

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 16-07-2008		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 1-Oct-2007 - 30-Jun-2008	
4. TITLE AND SUBTITLE Socially Coherent Negotiation for Cooperative Multiagent Systems			5a. CONTRACT NUMBER W911NF-07-1-0650		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 611102		
6. AUTHORS Wynn C. Stirling			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES Brigham Young University Brigham Young University C-100 BNSN Provo, UT 84602 -			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 53352-MA-II.2		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT This project presents a new mathematical formalization for the design of cooperative multiagent systems. Conventional decision-making concepts focus largely on the foundational premise of individual rationality: the doctrine that the preferences of each participant are concerned with its own welfare regardless of the effect on others. This solution concept may be inadequate for decision scenarios where notions of cooperation, compromise, and negotiation are critical to the success of the system. Cooperation can be enhanced if the system possesses the following properties: (a) a notion of sociality such that the sphere of concern of an individual extends beyond the self, (b) a notion					
15. SUBJECT TERMS Cooperation, Negotiation, Coherence, Conditional Preferences					
16. SECURITY CLASSIFICATION OF:		17. LIMITATION OF ABSTRACT		15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	SAR		Wynn Stirling
					19b. TELEPHONE NUMBER 801-422-7669

## Report Title

### Socially Coherent Negotiation for Cooperative Multiagent Systems

#### ABSTRACT

This project presents a new mathematical formalization for the design of cooperative multiagent systems. Conventional decision-making concepts focus largely on the foundational premise of individual rationality: the doctrine that the preferences of each participant are concerned with its own welfare regardless of the effect on others. This solution concept may be inadequate for decision scenarios where notions of cooperation, compromise, and negotiation are critical to the success of the system. Cooperation can be enhanced if the system possesses the following properties: (a) a notion of sociality such that the sphere of concern of an individual extends beyond the self, (b) a notion

of coherence such that the interests of no agent are categorically subjugated to the interests of the group, and (c) the classical solution concept of optimization is replaced by a more socially accommodating concept of being good enough, or satisficing. These properties provide a mathematical framework within which it is possible to characterize both group and individual interests and thereby to define a multiagent decision that is satisfactory to the group as a whole and to each of its members.

---

#### List of papers submitted or published that acknowledge ARO support during this reporting period. List the papers, including journal references, in the following categories:

##### (a) Papers published in peer-reviewed journals (N/A for none)

Matthew Nokleby and Wynn Stirling, "Attitude Adaptation in Satisficing Games", to appear in IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics.

Number of Papers published in peer-reviewed journals: 1.00

---

##### (b) Papers published in non-peer-reviewed journals or in conference proceedings (N/A for none)

Number of Papers published in non peer-reviewed journals: 0.00

---

##### (c) Presentations

Number of Presentations: 0.00

---

##### Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts): 0

---

##### Peer-Reviewed Conference Proceeding publications (other than abstracts):

Wynn C. Stirling and Matthew S. Nokleby, "Rational Coordination Under Risk: Coherence and the Nash Bargain," Proceedings of the 2008 Conference on Automation Science and Engineering, 2008.

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts): 1

---

##### (d) Manuscripts

Wynn C. Stirling and Matthew S. Nokleby, "Satisficing Coordination and Social Welfare for Robotic Societies", Submitted to the International Journal of Social Robotics.

Number of Manuscripts: 1.00

---

**Number of Inventions:**

**Graduate Students**

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
Matthew Nokleby	0.50
Yabing Luo	1.00
<b>FTE Equivalent:</b>	<b>1.50</b>
<b>Total Number:</b>	<b>2</b>

**Names of Post Doctorates**

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
<b>FTE Equivalent:</b>	
<b>Total Number:</b>	

**Names of Faculty Supported**

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	National Academy Member
Wynn Stirling	0.00	No
<b>FTE Equivalent:</b>	<b>0.00</b>	
<b>Total Number:</b>	<b>1</b>	

**Names of Under Graduate students supported**

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
Steffanie Kuehn	0.50
Steven Flygare	0.50
<b>FTE Equivalent:</b>	<b>1.00</b>
<b>Total Number:</b>	<b>2</b>

**Student Metrics**

This section only applies to graduating undergraduates supported by this agreement in this reporting period

- The number of undergraduates funded by this agreement who graduated during this period: ..... 0.00
- The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00
- The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00
- Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00
- Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00
- The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense ..... 0.00
- The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields: ..... 0.00

**Names of Personnel receiving masters degrees**

<u>NAME</u>	
Matthew Nokleby	
<b>Total Number:</b>	<b>1</b>

---

**Names of personnel receiving PhDs**

NAME

**Total Number:**

---

**Names of other research staff**

NAME

PERCENT SUPPORTED

**FTE Equivalent:**

**Total Number:**

---

**Sub Contractors (DD882)**

**Inventions (DD882)**

**Final Report for Grant Number  
W911NF-07-1-0650**

**Socially Coherent Negotiation for  
Cooperative Multiagent Systems**

Principal Investigator: Wynn C. Stirling  
Electrical and Computer Engineering Department  
459CB Brigham Young University, Provo, Utah 84602

# Contents

<b>1</b>	<b>Problem Statement</b>	<b>3</b>
<b>2</b>	<b>Summary of Results: Negotiations</b>	<b>5</b>
2.1	Conditional Utilities . . . . .	5
2.2	Coherence . . . . .	6
2.3	Satisficing . . . . .	8
2.4	Negotiation . . . . .	10
2.4.1	Optimal Compromise . . . . .	10
2.4.2	Nash Bargains . . . . .	12
<b>3</b>	<b>Summary of Results: Attitude Adaptation</b>	<b>12</b>

## List of Figures

1	An influence network for a five-member multiagent system. . . . .	6
2	A praxi-epistemic network for a four-member multiagent system. . . . .	8
3	A Satisficing network for a four-member multiagent system. . . . .	9

## List of Appendices

A	Rational Coordination Under Risk: Coherence and the Nash Bargain
B	Satisficing Coordination and Social Welfare for Robotic Societies
C	Attitude Adaptation in Satisficing Games

# 1 Problem Statement

The objective of this research grant is to develop a new mathematical formalism for cooperative multiagent system synthesis that is explicitly designed to accommodate sophisticated social relationships such as negotiation and compromise. To do so, we focus on the foundational assumptions that undergird multiagent decision making, and challenge the adequacy of the classical assumptions for the design of socially sophisticated artificial multiagent systems.

Classical approaches to multiagent decision making, such as von Neumann-Morgenstern game theory [19] and social choice theory [1, 4, 10], are founded upon two key assumptions.

- It is assumed that each member of a multiagent system possesses a well-defined total preference ordering over all of the feasible actions of the collective. Such preference orderings are *categorical* in the sense that they are unconditional — once defined, the preference orderings are immutable and are viewed as the selfish desires of the members even if, ostensibly, they express some notion of altruism by substituting the preferences of others for one’s own.
- It is assumed that each member will seek to maximize benefit to itself, regardless of the effect doing so has on other members.

These two assumptions form the basis of the classical doctrine of *individual rationality*. Perhaps the most well-known game-theoretic instantiation of this doctrine is the concept of Nash equilibria: a state of mutual constrained optimization for all members in the sense that any member who unilaterally deviates from an equilibrium state will be less satisfied. Individual rationality is appropriate for competitive social situations, but does not provide a framework within which sophisticated social relationships can be easily modeled and, hence, is not well suited as a model for cooperative multiagent systems.

The social choice solution to the multiagent system decision problem is to combine, or aggregate, the utilities of each individual to form a social welfare function to be maximized. As with the game-theoretic approach, however, classical social choice approaches use categorical utilities, and do not account for social relationships among the individuals.

We introduce a significant departure from classical approaches to multiagent decision making. Our approach differs from the classical formulation in three major ways.

1. *Conditioning.* We relax the assumption that each member of a multiagent system possesses a total preference ordering over all feasible actions of the collective. We assume, instead, that members of a multiagent system are able to modulate their preferences as a function of the preferences of others. To account for this change, we replace categorical utilities with *conditional utilities* that are designed to express the preferences of each individual as a function of the preferences of others, as appropriate.
2. *Coherence.* We invoke a weak notion of equity by assuming that a minimal condition for meaningful negotiations to take place is for each member of the system to have a “seat at the table” in the sense that its interests at least have a chance of being taken seriously by the group as a whole. Stated more formally, we require the system to be *coherent*, meaning that no individual can be *categorically subjugated* in the sense that every action that is acceptable to the collective requires the individual to be

disadvantaged. Such an individual would effectively be disenfranchised, and would not be in a position to undertake meaningful negotiations. This structure does not eliminate hierarchical systems; rather, it simply means that, even in master/slave relationships, the possibility exists (but not the guarantee) that the slave’s preferences can be acceptable to the master. For a slave to be categorically subjugated, every action that is good for the master would have to be bad for the slave.

3. *Satisficing*. We replace the notion of optimization with a concept of being adequate, or good enough. The terminology we use for this concept is *satisficing*. This term was initially introduced by Simon [11–13], who addressed the question of how a decision maker might choose in the presence of informational or computational limitations. Simon’s approach is to seek an optimal choice, but to terminate searching and once the decision maker’s aspiration level has been met. A slightly different notion of satisficing is to accept the best solution so far obtained, once the cost of continuing to search exceeds the expected improvement in value were the search to continue. Many other variations of this concept have appeared in the literature and it is not the intent of this report to review them in detail. Suffice it to say, however, that all of these approaches view satisficing as a species of bounded rationality: one settles for a solution that is deemed to be “good enough,” but which is not necessarily, and usually not, optimal in any meaningful sense. Satisficing *à la* Simon is an heuristic approximation to the ideal of being best (and is only constrained from achieving this ideal by practical limitations). The concept of satisficing we employ, however, differs from the afore-mentioned notions in several important ways.

- (a) In contrast to satisficing as advanced by Simon and others, it is not heuristic; rather, it is a concept that is as mathematically formalized and precise as is the notion of optimization.
- (b) It treats being good enough as the ideal (rather than an approximation) — it is *not* a species of bounded rationality.
- (c) It naturally extends to the multi-agent case, thereby providing a natural framework for multi-agent decision making.
- (d) It readily accommodates the extension of interests beyond the self, thereby accommodating more sophisticated social relationships than self-interest affords.

We retain the term “satisfice” because, even though our approach is not heuristic, we nevertheless seek solutions that are good enough, with the essential difference being that we provide a non-heuristic and mathematically precise definition of what it means to be good enough.

While optimization is intrinsically an individual concept (if a group is to optimize, it must act as an individual), satisficing, as we define it, is a social concept: what is best for you may be incompatible with what is best for me, but what is good enough for you can also be good enough for me, provided we each have some flexibility regarding what we view as good enough.

To motivate our concept of satisficing, we note that humans often invoke a systematic approach to decision making that, while still based on quantitative measures of performance, does not correspond to optimization. In the vernacular, the optimization paradigm corresponds to seeking “the best and only the best” solution. Also common, however, is the paradigm of “getting your money’s worth,” or ensuring that the benefits are greater than the costs. This notion of being good enough is the satisficing paradigm that we advocate. A comprehensive introduction to this perspective can be found in [15].

## 2 Summary of Results: Negotiations

A multiagent system comprises a collective of agents who must work cohesively to accomplish some fundamental objective. Typically, however, such systems are mixed-motive, in the sense that the interests of all individuals will not all coincide perfectly; hence, opportunities for both cooperation and competition will exist. The major contribution of this study is the development of a mathematical framework that accommodates both cooperative and competitive aspects of a multiagent system. In this section we briefly describe the three main components of our theory (conditioning, coherence, and satisficing) and show how they are used to define a framework within which to conduct negotiations. Publications arising from this research are [16, 17], which are included in Appendices A and B, respectively.

### 2.1 Conditional Utilities

Let  $\{X_1, \dots, X_n\}$ ,  $n \geq 2$ , denote a group of autonomous decision makers. Let  $\mathcal{A}_i$  denote a finite set of feasible actions available to  $X_i$ ,  $i = 1, \dots, n$ , let  $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$  denote the product action space, and let  $\mathbf{a} = (a_1, \dots, a_n)$  denote the action profile that obtains when each  $X_i$  instantiates  $a_i \in \mathcal{A}_i$ . A *categorical utility* for  $X_i$  is a mapping  $u_{X_i}: \mathcal{A} \rightarrow \mathbb{R}$  such that  $u_{X_i}(\mathbf{a}) > u_{X_i}(\mathbf{a}')$  if  $X_i$  strictly prefers  $\mathbf{a}$  to  $\mathbf{a}'$  and  $u_{X_i}(\mathbf{a}) = u_{X_i}(\mathbf{a}')$  if  $X_i$  is indifferent between  $\mathbf{a}$  and  $\mathbf{a}'$ . Classical decision-theoretic approaches, such as von Neumann-Morgenstern game theory, employ categorical utilities (i.e, they are the payoffs of a game).

A conditional utility differs from a categorical utility in that it is a hypothetical, rather than a concrete, expression. Before formally defining a conditional utility, we must first introduce the notion of a commitment. In the interest of clarity, we temporarily restrict our discussion to a two-agent system  $(X_1, X_2)$ . Now suppose, from the point of view of  $X_2$  that  $X_1$  views  $\mathbf{a} = (a_1, a_2)$  to be its most preferred joint action. We shall call this hypothetical constraint on  $X_1$  a *commitment*. A commitment, therefore, represents the antecedent of a hypothetical proposition, the consequent of which is a conditional utility denoted  $u_{X_2|X_1}(\cdot|\mathbf{a})$ . More generally, for an  $n$ -agent system, if  $X_i$  is influenced by the  $p_i$  element sub-collective  $\{X_{i_1}, \dots, X_{i_{p_i}}\}$ , then the conditional utility of  $X_i$  is of the form  $u_{X_i|X_{i_1}, \dots, X_{i_{p_i}}}(\mathbf{a}_i|\mathbf{a}_{i_1}, \dots, \mathbf{a}_{i_{p_i}})$ .

In contrast to categorical utilities, a conditional utility expresses  $X_i$ 's preferences over  $\mathcal{A}$  given the commitments of all other agents that influence it. In the most general case, each agent would be influenced by every other agent, but it is often the case that agents will be most heavily influenced by their immediate neighbors. For example, hierarchical organizations are organized so that superiors influence subordinates. Other multiagent systems are

organized into small loosely connected clusters. Thus, although a fully connected system is possible, many interesting multiagent systems are relatively sparsely connected. With this project we will focus on systems whose influence relationships can be represented graphically with a directed acyclic graph, or DAG. The vertices of the graph represent the various members of the collective, and the edges represent the conditional utilities. Figure 1 illustrates an influence network for a five-member multiagent system. We see that  $X_1$  influences  $X_2$ , who in turn influences  $X_3$  and  $X_5$ .  $X_5$  is also influenced by  $X_4$ . Finally,  $X_3$  is influenced by  $X_2$  and  $X_5$ . Since  $X_1$  and  $X_4$  are root vertices, they possess categorical utilities  $u_{x_1}$  and  $u_{x_4}$  (not shown on the graph). It should be noted that, if all utilities were categorical, then this graph would have no edges — it would consist of  $n$  isolated vertices, each possessing a categorical utility. With this more general model, only the root vertices possess categorical utilities, all others possess conditional utilities.




Figure 1: An influence network for a five-member multiagent system.

## 2.2 Coherence

A collective that possesses the property that none of its members can be categorically subjugated is said to be *coherent*. We have appropriated this term from probability theory, since the notion of avoiding sure subjugation is completely analogous to the probabilistic notion of avoiding sure loss. As the Dutch Book Theorem and its converse establish, the only way for a gambler to avoid a situation of sure loss (his payoff is less than his stake regardless of the outcome), is for him to place bets in accordance with the axioms of probability. One of the key results of our investigation is to demonstrate that, similarly, the only way for a member of a collective to avoid categorical subjugation is for all utilities to possess the mathematical structure of conditional or marginal mass functions. Under this constraint, the edges in Figure 1 are conditional mass functions, and the graph therefore possesses the mathematical structure of a Bayesian network (albeit with different semantics). Conventionally, Bayesian networks operate in the epistemological<sup>1</sup> domain; that is, involving random phenomena. To distinguish between the conventional probabilistic application of Bayesian networks and our praxeological<sup>2</sup> application, we shall refer to networks such as is depicted in Figure 1 as *praxeic networks*.

<sup>1</sup>Epistemology relates to the categorization of propositions in terms of knowledge and belief.

<sup>2</sup>Praxeology relates to the categorization of actions in terms of their effectiveness and efficiency.

The DAG structure, coupled with the fact that the edges are mass functions, permits a natural way to aggregate the preference orderings of the individuals to form a group preference ordering. As is well known from Bayesian network theory, the so-called Markov condition, which states that nondescendent nonparents of a vertex have no influence on the vertex, given the state of its parent vertices [2]. Accordingly, just as the multivariate probability mass function is formed as the product of the conditional and marginal mass functions of a Bayesian network, the multiagent utility of collective is formed as the product of the conditional and marginal utilities of the praxeic network. Thus, the multiagent utility associated with the network illustrated in Figure 1 is

$$u_{X_1 X_2 X_3 X_4 X_5}(\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3, \mathbf{a}_4, \mathbf{a}_5) = u_{X_1}(\mathbf{a}_1) u_{X_2|X_1}(\mathbf{a}_2|\mathbf{a}_1) u_{X_3|X_2 X_5}(\mathbf{a}_3|\mathbf{a}_2, \mathbf{a}_5) u_{X_4}(\mathbf{a}_4) u_{X_5|X_2 X_4}(\mathbf{a}_5|\mathbf{a}_2, \mathbf{a}_4). \quad (1)$$

More generally, let  $\text{pa}(X_i) = (X_{i_1}, \dots, X_{i_{p_i}})$  denote the  $p_i$  parents of  $X_i$ , and let  $u_{X_i|\text{pa}(X_i)}$  denote the conditional utility of  $X_i$  given its parents. If a vertex has no parents, then the conditional utility becomes a categorical utility; that is,  $u_{X_i|\text{pa}(X_i)} = u_{X_i}$  if  $\text{pa}(X_i) = \emptyset$ . The multiagent utility then becomes

$$u_{X_1 \dots X_n}(\mathbf{a}_1, \dots, \mathbf{a}_n) = \prod_{i=1}^n u_{X_i|\text{pa}(X_i)}[\mathbf{a}_i | \text{cp}(X_i)], \quad (2)$$

where  $\text{cp}(X_i) = \{\mathbf{a}_{i_1}, \dots, \mathbf{a}_{i_{p_i}}\}$ .

The probabilistic syntax of the utilities constructed in this way provides a natural way to link the praxeological and epistemological aspects of a decision problem into a common unifying framework. To illustrate, let us modify the network in Figure 1 by replacing vertex  $X_5$  with a random variable,  $\theta$ , as illustrated in Figure 2, where  $u_{X_3|X_2\theta}$  is a utility conditioned on the commitment of  $X_2$  and the value that  $\theta$  assumes and  $p_{\theta|X_2 X_4}$  is a probability mass function conditioned on the commitments of  $X_2$  and  $X_4$ . The resulting multiagent utility is of the form

$$u_{X_1 X_2 X_3 X_4 \theta}(\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3, \mathbf{a}_4, \vartheta) = u_{X_1}(\mathbf{a}_1) u_{X_2|X_1}(\mathbf{a}_2|\mathbf{a}_1) u_{X_3|X_2 \theta}(\mathbf{a}_3|\mathbf{a}_2, \vartheta) u_{X_4}(\mathbf{a}_4) p_{\theta|X_2 X_4}(\vartheta, \mathbf{a}_2, \mathbf{a}_4), \quad (3)$$

where  $\vartheta$  is the value assumed by the random variable  $\theta$ . The expected utility is then obtained by averaging over the values that  $\theta$  may assume, yielding

$$\hat{u}_{X_1 X_2 X_3 X_4}(\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3, \mathbf{a}_4) = \sum_{\vartheta} u_{X_1 X_2 X_3 X_4 \theta}(\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3, \mathbf{a}_4, \vartheta). \quad (4)$$

This result extends to the general  $n$ -dimensional case in the obvious way.

The most general formulation of this framework assumes that each individual's utility is defined over the product action space  $\mathcal{A}$ , given the commitments of each of its parents to action profiles in  $\mathcal{A}$ , as is presented above. For many applications, however, this full generality is not necessary, since it is often reasonable to assume that agents' utilities are defined with respect to their own actions, given commitments by others to only their own actions. Thus, we introduce the notion of decoupling.




Figure 2: A praxi-epistemic network for a four-member multiagent system.

**Definition 1** A multiagent system is conditionally decoupled if the conditional preference of each agent is a function only of its own actions, given the commitments of its parents to their own actions.

For a decoupled multiagent system, (3) becomes

$$u_{X_1 X_2 X_3 X_4 \theta}(a_1, a_2, a_3, a_4, \vartheta) = u_{X_1}(a_1) u_{X_2|X_1}(a_2|a_1) u_{X_3|X_2\theta}(a_3|a_2, \vartheta) u_{X_4}(a_4) p_{\theta|X_2 X_4}(\vartheta, a_2, a_4). \quad (5)$$

In general, the individual conditional utilities are of the form

$$u_{X_i|\text{pa}(X_i)}[a_i | \text{cp}(X_i)] = u_{X_i|\text{pa}(X_i)}(a_i | a_{i_1}, \dots, a_{i_{p_i}}) \quad (6)$$

where the action sub-profile  $\{a_{i_1}, \dots, a_{i_{p_i}}\}$  corresponds to the commitments by  $\text{pa}(X_i) = \{X_{i_1}, \dots, X_{i_{p_i}}\}$ , and the multiagent utility is of the form

$$u_{X_1 \dots X_n}(a_1, \dots, a_n) = \prod_{i=1}^n u_{X_i|\text{pa}(X_i)}(a_i | a_{i_1}, \dots, a_{i_{p_i}}), \quad (7)$$

For the remainder of this report we focus on decoupled systems.

## 2.3 Satisficing

Even though optimization is often taken as the *sine qua non* of formalized decision-making procedures, humans are often wont to evaluate propositions in terms of the upside versus the downside, the pluses versus the minuses, the benefits versus the costs, and so forth. One of the important omissions in the extant literature is a systematic formal treatment of this mode of evaluating possible choices. An important result of earlier research by the principal investigator is the introduction of a formalized mathematical treatment of this alternative mode of decision making. It should be noted that this approach has been inspired by the work of the philosopher Isaac Levi [6], who proposed a novel way, using the mathematics of probability theory, to improve one's knowledge. In [15], the principal investigator applied Levi's approach to the praxeological domain and extended it to the multiagent case.

Conventional utilities combine all costs and benefits of taking action into a single function. One common approach is to define utility as a linear combination of those aspects of taking

action that relate to the effectiveness (benefits) of taking an action and those aspects that relate to the inefficiency (costs) of taking the action. In practice, the weights of these two facets of taking action become tuning parameters to facilitate the design of a system that provides acceptable performance (at the end of the day, even optimization is subjective).

Many theorists (e.g., [1, 3, 5, 7]) have argued, however, that it is unwise to aggregate conflicting interests into a single preference ordering. Some have asserted that in a social setting individuals have multiple facets, as defined by Steedman and Krause [14], who maintain that an agent, although an indivisible unit, nevertheless is capable of considering its choices from different points of view, and that separate utilities may be defined to correspond to each facet of an individual. A natural way to classify attributes is according to their effectiveness and efficiency. Each individual may be viewed as being composed of two facets: the *selecting facet*, which evaluates actions in terms of effectiveness toward pursuing objectives without concern for efficiency, and the *rejecting facet*, who evaluates actions in terms of efficiency with respect to consuming resources without concern for effectiveness. We shall view these selecting and rejecting facets as the “atoms” of the system. Notationally, we define  $S_i$  and  $R_i$  as the selecting and rejecting facets, respectively, of  $X_i$

Accordingly, we define separate utilities for the selecting facet and the rejecting facet. In accordance with the conditioning and coherence properties, these utilities are conditional mass functions. Each agent has a unit of selecting utility to apportion among the feasible actions and a unit of rejecting inutility also to apportion. An  $n$ -agent system thus comprises  $2n$  atoms:  $n$  selecting facets and  $n$  rejecting facets, and the graph of such a system comprises  $2n$  praxeic vertices whose edges are conditional utilities. Figure 3 illustrates a refinement, in terms of the facets, of the influence relationships originally defined by Figure 2. This network reveals more explicitly just how the agents influence each other. We see that  $S_1$  influences  $R_2$ ,  $S_4$  influences  $\theta$ , and so forth. Also, facets  $R_1$  and  $R_4$  are not influenced by any other facets and hence, in addition to  $S_1$ ,  $S_2$ , and  $S_4$ , are root nodes.




Figure 3: A Satisficing network for a four-member multiagent system.

According to the fundamental property of Bayesian networks, we may form the multiagent utility as the product of all marginal and conditional utilities, yielding

$$\begin{aligned}
 u_{S_1 S_2 S_3 S_4 R_1 R_2 R_3 R_4 \theta}(a_1, a_2, a_3, a_4, a'_1, a'_2, a'_3, a'_4, \vartheta) = \\
 u_{S_1}(a_1) u_{S_2}(a_2) u_{S_3|R_2 \theta}(a_3|a'_2, \vartheta) u_{S_4}(a_4) \\
 u_{R_1}(a'_1) u_{R_2|S_1}(a'_2|a_1) u_{R_3|S_2}(a'_3|a_2) u_{R_4}(a'_4) p_{\theta|S_4 R_2}(\vartheta|a_4, a'_2), \quad (8)
 \end{aligned}$$

and the expected utility is

$$\hat{u}_{S_1 S_2 S_3 S_4 R_1 R_2 R_3 R_4}(a_1, a_2, a_3, a_4, a'_1, a'_2, a'_3, a'_4) = \sum_{\vartheta} u_{S_1 S_2 S_3 S_4 R_1 R_2 R_3 R_4 \theta}(a_1, a_2, a_3, a_4, a'_1, a'_2, a'_3, a'_4, \vartheta). \quad (9)$$

This expected utility is called the *aggregation function*. Analogous to the way a joint probability distribution captures all of the interdependencies among multiple random variables, the aggregation function captures all of the inter-relationships among the facets of a multiagent system.

## 2.4 Negotiation

### 2.4.1 Optimal Compromise

The three components of conditioning, coherence, and satisficing provide a framework within which members of multiagent system can negotiate and compromise. The key feature that enables this ability is that the satisficing approach provides a set of acceptable actions, rather than a singleton set comprising the optimal action.

For a non-decoupled system, the utilities are functions of the entire action profile, but for a decoupled system, the utilities are functions of individual actions. With this restriction, the multiagent utility function becomes

$$\hat{u}_{S_1 \dots S_n R_1 \dots R_n}(a_1, \dots, a_n, a'_1, \dots, a'_n) = \prod_{i=1}^n u_{S_j | \text{pa}(S_j)}[a_j | \text{cp}(S_j)] \prod_{j=1}^n u_{R_j | \text{pa}(R_j)}[a'_j | \text{cp}(R_j)], \quad (10)$$

where  $\text{cp}(S_i)$  and  $\text{cp}(R_i)$  denote the commitments by  $\text{pa}(S_i)$  and  $\text{pa}(R_i)$ , respectively.

The corresponding joint selectability and rejectability marginals are given by

$$u_{S_1 \dots S_n}(a_1, \dots, a_n) = \sum_{(a'_1, \dots, a'_n)} u_{S_1 \dots S_n R_1 \dots R_n}(a_1, \dots, a_n, a'_1, \dots, a'_n) \quad (11)$$

and

$$u_{R_1 \dots R_n}(a'_1, \dots, a'_n) = \sum_{(a_1, \dots, a_n)} u_{S_1 \dots S_n R_1 \dots R_n}(a_1, \dots, a_n, a'_1, \dots, a'_n). \quad (12)$$

We may now define a *social welfare function* as

$$W(a_1, \dots, a_n) = u_{S_1 \dots S_n}(a_1, \dots, a_n) - q_G u_{R_1 \dots R_n}(a_1, \dots, a_n) \quad (13)$$

where  $q_G \in [0, 1]$  regulates the threshold for rejecting elements of  $\mathcal{A}$ . Nominally,  $q_G = 1$ , but as we shall see, this parameter serves as negotiating parameter. The *jointly satisficing set* is the set of action profiles that are jointly satisficing for the system as a whole, and is defined as

$$\mathcal{S} = \{(a_1, \dots, a_n) \in \mathcal{A} : W(a_1, \dots, a_n) \geq 0\}. \quad (14)$$

This set, however, does not account for the possibility that the elements of  $\mathcal{S}$  may not be acceptable to all (or any) of the individuals. Thus, we must also compute the individual satisficing sets. To proceed, we must first compute the selectability and rejectability marginals as

$$u_{S_i}(a_i) = \sum_{\neg a_i} u_{S_1 \dots S_n}(a_1, \dots, a_n) \quad (15)$$

and

$$u_{R_i}(a_i) = \sum_{\neg a_i} u_{R_1 \dots R_n}(a_1, \dots, a_n), \quad (16)$$

respectively, where the notation  $\sum_{\neg a_i}$  is the so-called “not sum” notation meaning the sum is taken over all elements except  $a_i$ .

We define the individually satisficing sets as

$$\Sigma_i = \{a_i \in \mathcal{A}_i: u_{S_i}(a_i) - q_i u_{R_i}(a_i)\}, \quad (17)$$

where  $q_i \in [0, 1]$  is  $X_i$ 's individual negotiation index. This set includes all alternatives that are satisficing, or good enough, for  $X_i$  at the given negotiation index. The *satisficing rectangle* is the set of all action profiles such that each component is individually satisficing, and is given by

$$\mathcal{R} = \Sigma_1 \times \dots \times \Sigma_n. \quad (18)$$

The intersection of the jointly satisficing set and the satisficing rectangle yields the *compromise set*, comprising the action profiles that are simultaneously good enough for the group and for each individual.

$$\mathcal{C} = \mathcal{S} \cap \mathcal{R}. \quad (19)$$

If  $\mathcal{C} \neq \emptyset$ , then we may form a *best compromise* as

$$\mathbf{a}^* = \arg \max_{\mathbf{a} \in \mathcal{C}} W(\mathbf{a}). \quad (20)$$

If  $\mathcal{C} = \emptyset$ , then there are no action profiles that are simultaneously good enough for the group and each individual. However, the satisficing approach provides a natural and systematic negotiation framework by which each individual may control the degree to which it is willing to lower its standards in an attempt to reach a compromise. By lowering its  $q_i$ -value incrementally, each  $X_i$  increases the size of its satisficing set. By specifying the increment  $\Delta q_i$  that  $X_i$  is willing to reduce its standards, each participant can control the amount of compromise it is willing to offer others. If enough participants are willing to lower their  $q$ -values sufficiently, it is easy to see that, eventually, the consensus set will be non-empty, and a best compromise can be achieved. Although such negotiations may fail to reach a compromise that is acceptable to all members, the significant aspect of this type of negotiation is that no individual is *a priori* subjugated to the will of the collective in the sense that there is no possibility for that individual's preferences to receive consideration. Thus, every individual can be assured of receiving sufficient benefit, by its own definition, before agreeing to the compromise.

### 2.4.2 Nash Bargains

A bargaining game is a cooperative game in which each participant possesses a *disagreement point* that defines the benefit that is guaranteed to accrue to it if a compromise cannot be reached. A well-known bargaining concept that offers a clear definition of individual acceptability is the Nash bargain [8], which permits each participant to make maximal use of its strategic strength. Let  $d_{X_i}$  denote the disagreement point for  $X_i$ . The *negotiation set*, denoted  $\mathcal{N}$ , is the subset of action profiles such that every participant achieves at least its disagreement point. In terms of categorical utilities, the negotiation set is

$$\mathcal{N} = \{\mathbf{a} \in \mathcal{A}: u_{X_i}(\mathbf{a}) \geq d_{X_i}, i = 1, \dots, n\} \quad (21)$$

and the Nash bargain is

$$\mathbf{a}_N = \arg \max_{\mathbf{a} \in \mathcal{N}} \prod_{i=1}^n [u_{X_i}(\mathbf{a}) - d_{X_i}]. \quad (22)$$

The intuitive interpretation of a Nash bargain is that it defines a fair compromise. It enables each player to take advantage of the strategic strength endowed by its disagreement point. The higher  $X_i$ 's disagreement point, the more action profiles that are unfavorable to it are eliminated.

The structure of (22) suggests that the optimal group solution can be interpreted as a Nash bargain with unilateral utilities replaced by conditional utilities and all disagreement points set to zero. Analogously, therefore, we may define a *conditional Nash bargaining solution*. When decisions are made under certainty, the negotiation set is defined as

$$\mathcal{N} = \{\mathbf{a} \in \mathcal{A}: u_{X_i|pa(X_i)}(\mathbf{a}) | cp(X_i) \geq d_{X_i}, i = 1, \dots, n\}. \quad (23)$$

The conditional Nash bargaining solution is

$$\mathbf{a}_N = \arg \max_{\mathbf{a} \in \mathcal{N}} \prod_{i=1}^n [u_{X_i|pa(X_i)}[\mathbf{a} | cp(X_i)] - d_{X_i}]. \quad (24)$$

## 3 Summary of Results: Attitude Adaptation

An important benefit of the satisficing approach is that cooperation occurs much more readily than under standard utility-maximization. To examine this phenomenon more closely, we studied the emergence of cooperation using evolutionary game theory. Evolutionary game theory [20] studies large populations of players whose reproductive potential is determined by the payoff gained during play. For infinitely large, well-mixed populations, the evolution of the population is described by the *replicator dynamics* [18]. In the simplest case, all players have the same action space  $\mathcal{A}$ , and are paired with one other player each ‘‘round.’’ Let  $x_i(t)$  be the fraction of the population playing strategy  $a_i \in \mathcal{A}$  at time  $t$ . Then, the population shares evolve according to the following system of differential equations:

$$\dot{x}_i(t) = [u(a_i, \mathbf{x}(t)) - u(\mathbf{x}(t), \mathbf{x}(t))]x_i(t), \forall i, \quad (25)$$

where  $u(a_i, \mathbf{x}(t))$  is the expected utility of playing strategy  $a_i$  against a player randomly drawn from the population described by  $x(t)$ , and  $u(\mathbf{x}(t), \mathbf{x}(t))$  is the average expected utility.

Essentially, a strategy's population share grows or shrinks if it fares better or worse than average, respectively. Given appropriate initial conditions, the steady-state of the replicator dynamics is a Nash equilibrium.

To apply evolutionary dynamics to the satisficing case, we note that players' conditional utilities are often expressed in terms of tunable parameters that govern (for example) players' willingness to cooperate or defer to the preferences of others. We term these parameters *attitudes* and study how players might adapt their attitudes in order to increase payoff. Instead of running the replicator dynamics on players' actions (as in the classical case), we run the replicator dynamics on players' attitudes, allowing us to study the ecological fitness of exhibiting a particular attitude. This dynamics leads the players to an *attitude equilibrium*, a point at which no player can improve its payoff by changing its attitudes.

As a concrete example, we focus on the well-studied Stag Hunt game, which involves two players. They can catch a stag but cooperating, but each can catch a (much smaller) hare alone. That is, a player earns maximum payoff if both players cooperate, but risks failure if it attempts to cooperate while the other does not. Under the standard replicator dynamics, the population ends up entirely non-cooperative (hunting hare) unless a significant majority of the population initially hunts stag. So, it is impossible under this framework to evolve a cooperative population from non-cooperation.

We applied satisficing theory to see if we could do any better. We developed a simple satisficing model for the Stag Hunt and applied the replicator dynamics to the players' attitudes. Under the satisficing model, cooperation is significantly easier to achieve than under the standard model. Indeed, the population evolves toward cooperation even when only 10% of the initial population hunts stag. This study is detailed in [9], which is included in Appendix C.

## References

- [1] K. J. Arrow. *Social Choice and Individual Values*. John Wiley, New York, NY, 1951. 2nd ed. 1963.
- [2] F. G. Cozman. Credal networks. *Artificial Intelligence*, 120:199–233, 2000.
- [3] J. Elster, editor. *The Multiple Self*. Cambridge University Press, Cambridge, UK, 1985.
- [4] P. C. Fishburn. *The Theory of Social Choice*. Princeton University Press, Princeton, NJ, 1973.
- [5] J. Harsanyi. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, 63:315, 1955.
- [6] I. Levi. *The Enterprise of Knowledge*. MIT Press, Cambridge, MA, 1980.
- [7] H. Margolis. Dual utilities and rational choice. In J. J. Mansbridge, editor, *Beyond Self-Interest*, chapter 15, pages 239–253. Univ. of Chicago Press, Chicago, IL, 1990.
- [8] J. F. Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.

- [9] M. S. Nokleby and W. C. Stirling. Attitude adaptation in satisficing games. *IEEE Transaction on Systems, Man, and Cybernetics — Part B: Cybernetics*, 2008. To Appear.
- [10] A. K. Sen. *Collective Choice and Social Welfare*. North-Holland, Amsterdam, The Netherlands, 1970.
- [11] H. A. Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, 59:99–118, 1955.
- [12] H. A. Simon. Rational choice and the structure of the environment. *Psychological Review*, 63(2):129–138, 1956.
- [13] H. A. Simon. Theories of decision-making in economics and behavioral science. *Amer. Econ. Rev.*, XLIX:253–283, 1959.
- [14] I Steedman and U. Krause. Goethe’s *Faust*, Arrow’s Possibility Theorem and the individual decision maker. In J. Elster, editor, *The Multiple Self*, chapter 8, pages 197–231. Cambridge Univ. Press, Cambridge, UK, 1985.
- [15] W. C. Stirling. *Satisficing Games and Decision Making: with applications to engineering and computer science*. Cambridge Univ. Press, Cambridge, UK, 2003.
- [16] W. C. Stirling and M. S. Nokleby. Rational coordination under risk: Coherence and the nash bargain. In *Proceedings of Fourth IEEE International Conference on Automation Science and Engineering*, 2008.
- [17] W. C. Stirling and Matthew S. Nokleby. Satisficing coordination and social welfare for robotic societies. Submitted to *The International Journal of Social Robotics*, 2008.
- [18] P. Taylor and L. Jonker. Evolutionarily stable strategies and game dynaics. *Math. Biosc.*, 40(2):146–156, 1978.
- [19] J. von Neumann and O. Morgenstern. *The Theory of Games and Economic Behavior*. Princeton Univ. Press, Princeton, NJ, 1944. (2nd ed., 1947).
- [20] J. W. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1995.

# Rational Coordination Under Risk: Coherence and the Nash Bargain

Wynn C. Stirling and Matthew S. Nokleby *Student Member, IEEE*

**Abstract**—The design of automated multiagent cooperative systems can be greatly facilitated by the use of conditional utilities, which provide each individual the capability of modulating its interests as a function of the interests of others. Perhaps the weakest possible requirement for meaningful coordination is that the group be *coherent*: no individual is required, under all circumstances, to sacrifice its own welfare to benefit the group. When the influence relationships among the members of a group can be expressed via a directed acyclic graph, a group is coherent if and only if its utilities are conditional mass functions. This structure permits the performance aspects to be merged with the random aspects to form a unified mathematical framework for decision problems under risk. The resulting solution may be interpreted as the Nash bargaining solution when the disagreement points of all agents are set to zero. Coherence is shown to be operationally equivalent to the concept of symmetry for a cooperative game. The resulting theory is designed to account for both individual and group-level preferences.

## I. INTRODUCTION

Many multiagent decision problems require the decision makers to cooperate to achieve the goals of the collective. A purely cooperative collective is one in which the interests of all individuals coincide perfectly. Many collectives, however, are mixed-motive, and opportunities for both conflict and cooperation exist. A key question in such cases is the definition of what it means to be rational. Historically, this question has been addressed from two distinct points of view: game theory and social choice theory. Under game theory, each individual seeks to optimize its own performance, whereas in the social choice context, the goal is to maximize performance of the group as a whole. In the former case, the value of the individual decisions to the group is not explicitly considered, and in the latter case, although the value judgments of the individuals may be used to define group-level performance, there is no assurance that the resulting decision will maximize the performance of, or even be acceptable to, any given individual.

The reconciliation of these two extreme perspectives is an important theoretical objective, both for human decision making and for the design of artificial decision-making entities who must cooperate. An important design principle for such scenarios is that the agents function according to a mathematical framework that is *coherent* in the sense that no individual can be categorically subjugated; i.e., is required in all situations to sacrifice its own welfare to benefit the group. If an agent were so required, it would not enjoy even an exiguous sense of equity—it would effectively be disenfranchised. Coherence is a minimal, yet critical,

property of a collective that is capable of sophisticated social behaviors such as negotiation, compromise, and altruism.

## II. MODELING FUNDAMENTALS

Let  $\{X_1, \dots, X_n\}$ ,  $n \geq 2$ , denote a group of autonomous decision makers. Let  $\mathcal{A}_i$  denote a finite set of feasible actions available to  $X_i$ ,  $i = 1, \dots, n$ , let  $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$  denote the product action space, and let  $\mathbf{a} = (a_1, \dots, a_n)$  denote the action profile that obtains when each  $X_i$  instantiates  $a_i \in \mathcal{A}_i$ .

Classical multiagent decision theory assumes that each individual possesses a total preference ordering over all action profiles. Under this assumption, each  $X_i$  possesses a utility  $u_{X_i}: \mathcal{A} \rightarrow \mathcal{R}$  such that  $u_{X_i}(\mathbf{a}) > u_{X_i}(\mathbf{a}')$  if  $X_i$  prefers  $\mathbf{a}$  to  $\mathbf{a}'$ , and  $u_{X_i}(\mathbf{a}) = u_{X_i}(\mathbf{a}')$  if  $X_i$  is indifferent between  $\mathbf{a}$  and  $\mathbf{a}'$ . These utility functions are assumed to provide a complete and immutable description of the valuations of actions for the collective. Generally, they are provided as part of the problem statement and, once defined, the logic used to arrive at these orderings is assumed to be irrelevant to the actual decision-making enterprise. We will term these functions *categorical utilities* since they are unconditional valuations of decision-maker preferences.

The categorical model of preferences, however, restricts the ability of individuals to modulate their preferences by giving deference to others under specific situations. For example, consider a collective that possesses a hierarchical structure such that  $X_1$  dominates  $X_2$  in some functional way. In such a case,  $X_2$  may need to adjust its preferences according to the preferences of  $X_1$ , but could not do so with a categorical preference ordering. Instead,  $X_2$  would possess a set of conditional preference orderings, each depending on the hypothetical assumption of  $X_1$ 's preferences. We may represent this set of conditional preferences by a set of *conditional utilities*  $u_{X_2|X_1}$  such that  $u_{X_2|X_1}(a_2|a_1)$  is the utility that  $X_2$  ascribes to  $a_2 \in \mathcal{A}_2$  given the hypothetical assumption that  $X_1$  is committed to action  $a_1 \in \mathcal{A}_1$ . This hypothetical commitment serves as the antecedent to a hypothetical proposition whose consequent is the conditional utility. A hypothetical commitment may take many forms, but perhaps the most important one, with respect to the social interaction of the agents, is that, from the perspective of  $X_2$ ,  $X_1$  considers  $a_1 \in \mathcal{A}_1$  to be its most preferred action. Under this interpretation,  $X_2$  is in a position to give deference to  $X_1$  by adjusting its conditional utility in a way that benefits (or, in a malevolent scenario, injures)  $X_1$ .

In general, if every agent influences every other agent (a fully connected group), then every agent's utility would

be conditioned on every other agent’s hypothesized commitments. It is often the case, however, that individual members of a group are most strongly influenced by their neighbors (functionally, spatially, or temporally). For example, hierarchical groups possess a distinctive “top-down” structure, where the preferences of subordinate agents are influenced by their superiors. A hierarchical structure is a special case of more general “Markovian” structures that are amenable to graphical analysis. A graphical structure that has been shown to be effective in many situations is a *directed acyclic graph*, or DAG. DAGs provide a convenient and powerful language with which to encode influence relationships—the most well known being so-called *Bayesian networks*, which are used extensively for the design of artificially intelligent systems [1–3].

A directed graph is a pair  $\mathcal{G} = (\mathbf{X}, E)$ , where  $\mathbf{X} = (X_1, \dots, X_n)$  is a finite set set of *vertices* and  $E$  is a set of *directed edges* linking pairs of vertices. If  $X_j$  is directly influenced by  $X_i$ , then there is a directed edge, denoted “ $\rightarrow$ ” from  $X_i$  to  $X_j$ . A *path* from  $X_i$  to  $X_j$  is a sequence of vertices  $\{X_i, X_{k_1}, X_{k_2}, \dots, X_j\}$  such that  $X_i \rightarrow X_{k_1} \rightarrow X_{k_2} \rightarrow \dots \rightarrow X_j$ . We write  $X_i \mapsto X_j$  if there is a path from  $X_i$  to  $X_j$ . If there are no paths such that  $X_i \mapsto X_i$  for any  $i$ , the graph is said to be *acyclic*.

If  $X_i \rightarrow X_j$ , then  $X_i$  is called a *parent* of  $X_j$ , and  $X_j$  is a *child* of  $X_i$ . The set of parents of  $X_i$  is denoted  $\text{pa}(X_i) = \{X_{i_j} : X_{i_j} \rightarrow X_i, j = 1 \dots, p_i\}$ , and the set of *children* of  $X_i$  is denoted  $\text{ch}(X_i)$ . The descendants of  $X_i$ , denoted  $\text{de}(X_i)$ , is the subset of vertices  $\{X_{i_m} : X_i \mapsto X_{i_m}, m = 1 \dots, d_i\}$ .

A fundamental property of a DAG is the *Markov condition*: nondescendent nonparents of a vertex have no influence on the vertex, given the hypothesized commitments of its parent vertices. Suppose  $\text{pa}(X_i) = \{X_j\}$ . By the Markov condition,  $X_i$ ’s utility is therefore a function only of the pair  $(x_i, x_j)$ . In general, suppose  $X_i$  has  $p_i$  parents, denoted  $\text{pa}(X_i) = \{X_{i_1}, \dots, X_{i_{p_i}}\}$ . For any action profile  $\mathbf{a} = (a_1, \dots, a_n)$ , let  $\mathbf{a}_i = (a_{i_1}, \dots, a_{i_{p_i}})$  denote the sub-profile of  $\mathbf{a}$  corresponding to  $\text{pa}(X_i)$ . We may then express the utility of  $\mathbf{a}$  to  $X_i$  as

$$u_{X_i}(\mathbf{a}) = u_{X_i|\text{pa}(X_i)}(a_i|\mathbf{a}_i), \quad (1)$$

the conditional utility of  $X_i$  given the action sub-profile of its parents. If  $\text{pa}(X_i) = \emptyset$ , then its utility is not influenced by the commitments of any other agent. Its utility is then *marginal* and is of the form  $u_{X_i}(a_1, \dots, a_n) = u_{X_i}(a_i)$ . The conditional utilities constitute the edges of the DAG.

The conditional and marginal utility structures provide an important mechanism by which an agent may assess its preferences. Whereas the general structure  $u_{X_i}(a_1, \dots, a_n)$  requires  $X_i$  to specify its preferences over all action profiles  $(a_1, \dots, a_n)$ , the conditional approach requires  $X_i$  to specify its preferences only over its own action space, given each possible action of its parents. Thus, the agent is required to define its conditional preference ordering with respect only to the actions of itself for each hypothetical situation regarding its parents. Although this structure can be generalized, in this paper we restrict attention to collectives where at least one

agent’s preferences are not conditioned on the commitments of any other agent.

*Example 2.1:* Consider a collective involving three agents with the hierarchical structure illustrated in Figure 1.  $X_1$  is the primary agent (in the sense that it’s mission is most critical to the success of the enterprise);  $X_2$  and  $X_3$  are the secondary and tertiary agents, respectively. We observe that  $\text{pa}(X_1) = \emptyset$ ,  $\text{pa}(X_2) = \{X_1\}$ , and  $\text{pa}(X_3) = \{X_1, X_2\}$ . As a specific illustration, suppose  $X_1$ ’s concern is the appropriate market sector of a product to be manufactured (either  $a_1$ , the affluent customers, or  $a'_1$ , the less prosperous consumers). Given the sector,  $X_2$ ’s concern is to decide which product to manufacture (either widgets  $a_2$ , or gizmos  $a'_2$ ). Finally, given the sectors and the product,  $X_3$ ’s concern is to choose which grade of materials to use (either high quality  $a_3$ , or low quality  $a'_3$ ). Thus,  $\mathcal{A}_i = \{a_i, a'_i\}$ ,  $i = 1, 2, 3$ . The product action space  $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \times \mathcal{A}_3$  contains eight action profiles. The three agents must cooperate to achieve maximum productivity and hence must coordinate their choices. The corresponding utilities are  $u_{X_1} : \mathcal{A}_1 \rightarrow \mathcal{R}$ ,  $u_{X_2|X_1} : \mathcal{A}_2 \times \mathcal{A}_1 \rightarrow \mathcal{R}$ , and  $u_{X_3|X_1X_2} : \mathcal{A}_3 \times \mathcal{A}_1 \times \mathcal{A}_2 \rightarrow \mathcal{R}$ . The issue facing this group is to use these three utility structures to formulate a plan that is acceptable individually as well as for the group.




Fig. 1. The influence network for a three-agent hierarchy.

The conditional structure permits agents to exhibit *conditional altruism* by defining their preference orderings as a function of the preferences of others. For example, suppose  $u_{X_1}(\mathbf{a}) \gg u_{X_1}(\mathbf{a}')$ .  $X_2$  could reinforce this strong preference by setting  $u_{X_2|X_1}(\mathbf{a}|\mathbf{a}) \gg u_{X_2|X_1}(\mathbf{a}'|\mathbf{a})$ , thereby deferring to the preferences of  $X_1$ . This type of altruism, however is not categorical, since, conditioned on, say, a commitment by  $X_1$  to  $\mathbf{a}''$ ,  $X_2$  need not prefer  $\mathbf{a}$  to  $\mathbf{a}'$ . Conditional altruism thus provides decision makers with a natural vehicle with which to establish sophisticated social relationships that can enhance the possibilities for compromise and negotiation. For example,  $X_2$  can use its conditional utility as a parameter with which to adjust the amount of deference it is willing to grant  $X_1$  to effect a compromise. Conversely,  $X_2$  can use its conditional utility to threaten or punish  $X_1$  by reducing its utility of actions that are beneficial to  $X_1$  and, thereby, reducing the utility of that action to the group (e.g., through aggregation, as will be discussed shortly).

### III. GROUP-LEVEL RATIONAL DECISIONS

The study of how individual preferences are used to form a group decision is the central issue of social choice theory [4–6], and hence is relevant to the study of autonomous multiagent decision making groups. Social choice theory

has traditionally been applied to human societies, but the concepts are directly applicable to artificial societies as well, particularly those that are intended to function cooperatively. A key issue of this theory is how to aggregate the interests of individuals to form a group decision in a democratic fashion; i.e., in a way such that the interests of all individuals are respected and given equitable consideration.

Informally, a society is *coherent* if each member has a “seat at the table” in the sense that the possibility exists (although not the guarantee) that, for each of the individuals, a solution exists that is good for the group and is also good for that individual. Obviously, most voting schemes are transparently coherent (assuming each voter’s most preferred candidate is on the ballot), but when complex influence relationships exist among the members of a group, establishing coherence may not be obvious.

To formalize the notion of coherence, let us assume that  $X_i$  is able, after taking into consideration all social, economic, and political relationships between it and other agents, to define a utility  $u_{X_i}$  over its action space. We also assume that the group possesses a group-level utility  $u_{X_1, \dots, X_n}: \mathcal{A} \rightarrow \mathcal{R}$ .

*Definition 3.1:* Let  $u_{X_i}$  denote  $X_i$ ’s categorical utility,  $i = 1, \dots, n$ , and let  $u_{X_1, \dots, X_n}$  denote the utility of the group  $\mathbf{X} = \{X_1, \dots, X_n\}$ .  $\mathbf{X}$  is *coherent* if, given that  $u_{X_i}(a_i) > u_{X_i}(a'_i)$ , there exists an action sub-profile  $(a_1^*, \dots, a_{i-1}^*, a_{i+1}^*, \dots, a_n^*)$  such that

$$u_{X_1 \dots X_n}(a_1^*, \dots, a_{i-1}^*, a_i, a_{i+1}^*, \dots, a_n^*) \geq u_{X_1 \dots X_n}(a_1^*, \dots, a_{i-1}^*, a'_i, a_{i+1}^*, \dots, a_n^*).$$

If there does not exist such a sub-profile, then  $X_i$  is in a position of *categorical subjugation*: every action profile that contains its most preferred action is dominated by profiles that do not contain its most preferred action. In terms of voting, incoherence means that, no matter how the others vote,  $X_i$ ’s candidate will lose. Effectively,  $X_i$  is disenfranchised. Categorical subjugation is similar to the notion of *suppression* as discussed by [7] and [5].

The question thus becomes: what constraints must be placed on the utilities to ensure that a condition of categorical subjugation is impossible? To address this question, let us turn to an analogous issue. A Dutch book is a gambling situation such that, no matter what the outcome, the gambler will be worse off for having taken the gamble—a situation of *sure loss* (one’s reward is always less than one’s stake). To illustrate a Dutch book, Suppose  $Y$  can take one of two distinct values:  $y_1$  or  $y_2$ , and let  $q(y)$  denote a belief function<sup>1</sup> of  $y$ ; i.e.,  $q(y)$  measures the strength of belief that  $Y = y$ .

By convention, we will assume that we have full belief that exactly one of these values obtains—the disjunction of  $y_1$  and  $y_2$  must occur. We further assume that beliefs are additive, thus,  $q(y_1 \vee y_2) = q(y_1) + q(y_2) = 1$ . Now let  $Z$  take on one of two distinct values  $z_1$  or  $z_2$ , and let  $r(y, z)$  denote the belief that  $Y = y$  and  $Z = z$  simultaneously. Let

us assume that  $q(y_2) > q(y_1)$ , but  $r(y_1, z_1) > r(y_2, z_1)$  and  $r(y_1, z_2) > r(y_2, z_2)$ . Suppose you purchase a \$1 gamble  $Y = y_2$ , and deem a fair purchase price to be  $q(y_2)$ ; i.e., you pay  $\$q(y_2)$  for the gamble to win \$1. Now also suppose you sell the gamble  $(y_2, z_1) \vee (y_2, z_2)$ . By additivity of beliefs, a fair selling price for this bet would be  $r[(y_2, z_1) \vee (y_2, z_2)] = r(y_2, z_1) + r(y_2, z_2)$ . However, according to the above ordering, you must have  $q(y_2) > \frac{1}{2}$  and, since  $r(y_2, z_1) + r(y_2, z_2) < r(y_1, z_1) + r(y_1, z_2)$ , it follows that  $r[(y_2, z_1) \vee (y_2, z_2)] < \frac{1}{2}$ . After all gambles have been bought and sold, your net wealth is  $r[(y_2, z_1) \vee (y_2, z_2)] - q(y_2) < 0$ . To overcome this loss, you must make up the difference once the outcome of the gamble is known. But if neither  $y_2$  nor  $(y_2, z_1) \vee (y_2, z_2)$  occur, you win nothing and you pay nothing; if  $(y_2, z_1) \vee (y_2, z_2)$  occurs, then, of course,  $y_2$  occurs, so you win \$1 which you must pay to the buyer of your gamble. Thus, once the gambles have been bought and sold, your net wealth is invariant to whatever happens—you suffer a sure loss.

A belief system is said to be *coherent* if it is not possible to construct a Dutch book. The Dutch Book theorem [8, 9] and its converse [10] state that a belief system is coherent if and only if it complies with a probability measure that describes the degrees of belief regarding the propositions under consideration. The above example does not comply with the laws of probability theory, since  $q(y_2) \neq r(y_2, z_1) + r(y_2, z_2)$ ; i.e., marginalization fails.

Mathematically, a condition of categorical subjugation corresponds to a condition of sure loss. Therefore, to eliminate the possibility of categorical subjugation, the utilities must possess the same syntax as probability mass functions. We formalize this result as follows.

Let  $\mathbf{X} = \{X_1, \dots, X_n\}$  be a group of distributed decision-makers whose influence relationships can be expressed with a directed acyclic graph. For each  $X_i$ , let  $\text{pa}(X_i) = \{X_{i_1}, \dots, X_{i_{p_i}}\}$  denote the  $p_i$  parents of  $X_i$ , and let  $\mathcal{A}_i = \mathcal{A}_{i_1} \times \dots \times \mathcal{A}_{i_{p_i}}$  denote the  $p_i$ -dimensional product of the action spaces corresponding to the parents of  $X_i$ . If  $X_i$  has no parents, then  $\mathcal{A}_i = \emptyset$ . For each  $X_i$ ,  $u_{X_i | \text{pa}(X_i)}(a_i | \mathbf{a}_i)$  is the utility that  $X_i$  ascribes to  $a_i$ , conditioned on  $X_{i_j}$  committing to  $a_{i_j}$ ,  $j = 1, \dots, p_i$ . If  $X_i$  has no parents, the conditional utility is the marginal utility; i.e.,  $u_{X_i | \text{pa}(X_i)} = u_{X_i}$  if  $\text{pa}(X_i) = \emptyset$ .

*Theorem 3.1:* Categorical subjugation cannot occur if and only if the utilities  $u_{X_i | \text{pa}(X_i)}$  are conditional mass functions defined over  $\mathcal{A}_i \times \mathcal{A}_i$ ; i.e.,  $u_{X_i | \text{pa}(X_i)}(a_i | \mathbf{a}_i) \geq 0 \forall a_i \in \mathcal{A}_i$  and  $\sum_{a_i \in \mathcal{A}_i} u_{X_i | \text{pa}(X_i)}(a_i | \mathbf{a}') = 1 \forall \mathbf{a}' \in \mathcal{A}_i$ . Furthermore, the group-level utility of  $(a_1, \dots, a_n)$  for the group  $\mathbf{X} = \{X_1, \dots, X_n\}$  is

$$u_{\mathbf{X}}(\mathbf{a}) = u_{X_1 \dots X_n}(a_1, \dots, a_n) = \prod_{i=1}^n u_{X_i | \text{pa}(X_i)}(a_i | \mathbf{a}_i), \quad (2)$$

where  $\mathbf{a}_i$  is the sub-profile of a corresponding to  $\text{pa}(X_i)$ .

*Proof:* The Dutch Book Theorem and its converse establish that the utilities must be mass functions. Consequently, all of the edges of the DAG are utilities that possess the mathematical structure of conditional probability mass functions

<sup>1</sup>We refrain from using the term “probability” here, since we do not require  $q$  to possess all of the properties of a probability mass function.

(albeit with different semantics). Furthermore, the categorical utilities of each root vertex of the DAG possess the mathematical structure of marginal probability mass functions. Thus, the DAG satisfies all of the conditions of a Bayesian network, and we may apply the fundamental theorem of Bayesian networks; namely, that the joint probability mass function of the random variables associated with the vertices is the product of the conditional probability mass functions of all vertices with parents, and the marginal mass functions of all root vertices [1–3].  $\square$

The content of this theorem is that the mathematics of probability theory, which traditionally applies to *epistemological* situations involving assessments of belief and knowledge, also applies to *praxeological* situations involving assessments of expediency and efficiency. This result means that the mathematical notions of probability theory, such as independence, conditioning, marginalization, and so forth, can be given praxeological, as well as epistemological, interpretations.

Once the joint utility has been formed by the aggregation of individual utilities, the group-optimal action profile that maximizes the group utility is

$$\mathbf{a}_G = \arg \max_{\mathbf{a} \in \mathcal{A}} \prod_{i=1}^n u_{X_i | \text{pa}(X_i)}(a_i | \mathbf{a}_i). \quad (3)$$

#### IV. INDIVIDUALLY RATIONAL SOLUTIONS

Although the social choice-theoretic approach presented above possesses a weak notion of acceptability for the individuals (coherence), that does not imply that the group solution is acceptable to any given individual in terms of benefit to it. Simply having the opportunity for one’s interests to be equitably considered by the group does not imply that one’s interests are adequately represented in the group decision.

The most well known solution concept of non-cooperative game theory is Nash equilibria [11]. This solution concept is a reasonable approach under competitive scenarios, but when the agents are ostensibly to cooperate, it can lead to overly pessimistic results. For example, for scenarios where attempting to cooperate leaves one vulnerable to exploitation, such as Prisoner’s Dilemma-type games, the Nash equilibrium leads to the next-worst solution, rather than the Pareto solution. Particularly when the agents are disposed to communicate with each other, a more appropriate solution concept is one that permits some notion of equity or fairness to guide the decisions.

Cooperative game theory differs from non-cooperative theory in that players may enter into binding agreements regarding their behavior. For the players to forge an agreement, however, each must achieve an acceptable degree of satisfaction. A *bargaining game* is a cooperative game in which each participant possesses a *disagreement point* that defines the benefit that is guaranteed to accrue to it if a compromise cannot be reached. The disagreement point, therefore, is an indication of the strategic strength that is conferred on the participant as it participates in negotiations:

the higher the disagreement point, the greater bargaining strength of the participant.

A well-known bargaining solution concept that offers a clear definition of individual acceptability is the Nash bargain [12], which permits each participant to make maximal use of its strategic strength. The approach is based on four fundamental principles: (i) invariance to positive affine transformations; (ii) Pareto optimality; (iii) independence of irrelevant alternatives, and (iv) symmetry, which is the notion that no individual agent can expect that the other agents will grant it better terms than that individual itself would be willing to grant, were roles reversed.

Nash showed that these four conditions lead to a unique solution. Let  $d_{X_i}$  denote the disagreement point for  $X_i$ . The *negotiation set*, denoted  $\mathcal{N}$ , is the subset of action profiles such that every participant achieves at least its disagreement point. Although Nash’s theory pertains to categorical utilities, we may adapt the concept to the conditional case by replacing categorical utilities with conditional utilities. The negotiation set is defined as

$$\mathcal{N} = \{ \mathbf{a} \in \mathcal{A} : u_{X_i | \text{pa}(X_i)}(a_i | \mathbf{a}_i) \geq d_{X_i}, i = 1, \dots, n \}. \quad (4)$$

Following in the spirit of Nash’s result, the bargaining solution is

$$\mathbf{a}_N = \arg \max_{\mathbf{a} \in \mathcal{N}} \prod_{i=1}^n [u_{X_i | \text{pa}(X_i)}(a_i | \mathbf{a}_i) - d_{X_i}]. \quad (5)$$

We note that this solution is not, strictly speaking, a Nash bargain, since the utilities are not categorical. Nevertheless, the solution still possess the key feature of the Nash bargain; namely, that each participant takes full advantage of its strategic strength in that all action profiles that do not achieve at least its disagreement point are excluded from consideration.

#### V. RECONCILING GROUP AND INDIVIDUAL CHOICES

The above discussion demonstrates that, for groups whose social relationships can be represented by a directed acyclic graph, the bargaining solution and the coherent group-level optimal solution possess similar structure, differing mainly in the introduction of disagreement point for the bargaining solution concept. With the bargaining approach, the disagreement point is the value that the decision maker can guarantee for itself, regardless of whether or not a compromise can be reached. The justification for this approach is that it is possible for the agent to walk away from negotiations and go its own way without regard for others. While a go-it-alone option may be possible for human decision makers, an automated system that will not cooperate is likely to be dysfunctional. In fact, it may be necessary that they reach a compromise solution, regardless of the individual costs. If such a situation obtains, then the disagreement point for each agent will be its zero level, in which case the bargaining solution will coincide with the optimal solution for the group. This is an interesting result. Why should the best individual solution in the sense of a fair compromise

for each individual, as expressed by (5), also result in the best group-level solution, as expressed by (3)? The fact that there is such a close correspondence suggests that the notions of *symmetry* (no individual can expect that others will grant it better terms than it would be willing to grant, were roles reversed) and *coherence* (no individual interests can be categorically subjugated to the interests of the group) are operationally equivalent.

*Theorem 5.1:* For groups whose social relationships can be represented by a directed acyclic graph, (a) coherence implies symmetry, and (b) if symmetry applies and one individual is categorically subjugated, then all individuals are categorically subjugated—a condition of mutual categorical subjugation.

*Proof:* If coherence holds, then the utilities may be aggregated as the product of the conditional and marginal utilities, as given by (2), which, by changing the zero level, yields the bargain structure

$$\prod_{i=1}^n [u_{X_i | \text{pa}(X_i)}(a_i | \mathbf{a}_i) - d_{X_i}]. \quad (6)$$

Since the labeling of agents is arbitrary, exchanging  $X_i$  and  $X_j$  leaves this structure unaffected, hence symmetry holds.

Now suppose  $X_i$  can be categorically subjugated. By symmetry, if the roles of  $X_i$  and  $X_j$ ,  $j \neq i$ , are exchanged (including the utilities), then the solution is unchanged. Thus,  $X_j$  must be categorically subjugated as well. Since  $j$  is arbitrary, this means that all players are categorically subjugated.  $\square$

Mutual categorical subjugation is a pathological condition. It says that *all* agents must *always* sacrifice their own welfare to benefit the group. In fact, even if all individuals could agree that a given action profile were simultaneously best for all of them, that profile would not be best for the group—a violation of the Pareto principle. More generally, such a situation would mean that the interests of the individuals have only partial influence, at best (and perhaps no influence), on the interests of the group. Such a pathological situation would violate the most fundamental premise of social choice theory: “Democratic theory is based on the premise that the resolution of a matter of social policy, group choice or collective action should be based on the desires or preferences of the individuals in the society, group, or collective” [5, p. 3].

## VI. MULTIAGENT DECISION MAKING UNDER RISK

A decision is made under risk if the utility of actions is dependent upon random phenomena. When decisions are made under risk, the classical approach is a two-step procedure. First, the utilities are defined to correspond with the decision makers’ preferences; next, the expected value of the utility is computed. However, since Theorem 3.1 establishes that coherent utilities must possess the mathematical syntax of probability mass functions, the praxeological and epistemological aspects of a decision problem may be merged into a single praxi-epistemic structure. In particular, we may view the decision-making elements and the random

elements as vertices of a *praxeic-epistemic network*, whose edges are conditional mass functions.

Let  $\theta = \{\theta_1, \dots, \theta_m\}$  denote  $m$  random variables over the product sample space  $\Theta = \Theta_1 \times \dots \times \Theta_m$  associated with the decision problem which, when merged with  $\mathbf{X}$ , forms an  $(n+m)$ -dimensional DAG, called a *praxi-epistemic network*. Let  $\vartheta_i$  denote the realization of  $\theta_i$ , and let  $\vartheta = (\vartheta_1, \dots, \vartheta_m)$ . Then the joint *praxi-epistemic utility* is

$$u_{\mathbf{x}\theta}(\mathbf{a}, \vartheta) = \prod_{i=1}^n u_{X_i | \text{pa}(X_i)}(a_i | \mathbf{a}_i, \vartheta_i) \prod_{j=1}^m p_{\theta_j | \text{pa}(\theta_j)}(\vartheta_j | \mathbf{a}_j, \vartheta_j), \quad (7)$$

where  $p_{\theta_j | \text{pa}(\theta_j)}$  is the conditional probability of  $\theta_j$  given its parents,  $\mathbf{a}_i$  and  $\vartheta_i$  correspond to the praxeic and epistemic parents, respectively, of  $X_i$ , and  $\mathbf{a}_j$  and  $\vartheta_j$  correspond to the praxeic and epistemic parents, respectively, of  $\theta_j$ . The expected utility then becomes the praxeic marginal

$$u_{\mathbf{x}}(\mathbf{a}) = \sum_{\vartheta \in \Theta} u_{\mathbf{x}\theta}(\mathbf{a}, \vartheta). \quad (8)$$

Theorem 3.1 establishes that maximizing  $u_{\mathbf{x}}(\mathbf{x})$  yields the optimal joint action; i.e.,

$$\mathbf{a}_G = \arg \max_{\mathbf{a} \in \mathcal{A}} u_{\mathbf{x}}(\mathbf{a}). \quad (9)$$

Furthermore, under the assumption that no action can be taken unless all decision makers can agree on a joint action, the disagreement point for each decision maker is zero, and (8) constitutes the bargaining solution for the group. Thus, the group decision can also be viewed as individually optimal in the sense that each takes full advantage of its strategic strength.

In addition, this solution also satisfies the coherency property, as established by the following corollary.

*Corollary 6.1:* Let the *marginal* expected utility of  $X_i$  be given by

$$u_{X_i}(a_i) = \sum_{\neg a_i} u_{\mathbf{x}}(a_1, \dots, a_n) \quad i = 1, \dots, n, \quad (10)$$

where  $\sum_{\neg a_i}$  is the so-called “not-sum” notation meaning that the sum is taken over all elements of  $(a_1, \dots, a_n)$  *except*  $a_i$ .

If  $u_{X_i}(a_i) > u_{X_i}(a'_i)$ , then there exists  $(a_1^*, \dots, a_{i-1}^*, a_{i+1}^*, \dots, a_n^*)$  such that

$$u_{\mathbf{x}}(a_1^*, \dots, a_{i-1}^*, a_i, a_{i+1}^*, \dots, a_n^*) \geq u_{\mathbf{x}}(a_1^*, \dots, a_{i-1}^*, a_i, a_{i+1}^*, \dots, a_n^*).$$

*Proof:* Suppose  $u_{X_i}(a_i) > u_{X_i}(a'_i)$  holds, but there is no such  $(a_1^*, \dots, a_{i-1}^*, a_{i+1}^*, \dots, a_n^*)$ . Then

$$u_{X_i}(a_i) = \sum_{\neg a_i} u_{\mathbf{x}}(a_1, \dots, a_i, \dots, a_n) < \sum_{\neg a'_i} u_{\mathbf{x}}(a_1, \dots, a'_i, \dots, a_n) = u_{X_i}(a'_i), \quad (11)$$

a contradiction.  $\square$

*Example 6.1:* Figure 2 displays a praxi-epistemic network corresponding to the hierarchal manufacturing scenario with the network introduced in Example 2.1. The root node of

this DAG is  $X_1$ , the agent who decides which market sector to target, and is given by

$$u_{X_1}(a_1) = 0.6 \quad u_{X_1}(a_2) = 0.4.$$

This example also includes a random component,  $\theta$ , that characterizes the economic environment of the market sector. Let us take  $\Theta = \{\vartheta, \vartheta'\}$ , where  $\vartheta$  corresponds to a growing economy and  $\vartheta'$  corresponds to a shrinking economy. Thus, the probability of the economic status is conditioned on the market sector. The corresponding conditional probability functions are

$$\begin{aligned} p_{\theta|X_1}(\vartheta|a_2) &= 0.5 & p_{\theta|X_1}(\vartheta'|a_2) &= 0.5 \\ p_{\theta|X_1}(\vartheta|a'_2) &= 0.6 & p_{\theta|X_1}(\vartheta'|a'_2) &= 0.4. \end{aligned}$$

The utility of the product to be manufactured depends upon the market sector and the economic state, and is given as

$$\begin{aligned} p_{X_2|X_1\theta}(a_2|a_1, \vartheta) &= 0.7 & p_{X_2|X_1\theta}(a'_2|a_1, \vartheta) &= 0.3 \\ p_{X_2|X_1\theta}(a_2|a'_1, \vartheta) &= 0.5 & p_{X_2|X_1\theta}(a'_2|a'_1, \vartheta) &= 0.5 \\ p_{X_2|X_1\theta}(a_2|a_1, \vartheta') &= 0.4 & p_{X_2|X_1\theta}(a'_2|a_1, \vartheta') &= 0.6 \\ p_{X_2|X_1\theta}(a_2|a'_1, \vartheta') &= 0.2 & p_{X_2|X_1\theta}(a'_2|a'_1, \vartheta') &= 0.8 \end{aligned}$$

Finally, the utility of the grade of materials used in the manufacture is conditioned on the product and the sector is given by

$$\begin{aligned} p_{X_3|X_1X_2}(a_3|a_1, a_2) &= 0.6 & p_{X_3|X_1X_2}(a'_3|a_1, a_2) &= 0.4 \\ p_{X_3|X_1X_2}(a_3|a_1, a'_2) &= 0.6 & p_{X_3|X_1X_2}(a'_3|a_1, a'_2) &= 0.4 \\ p_{X_3|X_1X_2}(a_3|a'_1, a_2) &= 0.3 & p_{X_3|X_1X_2}(a'_3|a'_1, a_2) &= 0.7 \\ p_{X_3|X_1X_2}(a_3|a'_1, a'_2) &= 0.2 & p_{X_3|X_1X_2}(a'_3|a'_1, a'_2) &= 0.8 \end{aligned}$$

The praxi-epistemic utility is

$$\begin{aligned} u_{X_1X_2X_3\theta}(a_1, a_2, a_3, \vartheta) &= u_{X_1}(a_1)u_{X_2|X_1\theta}(a_2|a_1, \vartheta) \\ &u_{X_3|X_1X_2}(a_3|a_1, a_2)p_{\theta|X_1}(\vartheta|a_1). \end{aligned} \quad (12)$$




Fig. 2. The praxi-epistemic network for the hierarchy scenario.

The expected utility is the praxeic marginal

$$u_{X_1X_2X_3}(a_1, a_2, a_3) = \sum_{\vartheta \in \Theta} u_{X_1X_2X_3\theta}(a_1, a_2, a_3, \vartheta). \quad (13)$$

Straightforward calculations using the above utility values indicates that the optimal solution for the group is  $(a_1, a_1, a_3)$ , with a global utility value of  $u_{X_1X_2X_3}(a_1, a_2, a_3) = 0.216$ .

Upon computing the marginals, we obtain

$$\begin{aligned} u_{X_1}(a_1) &= 0.6 \\ u_{X_2}(a_2) &= 0.584 \\ u_{X_3}(a_3) &= 0.4624, \end{aligned}$$

indicating that the best group solution is also best for  $X_1$  and  $X_2$ , but is worst for  $X_3$ .

## VII. CONCLUSION

This paper provides a new theoretical approach to the modeling of distributed autonomous decision makers. Potential applications include mobile unmanned robotic systems such as coordinated UAV surveillance and reconnaissance missions, distributed decision making, scheduling and coordination for manufacturing enterprise automation, and man/machine decision making scenarios.

Conventional categorical preference orderings are not designed to account for sophisticated social relationships such as compromise and negotiation, since they do not easily permit individuals to expand their spheres of interest to account for the preferences of others. The introduction of conditional utilities is an important contribution to the theory of multiagent decision making, since it permits each agent to express its preferences as a function of the preferences of others. Individual conditional individual utilities can be aggregated to form a group utility that incorporates the social relationships that exist among the individuals, thereby providing a complete model of the community of decision makers.

A second contribution is the notion of coherence and the introduction of a mathematical structure for the utilities that ensures that no agent can be categorically subjugated. This structure permits the social relationships between individuals to be represented by a directed acyclic graph whose edges are conditional mass functions — a Bayesian network. This new syntax provides a natural vehicle with which to model sophisticated social relationships such as altruism.

A third contribution is the merging of the praxeic and epistemic components of a decision problem into a single praxi-epistemic utility that accounts for both utility and risk.

## REFERENCES

- [1] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kaufmann, 1988.
- [2] R. G. Cowell, A. P. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter, *Probabilistic Networks and Expert Systems*. New York, NY: Springer Verlag, 1999.
- [3] F. V. Jensen, *Bayesian Networks and Decision Graphs*. New York, NY: Springer Verlag, 2001.
- [4] K. J. Arrow, *Social Choice and Individual Values*. New York, NY: John Wiley, 1951, 2nd ed. 1963.
- [5] P. C. Fishburn, *The Theory of Social Choice*. Princeton, NJ: Princeton University Press, 1973.
- [6] J. Craven, *Social Choice*. Cambridge, UK: Cambridge University Press, 1992.
- [7] B. Hansson, "The independence condition on the theory of social choice," 1972, working paper no. 2, The Mattias Fremling Society, Department of Philosophy, Lund, 1972.
- [8] F. P. Ramsey, "Truth and probability," in *The Foundations of Mathematics and Other Logical Essays*, R. B. Braithwaite, Ed. New York, NY: The Humanities Press, 1950.
- [9] B. de Finetti, "La prévision: ses lois logiques, ses sources subjectives," *Annales de l'Institut Henri Poincaré*, vol. 7, pp. 1–68, 1937, translated as 'Forsight. Its Logical Laws, Its Subjective Sources', in *Studies in Subjective Probability*, H. E. Kyburg Jr. and H. E. Smokler (eds.), Wiley, New York, NY, 1964, pages 93–158.
- [10] J. Kemeny, "Fair bets and inductive probabilities," *Journal of Symbolic Logic*, vol. 20, no. 1, pp. 263–273, 1955.
- [11] J. F. Nash, "Equilibrium points in  $n$ -person games," in *Proceedings of the National Academy of Sciences*, vol. 36, USA, 1950, pp. 48–49.
- [12] —, "The bargaining problem," *Econometrica*, vol. 18, pp. 155–162, 1950.

# Satisficing Coordination and Social Welfare for Robotic Societies

Wynn C. Stirling · Matthew S. Nokleby

the date of receipt and acceptance should be inserted later

**Abstract** The design of robotic systems that are capable of sophisticated social behavior such as cooperation, compromise, negotiation, and altruism, requires more complex mathematical models than is afforded by the classical mechanisms for making value judgments and decisions. A new concept of multi-agent satisficing, defined in terms of relative effectiveness and efficiency, is an alternative to classical optimization-based decision making. Conditional utilities, which take into account the interests of others as well as the self, represent an alternative to the categorical utilities of classical decision theory. A multi-agent utility aggregation structure is developed that avoids the sure subjugation of the interests of any individual to the interests of the group. By expressing a society as a directed acyclic graph, Bayesian network theory is applied to artificial societies. A satisficing social welfare function accounts for the influence relationships among decision-making agents.

**Keywords** Multi-agent Systems, Game Theory, Social Choice Theory, Satisficing, Conditional Preferences, Coherence

## 1 Introduction

Multi-stakeholder decision problems arise in many contexts, including social choice theory, game theory, distributed control theory, multi-criterion/multi-objective decision theory, multi-agent systems theory, and social robotics. Although the particulars of these various contexts can differ widely, to be rational, each must possess two fundamental attributes: (a) an ability to make value judgments regarding alterna-

tives, and (b) procedures for using value judgments to make choices.

The field of social robotics, in particular, provides a rich environment for the application of decision-making logics that are able to accommodate sophisticated social behaviors such as compromise, negotiation, and altruism. Whether a social robot interfaces with humans, other robots, or both, it typically resides in a community that involves some notion of coordination (which may be either cooperative or competitive). In such an environment, value judgments can depend upon the desires and preferences of others, and procedures for making choices must take these complex social relationships into account.

This paper provides a mathematical framework within which to design and synthesize complex decision-making collectives that are able to accommodate socially complex decision making. Section 2 provides a brief history of classical multi-agent decision making and motivates our approach. Section 3 introduces the key components of the framework we are proposing, Section 4 introduces new concepts of social welfare, Section 5 reconciles our theory with classical approaches, Section 6 describes a special case of what we term *decoupled* social systems, and Section 7 offers conclusions.

## 2 Background

Cooperative robotics is an active research area. Of particular interest is the development of theories for decentralized control of multi-robot societies. Swarm-based approaches have demonstrated the emergence of cooperative behavior [24, 25]. Potential functions, consisting of constraints and goals that are imposed upon the system, have been used to address the mobile robot navigation problem [5]. Shannon information theory has been applied to the investigation of

diversity among heterogeneous agents, thereby enabling an assessment of the ability of the system to perform cooperatively [4]. Behavior-based approaches have been applied to the design of cooperative robotic teams, stressing minimalism, statelessness, and tolerance [53]. The variety displayed by these various of approaches is a strong indication of the complexity involved in the design of cooperative multiagent systems, and is in indication that there is no single approach that can be universally applied to the design and synthesis of such systems.

Because of the complexity of multiagent systems, it is important to review the fundamental principles that are exploited, either implicitly or explicitly, in their design. Accordingly, we provide a brief review of classical decision-theoretic foundations and a discussion of rationality.

## 2.1 Classical Decision-theoretic Foundations

The multi-stakeholder decision problem originated in the social sciences context, with foundations laid by Bergson [6], Samuelson [39], Arrow [1, 2], and others, who assert that individual values are the fundamental elements of a society. Arrow has provided what is perhaps the most clear definition of this concept: “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 standard he deems relevant” [1, p. 17]. Furthermore, Friedman argues that the process by which these preferences are obtained is irrelevant: “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. The legitimacy of any justification for this abstraction must rest ultimately, in this case as with any other abstraction, on the light that is shed and the power to predict that is yielded by the abstraction” [17, p. 13]. According to the Arrow/Friedman model, each participant in a multi-agent decision problem comes to the decision-making activity with pre-defined preference orderings, the origins of which are not germane to the decision problem. Such preference orderings are *categorical*. The assumption that each individual possesses a categorical preference ordering has been adopted almost universally in classical multi-stakeholder decision-making contexts.

The most common procedure for using value judgments to make choices is to invoke some notion of optimization — the *sine qua non* of classical decision theory. As put by Euler, “Since the fabric of the world is the most perfect and was established by the wisest Creator, nothing happens in this world in which some reason of maximum or minimum would not come to light” (quoted in [35]). What is optimal in multi-stakeholder settings, however, can depend

upon the point of view. In the classical game-theoretic context, each individual seeks to optimize value to itself, and a Nash equilibrium is a constrained mutually optimal solution for all players in the sense that no individual can unilaterally improve its welfare by changing its decision. On the other hand, in the social choice context, it is the “organization incarnate,” as Raiffa put it [37], who seeks to maximize value for the group considered as a whole. In the former case, the value of the individual decisions to the group is not explicitly considered, and in the latter case, although the value judgments of the individuals are used to define group-level decisions (e.g., a weighted sum of individual valuations), there is no assurance that the resulting decision will maximize the value to any individual. In fact, the decision that is best for the group can be extremely unfavorable to some members of the group.

## 2.2 Rationality

The classical approach to decision making in group settings is the doctrine of individual rationality: the notion that each individual should act in a way that maximizes its own satisfaction (without explicit regard for the satisfaction of others). This doctrine enjoys a central role in classical decision theory and game theory. As discussed by Tversky and Kahneman, “The assumption of [individual] rationality has a favored position in economics. It is accorded all of the methodological privileges of a self-evident truth, a reasonable idealization, a tautology, and a null hypothesis. Each of these interpretations either puts the hypothesis of rational action beyond question or places the burden of proof squarely on any alternative analysis of belief and choice. The advantage of the rational model is compounded because no other theory of judgment and decision can ever match it in scope, power, and simplicity” [52, p.89].

The uncritical application of individual rationality as a model for decision making in multi-agent contexts can be problematic. Arrow has observed that “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” [3, p. 203].

If all agents are indeed focused on, and only on, their narrow self-interest, then categorical preferences are appropriate. Difficulties arise, however, when the sphere of concern of an individual extends beyond its own narrow self-interest. The only way such an individual can use categorical preferences to accommodate the preferences of other individuals is to redefine its values by substituting (at least par-

tially) the values of the others for its own. Such behavior is a manifestation of *categorical altruism*, i.e., irrevocably sacrificing one's own welfare in an attempt to benefit another, thus fundamentally changing the nature of the association.

Considerable research, notably in the field of behavioral economics [8], has addressed the need for agents to define their preferences such that they consider social interactions. Fehr and Schmidt [15] discuss how individual preference orderings may be modified to take into account concepts such as fairness and cooperation by introducing a notion of inequity aversion. To account for this attribute, they include, in addition to a purely selfish component, an inequity aversion component in their utility. Consequently, they rely upon (re-defined) categorical preference orderings to model social interactions. All that changes is the definition of the individual's self-interest. This approach, however, has serious limitations, as acknowledged by Sen: "It is possible to define a person's interests in such a way that no matter what he does he can be seen to be furthering his own interests in every isolated act of choice . . . no matter whether you are a single-minded egoist or a raving altruist or a class-conscious militant, you will appear to be maximizing your own utility in this enchanted world of definitions" [40, page 19]. Categorical altruism simulates cooperation, compromise, and altruism with a regime that is explicitly designed to characterize selfishness, competition, and avarice, and does not offer a natural and intuitively pleasing framework within which to express sophisticated and complex social relationships. While such constructions may serve to explain some forms of human behavior, it is difficult to see how they can be used systematically to synthesize complex relationships between artificial agents.

The foundational assumptions of categorical preferences for each individual and optimization (either for individuals or for the group) undergird virtually all of classical formalized decision theory in both individual and group settings. These assumptions correspond to *analysis* tools that serve, with varying success, to explain and predict human behavior, but they are not causal: they do not govern human behavior. On the other hand, models that are used to design a system of artificial autonomous decision-making agents must be causal: they are *synthesis* tools that will indeed govern the behavior of the artificial society.

Many social science researchers argue, however, that the classical foundational assumptions do not provide an adequate model for human behavior (e.g., see [26, 47]). And if their adequacy to analyze human behavior is questioned, then we may rightly question their appropriateness as assumptions with which to synthesize the behavior of artificial societies that are expected to behave in ways that are can be understood and trusted by humans. As Shubik as acknowledged, "Economic man, operations research man and the game theory player were all gross simplifications. They

were invented for conceptual simplicity and computational convenience in models loaded with implicit or explicit assumptions of symmetry, continuity, and fungibility in order to allow us (especially in a pre-computer world) to utilize the methods of calculus and analysis. Reality was placed on a bed of Procrustes to enable us to utilize the mathematical techniques available [43]."

It is time to make the bed fit its occupant. Particularly in the context of artificial multi-agent system design and synthesis, a framework is needed to model explicitly the possibly complex value judgments the may exist among the members of an artificial society. It is time to account for situations where the conditions under which preferences are formed are relevant and cannot be summarily ignored; it is time to accommodate more complex and flexible criteria for making decisions. In other words, it is time to move beyond categorical preference orderings and optimization as the foundational components of multi-stakeholder decision making.

### 3 A Social Framework for Cooperative Decision Making

A social welfare function, as defined by Arrow, is "a process or rule which, for each set of individual orderings  $R_1, R_2, \dots, R_n$  for alternative social states (one ordering for each individual), states a corresponding social ordering of alternative social states  $R$ " [1, p. 23]. Classically, the individual orderings (either ordinal or cardinal), are categorical, in that they account only for the interests of the individuals. We wish to expand the spheres of interest of the individuals to include the interests of others as itself. However, once we move beyond restricting to individual interests, the notion of optimization becomes problematic. Optimization is an individual concept: for a group to optimize, it must act as a single unit, capable of making rational judgments and choices. Such a structure however, is not consistent with our assumption that the decision-makers are autonomous.

Our approach is to replace the twin assumptions of optimization and categorical preferences with two alternative concepts: satisficing and conditioning. Our goal is to create a satisficing social welfare function and individual welfare functions that can be used to construct compromise solutions that are simultaneously acceptable to the group and the individuals, thereby removing, or at least reducing, the wedge that separates classical concepts of group and individual interests.

#### 3.1 Satisficing

In a multi-agent setting it is not generally possible to maximize both individual and group preferences simultaneously. A potentially more socially accommodating concept is that

decisions are “good enough.” What is best for you may be different from what is best for me, but what is good enough for you may also be good enough for me, provided we have some flexibility in our respective notions of what it means to be good enough. The term “satisficing” has been advanced as a synonym for this alternative to strict optimization.

The first usage of the term “satisficing” in a decision-theoretic context is attributed to Simon [44–46], who addressed the question of how a decision maker might make a choice in the presence of informational or computational limitations. Simon’s approach is to seek an optimal choice, but to terminate searching and once the decision maker’s aspiration level has been met. Put another way, to satisfice is to accept the best solution so far obtained, once the cost of continuing to search exceeds the expected improvement in value were the search to continue. Many other variations of this concept have appeared in the literature [7, 14, 20, 23, 28, 29, 33, 36, 41, 50, 51, 54–56], and it is not the intent of this paper to review them in detail. Suffice it to say, however, that all of these approaches view satisficing as a species of bounded rationality: one settles for a solution that is deemed to be “good enough,” but which is not necessarily, and usually not, optimal in any meaningful sense. Satisficing *à la* Simon is an heuristic approximation to the ideal of being best (and is only constrained from achieving this ideal by practical limitations)

The concept of satisficing developed herein differs from the afore-mentioned notion in several important ways. First, in contrast to satisficing as advanced by Simon and others, it is not heuristic; rather, it provides a concept of satisficing that is as mathematically formalized and precise as is the notion of optimization. Second, it treats the notion of being good enough as the ideal (rather than an approximation) — it is *not* a species of bounded rationality. Third, it extends to the multi-agent case, thereby providing a natural framework for multi-agent decision making. Fourth, it readily accommodates the extension of interests beyond the self, thereby accommodating more sophisticated social relationships than self-interest affords. We retain the term “satisfice” because, even though our approach is not heuristic, we nevertheless seek solutions that are good enough, with the essential difference being that we provide a non-heuristic definition of what it means to be good enough.

Although it seems eminently reasonable at least to attempt (given sufficient resources) to seek an optimal decision, humans often invoke a systematic approach to decision making (even in single-agent decision problems) that, while still based on quantitative measures of performance, does not correspond to optimization. In the vernacular, the optimization paradigm corresponds to seeking “the best and only the best” solution. Also common, however, is the paradigm of “getting your money’s worth.” In an intuitively pleasing sense, this latter notion admits an interpretation as being

good enough, and it is this concept that we invoke as the satisficing paradigm that we develop in this paper. A comprehensive introduction to this perspective can be found in [49].

Many theorists (e.g., [1, 13, 18, 27]) have argued that it is unwise to aggregate conflicting interests into a single preference ordering. Some have asserted that in a social setting individuals have multiple facets, as defined by Steedman and Krause [48], who maintain that an agent, although an indivisible unit, nevertheless is capable of considering its choices from different points of view, and that separate utilities may be defined to correspond to each facet of an individual. A natural way to classify attributes is according to their effectiveness and efficiency. Each individual may be viewed as being composed of two facets: the *selecting facet*, which evaluates actions in terms of effectiveness toward pursuing objectives without concern for efficiency, and the *rejecting facet*, who evaluates actions in terms of efficiency with respect to consuming resources without concern for effectiveness. We shall view these selecting and rejecting facets as the “atoms” of the society,

When formulating a problem under the satisficing framework, it is essential that the selecting and rejecting criteria not be restatements of each other. The selecting criterion should correspond to the goals of the problem, and the rejecting criterion should correspond to the consumption of resources. This dual utilities approach is the basis for our notion of satisficing.

Under the optimization paradigm, all of the performance measures are combined into a single utility, whereas under the satisficing paradigm, the measures of effectiveness are encoded separately from the measures of efficiency. Under the optimization paradigm, the alternatives are compared against each other in order to identify the globally best one. By contrast, under the satisficing paradigm, the effectiveness and efficiency attributes are locally compared for each alternative separately, and all alternatives for which the effectiveness measures exceed the efficiency measures are considered to be satisficing. Thus, whereas the optimization paradigm is designed to identify a single best alternative, the satisficing paradigm is designed to identify all alternatives that are good enough. The non-uniqueness attribute is a key feature of satisficing in a multi-stakeholder environment, since it is amenable to flexibility on the part of the individuals and of the group.

To introduce the formalism of satisficing, let us first consider a single agent  $X$ , with selecting and rejecting facets denoted  $S$  and  $R$ , respectively, and let  $u_S$  denote the selecting utility, or *selectability*, which measures the progress toward the goal of  $X$ , and  $u_R$  denotes the rejecting inutility, or *rejectability* which measures the consumption of resources such as cost, exposure to hazard, loss of social reputation, and so forth.

**Definition 1** Let  $\mathcal{A}$  denote the set of actions available to  $X$ . An action  $a \in \mathcal{A}$  is *satisficing* if  $u_S(a) \geq qu_R(a)$  where  $q \in [0, 1]$  regulates the threshold for rejecting elements of  $\mathcal{A}$  as not satisficing. (Nominally,  $q = 1$ , but as we shall see,  $q$  can serve as a measure of how willing an agent is to negotiate.) The *satisficing set* is

$$\Sigma_q = \{a \in \mathcal{A}: u_S(a) - qu_R(a) \geq 0\}. \quad (1)$$

Satisficing as defined above is expressed in a single-agent context with categorical utilities. It is easily seen, in this simple context, that  $u_S$  and  $u_R$  can easily be combined to form a classical utility  $u_X(a) = u_S(a) - qu_R(a)$ , which is amenable to optimization. Optimization, however, is designed to produce a single best solution, whereas, by contrast, satisficing is designed to produce a (possibly) non-singleton set of solutions that are good enough in the sense that the effectiveness of the action is as least as great as its inefficiency. In the single-agent context, satisficing represents a novel approach, but if it is possible to optimize, then there may be little incentive to seek a satisficing solution. The real power of the satisficing concept, however, is manifest in the multi-agent case, as will be further developed below.

### 3.2 Conditioning

Let  $\mathbf{X} = (X_1, \dots, X_n)$  denote a collective of autonomous stake-holders (e.g., agents). More specifically, let  $\mathbf{S} = (S_1, \dots, S_n)$  denote the collective of selecting facets, and let  $\mathbf{R} = (R_1, \dots, R_n)$  denote the collective of rejecting facets. Notationally, we write  $\mathbf{V} = \mathbf{SR} = (S_1, \dots, S_n, R_1, \dots, R_n)$ , a system of  $2n$  facets. Since we will be dealing with the facets, rather than the agents, it is convenient to use the symbol  $V_i$ ,  $i = 1, \dots, 2n$ , to denote either a selecting facet or a rejecting facet.

Let  $\mathcal{A}_i$  denote a finite set of alternatives available to  $X_i$ . Of course, if  $X_i$  takes action  $a_i \in \mathcal{A}_i$ , then that action also applies to  $S_i$  and  $R_i$  (split personalities are not allowed, but this does not mean that  $S_i$  and  $R_i$  must always contemplate taking the same action). The product action space is denoted  $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ , and an *action profile*  $\mathbf{a} = (a_1, \dots, a_n) \in \mathcal{A}$  denotes the joint action taken by the collective.

A *categorical utility* for  $V_i$ , denoted  $u_{V_i}$ , is a mapping  $u_{V_i}: \mathcal{A} \rightarrow \mathbb{R}$ , and provides a total ordering of all action profiles for  $V_i$ . According to the conventional Arrow/Friedman model, categorical utilities for all participants in the multi-stakeholder decision problem are defined prior to the decision-making activity and, furthermore, the mechanisms that dictate the way they are defined are irrelevant. As an alternative, we introduce the notion of a *conditional utility*. To develop this concept, we must first define a *commitment*.

**Definition 2** Let  $V_i$  be an arbitrary element of  $\mathbf{V}$ , and let  $\mathbf{V}_j = (V_{j_1}, \dots, V_{j_k})$  be an arbitrary  $k$ -element subset of  $\mathbf{V}$  that does not include  $V_i$ . A *commitment profile*  $\{\mathbf{a}_{j_1}, \dots, \mathbf{a}_{j_k}\}$ ,  $\mathbf{a}_{j_\ell} \in \mathcal{A}$ , is a hypothetical statement by  $V_i$  that the action profile  $\mathbf{a}_{j_\ell}$  is the one that is most preferred by  $V_{j_\ell}$ ,  $\ell = 1, \dots, k$ .

**Definition 3** A *conditional utility* for  $V_i$  with respect to a commitment profile  $\{\mathbf{a}_{j_1}, \dots, \mathbf{a}_{j_k}\}$ , denoted  $u_{V_i|V_j}(\mathbf{a}|\mathbf{a}_{j_1}, \dots, \mathbf{a}_{j_k})$ , is a utility for  $V_i$  given that  $\mathbf{V}_j$  is committed to  $\{\mathbf{a}_{j_1}, \dots, \mathbf{a}_{j_k}\}$ .

Operationally, a conditional utility for  $V_i$  serves as the consequent of a hypothetical proposition whose antecedent is a commitment by  $\mathbf{V}_j$ . This expression does not represent  $V_i$ 's actual utility of  $\mathbf{a}$ , nor does it imply that  $V_{j_\ell}$  truly most prefers  $\mathbf{a}'_{j_\ell}$ . Instead, it means that, if  $(V_{j_1}, \dots, V_{j_k})$  were simultaneously to prefer  $\{\mathbf{a}'_{j_1}, \dots, \mathbf{a}'_{j_k}\}$  to all other action profiles, then  $V_i$  would define its utility of  $\mathbf{a}$  accordingly.

An attractive feature of a conditional utility is that it permits  $V_i$  to express *conditional altruism*. To illustrate, suppose  $u_{V_j}(\mathbf{a}) \gg u_{V_j}(\mathbf{a}')$ , that is,  $V_j$  were to ascribe much higher categorical utility to  $\mathbf{a}$  than to  $\mathbf{a}'$ , but  $V_i$  were to do the opposite, ascribing higher utility to  $\mathbf{a}'$  than to  $\mathbf{a}$ , i.e.,  $u_{V_i}(\mathbf{a}') \geq u_{V_i}(\mathbf{a})$ .  $V_i$  could give deference to  $V_j$  by replacing its categorical utility  $u_{V_i}$  with a conditional utility  $u_{V_i|V_j}$  such that  $u_{V_i|V_j}(\mathbf{a}|\mathbf{a}) \geq u_{V_i|V_j}(\mathbf{a}'|\mathbf{a})$  but  $u_{V_i|V_j}(\mathbf{a}'|\mathbf{a}') = u_{V_i|V_j}(\mathbf{a}|\mathbf{a}') = u_{V_i}(\mathbf{a}')$ , thus deferring to  $V_j$  if, but only if,  $V_j$  were to favor  $\mathbf{a}$  strongly over  $\mathbf{a}'$ .

### 3.3 Social Networks

Conditional preferences provide each individual with the ability to define its preferences as a function of the hypothetical preferences of all other subsets of the collective. This feature represents an important departure from the traditional categorical definitions of preference and provides the foundation for the modeling of a complex society that possesses sophisticated social relationships such as altruism (either benevolent or malevolent). Conditional preference relations permit the explicit modeling of such relationships, rather than merely simulating them by redefining categorical preference orderings. Although conditional preference relations are more complex than are categorical ones, as noted by Palmer, "Complexity is no argument against a theoretical approach if the complexity arises not out of the theory itself but out of the material which any theory ought to handle" [32, p. 176].

Nevertheless, the introduction of conditional utilities increases the complexity of the mathematical model of a collective. At one extreme, all of the members of the collective would be devoted to narrow self-interest, and all utilities would be categorical (the classical game-theoretic model). At the other extreme, each of the members would be influenced by the preferences of every other member, resulting

in a fully connected collective. Fortunately, however, many potentially interesting societies are such that the connections between the members are relatively sparse. Just as with human societies, it is likely that members will be organized into relatively small clusters of individuals that are somewhat loosely connected with other clusters. One such model is a hierarchical structure, where the preferences of superiors influence those of subordinates. Another, more parallel model, is one where the individuals are grouped into function, spatial, or temporal neighborhoods. A powerful and convenient way to represent such relationships is through graph theory, which provides a means to express directly the influence relationships that exist among the individuals. With such a formalism, the vertices of the graph represent the members of the collective, and the edges represent the influence flows among them as encoded in the conditional utilities. For the extreme case where all individuals possess categorical preferences, the graph would have no edges — each individual would be expressed by an isolated vertex. When conditional preferences exist, however, the graph will have edges, as illustrated in Figure 1.




Fig. 1 A directed acyclic graph

In this paper we concentrate on *directed acyclic graphs*, or DAGs. Formally, a directed graph is a pair  $\mathcal{G} = (\mathbf{V}, E)$ , where  $\mathbf{V} = (V_1, \dots, V_{2n})$  is a finite set of vertices and  $E$  is a set of *edges* linking pairs of vertices. If  $V_j$  is directly influenced by  $V_i$  but  $V_j$  does not directly influence  $V_i$ , then there is a directed edge, denoted “ $\rightarrow$ ” from  $V_i$  to  $V_j$ . A *path* from  $V_i$  to  $V_j$  is a sequence of vertices  $\{V_i, V_{k_1}, V_{k_2}, \dots, V_j\}$  such that  $V_i \rightarrow V_{k_1} \rightarrow V_{k_2} \rightarrow \dots \rightarrow V_j$ . We write  $V_i \mapsto V_j$  if there is a path from  $V_i$  to  $V_j$ . If there are no paths such that  $V_i \mapsto V_i$  for any  $i$ , the graph is said to be *acyclic*.

If  $V_i \rightarrow V_j$ , then  $V_i$  is called a *parent* of  $V_j$ , and  $V_j$  is a *child* of  $V_i$ . The set of parents of  $V_i$  is denoted  $\text{pa}(V_i) = \{V_{i_j}: V_{i_j} \rightarrow V_i, j = 1, \dots, p_i\}$ , and the set of *children* of  $V_i$  is denoted  $\text{ch}(V_i)$ . If  $V_i$  has no parents, then  $\text{pa}(V_i) = \emptyset$ . The descendants of  $V_i$ , denoted  $\text{de}(V_i)$ , is the subset of vertices  $\{V_{i_m}: V_i \mapsto V_{i_m}, m = 1, \dots, d_i\}$ .

Let  $\text{cp}(V_i) = \{\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_{p_i}}\}$  denote the commitment profile for  $\text{pa}(V_i)$ . For each  $V_i$ ,  $u_{V_i|\text{pa}(V_i)}[\mathbf{x}] \text{cp}(V_i)$  is the utility that  $V_i$  ascribes to  $\mathbf{x}$ , given that  $V_{i_j}$  commits to  $\mathbf{x}_{i_j}$ ,

$j = 1, \dots, p_i$ . If  $V_i$  has no parents, the conditional utility is the categorical utility; i.e.,  $u_{V_i|\text{pa}(V_i)} = u_{V_i}$  if  $\text{pa}(V_i) = \emptyset$ . Consider the DAG illustrated in Figure 1. By inspection,  $\text{pa}(S_1) = \text{pa}(R_3) = \text{pa}(R_4) = \emptyset$ ,  $\text{pa}(S_2) = \{R_1\}$ ,  $\text{pa}(S_3) = \{S_2, S_4, R_1\}$ ,  $\text{pa}(S_4) = \{R_3\}$ ,  $\text{pa}(R_1) = \{R_2\}$ , and  $\text{pa}(R_2) = \{S_1\}$ .

A fundamental property of a DAG is the *Markov condition*: nondescendent nonparents of a vertex have no influence on the vertex, given the state of its parent vertices [10]. Consequently, if a society can be represented as a DAG, the conditional utility of a facet is dependent only upon the commitments of its parents. Thus, for the DAG in Figure 1,  $R_2$  is influenced only by the commitments of  $S_1$ ,  $S_3$  is influenced by the commitments of  $S_2$ ,  $S_4$ , and  $R_1$ , and so forth. Thus, conditional utility of  $R_2$  is of the form  $u_{R_2|\text{cp}(R_2)}$ , where  $\text{cp}(R_2) = \{S_1\}$ , and the conditional utility of  $S_3$  is of the form  $u_{S_3|\text{cp}(S_3)}$ , where  $\text{cp}(S_3) = \{S_2, S_4, R_1\}$ . Categorical utilities are associated with the root nodes,  $S_2$ ,  $R_3$ , and  $R_4$ , since these nodes have no parents.

## 4 Social Welfare

### 4.1 Collective Preferences

The central question for a collective of autonomous decision makers is how they should function as a group. In the classical non-cooperative game-theoretic formulation, the notion of a group preference is irrelevant — each individual is committed to, and only to, its own satisfaction, and the emergence of a coherent notion of group welfare would be strictly coincidental. As observed by Shubik, “It may be meaningful, in a given setting, to say that group ‘chooses’ or ‘decides’ something. It is rather less likely to be meaningful to say that the group ‘wants’ or ‘prefers’ something” [42, p. 124]. Social choice theory, on the other hand, focuses on the aggregation of individual preferences to form a social welfare function that can be used to define what is best for the group. Classical social choice theory, however, as developed by Arrow [1], Debreu [12], Fishburn [16], and others, also relies upon categorical preferences, as does multi-objective decision theory [21]. The main classical result, attributed to Debreu, is that a necessary and sufficient for a group utility to be defined as the weighted sum of individual utilities is that the individual utilities must be categorical.

In the presence of conditional preferences, the issue of social welfare takes on added complexity. For example, the traditional axioms of social choice theory, such as the independence of irrelevant alternatives, becomes problematic. Thus, we must pursue a different course when aggregating conditional preferences. In the interest of clarity, we begin our discussion of this concept with the bi-agent case, with  $\mathbf{V} = (V_1, V_2)$ . Let us suppose that  $V_1$  possesses a categorical utility  $u_{V_1}$  and  $V_2$  possesses a conditional utility  $u_{V_2|V_1}$ .

The corresponding DAG is displayed in Figure 4.1. Given these utilities, the central questions are: (i) Can these two utilities be combined in a rational way to form a group utility? and, if so, (ii) How should they be combined?




Fig. 2 A two-agent DAG

To address this issue, we introduce the notion of a joint commitment. A *joint commitment* by  $(V_1, V_2)$  is a condition that, simultaneously,  $V_1$  is committed to  $(a_1, a_2) \in \mathcal{A}_1 \times \mathcal{A}_2$  and  $V_2$  is committed to  $(a'_1, a'_2) \in \mathcal{A}_1 \times \mathcal{A}_2$ . The utility of a joint commitment would provide a complete description of the way the collective views all possible consequence profiles (one for each decision maker). It would provide information regarding the degree of conflict and the possibilities for compromise, since only one profile can actually be implemented by the collective.

If there were no conflicts, then there would exist a joint commitment of the form  $[(a_1^*, a_2^*), (a'_1, a'_2)]$  that would simultaneously maximize benefit to  $V_1$  and  $V_2$  and, hence, by the Pareto principle, to the collective. In the presence of conflicts, however, joint commitments of the form  $[(a_1, a_2), (a'_1, a'_2)]$ , where both commit to the same profile, would represent a compromise solution. The issue, then, is to define an acceptable compromise.

To determine the utility of a joint commitment, consider the following argument. For  $[(a_1, a_2), (a'_1, a'_2)]$  to be a joint commitment, it is necessary that  $(a_1, a_2)$  be a commitment by  $V_1$ . But if  $(a_1, a_2)$  is a commitment by  $V_1$ , then for  $(a'_1, a'_2)$  to be a commitment by  $V_2$ ,  $(a'_1, a'_2)$  must be a commitment by  $V_2$  given that  $(a_1, a_2)$  is a commitment by  $V_1$ . Furthermore, if  $(a_1, a_2)$  is not a commitment by  $V_1$ , then  $[(a_1, a_2), (a'_1, a'_2)]$  is not a joint commitment, regardless of whether or not  $(a'_1, a'_2)$  is a commitment by  $V_2$ . Thus, when considering the utility of a joint commitment to  $[(a_1, a_2), (a'_1, a'_2)]$ , if the utility of a commitment to  $(a_1, a_2)$  by  $V_1$  is considered first, then the utility of a commitment to  $(a'_1, a'_2)$  by  $V_2$  will be relevant only if  $(a_1, a_2)$  is a commitment by  $V_1$ . Consequently, given the utility of a commitment to  $(a_1, a_2)$  by  $V_1$  and the conditional utility of a commitment to  $(a'_1, a'_2)$  by  $V_2$  given that  $(a_1, a_2)$  is a commitment by  $V_1$ , knowledge of the categorical utility of a commitment to  $(a'_1, a'_2)$  by  $V_2$  is not required in order to compute the utility of a joint commitment to  $[(a_1, a_2), (a'_1, a'_2)]$ . Thus, the utility of a joint commitment to  $[(a_1, a_2), (a'_1, a'_2)]$  is a function of the categorical utility of a commitment to  $(a_1, a_2)$  by  $V_1$  and the conditional utility of a commitment by  $V_2$  to  $(a'_1, a'_2)$  given that  $(a_1, a_2)$  is a commitment by  $V_1$ .

Let  $\hat{u}_{V_1 V_2}$  denote the *utility of a joint commitment*. By the above arguments, this function can be expressed as

$$\hat{u}_{V_1 V_2}[(a_1, a_2), (a'_1, a'_2)] = F[u_{V_1}(a_1, a_2), u_{V_2|V_1}(a'_1, a'_2|a_1, a_2)] \quad (2)$$

for some function  $F$ , called the *aggregation function*.

## 4.2 Aggregation

Obviously, there are many possibilities for  $F$ , and to narrow the choices, it is necessary to impose some additional constraints. One reasonable constraint is that the collective possess at least a weak sense of equity so that a meaningful notion of cooperation can occur. Specifically, we wish to avoid a condition of *categorical subjugation*. To introduce this concept, let us restrict interest to the collective  $S_1$  and  $V_2$ . We shall say that  $S_1$  is categorically subjugated to the collective if every consequence profile that is acceptable to the collective would require  $S_1$  to sacrifice its performance. Suppose that

$$u_{S_1}(a'_1, a'_2) > u_{S_1}(a''_1, a''_2), \quad (3)$$

but

$$\hat{u}_{S_1 V_2}[(a'_1, a'_2), (a_1, a_2)] < \hat{u}_{S_1 V_2}[(a''_1, a''_2), (a_1, a_2)] \quad (4)$$

for all  $(a_1, a_2) \in \mathcal{A}_1 \times \mathcal{A}_2$ . Then  $S_1$  would be categorically subjugated, since  $S_1$ 's preferred joint action can never be preferred by the society. Avoiding categorical subjugation ensures that all participants have a “seat at the table” when negotiating. Otherwise, the interests of some facets will be so contrary to the interests of the collective that, no matter what the collective decides, the interests of the affected individual facets will be suppressed. Unless the possibility (although not the guarantee) exists that the interests of the individual are compatible with the interests of the collective, the individual will be effectively disenfranchised. Although categorical subjugation is not always avoided in human societies (e.g., dictatorships), avoiding categorical subjugation is an important feature of an artificial society that must negotiate to reach a compromise.

If categorical subjugation is to be avoided, then there must exist an action profile  $(\tilde{a}_1, \tilde{a}_2)$  such that, if (3) holds, then

$$\hat{u}_{S_1 V_2}[(a'_1, a'_2), (\tilde{a}_1, \tilde{a}_2)] \geq \hat{u}_{S_1 V_2}[(a''_1, a''_2), (\tilde{a}_1, \tilde{a}_2)]. \quad (5)$$

A similar argument regarding the categorical subjugation of, say,  $R_1$  can be made with the inequalities reversed in (3), (4), and (5) when  $R_1$  replaces  $S_1$ .

The question now becomes: what conditions are necessary to impose upon the aggregation function  $F$  to ensure that categorical subjugation can never occur? To address this

question, let us turn to an analogous issue. A Dutch book is a gambling situation such that, no matter what the outcome, the gambler will be worse off for having taken the gamble — a situation of *sure loss* (one’s reward is always less than one’s stake). To illustrate a Dutch book, Suppose  $Y$  can take one of two distinct values:  $y_1$  or  $y_2$ , and let  $q(y)$  denote a belief function of  $y$ ; that is  $q(y)$  measures the strength of belief that  $Y = y$ . Without loss of generality, we may restrict belief functions to the unit interval; that is,  $0 \leq q(y) \leq 1$ . (We refrain from using the term “probability” here, since we do not require  $q$  to possess all of the properties of a probability mass function.)

By convention, we will assume that we have full belief that exactly one of these values obtains, that is, that the disjunction of  $y_1$  and  $y_2$  must occur, and that beliefs are additive, thus,

$$q(y_1 \vee y_2) = q(y_1) + q(y_2) = 1. \quad (6)$$

Now let  $Z$  take on one of two distinct values  $z_1$  or  $z_2$ , and let  $r(z, y)$  denote the belief that  $Z = z$  and  $Y = y$  simultaneously. Let us now assume that

$$q(y_2) > q(y_1) \quad (7)$$

$$r(z_1, y_2) < r(z_1, y_1) \quad (8)$$

$$r(z_2, y_2) < r(z_2, y_1). \quad (9)$$

The following example illustrates a Dutch book. Suppose you purchase a \$1 gamble that  $Y = y_2$ , and deem a fair purchase price to be  $q(y_2)$ ; that is, you pay  $\$q(y_2)$  for the gamble to win \$1. Now also suppose you sell the gamble  $(z_1, y_2) \vee (z_2, y_2)$ . By additivity of beliefs, a fair selling price for this bet would be  $r[(z_1, y_2) \vee (z_2, y_2)] = r(z_1, y_2) + r(z_2, y_2)$ . However, according to the above ordering, you must have  $q(y_2) > \frac{1}{2}$  and, since  $r(z_1, y_2) + r(z_2, y_2) < r(z_1, y_1) + r(z_2, y_1)$ , it follows that  $r[(z_1, y_2) \vee (z_2, y_2)] < \frac{1}{2}$ . After all gambles have been bought and sold, your net wealth is  $r[(z_1, y_2) \vee (z_2, y_2)] - q(y_2) < 0$ . To overcome this loss, you hope to make up the difference once the outcome of the gamble is known. But if neither  $y_2$  nor  $(z_1, y_2) \vee (z_2, y_2)$  occur, you win nothing and you pay nothing, and if  $(z_1, y_2) \vee (z_2, y_2)$  occurs, then, of course,  $y_2$  occurs, so you win \$1 which you must pay to the buyer of your gamble. Thus, once the gambles have been bought and sold, your net wealth is invariant to whatever happens — you suffer a sure loss.

A belief system is said to be *coherent* if it is not possible to construct a Dutch book. The Dutch Book Theorem [11, 38] and its converse [22] state that a belief system is coherent if and only if it complies with a probability measure that describes the degrees of belief regarding the propositions under consideration. The above example does not comply with the laws of probability theory, since  $q(y_2) \neq r(z_1, y_2) + r(z_2, y_2)$ ; that is, marginalization fails.

The above discussion illustrates the fact that categorical subjugation and sure loss are mathematically equivalent. Thus, if a multi-agent valuation system is to be coherent, in that it is not not possible to construct a situation where categorical subjugation can occur, then the valuation system must comply with the mathematical structure of probability theory.

**Definition 4** Let  $u_{V_i}$  denote a categorical utility for  $V_i$ . The collective  $\mathbf{V}$  is *coherent* if, for each  $i \in \{1, \dots, 2n\}$ , given that  $u_{V_i}(\mathbf{a}) > u_{V_i}(\mathbf{a}')$ , there exists a commitment sub-profile  $(\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_{i-1}, \tilde{\mathbf{a}}_{i+1}, \dots, \tilde{\mathbf{a}}_{2n})$  such that

$$\hat{u}_{V_1 \dots V_{2n}}(\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_{i-1}, \mathbf{a}, \tilde{\mathbf{a}}_{i+1}, \dots, \tilde{\mathbf{a}}_{2n}) \geq \hat{u}_{V_1 \dots V_{2n}}(\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_{i-1}, \mathbf{a}', \tilde{\mathbf{a}}_{i+1}, \dots, \tilde{\mathbf{a}}_{2n}) \quad (10)$$

if  $V_i$  is a selecting facet, with the inequalities reversed if  $V_i$  is a rejecting facet.

Let  $\mathbf{V}$  be a group of decision making facets whose influence relationships can be expressed with a directed acyclic graph. For each  $V_i$ , let  $\text{pa}(V_i) = (V_{i_1}, \dots, V_{i_{p_i}})$  denote the  $p_i$  parents of  $V_i$ , and let  $\mathcal{A}^{p_i} = \mathcal{A} \times \dots \times \mathcal{A}$  ( $p_i$  times) denote the  $p_i$ -fold product of the joint action space corresponding to the parents of  $V_i$ . If  $V_i$  has no parents, then  $\mathcal{A}^{p_i} = \emptyset$ . Let  $\text{cp}(V_i) = (\mathbf{a}_{i_1}, \dots, \mathbf{a}_{i_{p_i}})$  denote the commitment profile for  $\text{pa}(V_i)$ . For each  $V_i$ ,  $u_{V_i | \text{pa}(V_i)}[\mathbf{a} | \text{cp}(V_i)]$  is the utility that  $V_i$  ascribes to  $\mathbf{a}$ , given that  $V_{i_j}$  commits to  $\mathbf{a}_{i_j}$ ,  $j = 1, \dots, p_i$ . If  $V_i$  has no parents, the conditional utility is the categorical utility; i.e.,  $u_{V_i | \text{pa}(V_i)} = u_{V_i}$  if  $\text{pa}(V_i) = \emptyset$ .

**Theorem 1** *If a society can be represented as a directed acyclic graph, categorical subjugation cannot occur if and only if the utilities  $u_{V_i | \text{pa}(V_i)}$  are conditional mass functions. That is,*

$$u_{V_i | \text{pa}(V_i)}[\mathbf{a} | \text{cp}(V_i)] \geq 0 \quad \forall \mathbf{x} \in \mathcal{A} \quad (11)$$

and

$$\sum_{\mathbf{a} \in \mathcal{A}} u_{V_i | \text{pa}(V_i)}[\mathbf{a} | \text{cp}(V_i)] = 1 \quad (12)$$

for all  $(\mathbf{a}_{i_1}, \dots, \mathbf{a}_{i_{p_i}}) \in \mathcal{A}^{p_i}$ . Furthermore, the utility of a joint commitment to  $(\mathbf{a}_1, \dots, \mathbf{a}_{2n})$  is

$$\hat{u}_{\mathbf{V}}(\mathbf{a}_1, \dots, \mathbf{a}_{2n}) = \prod_{i=1}^{2n} u_{V_i | \text{pa}(V_i)}[\mathbf{a}_i | \text{cp}(V_i)] \quad (13)$$

or, more specifically,

$$\hat{u}_{SR}(\mathbf{a}_1, \dots, \mathbf{a}_n, \mathbf{a}'_1, \dots, \mathbf{a}'_n) = \prod_{i=1}^n \prod_{j=1}^n u_{S_i | \text{pa}(S_i)}[\mathbf{a}_i | \text{cp}(S_i)] u_{R_j | \text{pa}(R_j)}[\mathbf{a}'_j | \text{cp}(R_j)], \quad (14)$$

where  $\mathbf{a}_i$  is the commitment by  $S_i$ ,  $i = 1, \dots, n$  and  $\mathbf{a}'_j$  is the commitment by  $R_j$ ,  $j = 1, \dots, n$ .

*Proof:* Mathematically (albeit with different semantics), we may view  $V_i$  as random variables defined over the sample spaces  $\mathcal{V}_i$ ,  $i = 1, \dots, 2n$ . The Dutch Book Theorem and its converse establish that the necessary and sufficient condition to ensure that sure loss (categorical subjugation) cannot occur is that  $u_{\mathcal{V}_i | \text{pa}(V_i)}$  must correspond to the conditional probability mass functions of  $V_i$  given  $\text{cp}(V_i)$ . Thus, the categorical utilities of the root vertices must possess the mathematical structure of marginal probability mass function and the conditional utility of non-root vertices possesses the mathematical structure of conditional probability mass functions. Consequently, the vertices and edges of the DAG satisfy all of the conditions of a Bayesian network, and we may apply the fundamental theorem of Bayesian networks; namely, that the joint probability mass function of the random variables associated with the vertices is the product of the conditional probability mass functions of all non-root vertices, and the marginal mass functions of all root vertices [9, 19, 34]. Equation (13) is simply an application of the law of compound probability. Thus, coherence is established.  $\square$

We will term utilities that comply with Theorem 1 *praxeic utilities*. It should be noted that this formulation requires all utilities to be non-negative and sum to unity. This restriction, however, does not reduce the generality of the theorem, since utilities can be subjected to positive affine transformations without affecting the solution.

Equation (14) expresses the values of the selecting and rejecting facets simultaneously. Since parents of selecting facets may comprise both selecting and rejecting facets and similarly for the parents of rejecting facets, this function contains all of the possibilities for compromise and conflict. To be useful for decision making, however, it is necessary to compute the joint selectability for all joint commitments by the selecting facets, and the joint rejectability for all joint commitments by the rejecting facets. Since  $\hat{u}_{SR}$  is a multivariate mass function, we may compute the *joint selectability and rejectability marginals* as

$$\hat{u}_S(\mathbf{a}_1, \dots, \mathbf{a}_n) = \sum_{\mathbf{a}'_1, \dots, \mathbf{a}'_n} \hat{u}_{SR}(\mathbf{a}_1, \dots, \mathbf{a}_n, \mathbf{a}'_1, \dots, \mathbf{a}'_n) \quad (15)$$

$$\hat{u}_R(\mathbf{a}'_1, \dots, \mathbf{a}'_n) = \sum_{\mathbf{a}_1, \dots, \mathbf{a}_n} \hat{u}_{SR}(\mathbf{a}_1, \dots, \mathbf{a}_n, \mathbf{a}'_1, \dots, \mathbf{a}'_n) \quad (16)$$

Once the joint selectability and rejectability marginals have been computed, we are in a position to define a satisficing social welfare function. We first observe that, since only one consequence profile can be implemented, to make a decision, we must ascribe the same commitment to each facet, yielding the *joint praxeic selectability and joint praxeic rejectability*

$$\tilde{u}_S(\mathbf{a}) = \hat{u}_S(\mathbf{a}, \dots, \mathbf{a}) \quad (17)$$

and

$$\tilde{u}_R(\mathbf{a}) = \hat{u}_R(\mathbf{a}, \dots, \mathbf{a}). \quad (18)$$

We next define *satisficing social welfare function*

$$W(\mathbf{a}) = \tilde{u}_S(\mathbf{a}) - q\tilde{u}_R(\mathbf{a}) \quad (19)$$

and the *jointly satisficing set*

$$\Sigma_q = \{\mathbf{a}: W(\mathbf{a}) \geq 0\}. \quad (20)$$

The parameter  $q$  is a measure of caution. Nominally,  $q = 1$ , but as  $q$  decreases, the number of consequence profiles that are rejected decreases. As will subsequently become apparent, another interpretation of  $q$  is as an *index of negotiation*, since lowering  $q$  enlarges the satisficing set, thereby increasing the opportunities for reaching a compromise. We will define all consequence profiles such that the satisficing social welfare is non-negative as being satisficing.

We may also compute the *individual selectability and rejectability marginal utilities* as

$$\tilde{u}_{S_i}(\mathbf{a}_i) = \sum_{\neg \mathbf{a}_i} \hat{u}_S(\mathbf{a}_1, \dots, \mathbf{a}_n) \quad (21)$$

and

$$\tilde{u}_{R_i}(\mathbf{a}_i) = \sum_{\neg \mathbf{a}_i} \hat{u}_R(\mathbf{a}_1, \dots, \mathbf{a}_n), \quad (22)$$

where we have employed the so-called *not-sum* notation; namely,  $\sum_{\neg \mathbf{a}_i}$  to mean that the sum is taken over all  $\mathbf{a}_j$  for  $j \neq i$ .

The individual welfare function is

$$W_i(\mathbf{a}) = \tilde{u}_{S_i}(\mathbf{a}_i) - q_i \tilde{u}_{R_i}(\mathbf{a}_i) \quad (23)$$

and the *individually satisficing set* is

$$\Sigma_{q_i}^i = \{\mathbf{a}: W_i(\mathbf{a}) \geq 0\}. \quad (24)$$

The *compromise set* is the set of all joint actions that are simultaneously satisficing for the group and for the individuals; that is,

$$\mathcal{C} = \Sigma_q \cap \Sigma_{q_1}^1 \cap \dots \cap \Sigma_{q_n}^n. \quad (25)$$

A satisficing set (either for the group or individuals) constitutes the set of consequences for which effectiveness, as measured by the selectability utility, is at least as great as  $q_i$  times the inefficiency, as measured by the rejectability inutility. Rather than focusing on seeking the best and only the best solution, the satisficing methodology focuses on eliminating bad solutions. Since the satisficing set eliminates all alternatives whose effectiveness does not exceed their efficiency, it is optimal in the sense that it eliminates the maximum number of bad choices. If, in the extreme case, all but one choice are eliminated, then the satisficing solution

coincides with the optimal solution. Thus, far from being a boundedly rational solution, the set of satisficing solutions possess a well-defined notion of optimality (albeit different). Thus, we come to Euler’s conclusion through the “back door.”

*Example 1 The Social Prisoner’s Dilemma.* The conventional Prisoner’s Dilemma game is designed to characterize behavior between two decision makers in an environment where cooperation leads to better results than does defection but, if only one attempts to cooperate, that individual becomes vulnerable to being exploited by the other. Classically, this game is defined in terms of categorical utilities. Let  $C$  and  $D$  denote cooperation and defection, respectively. The corresponding categorical utilities are the entries of the payoff matrix displayed in Figure 1. The joint option  $(C, C)$  (next best for both) is the Pareto optimal solution, while  $(D, D)$  (next worst for both) is the Nash equilibrium solution. Notice that the game is symmetrical. The classical assumption for this game is that there is no social relationship between the players, and that each is intent on, and only on, maximizing its own welfare, regardless of the effect doing so has on the other.

**Table 1** The payoff matrix for the conventional Prisoner’s Dilemma game.


$X_1$	$X_2$	
	$C$	$D$
$C$	(3, 3)	(1, 4)
$D$	(4, 1)	(2, 2)

Now let us add some social context to this problem. Suppose a leader-follower relationship exists between them, with  $X_1$  being the leader and  $X_2$  the follower. We shall assume that  $X_1$  follows the conventional structure of maximizing payoff, but  $X_2$  is interested in (a) following the lead of  $X_1$ , (b) resisting exploitation, and (c) not offending  $X_1$  by taking advantage of the possible propensity for  $X_1$  to cooperate. We shall take the definition of selectability as the same as with the conventional formulation; namely, to seek to maximize payoff. For rejectability, however, we invoke a component that is not present in the conventional formulation; namely, to account for social issues, and assume that the players have a unit of social resource they may commit to each outcome. Since the leader has no social commitments, we take rejectability as the same for each outcome. Accordingly, the categorical selectability and rejectability values for the leader are provided in Table 2.

To account for the social context, we take the utilities for  $X_2$  to be conditional, and assume that both selectability and the rejectability of  $X_2$  are influenced by the selectability of  $X_1$ , as indicated in Figure 3.

**Table 2** The categorical selectability and rejectability for the Prisoner’s Dilemma leader.

	$(C, C)$	$(C, D)$	$(D, C)$	$(D, D)$
$u_{S_1}$	$\frac{3}{10}$	$\frac{1}{10}$	$\frac{4}{10}$	$\frac{2}{10}$
$u_{R_1}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$



**Fig. 3** The influence network for the social Prisoner’s Dilemma game

Table 3 displays the conditional selectability for  $X_2$  given the commitments of  $X_1$ . If  $X_1$  were to commit to  $(C, C)$ ,  $((D, C)$ , or  $(D, D)$ , then  $X_2$  would do likewise in the interest of maximizing its payoff. But if  $X_1$  were to commit to  $(C, D)$ , then  $X_2$  would resist being exploited by placing zero conditional utility on  $(C, D)$  and apportioning equally to the other outcomes.

**Table 3**  $R_2$ ’s conditional selectability for the social Prisoner’s Dilemma game.

	$(x_1, x_2)$			
	$(C, C)$	$(C, D)$	$(D, D)$	$(D, D)$
$u_{S_2 S_1}(x_1, x_2 C, C)$	1	0	0	0
$u_{S_2 S_1}(x_1, x_2 C, D)$	$\frac{1}{3}$	0	$\frac{1}{3}$	$\frac{1}{3}$
$u_{S_2 S_1}(x_1, x_2 D, C)$	0	0	1	0
$u_{S_2 S_1}(x_1, x_2 D, D)$	0	0	0	1

Table 4 displays the conditional rejectability for  $X_2$  given the commitments of  $X_1$ . If  $X_1$  were to commit to  $(C, C)$ , then  $X_2$  would place zero conditional rejectability on that outcome and apportion all of its conditional rejectability equally to the other outcomes. If  $X_1$  were to commit to  $(C, D)$ ,  $X_2$  would place all of its rejectability on that outcome to ensure it will not be exploited. If  $X_2$  were to commit to  $(D, C)$ , then  $X_2$  would not reject that outcome so as to not exploit  $X_1$ , and instead would reject exploitation by placing its conditional rejectability on  $(C, D)$ . Finally, if  $X_1$  were to commit to  $(D, D)$ ,  $X_2$  would not reject that outcome, but would instead reject  $(C, D)$  as before.

The utility of a joint commitment is given by

$$u_{S_1 S_2 R_1 R_2}[(x_1, x_2), (x'_1, x'_2), (y_1, y_2), (y'_1, y'_2)] = u_{S_1}(x_1, x_2)u_{S_2|S_1}(x'_1, x'_2|x_1, x_2)u_{R_1}(y_1, y_2)u_{R_1|S_1}[(y'_1, y'_2)]. \quad (26)$$

The joint praxeic selectability and rejectability functions, as defined by (17) and (18) are given in Table 5, and the

**Table 4**  $R_2$ 's conditional rejectability for the social Prisoner's Dilemma game.

	$(x_1, x_2)$			
	$(C, C)$	$(C, D)$	$(D, D)$	$(D, C)$
$u_{R_2 S_1}(x_1, x_2 C, C)$	0	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
$u_{R_2 S_1}(x_1, x_2 C, D)$	0	1	0	0
$u_{R_2 S_1}(x_1, x_2 D, C)$	0	1	0	0
$u_{R_2 S_1}(x_1, x_2 D, D)$	0	1	0	0

jointly satisficing set (with  $q = 1$ ) is  $\{(C, C), (D, D)\}$ . The individual selectability and rejectability marginal utilities, as defined by (21) and (22) are displayed in Table 6, from which it can be seen that the individually satisficing sets (with  $q_1 = q_2 = 1$ ) are  $\Sigma_1^1 = \{(C, C), (D, C)\}$  and  $\Sigma_1^2 = \{(C, C), (D, D)\}$ . Consequently, the compromise set is  $\mathcal{C} = \{(C, C)\}$ .

Under the classical formulation of the Prisoner's Dilemma, the only rationally justifiable solution is mutual defection, since that formulation does not take into consideration any social relationships. From the classical point of view, mutual cooperation, although Pareto optimal, cannot be justified. The social version of the game as developed here, however, indicates that mutual cooperation is the only justified solution.

**Table 5** The joint praxeic selectability and rejectability for the social Prisoner's Dilemma game.

$[(x_1, x_2), (x_1, x_2)]$	$\tilde{u}_{S_1 S_2}[(x_1, x_2), (x_1, x_2)]$
$[(C, C), (C, C)]$	0.3
$[(C, D), (C, D)]$	0.0
$[(D, C), (D, C)]$	0.0
$[(D, D), (D, D)]$	0.2

$[(x_1, x_2), (x_1, x_2)]$	$\tilde{u}_{R_1 R_2}[(x_1, x_2), (x_1, x_2)]$
$[(C, C), (C, C)]$	0.0
$[(C, D), (C, D)]$	0.2
$[(D, C), (D, C)]$	0.025
$[(D, D), (D, D)]$	0.025

**Table 6** The individual selectability and rejectability utilities for the social Prisoner's Dilemma game.

	$(x_1, x_2)$			
	$(C, C)$	$(C, D)$	$(D, D)$	$(D, C)$
$\tilde{u}_{S_1}(x_1, x_2)$	0.3	0.1	0.4	0.2
$\tilde{u}_{R_1}(x_1, x_2)$	0.25	0.25	0.25	0.25
$\tilde{u}_{S_2}(x_1, x_2)$	0.3	0.0	0.0	0.7
$\tilde{u}_{R_2}(x_1, x_2)$	0.25	0.25	0.25	0.25

## 5 Reconciliation with Classical Theory

Not all problems fit naturally into the dual-utility structure of satisficing theory. One way to deal with this situation, while still retaining some of the flavor of satisficing theory, is to invoke the assumption that all consequences are rejectability neutral, and ascribe all meaningful utility to selectability. Under this situation, we set the rejectability to a constant:  $u_{R_j|pa(R_j)}[\mathbf{a}_j | cp(R_j)] = K_j = \frac{1}{|\mathcal{A}_j|}$  ( $|\cdot|$  denotes cardinality) for all  $\mathbf{a}_j$ . We define the conditional utility of  $X_i$  as

$$u_{X_i|pa(X_i)}[\mathbf{a}_i | cp(X_i)] = u_{S_i|pa(S_i)}[\mathbf{a}_i | cp(S_i)]. \quad (27)$$

Thus, (14) becomes a function of  $\mathbf{a}_i$ ,  $i = 1, \dots, n$ , only, and we may write

$$\hat{u}_{\mathbf{x}}(\mathbf{a}_1, \dots, \mathbf{a}_n) = \prod_{i=1}^n u_{X_i|pa(X_i)}[\mathbf{a}_i | cp(X_i)], \quad (28)$$

and the marginals become

$$\tilde{u}_{X_i}(\mathbf{a}_i) = \sum_{\neg \mathbf{a}_i} \hat{u}_{\mathbf{x}}(\mathbf{a}_1, \dots, \mathbf{a}_n). \quad (29)$$

Once all of the valuations are concentrated in a single utility, we view the decision problem from the classical perspective of optimization. The most well-known solution concept for individuals is the non-cooperative game theoretic concept of Nash equilibria [31]. Let  $\mathbf{a}^* = (a_1^*, \dots, a_n^*)$ . The action profile  $\mathbf{a}^*$  is a *Nash equilibrium* if, were any single individual to alter its choice, its utility would decrease; i.e., if  $\mathbf{a}^\dagger = (a_1^*, \dots, a_i', \dots, a_n^*)$ , then, in terms of categorical utilities,

$$u_{X_i}(\mathbf{a}^*) \geq u_{X_i}(\mathbf{a}^\dagger) \quad (30)$$

for all  $a_i' \in \mathcal{A}_i \setminus \{a_i^*\}$  for  $i = 1, \dots, n$ .

When conditional utilities are involved, we may define two notions of equilibrium. First, let us define what might be called a conditional Nash equilibrium. The action profile  $\mathbf{a}^*$  is a *conditional Nash equilibrium* if

$$u_{X_i|pa(X_i)}(\mathbf{a}^* | \mathbf{a}^*, \dots, \mathbf{a}^*) \geq u_{X_i|pa(X_i)}(\mathbf{a}^\dagger | \mathbf{a}^\dagger, \dots, \mathbf{a}^\dagger) \quad (31)$$

for all  $a_i' \in \mathcal{A}_i \setminus \{a_i^*\}$  for  $i = 1, \dots, n$ .

We may also compute the Nash equilibrium in terms of the marginal utility defined by (29). The action profile  $\mathbf{a}^*$  is a Nash equilibrium if

$$\tilde{u}_{X_i}(\mathbf{a}^*) \geq \tilde{u}_{X_i}(\mathbf{a}^\dagger) \quad (32)$$

*Example 2 Prisoner's Dilemma, Continued.* Let us revisit the Prisoner's Dilemma discussed in Example 1 under the assumption of neutral rejectability, and set  $u_{X_1} = u_{S_1}$  and  $u_{X_2|pa(X_2)} = u_{S_2|pa(X_2)}$  as defined in Tables 2 and 3, respectively. By inspection, we see that the conditional Nash

**Table 7** The marginal utilities for the Prisoner's Dilemma.

	$(C, C)$	$(C, D)$	$(D, C)$	$(D, D)$
$u_{X_1}$	$\frac{3}{10}$	$\frac{1}{10}$	$\frac{4}{10}$	$\frac{2}{10}$
$u_{X_2}$	$\frac{3}{10}$	0	0	$\frac{7}{10}$

equilibrium is  $(D, D)$ , as with the conventional formulation. Furthermore, the marginal utilities are given in Table 7, and we see that, again the Nash equilibrium is  $(D, D)$ .

The Nash equilibrium is usually considered to be an appropriate solution concept for non-cooperative games. On the other hand, with a cooperative game (i.e., one where binding agreements are possible), it may be possible to enter into negotiations and bargain for a solution. For the players to forge an agreement, however, each must achieve an acceptable degree of satisfaction. A *bargaining game* is a cooperative game in which each participant possesses a *disagreement point* that defines the benefit that is guaranteed to accrue to it if a compromise cannot be reached. The disagreement point, therefore, is an indication of the strategic strength that is conferred on the participant as it participates in negotiations: the higher the disagreement point, the greater bargaining strength of the participant.

A well-known bargaining concept that offers a clear definition of individual acceptability is the Nash bargain [30], which permits each participant to make maximal use of its strategic strength. The approach is based on four fundamental principles: (i) invariance to positive affine transformations; (ii) Pareto optimality; (iii) independence of irrelevant alternatives, and (iv) symmetry, which is the notion that no individual agent can expect that the other agents will grant it better terms than that individual itself would be willing to grant, were roles reversed.

Nash showed that these four conditions lead to a unique solution. Let  $d_{X_i}$  denote the disagreement point for  $X_i$ . The *negotiation set*, denoted  $\mathcal{N}$ , is the subset of action profiles such that every participant achieves at least its disagreement point. In terms of categorical utilities, the negotiation set is

$$\mathcal{N} = \{\mathbf{a} \in \mathcal{A}: u_{X_i}(\mathbf{a}) \geq d_{X_i}, i = 1, \dots, n\} \quad (33)$$

and the Nash bargain is

$$\mathbf{a}_N = \arg \max_{\mathbf{a} \in \mathcal{N}} \prod_{i=1}^n [u_{X_i}(\mathbf{a}) - d_{X_i}]. \quad (34)$$

The intuitive interpretation of a Nash bargain is that it defines a fair compromise. It enables each player to take advantage of the strategic strength endowed by its disagreement point. The higher  $X_i$ 's disagreement point, the more action profiles that are unfavorable to it are eliminated.

The structure of (34) suggests that the optimal group solution can be interpreted as a Nash bargain with unilateral

utilities replaced by conditional utilities and all disagreement points set to zero. Analogously, therefore, we may define a *conditional Nash bargaining solution*. When decisions are made under certainty, the negotiation set is defined as

$$\mathcal{N} = \{\mathbf{a} \in \mathcal{A}: u_{X_i | \text{pa}(X_i)}(\mathbf{a}) | \text{cp}(X_i) \geq d_{X_i}, i = 1, \dots, n\}. \quad (35)$$

The conditional Nash bargaining solution is

$$\mathbf{a}_N = \arg \max_{\mathbf{a} \in \mathcal{N}} \prod_{i=1}^n [u_{X_i | \text{pa}(X_i)}[\mathbf{a} | \text{cp}(X_i)] - d_{X_i}]. \quad (36)$$

Referring again to the Prisoner's Dilemma example, it is easily seen that both the conditional Nash bargain is the same as the conventional Nash bargain for the Prisoner's Dilemma; namely, the Pareto optimal solution  $(C, C)$ .

## 6 Conditionally Decoupled Societies

### 6.1 The General Case

The approach developed above assumes that the conditional preferences are defined over the entire product action space. In this respect conditional preferences are generalizations of classical categorical preferences, the difference being that the preferences can be modulated by the commitments of others. Although increased complexity is associated with the introduction of conditional preferences, there are cases where this additional complexity is not justified. It can be the case that the only commitments that affect the preferences of an agent are the direct consequences to its parents. This situation motivates the notion of conditional decoupling.

**Definition 5** A society is *conditionally decoupled* if the conditional preference of each agent is a function only of its own actions, given the commitments of its parents to their own actions.

Whereas, for a non-decoupled system, the utilities are functions of the entire action profile, for a decoupled system, the utilities are functions of individual actions. To develop this concept, suppose  $\text{cp}(V_i) = \{\mathbf{a}_{i_1}, \dots, \mathbf{a}_{i_{p_i}}\}$ . Then the conditional utility

$$u_{V_i | \text{pa}(V_i)}[\mathbf{a} | \text{cp}(V_i)] = u_{V_i | \text{pa}(V_i)}(\mathbf{a} | \mathbf{a}_{i_1}, \dots, \mathbf{a}_{i_{p_i}}) \quad (37)$$

becomes

$$u_{V_i | \text{pa}(V_i)}[a_i | \text{cp}(V_i)] = u_{V_i | \text{pa}(V_i)}(a_i | a_{i_1}, \dots, a_{i_{p_i}}) \quad (38)$$

Then (14) becomes

$$\hat{u}_{SR}(a_1, \dots, a_n, a'_1, \dots, a'_n) = \prod_{i=1}^n u_{S_j | \text{pa}(S_j)}[a_j | \text{cp}(S_j)] \prod_{j=1}^n u_{R_j | \text{pa}(R_j)}[a'_j | \text{cp}(R_j)] \quad (39)$$

The corresponding joint selectability and rejectability marginals are given by

$$u_S(a_1, \dots, a_n) = \sum_{(a'_1, \dots, a'_n)} u_{SR}(a_1, \dots, a_n, a'_1, \dots, a'_n) \quad (40)$$

and

$$u_R(a'_1, \dots, a'_n) = \sum_{(a_1, \dots, a_n)} u_{SR}(a_1, \dots, a_n, a'_1, \dots, a'_n). \quad (41)$$

We may now define a *social welfare function* as

$$W(a_1, \dots, a_n) = u_S(a_1, \dots, a_n) - q_G u_R(a_1, \dots, a_n) \quad (42)$$

where  $q_G$  is the joint  $q$ -value for the group. The jointly satisficing set is the set of action profiles that are jointly satisficing for the society as a whole, and is defined as

$$\mathcal{S} = \{(a_1, \dots, a_n) \in \mathcal{A}: W(a_1, \dots, a_n) \geq 0\}. \quad (43)$$

This procedure, however, does not account for the possibility that the elements of  $\mathcal{S}$  may not be acceptable to all (or any) of the individuals. Thus, we must also compute the individual satisficing sets. To proceed, we must first compute the selectability and rejectability marginals as

$$u_{S_i}(a_i) = \sum_{\neg a_i} u_S(a_1, \dots, a_n) \quad (44)$$

and

$$u_{R_i}(a_i) = \sum_{\neg a_i} u_R(a_1, \dots, a_n), \quad (45)$$

respectively. We may then define the individually satisficing sets as

$$\Sigma_i = \{a_i \in \mathcal{A}_i: u_{S_i}(a_i) - q_i u_{R_i}(a_i)\}. \quad (46)$$

this set includes all alternatives that are satisficing, or good enough, for  $X_i$ . The *satisficing rectangle* is the set of all action profiles such that each component is individually satisficing, and is given by

$$\mathcal{R} = \Sigma_1 \times \dots \times \Sigma_n. \quad (47)$$

The intersection of the jointly satisficing set and the satisficing rectangle yields the *compromise set*, comprising the action profiles that are simultaneously good enough for the group and for each individual.

$$\mathcal{C} = \mathcal{S} \cap \mathcal{R}. \quad (48)$$

If  $\mathcal{C} \neq \emptyset$ , then we may form a *best compromise* as

$$\mathbf{a}^* = \arg \max_{\mathbf{a} \in \mathcal{C}} W(\mathbf{a}). \quad (49)$$

If  $\mathcal{C} = \emptyset$ , then there are no action profiles that are simultaneously good enough for the group and each individual. However, the satisficing approach provides a natural and systematic negotiation framework which which each individual can control the degree to which it is willing to lower its standards in an attempt to reach a compromise. By lowering its  $q_i$ -value incrementally, each  $X_i$  increases the size of its satisficing set. By specifying the increment  $\Delta q_i$  that  $X_i$  is willing to reduce its standards, each participant can control the amount of compromise it is willing to offer others. If enough participants are willing to lower their  $q$ -values sufficiently, it is easy to see that, eventually, the consensus set will be non-empty, and a best compromise can be achieved. Although such negotiations may fail to reach a compromise that is acceptable to all members, the significant aspect of this type of negotiation is that no individual is *a priori* subjugated to the will of the society in the sense that there is no possibility for that individual's preferences to receive consideration. Thus, every individual can be assured of receiving sufficient benefit, by its own definition, before agreeing to the compromise. If an individual could not enjoy at least that minimal assurance, it may not be inclined to join or remain affiliated with a society.

## 6.2 Social Choice

With the general multi-agent decision problem, each individual possesses its own action set. Some scenarios, however, are such that there is only one action set that applies to the group as a whole. Scenarios of this type are termed *social choice* problems. Thus, with a social choice problem, there is only one action space  $\mathcal{A}$ . The social welfare function (42) becomes

$$W(a) = u_S(a, \dots, a) - q_G u_R(a, \dots, a) \quad (50)$$

and the jointly satisficing set becomes

$$\mathcal{S} = \{a \in \mathcal{A}: W(a) \geq 0\}. \quad (51)$$

The individual selectability and rejectability marginals are computed according to (44) and (45), and the individually satisficing sets  $\Sigma_i$ ,  $i = 1, \dots, n$  are given by (46). The intersection of the individually satisficing sets and the jointly satisficing set forms the *social compromise set*

$$\mathcal{C}_S = \Sigma_1 \cap \dots \cap \Sigma_n. \quad (52)$$

If  $\mathcal{S} = \emptyset$ , then there is no group action that is good enough for the group and each individual. However, by reducing the  $q$ -values incrementally as discussed above, a consensus will eventually emerge. The *social consensus* set is the intersection of the social compromise set and the jointly satisficing set

$$\mathcal{G}_S = \mathcal{S} \cap \mathcal{C}_S. \quad (53)$$

The *best compromise* is the action in this set that maximizes the social welfare function; that is,

$$a^* = \arg \max_{a \in \mathcal{G}_S} W(a). \quad (54)$$

*Example 3 The Family Walk.* Suppose a family, consisting of a father, mother, and child, is take one of three possible nature walk, denoted  $\{w, w', w''\}$ . The father prefers long hikes, the mother prefers beautiful scenery, and the child prefers an easy walk.

The first order of business in framing this in the satisficing context is to settle on operational definitions for the notions of selectability and rejectability. From the point of view of each individual, the main goal of walk is enjoyment according to its own criterion. Thus, it is reasonable to associate selectability with the degree of narrow self-interest. Accordingly, we define the three selectability utilities in Table 8.

**Table 8** Individual selectability utilities.

$x$	$u_{S_1}(x)$	$u_{S_2}(x)$	$u_{S_3}(x)$
$w$	0.1	0.4	0.3
$w'$	0.3	0.4	0.6
$w''$	0.6	0.2	0.1

As the operational definition of rejectability, we assume that each agent has a unit of concern for the interests of others. Let us first consider the mother. Since she has concern for the interests of her child, she will encode this information in a rejectability function that is conditioned on the selectability commitment of her child, as illustrated in Table 9. To interpret this table, consider the first column, which corresponds to  $u_{R_2|S_1}(\cdot|w)$ ; that is, the child commits to selecting  $w$ . Since this walk is tied for the most preferred by the mother, she ascribes no conditional rejectability to that alternative, and places all of her conditional rejectability mass on  $w'$  and  $w''$  in inverse proportion to her her selectability. Similar arguments apply if the child commits to  $w'$  or  $w''$ .

**Table 9** Mother's conditional rejectability  $u_{R_2|S_1}$ .

$x$	$u_{R_2 S_1}(x w)$	$u_{R_2 S_1}(x w')$	$u_{R_2 S_1}(x w'')$
$w$	0.0	0.4	0.5
$w'$	0.4	0.0	0.5
$w''$	0.6	0.6	0.0

The father's role in this decision process is first to defer first to the commitments of his child, then to the commitments to his wife, and then, subject to those constraints, to reject the alternative that is least preferred in terms of his narrow self-interest. These values are provided in Table 10.

**Table 10** Father's conditional rejectability  $u_{R_1|S_1S_2}$ .

	$x$		
	$w$	$w'$	$w''$
$u_{R_3 S_1S_2}(x w, w)$	0	0	1
$u_{R_3 S_1S_2}(x w, w')$	0	0	1
$u_{R_3 S_1S_2}(x w, w'')$	0	1	0

	$x$		
	$w$	$w'$	$w''$
$u_{R_3 S_1S_2}(x w', w)$	0	0	1
$u_{R_3 S_1S_2}(x w', w')$	0	0	1
$u_{R_3 S_1S_2}(x w', w'')$	1	0	0


  

	$x$		
	$w$	$w'$	$w''$
$u_{R_3 S_1S_2}(x w'', w)$	0	1	0
$u_{R_3 S_1S_2}(x w'', w')$	1	0	0
$u_{R_3 S_1S_2}(x w'', w'')$	1	0	0

Finally, we must specify the child's rejectability. This rejectability is not conditioned, since the model does not call for the child's preferences to be influenced by the parents' preferences. Thus, the child's concern for the interests of others is neutral; that is, the child's rejectability function is uniform, as provided in (55).

$$u_{R_1}(w) = u_{R_1}(w') = u_{R_1}(w'') = \frac{1}{3} \quad (55)$$

Figure 4 illustrates the influence flows of the satisficing praxeic network for the family walk.



**Fig. 4** A satisficing praxeic network for the family walk

Using the values provided in the above tables, we may compute the social welfare function, yielding

$$W(w) = -0.05$$

$$W(w') = 0.36667$$

$$W(w'') = -0.052;$$

$$\text{hence } \mathcal{S} = \{w'\}.$$

We next compute the individually satisficing sets, yielding

$$u_{S_1}(w) - q_1 u_{R_1}(w) = -0.233$$

$$u_{S_1}(w') - q_1 u_{R_1}(w') = -0.033$$

$$u_{S_1}(w'') - q_1 u_{R_1}(w'') = 0.267,$$

$$\begin{aligned} u_{S_2}(w) - q_2 u_{R_2}(w) &= 0.14 \\ u_{S_2}(w') - q_2 u_{R_2}(w') &= 0.14 \\ u_{S_2}(w'') - q_2 u_{R_2}(w'') &= -0.06, \end{aligned}$$

and

$$\begin{aligned} u_{S_3}(w) - q_3 u_{R_3}(w) &= -0.12 \\ u_{S_3}(w') - q_3 u_{R_3}(w') &= 0.18 \\ u_{S_3}(w'') - q_3 u_{R_3}(w'') &= -0.32. \end{aligned}$$

Thus, we have  $\Sigma_1 = \{w''\}$ ,  $\Sigma_2 = \{w, w'\}$ , and  $\Sigma_3 = \{w'\}$ . Consequently,  $\mathcal{C} = \Sigma_1 \cap \Sigma_2 \cap \Sigma_3 = \emptyset$ , and the society has not reached a compromise that is acceptable to all participants. However, if the father reduces  $q_1$  to 0.9, then

$$\begin{aligned} u_{S_1}(w) - q_1 u_{R_1}(w) &= -0.2 \\ u_{S_1}(w') - q_1 u_{R_1}(w') &= 0.0 \\ u_{S_1}(w'') - q_1 u_{R_1}(w'') &= 0.3. \end{aligned}$$

Hence,  $\Sigma_1 = \{w', w''\}$ , and a consensus exists with  $\mathcal{G}_S = \{w'\}$ . An important feature of this example is that the father need only reduce its standards by a small amount to achieve a consensus. In terms of the narrow self-interest offering given by Table 8, we see, after taking into consideration the social dependencies that exist among the individuals, that the consensus alternative is best for the mother and the child and second best for the father.

## 7 Conclusions

Multi-stakeholder decision theory, and social robotics in particular, is in need of a mathematical framework that is designed to accommodate sophisticated social behaviors such as cooperation, compromise, negotiation, and altruism. The classical framework developed by the social sciences and operations research is based on categorical preference orderings and optimization, and is not sufficiently general to characterize these social behaviors. This research represents a significant departure from classical theory by incorporating three critical notions: conditioning, coherence, and satisficing.

In contrast to categorical utilities, which are designed to characterize self-interest, conditional utilities provide a means whereby individuals may extend their spheres of interest beyond the self. By modulating its preference structure to account for the preferences of others, an individual may account for sophisticated social relationships such as conditional altruism, and thereby give deference to others without categorically redefining its preferences.

In a homogeneous environment where decision makers are required to compromise and negotiate, it is important to ensure that no agent can be categorically subjugated. Coherence is a minimal notion of equity among the participants that can be ensured if and only if the mathematical

syntax of the utilities corresponds to probability mass functions (albeit with different semantics). For societies whose inter-agent influence relationships can be represented by a directed acyclic graph, coherence ensures that the edges are conditional mass functions, resulting in a structure that is mathematically identical to a Bayesian network. This structure permits individual utilities to be aggregated to form a group utility that accounts for social relationships between individuals, thereby providing a complete model of the community.

Satisficing, as defined herein, is an approach to decision making that is as mathematically precise and formalized as is the conventional notion of optimization. The essential advantage of satisficing is that it readily extends to the multi-agent case, whereas optimization is intrinsically a single-agent concept. Furthermore, since satisficing is designed to provide a set of acceptable solutions rather than a unique best solution, it provides a natural mechanism with which to design a negotiation protocol and reach a compromise.

**Acknowledgements** This material is based upon work supported by, or in part by, the U. S. Army Research Laboratory and the U. S. Army Research Office under grant number W911NF-07-1-0650.

## References

1. K. J. Arrow. *Social Choice and Individual Values*. John Wiley, New York, NY, 1951. 2nd ed. 1963.
2. K. J. Arrow. *Social Choice and Individual Values*. John Wiley, New York, NY, 1951. 2nd ed. 1963.
3. K. J. Arrow. Rationality of self and others. In R. M. Hogarth and M. W. Reder, editors, *Rational Choice*, pages 201–215. Univ. of Chicago Press, Chicago, IL, 1986.
4. T. Balch. Hierarchic social entropy: An information theoretic measure of robot group diversity. *Autonomous Robots*, 8(3):209–238, 2004.
5. T. Balch and M. Hybinette. Social potentials for scalable multi-robot formations. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 73–80, 2000.
6. A. Bergson. A reformulation of certain aspects of welfare economics. *Quarterly Journal of Economics*, 52:310–334, 1938.
7. H. Bestougeff and M. Rudnianski. Games of deterrence and satisficing models applied to business process modeling. In *Proceedings of the 1998 AAAI Symposium*, pages 8–14, 1998. March 23–25, Stanford California. Technical Report SS-98-05.
8. C. Camerer, G. Lowenstein, and M. Rabin, editors. *Advances in Behavioral Economics*. Princeton Univ. Press, Princeton, NJ, 2004.
9. R. G. Cowell, A. P. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter. *Probabilistic Networks and Expert Systems*. Springer Verlag, New York, NY, 1999.
10. F. G. Cozman. Credal networks. *Artificial Intelligence*, 120:199–233, 2000.
11. B. de Finetti. La prévision: ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincaré*, 7:1–68, 1937. translated as 'Forsight. Its Logical Laws, Its Subjective Sources', in *Studies in Subjective Probability*, H. E. Kyburg Jr. and H. E. Smokler (eds.), Wiley, New York, NY, 1964, pages 93–158.
12. G. Debreu. *Theory of Value*. Yale University Press, New Haven, CT, 1959.

13. J. Elster, editor. *The Multiple Self*. Cambridge University Press, Cambridge, UK, 1985.
14. R. F. Erlandson. The satisficing process: A new look. *IEEE Trans. Systems, Man, Cybernet.*, 11(11):740–752, 1981.
15. E. Fehr and K. Schmidt. A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114:817–868, 1999.
16. P. C. Fishburn. *The Theory of Social Choice*. Princeton University Press, Princeton, NJ, 1973.
17. M. Friedman. *Price theory*. Aldine Press, Chicago, IL, 1961.
18. J. Harsanyi. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, 63:315, 1955.
19. F. V. Jensen. *Bayesian Networks and Decision Graphs*. Springer Verlag, New York, NY, 2001.
20. B. E. Kaufman. A new theory of satisficing. *Journal of Behavioral Economics*, 19(1):35–51, 1990.
21. R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives*. Cambridge University Press, Cambridge, UK, 1993. First published by John Wiley & Sons, 1976.
22. J. Kemeny. Fair bets and inductive probabilities. *Journal of Symbolic Logic*, 20(1):263–273, 1955.
23. Y. Kim. Satisficing and optimality in 2 $\bar{A}$ -2 common interest games. *Economic Theory*, 13(2):365–375, 1999.
24. C. R. Kube and H. Zhang. Collective robotics: From social insects to robots. *Adaptive Behavior*, 2(2):189–218, 1993.
25. C. R. Kube and H. Zhang. Collective robotic site preparation. *Adaptive Behavior*, 14:5–19, 2006.
26. J. J. Mansbridge, editor. *Beyond Self-Interest*. Univ. of Chicago Press, Chicago, IL, 1990.
27. H. Margolis. Dual utilities and rational choice. In J. J. Mansbridge, editor, *Beyond Self-Interest*, chapter 15, pages 239–253. Univ. of Chicago Press, Chicago, IL, 1990.
28. T. Matsuda. Algebraic properties of satisficing decision criterion. *Information Sciences*, 17:221–237, 1979.
29. T. Matsuda. Characterization of satisficing decision criterion. *Information Sciences*, 17:131–151, 1979.
30. J. F. Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.
31. J. F. Nash. Equilibrium points in  $n$ -person games. In *Proceedings of the National Academy of Sciences*, volume 36, pages 48–49, USA, 1950.
32. F. R. Palmer. *Grammar*. Harmondsworth, Penguin, Harmondsworth, Middlesex, UK, 1971.
33. A. Pazgal. Satisficing leads to cooperation in mutual interest games. *International Journal of Game Theory*, 26(4):439–453, 1997.
34. J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, CA, 1988.
35. G. Polya. *Induction and Analogy in Mathematics*. Princeton Univ. Press, Princeton, NJ, 1954.
36. R. Radner. Satisficing. *Journal of Mathematical Economics*, 2:253–262, 1975.
37. H. Raiffa. *Decision Analysis*. Addison-Wesley, Reading, MA, 1968.
38. F. P. Ramsey. Truth and probability. In R. B. Braithwaite, editor, *The Foundations of Mathematics and Other Logical Essays*. The Humanities Press, New York, NY, 1950.
39. P. A. Samuelson. *Foundations of Economic Analysis*. Harvard Univ. Press, Cambridge, MA, 1948.
40. A. K. Sen. Rational fools: A critique of the behavioral foundations of economic theory. In J. J. Mansbridge, editor, *Beyond Self-Interest*, chapter 2. Univ. of Chicago Press, Chicago, IL, 1990.
41. S. Sen, editor. *Satisficing Models*. AAI Press, San Mateo, CA, 1998.
42. M. Shubik. *Game Theory in the Social Sciences*. MIT Press, Cambridge, MA, 1982.
43. M. S. Shubik. Game theory and operations research: Some musings 50 years later. Yale School of Management Working Paper No. ES-14, May 2001.
44. H. A. Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, 59:99–118, 1955.
45. H. A. Simon. Rational choice and the structure of the environment. *Psychological Review*, 63(2):129–138, 1956.
46. H. A. Simon. Theories of decision-making in economics and behavioral science. *Amer. Econ. Rev.*, XLIX:253–283, 1959.
47. E. Sober and D. S. Wilson. *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Harvard Univ. Press, Cambridge, MA, 1998.
48. I. Steedman and U. Krause. Goethe's *Faust*, Arrow's Possibility Theorem and the individual decision maker. In J. Elster, editor, *The Multiple Self*, chapter 8, pages 197–231. Cambridge Univ. Press, Cambridge, UK, 1985.
49. W. C. Stirling. *Satisficing Games and Decision Making: with applications to engineering and computer science*. Cambridge Univ. Press, Cambridge, UK, 2003.
50. S. Takatsu. Decomposition of satisficing decision problems. *Information Sciences*, 22:139–149, 1980.
51. S. Takatsu. Latent satisficing decision criterion. *Information Sciences*, 25:145–152, 1981.
52. A. Tversky and D. Kahneman. Rational choice and the framing of decisions. In R. M. Hogarth and M. W. Reder, editors, *Rational Choice*, pages 67–94. Univ. Chicago Press, Chicago, IL, 1986.
53. B. B. Werger. Cooperation without deliberation: A minimal behavior-based approach to multi-robot teams. *Artificial Intelligence*, 110(2):293–320, 1999.
54. A. P. Wierzbicki. A mathematical basis for satisficing decision making. In J. N. Morse, editor, *Organizations: Multiple Agents with Multiple Criteria*, pages 465–483. Springer-Verlag, 1981.
55. S. Winter. Satisficing, selection, and the innovating remnant. *Quart. J. Econ.*, 85:237–261, 1971.
56. S. Zilberstein. Satisficing and bounded optimality. In *Proceedings of the 1998 AAI Symposium*, pages 91–94, 1998. March 23–25, Stanford California. Technical Report SS-98-05.

# Attitude Adaptation in Satisficing Games

Matthew Nokleby and Wynn Stirling

**Abstract**—Satisficing game theory offers an alternative to classical game theory that describes a flexible model of players’ social interactions. Players’ utility functions depends on other players’ attitudes rather than simply their actions. However, satisficing players with conflicting attitudes may enact dysfunctional behaviors, resulting in poor performance. We present an evolutionary method by which a population of players may adapt its attitudes to improve payoff. Additionally, we extend the Nash equilibrium concept to satisficing games, showing that the method presented leads players toward an equilibrium in their attitudes. We apply these ideas to the Stag Hunt, a simple game in which cooperation does not easily evolve from non-cooperation. The evolutionary method presented provides two major contributions. First, satisficing players may improve their performance by adapting their attitudes. Second, numerical results demonstrate that cooperation in the Stag Hunt can emerge much more readily under the method presented than under traditional evolutionary models.

## I. INTRODUCTION

Game-theoretic models are often used to construct societies of artificial agents. Commonly, agents are modeled as players in a non-cooperative game in which players focus solely on the maximization of individual payoff. The players’ self-interest leads to Nash equilibria [3], which are strategy profiles such that no single player can improve its payoff by changing strategies. Unfortunately, self-interested behavior places significant limitations in terms of the players’ social interactions. For example, it is often difficult to engender cooperation and other social behaviors with self-interested players. Indeed, the self-interest hypothesis has come under nearly continuous criticism since the inception of game theory [4–7].

Satisficing game theory [8] offers an alternative to non-cooperative game theory. It was developed for the synthesis of artificial agents and specifically focuses on social interactions between players. Players utilities are expressed as conditional mass functions, allowing them to consider the preferences of others rather than focusing solely on individual self-interest. Satisficing models have previously been successful in overcoming the social hurdles presented by non-cooperative game theory, allowing players to exhibit sophisticated social behaviors such as altruism, negotiation, and compromise [9]. However, satisficing theory presents its own set of challenges. As in real-life social situations, satisficing communities may behave dysfunctionally. When players with incompatible attitudes are

grouped together, they can choose incoherent behaviors that lead to poor performance.

The Stag Hunt, a simple game originally suggested by Rousseau [10], underscores the difficulty of achieving cooperation under self-interest. As usually formalized, the game involves two hunters. They can catch a stag only if they hunt stag together, but each can catch a (much smaller) hare separately. That is, a player earns maximum payoff if both players cooperate, but risks failure if it attempts to cooperate while the other does not. Since each player must individually decide between cooperation and non-cooperation, it represents a useful model for the analysis of potentially cooperative behavior. For example, a group of workers choosing whether to strike loosely fall under the Stag Hunt model: a large number of workers may achieve a significant benefit by striking, while a single worker who “strikes” alone incurs significant loss.

Social dilemmas such as the Stag Hunt have been studied extensively by (among others) social scientists, economists, and biologists. A large body of recent work focuses on learning-based [11–13] and evolutionary [14–16] methods for achieving cooperation. In evolutionary game theory, pioneered by Maynard Smith [17, 18], populations of players make decisions by trial-and-error rather than by explicit utility maximization. Over time, natural selection favors individuals who earn higher payoff, altering the population’s makeup. Large, well-mixed populations are described by the replicator dynamics [19], which defines a system of ordinary differential equations governing the evolution of the population. Under suitable conditions, the replicator dynamics drives the population to a Nash equilibrium.

The Stag Hunt presents considerable difficulties from an evolutionary perspective. Under the standard replicator dynamics, a population composed primarily of hare hunters cannot evolve into a group of stag hunters, even though each player benefits from cooperation. Skyrms posits a compelling reason for this failure: “for the Hare Hunters to decide to be Stag Hunters, each must *change her beliefs* about what the others will do. But rational choice based on game theory as usually conceived, has nothing to say about how or why such a change might take place” [20, emphasis in the original].

Motivated by Skyrms’ conjecture, we explore methods by which “such a change” may take place in satisficing game theory. To do so, we attempt to bridge the gap between non-cooperative and satisficing game theory by incorporating elements of non-cooperative game theory into satisficing theory. In a manner similar to [21, 22], we present a method whereby a population of players may modify its attitudes according to the game structure and the attitudes of other players. In our method, which employs the standard replicator dynamics, players whose attitudes result in higher payoffs reproduce more readily, causing their attitudes to dominate the popula-

The authors are with the Electrical and Computer Engineering Department, Brigham Young University, Provo, UT 84602, USA (e-mail: {nokleby, wynn}@ee.byu.edu). Portions of this work were presented at the 2006 IEEE World Congress on Computational Intelligence [1] and the 2007 IEEE Symposium of Foundations of Computational Intelligence [2]. This material is based upon work supported by, or in part by, the U. S. Army Research Laboratory and the U. S. Army Research Office under grant number W911NF-07-1-0650.

tion. The resulting model blends the two decision theories: players retain the conditional utility structure of satisficing theory while improving payoff by evolutionary means. The dynamics leads the players toward a Nash equilibrium in players’ attitudes rather than in their actions.

In Section II we familiarize the reader with the basics of satisficing game theory. In Section III we review the classical formulation of the Stag Hunt and its evolutionary difficulties. We present a satisficing model for the Stag Hunt in Section IV. In Section V we define the attitude equilibrium and present the attitude dynamics. We present experimental results in Section VI and compare the satisficing approach to other recent methods in evolutionary game theory. We give our conclusions in Section VII.

## II. SATISFICING GAME THEORY

While the simple and seemingly reasonable assumption of self-interest—also called individual rationality—has given rise to a rich and successful theory of games, narrow maximization may be *too* simple, particularly in describing social situations. As observed by Luce and Raiffa, “general game theory seems to be in part a sociological theory which does not include any sociological assumptions. . . it may be too much to ask that any sociology be derived from the single assumption of individual rationality” [4, p. 196]. Satisficing game theory provides an alternative to the classical framework. It presents a more elaborate structure which may be more useful in modeling social behaviors. Players may directly concern themselves with the preferences of others, rather than explicitly attempting to maximize utility.

We construct the satisficing framework by altering the structure of the players’ utility functions. First, each player possesses *two* utilities: one to characterize the benefits associated with taking an action and one to characterize the costs. A satisficing player contents itself with a decision for which the benefits outweigh the costs is “good enough” or satisficing.<sup>1</sup> Second, the players’ utility functions share a common syntax with probability mass functions, allowing probabilistic concepts such as conditioning and independence to be applied to players’ preferences—albeit with a significantly different interpretation.

The use of probability mass functions to describe a player’s preferences rather than a random phenomenon is an unusual one, and warrants further explanation. A rigorous justification is given in [24], where it is shown that the use of mass functions as utilities guarantees several useful social properties regarding the reconciliation of group and individual preferences. Fortunately, however, the benefits of conditional utilities may also be appreciated intuitively. For two discrete random phenomena  $X$  and  $Y$ , where  $Y$  is dependent on  $X$ , we can express the probabilities for  $Y$  by the *conditional* mass function  $p_{Y|X}(y|x)$ . The conditional mass function gives hypothetical probabilities of  $Y$ : what would be the probability

<sup>1</sup>Although they share similarities, satisficing game theory should not be confused with the concept of “bounded rationality” satisficing introduced by Simon [23]. With satisficing *à la* Simon, individuals search for sub-optimal choices that meet a variable threshold or *aspiration level*, implicitly accounting for the cost of continued searching.

that  $Y = y$  if we knew that  $X$  took on some value  $x$ ? If we know the probabilities for  $X = x$ , we can compute the marginal mass function according to basic rules of probability theory:  $p_Y(y) = \sum_x p_{Y|X}(y|x)p_X(x)$ . The marginal probabilities for  $Y$  are influenced—but not entirely dictated—by the probabilities of  $X$ .

Similarly, players’ preferences may depend upon the preferences of others, allowing their utilities (which we call *social utilities*) to be expressed as conditional mass functions. The conditional mass functions allow for hypothetical expressions of utility: what would player 1’s utilities be if player 2 unilaterally preferred a particular action? We can compute player 1’s marginal utilities—which are the utilities used for decision-making—by summing the conditional utilities over player 2’s actual preferences. This structure allows players to consider not simply what actions other players may prefer, but *how strong* the preferences for action are. Their utilities are influenced by others’ preferences in a controlled manner which does not require that they discard their own preferences.

### A. Formalization

First, define the set of players  $\mathbf{X} = \{1, 2, \dots, n\}$ . Each player chooses a pure strategy  $u_i \in U_i$ , where  $U_i$  is player  $i$ ’s *pure-strategy set*. A *pure-strategy profile*, which describes the actions of all of the players, is an  $n$ -dimensional vector  $\mathbf{u} \in \mathbf{U}$ , where  $\mathbf{U} = U_1 \times U_2 \times \dots \times U_n$  is the *pure-strategy space*.

As mentioned in the previous subsection, each player possesses two social utilities. To describe these, we define two “selves” or perspectives from which each player may consider its actions [25]. The selecting self considers actions strictly in terms of their associated benefits, while the rejecting self considers actions only in terms of the costs incurred in implementing them. These selves are described by the *selectability function*  $p_{S_i}(u_i)$  and *rejectability function*  $p_{R_i}(u_i)$ , respectively.

Since social utilities are mass functions, they are normalized across the pure-strategy sets and therefore describe the *relative* benefits and costs associated with a pure strategy in  $U_i$ . They also provide players with a formal definition of “good enough.” A pure strategy is “good enough,” or satisficing, if the relative benefits are at least as great as the relative costs. In the vernacular, we may view satisficing as “getting one’s money worth,” as opposed to optimization, where players seek “the best and only the best.” While the former concept allows for a set of multiple actions that are “good enough,” the latter is designed to produce a unique solution. We therefore define the *individually satisficing set* for player  $i$  as

$$\Sigma_i = \{u \in U_i: p_{S_i}(u) \geq qp_{R_i}(u)\}, \quad (1)$$

where  $q$  is the *index of caution*. Typically,  $q = 1$ , but we may adjust a player’s definition of “good enough” by changing  $q$ . Setting  $q \leq 1$  ensures that  $\Sigma_i$  is not empty. We may combine the players’ individually satisficing sets by forming the *satisficing rectangle*  $\mathcal{R}_{12\dots n}$ , which is defined as the Cartesian product

$$\mathcal{R}_{12\dots n} = \Sigma_1 \times \Sigma_2 \times \dots \times \Sigma_n. \quad (2)$$

The satisficing rectangle is the set of all strategy profiles that are simultaneously satisficing to each player.

It is convenient to express the relationship between players' utilities graphically. In probability theory, relationships between random variables are expressed in Bayesian networks [26]. Similarly, in satisficing theory the relationship between players' utilities are expressed in *praxeic networks*.<sup>2</sup> The praxeic network consists of a directed acyclic graph (DAG), where the nodes are the selecting and rejecting perspectives of each player and the edges are the conditional utility functions. For example, consider the simple two-player community depicted in Figure 1. For each player, the rejecting preferences depend on the selecting preferences of the other player, while the selecting preferences are independent.




Fig. 1. A simple praxeic network.

Parenthetically, we note that praxeic networks also resemble the spatial evolutionary models discussed in [15, 16, 28–30]. In these models, graphical connections determine which players can interact during play. That is, individuals may only play with players to whom they are connected. In contrast, graphical connections in praxeic networks define how players influence each other in play. Both models describe, in some sense, players' social relationships. But, while spatial evolutionary models describe which players can pair up in a game, praxeic networks describe which players' utilities can influence the utilities of others.

In discussing the players' social utilities, we retain the terminology of probability theory. In the community from Figure 1, we refer to player 1's *conditional* rejectability function, denoted  $p_{R_1|S_2}(v_1|u_2)$ . As mentioned above, the conditional mass function expresses a hypothetical proposition, where the antecedent is the strategy favored by player 2, and the consequent is the utility of player 1. That is, if player 2's selecting preferences entirely favored strategy  $u_2$ , what would be player 1's rejectability for  $v_1$ ? As with probability mass functions, we may compute the *marginal* rejectability by summing over the conditionals:  $p_{R_1}(v_1) = \sum_{u_2 \in U_2} p_{R_1|S_2}(v_1|u_2)p_{S_2}(u_2)$ . The marginal utilities determine the individually satisficing sets and the satisficing rectangle. If a utility is independent (such as the selectability functions in this example), its marginal is expressed directly, without conditioning.

By allowing conditioning in the players' utilities, we implicitly assume that players have at least partial knowledge of each other's utilities. Each player must have sufficient knowledge of other players' utilities in order to compute its marginal utilities and find its individually satisficing set. In the example community, each player must know the other player's

selectability function in order to compute its own rejectability. However, since players do not consider each other's actions in determining the individually satisficing sets, they need not observe (or predict) each other's choices.

With the marginal and conditional utilities defined for the example community, we can form the *interdependence function*  $p_{S_1 \dots S_n R_1 \dots R_n}(u_1, \dots, u_n, v_1, \dots, v_n)$ , which is the joint mass function of all players' selecting and rejecting preferences. By the chain rule of probability theory, the interdependence function for this example is  $p_{S_1 S_2 R_1 R_2}(u_1, u_2, v_1, v_2) = p_{R_1|S_2}(v_1|u_2)p_{R_2|S_1}(v_2|u_1)p_{S_1}(u_1)p_{S_2}(u_2)$ .

Satisficing games are characterized by the triple  $(\mathbf{X}, \mathbf{U}, p_{S_1 \dots S_n R_1 \dots R_n})$ , where  $\mathbf{X}$  is the set of players,  $\mathbf{U}$  is the pure-strategy space, and  $p_{S_1 \dots S_n R_1 \dots R_n}$  is the interdependence function. From this information, all necessary marginal utilities can be computed and the satisficing rectangle can be determined.

Finally, it is often useful to specify the players' social utilities in terms of variable parameters, which we refer to as the players' *attitudes*. The interpretation of the attitudes, of course, depends on the specific game being played, but in general they express each player's temperament, which affects the degree to which its utilities depend on those of other players. For example, in the Stag Hunt, the players' attitudes will characterize their aversion to risk, which influences each player's willingness to engage in stag-hunting.

### B. Random Satisficing Games

Often, a player's utility will depend on random phenomena, resulting in expected utilities based on the distribution of the random event. With classical game theory, it is required that the probabilistic distributions of the random phenomena not be influenced by the preferences of the players. In other words, a player's belief regarding a random event may affect its utilities, but not vice versa. In most cases this restriction poses no difficulty. However, we may want to consider circumstances in which a player's subjective probability about an event depends on players' preferences.

The conditional structure of social utilities provides for such a possibility. Since the utilities are mass functions, we can combine both probabilistic and preferential information into a single model. Figure 2 illustrates a network implementing such a model. This praxeic network is similar to Figure 1 in that it contains the same four vertices associated with the players' selecting and rejecting selves. However, we also include two random variables  $\theta_1$  and  $\theta_2$ , which represent phenomena that are known to the players only probabilistically. This network describes both players whose preferences depend on random phenomena *and* random phenomena which depend on players' preferences. The dependencies from Figure 1 still persist.  $R_1$  still depends—albeit indirectly, through  $\theta_2$ —on  $S_2$ , and  $R_2$  still depends on  $S_1$ , which now depends  $\theta_1$ .

## III. THE STAG HUNT

In the Stag Hunt, players choose between two pure strategies: hunt stag or hunt hare, denoted  $s$  and  $h$ , respectively. The payoff for playing each pure strategy depends on the

<sup>2</sup>The term *praxeic* is derived from *praxeology*, which refers to “the science of human conduct” or “the science of efficient action.” [27]




Fig. 2. An praxeic network with “true” random variables.

action of the other player. If the other player hunts stag, the payoff for hunting stag is higher than that of hunting hare. However, if the other player hunts hare, stag hunting yields a low payoff. That is, the players must hunt together to catch the stag and obtain the higher payoff. The payoff for hunting hare, on the other hand, is independent of the other player’s choice. Each player can individually catch a hare, and therefore can always opt for the modest—but more secure—payoff associated with consuming a hare. We quantitatively express the players’ utilities in the payoff matrix of Table I.

TABLE I  
PAYOFF MATRIX FOR A TWO-PLAYER STAG HUNT.

Player 1	Player 2	
	$s$	$h$
$s$	(4, 4)	(0, 3)
$h$	(3, 0)	(3, 3)

There are two pure-strategy Nash equilibria for the Stag Hunt:  $(s, s)$  and  $(h, h)$ . If the players simultaneously hunt stag or hare, there is no incentive for either player to change actions. There is also a *mixed-strategy* equilibrium, in which each player invokes a randomized rule to choose between the two pure strategies. We will study the mixed-strategy equilibrium in more detail later. Each pure-strategy equilibrium has its benefits. The  $(s, s)$  equilibrium is optimal in that it maximizes both players’ payoffs. However, since successful stag hunting requires the cooperation of the other player, risk-averse players may choose instead to hunt hare. The  $(h, h)$  equilibrium is regarded as the *risk-dominant* equilibrium in the sense that the potential gains of deviating from hare hunting are less than the potential losses: at best, a hare hunter will increase its utility by one by switching to stag, but at worst, it will decrease its utility by three. Thus, conservative—yet fully rational—players might choose to hunt hare.

This dichotomy illustrates the fundamental issue of the Stag Hunt. Obviously, if each player had certain assurance that the other player would hunt stag, everyone would cooperate.<sup>3</sup> However, players do not have such an assurance under the usual model, but must choose their actions independently. The players’ actions then boil down to how much confidence each

<sup>3</sup>Interestingly, it is straightforward to show that if the game is played sequentially (i.e. player 1 makes its move, and then player 2—who observes player 1’s choice—moves), mutual stag-hunting becomes the unique *subgame perfect* Nash equilibrium. [31]

player has in the other’s willingness to cooperate and how risk-averse each player is. As mentioned by Skyrms, classical game theory has little to say about this topic. Indeed, the Nash equilibria do not tell us which actions the players will take. They simply imply that once a pair of players is in either of the pure-strategy equilibria, neither player will have incentive to deviate. To study which equilibrium will result under different circumstances, we turn to evolutionary game theory [32, 33].

### A. The Replicator Dynamics

The replicator dynamics is the classic instantiation of evolutionary game theory. It models the evolution of a population’s strategies according to their ecological fitness. Consider a large population of players who are “programmed” to play a particular strategy—regardless of the other player’s behavior—in a symmetric two-player game such as the Stag Hunt. The players are randomly paired up to play the game at each time step. Each player reproduces asexually<sup>4</sup> according to its payoffs; that is, the number of offspring that a player has is proportional to its payoff during the previous game. Players’ strategies also “breed true,” meaning that offspring are programmed to the same pure strategy as their parents. We assume that the population is well-mixed, giving each player an equal chance of being paired with any other player.

For a symmetric, two-player game where each player must choose some strategy in the pure-strategy set  $U$ , define the *mixed-strategy simplex*  $\Delta_U$  as the set of all mixed (randomized) strategies over  $U$ . If  $U$  contains  $m$  elements, we can characterize a mixed strategy as a nonnegative  $n$ -dimensional vector  $\mathbf{x}$  that obeys the constraint  $\sum_{i=1}^m x_i = 1$ . Each player’s mixed strategy is probabilistically independent of the other player’s. The interior of  $\Delta_U$  is the set of mixed strategies which assign nonzero probability to each pure strategy:

$$\text{int}(\Delta_U) = \{\mathbf{x} \in \Delta_U : x_i > 0, i \in \{1 \dots m\}\}.$$

In the replicator dynamics, we interpret each element  $x_i$  as the *population share*, or fraction of the population, playing pure strategy  $i$ . That is, if we randomly draw an individual from the population described by  $\mathbf{x}$ , the probability that it will be programmed to play  $i$  is  $x_i$ . At time  $t$ , the expected utility<sup>5</sup> of a player who plays pure strategy  $i$  against a random member of the population is  $u(i, \mathbf{x}(t)) = \sum_{j=1}^m \pi(i, j)x_j(t)$ , where  $\pi(i, j)$  represents the utility of playing pure strategy  $i$  against pure strategy  $j$ . As the players reproduce, the population shares described by  $\mathbf{x}(t)$  vary, and the more successful strategies tend to dominate over those which are poorly-adapted to the evolving community. As the population size approaches infinity we may invoke the law of large numbers, and the dynamics of the population shares becomes a system of  $m$  differential equations:

$$\dot{x}_i(t) = [u(i, \mathbf{x}(t)) - u(\mathbf{x}(t), \mathbf{x}(t))]x_i(t), i \in \{1 \dots m\}, \quad (3)$$

<sup>4</sup>This does not contradict the fact that the players must pair off to play the game. While they do play the game pairwise, each player earns its payoff individually. The number of offspring it produces is proportional only to its own payoff, and is entirely independent of the other player’s.

<sup>5</sup>We use  $\pi$  to represent the utility (or payoff) for when players use only pure strategies, while  $u$  represents the expected utility when mixed strategies are involved.

where  $u(\mathbf{x}(t), \mathbf{x}(t))$  is the population's average expected utility,

$$\begin{aligned} u(\mathbf{x}(t), \mathbf{x}(t)) &= \sum_{i=1}^m u(i, \mathbf{x}(t))x_i(t) \\ &= \sum_{i=1}^m \sum_{j=1}^m \pi(i, j)x_i(t)x_j(t). \end{aligned}$$

Intuitively, (3) tells us that a pure strategy's population share increases at time  $t$  if its expected utility is higher than the average expected utility across the population. It is shown in [32] that, if the initial conditions satisfy  $\mathbf{x}(0) \in \text{int}(\Delta_U)$  (all pure strategies are represented in the initial conditions), any steady state of the dynamics is a Nash equilibrium in the players' strategies.

It should be noted that the standard replicator model describes a selection dynamics rather than a mutation dynamics. Players do not change strategies under this model; instead, the offspring of players whose strategies are suboptimal are overwhelmed by the offspring of more successful players. As time continues, the fraction of the population playing suboptimal strategies becomes arbitrarily small.

To account for random factors such as mutation, migration, and payoff fluctuations, several stochastic replicator models have been proposed [13, 14, 34, 35]. We examine a model from [14], which augments the standard replicator dynamics by introducing fixed mutation probabilities into the dynamics. The mutation probabilities are contained in the matrix  $\mathbf{W} = [W_{ij}]$ , where  $W_{ij}$  represents the probability that an individual playing strategy  $j$  spontaneously switches to strategy  $i$ . The mutation dynamics differs from (3) by the addition of a mutation term:

$$\begin{aligned} \dot{x}_i(t) &= [u(i, \mathbf{x}(t)) - u(\mathbf{x}(t), \mathbf{x}(t))]x_i(t) \\ &\quad + \sum_{j=1}^m (W_{ij}x_j(t) - W_{ji}x_i(t)). \end{aligned} \quad (4)$$

The dynamics for  $x_i$  are altered by adding the rate at which players mutate into the population share  $x_i$  (described by  $\sum_j W_{ij}x_j$ ) and subtracting the rate at which players mutate out of the population share  $x_i$  (described by  $\sum_j W_{ji}x_j$ ). When mutation probabilities are zero ( $\mathbf{W} = \mathbf{I}$ ), (4) collapses to the standard replicator dynamics. In general, however, we are forced to give up the theoretical properties guaranteed under the standard replicator model. The steady-state behavior of the system no longer corresponds to Nash equilibria, regardless of initial conditions.

## B. Stag Hunt Replicator Dynamics

1) *Standard Dynamics*: For the Stag Hunt, the population is described by the two-dimensional vector  $\mathbf{x} = (x_s, x_h)$ . The payoff matrix (Table I) shows that the payoff for a stag hunter is four when paired with another stag hunter, and zero when paired with a hare hunter. A stag hunter therefore gains an expected utility of  $u(s, \mathbf{x}) = 4x_s$ . Since the utility for hunting hare is independent of the other player's actions,  $u(h, \mathbf{x}) = 3$ . The population's average expected payoff is given by  $u(\mathbf{x}, \mathbf{x}) = 4x_s^2 - 3x_s + 3$ . Since  $x_s = 1 - x_h$ , we can

characterize the dynamics by examining only the stag hunting share. Suppressing time arguments, we get

$$\dot{x}_s = [u(s, \mathbf{x}) - u(\mathbf{x}, \mathbf{x})]x_s = -4x_s^3 + 7x_s^2 - 3x_s. \quad (5)$$

While the nonlinearities prevent a closed-form solution, we can easily examine the qualitative behavior of the population. In Figure 3, we show a direction field for the replicator dynamics, which gives the sign of the derivative as a function of  $x_s$ . The stationary points, where  $\dot{x}_s = 0$ , occur at  $x_s = \{0, 3/4, 1\}$ . The point at  $x_s = 3/4$  corresponds to the mixed-strategy Nash equilibrium discussed previously. However, the mixed-strategy equilibrium is not stable; any deviation drives the dynamics to one of the pure-strategy points, which are asymptotically stable. We may regard  $x_s = 3/4$  as a boundary for the initial conditions of the population: if fewer than 75% of the population initially hunt stag, the dynamics quickly drives stag hunters to relative extinction. If more than 75% initially hunt stag, hare hunters die out. Although stag hunting prevails in a predominantly cooperative society, these dynamics cannot evolve cooperation from an initially non-cooperative population.




Fig. 3. Direction field for Stag Hunt replicator dynamics.

2) *Mutation Dynamics*: Using the replicator model in (4), we add a probability of mutation into the Stag Hunt dynamics in the hope that mutation may help evolve a cooperative population. We assume that the probability of mutating from stag to hare is identical to the probability of mutation from hare to stag. Consequently, we can parameterize the mutation matrix by a single mutation probability  $0 \leq \alpha \leq 1$ :

$$\mathbf{W} = \begin{bmatrix} 1 - \alpha & \alpha \\ \alpha & 1 - \alpha \end{bmatrix}.$$

The dynamics for  $x_s$  becomes

$$\begin{aligned} \dot{x}_s &= -4x_s^3 + 7x_s^2 - 3x_s + W_{sh}(1 - x_s) - W_{hs}x_s \\ &= -4x_s^3 + 7x_s^2 - 3x_s + \alpha(1 - 2x_s). \end{aligned} \quad (6)$$

The closed-form expression for the stationary points of the dynamics is quite unwieldy, so in Figure 4 we plot the direction field for the dynamics as a function of  $\alpha$  and  $x_s$ . When mutation probabilities are small, the qualitative behavior of the solution does not change: there remain two stable stationary points at which nearly all of the population hunts either stag or hare and an unstable stationary point which defines the boundary between the stag-hunting and hare-hunting basins of attraction. The boundary point increases with the mutation rate, suggesting that mutation exacerbates the evolutionary difficulties of the Stag Hunt.

For large mutation probabilities, the dynamics differs considerably, leaving a single stationary point to which the dynamics converges independent of initial conditions. Even with absurdly high mutation rates—in which evolution is governed more by mutation than by payoff—only a minority of the

population hunts stag. Of course, since the population size is infinite, the mutation replicator model defines a deterministic system as in the standard dynamics. Resultantly, finite populations, with random pairings and mutation, may spontaneously evolve cooperation from non-cooperation. But the moral of the story is that, on average, even finite populations rarely cooperate if they are large, well-mixed, and composed of players that are pre-programmed to play a particular pure-strategy.




Fig. 4. Direction field for Stag Hunt stochastic replicator dynamics.

Finally, as we have already discussed, there exist other evolutionary models than the replicator dynamics. In Section VI, we investigate the effects of more sophisticated evolutionary mechanisms on the Stag Hunt. For the time being, however, we focus on the underlying structure of the players' behavior. Our solution, based upon satisficing game theory, affords a flexible structure for players' social interactions, increasing the possibility for cooperation even under simple evolutionary dynamics.

#### IV. THE SATISFICING STAG HUNT

In a two-player Stag Hunt, the set of players is  $\mathbf{X} = \{1, 2\}$ , and each player has an identical pure-strategy set  $U_i = \{s, h\}$ ,  $i \in \mathbf{X}$ . In formulating a satisficing game, we are free to select an arbitrary structure for the praxeic network and specify the conditional utilities as we see fit. We are then constrained to carry out the rules of probability in computing the marginal utilities which determine the players' behavior. Thus, the formulation of a satisficing game is a process of "designing" the conditional structure and examining the results to see if the players' behavior makes sense.

First, we give conceptual definitions for the selectability and rejectability preferences, which we will further clarify as we mathematically define the players' social utilities. What do we mean by "benefits" and "costs" for the players in the Stag Hunt? In our treatment, we consider selectability in terms of successful cooperation. To the extent to which stag hunting can be successful, the selecting self prefers to hunt stag. We associate rejectability with the raw opportunity

cost of an action, tempered by risk-aversion. The opportunity cost of hunting hare is the payoff for catching a stag, and the opportunity cost of hunting stag is the payoff for catching a hare.

Next, we define the interconnections between the four selves and form the praxeic network. Our model is illustrated in Figure 5. In addition to the vertices corresponding to the selecting and rejecting selves, we include a vertex which corresponds to a binary random variable  $\theta_s$  which accounts for the possibility of failure. It is not necessarily certain, even if both players hunt stag, that they will succeed. We use  $\theta_s = 1$  to denote that a successful stag hunt is possible and  $\theta_s = 0$  to denote that stag hunting will result in failure.




Fig. 5. The praxeic network for the Stag Hunt.

To define the rejectability function for each agent, we first must define a normalized measure of opportunity cost. Let  $\phi_{s_i}$  and  $\phi_{h_i}$  denote the raw utility (in arbitrary units) of consuming stag and hare, respectively. Normalizing, the relative utility of hare-hunting becomes  $\mu'_i = \frac{\phi_{h_i}}{\phi_{h_i} + \phi_{s_i}}$  for  $i = 1, 2$ . The relative utility of stag-hunting is then  $1 - \mu'_i$ .

Given this definition, we may let  $\phi_{s_i} = 4$  and  $\phi_{h_i} = 3$ , the payoff values given in Table I, resulting in  $\mu'_i = \frac{3}{7}$ . However, we further wish to take into account the temperament of the players. As discussed in Section III, a central issue in the Stag Hunt is to determine what players of differing risk-aversion levels should do. Therefore, we introduce a parameter,  $\rho_i$ , which expresses the degree of player  $i$ 's risk aversion. A player with  $\rho_i = 1$  is risk-neutral, a player with  $\rho_i > 1$  is risk-averse, and a player with  $\rho_i < 1$  is payoff-seeking and tends to ignore risk. We then define  $\mu_i = \rho_i \frac{\phi_{h_i}}{\phi_{h_i} + \phi_{s_i}}$ . Thus,  $\mu_i$  reflects a player's willingness to take risks as well as the relative utility for stag and hare. A maximally risk-averse player will hunt stag only if success is certain, while a fully payoff-seeking player will hunt stag regardless of the odds. To ensure a meaningful game, we still require that both players will never prefer hare to stag, or  $\mu_i < \frac{1}{2}$  for  $i = 1, 2$ . For convenience, we will simply refer to  $\mu_i$  as player  $i$ 's *risk-aversion level*, which parameterizes the player's attitudes.

We define each player's rejectability function as

$$p_{R_i}(u_i) = \begin{cases} \mu_i, & \text{for } u_i = s \\ 1 - \mu_i, & \text{for } u_i = h \end{cases}, \quad (8)$$

an expression of normalized opportunity cost for each action. The cost of hunting stag is the relative hare hunting utility, and vice versa. Note that the players' rejecting selves are not

dependent on others' preferences, allowing us to define the marginal utilities directly.

We next define the conditional distribution for  $\theta_s$ . The distribution of this random variable, which is conditioned upon both players' rejecting selves, represents the probability that the players will successfully hunt stag. The distribution of  $\theta_s$  incorporates whether or not  $R_1$  and  $R_2$  reject cooperation *and* how likely the players are to catch a stag if they cooperate. We model the latter consideration by defining  $0 \leq \sigma \leq 1$ , which represents the probability of catching a stag given that the players cooperate. It may reflect the number of stag in the environment, the players' hunting skills, or other external factors. If  $R_1$  and  $R_2$  altogether reject hare hunting, then the players will cooperate and successfully capture a stag with probability  $\sigma$ . We characterize this by defining

$$p_{\theta_s|R_1 R_2}(\vartheta_s|h, h) = \begin{cases} \sigma, & \text{for } \vartheta_s = 1 \\ 1 - \sigma, & \text{for } \vartheta_s = 0 \end{cases}, \quad (9)$$

where  $\theta_s$  represents the random variable and  $\vartheta_s$  represents its realization. If, however, either player unilaterally rejects stag hunting, the probability of catching a stag is zero, yielding

$$p_{\theta_s|R_1 R_2}(\vartheta_s|s, s) = p_{\theta_s|R_1 R_2}(\vartheta_s|s, h) \quad (10)$$

$$= p_{\theta_s|R_1 R_2}(\vartheta_s|h, s) \quad (11)$$

$$= \begin{cases} 0, & \text{for } \vartheta_s = 1 \\ 1, & \text{for } \vartheta_s = 0 \end{cases}. \quad (12)$$

Notice that the players' preferences influence the probability of a random event as discussed in Section II-B. Since the players' rejecting preferences affect their willingness to hunt stag, the conditional structure is justifiable.

We compute the marginal mass function by summing over the conditional random variables, yielding

$$p_{\theta_s}(\vartheta_s) = \sum_{v_1, v_2} p_{\theta_s|R_1, R_2}(\vartheta_s|v_1, v_2) p_{R_1}(v_1) p_{R_2}(v_2) \quad (13)$$

$$= \begin{cases} \sigma(1 - \mu_1)(1 - \mu_2), & \text{for } \vartheta_s = 1 \\ 1 - \sigma(1 - \mu_1)(1 - \mu_2), & \text{for } \vartheta_s = 0 \end{cases}. \quad (14)$$

From (14) we see that as the risk-aversion levels decrease, the probability of a successful stag hunt increases. If both players are completely payoff-seeking ( $\mu_1 = \mu_2 = 0$ ), the probability of a successful stag hunt is  $\sigma$ . Either player can reduce the chances for a successful hunt. As the risk-aversion  $\mu_i$  increases for either player, the probability of a successful stag hunt decreases.

Finally, we define the conditional selectability. Each player's selectability is influenced by the probability of a successful stag hunt. The selectability, as discussed earlier, is tied to the benefits of cooperation: to the extent that a successful stag hunt is possible ( $\theta = 1$ ), selectability favors stag hunting. The higher the probability of successful stag hunting, the more beneficial it is to hunt stag. The corresponding conditional

selectability function is

$$p_{S_i|\theta_s}(u_i|\vartheta_s) = \begin{cases} 1 & \text{for } u_i = s|\vartheta_s = 1 \\ 0 & \text{for } u_i = h|\vartheta_s = 1 \\ 0 & \text{for } u_i = s|\vartheta_s = 0 \\ 1 & \text{for } u_i = h|\vartheta_s = 0 \end{cases}. \quad (15)$$

The simple form of the conditionals allows us to express the marginal selectability as

$$p_{S_i}(u_i) = \begin{cases} \sigma(1 - \mu_1)(1 - \mu_2) & \text{for } u_i = s \\ 1 - \sigma(1 - \mu_1)(1 - \mu_2) & \text{for } u_i = h \end{cases}. \quad (16)$$

#### A. The Satisficing Rectangle

With all of the social utilities defined, we have completely characterized the players' utilities and can solve for the pure-strategy profiles that form the satisficing rectangle. As discussed in Section II, the satisficing rectangle is the set of pure-strategy profiles that are simultaneously satisficing to each player individually. In Figure 6, we set  $q = 1$  and plot the regions of the satisficing rectangle as functions of  $\mu_1$  and  $\mu_2$ , which specify the players' attitudes. There are four possibilities. When both players have low risk-aversion,  $(s, s)$  is the unique strategy profile in the satisficing rectangle. If risk-aversion is high in both players,  $(h, h)$  results. In the  $(h, s)$  and  $(s, h)$  regions, however, one player is strongly risk-averse while the other strongly seeks payoff, resulting in one player that tries to cooperate while the other does not. On the boundaries of the four regions, the satisficing rectangle contains multiple strategy profiles.




Fig. 6. Satisficing rectangle regions for the Stag Hunt.

These last two regions illustrate a unique feature of satisficing models. In the  $(h, s)$  and  $(s, h)$  regions, one player chooses to hunt hare while the other player, who is aware of the first player's increased risk-aversion, nevertheless stands by its post and attempts to hunt stag. Such dysfunctional behavior is a

consequence of the structure of the utilities: players' utilities depend on the others' attitudes rather than the strategies they play.

We hasten to note that dysfunctional behavior is not a failure *per se* of the satisficing model. Dysfunctional societies do exist in practice, and we may interpret these regions as an acknowledgement that players with incompatible attitudes may act incoherently. However, in designing artificial systems, we typically prefer to avoid incoherent behaviors, sociologically justifiable or not. It seems unreasonable that incompatible players would continue to exhibit the same attitudes and to enact the same incoherent strategies. Thus, we introduce the attitude dynamics, which provides a way for players to adapt their attitudes and avoid such dysfunctional behavior.

## V. ATTITUDE DYNAMICS

To introduce the attitude equilibrium and the attitude dynamics, we first embellish the structure of the satisficing game. We endow each player with a classical utility function which is based solely on the strategy profile that the players implement.

*Definition 1:* An *augmented satisficing game* is a 5-tuple  $(\mathbf{X}, \mathbf{U}, p_{S_1 \dots S_n R_1 \dots R_n}, \mathbf{A}, \boldsymbol{\pi}(\mathbf{u}))$ . The first three elements are the set of players, the pure-strategy space, and interdependence function as normal. Additionally, we introduce the pure-attitude space  $\mathbf{A} = A_1 \times A_2 \times \dots \times A_n$  containing the attitudes that the players may exhibit. These attitudes are parameters in the players' social utilities, and are different for each satisficing game. We also introduce  $\boldsymbol{\pi}(\mathbf{u})$ , a vector payoff function which describes the raw payoff to the players for implementing the pure-strategy profile  $\mathbf{u} \in \mathbf{U}$ .

To augment a satisficing game, the players' attitudes must be specified as distinct parameters in the players' social utilities. Further, we must be able to construct a raw payoff function that is separate from the social utilities. Constructing raw payoff functions may be difficult in practice. In a system of artificial agents, for example, the agents' objectives may be sufficiently complicated that it is impossible to define a simple payoff function for each agent. In a simple game like the Stag Hunt, the extension is straightforward. Each player's attitudes are given by the risk-aversion level  $\mu_i$ , yielding a pure-attitude space of  $\mathbf{A} = [0, 1/2) \times [0, 1/2)$ . The payoff function  $\boldsymbol{\pi}(\mathbf{u})$  is described by the payoff matrix in Table I.

The augmented satisficing game describes a two-step mapping from attitudes to payoffs. The social utilities—determined by the interdependence function—map the players' attitudes to pure-strategy profiles.<sup>6</sup> The payoff function then maps the pure-strategy profile to raw payoffs. Thus, in an augmented satisficing game, we may evaluate the raw utility of exhibiting a particular attitude. To simplify notation, we will occasionally refer to  $\boldsymbol{\pi}(\mathbf{a})$ , the payoff to the players for implementing the pure-strategy profile determined by the pure-attitude profile  $\mathbf{a} \in \mathbf{A}$ . That is, we may think of an augmented satisficing game as a non-cooperative game where players' payoffs are

determined by the attitudes they exhibit rather than the strategies they play.

We may also discuss *mixed attitudes* which are probability distributions over the attitudes the players exhibit. Denoting the cardinality of  $U_i$  as  $k_i$ , the mixed attitude of player  $i$  is given by a (normalized and nonnegative)  $k_i$ -dimensional vector  $\mathbf{z}_i$ . The discussion of mixed strategies in Section III-A applies directly to mixed attitudes. We assume that players' mixed attitudes are probabilistically independent of each other. We define player  $i$ 's mixed attitude simplex  $\Delta_i^a$ . The mixed-attitude space is the Cartesian product  $\Theta^a = \Delta_1^a \times \Delta_2^a \times \dots \times \Delta_n^a$ . A mixed-attitude profile is a vector of mixed attitudes  $\mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_n) \in \Theta^a$ .

Since the players' mixed attitudes are independent, the probability that a pure-attitude profile is exhibited is equal to the product of the associated probabilities. Thus, player  $i$ 's expected utility  $u_i(\mathbf{z})$  when the players exhibit the mixed-attitude profile  $\mathbf{z} \in \Theta^a$  is:

$$u_i(\mathbf{z}) = \sum_{\mathbf{a} \in \mathbf{A}} \pi_i(\mathbf{a}) \prod_{i=1}^n z_{ia_i}, \quad (17)$$

where  $z_{ia_i}$  is the probability with which player  $i$  exhibits the pure-attitude  $a_i$ . Now, given complete knowledge of the satisficing game and the other players' utilities, a player may consider *changing* their attitudes to increase expected utility, which motivates the attitude equilibrium.

*Definition 2:* An *attitude equilibrium* is a mixed-attitude profile  $\mathbf{z}^* \in \Theta^a$  such that

$$u_i(\mathbf{z}_1^*, \dots, \mathbf{z}_i^*, \dots, \mathbf{z}_n^*) \geq u_i(\mathbf{z}_1^*, \dots, \mathbf{z}'_i, \dots, \mathbf{z}_n^*) \quad (18)$$

for each  $\mathbf{z}'_i \in \Delta_i^a$  and for each  $i \in \mathbf{X}$ .

The definition for the attitude equilibrium is essentially identical to that of the Nash equilibrium: no player can improve its expected utility by exhibiting a different mixed attitude. In fact, we may say that an attitude equilibrium is an equilibrium in players' attitudes, rather than in their strategies. Because of the analogy between the attitude equilibrium and the Nash equilibrium, many theoretical results apply.

*Theorem 1:* An attitude equilibrium exists for every augmented satisficing game with finite attitude spaces.

*Proof:* This result relies upon the fact that any augmented satisficing game defines a classical non-cooperative game where  $\mathbf{X}$  is the set of players,  $\mathbf{A}$  takes the role of the pure-strategy space and  $\boldsymbol{\pi}(\mathbf{a})$  is the payoff function. In [3], it is shown that any non-cooperative game with a finite pure-strategy space has at least one Nash equilibrium, although it may exist only in mixed strategies. Since an attitude equilibrium is simply a Nash equilibrium in the players' attitudes, one must exist for any augmented satisficing game with a finite pure-attitude space, even if it exists only in mixed attitudes. ■

Note that a finite attitude space is a sufficient, but not necessary, condition for the existence of an attitude equilibrium. Indeed, for the Stag Hunt, even though the attitude spaces are continuous, it is immediate that attitude equilibria exist in pure attitudes. In Figure 7, the attitude equilibria are shown for several values of  $\sigma$ . If the players' pure-attitude profile

<sup>6</sup>We have glossed over the fact that, in general, the satisficing rectangle contains multiple pure-strategy profiles. For the Stag Hunt, this presents no problem because the satisficing rectangle contains a single strategy profile almost everywhere. We will assume that, if necessary, the players employ a tie-breaking mechanism to select a unique strategy profile.

lies in these regions, there is no incentive for either player to change attitudes.




Fig. 7. Attitude equilibrium regions for the Stag Hunt.

Consider the  $(s, s)$  region of the satisficing rectangle. Here, both players receive maximum payoff and there is no incentive for either player to deviate. Notice, however, that only part of the  $(h, h)$  region is an equilibrium. This is because, when player  $i$ 's risk-aversion  $\mu_i$  is sufficiently low, it is possible for player  $j$  to move the group from mutual hare-hunting to stag-hunting by lowering its own  $\mu_j$ . Even though  $(h, h)$  is an equilibrium under the classical game, the satisficing model gives the players greater influence over each other's behavior, increasing the possibility for cooperation. As  $\sigma$  increases, the size of the  $(h, h)$  equilibrium decreases, disappearing entirely when  $\sigma = 1$ .

Finally, notice that the dysfunctional regions  $(s, h)$  and  $(h, s)$  do not contain equilibria. In these regions, each player can improve its payoff by changing  $\mu_i$  and forcing the game into either  $(s, s)$  or  $(h, h)$ . The attitude equilibrium concept provides a useful juxtaposition of satisficing theory and individual rationality: the social structure of the satisficing model decreases the attraction of mutual hare-hunting, while the introduction of the classical payoff function gives incentive for players to adapt their attitudes and avoid dysfunctional behaviors of the  $(s, h)$  and  $(h, s)$  regions.

If a large population of players adapts by trial-and-error experimentation, we can model the evolution of the players' attitudes by a straightforward application of the standard replicator dynamics. We again restrict our attention to symmetric, two-player games. Thus, both players are described by the pure-attitude set  $A$  and the payoff function  $\pi(\mathbf{a})$ . We require that  $A$  be finite, and we denote the cardinality of  $A$  as  $m$ . Define a normalized vector  $\mathbf{z}(t) = (z_1(t), z_2(t), \dots, z_m(t))$ , where  $z_i(t)$  represents the population share exhibiting the  $i$ th pure attitude. Just as with traditional games, we may describe the dynamics of the population shares by a system of  $m$

differential equations:

$$\dot{z}_i(t) = [\pi(i, \mathbf{z}(t)) - \pi(\mathbf{z}(t), \mathbf{z}(t))]z_i(t). \quad (19)$$

By analogy with the standard formulation,  $\pi(i, \mathbf{z}(t))$  is the expected payoff for exhibiting the  $i$ th attitude against a random sample from the population and  $\pi(\mathbf{z}(t), \mathbf{z}(t)) = \sum_i \sum_j \pi(i, j)z_i(t)z_j(t)$  is the average expected payoff.

Let  $\Delta_A$  be the mixed-attitude simplex of  $A$ . Just as with mixed strategies, the interior of  $\Delta_A$  is the set of all mixed attitudes which gives nonzero probability to each pure attitude.

*Theorem 2:* Let  $\xi(t, \mathbf{z}(0))$  denote the solution for the attitude dynamics in (19) at time  $t$  with initial conditions  $\mathbf{z}(0)$ . If  $\mathbf{z}(0) \in \text{int}(\Delta_A)$  and  $\lim_{t \rightarrow \infty} \xi(t, \mathbf{z}(0)) = \mathbf{z}^*$ , then  $\mathbf{z}^*$  is an attitude equilibrium.

*Proof:* This result follows directly from the fact that an augmented satisficing game can be thought of as a classical game where players choose attitudes rather than play strategies. As mentioned in Section III-A, it is shown in [32] that, when initialized with a mixed strategy on the interior of the mixed-strategy simplex, any steady state of the replicator dynamics is a Nash equilibrium. Since an attitude equilibrium is a Nash equilibrium in players' attitudes, the result holds for the attitude dynamics. ■

Note that Theorem 2 does not guarantee that a steady-state will occur, even under well-behaved initial conditions. Rather, if a steady-state results under suitable initial conditions, it must be an attitude equilibrium.

## VI. RESULTS

### A. Attitude Dynamics

To apply the attitude dynamics, we first quantize the values that  $\mu$  may assume. Define  $A = \{\nu_1, \nu_2, \dots, \nu_{100}\}$ , a set of 100 evenly-spaced values of  $\mu$  over  $[0, 1/2)$ . We initialize the population shares  $\mathbf{z}$  according to an exponential distribution so that most players hunt hare, or  $z_i(0) \propto e^{-\lambda(\frac{1}{2} - \nu_i)}$ . As we set  $\lambda$  higher, the initial population is more risk-averse and less willing to hunt stag.

We use the payoff matrix in Table I to determine the raw payoff for exhibiting a particular pure-attitude profile  $\mathbf{a} = (\mu_1, \mu_2) \in A \times A$ . If  $\mathbf{a}$  is in the  $(h, h)$  region of the satisficing rectangle (see Figure 6), then the payoff to the first player is  $\pi(\mu_1, \mu_2) = 3$ . Similarly, the payoffs are  $\pi(\mu_1, \mu_2) = 3$  and  $\pi(\mu_1, \mu_2) = 0$  if  $\mathbf{a}$  belongs to the  $(h, s)$  and  $(s, h)$  regions, respectively. Finally,  $\pi(\mu_1, \mu_2) = 4\sigma$  if  $\mathbf{a}$  is in the  $(s, s)$  region.<sup>7</sup>

Because of the high dimensionality of the state space and the complexity of the utility functions of the players' preferences, it is difficult to examine the attitude dynamics analytically. We cannot easily solve for stationary points or say much about the relative sizes of the basins of attraction as we could under the (much simpler) standard replicator dynamics. Fortunately, we can specify meaningful initial conditions and numerically approximate the solution to the system of differential equations. We examine several scenarios where the vast majority of the

<sup>7</sup>We multiply by  $\sigma$  in the payoff to account for the probability that the players succeed given that they both hunt stag.

population hunts hare and discuss when it is possible to evolve a cooperative community.

First, we examine the dynamics with  $\sigma = 1$ . We initialize the population with  $\lambda = 10$ , leaving over 85% of the population hunting hare. Figure 8(a) shows the initial joint probability mass function of the players plotted along with the four regions of the satisficing rectangle. The vertical axis shows the joint probability that a pair of players—randomly selected from the population—will end up at a particular point. Since the players are drawn randomly and independently from the infinite population, the joint probability is the product of the marginal probabilities given by  $\mathbf{z}$ . That is,  $Pr(\mu_1 = \nu_i, \mu_2 = \nu_j) = z_i(t)z_j(t)$ .

Initially, almost all of the joint probability mass is in the mutual hare-hunting region. The dynamics, however, quickly pushes the population towards stag-hunting. Within thirty iterations, almost the entire population is in the mutual stag-hunting region, the most common values of  $(\mu_i, \mu_j)$  close to zero (Figure 8(b)). This is due to the fact that mutual cooperation is the only attitude equilibrium when  $\sigma = 1$ . For any positive, finite  $\lambda$ , all steady-state population distributions will be entirely within the  $(s, s)$  region.

Next, we lower  $\sigma$  to see how the dynamics changes. Keeping the initial conditions the same, we let  $\sigma = 0.925$ , introducing the  $(h, h)$  attitude equilibrium region. Now, over 90% of the initial population hunts hare. This scenario yields a highly interesting result. The hare hunting equilibrium initially dominates and the population shares associated with the stag hunting regions quickly diminish (Figure 9(a)). We notice, however, that there are small migrations toward the boundaries of the decision regions. These players still predominantly hunt hare, but they are less risk-averse. As evolution continues, a small concentration of players emerges around the boundaries of the four regions, as illustrated in Figure 9(b). Players in this region are quite versatile: they hunt hare with risk-averse players, hunt stag with the payoff-seekers, and only very rarely will they end up hunting stag with a player who refuses to cooperate. The concentration of players slowly begins to dominate, causing more and more players to hunt stag. Figure 9(c) shows the population at  $t = 100$ . By this time, essentially all of the population is composed of moderately risk-averse but versatile players. This truly emergent result provides an interesting insight in defining “fitness” in a social system. In an uncertain scenario where both hare-hunting and stag-hunting are potentially dominant strategies, the most successful players are those who are flexible—those who can adapt their actions to the preferences of those around them.

If we lower  $\sigma$  much below 0.925, the dynamics fails to evolve the society toward cooperation for these initial conditions. This happens for two reasons: (1) the size of the  $(s, s)$  region becomes smaller with decreasing  $\sigma$ , and (2) the expected payoff for exhibiting attitudes in the  $(s, s)$  region decreases. However, even under the unfavorable conditions shown where a pair of stag hunters might fail, the satisficing model can evolve cooperation from noncooperation. Fewer than 10% of the initial population are required to hunt stag in the satisficing model, a significant improvement over the standard replicator model, where over 75% must initially hunt




Fig. 8. Joint attitude distribution for  $\sigma = 1, \lambda = 10$ .

stag.

## B. Spatial Evolutionary Models

For comparison, we also consider the Stag Hunt under the spatial evolutionary models discussed in [15, 16, 28–30], which have been proven effective in promoting cooperation in social dilemmas. In [15], the Stag Hunt is specifically studied in terms of the relative benefit for mutual stag hunting. Here, we examine the question in terms of initial population: what fraction of the population must initially hunt stag in order for cooperation to flourish?

Spatial evolutionary models are described by undirected graphs, where each vertex represents a player, and each edge represents a social link between two players. As with the replicator dynamics, each player is pre-programmed to play a particular pure strategy. But, in the spatial dynamics, a player may change strategies depending on the relative fitness of

its neighbors. At each generation, players accrue payoff by playing a single instance of the game with each neighbor. After play, each player randomly selects a neighbor (possibly itself) with a probability proportional to the payoff accrued in the current previous round, adopting that player's pure-strategy for the next round.

We may interpret the spatial dynamics as an imitation dynamics, where a player imitates the behavior of its neighbors, or as a death-birth dynamics, where players "die" and give rise to a new generation whose strategies depend on the neighbors' relative fitness. Regardless of interpretation, for fully connected graphs, the dynamics converges to the standard replicator dynamics as the population size becomes large and the time between generations becomes small.

In the Stag Hunt, let  $N_s(i)$  and  $N_h(i)$  be the set of player  $i$ 's neighbors (including itself) that hunt stag and hare, respectively, and let  $P(i)$  denote the payoff earned by player  $i$  during a single generation. Thus, letting  $|\cdot|$  denote the cardinality of a set, player  $i$  earns  $F(i) = 4(|N_s(i)| - 1)$  if it hunts stag, and  $F(i) = 3(|N_s(i)| + |N_h(i)| - 1)$  if it hunts hare.<sup>8</sup> Next, define  $F_s(i) = \sum_{j \in N_s(i)} F(j)$  and  $F_h(i) = \sum_{j \in N_h(i)} F(j)$ , the respective sum payoff of stag- and hare-hunting neighbors. Finally, since a neighbor is selected with a probability proportional to its fitness, player  $i$  hunts stag during the next generation with probability  $F_s(i)/(F_s(i) + F_h(i))$ .

The spatial dynamics is highly dependent on the structure of the graph used to model the population. We construct our graphs according to so-called "scale-free" models [36], in which the number of neighbors follows a power-law distribution. If  $K_i$  is the random variable describing the number of neighbors for player  $i$ , then each  $K_i$  is identically and independently distributed according to  $p_{K_i}(k) \propto k^{-\gamma}$  for some constant  $\gamma$ . This distribution describes a heterogeneous, and realistic, model of social connectivity: many players have only a few neighbors, while a few players are heavily connected to the rest of the population. Scale-free models have been shown to improve the possibility of cooperation in social dilemmas [15].

To evaluate the performance of the spatial dynamics, we construct graphs with 50 players, an average number of connections per player  $z = E(K)$ , and an initial fraction of the population  $x_s(0)$  hunting stag. For each  $(x_s(0), z)$  pair, we construct ten graphs, each of which is seeded with ten initial populations. After running the dynamics for 5000 generations, we record the steady state behavior by averaging the fraction of stag hunters over an additional 500 generations. Figure 10 shows the average results of our trials. For moderately low values of  $z$ , the spatial dynamics considerably improves the possibility for cooperation: a sizeable fraction of the steady-state population hunts stag even when only a quarter of the initial population cooperate. This result is consistent with previous studies of cooperation in spatial networks [15, 16]. When the average number of connections is small, cooperation emerges more readily. However, in contrast to the attitude dynamics, stag hunting does not consistently dominate the

population unless a solid majority of players initially cooperate.

## VII. CONCLUSION

In this paper, we have extended the theory of satisficing games by incorporating elements from non-cooperative game theory. We augment the satisficing game with a standard utility function that gives the raw payoff to a player for exhibiting particular attitudes. The augmented framework results in an attitude equilibrium in which no single player can improve its raw payoff by exhibiting different attitudes. The attitude equilibrium combines the merits of both satisficing and non-cooperative game theory. The conditional utility structure allows players to consider others' preferences in making decisions, and the standard payoff function allows players to adapt their attitudes to avoid dysfunctional behavior.

The non-cooperative elements of augmented satisficing games allow us to employ evolutionary game theory, where adaptation occurs by trial-and-error. We define an attitude dynamics by applying the standard replicator dynamics to the attitudes exhibited by the players, rather than the strategies play. The attitude dynamics models the evolution of players' attitudes according to the game and the attitudes of other players. Given appropriate initial conditions, the steady state of the dynamics is an attitude equilibrium.

We have presented a satisficing model for the Stag Hunt, a game under which it is difficult to evolve a cooperative population. Under the augmented satisficing framework, dysfunctional behaviors vanish: the attitude equilibria lie entirely within the regions where players either mutually hunt stag or mutually hunt hare. Also, the attitude dynamics facilitates the evolution of cooperation by introducing strategic complexity into the dynamics. Instead of simply choosing whether or not to hunt stag, a player chooses a risk-aversion level, which governs its interaction with the rest of the population. Under a wide variety of circumstances, the dynamics encourages the population to become less risk averse, allowing cooperation to flourish. Our results significantly outperform other evolutionary methods, including classic replicator models and recently-proposed spatial evolutionary models.

Finally, the theoretical properties that borrow from non-cooperative game theory suggest that our results will generalize to large classes of games. Specifically, any game with finite attitude spaces must have an attitude equilibrium, and any (properly initialized) steady state of the attitude dynamics is an attitude equilibrium. While we cannot, of course, guarantee any specific results, we may expect that the qualitative benefits of our approach will pertain to other games.

## REFERENCES

- [1] M. S. Nokleby and W. C. Stirling, "The stag hunt: A vehicle for evolutionary cooperation," in *Proc. of IEEE World Congress on Computational Intelligence*, Vancouver, BC, Canada, July 2006.
- [2] M. Nokleby and W. C. Stirling, "Attitude adaptation in satisficing games," in *Proc. of IEEE Symposium on Foundations of Computational Intelligence*, Honolulu, Hawaii, USA, April 2007.
- [3] J. F. Nash, "Non-cooperative games," *Annals of Mathematics*, vol. 54, no. 2, pp. 286–295, September 1951.
- [4] R. D. Luce and H. Raiffa, *Games and Decisions*. New York: John Wiley, 1957.

<sup>8</sup>The  $(-1)$  term in each payoff accounts for the fact that, although  $N_s(i)$  or  $N_h(i)$  includes player  $i$ , the player does not pair with itself during play.

- [5] A. K. Sen, "Rational fools: A critique of the behavioral foundations of economic theory," in *Scientific Models and Man*, H. Harris, Ed. Oxford: Clarendon Press, 1979, ch. 1.
- [6] A. Tversky and D. Kahneman, "Rational choice and the framing of decisions," in *Rational Choice*, R. M. Hogarth and M. W. Reder, Eds. Chicago: Univ. Chicago Press, 1986.
- [7] E. Sober and D. S. Wilson, *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard Univ. Press, 1998.
- [8] W. C. Stirling, *Satisficing Games and Decision Making: with applications to engineering and computer science*. Cambridge, UK: Cambridge Univ. Press, 2003.
- [9] J. K. Archibald, J. C. Hill, F. R. Johnson, and W. C. Stirling, "Satisficing negotiations," *IEEE Transactions on Systems, Man, and Cybernetics (Part C)*, vol. 36, no. 1, pp. 4–18, 2006.
- [10] P. Ordeshook, *Game Theory and Political Theory: An Introduction*. Cambridge, UK: Cambridge Univ. Press, 1986.
- [11] E. Fehr and S. Gächter, "Altruistic punishment in humans," *Nature*, vol. 415, pp. 137–40, 2002.
- [12] C. Wedekind and M. Milinski, "Cooperation through image scoring in humans," *Science*, vol. 288, pp. 850–2, 2000.
- [13] K. Tuyls and A. Nowé, "Evolutionary game theory and multi-agent reinforcement learning," *The Knowledge Engineering Review*, vol. 20, no. 1, pp. 63–90, 2005.
- [14] M. Ruijgrok and T. W. Ruijgrok, "Replicator dynamics with mutations for games with a continuous strategy space," *arXiv:nlin/0505032v2*, 2005.
- [15] F. C. Santos, J. M. Pacheco, and T. Lenaerts, "Evolutionary dynamics of social dilemmas in structured heterogeneous populations," *Proc. National Academy of Sciences*, vol. 103, no. 9, pp. 3490–3494, 2006.
- [16] H. Ohtsuki, C. Hauert, E. Lieberman, and M. A. Nowak, "A simple rule for the evolution of cooperation on graphs and social networks," *Nature*, vol. 441, pp. 502–505, 2006.
- [17] J. Maynard Smith, "The theory of games and the evolution of animal conflicts," *J. Theor. Biol.*, vol. 47, pp. 209–21, 1974.
- [18] ———, *Evolution and the Theory of Games*. Cambridge, UK: Cambridge University Press, 1982.
- [19] P. Taylor and L. Jonker, "Evolutionarily stable strategies and game dynamics," *Math. Biosc.*, vol. 40, no. 2, pp. 146–46, 1978.
- [20] B. Skyrms, "The stag hunt," in *Proceedings and Addresses of the APA*, vol. 75, 2001, pp. 31–41, presidential Address of the Pacific Division of the American Philosophical Association.
- [21] W. Güth, "An evolutionary approach to explaining cooperative behavior by reciprocal incentives," *International Journal of Game Theory*, vol. 24, no. 4, pp. 323–344, 1995.
- [22] W. Güth and H. Kliemt, "The indirect evolutionary approach," *Rationality and Society*, vol. 10, no. 3, pp. 377–399, 1998.
- [23] H. A. Simon, "A behavioral model of rational choice," *Quart. J. Econ.*, vol. 59, pp. 99–118, 1955.
- [24] W. C. Stirling, "Social utility functions, part 1 — theory," *IEEE Transactions on Systems, Man, and Cybernetics (Part C)*, vol. 35, no. 4, pp. 522–532, 2005.
- [25] I. Steedman and U. Krause, "Goethe's *Faust*, Arrow's Possibility Theorem and the individual decision maker," in *The Multiple Self*, J. Elster, Ed. Cambridge, UK: Cambridge Univ. Press, 1985, ch. 8.
- [26] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*. San Mateo, Ca: Morgan Kaufmann, 1988.
- [27] J. H. Murray, H. Bradley, W. A. Craigie, and C. T. Onions, Eds., *The Compact Oxford English Dictionary*. Oxford: Clarendon Press, 1991, second Edition.
- [28] M. A. Nowak and R. M. May, "Evolutionary games and spatial chaos," *Nature*, vol. 437, pp. 826–29, 1992.
- [29] T. Killingback and M. Doebeli, "Spatial evolutionary game theory: Hawks and Doves revisited," *Proc. R. Soc. Lond.*, vol. 263, pp. 1135–44, 1996.
- [30] B. Skyrms and R. Premantle, "A dynamic model of social network function," *arXiv:math/0404101v1*, 2004.
- [31] R. Selten, "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragertragheit," *Zeitschrift für die gesamte Staatswissenschaft*, vol. 12, pp. 301–324, 1965.
- [32] J. W. Weibull, *Evolutionary Game Theory*. Cambridge, MA: MIT Press, 1995.
- [33] J. Hofbauer and K. Sigmund, *Evolutionary Games and Population Dynamics*. Cambridge, UK: Cambridge Univ. Press, 1998.
- [34] D. Foster and P. Young, "Stochastic evolutionary game dynamics," *Theor. Population Biol.*, vol. 38, no. 2, 1990.
- [35] A. Cabrales, "Stochastic replicator dynamics," *International Economic Review*, vol. 41, no. 2, pp. 451–482, 2000.
- [36] A. Barabasi and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, pp. 509–512, 1999.





(a)  $t = 5$ (b)  $t = 50$ (c)  $t = 100$ Fig. 9. Joint attitude distribution for  $\sigma = 0.925$ ,  $\lambda = 10$ .

Fig. 10. Average steady-state stag-hunting fraction under spatial evolutionary dynamics.