

Role Assignment for Game-Theoretic Cooperation

Catherine Moon
Duke University
Durham, NC 27708, USA
csm17@duke.edu

Vincent Conitzer
Duke University
Durham, NC 27708, 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.

1 Introduction

Role assignment is an important problem in the design of multiagent systems. When multiple agents come together to execute a plan, there is generally a natural set of roles to which the agents need to be assigned. There are, of course, many aspects to take into account in such role assignment. It may be impossible to assign certain combinations of roles to the same agent, for example due to resource constraints. Some agents may be more skilled at a given role than others.

In this paper, we assume agents are interchangeable and instead consider another aspect: if the agents are self-interested, then the assignment of roles has certain *game-theoretic* ramifications. 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.¹ 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. This paper studies this game-theoretic role-assignment problem.

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 generally interact with each other in other settings as well. Firms often compete in several markets (e.g., on computers and phones, cameras, and displays); 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.

In this paper, we first formalize this setup. Subsequently, we give useful necessary and sufficient conditions for a role assignment to sustain cooperation. We then consider the computational problem of finding a role assignment satisfying

texts, such reassignment is infeasible or prohibitively costly due to agents having built up personalized infrastructure or specialized expertise for their roles, as is easily seen in some of the examples in the next paragraph.

¹Our model disallows reassigning agents because in many con-

these conditions. We show that this problem is NP-hard. We then give two algorithms for solving the problem nonetheless, both of which can be modified to find the minimal subsidy necessary to induce cooperation as well. One relies on an integer program formulation; the other relies on a dynamic programming formulation. We show the latter solves the problem in pseudopolynomial time when the number of agents is constant. However, in our experiments, the former algorithm is significantly faster, as shown at the end of our paper.

2 Background: Repeated Games and the Folk Theorem

When considering systems of multiple self-interested agents, behavior cannot be directly imposed on the agents. However, incentives for desirable behavior can be created by means of rewards and punishments. In one-shot games, there are no opportunities to reward cooperation or to punish defection. On the other hand, when the game is repeated and agents interact repeatedly over time, richer outcomes could arise in equilibrium. Intuitively, when there is potential for future interactions (given that agents care about future outcomes, that is, the discount factor δ is positive), outcomes not sustainable in the one-shot version can be attained given the right reward and punishment strategies.

The *folk theorem* characterizes equilibrium payoffs that can be obtained in infinitely repeated games, as agents become arbitrarily patient (see, e.g., [Fudenberg and Tirole, 1991]). The focus is on infinitely repeated games with discount factor δ (or, equivalently, repeated games that end after each period with probability $1 - \delta$). The possibility of future rewards and punishments makes certain outcomes besides the static Nash equilibria of the stage game sustainable, if the agents care “enough” about future payoffs. To characterize which outcomes are sustainable in equilibrium, an agent’s minimax payoff—the payoff that the other agents can guarantee she gets at most, if they set out to minimize her utility—is key. There can be no equilibrium where an agent receives less than her minimax payoff, because the agent could then deviate and receive more. The folk theorem states that basically all feasible payoff vectors that Pareto dominate the minimax payoff vector can be attained in equilibrium with the threat of sufficiently harsh punishment for deviating from the intended behavior.

Based on the ideas behind the folk theorem, we develop a characterization theorem (Theorem 1) that will serve as the foundation for analyzing our problem of bundling roles within (mini)games for game-theoretic cooperation. This theorem provides an easy-to-check condition for whether cooperation can be attained under a given role assignment.

3 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.)

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

Committee 1

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

Committee 2

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 mini-

games 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.

4 Definitions

When there is a risk of confusion, we distinguish between *minigames* (of which there are two in the example above, corresponding to the two committees) and the larger *metagame* that the minigames together constitute after roles have been assigned to agents. Note that technically the players in minigames are roles, not agents, and the metagame among the agents is not defined until roles have been assigned to them. Let $N = \{1, \dots, n\}$ be the set of agents and G the set of minigames. We assume each minigame has n roles (such as committee chair or committee member); if this is not the case, we can simply add dummy roles. For each minigame $g \in G$ and each role r in g , there is a set of actions A_r^g for the agent in that role to play, and a utility function u_r^g that takes as input a profile of actions \vec{a}^g in g , consisting of one action a_r^g for each role r in g , and as output returns the utility of the agent in role r from this minigame. An agent i ’s total payoff in round t is $\sum_{g \in G} u_{r(i,g)}^g(\vec{a}^g(t))$, where $r(i, g)$ denotes the role assigned to agent i in minigame g and $\vec{a}^g(t)$ is the profile played in minigame g in round t . For the repeated game, we consider both discounted payoff ($\sum_{t=0}^{\infty} \delta^t \sum_{g \in G} u_{r(i,g)}^g(\vec{a}^g(t))$) and limit average payoff ($\liminf_{T \rightarrow \infty} (1/T) \sum_{t=0}^{T-1} \sum_{g \in G} u_{r(i,g)}^g(\vec{a}^g(t))$). We assume perfect monitoring, i.e., agents observe all actions taken by all agents in prior rounds. All this is common knowledge, and so is the role assignment $r(i, g)$ once the agents play the game.

Our main interest is in assessing whether a particular outcome can be sustained in repeated play. We assume that for each game g , for each role r in g , there is a distinguished action $a_r^{*g} \in A_r^g$ that we call the *target* action. (There is no requirement that this action should be “cooperative” in the sense of increasing other agents’ utilities or maximizing the social utility; it can be any action, for example one that a principal assigning the roles would like to see happen for exogenous reasons. Alternatively, one can think of the principal having an extremely large utility for the target action being played, and including the principal’s utility in the social welfare.) Our main question is whether there exists a role assignment function $r(i, g)$ such that there is an equilibrium of the repeated game where every agent always plays the target action in every role assigned to her. 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. We postpone the definition of this question as a computational problem until we have done some further simplifying analysis.

5 Related Literature

The assignment of roles in multiagent systems has of course received previous attention, especially in domains such as RoboCup soccer. However, we are not aware of any multiagent systems literature on assigning roles across multiple games in a way to achieve game-theoretically stable cooperation. The work by Grossi and Turrini (2012) that combines the concepts of dependence theory and game theory is distantly related to our work, focusing on identifying coalitions such that agents mutually benefit within this coalition.

In the economics literature, some of the first work to recognize the effect of playing multiple games in parallel took place in research on industrial behavior. In 1955, Corwin Edwards proposed the possibility that multimarket contact between firms could allow them to reach strategically stable arrangements that could not be reached in a single market and thereby foster anticompetitive outcomes. Subsequent theoretical works explore the effect of multimarket contact on economic performance (e.g., [Bulow *et al.*, 1985]) and the degree of cooperation sustainable with repeated competition (e.g., [Bernheim and Whinston, 1990]).

Papers by Folmer and von Mouche (2007) and Just and Netanyahu (2000) concern how the structure of the linked component games affect the potential for cooperation. Both identify the expansion in the bargaining set through linkage as the key to potential increase in cooperation. The latter work further examines linking common game classes such as the prisoner’s dilemma, assurance, iterated dominance, and chicken games; they do not observe the chances of coming to a fully cooperative equilibrium increasing significantly except for the case of linking prisoner’s dilemma games.

The intuition that linking games relaxes incentive constraints and the results on what game structure allows linkage to lead to cooperation are both in accordance with the results by Jackson and Sonnenschein (2007). These authors formally address the relaxation of incentive constraints through linking decision problems and the resulting efficiency gains in the general context of social choice with preference announcements.

6 Theoretical Analysis

In this section, we use the ideas behind the folk theorem to 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* games, 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 for-

ever 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.) Target actions are limited to pure strategies (since agents cannot observe whether a specific mixed strategy was played), but we allow for mixed (even correlated, as we discuss below) actions in the punishment phase. 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. We now formalize this. Recall that a^{*g} is the profile of target actions for minigame g .

Definition 1. Given a minigame g and a role r in g , let c_r^g denote the (long-run) **cooperation value** for that role when the target actions are played by all players. With limit average payoffs, $c_r^g = u_r^g(\bar{a}^{*g})$. With discounted payoffs, $c_r^g = \sum_{t=0}^{\infty} \delta^t u_r^g(\bar{a}^{*g}) = u_r^g(\bar{a}^{*g})/(1 - \delta)$.

Next, we want to specify the defection value. This requires us to know what utility a player will get in rounds after defection, which depends on how effective the other players are in punishing. In a two-player game, the punishing player should play a minimax strategy—i.e., play as if she were playing a zero-sum game where her utility is the negative of that of the defecting player. With three or more players, an important question is whether the players other than the defector can coordinate (i.e., correlate) their strategies.² If not, this leads to NP-hardness [Borgs *et al.*, 2010]. Therefore, we assume that they can correlate, which allows polynomial-time computability [Kontogiannis and Spirakis, 2008]. Formally, when the player in role r has defected, the remaining players $-r$ will play $\arg \min_{\sigma_{-r}^g} \max_{a_r^g} u_r^g(a_r^g, \sigma_{-r}^g)$, where σ_{-r}^g is a mixed strategy for the players $-r$ (allowed to be correlated if there are two or more players in $-r$). In the case of discounted payoffs, we also need to know the utility a player will get in the first round she defects; this is $\max_{a_r^g} u_r^g(a_r^g, a_{-r}^{*g})$.

Definition 2. Given a minigame g and a role r in g , let d_r^g denote the (long-run) **defection value** for that role. With limit average payoffs, $d_r^g = \min_{\sigma_{-r}^g} \max_{a_r^g} u_r^g(a_r^g, \sigma_{-r}^g)$. With discounted payoffs, $d_r^g = \max_{a_r^g} u_r^g(a_r^g, a_{-r}^{*g}) + \sum_{t=1}^{\infty} \delta^t \min_{\sigma_{-r}^g} \max_{a_r^g} u_r^g(a_r^g, \sigma_{-r}^g) = \max_{a_r^g} u_r^g(a_r^g, a_{-r}^{*g}) + \frac{\delta}{1-\delta} \min_{\sigma_{-r}^g} \max_{a_r^g} u_r^g(a_r^g, \sigma_{-r}^g)$.

In the end, what matters is the net effect of defection.

Definition 3. Given a minigame g and a role r in g , let $m_r^g = c_r^g - d_r^g$ denote the **robustness measure** for that role.

²For 3+ players, it strengthens our hardness result (to be discussed in Section 7) that it holds even with correlated punishment. Computational issues aside, without correlated punishment there are some (known) conceptual problems because the minimax theorem fails with 2 players against 1. Briefly, consider a variant of rock-paper-scissors where player 3 just plays R, P, or S, but players 1 and 2, who play together in a sense that they have identical payoffs, both pick from $\{0, 1\}$; 00 means rock, 01 means paper, 10 means scissors, and 11 means a fourth action (fragile) that always loses. If 1 and 2 can correlate, they can play the regular RPS minimax strategy guaranteeing payoff 0. But if they cannot, there is no joint strategy for 1 and 2 guaranteeing 0. On the other hand, the best that player 3 can guarantee is 0. So there is no well-defined value of the game and it is arguably less clear what will happen.

The theorem now says that we can obtain the desired outcome if and only if there is an assignment that gives each agent a nonnegative sum of robustness measures.

Theorem 1. The repeated metagame with assignment $r(i, g)$ has an equilibrium (allowing correlated punishment) in which on the path of play, in every round t , in every minigame g , every role r plays a_r^{*g} if and only if the assignment $r(i, g)$ is such that for all i , $\sum_g m_{r(i, g)}^g \geq 0$.

Proof. We first prove the “if” direction, supposing that $\sum_g m_{r(i, g)}^g \geq 0$ for all i . Consider a grim trigger strategy profile where all players cooperate (play a_r^{*g}) in every role r to which they have been assigned as long as everyone else does so; if some player (say i) has deviated, the other players $-i$ switch to maximally punishing i via correlated punishment—that is, in every game g , they play a strategy in $\arg \min_{\sigma_{-r(i, g)}^g} \max_{a_r^g} u_r^g(a_r^g, \sigma_{-r(i, g)}^g)$. We must show this strategy profile is an equilibrium. Consider an arbitrary agent i . Not deviating will give i a long-term utility of $\sum_g c_{r(i, g)}^g$. What about deviating? Without loss of generality, suppose i deviates in the first round. The highest expected payoff i can obtain in that first round is $\sum_g \max_{a_r^g} u_r^g(a_r^g, a_{-r(i, g)}^{*g})$; in every remaining round, she can obtain at most $\sum_g \min_{\sigma_{-r(i, g)}^g} \max_{a_r^g} u_r^g(a_r^g, \sigma_{-r(i, g)}^g)$. Hence, her long-term utility is at most $\sum_g d_{r(i, g)}^g$. But by assumption, $\sum_g m_{r(i, g)}^g \geq 0$, which is equivalent to $\sum_g c_{r(i, g)}^g \geq \sum_g d_{r(i, g)}^g$. So agent i has no incentive to deviate.

We now prove the “only if” direction, supposing that $\sum_g m_{r(i, g)}^g < 0$ for some i . Consider a strategy profile where on the path of play all players cooperate (play a_r^{*g}). Again, not deviating will give player i a long-term utility of $\sum_g c_{r(i, g)}^g$. By the minimax theorem, $\min_{\sigma_{-r(i, g)}^g} \max_{a_r^g} u_r^g(a_r^g, \sigma_{-r(i, g)}^g) = \max_{\sigma_{r(i, g)}^g} \min_{a_{-r(i, g)}^g} u_r^g(a_{-r(i, g)}^g, \sigma_{r(i, g)}^g)$. Consider the deviating strategy where player i plays, in every minigame g , a strategy from $\arg \max_{\sigma_{r(i, g)}^g} u_r^g(a_{-r(i, g)}^g, \sigma_{r(i, g)}^g)$ in the first round, and then a mixed strategy from $\arg \max_{\sigma_{r(i, g)}^g} \min_{a_{-r(i, g)}^g} u_r^g(a_{-r(i, g)}^g, \sigma_{r(i, g)}^g)$. This guarantees her a long-term utility of at least $d_{r(i, g)}^g$ in each minigame g . By assumption, $\sum_g m_{r(i, g)}^g < 0$, which is equivalent to $\sum_g c_{r(i, g)}^g < \sum_g d_{r(i, g)}^g$. So player i has an incentive to deviate and the original strategy profile is not an equilibrium. \square

To determine whether cooperation can be attained, properties of the group of minigames as a whole matter, as opposed to those of individual minigames: even if there are many minigames with negative robustness measures to all agents, one more minigame could reverse the sum of the robustness measures to positive and result in cooperation. Because it is straightforward to compute the robustness measures from the minigames using the formulas above, Theorem 1 allows us to efficiently check whether a given role assignment has the desired behavior as an equilibrium. It also reduces the

problem of *finding* such a role assignment to the following computational problem

Definition 4. (ROLE-ASSIGNMENT) We are given the m_r^g (a vector of $n|G|$ numbers). We are asked whether there exists a function $r(i, g)$ that maps players one-to-one to the roles of the game g , such that for all i , $\sum_g m_{r(i,g)}^g \geq 0$.

The reader may be concerned that perhaps the formal ROLE-ASSIGNMENT problem is harder than the problem we actually need to solve, because perhaps some instances of this problem would not correspond to any actual game. Theorem 2 establishes that, conversely, given any vector of robustness measures m_r^g (a vector of $n|G|$ numbers), it is straightforward to construct a vector of g normal-form minigames that results in these robustness measures. This implies that if there is an efficient algorithm to solve the original problem given in the normal-form representation, then any ROLE-ASSIGNMENT instance can also be solved efficiently by solving the corresponding normal-form instance. We omit some proofs to save space.

Theorem 2. Given a length n vector of robustness measures m , we can efficiently construct an n -player normal-form minigame with two actions per player that generates exactly these robustness measures.

Proof. We refer to the i -th element of the vector m by m_i . From the robustness measure vector, we construct an n -player binary-action normal-form minigame, where each agent chooses whether to play the target action (cooperate) or not (defect).

We first show how to construct the normal-form game in the limit average payoff model. If a given m_i is nonnegative, then we let the i -th agent's payoff for when the action profile $\vec{a} = (\text{cooperate}, \dots, \text{cooperate})$ is played be m_i . If, on the other hand, a given m_i is negative, then we let the i -th agent's payoff for when $\vec{a} = (\text{cooperate}, \dots, \text{cooperate})$ is played be 0, and all of i 's payoffs where i 's action a_i is defect (regardless of the actions of others $-a_i$) be $-m_i$ (that is, there is a positive value for defection). All payoffs not yet specified are set to 0.

In the discounted payoff model, we go through the same operations, except instead of using the payoffs m_i and $-m_i$, we replace them with $(1 - \delta)m_i$ and $-(1 - \delta)m_i$, to account for the discounting.

Finally, we show that the normal-form game we constructed produces the original set of robustness measures. If m_i was nonnegative, $c_i = m_i$ and $d_i = 0$, and if m_i was negative, $c_i = 0$ and $d_i = -m_i$, for both limit average and discounted payoffs. Consequently, we get $c_i - d_i = m_i$ for all i , as desired. \square

7 Complexity of ROLE-ASSIGNMENT

In this section, we show that the ROLE-ASSIGNMENT problem is NP-complete. We first show that it is weakly NP-complete even in an extremely restricted special case, namely, the case where we have only two players and each minigame has the following structure.

Each minigame has two roles, Active and Passive. Passive has no choice. Active can choose to defect, in which

		Active	
		cooperate	defect
Passive	(no choice)	$-x_g$	0
		x_g	0

case both players get 0, or to cooperate, in which case Active gets $-x_g \leq 0$ and Passive gets $x_g \geq 0$. Hence, the robustness values are $m_{\text{Active}}^g = -x_g$ and $m_{\text{Passive}}^g = x_g$, for all $g \in G$. We call these games *Two-Player Active-Passive (2PAP)* games. Intuitively, cooperating in a game means giving the other player a specified gift, and enabling cooperation in all games requires that we balance the gifts exactly between the players.³ This suggests a reduction from the PARTITION problem (which is only weakly NP-hard, allowing pseudopolynomial-time algorithms).

Theorem 3. ROLE-ASSIGNMENT is (weakly) NP-complete for 2PAP games.

Proof. To prove NP-hardness, we reduce from the PARTITION problem, in which we are given a set of integers $\{w_1, \dots, w_q\}$, and are asked whether there exists a subset $S \subseteq \{1, \dots, q\}$ such that $\sum_{j \in S} w_j = \sum_{j=1}^q w_j / 2$. For an arbitrary instance of the PARTITION problem, we construct a ROLE-ASSIGNMENT instance by creating a 2PAP game $g(j)$ with $x_{g(j)} = w_j$ for each $j \in \{1, \dots, q\}$.

If a solution S to the PARTITION instance exists, then assign agent 1 to the Active role in all $g(j)$ with $j \in S$ and to the Passive role in all other games. Then, we have $\sum_g m_{r(1,g)}^g = -\sum_{j \in S} w_j + \sum_{j \notin S} w_j = 0$ and $\sum_g m_{r(2,g)}^g = -\sum_{j \notin S} w_j + \sum_{j \in S} w_j = 0$. Hence a solution to the ROLE-ASSIGNMENT instance exists.

Conversely, if a solution to the ROLE-ASSIGNMENT instance exists, let S be the set of all j such that agent 1 is assigned to the Active role in $g(j)$. We know $0 \leq \sum_g m_{r(1,g)}^g = -\sum_{j \in S} w_j + \sum_{j \notin S} w_j$ and $0 \leq \sum_g m_{r(2,g)}^g = -\sum_{j \notin S} w_j + \sum_{j \in S} w_j$. Hence $\sum_{j \in S} w_j = \sum_{j \notin S} w_j$ and S is a solution to the PARTITION instance. \square

In the coming subsection on the dynamic program algorithm, we will show that the ROLE-ASSIGNMENT problem can in fact be solved in pseudopolynomial time when there are at most a constant number of agents. We now proceed to exhibit strong NP-completeness for n -Player Active-Passive (n PAP) games, in which each minigame g has one Active and $n - 1$ Passive players, and the Active player can choose to make a specified gift x_g that will be equally divided among the other players (so they each receive $x_g / (n - 1)$). The reduction resembles the previous one but is based on the strongly NP-hard 3-PARTITION problem.

Theorem 4. ROLE-ASSIGNMENT is (strongly) NP-complete for n PAP games.

³One may wonder why cooperation is desirable at all in these games, but note that the reduction will work just as well if Passive receives a sufficiently small bonus ϵ when receiving a gift and Active does not have to pay this bonus. Alternatively, the principal may have an exogenous reason for preferring cooperation.

8 Algorithms for Role Assignment

Here, we present two algorithms for ROLE-ASSIGNMENT.

8.1 Integer Program

First, we reduce ROLE-ASSIGNMENT to the integer program (IP) in Figure 1. Combining this with any IP solver results in an algorithm for ROLE-ASSIGNMENT. The robustness measure m_r^g for each minigame g and each role r in g is a parameter of the integer program. We have an indicator variable $b(i, g, r) \in \{0, 1\}$ for each agent i , each minigame g , and each role r in g . ($b(i, g, r) = 1$ if and only if i is assigned role r in g .) There is another variable v which the solver will end up setting to the minimum aggregate robustness value any agent has, $\min_i \sum_g m_r^g b_{r(i,g)}$; maximizing this is the objective of the IP.

Note that this is not necessarily pushing things towards an equitable solution, as m_r^g is not the payoff to the agent. The point is that this lets the IP determine not only whether the ROLE-ASSIGNMENT instance has a solution (which is the case if and only if the optimal objective value is nonnegative), but also the “most robust” solution.⁴ The assignment does not affect the overall welfare of the agents as the aggregate payoffs are constant and predetermined from the prescription of the “cooperation” actions. This IP can easily be modified to determine the minimal subsidy necessary to induce cooperation, by adding a payment variable for each player, whose sum is then minimized (cf. cost of stability [Bachrach *et al.*, 2009]).

<p>maximize v subject to $(\forall i) v - \sum_g \sum_r \text{in } g m_r^g b_{i,g,r} \leq 0$ (min. robustness) $(\forall g, r \text{ in } g) \sum_i b_{i,g,r} = 1$ (one player per role) $(\forall i, g) \sum_r \text{in } g b_{i,g,r} = 1$ (one role per player per game)</p>
--

Figure 1: Integer program for ROLE-ASSIGNMENT.

8.2 Dynamic Program

Even though the general n -player ROLE-ASSIGNMENT problem is strongly NP-complete (Theorem 4), below we give a dynamic programming algorithm that solves it in pseudopolynomial time when the number of agents is constant. For the purpose of presenting this algorithm, we assume that the payoffs are integers. (Of course, any rational numbers could be scaled up to integers).

The algorithm takes as input the the vector of robustness measure m_r^g for each minigame $g \in G$ and each role r in g . Let $L = \sum_g \min\{0, \min_r(m_r^g)\}$ (the lowest possible aggregate robustness measure for an agent from any subset of the games), $U = \sum_g \max\{0, \max_r(m_r^g)\}$ (the highest), and $X = U - L$. Also, let P_g be a (size $r!$) set of vectors of length $|N|$, where each element is a permutation of

⁴The word “robustness” was chosen as a higher value implies the solution would be more robust to changes in the utilities.

$\{m_1^g, m_2^g, \dots, m_{|N|}^g\}$. That is, P_g is the set of all the possible robustness measure vectors from game g .

The algorithm fills up a table of size $|G| \times X^n$ containing Boolean values, with the first axis of the table ranging from 1 to $|G|$, and the other n axes (one for each agent) ranging from L to U . The table entry $T(g, k_1, k_2, \dots, k_n)$ represents whether it is possible for each player i to obtain an aggregate robustness measure of k_i from role assignments to the first g minigames only (arbitrarily labeling the games as $1, \dots, |G|$). We omit the formal description of the algorithm to save space.

The ROLE-ASSIGNMENT instance then has a solution if and only if the last row of the table (for $g = |G|$) has a 1 for an entry with nonnegative values for the other axes—i.e., $(\exists k_1, \dots, k_n \geq 0) T(|G|, k_1, \dots, k_n) = 1$. In this row we can also find the maxmin aggregate robustness level that the integer programming algorithm finds, i.e., $\max\{v : (\exists k_1, \dots, k_n \geq v) T(|G|, k_1, \dots, k_n) = 1\}$. We can also determine the minimal subsidy necessary to induce cooperation by a single pass through this row. All that is needed is to bring up aggregate robustness values that are below zero to zero.

Theorem 5. *ROLE-ASSIGNMENT can be solved in pseudopolynomial time for a constant number of agents n .*

Proof. The table has $|G| \times X^n$ entries, and filling in an entry requires up to $n!$ lookups. Moreover, $X \leq |G| \cdot d$, where d is the maximum difference between two robustness values in a minigame, which itself is $O(v - \lambda)$ where v (λ) is the largest (smallest) single payoff in a minigame. Hence, with constant n , the algorithm is polynomial in $|G|$ and $v - \lambda$. (Of course, the input size is polynomial in $|G|$ and $\log(v - \lambda)$, which is why the algorithm is only pseudopolynomial.) \square

9 Simulation Analysis

In this section, we evaluate the two algorithms on random instances, generated using GAMUT [Nudelman *et al.*, 2004]. For a given number n of players, a given number $|G|$ of minigames, and a given game generator in GAMUT, we generate an instance by drawing $|G|$ n -player games with payoffs in the interval $[-5, 5]$ from the generator. (Though we have done the simulation on many different families of games (dispersion, coordination, N-player chicken, etc.), the runtimes do not appear to depend on the family, so we omit them due to limited space.) Because the DP algorithm requires payoffs to be integers, we round all the payoffs in each game to integers in $\{-5, -4, \dots, 5\}$. We evaluate the IP algorithm on nondiscretized payoffs. (When we run it on the rounded payoffs, the IP algorithm is in fact even faster, and always returns the same solution as the DP.) CPLEX 12.6.0.0 and g++ 4.8.4 were used for IP and DP respectively.

We present the experimental results for when the target action is determined as the action profile maximizing the social welfare. (Similar patterns are observed when the target action is determined randomly.) We evaluate whether the target action can be sustained in an equilibrium. Both algorithms are guaranteed to return a solution, if one exists.

Figure 2 show the results for the DP algorithm. Predictably, the runtime of the DP algorithm closely tracks the number of




Figure 2: The plot represents the average runtime of solving an instance of ROLE-ASSIGNMENT through dynamic programming, given the social welfare maximizing target action profile. The x axis represents how many minigames are to be assigned ($|G|$), and the y -axis represents how long it took to solve the instances, in seconds. Each data point in the graph is an average of multiple instances (ranging from 4 to 200, due to cases such as $n = 6$, where we decide to timeout the program). While we studied many different game generators, as a representative case, we present the results for the case where G consists of uniformly random games.




Figure 3: The plot represents the average runtime of solving an instance of ROLE-ASSIGNMENT through integer programming, given the social welfare maximizing target action profile. The x axis represents how many minigames are to be assigned ($|G|$), and the y -axis represents how long it took to solve the instances, in seconds. Each data point in the graph is an average of 200 instances. The top of the range bar indicates the maximum time an individual instance required to be solved, and the bottom of the range bar indicates the shortest. The graphs presented here are from uniformly random games with no integral payoff restrictions.

table entries that need to be filled in ($|G| \times X^n$). The number of table entries blows up quickly when n increases.

Figure 3 shows the results for the IP algorithm. The IP algorithm scales much, much better.

10 Conclusion

In this paper, we have identified the problem of assigning roles to agents across multiple games in such a way that cooperative behavior becomes an equilibrium. We provided an easy-to-check necessary and sufficient condition for a given role assignment to induce cooperation and used this to obtain hardness results as well as algorithms for the problem of finding such a role assignment. Our integer programming algorithm significantly outperformed our dynamic programming algorithm in experiments, even though the latter is pseudopolynomial for constant numbers of agents.

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 [Klaus *et al.*, 2015] and hedonic games [Aziz and Savani, 2015]); side payments between agents (providing connections to matching with contracts [Hatfield and Milgrom, 2005]); not every minigame being played in each round; generalizing from repeated games to stochastic or arbitrary extensive-form games; and so on. An alternate formation of studying which *payoff profiles* can be sustained can be considered as well. In the

limit-average scenario generalizing to this is straightforward. In the discounted case, however, it is less obvious because the immediate payoff from deviation becomes significant, which depends on which specific entry is played (especially when $\delta < 1$ is considered). We believe that our paper provides a good foundation for such follow-up work.

The availability of a pseudopolynomial-time algorithm, when the number of agents is constant, also suggests that there may be potential for approximation algorithms. However, note that the problems as we have defined them are decision problems, and it is not immediately obvious what the right optimization variant would be. One possibility may be to consider approximate equilibria.

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

- [Aziz and Savani, 2015] Haris Aziz and Rahul 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.

- [Bachrach *et al.*, 2009] Yoram Bachrach, Edith Elkind, Reshef Meir, Dmitrii Pasechnik, Michael Zuckerman, Jörg Rothe, and Jeffrey Rosenschein. The cost of stability in coalitional games. In Marios Mavronicolas and Vicky Papadopoulou, editors, *Algorithmic Game Theory*, chapter 12. Springer, 2009.
- [Bernheim and Whinston, 1990] B. Douglas Bernheim and Michael D. Whinston. Multimarket contact and collusive behavior. *The RAND Journal of Economics*, 21(1):1–26, 1990.
- [Borgs *et al.*, 2010] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Adam Tauman Kalai, Vahab Mirrokni, and Christos Papadimitriou. The myth of the Folk Theorem. *Games and Economic Behavior*, 70(1):34–43, 2010.
- [Bulow *et al.*, 1985] Jeremy I. Bulow, John D. Geanakoplos, and Paul D. Klemperer. Multimarket oligopoly: Strategic substitutes and complements. *Journal of Political Economy*, 93(3):488–511, 1985.
- [Folmer and von Mouche, 2007] Henk Folmer and Pierre von Mouche. Linking of repeated games: When does it lead to more cooperation and Pareto improvements? Working paper 60.2007, FEEM Fondazione Eni Enrico Mattei, C.so Magenta, 63, 20123 Milano, Italy, May 2007.
- [Fudenberg and Tirole, 1991] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, October 1991.
- [Grossi and Turrini, 2012] Davide Grossi and Paolo Turrini. Dependence in games and dependence games. *Autonomous Agents and Multi-Agent Systems*, 25(2):284–312, 2012.
- [Hatfield and Milgrom, 2005] John William Hatfield and Paul R. Milgrom. Matching with contracts. *American Economic Review*, 95(4):913–935, September 2005.
- [Jackson and Sonnenschein, 2007] Matthew O. Jackson and Hugo F. Sonnenschein. Overcoming incentive constraints by linking decisions. *Econometrica*, 75(1):241–257, 2007.
- [Just and Netanyahu, 2000] Richard E. Just and Sinaia Netanyahu. The importance of structure in linking games. *Agricultural Economics*, 24(1):87–100, 2000.
- [Klaus *et al.*, 2015] Bettina Klaus, David Manlove, and Francesca 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.
- [Kontogiannis and Spirakis, 2008] Spyros C. Kontogiannis and Paul G. Spirakis. Equilibrium points in fear of correlated threats. In *Proceedings of the Fourth Workshop on Internet and Network Economics (WINE)*, pages 210–221, Shanghai, China, 2008.
- [Nudelman *et al.*, 2004] Eugene Nudelman, Jennifer Wortman, Kevin Leyton-Brown, and Yoav Shoham. Run the GAMUT: A comprehensive approach to evaluating game-theoretic algorithms. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 880–887, New York, NY, USA, 2004.