

CO-EVOLUTION IN ITERATED PRISONER'S DILEMMA WITH INTERMEDIATE LEVELS OF COOPERATION: APPLICATION TO MISSILE DEFENSE

PAUL J. DARWEN

*School of Information Technology and Electrical Engineering
The University of Queensland, Brisbane 4072, Australia*

XIN YAO

*School of Computer Science, The University of Birmingham
Edgbaston B15 2TT, United Kingdom*

There is a widespread perception that in conflict situations, more intermediate choices between full peace and total war makes full peace less likely. This view is a motivation for opposing the proposed National Missile Defense. This perception is partly due to research in the abstract game of Iterated Prisoner's Dilemma. This paper critically evaluates this perception.

Keywords: Co-Evolution; prisoner's dilemma; missile defense.

1. Introduction

In a world torn by conflict, how can mutually beneficial cooperation can emerge spontaneously, with no higher authority than the rule of the jungle?

History, economics, and especially biology¹ provide many real-world examples of mutual cooperation without any enforcement from a higher authority. This cooperation occurs despite a bloodthirsty option offering more payoff to those might exploit their cooperative partners. While it's clearly a good thing to keep such a partnership going, there are many question about how such cooperative partnerships are formed and maintained.

In order to study how this kind of cooperation comes about, many researchers use the abstract game of Iterated Prisoner's Dilemma (IPD). Normally, the 2-choice game is studied, where players can choose only between full cooperation or full defection. Some studies have looked at intermediate choices between these two extremes. They found that intermediate choices made full cooperation less likely to dominate, but did not explain *why* this is so.

This paper examines why intermediate choices make full cooperation less likely. This question is relevant to the proposed National Missile Defense, which (its opponents claim) would allow intermediate choices between full peace and all-out war, and thus make a partial nuclear war more likely.

1.1. *Overview*

Section 2 leads the reader from some broad aspects of history up to the abstract issues of the game of Iterated Prisoner's Dilemma (IPD). Section 3 describes the experimental setup, and Sec. 4 shows two causes for why cooperation is less likely, described in Secs. 4.3 and 4.4. Section 5 discusses these results.

To bookend Sec. 2, Sec. 6 leads the reader back from the abstract game of IPD and into issues about the proposed National Missile Defense (NMD). Those two mechanisms that make full cooperation unlikely, and how they may be relevant to missile defense, are covered in Secs. 6.3 and 6.4. The paper concludes in Sec. 7.

2. History and the Game of Iterated Prisoner's Dilemma

How can conflict in international relations be turned into peaceful cooperation? While this question may seem intractable, many nations who were once bitter enemies have become close friends.

For example, the world's longest undefended border is between Canada and the United States. But over several days in August 1814, troops from Canada and Britain captured and sacked Washington DC, in retaliation for the American destruction of Canada's then-capital city of York in April 1813.

Two bitter enemies can become close friends even after destroying each other's capital cities. It can be done. As a general question, how can conflicts be brought to a peaceful conclusion?

2.1. *Enforcing cooperation by coercion*

A recurring suggestion is to create a higher authority that enforces peace and cooperation by coercing people with the threat of lethal force. That is, if two groups argue, it results in conflict and destruction unless some higher authority resolves the argument. And if two of those higher authorities argue, then it has to go higher. The very highest authority is the *sovereign*, who has the last word and can enforce cooperation.

A classic example is the "tragedy of the commons".¹⁶ About the time that philosophers like Thomas Hobbes (1588-1679) were thinking about sovereignty, every village in England had a *common*, an area of land set aside for grazing by each villagers' cow.

The dilemma was that if a villager put an extra cow on the common, then there's less grass per cow. But two slightly thinner cows are better than one fat cow. So each villager has an incentive to add an extra cow.

So a village of self-interested maximizers will increase the number of cows without limit. The result is overgrazing, erosion, the destruction of the common as a resource, starvation, and the depopulation and destruction of the village.

With no higher authority to stop villagers from adding cows to the common, eighteenth-century England had numerous villages ruined in this way. Eventually the Parliament regulated grazing by privatizing the commons.

Experiences like the tragedy of the commons has convinced many people of the need for a higher authority to prevent such short-sighted greediness by self-interested maximizers.

To oppose this view, it is often argued that a higher authority to supposedly reduce conflict by controlling self-interested maximizers is, itself, a cause of conflict and death. Nazism, Communism, and other other utopian “isms” are often cited as examples of how higher authority can go wrong.

2.2. *Reciprocal altruism*

Robert Trivers³⁰ suggested an alternative to cooperation by coercive authority. He was motivated by many biological examples of cooperation without higher authority.

A startling example: fish that swim around the inside of a shark’s mouth, cleaning the scraps of meat left from the shark’s last meal. The shark gets free dental care, and the fish gets a free meal. Both benefit by cooperating. Even with no higher authority to enforce this cooperation, the shark nonetheless refrains from eating its dentist.

Reciprocal altruism raises many questions, most importantly: how did it come to be like this? If you or I reach into a hungry shark’s mouth with a toothbrush, we lose an arm. A human dentist floating in the ocean, armed only with a toothbrush, would appear to have a conflict with a hungry shark — a conflict not easy to resolve, but which this humble fish seems to have mastered.

Biological examples of cooperation, with no higher authority to enforce it, have motivated political scientists and economists to see if reciprocal altruism can be induced between human communities.

2.3. *Iterated prisoner's dilemma*

To better understand how cooperation comes about, the abstract mathematical game of Iterated Prisoner’s Dilemma (IPD) is widely studied in biology, economics, political science, and artificial intelligence.¹ IPD has been cited to explain the Watergate scandal of 1972-74,²⁴ the Cold War of 1945-1990,⁶ and life in the trenches of the First World War of 1914-1918.¹

In its basic form, the Prisoner’s Dilemma is a two-player game where each player has two choices, cooperate or defect. The payoff matrix satisfies the following conditions,¹ shown in Fig. 1:

- $T > R$ and $P > S$ (Defection always pays more)
- $R > P$ (Mutual cooperation beats mutual defection)
- $R > (S + T)/2$ (Alternating doesn’t pay)

Many parameters satisfy the conditions of an IPD shown in Fig. 1. This paper uses $T = 5$, $R = 4$, $P = 1$, and $S = 0$, as shown in Fig. 2.

In the *iterated* prisoner’s dilemma, the game in Fig. 1 is played not just once, but many times, with a memory of previous iterations. This allows the possibility

	Cooperate	Defect
Cooperate	R R	T S
Defect	S T	P P

Fig. 1. The payoff matrix for the 2-choice 2-player prisoner's dilemma game. The values T, R, P, S must satisfy $T > R > P > S$ and $R > (S + T)/2$.

	Cooperate	Defect
Cooperate	4 4	5 0
Defect	0 5	1 1

Fig. 2. The payoff matrix for the 2-choice 2-player prisoner's dilemma game as used here.

of retaliation and mutual cooperation. The iterated game is widely studied because it contains the basic elements of many real-world situations.

2.4. *The evolution of cooperation*

To gain insight into why mutual cooperation does or does not emerge in IPD, much research follows political scientist Robert Axelrod.¹ He imposed a co-evolutionary dynamic on a population of trial strategies playing against each other.² In this approach, a computer maintains a population of trial strategies. An *evaluation function* judges the quality of each trial strategy in the population.

Instead of a human programmer writing that evaluation function, in co-evolution a strategy's fitness is evaluated by its peers in the same population. To learn a game like IPD, a strategy's fitness is its average payoff from playing all the other strategies in the same evolving population. As the population improves, its members become more challenging opponents, and the evaluation function becomes more discerning. The aim is to set up an escalating arms race of innovation.

Co-evolutionary learning has been applied to board games like Checkers^{8,9} and Backgammon,^{10,25} as well as certain non-game tasks,^{21,23,26,27,29} such as schedule optimization^{18,20} and creating a sorting algorithm.¹⁹

Co-evolution applied to IPD shows how mutual cooperation can come to dominate, even without any central authority to enforce cooperation.² Figure 3 shows a typical run in which cooperative strategies eventually take over.

Axelrod^{1,2} found that cooperation tends to dominate because strategies can discern between other nice cooperators (with whom they cooperate) and mean nasty defectors (whom they punish). This discernment is summarized in the strategy of “Tit for Tat”, which simply does to the other player whatever that player did on the previous round of iterated prisoner’s dilemma.

Axelrod hinted that peace and cooperation could spontaneously come about if everybody followed a strategy along the lines of “Tit for Tat”. That is, to punish the exploiters. Recent empirical research on human subjects has demonstrated that “cooperation can flourish if the public-spirited majority can punish freeloaders”.¹⁴

In politics, many national leaders have followed the general idea of keeping the peace by the threat of retaliation against aggressors. Teddy Roosevelt (US President 1901–1909) pushed for a large navy with the slogan “Walk softly, but carry a big stick”. Ronald Reagan (US President 1981–1989) also pushed for a large navy with the slogan “Peace through strength”.

2.4.1. *Tit-for-Tat cannot make cooperation permanent*

This “evolution of cooperation” is occasionally punctuated by mass extinctions of the apparently stable community of cooperative strategies, and their sudden replacement by more ruthless strategies.^{4,11,17,22} This occurs without any external cause (like a meteor impact). Figure 4 shows such a collapse in cooperation.

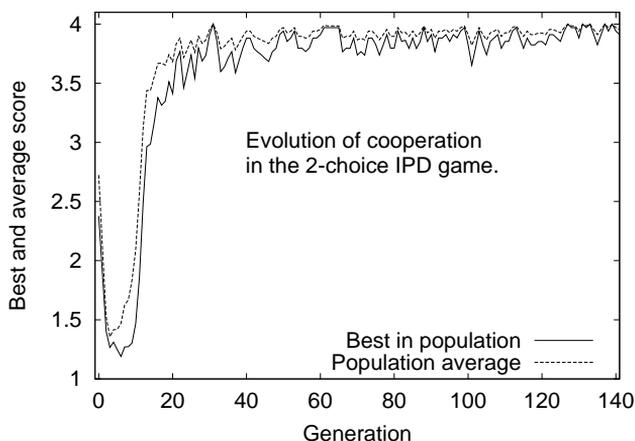


Fig. 3. A typical run with 2 choices of full cooperation or full defection, showing the “evolution of cooperation”.²

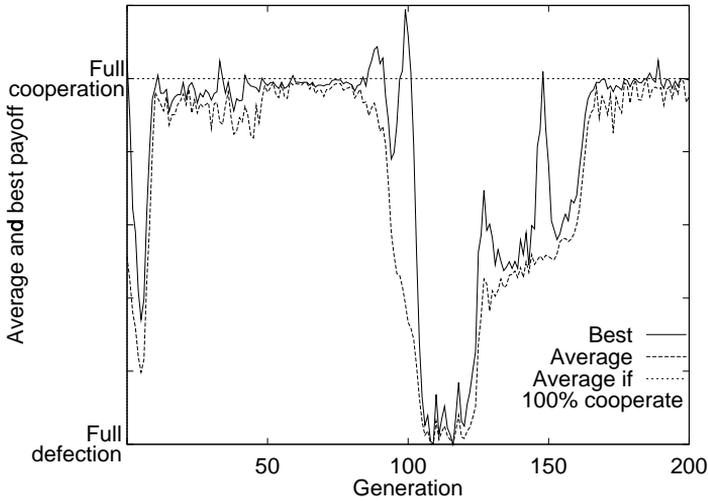


Fig. 4. Cooperative strategies don't always maintain their dominance after the evolution of cooperation. This is due to the atrophy of their ability to retaliate against exploiters.¹¹ There is no payoff to keeping your claws sharp if you never need them.

Cooperation falters in this way because the retaliatory capability embodied in Tit-for-Tat tends to atrophy.¹¹ Without strategies that defect, there is no payoff to maintaining the ability to retaliate against defectors. Why keep your claws sharp if you never need them? When that ability has withered, then defectors can invade because they no longer risk retaliation. “The price of peace is eternal vigilance”, according to Thomas Jefferson (US President 1801–1809).

If everyone (or sufficiently close to everyone) pursued a strategy like Tit-for-Tat, then cooperation is likely to dominate. The problem is, it doesn't last forever.

This begs the question, is there some strategy that, if nearly everyone pursues it, could bring about spontaneous mutual cooperation and world peace?

Alas, in general there is no strategy (or finite collection of strategies) that is uninvadeable in the noise-free 2-choice IPD game.^{5,13} That is, if most of the players in a population pursue a particular strategy, then it might induce cooperation for a time, but other strategies can always invade by emulating the dominant cooperators. Those emulators can then be exploited by a non-cooperative invader strategy that takes over. The collapse in payoff in Fig. 4 is always possible. So no cooperative community can be guaranteed to remain cooperative forever.

Of more interest to politics, is how to induce cooperation out of conflict in the first place?

2.5. *Number of choices*

The classic IPD game offers only two choices, full conflict or full cooperation, with nothing in between. This simple two-choice case is the most studied.^{1,2}

Being able to choose intermediate actions between full cooperation and full defection affects whether or not cooperation comes to dominate. The real world is often complicated, with intermediate degrees of cooperation, rather than all-or-nothing.^{15,28} Players can usually choose among varying degrees of cooperation.

Some recent papers add intermediate levels of cooperation to the IPD game.^{4,15,17,28,32} The general conclusion from those studies is that mutual cooperation still happens with intermediate choices, but not as often nor as stable as in the 2-choice case: "... the average level of cooperation remains high throughout, so although not stable, as a group these strategists are cooperative ... there is a restless interplay of invasions and re-invasions between strategists that are broadly cooperative".¹⁵ Similarly, "... any cooperative behavior that did arise did not tend toward complete cooperation".¹⁷

When talking about nuclear weapons, the phrase "... did not tend toward complete cooperation" has an ominous tone.

Previous work shows that intermediate choices tend to prevent full cooperation, but without explaining *why*. This paper verifies these earlier observations, and demonstrates two mechanisms for why.

Following previous work,^{15,17} intermediate payoffs are a linear interpolation of the 2-choice game. For example, if each player has four choices (instead of the usual two), the payoff matrix is in Fig. 5, which shows the payoff to the player on the left, player A.

Figure 5 illustrates two important points:

- the four corners of an n -choice $n \times n$ payoff matrix are the same payoffs as for the 2-choice game, and;
- any 2×2 sub-matrix of an $n \times n$ payoff matrix is itself a 2-choice IPD, satisfying the 2-choice conditions in Fig. 1.

		Player B			
		+1	$+\frac{1}{3}$	$-\frac{1}{3}$	-1
Player A	+1	4	$2\frac{2}{3}$	$1\frac{1}{3}$	0
	$+\frac{1}{3}$	$4\frac{1}{3}$	3	$1\frac{2}{3}$	$\frac{1}{3}$
	$-\frac{1}{3}$	$4\frac{2}{3}$	$3\frac{1}{3}$	2	$\frac{2}{3}$
	-1	5	$3\frac{2}{3}$	$2\frac{1}{3}$	1

Fig. 5. Four-choice IPD payoff matrix, showing payoff to player A, is interpolated from the 2-choice game. Here -1 is full defection and $+1$ is full cooperation. Note that any 2×2 sub-matrix of the $n \times n$ game is itself a 2-choice IPD.

This paper will mostly consider the case where each player has eight evenly discretized choices of cooperative level, giving a 8×8 payoff matrix.

For the choice of payoff values used in Fig. 5, if two players A and B choose to cooperate at levels c_A and c_B respectively, then player A's payoff is given by Eq. (1) (and symmetrically for player B).

$$p_A = 2.5 - 0.5c_A + 2c_B, (-1 \leq c_A, c_B \leq 1). \quad (1)$$

There is an n -choice analogy to the 2-choice conditions in Fig. 1. In the n -choice case, payoffs must satisfy Eqs. (2) through (4) to be a Prisoner's Dilemma game. These conditions generalize from the 2-choice conditions in Sec. 2.3. Here, c_A and c_B are the cooperative levels of two players A and B, and p_A is the payoff to player A.

- Ripping off a cooperator pays. For $c_A < c'_A$ and constant c_B :

$$p_A(c_A, c_B) > p_A(c'_A, c_B). \quad (2)$$

- Mutual cooperation pays more than mutual defection. For $c_B < c'_B$ and constant c_A :

$$p_A(c_A, c_B) < p_A(c_A, c'_B). \quad (3)$$

- It's not worth taking turns. For $c_A < c'_A$ and $c_B < c'_B$:

$$p_A(c'_A, c'_B) > \frac{1}{2}(p_A(c_A, c'_B) + p_A(c'_A, c_B)). \quad (4)$$

3. Experimental Setup

This paper uses a population of 100 trial strategies for 2-player IPD. Each generation, member of the current population (including itself). A player's fitness is its average score over all this generation's games. If a player makes it to the next generation with no change, it still gets a new fitness by playing all the members of that new generation.

3.1. How to represent a strategy

One convenient way to represent a strategy for IPD is as look-up table, with an action to take for every possibility of the opponent's earlier actions. However, for numerous intermediate levels of cooperation, such a look-up table would be large and cumbersome.

A more convenient way to represent an IPD strategy is as a feed-forward neural network. Each member of the population is a fixed-length array of floating-point numbers. These represent the weights and biases of a feed-forward neural network with a fixed architecture, with a single layer of hidden nodes and one output node. This paper uses 10 hidden nodes. More than that made no noticeable difference. Previous work found different behavior for 2 and 20 hidden nodes.¹⁷

In each iteration of IPD, players move simultaneously. They remember only a finite number of such moves into the past. Most previous work has used a memory of three previous rounds. This paper mostly uses a memory of just one round. This is partly because many strategies of interest (such as “Tit for Tat”) only use a memory of one previous round. Also, such a short memory makes it easy to visualize a strategy, as it is a function of only two variables: one’s own and the opponent’s previous cooperation level. This makes it easy to display strategies later in this paper, in Figs. 13 and 14.

For each past round that is remembered, the neural network has four inputs:

- One’s own previous level of cooperation, in $[-1, +1]$.
- The opponent’s earlier level of cooperation.
- An input which is 1 if the opponent exploited the player, and zero otherwise.
- An input which is 1 if the player exploited the opponent, and zero otherwise.

For the last two inputs, “exploited” is merely the sign of the difference in the previous cooperation levels. Even if a player is exploited by only a single fine-grained difference in cooperation, the corresponding input will still be 1. Thus, players can know immediately if they are being exploited, even by a small amount.

The last two inputs are just functions of the first two. So although each neural network has four input nodes, it is a function of two variables: one’s own previous level of cooperation and the opponent’s. This makes it easy to visualize a strategy.

Each individual is a neural network with one output node. The hidden nodes and the output node are all sigmoided, to stay in the range $[-1, +1]$, using the hyperbolic tangent $\tanh(x)$ as the sigmoid function.

After sigmoiding, the output is discretized so that all the values are equidistant from each other. For example, for 4 choices of cooperation, the output node’s value in the range $[-1, +1]$ would be adjusted to the closest of the values $(-1, -\frac{1}{3}, +\frac{1}{3}, +1)$. This discretized output is the player’s choice of cooperation for the current move. Thus, inputs for the next round are discretized.

On the very first round of a game of IPD, there are no previous rounds on which to base a decision about the next round. To fill this gap, each strategy’s genotype contains two extra values, that act as the imaginary pre-game “previous” levels.

Each of the ten hidden nodes has six weights: four input weights, an output weight, and a bias. An extra bias term for the final output, plus those two “pre-game” inputs, gives a genotype of 63 floating-point numbers.

3.2. Other parameters

The settings used here are quite conventional, and include the following.

Mutation is Gaussian, with a standard deviation of 0.5. The per-allele probability of mutation is 0.03, averaging around 2 mutations in a 63-allele genotype. Crossover is uniform, with probability of 0.5. An offspring suffers either mutation or crossover, but not both, so there is a 0.5 chance of being mutated.

Selection was stochastic universal selection with linear ranking.³ The best player in the population expects 1.2 offspring, and the worst expects 0.8 offspring. This weak selection pressure aims to prevent premature convergence.

Population size is 100. Replacement is 50%, so the worse half of the population is deleted and replaced by the offspring of the better half. The unchanged strategies in the best half are re-evaluated, the old fitness values don't get copied.

A modest number of runs at different parameters did not reveal any great sensitivity. An exception is that (as one would expect) high mutation combined with high selection pressure tended to cause premature convergence. To make this less likely, the replacement value is only 50% of the population.

3.2.1. *The shadow of the future*

The “shadow of the future” may sound like the next Star Wars sequel, but in the context of IPD it has a specific meaning.¹ If the number of iterations in a game of IPD is known in advance, then players have no incentive to cooperate on the very last move, because there is no risk of retaliation. So every player will want to defect on the last move.

But if every player thus plans to defect on the very last move, then there is no incentive to cooperate on the *second* last move (seeing as nobody will cooperate on the last move anyway). So every player will defect on the second-last move as well.

And so on back to the start of the game. Thus, cooperation only emerges if the game length is somehow kept uncertain.

A popular way to keep game length uncertain is to have a fixed probability of ending the game on each successive move.² That gives an average number of rounds per game, but any particular game's length is uncertain.

However, in this paper, individual strategies are represented in such a simple way (feed-forward neural networks) that they cannot count. This lets us use a fixed-length game, with no chance of the players figuring out when the last (or second-last) move will be. Here, game length is 150 iterations, comparable to the average length used by Axelrod.²

4. Results

4.1. *Classic two-choice evolution of cooperation*

For the two-choice game, as Axelrod² observed, starting out with a population of random strategies gives the behavior shown in Fig. 3. Average payoffs initially plummet as naïve cooperators are exploited. Exploiters soon run out of victims. Eventually, strategies that mutually cooperate but retaliate against defectors come to dominate. Recalling Fig. 5, when the whole population is mutually cooperating, then every player's payoff is $R = 4$.

Mutual cooperation remains dominant for long periods of time, punctuated by rare mass extinctions.^{4,11,17,22}

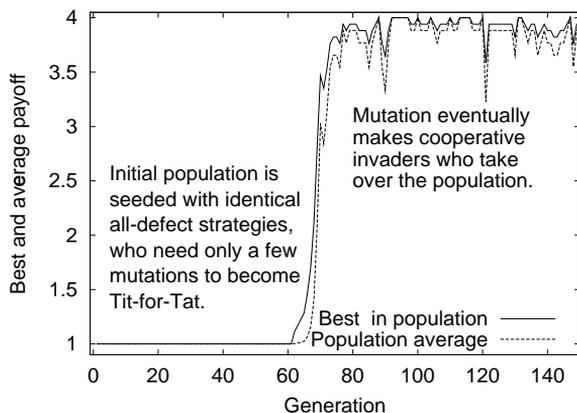


Fig. 6. This run's initial population are strategies that always defect, but which are only a few mutations away from the cooperative tit-for-tat strategy. Mutation eventually hits onto the right mutations to create a few successful invaders who take over the population, but it's a rather long wait.

Early in the run in Fig. 3, all four outcomes in the 2×2 payoff matrix are occurring. Later in the run, when mutual cooperation dominates, three of the four outcomes in the payoff matrix (those other than mutual cooperation) rarely occur.

Mutual cooperation is not inevitable, even in the 2-choice game. A population may get stuck on mutual defection for a long period of time. When this happens, only one of the four outcomes in the 2×2 payoff matrix is being sampled. This is not necessarily a permanent state of affairs, because the noise-free 2-choice IPD game does not have an uninvadeable strategy.^{5,13} Mutation may eventually produce an invader strategy that is more cooperative. However, this is an improbable and therefore rare event. Once a population has settled into full defection, mutation may take a long time (if ever) to produce a small number of suitable invader cooperators to take over.

To demonstrate, Fig. 6 shows a run where the initial population consists of strategies that always defect, but are only a few mutations away from the cooperative Tit-for-Tat strategy. Even in this case, where the all-defect strategy is specifically mutation-friendly, it takes some time (more than 60 generations) for mutation to produce some suitably cooperative invaders.

The point is, once behavioral diversity disappears, there is little variation for natural selection to increase cooperation. Increasing cooperation, when the whole population is already cooperating at a lower level, is rare. Mutation is a slow substitute for behavioral diversity.

4.2. Intermediate choice causes less cooperation

Earlier studies have found that allowing degrees of cooperation (rather than only full cooperation or full defection) still achieves some degree of mutual cooperation,

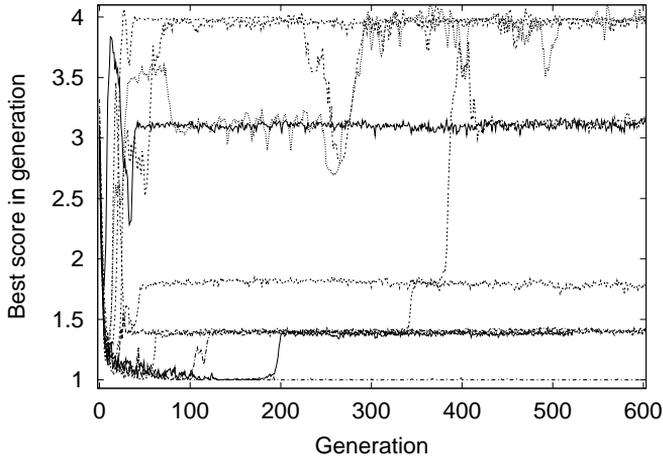


Fig. 7. Ten runs with 8 choices of cooperation. As previous studies have shown, there is usually some degree of mutual cooperation, but full cooperation is less likely.

but with a “restless interplay of invasions and re-invasions”,¹⁵ where “... any cooperative behavior that did arise did not tend toward complete cooperation”.¹⁷

With a choice of 8 levels of cooperation (instead of just 2), Fig. 7 shows the score of the best individual out of each generation for 10 different runs. This seems to agree with those earlier studies. Unlike the 2-choice case in Fig. 3, full mutual cooperation is less likely in Fig. 7.

Note that in Fig. 7, payoffs stick at a particular level for a while, before changing erratically and perhaps settling on another level. Only one of the ten runs is stuck at the lowest payoff, from full defection. But most runs stay stuck at an intermediate level for long periods.

The question is, if a population can usually increase its degree of cooperation from the bottom level, why doesn’t it keep increasing? Why not ratchet all the way up early in the run, and stay up?

4.3. *One cause: Lack of behavioral diversity*

As mentioned in Fig. 5, any 2×2 sub-matrix of the $n \times n$ payoff matrix of the n -choice IPD is itself a 2-choice IPD. One can verify (not shown here) that using a 2-choice payoff matrix whose values are a sub-matrix of the 8×8 game will still result in mutual cooperation, just like in Fig. 3.

This suggests that the population should work its way up one level at a time, since the top corner of one 2×2 sub-matrix is only the bottom corner of another 2×2 sub-matrix. But this doesn’t consistently happen with the 8-choice game.

Consider a run from Fig. 7, for which Fig. 8 shows the best and average payoff for each generation. Payoff remains at a constant level, often for long periods. Brief and rapid changes punctuate these stable periods.

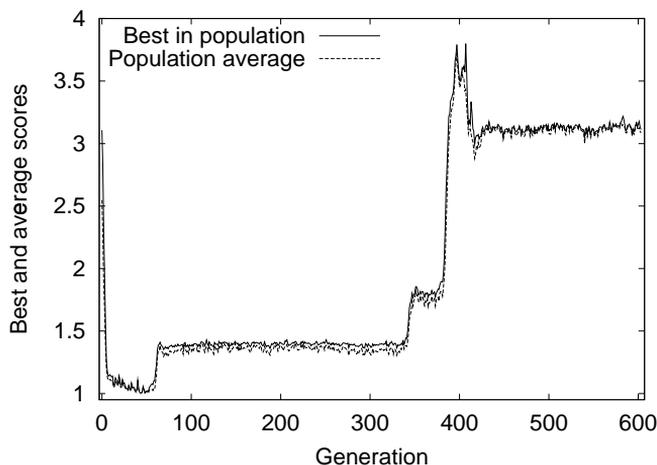


Fig. 8. This shows the population's best score in a run with 8 choices of cooperation, one of those in Fig. 7. The degree of cooperation is usually stuck on some intermediate level, with occasional rapid changes, up or down.

What outcomes are happening? From the 8×8 payoff matrix, one can plot a histogram of which outcomes occur in all the games in a generation. Consider the period from generations 340 to 350 in Fig. 8, where average payoff goes up a notch. What choices are being made then? Figure 9 shows histograms of which choices (from the 8×8 payoff matrix) are being made during the course of generations 342 to 350. It is during this brief period that the payoff level in Fig. 8 goes up a notch.

In Fig. 9, there is little behavioral (i.e. phenotypic) diversity in generation 342. Almost everyone is doing the same thing, and choosing the second-lowest level of cooperation. With no variation in behavior, evolution stagnates.

Eventually, mutation does produce an invader strategy that can cooperate at the next highest level. This is analogous to when mutual defection dominates in the 2-choice game, and a mutation eventually happens which produces a more cooperative strategy which can invade, as was shown in Fig. 6.

The reason the 8-choice game takes so long to move up to the next level is because behavioral diversity disappears so quickly. In the 2-choice game (Fig. 3), there is a great deal of diversity early on due to the random initial population, and it explores all 4 possible outcomes. But in the 8-choice game, it doesn't explore 64 possible outcomes for long. Payoff rises to a particular level by making use of that initial diversity, but once that diversity dies out it remains stuck at the same level for a long period, until mutation eventually produces a suitable invader.

Similarly, if the 2-choice game happens to settle on all-defection, it takes a long period time for mutation to move events to the next level up (the highest, there being only 2 choices). This was shown in Fig. 6.

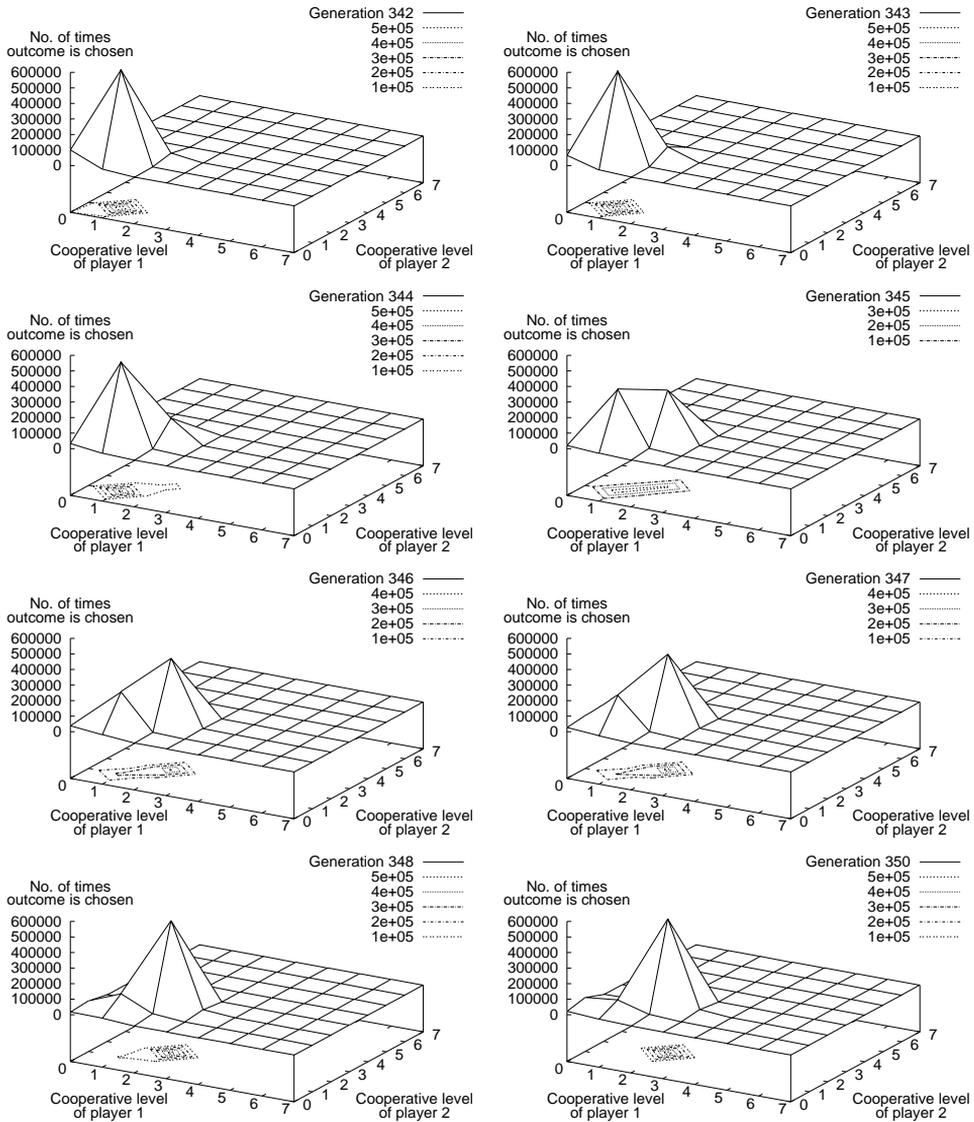


Fig. 9. Histograms of which choices (out of the 8×8 payoff matrix) are being made during generations 342 to 350 in the run shown in Fig. 8, during a brief change in payoff level.

4.3.1. Verification: More diversity promotes cooperation

We may test the notion that cooperation is less likely because a small population doesn't have the behavioral diversity to sample those $n \times n$ possible outcomes, for each round of IPD (and for each possible outcome, a player has n responses to choose from). More mutations (to create more behavioral diversity) should explore more outcomes, and make mutual cooperation more likely. Instead of a per-allele

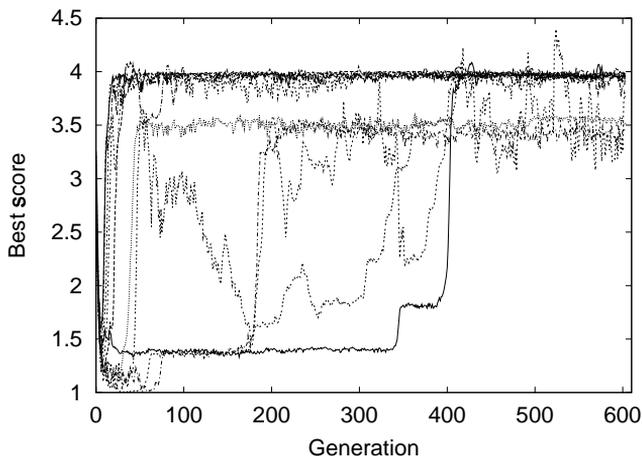


Fig. 10. Ten runs with 8 choices of cooperation, with more mutations, to cause extra behavioral diversity. All runs achieve the highest or second-highest level of cooperation.

mutation rate of 0.03 with Gaussian mutation step size of 0.5 (i.e. few big changes), this section uses a per-allele mutation rate of 0.2 with step size of 0.15 (i.e. many smaller changes).

Figure 10 shows 10 runs with this higher level of mutation. In contrast to the previous low-mutation runs in Fig. 7, here all ten runs achieve either the highest or second-highest level of mutual cooperation.

This verifies that more phenotypic diversity (simply by raising the mutation rate) makes mutual cooperation more likely in the 8-choice IPD. There are only 64 possible outcomes, so a population of a hundred (with rather high mutation) samples enough outcomes to discover mutual cooperation pays better.

However, more mutation is not a cure-all. With more possible choices, this simple expedient cannot hope to provide sufficient phenotypic diversity. Consider the case of an almost continuous IPD, where each player has 1024 choices. That makes more than a million possible outcomes to each round. A small population just cannot hope to sample that many possible outcomes. Even with the higher mutation rate, Fig. 11 shows that mutual defection is the most likely outcome.

4.4. Another cause: Stable non-symmetric strategies

Even with the higher mutation rate, one of the runs in Fig. 10 does not appear to settle on any particular level, but wanders around, as shown in Fig. 12. Why is that?

Closer inspection reveals that the population has formed two distinct species. Representatives from generation 520 of Fig. 12 are shown in Figs. 13 and 14. The entire population consists of a minority like that Fig. 13, and the majority are the other way around with a larger valley floor, as in Fig. 14.

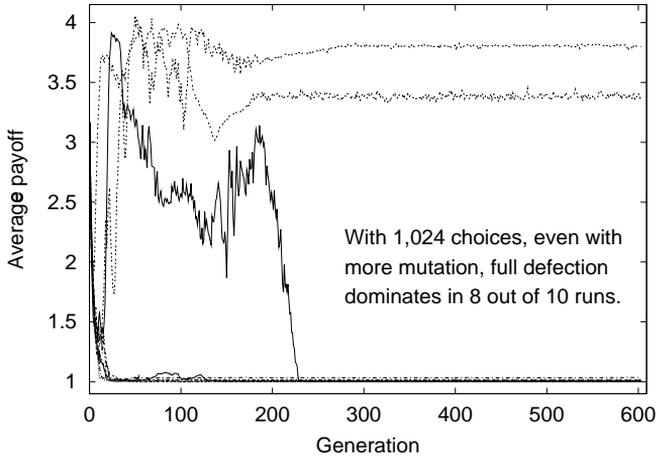


Fig. 11. Ten runs with 1024 choices of cooperation. Even with more mutations, just like in Fig. 10, there are simply too many possibilities to sample very many with a population of only 100. As a result of insufficient phenotypic (behavioral) diversity, mutual defection is once again the most common outcome, in 8 runs out of 10.

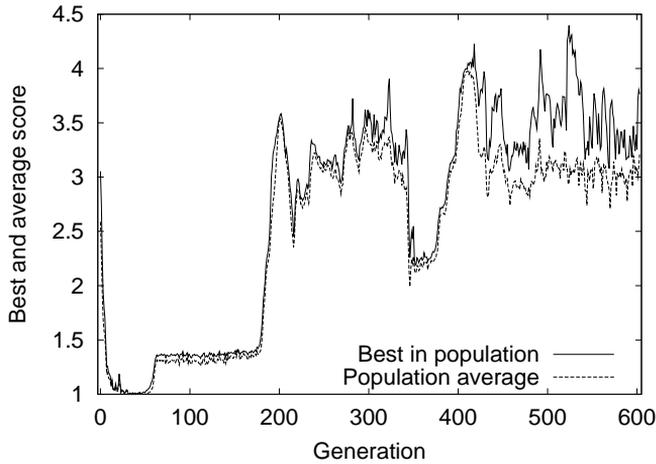


Fig. 12. This run, one of those in Fig. 10, does not appear to settle down to a constant level after generation 400.

Following each player’s moves round-by-round, each of these two types does not achieve full cooperation with its own type, but it does with the other type. As a result, they are in a stable symbiosis.

Following the first few moves, those 29% of individuals in generation 520 which look like Fig. 13 (plateau on the right) make their first move at the second-lowest level of cooperation. When both inputs are at that level, the output is also at the same level. So two such individuals achieve a low payoff against each other.

Generation 520, individual 0. 29% of the population is like this.

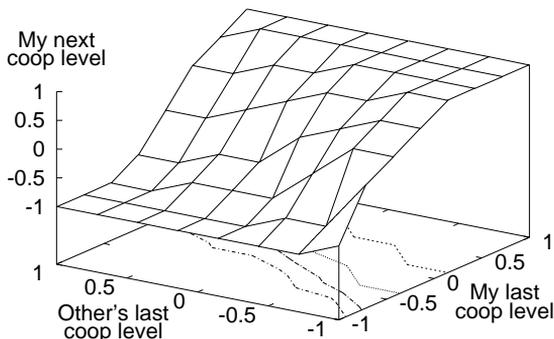


Fig. 13. Generation 520 of Fig. 12 consists of 29% of strategies like this.

Generation 520, individual 16. 71% of the population is like this.

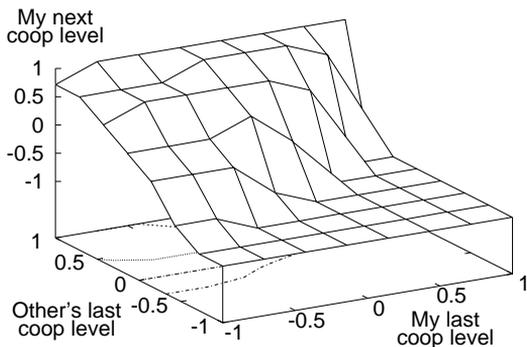


Fig. 14. Generation 520 of Fig. 12 consists of 71% of strategies like this.

Similarly, those 71% of individuals in generation 520 which resemble the individual in Fig. 14 (plateau on the left) make their first move at the second-highest level of cooperation. When both inputs are at that level, the output is also at the same level. So two such individuals achieve a rather high payoff against each other, but still not the highest and most profitable level of cooperation.

But when one type plays against the other, their first few rounds lead to one round of full mutual defection, followed by a gradual increase in the level of cooperation, until finally both play at full cooperation for the rest of the game.

The strategy in Fig. 13, after full mutual defection, responds not with another full defection, but with a level of cooperation one notch higher. That is, in Fig. 13, when both x and y is -1 (full defection), the z -axis is not also -1 but just above that. This element of forgiveness in the strategy of Fig. 13 allows it to increase cooperation with the other kind, but not against itself.

These two strategies are reminiscent of the hand-written strategy called “raise the stakes” that was created for an earlier study.²⁸ That strategy cooperated with itself. Here, in contrast, two symbiotic strategies have evolved that also raise their level of cooperation, but only against each other.

Why can’t mutation produce an invader, that cooperates with these two and also with itself, to take over? It turns out that the two strategies in Figs. 13 and 14 are surprisingly robust to mutant invaders. They are even robust against seeding the population with various hand-written strategies, such as Tit-for-Tat. Most of these would-be invaders are ruthlessly wiped out by the two symbiotic strategies already present.

Just like the 2-choice IPD,^{5,13} IPD with more choices does not allow a strategy (or a finite group of strategies) that are perfectly invasion-proof, unless there is random noise.³¹ Even so, there are still strategies which can survive for surprisingly long periods, such as “Tit for Tat”, because it is rare for mutation to create a suitable invader. The two strategies in Figs. 13 and 14 are similarly stable. The novelty is that they are asymmetric and symbiotic.

5. Discussion

5.1. *Too many outcomes to sample*

In IPD with only 2 choices, the 2×2 payoff matrix gives only 4 possible outcomes. Each possible outcome has only 2 choices, so with a memory of only one previous round, there are only 8 possible strategies. This small space is easy to search.

But for n choices, there are n^2 outcomes, each of which offers n choices. That gives n^{2n} possible behaviors. As soon as evolution weeds out the poor players and goes some way to converging, behavioral (phenotypic) diversity quickly becomes too low to sample any more than a tiny fraction of those n^{2n} possibilities.

When the evolving population converges on a particular level of cooperation, the population is as stable there as a 2-choice population would be if it were stuck on full defection. This is because any 2×2 sub-matrix of the $n \times n$ payoff matrix is itself an IPD. It takes a long time for mutation to create an invader who can push cooperation up a notch, as shown in Fig. 6.

Causing more behavioral diversity — implemented here by a higher mutation rate — can push a population away from such stagnation, as shown in Fig. 10. But this simple fix cannot help when there are simply too many choices, as shown in Fig. 11.

5.1.1. *Behavioral diversity is not genetic diversity*

It is important to distinguish between behavioral and genetic diversities. High genetic diversity, i.e. a high diversity in genotypes, does not necessarily result in high behavioral diversity, especially when there is no one-to-one mapping between genotypes and phenotypes.

Because we can only evaluate a phenotype's fitness, and use that as a proxy for the corresponding genotype's fitness, it is the phenotypical diversity that is most important here.

Our earlier work¹² showed how, on this same learning problem, a high genotypical (genetic) diversity can perversely correspond to a low phenotypical (behavioral) diversity.

5.2. *Non-symmetric ensembles of strategies*

Previous work on stability in IPD focused on a population of identical strategies, all doing the same thing with each other, and how well such a symmetric strategy can maintain itself in the face of invaders.

No finite group of strategies in IPD is invasion-proof, unless there is random noise,^{5,13} even for the n -choice game.³¹ Nonetheless, strategies can maintain themselves for long periods. For example, in the 2-choice game, strategies like Tit-for-Tat can maintain cooperation for long periods, and only occasionally does an invader take over, such as in Fig. 4.

With the richer possibilities of the n -choice game, there are symbiotic strategies that can stay dominant for considerable periods.

If a population was 100% of the strategy in Fig. 13, or 100% of that in Fig. 14, a non-cooperative invader would find it easy to take over. But with both together in the same population, the asymmetry of their symbiosis makes it very hard for an invader. They have a good cop/bad cop relationship. Such a situation is unforeseen by previous work.

Even though those strategies resemble the hand-written "raise the stakes" strategy described in previous work,²⁸ here they raise their cooperativeness only with each other, but not with themselves.

The conclusion is that the n -choice game offers many more possible ensembles of quite stable strategies, that an evolving population may find itself attracted to. Such islands of stability act as stop signs on the road to full cooperation. That these extra islands of stability exist only in the richer n -choice game is another reason more choices make full cooperation less likely.

5.2.1. *Aside: Sympatric and allopatric speciation*

Incidentally, the two strategies in Figs. 13 and 14 are an example of speciation in the absence of mating restrictions, such as those caused by geographic barriers.

In biology, speciation often occurs as a result of some restriction on mating, usually some geographic barrier such as an ocean or mountain range. Darwin's finches on the Galapagos archipelago are the archetypical example. Mating restrictions are the cause of speciation. This is called *allopatric* speciation.

Sometimes, however, species emerge within overlapping regions, in the absence of geographic barriers to mating. The four species of tern (a kind of sea bird) on the

Black Sea island of Jorilgatch⁷ give a good example. In this *sympatric* speciation, inter-species mating restrictions are a result of speciation, but not the cause.

6. IPD and the National Missile Defense

This section is rather speculative. The way humans think has little to do with these simplified, abstracted simulations. It is dangerous to read too much into the results of a computer model. Nonetheless, we are looking for useful insights. Sections 6.3 and 6.4 tentatively discuss the two mechanisms from Secs. 4.3 and 4.4.

6.1. *The anti-ballistic missile treaty*

The game of iterated prisoner's dilemma is a simplified version of many real-world situations that offer a choice between cooperation and exploitation. The use or non-use of nuclear weapons is one such situation that has attracted attention.⁶

Previous work shows that the n -choice game is less likely to reach full cooperation.^{15,17} This is verified in Fig. 7.

In the context of nuclear weapons, this suggests that allowing only two choices — between full peace or total all-out full-scale nuclear war — would make the peaceful alternative more likely. This view argues that a defense against incoming ballistic missiles, which would allow for limited nuclear war, would actually make such a war more likely, not less.

This kind of thinking is therefore in favor of the 1972 Anti-Ballistic Missile (ABM) treaty, which outlaws a missile defense. With no way to limit nuclear war, the stark choice between either complete non-use of nuclear weapons and all-out war makes any kind of war less likely. This notion was called Mutually Assured Destruction (MAD).

The ABM treaty allows each signatory to have two limited ABM systems, one to protect its capital city and another to protect an area where land-based missiles are based. The two protected areas must be at least 1300 kilometers (812 miles) apart, to prevent any effective nation-wide defense.

The treaty forbids a general, national defense against incoming nuclear missiles. This allows rapid escalation of a nuclear war to destroy cities, by protecting the land-based missiles that would do so.

The ABM treaty allows either side to withdraw from the treaty six months after giving notice, and the current Bush administration gave its notice to withdraw in December 2001. The argument is that, since the Cold War finished, Russia and the United States are unlikely to attack each other, so the only plausible threat is an attack from a small “rogue” nations.

6.2. *The big question*

Would the proposed National Missile Defense increase the chance of a partial nuclear war between the major powers? Even if a missile defense stops a small-scale one-off

attack from a “rogue” nation, this is cold comfort if it causes even a partial nuclear war with a big nation like Russia or China.

Recent research has found that intermediate levels of cooperation make full cooperation less likely.^{15,17} At first glance, that suggests a missile defense would make partial nuclear more likely. The problem is, those papers did not illuminate *why*.

This paper looks more closely at the reasons why full cooperation happens less often with intermediate choices. Two reasons found by this paper are described in Secs. 6.3 and 6.4. Of course, these are speculative as IPD is only a very simplified and abstract model.

6.3. *Getting stuck on a particular level*

Section 4.3 found that a small population of similar players cannot explore all $n \times n$ possible outcomes of the n -choice game. They tend to get stuck on an intermediate level of cooperation. As each 2×2 sub-matrix of the $n \times n$ payoff matrix is itself an IPD, after it's settled on a particular level of cooperation, it is just as difficult to shift it to a higher level as it is for the 2-choice game to go from full defection to full cooperation — not impossible (as shown in Fig. 6) but still very difficult.

A tendency to maintain the status quo is admirable when the status quo is a happy one. However, when the status quo is not happy, each additional intermediate level of cooperation thus becomes an additional obstacle when trying to increase the overall amount of cooperation.

As each 2×2 sub-matrix of the $n \times n$ payoff matrix is itself an IPD, it's just as hard to raise the status quo by just one notch, as it is for the 2-choice game to go from full defection to full cooperation. And similarly, on the way down, those intermediate levels act as obstacles to things getting worse.

So, at first glance, this suggests that allowing partial levels of cooperation will tend to slow down the escalation of conflict, and also to prolong an existing conflict.

In particular, if full cooperation is already in place, then introducing partial degrees of cooperation will not by itself be enough to cause cooperation to decline.

6.3.1. *Relevance to nuclear missile defense*

For nuclear weapons, one could argue that a missile defense would be a bad idea in the middle of a war, but a good idea during peaceful times. That way, it could prevent the rapid escalation of a nuclear envisaged by Mutually Assured Destruction (MAD) and the Anti-Ballistic Missile Treaty. But only if the prevailing level of cooperation is already at its best, on the day the missile defense is switched on.

If, instead, the missile defense is switched on in a mood of hostility, it may simply add more obstacles to improving that level of cooperation.

6.4. *Symbiotic strategies*

Most research on IPD has looked at a population of identical strategies. But with multiple choices, the richer environment allows for a population of different symbiotic groups to emerge. Section 4.4 describes this.

The two strategies in Sec. 4.4, whose actions are graphed in Figs. 13 and 14, dominate the population. Few strategies can invade that are good against both of these. Together, they maintain their dominance. So although no strategy or collection of strategies in noise-free IPD is ever uninvadeable,^{5,13} even with multiple choices,³¹ these two can still maintain their dominance for a long time.

Section 6.4.1 dissects those two strategies in more detail, and Sec. 6.4.2 discusses how they maintain themselves.

6.4.1. *The miserly raise-the-stakes, and grouchy tit-for-tat*

Looking back at the two strategies in Figs. 13 and 14, consider the diagonal across the $x - y$ plane when both players have just played the same level of cooperation.

Figure 13 resembles the “raise the stakes” strategy created for an earlier study,²⁸ except that Fig. 13 starts at that second-lowest level. If both it and the opponent do a round at that second-lowest level, then it responds by staying there. As a result, it keeps on that low level of cooperation when playing against itself, but otherwise it’s very forgiving and goes more cooperative.

To summarize Fig. 13, if the opponent is miserly then this strategy is just as miserly. But if the opponent gets above the bottom two levels, then this strategy raises the stakes to full cooperation. Call it Miserly-Raise-the-Stakes.

In contrast, the strategy in Fig. 14 starts out generously, making its first move at the second-highest level of cooperation. But take a look at the diagonal across the $x - y$ plane of Fig. 14 where both previous levels are the same — sure, it goes up evenly to full cooperation, but there’s a sudden drop to the lowest level if the opponent tries to exploit this strategy. It’s a cliff along the diagonal.

To summarize Fig. 14, it starts out generously and reciprocates whatever cooperation it receives, but at the first hint of being exploited it goes to full defection. Hard but fair, it retaliates massively at the slightest provocation.

In this sense, the touchy strategy in Fig. 14 is a paranoid version of Tit-for-Tat. It keeps a finger on the button. And strategies like this make up 71% of generation 520 in Fig. 12. Call it Grouchy-Tit-for-Tat.

In a game where nearly three-quarters of the players are so grouchy, one might expect a lot of conflict. But that doesn’t happen.

6.4.2. *Relevance to nuclear missile defense*

Thomas Jefferson claimed “the price of peace is eternal vigilance”. In Fig. 4, Tit-for-Tat’s retaliatory ability established full cooperation, but the peaceful environment

atrophied that ability. Why keep your claws sharp while there's full cooperation? With no fear of retribution, exploiters invaded and dominated.

In contrast, the minority Miserly-Raise-the-Stakes strategy in Fig. 13 keeps the majority Grouchy-Tit-for-Tat 14 extremely vigilant. The first strategy initiates a low level of cooperation, but quickly raises its cooperativeness when it gets punished.

In this way, it constantly probes the retaliatory ability of the majority strategy. The claws on Fig. 14 are thus kept razor sharp.

Nonetheless, it points towards a tense stand-off, with the possibility of a short, sharp conflict. This is little different from the ABM Treaty's policy of Mutually Assured Destruction (MAD) which also encourages a tense stand-off, except that the conflict is unlikely to be partial.

7. Conclusion: War no More Likely than with 2 Choices

The 1972 Anti-Ballistic Missile (ABM) treaty ensured that nuclear weapons had only two levels of cooperation: full peace or total war. The proposed National Missile Defense (NMD) would have the effect of providing intermediate degrees between those two extremes choices. Would that make a war more likely?

The abstract game of Iterated Prisoner's Dilemma (IPD) captures the basic elements of many conflict situations, including the nuclear arms race.⁶ Previous papers have found that for IPD with intermediate choices between full cooperation and full defection, achieving full mutual cooperation is less likely.

Some may interpret this to mean the ABM treaty is good, and missile defense is bad, because less than full cooperation suggests a partial nuclear war. Even a small nuclear war is unacceptable.

This paper studies the mechanisms for why intermediate choices make full cooperation less likely. This paper describes two mechanisms.

First, there is tremendous inertia to stay at the status quo. Any 2×2 sub-matrix of the $n \times n$ payoff matrix is itself an IPD game. So raising the level of cooperation by just one notch in the n -choice game is just as difficult as the 2-choice game going from full defection to full cooperation — difficult, but not impossible.

So introducing intermediate choices has the effect of adding obstacles to any change, either for better or for worse. This suggests that introducing a missile defense during happy times may act as an obstacle to any deterioration in the diplomatic climate. However, introducing one during a period of tension may act as one more block to improving relations. Therefore, the addition of intermediate choices is not, by itself, going to make conflict more likely.

Secondly, although no single strategy is invasion-proof in IPD,^{5,13} even for the n -choice game,³¹ there exist strategies that can dominate an evolving population for long periods of time. The best-known example in the 2-choice game is a strategy called "Tit-for-Tat", which creates and maintains full mutual cooperation thanks to its ability to retaliate against exploiters.²

Unfortunately, previous work¹¹ has shown that the cooperative climate thus created will itself tend to atrophy Tit-for-Tat's retaliatory ability. This eventually allows exploiters to invade and take over, such as in Fig. 4. This demonstrates Thomas Jefferson's observation that "the price of peace is eternal vigilance".

However, previous work has generally considered a population of identical strategies. Intermediate choices in IPD allows asymmetric but symbiotic strategies, each needing the other to resist invading exploiters.

The pair of strategies found here has a certain resemblance to the hand-written "raise the stakes" strategy.²⁸ But it more strongly resembles the old Tit-for-Tat strategy, except that the minority strategy of the pair merely keeps the claws sharp on the majority, which is always poised to retaliate against any exploiter.

This situation is qualitatively little different from the 2-choice dominance of Tit-for-Tat. Once again, introducing intermediate choices does not, by itself, make conflict more likely.

At first glance studies of IPD might seem to support the dire notion that intermediate choices between full peace and total war would make a partial war more likely. But this paper's closer look at some of the mechanisms involved do not support that notion. Although the abstract games of IPD captures only the basic elements of real situations, results here suggest that a missile defense would *not* make a major war any more likely than the old two-choice policy of Mutually Assured Destruction (MAD) embodied in the 1972 Anti-Ballistic Missile treaty.

References

1. R. M. Axelrod, *The Evolution of Cooperation* (Basic Books, New York, 1984).
2. R. M. Axelrod, The evolution of strategies in the iterated prisoner's dilemma, *Genetic Algorithms and Simulated Annealing*, Chap. 3, ed. L. Davis (Morgan Kaufmann, Los Altos, California, 1987) 32–41.
3. J. E. Baker, Reducing bias and inefficiency in the selection algorithm, *Proc. 2nd Int. Conf. Genetic Algorithms*, ed. J. J. Grefenstette (Lawrence Erlbaum Associates, July 1987) 14–21.
4. P. S. S. Borges, R. C. S. Pacheco, R. M. Barcia and S. K. Khator, A fuzzy approach to the prisoner's dilemma, *Biosystems* **41**, 2 (1997) 127–137.
5. R. Boyd and J. P. Lorberbaum, No pure strategy is evolutionarily stable in the repeated prisoner's dilemma game, *Nature* **327** (7 May 1987) 58–59.
6. S. J. Brams, *Superpower Games* (Yale University Press, 1985).
7. M. Braun, *Differential Equations and Their Applications* (Springer-Verlag, 4th edn., 1993).
8. K. Chellapilla and D. B. Fogel, Evolution, neural networks, games, and intelligence, *Proc. IEEE* **87**, 9 (September 1999) 1471–1496.
9. K. Chellapilla and D. B. Fogel, Anaconda beats Hoyle 6-0: A case study competing an evolved Checkers program against commercially available software, *Proc. 2000 Cong. Evolutionary Comput.* (IEEE Press, 6–9 July 2000) 857–863.
10. P. J. Darwen and J. B. Pollack, Co-evolutionary learning on noisy tasks, *Cong. Evolutionary Comput.* (IEEE Press, July 1999) 1724–1731.

11. P. J. Darwen and X. Yao, On evolving robust strategies for iterated prisoner's dilemma, *Progress in Evolutionary Computation, Vol. 956 of Lecture Notes in Artificial Intelligence*, ed. X. Yao (Springer, Berlin, 1995) 276–292.
12. P. J. Darwen and X. Yao, Does extra genetic diversity maintain escalation in a co-evolutionary arms race, *Int. J. Knowledge-Based Intell. Eng. Syst.* **4**, 3 (2000) 191–200.
13. J. Farrell and R. Ware, Evolutionary stability in repeated prisoner's dilemma, *Theoretical Population Biology*, **36** (1989) 161–166.
14. E. Fehr and S. Gächter, Altruistic punishment in humans, *Nature* **415** (2002) 137–140.
15. M. Frean, The evolution of degrees of cooperation, *J. Theoret. Bio.* **182**, 4 (1996) 549–559.
16. G. Hardin, The tragedy of the commons, *Science* **162** (1968) 1243–1248.
17. P. G. Harrald and D. B. Fogel, Evolving continuous behaviors in the iterated prisoner's dilemma, *Biosystems* **37** (1996) 135–145.
18. J. W. Herrmann, A genetic algorithm for minimax optimization problems, *Congress on Evolutionary Computation* (IEEE Press, July 1999) 1099–1103.
19. W. D. Hillis, Co-evolving parasites improve simulated evolution as an optimization procedure, *Artif. Life 2, Vol. 10 of Santa Fe Institute Studies in the Sciences of Complexity*, eds. C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen (Addison-Wesley, 1991) 313–323.
20. P. Husbands and F. Mill, Simulated co-evolution as the mechanism for emergent planning and scheduling, *Proc. 4th Int. Conf. Genetic Algorithms*, eds. R. K. Belew and L. B. Booker (Morgan Kaufmann, July 1991) 264–270.
21. H. Juillé and J. B. Pollack, Co-evolving intertwined spirals, *Proc. 5th Ann. Conf. Evolutionary Programming* (MIT Press, 1996) 461–468.
22. K. Lindgren, Evolutionary phenomena in simple dynamics, *Artif. Life 2, Vol. 10 of Santa Fe Institute Studies in the Sciences of Complexity*, eds. C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen (Addison-Wesley, 1991) 295–312.
23. H. A. Mayer and R. Schwaiger, Evolutionary and coevolutionary approaches to time series prediction using generalized multi-layer perceptrons, *Congress on Evolutionary Computation* (IEEE Press, July 1999) 275–280.
24. D. Muzzio, *Watergate Games: Strategies, Choices, Outcomes* (New York University Press, 1982).
25. J. B. Pollack and A. D. Blair, Co-evolution in the successful learning of Backgammon strategy, *Mach. Learning* **32**, 3 (1998) 225–240.
26. M. A. Potter, K. A. De Jong and J. J. Grefenstette, A coevolutionary approach to learning sequential decision rules, *Proc. 6th Int. Conf. Genetic Algorithms* (Morgan Kaufmann, July 1995) 366–372.
27. C. W. Reynolds, Competition, coevolution and the game of tag, *Artif. Life 4*, eds. R. A. Brooks and P. Maes (MIT Press, 1994) 59–69.
28. G. Roberts and T. N. Sherratt, Development of cooperative relationships through increasing investment, *Nature* **394** (9 July 1998) 175–179.
29. C. D. Rosin and R. K. Belew, Methods for competitive co-evolution: Finding opponents worth beating, *Proc. 6th Int. Conf. Genetic Algorithms* (Morgan Kaufmann, July 1995) 373–380.
30. R. Trivers, The evolution of reciprocal altruism, *Quar. Rev. Bio.* **46** (1972) 35–57.
31. X. Yao, Evolutionary stability in the N-person iterated prisoner's dilemma, *Bio-Systems* **37**, 3 (1996) 189–197.
32. X. Yao and P. J. Darwen, How important is your reputation in a multi-agent environment, *Proc. 1999 IEEE Int. Conf. Syst. Man Cybern.*, Tokyo, October 1999, 575–580.