

Exercise on the Repeated Prisoners' Dilemma

Suppose the Prisoners' Dilemma game

	D	C
D	3,3	5,1
C	1,5	4,4

will be played twice in succession by the same pair of players, and that each player will observe the other's first-stage play before the second stage is played. Each player's payoff in the repeated game is the sum of the payoffs he receives at the two stages of play.

- (a) Draw the game tree for the repeated game, *i.e.*, the game that consists of two successive plays of the PD game given above.
- (b) In the normal form (or strategic form) of this RPD, indicate why each player has $2^5 = 32$ available strategies. (You needn't actually write down each of the strategies.)

Since there are 32 strategies for each player, the payoff matrix for the normal form description of this RPD has $32 \times 32 = 1024$ cells, or strategy profiles. But we can simplify things. The fact that a strategy tells a player how his play at stage 2 should depend on *his own* play at stage 1 is redundant: a strategy needs to specify only what to choose at stage 1 and then what to do at stage 2 as a function of what *the other player* did at stage 1. (To put it another way, if my strategy tells me to choose C at stage 1, for example, then there's no need for it to also tell me what to do at stage 2 as a function of what I did at stage 1.)

- (c) Make use of the observation above to enumerate the eight behaviorally distinct strategies for each player. Write down the 8×8 normal form payoff bi-matrix for the game, placing the two players' payoffs in each of the 64 cells.
- (d) Can you identify any of a player's eight strategies as the Tit-for-Tat strategy? Can you give a similar informal descriptive name to any of the other strategies?
- (e) If a player is using the Tit-for-Tat strategy, what is the best strategy for the other player to use in response?
- (f) Determine each player's best-response correspondence in the payoff bi-matrix.
- (g) Use the best-response correspondences in (f) to determine whether either player has a dominant strategy.

(h) Determine all equilibrium profiles of strategies. (You should be able to do this by finding all intersections of the two best-response correspondences in (f).) What are the players' payoffs in each of the equilibria? How will the path of play proceed in the various equilibria — *i.e.*, in each equilibrium, what profile of actions will we observe the players choosing at each stage of play? Are any of the equilibrium strategies weakly dominated?

(i) Suppose Player 1 believes that with probability p Player 2 will play the Tit-for-Tat strategy, and that with probability $1 - p$ Player 2 will play the strategy “always defect” (*i.e.*, will always choose D). What values of p would induce Player 1 to play the “reputation strategy” in which she plays C at the first play and then D at the second play?

(j) Now suppose the players are going to play this same PD game three times. Indicate how you would go about extending the two-play game tree to yield the three-play tree. (Don't actually draw the tree. Just explain how you would extend the two-play tree to get it.) Suppose again that Player 1 believes that with probability p Player 2 will play the Tit-for-Tat strategy, and that with probability $1 - p$ Player 2 will play “always defect.” Player 1 is trying to decide whether to always defect or to instead play Tit-for-Tat until the last play (at which Player 1 will defect). Thus, Player 1 has to consider four possible profiles of strategies. For each of the four profiles, determine the resulting path of play and the resulting payoff for Player 1. Determine how large p would have to be to induce Player 1 to play Tit-for-Tat until defecting on the last play instead of simply always defecting.

(k) Now suppose the players are going to play this same PD game T times. (If you can't see how to answer this question for general values of T , try it for $T = 4$. You've already done it for $T = 2$ and for $T = 3$, above.) For each of the four strategy profiles in (j), determine the resulting path of play and the resulting payoff for Player 1. Determine how large p would have to be to induce Player 1 to play Tit-for-Tat until defecting on the last play instead of simply always defecting. (You might find it easier if you take as a benchmark Player 1's total payoff when both players always defect, and then figure out how much larger or smaller Player 1's total payoff will be under each of the other three strategy profiles.)

(l) As the repeated play of the PD game proceeds, will Player 1 be able to “update her belief” p about which strategy Player 2 might be using, on the basis of her observation of Player 2's actions?