Games with Minimalistic Agents

Asrar Ahmed and Kamalakar Karlapalem

Center For Data Engineering,
International Institute of Information Technology,
Hyderabad, India
ahmed@students.iiit.ac.in,kamal@iiit.ac.in
http://cde.iiit.ac.in

Abstract. In this paper we are study solution concepts when agents are interested to have a threshold utility or a cutoff above which they choose to benefit the system. Such behavior would be more relevant when we want agents to make socially responsible decisions. For example when agents mediate on behalf of humans or with humans themselves we would prefer agents to have such an attribute. We consider such a behavior to be more closer to human nature rather than maximizing ones own utility- in case of self interested agents, or always choose actions for the benefit of the system- in case of altruistic agents. To this end we have extended the notion of satisficing and present a formal analysis of games when agents preferences reflect such characteristic. Apart from discussing the solution concept for n player normal form game, we also consider the issues when not all agents can satisfy their minima. We then discuss the case when agents defect from the solution concept.

1 Introduction

The conventional game theory have discussed the solution concepts wherein agents are utility maximizers. Maximizing utility is considered as a rational choice [1]. In cooperative games agents are still utility maximizers, but the modeling unit is the coalition rather than an individual agent. Altruistic agents, which benefit the system at a cost to themselves have also been discussed [7].
In this paper, we present preliminary work which is motivated by the fact that neither of the above models capture the fact that we are maximizers under some conditions, altruistic under other conditions. If we map agent type to social awareness scale, we will have altruistic agents on one end and maximizers on the other end. We are interested in an agent type whose action preference profile reflects a right balance between its own performance and that of the system. By system, we mean the environment as well as other agents in it. Such a behavior *might* amount to ethical behavior in multi agent environment. We do not intend to digress into the issues related to ethical behavior at this stage. Discussion on what amounts to ethical behavior and issues arising from it have been discussed elsewhere. [2] gives a brief summary of machine ethics. [3] discusses the need to have machine ethics.

To study games involving agents with such an attribute, we incorporate an additional constraint, that agents have to reveal their minimum utility which they would want from a game. We now briefly discuss the concept of satisficing to highlight the underlying difference between minimalistic and satisficing agents.

Satisficing is a decision making approach whereby agents choose action which satisfy some minimal conditions. The argumentation by Simon [4] was human beings do not always have well defined pay-offs, or we can rarely evaluate all possible outcomes with accuracy. In such cases we tend only to suffice some minimal conditions rather than adopt an optimization approach which comes at a cost. Maximizing, as Simon argues is too hard for us. His argument was 1. Complexity involved in assigning utility to all possible outcomes. 2. Probability distribution over outcomes, given our actions. 3. Computational complexity involved in calculating the best outcome.

As is obvious, the similarity between minimalistic agent and satisficing is both are interested in a minimum utility. But once a minimalistic agent is ensured his minimum, the excess utility is distributed to other agents. Satisficing agent stops once it gets its minima whereas minimalistic agents will benefit the system once it gets its minima. Speaking in terms of the mouse example which Simon used, once the mouse is guaranteed with a piece of cheese, say cheddar, it tries to opt for solution concepts which will guarantee a minimum to other mice. Additionally, minimalistic agents satisfice their utility by choice and not because of complexity of calculating more optimal solution. [5] gives indepth analysis of satisficing games.

We can then safely categorize the agent behavior as follows

Maximizers: Agents who tend to maximize their utility.

Altruistic: Agents who prioritize system utility above their.

Minimalistic: Agents who would help other agents satisfy their minimum, once they are guaranteed their own minimum.

[6] has discussed the notion of ethical solution concept. But they consider the case when agents would choose to maximize their utility once they ensure minimum utilities to themselves and their opponent. Additionally, their discussion is limited to only two-player games.

With the above definition of what 'minimalistic' agent is, we discuss a solution concept we refer as minimal equilibrium. We discuss the solution concept for transferable and non transferable utility. We also define agent behavior for the case when not all agents can satisfy their minimum.

The rest of the paper is organized as follows. In section 2 we give an overview of the solution concept and discuss the example problems where it would be a 'rational or ethical' choice. In section 3 we formalize the solution concept and analyze different cases that might arise. In section 4 we analyze under what conditions would agents truthfully reveal their minima and in section 5 we briefly discuss applications and conclude.

2 Example Problem Scenario

Consider two agents p_1, p_2. Each agent requires a resource R and has preference over the action set A. In the following example, we consider R to be a natural resource.

1. Two players - Using Shared Natural Resource $p1, p2$ share a ground-water supply. Each needs water on daily basis for harvesting. Each agent has a minimum cut-off defined which it requires for that day. We assume for now that each agent reveals its true minimum value. The following normal form game reflects their preferences, agent-1 along row and agent-2 along column. The minimum utilities for each players are say $3, 3.5$ respectively. Each agent chooses between actions *minimum* which gives utility closer to its minimum utility or *excess* which gives it much more utility. In the following example the Nash equilibrium is *(minimum,excess),(excess,minimum)*.

	minimum	excess
minimum	(3,3)	(2,6)
excess	(6,2)	(1,1)

Although there is no pure strategy that satisfies minimum utilities of both agents, there are mixed strategies which can guarantee both agents their minimum, as can be seen in figure-1. If agents are ethical enough to care only for minimum

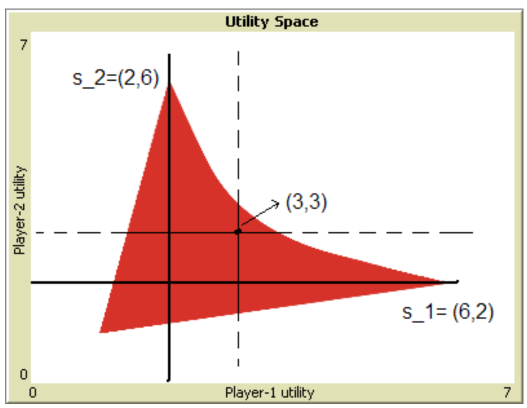

Fig. 1. The shaded region is the utility space for the example problem

each would be able to satisfy their minimum.

3 Solution Concept for Minimalistic Equilibrium

In this section we give a formal definition of the solution concept and analyze different cases that may arise.

From the example discussed above it is clear that we are interested in solution concepts which satisfy the agents minimum utility. If all the agents in the group adopt such a strategy the resultant solution would be the point which is closest to minimum utility vector and dominates it. Hence for multiple solution concepts that satisfy the minimum utility, the agents choose one which is closest to minimum utility.

3.1 Formalism

We consider the n player normal form game which is defined as a tuple (N, A, u, b) where N is set of players, such that $|N| = n$, $A = A_1 \times ...A_n$ where A_i is set of actions available to player i, and $u = (u_1, u_2...u_n)$ where $u_i : A \longmapsto R$ is a real valued utility function for player i. Additionally we define minimum utility vector for agent a_i as b_i. Let $b = b_1, b_2...b_n$ be the minimum utility vector for n agents. Let S denote the set of all available strategies. Let S' denote the set of pure or mixed strategies which dominate B. The minimalistic equilibrium is one which satisfy the following conditions.

- For each strategy, $S_i \in S'$ is a minimalistic equilibrium if

$$distance(S_i, b) < argmin_{S_{-m}} distance(S_{-m}, b)$$

  where $S_{-i} = S' - S_i$

Given that the solution concept lies in strategy space which dominates b, it might happen that it might not exist always. We study such case in the subsequent sections.
We consider the case with transferable and non-transferable utility and in each case we consider what strategy agents can adopt. Figure-2 illustrates the solution concept for 2-players. b_1, b_2 are minimum utilities for agent-1, agent-2 respectively. x_{min}, y_{min} is the solution concept.

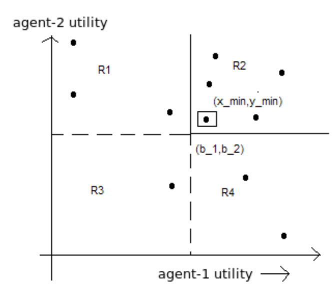

Fig. 2. Solution Concept

3.2 Transferable Utilities

Case-1 We first consider the case when there exists atleast one strategy s_j such that $\sum_i u_i(s_j) \geq \sum_i b_i$
In this case the agents choose to take actions such that after transferring utilities to each other, the resultant distance to minimum utility vector b is minimized. Consider 2-player environment. The solution strategy is the point closest to minimum (b_1, b_2). The following theorem gives the re-assignment of utilities which minimizes the distance.
Theorem-1 *For transferable utilities, let players $p_1, p_2 ... p_n$ with minimum utilities $(b_1, b_2, ..b_n)$ and utilities $(u_1, u_2, ...u_n)$ for given strategy, the resultant re-assignment of utilities so as to minimize the distance from b should be as follows.*

$$u'_i = b_i + (excess/n)$$

where

$$excess = \sum_i u_i - \sum_i b_i$$

and u'_i is new utilities for p_i
Resultant Utilities Proof Let new point after transferring utilities be $u' = (u'_1, u'_2 ... u'_n)$. Since, the agents only transfer utilities, we have $\sum_i u'_i = \sum_i u_i$ that is, net utility is same. Now, we need to

$$minimize\ (distance(u', b))$$

That is, to find a point on plane $u'_1 + u'_2 + ... + u'_n = C$ so that the point is closest to b. Using Lagrange multipliers,

$$F(u) = distance(u', B) - \lambda * (\sum_i u'_i - C)$$

Differentiating we get the required result.

$$u'_i = b_i + excess/n$$

where

$$excess = \sum_i u_i - \sum_i b_i$$

After the reassignment, the point with the least distance to B is the solution concept. Additionally as the following proof shows, it cannot happen that, after re-assignment, two distinct points,both in region R_2, are at the same distance.
Uniqueness Theorem Equally sharing the excess utility gives a unique solution concept.
Uniqueness Proof The proof follows from the previous theorem. Since a point on a plane such that its distance from a fixed point b, is always unique.

Case-2 In the above section we assumed that there is atleast one strategy which gives all the agents their required minimum utility. Here we consider the case when the dominant strategy space is empty. That is, minimum utilities of all agents cannot be satisfied simultaneously . The figure-3 illustrates the same. For all j $\sum_i u_i(s_j) < \sum_i b_i$.

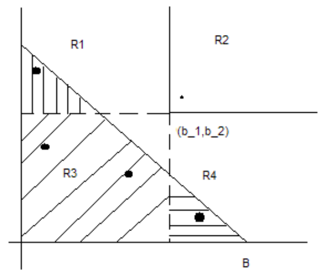

Fig. 3. Dominant strategy space is empty

If the dominant strategy space is empty, agents choose to form coalitions to ensure that a maximal subset of agents will get their minimum. It should be noted that when the dominant strategy space is empty agents compete and thus have an incentive to deviate from cooperative strategy.
We consider the n-player normal form game which is defined as a tuple $G = (N, A, u, b)$ where N is set of players, such that $|N| = n$, $A = A_1 \times ...A_n$ where A_i is set of actions available to player i, and $u = (u_1, u_2...u_n)$ where $u_i : A \longmapsto R$ is a real valued utility function for player i. We define the minimum utility for agent a_i as b_i. Let $b = b_1, b_2...b_n$ be the minimum utility vector for n agents. Let S denote the set of all available strategies.
Define game G' where $N_i \in N$ be the subset of agents which form a coalition to attain their minima and N_{-i} be the remaining players. Let s' be the Nash equilibrium for G'. The set N_i will form a coalition if in the game G' players $\forall p_j \in N_i$ Nash-equilibrium s' is such that $u(s'_j) \geq b_j$, where $u(s'_j)$ is the utility which player j gets in s'. That is, for all players in set N_i the utility they get in s' is at least equal to their minimum utilities. We illustrate it by the following 3-player game

Let the minimum of each player be $(3, 4, 5)$. It can be seen that none of the strategies satisfies the minimum of all players simultaneously. But, if player-1 and player-2 form a coalition the resultant game will ensure that they get their minimum utilities. As can be seen in the table 1.

The Nash equilibrium in the game G' in the above example guarantees p_1, p_2 their minima. Hence they have an incentive to form a coalition. This case can

u(1x1x1) = (4,3,4)	u(1x1x2) = (1,2,3)
u(1x2x1) = (3,1,6)	u(1x2x2) = (1,2,6)
u(2x1x1) = (5,1,3)	u(2x1x2) = (3,1,1)
u(2x2x1) = (2,2,1)	u(2x2x2) = (2,4,3)

Fig. 4. 3-player game

	a_1	a_2
a_1	**(7,4)**	(3,3)
a_2	(4,6)	(3,5)
a_3	(6,3)	(4,1)
a_4	(4,1)	**(6,3)**

Table 1. Game G' with p_1, p_2 forming a coalition

be extended to model the case when players have preferences over other players. Players can form subset with other players which they prefer.

3.3 Non-Transferable Utilities

Case-1 Consider the dominant strategy space is non-empty. Since the utilities are non-transferable the point closest to B would be a solution concept. If multiple points satisfy the condition, agents can choose randomly, since each player is guaranteed a minimum and they cannot transfer the excess utility

Case-2 Consider the case when dominant Strategy space is empty. Unlike the transferable utility case where agents could form subset so as to ensure that a maximal subset gets their minimum utility, here since agents cannot transfer utilities, they play a competitive strategy, Nash-equilibrium or other strategies.

4 Will agents reveal their minimum truthfully?

Here, we discuss the condition under which agents would reveal their minimum truthfully. We have discussed here results for two players game only.
Let (m_1, m_2) are minimum utilities for players p_1, p_2. As has been discussed in previous sections, the ethical solution lies in the utility space which dominates (m_1, m_2). We make following general assumptions in the discussion that follows.

1. Agents, once they reveal their minimum, will adopt ethical equilibrium. That is, no agent would adopt a maximizing strategy. The agents in the group can adopt costly punishment strategy for such a behavior.

2. Agents adopt a non-increasing cooperative function f_c. It implies that there is an upper limit to the minimum utility which one can demand from a game and which is always less than or equal to the maximum utility it can get from the game.
3. Agents do not know each others minimum utility apriori and they reveal them simultaneously.

Assume f_c is a step function for all players and agents prefer an ethical solution to any competitive solution concept. This could be because since we have considered a repeated game environment agents reputation might decrease if they fail to come to a cooperative solution. Then *players would scale up their minimum(lie), if by scaling up their minimum value, for any minimum utility value of other player, there is at least one point in utility space which dominates* (m_1, m_2).
Let $s_1 = (s_{11}, s_{12})$ be the strategy which gives maximum utility to p_1. s_{11} is utility which p_1 gets in s_1 and p_2 gets s_{12}. Let $s_2 = (s_{21}, s_{22})$ be the strategy which gives maximum utility to p_2. s_{21} is utility which p_1 gets in s_2 and p_2 gets s_{22}.
Theorem : If $m_1 \geq s_{21}$ *and* $m_2 \geq s_{12}$*, then agents would truthfully reveal their minimum.*
Proof: The theorem says that agents would be truthful about their minimum if the minimum utility which they want is greater than equal what they get if they allow the other player its maximum utility.
We assume p_2 is selfish. The condition $m_2 \geq s_{12}$ must be satisfied for p_2 to be

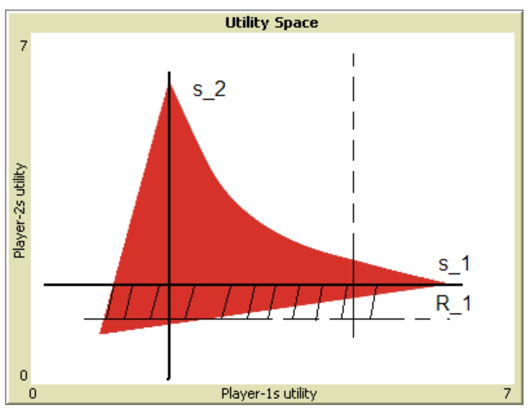

Fig. 5. $m_2 < s_{12}$ p_2 can scaleup its minimum to s_{12} ans still ensure a non empty dominant strategy space.

truthful. Consider the case when $m_2 < s_{12}$. Figure-5 illustrates the condition. Player-2 can scale-up its minimum utility from $m_2 = s_{12} - \alpha$ to s_{12} and still guarantee that for any value of minimum utility that player-1 sets, (m_1, m_2) is dominated by atleast one point(s_1). For $m_2 > s_{12}$ there always exists some minimum utility for player-1 which can make the dominating utility space empty. Moreover, as can be seen in figure-6, for $m_2 \geq s_{12}$, if player-2 scales up its minimum, the dominant strategy space decreases proportionately. This holds

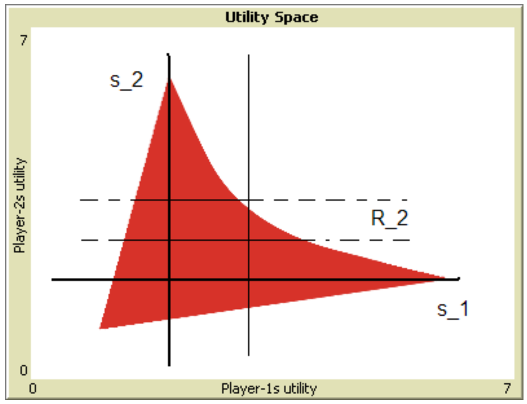

Fig. 6. $m_2 \geq s_{12}$ Dominating strategy space decreases as p2 scales up its minimum

for any utility space.
We can now define fair and unfair games as follows
A game is said to fair to both players if it allows both players to either be truthfull or scale-up their minimum utilities equally. Unfair to one of them if one can scaleup its minimum more than the other.
A game can be considered as zero-fairness game if there is a strategy which strictly dominates every other strategy, in which case none of the agents, if they are maximizing, would remain truthful.
We might want to relax the condition that agents prefer an ethical equilibrium to any competitive solution concept. A weak condition for players to be truthful is discussed here.
Let $V = (v_1, v_2)$ be the utility vector where v_i denotes the worst case payoff which player i gets if the ethical equilibrium is not reached. Then the condition for players to be truthful is given by the following theorem.
Theorem : If $m_1 \geq max(v_1, s_{21})$ *and* $m_2 \geq max(v_2, s_{12})$, *then agents would truthfully reveal their minimum.*
Proof: The proof follows from the previous theorem.

5 Summary

Some applications of the minimalistic games are
Resource usage: Resources whose excessive usage is not appreciated, players can adopt an ethical equilibrium concept. Players would reveal their minimum requirement and are allotted resources accordingly. Also from the above theorems it is clear that there is an upper limit above which an agent would not scale

up his minimum.
For example consider performance of nodes in a network. Each node can be assigned a minimum performance value below which it competes and above which it assists other nodes get their minimum performance. Such a system would implicitly guarantee that nodes cooperate with each other as well as compete with those which set their minimum performance value high.

This paper is an attempt to introduce a different class of agents. Agents which can make socially responsible decisions.With the introduction of virtual agents which interact on behalf or with humans the need to incorporate ethics in decision making by agents has increased. We would expect agents actions to reflect our preferences towards individuals. Minimalistic games ensure that such preferences can be captured even without having to quantify our preferences towards individuals. All that is required is ordinal preferences, and then such preferences will be handled by the case in which agents form subsets to ensure minima.
An indepth analysis of different cases still needs to be done. We could, instead of equally sharing the excess, associate a cost with transferring utilities and minimize the cost. A formal experimental analysis needs to be done, comparing the performance of minimalistic agents against maximizers, and altruistic agents. Our contention is that, minimalistic agents can be classified as 'evolutionary stable class' of agents, in that they maintain a right balance between their own utility and that of the system. Additionally, computational complexity associated with different cases needs to be studied.

References

1. von Neumann. J & Morgenstein. O, (1944) Theory of Games and economic behavior. Princeton University Press
2. Michael Anderson & Susan Leigh Anderson, The Status of machine ethics: A report from the AAAI symposium Minds and Machines(2007).
3. Colin Allen, Wendell Wallach & Iva Smit, Why Machine Ethics?(2006), IEEE Intelligent Systems.
4. Simon, H. A behavioral model of rational choice (1955), Quarterly Journal of Economics.
5. Wynn C. Stirling Satisficing Games and Decision making(2003), Cambridge University Press.
6. Joshua Letchford, Vincent Conitzer, Kamal Jain: An "Ethical" Game-Theoretic Solution Concept for Two-Player Perfect-Information Games, WINE 2008.
7. Bergstrom, Theodore C. and Oded Stark. 1993. How altruism can prevail in an evolutionary environment. The American Economic Review.