Played By the System: Chaotic Games and the Death of Foresight

Ethan Klein '16

November 24, 2015

The study of game theory focuses on "games" between players which may be anything from individuals, businesses, governments, or competing populations of animals. Each player has a set of n "moves" $\{E_1, E_2, \ldots, E_n\}$ from which they form a strategy. These are either pure strategies in which a player always plays one move, or mixed strategies defined as a probability distribution over their moves, in which players randomly decide between their moves according to the distribution. What each player gains or loses given their move and the moves of the other players is held as a matrix. As the example this paper will mainly focus on, we look to Rock, Paper, Scissors(RPS). RPS is played by 2 players A and B, both of which share a moveset $\{Rock, Paper, Scissors\}$. Their payoffs are held by the payoff matrices A for Alice and B for Bob, where

$$A = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix} B = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}$$

Note that in this case (and, in fact, in any zero-sum game between 2 players), $A = -B^T$. A natural and useful interpretation, of particular interest to ecologists, is to consider the "players" as populations of one or more species and each move E_i to represent a phenotype. Here, the frequency of $E_i$in a mixed strategy is the proportion of population 1 that is type i, and A_{ij}in payoff matrix A holds the payoff value to a member of species 1 of type i in competition with a type j member of population 2.

A primary focus of classical game theory is to find strategy sets in games such that no player would wish to change their strategy, called Nash Equilibria, often with the sense that these strategies are ones which players will naturally gravitate towards as the most rational strategies. However, this leaves unclear how players actually arrive at these strategies, if they do so at all. It is entirely possible that in the face of irrational play, this supposedly rational play is in fact not optimal. To approach this issue, we develop a dynamical system to represent the game and strategy shifts over time.

To see how players update their strategies with time, we must consider the strategies as a vector space. Each strategy becomes a vector $\vec{\mathbf{x}} = (x_1, \ldots, x_n)$ for player 1 and $\vec{\mathbf{y}} = (y_1, \ldots, y_m)$ where x_i is the chance that $\bar{\mathbf{x}}$'s player plays move x_i. Because of this probabilistic restriction, $\sum_{i=1}^{n} x_i = 1$, $\sum_{i=1}^{m} y_i = 1$, and all compenents of $\vec{\mathbf{x}}$ and $\vec{\mathbf{y}}$are positive. We graph the strategies for players 1 and 2 on separate simplexes embedded in $\mathbb{R}^n$ and $\mathbb{R}^m$respectively, with each pure strategy at a vertex. The cartesian product of these simplexes forms our total strategy space. Now we can calculate the payoff to player 1 as $\mathbf{x}A\mathbf{y}^{\mathbf{T}}$. Such a view allows us to analyze strategy shifts dynamically.

Dynamic Strategy Evolution: The Replicator Equations

A particularly popular and useful model is that of "replicator equations". Here, each player continually updates their strategies towards the highest payoff along a flow which naturally arises from evolutionary reasoning as in Hofbauer(1998). As I will then show, this is a version of gradient ascent scaled to the boundaries set by game theory.

The intuition for the replicator equations is simple in evolutionary terms. For a given phenotype E_i with frequency x_i, we measure the "evolutionary success" as how fast the extant group of that phenotype is expanding, expressed as rate of increase$\frac{\dot{x_i}}{x_i}$. By basic Darwinian reasoning, such success is the difference between our phenotype's success and the average success of the population. The" average succes" is the sum of the possible outcomes weighted by the probability with which each outcome happens, $\sum_{i=1}^{n} \sum_{j=1}^{m} x_i A_{ij} y_j$, which is just $\mathbf{x}A\mathbf{y}^{\mathbf{T}}$. Thus we write:

$$\dot{x_i} = x_i((A\mathbf{y}^{\mathbf{T}})_i - \mathbf{x}A\mathbf{y}^{\mathbf{T}})$$

for each phenotype x_i, where $(A\mathbf{y}^{\mathbf{T}})_i$ is the i_{th}element of $(A\mathbf{y}^{\mathbf{T}})$.

Note that $\sum_{i=1}^{n} \dot{x}_i = \sum_{i=1}^{n} (xA\mathbf{y}^{\mathbf{T}})_i - \sum_{i=1}^{n} x_i \mathbf{x}A\mathbf{y}^{\mathbf{T}} = 0$, so the sum total of $x_i$over all i is still 1, as the probabilistic interpretation demands. To look from a more purely mathematical perspective appropriate for single players in addition to whole population, observe that

$$\sum_{i=1}^{n}(A\mathbf{y}^T)_i = \sum_{i=1}^{n}\frac{\delta}{\delta x_i}xAy^T = \vec{\nabla}_{\mathbf{x}}\mathbf{x}A\mathbf{y}^{\mathbf{T}}$$

. Thus, players follow the gradient of the payoff matrix ever towards the maximum with rescaling to ensure that the strategy vector for each player does not leave their strategy space.

Hamiltonian Dynamics of Replicator Equation

Here we show that RPS displays Hamiltonian dynamics. Following Hofbauer (1996), we change variables to $u_i = log(\frac{x_i}{x_1})$and $v_j = log(\frac{y_i}{y})$ for $(\mathbf{u}, \mathbf{v}) \in \mathbb{R}^4$. We can reduce to a 4-dimensional space in this way because each player's strategy space only has 2 degrees of freedom: $x_3 = 1 - x_1 - x_2$. This coordinate transform projects the the whole strategy space, a 4-dimensional simplex in 6-dimensional space, down to 4-dimensional space. After performing some algebra, we obtain

$$\dot{u_i} = \frac{\sum_{j=1}^{m}\tilde{A}_{ij}e^{v_j} + \tilde{A}_{i0}}{1 + \sum_{j=1}^{m}e^{v_j}}$$

where $\tilde{A}_{ij} = A_{ij} - A_{0j}$. This transformation expresses the replicator equation as a linear bipartite system of the form $\dot{u} = f(v)$ and $\dot{v} = g(u)$. Via further manipulations, we can express the system as a Hamiltonian system

$$\dot{u} = P\nabla_v H$$

$$\dot{v} = -P^T\nabla_u H$$

Where

$$H = \sum_{1}^{n} p_i u_i - log\left(1 + \sum_{1}^{n} e^{u_i}\right) + \sum_{1}^{n} q_i v_i - log\left(1 + \sum_{1}^{n} e^{v_i}\right)$$

for interior equilibrium (p, q) and the matrix P is a skew-symmetric matrix where $P_{ij} = \tilde{A}_{i0} - \tilde{A}_{ij} = A_{i0} + A_{0j} - A_{00} - A_{ij}$. Serendipitously, this allows us to easily extend to a case we wish to analyze later; Rock Paper Scissors where player 1 wins ϵ in a tie while player 2 loses ϵ. In that case,

$$P = \begin{pmatrix} 0 & 0 & 2\epsilon & 3+\epsilon \\ 0 & 0 & -3+\epsilon & 2\epsilon \\ -2\epsilon & 3-\epsilon & 0 & 0 \\ -3-\epsilon & -2\epsilon & 0 & 0 \end{pmatrix}$$

Consequences of Hamiltonian Dynamics and Chaos

As is immediately apparent from RPS's Hamiltonian form as a linear bipartite system, Liouville's theorem applies. Thus, the volume in the RPS state-space (a Euclidean volume in $\mathbb{R}^4$under the coordinate transformation above, a less concrete volume form otherwise) is invariant. This implies that the Nash Equilibrium at $\vec{\mathbf{x}} = \vec{\mathbf{y}} = \{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\}$is not asymptotically stable. Thus, when players begin with any strategy pair which is not the Nash equilibrium, they never move to the Nash equilibrium. However, the equilibrium still tells us something about the game's payoffs. It can be shown that the time averages of the payoffs for each player in any orbit are the same as if both players played the equilibrium. (Hofbauer 1998)

According to the replicator equation players move along a 4-dimensional toroidal orbit, which, projected onto a shared 2-dimensional simplex, is pictured below

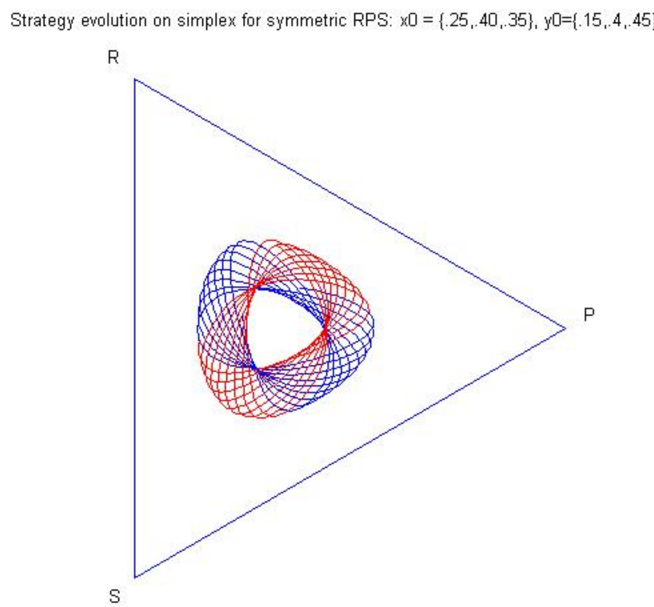

However, when $\epsilon \neq 0$, many of these torii can collapse into chaotic orbits. According to KAM theory, many toroidal orbits remain , but for higher and higher ϵ more orbits collapse. In fact, chaotic orbits finely interweave with the toroidal paths such that the orbits are dense in each other (Sato et. al, 2002). This chaotic behavior is shown in a hyperplanar Poincare section of the 4-dimensional simplex with $\epsilon = 0.25$ over many initial conditions from Sato et. al.

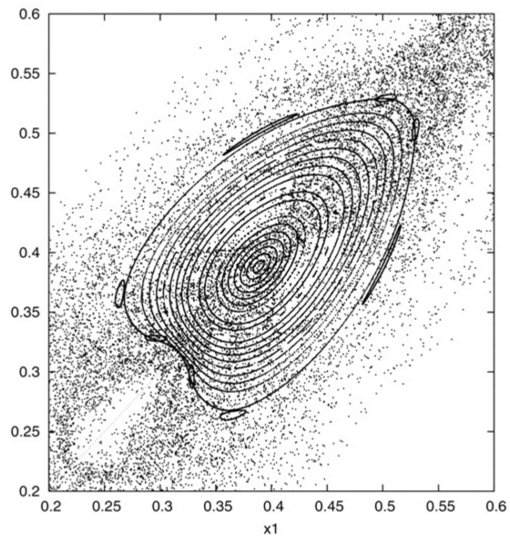

Here, the circles rare sections of toroidal Hamiltonian orbits in 4-space while the "snow" around and between them is where chaotic orbits cross through the Poincare section. Indices do not correspond directly to frequencies due to variable transformations which project the strategy space into 4 dimensions.

When $\epsilon = 0$, the integrable nature of the motion implies that all Lyapunov exponents are precisely 0. However, when $\epsilon > 0$, Hamiltonian mechanics requires only that $\lambda_2 = \lambda_3 = 0$ and that $\lambda_1 = -\lambda_4$. Numerical computation (Sato et. al) shows that $\lambda_1$is positive for many initial conditions, which shows these orbits are in fact chaotic.

the consequences of chaotic orbits seriously affect predictions in game theoretical models. While, as stated above, the average expected payoffs for both players are the same as if both played the Nash equilibrium, the divergences from the equilibrium payoff along chaotic orbits are larger than the toroidal orbits, and so a risk-averse player would prefer to avoid chaotic orbits. Furthermore, the essential unpredictability of chaotic orbits means that no player can use a prediction algorithm to gain an advantage over another player in the game as well. Thus, chaos in such a simple game demands we exercise great caution as to when a lack of agency may undercut entirely the predictions of game theory.

Works Cited

Hofbauer, Josef, and Karl Sigmund. Evolutionary Games and Population Dynamics. Cambridge: Cambridge UP, 1998. Web. J. Math. Biol. (1996) 34: 67. (Hofbauer 1998)

Hofbauer.J. Math. Biol. (1996) 34: 675—688 Evolutionary Dynamics for Bimatrix Games: A Hamiltonian System? (n.d.): n. pag. (Hofbauer 1996)

Web. Sato, Y., E. Akiyama, and J. D. Farmer. "Chaos in Learning a Simple Two-person Game." Proceedings of the National Academy of Sciences 99.7 (2002): 4748-751. Web. (Sato et. al.)

Tadelis, Steve. Game Theory: An Introduction. Princeton: Princeton UP, 2013. Print.