}} \cdots \sum_{\neg a_{nn}} u_{1:n}[(a_{11},\ldots,a_{1n}), \ldots,(a_{n1},\ldots,a_{nn})], \quad (12)$$

where the *not-sum* notation $\sum_{\neg a_{11}}$ means that the sum is taken over all elements in the argument list except a_{ii}.

Social utility as a measure of coordination serves as an operational manifestation of the sociologic notion of utility introduced by Pareto, and is distinct from the economic concept of utility expressed via individual payoffs. The relation

$$w_{1:n}(\mathbf{a}) > w_{1:n}(\mathbf{a}') \quad (13)$$

means that the degree to which the set of self-conjectures $\{a_{ii}, i = 1,\ldots,n\}$ (the parts) fit together to form systematic group-level behavior (a whole) is greater than the degree to which $\{a'_{ii}, i = 1,\ldots,n\}$ generates a whole.

Once the coordination function has been defined, the final deduction step is to identify the payoffs for each member of the collective. The *coordinated payoff* for X_i is the i-th marginal of $w_{1:n}$, that is,

$$w_i(a_{ii}) := \sum_{\neg a_{ii}} w_{1:n}(a_{11},\ldots,a_{1n}). \quad (14)$$

The relationship between social utility $w_{1:n}$ and coordinated payoffs w_i for a collective of agents $\{X_1,\ldots,X_n\}$ is analogous to the relationship between a joint probability mass function $p_{1:n}$ and marginal mass functions p_i for a collective of random variables $\{Y_1,\ldots,Y_n\}$. $p_{1:n}(y_1,\ldots,y_n)$ is the degree of probability of the simultaneous realization of the joint event $\{Y_i = y_1,\ldots,Y_n = y_n\}$, and $p_i(y_i)$ is the probability of the single event $\{Y_i = y_i\}$ for each Y_i. If the Y_i's are mutually independent, then $p_{1:n}(y_1,\ldots,y_n) = \prod_{i=1}^n p_i(y_i)$. The "difference" between $p_{1:n}(y_1,\ldots,y_n)$ and $\prod_{i=1}^n p_i(y_i)$ is a measure of the degree of statistical dependence that exists among the random variables.

Similarly, for $(a_{11},\ldots,a_{nn}) \in \mathcal{A}_1, \times\cdots\times \mathcal{A}_n$, the social utility $w_{1:n}(a_{11},\ldots,a_{nn})$ is the degree of coordination of the simultaneous actualization of the joint event $\{X_i \vDash a_{11},\ldots,X_i \vDash a_{nn}\}$, and $w_i(a_{ii})$ is the coordinated payoff of the single event $\{X_i \vDash a_{ii}\}$ for each X_i. If the X_i's are mutually socially independent, then $w_{1:n}(a_{11},\ldots,a_{nn}) = \prod_{i=1}^n w_i(a_{ii})$. The "difference" between $w_{1:n}(a_{11},\ldots,a_{nn})$ and $\prod_{i=1}^n w_i(a_{ii})$ is a measure of the degree of social dependence that exists among the agents. Intuitively, the greater the social dependence, the more the group is able to coordinate.

5 RELATION TO PREVIOUS RESEARCH

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 loosely coupled collectives of mutually interacting individuals (Weick, 1995). Specifically, much of the emphasis has been placed on the structure of such organizations (cf. (French, 1956; DeGroot, 1974; Friedkin, 1986; Arrow et al., 2000; Friedkin and Johnson, 2011)). A basic model is that an individual's socially adjusted payoff is a convex combination of its own categorical payoff and a weighted sum of the categorical payoffs of those agents who influence it. (Hu and Shapley, 2003a; Hu and Shapley, 2003b) apply a command structure to model player interactions by simple games. The subject of influence has also been extensively studied in the context of voting games where the individuals must vote yes or no on a given proposition. (Hoede and Bakker, 1982) introduce the concept of decisional power as a measure of the degree of influence of an individual or coalition of other voters to alter their vote from their original inclination (cf. (Grabisch and Rusinowska, 2010)). (Galeotti et al., 2013) establish conditions for reaching an equilibrium for social influence networks.

Other approaches to the issue of coordination focus on models drawn from biological and social evolutionary processes ((Axelrod, 1984; Bicchieri, 2003; Fefferman and Ng, 2007; Goyal, 2007; Gintis, 2009; Bossert et al., 2012)). Coordination is addressed by studying repeated games, where players replay the same game multiple times. The argument supporting these approaches is that players gain insight regarding the social dispositions of the other players through repeated interaction. They may learn to recognize behavioral patterns and predict the behavior of others. Through this process, they can establish their own reputations and gain the trust of others. Coordination, therefore, is viewed as the end result of social evolution. Such approaches provide important models of the emergence of social relationships in repeated-play environments where individual fitness for long-term survival is taken into consideration in addition to short-term material payoffs. Coordination issues are also central to the study of multiagent systems and general network theory (Jackson, 2008; Shoham and Leyton-Brown, 2009; Easley and Kleinberg, 2010).

Social scientists have long recognized the need to expand notions of preference beyond egocentric interest. Behavioral game theory (cf. (Bolton and Ockenfels, 2005; Fehr and Schmidt, 1999; Henrich et al., 2004; Camerer et al., 2004; 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 preferences in addition to considerations of material benefit. The closely related field of psychological game theory (cf. (Geanakoplos et al., 1989; Dufwenberg and Kirchsteiger, 2004; Colman, 2003; Battigalli and Dufwenberg, 2009; Gilboa and Schmeidler, 1988)) also employs preferences that account for beliefs as well as actions and takes into consideration belief-dependent motivations such as guilt aversion, reciprocity, regret, and shame. The concept of "team-reasoning" has been promoted by (Sugden, 2015) and (Bacharach, 2006), where individuals view themselves as members of a team, and therefore are motivated to modify their behavior to conform with team aspirations. (Hedahl and Huebner, 2018) focus on value sharing and discuss processes for providing normative grounding for pursuing shared ends. (Reischmann and Oechssler, 2018) introduce a mechanism for public good provision using conditional offers based on the willingness of others to contribute.

A thread common to these approaches is that they rely on ex ante linear preference orderings that are static, immutable, global, and unconditional—they are *categorical*. We argue that this single thread must be replaced by a richer interweave of preference relationships that involve explicit social influence.

The perspectives that comport most closely with this paper are the views held by (Ross, 2014) and (Bratman, 2014). Ross asserts that individual preferences are not formed in a social vacuum; rather, they are the consequence of social processes, and must therefore be dependent on the social environment. Bratman argues similarly, and introduces a notion of *augmented individualism*, where the intentions of an individual are composed of relevant interrelated attitudes, leading to a notion of shared agency. Essentially, conditional game theory is the operationalization of these two perspectives.

6 CONCLUSIONS

Conditional game theory offers a significant extension of standard game theory as a framework for both the analysis of human networks and the design and synthesis of artificial social influence networks.

- Social influence is ex ante incorporated endogenously into the payoffs rather than exogenously imposed via an ex post solution concept.
- An operational definition of coordination is generated as a group-level attribute that is considered parallel to the individual-level attribute of preference.
- Individual coordinated decisions are deduced as a consequence of the diffusion of social influence throughout the network.

REFERENCES

Arrow, H., McGrath, J. E., and Berdahl, J. L. (2000). *Small Groups as Complex Systems*. Sage Publications, Inc., Thousand Oaks, CA.

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

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

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

Bicchieri, C. (2003). *Rationality and Coordination*. Cambridge University Press, Cambridge, UK.

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

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

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

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

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

Coser, L. A. (1971). *Masters of Sociological Thought*. Harcourt Brace Jovanovich, Inc., New York, NY.

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

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

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

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

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

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

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

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

Galeotti, A., Ghiglino, C., and Squintani, F. (2013). Strategic information transmission networks. *Journal of Economic Theory*, 148(5):1751–1769.

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

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

Gintis, H. (2009). *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*. Princeton University Press, Princeton, NJ.

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

Grabisch, M. and Rusinowska, A. (2010). Different approaches to influence based on social networks and simple games. In Deeman, A. V. and Rusinowska, A., editors, *Collective Decision Making: Views from Social Choice and Game Theory*. Springer-Verlag, Berlin.

Hedahl, M. and Huebner, B. (2018). Sharing values. *The Southern Journal of Philosophy*, 56:240–272.

Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., and Gintis, H., editors (2004). *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-scale Societies*. Oxford University Press, Oxford, UK.

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

Hoede, C. and Bakker, R. (1982). A theory of decisional power. *Journal of Mathematical Sociology*, 8:309–322.

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

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

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

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

Pareto, V. (1935). *A Treatise on General Sociology, Volumes Three and Four*. Dover, New York, NY. edited by A. Livingston.

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

Reischmann, A. and Oechssler, J. (2018). The binary conditional contribution mechanism for public good provision in dynamic settings—theory and experimental evidence. *Journal of Public Economics*, 159:104–115.

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

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

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

Stirling, W. C. and Felin, T. (2013). Game theory, conditional preference, and social influence. PLos ONE 8(2): e56751. doi:10.1371/journal.pone.0056751.

Sugden, R. (2015). Team reasoning and intentional cooperation for mutual benefit. *Journal of Social Ontology*, 1(1):143–166. ISSN (Online) 2196-9963, ISSN (Print) 2196-9655, DOI: 10.1515/jso-2014-0006, November 2014.

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