

Attitude Adaptation in Satisficing Games

Matthew Nokleby, *Student Member, IEEE*, 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 depend on other players' attitudes rather than simply their actions. However, satisficing players with conflicting attitudes may enact dysfunctional behaviors, which results in poor performance. We present an evolutionary method by which a population of players may adapt their attitudes to improve payoff. In addition, we extend the Nash-equilibrium concept to satisficing games, showing that the method leads players toward the equilibrium in their attitudes. We apply these ideas to the stag hunt—a simple game in which cooperation does not easily evolve from noncooperation. The evolutionary method 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 we present than under traditional evolutionary models.

Index Terms—Adaptive systems, cooperative systems, game theory, replicator dynamics, satisficing games.

I. INTRODUCTION

GAME-THEORETIC models are often used to construct societies of artificial agents. Commonly, agents are modeled as players in a noncooperative game in which players solely focus 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 noncooperative game theory. It was developed for the synthesis of artificial agents and particularly focuses on social interactions between players. The players' utilities are expressed as conditional mass functions, allowing them to consider the prefer-

ences of others rather than solely focusing on individual self interest. Satisficing models have previously been successful in overcoming the social hurdles that were presented by noncooperative game theory, allowing players to exhibit sophisticated social behaviors such as altruism, negotiation, and compromise [9]. However, satisficing game theory presents its own set of challenges. As in real-life social situations, satisficing communities may dysfunctionally behave. When players with incompatible attitudes are grouped together, they can choose incoherent behaviors that lead to poor performance.

The stag hunt, a simple game that was 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 the stag together, but each can separately catch a (much smaller) hare. That is, a player earns maximum payoff if both players cooperate but risks failure if it attempts to cooperate whereas the other does not. Each player must individually decide between cooperation and noncooperation; thus, it represents a useful model for the analysis of potentially cooperative behavior. For example, a group of workers who choose whether to loosely strike fall under the stag-hunt model: a large number of workers may achieve a significant benefit by striking, whereas a single worker who "strikes" alone incurs significant loss.

Social dilemmas such as the stag hunt have extensively been studied by (among others) social scientists, economists, and biologists. A large body of recent work has focused on learning-based [11]–[13] and evolutionary [14]–[16] methods for achieving cooperation. In evolutionary game theory, which was 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 that govern the evolution of the population. Under suitable conditions, the replicator dynamics drives the population to a Nash equilibrium.

The stag hunt presents considerable difficulties based on an evolutionary perspective. Under the standard replicator dynamics, a population with primarily hare hunters cannot evolve into a group of stag hunters, although 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

Manuscript received May 11, 2007; revised April 3, 2008. First published June 23, 2009; current version published November 18, 2009. This work was supported in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under Grant W911NF-07-1-0650. This work was presented in part at the 2006 IEEE World Congress on Computational Intelligence and the 2007 IEEE Symposium on Foundations of Computational Intelligence. This paper was recommended by Associate Editor E. Santos.

M. Nokleby was with the Department of Electrical and Computer Engineering, Brigham Young University, Provo, UT 84602 USA. He is now with the Department of Electrical and Computer Engineering, Rice University, Houston, TX 77251-1892 USA (e-mail: nokleby@rice.edu).

W. Stirling is with the Department of Electrical and Computer Engineering, Brigham Young University, Provo, UT 84602 USA (e-mail: wynn@ee.byu.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSMCB.2009.2021013

elements of non-cooperative game theory into satisficing game theory. In a manner similar to [21] and [22], we present a method where 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 more readily reproduce, causing their attitudes to dominate the population. The resulting model blends the two decision theories: players retain the conditional utility structure of satisficing game theory while improving the payoff by evolutionary means. The dynamics leads the players toward the 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

Although the simple seemingly reasonable assumption of self interest—also called individual rationality—has given rise to a rich 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 that 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: 1) one utility for characterizing the benefits associated with taking an action and 2) another utility for characterizing the costs. A satisficing player contents itself with a decision for which the benefits outweighing the costs is “good enough” or satisficing.¹ Second, the players’ utility functions share a common syntax with probability mass functions, allowing probabilistic concepts, e.g., 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 unusual 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 in terms of the reconciliation of group and individual preferences. Fortunately,

¹Although they share similarities, satisficing game theory should not be confused with the concept of “bounded rationality” satisficing as introduced by Simon [23]. With Simon’s satisficing, individuals search for suboptimal choices that meet a variable threshold or *aspiration level*, implicitly accounting for the cost of continued searching.

however, the benefits of conditional utilities may also intuitively be appreciated. 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 the hypothetical probabilities of Y : What would be the probability 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 the basic rules of probability theory, i.e., $p_Y(y) = \sum_x p_{Y|X}(y|x)p_X(x)$. The marginal probabilities for Y are influenced—but not entirely dictated—by the probabilities of X .

Similarly, players’ preferences may depend on 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, i.e., the utilities that were used for decision making, by summing the conditional utilities over Player 2’s actual preferences. This structure allows players to consider not only what actions other players may prefer but also *how strong* the preferences for action are. Their utilities are influenced by others’ preferences in a controlled manner, which does not require the players to 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 section, each player possesses two social utilities. To describe these social utilities, 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, whereas 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.

Social utilities are mass functions; thus, 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.” Although the former concept allows for a set of multiple actions that are “good enough,” the latter concept 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 Σ_i is non-empty. We may combine the players’ individually satisficing sets by forming the *satisficing*

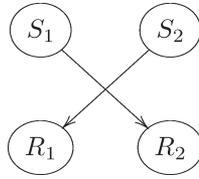


Fig. 1. Simple praxeic network.

rectangle $\mathfrak{R}_{1,2,\dots,n}$, which is defined as the Cartesian product, i.e.,

$$\mathfrak{R}_{1,2,\dots,n} = \Sigma_1 \times \Sigma_2 \times \dots \times \Sigma_n. \quad (2)$$

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

It is convenient to graphically express the relationship between players' utilities. In probability theory, relationships between random variables are expressed in Bayesian networks [26]. Similarly, in satisficing game theory, the relationships between players' utilities are expressed in *praxeic networks*.² The praxeic network consists of a directed acyclic graph, 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 in Fig. 1. For each player, the rejecting preferences depend on the selecting preferences of the other player, whereas the selecting preferences are independent.

Parenthetically, we note that praxeic networks also resemble the spatial evolutionary models in [15], [16], and [28]–[30]. In these models, graphical connections determine which players may 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, the players' social relationships. However, spatial evolutionary models describe which players can pair up in a game, whereas 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 in Fig. 1, we refer to Player 1's *conditional* rejectability function, which is denoted as $p_{R_1|S_2}(v_1|u_2)$. As aforementioned, 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, i.e., $p_{R_1}(v_1) = \sum_{u_2 \in U_2} p_{R_1|S_2}(v_1|u_2)p_{S_2}(u_2)$. The marginal utilities determine the individually satisficing sets and the satisficing rectangle. If a utility is independent (e.g., the selectability functions), its marginal is directly expressed 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 to compute its marginal utilities and find

²The term *praxeic* is derived from *praxeology*, which refers to "the science of human conduct" or "the science of efficient action" [27].

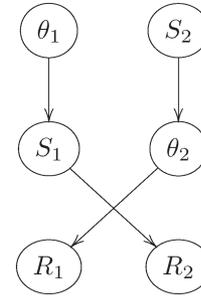


Fig. 2. Praxeic network with "true" random variables.

its individually satisficing set. In the example community, each player must know the other player's selectability function to compute its own rejectability. However, players do not consider each other's actions in determining the individually satisficing sets; thus, 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

$$\begin{aligned} p_{S_1 S_2 R_1 R_2}(u_1, u_2, v_1, v_2) \\ = p_{R_1|S_2}(v_1|u_2)p_{R_2|S_1}(v_2|u_1)p_{S_1}(u_1)p_{S_2}(u_2). \end{aligned}$$

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. Based on 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

Oftentimes, a player's utility will depend on random phenomena, which result 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 are not influenced by the preferences of the players. In other words, a player's belief on 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. Fig. 2 illustrates a network that implements such a

TABLE I
PAYOFF MATRIX FOR A TWO-PLAYER STAG HUNT

	Player 2	
Player 1	<i>s</i>	<i>h</i>
<i>s</i>	(4, 4)	(0, 3)
<i>h</i>	(3, 0)	(3, 3)

model. This praxeic network is similar to the network in Fig. 1 in that it contains the same four vertices that were associated with the players’ selecting and rejecting selves. However, we also include two random variables θ_1 and θ_2 , which represent phenomena that are only probabilistically known to the players. This network describes both players whose preferences depend on random phenomena and random phenomena that depend on players’ preferences. The dependencies in Fig. 1 still persist. R_1 still depends—but indirectly, through θ_2 —on S_2 , and R_2 still depends on S_1 , which now depends on θ_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 action of the other player. If the other player hunts a stag, the payoff for hunting a stag is higher than that of hunting a hare. However, if the other player hunts a 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 a 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 in Table I.

There are two pure-strategy Nash equilibria for the stag hunt: 1) (s, s) and 2) (h, h) . If the players simultaneously hunt stag or a 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, successful stag hunting requires the cooperation of the other player; thus, risk-averse players may instead choose 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 hunt 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 a certain assurance that the other player would hunt stag, everyone would cooperate.³ However, players do not have such an assurance under the usual model but must independently choose their actions. The players’ actions then boil down to how much confidence each

³Interestingly, it is straightforward to show that, if the game is sequentially played (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, the 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 pure-strategy equilibrium, neither player will have an incentive to deviate. To study which equilibrium will result under different circumstances, we turn to evolutionary game theory [32], [33].

A. Replicator Dynamics

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 asexually⁴ reproduces according to its payoffs, i.e., the number of offspring that a player has is proportional to its payoff during the previous game. Players’ strategies also “breed true,” which means 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* Δ_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 Δ_U is the set of mixed strategies that assign nonzero probability to each pure strategy, i.e.,

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

In the replicator dynamics, we interpret each element x_i as the *population share* for i , or the fraction of the population that plays the pure strategy i . That is, if we randomly draw an individual from the population that was described by \mathbf{x} , the probability that it will be programmed to play i is x_i . At time t , the expected utility⁵ 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 strategies that 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. We have

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

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

⁵We use π to represent the utility (or payoff) when players use only pure strategies, whereas 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, i.e.,

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

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 that plays 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 the model in [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 that plays strategy j spontaneously switches to strategy i . The mutation dynamics differs from (3) by the addition of a mutation term, i.e.,

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

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 that were guaranteed under the standard replicator model. The steady-state behavior of the system no longer corresponds to the Nash equilibria, regardless of the initial conditions.

B. Stag-Hunt Replicator Dynamics

1) *Standard Dynamics*: For the stag hunt, the population is described by the 2-D vector $\mathbf{x} = (x_s, x_h)$. The payoff matrix (see Table I) shows that the payoff for a stag hunter is four when paired with another stag hunter, and it is zero when paired with a hare hunter. A stag hunter therefore gains an expected utility of $u(s, \mathbf{x}) = 4x_s$. The utility for hunting a hare is independent of the other player’s actions; thus, $u(h, \mathbf{x}) = 3$. The population’s average expected payoff is given by $u(\mathbf{x}, \mathbf{x}) =$

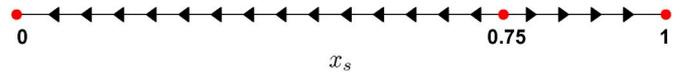


Fig. 3. Direction field for the stag-hunt replicator dynamics.

$4x_s^2 - 3x_s + 3$. $x_s = 1 - x_h$; thus, we can characterize the dynamics by examining only the stag hunting share. Suppressing the time arguments, we get

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

Although non-linearities prevent a closed-form solution, we can easily examine the qualitative behavior of the population. In Fig. 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 aforementioned mixed-strategy Nash equilibrium. 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 noncooperative population.

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

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

The dynamics for x_s becomes

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

The closed-form expression for the stationary points of the dynamics is quite unwieldy. So, in Fig. 4, we plot the direction field for the dynamics as a function of α and x_s . When the 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 a stag or a hare and one unstable stationary point that defines the boundary between the stag-hunting and hare-hunting basins of attraction. The boundary point increases with the mutation rate, which suggests that mutation exacerbates the evolutionary difficulties of the stag hunt.

For large mutation probabilities, the dynamics considerably differs, leaving a single stationary point to which the dynamics converges, independent of the 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. The population size is infinite; thus, the

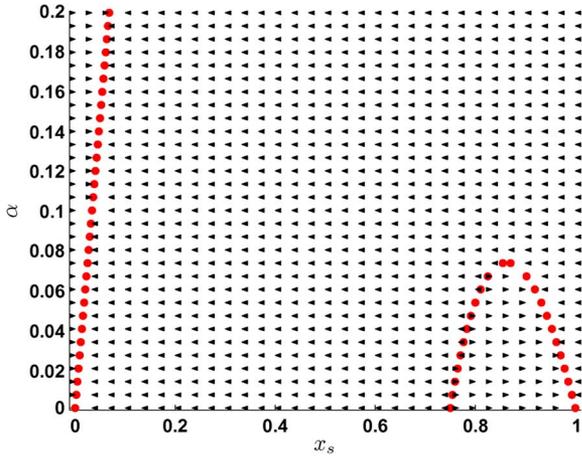


Fig. 4. Direction field for the stag-hunt stochastic replicator dynamics.

mutation replicator model defines a deterministic system as in the standard dynamics. Consequently, finite populations, with random pairings and mutation, may spontaneously evolve cooperation from noncooperation. However, 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 preprogrammed to play a particular pure strategy.

Finally, as we have already discussed, there exist evolutionary models other 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, which is based on 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 that 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 the 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, which is tempered by risk aversion. The opportunity cost of hunting a hare is the payoff for catching a stag, and the opportunity cost of hunting a 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 Fig. 5. In addition to the vertices that correspond to the selecting and

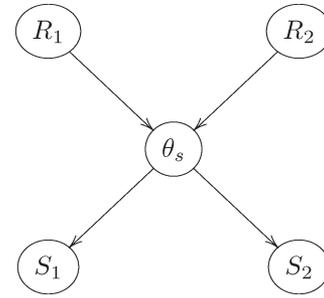


Fig. 5. Praxeic network for the stag hunt.

rejecting selves, we include a vertex that corresponds to a binary random variable θ_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.

To define the rejectability function for each agent, we must first define a normalized measure of opportunity cost. Let ϕ_{s_i} and ϕ_{h_i} denote the raw utility (in arbitrary units) of consuming a stag and a hare, respectively. By normalizing, the relative utility of hare hunting becomes $\mu'_i = (\phi_{h_i}/(\phi_{h_i} + \phi_{s_i}))$ 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 $\phi_{h_i} = 3$, i.e., the payoff values in Table I, resulting in $\mu'_i = (3/7)$. However, we further wish to take into account the temperament of the players. As discussed in Section III, one central issue in the stag hunt is to determine what players of different risk-aversion levels should do. Therefore, we introduce a parameter, i.e., ρ_i , that 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 $\mu_i = \rho_i(\phi_{h_i}/(\phi_{h_i} + \phi_{s_i}))$. Thus, μ_i reflects a player’s willingness to take risks and the relative utility for a stag and a hare. A maximally risk-averse player will hunt stag only if success is certain, whereas 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 a hare to a stag, i.e., $\mu_i < (1/2)$ for $i = 1, 2$. For convenience, we will simply refer to μ_i as player i ’s *risk-aversion level*, which parameterizes the player’s attitudes.

We define each player’s rejectability function as

$$p_{R_i}(u_i) = \begin{cases} \mu_i, & \text{for } u_i = s \\ 1 - \mu_i, & \text{for } u_i = h \end{cases} \quad (7)$$

which is an expression of the normalized opportunity cost for each action. The cost of hunting a 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 directly define the marginal utilities.

We next define the conditional distribution for θ_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 θ_s incorporates whether 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 σ . We characterize this condition by defining

$$p_{\theta_s|R_1R_2}(\vartheta_s|h, h) = \begin{cases} \sigma, & \text{for } \vartheta_s = 1 \\ 1 - \sigma, & \text{for } \vartheta_s = 0 \end{cases} \quad (8)$$

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

$$\begin{aligned} p_{\theta_s|R_1R_2}(\vartheta_s|s, s) &= p_{\theta_s|R_1R_2}(\vartheta_s|s, h) \\ &= p_{\theta_s|R_1R_2}(\vartheta_s|h, s) \\ &= \begin{cases} 0, & \text{for } \vartheta_s = 1 \\ 1, & \text{for } \vartheta_s = 0. \end{cases} \end{aligned} \quad (9)$$

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

$$\begin{aligned} p_{\theta_s}(\vartheta_s) &= \sum_{v_1, v_2} p_{\theta_s|R_1, R_2}(\vartheta_s|v_1, v_2)p_{R_1}(v_1)p_{R_2}(v_2) \\ &= \begin{cases} \sigma(1 - \mu_1)(1 - \mu_2), & \text{for } \vartheta_s = 1 \\ 1 - \sigma(1 - \mu_1)(1 - \mu_2), & \text{for } \vartheta_s = 0. \end{cases} \end{aligned} \quad (10)$$

Based on (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 σ . Either player can reduce the chances for a successful hunt. As the risk-aversion μ_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 we have previously discussed, 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 becomes, the more beneficial it is to hunt stag. The corresponding conditional selectability function is

$$p_{S_i|\theta_s}(u_i|\vartheta_s) = \begin{cases} 1, & \text{for } u_i = s|\vartheta_s = 1 \\ 0, & \text{for } u_i = h|\vartheta_s = 1 \\ 0, & \text{for } u_i = s|\vartheta_s = 0 \\ 1, & \text{for } u_i = h|\vartheta_s = 0. \end{cases} \quad (11)$$

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

$$p_{S_i}(u_i) = \begin{cases} \sigma(1 - \mu_1)(1 - \mu_2), & \text{for } u_i = s \\ 1 - \sigma(1 - \mu_1)(1 - \mu_2), & \text{for } u_i = h. \end{cases} \quad (12)$$

A. 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 dis-

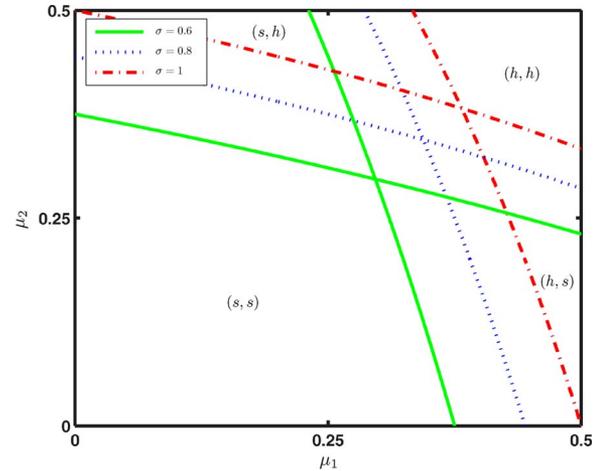


Fig. 6. Satisficing rectangle regions for the stag hunt.

cussed in Section II, the satisficing rectangle is the set of pure-strategy profiles that are simultaneously satisficing to each player. In Fig. 6, we set $q = 1$ and plot the regions of the satisficing rectangle as functions of μ_1 and μ_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, whereas the other player strongly seeks payoff, thus resulting in one player that tries to cooperate, whereas the other does not. On the boundaries of the four regions, the satisficing rectangle contains multiple strategy profiles.

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, whereas 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: the players' utilities depend on the others' attitudes rather than the strategies that they play.

We hasten to note that dysfunctional behavior is not a failure per se of the satisficing model. Dysfunctional societies exist in practice, and we may interpret these regions as an acknowledgement that players with incompatible attitudes may incoherently act. However, in designing artificial systems, we typically prefer to avoid incoherent behaviors, regardless of whether they are sociologically justifiable. 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 that 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. In addition, we introduce the pure-attitude space, i.e., $\mathbf{A} = A_1 \times A_2 \times \dots \times A_n$, which contains 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 $\pi(\mathbf{u})$, a vector payoff function that 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. Furthermore, we must 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 sufficiently be complicated such 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 μ_i , yielding a pure-attitude space of $\mathbf{A} = [0, 1/2] \times [0, 1/2]$. The payoff function $\pi(\mathbf{u})$ 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, which are determined by the interdependence function, map the players' attitudes to pure-strategy profiles.⁶ 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 the notation, we will occasionally refer to $\pi(\mathbf{a})$, i.e., 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 noncooperative game, where the players' payoffs are determined by the attitudes that they exhibit rather than the strategies that 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 directly applies to mixed attitudes. We assume that the players' mixed attitudes are probabilistically independent of each other. We define player i 's mixed attitude simplex Δ_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

$$u_i(\mathbf{z}) = \sum_{\mathbf{a} \in \mathbf{A}} \pi_i(\mathbf{a}) \prod_{i=1}^n z_{ia_i} \quad (13)$$

where z_{ia_i} is the probability with which player i exhibits the pure attitude a_i . Now, given complete knowledge of the satisficing

⁶We have glossed over the fact that, in general, the satisficing rectangle contains multiple pure-strategy profiles. For the stag hunt, this fact 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.

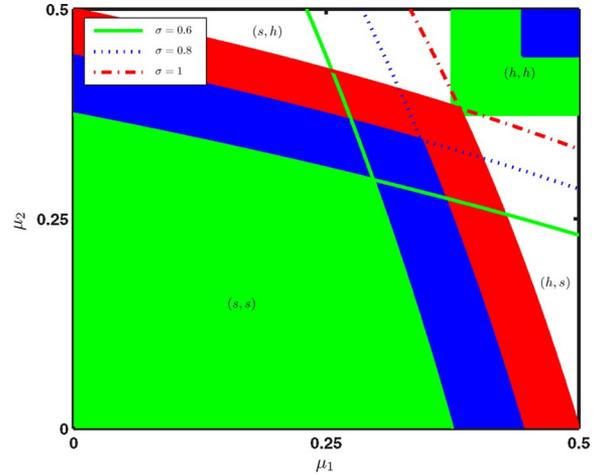


Fig. 7. Attitude equilibrium regions for the stag hunt.

ing game and the other players' utilities, a player may consider *changing* their attitudes to increase the expected utility, which motivates the attitude equilibrium.

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

$$u_i(\mathbf{z}_1^*, \dots, \mathbf{z}_i^*, \dots, \mathbf{z}_n^*) \geq u_i(\mathbf{z}_1^*, \dots, \mathbf{z}'_i, \dots, \mathbf{z}_n^*) \quad (14)$$

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 the attitude equilibrium is the equilibrium in the 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 on the fact that any augmented satisficing game defines a classical noncooperative game, where \mathbf{X} is the set of players, \mathbf{A} takes the role of the pure-strategy space, and $\pi(\mathbf{a})$ is the payoff function. In [3], it is shown that any noncooperative game with a finite pure-strategy space has at least one Nash equilibrium, although it may exist only in mixed strategies. One attitude equilibrium is simply one Nash equilibrium in the players' attitudes; thus, one must exist for any augmented satisficing game with a finite pure-attitude space, although 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. For the stag hunt, although the attitude spaces are continuous, it is immediate that attitude equilibria exist in pure attitudes. In Fig. 7, the attitude equilibria are shown for several values of σ . If the players' pure-attitude profile lies in these regions, there is no incentive for either player to change attitudes.

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, because when player i 's risk aversion μ_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 μ_j . Although (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 σ increases, the size of the (h, h) equilibrium decreases, entirely disappearing 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 μ_i and forcing the game into either (s, s) or (h, h) . The attitude equilibrium concept provides a useful juxtaposition of satisficing game 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 an 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 is 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 that exhibits the i th pure attitude. Similar to traditional games, we may describe the dynamics of the population shares by a system of m differential equations, i.e.,

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

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 Δ_A be the mixed-attitude simplex of A . Similar to mixed strategies, the interior of Δ_A is the set of all mixed attitudes that 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 directly follows from the fact that an augmented satisficing game can be considered 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. An attitude equilibrium is a Nash equilibrium in players’ attitudes; thus, 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 μ may assume. Define $A = \{\nu_1, \nu_2, \dots, \nu_{100}\}$, which is a set of 100 evenly spaced values of μ over $[0, 1/2)$. We initialize the population shares \mathbf{z} according to an exponential distribution

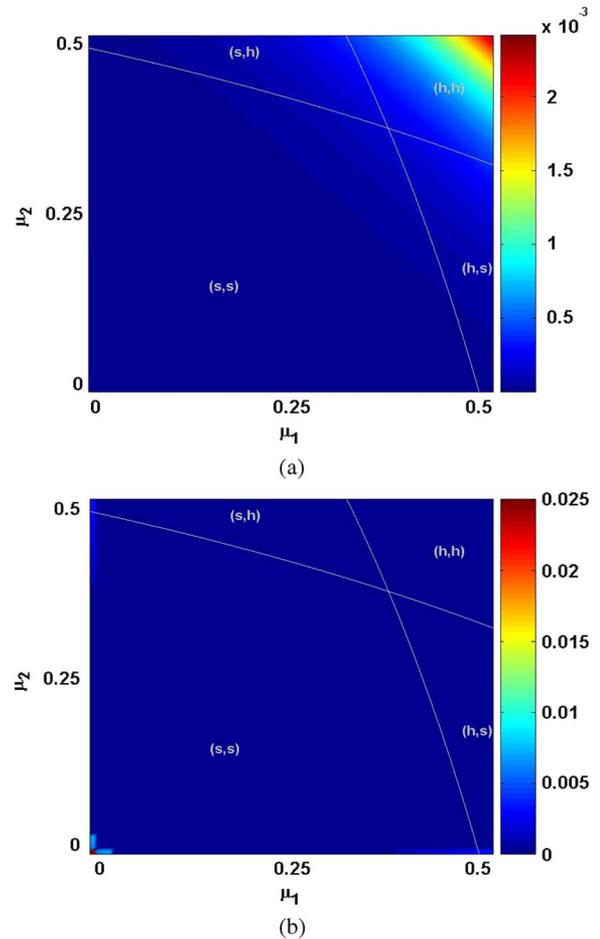


Fig. 8. Joint attitude distribution for $\sigma = 1$ and $\lambda = 10$. (a) $t = 0$. (b) $t = 30$.

so that most players hunt hare, i.e., $z_i(0) \propto e^{-\lambda((1/2)-\nu_i)}$. As we set λ 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 Fig. 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.⁷

Because of the high dimensionality of the state space and the complexity of the utility functions of the players’ preferences, it is difficult to analytically examine the attitude dynamics. 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 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 more than 85% of the population that hunts a hare. Fig. 8(a) shows the initial joint

⁷We multiply by σ in the payoff to account for the probability that the players succeed, given that they both hunt stag.

probability mass function of the players and the four regions of the satisficing rectangle. The vertical axis shows the joint probability that a pair of players, which were randomly selected from the population, will end up at a particular point. The players are randomly and independently drawn from the infinite population; thus, 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 toward stag hunting. Within 30 iterations, almost the entire population is in the mutual stag-hunting region, with the most common values of (μ_i, μ_j) being close to zero [see Fig. 8(b)]. This result is due to the fact that mutual cooperation is the only attitude equilibrium when $\sigma = 1$. For any positive finite λ , all steady-state population distributions will entirely be within the (s, s) region.

Next, we lower σ 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, more than 90% of the initial population hunts hare. This scenario yields a highly interesting result. The hare-hunting equilibrium initially dominates, and the population shares that were associated with the stag-hunting regions quickly diminish [see Fig. 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 Fig. 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 end up hunting stag with a player who refuses to cooperate. The concentration of players slowly begins to dominate, causing much more players to hunt stag. Fig. 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, i.e., players who can adapt their actions to the preferences of those around them.

If we lower σ much below 0.925, the dynamics fails to evolve the society toward cooperation for these initial conditions. This case happens for two reasons: 1) the size of the (s, s) region becomes smaller with decreasing σ , 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. Less than 10% of the initial population is required to hunt stag in the satisficing model, which is a significant improvement over the standard replicator model, where more than 75% must initially hunt stag.

B. Spatial Evolutionary Models

For comparison, we also consider the stag hunt under the spatial evolutionary models in [15], [16], and [28]–[30], which have been proven effective in promoting cooperation in social dilemmas. In [15], the stag hunt is particularly studied in terms

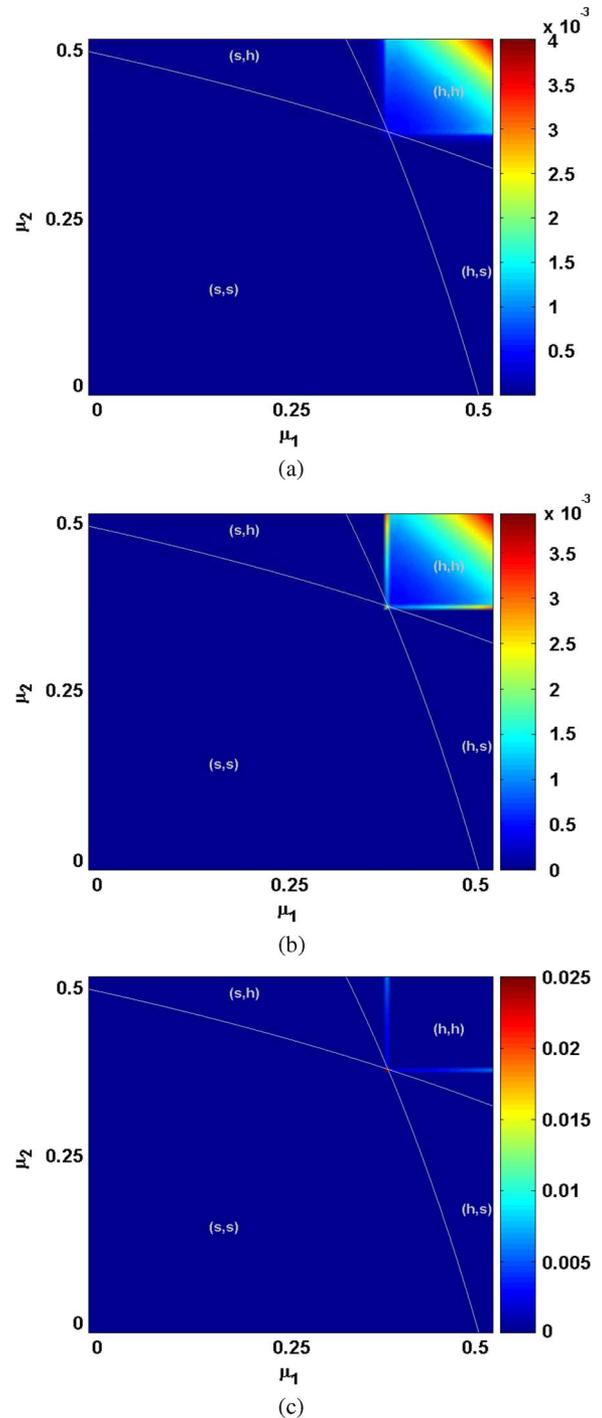


Fig. 9. Joint attitude distribution for $\sigma = 0.925$ and $\lambda = 10$. (a) $t = 5$. (b) $t = 50$. (c) $t = 100$.

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 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 preprogrammed to play a particular pure strategy. However, 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 that was 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 a hare, respectively, and let $P(i)$ denote the payoff that was 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.⁸ Next, define $F_s(i) = \sum_{j \in N_s(i)} F(j)$ and $F_h(i) = \sum_{j \in N_h(i)} F(j)$, i.e., the respective sum payoff of stag- and hare-hunting neighbors. Finally, a neighbor is selected with a probability proportional to its fitness; thus, 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 that was used to model the population. We construct our graphs according to the so-called “scale-free” models [36], in which the number of neighbors follows a power-law distribution. If K_i is the random variable that describes 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 γ . This distribution describes a heterogeneous and realistic model of social connectivity, i.e., many players have only a few neighbors, whereas 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. Fig. 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, although only a quarter of the initial population cooperates. This result is consistent with previous studies of cooperation in spatial networks [15], [16]. When the average number of connections is small, cooperation more readily emerges. However, compared to the attitude dynamics, stag hunting does not consistently dominate the population unless a solid majority of players initially cooperate.

⁸The (-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.

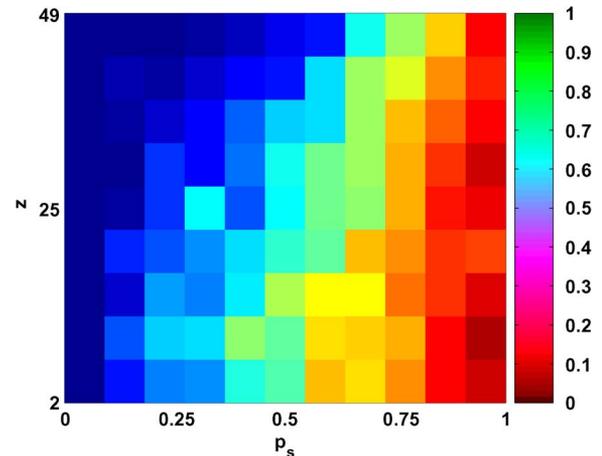


Fig. 10. Average steady-state stag-hunting fraction under spatial evolutionary dynamics.

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 describes 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 the satisficing and non-cooperative game theories. 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.

Non-cooperative elements of augmented satisficing games have allowed 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 that the players exhibited rather than the strategies that they 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 behavior vanishes: the attitude equilibria entirely lie within the regions where players either mutually hunt stag or mutually hunt hare. In addition, the attitude dynamics facilitates the evolution of cooperation by introducing strategic complexity into the dynamics. Instead of simply choosing whether 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 borrowed from non-cooperative game theory suggest that our results will generalize to large classes of games. In particular, 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. Although we cannot guarantee any specific results, we 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. IEEE World Congr. Comput. Intell.*, Vancouver, BC, Canada, Jul. 2006, pp. 348–355.
- [2] M. Nokleby and W. C. Stirling, "Attitude adaptation in satisficing games," in *Proc. IEEE Symp. Foundations Comput. Intell.*, Honolulu, HI, Apr. 2007, pp. 331–338.
- [3] J. F. Nash, "Noncooperative games," *Ann. Math.*, vol. 54, no. 2, pp. 286–295, Sep. 1951.
- [4] R. D. Luce and H. Raiffa, *Games and Decisions*. New York: Wiley, 1957.
- [5] A. K. Sen, "Rational fools: A critique of the behavioral foundations of economic theory," in *Scientific Models and Man*, H. Harris, Ed. Oxford, U.K.: Clarendon, 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, IL: Univ. of 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, U.K.: Cambridge Univ. Press, 2003.
- [9] J. K. Archibald, J. C. Hill, F. R. Johnson, and W. C. Stirling, "Satisficing negotiations," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 36, no. 1, pp. 4–18, Jan. 2006.
- [10] P. Ordeshook, *Game Theory and Political Theory: An Introduction*. Cambridge, U.K.: Cambridge Univ. Press, 1986.
- [11] E. Fehr and S. Gächter, "Altruistic punishment in humans," *Nature*, vol. 415, no. 6868, pp. 137–140, Jan. 2002.
- [12] C. Wedekind and M. Milinski, "Cooperation through image scoring in humans," *Science*, vol. 288, no. 5467, pp. 850–852, May 2000.
- [13] K. Tuyls and A. Nowé, "Evolutionary game theory and multiagent reinforcement learning," *Knowl. Eng. Rev.*, vol. 20, no. 1, pp. 63–90, Mar. 2005.
- [14] M. Ruijgrok and T. W. Ruijgrok, *Replicator Dynamics With Mutations for Games With a Continuous Strategy Space*, 2005. arXiv:nlin/0505032v2.
- [15] F. C. Santos, J. M. Pacheco, and T. Lenaerts, "Evolutionary dynamics of social dilemmas in structured heterogeneous populations," *Proc. Nat. Acad. Sci.*, vol. 103, no. 9, pp. 3490–3494, Feb. 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, no. 7092, pp. 502–505, May 2006.
- [17] J. Maynard Smith, "The theory of games and the evolution of animal conflicts," *J. Theor. Biol.*, vol. 47, no. 1, pp. 209–221, Sep. 1974.
- [18] J. Maynard Smith, *Evolution and the Theory of Games*. Cambridge, U.K.: Cambridge Univ. Press, 1982.
- [19] P. Taylor and L. Jonker, "Evolutionarily stable strategies and game dynamics," *Math. Biosci.*, vol. 40, no. 2, pp. 145–156, 1978.
- [20] B. Skyrms, "The stag hunt," in *Proc. Addresses APA*, 2001, vol. 75, 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," *Int. J. Game Theory*, vol. 24, no. 4, pp. 323–344, Dec. 1995.
- [22] W. Güth and H. Kliemt, "The indirect evolutionary approach," *Ration. Soc.*, vol. 10, no. 3, pp. 377–399, 1998.
- [23] H. A. Simon, "A behavioral model of rational choice," *Q. J. Econ.*, vol. 69, no. 1, pp. 99–118, Feb. 1955.
- [24] W. C. Stirling, "Social utility functions—Part 1—Theory," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 35, no. 4, pp. 522–532, Nov. 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, U.K.: Cambridge Univ. Press, 1985, ch. 8.
- [26] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kaufmann, 1988.
- [27] *The Compact Oxford English Dictionary*, 2nd ed. J. H. Murray, H. Bradley, W. A. Craigie, and C. T. Onions, Eds. Oxford, U.K.: Clarendon, 1991.
- [28] M. A. Nowak and R. M. May, "Evolutionary games and spatial chaos," *Nature*, vol. 437, no. 6398, pp. 826–829, Oct. 1992.
- [29] T. Killingback and M. Doebeli, "Spatial evolutionary game theory: Hawks and Doves revisited," *Proc. R. Soc. Lond.*, vol. 263, no. 1374, pp. 1135–1144, Sep. 1996.
- [30] B. Skyrms and R. Premantle, *A Dynamic Model of Social Network Function*, 2004. arXiv:math/0404101v1.
- [31] R. Selten, "Spieltheoretische behandlung eines oligopolmodells mit nachfragetragheit," *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, U.K.: Cambridge Univ. Press, 1998.
- [34] D. Foster and P. Young, "Stochastic evolutionary game dynamics," *Theor. Popul. Biol.*, vol. 38, no. 2, pp. 219–232, 1990.
- [35] A. Cabrales, "Stochastic replicator dynamics," *Int. Econ. Rev.*, vol. 41, no. 2, pp. 451–482, 2000.
- [36] A. Barabasi and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, Oct. 1999.



Matthew Nokleby (S'04) received the B.S. (*cum laude*) and M.S. degrees in electrical engineering from the Brigham Young University, Provo, UT, in 2006 and 2008, respectively. He is currently working toward the Ph.D. degree in electrical engineering at Rice University, Houston, TX.

His research interests include game theory and its applications to wireless communications.



Wynn Stirling received the B.A. (*magna cum laude*) degree in mathematics and the M.S. degree in electrical engineering from the University of Utah, Salt Lake City, in 1969 and 1971, respectively, and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, in 1983.

From 1972 to 1975, he was with the Rockwell International Corporation, Anaheim, CA. From 1975 to 1984, he was with the ESL, Inc., Sunnyvale, CA, where he was responsible for the development of multivehicle trajectory reconstruction capabilities. In 1984, he joined the faculty of the Department of Electrical and Computer Engineering, Brigham Young University, Provo, UT, where he is currently a Professor. He is the author or coauthor of more than 70 publications. He is a coauthor of *Mathematical Methods and Algorithms for Signal Processing* (Prentice-Hall, 2000) and the author of the monograph *Satisficing Games and Decision Making: With Applications to Engineering and Computer Science* (Cambridge University Press, 2003). His research interests include multiagent decision theory, estimation theory, information theory, and stochastic processes.

Dr. Stirling is a member of Phi Beta Kappa and Tau Beta Pi. He has served on the program committees of conferences on imprecise probability theory and multiagent decision theory.