Games with Real Talk

Benjamin Bachi

April 8, 2010

Abstract

When players in a game can communicate they may learn each other's strategy. In such situations, it is natural to define a player's (pure) strategy as a mapping from what he has learned about the other players' strategies into actions. In this paper we investigate the consequences of this possibility in two player games and show that it expands the set of equilibrium outcomes the players can reach. When strategies are completely observable, any feasible and individually rational outcome can be sustained in equilibrium. If communication fails to reveal the players' strategies with some positive probability, the set of equilibria may be smaller. We demonstrate this in the prisoner's dilemma and find the exact level of cooperation the can be sustained in equilibrium for any set of parameters.

Contents

1 Introduction

Standard game theory usually ignores the fact that people in strategic situation often have the opportunity to communicate before they choose their actions. During this communication, players may reveal their true strategies to their opponents, sometime intentionally and sometimes despite their will. When this information is detected by the other players, it may have some effect on the action they choose in the game itself. For this reason, a player may want to condition his own actions on the information he learns from his opponents during the interaction they have.

To keep things simple we focus on two player games. We assume that during a conversation, or some other form of interaction, each player either receives a correct signal about the other player's strategy (with some known probability), or learns nothing. Since it's not implausible to assume that the probability of detecting the opponent's strategy depends on the quality (or some other aspect) of the interaction the players have, we allow for correlation between the probabilities of the signals received.

In a game with Real Talk we define a strategy to be a function from the set of the other player's strategies and the empty set, which represents the possibility of learning nothing, into actions. Suppose that two players play the game G with sets of actions A_1 and A_2, respectively. Then, if we denote the set of player 2's strategies by S_2, it follows that a (pure) strategy for player 1 is given by a mapping $s_1 : S_2 \cup \{\emptyset\} \to A_1$, and similarly, a (pure) strategy for player 2 is given by a mapping $s_2 : S_1 \cup \{\emptyset\} \to A_2$. Thus, the introduction of communication implies that the players' strategies belong to a richer set than under "standard games", where players' pure strategies are given by $S_1 = A_1$ and $S_2 = A_2$, respectively.

Clearly, our definition of a strategy is circular: in order to define a strategy for player 1 we need to know what is the set of strategies of player 2, but in order to define a strategy for player 2 we need to know what is the set of player 1's strategies. We do not attempt to construct, or characterize, strategy sets for the two players in this paper. However, such pairs of mutually consistent strategy sets exist. One trivial example is defining for any action $a_i \in A_i$ a constant strategy $s_i \in S_i$ which always plays a_i regardless of what player i learns about player j's strategy. In this case the game with Real Talk is equivalent to the original game G.

It should be emphasized however, that regardless of the way in which we define the sets of the players' strategies, at least for one player it is a proper subset of all the possible functions from his opponent's strategies to his own actions. To see this, suppose that each player has at least two possible actions and that player 2 has M strategies (M can be finite or not). Thus, if player 1 had all possible functions from player 2's strategies into actions as pure strategies, he would have had a minimum of 2^{M+1} strategies. The same logic implies that player 2 should have at least $2^{2^{M+1}}$ pure strategies. This implies a contradiction, because there is no M, finite or not, that satisfies the equation $M = 2^{2^{M+1}}$.

In the range between games with Real Talk with strategy sets that include only trivial strategies, and games with Real Talk with strategy sets that are too large to exist, there exist some interesting examples. To illustrate how the possibility of real talk can expand the set of equilibrium outcomes, consider the following example. Suppose that the game G is the prisoner's dilemma (PD), where the actions for each player are either to cooperate, C, or defect, D. Let s be a strategy that plays the action C if the other player chooses s as well, and plays D otherwise. Let $S_1 = S_2 = \{s\}$. Clearly, playing s by both players is an equilibrium, because s is the only strategy. However, even if S_1 and S_2 are enlarged so they include any other strategies, playing s still remains an equilibrium, which means that the profile action selected by the players would be (C, C). This example shows that our setting broadens the set of equilibria, since cooperation is not an equilibrium in the regular game of PD.

However, the ability of s in the above example to reach full cooperation relied on the fact that both players know exactly what strategy their opponent is about to play. When this is not the case, that is, when the players may fail to learn their opponent's strategy, this might not work. If the probability for not learning anything is large enough, the players might exploit this and not cooperate, hoping that their opponent will not know that they defected. Therefore, the ability of Real Talk to expand the set of equilibria depends both on the specific game the players are playing, and both on the probability for detecting the opponent's strategy.

Similarly, the correlation between the signals also has some effect on the game's outcome. Knowing that if a one player learns his opponent strategy, then his opponent learns his, may allow players to choose strategies that they could not have chosen otherwise. Of course, this is true also in the other direction - outcomes that are possible when there's no correlation might not be an equilibrium when the correlation is higher.

The plan of the paper is as follows. Section 2 consists of a review of literature that relates to this paper. In Section 3 we present the Real Talk model and basic definitions. In section 4 we provide several folk theorem like propositions, analyzing the possible payoffs that may be sustained in equilibrium. Section 5 analyzes the prisoner's dilemma with Real Talk. Section 6 concludes.

2 Related literature

This paper concerns Real Talk - true information that is transferred between two people in a conversation. This is very different from Cheap Talk, which Farrell and Rabin (1996) describe as "costless, nonbinding, nonverifiable messages that may affect the listener's beliefs", in their review of this literature. Even though Real Talk is costless, the messages (or signals) that pass between the players are true, and in that sense are both binding and verifiable. It is also important to point out that adding Real Talk to a game may expand the set of equilibria in games where Cheap Talk fails to. The best example is the PD: Cheap talk does not add a single equilibrium to the game, whereas in Real Talk player can achieve full cooperation.

Real Talk is also different from Aumann's correlated equilibrium (1974), in which players condition their actions on a public signal: in Real Talk a strategy assigns an action to every strategy of the other player, while in correlated equilibrium it assigns one to any observation a player can make of the public signal. Moreover, as in Cheap Talk, cooperation cannot emerge in a correlated equilibrium of the PD.

The rest of this section describes other works which relate to the current one.

2.1 Computer Programs

Our notion of strategy can be interpreted in several ways, one of which is as thought procedure, or way of thinking. Under this interpretation, each player receives some information about the way the other player thinks, and based on this information chooses an action. This procedure, receiving a signal and choosing an action, is the player's own thought procedure, and in turn may be discovered by the other player. If one adopts this interpretation, it is possible to think about each player as a computer that runs a certain program, or a Turing machine. Each computer receives the program that runs on the other computer as input, and it's output is an action.

Howard (1988) analyzes a game in which two computer programs play the prisoner's dilemma against each other. He shows that it is possible to write a program that receives the code of the program running on the other computer as input, and tells if it is identical to itself or not. This program can be slightly modified to also choose an action as output. In such a way, Howard constructs a program that plays C when receiving itself as an input, and D otherwise. Clearly, if both computers run this program, it will lead to an equilibrium in which both computers cooperate. Moreover, it is possible to write two different programs, P and Q, such that P recognizes Q, and vice versa. By doing so other equilibria may be sustained.

Tennenholtz, in his paper "Program Equilibrium" (2004), shows that in this setting any payoff which is both individually rational and feasible can be achieved in equilibrium. This result is similar to the Folk Theorem in repeated games, with one distinction: since in Tennenholtz's model each player plays a mixed strategy that is independent of the other player's strategy, it might be impossible to support a payoff profile that is achieved by a correlated distribution over the game's possible outcomes.

Fortnow (2009) extends Tennenholtz's program equilibrium to an environment in which the player's payoffs are discounted based on the computation time used. He also proves a full Folk Theorem (that is, including payoffs obtained only by correlated mixed strategies), in which for any probability distribution over the game's outcomes, D, and for any ϵ, an there is a Nash equilibrium in the mixed program equilibrium game where each player's expected payoff is within ϵ of his expected payoff over D.

2.2 Critique of This Literature

One reason for not thinking of players as computer programs, or Turing machines, is the existence of some important impossibility results. Binmore (1987) shows that a computer program cannot "predict its opponent's behavior and simultaneously participate in the action of the game". In more detail, Binmore describes a situation in which each program receives an input that includes its own program, the opponent's program and the rules of the game. Its expected output is both a recommendation for an action to be played, and a prediction of what would be the recommendation given by the other program. The result is that it's impossible to write a program that will always predict correctly the recommendation of the other program.

Note that in Binmore's setting, the program was only expected to predict correctly the other program's recommendation - it wasn't required that the recommendation would be optimal given the program's prediction. That is, the programs were allowed to recommend actions that are not best response to their prediction, thus making it difficult for the other program to predict them correctly. However, Anderlini (1990) showed that even if we restrict ourselves to programs that always recommend an action that is a best response to the prediction, it is impossible to construct a program that will always predict correctly the other program's recommendation. The positive result is that if one allows the programs to run indefinitely (i.e., entering a never-ending loop), then it is possible to write a program that gives a correct prediction when it does halt. See also Rubinstein (1998).

2.3 Delegation Models

Kalai et al. (2009) offer a similar model to the program equilibrium, in which instead of computer programs each player chooses a commitment device. A commitment device is a function that takes the other player's commitment device as its input, and returns a strategy (a probability distribution over the player's actions). In equilibrium, both players choose the optimal commitment device given the other player's choice. The authors prove a full Folk Theorem for two players games. That is, they show that it is possible to achieve any individually rational and feasible payoff, including payoffs that are obtained through correlated mixed strategies.

Peters and Szentez (2009) explore games in which instead of a commitment device, each player writes a contract that specifies what action to be played given any of the opponent's contracts. Peters and Szentez go a step further and prove a similar folk theorem for any (finite) number of players. Moreover, they show that in an environment with incomplete information, this result does not hold. That is, there are payoff vectors that can only be obtained by a centralized mechanism designer.

The difference between the type of models mentioned above and the current one can be best seen in Fershtman et al. (1991). In their model, players can use agents to play on their behalf. If this delegation is done by an observable contract, they show cooperation can emerge and prove a full folk theorem.

In all these models, the players don't actually play the game themselves - the game is played by a computer program, a commitment device or an agent. Whereas in our game, the players themselves play, without the aid of an additional factor. Another important difference is that while the computer programs or contracts are completely visible, we assume noisy signals. That is, with some probability players might not know the strategy of their opponent.

2.4 Informal Models

One informal commitment model that does not involve any external mechanisms is Frank's (1998) Commitment Model. In this model the commitment devices are emotions. It is argued that feelings, such as love, anger or revenge, can sometimes make people act in way that are not in their best interests. Hence, a person's feelings commit him to act in a certain way. Since psychological research shows that emotions are both observable and hard to fake (see Frank (1988) and references within), an agent can use them as signals in a game. This enables each player to discern his opponents emotional predispositions through physical and behavioral clues, and play accordingly.

Frank reports some experimental results that show that when subjects are allowed to interact for 30 minutes before playing the PD, they are able to predict quite accurately their opponent's behavior. Moreover, roughly 84% of the subjects that predict that their opponent will cooperate (defect), cooperate (defect) themselves. He also reports that when players are allowed to interact only for 10 minutes, or when not allowed to make any promises, the level of cooperation drops, as does the accuracy of the predictions.

A similar approach is explored by Gauthier (1986). He proposes an environment in which there are two types of agents: straightforward maximizers (SM) and constrained maximizers (CM). SM simply maximize their utility. CM, however, are more sophisticated. They take into account also the utilities of the other players and base their actions on a joint strategy: "A CM is conditionally disposed to cooperate in ways that, followed by all, would yield nearly optimal and fair outcomes, and does cooperate in such ways when she may actually expect to benefit". Gauthier assumes that an agent's type is known to everybody else (or at least with some positive probability). Thus, in the PD, when a CM meets another CM, they will both cooperate. In any other interaction between two players, they will both defect.

These last two models are very similar to the model we propose. The main contribution of this paper is to take these ideas and to incorporate them in a formal model.

2.5 Non Simultaneous Models

A different line of research that relates to this work is one in which players do not play simultaneously. That is, one player chooses a strategy and only then the second chooses his, conditional on the first's.

One example for such a model is Howard's (1971) Metagame model. A 2-Metagame is a game in which player 1 chooses a "regular" strategy (an action), while player 2 chooses a function from player 1's actions to his own action space. For instance, in the PD, player 1 can either play C or D, and player 2 can play CC, DD, CD, DC where the first letter describes the action he plays if player 1 plays C and the second is the action to be played if player 1 plays D. The strategy CD can be interpreted as "I will cooperate if, and only if, you will". However, (C, CD) is not an equilibrium since given the fact that player 1 plays C, player 2 will deviate to DD.

Similarly, a 1-2-Metagme is a game in which player 2's strategies are functions from player 1's actions to his own, and player 1's strategies are functions from player 2's strategies, as just defined, into actions. In the PD example, since player 2 has 4 strategies, player 1 now has 16. Interestingly, now $(DDCD, CD)$ is an equilibrium yielding cooperation by both players. In Howard's words: "Player 2 says "I'll cooperate if you will" (implying "not if you won't," i.e., the policy CD), and 1 replies "in that case (meaning if CD is your policy) I'll cooperate too (implying "not otherwise,", i.e., the policy $DDCD$).".

Another model of that nature is Solan and Yariv's (2004) model of espionage. In their model, first player 2 chooses an action and then player 1 can purchase information about that strategy. In the next step player 1 receives some signal with some probability about the action player 2 had chosen. Finally, depending on the information he received, player 1 chooses an action. Solan and Yariv provide a few examples for games with espionage and prove the existence of an equilibrium with espionage under certain restrictions.

There are two main differences between the espionage model and this paper, in addition to the fact that the strategies are not chosen simultaneously. First, in Solan and Yariv only one player can obtain information about the other player's strategy, while in this model both can. Second, in our model the players bear no cost in obtaining the information, unlike the costly espionage.

Espionage also appears in Matsui (1989), but in a setting of an infinitely repeated game. In the beginning, each player chooses a strategy for the entire game. Then, information may leak so that one of the players might be informed of the other player's strategy. If this happens, the player who received the information (and only him) may revise his strategy. Then, the actual repeated game begins. The fact that the game repeats enables the player who acquired the information to signal this fact to his opponent, and then the two players can switch to Pareto efficient strategies. Matsui shows that any subgame perfect equilibrium pair of payoffs is Pareto efficient as long as the probability of information leakage is small enough. He also illustrates this in the PD, and shows that full cooperation can be achieved.

Matsui's model resembles the current model in that the information about one's strategy may be used by the other player. Other then that, the models a very different (one shot vs. repeated game, one sided espionage vs. two sided, simultaneous choice of strategies vs the possibility to revise one's strategy after obtaining the information).

2.6 Additional Experimental Evidence

The prisoner's dilemma in general, and the affect of communication on cooperation in particular, have been the subject of many experiments in past half a century. Sally (1995) has conducted a meta-analysis of experiments from 1958 to 1992. He had combined data from 37 different experiments and showed that communication increases the rate of cooperation by roughly 40%. Interestingly, communication was one of a very small group of variables that has a significant affect on cooperation (see Sally (1995) and references within).

Kalay et. al. (2003) have analyzed data obtained from a TV game, similar to the PD, in which two players accumulate together a substantial amount of money, and then have to divide it between them. The division process is as follows: the players are allowed to communicate which each other for several minutes, and afterwords each has to choose an action - cooperate or defect .If they both cooperate, each obtains half of the money they accumulated. If one cooperates and the other defects, the one that defected receives everything and the other nothing. In case both defect, both receive nothing. Like in the PD, the dominant strategy is to defect. However, 42% of the time the players cooperated. Moreover, the data reveals a correlation between the actions chosen by the two players: 21% of the time both players cooperated, compared to 17.64% if there had been no correlation. This implies a correlation coefficient of 0.14.

As mentioned earlier, these results cannot be explained by traditional game theory, nor by cheap talk or correlated equilibrium. Moreover, these results are obtained without a mediator, such as a contract, commitment device or a computer program. This paper provides a model that may explain such results in a simple environment where no external mechanism are use, other that communication.

3 The Real Talk Model

Let $G =< A, \Pi >$ be a two-person game in normal form where:

- $A = A_1 \times A_2$: A_i is a non-empty finite set of actions for player i, $(i = 1, 2)$.
- $\Pi = \Pi_1 \times \Pi_2$: Π_i is the payoff function for player i, $\Pi_i : A_1 \times A_2 \to \mathbb{R}$ $(i = 1, 2)$.

Let α_i be a mixed action for player i, i.e., a probability distribution over A_i, and $\Pi_i(\alpha_1, \alpha_2)$ be the expected payoff for player i when the mixed actions (α_1, α_2) are played. A mixed action Nash equilibrium in G is a pair of mixed actions, (α_1^*, α_2^*) such that neither player can increase his expected payoff by deviating to another (mixed) action. Formally:

Definition 1 *A mixed action Nash equilibrium in G is a pair of mixed actions, (α_1^*, α_2^*) such that for $i = 1, 2$: $\Pi_i(\alpha_i^*, \alpha_{-i}^*) \geq \Pi_i(\alpha_i, \alpha_{-i}^*)$ for any $\alpha_i \in \Delta(A_i)$.*

A game with real talk, $\widehat{G}$, that is induced by the game G, consists of 3 stages and is played as follows:

1. Both players choose a strategy simultaneously.
2. Each player may, or may not, observe his opponent's chosen strategy.
3. Each player uses its own strategy to choose an action in G.

The strategies that the players choose in the first stage will determine what action they play in the last stage given the information they have, or have not, learned in the middle stage. Note that the strategies themselves are fixed, and cannot be changed once chosen. Formally:

Definition 2 *A pure strategy for player i in the game with Real Talk $\widehat{G}$ that is induced by G is a function from $S_{-i} \cup \phi$ to $\Delta(A_i)$, where S_{-i} is the opponent's strategy set, ϕ represents learning nothing and $\Delta(A_i)$ is a probability distribution over A_i. That is,*

$$S_i \subseteq \{f : S_{-i} \cup \phi \to \Delta(A_i)\}$$

Since $s_i \in S_i$ is a function, $s_i(s_{-i})$ denotes the action that player i plays when receiving the signal s_{-i}, and $s(\phi)$ denotes the action player i plays when receiving no signal.

Proposition 1 *For at least one player $S_i \neq \{f : S_j \cup \phi \to \Delta(A_i)\}$.*

Proof. Otherwise $S_1 = |\Delta(A_1)|^{|S_2|+1}$ and $S_2 = |\Delta(A_2)|^{|S_1|+1}$ which is impossible by Cantor's Theorem. ■

Note that:

1. It is possible to construct finite strategy spaces, which contain as few as just one strategy for each player. For example, $S_i = \{s_i\}$ where for each player, s_i is a strategy that always plays some pure action $a_i \in A_1$.

2. It is possible to construct infinite strategy spaces such that S_i includes all function from $S_{-i} \cup \phi$ to $\Delta(A_i)$ that can be described in finite sentences (using Gődel encoding). See Peters and Szentes (2009).

We assume that the probability that each player observes his opponent's strategy in the second stage is the same for both players, and that the two probabilities may be correlated.

Definition 3 *The game with Real Talk, $\widehat{G}$, that is induced by G, is a tuple (G, S, p, ρ) where:*

- *G is a two-person game in normal form.*
- *$S = S_1 \times S_2$ are the sets of the players' feasible pure strategies.*
- *$p \in [0,1]$ is the probability that each player observes the other player's strategy.*
- *$\rho \in [0,1]$ is the correlation coefficient between the two probabilities.*

Note that even though technically the correlation coefficient could be also negative, under our interpretation it makes less sense: if one player detects his opponent's strategy in a conversation, then it usually increases the probability that the opponent detects his, rather than decreases it.

After a strategy profile is chosen, there are four possibilities for the information the players have: both player receives a signal, player 1 receives a signal and 2 does not, player 2 receives a signal and 1 does not, and none of the players receives a signal. These four cases can be regarded as states of nature in games with incomplete information. The following table shows the probabilities for the four cases:

	signal	no signal
signal	$p^2 + \rho p\,(1-p)$	$p(1-p) - \rho p\,(1-p)$
no signal	$p(1-p) - \rho p\,(1-p)$	$(1-p)^2 + \rho p\,(1-p)$

As would be expected, correlation increases the probabilities along the main diagonal, and decrease those in the secondary diagonal, by $\rho p\,(1-p)$.

If a strategy profile (s_1, s_2) is chosen by the two players, each one plays one of two possible actions, according to the signal he receives. The action profiles for the four different possibilities are shown in the following table:

	signal	no signal
signal	$(s_1(s_2), s_2(s_1))$	$(s_1(s_2), s_2(\phi))$
no signal	$(s_1(\phi), s_2(s_1))$	$(s_1(\phi), s_2(\phi))$

Let $\widehat{\Pi}_i(s_1, s_2)$ be the expected payoff for player i if the strategies chosen by player 1 and player 2 are s_1 and s_2, respectively. Using the above two tables and the action payoff function, Π, we obtain:

$$\widehat{\Pi}_i(s_1, s_2) =$$

$$= \left[p^2 + \rho p\,(1-p)\right] \cdot \Pi_i(s_1(s_2), s_2(s_1)) + \left[p(1-p) - \rho p\,(1-p)\right] \cdot \Pi_i(s_1(s_2), s_2(\phi)) +$$
$$+ \left[p(1-p) - \rho p\,(1-p)\right] \cdot \Pi_i(s_1(\phi), s_2(s_1)) + \left[(1-p)^2 + \rho p\,(1-p)\right] \cdot \Pi_i(s_1(\phi), s_2(\phi)).$$

As mentioned before, the simplest possible strategies are constant mixed actions, i.e. ones that always play an action $a_i \in \Delta(A_i)$ regardless of what player i learns about player j's strategy. If both players choose such strategies, then their payoff would be: $\widehat{\Pi}_i(s_1, s_2) = \Pi_i(a_1, a_2)$.

Definition 4 *We say that a strategy space S is natural if it contains all constant mixed actions. That is, for any action $\alpha_i \in \Delta(A_i)$, the strategy set S_i contains a strategy that always plays the (mixed) action α_i, regardless of the opponent's strategy, for $i = 1, 2$.*

Similarly to other models, the solution concept we will use is a Bayesian Nash equilibrium, which is simply a Nash equilibrium in our strategies space. Each player responds optimally to his opponent's strategy choice, taking into account p and ρ, and maximizes his expected payoff over his possible strategy choices. Formally:

Definition 5 *A pure strategy Nash equilibrium in the game with Real Talk $\widehat{G}$ is a pair of strategies (s_1^*, s_2^*) such that $\widehat{\Pi}_1(s_1^*, s_2^*) \geq \widehat{\Pi}_1(s_1, s_2^*)$ for any $s_1 \in S_1$ and $\widehat{\Pi}_2(s_1^*, s_2^*) \geq \widehat{\Pi}_2(s_1^*, s_2)$ for any $s_2 \in S_2$.*

The following proposition follows immediately from these definitions and requires no proof:

Proposition 2 *If S is natural then the players in the Real Talk game, $\widehat{G}$, that is induced by G, can play the original game G. That is,*

1. *Every strategy in G has a corresponding (constant) strategy in $\widehat{G}$.*
2. *Any mixed action Nash equilibrium in the original game G, (α_1^*, α_2^*), has a corresponding Nash equilibrium in the game with Real Talk $\widehat{G}$, (s_1^*, s_2^*), such that $s_i^* \equiv \alpha_i^*$.*

Note that this does not imply that all Nash equilibria in $\widehat{G}$ are in constant strategies.

4 A Real Talk Folk Theorem

Let G be a two person game and let v_i be the minmax value for player i G, i.e.,

$$v_i = \min_{\alpha_{-i} \in \Delta A_{-i}} \max_{a_i \in A_i} \Pi_i(a_i, \alpha_{-i}).$$

Definition 6 *We say that a payoff for player i, π_i, is Individually Rational if $\pi_i \geq v_i$.*

Let ψ_i be player i's minmax strategy. That is, when ψ_i is played, player $-i$ can achieve a payoff of at most v_{-i}. Formally, ψ_i is a member of $\arg\min_{i \in A_i} \max_{a_{-i} \in A_{-i}} \Pi_{-i}(a_i, a_{-i})$.

Let α_i be a probability distribution over A_i, and $\alpha_i(a_i)$ the probability of the action a_i, determined by α_i.

Definition 7 *We say that a payoff profile (π_1, π_2) is feasible if $(\pi_1, \pi_2) = \sum_{a_1 \in A_1} \sum_{a_2 \in A_2} \alpha_1(a_1)\alpha_2(a_2)\Pi(a_1, a_2)$.*

Note that this definition is not standard because of independent mixing of actions, and it does not always coincide with the standard definition of a feasible payoff profile. (The standard definition requires (π_1, π_2) to be a convex combination of all outcomes in G, that is $\sum_{\alpha \in A} \alpha(a_1, a_2)\Pi(a_1, a_2)$ where α is a probability distribution over the joint action space A.)

For example, consider the following payoff matrix:

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

In the standard definition, a feasible payoff profile is a convex combination of these four payoffs, and the set of all feasible payoffs is the unit square. However, according to our definition the feasible payoffs are depicted in the following graph:

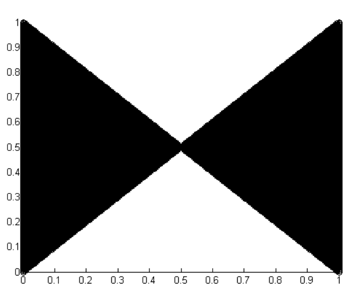

which is only a subset of the unit square.

Proposition 3 *(Folk Theorem) For any game G there exists a game with Real Talk, $\widehat{G}$, such that any individually rational and feasible payoff profile (π_1, π_2) in G is the payoff profile of some pure Nash equilibrium in $\widehat{G}$.*

Proof. Let (π_1, π_2) be an individually rational and feasible payoff profile. Let α_1 and α_2 be probability distributions over A_1 and A_2, respectively, for which $(\pi_1, \pi_2) = \sum_{a_1 \in A_1} \sum_{a_2 \in A_2} \alpha_1(a_1)\alpha_2(a_2)\Pi(a_1, a_2)$ (note that there can be more then one pair of probability distributions yielding the same payoff). Define $s_1^{(\pi_1,\pi_2)}$ and $s_2^{(\pi_1,\pi_2)}$ in the following way:

$s_1^{(\pi_1,\pi_2)}(s_2^{(\pi_1,\pi_2)}) = \alpha_1$ and ψ_1 otherwise (that is, if player 1 receives a signal which is $s_2^{(\pi_1,\pi_2)}$, he plays α_1. If he receives a different signal, or no signal at all, he plays ψ_1). In a similar way, define $s_2^{(\pi_1,\pi_2)}(s_1^{(\pi_1,\pi_2)}) = \alpha_2$ and ψ_2 otherwise.

Let S_1 and S_2 be arbitrary mutually consistent strategy sets such that $s_1^{(\pi_1,\pi_2)} \in S_1$ and $s_2^{(\pi_1,\pi_2)} \in S_2$ for all individually rational and feasible payoffs (π_1, π_2). Note that if $S = S_1 \times S_2$ contains only these strategies it is mutually consistent. However, it can be much larger. A trivial example is including all constant strategies, making S natural.

Let $\widehat{G} = (G, S, 1, 1)$. That is $\widehat{G}$ is the Real Talk game that is induced from G when $S_1 \times S_2$ is the strategy set, $p = 1$ and $\rho = 1$.

The strategy profile $(s_1^{(\pi_1,\pi_2)}, s_2^{(\pi_1,\pi_2)})$ is a Nash equilibrium in $\widehat{G}$ for any (π_1, π_2). To see this, assume that player 2 plays $s_2^{(\pi_1,\pi_2)}$. If player 1 plays $s_1^{(\pi_1,\pi_2)}$ then the players will play (α_1, α_2), yielding a payoff of π_1 for player 1. If player 1 deviates to any other strategy, player 2 will play ψ_2 against him, yielding player 1 a payoff of no more then v_1. However, since (π_1, π_2) is individually rational, $\pi_1 \geq v_1$ and therefore π_1 is at least as good as player 1's payoff if he chooses to deviate. The same argument holds for player 2. Since no player has an incentive to deviate, $(s_1^{(\pi_1,\pi_2)}, s_2^{(\pi_1,\pi_2)})$ is a Nash equilibrium. ■

Consider for example the following game matrix:

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

The set of feasible payoffs according to standard definition is:

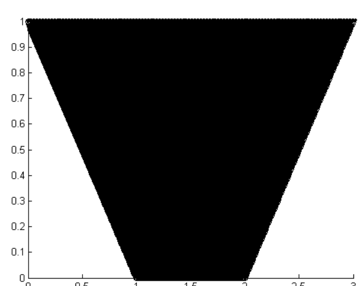

But according to our definition, it is only this subset:

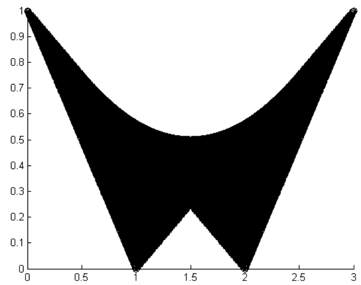

Clearly, in this game T is the dominant strategy for player 1, and if he does indeed play T, then player 2 will play L. Therefore, (T, L) is the only Nash equilibrium in this game, and the payoffs are $(3, 1)$.

The minmax values are $v_1 = 1$ and $v_2 = \frac{1}{2}$. The minmax strategies (minimizing the opponent's payoffs) are playing T and B with probability $\frac{1}{2}$ for player 1 and R for player 2.

Hence the set of payoffs that can be supported in a Real Talk Nash equilibrium (feasible and individually rational) is:

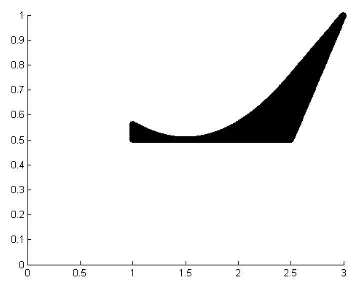

which of course includes the 'regular' Nash equilibrium payoff profile.

Note that the regular Nash equilibrium is Pareto optimal - it yields both player the maximal possible payoffs. All the other Real Talk equilibria actually lower the payoffs, and since they are also more complicated, they seem not to make much sense. However, an example without this problem can be easily obtained if one reverses the payoffs for player 2:

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

In this case the only Nash equilibria is (T, R) and the payoffs are $(1, 1)$. The set of feasible outcomes is:

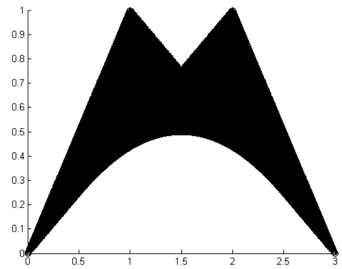

and the set of payoffs that can be supported in a Real Talk Nash equilibrium is:

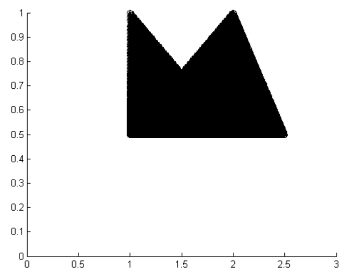

Clearly, this set includes a payoff vector that weakly dominates the original one, $(2, 1)$, as well as other payoffs that increase the payoff of player 1 but decrease player 2's.

If we require the induced Real Talk game $\widehat{G}$ to have $p < 1$, then the proposition doesn't hold since the threat of punishing a deviation is weaker. In the extreme case, when $p = 0$, no signals are ever received, which means players can't condition their actions on their opponent's strategy. In other words, the players simply choose a probability distribution over the actions in G. This makes the game with Real Talk $\widehat{G}$ very similar to the original game G, even though a strategy is still a function from the opponent's strategies to actions, rather than simply an action. In some cases the game with Real Talk $\widehat{G}$ becomes equivalent to G:

Proposition 4 *Any game with Real Talk, $\widehat{G}$, that is induced by G, that consists of a natural strategy set S and $p = 0$, is equivalent to the original game G.*

Proof. If $p = 0$ and a player chooses a strategy s, then the only possible input that the strategy would receive is ϕ, and the action that will be played in G is $s(\phi)$ with probability 1 (players play $s(\phi)$ regardless of the opponent's strategy). Therefore, by choosing a strategy all the players do is simply choose a probability distribution over their own action space, A_i. Since S is natural, for every probability distribution in A_i, player i has a constant strategy that always plays it (of course, there may be many other strategies that play the same mixed action when receiving no signal, but they are all equivalent in this case). Hence, strategically the players face exactly the same choices in $\widehat{G}$ as they do in G. Clearly, the feasible payoff profiles and Nash equilibria are the same in the two games. ∎

This proposition is very similar to proposition 1, except for the fact that in the general case, if S is natural, then the players *can* play the original game G, whereas when $p = 0$ they have no other choice.

Note that if $p = 0$ and S is *not* natural, then each player may have a smaller set of actions, which is the union of all probability distributions that are played by his strategies when they do not receive a signal. That is: $\cup_{s \in S_i} s(\phi)$. In this case, instead of adding more sophistication to the game, Real Talk may actually produce a more degenerate version of G, depending on S_i.

When $0 < p < 1$ things become more complicated. The set of feasible payoff profiles and equilibria depend on the exact values of p and ρ (which plays no role when $p \in \{0, 1\}$), and on the game itself. However, the following two propositions hold:

Proposition 5 *(Reverse Folk Theorem 1) Let $\widehat{G}$ be a game with Real Talk that consists of a natural strategy space S. If (π_1, π_2) is the payoff profile of a pure Nash equilibrium in $\widehat{G}$, then it is individually rational.*

Proof. Suppose that (π_1, π_2) is the payoff profile of a pure Nash equilibrium in $\widehat{G}$ that is not individually rational. Then for some i, $\pi_i < v_i$. Consider a deviation of player i to a constant strategy which plays his maxmin strategy in G. By definition, his payoff would be at least as high as his maxmin payoff. But since the maxmin and minmax payoffs are identical, his payoff would be at least v_i. Since a deviation is profitable, the original payoff profile was not an equilibrium. ■

If the game is not natural, however, the strategy sets might not be rich enough to enable the players to deviate to a strategy that would induce the minmax payoff. As a trivial example we can consider the game mentioned above, where S_1 contains only the constant strategy T and S_2 contains only the constant strategy L. The only possible payoff is $(3, 0)$. This is a Nash equilibrium but not individually rational for player 2. Naturally, this is an extremely degenerate example, but more elaborate examples that have this property may be easily created. If, however, the players do have a constant strategy that plays their maxmin strategy in G, then the proposition holds even if S is not natural.

Proposition 6 *(Reverse Folk Theorem 2a) For any game with Real Talk $\widehat{G}$ in which $p = 1$, if (π_1, π_2) is the payoff profile of some pure Nash equilibrium in $\widehat{G}$, then it is feasible.*

Proof. When $p = 1$ both players always detect their opponent's strategies. Each player plays the action determined by the strategies with probability 1. Since the players' actions are simply a probability distribution over their own action spaces, they are independent. By definition, this induces a feasible payoff profile. ■

Proposition 7 *(Reverse Folk Theorem 2b) For any game with Real Talk $\widehat{G}$ in which $\rho = 0$, if (π_1, π_2) is the payoff profile of some pure Nash equilibrium in $\widehat{G}$, then it is feasible.*

Proof. When $\rho = 0$ the signals the players receive are independent. Since the players' actions are simply a probability distribution over their own action spaces, they are independent. Hence, each player plays an independent lottery conditioned on an independent signal. These are compound independent lotteries, which in turn are also independent. By definition, this induces a feasible payoff profile. ■

When $\rho \neq 0$ this argument does not hold, since the players condition their (independent) lotteries on a correlated signal, which induces a correlated compound lottery. And indeed, there are games and payoff profiles, that are obtained in equilibrium, which are not feasible according to our definition (clearly, any payoff profile is feasible according to the standard definition).

To see that, consider the game usually referred to as the 'battle of the sexes':

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

This game has three Nash equilibria: (T,L), (B,R), and $((\frac{1}{3},\frac{2}{3}),(\frac{2}{3},\frac{1}{3}))$. The payoffs are $(1,2)$, $(2,1)$ and $(\frac{2}{3},\frac{2}{3})$, respectively. The set of feasible payoffs is:

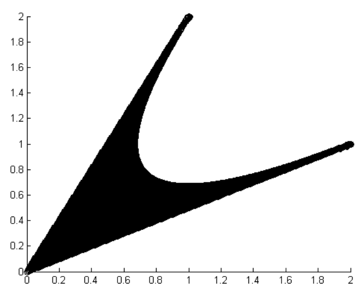

The minmax value for both players is $\frac{2}{3}$, hence the set of feasible and individually rational payoffs is:

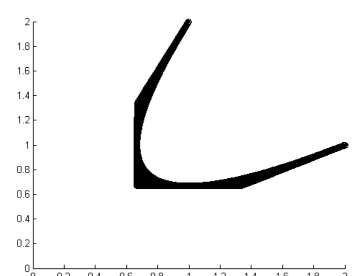

Let $\widehat{G}$ be the induced Real Talk game with $p \in [0,1]$ and $\rho = 1$. In this case *both* players see each other with probability p, and with probability $(1-p)$ *neither one of them* does. Consider the pair of strategies (s_1, s_2) defined below:

$s_1(s_2) = T, s_1(\phi) = B$ and otherwise $s_1(\cdot) = T$.

$s_2(s_1) = L, s_2(\phi) = R$ and otherwise $s_2(\cdot) = L$.

Let S be a strategy space that contains (s_1, s_2). If (s_1, s_2) is played, then the payoff profile is $p \cdot (1,2) + (1-p) \cdot (2,1)$. This is a point on the line connecting $(1,2)$ and $(2,1)$, which clearly is not feasible unless $p \in \{0.1\}$. However, this is indeed a Nash equilibrium. Assume player 2 plays s_2 and consider a possible deviation for player 1. If he deviates to a strategy that plays differently when receiving a signal, then he would play B with some positive probability while player 2 continues to play L with probability 1. If he deviates to a strategy that plays differently when not receiving a signal, then he would play T with some positive probability while player 2 continues to play R with probability 1. In both cases, his payoff would be strictly lower. The same argument holds for player 2.

This example shows us that when p is less than 1 and there is a positive correlation between the signals the players receive, they can use the signal to correlate between their actions and obtain higher payoffs than they could obtain otherwise (in this case, as high as 1.5 each when $p = \frac{1}{2}$).

5 Cooperation in the Prisoner's Dilemma

The general form of a Prisoner's Dilemma payoff matrix is:

	C	D
C	b, b	d, a
D	a, d	c, c

where $a > b > c > d$. However, without loss of generality, one can subtract d from all payoffs and divide by c in order to obtain a matrix of the form:

	C	D
C	b, b	$0, a$
D	$a, 0$	$1, 1$

where $a > b > 1$. We will consider the last version as the general case of the Prisoner's Dilemma.

Clearly, the only (mixed) action Nash equilibrium in this game is (D, D) regardless of the exact values of a and b. However, for reason that will become clear shortly, assume that the players use mixed actions: player 1 plays C with probability $x \in [0,1]$ and player 2 plays C with probability $y \in [0,1]$. The expected payoff for player 1 is therefore:

$$\Pi = x \cdot y \cdot b + (1-x) \cdot y \cdot a + (1-x) \cdot (1-y).$$

If we rearrange this equation we obtain:

$$\Pi = x \cdot y \cdot (b - a + 1) - x + y \cdot (a-1) + 1.$$

Note that $\frac{\partial \Pi}{\partial x} = y \cdot (b - a + 1) - 1$ and $\frac{\partial \Pi}{\partial y} = x \cdot (b - a + 1) + (a - 1)$. It is easy to verify that under the assumptions on the values of a and b, $\frac{\partial \Pi}{\partial x} < 0$ and $\frac{\partial \Pi}{\partial y} > 0$, which means that player 1 is interested in choosing x as low as possible, and would like player 2 to choose y as high as possible. Moreover, given the value of y, the incentive of player 1 to reduce x depends on the value of $b - a + 1$: the higher this value is, the less player 1 has to loose by playing cooperatively. The same is true for player 2, and thus the value $b - a + 1$ can be seen as the strength of the incentive to play cooperatively (or the inverse of the temptation to defect), and will be denoted hereinafter by c.

5.1 Equilibria and possible payoffs

In this section we discuss the criteria for an equilibrium in the PD and show what are the possible payoffs in equilibria, using a few examples.

The following two observations stem directly from the payoff matrix:

1. The mimax action for both players is D. If a player deviates to it, his opponent's payoff is at most 1.

2. Assuming that S is natural, each player can guarantee a payoff of 1 by choosing the constant strategy that always plays D. Let d be such a strategy.

For any strategy profile (s_1, s_2), let (s_1^d, s_2^d) be strategies such that $s_i^d(s_{-i}^d) = s_i(s_{-i})$, $s_i^d(\phi) = s_i(\phi)$ and otherwise $s_i^d(\cdot) = D$. In words, when playing against each other, (s_1^d, s_2^d) play exactly like (s_1, s_2). However, against any other strategy, they play the minmax action D. Clearly, the payoffs when (s_1, s_2) and (s_1^d, s_2^d) are played are exactly the same.

Proposition 8 *If the strategy set is natural and (s_1, s_2) is an equilibrium, then (s_1^d, s_2^d) is also an equilibrium in any strategy set that contains it.*

Proof. By (s_1, s_2) being an equilibrium, $\widehat{\Pi}_1(d, s_2) \leq \widehat{\Pi}_1(s_1, s_2)$ and $\widehat{\Pi}_2(s_1, d) \leq \widehat{\Pi}_2(s_1, s_2)$. By the definition of (s_1^d, s_2^d), $\widehat{\Pi}_1(d, s_2^d) \leq \widehat{\Pi}_1(d, s_2)$ and $\widehat{\Pi}_2(s_1^d, d) \leq \widehat{\Pi}_2(s_1, d)$. Thus, $\widehat{\Pi}_1(d, s_2^d) \leq \widehat{\Pi}_1(s_1, s_2)$ and $\widehat{\Pi}_2(s_1^d, d) \leq \widehat{\Pi}_2(s_1, s_2)$. But $\widehat{\Pi}_1(s_1, s_2) = \widehat{\Pi}_1(s_1^d, s_2^d)$ and $\widehat{\Pi}_2(s_1, s_2) = \widehat{\Pi}_2(s_1^d, s_2^d)$, which means that when (s_1^d, s_2^d) is played, neither player has an incentive to deviate to d. What remains to be shown is that deviating to d is the most profitable deviation. This completes the proof, since if the players do not have an incentive to deviate to the most profitable deviation, they do not have an incentive to deviate at all, which means (s_1^d, s_2^d) is in fact an equilibrium.

With out loss of generality, assume that player 2 plays the strategy s_2^d and that player 1 deviates to some strategy. If player 2 receives a signal, the deviation is detected, which results in player 2 playing D, regardless of the chosen deviation. If player 2 does not receive a signal, the deviation is not detected, and player 2's action is not at all affected by the deviation. Since in both cases, all deviations result in the same action played by player 2, playing the strictly dominant action D is optimal. $\blacksquare$

If we are interested only in what payoffs can be sustained in equilibrium, this proof allows us to restrict our attention only the strategies of the type (s_1^d, s_2^d). Each strategy, s_1^d and s_2^d has to specify the action to be played against each other, and also what action to play when not receiving a signal. Since there are only two pure actions in this game, each action is a probability distribution over (C, D) and it can be described as a number in the interval $[0, 1]$. In total there are four such probabilities. Thus, the space of all strategies (s_1^d, s_2^d) is equivalent to $[0, 1]^4$.

In what follows we used a computer simulation in order to check which strategies (s_1^d, s_2^d), that is what vectors in $[0, 1]^4$, are an equilibrium, and draw the set of all payoffs that can be achieved in these equilibria. Since d is the most profitable deviation, checking that the players loose by deviating to it is a sufficient condition for the optimality of playing s_1^d against s_2^d. and vice versa. Clearly this condition is also necessary if the strategy set is natural, which we assume. We do not construct a full strategy space that contains (s_1^d, s_2^d). This is not necessary in order to prove that (s_1^d, s_2^d) is an equilibrium because of the way the strategies are defined (they play D against any other strategy).

Assume, for example, that $a = 4$ and $b = 3$. The set of feasible payoffs is:

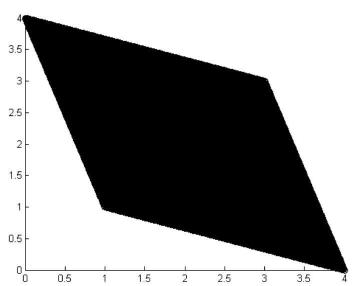

Note that in this case the standard definition and ours coincide. The set of payoffs which are also individually rational is:

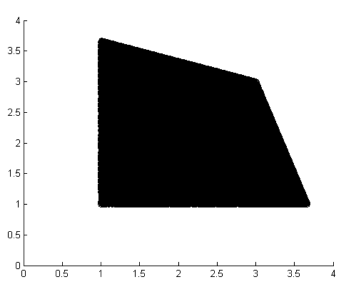

When $p = 1$ this is also the set of payoffs that can be obtained in a Nash equilibrium. As a matter of fact, this remain to be the equilibria payoff set for any $p \geq \frac{1}{3}$ and for any ρ. For $p < \frac{1}{3}$, however, the set is smaller. For instance, when $p = \frac{1}{10}$ it becomes :

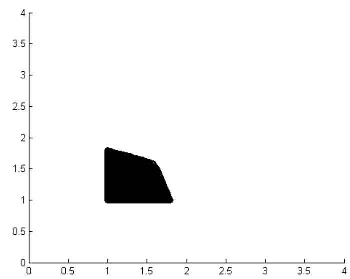

regardless of the value of ρ.

The set of feasible and individually rational payoffs looks similar to the one in the above example as long as $a < \frac{b}{2}$. Therefore, by the folk theorem, when $p = 1$ the set of equilibria payoffs also looks qualitatively the same. However, when p and ρ change, the set of equilibria can look quite different, depending on the exact values of a and b. Take for example $a = 4$ and $b = 3.9$. When $p = \frac{1}{10}$ and $\rho = 0$. the set is:

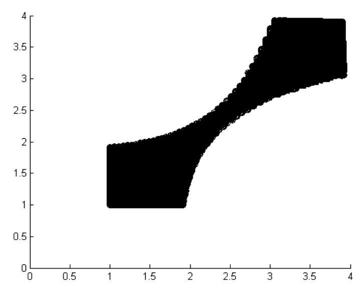

When $a > \frac{b}{2}$ the set of feasible payoffs is different from the one before. The following example, when $a = 10$ and $b = 3$, is typical:

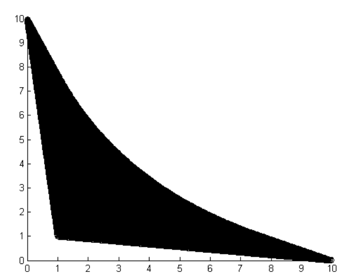

Here, under the standard definition, the feasible payoffs set is larger.

If we add the individual rationality restriction, it becomes the set of equilibria payoffs when $p = 1$:

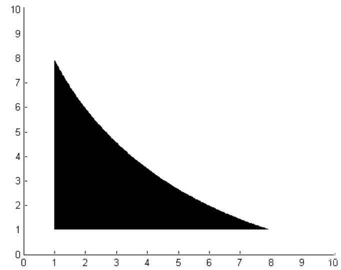

When p is smaller, the set of equilibria payoffs may also be smaller. For instance, for $p = 0.2$ and $\rho = 0$, the set is:

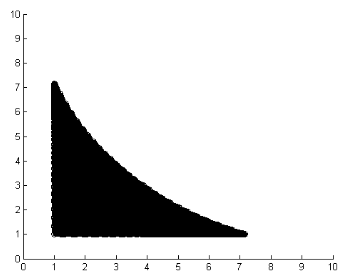

5.2 The Strategy Nice_q

We define the pure strategy *nice_q* in the following way: if the opponent's strategy is detected, *nice_q* plays C against the strategy *nice_q*, and D against any other strategy. In case it receives no signal, it plays C with probability q and D otherwise.

This is a slightly more complicated version of the reciprocal strategy s mentioned in the introduction. *nice_q* is "nice" as long as it detects that the opponent is nice as well. If, however, it detects that the other player is not "nice", then it punishes him by defecting. In the event that it doesn't receive a signal at all, it tosses a coin and cooperates with probability q.

It should be noted that although we interpret this strategy as "nice", it is only nice if the opponent plays exactly the same strategy. Therefore, if for example player 1 plays *nice_*$\frac{1}{2}$ and player 2 plays *nice_*$\frac{1}{3}$ the result would be that nobody will cooperate if they detect each other, even though they are both "nice". Clearly, it is possible to solve this problem by defining two other strategies such that when detecting each other they cooperate, and when they receive no signal, one cooperates with probability $\frac{1}{2}$ and the other with probability of $\frac{1}{3}$.

This disadvantage, the fact that *nice_q* reacts nicely only to one specific strategy, is also an advantage: once a pair of such strategies belongs to a strategy set S, more strategies can be added without the need of re-defining the strategy *nice_q*.

Going back to the symmetric case, assume that the two players choose the strategy *nice_q*. Consider the event "both players play C". This event is the union of the following three events:

1. Both players receive a signal about the opponent's strategy.

2. One player receives a signal and the other does not, but chooses to play C anyway.

3. Both players don't receives a signal but choose to play C nonetheless.

The corresponding probabilities for these events are:

1. $p^2 + \rho p\,(1-p)$

2. $2 \cdot q \cdot [p(1-p) - \rho p\,(1-p)]$ (Either player may be the one receiving the signal, hence the 2.)

3. $q^2 \cdot \left[(1-p)^2 + \rho p\,(1-p)\right]$

Since they are disjoint, the event "both players play C" occurs with probability equal to the some of these three probabilities. That is,

$$[p + q(1-p)]^2 + \rho p(1-p)(1-q)^2.$$

Similarly, the probabilities for the other two possible action profiles (which can be calculated in a similar way) are:

$$p(1-p)(1-q) + (1-p)q(1-p)(1-q) - \rho p(1-p)(1-q)^2$$

for one player playing C and the other D, and

$$[(1-p)(1-q)]^2 + \rho p(1-p)(1-q)^2$$

for both players playing D.

In addition to that, note that the marginal probability for each player to cooperate is $p+q(1-p)$ and to defect is $(1-p)(1-q)$.

By multiplying the payoffs of the game by these probabilities we obtain the expected payoff for each of the two players:

$$\begin{aligned}
\widehat{\Pi}_i(nice_q, nice_q) \quad = \quad & b \cdot \left[[p + q(1-p)]^2 + \rho p(1-p)(1-q)^2\right] \\
& +0 \cdot [p(1-p)(1-q) + (1-p)q(1-p)(1-q) - \rho p(1-p)(1-q)^2] + \\
& +a \cdot [p(1-p)(1-q) + (1-p)q(1-p)(1-q) - \rho p(1-p)(1-q)^2] + \\
& +1 \cdot \left[[(1-p)(1-q)]^2 + \rho p(1-p)(1-q)^2\right].
\end{aligned}$$

By rearranging and replacing $b - a + 1$ by c we obtain;

$$\begin{aligned}
\widehat{\Pi}_i(nice_q, nice_q) \quad = \quad & c \cdot [p + q(1-p)]^2 + (a-2)\,[p + q(1-p)] + 1 + \\
& +c \cdot \rho p(1-p)(1-q)^2.
\end{aligned}$$

Note that when holding the other parameters fixed, the higher the value of the incentive to cooperate, c, the higher the payoff for both players.

If, for example, a is chosen to be 4 and b to be 3, then this expression is reduced to the simple expression:

$$\widehat{\Pi}_i(nice\_q, nice\_q) = 2[p + q(1-p)] + 1.$$

5.3 Conditions for $(nice\_q, nice\_q)$ to be a Nash Equilibrium

In this section we find for which values of q the strategy profile $(nice\_q, nice\_q)$ is a Nash equilibrium. These calculation will become useful in the following section. For convenience, we analyze the conditions for $(nice\_q, nice\_q)$ to be a Nash Equilibrium from player 1's perspective. Since $(nice\_q, nice\_q)$ is a special case of (s_1^d, s_2^d), deviating to d is the most profitable deviation. In what follows we assume that $d \in S_1$ (Note that many other strategies that obtain the same payoff as d may exist . For example, a strategy that plays D if it detects $nice\_q$, C if it detects any other strategy, and D if it receives no signal. However, this, and other strategies of this type, are less plausible, tend to be more complicated and are not necessarily part of the strategy set S_1, whereas it is natural to assume that S_1 contains the constant strategy that plays always D, if not all constant strategies.)

Since d is the most profitable deviation, by checking that players loose by deviating to it, we can obtain a sufficient condition for the optimality of playing $nice\_q$ against $nice\_q$. Clearly this condition is also necessary.

In the example mentioned before ($a = 4$, $b = 3$), the payoff for the deviating player, if he deviates to playing d, is 4 in case he is not detected and the other player plays C, (probability of $(1-p)q$) and 1 otherwise:

$$\widehat{\Pi}_i(d, nice\_q) = 4(1-p)q + 1[p + (1-p)(1-q)] = 3(1-p)q + 1.$$

Thus, $(nice\_q, nice\_q)$ is an equilibrium iff

$$2[p + q(1-p)] + 1 \geq 3(1-p)q + 1.$$

That is, a sufficient condition for $(nice\_q, nice\_q)$ to be an equilibrium is $\frac{2p}{1-p} \geq q$. Since $q \leq 1$, for any p greater then $1/3$, playing $(nice\_q, nice\_q)$ is an equilibrium regardless of q (including the case $p = 1$).

In the general case, the payoff for the deviating player, if he deviates to playing d, is:

$$\widehat{\Pi}_i(d, nice\_q) = a(1-p)q + 1[p + (1-p)(1-q)].$$

Thus, $(nice\_q, nice\_q)$ is an equilibrium iff

$$c[p + q(1-p)]^2 + (a-2)\,[p + q(1-p)] + 1 + c\rho p(1-p)(1-q)^2 \geq a(1-p)q + [p + (1-p)(1-q)]$$

or:

$$c[p + q(1-p)]^2 + (a-2)\,p - q(1-p) + c\rho p(1-p)(1-q)^2 \geq 0.$$

Clearly, if for a certain set of parameters $(nice\_q, nice\_q)$ is an equilibrium, increasing the incentive to play cooperatively, c, does not reverse the inequality, and $(nice\_q, nice\_q)$ remains an equilibrium.

Since a, c, p and ρ are the parameters of the game, it is convenient to analyze this inequality as a polynomial of q:

$$\left((p-1)^2 + \rho p(1-p)\right) cq^2 + (1-p)\left(2cp\left(1-\rho\right) - 1\right) q + p\left(c\left(p + \rho(1-p)\right) + (a-2)\right) \geq 0.$$

In order to solve this inequality, we consider the following three cases:

1. $c = 0$:

The following linear inequality is obtained:

$$p\,(a-2) - (1-p)\,q \geq 0.$$

Hence, the condition for $(nice\_q, nice\_q)$ to be an equilibrium is $q \leq \frac{p(a-2)}{1-p}$. Since $q \leq 1$, this condition is satisfied for every q if $p \geq \frac{1}{a-1}$.

In order to simplify the computations, in what follows we assume $\rho = 0$, which implies the following condition for equilibrium:

$$(p-1)^2\, cq^2 + (1-p)\,(2cp-1)\,q + p\,(cp + (a-2)) \geq 0.$$

When $c \neq 0$ the quadratic equation $(p-1)^2\, cq^2 + (1-p)\,(2cp-1)\,q + p\,(cp + (a-2)) = 0$ may have two, one or no roots, depending on the discriminant:

$$[(1-p)\,(2cp-1)]^2 - 4\,(p-1)^2\, c \cdot p\,(cp + (a-2))\,.$$

2. $c < 0$:

Since by assumption $a > 1$ and p is not negative, we obtain that $p > \frac{1}{4c(a-1)}$. Therefore, the discriminant is positive and the equation:

$$(p-1)^2\, cq^2 + (1-p)\,(2cp-1)\,q + p\,(cp + (a-2)) = 0$$

has 2 real valued roots for any feasible parameters (a, c, p). Denote the smaller root by r_1 and the larger by r_2. Explicitly:

$$r_1(a,c,p) = \frac{1 - 2cp + \sqrt{4cp - 4acp + 1}}{2c - 2cp},\;\; r_2(a,c,p) = \frac{1 - 2cp - \sqrt{4cp - 4acp + 1}}{2c - 2cp}.$$

The condition for equilibrium holds for any $r_1 \leq q \leq r_2$. Since q denotes a probability, the relevant range for q is $[0, 1]$. It is possible to show that:

- $r_1 < 0$.

- $r_2 > 0$.
- $r_2 \geq 1$ iff $p \geq \frac{1-c}{a-1}$.

We can now check when $(nice\_q, nice\_q)$ is an equilibrium, depending on the possible locations of r_1 and r_2:

1. $r_1 < 0$ and $0 < r_2 < 1$:

$(nice\_q, nice\_q)$ is an equilibrium only for q smaller or equal to r_2.

For example, if $a = 5$, $c = -2$ and $p = 0.4$, then $(nice\_q, nice\_q)$ is an equilibrium only for $q \leq 0.464$.

2. $r_1 < 0$ and $1 \leq r_2$:

$(nice\_q, nice\_q)$ is an equilibrium for any q.

For example, if $a = 9$, $c = -2$ and $p = 0.4$.

3. $c > 0$**:**

If $p \geq \frac{1}{4c(a-1)}$ the discriminant is non-positive and the equation:

$$(p-1)^2 cq^2 + (1-p)(2cp-1)q + p(cp + (a-2)) = 0$$

has one or no solutions. Therefore, the condition for equilibrium always holds, which implies that $(nice\_q, nice\_q)$ is an equilibrium for any q.

However, if $p < \frac{1}{4c(a-1)}$, the equation has two real valued roots. As before, denote:

$$r_1(a,c,p) = \frac{1 - 2cp + \sqrt{4cp - 4acp + 1}}{2c - 2cp},\ r_2(a,c,p) = \frac{1 - 2cp - \sqrt{4cp - 4acp + 1}}{2c - 2cp}.$$

It should be noted that since the denominator is now positive, r_1 becomes the larger root.

In this case, the inequality hold for $q \leq r_2$ or $q \geq r_1$. Once again, since q denotes a probability, the relevant range for q is $[0,1]$. It is possible to show that:

- $r_1 \leq 0$ iff $p \geq \max\{\frac{1}{2c}, \frac{2-a}{c}\}$.
- $r_1 \geq 1$ iff, $c \leq 0.5$ or $p \leq \frac{1-c}{a-1}$.
- $r_2 \leq 0$ iff $p \geq \frac{1}{2c}$ or $p \leq \frac{2-a}{c}$.
- $r_2 \geq 1$ iff $c \leq 0.5$ and $p \geq \frac{1-c}{a-1}$.

We can now check when $(nice\_q, nice\_q)$ is an equilibrium, depending on the possible locations of r_1 and r_2:

1. $1 < r_1$ and $1 \leq r_2$:

   $(nice\_q, nice\_q)$ is an equilibrium for any q.

   For example, $a = 2$, $c = 0.4$ and $p = 0.61$.

2. $1 < r_1$ and $0 \leq r_2 < 1$:

   $(nice\_q, nice\_q)$ is an equilibrium only for q smaller or equal to r_2.

   For example, if $a = 4$, $c = 0.4$ and $p = 0.1$, then $(nice\_q, nice\_q)$ is an equilibrium only for $q \leq 0.28$.

3. $1 < r_1$ and $r_2 < 0$:

   $(nice\_q, nice\_q)$ is not an equilibrium for any q.

   For example, $a = 1.7$, $c = 0.4$ and $p = 0.4$.

4. $0 < r_1 \leq 1$ and $0 \leq r_2 < 1$:

   $(nice\_q, nice\_q)$ is an equilibrium only for q smaller or equal to r_2 or larger or equal to r_1.

   For example, if $a = 4$, $c = 0.75$ and $p = 0.1$, then $(nice\_q, nice\_q)$ is an equilibrium only for $0.4 \leq q \leq 0.86$.

5. $0 \leq r_1 \leq 1$ and $r_2 < 0$:

   $(nice\_q, nice\_q)$ is an equilibrium only for q larger or equal to r_1.

   For example, if $a = 1.5$, $c = 0.9$ and $p = 0.33$, then $(nice\_q, nice\_q)$ is an equilibrium only for $q \geq 0.86$.

6. $r_1 < 0$ and $r_2 < 0$:

   $(nice\_q, nice\_q)$ is an equilibrium for any q.

   For example, $a = 1.4$, $c = 0.9$ and $p = 0.67$.

As can be seen in the examples above, each of these six cases is obtained by some vector of parameters.

5.4 Achieving maximal cooperation

This section discusses the probability of the event that both players cooperate, i.e. the event that both players play the action C, in a Real Talk Nash equilibrium. We refer to this probability as the probability for cooperation.

Given a PD game G (i.e. the values of a and b) and the values of p and ρ, denote the maximal probability for cooperation in a symmetric Real Talk Nash equilibrium by $P_{\max}$. That is, in any Real Talk game $\widehat{G}$ that is induced by G there is no strategy set S and a strategy $(s, s) \in S$ that yields a probability for cooperation that is higher than $P_{\max}$.

Denote the maximal probability that each player plays C (separately) in a symmetric real talk Nash equilibrium of the Prisoner's Dilemma by P^*. If $\rho = 0$ (that is, when there is no correlation between the signals the players receive), then $P_{\max} = P^{*2}$.

In what follows we find the value of $P_{\max}$ as a function of the parameters of the game, a, b and p assuming that $\rho = 0$. Furthermore, given the parameters of the game we also provide a strategy that achieves that probability in any strategy set it belongs to. The first step would be to show that under very minor assumptions we can restrict our attention to strategy sets containing strategies of the form $nice\_q$.

Proposition 9 *Let G be a PD game and let $\widehat{G} = (G, S, p, 0)$ be a Real Talk game induced by G. Assume that S contains the strategy d for each player and a strategy s such that (s, s) is a Nash equilibrium. If P is the probability for cooperation when (s, s) is played, then there exist $q \in [0, 1]$ such that for any strategy set S' containing $nice\_q$ for both players, $(nice\_q, nice\_q)$ is a Nash equilibrium in $\widehat{G}' = (G, S', p, 0)$ and the probability for cooperation is at least P.*

In words, this proposition says that, keeping all the parameters of the games G and $\widehat{G}$ fixed, any level of cooperation that can be achieved with some strategy s can also be achieved with strategies of the form $nice\_q$. Naturally, since the original strategy set S might not contain $(nice\_q, nice\_q)$, we can only prove that $(nice\_q, nice\_q)$ has this property in any strategy set S' that contains it.

Proof. Let $s \in S$. The strategy s has to specify what action to play when it receives the signal s, and also what action to play if it receives no signal. Since an action is a probability distribution over C and D, it is enough to characterize an action by the probability of playing C. Denote the probabilities for playing C given s and ϕ by q_1 and q_2 ($q_1 = s(s)$ and $q_2 = s(\phi)$). The marginal probability that a player cooperates is thus $pq_1 + (1-p)q_2$. Denote this probability by P_s. Since the signals that the players receive are independent, each player plays C with probability P_s independently. The probability for cooperation is therefore $P = P_s^2$. The payoff for each player is:

$$\pi_s = P_s^2(b - a + 1) + P_s(a - 2) + 1.$$

Clearly, s also specifies some action to play against $\dot{d}$. Denote by $q_3 = s(d)$ the probability that s plays C facing the signal d. The probability that s plays C against d is $pq_3 + (1-p)q_2$. Denote this probability by P_d. The payoff for a player who plays d against s is

$$\pi_d = P_d \cdot a + (1 - P_d) = P_d(a - 1) + 1.$$

If (s, s) is a Nash equilibrium, then π_s is greater then any other payoff that a player can receive by deviating, including π_d. Thus, $\pi_s \geq \pi_d$ and moreover, since $\pi_d \geq 1$ also $\pi_s \geq 1$.

We consider two cases:

1. $p \geq P_s$.

Consider the strategy $nice\_0$. This strategy plays C against itself and D against any other strategy, including when receiving no signal. Let S' be a strategy space containing $(nice\_0, nice\_0)$. We need to show that a) $(nice\_0, nice\_0)$ yields a probability of cooperation of at least P and b) that it is a Nash equilibrium.

a) If both players play $nice\_0$, the probability for cooperation is p^2 and by assumption $p^2 \geq P_s^2 = P$.

b) The payoffs for $(nice\_0, nice\_0)$ are:

$$\pi_{nice\_0} = p^2 \cdot (b - a + 1) + p \cdot (a - 2) + 1.$$

However, if a player deviates to any other strategy he receives a payoff of exactly 1. Hence, $(nice\_0, nice\_0)$ can be a Nash equilibrium iff

$$p^2 \cdot (b - a + 1) + p \cdot (a - 2) + 1 \geq 1.$$

Consider the function $f = x^2(b-a+1)+x(a-2)$. We will show that it is positive for $x \in [0, 1]$. There are three cases to analyze:

1. $(b - a + 1) < 0$. It is easy to verify that $f = 0$ for $x_1 = 0$ and $x_2 = \frac{2-a}{b+1-a}$. Note that by the construction of the game $b > 1$ and thus $2 - a < (b - a + 1)$. Therefore also $2 - a < 0$ and $x_2 > 1$. Thus the function is not negative for any $x \in [0, 1] \subset [x_1, x_2]$, including $x = p$.

2. $(b - a + 1) > 0$. Once again, $f = 0$ for $x_1 = 0$. The other root can be either negative or positive, and the function itself is negative only for x between the two roots. If $x_2 < 0$, then clearly f is positive for any $x > 0$, including $x = p$. If $x_2 > 0$ then f is positive only for $x > x_2$. Since we know that it is positive for $x = P_s$, and since $p \geq P_s$ then it is positive also for $x = p$.

3. $(b - a + 1) = 0$. Since $b > 1$ it implies that $a > 2$. Thus $a - 2 > 0$ and f is not negative for $x \in [0, \infty)$ including $x = p$.

Hence, in all cases for any $p \in [0, 1]$, $p^2(b - a + 1) + p(a - 2) \geq 0$ and the inequality above holds. Therefore $(nice\_0, nice\_0)$ is a Nash equilibrium with a probability for cooperation $p^2 \geq P$.

2. $p < P_s$.

   Choose $q \in [0, 1]$ such that $p + (1 - p)q = P_s$. Note that by construction $q < q_2$. Let S' be a strategy space containing $(nice\_q, nice\_q)$. The strategy $nice\_q$ plays C when recognizing itself (an event with probability p) and plays C with probability q when not receiving a signal at all (an event with probability $1 - p$). Hence the probability that each player cooperates is $p + (1 - p)q = P_s$, and therefore the payoff for each player is exactly π_s.

   Since the most profitable deviation against $nice\_q$ is to play d, if one player deviates, the expected payoff for the deviator, π_{dev}, is not greater than $(1 - p)qa + p + (1 - p)(1 - q) =$

$(1-p)q\,(a-1)+1$. (Note that we are not assuming that d exists in $S^{'}$). Thus we get:

$$\pi_{dev} \leq (1-p)q\,(a-1)+1 \leq (1-p)q_2\,(a-1)+1 \leq P_d(a-1)+1 = \pi_d \leq \pi_s = \pi_{nice\_q}.$$

This shows that even the most profitable deviation yields a payoff which is not greater than π_s and therefore $(nice\_q, nice\_q)$ is a Nash equilibrium, with a probability for cooperation $P_s^2 = P$.

■

Since the maximal probability of cooperation can be achieved by using strategies of the $nice\_q$ type, we will now check what is the exact value of this probability given the different parameters of the game. As seen earlier, the probability that both players play C when they both choose the strategy $nice\_q$ is $[p+q(1-p)]^2$.

Note that

$$\frac{\partial[p+q(1-p)]^2}{\partial q} = 2[p+q(1-p)]\,(1-p) \geq 0,$$

and

$$\frac{\partial[p+q(1-p)]^2}{\partial p} = 2[p+q(1-p)]\,(1-q) \geq 0.$$

That is, the probability for cooperation increases both in p and in q. Since p is a parameter of the game, we are interested in finding the maximal q such that $(nice\_q, nice\_q)$ is a Nash equilibrium, given p, i.e., $p_{\max}$ is achieved by maximizing q.

In the example where $a = 4$ and $b = 3$, for any $p \leq \frac{1}{3}$, we can maximize the probability for cooperation by increasing q as much as possible, which means choosing $q = \frac{2p}{1-p}$. Substituting q in the probability for cooperation equation yields $[p + \frac{2p}{1-p}(1-p)]^2$, or $P_{\max} = (3p)^2$.

For $p > 1/3$, the maximization is achieved by choosing $q = 1$ which induces cooperation with probability 1.

In sum, $P_{\max} = \min\left((3p)^2, 1\right)$:

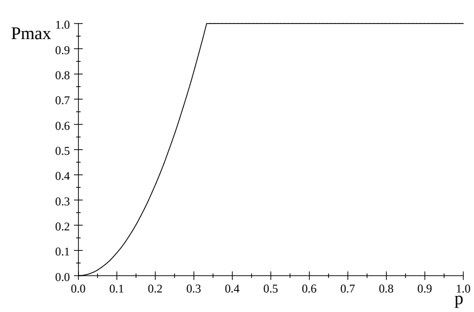

In the general case, the analysis follows the same division to cases as in the previous section:

1. $c = 0$**:**

Similarly to the example before, for $p < \frac{1}{a-1}$ maximal cooperation occurs when $q = \frac{p(a-2)}{1-p}$. Substituting q yields $[p + \frac{p(a-2)}{1-p}(1-p)]^2$, or $P_{\max} = (p\,(a-1))^2$.

For $p \geq \frac{1}{a-1}$ maximal cooperation occurs when $q = 1$.

$P_{\max} = \min\left((p\,(a-1))^2, 1\right)$.

In this case, the graph of $P_{\max}$ as a function of p looks similar the one above.

2. $c < 0$**:**

1. $r_1 < 0$ and $0 < r_2 < 1$:

$(nice\_q, nice\_q)$ is an equilibrium only for q smaller or equal to r_2. Therefore, maximal cooperation can be achieved at $q = r_2$. Substituting q yields $[p+r_2(1-p)]^2$, or $(p+\frac{1-2cp-\sqrt{4cp-4acp+1}}{2c-2cp}(1-p))^2$.

$P_{\max} = (\frac{1-\sqrt{4cp-4acp+1}}{2c})^2$.

2. $r_1 < 0$ and $1 \leq r_2$:

$(nice\_q, nice\_q)$ is an equilibrium for every q and maximal cooperation can be achieved by choosing $q = 1$.

$P_{\max} = 1$.

In sum,

$$P_{\max} = \min\left((\frac{1-\sqrt{4cp-4acp+1}}{2c})^2, 1 \right).$$

For example, when $a = 5$ and $c = -2$ the graph of $P_{\max}$ as a function of p is:

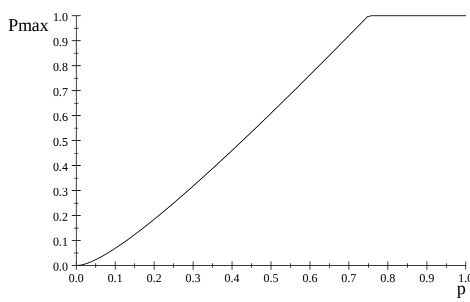

3. $c > 0$**:**

If $p \geq \frac{1}{4c(a-1)}$ $(nice\_q, nice\_q)$ is an equilibrium for all $0 \leq q \leq 1$, thus maximal cooperation can be achieved by choosing $q = 1$.

$P_{\max} = 1$.

If $p < \frac{1}{4c(a-1)}$ maximal cooperation depends on r_1 and r_2, as defined earlier.

1. $1 < r_1$ and $1 \leq r_2$:

   $(nice\_q, nice\_q)$ is an equilibrium for every q and maximal cooperation can be achieved by choosing $q = 1$.

   $P_{\max} = 1$.

2. $1 < r_1$ and $0 \le r_2 < 1$:

   $(nice\_q, nice\_q)$ is an equilibrium only for q smaller or equal to r_2. Therefore, maximal cooperation can be achieved at $q = r_2$. Substituting q yields $[p + r_2(1-p)]^2$, or

   $(p + \frac{1-2cp-\sqrt{4cp-4acp+1}}{2c-2cp}(1-p))^2$.

   $P_{\max} = (\frac{1-\sqrt{4cp-4acp+1}}{2c})^2$.

3. $1 < r_1$ and $r_2 < 0$:

   $(nice\_q, nice\_q)$ is not an equilibrium for any q, hence there is no cooperation.

   $P_{\max} = 0$.

4. $0 < r_1 \le 1$ and $0 \le r_2 < 1$:

   $(nice\_q, nice\_q)$ is an equilibrium only for q smaller or equal to r_2 or larger or equal to r_1. Maximal cooperation is reached when $q = 1$.

   $P_{\max} = 1$.

5. $0 \le r_1 \le 1$ and $r_2 < 0$:

   $(nice\_q, nice\_q)$ is an equilibrium only for q larger or equal to r_1. Once again, maximal cooperation is reached when $q = 1$.

   $P_{\max} = 1$.

6. $r_1 < 0$ and $r_2 < 0$:

   $(nice\_q, nice\_q)$ is an equilibrium for every q and maximal cooperation can be achieved by choosing $q = 1$.

   $P_{\max} = 1$.

Combining cases 1 through 6 we can see that cooperation can be achieved with probability 1 if $r_1 \le 1$ or $r_2 \ge 1$. Since the condition for $r_1 \le 1$ is $\left(c \ge 0.5 \textit{ and } p \ge \frac{1-c}{a-1}\right)$, and the condition for $r_2 \ge 1$ is $\left(c \le 0.5 \textit{ and } p \ge \frac{1-c}{a-1}\right)$, we simply get the condition $p \ge \frac{1-c}{a-1}$ In short:

$$P_{\max} = 1 \text{ iff } p \ge \frac{1-c}{a-1}.$$

Otherwise, that is when $r_1 > 1$ and $r_2 < 1$, the probability for cooperation is reduced:
The case $r_2 \ge 0$:

$$P_{\max} = (\frac{1-\sqrt{4cp-4acp+1}}{2c})^2 \text{ iff } \frac{2-a}{c} \le p \le \frac{1}{2c}.$$

The case $r_2 < 0$:

$$P_{\max} = 0 \text{ iff } \left(p > \frac{1}{2c} \text{ or } p < \frac{2-a}{c}\right).$$

A note is due regarding the last two cases, where $p < \frac{1-c}{a-1}$. On first glance it looks as if $P_{\max}$ is not monotonic in p, because for small and large ps $P_{\max}$ is zero, and in between it is positive. This is not the case, however. For given parameters a, c the following options are possible.

1. $a < 1.5$, which implies $\frac{1}{2c} < \frac{2-a}{c}$. In this case, for every p (under the assumption $p < \frac{1-c}{a-1}$) we are in the last case, where $P_{\max} = 0$.

2. $a \geq 1.5$, which implies $\frac{1}{2c} \geq \frac{2-a}{c}$ Note that $a \geq 1.5$ also implies that $\frac{1-c}{a-1} < \frac{1}{2c}$. Thus, if $p < \frac{1-c}{a-1}$,then $p < \frac{1}{2c}$. This leaves only two options: if $\frac{2-a}{c} < p$ we get some positive probability for cooperation, and for $p \leq \frac{2-a}{c}$ there is none.

In sum:

$$P_{\max} = 1 \text{ iff } \left(p \geq \frac{1}{4c(a-1)} \text{ or } p \geq \frac{1-c}{a-1}\right),$$

$$P_{\max} = \left(\frac{1-\sqrt{4cp-4acp+1}}{2c}\right)^2 \text{ iff } \left(p < \frac{1}{4c(a-1)} \text{ and } p < \frac{1-c}{a-1} \text{ and } \frac{2-a}{c} \leq p \leq \frac{1}{2c}\right) \text{ and}$$

$$P_{\max} = 0 \text{ iff } p < \frac{1}{4c(a-1)} \text{ and } \left(p < \frac{1-c}{a-1} \text{ and } \left(p > \frac{1}{2c} \text{ or } p < \frac{2-a}{c}\right)\right).$$

Following are three possible graphs for $P_{\max}$ as a function of p, given the parameters a and b:
For $a = 1.7$ and $c = 0.4$:

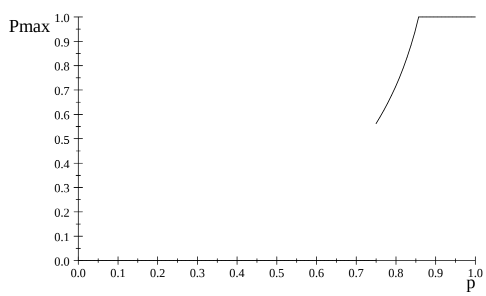

For $a = 4$ and $c = 0.75$:

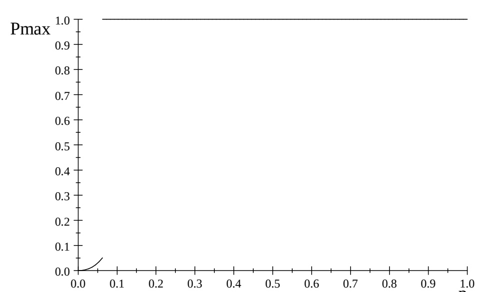

and for $a = 4$ and $c = 0.4$:

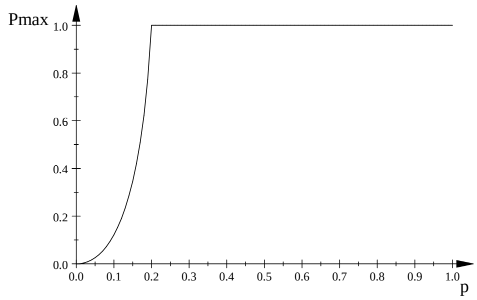

As can be seen, $P_{\max}$ is not necessarily a continuous function of p, but it is (weakly) monotonic increasing.

6 Conclusion

This paper presented a new way of introducing communication to strategic games. Moreover, an equilibrium concept was defined, and analyzed in specific cases. This was done without constructing strategy sets for the players. In some sense, the lack of explicit sets is one of the week points of this work. However, this is also its strength: we were able to analyze possible equilibria and possible payoffs without the need of constructing specific strategy sets. This was done in very general cases, with little assumptions on the strategy sets.

Communication, as modeled here, replaces other, more complicated, external mechanisms that allow players to reach similar outcomes. Even though some of the results in this paper were obtained in those other models as well, here they are obtained in a simpler, more natural, framework.

Two major assumptions were made about the players' signals. The first is that the signals the players receive during the conversation are very simple: receiving a correct signal or not. The possibility of receiving a wrong signal was not allowed, and the effect such signals might have on the set of equilibria was not examined.

The second assumption we made, in order to simplify computations, is that the probability that the players receive a signal is identical for both. However, having a separate probability for each player is not implausible: some people are better in detecting their opponent's character, not to mention that some people are better in hiding their own. Of course, having multiple probabilities is very likely to change the outcomes of the game.

We analyzed Real Talk only in two player games. Generalizing this to any n player games is possible. The players' strategies would have to be functions from all the other players' strategies into actions. In addition to that, probabilities for receiving a signal and correlations would have to be redefined.

The effect of Real Talk on the possible equilibria was presented using a few examples, such as the battle of the sexes and the prisoner's dilemma. Achieving cooperation in the prisoner's dilemma is considered as one of the hardest tests for a new model or a solution concept. This is why the prisoner's dilemma was thoroughly analyzed in this paper, and in a relatively general environment. However, there might be other interesting examples of the effect Real Talk has on equilibria outcomes that are yet to be found.

Analyzing the prisoner's dilemma, we have shown that a significant level of cooperation can emerge, and found the maximal probability for cooperation that can be sustained in equilibrium, when the players' signals are independent. This was done using $nice\_q$ strategies. The question whether the maximal probability of cooperation can be achieved by using $nice\_q$ strategies when there is correlation between the signals the players receive remains open, and requires further research. Moreover, the exact value of this probability is still unknown, although it is plausible that it is, at least weakly, higher.

Communication between people is quite complex, certainly more complicated then transferring a single signal between the players. Despite this fact, this work models at least some aspects of the interaction between the players, and it can help explain people's behavior, both in a laboratory and in real life situations.

References

[1] Anderlini, L. (1990). "Some Notes on Church's Thesis and the Theory of Games." *Theory and Decision 29*, 19-52.

[2] Binmore, K. (1987). "Modeling Ration Players, Part 1." *Economics and Philosophy 3*, 179-214.

[3] Aumann, R. (1974). "Subjectivity and Correlation in Randomized Strategies." *Journal of Mathematical Economics 1*, 67-96.

[4] Fershtman, C., K. L. Judd, and E. Kalai (1991). "Observable Contracts: Strategies Delegation and Cooperation." *International Economics Review 32*, 551-559.

[5] Farrell, J., and M. Rabin (1996). "Cheap Talk." *The Journal of Economic Perspectives 10*, 103-118.

[6] Fortnow, L. (2009). "Program Equilibria and Discounted Computation Time." *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, 128-133.

[7] Frank, R. H. (1998). "Passions within Reason." New York: Norton.

[8] Gauthier, D. (1986). "Morals by Agreement." Oxford: Clarendon Press.

[9] Howard, J. (1988). "Cooperation in the Prisoner's Dilemma." Theory and Decision 24, 203-213.

[10] Howard, N. (1971). "Paradoxes of Rationality: Theory of Metagames and Political Behavior." Cambridge: MIT Press.

[11] Kalay A., A. Kalay, and A. Kalay (2003). "Friends or Foes? Empirical Test of a Simple One Period Nash Equilibrium." Mimeo.

[12] Kalai, A. T., E. Kalai, E. Lehrer, and D. Samet (2009). "A commitment Folk Theorem." *Games and Economic Behavior* (forthcoming).

[13] Matsui, A. (1989). "Information Leakage Forced Cooperation." *Games and Economic Behavior 1*, 94-115.

[14] Peters, M., and B. Szentez (2009). "Definable and Contractible Contracts." Mimeo.

[15] Rubinstein, A. (1998). "Modeling Bounded Rationality." Cambridge: MIT Press.

[16] Sally, D. (1995). "Conversation and Cooperation in Social Dilemmas." *Rationality and Society* 7, 58-92.

[17] Solan, E., and L. Yariv (2004). "Games with Espionage." *Games and Economics Behavior 47*, 172-199.

[18] Tennenholtz, M., (2004). "Program Equilibrium." *Games and Economic Behavior 49,* 363-373.