COOPERATION

- Creation of economic surplus often requires cooperation of involved parties. This situation has an in-built conflict between taking a privately costly action that improves social welfare and a non-costly action that only suits private interests of the player.
- We already looked at two examples of *sequential games* that illustrate the point:
  - **Trust Game**: surplus is created by the first-mover (trustor)
  - **Gift-Exchange Game**: surplus is created by the second-mover (worker)
- The most traditional workhorse for studying cooperation, though, is the *simultaneous-move game* called **Prisoner's Dilemma (PD)**. In this game of two players, regardless of the strategy pursued by the opponent, one is always strictly better off to take the *non-cooperative action* ("defect", "confess","cheat", "free-ride") rather than the *cooperative action* ("cooperate", "deny", "play honest", "contribute"). In the jargon of game theory, the non-cooperative action is the **dominant strategy** for each player.
- Classic example of two prisoners (with sentences in years, with negative signs, being the payoffs):

Player 1\Player 2	Confess	Deny
Confess	$(-5, -5)$	$(0, -10)$
Deny	$(-10, 0)$	$(-1, -1)$

- In general, the payoff structure of a symmetric 2×2 PD is given by

Player 1\Player 2	Cooperate	Defect
Cooperate	(a, a)	(d, c)
Defect	(c, d)	(b, b)

  with $c > a > b > d$.
- This game in fact has what is known as a **Dominant Strategy Equilibrium**: (Defect, Defect). In fact, this is also the only **Nash Equilibrium** of the game. This is the unique prediction of game theory (with self-regarding preferences) for this game.
- One would expect, though, that is real life this game is repeated between a fixed pair of players, and hence there is an opportunity to punish previous non-cooperation. But again, the unique *subgame-perfect equilibrium* of a finitely many times repeated PD is eternal non-cooperation. This result can easily be obtained by backward induction.
- It is only when there is no upper bound on how many times the game may be repeated that things change. In that case, there is no last-stage effect, and virtually any level of cooperation (in terms of average realized payoff up to the cooperative outcome) can be realized by some strategy combination as long as players are sufficiently patient (**Folk Theorem**). This can be illustrated by the use of grim or trigger strategies.
- However, **Dawes and Thaler (1988)**, in their survey article, document that many subjects are at least somewhat cooperative even if the game is repeated only finitely many times, sometimes even only once.
- There are two leading explanations of this finding:

1. **Kreps et al. (1982)** propose a model of **reputation building** in the repeated play of Prisoner's Dilemma. If there is a chance (even a relatively small one) that an opponent is a conditional cooperator (by, for example, playing a tit-for-tat strategy), it may be in a player's interest to cooperate, at least for some time, in order to reap higher payoffs. For example, if the opponent plays tit-for-tat for sure, then it pays off to cooperate in all except for the last period. If the opponent plays tit-for-tat with probability of less than 1, the cooperation will start decaying sooner. However, in the one-shot version of Prisoner's Dilemma this model predicts defection.

2. An alternative explanation is **altruism**. As a simplest version, suppose that Player 1 enjoys a "warm glow" feel from cooperation captured in money-equivalent premium of δ, whereas Player 2 is selfish. The payoff structure is then changed to

Player 1\Player 2	Cooperate	Defect
Cooperate	$(a+\delta, a)$	$(d+\delta, c)$
Defect	(c, d)	(b, b)

   If $\delta < \min(b-d, c-a)$, Player 1 is labeled as *egoist.* If $\delta > \max(b-d, c-a)$, Player 1 is labeled as *dominant strategy altruist.* Player 1 with intermediate values of δ is called a *best response altruist.* As long as there are players of the later two types, cooperation will be observed in both repeated and one-shot versions of PD.

Andreoni and Miller (1993)

- Andreoni and Miller (1993) experimentally implement prisoner's dilemma with 200 repetitions and payoffs given by $a = 7, b = 4, c = 12, d = 0$.
- They consider four conditions with 14 subjects each:
  1. *partner*: same partner for 10 consecutive periods, then random rematch
  2. *stranger*: random rematch in each period
  3. *computer50*: same as partner, but subjects have a 50% chance of playing a computer that plays tit-for-tat (this is intended to test for the idea proposed by Kreps et al.)
  4. *computer0*: same as computer 50, but subjects have a 0.1% chance of playing a computer that plays tit-for-tat (the reason for including this treatment is that computer50 not only introduces altruistic players, but also the notion of tit-for-tat play; computer0 is an attempt to separate the two effects)
- Here are their results:

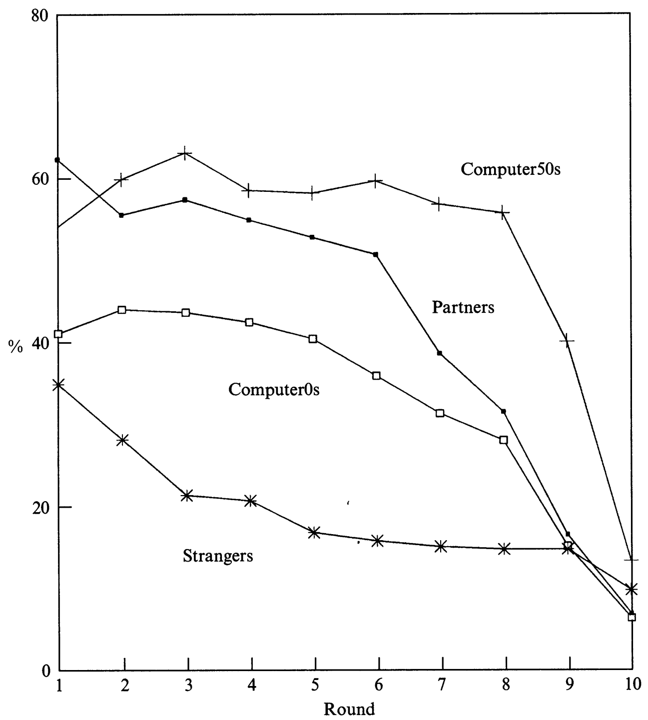

Fig. 2. Percent cooperation by round. Averaged over all 20 10-period games.

- We observe that:
  1. Cooperation is positive and decreasing over time in all four treatments.
  2. Subjects in *partners* and *computer50* treatments are more cooperative than subjects in *computer0* treatment, who are on turn more cooperative than subjects in the *strangers* treatment.
- The fact that there is a positive amount of cooperation in the *strangers* is consistent with existence of **altruism**.

- Likewise, results for the other three treatments lend support to the hypothesis of **reputation building**.
- Also consider the timing of the first defection:

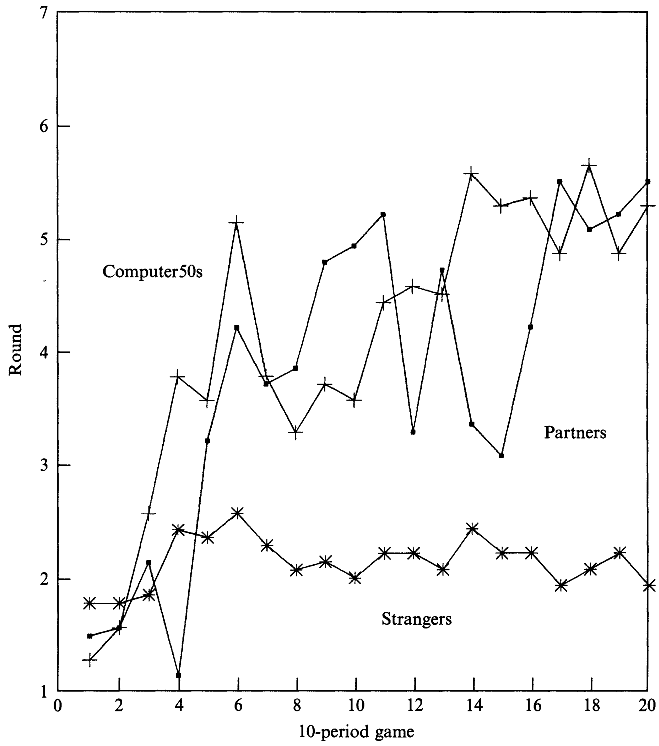

Fig. 4. Mean time until first defection.

- The authors sometimes try to sell cooperation in the repeated play treatments as "altruism", but this is disputable. If a player believes that his opponent uses tit-for-tat, then even if the player is completely selfish, he should cooperate in all except for the last period (use backward induction to see this result).

Cooper, DeJong, Forsythe and Ross (1996)

- CDFR implement both one-shot (20 periods with random rematching in each period) and repeated versions (two blocks of fixed pairs for 10 periods) of PD with $a = 800, b = 350, c = 1000, d = 0$. Their findings:

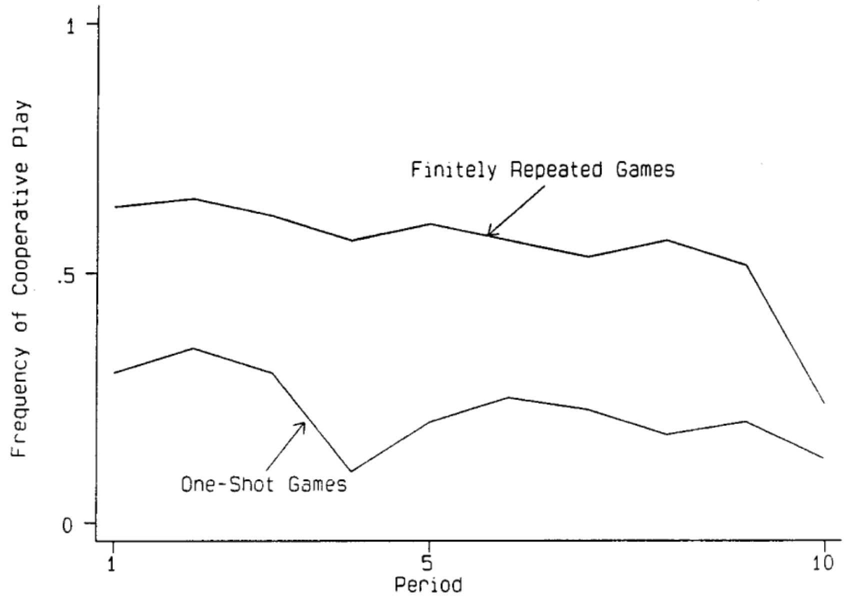

- Similarly to Andreoni and Miller (1993), we observe that:

  1. Cooperation rate is positive and decreasing over time in both repeated and one-shot play.
  2. Cooperation rate is higher in repeated play compared to one-shot play.

- Again, the fact that there is a positive amount of cooperation even in the one-shot game suggests that **altruism** must be present in some subjects. Moreover, since most of cooperative play in one-shot games comes from subjects who do not cooperate all the time, this suggests that subjects tend to be best response altruists rather than dominant strategy altruists:

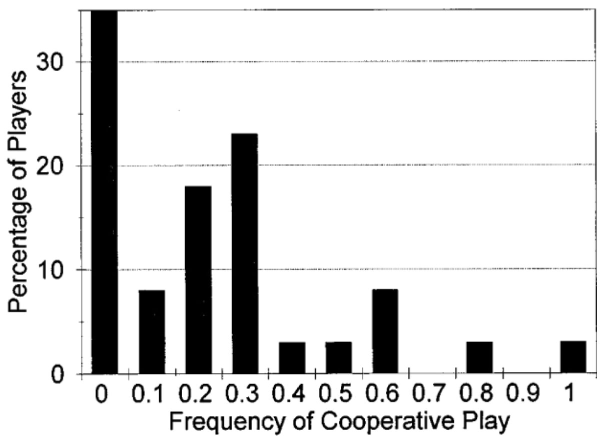

- Looking at behavior at the individual level, CDFR estimate the fraction of best response altruists to be between 12.5% and 15%.

- CDFR point out that the results from repeated play are, on aggregate, indeed consistent with the **reputation building** hypothesis. However, when going down to pair or individual level, this is no longer so because:

    1. the reputation building model suggests that a player initially cooperates and then switches to non-cooperation and that there should be no cooperation in the last period; only 10% of pairs of players match these two predictions
    2. using the same criteria, at the individual level only 25% of subjects match these two predictions
    3. the time pattern of play observed is rather different than predicted by the reputation building model even though qualitatively the observed cooperation rates have the predicted pattern of decline over time; also the observed cooperation rates are a concave function of time while the predicted cooperation rates are convex:

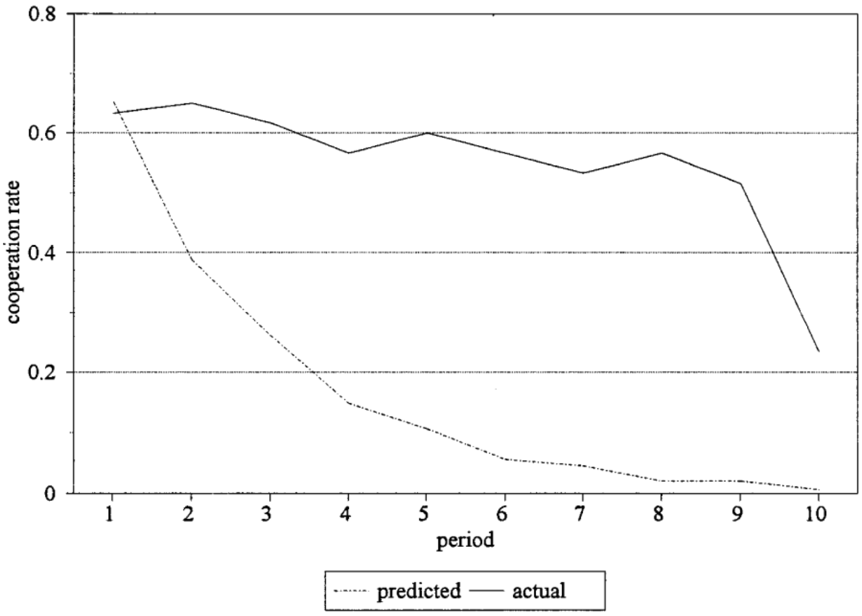

FIG. 3. Predicted and actual cooperation rates for PD-FR.

- The authors do not exploit in detail a mixed model of altruism and reputation building in repeated play.

Dal Bo (2005)

- The author considers the infinitely repeated PD. The way to implement this in a lab is to have a random continuation rule in each fixed pair. Although this is not the first paper to do this, it improves the methodology in two important respects:

  1. in runs analogous finitely repeated PD as a control; this is important because it allows distinguishing between the effect of playing for more rounds in expectation as opposed to players always having a "shadow of the future" hanging over their heads; lengths of the finitely repeated games were chosen to coincide with the expected lengths of the infinitely repeated ones
  2. it uses two different stage games, one more conducive to cooperation (PD1 with $a = 65, b = 35, c = 100, d = 10$) and one less conducive to cooperation (PD2 with $a = 75, b = 45, c = 100, d = 10$) at the discount factor of $\delta = 1/2$.

TABLE 3—EQUILIBRIUM OUTCOMES

	PD1	PD2
0	DD	DD
½	DD, CD, DC	DD, CC
¾	DD, CD, DC, CC	DD, CD, DC, CC

  3. it uses a large number of subjects per session and two-group rotation matching procedure to limit contagion effects. Hence the decisions one subject made in one match could not affect, in any way, the decision of subjects he or she would meet in the future.

- The author uses three different continuation probabilities: $\delta = 0$, i.e., one-shot, $\delta = 1/2$ and $\delta = 3/4$.

- Results: we observe that:

  1. The level of cooperation increases with the shadow of the future (higher δ): on average, it is 9% for $\delta = 0$, 27% for $\delta = 1/2$ and 37% for $\delta = 3/4$.

TABLE 5—PERCENTAGE OF COOPERATION BY MATCH AND TREATMENT*

		Match									
		1	2	3	4	5	6	7	8	9	10
Dice	$\delta = 0$	26.26	18.18	10.61	11.62	12.63	12.63	5.56	5.26	5.26	5
	$\delta = ½$	28.36	27.12	34.58	35.53	21.60	19.08	29.84	35.96	28.16	50
	$\delta = ¾$	40.44	28.57	27.78	32.92	46.51	33.09	44.05	53.51	42.26	45.83
Finite	$H = 1$	26.56	18.23	16.67	17.19	11.98	8.02	6.79	10.49	6.14	6.67
	$H = 2$	19.79	15.89	14.84	9.64	11.46	10.80	12.04	10.19	6.58	6.67
	$H = 4$	31.64	30.34	30.47	25.52	25.13	23.77	16.36	19.75	14.91	20.83

* *All rounds and sessions.*

  2. The level of cooperation is larger in the first round of an infinitely repeated game than in the first round of the finitely repeated game. **Question:** this should be true for any generic round of infinitely repeated play, but is not. In fact, cooperation decreases over time even in the infinitely repeated game. Why?
  3. For $\delta = 1/2$, PD2 results in more outcomes CC than PD1.

Table 6—Percentage of Cooperation by Round and Treatment*

		Round											
		1	2	3	4	5	6	7	8	9	10	11	12
Dice	$\delta = 0$	9.17											
	$\delta = 1/2$	30.93	26.10	19.87	12.50	12.96							
	$\delta = 3/4$	46.20	40.76	38.76	34.58	33.04	27.27	24.75	26.28	29.17	26.04	32.29	31.25
Finite	$H = 1$	10.34											
	$H = 2$	13.31	6.90										
	$H = 4$	34.58	21.55	18.97	10.63								

* *All sessions, matches four through ten.*

Table 7—Distribution of Outcomes by Stage Game and Treatment*

	$\delta = 0$		$\delta = 1/2$		$\delta = 3/4$	
	PD1	PD2	PD1	PD2	PD1	PD2
CC	2.98	0.27	**3.17**	**18.83**	20.68	25.64
CD & DC	20.83	13.98	**28.57**	**25.50**	30.34	26.03
DD	76.19	85.75	68.25	55.67	48.98	48.33

* *Matches four through ten, and all rounds.*

Guide to Further Reading

- **Ellison (1994)** (a theoretical paper) considers an environment with random rematching in a large population and shows that there is an equilibrium in which cooperations is sustained by "contagious punishments".
- **Duffy and Ochs (2009)** implement this environment in the lab but find that a norm of cooperation and punishment does not arise.