

Effects of Evolution on a Nonzero-Sum Game

Ruiyuan Chen

Science One Program, University of British Columbia, Vancouver, Canada

May 28, 2010

Abstract

An N -player nonzero-sum game was modelled on a computer and evolutionary algorithms were used to determine whether the evolutionarily stable strategy coincides with the group optimum behavior.

Introduction

This project aims to model and study the following game:

N players sit in a line and take turns standing up. Each player seeks to stand up at a unique time. If two or more players stand up at the same time, they are all punished (for example, made to do pushups) and the game ends. If the game continues until only one player is left sitting, that last player is punished instead.

This game is *nonzero-sum*, meaning that the total gain (or loss, in this case) of all the players depends on the choices that the players make: if all players stand up as soon as the game starts then all will be punished, whereas if each player stands at a random time then probably at least one player will pass unscathed. This is unlike a zero-sum game such as gambling, in which total wealth is conserved and redistributed among the players. In contrast, the players of a nonzero-sum such as that described above together determine how much joy or pain is created and shared.

The goal of this study is to determine the effects of evolution on the players of the game described above. After having an initial population play a certain number of games, we mimic natural selection, killing the losers and allowing the winners to reproduce. Many generations later, we expect to be left with a homogeneous population consisting of the *evolutionarily stable strategy* (*ESS*), that will be able to outcompete any minority groups in the population [1]. However, such a selection procedure works only on an individual level. Well-known nonzero-sum games like the prisoner's dilemma exemplify a key feature of this class of games: maximization of *individual gain* may result in suboptimal *group gain*. Thus the effective goal of this study is to determine whether such a conflict between individualism and collectivism exists in the game described above.

Methods

In order to simulate the game on a computer, we created the following simplified model. The times at which players may stand are represented as nonnegative integers, and each strategy is

Type and Parameters	pmf(k)
Fixed(t): always stand at time $t \in \{0, 1, 2, \dots\}$	1 if $k = t$, 0 otherwise
Uniform(n): stand before time $n \in \{1, 2, 3, \dots\}$	$1/n$ if $k < n$, 0 otherwise
Geometric(p): at any time, if still seated, stand with probability $p \in (0, 1]$	$p(1 - p)^k$
Bell(q): probability of remaining seated starts at 1 and decreases by a factor of $q \in [0, 1)$ at each time	$q^{k(k-1)/2}(1 - q^k)$
Uniform2(a, b): stand at or after time $a \in \{0, 1, 2, \dots\}$ but before time $b > a$	$1/(b - a)$ if $a \leq k < b$, 0 otherwise

Table 1: The types of strategies used in the initial populations, their parameters, and their probability mass functions (probabilities of standing at each possible time k)




Figure 1: pmf curves for strategies (a) Fixed(10), (b) Uniform(4), (c) Geometric(0.5), (d) Bell(0.9), and (e) Uniform2(14, 17)

represented by a probability distribution on the set of nonnegative integers. In a game, each player stands at a time randomly chosen according to his strategy’s distribution. If two or more players choose the same time, the earliest such group of players is “punished” by a single point added to their score; otherwise the player who chooses the latest time gains a single point.

Note that these assumptions fail to reflect many aspects of the real-life game, such as the ability to change one’s strategy mid-game depending on when other players stand, the potential to learn about other players’ strategies and improve one’s own strategy for future games, and various psychological effects. While our assumptions were necessary to allow efficient simulation of the game on a computer, they do mean that the results of this study are best interpreted only in the context of our simplified model. Any extrapolation to the original game must take the above factors into account.

For each simulation, we set up an initial population of strategies with simple discrete probability distributions, listed in Table 1, with probability mass functions shown in Figure 1. Multiple ($M \in [100, 1000]$) games were then played in order to reduce the random effects of the choices of standing times. At the end of all the games, the total score of all the players was divided by the number of games, M , to obtain the *total loss per game* L . The players were then ranked in order of increasing score, and a certain fraction (E) of the highest-scoring strategies were removed. An evolutionary algorithm (with parameter P) was used to produce “children” of the surviving strategies which fill the vacancies in the population; the algorithms used are shown in Table 2. All of this constitutes one generation of the simulation. We call the initial population “generation 0”, the population after one evolution “generation 1”, and so on. A minimum of 8 generations were run, and then more generations were run until the population had “converged”. “Convergence” is defined as the earliest generation when the mean standing time of the entire population has not changed by more than 0.01 over the last half of all generations up to that point.

Type and Parameter (P)	Description
evolve P : # of parents per child	P parents are selected randomly (with replacement) and combined into a child that plays according to each parent with a uniformly randomly chosen probability
evolve2 P : # of children per parent	Parents' pmfs are split at P -quantiles (median for $P = 2$, quartiles for $P = 4$, and so on) into P new pmfs; the lowest-scoring parent gives rise to the first P children, the second-lowest-scoring gives rise to the next P , etc. until the population is filled again
evolve3 P : # of parents per child	Same as evolve , except parents are not combined at random; rather, the first P are combined, the next P are combined, etc. such that each group of P gets to procreate at least once (assuming enough vacancies in the population)

Table 2: The types of evolutionary algorithms used

In order to determine the average total loss per game L of the convergent population, we took the mean and standard error of the L of the last half of all generations. There were two exceptions. First, if the last generation consisted of a single deterministic strategy (i.e. with a probability mass function zero except at a single point), we took the L of that generation, with an uncertainty of 0, since all future generations will have the same L . Second, for each individual population we ran a control simulation in which no evolution occurred; in these control simulations, 100 generations were always run, and the L of all 100 generations were averaged.

We conducted the simulations in two stages, with each simulation using an initial population of size 64. In stage 1, we conducted 8 simulations, consistently using evolutionary algorithm **evolve** (see Table 2) while varying the composition of the initial population. In any simulation, each type of strategy present in the initial population was present in equal numbers, and the parameters of the strategies of each type were distributed uniformly across the parameters' domains. For example, an initial population of **Fixed**s and **Geometric**s is the set $\{\text{Fixed}(t) : t = 0, \dots, 31\} \cup \{\text{Geometric}(p) : p = 1/32, \dots, 32/32\}$. In stage 2, we conducted 12 simulations, keeping an initial population of $\{\text{Uniform2}(a, b) : 0 \leq a < 8, a < b \leq a + 8\}$ while changing both the type and parameter P of the evolutionary algorithm used.

Results and Discussion

The results of the stage 1 simulations are shown in Table 3. A quick glance at the L s at convergence and the ratios to the L s of the controls reveals that while evolution effectively reduced the L value in most of the simulations, there are three blatant exceptions to this trend. Note that these are exactly the simulations that began with **Fixed**s in the initial population. Figure 2 shows the pmfs of the strategies in the **Fixed**+**Uniform** simulation at four generations. In Fig. 2a we see the “spikes” that are the pmfs of $\text{Fixed}(0), \dots, \text{Fixed}(31)$, as well as the pmfs of the **Uniform**s clustered around time 0. After 3 generations, as shown in Fig. 2b, the strategies that stand up early have killed each other through competition, while the surviving strategies have successfully reproduced. However, because children inherit from parents, this reproduction is effectively suicidal, for the parents end up competing with their own offspring. Meanwhile, the late-standing strategies proliferate, such that by generation 20 (Fig. 2c), the once-sparsest region on the pmf graph has become the densest. This continues until the latest-standing strategy (Fig. 2d) is the only strategy left.

Initial population	L at convergence	L of control	L ratio
Fixed	44.8 ± 0.5	1 ± 0	44.8 ± 0.5
Uniform	2.413 ± 0.004	4.805 ± 0.006	0.5022 ± 0.0010
Geometric	2.394 ± 0.004	32.50 ± 0.02	0.07366 ± 0.00013
Bell	3.021 ± 0.009	32.49 ± 0.02	0.0930 ± 0.0003
Fixed + Uniform	40.4 ± 1.5	5.053 ± 0.005	8.0 ± 0.3
Uniform + Geometric	2.931 ± 0.005	3.17 ± 0.18	0.92 ± 0.05
Geometric + Bell	3.502 ± 0.008	16.495 ± 0.007	0.2123 ± 0.0005
all 4 types	39.5 ± 1.4	12.884 ± 0.006	3.07 ± 0.11

Table 3: Stage 1 simulation results, showing the types of strategies in the initial population, the average total loss L at convergence, the L of the control, and the ratio of these two L s. The exact compositions of the initial populations were as described in the Methods section.




Figure 2: pmfs of strategies in **Fixed + Uniform** at generations 0, 3, 20, and 60

We can see, then, that in the simulations involving **Fixed** strategies, the ESS is the latest-standing strategy because such a strategy avoids competition until all other strategies have died. In fact, this also applies to the simulations in Table 3 *not* involving **Fixed**s, and explains the reduction of L in those simulations just as well: for **Uniform**s, **Geometric**s, and **Bell**s are all strategies where the latest-standing strategy also has the most spread-out pmf curve (see Figure 1). Thus in an initial population lacking **Fixed**s, the ESS is far less likely to compete with itself, which explains the effectiveness of evolution in reducing L .

Suppose that in some initial population including **Fixed**s, a lone early-standing strategy somehow survives without competition until the rest of the population has converged to an ESS similar to that shown in Figure 2d. The early-standing strategy has an enormous advantage, a near-guarantee of neither competing with another strategy, nor being punished for being the last player sitting – for a single turn. As soon as the early-standing strategy is rewarded with like-minded offspring, competition kicks in and kills both parent and children. This ability of the equilibrium to withstand invasion allows us to fully justify labelling the latest-standing strategy as the ESS of any given initial population.

The results of the stage 2 simulations are shown in Table 4. There are a number of trends, the most apparent being the common L value of 64 ± 0 for all of the simulations using **evolve2**. Since the size of the population was 64, this was the maximum L possible. The uncertainty is 0 because in each of the 4 simulations, the final population was completely homogeneous and consisted of a single deterministic strategy. This is the worst possible collectivistic outcome, since the population must maintain the maximum L for all future generations. This result makes theoretical sense, for recall the pmf-splitting mechanism of **evolve2** (Table 2): the only populations on which **evolve2** does *not* act must consist of identical deterministic strategies.

Algorithm	P	E	L at convergence	L of control	L ratio
evolve	2	0.5	9.90 ± 0.05	3.279 ± 0.014	3.02 ± 0.02
	5		8.83 ± 0.02		2.693 ± 0.013
	8		9.26 ± 0.02		2.824 ± 0.014
	2	0.25	8.66 ± 0.04		2.641 ± 0.017
evolve2	2	0.5	64 ± 0		19.52 ± 0.08
	5		64 ± 0		19.52 ± 0.08
	8		64 ± 0		19.52 ± 0.08
	2	0.25	64 ± 0		19.52 ± 0.08
evolve3	2	0.5	8.50 ± 0.05		2.59 ± 0.02
	5		8.77 ± 0.03		2.675 ± 0.015
	8		10.61 ± 0.05		3.24 ± 0.02
	2	0.25	8.51 ± 0.04		2.595 ± 0.016

Table 4: Stage 2 simulation results, showing the evolutionary algorithm used and its parameter P , the selection rate E , the L at convergence, the control L , and the L ratio.

Examining the L ratios for the `evolve` and `evolve3` simulations, we see that these two evolutionary algorithms produced net increases in L that were largely independent of the evolutionary parameters P and E . Furthermore, the effect of `evolve3` on the pmf curves of the population looks very similar to the sequence shown in Figure 2: a gradual but steady recession of early-standing strategies that leads to a late-standing ESS. Thus it appears that the key features of both algorithms are the same: the blessing/cursing of well-scoring parents with like-minded children. And just as the detailed composition of the initial populations in stage 1 did not affect the general “evolutionarily stable” = “late-standing” trend, such quantitative fine print as the parameter P or the selection rate E seem unimportant to describing the general effects of evolution.

Conclusion

The effects of evolution on our modelled game were studied using different initial populations and evolutionary algorithms and were found to optimize neither individualistic nor collectivistic behavior. However, the latest-standing strategy became the ESS in all cases, irrespective of the initial population and the evolutionary algorithm. It seems remarkable that a complex system with many variables always behaves in the same way, independently of those variables.

Acknowledgments

I would like to thank Fok-Shuen Leung for allowing me to proceed with this project and giving me some early feedback on my results, Josh Zukewich for a number of useful suggestions for improvement, and Andy Tan for reviewing and commenting on my paper.

References

- [1] Maynard Smith, J. and Price, G.R. The Logic of Animal Conflict. *Nature* **246**, 15-18 (1973).