

Contents

1	Prisoner's Dilemma	2
2	Elements of Repeated Games	3
3	Example: Finitely Repeated Prisoners Dilemma	5
4	Example: Infinitely Repeated Prisoners Dilemma	6
5	Folk Theorem	9
6	Folk Theorem Applications	11
7	Folk Theorem Extensions	15
8	Behavioral Considerations	17
9	Further Reading	18

Chapter 6: Repeated Games

This chapter considers repeated games, which are dynamic games in discrete time. We first begin with a Prisoners' Dilemma and show that equilibrium outcomes in a repeated game can be very different than the outcomes in the underlying stage game. After laying out the elements of repeated games, we present the fundamental result in repeated games known as the Folk Theorem. It describes the set of outcomes that can be supported as a SPNE of the repeated game; we offer a proof sketch based on the Grim Trigger strategy. We then sparsely sample the wide variety of applications of the Folk Theorem, and mention some extensions based on other criteria and different repeated game strategies.

1 Prisoner's Dilemma

- Consider the following symmetric normal form game, where C represents cooperate and D represents defect.

	(q)	$(1 - q)$
	C	D
C	5, 5	-3, 8
D	8, -3	0, 0

- In a one-shot game D is a (strictly) dominant strategy since it is a best-response to all possible strategies by the other player. The unique NE (D, D) is therefore a dominant strategy solution with payoffs $(0, 0)$.
 - The NE is very inefficient. Both players are better off at (C, C) , where the payoff sum is 10, and even at (C, D) , or (D, C) the payoff sum is 5, versus 0 in NE. Indeed, every non-NE pure strategy profile is Pareto optimal, which is defined as not being able to make one player better off without making someone else worse off.

– However, neither player can credibly commit to cooperate.

* Let player 1 place belief q that player 2 would honor their promise and play C . [I think there is a scorpion crossing a river analogy here!]

$$\begin{aligned} E\pi_1(C) - E\pi_1(D) &= [5q - 3(1 - q)] - [8q + 0(1 - q)] \\ &= -3 \end{aligned}$$

* Thus, for all possible beliefs $q \in [0, 1]$ regarding player 2's actions, player 1's BR is always to play D .

- Is there any way to rectify this unfortunate situation, and to sustain efficient cooperation?
- Maybe repeated interaction will do the trick. Let's have a look.

2 Elements of Repeated Games

Repeated games are composed of the following six elements.

1. A stage game, usually a NFG with 2 or more players.
 - For example, a Prisoners' Dilemma choosing (C or D) at each stage.
 - Or the centipede game, where each player chooses between pocket or push at each stage.
 - Later we may consider simple EFGs as stage games, e.g., a Stackelberg duopoly.
2. A set of stages $t = 1, \dots, T \leq \infty$, in discrete time.
3. An ending rule, where either
 - (i) $T \leq \infty$ is known with certainty to all players, or

- (ii) random termination. With probability $q \leq 1$ the game continues to the next stage; in simpler versions that event is independently and identically distributed.

4. Observed history of play.

- Often, but not always, it is assumed that all players observe everyone's actions.
- This does not mean that players observe others' *strategies*. Only realizations of mixed strategies are typically observed by other players, not the mixing probabilities.
- An opponent's mixed strategy may be learned over time, if it is played repeatedly, by keeping track of historical frequency of any given action.
- A pure or mixed strategy that responds only to play at the previous stage might be learned in a similar fashion.
- Even pure history-contingent strategies are not generally observable since only one contingency occurs as the game is played out.

5. A set of strategies.

- A strategy for player i maps history (e.g., the sequence of actions observed so far) to the set of actions.
- A strategy is a complete contingency plan, specifying an action for all possible histories of play.
- Ex: a two action, two player game played two periods has 8 pure strategies in the reduced NFG, after eliminating unrealized contingencies.
- Note that cardinality of the pure strategy set increases exponentially in the number of previous stages; see problem 1 of the current problem set.

- Typically we focus on a small subset of history dependent strategies. For example, the dependence may only be on
 - the other player’s choice in the previous period, or
 - the number of times so far that he has chosen a particular strategy.

6. A stream of payoffs earned in each stage and accumulated.

- Usually it is assumed that players seek to maximize the expected present value of the payoff stream.
- Game theorists usually write the discount factor as d (or δ), so the PV of the payoff stream (π_1, \dots, π_T) is $\sum_{t=1}^T d^t \pi_t$, for $T \leq \infty$.
- If the interest rate (or marginal rate of time preference) is $r \geq 0$ and $T = \infty$ for sure, then the discount factor is $d = \frac{1}{1+r}$.
- If instead the continuation probability is constant at $q \in (0, 1]$ then the discount factor is $d = \frac{q}{1+r}$.

For repeated games we keep the basic definition of a Nash equilibrium as a simultaneous best-response, albeit here in terms of history-contingent strategies.

3 Example: Finitely Repeated Prisoners Dilemma

To illustrate ideas, and because it is of considerable interest in its own right, let’s find NE for a finitely repeated Prisoners Dilemma in which $T < \infty$ is known with certainty by all players. For example, everyone knows that $T = 80$. We’ll use backwards induction (BI).

- At stage T we have a 1-shot game.
- D is a dominant strategy for both players.

- Thus, in any NE, regardless of the history, the profile in the last stage is (D, D) .
- Now consider stage $T - 1$. Since (D, D) will be played in stage T no matter what happens in the current stage, rational players will pick their stage game BR in the current stage.
- Thus in any NE, regardless of previous history, (D, D) will be played at stage $T - 1$.
- ... and so on. Cooperation unravels under BI, from each stage to the previous stage.
- The unique SPNE of the repeated game, and indeed its unique NE, is to play (D, D) in all stages.

Too bad. It looks like repetition doesn't help us escape the trap of inefficient defection. But what about an infinite horizon?

- Suppose $T = \infty$ for sure. There is no last period, so BI can't get started. It has no traction, so perhaps cooperation does not unravel.
- Suppose instead that there is a $q < 1$ iid probability that the game continues to the next stage.
- If you don't know the final period, then you can't start BI, so cooperation does not necessarily unravel.

So, with either an uncertain or infinite time horizon there is the possibility that cooperation can be sustained if the players are sufficiently patient. The Folk Theorem tells us how and allows us to define sufficiently patient.

4 Example: Infinitely Repeated Prisoners Dilemma

Define the following repeated game strategy:

- *Grim Trigger*: Play C until the other player plays D , then play D forever after. The strategy is called Grim since the player never forgives another player who ever plays D .
- Note that GT is a complete contingency plan. It is based on the full history of play at any given point in time and specifies an action for all future time periods.

Define a discount factor $d \equiv \frac{q}{1+r} \in (0, 1)$ as usual.

- The smaller is $r > 0$, the closer d is to one and the less the future is discounted.
- Let q be the iid probability that the game continues to the next period. The larger is q , the more likely it is that the game will continue, and the closer d is to one. So larger q implies less discounting or more “patience.”
- Rational players seek to maximize the present discounted value of the payoff stream

$$\Pi_i = \sum_{t=0}^{\infty} d^t \pi_i(t)$$

The game is symmetric, so we can look at the opponent (player 2) also playing Grim Trigger and ask if it is a best-response to itself.

- If so, then it is a NE.

The present value (PV) of the payoff stream to playing Grim against Grim is:

$$PV(5, 5, 5, \dots, 5, \dots) = \sum_{t=0}^{\infty} d^t 5 = 5 \sum_{t=0}^{\infty} d^t \quad (1)$$

because neither player is the first to defect.

- The infinite geometric series has sum $S \equiv \sum_{t=0}^{\infty} d^t = \frac{1}{1-d}$, since:

$$\begin{aligned} S &= d^0 + d^1 + d^2 + d^3 + \dots + d^n + \dots \\ dS &= d^1 + d^2 + d^3 + d^4 + \dots + d^{n+1} + \dots \\ S - dS &= d^0 = 1 \\ S &= \frac{1}{1-d}. \end{aligned}$$

- Thus, $PV(5, 5, 5, \dots, 5, \dots) = \frac{5}{1-d}$.

The present value of the payoff stream from deviating at time zero and choosing the obvious BR thereafter against a player playing Grim is:

$$PV(8, 0, 0, \dots, 0, \dots) = 8 \tag{2}$$

Therefore cooperation is sustainable if C gives a higher payoff than D against Grim, i.e., if

$$\begin{aligned} \frac{5}{1-d} &\geq 8 \\ 5 &\geq 8(1-d) \\ d &\geq \frac{3}{8} \end{aligned}$$

The same comparison (with both sides discounted by the same factor) applies to a first defection at any later time: it reduces payoff as long as $d \geq \frac{3}{8}$.

We conclude that cooperation is sustainable as long as the players are sufficiently patient and don't discount the future too heavily, i.e., $d \geq \frac{3}{8}$. If $q = 1$ then the condition is $r \leq \frac{5}{3}$.

Note that the net loss of payoff from playing D (due to "punishment" by the other player who plays D instead of C in future periods) is larger when the future is discounted less heavily,

- either because the probability of continuing (and defection being punished) is greater,
or
- because interest rates are lower, or periods are shorter in clock time, or players really are more patient in the sense of lower marginal rate of time preference.
- If there are two players with different discount rates in a symmetric game then Grim can sustain cooperation up to the critical discount rate of the less patient player.

5 Folk Theorem

- For infinitely repeated games, a subgame is what remains to be played at any point in time, thus subgame perfection (SGP) is easily met.
 - Each subgame is essentially identical to the game that begins at the first period since there is no determined terminal period.
 - Of course, in a repeated game, each later subgame contains a longer set of possible histories that subsequent actions can be conditioned upon.

We just saw that cooperation can be sustained in the repeated Prisoners Dilemma as long as the final period is unknown and the players are sufficiently patient.

However, the problem is that the solution was “too powerful.”

- Full cooperation can be sustained as a NE (indeed, SPNE) via Grim Trigger strategies, but is that the only NE?
- The Folk Theorem implies that *almost anything* can be sustained as a SPNE !

	C	D
C	5, 5	-3, 8
D	8, -3	0, 0

- For the Prisoners’ Dilemma example, with stage game copied above, Figure 1 below shows the payoffs for pure strategy profiles as the corners of a parallelogram.
- By weighted averaging these four payoff vectors over numerous stages, we can generate all possible convex combinations (i.e. the convex hull) as normalized payoffs in the repeated game.
- For example, with a discount factor close to 1.0, playing (C,C) 60% of the time, (D,C) 20%, and (C,D) and (D,D) 10% of the time, we’d obtain average payoff near the point at the arc labelled SGPNE.

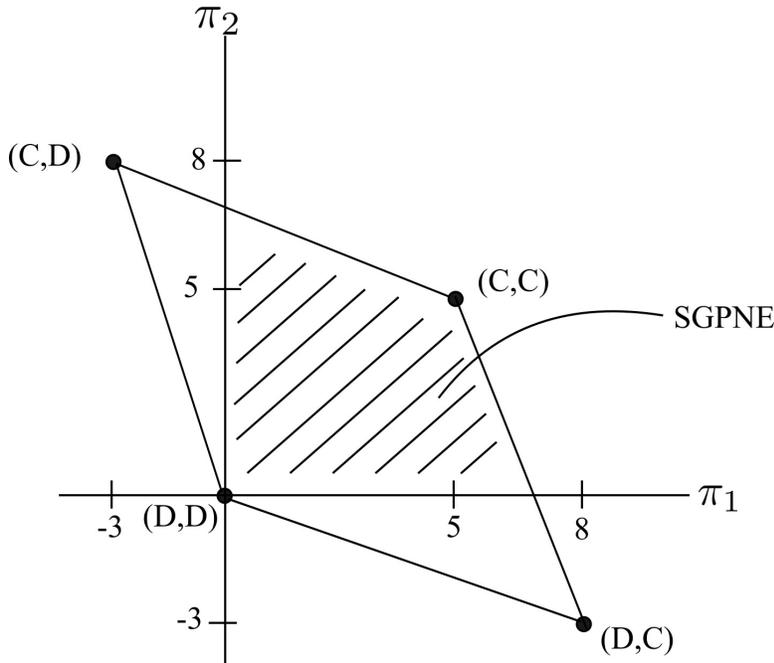


Figure 1: Feasible payoffs in a Prisoner's Dilemma stage game. The shaded region shows payoff vectors sustainable for sufficiently patient players.

- Anything in the shaded region of Figure 1 can be sustained as a NE (indeed, SPNE) of the infinitely repeated game.
- To see what is excluded, consider that the player can always revert to all-D, i.e., for each future history, just play D. This will guarantee a payoff of 0.
- A rational player will not accept a lower payoff than that.

We are now ready to state the Folk Theorem more formally and generally.

Given a stage game G with $n \geq 2$ players, let $F[G] \subset R^n$ be the feasible payoff vectors, i.e., the convex hull of the payoff vectors arising from all pure strategy profiles. Let $z \in R^n$ be the payoff vector for the least efficient Nash equilibrium of G , and let $IR[G] = \{x \in R^n : x_i \geq z_i, i = 1, \dots, n\}$ be the individually rational payoff vectors.

Folk Theorem (*J. Friedman, 1971, among others*). *Any payoff vector $u \in F[G] \cap IR[G]$ can be approximated arbitrarily closely as the average payoff in SPNE of an infinite hori-*

zon repeated game with stage game G , for a discount factor sufficiently close to 1.0.

That is, there are zillions of NE, and even of SPNE, in the infinitely repeated game. They generate average payouts that range from the most inefficient stage game NE to the most efficient possible payoff vector. The main caveat is that, to get some of them (especially those that give a player just a bit more than her stage game NE payoff), all the players must be patient enough.

proof sketch.

- All players adopt a generalized Grim Trigger (or “Nash reversion”) strategy: play a script that generates the desired average payoff vector. If anyone deviates from the script, then revert to the specified stage game Nash equilibrium strategy.
- The same sort of argument used above for the Prisoner’s dilemma shows that sticking with the script yields, at every stage, a higher expected payoff than deviating, as long as d is sufficiently close to 1.0.

6 Folk Theorem Applications

The applications all involve the same dynamic tension: the temptation to grab a higher immediate payoff is opposed by anticipating lower future payoffs. This “shadow of the future” looms larger when the discount factor is closer to 1.0, e.g., when there is a higher probability of continuing the game, or when the next stage comes sooner.

1. Sovereign Debt [cite?]. Developing nations with large amounts of debt face high cost of servicing this debt (rD).
 - The temptation is to avoid that cost by defaulting on the debt. The shadow of the future is that the country will likely be “punished” by creditors in interna-

tional capital markets, who will be unwilling to loan to them at a reasonable interest rate for many years.

- A currency crisis can increase the debt servicing cost (rD) so much that the inequality changes sign and the country rationally defaults after all.
- An increase in international interest rates can also can trigger a default since it will increase both debt servicing cost and the discount rate.

2. Central Banks. [cite Carl Walsh's original article.]

- The temptation is to create a surprise burst inflation, which generates an output boost. Everyone (including the central bank) likes higher output.
- The shadow of the future is that the surprise burst harms the reputation of the central bank, and firms and households increase their expectations of inflation.
- High expected inflation, it turns out, lowers the value of the central bank's objective function in later periods.
- Independent central banks can announce inflation targets, and improve their reputation by hitting them. This keeps expected inflation down.

3. Oligopoly models of collusion.

- An OPEC member can increase profit today by exceeding its output quota.
- Other OPEC members, particularly Saudi Arabia, can punish the defector by substantially increasing their own output in subsequent periods, causing oil prices to plummet.
- They generally don't play "grim trigger," but instead usually limit the punishment period to a year or less. Usually that is enough to fully offset the temptation.

4. Incentive contracts.

- CEOs with stock options have an incentive to inflate short-run earnings (or encourage creative accounting, or just plain falsify of reports) but there are long-run consequences if caught.
- The temptation is more severe when the job is less secure (higher termination probability q).

5. Cold war arms race could be elaborated this way. Also vehicle size “arms race” (think of Hummers around 2006). Standardized bumpers at the same height is an analogue to an arms treaty.

6. Preemption games are a variant on this theme. For example, Mendelsohn, R. (1994). “Property Rights and Tropical Deforestation” develops a model with a common resource whose value increases over time, e.g., a young forest. With poorly defined property rights each individual has an incentive to poach inefficiently early.

We’ll spell out one application, from Lin (1990) *Journal of Political Economy*, “Collectivization and China’s Agricultural Crisis in 1959-1961.”

- During the ten years that followed the communist takeover (1948-1958) agricultural output increased rapidly as small family farms were encouraged to take advantage of economies of scale associated with joining an agricultural collective (commune).
- Then during 1959-1961 there was a huge drop in agricultural output.
 - Estimates are 30 million deaths as a result
 - The official explanation was “bad weather.”
 - Lin argues that this was one of the largest catastrophes in human history, and that it was a result of turning a repeated PD into a 1-shot game.

- Argument is based on the following stage game (the model ours; it is not spelled out in the paper)

– Each family chooses a strategy s_i from the following set $S_i = \{B, C, D\}$:

B : join commune, but shirk

C : join commune, and work hard

D : depart from commune

- Assumptions:

1. $\pi_i(B, s_{-i})$ and $\pi_i(C, s_{-i})$ are both strictly increasing in the number that play C .

- This shows the gains to cooperation and the public goods nature of commune membership.

- Output is increasing in hard work, and output is “non-excludable,” or “to each according to their needs (not effort).”

2. $\pi_i(B, s_{-i}) > \pi_i(C, s_{-i}) \forall s_{-i}$.

- Shirk has a higher payoff than working hard, due to the individual effort cost of hard work (i.e., shirk strictly dominates working hard, making it a PD game). This is the “public goods” nature of the agricultural output of the commune. There is a $\frac{1}{n}$ problem where each individual gets $\frac{1}{n}$ of the output even if they shirk, but bear 100% of the effort cost from hard work.

3. $\pi_i(D, s_{-i}) = d_o \forall s_{-i}$.

- Each family has an outside payoff (i.e., threat point, minimax payoff or reservation payoff) that is independent of what is happening on the commune.

4. $\max_{s_{-i}} \pi_i(C, s_{-i}) \gg d_o \gg \min_{s_{-i}} \pi_i(B, s_{-i})$, where the very strict inequality \gg means “much greater than.”

Assumptions 1. through 4. imply that it is better to be in a commune where everyone works hard than to be in a commune where everyone shirks, with d_o in between.

- Timing: at each stage (1 year, or growing season) crops are planted, those in the commune work hard or shirk, then harvest and $\pi_i(s_i, s_{-i})$ is realized.
 - Decision rule: If $\pi_i \geq d_o$ then remain in commune, if $\pi_i < d_o$ then leave.
 - Game proceeds to the next stage.
 - Thus, we have a repeated game, no known final period and if players are sufficiently patient Folk Theorem results are possible, i.e., cooperation where all play C .
- From 1948-1958 commune membership was voluntary and the data suggests that cooperation was sustained.
- In 1959 commune membership became mandatory.
 - Depart is no longer an option.
 - Only available strategies are B and C , and what was a repeated game, essentially became a 1-shot game. (“Punishment” for B was D , but that became unavailable.
 - By assumption 2) above shirk was a dominant strategy and production collapsed. NE is everyone shirking.
 - The evidence for this hypothesis, according to Lin, was that if this was not the case then output should never fall below d_o per household.

7 Folk Theorem Extensions

1. Fudenberg and Maskin (1986) *Econometrica* 54: 533-54.

- They show that the minimax payoff can replace the NE of the 1-shot game in the Folk Theorem.
- One application is that in a Cournot duopoly Grim Trigger can sustain tacit (no communication) collusion in a repeated game.
 - Can the players guarantee themselves the 1-shot NE payoffs?
 - * No, if the other player produces $q_j > q^{NE}$ then $\pi_i < \pi^{NE}$.
 - * The reservation payoff is 0, or $-FC$ (fixed costs), and therefore a SGPNE can be obtained with payoffs less than the NE payoffs from the stage game.

2. Abreu (1988) *Journal of Economic Theory* 39: 191-225.

- The most effective strategy is to administer the strongest credible punishment.
 - That is, choose that SGPNE which yields the lowest payoff to a defector.

3. Myerson (1991) *Game Theory: Analysis of Conflict* (Harvard University Press), via Barrett Environment and Statecraft chapter 10.

- The problem with Grim when there are more than two players is that it may not be collectively rational to inflict the punishment of defect forever. Are the punishments collectively credible to carry out? SGP requires that the punishment be individually rational, but renegotiation-proof equilibria (Farrell and Maskin, 1989) requires that the punishment be collectively rational to carry out.
 - Put differently, nothing about SGP with more than two players stops a reversion to cooperation.
- To address this issue, Myerson proposes the strategy of Getting Even.

- Getting Even states play C if I have played D less often than other player in the past.
 - * Consider a situation where the other players have played defect for three periods, and I have played defect the last two periods after playing C the first period. If I play D then they would all have to play C for cooperation to re-emerge.
- Thus, Getting Even played against itself can SGPNE and re-negotiation proof since it is both individually and collectively rational to administer the punishment, and cooperation can resume following defection.

Historical notes.

The Folk Theorem was given its name because it was “folk wisdom,” known by most leading game theorists, before a version was formally published by James Friedman (1971).

There are many versions of the Folk Theorem. Some define IR in terms of maximin strategies rather than stage game NE strategies. See Gibbons (1992) for a fairly accessible proof.

Repeated games sometimes used to be called supergames, but that terminology has fallen into disuse.

Repeated games can be formulated in continuous time, not just discrete time. See Friedman and Oprea (2012).

8 Behavioral Considerations

We’ve seen that in the finitely repeated PD, the only NE is the very inefficient all-D profile. Does this describe actual human behavior in the lab? Not entirely. People tend to cooperate a lot in early periods, not so much in later periods. As they gain experience with a particular horizon length, they tend to cooperate more in early periods! But as

they gain experience, they tend to defect slightly earlier (say at period 6 or 7 of 10, instead of period 8 or 9), at least with favorable payoff parameters. See Selten and Stoecker (1986) for the first study, and see Embry et al (forthcoming) for a recent study.

There are two main theoretical explanations for the finding that people tend to cooperate at least for a while in the finitely repeated PD. The most famous is Kreps et al (1982), known as the “Gang of Four” paper. The authors assume that players believe that, with small probability, their opponents have an irrational urge to cooperate, and that their opponents believe (with the same small probability) that the given player is similarly irrational. In BNE, it pays to pretend to be such a crazy type for quite a while when the horizon is long but finite.

An alternative explanation, less famous but (some argue) more consistent with the evidence was developed in Radner (1986). He points out that the cost of cooperating one more period – to suffer the sucker payoff for one period only – is outweighed by the cost of losing many periods of potential cooperation. Thus it is almost (not quite) rational to employ a threshold strategy of playing Grim until a fixed period late in the game, and playing all-D thereafter. Friedman and Oprea (2012) extend this idea to explain the high degree of cooperation they see in their continuous time finite horizon experiments.

There is also an extensive literature on programmed strategies playing finite (or infinite) horizon repeated prisoner’s dilemma. A high profile public tournament by political scientist Robert Axelrod found that the most successful programmed strategy entered was also the simplest: Tit-for Tat. (This strategy plays C in the first period, and thereafter imitates the opponent’s most recent action.) See Axelrod (2006) for a summary.

9 Further Reading

Abreu D. (1988). “On the Theory of Infinitely Repeated Games with Discounting,” *Econometrica*, 56(2): 383-396.

- Axelrod, Robert M. The evolution of cooperation. Basic books, 2006.
- Embrey, Matthew, Guillaume R. Frchette, and Sevgi Yuksel. "Cooperation in the finitely repeated prisoner's dilemma." (forthcoming).
- Farrell, J. and Maskin, E. (1989). "Renegotiation in Repeated Games," *Games and Economic Behavior*, 1: 327-60.
- Friedman, Daniel, and Ryan Oprea. "A continuous dilemma." *The American Economic Review* 102.1 (2012): 337-363.
- Friedman, J. (1971). "A Non-cooperative Equilibrium for Supergames," *Review of Economic Studies*, 38(113): 1-12.
- Fudenberg, D. and E. Maskin (1986). "The Folk Theorem in Repeated Games with Discounting or Incomplete Information," *Econometrica*, 54(3): 533-54.
- Gibbons, Robert. Game theory for applied economists. Princeton University Press, 1992.
- Kreps, David M., et al. "Rational cooperation in the finitely repeated prisoners' dilemma." *Journal of Economic theory* 27.2 (1982): 245-252.
- Lin, J. (1990). "Collectivization and China's Agricultural Crisis in 1959-1961," *The Journal of Political Economy*, 98(6): 1228-52.
- Radner, Roy. "Can bounded rationality resolve the prisoners dilemma?." *Essays in honor of Gerard Debreu* (1986): 387-399.
- Mendelsohn, R. (1994). "Property Rights and Tropical Deforestation," *Oxford Economic Papers*, 46: 750-6.
- Selten, Reinhard, and Rolf Stoecker. "End behavior in sequences of finite Prisoner's Dilemma supergames A learning theory approach." *Journal of Economic Behavior & Organization* 7.1 (1986): 47-70.