Role Assignment for Game-Theoretic Cooperation

(Extended Abstract)

Catherine Moon
Duke University
Durham, NC, USA
csm17@duke.edu

Vincent Conitzer
Duke University
Durham, NC, USA
conitzer@cs.duke.edu

ABSTRACT

In multiagent systems, often agents need to be assigned to different *roles*. Multiple aspects should be taken into account for this, such as agents' skills and constraints posed by existing assignments. In this paper, we focus on another aspect: when the agents are self-interested, careful role assignment is necessary to make cooperative behavior an equilibrium of the repeated game. We formalize this problem and provide an easy-to-check necessary and sufficient condition for a given role assignment to induce cooperation. However, we show that finding whether such a role assignment exists is in general NP-hard. Nevertheless, we give two algorithms for solving the problem. The first is based on a mixed-integer linear program formulation. The second is based on a dynamic program, and runs in pseudopolynomial time if the number of agents is constant. Minor modifications of these algorithms also allow for determination of the minimal subsidy necessary to induce cooperation. In our experiments, the IP performs much, much faster.

Keywords

role assignment, cooperation, repeated games

1. INTRODUCTION

In this paper, we study the *game-theoretic* ramifications in role assignment when agents are self-interested. A careful assignment of roles might induce cooperation whereas a careless assignment may result in incentives for an agent to defect. Specifically, we consider a setting where there are multiple *minigames* in which agents need to be assigned roles. These games are then infinitely repeated, and roles cannot be reassigned later on. It is well known, via the folk theorem, that sometimes cooperation can be sustained in infinitely repeated games due to the threat of future punishment. Nevertheless, some infinitely repeated games, in and of themselves, do not offer sufficient opportunity to punish certain players for misbehaving. If so, cooperation may still be attained by the threat of punishing the defecting agent in *another* (mini)game. But for this to be effective, the defecting agent needs to be assigned the right role in the other minigame.

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

Our work contrasts with much work in game theory in which the model zooms in on a single setting without considering its broader strategic context. In such models, firms make production and investment decisions based on competition in a single market; teammates decide on how much effort to put in on a single project; and countries decide whether to abide by an agreement on, for instance, reducing pollution. In reality, however, it is rare to have an isolated problem at hand, as the same agents interact with each other in other settings as well. Firms often compete in several markets (e.g., computers and phones); members of a team usually work on several projects simultaneously; and countries interact with each other in other contexts, say trade agreements, as well.

Looking at a problem in such an isolated manner can be limiting. There are games where a player has insufficient incentive to play the "cooperative" action, as the payoff from that action and the threat of punishment for defection are not high enough. In such scenarios, putting two or more games with compensating asymmetries can leave hope for cooperation. A firm may allow another firm to dominate one market in return for dominance in another; a team member may agree to take on an undesirable task on one project in return for a desirable one on another; and a country may agree to a severe emissions-reducing role in one agreement in return for being given a desirable role in a trade agreement.

2. MOTIVATING EXAMPLE

Consider two individuals (e.g., faculty members or board members) who together are to form two distinct committees. Each of the committees needs a chair and another member; these are the roles we need to assign to the two individuals. Each committee's chair can choose to behave selfishly or cooperatively. Each committee's other member can choose to sabotage the committee or be cooperative. The precise payoffs differ slightly across the two committees because of their different duties. (For example, acting selfishly as the chair of a graduate admissions committee is likely to lead to different payoffs than acting selfishly as the chair of a faculty search committee.)

Chair \ Member	sabotage	cooperate
selfish	2, 1	3, 0
cooperate	1, 1	2, 2

Chair \ Member	sabotage	cooperate
selfish	2, 1	4, 0
cooperate	0, 1	2, 3

Let us first consider each of these two minigames separately. If the minigame is only played once, the chair has a strictly dominant strategy of playing selfishly (and hence, by iterated dominance, the other member will sabotage the committee). Even if the game is repeated (with a discount factor $\delta < 1$), we cannot sustain the (cooperate, cooperate) outcome forever. This is because the chair would receive a payoff of 2 in each round from this outcome—but defecting to playing selfishly would give her an immediate utility of 3 or 4 in that round after which she can still guarantee herself a utility of at least 2 in each remaining round by playing selfishly.

Now let us consider the minigames together. If the same agent is assigned as chair in each minigame, again we could not sustain the (cooperate, cooperate) outcome in both minigames, because the chair would gain $3+4$ immediately from defecting and still be able to obtain $2+2$ in each round forever after. On the other hand, if each agent is chair in one game, then with a reasonably high discount factor, (cooperate, cooperate) can be sustained. For suppose the chair of the second committee deviates by acting selfishly in that committee. This will give her an immediate gain of $4-2=2$. However, the other agent can respond by playing selfishly on committee 1 and sabotaging committee 2 forever after. Hence in each later round the original defector can get only $1+2=3$ instead of the $2+2=4$ from both agents cooperating, resulting in a loss of 1 in each round relative to cooperation. Hence, if δ is such that $2 \leq \delta/(1-\delta)$, the defection does not benefit her in the long run. This shows that linking the minigames allows us to attain cooperative behavior where this would not have been possible in each individual minigame separately. It also illustrates the importance of assigning the roles carefully in order to attain cooperation.

One may wonder what would happen if we link the minigames in a single-shot (i.e., not repeated) context. This would correspond to the case $\delta = 0$, so that the above formula indicates that cooperation is not attained in this case. In fact, linking minigames cannot help in single-shot games in general: in a single-shot model, any equilibrium of the (linked) game must consist simply of playing an equilibrium of each individual minigame. (Otherwise, a player could improve her overall payoff by deviating in a minigame where she is not best-responding.) Linking becomes useful only when the game is repeated, because then one's actions in one minigame can affect one's future payoffs in other minigames, by affecting other players' future actions. This is why the repeated game aspect is essential to our model.

3. THEORETICAL OUTLINE

We are interested in how to assign roles within minigames (of which there are two in the example above, corresponding to the two committees) to assess whether a particular outcome can be sustained in repeated play. For this question, the key issue is which roles (from different minigames) are bundled together, rather than which particular agent is assigned this bundle of roles.

Using the ideas behind the folk theorem, we can analyze whether our problem has a solution or not. We show that whether it does comes down to a single number per minigame role. The intuition that allows us to show this is as follows. To determine whether a given agent i will defect (i.e., play something other than the target action in some role assigned to her), by the folk theorem, we may assume that all other agents will play their target actions until some defection has taken place, after which they maximally punish agent i (in *all* minigames, not just the ones in which she defected). Thus, in the round in which agent i defects, she may as well play the single-round best-response to the target actions in every role assigned to her; afterwards, she will forever receive the best she can do in response to maximal punishment. (Since we only consider Nash equilibrium, we do not have to worry about multiple agents deviating.) The net effect of the defection on i's utility may be positive or negative for any given role; whether i will defect depends solely on the sum of these effects.

Because it is straightforward to compute the net effect of the defection on agent utilities, the problem of *finding* a role assignment reduces to a computational problem. We show that this problem is weakly NP-complete even in an extremely restricted special case, and strongly NP-complete for n player games. The problem can be solved in pseudopolynomial time using dynamic programming when there are at most a constant number of agents, though an integer programming approach performs much better empirically.

4. CONCLUSION

We believe that there are many other important directions that can be studied in the context of game-theoretic role assignment. Our model can be extended to allow (perhaps costly) reassignment of roles as time progresses; different agent types that value roles differently, and preferences not only over roles but also over which type of agent one is matched with (providing connections to matching [3] and hedonic games [1]); side payments between agents (providing connections to matching with contracts [2]); not every minigame being played in each round; generalizing from repeated games to stochastic or arbitrary extensive-form games; and so on. We believe that our paper provides a good foundation for such follow-up work.

Acknowledgments

We are thankful for support from ARO under grants W911NF-12-1-0550 and W911NF-11-1-0332, NSF under awards IIS-1527434, IIS-0953756, CCF-1101659, and CCF-1337215, and a Guggenheim Fellowship. This work was done in part while Conitzer was visiting the Simons Institute for the Theory of Computing. We thank the anonymous referees for their comments and suggestions.

REFERENCES

[1] H. Aziz and R. Savani. Hedonic games. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 15. Cambridge University Press, 2015.

[2] J. W. Hatfield and P. R. Milgrom. Matching with contracts. *American Economic Review*, 95(4):913–935, September 2005.

[3] B. Klaus, D. Manlove, and F. Rossi. Matching under preferences. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 14. Cambridge University Press, 2015.