Evolutionary Game Theory – From Cooperation to Topological Phases

Marc S. Klinger

May 2021

Abstract

Evolutionary Game Theory extends the tools of Rational Game Theory to the study of (typically infinite) Repeated Games in which assumptions about the hyperrationality of players are relaxed to accommodate "Bounded Rationality". When Evolutionary Games are placed on a lattice they exhibit the formation of complex patterns. Of the most interesting of these are the emergence of cooperation between players in spite of the manifestly non-cooperative paradigm of Game Theory, and the emergence of topological phases despite the manifestly non-physical nature of Game Theory. In this essay, we will provide an introduction to the machinery of Rational and Evolutionary Game Theory, and observe how these complex and even potentially counter-intuitive phenomena can emerge therein from a relatively simple set of strategic interactions.

1 The Game Theory Dilemma

When many people think of Game Theory they immediately think of one game in particular – the Prisoner's Dilemma. This is not without good reason; the Prisoner's Dilemma serves to demonstrate many of the essential features of Game Theory, all wrapped in a compelling storyline that is surprisingly generalizable despite its simplicity. Depending on who is telling it the characters and setting may change, but for simplicity let us tell the story as it was originally conceived:

A pair of individuals are captured by the police under suspicion of having committed a crime together. The police do not have enough evidence to convict either criminal completely unless one or both of them provides additional evidence. So, they place the two accomplices in solitary without the ability to communicate with one another and approach each individually with a bargain: "If you are willing to testify against your accomplice you will receive a commuted sentence and he will receive the full sentence." Of course, there is a catch: if both criminals agree to testify against each other the police will have sufficient evidence to convict both in which case they will share the full sentence equally. Each prisoner must therefore choose between two options: Cooperating with their accomplice by staying silent, or Defecting by agreeing to testify.

So, what should the prisoners do? Let's analyze the situation more systematically. Each prisoner can choose between two different actions: Cooperate (C) or Defect (D). Thus, there are four possible states in this game: (C, C) in which each of the prisoners cooperate, (D, D) in which each of the prisoners defect and $(C, D), (D, C)$ in which one of the prisoners cooperates while the other defects. To each of these states we assign a pair of values corresponding to the "payoff" received by each of the prisoners, in this case equal to the length of their prison sentence in each state. Denote this payoff $\pi(S) = (\pi_1(S), \pi_2(S))$, with $\pi_i(S)$ the payoff to prisoner i. The Prisoner's Dilemma is therefore represented by the payoff bi-matrix seen in Table 1.

When each prisoner is making his decision about whether or not to cooperate they have no contact at all with their accomplice. The objective of each prisoner is thus to minimize his prison sentence taking as granted the fact that their comrade will do the same. This leads the pair into what is called a "Nash Equilibrium". In the prisoner's dilemma the Nash Equilibrium is especially stark. The problem is symmetric, so let us consider the problem from the perspective of prisoner number one:

- **Case 1: Prisoner 2 Defects**: Then, prisoner number one need only consider the reduced payoff function $\pi_1(a_1, D)$ which is equal to P_4 if he cooperates and P_3 if he defects. Since $P_3 < P_4$, in this case prisoner number one should **always defect**.

- **Case 2: Prisoner 2 Cooperates**: Then, prisoner number one need only consider the reduced payoff function $\pi_1(a_1, C)$ which is equal to P_2 if he cooperates and P_1 if he defects. Since $P_1 < P_2$, in this case prisoner number one should **always defect**.

	Prisoner 2 Cooperates	**Prisoner 2 Defects**
Prisoner 1 Cooperates	$\pi(C, C) = (P_2, P_2)$	$\pi(C, D) = (P_4, P_1)$
Prisoner 1 Defects	$\pi(D, C) = (P_1, P_4)$	$\pi(D, D) = (P_3, P_3)$

Table 1: Payoff bi-matrix for the generalized Prisoner's Dilemma. Here $P_1 < P_2 < P_3 < P_4$.

Thus, we arrive at the startling conclusion: No matter what the prisoner believes his accomplice will do, he should always defect. This holds for both prisoners, and thus the unique Nash Equilibrium of the Prisoner's Dilemma is for *both* Prisoners to defect. Under the assumptions of Rational Game Theory this conclusion is true for any finite values of P_i provided they satisfy the stated inequality. So, even if $P_1 = 1, P_2 = 2, P_3 = 50$, and $P_4 = 100$, the prisoners will *still* choose to defect, despite the fact that had they both simply cooperated with one another they could have reduced their sentences from 50 years each to 2 years each!

I like to refer to this feature of Rational Game Theory – the fact that the Nash Equilibrium is inefficient from the perspective of aggregate social welfare – as the Game Theory Dilemma. The Game Theory Dilemma poses many real world conflicts in terms of the inability to compel individuals to cooperate. For one Political Scientist, the inability to ensure cooperation was more than just an inconvenience – it was potentially catastrophic!

1.1 The Emergence of Cooperation

When Political Scientist Robert Axelrod looked at the Prisoner's Dilemma he didn't see two criminals apprehended by the police asked to turn on each other, he saw two world superpowers with their fingers on the nuclear button [3]. He feared that the lack of an intrinsic mechanism for enforcing cooperation could lead to a scenario of mutually assured detonation. The thing was, to wit, such a detonation had not occurred. This got Axelrod to thinking – what was there stopping either side from pulling the trigger? His realization can be phrased in two different ways: Two competing entities can be compelled into cooperation with one another without an external agency to uphold that cooperation provided

1. The cost of not cooperating is exorbitantly large.
2. The two entities engage with one another exorbitantly many times.

Put simply, not cooperating in the nuclear detonation example would mean bringing about the end of all days – a cost that is effectively infinite in size. Or, equivalently, even if the cost of defecting in a single interaction is finite, if the two players will interact many many times this cost will pile up to become effectively infinite. Thus, the two factors that inspire cooperation are essentially identical: players would rather cooperate than endure unmitigated losses. Axelrod was therefore lead to the gratifying conclusion: If a single game – called a "Stage Game" – is played an infinite number of times it can become advantageous, even from the perspective of game theory, to cooperate over large time horizons, even if the Nash Equilibrium of the Stage Game is inherently non-cooperative! In fact, cooperation can even be advantegous if the game is repeated only a finite number of times.

To test his hypotheses and explore further the question of how, when and what kinds of cooperation can emerge in repeated stage games, Axelrod organized a tournament [2]. He asked his colleagues to devise strategy profiles that would choose to either cooperate or defect in a given round of some finite repeated Prisoner's Dilemma game based on the results of all previous rounds. He then pit these strategies against each other randomly and recorded the results. The winner was a strategy fittingly titled "TIT-FOR-TAT". The TIT-FOR-TAT strategy uses the following simple strategy profile: In round 1 the player cooperates. Thereafter, the player simply mirrors the strategy of her opponent. In this respect the player is willing to cooperate only insofar as her opponent is willing to as well.

1.2 The Birth of Evolutionary Game Theory

Axelrod's tournament gave experimental verification to the idea that cooperation could emerge naturally as a preferred alternative to defection in Game Theoretical contexts. It also inspired a follow-up question – in a population of individuals simultaneously playing a series of repeated games with one another, how will the distribution of different strategies change over time? Or, even more ambitiously, if those agents are allowed to disperse themselves over a spatial lattice in which they are more or less likely to interact with other players who play different strategies, how will the distribution of different strategies change over space? These are the primary questions that concern the study of *Evolutionary Game Theory*.

As the name implies, evolutionary game theory borrows heavily from biology and the theory of evolution for both its analytical tools and interpretations. However, due to the large degree of abstraction possessed by Game Theory, evolutionary games have the generality to apply to almost any system in which a large number of entities maximize their individual welfare across a large set of possible states. The mathematical structure of Evolutionary Game Theory is exceptionally rich, building on ideas from probability theory, dynamical systems theory, optimization theory, and information theory to name just a few. Thus, evolutionary game theory offers the tantalizing opportunity to work in an area that is vibrant both in terms of its pure theory and its applications.

What's more, Evolutionary games are capable of producing very intricate behavior and complex patterns even if the game dynamics or evolutionary dynamics are taken to be profoundly simple. As such, evolutionary game theory proves also to be a prime area for studying the emergence of large scale patterns very much in the spirit of non-equilibrium dynamical systems. The kinds of emergent properties observed in evolutionary games can be brought into stark relation with emergence occurring in physical systems, for example in the observation of long range order, power law behavior, and symmetry breaking. Remarkably, evolutionary games with a distinctly non-physical character are even capable of reproducing physically motivated emergent behavior.

In this paper I will provide an introduction to the mathematical machinery of Rational and Evolutionary Game Theory that concisely examines its most crucial ideas. Then, I will discuss a beautiful example which demonstrates how the physical concept of topologically protected phases can be realized in an evolutionary game of Rock-Paper-Scissors placed on a simple lattice. Working through this result serves well to illustrate the general theory built in the first section, as well as the power that game theory holds as a tool for understanding physical systems.

2 Rational Game Theory

Before we begin our study of Evolutionary Game Theory, let us provide a lightning review of Rational Game Theory – or the theory of single games. As the reader will see, this involves simply generalizing all of the features we observed in the Prisoner's Dilemma Game. The discussion in this and the following sections is inspired by the wonderful review [11].

A Game, $G = (\mathcal{P}, S_i, \pi)$, is defined by three items:

1. A set of players, represented by the index set: $\mathcal{P} = \{1, ..., N\}$
2. A set of feasible actions for each player: $S_i = \{a_i^1, ..., a_i^{N_i}\}$ for each $i \in \mathcal{P}$
3. A payoff function specifying the value received by each player given a combined strategy profile $a = (a_1, ..., a_N) \in \bigtimes_{i\in\mathcal{P}} S_i$: $\pi : \bigtimes_{i\in\mathcal{P}} S_i \to \mathbb{R}^N$.

In the most general case each player is allowed to choose a mixed strategy which corresponds to a probability distribution over his action set. Thus, the states of the game are labelled by an N-tuple of probability distributions: $P = (P_1, ..., P_N)$ with $P_i(a_i)$ giving the probability that player i will play the strategy $a_i \in S_i$. The Game Utility function is defined to be the expected payoff received by each player as a function of the state P:[1]

$$U_i(P) = \sum_{a_j \in S_j \forall j \in \mathcal{P}} \prod_{k \in \mathcal{P}} P_k(a_1, ..., a_N)\pi_i(a_1, ..., a_N) \tag{1}$$

Let us introduce the notation $P_{-i} = (P_1, ..., P_{i-1}, P_{i+1}, ..., P_N)$. The objective of each individual player $i \in \mathcal{P}$ is to choose the P_i that optimizes $U_i(P_i, P_{-i})$ viewed as a function of some fixed opponent strategy profile P_{-i}. That is, given an opponent strategy profile P_{-i} agent i will play the so-called best response strategy defined by:

$$BR_i(P_{-i}) = \arg\max_{P_i} U_i(P_i, P_{-i}) \tag{2}$$

The fundamental concept in Rational Game Theory is that of Nash Equilibrium.

Definition 2.1. Nash Equilibrium:[2] A Nash Equilibrium is realized if every agent is best responding to every other agent's best response. A strategy profile P^* is a Nash Equilibrium if, for all $P \neq P^*$ and each player, i:

$$U_i(P_i^*, P_{-i}^*) \geq U_i(P_i, P_{-i}^*) \tag{3}$$

Symbolically: $P_i^* = BR_i(P_{-i}^*)$.

It's not difficult to see that there is a very stark analogy between Game Theory and the Physics of Many-Body Systems. Each includes a number of individual agents seeking to extremize an objective function. The major difference between Many-Body Physics and Game Theory is that in Game Theory each individual seeks to extremize their own unique objective function, while in Many-Body Physics there is a single collective objective function to be extremize – e.g. the Free Energy. It turns out, however, that there does exist a subset of games which can be reduced to the extremization over a single energy landscape:

Definition 2.2. Potential Games: A Game is called a potential game if there exists a function $\mathcal{F} : \mathcal{S} \to \mathbb{R}^N$ mapping the set of all collective strategy profiles into utilities with the property that:

$$U_i(P_i', P_{-i}) - U_i(P_i, P_{-i}) = \mathcal{F}(P_i', P_{-i}) - \mathcal{F}(P_i, P_{-i}) \ \forall i \tag{4}$$

The Nash Equilibrium of such a game therefore corresponds to extremizing the single function $\mathcal{F}$. If we regard the P_i as configurations for a set of fields over the space of pure strategies we can map this directly into a problem in Statistical Mechanics.

[1]Notice $\prod_{k \in \mathcal{P}} P_k(a_1, .., a_N)$, which is the product of the marginal probability distributions for each player, can be regarded as the joint probability distribution over the joint set of strategies since individual agents choose strategies independently.

[2]John Nash of *A Beautiful Mind* fame proved the seminal theorem: Every normal form game with a finite number of players and a finite number of pure strategies possesses at least one Nash Equilibrium.

3 Evolutionary Game Theory

When studying Many-Body Systems one makes use of a hierarchy of scales – choosing the level of description that is most advantageous for the analysis of a particular system. For statistical physical systems, the most basic distinction that can be made is between the Thermodynamic Perspective which takes into account the mean properties of a system, and the Statistical Field Theoretical Perspective which allows for distributions to be formed around these mean values. Thermodynamics is useful in the limit where the number of interacting agents becomes very large, but, by successive averaging, does away with important information at finer scales. By contrast the Field Theoretical approach gives a more fine picture of the system, and – quite crucially – allows for the analysis of spatial variation. However, this comes at the cost of a marked increase in the complexity of computations.

In the study of Evolutionary Game Theory there are two different levels of analysis which map perfectly onto the Macroscopic, Thermodynamic Picture and the Microscopic Field Theoretic Picture. These are the Population Dynamics and the Agent Based Dynamics, respectively.

- **Population Dynamics** track only the frequency with which each pure strategy in a stage game is played throughout a very large population of players. The strategy profile of the population changes according to a set of global dynamical rules.

- **Agent Based Dynamics** track the individual strategy profiles of each player, distributed throughout a spatial region usually taken to be some kind of lattice. Strategy profiles change at the level of the individual player according to a locally defined dynamical rule.

3.1 Population Dynamics

Population Dynamics begin with a Stage Game – some one shot Rational Game for two players: $G = (\{1, 2\}, \{S_1, S_2\}, \pi)$ – which will be played T times. We take the game to be either symmetric in which case $S_1 = S_2 = S$ and $\pi_1(a_1, a_2) = \pi_2(a_2, a_1) = \pi(a_1, a_2)$, or anti-symmetric in which case $\pi_1(a_1, a_2) = -\pi_2(a_2, a_1)$.

If the game is symmetric we suppose that there is a population, $\mathcal{P}$, of $N \to \infty$ identical players who can choose from the set of pure strategies S. In each round every player is randomly matched with another player with equal probability.

Suppose that N_i is the number of players playing strategy $a^i \in S$. Then, we define:

$$\rho(a_i) = \frac{N_i}{N} \tag{5}$$

Which is nothing but the population frequency for the strategy a_i. Since each agent is randomly assigned an opponent from the rest of the population with equal probability, the expected utility for playing a given strategy is given by:

$$U(p, \rho) = \sum_{a_i \in S, a_j \in S} p(a_i)\rho(a_j)\pi(a_i, a_j) \tag{6}$$

This is analogous to each individual player coming up against another player with a mixed strategy equivalent to the probability profile of the population e.g. the Population Representative Agent. We refer to $\rho(a)$ as the state of the system.

It is very straightforward to extend the concept of Nash Equilibirum to Population Dynamics. A state $\rho^*(a)$ is a Population Nash Equilibrium if it outperforms any other state, ρ, in the sense:

$$U(\rho^*, \rho^*) \geq U(\rho, \rho^*) \tag{7}$$

The first place where Evolutionary Game theory begins to differ from Rational Game Theory is in the evaluation of the Evolutionary Stability of a state. Similarly to how one would analyze the stability of an equilibrium configuration, we imagine perturbing a population slightly by introducing a small number of "mutants" who play some new strategy $p(a)$ while the remainder of the population continues to play the incumbent strategy $\rho^*(a)$. We then ask whether the incumbent strategy remains superior to the mutant strategy. If a proportion ϵ of the total population plays the mutant strategy, the population state is given by $\rho(a) = (1-\epsilon)\rho^*(a) + \epsilon p(a)$. We therefore conclude that $\rho^*(a)$ is evolutionarily stable with critical size ϵ_c if for all $0 < \epsilon < \epsilon_c$ and any p:

$$U(\rho^*, \rho) \geq U(p, \rho) \tag{8}$$

If the game is antisymmetric we imagine that there are two populations, $\mathcal{P}_1$ and $\mathcal{P}_2$, each of asymptotically large sizes, $N_1 \to \infty$ and $N_2 \to \infty$, in which all of the players from $\mathcal{P}_1$ choose from S_1 and all of the players from $\mathcal{P}_2$ choose from S_2. In each period every player from $\mathcal{P}_1$ is randomly matched with a player from $\mathcal{P}_2$.

The analysis of this situation is identical to the symmetric case, except that the population state is specified by two probability distributions $\rho_1(a)$ for $a \in S_1$ and $\rho_2(b)$ for $b \in S_2$ corresponding to the frequency with which strategies are played in each population. The utility gained by an agent in $\mathcal{P}_1$ who plays a mixed strategy p_1 is given by:

$$U_1(p_1, \rho_2) = \sum_{a \in S_1, b \in S_2} p_1(a)\rho_2(b)\pi_1(a, b) \tag{9}$$

and similiarly for an agent in $\mathcal{P}_2$ playing a mixed strategy p_2:

$$U_2(p_2, \rho_1) = \sum_{a \in S_1, b \in S_2} \rho_1(a)p_2(b)\pi_2(a, b) \tag{10}$$

Population Dynamics are completed by defining a dynamical rule for how the population state changes over time. This rule will take the form:

$$\frac{d\rho(a;t)}{dt} = \mathcal{R}[\rho(a;t)] \tag{11}$$

where $\mathcal{R}$ is some, potentially nonlinear operator acting on ρ.

A particularly popular rule for updating the mean population is the replicator dynamics which are inspired by Biology. The Replicator Dynamics take the form:

$$\frac{d\rho(a;t)}{dt} = (f(a) - \langle f \rangle)\rho(a;t) \tag{12}$$

Here $f(a) = U(a, \rho)$ is the utility earned by the pure strategy a in background of the state ρ. This is interpreted as the "fitness" of the state a in analogy with biological evolution. $\langle f \rangle = \sum_{a \in S} \rho(a) f(a)$ is the mean fitness. Hence, the interpretation of the replicator equation is that strategies grow or shrink in proportion to how much more or less successful they are as compared to the average fitness.

3.2 Agent Based Dynamics

In agent based dynamics we regard individual players as the fundamental degree of freedom. Let $G = (\{1,2\}, \{S_1, S_2\}, \pi)$ denote a stage game, and, for simplicity, we take it to be symmetric so that $S_1 = S_2 = S$, and $\pi_1(a,b) = \pi_2(b,a) = \pi(a,b)$. Agent based dynamics take place on a lattice, L, with sites labeled by $x \in L$. At each site there is an agent who plays a (potentially mixed) strategy on the set S. The properties of the lattice are defined by specifying the neighborhood of each point $x \in L$ corresponding to the set of players that the player at x plays with. We denote this set as $\Omega(x) \subset L$. Let p_x denote the strategy profile of the player at the lattice position x. Then, the utility obtained by the player at x is simply:

$$U(x) = \sum_{y \in \Omega(x)} \sum_{a_x \in S, a_y \in S} p_x(a_x) p_y(a_y) \pi(a_x, a_y) \tag{13}$$

The total utility of the lattice system (up to double counting) is simply the sum over each lattice position:

$$\mathcal{H} = \sum_{x \in L} U(x) = \sum_{x \in L} \sum_{y \in \Omega(x)} \sum_{a_x \in S, a_y \in S} p_x(a_x) p_y(a_y) \pi(a_x, a_y) \tag{14}$$

which, as the notation suggests, can be regarded as the Hamiltonian of a Lattice Statistical Mechanics system. If the Stage Game is a Potential Game this analogy is manifest and one can use the tools of Statistical Mechanics to analyze the agent based game.

Agent Based Dynamics are derived from site specific dynamical rules. In particular, let $p_x(a;t)$ denote the strategy profile at the lattice location x in the period t of the repeated game. Then, an arbitrary Dynamical Rule is a stochastic map in the form of a conditional probability distribution which can depend on information from the neighborhood of the point x:

$$p_x(a; t+1) = \sum_{a' \in S} \omega_x \left(a \mid a', \{p_y(a;t), U(y)\}_{y \in \Omega(x)}\right) p_x(a'; t) \tag{15}$$

Even the most ostensibly simplistic forms of Agent Based Evolutionary Games are capable of producing wildly intricate complex dynamics. For example, Agent Based Dynamics with deterministic dynamical rules are equivalent to Cellular Automata. In [13], Steven Wolfram classifies the spatio-temporal patterns occurring in Cellular Automata into four classes. The most interesting of these classes is Class Four which he characterizes by the formation of highly complex patterns in space and time included so-called "nested" or repeating elements that exhibit power law behavior. Killingback and Doebali have further demonstrated that Game Theoretic Cellular Automata can exhibit long range order in the form of spatial and temporal correlations as measured, for example, by the mutual information between strategy profiles of players in fixed positions at separate times or separate positions at fixed times [9].

We can recognize all of these features as the hallmarks of emergence. Indeed, the kinds of interesting behaviors that can be produced by studying Evolutionary Games seems almost infinite. Recently, this fact has been given further credence by the rather exciting discovery that a particularly simple class of Evolutionary Games is capable of reproducing behavior previously believed to be restricted to physical systems.

4 Rock-Paper-Scissors says Topological Phases

It is likely that anyone reading this is familiar with the game of Rock-Paper-Scissors. In the language of Game Theory this is a two player game with three strategies for each player: $S = \{R, P, S\}$. Letting $+1$ denote victory, -1 denote loss, and 0 denote draw, the payoff matrix for this game is given by:

$$\begin{array}{lll} \pi(R,R) = 0 & \pi(R,P) = -1 & \pi(R,S) = 1 \\ \pi(P,R) = 1 & \pi(P,P) = 0 & \pi(P,S) = -1 \\ \pi(S,R) = -1 & \pi(S,P) = 1 & \pi(S,S) = 0 \end{array}$$

We begin by evaluating the mean field properties of this game in a large population of identical players. Let $\rho(a;t)$ denote the proportion of players playing the strategy $a \in \{R, P, S\}$ at the time t. Further, let us assume that the population dynamics of this game are governed by the Replicator Dynamics 12. Rock-Paper-Scissors has the property that it is a zero-sum game – meaning that whenever one player wins another player loses. Hence, the total fitness, and by extension the mean fitness must be equal to zero. The replicator equation therefore takes the following form:

$$\frac{d\rho(a;t)}{dt} = \left(\sum_{a' \in S} \rho(a')\pi(a,a') \right) \rho(a;t) \tag{16}$$

This is, in fact, nothing by the Lotka-Volterra equation for a three species antisymmetric predator-prey model!

The solution to this equation is therefore a set of three off-set oscillating functions which can be seen in figure 1. This solution is totally intuitive. Begin at the time $t = 0$:

1. Scissors is the most prevalent strategy. This means that Rock, which beats scissors becomes the most advantageous strategy and therefore grows the fastest. We can think of this in terms of the evolutionary stability argument 8 – predominantly scissors is not an evolutionarily stable strategy because an influx of mutant players who specialize in the rock strategy will always outperform the scissors incumbants.

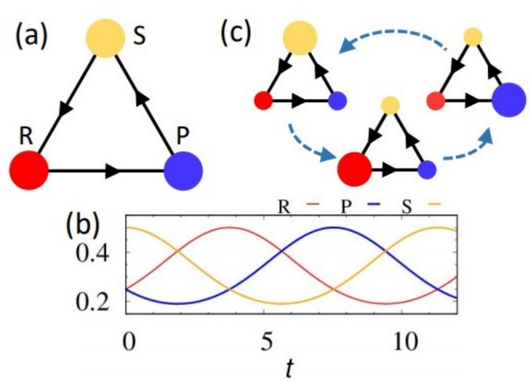

Figure 1: The Mean Field Behavior of the Rock Paper Scissors evolutionary game exhibits cyclotron behavior [14].

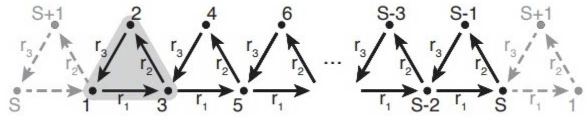

Figure 2: 1D lattice of coupled RPS cycles [10]

2. As Rock becomes a more prevalent strategy, Scissors begins to dwindle and eventually starts to become a less prevalent strategy. As a result, Paper, which beats Rock and loses to Scissors, becomes incrementally more advantageous and therefore lags the growth of Rock and the shrinking of Scissors.

3. Eventually, Rock saturates and becomes the most prevalent strategy at which point we can return to step 1 with Rock taking the place of Scissors, Paper taking the place of Rock, and Scissors taking the place of Paper.

In [14], Yoshida et. al. recognized that this oscillatory motion mimics the cyclotron motion of fermions between Landau Levels in systems exhibiting the Integer Quantum Hall Effect. As a result, they hypothesize and later demonstrate that the associated lattice model of the Rock-Paper-Scissors game demonstrates the emergence of Chiral-Edge Modes corresponding to the propagation of population density waves around the edge of the lattice.

Following the analysis of [10] we can consider the lattice system for a Rock-Paper-Scissors stage game in earnest and in doing so uncover a fascinating analogy between agent based dynamics and quantum mechanics. We imagine a one dimensional lattice of coupled Rock-Paper-Scissors cycles as drawn in figure 2. Then, we let the mass at each lattice point, m_x, evolve according to an extended RPS game with payoff matrix:

$$\pi_{xy} = \begin{pmatrix} 0 & r_3 & -r_1 & 0 & 0 & ... & 0 & 0 & 0 \\ -r_3 & 0 & r_2 & 0 & 0 & ... & 0 & 0 & 0 \\ r_1 & -r_2 & 0 & r_3 & -r_1 & ... & 0 & 0 & 0 \\ 0 & 0 & -r_3 & 0 & r_2 & ... & 0 & 0 & 0 \\ 0 & 0 & r_1 & -r_2 & 0 & ... & 0 & 0 & 0 \\ & & & & . & & & & \\ & & & & . & & & & \\ & & & & . & & & & \\ 0 & 0 & 0 & 0 & 0 & ... & 0 & r_3 & -r_1 \\ 0 & 0 & 0 & 0 & 0 & ... & -r_3 & 0 & r_2 \\ 0 & 0 & 0 & 0 & 0 & ... & r_1 & -r_2 & 0 \end{pmatrix} \tag{17}$$

Here r_1, r_2, r_3 control the rate at which mass is transferred between lattice sites. The lattice dynamics of this game can therefore be regarded as the Lotka-Volterra equation:

$$\frac{dm_x}{dt} = \left(\sum_{y \in L} m_y \pi_{xy} \right) m_x \tag{18}$$

This approach has a somewhat different character from agent based dynamics discussed, for instance, in 15. Rather than considering the strategy profile of each agent based at a lattice site

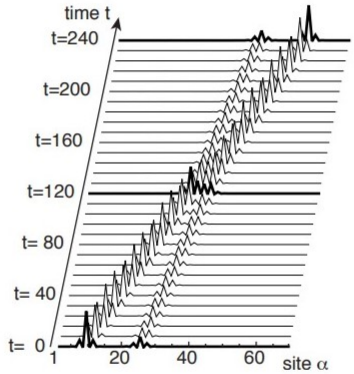

Figure 3: One solution to 18 is a solitary wave which transports mass up the RPS-lattice [10].

as changing over time, we take the strategy at each site to be fixed and allow the profitability of each site to change as the climate of the game changes. In this regard 18 acts like a kind of hybrid between the agent based, lattice evolutionary game theory and the population based mean field theory.

18 is the natural extension of 16 to a simple lattice model, and we can very easily intuit some of the most pertinent characteristics of the solution. Indeed, since the lattice is nothing but a series of coupled RPS-cycles which each elicit the cyclotron dynamics discussed above, the total RPS-Lattice will admit traveling wave solutions with the general profile of figure 1. By simulating the equation 18 Knebel et. al. were able to observe these solitary waves quite emphatically as seen in figure 3. These are related to the Chiral-Edge modes predicted by Yoshida et. al. in their paper.

Looking at equation 18 one cannot help but be reminded of Schrodinger's equation for a chain of non-linearly coupled oscillators. The mass at each lattice site is coupled through a non-linear interaction term $\sim m_x m_y$ which is facilitated by the anti-symmetric matrix π_{xy}. We can make this analogy manifest by defining the RPS-Hamiltonian as an analytic continuation of the RPS-Payoff Matrix: $\mathcal{H}_{xy} = i\pi_{xy}$. This is a standard trick that hearkens to mind the method for relating standard problems in Statistical Mechanics (of which we have argued Game Theoretical problems are most alike) with standard problems from Quantum Mechanics. By the anti-symmetry of π, $\mathcal{H}$ is a hermitian matrix: $\mathcal{H}^\dagger = \mathcal{H}$, and we can rewrite the Lotka-Volterra equation:

$$i\frac{dm_x}{dt} = \left(\sum_{y\in L} m_y \mathcal{H}_{xy}\right) m_x \tag{19}$$

This is precisely the Schrodinger equation (with $\hbar = 1$) for a $1D$-chain of non-linearly coupled oscillators.

Recognizing the repeating cell structure of 17 which is inherited by $\mathcal{H}$, we can use Fourier Transform techniques to simplify the analysis of 19 to the solution of an eigenvalue problem in the Brillouin Zone of the RPS-Lattice. In particular, the topological band structure of the RPS lattice can be determined by solving the eigenvalue problem:

$$\tilde{\mathcal{H}}(k)\tilde{u}(k) = \lambda(k)\tilde{u}(k) \tag{20}$$

Where the cell Hamiltonian takes the form:

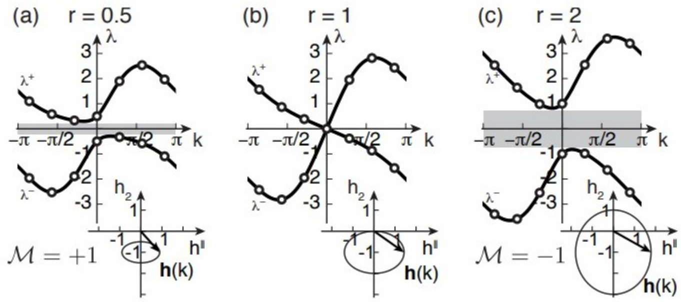

Figure 4: The band structure of RPS lattice depends on the control parameter r [10].

$$\tilde{\mathcal{H}}(k) = \begin{pmatrix} 2r_1 \sin(k) & -r_2 \sin(k) + i(r_3 - r_2 \cos(k)) \\ -r_2 \sin(k) - i(r_3 - r_2 \cos(k)) & 0 \end{pmatrix} \tag{21}$$

Investigating the symmetry properties of the Hamiltonian 21, we recognize that it breaks both Time Reversal and Chiral Symmetry, but satisfies the so-called "Particle-Hole" Symmetry:[3] $\mathcal{C}\tilde{\mathcal{H}}(k)\mathcal{C}^{-1} = -\tilde{\mathcal{H}}(-k)$. According to the so-called "Ten-Fold Way" scheme for classifying topological fermionic systems this implies that the RPS-lattice can be understood as a gapped free-fermion system of symmetry class D [14]. Remarkably, this means that the RPS-lattice falls into the same symmetry class as the Bogoliubov-de Gennes Hamiltonian which provides a mean field description of superconductivity in the majorana fermion basis.

Normalizing $r_1 = 1$ to set the mass scale in the problem, one finds that the properties of the system are governed by the control parameter $r = \frac{r_2}{r_3}$ referred to as the "skewness". Solving the eigenvalue problem 20 directly, one finds two bands seen in figure 4. Depending on the value of the skewness these bands fall into two topologically distinct phases each of which exhibit a gapped band structure. We can confirm that the phases are topologically distinct because they have different values of the homotopic invariant $\mathcal{M} = \text{Pfaf}(\pi) = \text{sgn}(1 - r^2)$. One cannot continuously transition between these phases without closing the spectral gap at $r = 1$. Thus, we arrive at the remarkable conclusion that, starting from a profoundly simple, manifestly non-physical Evolutionary Game Theory we are able to observe the emergence of Topological Order!

5 Conclusion and Looking Forward

Game Theory formalizes the study of strategic interactions between autonomous decision making agents. Because of its generality, it provides a powerful tool for modeling a wide variety of different interactions with applications across the physical, biological, and social sciences and in society. Conceptually, the theory of Games bears a strong resemblance to the theory of many-body interacting particle systems in physics – an analogy that can made precise in certain special cases.

One of the main conclusions from Rational Game Theory is that it is often optimal – in the game theoretical sense – for players to choose action profiles that lead to outcomes that are sub-optimal from the perspective of social welfare. This has lead to the conclusion that Game Theory prevents, or at least strongly discourages cooperation. Yet, this hypothesis is often proven

[3]Here $\mathcal{C}$ is the charge conjugation operator.

false in Social Psychological and Behavioral Economic experiments where individuals are observed as engaging in cooperation, especially when games are played multiple times or between known adversaries [6] [1] [5].

These observations, among others, inspired the founding of a new branch of Game Theory called Evolutionary Game Theory. Evolutionary Game Theory studies a large population of players playing a single game amongst each other many times. It is primarily concerned with how the distribution and effectiveness of different strategies change over time and throughout the population as players learn from previous experiences playing the game.

The most fascinating and useful feature of Evolutionary Game Theory is the sense in which it abstracts away from all of the nitty gritty details of a scenario to produce minimal models from which an impressive array of interesting and complex behavior can emerge. The emergence of topological order in the evolutionary game theory of Rock-Paper-Scissors is a perfect example of this. The theory of topological phases in Physics involves a host of different considerations from Homotopy to Quantum Field theory, yet we were able to reproduce many of the most important hallmarks of topological order from dynamics as simple as those found in a children's game.

It is interesting to note that many of the fundamental features of Evolutionary Games can be recast as the dynamics of information and studied form the perspective of information theory [7] [8]. This includes retooling our understanding of biological evolution in these terms, such as the reformulation of the Genetic Code in the language of Rate-Distortion and Noisy Channels [12], and the study of replicator dynamics as open Markov-Processes [4]. Although there was not enough room to discuss these ideas in this essay, they represent an arena which is greatly appealing for future work. It is particularly gratifying how this field highlights the sense in which biology and physics can be married through the concept of information.

Perhaps the takeaway from all of this is that the phenomena we observe in nature, even those that are apparently the most intractable and complicated, can truly be traced back to a set of comparatively simple rules which, when placed in the right context, allow for the emergence of more intricate modes. Or, that even if the mechanisms underlying complex phenomena are complex, one can employ clever minimal descriptions and still make great progress in understanding their essential features. Either way, it is remarkable to reflect on how far we have come without ever leaving the confines of Game Theory: – from the Prisoner's Dilemma all the way to the Emergence of Topological Phases. It is thrilling to think about what other complex phenomena may be accessible to the analyses of Evolutionary Game Theory, making it an exciting new frontier for the study of dynamical systems.

References

[1] James Andreoni and John H Miller. “Rational cooperation in the finitely repeated prisoner’s dilemma: Experimental evidence”. In: *The economic journal* 103.418 (1993), pp. 570–585.

[2] Robert Axelrod. “Effective choice in the prisoner’s dilemma”. In: *Journal of conflict resolution* 24.1 (1980), pp. 3–25.

[3] Robert Axelrod and William D Hamilton. “The evolution of cooperation”. In: *science* 211.4489 (1981), pp. 1390–1396.

[4] John C Baez and Blake S Pollard. “Relative entropy in biological systems”. In: *Entropy* 18.2 (2016), p. 46.

[5] Jeannette Brosig. “Identifying cooperative behavior: some experimental results in a prisoner’s dilemma game”. In: *Journal of Economic Behavior & Organization* 47.3 (2002), pp. 275–290.

[6] Russell Cooper et al. “Cooperation without reputation: experimental evidence from prisoner’s dilemma games”. In: *Games and Economic Behavior* 12.2 (1996), pp. 187–218.

[7] Marc Harper. “Information geometry and evolutionary game theory”. In: *arXiv preprint arXiv:0911.1383* (2009).

[8] Marc Harper. “The replicator equation as an inference dynamic”. In: *arXiv preprint arXiv:0911.1763* (2009).

[9] Timothy Killingback and Michael Doebeli. “Self-organized criticality in spatial evolutionary game theory”. In: *Journal of theoretical biology* 191.3 (1998), pp. 335–340.

[10] Johannes Knebel, Philipp M Geiger, and Erwin Frey. “Topological Phase Transition in Coupled Rock-Paper-Scissors Cycles”. In: *Physical Review Letters* 125.25 (2020), p. 258301.

[11] György Szabó and Gabor Fath. “Evolutionary games on graphs”. In: *Physics reports* 446.4-6 (2007), pp. 97–216.

[12] Tsvi Tlusty. “A colorful origin for the genetic code: Information theory, statistical mechanics and the emergence of molecular codes”. In: *Physics of life reviews* 7.3 (2010), pp. 362–376.

[13] Stephen Wolfram. “Universality and complexity in cellular automata”. In: *Physica D: Nonlinear Phenomena* 10.1-2 (1984), pp. 1–35.

[14] Tsuneya Yoshida, Tomonari Mizoguchi, and Yasuhiro Hatsugai. “Chiral edge modes in game theory: a kagome network of rock-paper-scissors”. In: *arXiv preprint arXiv:2012.05562* (2020).