CPSC 532a Project
Better Responses

Anthony Lee
awll@cs.ubc.ca

December 27, 2006

Abstract

A Nash equilibrium refinement is introduced that forms a nonempty subset of the Nash equilibria of a game. This refinement is attained via removal of mutually dominated strategy subprofiles, a process justified by every agent being better off if every agent follows it. The new solution concept is applied to various example games and its usefulness discussed for general games.

1 Introduction

The Nash equilibrium is an important solution concept and theoretical construct, with a large number of fields of study using Nash equilibria to provide insight into systems where multiple agents make decisions [1]. However, the notion of stability in a Nash equilibrium relies primarily on each agent being worse off by deviating, regardless of whether or not other agents would decide to play their components. In particular, a subset of agents could be forcing each other to play certain strategies even though it is not in any of their best interest to do so.

Consider the two-player coordination game in Figure 1. There are two pure strategy Nash equilibria: TL and BR. However, in the TL Nash equilibrium, each player is forced to best respond with T or L simply because the other player is playing L or T. Both players would prefer to find the BR equilibrium, and should safely be able to assume that the other player will play his part of that pure strategy. Therefore, in this game, the TL equilibrium is less appealing. Another way of looking at this is to notice that a joint deviation from TL to BR would be mutually beneficial for both players.

The susceptibility of Nash equilibria to joint deviations has been discussed in [2] and [3], which contributed the strong Nash and coalition-proof equilibria

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

Figure 1: A Simple Coordination Game

respectively. However, these solution concepts are most interesting when communication is allowed and neither are guaranteed to exist for a general game [4]. Furthermore, it is unnecessary to move to the realm of games with free communication when addressing the problem of joint deviations. In fact, in games without free communication the formation of coalitions is not as much of a concern as is the difficulty in finding an equilibrium.

2 A New Concept: Mutual Domination

It is a combination of factors that makes the TL equilibrium in Figure 1 so unappealing. In particular, the BR joint strategy is better for both players and for each player, if they assume that the other player does not play their part in the TL strategy profile, they are strictly better off playing their part in the BR strategy. This intuition can be used to define the concept of mutual domination more formally.

2.1 Definition of Mutual Domination

Let a game be denoted by $G = \langle N, A, O, \mu, u \rangle$, where N is a finite set of n agents, $A = (A_1, ..., A_n)$ with A_i equal to the set of actions available to agent i, O is a set of outcomes, $\mu : A \to O$ maps action profiles to outcomes and $u = (u_1, ..., u_N)$ with $u_i : O \to \mathbb{R}$ equal to the real-valued utility function for agent i.

We want to consider a set of agents $M \subseteq N$. Without loss of generality, let this set be $M = \{1, ..., m\}$. In addition, let the set of all possible joint strategies by agents in $N - M$ be denoted by S_{-M} and the set of all possible joint strategies by agents other than i be denoted by S_{-i}. This subset M is of interest when there exists a pure strategy subprofile $(a_1, ..., a_m)$ for these agents for which there exists another (possibly mixed) strategy subprofile $(s_1, ..., s_m)$ with the following properties:

- $u_j(a_1, ..., a_m, s_{-M}) < u_j(s_1, ..., s_m, s_{-M}) \forall j \in M, \forall s_{-M} \in S_{-M}$
- $u_j(a_j, s_{-j}) < u_j(s_j, s_{-j}) \forall j \in M, \forall s_{-j} \in S_{-j} - \{(a_1, ..., a_m, s_{-M}) | s_{-M} \in S_{-M}\}$

The first condition states that every agent in M prefers the strategy subprofile $(s_1, ..., s_m)$ over $(a_1, ..., a_m)$. The second condition implies that action a_j for agent $j \in M$ can only be a best response to a strategy subprofile in which all other agents $i \in M$ play action a_i. In particular, if any agent $i \in M$ plays a strategy not involving a_i every other agent $j \in M$ is strictly better off playing s_j than a_j.

If these conditions hold, we consider $(a_1, ..., a_m)$ to be *mutually dominated* by $(s_1, ..., s_m)$.

2.2 Using Mutual Domination

The proposed new solution concept is any Nash equilibrium that survives iterated removal of mutually dominated strategy subprofiles, where the removal of a profile amounts to removing each agent's action in the subprofile from the set of available actions for that agent. In order for this solution concept to be interesting, the removal of the strategy subprofiles must be justified and the solution concept should have existence properties that are comforting. For the latter, it will be shown that Nash equilibria that survive iterated removal of mutually dominated strategy subprofiles are also Nash equilibria of the original game and that Nash equilibria of the original game that do not involve actions that are removed are preserved. In other words, this solution concept forms a nonempty subset of the Nash equilibria of a game.

2.2.1 Justifying the Removal

The most contentious issue with respect to the utility of this refined Nash equilibrium is the removal of the actions that make up a mutually dominated subprofile. However, the motive behind such a removal is that all the agents that remove their actions are better off by doing so, since they are better off by playing s_i than a_i, as long as at least one other agent does the same. Furthermore, from the point of view of an agent who is not in M, it can be useful to know which actions in other agent's action spaces take part in a mutually dominated strategy subprofile. In other words, the only agents who can possibly be worse off because of a successful[1] removal are agents not in M and even then there are advantages to knowing what will not be played.

2.2.2 Correspondence of Other Nash Equilibria

Let G be the original game and G' be the game after a removal of a mutually dominated strategy subprofile $(a_1, ..., a_m)$.

[1]Successful here means that more than one agent in M removes the mutually dominated strategy subprofile

Assume that we have a Nash equilibrium in G that does not involve any of the a_i's, $(s'_1, ..., s'_m, s'_{-M})$. Then, since we have only removed an action from each agent's set of available actions, this must be a Nash equilibrium in G'. This means that all Nash equilibria in G that do not involve the a_i's are preserved in G'.

Assume that we have a Nash equilibrium $(s'_1, ..., s'_m, s'_{-M})$ in G'. Let G'_i denote the game which is equivalent to G' except that action a_i made available to agent i. Since any agent has the all the actions available in G and G'_i that they do in G', $(s'_1, ..., s'_m, s'_{-M})$ is a valid strategy profile in G and G'_i. Note that G'_i is strategically equivalent to G' since action a_i is strictly dominated by s_i in G'_i. Thus $(s'_1, ..., s'_m, s'_{-M})$ is also a Nash equilibrium in G'_i. This means that given the strategy profile $(s'_1, ..., s'_m, s'_{-M})$, agent i is strictly worse off by playing a_i than s'_i (s'_i must be a weakly better response than s_i with is a strictly better response than a_i). Since this is true for all agents $i \in M$, this means that $(s'_1, ..., s'_m, s'_{-M})$ is a Nash equilibrium in G. Therefore, all Nash equilibria in G' are also Nash equilibria in G.

Furthermore, since every game has a Nash equilibrium, G' has a Nash equilibrium. So even after removing the a_i's from the game, we are guaranteed to have a Nash equilibrium in G' that is also a Nash equilibrium of G.

3 The Equilibrium in Practice

What follows is an exploration of how this new solution concept applies to some simple games. In the simple examples described below, the strategy subprofile $(s_1, ..., s_m)$ also acts as a suggested strategy subprofile for players in M. However, this is not necessarily the case in general. The significance of a pure strategy subprofile $(a_1, ..., a_m)$ being dominated by a strategy subprofile $(s_1, ..., s_m)$ is very similar to the significance of an action being strictly dominated. As such, whenever the agent considers playing a_j he should play s_j instead. In general, the removal of actions only corresponds to telling agents what *not* to play; the $(s_1, ..., s_m)$ subprofile is unimportant except for establishing which actions should be removed from consideration.

3.1 A Simple Coordination Game

In Figure 1, TL and BR are the pure strategy Nash equilibria. However, TL is mutually dominated by BR according to the above criteria and as such both T and L can be removed from consideration, leaving only B and R as available actions. Both agents are better off in this situation than in TL or in the mixed Nash equilibrium.

3.2 A Minimum-Effort Coordination Game

This game has been discussed in [1]. In Figure 2 it is reproduced with effort levels as any integer in $\{1, 2, 3, 4\}$ and an effort cost value of 0.75. In the above game,

	1	2	3	4
1	$0.25, 0.25$	$0.25, -0.5$	$0.25, -1.25$	$0.25, -2$
2	$-0.5, 0.25$	$0.5, 0.5$	$0.5, -0.25$	$0.5, -1$
3	$-1.25, 0.25$	$-0.25, 0.5$	$0.75, 0.75$	$0.75, 0$
4	$-2, 0.25$	$-1, 0.5$	$0, 0.75$	$1, 1$

Figure 2: A minimum-effort coordination game

the pure strategy Nash equilibria are all the common effort levels. However, $(1, 1)$ is mutually dominated by $(2, 2)$, which is in turn mutually dominated by $(3, 3)$, which is finally mutually dominated by $(4, 4)$. Again, the game is reduced to a single action per agent. Furthermore, playing $(4, 4)$ yields the best possible payoff for both players.

The unique refined Nash equilibrium is not always the strategy profile that arises most often when this game is played with human subjects. In particular, it seems that human behaviour in this game has effort levels inversely related to effort costs [1]. However, if all agents were to play their part in the refined Nash equilibrium they would all be strictly better off.

3.3 The Prisoner's Dilemma

	Cooperate (C)	Defect (D)
Cooperate (C)	$2, 2$	$0, 3$
Defect (D)	$3, 0$	$1, 1$

Figure 3: A prisoner's dilemma game

Since DD is the unique Nash equilibrium in the prisoner's dilemma (Figure 3) we know that the unique refined Nash equilibrium is also DD. However, it is interesting to notice where mutual domination falls apart in this case. CC is a candidate for (s_1, s_2) with $(a_1, a_2) = DD$. The first criteria is met but the second is not. In particular, if it is assumed that agent 2 will not play D then agent 1 is better off playing D than C, ie. $u_1(a_1, s_2) > u_1(s_1, s_2)$.

3.4 The Battle of the Sexes

The battle of the sexes game in Figure 4 is another example of a game in which mutual domination is not helpful. For each of the pure strategy Nash equilibria,

	Football (F)	Ballet (B)
Football (F)	$2, 1$	$0, 0$
Ballet (B)	$0, 0$	$1, 2$

Figure 4: A battle of the sexes game

FF and BB, there does not exist another strategy profile that pareto-dominates it. Intuitively, the utility values do not suggest a mutually beneficial alternative because none exist.

3.5 The Stag Hunt

	Stag (S)	Hare (H)
Stag (S)	$9, 9$	$0, 8$
Hare (H)	$8, 0$	$7, 7$

Figure 5: A stag hunt game

This famous game [5, p. 309] also has an interesting unique refined Nash equilibrium, SS, since this Nash equilibrium is often seen as unsafe. The difference between this game and the simple coordination game is that playing H is a safer bet than playing S, while the penalty of miscoordination in the game in Figure 1 is symmetric.

In fact, this game and the minimum-effort coordination game highlight the biggest concern with the process of removing mutually dominated strategies: the definition of a mutually dominated strategy is insensitive to the magnitude of utility losses when all other players j play a_j.

3.6 A Three-Player Game

	L	R
T	$1, 1, 3$	$0, 0, 1$
B	$0, 0, 1$	$4, 4, 0$

U

	L	R
T	$2, 2, 0$	$0, 0, 1$
B	$0, 0, 1$	$3, 3, 1$

D

Figure 6: An example three-player game

Although already interesting in two-player games, the refined Nash equilibrium is especially interesting in games with more players. Of particular importance is the fact that with more than two players, there is the possibility that $M \neq N$ and so the strategy profiles that are removed need not be pareto-dominated for *all* agents.

This game has two pure strategy Nash equilibria: TLU and BRD. However, consider $M = 1, 2$, $(a_1, a_2) = TL$ and $(s_1, s_2) = BR$. T and L are mutually dominated in this game. With T and L removed, BRD is the unique Nash equilibrium. Here is an example of a situation where it is in every agent's best interest to remove mutually dominated subprofiles. Agents 1 and 2 are strictly better off by coordinating on BR than anything else. Furthermore, agent 3 is better off playing D than anything else once he predicts that agents 1 and 2 are going to play BR. Notice that even though TLU is a pareto-optimal outcome in this game, it is not a refined Nash equilibrium.

4 Discussion

There are two ways to interpret the removal of mutually dominated strategies. The first is that this process helps agents *find* an equilibrium by inspection of utility values. The other is that it simply highlights equilibria that are less susceptible to joint deviations than others.

The Nash equilibria that survive iterated removal of mutually dominated strategies are not immune to joint deviations. However, such deviations are riskier in the sense that they would probably require communication to coordinate on. With mutual domination, the utility values for agents are all that are needed to suggest which strategies should not be played. The reason that this solution concept exists for all games is that it is considerably weaker than the strong Nash and coalition-proof equilibrium concepts.

One possible problem with the concept of mutual domination is that it involves trusting that at least one other player j won't play a_j. However, the redeeming quality of the new solution concept is that if every player, including players in $N - M$, use it, they are all better off by doing so. It is clear that every player in M is better off. In addition, players in $N - M$ are better off by being aware of what agents in M will not play. For example, in Figure 6 player 3 is better off playing D than U. The issue of insensitivity to the magnitude of possible utility losses could possibly be removed with a definition of mutual domination that takes into account this risk[2].

The computational power required to decide which strategy profiles are mutually dominated has been ignored in this paper. However, simple approaches can be expensive. For the solution concept to be useful, it must be possible for all agents to remove actions in mutually dominated subprofiles from consideration. As such, this solution concept might be less useful in the context of computationally bounded agents.

This solution concept also applies to correlated equilibria. In this case, if a

[2]One way to do this could be to bound the ratio of the minimum gain of playing $(s_1, ..., s_m, s_{-M})$ to the maximum loss for a player who is the only one to play s_i.

signal tells an agent j to play an action a_j that is part of a mutually dominated strategy subprofile, the agent is better off if he plays s_j, his component of a strategy subprofile that dominates $(a_1, ..., a_m)$.

5 Conclusions & Future Work

The removal of actions in mutually dominated strategy subprofiles leads to a new solution concept that is forms a nonempty subset of Nash equilibria for a game. As a Nash equilibrium refinement, this concept gives interesting results when applied to many games. Of particular interest is the fact that if every agent removes mutually dominated subprofiles when analyzing the game, they are better off by doing so.

The primary concern with the proposed removal strategy is that it is insensitive to the magnitude of utility losses when every other agent plays their part in a mutually dominated subprofile. It is possible that a characterization of agents could be useful in quantifying the risk that an agent would be willing to take when removing strategies. For example, there might be stag hunt games for which a type of agent would remove HH from consideration and others for which the same type of agent would not.

References

[1] Jacob K. Goeree and Charles A. Holt. Ten little treasures of game theory and ten intuitive contradictions. *The American Economic Review*, pages 1402–1422, 2001.

[2] R. Aumann. *Acceptable Points in General Cooperative n-Person Games*, volume IV. Princeton University Press, 1959.

[3] B. Douglas Bernheim, Bezalel Peleg, and Michael D. Whinston. Coalition-proof Nash equilibria I. Concepts. *Journal of Economic Theory*, 42:1–12, 1987.

[4] Nikolai S. Kukushkin. An existence result for coalition-proof equilibrium. *Economics Letters*, 57:269–273, 1997.

[5] Yoav Shoham and Kevin Leyton-Brown. *Multi Agent Systems*. Draft copy, 2006.