

CONDITIONAL GAME THEORY

A Generalization of Game Theory for Cooperative Multiagent Systems

Wynn C. Stirling

Electrical and Computer Engineering Department, Brigham Young University, Provo, 84602 Utah, U.S.A.

Keywords: Game theory, Multiagent systems, Group rationality, Cooperation, Negotiation.

Abstract: Game theory provides a framework within which to model multiagent systems. The conventional neoclassical theory is well suited for competitive scenarios where self-interest is the dominant concept of rational behavior, but is less appropriate for scenarios where opportunities for such complex social behavior as cooperation, compromise, and unselfishness are significant. Conditional game theory is an extension of the conventional neoclassical theory that permits agents to extend their spheres of interest beyond the self and enables them to condition their preferences on the preferences of other agents, thereby providing a mechanism with which to characterize complex social behavior. As these conditional preferences propagate through the system, social bonds are created among the players that permit notions of both group and individual preferences to emerge and, hence, for concepts of both group rationality and individual rationality to coexist. Computational complexity can often be mitigated by exploiting the sparseness of influence relationships among the members of the system.

1 INTRODUCTION

Game theory provides a mathematical framework within which to model decisions by multiple entities where the outcome for each depends on the choices of all. Game theory is increasingly invoked by engineering and computer science as a framework for multiagent systems (Shoham and Leyton-Brown, 2009; Nisan et al., 2007; Vlassis, 2007; Parsons et al., 2002; Weiss, 1999).

A noncooperative, single-stage, strategic-form game consists of (i) a set of autonomous decision makers, or *players*, denoted $X_n = \{X_1, \dots, X_n\}$ where $n \geq 2$, (ii) an action set \mathcal{A}_i for each X_i , and (iii) a utility u_{X_i} : $\mathcal{A} \rightarrow \mathbb{R}$ for each X_i , $i = 1, \dots, n$, where $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ is the product action space. For any *action profile* $\mathbf{a} = (a_1, \dots, a_n) \in \mathcal{A}$, the utility $u_{X_i}(\mathbf{a})$, defines the benefit to X_i as a consequence of the instantiation of \mathbf{a} . These utilities are *categorical* in the sense that $u_{X_i}(\mathbf{a})$ unconditionally defines the benefit to X_i of the group instantiating the action profile \mathbf{a} , ostensibly without regard for the benefit that instantiating \mathbf{a} offers to other agents.

In addition to the categorical structure of the utilities, it is usually assumed that each X_i possesses a logical structure that defines how it should play the game. The most widely used logical structure is the

doctrine of *individual rationality*: each X_i should act in a way that maximizes its own utility, regardless of the effect doing so has on others. Under the assumption that each player subscribes to this notion and believes that all others do so as well, they each will solve their corresponding constrained optimization problem, resulting in a Nash equilibrium.

The mathematical structure of categorical utilities and the logical structure of individual rationality are ideally matched to each other. Given categorical utilities, the only compatible notion of rationality is self-interest, since the utility is restricted to, and only to, the individual's interests. Conversely, given individual rationality, any structure other than categorical utilities would extend interest beyond the self.

These mathematical and logical structures may provide an appropriate vehicle with which to model behavior in an environment of competition and market driven expectations since, in that environment, the dominant notion of rational behavior clearly is self-interest. It is less clear, however, that self-interest is the dominant notion in mixed-motive environments, such as those that contain opportunities for cooperation, compromise, and unselfishness as well as for competition, intractability, and avarice. Arrow clearly delimits the context in which individual rationality applies: "rationality in application is not merely a

property of the individual. Its useful and powerful implications derive from the conjunction of individual rationality and other basic concepts of neoclassical theory—equilibrium, competition, and completeness of markets . . . When these assumptions fail, the very concept of rationality becomes threatened, because perceptions of others and, in particular, their rationality become part of one’s own rationality” (Arrow, 1986, p. 203).

Despite Arrow’s caution, the mathematical and logical structures of game theory are routinely applied to mixed-motive situations, often producing results that are at variance with observed behavior. Behavioral economics (e.g., see (Camerer, 2003)) seeks to mitigate this problem by inserting parameters to model fairness, loss aversion, and other such issues into the utilities to provide more psychological realism. Once included, however, the game is still solved according to conventional individual rationality and categorical utilities.

What is missing with conventional game theory is a notion of group benefit. Multiagent systems are typically designed such that the individuals work in a cooperative manner to accomplish some task, but, unfortunately, relying on individual rationality does not foster group rationality. As observed by Luce and Raiffa, “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 . . . general game theory seems to be in part a sociological theory which does not include any sociological assumptions, and, although one might hope one day to derive sociological theory from individual psychology, it may be too much to ask that any sociology be derived from the single assumption of individual rationality” (Luce and Raiffa, 1957, p. 193, 196). Consequently, game theory has proceeded by making assumptions about individual preferences only and then using those preferences to deduce information about the choices (but not the values) of a group.

It might be expected that cooperative game theory possesses some notion of group rationality. This version of game theory permits a subset of players to enter into a coalition such that each receives a payoff that is greater than it would receive if it acted alone. However, cooperative game theory employs categorical utilities and its solutions concepts are based squarely upon the assumption of individual rationality. Each player enters into a coalition solely on the basis of benefit to itself and, even though each may be better off for having joined, a notion of “group benefit” is not an issue when forming the coalition.

The need to extend beyond individual rationality is critical to the design of multiagent systems. With

the social and behavioral sciences, game theory models are used as *analysis* tools to explain, predict, justify, and recommend courses of action, but the models are not causal; they are approximations to reality, and do not dictate behavior. With the engineering and computer science applications, however, game theory models are used as *synthesis* tools to design artificial autonomous entities. Such models are causal; they do dictate behavior; they create reality. Thus, when synthesizing social behavior, all social relationships must explicitly be part of the mathematical and logical models.

Reliance on categorical utilities and individual rationality limits the application of conventional game theory for the design and synthesis of multiagent systems that are intended to be cooperative. The contributions of this paper are (i) to present a new utility structure that overcomes the limitations of categorical utilities as a model of complex social relationships, (ii) to offer a more general concept of rational behavior that simultaneously accounts for both group and individual welfare, and (iii) to address and control the computational complexity of the resulting model.

2 PREFERENCE MODELS

2.1 Neoclassical Preference Models

The most prevalent assumption employed by game theory when considering preference orderings is also the most simple: a preference ordering over alternatives is defined for each individual agent. Arrow put it succinctly: “It is assumed that each individual in the community has a definite ordering of all conceivable social states, in terms of their desirability to him . . . It is simply assumed that the individual orders all social states by whatever standards he deems relevant” (Arrow, 1951, p. 17). According to this view, each agent’s preference ordering is completely defined and immutable before the game begins. Thus, from the conventional point of view, the starting point of a game is the definition of categorical utilities for each player. Furthermore, as Friedman argues, it is not necessary to consider the process by which the agents arrive at their preference orderings. “The economist has little to say about the formation of wants; this is the province of the psychologist. The economist’s task is to trace the consequences of any given set of wants” (Friedman, 1961, p. 13).

If we take the Arrow/Friedman division of labor as the starting point when defining a game, we must assume that the individual is able to reconcile all internal conflicts to the point that a unique categorical

preference ordering can be defined that corresponds to its own self interest and which is not susceptible to change as a result of social interaction. This is a tall order, but nothing less will do if we are restricted to categorical preference orderings.

2.2 Social Influence Preference Models

When complex social relationships exist for which categorical preferences are not adequate or appropriate, a natural way for a player to take them into account is by the notion of *influence*. There are many ways to account for social influence, but the approach presented in this paper is to apply a set of principles to define a systematic and logically defensible mathematical model that leads to the definition and implementation of a multiagent decision methodology that accounts for influence relationships when they exist and treats conventional game theory as a special case when such relationships are absent.

Principle 1 (Conditioning). *Agents' preferences may be influenced by the preferences of other agents.*

X_j influences X_i if X_i 's preferences are affected by X_j 's preferences. Without knowledge of X_j 's preferences, X_i is in a state of suspense with respect to its own preferences. Essentially, X_j 's preferences propagate through the group to affect X_i 's preferences, thereby generating a social bond between the two agents. Once such a bond exists, it is possible to define a notion of joint preference for the two agents viewed simultaneously, and it is possible to extract individual preference orderings from this joint preference ordering since, once X_j 's preferences are revealed, X_i need no longer remain in suspense. It is thus possible for both group and individual preferences to co-exist.

Principle 1 represents an important shift in perspective from conventional game theory. With the conventional approach, the utility of an individual defines its self-interest with respect to the *instantiation* of actions taken by all players. By contrast, we view the utility of an individual as the consequent of a hypothetical proposition whose antecedent is the assumption that those who influence it have identified their most preferred outcomes.

Definition 1. A *conjecture* for X_j is an action profile, denoted \mathbf{a}_j , that is hypothesized as X_j 's most preferred outcome. Let $\mathbf{x}_m = \{X_{j_1}, \dots, X_{j_m}\}$ be a subgroup of \mathbf{x}_n . A *joint conjecture* for \mathbf{x}_m , denoted $\boldsymbol{\alpha}_m = \{\mathbf{a}_{j_1}, \dots, \mathbf{a}_{j_m}\}$, is a collection of action profiles in $\mathcal{A}^m = \mathcal{A} \times \dots \times \mathcal{A}$ (m times), where \mathbf{a}_{j_l} is a conjecture for X_{j_l} , $l = 1, \dots, m$. \square

Now suppose X_i is influenced by a subgroup \mathbf{x}_m . Given a joint conjecture $\boldsymbol{\alpha}_m$, the consequent of the hypothetical proposition is a conditional utility for X_i .

Definition 2. Let $\mathbf{x}_m = \{X_{j_1}, \dots, X_{j_m}\}$ be a subgroup of \mathbf{x}_n that influences X_i and let $\boldsymbol{\alpha}_m = \{\mathbf{a}_{j_1}, \dots, \mathbf{a}_{j_m}\}$ be a joint conjecture for \mathbf{x}_m . A *conditional utility* $u_{X_i|\mathbf{x}_m}(\cdot|\boldsymbol{\alpha}_m)$ is a real-valued function defined over \mathcal{A} that specifies the preference ordering for X_i given the joint conjecture $\boldsymbol{\alpha}_m$. That is, $u_{X_i|\mathbf{x}_m}(\mathbf{a}|\boldsymbol{\alpha}_m) > u_{X_i|\mathbf{x}_m}(\mathbf{a}'|\boldsymbol{\alpha}_m)$ means that X_i prefers \mathbf{a} to \mathbf{a}' , given that X_{j_l} conjectures \mathbf{a}_{j_l} , $l = 1, \dots, m$. \square

Each X_i must define a conditional utility for the joint conjectures of the subgroup that influences it. This requirement increases the complexity of a problem statement over the conventional requirement of specifying only one categorical utility for each X_i . However, as we shall explore in Section 5, there often will be ways to simplify the specification that keeps the complexity under control. Nevertheless, the inclusion of social influence will generally result in increased complexity.

2.3 Group Preference

Conventional game theory eschews the notion of group preference, primarily on the grounds that the group is not a "superplayer" that possesses the power to make decisions. Shubik put it this way: "It may be meaningful, in a given setting, to say that a group 'chooses' or 'decides' something. It is rather less likely to be meaningful to say that the group 'wants' or 'prefers' something" (Shubik, 1982, p. 124). The only exception to this dictum is the *principle of unanimity*: if all members of a group most prefer the same outcome, then the group most prefers that outcome.

Once we extend beyond self-interest via social bonds induced by conditional utilities, however, it becomes possible to consider a more general notion of group preference. It may happen that the social bonds are so strong that unanimity will result, but that situation will not generally obtain. If agents disagree regarding what is best, then some degree of conflict, or discord, will exist within the group. Thus, when designing a system whose members must coordinate, a critical issue is the *concordance*, or the degree of harmony, among its members.

Definition 3. Let $\mathbf{x}_k = \{X_{i_1}, \dots, X_{i_k}\}$ be a subgroup of \mathbf{x}_n . A *concordance utility* $U_{\mathbf{x}_k}$ is a real-valued function defined over \mathcal{A}^k such that, for each joint conjecture $\boldsymbol{\alpha}_k = (\mathbf{a}_{i_1}, \dots, \mathbf{a}_{i_k}) \in \mathcal{A}^k$, $U_{\mathbf{x}_k}(\boldsymbol{\alpha}_k)$ defines the concordance of $\boldsymbol{\alpha}_k$. When $k = 1$, the concordance utility becomes a conventional categorical utility for X_k , that is, $U_{X_k} = u_{X_k}$. \square

When $k > 1$, the concordance utility is a generalization of individual utility which, rather than providing a preference ordering for a single agent over the space \mathcal{A} of action profiles, provides a concordance ordering for a k -member subgroup over the product space \mathcal{A}^k of joint conjectures. When $\mathbf{a}_{i_1} = \dots = \mathbf{a}_{i_k}$, the concordance utility measures the degree of harmony if all members conjecture the same action profile. When the conjectures are different, $U_{x_k}(\mathbf{a}_{i_1}, \dots, \mathbf{a}_{i_k})$ measures the degree of concordance that exists among the members of the subgroup if each X_{i_l} were to view \mathbf{a}_{i_l} as its most-preferred outcome. The expression $U_{x_k}(\boldsymbol{\alpha}_k) > U_{x_k}(\boldsymbol{\alpha}'_k)$ means that $\boldsymbol{\alpha}_k$ causes a more severe conflict for the group than does $\boldsymbol{\alpha}'_k$.

Definition 4. Let $\mathcal{X}_k = \{X_{i_1}, \dots, X_{i_k}\}$ and $\mathcal{X}_m = \{X_{j_1}, \dots, X_{j_m}\}$ be two disjoint subgroups of \mathcal{X}_n . For each $\boldsymbol{\alpha}_m \in \mathcal{A}^m$, a *conditional concordance utility* given $\boldsymbol{\alpha}_m$ is a real-valued function $U_{x_k|x_m}(\cdot|\boldsymbol{\alpha}_m)$ defined over \mathcal{A}^k that defines a concordance utility for \mathcal{X}_k given that \mathcal{X}_m jointly conjectures $\boldsymbol{\alpha}_m$. When $k = 1$, the concordance utility becomes a conditional utility for X_k , that is, $U_{x_k|x_m} = u_{x_k|x_m}$ (see Definition 2). \square

Example 1. Suppose the group $\{X_1, X_2, X_3\}$ is to purchase an automobile. X_1 is to choose the model, either a convertible (C) or a sedan (S), X_2 is to choose the manufacturer, either domestic (D) or foreign (F), and X_3 is to choose the color, either red (R) or green (G). The action spaces are $\mathcal{A}_1 = \{C, S\}$, $\mathcal{A}_2 = \{D, F\}$, $\mathcal{A}_3 = \{R, G\}$. The expression

$$U_{x_1|x_2}[(C, F, R), (S, D, G)] > U_{x_1|x_2}[(S, D, G), (C, F, R)] \quad (1)$$

means that concordance is higher (i.e., it is less conflictive) for the sub-collective $\{X_1, X_2\}$ if X_1 were to most-prefer a foreign-made red convertible and, simultaneously, X_2 were to most-prefer a domestic-made green sedan, than if X_1 were to most-prefer a domestic green sedan and, simultaneously, X_2 were to most-prefer a foreign-made red convertible. Thus, even though the two stakeholders do not have the same preferences in either case, the severity of the differences in opinion is less for the $(C, F, R), (S, D, G)$ combination than for the $(S, D, G), (C, F, R)$ combination.

Next, the expression

$$U_{x_1|x_2}(C, F, R|C, F, R) > U_{x_1|x_2}(S, D, G|C, F, R) \quad (2)$$

means that X_1 prefers a foreign-made red convertible to a domestic-made green sedan, given the hypothesis that X_2 most-prefers a foreign-made red convertible. Notice that, since the consequent involves only one stakeholder, the conditional joint conjecture ordering becomes a conditional ordering, and we may

more properly replace (2) with the expression

$$u_{x_1|x_2}(C, F, R|C, F, R) > u_{x_1|x_2}(S, D, G|C, F, R). \quad (3)$$

Continuing with this example, the expression

$$u_{x_1|x_2|x_3}[S, D, G|(C, F, R), (S, D, G)] > u_{x_1|x_2|x_3}[C, F, R|(C, F, R), (S, D, G)] \quad (4)$$

means that X_1 prefers a domestic-made green sedan to a foreign-made red convertible, given the hypothesis that X_2 most-prefers a foreign-made red convertible and that X_3 most-prefers a domestic-made green sedan.

We conclude this example by examining the following expression.

$$U_{x_2|x_3|x_1}[(C, F, R), (S, F, R)|S, D, G] > U_{x_2|x_3|x_1}[(S, D, G), (C, D, G)|S, D, G]$$

means that the sub-collective $\{X_2, X_3\}$ is less conflicted, given that X_1 most-prefers (S, D, G) , for X_2 and X_3 to prefer (C, F, R) and (S, F, R) , respectively, than respectively to most-prefer (S, D, G) and (C, D, G) . \square

Computing the concordance utility of a group $\mathcal{X}_n = \{X_1, \dots, X_n\}$ is a key component of our approach, since that function captures all social relationships that exist in the group. Using this function, we can define notions of rational behavior both for the group as a whole and for each of its members. To proceed, we require the following principle.

Principle 2 (Endogeny). *If a concordance ordering for a group of agents exists, it must be determined by the social interactions among the subsets of the group.*

This principle precludes the exogenous imposition of aggregation structures. For example, a common conventional aggregation procedure is to form the weighted sum of individual utilities. Such a structure, however, is appropriate only under conditions of preferential independence (e.g., see (Debreu, 1959; Keeney and Raiffa, 1993)). When preferential dependencies exist, however, we seek an aggregation structure that naturally emerges from within the group.

Given the existence of a concordance utility U_{x_m} and a conditional concordance utility $U_{x_k|x_m}$, our goal is to compute the concordance utility of the union of the two subgroups; i.e., to form $U_{x_m x_k}$, the concordance utility for $\mathcal{X}_m \cup \mathcal{X}_k$.

Definition 5. Let $\mathcal{X}_k = \{X_{i_1}, \dots, X_{i_k}\}$ and $\mathcal{X}_m = \{X_{j_1}, \dots, X_{j_m}\}$ be two disjoint subgroups of \mathcal{X}_n such that U_{x_m} and $U_{x_k|x_m}$ are defined. These utilities are

endogenously aggregated if there exists a function F such that

$$U_{x_m x_k}(\alpha_m, \alpha_k) = F[U_{x_m}(\alpha_m), U_{x_k|x_m}(\alpha_k|\alpha_m)], \quad (5)$$

□

When social relationships exist among the members of a collective, there may not be a unique way to represent them mathematically. Let us illustrate this situation by returning to the car-buying example introduced above.

Example 2. Let X_1 , X_2 , and X_3 be as defined in Example 1. Let us suppose that X_1 possesses a categorical utility u_{x_1} over \mathcal{A} , but that X_2 possesses a conditional utility of the form $u_{x_2|x_1}$; that is, X_2 conditions its preferences on the preferences of X_1 . Ignoring, for the time being, the presence of X_3 , let us consider the aggregation of the utilities u_{x_1} and $u_{x_2|x_1}$ to form the sub-collective utility $u_{x_1 x_2}$. As thus framed, the goal is to define a function F such that

$$U_{x_1 x_2}(\mathbf{a}_1, \mathbf{a}_2) = F[u_{x_1}(\mathbf{a}_1), u_{x_2|x_1}(\mathbf{a}_2|\mathbf{a}_1)].$$

Now let us suppose that there is a well-defined social relationship between X_1 and X_2 such that, when defining their preferences, they both take into consideration that, ultimately, they will be operating in a group environment, and not in isolation. Under these conditions, it is possible to re-frame the scenario by X_2 defining a categorical utility u_{x_2} and X_1 defining a conditional utility $u_{x_1|x_2}$. Under this framing, the aggregation problem requires

$$U_{x_2 x_1}(\mathbf{a}_2, \mathbf{a}_1) = F[u_{x_2}(\mathbf{a}_2), u_{x_1|x_2}(\mathbf{a}_1|\mathbf{a}_2)].$$

□

Principle 3 (Consistency). *If a multiagent decision problem can be framed in more than one way using exactly the same information, all such framings should yield the same aggregated concordance ordering.*

Definition 6. Let \mathcal{X}_k and \mathcal{X}_m be two disjoint subgroups of \mathcal{X}_n and suppose there exist two framings of the preferences and relationships between the two subgroups of the forms $\{U_{x_k}, U_{x_m|x_k}\}$ and $\{U_{x_m}, U_{x_k|x_m}\}$. The endogenous aggregation is *consistent* if

$$F[U_{x_k}(\alpha_k), U_{x_m|x_k}(\alpha_m|\alpha_k)] = F[U_{x_m}(\alpha_m), U_{x_k|x_m}(\alpha_k|\alpha_m)]. \quad (6)$$

□

3 AGGREGATION OF SOCIAL PREFERENCES

One of the major aims of our development is to define a mechanism from which a notion of group preference ordering can emerge from the aggregation of conditional individual preference orderings. An essential characteristic of any such mechanism is that it must possess the following property.

Principle 4 (Monotonicity). *If a subgroup prefers one alternative to another and the complementary subgroup is indifferent with respect to the two alternatives, then the group as a whole must not prefer the latter alternative to the former one.*

Principle 4 invokes the common sense concept that, in the absence of opposition, the group must not arbitrarily override the wishes of individuals. Thus, if X_1 prefers \mathbf{a} to \mathbf{a} and X_2 is indifferent between the two profiles, the group $\{X_1, X_2\}$ should not prefer \mathbf{a} to \mathbf{a} . In terms of utilities, this condition means that F must be nondecreasing in both arguments.

When modeling influence relationships, it is critical that we delimit generality to ensure computational tractability. We thus propose the following principle.

Principle 5 (Acyclicity). *No cycles occur in the influence relationships among the agents.*

Given two disjoint subgroups \mathcal{X}_k and \mathcal{X}_m of \mathcal{X}_n , acyclicity means that it cannot happen that, simultaneously, \mathcal{X}_m directly influences \mathcal{X}_k and \mathcal{X}_k directly influences \mathcal{X}_m . The fact that cycles are not permitted does reduce the generality of the model. Nevertheless, restricting to one-way influence relationships is a significant generalization of the neoclassical approach, which assumes that all utilities are categorical and, hence, are trivially acyclical.

3.1 The Aggregation Theorem

It remains to define a function F that complies with the above-mentioned principles. Since positive affine transformations preserve the mathematical integrity of von Neumann-Morgenstern utilities, we may assume, without loss of generality, that all utilities are non-negative and normalized to sum to unity; that is,

$$U_{x_k}(\alpha_k) \geq 0 \quad \alpha_k, \quad (7)$$

$$U_{x_m|x_k}(\alpha_m|\alpha_k) \geq 0 \quad \alpha_m, \alpha_k, \quad (8)$$

$$\sum_{\alpha_k} U_{x_k}(\alpha_k) = 1, \quad (9)$$

$$\sum_{\alpha_m} U_{x_m|x_k}(\alpha_m|\alpha_k) = 1 \quad \alpha_k. \quad (10)$$

Theorem 1 (The Aggregation Theorem). Let $\mathbf{x}_n = \{X_1, \dots, X_n\}$ be an n -member multiagent system and let \mathcal{B}_m denote the set of all m -element subgroups of \mathbf{x}_n . That is, $\mathbf{x}_m \in \mathcal{B}_m$ if $\mathbf{x}_m = \{X_{i_1}, \dots, X_{i_m}\}$ with $1 \leq i_1 < \dots < i_m \leq n$. Let $\{U_{\mathbf{x}_m}: \mathbf{x}_m \in \mathcal{B}_m, m = 1, \dots, n\}$ be a family of normalized non-negative concordance utilities and let

$$\{U_{\mathbf{x}_m|\mathbf{x}_k}: \mathbf{x}_m \cap \mathbf{x}_k = \emptyset, \mathbf{x}_m \in \mathcal{B}_m, \mathbf{x}_k \in \mathcal{B}_k, m+k \leq n\} \quad (11)$$

be a family of normalized non-negative conditional concordance utilities associated with all pairs of disjoint subgroups of \mathbf{x}_n . These utilities are endogenously aggregated if and only if, for every pair of disjoint subgroups \mathbf{x}_m and \mathbf{x}_k ,

$$U_{\mathbf{x}_m\mathbf{x}_k}(\boldsymbol{\alpha}_m, \boldsymbol{\alpha}_k) = F[U_{\mathbf{x}_k}(\boldsymbol{\alpha}_k), U_{\mathbf{x}_m|\mathbf{x}_k}(\boldsymbol{\alpha}_m|\boldsymbol{\alpha}_k)] \quad (12)$$

$$= U_{\mathbf{x}_m\mathbf{x}_k}(\boldsymbol{\alpha}_m|\boldsymbol{\alpha}_k)U_{\mathbf{x}_k}(\boldsymbol{\alpha}_k). \quad (13)$$

This theorem was originally introduced by (Cox, 1946) as an alternative development of the mathematical syntax of probability theory. The proof below follows (Jaynes, 2003).

Proof of the Aggregation Theorem. Let \mathbf{x}_i , \mathbf{x}_j , and \mathbf{x}_k be arbitrary pairwise disjoint subgroups of \mathbf{x}_n , and let $U_{\mathbf{x}_i\mathbf{x}_j\mathbf{x}_k}$, $U_{\mathbf{x}_i|\mathbf{x}_j\mathbf{x}_k}$, $U_{\mathbf{x}_j|\mathbf{x}_i\mathbf{x}_k}$, $U_{\mathbf{x}_i\mathbf{x}_j}$, $U_{\mathbf{x}_i|\mathbf{x}_j}$, and $U_{\mathbf{x}_j}$ be endogenously aggregated concordance utilities. That is,

$$U_{\mathbf{x}_i\mathbf{x}_j\mathbf{x}_k}(\boldsymbol{\alpha}_i, \boldsymbol{\alpha}_j, \boldsymbol{\alpha}_k) = F[U_{\mathbf{x}_j\mathbf{x}_k}(\boldsymbol{\alpha}_j, \boldsymbol{\alpha}_k), U_{\mathbf{x}_i|\mathbf{x}_j\mathbf{x}_k}(\boldsymbol{\alpha}_i|\boldsymbol{\alpha}_j, \boldsymbol{\alpha}_k)] \quad (14)$$

$$= F[U_{\mathbf{x}_k}(\boldsymbol{\alpha}_k), U_{\mathbf{x}_i\mathbf{x}_j|\mathbf{x}_k}(\boldsymbol{\alpha}_i, \boldsymbol{\alpha}_j|\boldsymbol{\alpha}_k)]. \quad (15)$$

But

$$U_{\mathbf{x}_i\mathbf{x}_k}(\boldsymbol{\alpha}_i, \boldsymbol{\alpha}_k) = F[U_{\mathbf{x}_k}(\boldsymbol{\alpha}_k), U_{\mathbf{x}_i|\mathbf{x}_k}(\boldsymbol{\alpha}_i|\boldsymbol{\alpha}_k)] \quad (16)$$

and

$$U_{\mathbf{x}_i\mathbf{x}_j|\mathbf{x}_k}(\boldsymbol{\alpha}_i, \boldsymbol{\alpha}_j|\boldsymbol{\alpha}_k) = F[U_{\mathbf{x}_j|\mathbf{x}_k}(\boldsymbol{\alpha}_j|\boldsymbol{\alpha}_k), U_{\mathbf{x}_i|\mathbf{x}_j\mathbf{x}_k}(\boldsymbol{\alpha}_i|\boldsymbol{\alpha}_j, \boldsymbol{\alpha}_k)]. \quad (17)$$

Substituting (16) into (14) and (17) into (15) yields

$$F\left[F[U_{\mathbf{x}_k}(\boldsymbol{\alpha}_k), U_{\mathbf{x}_j|\mathbf{x}_k}(\boldsymbol{\alpha}_j|\boldsymbol{\alpha}_k)], U_{\mathbf{x}_i|\mathbf{x}_j\mathbf{x}_k}(\boldsymbol{\alpha}_i|\boldsymbol{\alpha}_j, \boldsymbol{\alpha}_k)\right] = F\left[F[U_{\mathbf{x}_k}(\boldsymbol{\alpha}_k), F[U_{\mathbf{x}_j|\mathbf{x}_k}(\boldsymbol{\alpha}_j|\boldsymbol{\alpha}_k), U_{\mathbf{x}_i|\mathbf{x}_j\mathbf{x}_k}(\boldsymbol{\alpha}_i|\boldsymbol{\alpha}_j, \boldsymbol{\alpha}_k)]]\right]. \quad (18)$$

In terms of general arguments, this equation becomes

$$F[F(x, y), z] = F[x, F(y, z)], \quad (19)$$

called the *associativity equation*. By direct substitution it is easy to see that (19) is satisfied if

$$f[F(x, y)] = f(x)f(y) \quad (20)$$

for any function f . It has been shown by (Cox, 1946) that if F is differentiable in both arguments, then (20) is the general solution to (19). Taking f as the identity function, $F(x, y) = xy$, and

$$\begin{aligned} U_{\mathbf{x}_i\mathbf{x}_j}(\boldsymbol{\alpha}_i, \boldsymbol{\alpha}_j) &= F[U_{\mathbf{x}_i}(\boldsymbol{\alpha}_i), U_{\mathbf{x}_j|\mathbf{x}_i}(\boldsymbol{\alpha}_j|\boldsymbol{\alpha}_i)] \\ &= U_{\mathbf{x}_i}(\boldsymbol{\alpha}_i)U_{\mathbf{x}_j|\mathbf{x}_i}(\boldsymbol{\alpha}_j|\boldsymbol{\alpha}_i). \end{aligned} \quad (21)$$

To prove the converse, we note that F given by (13) is nondecreasing in both arguments since $U_{\mathbf{x}_k}$ and $U_{\mathbf{x}_m|\mathbf{x}_k}$ are nonnegative. Also, since the subgroups \mathbf{x}_m and \mathbf{x}_k are arbitrary, (13) holds if we reverse the roles of m and k . Thus, consistency is satisfied and the aggregation is endogenous. \square

The aggregation theorem establishes that, upon compliance with the aforementioned principles, utility aggregation conforms to the same mathematical syntax as does probability. Consequently, the various epistemological properties of probability theory can be accorded analogous interpretations in the praxeological context. Key concepts in this regard are marginalization, independence, and the chain rule.

Marginalization. Let $\mathbf{x}_m = \{X_{j_1}, \dots, X_{j_m}\}$ and $\mathbf{x}_k = \{X_{i_1}, \dots, X_{i_k}\}$ be disjoint subgroups of \mathbf{x}_n . Then the marginal concordance utility of \mathbf{x}_m is obtained by summing over \mathbf{x}_k , yielding

$$U_{\mathbf{x}_m}(\boldsymbol{\alpha}_m) = \sum_{\boldsymbol{\alpha}_k} U_{\mathbf{x}_m\mathbf{x}_k}(\boldsymbol{\alpha}_m, \boldsymbol{\alpha}_k). \quad (22)$$

Independence. Let \mathbf{x}_m and \mathbf{x}_k be disjoint subgroups of \mathbf{x}_n . These subgroups are *praxeologically independent* if neither subgroup influences the other; that is,

$$U_{\mathbf{x}_m\mathbf{x}_k}(\boldsymbol{\alpha}_m, \boldsymbol{\alpha}_k) = U_{\mathbf{x}_m}(\boldsymbol{\alpha}_m)U_{\mathbf{x}_k}(\boldsymbol{\alpha}_k). \quad (23)$$

The Chain Rule. Let \mathbf{x}_m , \mathbf{x}_k , and \mathbf{x}_l be pairwise disjoint subgroups of \mathbf{x}_n . Then

$$\begin{aligned} U_{\mathbf{x}_m\mathbf{x}_k\mathbf{x}_l}(\boldsymbol{\alpha}_m, \boldsymbol{\alpha}_k, \boldsymbol{\alpha}_l) &= \\ U_{\mathbf{x}_m|\mathbf{x}_k\mathbf{x}_l}(\boldsymbol{\alpha}_m|\boldsymbol{\alpha}_k, \boldsymbol{\alpha}_l)U_{\mathbf{x}_k|\mathbf{x}_l}(\boldsymbol{\alpha}_k|\boldsymbol{\alpha}_l)U_{\mathbf{x}_l}(\boldsymbol{\alpha}_l). \end{aligned} \quad (24)$$

The chain rule is the mechanism by which individual conditional utilities can be aggregated to form the concordance utility. To see, let us first recall that the acyclicity principle ensures that at least one agent possesses a categorical utility. Without loss of generality, let us assume this condition holds for X_1 . Successively applying the chain rule, we obtain

$$\begin{aligned} U_{X_1 \dots X_n}(\mathbf{a}_1, \dots, \mathbf{a}_n) &= u_{X_n}(\mathbf{a}_n|\mathbf{a}_{n-1}, \dots, \mathbf{a}_1) \\ &u_{X_{n-1}}(\mathbf{a}_{n-1}|\mathbf{a}_{n-2}, \dots, \mathbf{a}_1) \dots u_{X_1}(\mathbf{a}_1). \end{aligned} \quad (25)$$

3.2 Utility Networks

Since the influence flows are acyclic, we may represent the multiagent system as a *directed acyclic graph* (DAG). Furthermore, since the utilities that comply with the aggregation theorem possess the mathematical syntax of probability mass functions, the edges of the DAG are conditional utilities. We shall term such a graph a *utility network*, and note that it possesses all of the properties of a Bayesian network, albeit with different semantics (e.g., see (Pearl, 1988; Cowell et al., 1999)).

Definition 7. The *parent set* for X_i , denoted $\text{pa}(X_i)$, is the subgroup of agents whose preferences directly influence X_i . The *child set* of X_i , denoted $\text{ch}(X_i)$, is the subgroup that is directly influenced by X_i . \square

Without loss of generality, we may assume that the vertices of the network are enumerated such that all children of any given node have a higher-numbered index, otherwise the indexing is arbitrary. We may then rewrite (25) as

$$U_{X_1 \dots X_n}(\mathbf{a}_1, \dots, \mathbf{a}_n) = \prod_{i=1}^n u_{X_i | \text{pa}(X_i)}[\mathbf{a}_i | \text{pa}(\mathbf{a}_i)]. \quad (26)$$

where $\text{pa}(\mathbf{a}_i) = \{\mathbf{a}_{i_1}, \dots, \mathbf{a}_{i_{p_i}}\}$ is the joint conjecture corresponding to $\text{pa}(X_i) = \{X_{i_1}, \dots, X_{i_{p_i}}\}$. If $p_i = 0$, then $u_{X_i | \text{pa}(X_i)} = u_{X_i}$, a categorical utility.

To illustrate, consider the network illustrated in Figure 1. X_1 is a root vertex, and possesses a categorical utility u_{X_1} , $\text{pa}(X_2) = \{X_1\}$, and $\text{pa}(X_3) = \{X_1, X_2\}$. The concordance utility is

$$U_{X_1 X_2 X_3}(\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3) = u_{X_1}(\mathbf{a}_1) u_{X_2 | X_1}(\mathbf{a}_2 | \mathbf{a}_1) u_{X_3 | X_1 X_2}(\mathbf{a}_3 | \mathbf{a}_1, \mathbf{a}_2). \quad (27)$$




Figure 1: The DAG for a three-agent system.

If the utilities of all agents are categorical, then no social influence exists, the corresponding DAG has no edges, and, hence, no social bonds are generated. The game reverts to its traditional neoclassical form. The aggregation formula defined by (26) becomes analogous to the creation of the joint distribution of independent random variables as the product of the marginal distributions, and aggregation sheds no additional light on group behavior.




Figure 2: Flow of social influence.

4 CONDITIONAL GAMES

A *conditional game* is a triple $\{\mathbf{x}_n, \mathcal{A}, U_{X_n}\}$ where $\mathbf{x}_n = \{X_1, \dots, X_n\}$ is a group of n agents with product action space $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ and $U_{X_n} = U_{X_1 \dots X_n}$ is a concordance utility. Equivalently, by application of (26), a conditional game can be defined in terms of the conditional utilities $u_{X_i | \text{pa}(X_i)}$, $i = 1, \dots, n$. If all utilities are categorical, a conditional game becomes a conventional game.

With a conditional game, the possibility exists for an expanded notion of rational behavior. To proceed, we observe that, since each agent can control only its own actions, what is of interest is the utility for the group if all agents *make conjectures over, and only over, their own action spaces*.

Definition 8. Consider the concordance utility $U_{X_1 \dots X_n}(\mathbf{a}_1, \dots, \mathbf{a}_n)$. Let a_{ij} denote the j th element of \mathbf{a}_i ; that is, $\mathbf{a}_i = (a_{i1}, \dots, a_{in})$ is X_i 's conjecture. Next, form the action profile (a_{11}, \dots, a_{nn}) by taking the i th element of each X_i 's conjecture, $i = 1, \dots, n$. Now let us sum the concordance utility over all elements of each \mathbf{a}_i except the ii -th elements to form the *social welfare function* for $\{X_1, \dots, X_n\}$, yielding

$$w_{X_1 \dots X_n}(a_{11}, \dots, a_{nn}) = \sum_{\neg a_{11}} \dots \sum_{\neg a_{nn}} U_{X_1 \dots X_n}(\mathbf{a}_1, \dots, \mathbf{a}_n), \quad (28)$$

where $\sum_{\neg a_{ii}}$ means the sum is taken over all a_{ij} except a_{ii} . The *individual welfare function* of X_i is the i -th marginal of $w_{X_1 \dots X_n}$, that is,

$$w_{X_i}(a_{ii}) = \sum_{\neg a_{ii}} w_{X_1 \dots X_n}(a_{11}, \dots, a_{nn}). \quad (29)$$

\square

The social welfare function provides a complete *ex post* description of the relationships between the members of a multiagent system as characterized by their *ex ante* conditional utilities. Unless the members of the system are praxeologically independent, the *ex post* utility is not simply an aggregation of individual utilities, as is the case with classical social choice theory. Rather, it constitutes a meso to macro/micro

propagation of preferences: from the intermediate, or meso, level, derived from local influences between the agents in the form of conditional preferences, up to the global, or macro, level and down to the individual, or micro, level, as illustrated in Figure 2.

We define the *maximum social welfare* solution as

$$(a_1, \dots, a_n) = \arg \max_{\mathbf{a}} w_{X_1 \dots X_n}(a_1, \dots, a_n). \quad (30)$$

Also, the *maximum individual welfare* solution is

$$a_i^\dagger = \arg \max_{a_i} w_{X_i}(a_i). \quad (31)$$

If $a_i^\dagger = a_i$ for all $i \in \{1, \dots, n\}$, the action profile is a *consensus* choice. In general, however, a consensus will not obtain, and negotiation may be required to reach a compromise solution.

The existence of group and individual welfare functions provides a rational basis for meaningful negotiations; namely, that any compromise solution must at least provide each agent with its security level; that is, the maximum guaranteed benefit it could receive regardless of the decisions that others might make. The security level for X_i is

$$s_{X_i} = \max_{a_i} \min_{\neg a_i} \sum_{\neg \mathbf{a}_i} U_{X_1 \dots X_n}(\mathbf{a}_1, \dots, \mathbf{a}_n). \quad (32)$$

In addition to individual benefit, we must also consider benefit to the group. Although a security level, per se, for the group cannot be defined in terms of a guaranteed benefit (after all, the group, as a single entity, does not actually make a choice), a possible rationale is that the benefit to the group it should never be less than the smallest guaranteed benefit to the individuals. This approach is consistent with the principles of justice espoused by (Rawls, 1971), who argues, essentially, that a society as a whole cannot be better off than its least advantaged member. Accordingly, let us define a security level for the group as $s_{X_1 \dots X_n} = \min_i \{s_{X_i}\}/n$, where we divide by the number of agents since the utility for the group involves n players.

Now define the *group negotiation set*

$$\mathcal{N}_{X_1 \dots X_n} = \{\mathbf{a} \in \mathcal{A} : w_{X_1 \dots X_n}(\mathbf{a}) \geq s_{X_1 \dots X_n}\}, \quad (33)$$

the *individual negotiation sets*

$$\mathcal{N}_{X_i} = \{a_i \in \mathcal{A}_i : w_{X_i}(a_i) \geq s_{X_i}\}, \quad i = 1, \dots, n, \quad (34)$$

and the *negotiation rectangle*

$$\mathcal{R}_{X_1 \dots X_n} = \mathcal{N}_{X_1} \times \dots \times \mathcal{N}_{X_n}. \quad (35)$$

Finally, define the *compromise set*

$$C_{X_1 \dots X_n} = \mathcal{N}_{X_1 \dots X_n} \cap \mathcal{R}_{X_1 \dots X_n}. \quad (36)$$

If $C_{X_1 \dots X_n} = \emptyset$, then no rational compromise is possible at the stated security levels. One way to overcome

this impasse is to decrement the security level of the group iteratively by a small amount, thereby gradually enlarging $\mathcal{N}_{X_1 \dots X_n}$ until $C_{X_1 \dots X_n} = \emptyset$. If $C_{X_1 \dots X_n} = \emptyset$ after the maximum reduction in group security has been reached, then no rational compromise is possible, and the system may be considered dysfunctional. Another way to negotiate is for individual members to decrement their security levels iteratively, thereby enlarging the negotiation rectangle.

Once $C_{X_1 \dots X_n} = \emptyset$, any element of this set provides each member, as well as the group, with at least its security level. One possible tie-breaker is

$$\mathbf{a}_c = \arg \max_{\mathbf{a} \in C_{X_1 \dots X_n}} w_{X_1 \dots X_n}(\mathbf{a}), \quad (37)$$

which provides the maximum benefit to the group such that each of its members achieves at least its security level.

5 PARTIAL SOCIATION

Our development thus far has assumed the full generality of conditioning; namely, that (i) a conditional utility depends on the entire conjecture profiles of all of the parents, and (ii) an agent's conditional utility is a function of all elements of the action profile. If maximum complexity is required to define social relationships properly, the full power of conditional game theory may be necessary. It is often the case, however, that the influence relationships are sparse, in that an agent does not condition its preferences on the entire conjecture profiles of its parents. It can also be the case that an agent's utility does not depend upon the entire action profile. To account for such situations, we introduce the notion of *sociation*.

Suppose X_i has $p_i > 0$ parents, denoted $\text{pa}(X_i) = \{X_{i_1}, \dots, X_{i_{p_i}}\}$, with conditional utility $u_{X_i|\text{pa}(X_i)}(\text{pa}(\mathbf{a}_i))$, where $\text{pa}(\mathbf{a}_i) = \{\mathbf{a}_{i_1}, \dots, \mathbf{a}_{i_{p_i}}\}$ is the joint conjecture for $\text{pa}(X_i)$.

Definition 9. A *conjecture subprofile*, for X_{i_k} , denoted $\hat{\mathbf{a}}_{i_k}$, is the subprofile comprising the elements of \mathbf{a}_{i_k} that influence X_i . We then have

$$u_{X_i|\text{pa}(X_i)}[\mathbf{a}_i|\text{pa}(\mathbf{a}_i)] = u_{X_i|\text{pa}(X_i)}[\mathbf{a}_i|\text{pa}(\hat{\mathbf{a}}_i)], \quad (38)$$

where $\text{pa}(\hat{\mathbf{a}}_i) = \{\hat{\mathbf{a}}_{i_1}, \dots, \hat{\mathbf{a}}_{i_{p_i}}\}$. $\{X_1, \dots, X_n\}$ is *completely conjecture sociated* if $\hat{\mathbf{a}}_{i_k} = \mathbf{a}_{i_k}$ for $k = 1, \dots, p_i$ and $i = 1, \dots, n$. It is *completely conjecture dissociated* if $\hat{\mathbf{a}}_{i_k} = a_{i_k}$ for $k = 1, \dots, p_i$ and $i = 1, \dots, n$, in which case, $\text{pa}(\hat{\mathbf{a}}_i) = (a_{i_1}, \dots, a_{i_{p_i}})$. Otherwise, the group is *partially conjecture sociated*. \square

Definition 10. A *utility subprofile*, denoted $\tilde{\mathbf{a}}_i$, comprises the subprofile of \mathbf{a}_i that affects X_i 's utility. We

then have

$$u_{x_i|pa(x_i)}[\mathbf{a}_i|pa(\mathbf{a}_i)] = \tilde{u}_{x_i|pa(x_i)}[\tilde{\mathbf{a}}_i|pa(\hat{\mathbf{a}}_i)], \quad (39)$$

where \tilde{u} denotes u with the dissociated arguments removed. $\{X_1, \dots, X_n\}$ is *completely utility sociated* if $\tilde{\mathbf{a}}_j = \mathbf{a}_j$ for $i = 1, \dots, n$. It is *completely utility dissociated* if $\tilde{\mathbf{a}}_i = a_i$ for $i = 1, \dots, n$, in which case

$$u_{x_i|pa(x_i)}[\mathbf{a}_i|pa(\mathbf{a}_i)] = \tilde{u}_{x_i|pa(x_i)}[a_i|pa(\hat{\mathbf{a}}_i)]. \quad (40)$$

Otherwise, the group is *partially utility sociated*. \square

Definition 11. A group $\{X_1, \dots, X_n\}$ is *completely dissociated* if it is both completely conjecture dissociated and completely utility dissociated, in which case, $pa(\mathbf{a}_i) = pa(a_i) = (a_{i_1}, \dots, a_{i_{p_i}})$, the profile of conjecture actions of the members of $pa(X_i)$. \square

For a partially sociated system, the concordance utility assumes the form

$$U_{X_1 \dots X_n}(\mathbf{a}_1, \dots, \mathbf{a}_n) = \tilde{U}_{X_1 \dots X_n}(\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_n) \quad (41)$$

$$= \prod_{i=1}^n \tilde{u}_{x_i|pa(x_i)}[\tilde{\mathbf{a}}_i|pa(\hat{\mathbf{a}}_i)], \quad (42)$$

where \tilde{U} is U with the dissociated arguments removed. For a completely dissociated group, the concordance utility coincides with the social welfare function and assumes the form

$$w_{X_1 \dots X_n}(a_1, \dots, a_n) = \prod_{i=1}^n \tilde{u}_{x_i|pa(x_i)}[a_i|pa(a_i)]. \quad (43)$$

Example 3. Let us now reconsider the automobile buying example introduced in Example 1. We shall assume that the influence flows are as depicted in Figure 1, with the corresponding concordance utility of the form expressed by (27), yielding

$$\begin{aligned} & U_{x_1 x_2 x_3}[(a_{11}, a_{12}, a_{13}), (a_{21}, a_{22}, a_{23}), (a_{31}, a_{32}, a_{33})] = \\ & u_{x_1}(a_{11}, a_{12}, a_{13}) u_{x_2|x_1}(a_{21}, a_{22}, a_{23} | a_{11}, a_{12}, a_{13}) \\ & u_{x_3|x_1 x_2}[a_{31}, a_{32}, a_{33} | (a_{11}, a_{12}, a_{13}), (a_{21}, a_{22}, a_{23})]. \quad (44) \end{aligned}$$

If all agents have opinions about the model, manufacturer, and color attributes, then each agent must specify utility valuations for each of the eight possible outcomes. Thus, X_1 would specify eight utility values. X_2 , however, would need to define eight conditional utilities, each with eight valuations, and thus would make 64 utility specifications. Finally, X_3 would define 64 conditional utilities, each with eight valuations, and thus would make 512 utility specifications — a formidable task.

Now let us suppose that X_1 is concerned only about the model and manufacturer, but has no opinion about the color. Also, we assume that X_2 is concerned only about the manufacturer given X_1 's conjecture about the model and manufacturer. Finally, let us assume that X_3 is concerned only about the color given X_1 's conjecture about the model and X_2 's conjecture about the manufacturer.

As a result of these simplifications, we see that X_1 is partially utility sociated, thus $\tilde{\mathbf{a}}_1 = (a_{11}, a_{12})$. We also see that X_2 is completely utility dissociated and partially conjecture sociated, hence $\tilde{\mathbf{a}}_2 = a_{22}$ and $pa(a_{22}) = (a_{11}, a_{12})$. Finally, X_3 is also completely utility dissociated and partially conjecture sociated, thus $\tilde{\mathbf{a}}_3 = a_{33}$ and $pa(a_{33}) = (a_{11}, a_{22})$. Thus, the concordance utility simplifies to

$$\begin{aligned} \tilde{U}_{x_1 x_2 x_3}(a_{11}, a_{12}, a_{22}, a_{33}) &= \tilde{u}_{x_1}(a_{11}, a_{12}) \\ & \tilde{u}_{x_2|x_1}(a_{22} | a_{11}, a_{12}) \tilde{u}_{x_3|x_1 x_2}(a_{33} | a_{11}, a_{22}). \quad (45) \end{aligned}$$

Under these simplifications, we see that X_1 need only make four specifications when defining \tilde{u}_{x_1} , X_2 need make two specifications for each of X_1 's four specifications, resulting in eight specifications when defining $\tilde{u}_{x_2|x_1}$, and X_3 need make two specifications for each of the four joint specifications of X_1 and X_2 , resulting in eight specifications when defining $\tilde{u}_{x_3|x_1 x_2}$, yielding a grand total of 20 utility specifications — a considerable reduction in complexity from the 584 specifications required under the condition of complete sociation (and even less than the 24 specifications needed to define categorical utilities).

Tables 1, 2, and 3 respectively tabulate X_1 's categorical utility, X_2 's conditional utilities given X_1 's conjectures, and X_3 's conditional utilities given the conjectures for X_1 and X_2 .

Table 1: The categorical utility $\tilde{u}_{x_1}(a_{11}, a_{12})$.

	a_{12}	
a_{11}	D	F
C	0.1	0.4
S	0.3	0.2

Table 2: The conditional utility $\tilde{u}_{x_2|x_1}(a_{22} | a_{11}, a_{12})$.

	(a_{11}, a_{12})			
a_{22}	(C, D)	(C, F)	(S, D)	(S, F)
D	0.3	0.5	0.6	0.4
F	0.7	0.5	0.4	0.6

The social welfare function $w_{x_1 x_2 x_3}$ is illustrated in Table 4, and the individual welfare functions are as

Table 3: The conditional utility $\tilde{u}_{x_3|x_1x_2}(a_{33}|a_{11}, a_{22})$.

a_{11}	(a_{11}, a_{12})			
	(C, D)	(C, F)	(S, D)	(S, F)
R	0.1	0.3	0.5	0.8
G	0.9	0.7	0.5	0.2

follows.

$$w_{x_1}(C) = 0.5 \quad w_{x_1}(S) = 0.5 \quad (46)$$

$$w_{x_2}(D) = 0.49 \quad w_{x_2}(F) = 0.51 \quad (47)$$

$$w_{x_3}(R) = 0.426 \quad w_{x_3}(G) = 0.574. \quad (48)$$

 Table 4: The social welfare function $w_{x_1x_2x_3}(a_{11}, a_{22}, a_{33})$.

a_{11}	(a_{22}, a_{33})			
	(D, R)	(D, G)	(F, R)	(F, G)
C	0.023	0.207	0.081	0.189
S	0.130	0.130	0.192	0.048

The group negotiation set is $\mathcal{N}_{x_1x_2x_3} = \{(C, D, G), (C, F, G), (S, F, R)\}$ and the negotiation rectangle is $\mathcal{R}_{x_1x_2x_3} = \{(C, F, G), (S, F, G)\}$, yielding the compromise set $\mathcal{C}_{x_1x_2x_3} = \{(C, F, G)\}$: a green foreign-made convertible. \square

6 CONCLUSIONS

As acknowledged by many decision theorists (Arrow, 1986; Luce and Raiffa, 1957; Shubik, 1982), neoclassical game theory is an appropriate model for competitive and market-driven scenarios, but it offers limited capacity for the design and synthesis of multiagent systems that are intended to cooperate, compromise, and negotiate.

This paper (i) presents a principle-based extension to neoclassical game theory that replaces categorical utilities with conditional utilities that encode the social influence relationships that exist among the agents; (ii) develops notions of rational multiagent decision making to define rational behavior simultaneously for groups and for individuals; and (iii) addresses computational complexity by maximally exploiting influence sparseness among the agents. Conditional game theory provides a powerful framework within which to design and synthesize cooperative multiagent systems.

REFERENCES

- Arrow, K. J. (1951). *Social Choice and Individual Values*. John Wiley, New York. 2nd ed. 1963.
- Arrow, K. J. (1986). Rationality of self and others in an economic system. In Hogarth, R. M. and Reder, M. W., editors, *Rational Choice*. University of Chicago Press, Chicago.
- Camerer, C. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton Univ. Press, Princeton, NJ.
- Cowell, R. G., Dawid, A. P., Lauritzen, S. L., and Spiegelhalter, D. J. (1999). *Probabilistic Networks and Expert Systems*. Springer Verlag, New York, NY.
- Cox, R. T. (1946). Probability, frequency, and reasonable expectation. *American Journal of Physics*, 14:1–13.
- Debreu, G. (1959). *Theory of Value*. Yale University Press, New Haven, CT.
- Friedman, M. (1961). *Price theory*. Aldine Press, Chicago, IL.
- Jaynes, E. T. (2003). *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge.
- Keeney, R. L. and Raiffa, H. (1993). *Decisions with Multiple Objectives*. Cambridge University Press, Cambridge, UK. First published by John Wiley & Sons, 1976.
- Luce, R. D. and Raiffa, H. (1957). *Games and Decisions*. John Wiley, New York.
- Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V., editors (2007). *Algorithmic Game Theory*. Cambridge Univ. Press, Cambridge, UK.
- Parsons, S., Gmytrasiewicz, P., and Wooldridge, M., editors (2002). *Game Theory and Decision Theory in Agent-Based Systems*. Kluwer, Boston.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, CA.
- Rawls, J. B. (1971). *A Theory of Justice*. Harvard University Press, Cambridge, MA.
- Shoham, Y. and Leyton-Brown, K. (2009). *Multiagent Systems*. Cambridge University Press, Cambridge, UK.
- Shubik, M. (1982). *Game Theory in the Social Sciences*. MIT Press, Cambridge, MA.
- Vlassis, N., editor (2007). *A Concise Introduction to Multiagent Systems and Distributed Artificial Intelligence*. Morgan & Claypool Publishers, San Rafael, CA.
- Weiss, G., editor (1999). *Multiagent Systems*. MIT Press, Cambridge, MA.