To Promote Cooperation: Reduce It
The Prisoner's Dilemma and Space

Hugo Taylor

April 2025

1 Introduction

The prisoner's dilemma is a simple two-player game with two options—cooperate or defect. Its core is that any player, regardless of what choice their opponent makes, is always better off defecting. For example, the points could be assigned as follows: 3 points each if they both cooperate, 5 points for defecting when the other cooperates, 0 points for cooperating when the other defects, and 1 point each for both defecting. Placing these values in a payoff matrix Figure 1 makes the inevitable conclusion of mutual defection clear. Despite the fact that both would be better off with 3 points each rather than 1, if the players are rational self-interested agents, they will always know they are better off with the 5 points rather than 3 if the opponent cooperates or the 1 point rather than 0 if the opponent defects. This is the Nash Equilibrium, a strategy that neither player can benefit from changing. The game can be further generalised by assigning variables to the payoffs seen in Figure 2 provided they are supported by the inequalities $T > R > P > S$, ensuring defective dominance, and $2R > T + S$, ensuring mutual cooperation is in sum optimal.

$3, 3$	$5, 0$
$0, 5$	$1, 1$

Figure 1: Payoff Matrix Example

R, R	T, S
S, T	P, P

Figure 2: Payoff Matrix Genralised

The idea of a prisoner's dilemma is as mathematically flexible as it is in its application to the real world. Problems from nuclear disarmament to corporate deals to prisoners ratting one another out, from which the game gets its namesake, can be explained through the dilemma. Yet despite its vast application in society, in its initial state the dilemma's ultimate outcome is still mutual defection. There have been a number of methods to explain the emergence of cooperation in society, but the one I will be exploring is space.

2 The Setup

Let us consider some n-by-n length grid. On each square of the grid lives a strategy which will only interact with the strategies around it. Being a simple strategy, it may only either cooperate or defect with all of its neighbours. Yet despite being a simple strategy, it still yearns to learn and grow, and so it will imitate success it sees. After each round of play, it will tally up its score and compare that to the scores of every strategy surrounding it. It will then swap its strategy to the highest scoring strategy it can see, unless it itself is the highest scoring. Running a simulation to follow this logic, you will see cooperation and defection grow and shrink across the grid like bacteria on a petri dish until they reach a final semi-stable conclusion, typically one iterating back and forth between the same two states. In contrast to the typical prisoner's dilemma where cooperation is entirely unable to survive, here it can cling on to life or even flourish if given the right circumstances.

Whether cooperation thrives or dies is determined by a couple of factors but most importantly the payoff matrix the game is simulated with. Despite the endless possible values we could choose to input into our payoff matrix, functionally there are only 63. This is because there can only be so many possible interactions between strategies. Indeed, you can calculate the outcome of any contest between two squares of opposite strategy by considering how many of that same strategy are surrounding it. For example, consider a cooperator surrounded by 3 other cooperators and a defector surrounded by 2 other defectors. They would therefore have a self-interaction score of 3/4 and 2/4, respectively. Let those scores be known as C, for cooperators, and D, for defectors. From that we can see that their total payoffs would be $CR + S(1 - C)$ for cooperators and $DP + T(1 - D)$ for defectors. If those values are the same, then neither strategy changes, but if one scores greater than the other, it will convert the other to that strategy.

Considering as well that any strategy may well be neighbours with another of the same strategy in a different circumstance, the values of self-interaction that are possible are 0, 1/4, 2/4, 3/4 and 1 (in this simulation). With both cooperators and defectors being capable of any of these 5 values, the number of unique interactions becomes 25. Of course, not all of these interactions will occur in reality, and many that do will give the same result regardless of payoff. For example, a C of 2/4 vs a D of 2/4 will always become defective as $2R + 2S < 2T + 2P$ is established by our fundamental inequalities. In fact, the majority of possible interactions are inevitably defective by nature, with only 1 being inherently cooperative ($C = 1, D = 1$). This leaves 9 inequalities to be determined by the payoff matrix. The number of unique possible strategies from this set are not immediately obvious, however, as many of them overlap, one being true proveing others true. To simplify this process, I used a variety of logic and programs to arrive at my final count of 63 unique strategies. (Though this number is variable depending on how rigid your definition of unique is) I then put them all through a simulation to find the final fraction of cooperation present.

3 Results

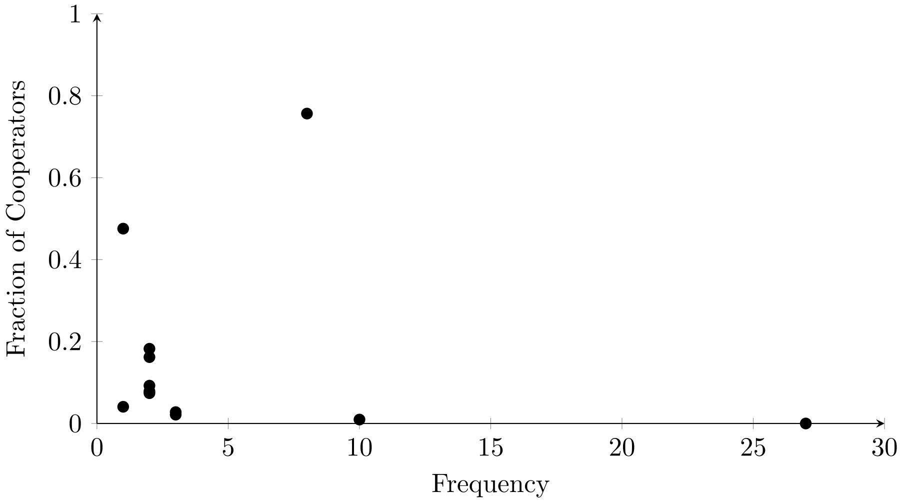

The results were shockingly clear. They split the payoffs into four clear camps: growth, unstable, stalemate and failure. The game starts on a 20-by-20 board with 200 cooperators and 200 defectors randomly distributed. Quickly the cooperators in this random noise are 'culled', replaced with defectors. Being spread out sporadically instead of concentrated in groups where they are strongest, they are very vulnerable to attack. For the most common occurrence, failures, this means total annihilation. Despite drawing or even winning in some inequalities, those that they lost were so abundant they were unable to sustain themselves and were completely wiped out.

The next step up are the stalemates. While most of their numbers were wiped out in the initial culling, some small clusters of them were able to remain. Their payoffs are barely high enough to create a stalemate with the surrounding defectors, but they are entirely unable to grow, resigning them to relatively low cooperation rates. With higher payoffs the culling phase may turn out more favourably for stalemates with more, often larger, clusters remaining, but none exceed 0.2 cooperation.

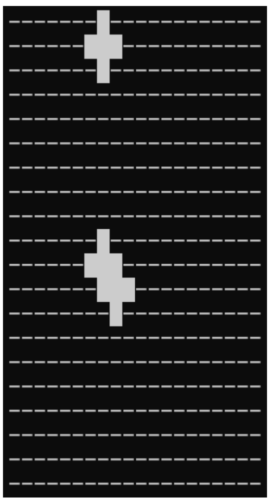

Figure 3: Stalemate with Low Payoff

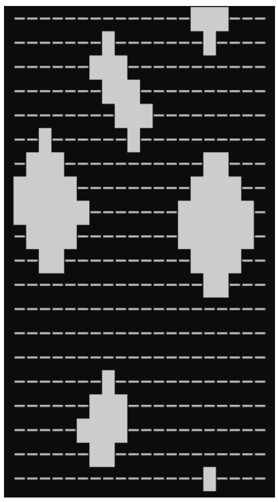

Figure 4: Stalemate with High Payoff

Next is a single payoff in a category entirely of its own. The unstable payoff is superior to the stalemate in its ability to grow. However, unlike pure growth payoffs, which are able to sustain their growth, due to certain unfavourable payoffs, the unstable strategy loses as many from its population as it gains each round. While every other payoff reaches a stable or semi-stable state within 30 iterations, often less, the unstable payoff continues to change for well over 100 iterations with no immediately obvious pattern or signs of stopping. It averages around 0.48 cooperation, being seemingly no better or worse than random noise.

Finally, the growth payoffs. As mentioned before, these payoffs are capable of growing their population and sustaining that growth. While most are initially reduced to small clusters after the culling, these clusters grow to fill the entire grid. The only places defectors are able to survive are on the edges between clusters. For this reason, growth payoffs have a unique property. Higher payoffs can reduce total cooperation.

Despite initially seeming unintuitive, this trait has a firm logic. All growth payoffs are capable of covering the entire board with cooperation from just a single cluster. However, if more than one cluster survives the culling, then edges between the two will begin to form. While defectors cannot survive being faced with a single edge of cooperators, being faced with two edges from both sides supplies them with enough cooperators to exploit to survive. Hence, the more edges, the more defection survives and the less cooperation. Therefore, it is better for growth strategies to survive the culling with as few clusters as possible, meaning that higher payoffs that bring greater culling survival become a burden, not a boon.

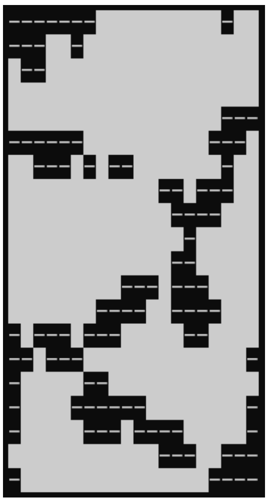

Figure 5: Growth with High Payoff

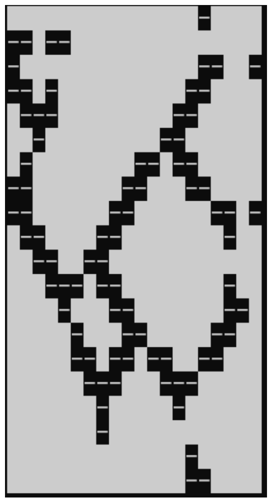

Figure 6: Growth with Low Payoff

For the same reason, starting with a lower number of cooperators is often beneficial to growth payoffs. Figure 4 shows this relationship with the final cooperative fraction against the starting cooperative fraction across grid sizes from 5 by 5 to 60 by 60. As the size of the grid grows arbitrarily large, it appears that the worst possible starting fraction of cooperators is exactly 0.5, while the most efficient for cooperators returned vs inputted converges on near 0.

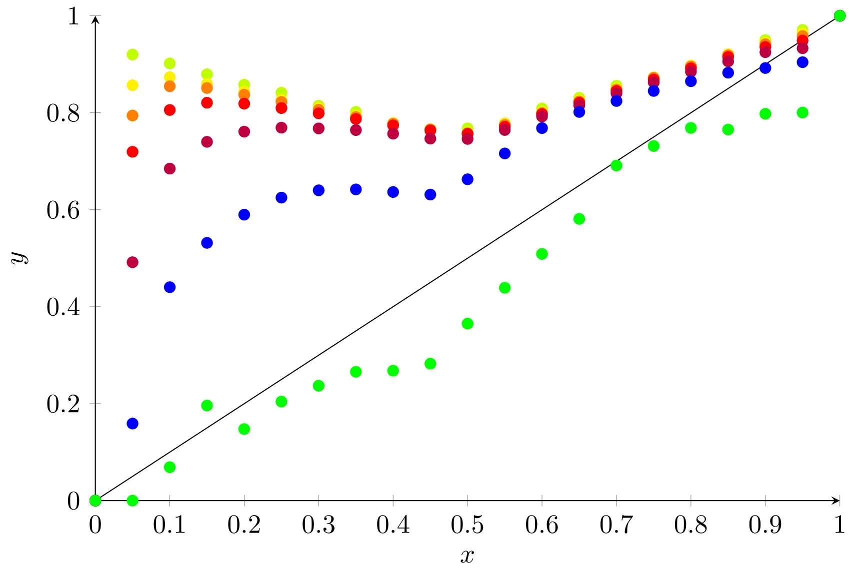

4 Conclusion

These unintuitive results are interesting but also just the surface of all this model has to explore. Each of the separate payoffs has endless nuance to their returns, which I could spend hours describing and exploring, and that without modification! Consider how the results would change if we tiled the plane with triangles or hexagons instead. I neglected corners in my simulations, having each edge wrap back around to the opposite side, but considering walls and other obstacles in the simulation would highlight an entirely new side to the dilemma. That's without adding other separate factors such as noise or signalling to the mix.

However, ultimately these results do not have much application to the real world, as they are entirely the product of odd interactions created by a niche setup. Despite their lack of use, they may still give us a good reminder that things which may seem obviously helpful to growth can, in fact, be instead a hindrance.

Thank you for your time in reading.