More than envy-free

Sandip Sen
Mathematical & Computer Sciences Department
University of Tulsa
sandip@kolkata.mcs.utulsa.edu

Abstract

There have been several interesting results in the game theory literature on dividing up goods between self-interested parties such that the allocation is envy-free (Brams and Taylor, 1996). An allocation is deemed envy-free when every party (agent) believes that its share is not less than anyone else's share. These procedures, however, are not efficient (in the sense of pareto optimality) in general. Envy procedures allow agents to ignore the utility metrics of other agents if they are satisfied with what can be called a fare share of the goods being divided. From the multiagent systems research perspective, however, we may be interested in studying augmentations of these procedures in which agents use models of the decision strategies or utility metrics of other agents to try to obtain more than their fare share. For example, it may be possible to improve the allocation to our modeling agent without decreasing the valuation of another agent if they trade things that one considers useless but is of value to the other agent. In particular, we are investigating the problem of dividing up a divisible good among two agents. The good can be discrete or be continuously divisible. We assume that one agent have a model of the utility function of the other agent. This model need not be accurate. We have adapted an envy-free division scheme for this problem to obtain a procedure by which the modeling agent can get more than its fare share of the allocation. The procedure also has the desired property of envy-freeness. So, even if the model being used is inaccurate, or both the agents have the same utility metrics, the allocation will still give the modeling agent at least its fare share of the good.

1 Introduction

Agents inhabit environments containing other agents. These agents interact in a variety of circumstances. Often agents conflict with other agents. Such conflicts are rooted in the goals and desires of agents. In a shared environment with limited resources, it is likely that there are not enough resources to provide for all goals of all agents. Just as an agent has to tradeoff its own goals under time constraints, multiple agents in such shared environments have to tradeoff goals with others because of resource constraints and goal conflicts. Artificial intelligence and related disciplines provide us techniques like planning, constraint satisfaction, decision analysis, etc. by which an agent can make reasoned tradeoffs under time constraints. Likewise, we need mechanisms and reasoning procedures by which two or more agents can tradeoff their goals under resource contentions and goal conflicts.

Negotiation is a process by which agents decide on mutually agreeable behaviors to resolve or avoid conflicts. Agents negotiate under a variety of information, time, and computational restrictions. A key research issue in agents and multiagent research is to develop negotiation procedures by which agents can efficiently and effectively negotiate solutions to their conflicts. Effectiveness requires that the outcome is fair, acceptable, or desirable to the parties involved in the negotiation process. Efficiency requires that the procedure is not excessively time-consuming or computing-intensive.

In this paper, we focus on the problem of agents vying for portions of a good to be divided up among them. The negotiation process will produce a partition and allocation of the goods among the agents. We will discuss both protocols by which agents interact and appropriate decision procedures to adopt given a particular procedure. Our discussion will be narrow in that our goal is not to enumerate all possible protocols and corresponding decision procedures. On the contrary, we analyze a particular protocol and suggest a class of effective decision procedures for the same.

2 Dividing a good

The problem of coalition formation by self-interested agents have received considerable attention among researchers in cooperative game theory. A key research issue in coalition formation is the distribution of the payoff received by a coalition among its constituent members. Centralized, fair payoff distribution schemes like the Shapley value metric has been proposed that require considerable amount of information about all possible coalitions. Multiagent systems researchers have also addressed the formation of coalition formation and payoff distribution among agents under different assumptions about computational resource and knowledge limitations.

In the following, we will discuss a related problem and some possible approaches to solving it. The problem is that of dividing up a good (this can be a payoff to a coalition, but not necessarily so) between multiple individuals. We will assume that the solution procedure is of a decentralized nature. This means that the agents will be required only to a protocol by which the division is to be made. They can freely choose any strategy to use to determine their actions within the accepted protocol.

For example, a protocol might be that every agent submits a sealed bid for the good. After every bid is collected, the good is divided up among the bidders in proportion to their bids. Our assumption is that once the agents agree to such a protocol, they are free to choose their bids following any strategies they adopt. We require, however, that agents will agree to the division of the good as specified once they have placed their bids.

From a designer's point of view, the choice of a protocol provides a platform for agents to negotiate an agreeable division. The choice of a strategy will be dictated by concerns for arriving at a preferred share of the good being divided. The protocol designer should then provide protocols which can be used by agents to successfully negotiate agreeable divisions with reasonable computational costs. We will now look at properties of divisions that can make them agreeable to self-interested agents.

3 Desired characteristics of divisions

As mentioned before, we assume that a single divisible good is to be divided among n agents. Any number of decision procedures can be adopted to make the division of the whole into parts and subsequently allocate the parts to different agents. We will be interested in evaluating both decision procedures or protocols which agents agree to abide by for making the division, as well as the strategies used by the agents to decide their negotiation behavior once the protocols have been decided.

The following criteria has been espoused by various researchers to be desirable characteristics of decision procedures or outcomes from such procedures:

Proportionality: Each agent believes that it received, by its own estimate, at least $\frac{1}{n}$ of the goods being allocated.

Envy-free: Each agent believes that it received, by its own estimate, at least as much as the share received by any other agent. This also implies that an agent has no incentive to trade its share with anyone else.

Equitable: A solution, i.e., a partition of the good among the n agents, is equitable, when the share received by each agent is identical in terms of their individual utility functions.

Efficiency: A solution,is said to be Pareto optimal or efficient if there is no other partition which will improve the perceived share of at least one agent without decreasing the perceived share of any other agent.

Here are some observations on the above-mentioned metrics:

- If all agents have identical perception or utility function, any partition is efficient. Also, in this case the only proportional solution is also envy-free and equitable.

- For two agents, proportionality implies envy-freeness. This implication, however, does not hold for arbitrary n. For example, an agent can believe it has got more than its fare share, but can still be envious of another agent which it believes has received an even larger share.

4 Improving on envy-freeness

Envy-freene procedures provide a very desirable performance guarantee. The solutions generated by such procedures guarantee that everyone considers itself a winner. While such guarantees are indeed extremely useful to bring parties to the discussion table, often a self-interested agent may want more. There is no reason a priori to justify why a self-interested agent should be happy with just the largest share in the group when it can possibly get even more. Put simply, a rational agent wants to maximize its utility, and if there is scope for cornering a larger share of the good being divided, even an envy-free procedure may not be satisficing!

For example, a procedure by which an agent can possibly improve on its share received by an envy-free procedure without losing the guarantee of envy-freeness, would be extremely attractive. The question then is what information needs to be known about others for an agent to corner a larger than envy-free share.

Though it can be practically possible for a rational agent to risk losing the envy-free guarantee in the hope of gaining substantial benefit most of the time, we will not pursue this line of investigation in the current paper. For the rest of the paper, we concentrate exclusively on dividing a good between two agents. We assume that the good being divided is possibly heterogeneous and the preference of an agent for various parts of the good is represented by a utility function. For historical reasons and to expedite the description of the procedures we have adapted from literature, we will represent the good as a rectangular piece of cake. Actually for all practical purposes, we are only interested in the length of the cake. In the following, and unless otherwise noted, we will assume the good to be continuously divisible.

We want to develop procedures that are resistant to spiteful agents. That is even if the agent being modeled realizes that the other agent is trying to exploit a model of its utility function, it cannot force a division that hurts itself but achieves the goal of making the modeling agent envious too. To illustrate a procedure without such a guarantee consider the "divide and choose" procedure, in which one agent cuts the cake into two halfs and the other agent gets to choose which half it wants for itself. In the non-modeling case, the basic strategy of the cutting agent will be to divide the cake into two equal halfs, so that it is envy-free irrespective of which half the other agent chooses. The choosing agent chooses whichever half it considers is larger, or at least as large as the other one, and thus is envy-free too. The choosing agent, however, may believe that the cutting agent is using a model of its utility function to make an unequal cut from the cutting agent's perspective and is hoping to pocket larger than its estimate of half of the cake. A spiteful choosing agent may choose to hurt both itself and the cutting agent by choosing the part it considers smaller by its estimate and is bigger by the estimate of the cutting agent.

The utility to the ith agent of a piece of the cake cut between points a and b (where $a < b$) is given by $\int_a^b U_i(x)dx$, where $U_i(x)$ is the utility function of the ith agent. Without loss of generality, we will assume that the first agent is having a model, $\bar{U}_2$, of the utility function of the second agent. This model need not be accurate, but the better the model, the more benefit the modeling agent stands to gain by using it.

5 An envy-free procedure

Division procedures will, in general, have two components: a rule set to be followed in arriving at an allocation, and strategies to be used by the players to obtain fair shares. For example, in the above-mentioned "divide and choose" scheme, the procedure is that one player divides the cake, and the other gets to choose the part it wants. To guarantee envy-freeness, the divider must divide the cake into half by its estimates, and the chooser chooses the part that it considers to be at least as big as the other part.

We first narrate Austin's moving knife procedure for arriving at an envy-free division of a cake. A referee moves a knife to the right starting from the left edge of the cake and always keeps the knife parallel to that edge (see Figure 1. At any time, either player can call "cut". When one player makes such a call, it gets the portion of the cake to the left of the knife, the other part going to the other player. To guarantee envy-freeness, a player has

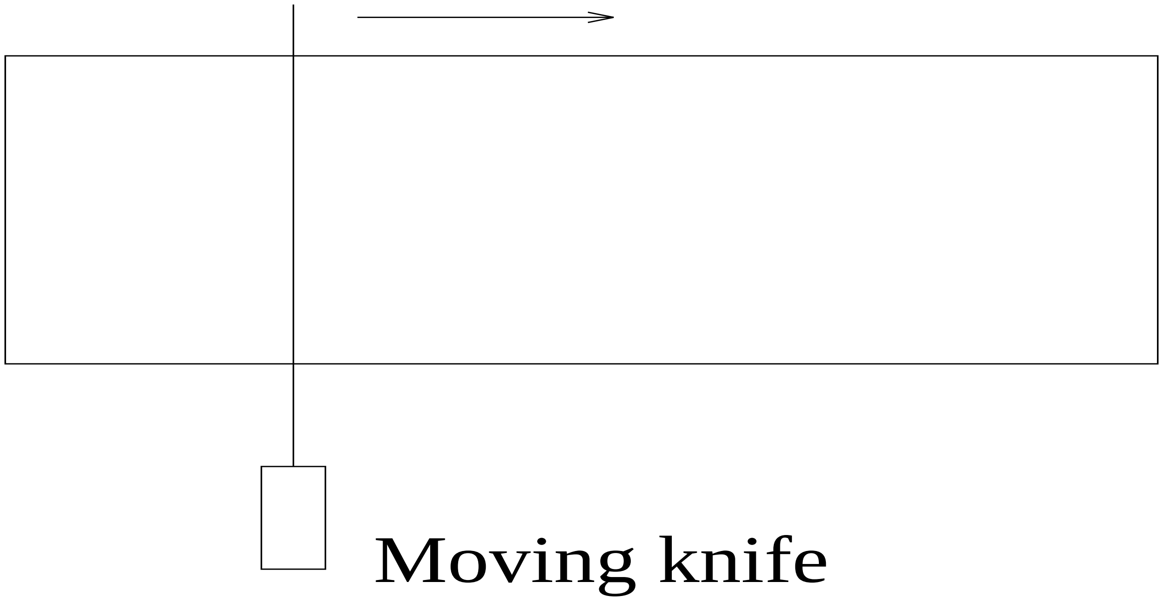

Figure 1: Austin's moving knife procedure.

to call "cut" if the portion of the cake to the left of the knife becomes at least as large as half of the cake in its estimate. If this player waits any longer, it risks losing more than half of the cake if the other player calls cut! Note that the above-mentioned strategy does leave open the possibility that the player calling cut will think that the other player, say B, also thinks that it, i.e., B, got more than half of the cake by B's estimate. This, however, should not lead to envy. Envy-freeness only requires that each player thinks that it got at least the largest share by its own estimate. This leaves open the possibility that each player thinks it got more than everyone else. This will be true except in the special cases where one or both players think that they got exactly half of the cake.

Now we analyze this procedure for efficiency. Suppose that agent A calls cut and gets the portion to the left of the cake. It might be the case that there is a portion immediately to the right of the knife that has no utility in the view of agent B, but is of value to agent A. Let us call this region R. Now, region R can be transferred to agent A, thereby increasing its utility. At the same time this transfer does not reduce the utility of agent B. Thus the strategy to guarantee envy-free division in the Austin's moving knife procedure does not guarantee efficiency. Note, however, that if no part of the cake has zero utility to any of the agent's, Austin's procedure results in an efficient allocation.

6 Using information to obtain a larger share

We allow agent A to use an approximate model, $\hat{U_B}$, of the utility function of the agent B. The actual utility functions of the two agents are U_A and U_B respectively. In our procedure, the modeling agent start by holding two knives parallel to the side edges of the cake and then move them to the right allowing for wrap around. Let the positions of the left and right knives at time t be l_t and r_t respectively. At any point in time, agent B can call cut, at which point it receives the portion of the cake between the two knives (see Figure 2. If

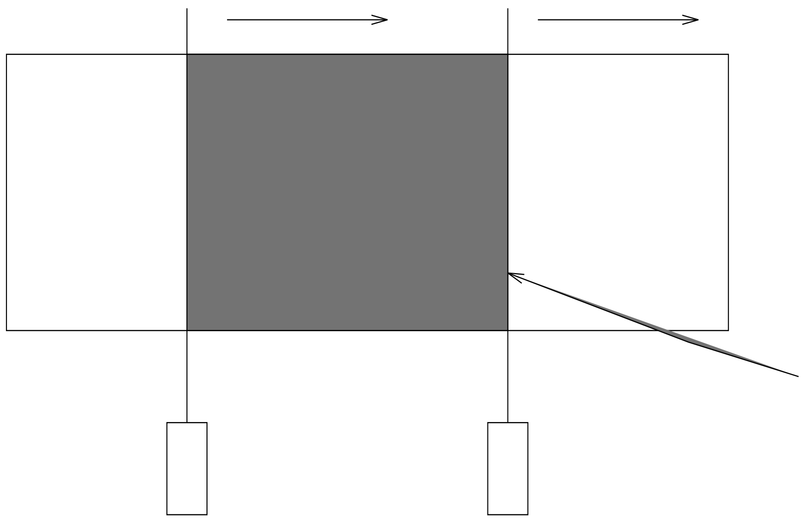

Figure 2: Our aumentation of Austin's procedure with the modeling agent moving two knives.

when B calls cut, $r_t < l_t$, i.e., the right knife is to the left of the left knive (this can happen since we allow wrap around), then B gets both the portions between the left knife and the right edge and from the left edge to the right knife. An additional requirement is that if B does not call cut, the knives stop when the right knife reaches the original position of the left knife, and the left knife reaches the original position of the left knife. None of the knives are allowed to move to the original positions of the other knife. When the knives finally stop moving, the portion between the left and right knives are given to B with the rest of the cake going to A.

Given the above rules, what should be the strategy of the modeling agent? Its goal is to give B just enough to make it envy-free. Intuitively, A should try to give B half of the cake according to B's estimate without giving up the parts that it values itself. More formally, A would have B call cut at time t when the following expression is maximized:

$$\int_{l_t}^{r_t} (\hat{U_B}(x) - U_A(x))dx. \tag{1}$$

On first thought, then, A should start with the left and right knives at the positions where the above integral is maximized. But this will not guarantee envy-freeness for A. To see why this is the case let m and n be the initial positions of the left and right knives. Now, it may be the case, that B calls cut immediately or never calls cut. In the first case, A gets a part of the cake of utility $\int_n^m U_A(x)dx$; in the second case, A gets a part of the cake of utility $\int_m^n U_A(x)dx$. Since these two parts are complementary, it follows that to guarantee envy-freeness, the initial knife placements should be such that they divide the cake exactly in half by A's estimate. So, the initial knife placements should be such that

$$\int_n^m U_A(x)dx = \frac{1}{2}\int_0^L U_A(x)dx,$$

where L is the length of the cake.

It is also required that the initial knife placements such that the region in between does divide the cake in half for agent B. Otherwise it will call cut immediately or may get more than proportional share by waiting untl the end. In either case, A will not be able to exploit its model of B's utility function and would have to be satisfied with just half of the cake by its estimate.

So, the initial placement of the knives should be such that in addition to the above constraint it should also satisfy the following inequality:

$$\int_n^m U_B(x)dx = \frac{1}{2}\int_0^L U_B(x)dx.$$

[1] Since A has only an approximate model of B's utility function, the best it can do is to choose the initial placements such that

$$\int_n^m \hat{U_B}(x)dx = \frac{1}{2}\int_0^L \hat{U_B}(x)dx.$$

Then B is guaranteed to get half the cake at the beginning or the end, and can be tempted to see the intermediate offerings. The latter is precisely what A hopes for.

The knives should then be moved in a manner such that the left and right knives will be placed any pair of points p and q points (p, q) that optimizes the expression in Equation 1, and satisfies both of the following inequalities:

$$\int_p^q U_A(x)dx < \frac{1}{2}\int_0^L U_A(x)dx,$$

$$\int_p^q \hat{U_B}(x)dx \geq \frac{1}{2}\int_0^L \hat{U_B}(x)dx.$$

The first of the above says that A will be envy-free in giving up the region between p and q, and the second inequality says that B will be envy-free in receiving that region. Put together, these three inequalities identify regions which A will most prefer to give up assuming that B will not accept a non-envy-free share.

While moving between pairs of points as identified in the previous paragraph, the knife locations (u, v) should satisfy the following inequality:

$$\int_u^v \hat{U_B}(x)dx < \frac{1}{2}\int_0^L \hat{U_B}(x)dx.$$

That is, for regions where A's utility is high, it will reduce the spacing between the knives to make that region non-envy-free for B, thereby preventing B from calling "cut".

A lurking concern in the mind of the reader can be the question of the effect of an inaccurate model used by A. By enforcing the constraint that at no time during the procedure the region between the two knives is greater than half the cake's worth in A's estimate, we have completely eliminated the possibility of a division which is non-envy-free for A. The

[1] We do not require $\int_0^L U_A(x)dx = \int_0^L U_B(x)dx$, i.e., the valuations of the entire cake need not be identical for both agents.

worst case scenario is that A receives exactly half of the cake in its estimate. Interestingly, this would also be the result if both agents have identical utility functions irrespective of whether A has an accurate model of B's utility function.

At this point, we have completed a coarse description of a procedure to be followed by our modeling agent to exploit its model of the other agent. Another way of looking at the above description is a characterization of a class of strategies that guarantees envy-freeness but can provide a significantly larger share of the cake than that obtained by Austin's moving knife procedure.

We have also developed particular instances of this class of strategies. This requires solving a function optimization problem. The function in question is a 2-d function, $f(x, y)$, with the variables being the positions of the left and right knives and is defined as

$$f(x, y) = \int_x^y (\hat{U_B}(z) - U_A(z))dz.$$

The procedure is as follows. Say (p, q) optimizes this function. A region of the cake to the left of p that is an exact division of the cake from both agents estimated viewpoint is chosen for initial knife placement. Such a region can always be found for finite-valued utility functions. The knives are then moved such that until the right knive coincides with p, the region between the two knives constitute a non-envy-free allocation to B. At this point the left knive is moved to p without moving the right knife. So, after some time, both knives are co-located at p. Following this, the right knife is moved without moving the left knife. When the right knife reaches q, the region between the two knives would constitute an envy-free allocation to B. If B calls cut, A's purpose have been served, i.e., it receives the optimal envy-free allocation under its estimate of B's utility metric. If B does not call cut, the knives are moved so that the region outside them constitute an envy-free allocation for A. The other constraint is that, if B never calls cut, the final position of the left and right knives mirror their initial positions. At this point, allocations are made, and A receives a proportional and envy-free allocation.

7 Discussion

We have presented a decentralized protocol and a class of decision strategies to be used by two agents to divide up a good between them. The goal of this work is to provide a mechanism by which an agent can use its knowledge about the preferences and utilities of the other agent to obtain a larger than proportional share of the good. The mechanism should still guarantee envy-freeness for both agents. In the case of two agents, this would mean that each agent believes that it got at least half of the good by its own estimate. As a result, neither agent would be ready to trade its allocation with the other agent.

We are looking to extend this analysis to more than two agents. In particular, the next step would be to augment the procedure for the case of three agents. It is unlikely, however, that this moving knife procedure can be extended for the case of an arbitrary number of agents. For the general case of n agents, we would be interested in augmenting approximately envy-free division procedures, where allocations are envy-free within some pre-specified error-bounds.

Reference

- Fair Division: From cake-cutting to dispute resolution. Steven J. Brams and Alan D. Taylor. Cambridge University Press, Cambridge: UK, 1996.

  The book contains a sequence of envy-free and proportional division procedures of both discrete and continuously divisible goods between different number of parties. Properties of these allocation procedures analyzed include: envy-freeness, equitability, efficiency, manipulability, etc.