An Essay on the Application of Rationality to Game Playing
by Ryan M McCabe, October 2006

Applying a hypothetical rationality to any situation can quickly become a funny thing. Afterall, it should be a process without much debate among rational people. But that is clearly not the case in practice. Strengthened and limited by our own personal powers of rationality and distracted by many other considerations, we often tend to disagree exactly on just what is rational behavior. That does not mean progress cannot be made. Rationality does contain a sort of self-governing set of rules, and given a certain set of circumstances even the most antagonistic bunch can usually come to some sort of agreement on the relative merits of a rational progression.

Game theory is an interesting area of study, since everyone knows how to play games and likes to play at least some types of games sometimes. As a people, we have been playing games longer than we have been intelligently using tools. Yet the dawn of the rigorous formalism of mathematics and the recent capacity of technology have seemingly increased the debate over ways to play games. We now try to play games not just as rational people but, in a sense, as rational computers. And although a computer could never be more rational than the set of theories telling it what to do, we often tend to be surprised by the results of our implementations.

At the heart of much of game theory is the assumption that the players are rational and only self-interested. Built in to this assumption is the notion that to play a game rationally is to try to win that game in some sense. Of course there are many ways to win a game, but for the purposes of rational discourse, it is convenient to reduce all such considerations to a utility matrix that summarily describes the entire reward structure for each player playing the game. Methodically, then, rationally, we may march through this utility matrix and all the possible actions in a game to find some result of play. If we don't like the result, sometimes we back-track and try to find where we went wrong. At what point did we stop progressing rationally?

Mathematicians are often obsessed with various things of which many rational people are not aware. Most rational people have never seen their own asymptote. And although many people are sure of their own greatness at a game, most do not know how to consider if their play is optimal or if they can guarantee a lower or upper bound on their losses. So it is often very fun to look to mathematics to find ways of beating otherwise rational people at different games.

But even by using the capacity of computers to leverage the power of mathematic expressions over games and other rational people, we cannot escape the debate over the application of this rationality. Any such theory we pluck to guide rational play was thought of by some supposedly rational person and is, therefore, limited to some extent. Furthermore, a necessary trait of rationality is maintaining an open mind *a priori*. Which is to say if a particular theory has been singled out beforehand to describe the rational play of a game, you can only guarantee that the ensuing debate will intensify as a result of such an initial irrational step. Obviously, there are plenty of simple games to play rationally and with conclusive results, such as tic-tac-toe. But such games are not interesting. As soon as rational people learn how to play it, they immediately stop. So, the interesting games are just the ones rational people keep talking about and playing.

The prisoner's dilemma is a game that gets talked about a lot. The game is very simple and is cast in an interesting context. Each player has to make one decision to either cooperate or defect relative to the other person playing the game. The players are assumed only to be rational and to in no way have the ability to communicate with each other or to be concerned with each other. The only thing that is known for sure is that someone is going to jail. No rational person wants to go to jail.

The utility matrix for the game is defined in a particular way to give the game some peculiar dynamics. Furthermore, the game is said to be a game of simultaneous decision, meaning each player has to make their decision before the other player's decision is known. Enter John Nash, a mathematician who gets talked about a lot as well. Nash defined a particular sort of equilibrium that can be found in many games, based on his application of a type of rationality to a given utility matrix of action outcomes.

Nash imposes a strong condition on his equilibrium and on the assumed type of rationality used by players to make decisions. Namely, a player must assume that their opponent's strategy is given and can therefore pick among the best available strategies for their own benefit. In this manner, a Nash equilibrium can be identified with the quality that no player will benefit by changing their strategy *unilaterally*. That is certainly a strong assumption about the way a game should be played and one that can easily be challenged under certain circumstances.

The Nash equilibrium of the prisoner's dilemma is that both players defect, and so both players go to jail for a decent amount of time. This is called by many people to be a stable equilibrium and is defended by the notion that either player will *always* do better by defecting. A quick glimpse at the utility matrix and walk-through of the if-then cases as guided by Nash's theory makes this result immediately clear. The somewhat surprising result of this analysis is that both players are forced to harm themselves through apparently pure rational play.

At this point, we might start to get fidgety about our application of rationality. As I mentioned earlier, sometimes it is necessary to back up and ask again, at what point did we stop behaving rationally? In this case, however, it may not be that we lost our rationality along the way; it may be that we have simply not taken it far enough. Two rational players of this game will quickly find the Nash equilibrium in an almost downhill manner. What is not obvious to a rational person, though, is why the decision-making process ends there.

For the sake of argument and illustration, I like to imagine that Spock and Ayn Rand robbed a bank together. If you are unfamiliar with either of these people, well, Spock is not a real person, but they are both imminently logical and self-preserving people. To say the least, each of these two is considered very rational, indeed. As rational people, they would each quickly come to the conclusion that they must rat each out and spend too much time in jail. However, as an obvious extension of rationality and self-preservation, it is clear that the state of mutual cooperation is preferred by each player to the state of mutual defect. Perhaps other states are preferable as well, but let's save those for later examination. So, using nothing but rationality of thought and self-preserving principles and using no direct communication or former contracts, could it be possible for these two to make a sort of quantum jump from the Nash equilibrium to the "unstable equilibrium" of mutual cooperation? Of course, the outcome itself is possible through the logical gates of the game, but Nash's theories remove any relationship among outcome states that do not share a row or column of the utility matrix. It is therefore necessary to show how the game may progress to achieve that point using only the first assumptions themselves as foundational truths.

In the first place, it is reasonable for a rational player to assume that another rational player with an identical utility matrix will play a game the same way they will. Why? Because they are both rational and will make only rational choices. This makes any outcomes away from the identity diagonal relatively unprobable options. In this context, the rational players now find themselves choosing from action outcomes along the identity diagonal only. Now, I have not made this a constraint on the game, I have only chosen to apply rationality in a more general way, yet have considered nothing more than my own outcome and the rationality of myself and my opponent. (I am pretty sure that neither Spock nor Ayn Rand cares what happens to the other.) The conclusion then, extending Nash's own logic in a metaform of self-play, is that each player will choose unilaterally to cooperate.

Of course, it is easy for such a conclusion to be tempted by further iterations of rational thought. Since we have already stepped forward through Nash, why not step back into Nash and take advantage of the other player's rationality to put them away for a long time and get ourselves out of a jam? That is a sort of myopic application, clearly, but it is limited by the general application of rationality. In a simultaneous decision game, it is difficult not to think of it as a sequential game with all the possible sequences sort of happening at once. That seems to be how Nash and many others thought of it. Certainly, the opportunity to exploit the rationality of others remains an available option in any game and with any application either through the myopic rationality used by Nash or the general rationality shown above, but is nonetheless limited by our own rational ability to avoid exploitation. The rational threat of defecting is a natural defense mechanism against it. Thus ends the temptation to take advantage of your opponent by any dimly assumed momentary state of that rationality. You assume they are no more of a sucker than you are.

This application of rationality becomes easier to accept and has already been explored through the iterative version of this game. However, the relationships between asymptotic, limited-iteration and single-instance theories of play remain murky. Accordingly, it is often easy to find examples of this general type of rational conclusion, such as in accepting a cold war over a nuclear war, but equally difficult to justify exactly what style of game is being played. That, ultimately, is the power of the general application of rationality: it relaxes all other assumptions the players need to make about the game.

In any case, it feels much more comfortable now to our own use of rationality to see that a purely rational approach even to the single chance form of this game can result in both an individual and shared optimal outcome rather than the predicted large cost and stability of the punitive state.

Of course this will have many game theorists seeing red by now, regardless of their own rationality, since rational play has already been clearly defined by other rational people. And that is the fun part of the method, reconciling the essential need for debate to enable the usefulness of rationality with the silencing effect that rationality can have on the debate itself.