

populations may decrease cooperation and lead to more frequent escalations of conflicts in situations in which cooperation persists in well-mixed populations. Thus, spatial structure may not be as universally beneficial for cooperation as previously thought. □

Methods

Spatial structure

In our spatially structured populations, individuals are confined to sites on regular 100×100 lattices with periodic boundary conditions, and interact with their neighbours. We used square lattices with $N = 4$ and $N = 8$ neighbours, hexagonal lattices ($N = 6$) and triangular lattices ($N = 3$). Whenever a site x is updated, a neighbour y is drawn at random among all N neighbours; the chosen neighbour takes over site x with probability $w_y = f(P_y - P_x)$, where the function f translates payoff differences into reproductive success, reflecting natural selection based on relative fitness. The site x remains unchanged, with probability $1 - w_y$. Lattice updating can be either synchronous or asynchronous. For synchronous updates, first all individuals interact in their respective neighbourhood and then all sites are updated simultaneously through competition with a randomly chosen neighbour. For asynchronous updates, only a single, randomly selected focal site is updated at each simulation step: first the payoffs of the focal individual and a random neighbour are determined, after which these two individuals compete to re-populate the focal site. See Supplementary Information for the case where competition involves all neighbours, rather than just a randomly chosen one.

Pure strategies

With pure strategies, each individual is either a cooperator or a defector. Lattices are initialized randomly with equal proportions of the two strategies. $f(z) = z_+/\alpha$ determines the transition probabilities, where z_+ is equal to z if $z > 0$ and 0 otherwise, and where $\alpha = T - P$ in the snowdrift game and $\alpha = T - S$ in the Prisoner's Dilemma, ensuring that $f(P_y - P_x) \leq 1$. In well-mixed populations this implements the replicator dynamics¹⁷. Equilibrium frequencies of cooperators and defectors are obtained by averaging over 1,000 generations after a relaxation time of 5,000 generations.

The individual-based spatial models are complemented by deterministic pair-approximation (ref. 28 and see Supplementary Information). This approach correctly predicts a decrease in the frequency of cooperators in spatially structured populations, but it underestimates the effects of local correlations: for larger r the fragility of cooperative clusters is underrated, as is the ability of cooperators to displace defectors for small r (Fig. 1). Near the extinction thresholds, interesting symmetrical dynamics occur: tiny patches of defectors (cooperators) meander in a sea of cooperators (defectors). Occasionally they divide into pairs or collide and vanish. This resembles a branching and annihilating random walk, which suggests that there are critical phase transitions and points to interesting relationships between game theory and condensed matter physics²⁹.

Mixed strategies

For mixed strategies in the hawk–dove game, an individual is characterized by the probability p to show dove-like behaviour. Exploration of this continuous strategy space requires mutations. Whenever an individual with strategy p reproduces, a mutation occurs with a small probability (0.01) that assigns the offspring the strategy $p + \xi$, where ξ denotes a gaussian-distributed random variable with a mean of 0 and an s.d. of 0.002. To speed up simulations, the lattice is initialized with random strategies drawn from a normal distribution with a mean corresponding to the equilibrium strategy in well-mixed populations and an s.d. of 0.02. The simulation results are insensitive to the initialization details.

An individual in x with strategy p interacting with a neighbour with strategy q gets an average payoff $P_x = pqR + p(1 - q)S + (1 - p)qT + (1 - p)(1 - q)P$. The small difference in the strategies of parents and mutant offspring leads to small payoff differences $P_y - P_x$ between neighbouring individuals. Thus, the update rule for pure strategies returns small probabilities for a strategy change, which slows down the simulations. We therefore used the nonlinear function $f(z) = [1 + \exp(-z/\kappa)]^{-1}$, in which κ is a noise term that reflects uncertainties in assessing the payoffs. This nonlinearity greatly speeds up the simulations and introduces an interesting and realistic form of error, whereby a worse performing player occasionally manages to reproduce. For $\kappa \rightarrow \infty$, errors in assessing the payoffs increase until no information is left and the players randomly adopt neighbouring strategies. We used $\kappa = 0.1$ in our simulations. The equilibrium levels of dove-like behaviour were determined by evolving the lattice over 10,000 generations and then averaging the mixed strategies over another 1,000 generations.

Received 11 September 2003; accepted 21 January 2004; doi:10.1038/nature02360.

1. Maynard Smith, J. & Szathmáry, E. *The Major Transitions in Evolution* (W. H. Freeman, Oxford, UK, 1995).
2. von Neumann, J. & Morgenstern, O. *Theory of Games and Economic Behaviour* (Princeton Univ. Press, Princeton, 1944).
3. Maynard Smith, J. & Price, G. The logic of animal conflict. *Nature* **246**, 15–18 (1973).
4. Axelrod, R. & Hamilton, W. D. The evolution of cooperation. *Science* **211**, 1390–1396 (1981).
5. Sugden, R. *The Economics of Rights, Co-operation and Welfare* (Blackwell, Oxford, UK, 1986).
6. Nowak, M. A. & May, R. M. Evolutionary games and spatial chaos. *Nature* **359**, 826–829 (1992).
7. Doebeli, M. & Knowlton, N. The evolution of interspecific mutualisms. *Proc. Natl Acad. Sci. USA* **95**, 8676–8680 (1998).
8. Killingback, T., Doebeli, M. & Knowlton, N. Variable investment, the continuous prisoner's dilemma, and the origin of cooperation. *Proc. R. Soc. Lond. B* **266**, 1723–1728 (1999).
9. Wilkinson, G. S. Reciprocal food-sharing in the vampire bat. *Nature* **308**, 181–184 (1984).

10. Clutton-Brock, T. H. *et al.* Selfish sentinels in cooperative mammals. *Science* **284**, 1640–1644 (1999).
11. Seyfarth, R. M. & Cheney, D. L. Grooming alliances and reciprocal altruism in vervet monkeys. *Nature* **308**, 541–543 (1984).
12. Milinski, M. Tit for tat in sticklebacks and the evolution of cooperation. *Nature* **325**, 433–435 (1987).
13. Milinski, M., Lüthi, J. H., Egger, R. & Parker, G. A. Cooperation under predation risk: experiments on costs and benefits. *Proc. R. Soc. Lond. B* **264**, 831–837 (1997).
14. Turner, P. E. & Chao, L. Prisoner's dilemma in an RNA virus. *Nature* **398**, 441–443 (1999).
15. Heinsohn, R. & Parker, C. Complex cooperative strategies in group-territorial African lions. *Science* **269**, 1260–1262 (1995).
16. Clutton-Brock, T. Breeding together: kin selection and mutualism in cooperative vertebrates. *Science* **296**, 69–72 (2002).
17. Hofbauer, J. & Sigmund, K. *Evolutionary Games and Population Dynamics* (Cambridge Univ. Press, Cambridge, UK, 1998).
18. Killingback, T. & Doebeli, M. Spatial evolutionary game theory: hawks and doves revisited. *Proc. R. Soc. Lond. B* **263**, 1135–1144 (1996).
19. Maynard Smith, J. *Evolution and the Theory of Games* (Cambridge Univ. Press, Cambridge, UK, 1982).
20. Hamilton, W. D. The evolution of altruistic behaviour. *Am. Nat.* **97**, 354–356 (1963).
21. Nowak, M. A. & Sigmund, K. Evolution of indirect reciprocity by image scoring. *Nature* **393**, 573–577 (1998).
22. Mitteldorf, J. & Wilson, D. S. Population viscosity and the evolution of altruism. *J. Theor. Biol.* **204**, 481–496 (2000).
23. Wilson, D. S., Pollock, G. B. & Dugatkin, L. A. Can altruism evolve in purely viscous populations? *Evol. Ecol.* **6**, 31–34 (1992).
24. Taylor, P. D. Altruism in viscous populations—an inclusive fitness approach. *Evol. Ecol.* **6**, 352–356 (1992).
25. Hauert, C. Fundamental clusters in spatial 2×2 games. *Proc. R. Soc. Lond. B* **268**, 761–769 (2001).
26. Turner, P. E. & Chao, L. Escape from prisoner's dilemma in RNA phage $\phi 6$. *Am. Nat.* **161**, 497–505 (2003).
27. Wilkinson, G. S. & Shank, C. C. Rutting-fight mortality among musk oxen on banks island, Northwest Territories, Canada. *Anim. Behav.* **24**, 756–758 (1977).
28. Matsuda, H., Ogita, N., Sasaki, A. & Sato, K. Statistical mechanics of populations. *Prog. Theor. Phys.* **88**, 1035–1049 (1992).
29. Szabó, G. & Hauert, C. Phase transitions and volunteering in spatial public goods games. *Phys. Rev. Lett.* **89**, 118101 (2002).

Supplementary Information accompanies the paper on www.nature.com/nature.

Acknowledgements We thank M. Ackermann and G. Szabó for comments. C.H. acknowledges support from the Swiss National Science Foundation. M.D. was supported by the National Science and Engineering Research Council (NSERC), Canada, and by the James S. McDonnell Foundation, USA.

Competing interests statement The authors declare that they have no competing financial interests.

Correspondence and requests for materials should be addressed to C.H. (hauert@zoology.ubc.ca).

Emergence of cooperation and evolutionary stability in finite populations

Martin A. Nowak^{1,2}, Akira Sasaki⁴, Christine Taylor^{1,5} & Drew Fudenberg³

¹Program for Evolutionary Dynamics, ²Department of Organismic and Evolutionary Biology and Department of Mathematics, ³Department of Economics, Harvard University, Cambridge, Massachusetts 02138, USA

⁴Department of Biology, Kyushu University, Fukuoka 812-8581, Japan

⁵Department of Mathematics, MIT, Cambridge, Massachusetts, 02139, USA

To explain the evolution of cooperation by natural selection has been a major goal of biologists since Darwin. Cooperators help others at a cost to themselves, while defectors receive the benefits of altruism without providing any help in return. The standard game dynamical formulation is the 'Prisoner's Dilemma'^{21–11}, in which two players have a choice between cooperation and defection. In the repeated game, cooperators using direct reciprocity cannot be exploited by defectors, but it is unclear how such cooperators can arise in the first place^{12–15}. In general, defectors are stable against invasion by cooperators. This understanding is based on traditional concepts of evolutionary stability and

dynamics in infinite populations^{16–20}. Here we study evolutionary game dynamics in finite populations^{21–25}. We show that a single cooperator using a strategy like ‘tit-for-tat’ can invade a population of defectors with a probability that corresponds to a net selective advantage. We specify the conditions required for natural selection to favour the emergence of cooperation and define evolutionary stability in finite populations.

In the Prisoner’s Dilemma, two players are offered a certain

Box 1
Game dynamics in finite populations

The fitness of strategies A and B with payoff matrix (1) is, respectively, given by:

$$\begin{aligned} f_i &= 1 - w + w[a(i - 1) + b(N - i)]/[N - 1] \\ g_i &= 1 - w + w[ci + d(N - i - 1)]/[N - 1] \end{aligned} \quad (4)$$

Here i denotes the number of individuals using strategy A, and $w \in [0, 1]$ specifies the contribution of the game to fitness. Selection dynamics can be formulated as a Moran process with frequency-dependent fitness. At each time step, an individual is chosen for reproduction proportional to its fitness. One identical offspring is being produced that replaces another randomly chosen individual. Thus N is strictly constant. The probability of adding an A-offspring is $f_i/[f_i + (N - i)g_i]$. At each time step, the number of A individuals can either increase by one, stay the same, or fall by one. Therefore, the transition matrix of the Markov process is tri-diagonal and defines a birth–death process given by:

$$\begin{aligned} P_{i,j+1} &= \frac{f_i}{f_i + (N - i)g_i} \frac{N - i}{N} \\ P_{i,j-1} &= \frac{(N - i)g_i}{f_i + (N - i)g_i} \frac{i}{N} \end{aligned} \quad (5)$$

We have $P_{i,i} = 1 - P_{i,j+1} - P_{i,j-1}$. All other entries of the transition matrix are 0.

The process has two absorbing states, $i = 0$ and $i = N$: if the population has reached either one of those states, then it will stay there forever. We denote by x_i the probability of ending up in state $i = N$ when starting in state i . We have:

$$x_i = P_{i,j+1}x_{i+1} + P_{i,j}x_i + P_{i,j-1}x_{i-1} \quad (6)$$

with boundary conditions $x_0 = 0$ and $x_N = 1$.

Let us calculate the probability, $\rho_A (= x_1)$, that a single individual A can invade and take over a population of B players. Solving equation (6), we obtain²⁷:

$$\rho_A = 1 / \left(1 + \sum_{k=1}^{N-1} \prod_{i=1}^k \frac{g_i}{f_i} \right) \quad (7)$$

If $\rho_A > 1/N$, then selection favours A replacing B.

The rate of evolution from all-B to all-A is given by $r = N\rho_A u$, where u is the mutation rate. We can rescale the rate of evolution in units of u . Thus, we set $u = 1$. The rate of evolution, r , can be an increasing or decreasing function of w . There can also be a maximum or minimum value of r for some intermediate $w \in (0, 1)$. It is possible that $r > 1$ for small w , but $r < 1$ for large w , or vice versa. For $w = 0$, we have $r = 1$.

In the limit of weak selection, $w \ll 1$, we find that:

$$N\rho_A \approx 1/[1 - (\alpha N - \beta)(w/6)] \quad (8)$$

with $\alpha = a + 2b - c - 2d$ and $\beta = 2a + b + c - 4d$. From this equation, we see that $N\rho_A > 1$ if $\alpha N > \beta$, which leads to equation (2). If $\alpha > 0$ then there is a minimum N for which $N\rho_A > 1$. It is given by $N_{\min} = \beta/\alpha$.

We define a strategy to be an ESS_N if selection opposes the invasion and fixation of any other strategy for any $w > 0$. Thus a necessary condition for ESS_N is that $N\rho_A$, as given by equation (8), is less than one. If instead we want to know whether a strategy is evolutionarily stable for a given N and a given w , then we cannot necessarily use the simplified equation (8), which only holds in the limit of small w . Instead, we have to check that the exact expression $N\rho_A$, as given by equation (7), is less than one.

payoff, R , for mutual cooperation and a lower payoff, P for mutual defection. If one player cooperates while the other defects, then the cooperator gets the lowest payoff, S , while the defector gains the highest payoff, T . Thus, we have $T > R > P > S$. In the non-repeated Prisoner’s Dilemma, defectors dominate cooperators, which means that in any mixed population, defectors have a higher fitness. In the repeated Prisoner’s Dilemma, the same two players meet more than once, and there are many conceivable strategies that allow cooperative behaviour which cannot be defeated by defectors. The most famous such strategy is tit-for-tat (TFT), in which the player cooperates in the first round and then does whatever the opponent did in the previous round. Another strategy is always to defect (AllD). If the number of rounds is sufficiently large, then AllD and TFT resist invasion attempts by the other strategy. Thus, TFT can maintain cooperation, but how it can become established is unclear.

In the standard evolutionary model of the finitely repeated Prisoner’s Dilemma, TFT cannot invade AllD, because if everyone in an infinitely large population uses AllD, then a small fraction of TFT players have a lower payoff. Every invasion attempt by TFT is therefore eliminated by natural selection. Past work has proposed several modifications to this model that allow TFT to invade successfully: (1) a mass of TFT players arises simultaneously to overcome an invasion barrier¹²; (2) TFT players form spatial clusters^{13,14}; or (3) aggregate payoffs are stochastic¹⁵. Here we show that none of these modifications are necessary to explain the emergence of cooperation in finite populations.

Consider a game between two strategies, A and B, with payoff matrix:

$$\begin{array}{cc} & \begin{array}{c} A \quad B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \begin{pmatrix} a & b \\ c & d \end{pmatrix} \end{array} \quad (1)$$

If A and B denote, respectively, TFT and AllD, then we have $a > c > d > b$ provided the number of rounds is finite and greater than $(T - P)/(R - P)$. In this case, both TFT and AllD are strict Nash equilibria and evolutionarily stable strategies (ESS).

Box 2
Invading AllD

Let strategies A and B denote, respectively, TFT and AllD in a Prisoner’s Dilemma which is repeated for n rounds on average. The payoff matrix is $a = Rn$, $b = S + P(n - 1)$, $c = T + P(n - 1)$ and $d = Pn$. If the average number of rounds, n , exceeds $(T - P)/(R - P)$ then we have $a > c > d > b$. In this case, there is an unstable equilibrium between TFT and AllD, and neither strategy can invade the other in the context of deterministic dynamics of infinitely large populations.

Condition (2) implies $n(R - P)(N - 2) > T(N + 1) - S(2N - 1) + P(N - 2)$. This inequality determines the minimum number of rounds required for (weak) selection to favour TFT replacing AllD for a given N . We need at least $N = 3$. For large N , the number of rounds must fulfill $n > (T + P - 2S)/(R - P)$. Let $R = 3$, $T = 5$, $P = 1$, $S = 0$. For $N = 3$ we have $n > 10.5$. For $N = 4$ we have $n > 6.75$. For large N we have $n > 3$.

Among all strategies of the repeated Prisoner’s Dilemma, TFT maximizes the probability of invading AllD in any finite population. More precisely, TFT belongs to a set of strategies that maximize this probability. All strategies of this set have the following property: (1) when playing themselves they cooperate all the time; (2) when playing AllD they cooperate on the first move and then never again. ‘Win-stay, lose-shift’ and ‘generous tit-for-tat’ have lower invasion probabilities because they attempt to cooperate with AllD too often. They work best once cooperation has been established^{12,30}.

Deterministic replicator dynamics of infinite populations admit an unstable equilibrium at a frequency of TFT given by $x^* = (d - b)/(a - b - c + d)$. If the initial frequency of TFT is less than this value, then it will be eliminated by natural selection. TFT can only replace AllD if its initial frequency exceeds this invasion barrier. The same evolutionary dynamics hold for AllD competing with other cooperative strategies such as ‘generous tit-for-tat’, ‘contrite tit-for-tat’ or ‘win-stay, lose-shift’.

We now study evolutionary game dynamics in finite populations. In this case, there is always a non-zero probability that a mutant strategy can invade and take over the population even though it is opposed by selection. TFT invading AllD has a lower fitness if its frequency is less than x^* but has a higher fitness if its frequency is greater than this threshold. What is the probability that a single TFT player will take over a population of AllD players?

Consider a stochastic process describing a finite population of size N . At each time step, one individual is chosen for reproduction proportional to its fitness, and its offspring replaces a randomly chosen individual²⁶. The population size is constant. The fitness of each player depends on the number of TFT or AllD players. In addition, we introduce a parameter, w , which determines the contribution of the game’s payoff to fitness. This parameter, quantifying the intensity of selection, cancels itself out in deterministic replicator dynamics of infinite populations, but plays a crucial role in the stochastic process that we study here.

We calculate the probability, ρ_A , that a single individual using strategy A will invade and take over a population of B players²⁷ (Box 1). For a neutral mutant²⁸ this fixation probability is $\rho_A = 1/N$. If $\rho_A > 1/N$ then selection favours A replacing B. In Fig. 1, we show that in the case of TFT and AllD, $N\rho_A$ is a one-humped function of N . For a wide choice of parameter values a, b, c, d and w there is an intermediate range of population sizes which ensure $N\rho_A > 1$. Thus, the invasion and replacement of AllD by TFT, starting from a single individual of TFT, can be favoured by natural selection. Interestingly, there are critical minimum and maximum population sizes that allow positive selection of TFT. In very small populations, there is a strong effect of spite: helping another individual leads to a significant disadvantage. For example, in a population of only two players, TFT always has a lower fitness than AllD. In a very large population it is extremely unlikely to reach the invasion barrier when starting with a single TFT player. Thus, neither small nor large but intermediate population sizes are optimum for initiating cooperation.

Can we derive the underlying principle that determines whether a particular payoff matrix (1) allows selection for TFT replacing

AllD? The exact expression for ρ_A is complicated. The condition $\rho_A > 1/N$ requires the solution of N th-order polynomials, and a diffusion approximation yields transcendental equations. Nevertheless, the following surprisingly simple theorem holds. For a given N and sufficiently weak selection (small w), selection favours TFT replacing AllD if:

$$a(N - 2) + b(2N - 1) > c(N + 1) + d(2N - 4) \quad (2)$$

For the smallest possible population size, $N = 2$ (it takes two to play), inequality (2) yields $b > c$ (which is not possible for the game between TFT and AllD). For the limit of large N , we obtain $a + 2b > c + 2d$. The latter condition is equivalent to $x^* < 1/3$. Therefore, if the invasion barrier of TFT is less than $1/3$, there can be positive selection for TFT to replace AllD in a finite population (Box 2).

In general, for any two strategies which are the best replies to themselves, we find that selection can favour A replacing B for some N and w , if $b > c$ or $x^* < 1/3$ (Fig. 2).

Our results have immediate consequences for the concept of evolutionary stability. The well-known definition of an ESS is motivated by selection dynamics in infinite populations^{16–20}. Strategy B is an ESS if either (1) $d > b$ or (2) $d = b$ and $a < c$. These conditions imply that selection opposes the spread of infinitesimally small fractions of A in infinitely large populations of B.

For finite N , we propose that B is an ESS, denoted ESS_N , if two conditions hold: (1) selection opposes A invading B, which means that a single mutant A in a population of B has a lower fitness; and (2) selection opposes A replacing B, which means $\rho_A < 1/N$, for any $w > 0$. The first condition is equivalent to $b(N - 1) < c + d(N - 2)$. The second condition, for small w , is equivalent to $a(N - 2) + b(2N - 1) < c(N + 1) + d(2N - 4)$. For $N = 2$, both conditions reduce to $b < c$. For large populations, the two conditions lead to $b < d$ and $x^* > 1/3$, respectively. Hence, for small populations, the traditional ESS concept is neither necessary nor sufficient; for large populations, it is necessary but not sufficient (Fig. 3). If we consider a game with many different strategies, then the two conditions must hold in pairwise comparison with every other strategy.

The motivation of the ESS_N concept is as follows. If a strategy is an ESS_N , then a single mutant of any other strategy must have a lower fitness. Therefore, selection opposes the initial spread of any other strategy. As we have seen in the case of AllD and TFT, however, in a finite population it is possible that the fixation of a strategy is favoured by selection even though its initial increase is opposed by selection. Thus, the second condition states that a strategy can only be an ESS_N if the fixation probability of all other strategies is less

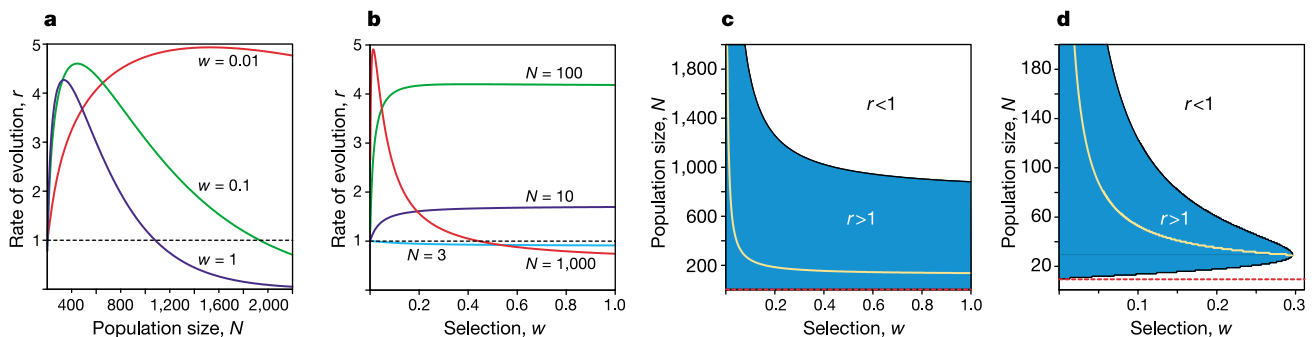


Figure 1 Selection can favour the replacement of AllD by TFT in finite populations. **a**, The rate of evolution, $r = N\rho_A$, is a one-humped function of N . There is an intermediate range of N which leads to positive selection of TFT, $N\rho_A > 1$. **b**, $N\rho_A$ is shown as function of w . For small N , we have $N\rho_A < 1$ for all w . For larger N we have $N\rho_A > 1$ for all w . For even larger N we have $N\rho_A > 1$ as long as w is below a certain threshold. **c, d**, The blue-

shaded region indicates the parameter region where $N\rho_A > 1$. The yellow lines show the optimum value N for a given w maximizing $N\rho_A$. The broken red line indicates $N_{\min} = (2a + b + c - 4d)/(a + 2b - c - 2d)$, which is the predicted minimum population size required for positive selection of TFT in the limit of weak selection. Parameter choices are: $R = 3, T = 5, P = 1, S = 0$; and $n = 10$ rounds for **a–c** and $n = 4$ rounds for **d**.




Figure 2 The 1/3-law of frequency-dependent evolution. **a**, Suppose A and B are the best replies to themselves, meaning $a > c$ and $d > b$ in payoff matrix (1). In this case all-A and all-B are stable equilibria of the replicator dynamics for infinite population size. The unstable equilibrium is located at a frequency of A given by $x^* = (d - b)/(a - b - c + d)$. If $x^* < 1/3$ then selection favours A replacing B for a sufficiently large population and weak selection. The minimum population size is given by $N_{\min} = (2a + b + c - 4d)/(a + 2b - c - 2d)$. **b**, The relationship between the 1/3-law and risk dominance. If $x^* < 1/3$ then $\rho_A > 1/N > \rho_B$. If $1/3 < x^* < 1/2$ then $1/N > \rho_A > \rho_B$. If

$1/2 < x^* < 2/3$ then $1/N > \rho_B > \rho_A$. If $2/3 < x^*$ then $\rho_B > 1/N > \rho_A$. We note that if A is risk dominant, meaning $x^* < 1/2$, then $\rho_A > \rho_B$. These results hold in the limit of weak selection and large population size. The location of x^* as shown implies that $1/N > \rho_A > \rho_B$. **c**, AIID and TFT in the repeated Prisoner's Dilemma. Although AIID is evolutionarily stable against invasion by TFT for deterministic dynamics of infinite populations, in a finite population the probability that a single mutant of TFT takes over an AIID population can exceed $1/N$. Selection can favour TFT replacing AIID, if the unstable equilibrium occurs at a frequency of less than $1/3$ TFT.

than the neutral threshold, $1/N$. In summary, we simply ask that a homogeneous ESS_N population is protected by selection against invasion and replacement. These requirements represent a natural extension of the original ESS concept formulated by Maynard Smith for infinitely large populations and deterministic evolutionary dynamics¹⁸.

Schaffer²² has proposed that a strategy is evolutionarily stable in a finite population if a single mutant of any other strategy has lower fitness. This is the first of our two conditions. Schaffer²² also proposes a global stability condition, namely that the ESS strategy must have a greater fitness than the other strategy for any composition of the population. This very stringent condition is a finite- N extension of Hamilton's unbeatable strategy²⁹, which dominates every other strategy. We note that every unbeatable strategy is an ESS, but the reverse is not true. Unfortunately, unbeatable strategies are rare. Many biological games admit ESS, but not unbeatable strategies.

Sometimes it is of interest to ask whether strategy A is more likely to replace B than vice versa. Let ρ_A and ρ_B denote the respective fixation probabilities. In the case where both A and B are the best replies to themselves and in the limit of weak selection and large populations, we find that $\rho_A > \rho_B$ if A is risk-dominant, meaning that $a + b > c + d$. For general N and w , however, risk dominance does not determine the ordering of the fixation probabilities.

Here we have studied a frequency-dependent Moran process, which is a natural finite- N analogue to the replicator equation. One can envisage many different stochastic processes that describe game dynamics in finite populations. An interesting possibility is the following. Pick two players at random. One is chosen for reproduction, the other for elimination. Hence, only mixed pairs can change the population. Suppose that player A is chosen for reproduction with probability $f_i/(f_i + g_i)$ and player B with probability $g_i/(f_i + g_i)$. In this case, we obtain the same process as we have analysed here, up to rescaling time.

If, instead, the fitter player is always chosen for reproduction, then the resulting process is stochastic in speed, but deterministic in direction: it will always follow the gradient of selection. But if player A is chosen for reproduction with probability $1/(1 + \exp[-(f_i - g_i)/\tau])$, then parameter w cancels out. There is, however, a new parameter, τ , which has a similar role and is equivalent to temperature in statistical physics. If $\tau \rightarrow 0$ then the fitter player is always chosen; selection is strong. If $\tau \rightarrow \infty$ then selection is weak, and the process is dominated by random drift. In the limit of large τ , we obtain exactly the same results as are presented here.

Another possibility is studying a frequency-dependent Wright-Fischer process with discrete generations. Furthermore, in all of those models, stochasticity could arise in evaluating the payoffs of individual players. We expect that all these processes (as long as they are not deterministic in following selection) will have a similar behaviour to what we have described here.

To sum up, (1) in finite populations, natural selection can favour the invasion and replacement of the AIID strategy by a cooperative strategy, when starting from a single individual using that strategy. No specific mechanism of invasion is required. (2) For any two strategies A and B, natural selection can favour A replacing B in a finite population provided $b > c$ or $a - c > 2(d - b)$. If A and B are the best replies to themselves then the latter condition implies that the frequency of A at the unstable equilibrium, x^* , must be less than $1/3$. (3) Our analysis leads to natural conditions for evolutionary stability in finite populations. These conditions specify whether a given resident strategy is protected by selection against invasion and replacement of any mutant strategy. □

a		A	B	
A	20	0		A is an ESS_N for $N > 12$
B	17	1		B is an ESS_N for $N < 53$
				$\rho_A > \rho_B$ for $N > 19$

b		A	B	
A	1	28		A is an ESS_N for $N < 22$
B	2	30		B is an ESS_N for $N > 17$
				$\rho_A > \rho_B$ for $N < 20$

Figure 3 A strategy is ESS_N if it is protected by selection against invasion and replacement by another strategy for given N and any $w > 0$. **a**, Both A and B are classical ESS, but for $2 \leq N \leq 12$ only B is ESS_N ; for $12 < N < 53$ both A and B are ESS_N , for $N \geq 53$ only A is ESS_N . We note that strategy B is a classical ESS, but is not ESS_N for large N ; in large populations there is selection for A replacing B. **b**, B dominates A. Therefore only B is a classical ESS. For $2 \leq N \leq 17$, however, we obtain that only A is ESS_N . For $17 < N < 22$ both A and B are ESS_N . For $N \geq 22$ only B is ESS_N . Examples **a** and **b** illustrate that for small populations the traditional ESS concept is neither necessary nor sufficient to imply ESS_N , and for large populations it is necessary but not sufficient.

Methods

Remarks on ESS

If $d > b$ then B is both a strict Nash equilibrium and an ESS in comparison with A. A strict Nash equilibrium implies protection by selection against replacement in the following sense: for a given payoff matrix (a, b, c, d) with $d > b$ and for any given intensity of selection, $0 < w \leq 1$, we have $\rho_A \rightarrow 0$ as $N \rightarrow \infty$.

For every finite population size, N , however, we can calculate the maximum net selective advantage for a mutant replacing a strict Nash equilibrium. Given b, d with $d > b$, what is the maximum probability ρ_A of A replacing B? We are free to choose a and c .

To maximize ρ_A , we set $a \rightarrow \infty$ and $c = 0$. For large populations, we obtain $\rho_A = [1 - w(1 - b)]/[2 - w(2 - b - d)]$. For $w \rightarrow 0$ we have $\rho_A = 1/2$. For $w = 1$ we have $\rho_A = b/(b + d)$. This fixation probability of A corresponds to a constant relative fitness of $1 + (b/d)$ or a net selective advantage of b/d . Hence there can be enormous selection pressure for replacement of a strict Nash equilibrium in arbitrarily large, finite populations (when the other equilibrium is much more efficient).

The calculation here uses the fact that from state $i = 1$ the system can go either to $i = 0$ or $i = 2$. Because $a \rightarrow \infty$ and $c = 0$, fixation of strategy A is certain from state $i = 2$. Hence, the fixation probability from $i = 1$ to $i = N$ is just the probability $P_{12}/(P_{12} + P_{10}) = (1 - w + wb)/(1 - w + wb + 1 - w + wd(N - 2)/(N - 1))$. This holds for all w . For large N , we obtain the above formula for ρ_A .

Risk dominance

Let ρ_A denote the probability that a single A player reaches fixation in a population of B. Let ρ_B denote the probability that a single B player reaches fixation in a population of A. We obtain:

$$\frac{\rho_A}{\rho_B} = \prod_{i=1}^{N-1} \frac{f_i}{g_i} \quad (3)$$

For weak selection (small w) we find $\rho_A/\rho_B = 1 + w[(N/2)(a + b - c - d) + d - a]$. It follows that $\rho_A > \rho_B$ is equivalent to $(N - 2)(a - d) > N(c - b)$. For large N this means $a - c > d - b$. Hence, if both A and B strategies are strict Nash equilibria then the risk-dominant equilibrium has a higher fixation probability when starting from a single player using that strategy. For general N and w , risk-dominance does not decide whether ρ_A is greater than ρ_B .

More general strategies

We have mostly studied the dynamics between ALLD and TFT. The repeated Prisoner's Dilemma, like other repeated games, admits a huge set of possible strategies, which makes it difficult to explicitly analyse the dynamics of evolution. In general, a strategy for playing the repeated Prisoner's Dilemma is a mapping from any history of the game between two players into the interval $[0, 1]$, denoting the probability of cooperation on the next move. However, we note that for the finitely repeated game, ALLD is a strict Nash equilibrium in comparison with all cooperative strategies, where we define a 'cooperative strategy' as a strategy which cooperates on the first move. Let us divide cooperative strategies into two subsets: (1) those that are dominated by ALLD and (2) those that are bistable with ALLD. In an infinitely large population, no cooperative strategy can ever invade ALLD. In a finite population of size N , strategies of the second subset can invade and replace ALLD provided inequality (2) holds and selection is sufficiently weak.

In an infinitely repeated Prisoner's Dilemma with time-average payoffs, it turns out that TFT dominates ALLD. In this case it can be shown that the 'win-stay, lose-shift'³⁰ strategy (also known as 'Pavlov' or 'perfect tit-for-tat') is the only simple strategy which cannot be invaded by any other strategy, and that it is the only strategy that is evolutionarily stable in an infinite population when actions are taken with a vanishingly small probability of error³¹. Moreover, this strategy is also the unique ESS in a model where strategies are encoded by finite-state automata, and the complexity of the automaton represents an evolutionary cost³¹.

Received 10 December 2003; accepted 11 February 2004; doi:10.1038/nature02414.

- Trivers, R. The evolution of reciprocal altruism. *Q. Rev. Biol.* **46**, 35–57 (1971).
- Axelrod, R. & Hamilton, W. D. The evolution of cooperation. *Science* **211**, 1390–1396 (1981).
- Axelrod, R. *The Evolution of Cooperation* (Basic Books, New York, 1984).
- Milinski, M. Tit for tat in sticklebacks and the evolution of cooperation. *Nature* **325**, 433–435 (1987).
- May, R. M. More evolution of cooperation. *Nature* **327**, 15–17 (1987).
- Dugatkin, L. A. *Cooperation Among Animals* (Oxford Univ. Press, Oxford, UK, 1997).
- Pfeiffer, T., Schuster, S. & Bonhoeffer, S. Cooperation and competition in the evolution of ATP-producing pathways. *Science* **292**, 504–507 (2001).
- Fehr, E. & Fischbacher, U. The nature of human altruism. *Nature* **425**, 785–791 (2003).
- Hammerstein, P. (ed.) *Genetic and Cultural Evolution of Cooperation* (MIT Press, Cambridge, Massachusetts, 2003).
- Boyd, R., Gintis, H., Bowels, S. & Richerson, P. J. The evolution of altruistic punishment. *Proc. Natl Acad. Sci. USA* **100**, 3531–3535 (2003).
- Fudenberg, D. & Maskin, E. Evolution and cooperation in noisy repeated games. *Am. Econ. Rev.* **80**, 274–279 (1990).
- Nowak, M. A. & Sigmund, K. Tit for tat in heterogeneous populations. *Nature* **355**, 250–253 (1992).
- Nowak, M. A. & May, R. M. Evolutionary games and spatial chaos. *Nature* **359**, 826–829 (1992).
- Killingback, T. & Doebeli, M. Self-organized criticality in spatial evolutionary game theory. *J. Theor. Biol.* **191**, 335–340 (1998).
- Fudenberg, D. & Harris, C. Evolutionary dynamics with aggregate shocks. *J. Econ. Theor.* **57**, 420–441 (1992).
- Maynard Smith, J. & Price, G. R. Logic of animal conflict. *Nature* **246**, 15–18 (1973).
- Taylor, P. D. & Jonker, L. B. Evolutionary stable strategies and game dynamics. *Math. Biosci.* **40**, 145–156 (1978).
- Maynard Smith, J. *Evolution and the Theory of Games* (Cambridge Univ. Press, Cambridge, UK, 1982).
- Hofbauer, J. & Sigmund, K. *Evolutionary Games and Population Dynamics* (Cambridge Univ. Press, Cambridge, UK, 1998).
- Hofbauer, J. & Sigmund, K. Evolutionary game dynamics. *Bull. Am. Math. Soc.* **40**, 479–519 (2003).
- Riley, J. G. Evolutionary equilibrium strategies. *J. Theor. Biol.* **76**, 109–123 (1979).
- Schaffer, M. Evolutionary stable strategies for a finite population and a variable contest size. *J. Theor. Biol.* **132**, 469–478 (1988).
- Fogel, G., Andrews, P. & Fogel, D. On the instability of evolutionary stable strategies in small populations. *Ecol. Model.* **109**, 283–294 (1998).
- Ficci, S. & Pollack, J. *Effects of Finite Populations on Evolutionary Stable Strategies*. *Proc. 2000 Genetic and*

- Evolutionary Computation Conf.* (ed. Whitley, D.) 927–934 (Morgan-Kaufmann, San Francisco, 2000).
- Schreiber, S. Urn models, replicator processes, and random genetic drift. *Siam. J. Appl. Math.* **61**, 2148–2167 (2001).
 - Moran, P. A. P. *The Statistical Processes of Evolutionary Theory* (Clarendon, Oxford, UK, 1962).
 - Karlin, S. & Taylor, H. M. *A First Course in Stochastic Processes* 2nd edn (Academic, London, 1975).
 - Kimura, M. Evolutionary rate at the molecular level. *Nature* **217**, 624–626 (1968).
 - Hamilton, W. D. Extraordinary sex ratios. *Science* **156**, 477–488 (1967).
 - Nowak, M. A. & Sigmund, K. A strategy of win-stay, lose-shift that outperforms tit for tat in Prisoner's Dilemma. *Nature* **364**, 56–58 (1993).
 - Binmore, K. & Samuelson, L. Evolutionary stability in repeated games played by the finite automata. *J. Econ. Theor.* **57**, 278–305 (1992).

Acknowledgements The Program for Evolutionary Dynamics is supported by J. Epstein.

Competing interests statement The authors declare that they have no competing financial interests.

Correspondence and requests for materials should be addressed to M.A.N. (martin_nowak@harvard.edu).

Pre-social benefits of extended parental care

Jeremy Field & Selina Brace

Department of Biology, University College London, Wolfson House, 4 Stephenson Way, London NW1 2HE, UK

The evolution of helping, in which some individuals forfeit their own reproduction and help others to reproduce, is a central problem in evolutionary biology. Recently proposed insurance-based mechanisms rely on a pre-existing life history with a long period of offspring dependency relative to the short life expectancies of adult carers^{1–4}: a lone mother's offspring are doomed if she dies young, whereas after a helper dies, other group members can finish rearing the offspring^{5,6}. A critical question, however, is how this life history could evolve in ancestral non-social populations, as offspring survival would then depend on a single, short-lived carer. Here, we resolve this paradox by focusing on the extended parental care inherent in prolonged dependency. We show experimentally that in non-social wasps, extended care can significantly reduce the impact of interspecific parasites. Under extended care, offspring are less vulnerable by the time they are exposed to parasites, and costs of parasitism are reduced because mothers have the option to terminate investment in failing offspring. By experimentally simulating aspects of extended care in a species where it is lacking, we demonstrate that neither benefit requires specialized behaviour. Such benefits could therefore offset the disadvantage of prolonged dependency in non-social species, thereby facilitating the evolution of helping.

Immature nest-building wasps are helpless larvae that are entirely dependent on adult carers for food. The duration of parental care is minimized in 'mass provisioning' wasps, including most non-social taxa^{7,8}: before it even hatches from the egg, each offspring is sealed into a cell containing all of the food that it will require for maturation, so that it is nutritionally independent of its mother. In contrast, almost all eusocial and a few non-social wasps have extended parental care. These 'progressive provisioners' feed each developing larva gradually as it grows^{7,8}. Whereas a single mass provisioner can fully provision about 1 offspring per day, provisioning is extended over 5–70 days in progressive provisioners^{4,9,10}. Even if they provision multiple offspring simultaneously, non-social progressive provisioners will, on average, leave fewer independent offspring than mass provisioners, because mothers are more likely