The Evolution of Cooperation in Finite, Growing Populations

Kevin Loope (University of Wisconsin, Madison) and Allison Bishop (Princeton University)

Nebraska REU, 2004

We begin with a population of individuals whose reproductive fitnesses partially depend on the results of a game that each individual plays repeatedly against every other individual. Suppose this population consists of two types of strategists, a portion of the population always choosing to play strategy A and the rest of the population always choosing to play strategy B. In any given generation, the fitness of each player is then determined by which strategy they are playing, how that strategy fares against itself and against the other strategy, and how many A and B strategists there are in the population at the given time. For example, if A is a strategy that does quite well against itself but poorly against strategy B, then an A strategist in a population of predominantly B strategists would do poorly, whereas an A strategist in a population of predominantly A strategists would do well. We define the payoff matrix for the game as follows:

$$\begin{array}{c c} & \begin{array}{cc} A & B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \begin{pmatrix} a & b \\ c & d \end{pmatrix} \end{array}$$

Here, a is the payoff for an A strategist against another A strategist, b is the payoff for an A strategist against a B, c is the payoff for a B against an A, and d is the payoff for a B against a B. These payoffs may not be constants, as they may be dependent upon the number of rounds of the game played or on the probability of playing another round. As an alternative to assuming a fixed, constant number of rounds between each pair of players, we assume a constant probability $0 < n < 1$ of playing another round. The first round is always assumed to occur. Thus, the probability of playing round t is n^{t-1}. We can then calculate total payoffs as geometric sequences. For example, if a player is receiving a payoff of R in every round against an opponent, the player's total payoff is:

$$\sum_{t=0}^{\infty} Rn^t = \frac{R}{1-n}$$

In order to define the fitness of an A or B strategist in a population of N total players with i players of strategy A, we introduce a parameter $0 < w < 1$ which controls the fraction of reproductive fitness which depends on the game. We assume that in matters unrelated to the game, A and B strategists are exactly alike, and thus have some fitness due to other factors which is constant, and we normalize this amount to be 1. We then define the total fitness as a weighted average of this constant fitness and the fitness due to the game, giving weight w to the game. If we define f to be the fitness of an A strategist in the given population and g to be the fitness of a B strategist, then f and g are the following functions of N,i,w, and n:

$$f(N, i, w, n) = 1 - w + w\frac{a(n)(i-1) + b(n)(N-i)}{N-1}$$

$$g(N, i, w, n) = 1 - w + w\frac{c(n)i + d(n)(N-i-1)}{N-1}$$

We note that players do not play against themselves, hence they have $N - 1$ opponents. Now, we will model the evolution of such a finite population using phases of reproduction and death in each generation. In their article, **Emergence of cooperation and evolutionary stability in finite populations**, Nowak, Sasaki, Taylor, and Fudenberg define and analyze such a model for the case of a population at carrying capacity. We will extend their model slightly to allow for population growth up to the upper bound of carrying capacity. Our reproductive phase we will define to coincide with theirs precisely: in each generation, which will be a time step, a single individual is chosen to reproduce an identical offspring which will be added to the population for the next time step. The probability of choosing any particular individual to reproduce is proportional to the fitness of that individual, and thus the probability that an A offspring will result is the number of A strategists multiplied by the fitness of an A strategist, divided by the total fitness of the population. The probability of acquiring a B offspring is similarly defined.

The death phase of our model we define differently. In Nowak, Sasaki, Taylor, and Fundenberg's model, an individual is randomly chosen for death and promptly is eliminated from the population. Therefore, the total population is always constant. In our model, an individual is also randomly chosen (so the probability of choosing an A individual is i/N for a population of N with i individuals playing A) but that individual is only eliminated with probability N/C, where C is the carrying capacity. Thus the individual survives into the next generation with probability $1 - N/C$. This probability is not dependent on fitness, and ensures that the population will always reach carrying capacity eventually, but will never exceed it. Once the population reaches carrying capacity, our model becomes exactly like the other model. As noted in the article, this process has only two possible end states if carried on long enough: either a homogeneous population of A strategists at carrying capacity or a homogeneous population of B strategists at carrying capacity will result.

The state of the population at any point can be completely described by two parameters, N and i. After one time step, one reproduction and possible death, there are five possible outcomes for a population initially in state (N,i): (N+1,i),(N+1,i+1),(N,i+1),(N,i-1),(N,i). Since $0 \leq i \leq N$, it is possible to visualize all of the possible states as a triangle, with the row number corresponding to N and the entry across the row corresponding to i. This allows us to describe each state by a single parameter j, starting at the top and counting off j's across the row and then jumping to the next row. Thus,

$$j - t - 1 = i$$

$$\frac{N(N+1)}{2} = t$$

for t the largest triangular number less than j. Using this order convention, we can create a *transition matrix*, a square matrix, $(C+1)(C+2)/2$ on a side, with the $P_{j,j'}$ entry containing the probability of moving from state j to state j'. Each row of this matrix will contain five non-zero entries, one for each of the five possible outcomes of a time step. We define Y_j to be the probability of strategy A taking over the population (i=C) starting at state j. We have:

$$Y_j = P_{j,j-1}Y_{j-1} + P_{j,j}Y_j + P_{j,j+1}Y_{j+1} + P_{j,j+N+1}Y_{j+N+1} + P_{j,j+N+2}Y_{j+N+2}$$

with boundary conditions $Y_j = 0$ for all j corresponding to a state (N,0) for

$N \geq 0$ (the right edge of the triangle) and $Y_j = 1$ for all j corresponding to a state (N,N) $N \geq 1$ (the left edge of the triangle). This set of equations can be represented by a matrix equation $P\overline{Y} = \overline{Y}$ with the added boundary conditions and solved for $\overline{Y}$. We further claim that the solution is unique.

Proof of Unique Solution. First, we define our transition probabilities as follows:

$$x_1(N,i) = P_{j,j+1} = \frac{if[N,i]}{if[N,i] + (N-i)\,g[N,i]} \frac{(N-i)}{C}$$
$$x_2(N,i) = P_{j,j-1} = \frac{(N-i)\,g[N,i]}{if[N,i] + (N-i)\,g[N,i]} \frac{i}{C}$$
$$x_3(N,i) = P_{j,j+N+1} = \frac{(N-i)\,g[N,i]}{if[N,i] + (N-i)\,g[N,i]} \left(1 - \frac{N}{C}\right)$$
$$x_4(N,i) = P_{j,j+N+2} = \frac{if[N,i]}{if[N,i] + (N-i)\,g[N,i]} \left(1 - \frac{N}{C}\right)$$

We stipulate that payoffs are non-negative, so fitness is always positive. Consider the triangle of all values of Y(N,i). On the interior of this triangle, $0 < x_l(N,i) < 1$ for $l = 1,2,3,4$. The proof will be inductive; we know a unique solution exists for the row N=C (Fudenberg, Nowak, Sasaki, and Taylor), so we solve one row at a time, assuming we already know the solutions on the subsequent row. We also assume the solutions Y(N,i) of this row are non-negative, since there is always some chance that a single A strategist will ultimately defeat all opponents. Now, let Y(N,i) denote the solution at state (N,i). We know

$$\begin{aligned} Y(N,i) &= (1 - x_1 - x_2 - x_3 - x_4)\,Y(N,i) + x_1(N,i)Y(N,i+1) \\ &\quad + x_2(N,i)Y(N,i-1) + x_3(N,i)Y(N+1,i) + x_4(N,i)Y(N+1,i+1) \end{aligned}$$

Solve this for Y(N,i+1):

$$Y(N,i+1) = \frac{\sum_{l=1}^{4} x_l(N,i)}{x_1(N,i)} Y(N,i) - \frac{x_2(N,i)}{x_1(N,i)} Y(N,i-1) - \frac{x_3(N,i)}{x_1(N,i)} Y(N+1,i)$$

$$-\frac{x_4(N,i)}{x_1(N,i)}Y(N+1,i+1) \tag{1}$$

We are solving for row N, and we stipulate that $Y(N,0) = 0$. We can then express all other solutions along the row in terms of Y(N,1), assuming Y(N+1,j) for all j to be non-negative constants. We further assume all interior solutions are positive, so the only positive term above, $\frac{\sum_{l=1}^{4} x_l(N,i)}{x_1(N,i)}Y(N,i)$ must be strictly greater in magnitude than the negative terms. We will ignore the x_3 and x_4 terms since they are just negative constants from this perspective. We will trace the coefficient of Y(N,1) in the expression of each Y(N,j): for example, since $Y(N,0) = 0$, we see that

$$y(N,2) = \frac{\sum_{l=1}^{4} x_l(N,1)}{x_1(N,1)}y(N,1) - constants$$

and

$$\begin{aligned} Y(N,3) \quad = \quad & \frac{\sum_{l=1}^{4} x_l(N,2)}{x_1(N,2)}\left(\frac{\sum_{l=1}^{4} x_l(N,1)}{x_1(N,1)}Y(N,1) - \frac{x_3(N,1)}{x_1(N,1)}Y(N+1,1) - \right. \\ & \left. -\frac{x_4(N,1)}{x_1(N,1)}Y(N+1,2)\right) - \frac{x_2(N,2)}{x_1(N,2)}Y(N,1) - \\ & -\frac{x_3(N,2)}{x_1(N,2)}Y(N+1,2) - \frac{x_4(N,2)}{x_1(N,2)}Y(N+1,3) \end{aligned}$$

So the coefficient of Y(N,1) is:

$$\frac{\left(\sum_{l=1}^{4} x_l(N,2)\right)\left(\sum_{l=1}^{4} x_l(N,1)\right)}{x_1(N,2)x_1(N,1)} - \frac{x_1(N,1)x_2(N,2)}{x_1(N,2)x_1(N,1)}$$

Notice this is strictly positive because all x_l are strictly between 0 and 1. In general, we note

$$\frac{\sum_{l=1}^{4} x_l(N,i)}{x_1(N,i)} > \frac{x_2(N,i)}{x_1(N,i)}$$

whenever (N,i) is in the interior of the triangle. Reconsidering the recursive equation (without the terms depending on the next row):

$$Y(N,i+1) = \frac{\sum_{l=1}^{4} x_l(N,i)}{x_1(N,i)} Y(N,i) - \frac{x_2(N,i)}{x_1(N,i)} Y(N,i-1)$$

we claim the coefficient of the parameter Y(N,1) is always strictly positive. Since the coefficient of the parameter in the negative term, $-\frac{x_2(N,i)}{x_1(N,i)}Y(N,i-1)$, can be put over the common denominator $x_1(N,i)...x_1(N,2)x_1(N,1)$ and compared to the coefficient of the parameter in the positive term, $\frac{\sum_{l=1}^{4} x_l(N,i)}{x_1(N,i)}Y(N,i)$, we see that $\frac{\sum_{l=1}^{4} x_l(N,i)}{x_1(N,i)} ... \frac{\sum_{l=1}^{4} x_l(N,1)}{x_1(N,1)}$ contains all negative terms, and so the coefficient is ultimately strictly positive. To illustrate this further, we consider the expression for Y(N,4) in terms of Y(N,1), and note that the coefficient of Y(N,1) is:

$$\begin{aligned} &\left(\frac{\sum_{l=1}^{4} x_l(N,3)}{x_1(N,3)}\right)\left(\frac{\sum_{l=1}^{4} x_l(N,2)}{x_1(N,2)}\right)\left(\frac{\sum_{l=1}^{4} x_l(N,1)}{x_1(N,1)}\right) \\ -&\left(\frac{\sum_{l=1}^{4} x_l(N,3)}{x_1(N,3)}\right)\frac{x_2(N,2)}{x_1(N,2)} - \left(\frac{\sum_{l=1}^{4} x_l(N,1)}{x_1(N,1)}\right)\frac{x_2(N,3)}{x_1(N,3)} \end{aligned}$$

We could easily put this expression over a common denominator, expand terms, and then subtract, noting that our subtracted terms merely cancel some, but never all of our positive terms. If we continued to consider the coefficient of Y(N,1) in the expression for Y(N,6), or Y(N,7), or so on, we would find the similar situation of subtracting some, but not all of the positive terms over a common denominator, leaving us with a strictly positive coefficient. It is essential to note that we never subtract the same term twice, since to get a common denominator we multiply by x_1 expressions, whereas the numerator in the negative recursive term is an x_2 expression. Thus, when we solve the equation Y(N,N) = 1, we uniquely determine our parameter Y(N,1), hence demonstrating a unique solution to our system with the proper boundary conditions.

This proof is valuable in that it simultaneously illustrates a process for actually calculating our solutions, though there is a much simpler proof of the fact that our solution in unique. We note that our recursive equation has the property that if two solutions exist, say Z(N,i) and V(N,i), their difference, Z(N,i)-V(N,i) also satisfies the recursive equation. Thus, if Z(N,i) and V(N,i) are identical on the boundary, their difference is zero on the boundary. Now, if we assume their difference is non-zero at some point on the interior of the triangle, it must achieve a maximum on the interior. Since the value at any given point is determined as a weighted average of the values of its adjacent points, this maximum must be achieved not only at the point itself, but also at its neighbors. Hence, the function must in fact be constant, and hence Z(N,i)-V(N,i) is identically zero for every point on the triangle, hence the solution is unique.

Proof of Drift Case

From equation (1), we can compute:

$$\begin{aligned} y(N,i) \quad = \quad & \left(1 + \frac{g(N,i-1)}{f(N,i-1)} + \frac{N\left(1-\frac{N}{C}\right)}{(i-1)\frac{N}{C}}\frac{g(N,i-1)}{f(N,i-1)} + \frac{N\left(1-\frac{N}{C}\right)}{(N-i+1)\frac{N}{C}}\right) y(N,i-1) \\ & -\frac{g(N,i-1)}{f(N,i-1)}y(N,i-2) - \frac{N\left(1-\frac{N}{C}\right)}{(i-1)\frac{N}{C}}y(N+1,i-1) - \frac{N\left(1-\frac{N}{C}\right)}{(N-i+1)\frac{N}{C}}y(N+1,i) \end{aligned}$$

For the drift case, the weight on the game (w) is zero. We define drift as a mutation which has no effect on the payoffs received from the game, and therefore has no effect on fitness. In this case:

$$\frac{f(N,i)}{g(N,i)} = 1$$

So

$$
\begin{aligned}
y(N,i) \quad = \quad & \left(1+1+\frac{N\left(1-\frac{N}{C}\right)}{(i-1)\frac{N}{C}}+\frac{N\left(1-\frac{N}{C}\right)}{(N-i+1)\frac{N}{C}}\right)y(N,i-1)-y(N,i-2) \\
& -\frac{N\left(1-\frac{N}{C}\right)}{(i-1)\frac{N}{C}}y(N+1,i-1)-\frac{N\left(1-\frac{N}{C}\right)}{(N-i+1)\frac{N}{C}}y(N+1,i) \qquad (2)
\end{aligned}
$$

When $y(N,i)=\frac{i}{N}$, the right hand side of equation (2) is:

$$
\begin{aligned}
& \left(1+1+\frac{N\left(1-\frac{N}{C}\right)}{(i-1)\frac{N}{C}}+\frac{N\left(1-\frac{N}{C}\right)}{(N-i+1)\frac{N}{C}}\right)\frac{i-1}{N}-\frac{i-2}{N} \\
& -\frac{N\left(1-\frac{N}{C}\right)}{(i-1)\frac{N}{C}}\frac{i-1}{N+1}-\frac{N\left(1-\frac{N}{C}\right)}{(N-i+1)\frac{N}{C}}\frac{i}{N+1}
\end{aligned}
$$

which simplifies to $\frac{i}{N}$, satisfying the system. To calculate the solutions for non-neutral mutations, we can follow the parametrization procedure outlined earlier, solving along each row in terms of the parameter Y(N,1), solving for that parameter using Y(N,N)=1, and also requiring that Y(N,0)=0. Thus, for any carrying capacity, we can graph the solutions Y(N,1) for N ranging from 1 to carrying capacity, limited only by computing resources of processing time and memory. In the following graphs, we specify values of C (carrying capacity), n (probability of playing another round of the game), and w (weight of the game) and graph our solutions Y(N,1) in green, representing the probability that a single mutant playing strategy A will invade a total population of N players, the others playing strategy B, for the fixed values of w, n, and C. For comparison, we also plot red and blue lines, the red line representing the probability that a neutral mutation arising in the population would invade, and the blue line representing the same mutation as the green line, arising in the same population, but instead of assuming a carrying capacity of C we assume that N, the total population at the state, is in fact the carrying capacity. The x-axis represents values of N, and the y-axis the solutions Y(N,1). In other words, the value of the green line at the point N=5 represents the probability that a single A player arising in a total population of 5 will eventually invade, assuming a carrying capacity of C, while the value of the blue line at N=5 represents the probability that a single A mutant arising in a population of 5 will invade, assuming carrying capacity to be 5. All lines are calculated for the same values of n and w, though of course these values have no effect on the red line.

We considered the game of Prisoner's Dilemma, a standard model for biological simulations. In this game, a player faces off in a series of rounds against one opponent. A round consists of each player choosing to cooperate with his partner or to defect (not cooperate). Payoff values are represented by T, R, S, and P, generally with $T > R > P > S$ and $2R > S + T$. T is the payoff for defecting when your partner cooperates, R is the payoff for cooperating when your partner cooperates, P is the payoff for defecting when your partner defects, and S is the payoff for cooperating when your partner defects. Your payoffs for one round can easily be represented in the following matrix, with your moves on the left and your partner's on the top.

$$\begin{array}{cc} & \begin{array}{cc} C & D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \begin{pmatrix} R & S \\ T & P \end{pmatrix} \end{array}$$

Common values for (T, R, P, S) are (5, 3, 1, 0). We have used these values for most of our calculations, though we will later compare our results to the case of R=4, all other values remaining the same. In a single round game or a finite, known number of rounds, Always Defect (AllD) is the obvious best strategy because no matter your opponent's move, a defect move will always pay more than a cooperate. This is called *domination* and can easily be seen in the above matrix: each value in the second row is higher than the corresponding value in the first row. This result is rather uninteresting, and gives little insight into the occurrence and evolution of cooperation. However, when the number of rounds is not finite or is finite but not known to the players, the game becomes more interesting. Because total payoff is more important than relative payoff (scoring highly is more important than winning), strategies other than AllD can arise. There is the possibility that cooperating strategies can be successful since they, while playing each other, will score highly. For our model we pitted Tit For Tat (TFT), a much studied cooperator strategy, vs. AllD. TFT's strategy is to copy the previous move of his opponent, and his first move is always to cooperate. In the following graphs, the green line represents the probability that a single TFT player will invade a population of ALLD players with a carrying capacity of C as described above, and the blue line represents a single TFT player in the same population, but with carrying capacity N. The red line remains the drift case. The following graphs both have C=100, n=.75, and R=3, the only difference being that w=.2 and w=.1 respectively:

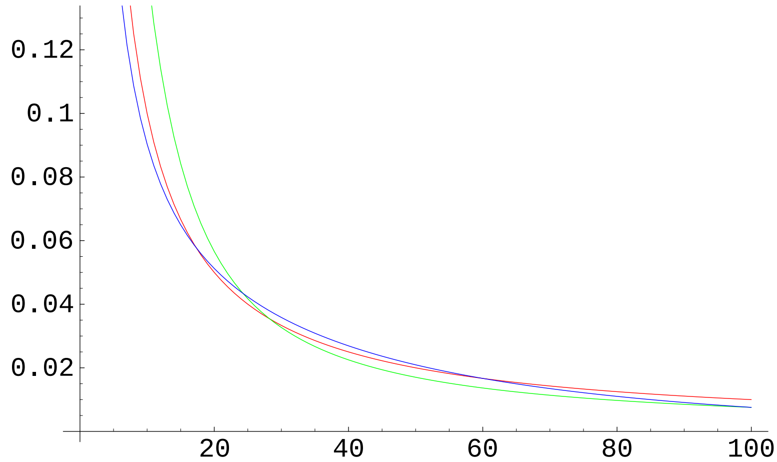

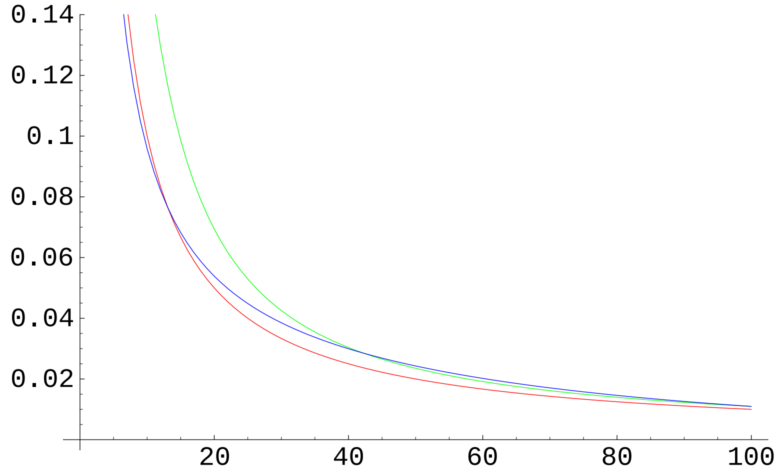

One can see from these that lower values of the selection pressure, w, will benefit the TFT strategy in the growth model. Lower w also seems to benefit TFT in the non-growth model, since the blue line is also higher in the second graph. What is perhaps most interesting about these graphs, however, is that they illustrate that it is possible for the red line to sometimes come between the green and blue lines, meaning that adding growth to the model can actually change when a mutation is evolutionarily favored, that is, if we use the intuitive definition of favored as simply more likely to take over the population than a neutral mutation. This effect is even more noticeable at small carrying capacities, such as C=30. The following graph has parameters C=30, w=.1, n=.75, R=3:

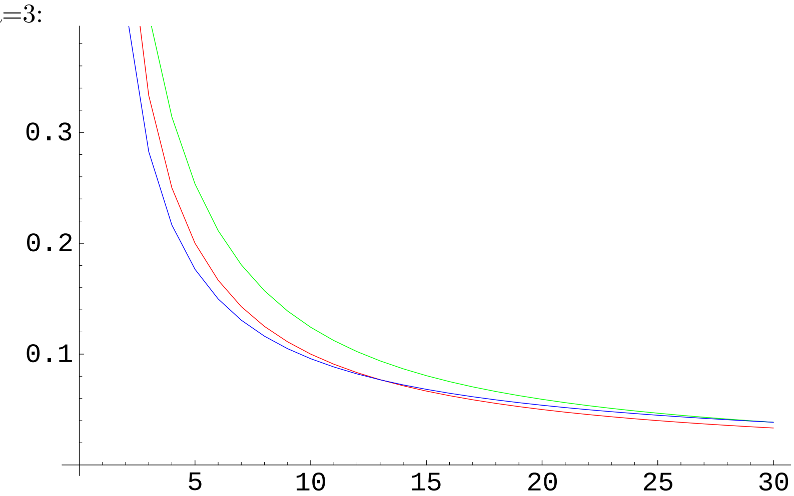

Here the growth case is clearly above drift throughout, while the non-growth case crosses drift between N=10 and N=15. We compare to the case C=30, w=.1, n=.6, R=3:

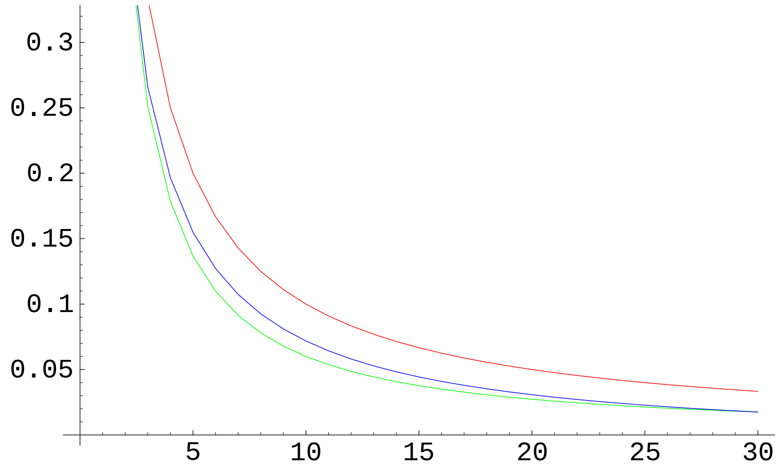

The above graph illustrates that population growth can hurt the chances of a TFT take-over when the probability of playing another game is not sufficiently high. This is not surprising, since lower values of n can result in domination. As long as $n < 1$, the TFT payoff vs. AllD is strictly less than the AllD payoff vs. itself, which is strictly less than the AllD payoff vs. TFT. Therefore, in order to avoid ALLD being the obviously better strategy, we must have the TFT payoff against itself be strictly greater than the AllD payoff vs. TFT. This requires $R/(1-n) > T + Pn/(1-n)$, which implies $(T-R)/(T-P) < n$. So for our values of T=5,P=1, and R=3, this becomes $n > 1/2$. We now consider the larger carrying capacity of C=125 to see how changes in the value of R might effect our results. Both graphs have n=.75, w=.2, and C=125. The only difference is the first graph has R=3, and the second graph has R=4:

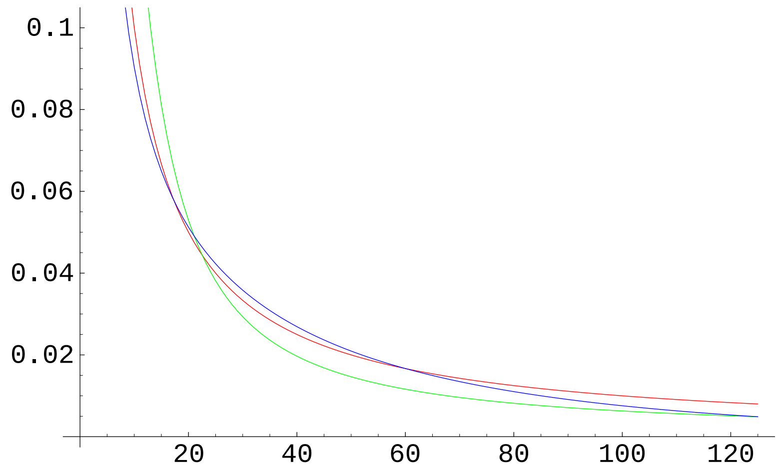

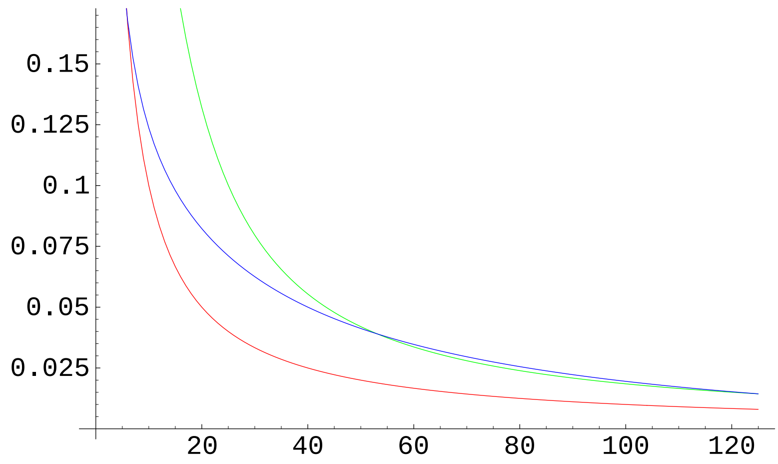

Comparing these graphs, we find that increasing the value of R in comparison to the other payoffs (though still maintaining $T > R > P > S$) allows TFT in a growing population to remain evolutionarily favored for a larger range of values of N. We now compare to graphs with all parameters being equal except carrying capacity. For the first graph, C=75, R=3, n=.75, and w=.2. The second graph, which we have already seen but include for convenience of comparison, has C=100:

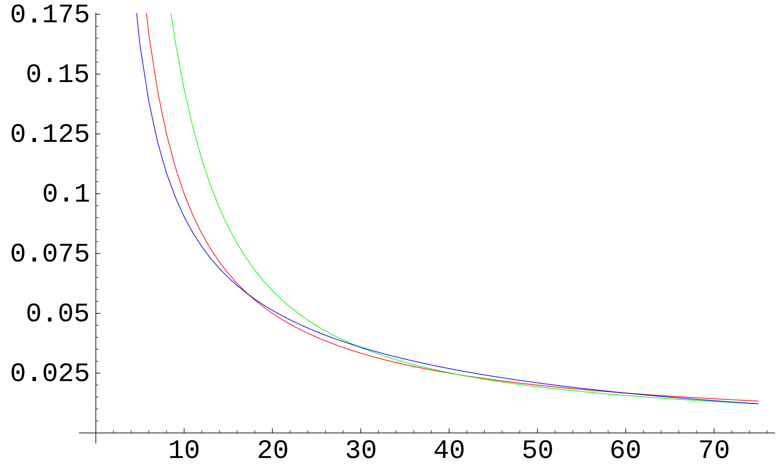

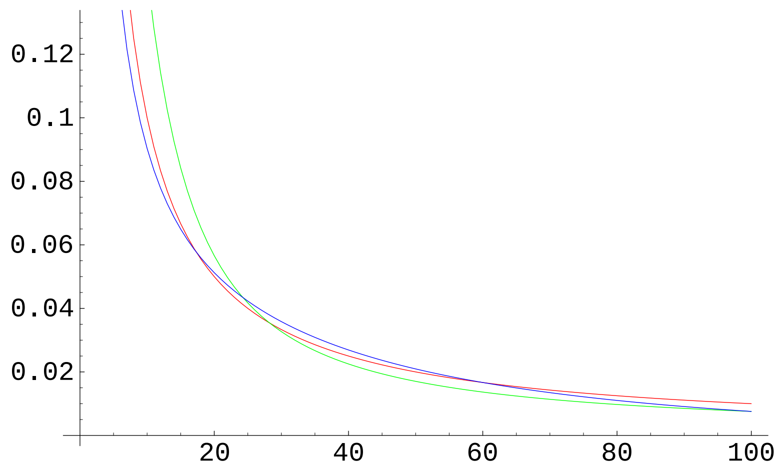

We notice that in the first graph, the growth case crosses below drift roughly at N=40, whereas in the second graph, this occurs closer to N=30. Thus, increasing C seems to make the green line steeper, causing the range of values of N for which growth benefits TFT to shrink. We conclude from all these comparisons that growth is most helpful to TFT's attempt at invading ALLD when carrying capacity is relatively low, selection pressure is relatively low, the probability of playing another round is relatively high, and the value of R is relatively high. To explain why growth might actually hurt TFT at lower and lower values of N for higher carrying capacities, we consider that death is a random process in our model, in that it is independent of fitness. Thus, for a single TFT player in a reasonably sized population of ALLD players, death mostly affects the ALLD players, since the chance of the single TFT player getting chosen for elimination is very low. When carrying capacity is relatively high in comparison to N, even the ALLD players will be very unlikely to die off, making it harder for a TFT player to increase fitness through reproduction, since though the TFT player may reproduce, the population will still be full of ALLD players who are dying very slowly, and whose fitnesses are increased by the presence of more TFT players. Thus, for TFT's proportion of the population to increase, a very delicate balance must exist between the benefit that TFT receives from playing itself and the number and benefit of ALLD players against TFT. As a result, it is helpful for ALLD players to be dying off more quickly and for R to be increased, as our results show.

We now examine a completely new model for a population of TFT and ALLD players. We similarly define w to measure the effect of the game on fitness, but this time we approximate the game to be infinite and take a time-averaged payoff, meaning that the payoff matrix is simply

$$\begin{array}{cc} & \begin{array}{cc} TFT & ALLD \end{array} \\ \begin{array}{c} TFT \\ ALLD \end{array} & \left(\begin{array}{cc} R & P \\ P & P \end{array} \right) \end{array}$$

This will obviously be to the benefit of TFT, but we can consider it as a reasonable approximation to very long, finite games. Under this assumption, the fitness f of a TFT player is $1 - w + w(R(i-1) + P(N-i))/(N-1)$ where N is the total population and i is the number of TFT players, as before. The fitness g of an ALLD player is simply $1 - w + wP$. We can then calculate the total fitness of a population N,i as:

$$T(N,i,w) = \frac{(i^2 - i)(Rw - Pw) + (N^2 - N)(1 - w + Pw)}{N - 1}$$

In this model, we define carrying capacity to be dependent on the population's fitness, with U representing our absolute upper bound for carrying capacity, regardless of fitness. More precisely, we define

$$C(N,i,U,w) = U\frac{T(N,i,w)}{N(1-w+wR)}$$

The denominator of this expression represents the maximum possible fitness for a population of size N, which occurs when all players are playing TFT. Next, we define a population adjusment term, notated ad:

$$ad(N,i,U,w) = Floor(N(1 - \frac{N}{C(N,i,U,w)}))$$

The Floor function returns the greatest integer less than or equal to its argument, so for example $Floor(2) = 2$, $Floor(2.8) = 2$, and $Floor(-.01) = -1$. The value of ad is always an integer, and may be zero. At every time step, we calculate the adjustment term, and if it is positive, we add the given number of players to the population, and if it is negative, we eliminate the given number of existing players. We choose which players to add or eliminate stochastically, with reproduction dependent on fitness and death random. If we are adding a player, the probability that it will be a TFT is $\frac{if(N,i,w)}{T(N,i,w)}$, and if we are eliminating a player, the probability that it will be a TFT is i/N. We make each of our choices for reproduction and death independently. We define a stable state to be a collection of values (N,i,U) for which ad(N,i,U,w)=0, given a fixed w. To be stable, a state's adjustment term without the floor function must then satisfy $0 \leq ad < 1$. By solving $ad < 1$ for i, we obtain the following condition for stablility:

$$(i^2 - i) < \frac{\frac{N^3(1-w+Rw)}{U} - (N^2 - N)(1 - w + Pw)}{Rw - Pw}$$

We note that as i takes integer values between 0 and N, the function $(i^2 - i)$ increases, and its minimum value is 0. When the right hand side of this inequality is negative, no value of i can be stable. By setting $P = 1$ and $R = 3$, assuming N to be positive, and solving the quadratic equation $(1 + 2w)N^2 - UN + U = 0$, we find that if

$$N < \frac{U + \sqrt{U^2 - 4U(1 + 2w)}}{2(1 + 2w)}$$

there is no stable value of i, and hence no stable states at that value of N. We have disregarded the root at

$$N = \frac{U - \sqrt{U^2 - 4U(1 + 2w)}}{2(1 + 2w)}$$

because this value is always less than 2 as long as $U > 12$. We now solve $ad \geq 0$, and find:

$$(i^2 - i) \geq \frac{\frac{(N^3 - N^2)(1 - w + Rw)}{U} - (N^2 - N)(1 - w + Pw)}{Rw - Pw}$$

Thus stable states must satisfy both these inequalities for $i^2 - i$. We now claim that this process can oscillate indefinitely for certain values of the parameters, but will always reach a stable state in the limit as the number of time steps goes to infinity. Oscillations can occur because adding a defector can sometimes allow you to exceed carrying capacity. To see an example of this, we consider the case U=12, N=10, and i=9. We calculate the adjustment to be 1, so there is a small probability that we will add a defector, in which case we have N=11, i=9, and carrying capacity is 10.8156, so the adjustment at the next step will be -1, meaning that we may kill a defector and end up where we started, hence illustrating the possibility of arbitrarily long oscillation. But of course, if we let the process continue infinitely long, every adjustment that has a non-zero probability of occurring will eventually occur.

First of all, it is obvious that we cannot stay above carrying capacity indefinitely, since adjustments for those states are always negative integers, and we cannot subtract individuals indefinitely. Therefore, it is sufficient to establish that when a state is below carrying capacity and the adjustment term is positive, there is always a possible choice of individuals that will cause the population to approach carrying capacity but never exceed it. Thus, if we keeping making such changes, the adjustment term will approach and eventually reach 0 (since adjustments occur by integers in discrete steps, the approach must be a finite process). Since a TFT player's fitness is always greater than or equal to the average fitness of the population when the game is infinite, adding a TFT player can never decrease the carrying capacity. We can then assume, in the worst case, that C remains constant. Thus, adding all TFT players in the adjustment phase can never cause a population to exceed carrying capacity at the next time step because $N(1 - N/C) \leq C - N$. To see this, we multiply both sides by C: $N(C - N) \leq C(C - N)$ holds, because $C \geq N$. Therefore, regardless of the starting conditions, the process will eventually reach a stable state. If we consider values of N for which stable values of i exist, we note that the lowest i satisfying

$$(i^2 - i) \geq \frac{\frac{(N^3 - N^2)(1 - w + Rw)}{U} - (N^2 - N)(1 - w + Pw)}{Rw - Pw}$$

is guaranteed to be stable. We can replace this inequality with equality, and solve for i:

$$i = \frac{1 \pm \sqrt{1 + 4\frac{\frac{(1+2w)(N^3-N^2)}{U} - (N^2-N)}{2w}}}{2}$$

Consider the condition $\frac{i}{N} \geq d$ for d a real number satisfying $0 \leq d \leq 1$. Then when $\frac{i}{N} = d$ the left hand side of the above equation is real, so the expression under the square root on the right must be positive. Thus, the lesser root is strictly less than one, and we can disregard it. Since the expression under the square root is strictly increasing with respect to N when the expression is positive, no imaginary solutions appear under this condition. Therefore, we can solve for N such that:

$$\frac{1 + \sqrt{1 + 4\frac{\frac{(1+2w)(N^3-N^2)}{U} - (N^2-N)}{2w}}}{2} = Nd$$

This yields two values of N, the lesser of which is always less than our lower bound of

$$N > \frac{U + \sqrt{U^2 - 4U(1+2w)}}{2(1+2w)}$$

and the greater of which is always above this value. Therefore, as long as N is greater than the larger root, all stable states for that N will have at least Nd TFT players. We can therefore describe the ratio of the number of N with this property over the number of N with any stable values of i as:

$$\frac{U - Floor[\frac{1}{2(1+2w)}(1 + U + 2w + 2d^2Uw + \sqrt{4U(1+2w)(-1-2dw) + (-1-U-2w-2d^2Uw)^2}]}{U - Floor[\frac{U+\sqrt{U^2-4U(1+2w)}}{2(1+2w)}]}$$

If we allow random starting conditions of N and i, we would expect this value to be a reasonable estimate of the proportion of trials that will end in stable states with at least Nd TFT players. This estimate needs to be adjusted, however, since it may be that some stable states are much more likely to be reached than others. By running randomized trials and comparing our results, we found that multiplying this ratio by 3/4 and adding certain small constants to it in specific ranges of w,d,and U, we were able to improve our estimate to be within a $\pm .07$ range of absolute error as long as $U \geq 100$, w is $\geq .02$ but still reasonably small, and $.3 \leq d \leq .8$. These adjustments are included in the following code, which returns our estimate:

$If(d \geq 0.8 \wedge d < 0.9, Return(\frac{3}{4}\, ratio(U, d, w) - 0.06)),$

$If(d \geq 0.7 \wedge d < 0.8, Return(\frac{3}{4}\, ratio(U, d, w) - 0.035)),$

$If(d \geq 0.6 \wedge d < 0.7, Return(\frac{3}{4}\, ratio(U, d, w) - 0.03)),$

$If(d \geq 0.5 \wedge d < 0.6, If(w < 0.03 \wedge w \geq 0.02 \wedge U \geq 100 \wedge U < 150,$

$Return(\frac{3}{4}\, ratio(U, d, w) + 0.07), Return(\frac{3}{4}\, ratio(U, d, w) - 0.01))),$

$If(d \geq 0.4 \wedge d < 0.5, If(w < 0.035 \wedge w \geq 0.02, Return(\frac{3}{4}\, ratio(U, d, w) + 0.06),$

$Return(\frac{3}{4}\, ratio(U, d, w) + 0.01))),$

$If(d \geq 0.3 \wedge d < 0.4, Return(\frac{3}{4}\, ratio(U, d, w) + 0.07)))$

We can also apply this model to a population of defectors and cooperators, each playing one round of prisoner's dilemma against each other. In this case the total fitness of a population of N total players with i cooperators is:

$$\frac{w}{N-1}(-i^2 + (3N-2)i) + N$$

(Here we have used $S = 0$, $P = 1$, $R = 3$, $T = 5$). We note that for $0 \leq i \leq N$

and $N \geq 2$, this is an increasing function of i and N. Hence, adding cooperators to a population will always raise the average fitness, and so if we define carrying capacity as dependent on fitness and an upper bound U as before, and the adjustment term as before, we see that once again, the system will always reach a stable state. By solving for adjustment < 1, we find:

$$-i^2 + (3N-2)i < \left(\frac{N^3(1+2w)}{(N-1)U} - N\right)\left(\frac{N-1}{w}\right)$$

So, there can be no stable values of i when the right-hand side of this expression is negative, which occurs, like in the TFT vs. ALLD case when:

$$N < \frac{U + \sqrt{U^2 - 4U(1+2w)}}{2(1+2w)}$$

(Again we have assumed $U > 12$ and $N \geq 2$ as an expedient means of discarding the negative root). We now solve adjustment ≥ 0 and obtain:

$$\frac{N^2(1+2w) - UN}{Uw}(N-1) < -i^2 + (3N-2)i$$

From which we conclude that for values of N with stable values of i, the lowest value of i satisfying $ad \geq 0$ is:

$$i = \frac{3N - 2 - \sqrt{(3N-2)^2 - \frac{4(N-1)}{Uw}(N^2(1+2w) - UN)}}{2}$$

So if we set this equal to d (again, a percentage between 0 and 1) and solve for N, we find two solutions, one of which is less than $\frac{U}{1+2w}$ and the other is greater than this value. (This is a critical value, since for N less than this, $i = 0$ will satisfy the condition that the adjustment is non-negative.) The greater root is:

$$\frac{1 + U + 2w + 3dUw - d^2Uw + \sqrt{-4U(1+2w)(1+2dw) + (1 + U + 2w + 3dUw - d^2Uw)^2}}{2+4w}$$

So we can define our ratio like before as:

$$\frac{U - Floor[root]}{U - Floor[\frac{U+\sqrt{U^2-4U(1+2w)}}{2(1+2w)}]}$$

where root refers to the greater root given above. By running random trials, we found that multiplying this ratio by 5/4 and similarly adjusting by small constants we can arrive at a reasonable estimate of the percentage of random trials resulting in at least Nd cooperators (again, a range of $\pm.07$ absolute error). To obtain our estimate, one would use the following code (where ratio2 refers to the ratio above):

$If(d \geq 0.8 \wedge d < 0.9, If(w \geq 0.04, Return(\frac{5}{4}\, ratio2(U, w, d) + 0.04),$

$\quad If(w \geq 0.02 \wedge w < 0.03, Return(\frac{5}{4}\, ratio2(U, w, d) + 0.02), Return(\frac{5}{4}\, ratio2(U, w, d)))))$

$If(d \geq 0.7 \wedge d < 0.8, If(w \geq 0.04, Return(\frac{5}{4}\, ratio2(U, w, d) + 0.04),$

$\quad If(w \geq 0.03 \wedge w < 0.04, Return(\frac{5}{4}\, ratio2(U, w, d) + 0.02)\ , Return(\frac{5}{4}\, ratio2(U, w, d)))))$

$If(d \geq 0.6 \wedge d < 0.7, If(w \geq 0.03, Return(\frac{5}{4}\, ratio2(U, w, d) + 0.04)))$

$If(d \geq 0.5 \wedge d < 0.6, Return(\frac{5}{4}\, ratio2(U, w, d) + 0.025))$

$If(d \geq 0.4 \wedge d < 0.5, Return(\frac{5}{4}\, ratio2(U, w, d) + 0.02))$

$If(d \geq 0.3 \wedge d < 0.4, Return(\frac{5}{4}\, ratio2(U, w, d) + 0.03)))$

If we consider games between TFT and ALLD with a parameter n that represents the probability of playing another round, then we can estimate the percentage of random trials that will end with at least Nd TFT players by taking a weighted average of our two estimates of the one-round and infinite game results, using n as the weight on the infinite game. We would expect this estimate to be within about $\pm.1$ absolute error, since the errors on each estimate seem to be independent, and therefore will not likely compound.

These estimates reveal an interesting tendency towards stable, mixed populations, even in the extreme cases of the one-round and infinite games, where one strategy is clearly superior to the other. For example, in the infinite game for d approximately 1/3, U between 100 and 500, and $.02 \leq w \leq .1$, the estimated percentage of random trials that will result in TFT reaching d percent of the population is significantly greater than a half, but less than 3/4. So at least 1/4 of the time the population will stabilize with less than a third TFT players, which is surprising since TFT is dominant. Similarly, if d equals .8, for the same range of U and w, we find the estimated ratio to be slightly less than 1/4, so TFT is often not able to reach extremely high percentages of the population, despite its dominance. In the case of the one round game, we find that cooperators can similarly persist as significant portions of stable populations, despite the fact that the fitness of a defector is strictly greater. Thus, the model has the curious feature that in the one round game, neither the strategy that is best for the individual (defect), nor the strategy that is best for the population (cooperate) is allowed to truly dominate. Even in the infinite game where TFT is the superior strategy both for the population and the individual, ALLD often survives by lowering the carrying capacity and therefore checking the growth of TFT.

The ability to reach stable states of coexistence between cooperative and non-cooperative strategies is a characteristic this model shares with some spatial models of populations playing Prisoner's Dilemma, which often allow cooperators to survive by clustering together. Our model illustrates that non-spatial mechanisms can also lead to stable, mixed populations, and it is not necessary to limit a cooperator's opponents to its spatial neighbors in order to see cooperators survive even in an atmosphere of dominance.

Bibliography:

Fudenberg,Drew, Martin A. Nowak, Akira Sasaki, and Christine Taylor,**Emergence of cooperation and evolutionary stability in finite populations**. In **NATURE**, Vol. 428, April 2004.