Coalitions and Mediators in Multiagent Games

Abstract

Nash equilibria are not immune from defections by a coalition of two or more players in order to improve their own payoffs. We describe this shortcoming of Nash equilibria, and survey the various approaches used to address it, including Aumann's strong Nash equilibria, and Bernheim et al's coalition-proof Nash equilibria. The use of mediators, or entities that can take actions on behalf of participating agents, is also examined. Finally, we investigate the connection between coalitions and mediators, and point to some areas of future research.

1 Introduction

We begin by defining the Nash equilibrium, and show how it is susceptible to deviations by coalitions of players. We continue by presenting other forms of Nash equilibrium that have been identified to address this drawback. We then will discuss mediators, including a particular variant of mediator whose power offers a possible solution to the coalition deviation problem.

To begin, we examine the normal form game of figure 1. Clearly, the ideal outcome of for player 1 (the row player) is TL; however, a little analysis reveals that if both players play rationally (i.e., to maximize their payoffs), this outcome will never occur, since player 2 (the column player) will always get a higher payoff by playing R than L. Player 1 can then ignore the first column, and his optimal choice becomes B.

	L	R
T	7,2	3,5
B	1,3	4,4

Figure 1: A normal form game

The *strategy profile* BR is in some sense an optimal strategy for the players. We define this optimality more rigorously in the next section.

2 Nash equilibria

Note that in the game in figure 1, player 1 cannot improve his payoff by selecting T, if he believes player 2 will play R. Similarly, player 2 cannot gain by playing L if he believes player 1 will play B. Strategy profile BR is called a Nash equilibrium[6].

Definition A Nash equilibrium is a strategy profile in which no player can improve his payoff by deviating from the strategy, if no other player deviates as well.

Nash equilibria are thus strategy profiles that are resistant to deviations by any single player.

A game may have more than one Nash equilibrium. Consider the game shown in figure 2. The strategies TL and BR are both Nash equilibria[1].

	L	R
T	1,1	0,0
B	0,0	2,2

Figure 2: A game with more than one Nash equilibrium

The reader might wonder why either player would deviate from the clearly optimal strategy BR. Indeed, this property of BR endows it with a special name.

Definition A Pareto optimal Nash equilibrium[4] in a game G is a Nash equilibrium in which every player's payoff is not less than his payoff in any other Nash equilibrium in G.

Since TL of the game in figure 2 is a Nash equilibrium, no single player can deviate to improve his payoff. However, *both* players can deviate to BR to their mutual advantage. This suggests another way of defining Pareto optimal Nash equilibria: strategy profiles that are resistent to deviations by all players at once.

The Nash equilibrium is a powerful concept, but has at least one major flaw: such equilibria can be susceptible to deviations of coalitions of two or more players. We illustrate this with the game in figure 3 (from [3]). As before, player 1 is the row player and player 2 is the column player, but now

[1]In this game, there exists a third Nash equilibrium: player 1 plays T with probability $\frac{2}{3}$, and player 2 plays L with probability $\frac{2}{3}$. For the moment, we ignore such *mixed strategy* equilibria.

	C_1: B_1	C_1: B_2
A_1	1,1,-5	-5,-5,0
A_2	-5,-5,0	0,0,10

	C_2: B_1	C_2: B_2
A_1	-1,-1,5	-5,-5,0
A_2	-5,-5,0	-2,-2,0

Figure 3: A three-player game

there is a third player who decides which of the left or right payoff tables are in effect by choosing an action from $\{C_1,C_2\}$.

This game has two Nash equilibria: $A_2B_2C_1$, which is Pareto optimal, and $A_1B_1C_2$, which is not. By definition, we know that if the players are following one of these strategy profiles, then no single player can benefit by deviating. Suppose the players are inclined to play strategy profile $A_2B_2C_1$. What if the players have some way of communicating with each other before deciding upon their strategies? In that case, $A_2B_2C_1$ is unlikely, since players 1 and 2 can agree among themselves to form a coalition and deviate by playing A_1B_1, increasing both their payoffs to player 3's detriment.

Definition A strong Nash equilibrium[1] is a strategy profile where no coalition of players can deviate so as to increase the payoff of every player within the coalition.

It is clear that every strong Nash equilibrium is a Pareto optimal Nash equilibrium. The converse of this is not true, as we shall see.

The strong Nash equilibrium concept seems appealing: such an equilibrium is safe from any deviations by rational players. However, as noted by Aumann when he introduced the concept, some games exist that do not have any strong Nash equilibria. For example, the famous Prisoner's Dilemma of figure 4 has a single Nash equilibrium, DD, but no strong Nash equilibrium. To see this, observe that in strategy profile DD, technically a coalition of every player (the 'grand' coalition) could deviate to strategy profile CC to improve both players' payoffs, however unlikely this would be (since a further deviation by one of the players would then be profitable).

	C	D
C	3,3	0,4
D	4,0	1,1

Figure 4: Prisoner's Dilemma

Indeed, strong Nash equilibria are 'too strong'[3]: excluding all strategy profiles that allow a profitable deviation of a coalition is unrealistic, and so strong Nash equilibria often do not exist.

	C_1: B_1	C_1: B_2	C_2: B_1	C_2: B_2
A_1	1,1,1	3,0,0	0,0,0	0,0,0
A_2	0,3,0	2,2,0	0,0,0	0,0,0

Figure 5: A game with no strong Nash equilibrium

Consider the three player game in figure 5. There is a single Nash equilibrium, $A_1B_1C_1$, which is Pareto optimal. It is not a strong Nash equilibrium, since a coalition of players 1 and 2 can deviate to play A_2B_2 to increase their payoffs. This deviation is not realistic, though, since if players 1 and 2 came to an arrangement to deviate in this way, then either of them would clearly have incentive to deviate again; for instance, player 1 could choose to play A_1, winning a payoff of 3 and leaving the other players both with 0.

For a coalition to be a threat to a Nash equilibrium, they must be able to make a *self-enforcing* deviation: one that is safe from further deviations by a subset of the coalition that would benefit every member of the subset. The definition of self-enforcing is actually a recursive one, since any of these further deviations must themselves be self-enforcing.

Definition A coalition-proof Nash equilibrium[3] is a strategy profile where no coalition exists that has a self-enforcing deviation.

Revisiting the game in figure 5, we see that $A_1B_1C_1$ is indeed a coalition-proof Nash equilibrium.

Every strong Nash equilibrium is also a coalition-proof Nash equilibrium, since by definition, the set of valid deviations in the latter is a subset of that of the former. The relationship between Nash equilibria and coalition-proof Nash equilibria is less clear. As pointed out by Bernheim et al [3], for two-player games, these two classes are equivalent. For multiplayer games, this is not the case; they provide an example of a three-player game that has no coalition-proof Nash equilibria. No other inclusion property between these two classes has yet been established for multiplayer games.

3 Mediators

Mediators are entities that can communicate with and make recommendations to the players in a game. Many types of mediators have been discussed in game theory, with varying abilities to interact in a game. Aumann's correlated equilibrium[2] is perhaps the simplest form of mediator[5]. Consider the Battle of the Sexes game of figure 6. Player 1 clearly prefers outcome

	A	B
A	2,1	0,0
B	0,0	1,2

Figure 6: Battle of the Sexes

AA, while player 2 prefers BB. These are, fact, both Nash equilibria. If we allow the players to play mixed strategies (in which they randomly choose an action based on a probability distribution), then their exists a third equilibrium: player 1 plays A with probability $\frac{2}{3}$, while player 2 plays A with probability $\frac{1}{3}$. The expected payoff for both players is thus $\frac{2}{3}$.

None of these three payoffs seems satisfying. The pure (i.e., non-mixed) strategy profiles are unfair to one of the players, while the mixed strategy profile is fair, but not optimal, as over half of the time the players will miscoordinate and receive a payoff of zero.

In a correlated equilibrium, the randomization that occurs in the latter case can be synchronized to improve the payoffs. If both players have access to the same random signal (i.e., a flip of a coin labelled 'A' and 'B'), and act accordingly, they will receive an average payoff of $\frac{3}{2}$ that is both fair and optimal.[2] The coin is, in a sense, acting as a mediator: an independent agent that communicates with the players, and makes recommendations.

In a correlated equilibrium, the mediator has somewhat limited powers: it has access to information not available to the players (i.e., the result of a coin flip) and can communicate this information, but cannot take actions on the part of a player. Monderer and Tennenholtz[5] have considered a more powerful form of mediator, one that can act on behalf of those players that choose to enlist the mediator's services. Recall that the Prisoner's Dilemma game has a single Nash equilibrium, DD. If we add a mediator to the game, and allow each player to optionally enlist the mediator to act on his behalf (action M), then we have the game shown in figure 7.

If neither player uses the mediator, the game is the same as before. If only one player uses the mediator, the mediator will defect on behalf of that player. If both players use the mediator, then each player receives the optimal 'fair' payoff of 3. This last case represents the new game's Pareto

[2]This presentation of correlated equilibrium is a simplified one, but is suitable for our purposes.

	M	C	D
M	3,3	4,0	1,1
C	0,4	3,3	0,4
D	1,1	4,0	1,1

Figure 7: Mediated Prisoner's Dilemma

optimal Nash equilibrium, MM.

We stress that the CC payoff can never occur in any equilibrium of the Prisoner's Dilemma game, even if some correlating signal is present. Regardless of what any such signal suggests, a player's best move is to defect. To achieve the CC payoff, a mediator must be endowed with the powers to act on behalf of a player; only then can the mediator be sure that the suggested strategy will be carried out.

Another advantage that this more powerful form of mediator provides is motivated by the following *extensive form* game (from [7]):

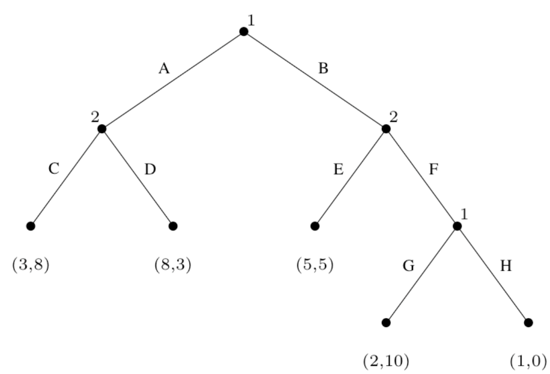

Figure 8: Extensive form game (from [7])

In this game, the players take turns, choosing actions that cause the game to progress down the tree to a root node, which contains the payoff vector. A player's strategy in this game is to describe a choice to be made at each node of the tree (or at least, at each node where it is the player's turn). Player 1 has four pure strategies: {AG, AH, BG, BH}. Player 2 also has four pure strategies: {CE, CF, DE, DF}.

One of the three Nash equilibria of this game is {BH,CE}. As described in [7], this equilibrium has the problem that if player 1 plays B, then player 2 might play F, by reasoning that player 1 will not follow through with H, as this would bring him a lower payoff than the choice of G. It can be argued that player 1's intention to play H, should the game reach the G/H decision point, is a *non-credible threat*: a rational player, when reaching that point, should choose G instead (thus giving player 2 a payoff of 10, higher than if he had not deviated from the equilibrium strategy).

Let us now augment this game with a mediator, one empowered to act on behalf of its participating players. The action that invokes the mediator is available in the root node only, and if selected, the mediator takes over the play for the player for the rest of the game. We will structure the mediator's actions so that it plays the {BH,CE} strategy. This game now has an *induced normal form* that includes, in addition to each player's previous four pure strategies, a mediator strategy M; see figure 9.

	M	CE	CF	DE	DF
M	5,5	5,5	1,0	5,5	1,0
AG	3,8	3,8	3,8	8,3	8,3
AH	3,8	3,8	3,8	8,3	8,3
BG	5,5	5,5	2,10	5,5	2,10
BH	5,5	5,5	1,0	5,5	1,0

Figure 9: Mediated extensive form game

If the mediator is *committed to playing its strategy once the game begins*, and player 1 uses the mediator, then the non-credible threat from the non-mediated version is weakened: player 2 will not defect by playing F if the mediator action was just played by player 1, since player 1 cannot affect the mediator's decision to then play H should this occur. We argue that the mediator is not behaving irrationally by playing H, since his primary motivation is to follow through with a known and predefined plan of action, which includes following through on any threats that the strategy may rely on. In this sense, the mediator can be thought of as an unthinking piece of machinery that once activated, has no ability to change its behaviour. Knowing this, player 2 will not deviate by playing F, since H cannot then be avoided.

Any normal form game can be augmented with a mediator, though the strategies available to the mediator cannot necessarily be represented by a single new row/column as was the case in the last example. Instead, we can represent the *expected* payoff for the strategy within the appropriate cell of the table. Consider the Battle of the Sexes game, with a mediator that mixes equally between AA and BB if both players use the mediator. If only player 1 uses the mediator, its strategy will be to play A, and if only player 2 uses the mediator, its strategy will be to play B.

This game has a Pareto optimal Nash equilibrium strategy MM, with an

	M	A	B
M	$\frac{3}{2},\frac{3}{2}$	2,1	0,0
A	0,0	2,1	0,0
B	1,2	0,0	1,2

Figure 10: Mediated Battle of the Sexes

expected payoff ($\frac{3}{2}$) for both players that was previously achievable only in a correlated strategy.

4 Mediators as Coalitions

We now investigate some ways in which coalitions and mediated games are related. In the mediated games seen in the previous section, the players who choose to use the mediator's services are acting as a coalition.

Definition The *mediator coalition* is the set of players that are employing the mediator's services in a mediated game.

One attractive property of mediated games is that any Nash equilibrium involving only mediator actions is by definition coalition-proof. The mediated Prisoner's Dilemma game is an example; once both players choose M, neither can deviate to D to obtain a higher payoff.

There are several issues to consider when designing a mediated game. Perhaps the most important is to specify the mediator's strategies. These are a set of strategy profiles, one for each possible subset of players that make up the mediator coalition. For instance, in the Prisoner's Dilemma, the mediator strategy profile is different if the coalition consists of both players, versus the case where it consists of only one.

Lemma 4.1 *It must not be the case that all of the mediator's strategies for a single member i of the coalition are (even very weakly) dominated[7] by any other strategy for i.*

Proof If a player can guarantee an average payoff for himself by acting alone that is at least as high as that attainable when he is within the mediator coalition, regardless of who else is in the coalition, then the player has no incentive to join the coalition. ∎

The way the mediator strategies are designed can influence a player's willingness to join the mediator coalition.

Definition An *R-Mediator* is a mediator whose strategies are designed to reward the members of the mediator coalition.

Definition A *P-Mediator* is a mediator whose strategies are designed to punish nonmembers of the mediator coalition.

In the three-player game in figure 11, action C_1 weakly dominates C_2 (player 3). In the non-mediated game, a nonzero payoff for player 3 is very unlikely, since the game reduces to the Prisoner's Dilemma if player 3 always plays C_1. (In this game, a single value in a table cell indicates the common payoff for all three players.)

	C_1		C_2	
	B_1	B_2	B_1	B_2
A_1	3	0,4,0	0	0
A_2	4,0,0	1,1,0	0	0

Figure 11: Three player game

Now consider the P-Mediated version of the game (figure 12). The mediator punishes any players that do not participate by playing A_2, B_2, or C_2 where possible. This game has a single coalition-proof Nash equilibrium {MMM}, which achieves the Pareto optimal payoff for each player of 3. The effect of the P-Mediator is apparent: by using the mediator, player 3 now has significant leverage over the other players, who must play M to avoid a payoff of zero.

	M			C_1			C_2		
	M	B_1	B_2	M	B_1	B_2	M	B_1	B_2
M	3	0	0	1,1,0	4,0,0	1,1,0	0	0	0
A_1	0	0	0	0,4,0	1,1,0	0,4,0	0	0	0
A_2	0	0	0	1,1,0	4,0,0	1,1,0	0	0	0

Figure 12: P-Mediated game

An R-Mediated version of the game is shown in figure 13. The difference between this and the previous game is that the mediator plays C_1 for player 3 if at least one other player is in the coalition, since that other player should not be punished for participating. The second difference is that if players 1 and 2 are both in the coalition, the mediator always strives for the payoff of 3. This game also has a single coalition-proof Nash equilibrium {MMM}.

	M			C_1			C_2		
	M	B_1	B_2	M	B_1	B_2	M	B_1	B_2
M	3	4,0,0	1,1,0	3	4,0,0	1,1,0	0	0	0
A_1	0,4,0	3	0,4,0	0,4,0	3	0,4,0	0	0	0
A_2	1,1,0	4,0,0	1,1,0	1,1,0	4,0,0	1,1,0	0	0	0

Figure 13: R-Mediated game

5 Future Research

To close, we describe some elements of the use of mediators that require further investigation.

We have left the structure of R-Mediator and P-Mediator strategy profiles a little vague. For instance, what abilities does a coalition of players have to encourage non-coalition players to participate? Do there exist rules that can predict whether such coercion can occur?

Is it always possible to structure a mediator strategy that satisfies every participant in the coalition (as in lemma 4.1), or are there instances where two players will always have opposing interests?

Maxmin and minmax strategies[7], which are designed with respect to single players against subsets of players, need to be extended to cases involving sets against sets, as this is key to the design of R-Mediator and P-Mediator strategies.

References

[1] R. Aumann. Acceptable points in general cooperative n-person games. In R.D. Luce and A. W. Tucker, editors, *Contributions to the Theory of Games IV*, pages 287–324. Princeton University Press, 1959.

[2] R. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, March 1974.

[3] B. D. Bernheim, B. Peleg, and M. D. Whinston. Coalition-proof Nash equilibria I. concepts. *Journal of Economic Theory*, 42(1):1–12, June 1987.

[4] J. Case. On Ricardo's problem. *Journal of Economic Theory*, 3(2):134–145, June 1971.

[5] D. Monderer and M. Tennenholtz. Strong mediated equilibrium. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, 2006.

[6] J. Nash. Non-cooperative games. *Annals of Mathematics*, 42:286–295, 1951.

[7] Y. Shoham and K. Leyton-Brown. Multiagent systems: Algorithmic, game-theoretic, and logical foundations. Manuscript in preparation, 2008.