

Evolutionary Game Theory- Evolutionary Stable Strategies

Michael Levet

March 9, 2016

1 Introduction

Neoclassical models of Game Theory consider fully rational agents, where each agent is perfectly aware of the other agents, their strategy sets, and the payoffs of each strategy profile. Each player is fully rational, and so is able to examine all possibilities and select a best option, taking into account the other players. The solution concept in Game Theory is the Nash Equilibrium, which intuitively is an outcome where no individual player can unilaterally deviate and strictly improve its outcome. However, more traditional models of Game Theory only examine the equilibrium outcomes, not the underlying dynamics driving them.

Evolutionary Game Theory seeks to provide a more realistic model of human behavior, while encompassing the notion of Darwinian dynamics or survival of the fittest. It has two key differences which separates it from traditional models of Game Theory. First, Evolutionary Game Theory models the dynamics. Rather than focusing solely on the equilibrium, the focus is shifted towards how such an equilibrium comes about, if one even exists. Realistically, there are always dynamics, even if they are negligible. In this sense, the evolutionary framework captures this realism better than the neoclassical framework. The second key difference is the notion of bounded rationality. In Evolutionary Game Theory, individuals are not perfectly rational. This means that not every agent computes every possible payoff, making the best decision. Non-rational agents may ignore certain players or possibilities, or opt to imitate the majority rather than maximize their payoffs. In many competitive situations, it may not be feasible to compute all possible outcomes even in the case of perfect information. Instead, agents start with initial strategies and adapt. The outcomes of the game reward the winners and punish the losers. Agents slowly adapt new and beneficial strategies, factoring in what actions they think the other players might choose. In this sense, agents are more recognizably human. Furthermore, the evolutionary framework yields similar equilibria in many cases as the neoclassical Game Theory model.

There are several important applications of Evolutionary Game Theory, outside of economics and biology. Sociologists and psychologists find applications in studying the prevalent social customs, both in small groups and large societies. Researchers in artificial intelligence, particularly multiagent systems, use Evolutionary Game Theory to study how repeated local interactions amongst arbitrary subsets of players drive the overall state of the system. Within the context of economic networks, Evolutionary Game Theory proves to be a useful tool in studying local interaction and learning.

This blog entry will introduce the evolutionary framework, as well as explore the solution concept of the Evolutionary Stable Strategy, which is a refinement of the Nash Equilibrium. The goal is to introduce textbook formulations of the Evolutionary Stable Strategy, while the next blog entry will discuss adaptations of this concept that frequently arise in the literature. Some familiarity with the basics of Game Theory is advisable, such as my previous blog entry on the subject.

2 Evolutionary Framework

Recall the definition of a normal form game:

Definition 1 (Normal Form Game). A normal form game is a three-tuple $\Gamma = [N, (S_i)_{i \in N}, (u_i)_{i \in N}]$ where N is the set of players, S_i is player i 's strategy set, and $u_i : \prod_{i=1}^n S_i \rightarrow \mathbb{R}$ is player i 's utility function.

In the neoclassical model, each player in the game is fully rational. Each player selects its strategy simultaneously, and the game is played exactly once. The evolutionary framework, in contrast, has three key assumptions:

1. There is a large population of non-fully rational agents.
2. The game is played multiple times. In a given round, each player selects its strategy simultaneously.
3. Some agents are drawn at random each period to play the game.

We denote the number of players in a given round as the contest size, while the population size is the number of players from which we can draw. The goal of the evolutionary framework is to examine which types of players or strategies survive repeated play, shedding light on which equilibrium is more realistic. Using a biological intuition, we compare the fitness of strategies by assuming certain proportions of the population are programmed to play a given strategy exclusively. In this way, the agents playing the "fittest" strategy exhibit a comparative advantage relative to those agents playing other strategies. Let's examine a coordination game.

Example 1: Consider the following payoff matrix.

		Player 2	
		A	B
Player 1	A	5, 5	0, 3
	B	3, 0	3, 3

The two pure strategies Nash equilibria of this game are (A, A) and (B, B)- it is easy to check that no player can unilaterally deviate from one of these strategy profiles and improve its payoff. Furthermore, it is clear that players prefer the equilibrium (A, A) to (B, B). However, which of these equilibria is more realistic?

Let's now apply the evolutionary framework to the coordination game. Let N be the finite set of players. At each round, two players are selected at random from N to play the coordination game. The players with the highest payoffs are most likely to survive.

If two players both play A, then their preferred Nash equilibrium is achieved. Suppose player i chooses to play B instead, while player $-i$ maintains the strategy A. Then player i has utility 3 while player $-i$ has utility 0. As the game repeats, agents playing the strategy B have a comparative advantage over those playing A. Agents playing the strategy A will eventually die out when faced with agents playing strategy B, and the equilibrium (B, B) will arise even though (A, A) is the preferred equilibrium. This phenomenon is captured by Evolutionary Game Theory with the solution concept known as the Evolutionary Stable Strategy (ESS).

3 Evolutionary Stable Strategy

Recall that a strategy profile is a Nash equilibrium if no player can unilaterally deviate and improve its payoff. The Evolutionary Stable Strategy is intuitively a strategy that, if played by a sufficiently large proportion of the population, is resistant to the adoption of a mutant strategy. Recall from Example 1 in the previous section that players adoption the strategy A in the coordination game ended up dying out if the strategy B was adopted by a fraction of the players. In this sense, the strategy A is not resistant to mutations.

Prior to engaging in the technical details, some motivation is first in order. Evolutionary Game Theory originated from mathematical biology rather than economics. This is a reasonable starting point. Suppose we have a population of small insects. Anytime two insects come across a food source, they compete for it, enabling each insect to lay five eggs. Now suppose a mutation is introduced, allowing mutant insects to grow to a larger size. So now we have small insects and large insects. A large insect is only able to lay three eggs after consuming food, as it requires more energy to maintain its size. When two large insects compete, they are each able to lay three eggs. However, when a large insect and a small insect compete, the large insect is able to consume all the food. Even though the small insects are able to reproduce more efficiently, will they survive in the long run with the mutation introduced? The answer is no. Observe that this is the game from Example 1.

While the notion of mutation has clear roots in biology, it also applies to social situations. Language coordinates communication, and it evolves by sufficiently many individuals adopting new terminology. In a market setting,

suppose a producer introduces a new good. The profitability of this good is determined by whether sufficiently many customers purchase it. Here, the adopting customers can be viewed as mutants in the population. Other societal conventions such as etiquette, law, and currency have evolved based on the natural dynamics arising from individual interactions amongst agents. The Evolutionary Stable Strategy seeks to capture conventions that arise from such interactions.

We now formalize the notion of the Evolutionary Stable Strategy. A Darwinian and biological intuition suggests that evolutionary forces will select against the mutant strategy if and only if the incumbent strategy has a higher payoff. This is the fitness condition. Formally, consider a two-player game Γ with the mixed strategies set Δ . Suppose that the incumbent strategy is $x \in \Delta$ and the mutant strategy is $y \in \Delta$. Let $\epsilon \in (0, 1)$ be the proportion of players that are programmed to play y . So an agent playing strategy $a \in \{x, y\}$ has expected payoff $(1 - \epsilon)u(a, x) + \epsilon u(a, y)$, which is the same expected payoff as if the opponent was playing mixed strategies x with probability $(1 - \epsilon)$ and y with probability ϵ . So evolutionary forces favor strategy x if and only if:

$$u[x, (1 - \epsilon)x + \epsilon y] > u[y, (1 - \epsilon)x + \epsilon y] \quad (1)$$

This is precisely the condition for a strategy to be evolutionary stable.

Definition 2 (Evolutionary Stable Strategies). A strategy $x \in \Delta$ is an Evolutionary Stable Strategy if for every strategy $y \neq x$, there exists an $\epsilon_y \in (0, 1)$ such that $u[x, (1 - \epsilon)y + \epsilon x] > u[y, (1 - \epsilon)y + \epsilon x]$ for every $\epsilon \in (0, \epsilon_y)$.

Intuitively, ϵ_y is the barrier for invasion. If there are not a sufficient number of mutants, then the incumbent agents dominate and the mutants die out. The evolutionary stable strategies can also be characterized as a subset of strict Nash equilibria. Formally, we have the following theorem.

Theorem 3.1. Let Δ^{ESS} be the set of evolutionary stable strategies for a two-player game Γ , Δ^{NE} be the set of Nash equilibria for Γ , and $\beta(x)$ be the best response correspondence of strategy x . We have:

$$\Delta^{ESS} = \{x \in \Delta^{NE} : u(y, y) < u(x, y), \forall y \in \beta(x), y \neq x\} \quad (2)$$

Proof. Let $x \in \Delta^{ESS}$, and let y be a strategy other than x . Let $\epsilon_y \in (0, 1)$ be the barrier to invasion of y -mutants against x . As x is an Evolutionary Stable Strategy, we have for every $\epsilon \in (0, \epsilon_y)$.

$$u[x, (1 - \epsilon)x + \epsilon y] > u[y, (1 - \epsilon)x + \epsilon y] \quad (3)$$

By multilinearity of the expected utilities, we have:

$$(1 - \epsilon)u[x, x] + \epsilon u[x, y] > (1 - \epsilon)u[y, x] + \epsilon u[y, y] \quad (4)$$

Taking $\epsilon \rightarrow 0$, we have either $u(x, x) > u(y, x)$; or $u(x, x) = u(y, x)$ and $u(x, y) > u(y, y)$. Note that x is necessarily a Nash equilibrium; otherwise, there would exist a strategy z resulting in higher expected payoff than x , contradicting the assumption that $x \in \Delta^{ESS}$.

Conversely, let s be a Nash equilibrium satisfying $u(a, a) < u(s, a)$ for all $a \in \beta(s) - \{s\}$, and let $t \in \beta(s)$. Then we have $u[s, s] > u[t, s]$; or $u[s, s] = u[t, s]$ and $u[s, t] > u[t, t]$. If $u[s, s] > u[t, s]$, then for $\epsilon = 0$ we have:

$$(1 - \epsilon)u[s, s] + \epsilon u[s, t] > (1 - \epsilon)u[t, s] + \epsilon u[t, t] \quad (5)$$

And so $s \in \Delta^{ESS}$. Now suppose instead that $u[s, s] = u[t, s]$ and $u[s, t] > u[t, t]$. We have:

$$u[s, s] + u[s, t] > u[t, s] + u[t, t] \quad (6)$$

Multiplying $u[s, t]$ and $u[t, t]$ by a fixed $\epsilon \in (0, 1)$, and multiplying $u[s, s]$ and $u[t, s]$ by $(1 - \epsilon)$ yields the result. \square

While every game has a Nash equilibrium (possibly in mixed strategies), not every game has an Evolutionary Stable Strategy. In fact, deciding if a game has an Evolutionary Stable Strategy is NP-Complete. We consider a couple examples to illustrate this point.

Recall Example 1, shown below. Intuitively, we discussed why (B, B) is evolutionary stable. It is clear that (B, B) is a Nash equilibrium. As $u[B, B] = 3 > 0 = u[A, B]$, we have that B is evolutionary stable. Now while (A, A) is a Nash equilibrium, we have $u[A, B] = 0 < 5 = u[A, A]$, so A is not stable.

		Player 2	
		A	B
Player 1	A	5, 5	0, 3
	B	3, 0	3, 3

Example 2: Consider a second game:

		Player 2	
		A	B
Player 1	A	3, 3	1, 2
	B	2, 1	0, 0

It is easy to verify that (A, A) is the unique Nash equilibrium of this game. In fact, (A, A) is a strict Nash equilibrium, which means that any deviation strictly decreases an agent's payoff. Now we have $u[A, A] > u[B, A]$ and $u[A, B] > u[B, B]$. So for any $\epsilon \in (0, 1)$, we certainly have:

$$(1 - \epsilon)u[A, A] + \epsilon u[A, B] > (1 - \epsilon)u[B, A] + \epsilon u[B, B] \tag{7}$$

Thus, A is the Evolutionary Stable Strategy for this game. In fact, for any game, its strict Nash equilibria must be Evolutionary Stable Strategies.

Example 3: Consider the matching pennies game, shown below. The unique Nash equilibrium is for each agent to play the mixed strategy σ of A and B with frequency $\frac{1}{2}$ each. However, it fails to meet the stability condition. Suppose player 1 plays A instead of σ . We have $u(\sigma, \sigma) = u(A, \sigma) = 0$. However, $u(\sigma, A) = 0 < u(A, A) = 1$.

		Player 2	
		A	B
Player 1	A	1, -1	-1, 1
	B	-1, 1	1, -1

4 Conclusion

In this blog entry, we explored the Evolutionary Stable Strategy solution concept, which is a refinement of the Nash equilibrium. In addition to finding Evolutionary Stable Strategies in games, we discussed motivations for this concept in understanding the dynamics that drive equilibria in games. The next blog entry will further discuss refinements and adaptations of the Evolutionary Stable Strategy, with important applications and results.