

Evolutionary stability for large populations

Daniel B. Neill

Department of Computer Science, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213-3891, USA

Received 29 August 2003; received in revised form 7 November 2003; accepted 9 November 2003

Abstract

We present a revision of Maynard Smith's evolutionary stability criteria for populations which are very large (though technically finite) and of unknown size. We call this the *large population ESS*, as distinct from Maynard Smith's *infinite population ESS* and Schaffer's *finite population ESS*. Building on Schaffer's finite population model, we define the large population ESS as a strategy which cannot be invaded by any finite number of mutants, as long as the population size is sufficiently large. The large population ESS is not equivalent to the infinite population ESS: we give examples of games in which a large population ESS exists but an infinite population ESS does not, and vice versa. Our main contribution is a simple set of two criteria for a large population ESS, which are similar (but not identical) to those originally proposed by Maynard Smith for infinite populations.

© 2003 Elsevier Ltd. All rights reserved.

Keywords: Evolutionary stability; ESS; Large populations

1. Introduction

The concept of an evolutionarily stable strategy (ESS), proposed by John Maynard Smith (Maynard Smith and Price, 1973; Maynard Smith, 1974, 1982) has been perhaps more influential than any other notion in the field of evolutionary games. As has been established by a number of papers in the 1970s and 1980s, the original mathematical definition of an ESS (Maynard Smith and Price, 1973, p. 17) is only correct under the assumption of an infinite population size. Riley (1979) and Vickery (1987) show that a *finite* population that is ESS in the Maynard Smith sense can be invaded by a mutant (see also Maynard Smith's (1988) response to Vickery and Vickery's (1988) reply). These results led to the concept of a *finite population ESS* (Schaffer, 1988). In a finite population, the probability that a mutant plays against another mutant is smaller than the probability that a player of the common strategy plays against a mutant, since the mutant player cannot face itself in a contest. Thus the ESS for a finite population game will not be identical to the ESS for an infinite population game, and this is the source of Riley's and Vickery's results.

Current work that applies the ESS concept typically falls into two camps: either the population size is *small*

and *known* (in which case, Schaffer's ESS concept is used) or the population size is *large* and *unknown*. In the latter, more common case, the population is assumed to be infinite, and Maynard Smith's ESS concept is used. The general rationale behind this approximation is that, *if the population is sufficiently large*, Schaffer's correction will not affect the results of the evolutionary process. This is a common argument in the evolutionary games literature: for example, see the discussion of the well-known Kandori–Mailath–Rob model of stochastic stability (Kandori et al., 1993) in the textbook of Fudenberg and Levine (1998, p. 142). Fudenberg and Levine note that the KMR model assumes a finite population of N players, and they argue that if players only play against the others in the population (and do not play against themselves), the distribution of opponents' play should depend on which strategy the player is currently using. However, they go on to claim that, “for reasonably large N this should not matter.”

This assumption, that Schaffer's correction can be neglected for finite but *large enough* populations, is suggested by his results on the hawk–dove game (Schaffer, 1988). Schaffer calculates the infinite population and finite population ESS strategies for the hawk–dove game, and finds that “in finite populations the ESS is to play hawk more often than in infinite populations, and that the smaller the population the greater the

E-mail address: neill@cs.cmu.edu (D.B. Neill).

probability of playing hawk.” However, he notes that as the population size increases, the difference between the infinite and finite population mixed strategy ESS goes to zero.

In cases where the finite population ESS converges in the limit to the infinite population ESS, as in the hawk–dove game, it makes sense to treat a large and unknown population as infinite: in these cases, we can use Maynard Smith’s criteria to test for evolutionary stability. However, there exist some games in which a strategy that meets Maynard Smith’s criteria is not evolutionarily stable for any finite population, no matter how large; in these cases, Schaffer’s correction cannot be neglected. However, Schaffer’s criteria cannot be applied when the population size is unknown, and thus neither concept applies to a population that is very large (though technically finite) and of unknown size. This requires a revision of Maynard Smith’s ESS criteria for large but finite populations, which we call the *large population ESS*.

2. The model

We consider a monomorphic population of a strategy X , which is visited periodically by a small number of mutants who play an alternative strategy $A \neq X$. We assume that the population is well mixed: the payoff to a given player is his average payoff against all other players in the population. If the payoff to players of the common strategy X is higher than the payoff to players of the alternative strategy A , then X resists invasion by A , and the mutant strain cannot invade the population. If X can resist invasion by any alternative strategy, we say that X is *evolutionarily stable*.

For any given strategies P and Q , we define $w(P|Q)$ as the payoff to a player of strategy P against a player of strategy Q . We can then define an evolutionarily stable strategy in terms of four quantities: the self-payoff $w(X|X)$ of the common strategy, the self-payoff $w(A|A)$ of the alternative strategy, the payoff $w(X|A)$ of the common strategy against the alternative strategy, and the payoff $w(A|X)$ of the alternative strategy against the common strategy.

The invasion criteria of Maynard Smith (Maynard Smith and Price, 1973; Maynard Smith, 1974, 1982) state that strategy X is evolutionarily stable if and only if, for any alternative strategy $A \neq X$:

1. $w(A|X) \leq w(X|X)$.
2. If $w(A|X) = w(X|X)$, then $w(X|A) > w(A|A)$.

We call a strategy meeting these criteria an *infinite population ESS*. This is because Maynard Smith’s criteria assume a continuum of individuals playing each strategy: thus an individual of type X and an individual of type A play against an identical mixture of strategies

X and A . Moreover, they assume that the population share of the alternative strategy A is approximately zero. Thus the payoff to each strategy is approximately its payoff against the common strategy X , and the payoff against the alternative strategy A is only relevant in the case of a tie.

When populations are large but finite, however, the Maynard Smith criteria may fail to give the correct results. Consider a population of N individuals, $N - 1$ of type X and a single mutant of type A . Let us assume that $w(A|X) = w(X|X) = 2$, $w(X|A) = 0$, and $w(A|A) = -1$, as shown in Table 1. Clearly, X is an infinite population ESS as defined above. However, for a finite population, the single mutant player interacts only with players of strategy X , and thus its average payoff is 2. Players of strategy X , on the other hand, interact both with players of X and the mutant, resulting in a payoff of slightly less than 2. For any finite N , no matter how large, the payoff to the mutant is higher than the payoff to the common strategy, and we would expect the mutant to survive and prosper. In fact, we find that the “break-even point” in this game is three mutants: if there are less than three mutants, the mutants have higher payoffs, and if there are more than three mutants, players of the common strategy have higher payoffs. Hence we would expect a small number of mutants to survive in the population. It is true that as the population size goes to infinity, the population share of mutants in the polymorphic population goes to zero; nevertheless, the important fact is that the mutants survive. This can make a huge difference if the combined population is then invaded by a third strategy. Consider the interaction given in Table 2, where alternative strategies A_1 and A_2 achieve high payoffs against each other. An invasion of X by either alternative strategy would result in a few mutants surviving; then invasion by the other alternative strategy would result in the population being taken over by a mixture of the two alternative strategies, and X would die off. For the interaction given in Tables 3 and 4, on the other hand, X can resist invasion by either alternative strategy, and

Table 1
 X is an infinite population ESS, but not a large population ESS

	vs. X	vs. A
X	2	0
A	2	-1

Table 2
 X is an infinite population ESS, but not a large population ESS

	vs. X	vs. A_1	vs. A_2
X	2	0	0
A_1	2	-1	1000
A_2	2	1000	-1

Table 3
X is ESS, for infinite or large populations

	vs. X	vs. A
X	2	4
A	2	3

Table 4
X is ESS, for infinite or large populations

	vs. X	vs. A ₁	vs. A ₂
X	2	4	4
A ₁	2	3	1000
A ₂	2	1000	3

thus cannot be taken over by a sequence of isolated mutations.

When is a strategy evolutionarily stable in large but finite populations? To answer this question, we turn to the finite population model of Schaffer (1988). Let us assume a population of N individuals, of which M are mutants playing the alternative strategy A , and $N - M$ play the common strategy X . Recall that an individual plays against all other individuals in the population, not including himself: a mutant player faces $M - 1$ mutants and $N - M$ common players, while a common player faces M mutants and $N - M - 1$ common players. Thus the average payoff to the common strategy X , given a population of N players with M mutants, is

$$\pi_X(N, M) = \frac{N - M - 1}{N - 1} w(X | X) + \frac{M}{N - 1} w(X | A).$$

The average payoff to the alternative strategy A , given a population of N players with M mutants, is

$$\pi_A(N, M) = \frac{N - M}{N - 1} w(A | X) + \frac{M - 1}{N - 1} w(A | A).$$

As in the models of Maynard Smith and Schaffer, we assume a deterministic population dynamics, neglecting the effects of random drift due to demographic stochasticity. Thus the proportion of mutants in the combined population is assumed to increase when $\pi_A > \pi_X$, decrease when $\pi_A < \pi_X$, and remain constant (with possible small fluctuations) when $\pi_A = \pi_X$. (In real, finite populations, random drift may cause the proportion of an inferior strategy to increase with some small but non-zero probability. For very small populations, this may even allow the inferior strategy to take over the population; however, this effect can be neglected when the population is large.)

A strategy is evolutionarily stable, for large but finite populations, if it cannot be invaded by any finite number of mutants, as long as the population size is sufficiently large. Formally:

Definition. Strategy X is a *large population ESS* if, for every alternative strategy $A \neq X$ and every positive integer M_0 , there exists a *population threshold* $N_0(M_0, A)$ such that, for all population sizes $N \geq N_0$ and for all positive integers $M \leq M_0$, we have $\pi_A(N, M) < \pi_X(N, M)$.

In other words, for a given number of mutants M_0 and alternative strategy A , if the population size is large enough ($N \geq N_0$), then the mutants' payoff is always less than the payoff to the common strategy, and thus the number of mutants M will decrease monotonically to zero. Note that the population threshold N_0 is allowed to depend on the alternative strategy A ; we can also define the stronger notion of a "uniform large population ESS" by requiring N_0 to apply uniformly to all $A \neq X$, as follows:

Definition. Strategy X is a *uniform large population ESS* if, for every positive integer M_0 , there exists a *population threshold* $N_0(M_0)$ such that, for all alternative strategies $A \neq X$, for all population sizes $N \geq N_0$, and for all positive integers $M \leq M_0$, we have $\pi_A(N, M) < \pi_X(N, M)$.

Note that, if the strategy space is finite, the two definitions are equivalent: we can set $N_0(M_0) = \max_A N_0(M_0, A)$. For an infinite strategy space, every uniform large population ESS is a large population ESS, but the converse does not necessarily hold.

Next we provide a simple set of two criteria for a large population ESS, which are similar (but not identical) to those originally proposed by Maynard Smith for infinite populations. These are given by the following proposition:

Proposition 2.1. *A strategy X is a large population ESS if and only if, for any alternative strategy A ≠ X:*

1. $w(A | X) \leq w(X | X)$.
2. *If $w(A | X) = w(X | X)$, then $w(X | A) \geq w(A | A)$ and $w(X | A) > w(A | X)$.*

Proof. We first compute the difference $\pi_X - \pi_A$. Subtracting the two expressions given above (and multiplying by $N - 1$), we find that the payoff for the common strategy is higher when $(N - M - 1)(w(X | X) - w(A | X)) + (M - 1)(w(X | A) - w(A | A)) + w(X | A) - w(A | X) > 0$. For a given alternative strategy $A \neq X$, there are three possibilities. If $w(X | X) > w(A | X)$, we know that the inequality holds for all sufficiently large N . If $w(X | X) < w(A | X)$, the inequality fails to hold for large N . Finally, if $w(X | X) = w(A | X)$, the inequality reduces to $(M - 1)(w(X | A) - w(A | A)) + w(X | A) - w(A | X) > 0$. For this latter expression to be true for all positive integers M , we must have $w(X | A) > w(A | X)$, and $w(X | A) \geq w(A | A)$. □

For a large population ESS to be a uniform large population ESS, one additional condition must be met: for all alternative strategies $A \neq X$ such that $w(X|X) > w(A|X)$, the quantities $[w(A|X) - w(X|A)]/[w(X|X) - w(A|X)]$ and $[w(A|A) - w(X|A)]/[w(X|X) - w(A|X)]$ must be uniformly bounded. As above, this is true for any finite strategy space, but is not necessarily true when the strategy space is infinite.

3. Discussion

We now consider the two cases in which the evolutionary dynamics differ for the infinite population and large population models. The first case occurs when $w(A|X) = w(X|X)$ and $w(X|A) > w(A|A)$. For an infinite population, the Maynard Smith criteria state that X resists invasion by A : if mutants of type A attempt to invade a monomorphic population of strategy X , then the population will converge (on a population dynamical time-scale) to a state in which no mutants survive. For a large but finite population, there are two possibilities: either $w(X|A) > w(A|X)$, in which case X resists invasion by A (as for an infinite population), or $w(X|A) \leq w(A|X)$, in which case a finite number of mutants can survive. As discussed above, Table 1 gives an example where some mutants survive, and Table 3 gives an example where the mutants are wiped out. To compute the number of surviving mutants, we again consider the quantities $\pi_X(N, M)$ and $\pi_A(N, M)$ given above. If $w(A|X) = w(X|X)$, we find that the payoffs to the common and alternative strategies are equal when $Mw(X|A) - (M - 1)w(A|A) - w(A|X) = 0$. Solving for M , we find that the number of mutants at the break-even point is:

$$M = \frac{w(A|X) - w(A|A)}{w(X|A) - w(A|A)}$$

Thus $M = 3$ for the example in Table 1. It is possible, of course, that a small number of mutants will be wiped out eventually by random drift; however, we assume as above that the amount of drift is negligible as compared to the selective pressures favoring the mutants' survival.

The second case where the large population dynamics differ from the infinite population dynamics occurs when $w(A|X) = w(X|X)$ and $w(X|A) = w(A|A)$, as in Table 5. For an infinite population, A and X obtain exactly the same payoff: hence there is no evolutionary pressure, and the proportions of A and X either remain constant, or fluctuate due to random drift. For a large but finite population, the common strategy X will have higher payoffs if $w(X|A) > w(A|X)$; in this case it will resist the invasion of the alternative strategy A . If $w(X|A) < w(A|X)$, on the other hand, strategy A will have higher payoffs, allowing it to invade and eventually to replace the common strategy. Finally, if

Table 5
 X is a large population ESS, but not an infinite population ESS

	vs. X	vs. A
X	2	3
A	2	3

$w(X|A) = w(A|X)$, the payoffs are identical, so (as in the infinite population case) there is no evolutionary pressure. It is interesting to note that, in situations where there is no evolutionary pressure in infinite populations, the large population results show that the strategy with lower self-payoff wins. This is again because a strategy cannot play against itself in an evolutionary interaction: the strategy with the lower self-payoff plays more of the high-payoff opponents than the strategy with the higher self-payoff, and thus achieves a higher average score.

We also briefly examine three other stability concepts: neutral stability (Maynard Smith, 1982), collective stability (Axelrod, 1984), and Boyd and Lorberbaum's (1987) stability. All of these concepts assume an infinite population size, as in Maynard Smith's infinite population ESS. A strategy X is a collectively stable strategy (CSS) if Maynard Smith's first criterion is satisfied. A strategy X is a neutrally stable strategy (NSS) if Maynard Smith's first criterion is satisfied, and his second criterion is satisfied with weak (instead of strict) inequality. Both of these concepts are weaker than the large population and infinite population ESS: every large population or infinite population ESS is an NSS, and every NSS is also a CSS. Boyd and Lorberbaum require a stable strategy to satisfy Maynard Smith's first criterion, but also a second criterion that, if $w(A|X) = w(X|X)$ then $w(X|Y) \geq w(A|Y)$ for all strategies Y . This allows a strategy to resist a simultaneous invasion by mutants of different types, which is a stronger criterion than resisting isolated invasion by a single type of mutant. For example, strategy X in Table 4 would not be stable in the Boyd and Lorberbaum sense. However, Boyd and Lorberbaum do not require a strategy to be robust against evolutionary drift, and thus a strategy that is stable in the Boyd and Lorberbaum sense could be displaced by another strategy with identical payoffs.

Thus the large population ESS differs from Maynard Smith's infinite population ESS in two ways: in certain cases, strategies are less resistant to weakly dominated strategies, and more resistant to evolutionary drift. Because the two sets of criteria differ, and because evolutionary biologists are typically interested in models where the population is finite (though it may be arbitrarily large), we claim that the large population ESS is a better test of evolutionary stability in these interactions. Unlike the original Schaffer model (which

these results extend to the large population case), neither the exact size of the population, nor the exact number of mutants, need to be considered. Instead, we have proposed two simple criteria (similar, but not identical to those originally proposed by Maynard Smith for infinite populations) under which a strategy is evolutionarily stable for large populations. We propose that these criteria be used instead of Maynard Smith's in models where the population is technically finite but arbitrarily large.

References

- Axelrod, R., 1984. *The Evolution of Cooperation*. Basic Books, New York.
- Boyd, R., Lorberbaum, J., 1987. No pure strategy is evolutionarily stable in the repeated prisoner's dilemma game. *Nature* 327, 58–59.
- Fudenberg, D., Levine, D.K., 1998. *The Theory of Learning in Games*. MIT Press, Cambridge.
- Kandori, M., Mailath, G.J., Rob, R., 1993. Learning, mutation, and long run equilibria in games. *Econometrica* 61 (1), 29–56.
- Maynard Smith, J., 1974. The theory of games and the evolution of animal conflict. *J. Theor. Biol.* 47, 209–221.
- Maynard Smith, J., 1982. *Evolution and the Theory of Games*. Cambridge, UP, Cambridge.
- Maynard Smith, J., 1988. Can a mixed strategy be stable in a finite population? *J. Theor. Biol.* 130, 247–251.
- Maynard Smith, J., Price, G.R., 1973. The logic of animal conflict. *Nature* 246, 15–18.
- Riley, J.G., 1979. Evolutionary equilibrium strategies. *J. Theor. Biol.* 76, 109–123.
- Schaffer, M.E., 1988. Evolutionary stable strategies for a finite population and variable contest size. *J. Theor. Biol.* 132, 469–478.
- Vickery, W.L., 1987. How to cheat against a simple mixed strategy ESS. *J. Theor. Biol.* 127, 133–139.
- Vickery, W.L., 1988. Reply to Maynard Smith. *J. Theor. Biol.* 132, 375–378.