An Application of Evolutionary Game Theory to Social Dilemmas

Joshua A. Reyes

Department of Computer Science
University of Massachusetts Boston
E-mail address: jreyes@cs.umb.edu
Last updated: July 29, 2009

Contents

Introduction

Game theory is the study of interactive decision-making by individuals. It aims to investigate and advise useful courses of action. In some circumstances game theory has been very successful. It is a staple in the study of modern economics, and finds good application in a diverse number of fields—from sociology to cellular biology. In this thesis, we will explore some situations that have been a stumbling block for classical game theory. In particular, we will examine contests which pit the immediate interests of the individual against the long-term benefits to the an entire group, known collectively as *social dilemmas*.

In order to write down a description of these games in a form that is suitable for mathematical analysis we will build up the machinery of classical game theory, evolutionary game theory, and adaptive dynamics. Each topic is in itself quite broad, and the corresponding literatures are large. I have avoided many of the classic results from evolutionary game theory. You will not find Axelrod's tournament or TIT FOR TAT here. Even the Prisoner's Dilemma—a celebrity in classical game theory—makes only a brief cameo appearance. The discussion that follows is stream-lined so that the reader can quickly pick up the basic concepts that are necessary to understand the simulations in the final chapter. By the end of the thesis, though, the reader should be in a position to read and understand contemporary papers in evolutionary game theory with only a little outside investment.

Outline. In Chapter 1, we begin our journey though game theory where the field itself began, with the classical theory of games. Here so-called *rational* actors will individually seek to maximize their amount of *utility*. Utility is a squishy concept. Simply stated utility is whatever the actors want it to be, literally. It is the object of desire, and classical game theory demands that you can measure it. In modern discussions utility is often equated with money and bars of Swiss chocolate.

The agents in this theory have precise knowledge about every conceivable option and outcome, an infinite store of computing power to weigh all options, and unlimited time to make their decision. Individuals use the concept of a *Nash equilibrium*, or mutually best response, to guide their choices. Fortunately, all games exhibit at least one such equilibrium. Finding such a strategy, however, is not always (or even often) easy.

The players in classical game theory never make mistakes. They act perfectly and have access to perfect information. Both of these assumptions are cartoons of common experience. Ideally people would have all the facts ready and disposable to careful examination. The sad reality is that we tend to do the best we can with what we have. No more, and sometimes less. Just as 'no plan survives contact with the enemy,' no complicated game strategy survives contact with the player playing it. Individuals are well-meaning, but they frequently make mistakes executing their plans, no matter how carefully drawn up.

The invisible hand is a metaphor by founding economist Adam Smith in 1776. It took nearly two hundred years for World Bank economist Herman Daly to recognize that it is, in fact, a foot [**Dal73**].

Evolutionary game theory relaxes the assumptions of its classical counterpart. If the markets of classical game theory are directed by an invisible hand, then populations of evolutionary game theory are guided by an invisible foot. In Chapter 2, we will trade in agents bound by rationality for agents which exhibit *bounded rationality*. These individuals are not the hard-working, one-time, single strategy, perfectionist optimizers of yesteryear. No, these players lack foresight and complete knowledge of their situation. Accordingly, they will make mistakes, ask for advice, and satisfice.

Instead of studying the behavior of a lone individual, we will observe the aggregate behavior of a whole group of players all at once. We will position ourselves to ask how and when an innovative behavior can spread across a population of bounded rational players. Evolutionary game theory was developed to reconcile micromotives and large-scale macrobehavior in animal populations. We will borrow concepts from biology such as evolution, mutants, and the replicators from selfish gene theory—with due apologies to biologists—and give them a social spin, a practice that is now almost forty years old.

The chapter closes with a proof that a particularly important system of ordinary differential equations known as the *replicator dynamic* can be faithfully simulated by a simple agent-based model. Computer simulations are more pliable than the current mathematical formulations of evolution. Simulations provide evolutionary game theorists with a quick, inexpensive laboratory to test and toy with new ideas. In some fields where evolutionary game theory has been applied, such as urban planning, there are few reasonable and reliable alternatives to simulation.

The third and final chapter introduces adaptive dynamics, which generalize of the replicator dynamic of Chapter 2. Now it is time to shift focus from the population at large to the innovative, mutant individuals. Adaptive dynamics describes a local search process that is akin in spirit and mathematical form to hill climbing techniques used in artificial intelligence. We will apply an adaptive dynamical analysis to two social dilemmas, the Traveler's Dilemma and the Minimum Effort Coordination Game. Before doing so, we will complicate the games slightly. As a result, the games will be more realistic, easier to analyze, and produce results that are in accord with the behavior observed in real, human subjects.

Notation in evolutionary game theory has not been entirely standardized. In this account I have tried to maintain clear and consistent notation. At times my notation and terminology deviates strongly from the accepted norm. In particular I have stripped most of the biology from the biological concepts, steered clear of the original motivating biological examples, and in some cases completely changed the names of concepts.

Personally I believe that it is crucially important to hold at least a casual understanding of the biology that inspired evolutionary game theory. The original examples are wildly different than the ones I present here and provide a wider perspective than I can give in thirty pages. I encourage all readers to take this thesis as an initial point of departure, and by no means a destination.

CHAPTER 1

Classical Game Theory

Game theory is the study of interactive decision making. The classical theory studies the systematic, strategic play by so-called *rational* agents; that is, we assume that the players in classical contexts seek to optimize a quantity known as utility. Often utility can be thought in terms of money, but it need not be. Because optimization depends not only on what one agent chooses to do, but on what all the other players decide to do, game theoretic problems are typically hard to analyze. This chapter will briefly introduce the foundational concepts of game theory and define its basic terms. As an example, we use the theory to produce a simple, well-known model: the Prisoner's Dilemma.

The Prisoner's Dilemma is an example of a class of games called *social dilemmas*. Collectively they form an important class of situations that have been used to model and make decisions in a range of decision problems—from how much work to put into a group project to whether to initiate a preemptive nuclear strike against another country [**Pou93**]. We will use the prisoner's dilemma as an example to springboard into the central solution concept in game theory, the *Nash equilibrium*.

We explain several different kinds of Nash equilibria, mention their reliability and ubiquity, and then review their computational complexity. The chapter ends with the introduction of two more social dilemmas. The first is a generalization of the prisoner's dilemma called the Traveler's Dilemma; the second is the Minimum Effort Coordination Game.

1. Basic Terms

A *game* is a collection of rules and strategies. The *rules* specify who can do what and when. The *strategies*, on the other hand, explain what a particular player actually does and in which situations. We say that each decision a player makes in a game is a *move*. The rules provide the underlying structure and available moves, and the strategies pick out which moves to make. Even in games where there are only a few permissible moves, the number of strategies can be quite large. The staggering number of available strategies stems from the fact that there are no restrictions on how to choose among the current set of moves. In general, a strategy can be very complicated, or even arbitrary. As long as it picks out a permissible move for every possible state of the game, it is valid.

Game, rules

Strategy

Move

Consider the simple game called Menu. In this game, a patron must select a single entrée from the three hundred different kinds of hot dogs at a local restaurant. The list of choices alone is overwhelming but finite. Despite there being a large but limited number of items to select from, the number of strategies for this game is infinite. To get an idea of why this is the case, take, for example, the strategy that one particular player employed to the simultaneous amusement and annoyance of his peers:

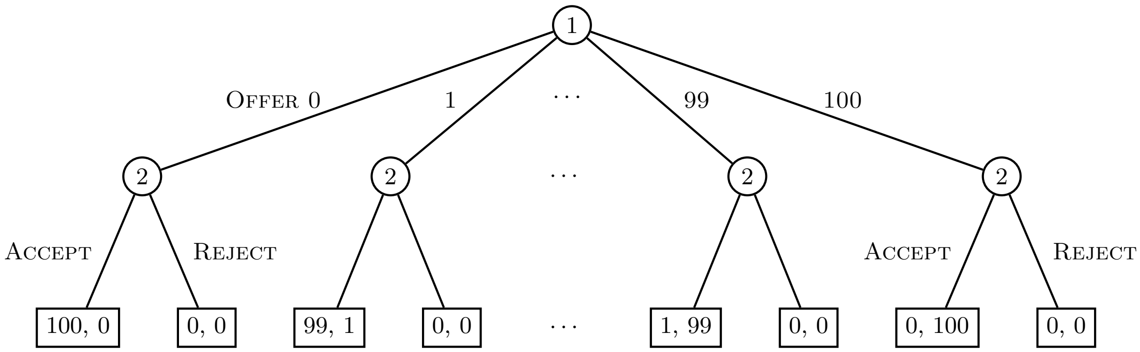

Figure 1. An extensive form representation of the Ultimatum Game.

1. Pick up a copy of the take-out menu.
2. Tear the menu into several small pieces and throw them into the air.
3. Order the dog printed on the piece that lands the farthest away.

While this strategy may seem foolish, it does the trick. It successfully makes a legal move in Menu, and so, by definition, is completely valid. Strategies may be probabilistic, as this one is, or deterministic. Because of the large number of alternative strategies, game theoretic problems are often difficult to analyze.

In order to study games in a systematic way, it is worth observing that a game is really just a sequence of decisions. Therefore we can simply list each possible move and connect it to its associated outcome, one move at a time, in the order dictated by the rules of the game. The resultant tree structure is known as the *extensive form* of the game.

Extensive form

Normally we think of the extensive form tree as being rooted at some node that represents what happens just before the game begins. Unless the rules explicitly mention which player moves first, a so-called act of God usually assigns the first move to one of the participants at random. The root node encodes this initial choice. Often the internal nodes carry labels with either the name or number of the player whose turn it is to move during that state of the game. Directed edges point to the state of the game that results from the action taken at the node. Terminal nodes carry additional information such as the *payoff* values at the end of the game along with the name of the players to receive the prizes. Payoffs can be positive, negative, or zero. A walk from the root to a terminal node is called a *play* of the game. In much the same way that a strategy describes what a player would possibly do in a given situation, a play describes which situations actually happened and their outcome in a real match.

Some games specify how to start explicitly; *e.g.*, the player to the right of the dealer goes first. When the rules are silent, an act of God settles the matter.

Payoff

Play

Since the extensive form captures the entire history of every conceivable play, the tree can be an unwieldy structure to analyze. Even after pruning for redundancies, the game tree for Tic-Tac-Toe includes more than thirty-four thousand nodes. It is not even known how many nodes are required to represent Chess in extensive form, but it is estimated that there must be about 10^{123} terminal nodes [**All94**]. For this reason, we will usually represent games in another, more compact way, known as normal form.

Some people have tried to develop an evolutionary theory of extensive form games. See Cressman [**Cre03**], for example.

Normal form

The *normal form* representation does not include the history of the game. In this case the strategies are given simple labels like A, B, or Cooperate, and the details of their exact specification are left out of the picture. The normal form details which strategies can be used, what their outcomes are, and nothing more. It tends exclusively to the bottom-line. More formally, a normal form representation of any game must include:

You can pretend that each strategy has been meticulously drawn up and housed in a very safe place, so that if someone asked you to play the game in real life he could. But since we're modeling behavior, and playing the game for real takes a lot of time, it's best left somewhere else.

(1) The number of players in the game,
(2) An exhaustive list of all the strategies available to each player, and
(3) A catalogue of payoffs that accrue to each player using each of her strategies against any other player using any of his strategies.

While this representation may seem too sparse to be as useful as extensive form, we routinely describe situations in normal form. Newspapers headlines exclaim that the Red Sox beat Tampa Bay 13–2 and that Lance Armstrong won the Tour de France without further detail. Often that level of detail is all that is needed. In extensive form, we would include the exciting and decisive moves that happened along the way—that the Sox were tied until the bottom of the ninth, or that Armstrong took a hard fall in the middle of the race. But normal form is like that really strict teacher who only pays attention to the fact that your homework is late and won't listen to excuses.

For two player games it is customary to write down the normal form as a table. The rows of the table list the strategies available to Player i, and the columns give the strategies of Player ii. The entries of the table are the payoffs to each player, assuming that they had adopted the corresponding row and column strategies. The resultant table has two numbers in each entry, one for the first player and another for the second player. Such a table is called a *bimatrix*. We denote the payoff to Player i playing strategy s against Player ii choosing strategy t as $\pi_1(s,t)$. The corresponding payoff to Player ii is notated $\pi_2(s,t)$. According to this notation, a typical payoff matrix π is given by $[(\pi_1(s,t), \pi_2(s,t)]$ for row strategies s and column strategies t.

Bimatrix

There is no reason in advance to assume that players draw from the same pool of strategies. In many games different players fill different roles, and therefore have different sets of strategies at their disposal. Even if two players do happen to adopt a common strategy, their individual payoffs need not be equal because their roles aren't. That is, if Player i plays strategy x while Player ii adopts strategy y in one match, and then they switch strategies and play again, the payoffs that the two players receive in each instance could differ:

In the pool party classic Marco Polo, only one player can shout "Marco!" whereas everyone else must respond with a loud "Polo!"

$$\pi_1(x,y) \neq \pi_2(y,x).$$

Such games are called *asymmetric games*. When the game is symmetric, we may use a matrix instead of a bimatrix to keep track of payoffs. In this thesis, I will only consider two-person symmetric games directly, though in general there could be more players and they could assume asymmetric roles. All of the payoff matrices I present will employ the simplified notation.

Asymmetric game

The payoff matrix of a symmetric game is not necessarily symmetric. A game is *symmetric* if $\pi_2 = \pi_1^\top$ as matrices.

2. Modeling Altruism

We want to make our models as simple as possible—that is, if we want to analyze them with any rigor. But this comes with the price of fidelity to real-life situations. Here we will develop the simplest, though most popular model of altruistic behavior in evolutionary game theory. In this game there aren't a lot of options. Each player can choose to help the other, but at a cost to himself. It captures a simple choice: to help, which consumes some amount of resources, time, or energy—or not. In this game there are only two acceptable moves, Help and Ignore.

A = Help
B = Ignore

Only two players can play this game at a time. A move of Help benefits the other player an amount b, but at a personal cost of c. Ignore does not benefit the other player, but it incurs no personal cost either. To make the transaction worthwhile we require that $b > c > 0$. Help describes a charitable act. Ignore expresses apathy.

$$\begin{array}{c} \\ \text{A} \\ \text{B} \end{array} \begin{array}{c} \begin{array}{cc} \text{A} & \text{B} \end{array} \\ \begin{bmatrix} b-c & -c \\ b & 0 \end{bmatrix} \end{array}$$

Dominant strategy

In this case strategy B (Ignore) *dominates* strategy A (Help); that is, no matter what the opposing player chooses to do, you will certainly obtain a higher payoff if you choose B. Because this game is symmetric, it makes sense for any two players to both play Ignore. The tenets of classical game theory suggest that anyone playing this game should look the other way.

But this Good Samaritan-take is not the only way to outfit this particular payoff matrix. In game theory, it is common to attach a little story with a payoff matrix like the one above in order (hopefully) to clarify the model. Many different stories can be made up from a single payoff matrix. As I've framed it, this model captures the story of the Good Samaritan (perhaps more aptly named the Mutually Good Samaritans). However, this instance is a special case of a more general class of games generally known as the PRISONER'S DILEMMA. Its story goes like this:[1]

DISCLAIMER: The characters mentioned in the following game theoretic narratives are fictitious, even when they represent real people.

Many years ago Matt and Zhimin planned to rob a bank. Unfortunately for them, the robbery was not as clean as they would have hoped, and they quickly found themselves engaged in a high-speed chase covering an excess of 400 miles over the rural roads of Nebraska. After several hours and considerable damage to the local corn fields, the Nebraska state police, in a singular act of courage and quick maneuverability, surrounded the car that contained the would-be robbers, and arrested them. Back at the police station, Matt and Zhimin were placed in two separate interrogation rooms and were made identical offers.

The Nebraska police could have equally well handed Matt the following payoff matrix

$$\begin{array}{c} \\ \text{A} \\ \text{B} \end{array} \begin{array}{c} \begin{array}{cc} \text{A} & \text{B} \end{array} \\ \begin{bmatrix} -1 & -20 \\ 0 & -10 \end{bmatrix}. \end{array}$$

"Matt, here's the offer that we're making to both you and Zhimin. If you both hold out on us, and don't confess to the bank robbery, then we admit that we don't have enough proof to convict you. However, we *will* be able to jail you both for one year for reckless driving and endangerment to corn. If you turn state's witness and help us convict Zhimin (assuming he doesn't also confess), then you will go free, and Zhimin will get twenty years in prison. On the other hand, if you don't confess but Zhimin does, then *he* will go scott free and *you* will get twenty years in the slammer."

"What happens if both Zhimin and I confess?" asked Matt.

"Then we'll lock you both up for ten years each," answered the interrogator.

[1]This version of the PRISONERS DILEMMA has been shamelessly adapted from a problem set that accompanies the *Structure and Interpretation of Programming Languages* by Sussman and Abelson on the book website `http://mitpress.mit.edu/sicp/psets/ps4prs/readme.html`; at least it did at 11:13am, 7 June 2008.

Matt, who picked up a thing or two about game theory from his labmate, reasoned this way: "Suppose Zhimin intends to confess. Then if I don't confess, too, I'll get twenty years. But if I do confess, I'll only get ten years. On the other hand, suppose Zhimin decides to hold out on the cops. Then if I don't confess, I'll only go away for a year. But if I do confess, I won't go to jail at all. So no matter what Zhimin does I'm better off confessing than holding out. So I'd better confess."

Naturally Zhimin, who had worked in the same lab and had learned game theory from the very same labmate, employed exactly the same reasoning. Both criminals confessed and both were sentenced a ten year term in jail. (Actually, they never wound up in prison. When they were in court and heard that they both had turned state's witness, they strangled each other. But that's another story.) The police, of course, were triumphant since the criminals would have been free in only a year had they both remained silent.

The payoff matrix for the general game gives a reward R for mutual cooperation, a punishment P for mutual defection, offers a high temptation T to defect unilaterally, and a sucker's payoff S to unilaterally cooperate. Here we require that $T > R > P > S$. It is customary to assume additionally that $2R > T + S$ to maintain the dilemma in iterated games. Without this restriction, colluding players could flout cooperation and still profit by taking turns receiving the temptation and sucker's payoffs in alternate rounds. The Mutually Good Samaritans game payoffs satisfy these constraints.

A = Stay mum
B = Fink

$$\begin{array}{c c} & \begin{array}{cc} \text{A} & \text{B} \end{array} \\ \begin{array}{c} \text{A} \\ \text{B} \end{array} & \begin{bmatrix} R & S \\ T & P \end{bmatrix} \end{array}$$

3. Solution Concepts: Nash Equilibria

In the Prisoner's Dilemma story, Matt and Zhimin came to the same conclusion—to fink. In this game that pair of strategies (Fink, Fink) is in some sense the only non-stupid move (that is, unless you happen to know some extra information about your opponent). Once both players choose this strategy neither can do better by *unilaterally* switching to some other strategy. This property will likely hold true for any "good" pair of strategies that we would want to look for.

A general pair of strategies (s, t) is good if

(1) Strategy s is an optimal choice for Player I given that Player II will choose t; s is a best response to t.
(2) Strategy t is an optimal choice for Player II given that Player I will choose s; t is a best response to s.

If both properties hold simultaneously, then the pair (s, t) is called a *Nash equilibrium* (NE). We can cast these requirements in a more mathematical language by noting that for all strategies $\tilde{s}$ and $\tilde{t}$,

Nash equilibrium

(1) s is a best response to t if and only if $\pi_1(s, t) \geq \pi_1(\tilde{s}, t)$,
(2) t is a best response to s if and only if $\pi_2(s, t) \geq \pi_2(s, \tilde{t})$.

If the game in question is symmetric, then the requirement for a Nash equilibrium simplifies. Then a pair (s, s) is a *symmetric Nash equilibrium* if it is a best

Symmetric Nash equilibrium

response to itself; that is, given the strategy s and any other strategy t,

$$\pi(s, s) \geq \pi(s, t).$$

Strict Nash equilibrium

If each of these inequalities is strict, then the corresponding pair is called a *strict Nash equilibrium*. Any Nash equilibrium is a situation of stable mutual adjustment. Both players anticipate what the other will do, and both of these guesses turn out to be right. These equilibria are a set of self-filling prophecies that players form about each other. They are often extremely conservative and can lead to suboptimal payoffs for everyone, as is witnessed by the Prisoner's Dilemma. While strict Nash equilibria are unique, an arbitrary two-person game with finite strategies may exhibit zero, one, or many NE. In the latter case, classical game theory gives no way to choose among them, making it difficult to coordinate on any one equilibrium.

Pure strategy

Mixed strategy

Profile

Support

Totally mixed strategy

Mixed Strategies. Until now we have only been considering *pure strategies*, those labeled on the sides of the payoff matrix. A *mixed strategy* σ selects to play pure strategy $s_1, \ldots, s_n$ with probability $p_1, \ldots, p_n$. Since the p_i are probabilities, we require that each p_i be non-negative and that they sum to one. The vector p of probabilities associated to a mixed strategy is the strategy's *profile*. The *support* of a mixed strategy σ is the collection of pure strategies that occur with strictly positive probability. If the support of σ contains all possible pure strategies, then σ is *totally mixed*. Mixed strategies generalize pure strategies, since a pure strategy s_i corresponds to the strategy profile with $p_i = 1$ and all other $p_j = 0$. Mixed strategies capture the fact that sometimes people decide what to do with an element of chance, but that some choices are more likely than others.

Let $s = (s_1 \cdots s_n)^\top$ and $p = (p_1 \cdots p_n)^\top$. Then the mixed strategy σ is nothing more than their inner product $\sigma = \langle s, p \rangle$.

It is common to write a mixed strategy as a linear combination of pure strategies

$$\sigma = p_1 s_1 + \cdots + p_n s_n = \sum_{i=1}^{n} p_i s_i.$$

To accommodate mixed strategies, we need to expand our notion of payoff accordingly. To do so, we'll sum over the payoffs of the pure strategies—because we already know how to do that—but weight them according to how often two strategies are (probably) played against each other. Therefore, let $\sigma = \sum_{i=1}^{n} p_i s_i$ and $\tau = \sum_{i=1}^{m} q_j t_j$. Then the payoff to PLAYER I playing σ against PLAYER II playing τ is

The product $p_i q_j$ measures how likely it is that s_i plays against t_j.

$$\pi_1(\sigma, \tau) = \sum_{i=1}^{n} \sum_{j=1}^{m} p_i q_j \pi_1(s_i, t_j) = \langle p, \pi_1(s, t) \cdot q \rangle,$$

where s and t are vectors of the pure strategies available to PLAYERS I and II, and the raised dot also denotes the standard inner product. We define the payoff to PLAYER II similarly.

Nash equilibrium for mixed strategies

Likewise, we need to generalize our solution concepts to include mixed strategies. We say that a pair of mixed strategies (σ, τ) is a *Nash equilibrium* if

$$\pi_1(\sigma, \tau) \geq \pi_1(\tilde{\sigma}, \tau) \text{ and } \pi_2(\sigma, \tau) \geq \pi_2(\sigma, \tilde{\tau})$$

for any other mixed strategies $\tilde{\sigma}$ and $\tilde{\tau}$. Strict NE over mixed strategies are defined similarly.

Armed with a new set of strategies and potential equilibria, we can ask whether we have really gained anything. The answer turns out to be yes. The Fundamental Theorem of Classical Game Theory, sometimes known as Nash's Existence Theorem, guarantees that any two-person game with a finite number of pure strategies admits at least one Nash equilibrium in (possibly) mixed strategies. Remember that when we restrict ourselves merely to pure strategies, it is possible that no equilibrium exists. A proof of the existence theorem can be found in any standard text on game theory, for example [Gib92].

MATCHING PENNIES is a simple two-player game with no NE in pure strategies. Its payoff matrix is $\begin{bmatrix}(-1, \;\; 1) & (\;\; 1, -1)\\ (\;\; 1, -1) & (-1, \;\; 1)\end{bmatrix}$.

Computational Complexity. Every day we are bombarded with decisions to make. Nash equilibria present attractive moves. However, we have not discussed how easy they are to come by. The classical rational agent has an infinite computational power and time at easy disposal, and so has little trouble spotting equilibria. Yet in the trenches of reality we seldom have unrestricted access to such resources. Even professional game theorists do not always play the Nash equilibrium [BCN05]. To measure just how obtainable NE are, we turn to complexity theory.

The standard notions of computational complexity, such as NP-completeness, were created to characterize computational problems such as SATISFIABILITY and SUBSET-SUM that seemed to resist easy solutions. The problem NASH, to find any equilibrium given a game in normal form, is a different kind of intractable problem, and NP-completeness is not an appropriate notion of complexity to describe it.

What separates NASH from other difficult problems is that *every* game is guaranteed to have a Nash equilibrium. In problems like SATISFIABILITY, a solution to individual instances may or may not exist. In fact, typical NP-completeness reduction arguments rely on the fact that solutions may not exist. The original proof Nash existence theorem is itself a reduction of the existence of a mixed strategy equilibrium to the existence of a Brouwer fixed point. The Brouwer fixed point theorem is itself notoriously non-constructive, and finding such a fixed point is also known to be a hard problem—again, because a solution is always known to exist.

THEOREM (Brouwer Fixed Point Theorem). Every continuous function from the unit n-dimensional ball to itself has a fixed point.

If NASH is slightly tweaked so that it becomes a standard decision problem, then NP-completeness appears almost immediately. For example, the following theorem about several modified NASH has been known since 1989 [GZ89]. The complexity class of the original problem was not proven until 2005 [DGP06].

THEOREM 3.1. *The following are NP-complete problems: given a two-player game in normal form, does it have*

1. *at least two Nash equilibria?*
2. *a Nash equilibrium in which* PLAYER I *has a payoff of a minimum, given value?*
3. *a Nash equilibrium in which the combined payoffs are at least a given amount?*
4. *a Nash equilibrium whose support contains the pure strategy s?*
5. *a Nash equilibrium whose support does not contain the pure strategy s?*

Complexity Classes PPAD and TNFP. NASH lies in the complexity class of polynomial parity arguments on directed graphs (PPAD). To understand PPAD, it is first easier to understand its superclass TFNP, the set of so-called Total Function Nondeterministic Polynomial problems. Here it is more useful to keep in mind the verifier interpretation of NP than of non-determinism.

Consider a polynomial-time recognizable predicate $P(x, y)$ such that for every x there is at least one y which renders $P(x, y)$ true.[2] A problem in TFNP follows

TNFP

Corollary 2.1.1 provides a proof that the predicate of NASH is polynomial-time decidable.

[2]Additionally, one requires that P be *polynomially-balanced*; that is, the truth of $P(x, y)$ implies that $|y| \leq p(|x|)$ for some polynomial p.

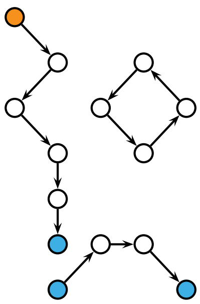

A typical problem in PPAD. The source vertex s is colored orange. Solutions are marked in blue. (*Adapted from Algorithmic Game Theory* [**Nis07**].)

this simple form: given an input x, find a y such that $P(x, y)$ holds. Even though a solution is assured to exist; it may be still difficult to find one. Famous inhabitants of TFNP include prime factorization, the antipodal points of the Borsuk-Ulam theorem, and the traveling salesperson problem. In fact, all polynomial local search problems lie within TFNP by the finiteness of the search space.

Many TFNP problems, like Nash, come from combinatorial theorems that guarantee solutions. Subclasses of TFNP come from type of mathematical proof used to prove that a solution always exists. The class PPAD where the guarantee of a solution is based on the parity argument on a directed graph. To define the class formally, we need only to specify one of its complete problems.

PPAD

> End-of-the-Line: Let G be a directed graph with no isolated vertices such that every vertex has at most one predecessor and one successor. Furthermore, let G be specified by a polynomial-time computable function $f(v)$ that returns the predecessor and successor (if they exist) of the vertex v. Given a vertex s in G with no predecessor, find a vertex $t \neq s$ with either no predecessor or no successor. In other words, find any source or sink of the directed graph other than s.

The fact that such a vertex t always exists can be used to prove the Brouwer fixed point theorem, and therefore the Nash existence theorem, consigning Nash to PPAD. What is more surprising is that Nash is complete for this complexity class, which means that it tantamount in difficulty to finding a Brouwer fixed point [**Pap07**].

PPAD-completeness shows that Nash is hard for arbitrary games. We may take this result as a reason to question the assumptions underlying rational behavior. In subsequent chapters, we will relax some of the requirements of rationality and refine our concept of equilibrium. To discover these new equilibria we will need to alter the classical theory by injecting some biology and dynamical systems theory. While these new equilibria will (often) be easier to find, they do not always exist.

Finding one of the refined equilibria is both an NP-hard and coNP-hard task.

4. Social Dilemmas

In the prisoner's dilemma, we saw that selfish individuals do better than cooperative ones. One the other hand, cooperative pairs do better than selfish pairs. Even though the noncooperative, rational strategy in the prisoner's dilemma always yields a higher payoff, if everyone were to adopt the rational strategy the entire group would be worse off than if they had cooperated. These observations are true more generally, across a wide variety of games. Games which pit short-term individual interests against shared group goals form an appealing category called *social dilemmas*.

Social dilemma

What makes social dilemmas fascinating is that real human players display a baseline of cooperation well above what rational choice theory would predict. In fact, many people cooperate despite holding the wide-spread belief that everyone else is selfish! The studies in [**Wut93**], for example, suggest that individual's rate of donating blood does not increase with monetary incentives. However the very same individuals who give for free assume that most everyone else would donate only when offered sufficient financial reward. Likewise, when asked whether they would rent to an unmarried couple, all landlords interviewed in the early 1970s in Oregon responded positively. Yet they estimated that the general acceptance rate at only about fifty percent [**Daw72**]. Such regularly high levels of cooperative behavior are puzzling for classical game theory.

The remainder of this thesis will try to reconcile the predictions of classical game theory with actual behavior of individuals observed both in and outside of the laboratory. We will pay special attention to the TRAVELER'S DILEMMA, a generalization of the prisoner's dilemma, and the MINIMUM EFFORT COORDINATION GAME. In a sense these games exhibit opposite pathologies: one admits a unique but seemingly irrational Nash equilibrium; the other contains an infinite number of indistinguishable ones.

Traveler's Dilemma. The Traveler's Dilemma was introduced by Cornell economist Kaushik Basu as a challenge to economic rationality, and, in particular, to the Nash equilibrium [**Bas07**], [**Bas94**]. In this game there is a unique, pure strategy Nash equilibrium. However almost no one ever plays it and they're better off for it. The narrative goes something like this:

Two travelers, Bob and Rob, returning from a biodiversity conference in Costa Rica find that their airline has damaged the identical antiques they had purchased from a local shop during their stay. The airline manager says that he would be happy to compensate the travelers for their losses, but that he has no idea how to price these strange objects fairly. He cannot simply ask Bob and Rob the value because he can tell that they have an eye for profit and will artificially inflate the cost.

Instead the manager devises a more complicated plan to coax out the true price of the antiques. "Because the luggage was not insured," began the manager, "the airline is regrettably only liable for losses worth up to $100, but you may certainly claim less. Any whole dollar value as low as $2 will do." Immediately Bob and Rob turned to each other and smiled. They were now sure that they'd make a killing on those lousy touristy trinkets they had picked up last minute. But the manager continued.

"If you both make the same claim," he explained, "I will assume that you have given me the honest price and reimburse each of you that amount. But if the claims are different, then I will assume that the lower of the numbers is the actual price and that the person who wrote the higher price is cheating. In that case I will only pay out the lower amount to both of you with a small penalty to the cheater. I will take $2 away from the one who wrote the higher number and give it to the person who wrote the lower number." He then escorted Bob and Rob, who were now visibly less excited, to opposite ends of the counter and handed them each a claim form to fill out without conferring.

Bob, who had learned some game theory from Matt and Zhimin before their fated chase in Nebraska, reasoned this way: "I could put down the maximum claim of $100, knowing that Rob is as greedy as I am and will likely do the same. But since Rob will put down $100, I can do slightly better by claiming $99. Then he'll receive $97 and I'll get a whopping $101. But since Rob will know that I'll do that, he'll claim $98. In that case I should claim $97." Following this train of logic to its bitter end, Bob eventually decided his best bet was to claim the minimum, a paltry two dollars.

Naturally, Rob, who had also learned his game theory from Matt and Zhimin, reasoned in an identical fashion. Both men claimed the two-dollar minimum and therefore both received the two-dollar minimum. For his cleverness and unwavering loyalty to the airline, the manager was awarded a raise and an all-expenses paid vacation to Detroit.

In this version, the reward/punishment parameter R was 2, though this value was somewhat arbitrary. As long as it is positive but smaller than the maximum claim, it provides an incentive for a player to undercut her opponent in order to increase individual gain. Then an identical backwards induction argument shows that (R, R) is the unique Nash equilibrium of this dilemma. This game is symmetric, and so it can be characterized by a single payoff function.

$$\pi(s, t) = \min(s, t) - R \operatorname{sgn}(s - t).$$

Minimum Effort Coordination Game. The minimum effort coordination game suffers more or less from the opposite problem of the traveler's dilemma. Instead of exhibiting a single, faulty Nash equilibrium, the minimum effort coordination game displays an uncountably infinite number of viable equilibria [**GH01**]. Classical game theory is silent about how to choose among them. Shockingly, in various other coordination games strangers inexplicably hone in on a handful of potential equilibria without communicating. Because they are similar in spirit and form, we present one of them here.

In MEET ME IN NEW YORK, an eccentric but wealthy computer scientist has selected independently two strangers to play. If the two individuals can meet at the same location and at the same time anywhere in New York City on April 1, they will each win twelve billion dollars and a lifetime supply of Tic Tacs. So that the players can recognize each other, the computer scientist has supplied them with unmistakably tacky jester's suits and matching cap. However, if they fail to meet at the same place or same time, then both go home empty-handed. Worse, they must decide where and when to show up alone. Despite the dizzying number of possible choices, when asked people usually choose to meet at a small number of places—the Statue of Liberty, Times Square, or the Empire State Building—and almost always at noon.

Unlike Meet Me in New York, the minimum effort coordination game only requires one decision (given by a real number between zero and one). But it also lacks the context and familiarity of New York City. In this game, you can imagine going on a road trip with a friend. Even though you have planned to be in the car and on the road by 11AM, you still have the choice to arrive early—or in the case of your friend, show up late. If you both show up early, then you can leave early. If you both show up late, then you leave late. No matter what, neither of you can leave until you are both there. Meeting at the same time is a Nash equilibrium of this game, early, late, or on time. This sort of situation is captured neatly by the payoff function

$$\pi(s, t) = \max(\min(s, t) - ks, 0)$$

where s, t, and $k \in (0, 1)$ and k represents the effort necessary to show up s units early. This is the minimum effort coordination, so-called because the payoff is at most what the lazier player puts in, but usually less. As noted before, any pair of the form (s, s) is a Nash equilibrium.

CHAPTER 2

Evolutionary Game Theory

1. Redefining Terms

The Nash existence theorem promises that at least one equilibrium exists in any game, but that it might be a mixed strategy. To get some practice playing mixed strategies in under your belt, follow this set of directions taken from the Nash equilibria of three hypothetical games:

- □ Color the square solid red with 15% probability.
- △ Color the triangle solid blue with 40% probability.
- ◯ Color the circle solid yellow with 70% probability.

In a one-shot game, how can you be sure that you've actually implemented the strategy—how would you prove to someone that your coloring was a faithful execution of the instructions? After all, the square will either be colored red or not. You were asked to perform a mixed strategy, yet the result is indistinguishable from a pure strategy.

When discussing mixed strategies, people sometimes introduce a handy randomizing device that will select a number according to a given profile of probabilities. At the end of the day, this number corresponds to a single pure strategy from the support of the associated mixed strategy. This device seems to have solved the problem, unless the decision is a bit more serious. Next imagine the mixed strategy's directive were to preemptively strike an opposing country with three-fifths probably. How do you feel when the device suggests to escalate a war?

There are two standard interpretations of the probability of an event. One counts the frequency of outcomes over a long series of trials. As the number of trials increases, the frequency of occurrences approaches the true probability of the event's happening. In the second, several experimenters each perform a single trial—such as coloring the triangle blue with probability two-fifths—all at the same time. The results are tallied, and again frequencies computed. Evolutionary game theory adopts the second portrayal for its working definition of probabilities.

The second reading of probability invites parallelization if simulated on a computer.

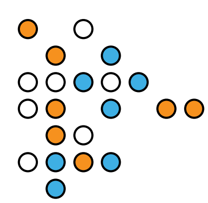

Population

Strategy. In evolutionary game theory, the concept of an individual player is not as flexible as its classical counterpart. Individuals can play just one fixed pure strategy to use during all of strategic encounters. This restriction may appear a step backwards at first glance. However in evolutionary game theory, the fundamental player is **not** an individual. Instead it is an entire *population* of individuals considered all together at once. The strategy profile σ of a population with access to pure strategies $s_1, \ldots, s_n$ is the vector of frequencies $p_1, \ldots, p_n$ with which the constituent individuals of the population have adopted each strategy. The figure on the lower right of the previous page depicts a population playing the mixed strategy $\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$.

The only difference between classical and evolutionary payoffs lies in their understanding. Rather than representing monetary values or ill-defined units of utility, payoffs to one individual represent the change in Darwinian fitness of its strategy. Individuals may ask for advice on which strategy to take from neighboring individuals from one generation to the next. Individuals weight the advice they receive by comparing relative payoffs before deciding whether to update their strategy. Therefore the strategy profile of a population can change over time. The resultant change is called *evolution*. Notice that individuals do **not** evolve themselves, only populations do. If an individual spontaneously switches strategies or accidentally implements advice to change incorrectly, we say a *mutation* has an occurred and that individual is deemed a *mutant*.

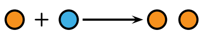
Evolution

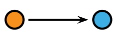
Mutation

Equilibrium. It is not common to discuss one population's playing against another. Contests still take place between individuals. Here, too, the evolutionary theory departs from the classical one. Individuals play several games against opponents randomly selected from the population. Because individuals may update their strategies, we seek those strategy profiles that are somehow stable with regard to switching. In this case, we call a strategy *evolutionarily stable* if, whenever all members of the population adopt it, no mutant strategy can invade the population under the influence of natural selection.[1]

Evolutionarily stable strategy

Imagine a population comprised of A-type individuals all playing strategy σ. Next introduce a small number of B-types playing a different strategy τ. For notation, let ϵ denote the frequency of the mutant B-types in the newly mixed population, and $1-\epsilon$ the frequency of the incumbent A-types. Then in total we have a population strategy profile of $p = \epsilon\tau + (1-\epsilon)\sigma$ in all. A population consisting of A-types will be evolutionarily stable if the old type does better than the newcomer B-types. Then the strategy σ is an *evolutionarily stable strategy* (or ESS) if for any $\tau \neq \sigma$,

Evolutionarily stable strategy

$$(\star) \qquad \pi(\sigma, p) > \pi(\tau, p),$$

holds for all $\epsilon > 0$ that are sufficiently small—that is, smaller than some suitable *invasion barrier* $\bar{\epsilon}(\tau) > 0$ which depends on the mutant invader.

Invasion barrier

The inequality $(\star)$ encodes two useful conditions that characterize evolutionarily stable strategies. To tease them out, we need only unpack the inequality and group like terms. Expanding the each side one at a time gives

$$\pi(\sigma, p) = \pi(\sigma, \epsilon\tau + (1-\epsilon)\sigma) = \epsilon\pi(\sigma,\tau) + (1-\epsilon)\pi(\sigma,\sigma)$$
$$\pi(\tau, p) = \pi(\tau, \epsilon\tau + (1-\epsilon)\sigma) = \epsilon\pi(\tau,\tau) + (1-\epsilon)\pi(\tau,\sigma).$$

When put together, two other familiar inequalities begin to emerge.

$$(1-\epsilon)\,[\underbrace{\pi(\sigma,\sigma) - \pi(\tau,\sigma)}_{\text{Equilibrium}}] + \epsilon\,[\underbrace{\pi(\sigma,\tau) - \pi(\tau,\tau)}_{\text{Stability}}] > 0.$$

Therefore we now have an alternative definition of evolutionarily stable strategies. A strategy σ is an ESS if and only if

(1) *Equilibrium.* It is a best response to itself: $\pi(\sigma,\sigma) \geq \pi(\tau,\sigma)$ for all other strategies $\tau \neq \sigma$.

[1] Most of the technical results of the following two sections are well-known, and can be found in [**Now06**] and [**Cre03**]. The ordering of results but not the exposition follows [**HS98**] most closely.

(2) *Stability.* If there is an alternative best response τ, then σ fares better against τ than τ plays against itself: if $\sigma \neq \tau$ but $\pi(\sigma,\sigma) = \pi(\tau,\sigma)$, then $\pi(\sigma,\tau) > \pi(\tau,\tau)$.

The first condition ensures that all ESS are also Nash equilibria. A strict Nash equilibrium is always an ESS. Be wary, though. Not all Nash equilibria are evolutionarily stable strategies. Evolutionarily stables strategies refine the notion of Nash equilibrium because they take into account strategies which do just as well as a standard NE and could therefore take over a population. ESS guard against neutral drift. In fact, ESS are protected by a buffer of less successful strategies.

THEOREM 1.1. *If σ is an evolutionarily stable strategy, then for every other strategy $\tau \neq \sigma$ in some neighborhood of σ,*

$$\pi(\sigma,\tau) > \pi(\tau,\tau).$$

PROOF. Assume that σ is evolutionarily stable. We will construct a neighborhood about σ such that the desired inequality holds. In fact, we will construct N in such a way that all nearby strategies τ are of the form $\epsilon\rho + (1-\epsilon)\sigma$ for suitably small ϵ and appropriate strategy ρ.

First choose ρ from the compact set $K = \{p \,|\, p_i = 0 \text{ for some } i \in \operatorname{supp}\sigma\}$. Then for any $\rho \in K$ the defining inequality ($\star$) holds, provided that $\epsilon < \bar{\epsilon}(\rho)$. As the invasion barrier $\bar{\epsilon}(\rho)$ varies continuously, it attains a strictly positive minimum $\bar{\epsilon}$ on K. Then after multiplying by $\epsilon \in (0, \bar{\epsilon})$,

$$\epsilon^2\pi(\sigma,\rho) + \epsilon(1-\epsilon)\pi(\sigma,\sigma) > \epsilon^2\pi(\rho,\rho) + \epsilon(1-\epsilon)\pi(\rho,\sigma).$$

Next add $\pi((1-\epsilon)\sigma, (1-\epsilon)\sigma + \epsilon\rho)$ to both sides. Assigning $\tau = (1-\epsilon)\sigma + \epsilon\rho$ produces the desired inequality for all τ in a neighborhood of σ. □

Computational Complexity. Evolutionarily stable strategies appear to be a resilient collection of behaviors for populations to adopt. As before, it is reasonable to wonder how easy it is to discover and verify that a strategy is indeed an ESS. Even though it looks as if the problem of verifying a Nash equilibrium resides in continuous mathematics, it is at its core a matter of combinatorics.

The number of pure strategies available to each individual is finite. Mixed strategies are nothing more than a weighted average of the pure strategies in its support. All of the strategies in the support of a Nash equilibrium (and therefore of an ESS) must give the same constant payoff when played against each other. If they resulted in unequal payoffs, then we could do better by playing the pure strategy that gives a higher payoff more often and playing the one that pays out a lower value less frequently.

This line of reasoning leads to the Bishop-Cannings theorem.

THEOREM 1.2. *Let σ be a Nash equilibrium. Then for all pure strategies s_i in the support of σ,*

$$\pi(s_i,\sigma) = \pi(\sigma,\sigma).$$

Furthermore, no individual $s_i \in \operatorname{supp}(\sigma)$ is a pure evolutionarily stable strategy.

PROOF. Suppose otherwise, and assume that $\pi(s_i,\sigma) < \pi(\sigma,\sigma)$ for some pure strategy $s_i \in \operatorname{supp}(\sigma)$. Then we may write $\sigma = ps_i + (1-p)\sigma_{-i}$ as a combination of s_i and the strategy σ_{-i} that remains when s_i is omitted. Both components have been weighted by their associated probabilities of being played. Then by definition,

$$\begin{aligned}\pi(\sigma,\sigma) &= p\pi(s_i,\sigma) + (1-p)\pi(\sigma_{-1},\sigma)\\ &< p\pi(\sigma,\sigma) + (1-p)\pi(\sigma_{-i},\sigma).\end{aligned}$$

Subtracting $p\pi(\sigma,\sigma)$ from both sides and dividing by $(1-p)$ yields the blasphemous relationship

$$\pi(\sigma,\sigma) < \pi(\sigma_{-i},\sigma).$$

The above inequality contradicts the fact that σ is a Nash equilibrium. Likewise, s_i cannot fare better against σ than σ itself, either. Therefore for any $s_i \in \mathrm{supp}(\sigma)$, the two payoffs must be equal.

Additionally, none of the s_i can be a pure ESS due to the stability criterion. All of them perform equally well against each other. □

The Bishop-Cannings theorem gives another characterization of Nash equilibria which makes verification a simple, polynomial-time affair.

COROLLARY 1.1. *A strategy σ is a Nash equilibrium if and only if its strategy profile satisfy the linear system of equations*

$$[\pi\cdot\sigma]_1 = \cdots = [\pi\cdot\sigma]_n \text{ and } p_1 + \cdots + p_n = 1,$$

where a single raised dot denotes normal matrix-vector multiplication and $[v]_i$ represents the ith coordinate of the vector v.

PROOF. The expected payoff to an individual playing s_i in a population described by the strategy profile σ is $\pi(s_i,\sigma)$. As a mixed strategy, s_i corresponds to the ith standard basis vector. Therefore,

$$\pi(s_i,\sigma) = \langle s_i, \pi\cdot\sigma\rangle$$

picks out the ith coordinate of the vector $\pi\cdot\sigma$. This value is precisely $[\pi\cdot\sigma]_i$. If the population profile σ is an ESS, then the Bishop-Cannings theorem applies and the result follows immediately. □

For example, the symmetric game with payoff matrix

$$\begin{bmatrix}0 & 2 & 0\\ 2 & 0 & 0\\ 1 & 1 & 0\end{bmatrix}$$

admits no ESS [**HS98**].

Unlike standard Nash equilibria, evolutionarily stable strategies do not always exist. It is relatively easy to produce a symmetric two-player game that lacks an ESS. The computational complexity of finding an ESS in an arbitrary two-player symmetric game is both NP-hard and coNP-hard. A proof by Nisan [**Nis07**] relies on a reduction from the problem of checking whether a graph has a maximum clique of size exactly k, which is already known to be both NP-hard and coNP-hard. In the same proof, he shows that the problem of recognizing whether a given strategy is an ESS is also coNP-hard.

These complexity results should be read as a warning, given a few provisos. If finding an ESS for a class of games is NP-hard, then it is unlikely that that a finite population of individuals obeying a simple learning rule will converge to it. However, this observation does not mean that all hope is lost. It simply means that we should check the computationally tractability when using finite populations to model situations. Moreover, the results found in [**Nis07**] do not directly imply that an *infinite* population cannot converge to an equilibrium. In fact, the replicator dynamic introduced in the next section explains one simple learning mechanism that does converge to ESS.

2. The Replicator Dynamic

Behaviors spread by imitation. Here, imitation means either copying another individual's actions or actively deciding to maintain your own. Individuals in evolutionary game theoretic populations can update strategies from time to time. Individuals will tend to adopt strategies that award them better than average payoffs and eschew those that do worse. In this model, we assume that individuals only compare their current payoff with the value of a neighbor's current payoff and do not track how well they have done over time. Our individuals are forgetful, living only in the here and now. As a population evolves, the frequencies of its profile will change accordingly.

To model the evolution of a general *idealized population*, we will make two assumptions: Idealized population

(1) The population contains an infinite number of individuals.
(2) The population is *well-mixed*; that is, every individual is equally likely to interact with every other individual. Well-mixed population

Idealized populations are very large but lack a sense of geography. Possible real-world examples might include people who have access to a telephone, email, or an internet-accessible computer. Everyone is located close to each other in a well-mixed population. In later sections we will consider more structured populations. We work with very large populations so that we may treat frequencies as probabilities. Combined with the continuous blending of a well-mixed population, we may assume that the strategy profile of an idealized population evolves as a differentiable function of time.

Let $p_i(t)$ denote the frequency of individuals playing strategy s_i at a time t. The success of s_i is the difference between the expected payoff of s_i in the current population $\pi(s_i, \sigma(t))$ and the average success $\bar{\pi}(\sigma(t))$.[2] The average success is calculated by summing individual payoffs weighted by their presence in the population. Average success

$$\bar{\pi}(\sigma) = \sum_{i=1}^{n} p_i \pi(s_i, \sigma) = \sum_{i=1}^{n} p_i \langle s_i, \pi \cdot \sigma \rangle = \langle \sigma, \pi \cdot \sigma \rangle = \pi(\sigma, \sigma).$$

Then the frequency p_i of an individual strategy s_i will change according to its current presence scaled by its relative success in the population. The resulting dynamic

$$(\star\star) \qquad \dot{p}_i = p_i \left[\pi(s_i, \sigma) - \bar{\pi}(\sigma)\right]$$

is the so-called *replicator dynamic*. It is a dynamical systems description of the learning rule "imitate the better strategy proportionally to the expected gain." Replicator dynamic

In order to understand the behavior of the replicator dynamics, we will investigate the long-term, rest point behavior of the replicator dynamic. In particular, the rest points of ($\star\star$) satisfy

$$\pi(s_i, \sigma) - \bar{\pi}(\sigma) = 0.$$

By expanding the term for average success, we reach an equivalent but familiar requirement.

$$\pi(s_i, \sigma) = \pi(\sigma, \sigma).$$

Therefore we have proved the following theorem.

[2]For notational convenience and clarity I will usually suppress direct dependence on the time variable t.

THEOREM 2.1. *If the population profile σ is a Nash equilibrium of the game described by the payoff matrix π, then it is a rest point of the replicator dynamic.*

Theorem 2.1 links Nash equilibria with the rest points of a dynamical system. It shows that a population which as a group has discovered a Nash equilibrium will not stray from it. However, being a Nash equilibrium does not necessarily offer protection from a small invasion of dissident individuals. If mutations are introduced, the innovative behavior may spread. Fortunately, we can prove another result about regarding convergence under the replicator dynamic.

Asymptotically stable rest point

Before stating and proving the theorem, we will need to take a short detour into the field of dynamical systems theory to pick up a few useful concepts. A rest point p of a dynamical system is *asymptotically stable* if all of the points in a neighborhood of p eventually converge to it under the system dynamics. The asymptotic features of a dynamic are encapsulated in the so-called ω-limit. Let $\dot{x} = f(x)$ be a time-independent ordinary differential equation (ODE) and $x(t)$ be a solution defined for all time $t \geq 0$. The *ω-limit* of the dynamical system is the collection of accumulation points

ω-limit

$$\omega(x, f) = \{p \,|\, x(t_k) \to p \text{ for some sequence } t_k \to \infty\}.$$

Even though the existence and uniqueness theorems certify that a solution to a system of ODEs exists, it is generally difficulty to produce one explicitly.

The set $\omega(x, f)$ is an invariant of the dynamical system. Moreover, rest points form their own ω-limits. It is possible to prove general properties about ω-limits even if we cannot explicitly compute them. The following theorem (stated without proof) gives sufficient criteria for asymptotic stability. It is a cornerstone in the qualitative study of dynamical systems theory.

THEOREM 2.2 (Lyapunov Stability Theorem [**Dev89**]). *Let $\dot{x} = f(x)$ be a time-independent ODE defined on some subset $U \subset \mathbb{R}^n$. Let $V : U \to \mathbb{R}$ be a smooth function. If for some solution $x(t)$, the inequality $\dot{V}(x(t)) \geq 0$ holds, then $\omega(x, f) \cap U$ is contained in the set $\{p \in U \,|\, \dot{V}(p) = 0\}$ of critical points of V.*

Lyapunov function

The stability theorem formalizes a commonplace observation. A rock rolling down a hill must settle somewhere at the bottom. The function V in the theorem is called a *Lyapunov function*. It captures the energy dissipating from the system. When the rock runs out of energy, it comes to rest. The ω-limit represents the resting points. Unfortunately the stability theorem does not give any hint as how to construct such a function. Worse still, there is no general prescription. When a Lyapunov function can be found, though, the stability theorem supplies a strong statement about the asymptotic behavior of a system. Such is the case for the replicator dynamic.

THEOREM 2.3. *If the strategy profile σ is an ESS of the game described by the payoff matrix π, then it is an asymptotically stable rest point of the replicator dynamic.*

PROOF. Let $\sigma = (p_1, \ldots, p_n)$ be the frequencies of σ. Then assemble the real-valued function of strategy profiles

$$P(x) = \prod_{i=1}^{n} x_i^{p_i}.$$

Postpone as a lemma the fact that P achieves a global maximum at the point σ. Moreover, $P(x)$ is non-zero only for those points x whose support is the same as the support of σ. If $\tau = (q_1, \ldots, q_n)$ is a population profile and $P(\tau) > 0$, then

$$\frac{\dot{P}}{P}(\tau) = \frac{d}{dt}(\log P)(\tau) = \frac{d}{dt}\left(\sum_{i=1}^{n} p_i \log q_i\right) = \sum_{p_i>0} p_i \frac{\dot{q}_i}{q_i} = \sum_{p_i>0} p_i[\pi(s_i,\tau) - \pi(\tau,\tau)]$$
$$= \pi(\sigma,\tau) - \pi(\tau,\tau).$$

Since σ is evolutionarily stable, it is guarded by a buffer of less successful strategies. Therefore $\dot{P} > 0$ for all τ in a neighborhood of σ. Thus the function P serves as a strict Lyapunov function for the replicator dynamic, and all trajectories near σ converge to it. □

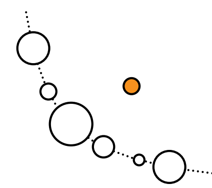

If chunks of a unit mass weighing p_i are placed on the convex curve at $(x_i, f(x_i))$ then their center of mass lies above the curve.

LEMMA 2.1. *The function P of Theorem 2.3 attains a global maximum over strategy profiles at the point σ.*

PROOF. This proof relies on a result that is well-known in the theory of convex functions known as *Jensen's inequality* [**Mac03**]: *if f is a strictly convex function defined on some interval I, then*

$$f\left(\sum_{i=1}^{n}(p_i x_i)\right) \le \sum_{i=1}^{n} p_i f(x_i),$$

for all $x_1, \ldots, x_n \in I$ and every probability weight vector $p = (p_1, \ldots, p_n)$, with equality if and only if all x_i coincide.

Apply Jensen's inequality to $-\log(x)$ on the non-zero real-line, setting $0 \log 0 = 0 \log \infty = 0$. Trading in a minus sign at the cost of a less-than sign, we obtain

$$\log\left(\prod_{i=1}^{n} x_i^{p_i}\right) = \sum_{i=1}^{n} p_i \log x_i \le \log\left(\sum_{i=1}^{n} p_i x_i\right).$$

To see how a general profile $\tau = (q_1, \ldots, q_n)$ compares against $\sigma = (p_1, \ldots, p_n)$, we calculate the difference

Keep in mind that these are the same p_i from the definition of P.

$$\log P(\tau) - \log P(\sigma) = \sum_{i=1}^{n} p_i \log q_i - \sum_{i=1}^{n} p_i \log p_i = \sum_{i=1}^{n} p_i \log \frac{q_i}{p_i} = \sum_{p_i>0} p_i \log \frac{q_i}{p_i}$$
$$\le \log\left(\sum_{p_i>0} q_i\right) \le \log\left(\sum_{i=1}^{n} q_i\right) = \log 1 = 0.$$

Therefore $P(\tau) < P(\sigma)$ unless $\tau = \sigma$. □

As an aside, it is interesting to note that the quantity

$$\sum_{i=1}^{n} p_i \log \frac{q_i}{p_i} = \underbrace{H(\tau,\sigma)}_{\text{Cross entropy}} - \underbrace{H(\sigma)}_{\text{Entropy}}$$
$$= D_{\mathrm{KL}}(\sigma \| \tau)$$

is the *Kullback-Leibler divergence* of information theory. A Lyapunov function can be interpreted as an approximate description of the energy in a physical system that is losing energy without being replaced over time. Still, it is not clear whether the appearance of the Kullback-Leibler divergence is anything more than a coincidence in this context.

Kullback-Leibler divergence

The results in this section can be neatly summarized in what is sometimes dubbed the Folk Theorem of Evolutionary Game theory.

THEOREM 2.4 (Folk Theorem of Evolutionary Game Theory). *Given a population playing a symmetric, two-player game evolved by the replicator dynamic the following statements are true.*

(1) *If a rest point is stable, then it is a Nash equilibrium.*
(2) *If the replicator dynamic converges to totally mixed strategy σ rest point, it is a Nash equilibrium.*
(3) *A strict Nash equilibrium is asymptotically stable.*

PROOF.

(1) Suppose that σ is stable rest point but not a Nash equilibrium. Then for all τ in a neighborhood about σ there exists a pure strategy s_i that does better against τ than τ does itself But then the frequency of s_i would increase, which contradicts the stability of σ.
(2) This is the conclusion of Corollary 1.1.
(3) A strict Nash equilibrium is also an ESS. Therefore, Theorem 2.3 applies.

□

The strongest case for evolutionary stability—and arguably most important result of this section—though, comes from as a straight-forward corollary of Theorem 2.3. Namely, a mix of individuals learning by imitation will converge on an ESS.

THEOREM 2.5. *If the totally mixed strategy σ is evolutionarily stable, then it is a globally stable rest point of the replicator dynamic in the sense that the trajectory of every totally mixed strategy converges to it.*

PROOF. $\dot{P}(\tau) > 0$ for all totally mixed strategies τ. □

3. Agent-based approximations

The replicator dynamic is a satisfying model of behavior because it shows that an uncoordinated collection of individuals can settle on a the same Nash equilibria that economists suppose that rational players would choose. Yet the replicator dynamic does so with only a modest set of assumptions based on imitative learning. It dispenses with the need for infinite computational power or time. Moreover, translating game theoretic concepts into the language of dynamical systems theory put problems like the social dilemma mentioned at the end of the previous chapter in a form that is vulnerable to investigation by the techniques in another well-developed field.

Yet we began our discussion of evolutionary game theory in terms of individuals. The replicator has abstracted away the individual players in favor of frequencies. In this section we aim to reintroduce the individual agents explicitly. Following [**San08**], we will develop a dynamic model of choice that can be easily implemented in any modern programming language. In particular, the agents in our model will act more or less independently of each other, leaving ample room for parallelization. This model is not meant to replace the replicator dynamic. In fact, we will show that these two seemingly distinct approaches really do arrive at the same result.

Imagine a large population of N individuals, each playing one of the pure strategies in the symmetric, two-player game represented by the payoff matrix π. Each individual comes equipped with a "stochastic alarm clock" that periodically rings. Once an alarm goes off, the owner of the clock has the option to revise her strategy. The method ρ that all individuals use to update their strategies is called a *revision protocol*. The time between rings for any one clock is, on average, constant. Therefore alarm times are drawn from an exponential distribution with ringing rate $R > 0$. Usually, we set $R = 1$. All of the clocks run and ring independently of one another. These special devices are sometimes called *Poisson alarm clocks* and are common in stochastic models.

Revision protocol

Changing the clock rate so that $R \neq 1$ only amounts to scaling the switching probabilities by a factor of $1/R$.

A revision protocol is simply a procedure an individual follows to transition from one strategy to another. In our model, individuals will use imitation to motivate a change in their behavior. This framework is slightly more general and could accommodate other reasons to switch strategies; *e.g.,* dissatisfaction, boredom, or more sophisticated models of learning. Typically an individual playing strategy s_i will observe how well her own strategy is doing compared to other strategies available to her. In particular, the *conditional switching probability* from strategy s_i to s_j is given by $\rho_{ij}(\pi, \sigma)$, where σ is the current profile of the population. To fix ρ_{ij} as an honest probability distribution, we require that

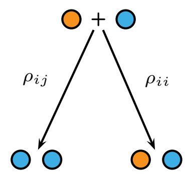

Conditional switching probability

$$\sum_j \rho_{ij}(\pi, \sigma) = 1.$$

The Poisson alarm clocks describe a *Poisson process*—a situation in which events occur continuously and independently at a constant average rate [**GS01**]. In this case, each event is an opportunity for an individual to switch strategies. Therefore the number of switching opportunities in the next dt units of time follows a Poisson distribution with mean $R\,dt$. The number of individuals playing s_i to receive revision opportunities during the next dt time units is approximately

Poisson process

$$Np_iR\,dt,$$

where p_i represents the frequency of i-strategists in σ. The above expression is only approximate because the value of p_i may change during $[0, dt]$, but if dt is very small then p_i will likely remain constant. The number who switch to strategy s_j is then about

$$Np_i\rho_{ij}\,dt.$$

The over-all change in the use of s_i during the next dt time units is roughly

$$N\left(\sum_j p_j\rho_{ji} - p_i\sum_j \rho_{ij}\right)dt.$$

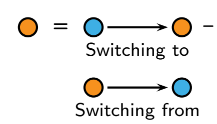

Dividing by N yields the expected change in the frequency p_i of individuals choosing strategy s_i. In the limit that $dt \to 0$, we obtain a system of differential equations collectively referred to as the *mean dynamic* of the revision protocol ρ.

Mean dynamic

$$\dot{p}_i = \sum_j p_j\rho_{ji} - p_i\sum_j \rho_{ij}.$$

Having developed a general framework, we need only to plug in a specific revision protocol. There are many ways to formulate the pairwise proportional imitation mentioned throughout. For the class of *imitative protocols*

Imitative protocol

$$\rho_{ij}(\pi, \sigma) = p_jr_{ij}(\pi(\sigma))$$

an individual playing s_i asks another playing s_j how well her strategy is presently doing as measured by his payoff. Then the first individual takes the advice to switch to strategy s_j with probability proportional to r_{ij}. The recommendation weight r_{ij}

is not directly a function of the current population profile. The factor p_j takes into account how likely a chance meeting between s_i and s_j occurs.

As we have assumed before, suppose that after a random pairing with another player, an individual decides to switch proportionally to the difference of their payoffs. Then we define

$$\rho_{ij}(\pi, \sigma) = p_j H[\pi(s_j, \sigma) - \pi(s_i, \sigma)],$$

where H is the Heaviside step function

$$H(x) = \begin{cases} 1 & x > 0 \\ \frac{1}{2} & x = 0 \\ 0 & x < 0. \end{cases}$$

The mean dynamic for this protocol is then

$$\begin{aligned} \dot{p}_i &= \sum_j p_j \rho_{ji} - p_i \sum_j \rho_{ij} \\ &= \sum_j p_j p_i H[\pi(s_i, \sigma) - \pi(s_j, \sigma)] - p_i \sum_j p_j H[\pi(s_j, \sigma) - \pi(s_i, \sigma)] \\ &= p_i \sum_j p_j [\pi(s_i, \sigma) - \pi(s_j, \sigma)]. \end{aligned}$$

That is, the mean dynamic of pairwise, proportional imitation is nothing more than the replicator dynamic,

$$\dot{p}_i = p_i[\pi(s_i, \sigma) - \bar{\pi}(\sigma)].$$

This result shows that agent-based simulations will converge to the replicator dynamic. As a bonus, the model itself gives a straight-forward implementation. The one given below relies on a sigmoidal function to calculate the switching probability. The motivation for this particular function is given an appendix. It appears throughout mathematics, computer science, and the physical sciences. It has the same qualitative form as a step function and offers the added convenience of differentiability.

Of practical importance is the function `Mutate`. Recall that the replicator dynamic will hone in on an interior ESS only if the population is totally mixed. Mutations introduce innovation to population. They broaden the sweep of the search and help the dynamic escape local optima when it has narrowed its scope prematurely. Since evolutionary stability can resist invasion, mutations are generally quite helpful. That said, it is important to keep the rate and effect of mutations relatively low. Mutation should occur fairly infrequently, and when they do, they should test nearby strategies. If either values are too high, the dynamic can cause the population to oscillate wildly.

Structured Populations. The property NEIGHBORS mentioned in the algorithm above is also of crucial importance. It defines the topology of the population. The standard replicator dynamic encodes a well-mixing assumption—that is, any individual can be pairwise with any other individual with equal probability. Most observed network structures are not well-mixed. Friendship networks tend to be local, for example. Most of you friends likely live in the same part of the country that you live in; moreover, most of your friends live in the same country you do. Other networks, ranging from cellular regulatory pathways to highway systems to

Data: A normal form game π and a strategy switching parameter *laziness*.
Results: A population evolved by the replicator dynamic.

▷ **Initialize the population.**
For each *individual* in *population* do
  $\textsc{Strategy}[individual] \leftarrow$ `Choose`$(\textsc{Strategies}[\pi])$
For each *individual* in *population* do
  $other \leftarrow$ `Choose`$(\textsc{Neighbors}[individual])$
  $\textsc{Payoff}[individual] \leftarrow \pi(individual, other)$

▷ **Evolve the population.**
Until $\langle stopping\ condition \rangle$ do
  $individual \leftarrow$ `Choose`$(\textsc{Individuals}[population])$
  $other \leftarrow$ `Choose`$(\textsc{Neighbors}[individual])$

  ▷ **Sigmoidal imitative learning.**
  $gain \leftarrow \textsc{Payoff}[individual] - \textsc{Payoff}[other]$
  $gumption \leftarrow 1/(1 + \exp(laziness \cdot gain))$
  If `Random()` $\leq gumption$ then do
    $\textsc{Strategy}[individual] \leftarrow$ `Mutate`$(\textsc{Strategy}[other])$
  $opponent \leftarrow$ `Choose`$(\textsc{Neighbors}[individual])$
  $\textsc{Payoff}[individual] \leftarrow \pi(individual, opponent)$

Algorithm. Pseudocode for the stochastic model introduced in this section.

the nation's electrical grid and the world-wide web show a type of preferential attachment. Some individuals are more connected than others. In these examples a few hubs dominate in number of connections. The implementation above provides space to test game dynamics in these structured populations.

Structured population

In general a *structured population* is graph $G(V, E)$. The vertices of the graph represent the individuals of the population. They carry two labels, a the individual's current strategy and the individual's most recent payoff. Individuals interact only with neighboring individuals connected by an edge in the graph. Populations with a complete graph structure are the well-mixed populations introduced in the original formulation of the replicator dynamic. The behavior of games on graphs is difficult and few general results are known ([**ON06**], [**ON08**], and [**SF07**]). The small number of neighboring individuals makes structured populations vulnerable to sampling bias. The fewer neighbors an individual has, the more influential the adjacent nodes become.

An evolutionary reformulation of the old adage: Don't put all your eggs in one basket.

Simulation offers a good approach to explore the effects structure has on the dynamics of a game. They provide a lab for experimentation and reflection. It was once commonly believed that structured populations offered pockets of protection that could nurture cooperative behavior, for example. In the classic Traveler's Dilemma, structured populations actually accelerate undercutting behavior. Simulations can help confirm intuition, suggest new directions for theory to explore, and help to distill and explain phenomena that we observe. In the remainder of this thesis, we see how theory and simulation, when used together, can deliver a sharper look into social dilemmas than either can on its own.

CHAPTER 3

Social Dilemmas Redux

Even with the machinery we have developed so far, we cannot analyze the social dilemmas mentioned at the end of the first chapter. The populations from the previous chapter are comprised of n distinct and discrete types. Many behaviors—like the amount of energy to invest in a group project, for example—are best categorized as continuous behaviors. The players minimum effort coordination game assume strategies that correspond to real numbers. In order to cope with continuous strategies, we will need to generalize evolutionary game theory, and the the replicator dynamic specifically, further still.

Moreover, the the replicator dynamic alone does not resolve the traveler's dilemma. Populations playing TD evolve to the unique but faulty Nash equilibrium. Both in one-shot ([**GH01**], [**BCN05**]) and repeated games [**CGGH99**] with human subjects, people will play fast and loose, offering high bids, when the risk of punishment is low. They tend to adopt more conservative, small bid strategies when the risk of penalty is high. While this sort of behavior seems obvious when cast in terms of risk, it is not easy to reproduce it within the realm of game theory. Common ways of explaining away the deviation between observed and idealized actions are unsatisfying. These arguments often appeal to the limited cognitive abilities of human players to perform the backwards induction argument presented in chapter one successfully.

In this chapter we will put forward another, perhaps more intuitive explanation that solves the traveler's dilemma and the minimum effort coordination game simultaneously. The technique relies on a simple observation: people act as if cause and effect is a continuous relationship [**Mea08**]—that a small change in the cause will result in a small change in the effect. If I work a little harder while running, my speed will increase a little bit. If I plant a few extra vegetables in my garden, I will get a few extra vegetables in my harvest. The payoff functions in the traveler's dilemma and the minimum effort coordination game (and a host of other social dilemmas) are exhibit discontinuous behavior. As stated, the payoffs of these games are, in a sense, 'unpsychological.' To settle the dilemmas, we will remove the discontinuities from the payoff functions as unobtrusively as possible.

1. Adaptive Dynamics

For this model of evolution, imagine a population of that is mostly homogenous. All individuals employ the same (continuously varying) strategy x, except for a small mutant community using a similar but distinct strategy $y = x + h$, where h is a small displacement in behavior. As per usual, the members of the population all interact in pairwise competition in the game π. This time, however, π is any differentiable function of two real-valued variables. The mutant strategy y has a *relative advantage* $A_x(y)$ in a population of x-strategists given by the difference in

Relative advantage

payoffs

$$A_x(y) = \pi(y, x) - \pi(x, x).$$

Not all mutant strategies will fare better than the resident population. Consequently the relative advantage can assume negative, positive and zero values.

Adaptive dynamics

Adaptive dynamics includes the replicator dynamic as a special case. To see so requires altering the geometry of the strategy space and would take us too far afield. See, for example, [**HS98**].

The *adaptive dynamics* of π describes evolution from the point of view of a mutant individual. This dynamic points in the direction that is most promising for an innovative individual testing new behavior against the backdrop of a uniformly acting group. When the strategy x is represented by a single real number, the direction of greatest increase in relative advantage is just the derivative, and the adaptive dynamics is

$$\dot{x} = \lim_{h \to 0} \frac{\pi(x+h, x) - \pi(x, x)}{h} = \frac{d}{dy} A_x(y) \bigg|_{y=0}.$$

For multi-trait behaviors, the gradient of the function sending y to its payoff $\pi(y, x)$ points in the direction of greatest increase in relative advantage. The adaptive dynamics in this case is

$$(\star) \qquad \dot{x}_i = \frac{\partial}{\partial y_i} \pi(y, x) \bigg|_{y=x}.$$

Adaptive dynamics portrays the change in behavior of opportunistic individuals who enjoy the safety of the status quo. They want to improve their lot, but are only willing to change their behavior incrementally. These individuals cannot step back and look at the big picture. Instead they go with what works best now, given the present behavioral climate. When individuals act independently of the rest of the population—when π is a function only of y, adaptive dynamics reduces to a group of hill climbers following the path of steepest ascent in search of a local optimum.

Selection gradient

Evolutionarily singular strategy

The right-hand side of ($\star$) is called the *selection gradient* $D(x)$. The vanishing points of $D(x)$ are termed *evolutionarily singular strategies*. As π is just a smooth function, the zeroes of its gradient represent local extrema. The maxima are strategies that resist invasion and are nothing more than the evolutionarily stable strategies from before. At minima, the population can split into two distinct behavioral groups, giving rise to so-called *evolutionary branching*.

As an example of branching, adaptive dynamics has been used to analyze the Continuous Snowdrift Game. In this game, two very tired motorists passing through a frozen, remote New England town both run into opposite sides of a giant snow drift during their commute home. Both commuters would like to get home, but neither would like to get out and clear the snow from the road in the bitter winter cold. Let x, y denote the amount of work each puts in to shovel the snow drift. Then the payoff the x-strategist receives is the combined benefit that she and other driver put in $b(x + y)$, but still at a personal cost $c(x)$. The payoff function is then

$$\pi(x, y) = b(x + y) - c(x),$$

for smooth functions b and c and efforts x, $y \in [0, 1]$.

Initially a population playing the continuous snowdrift game will converge to a singular strategy x^* that lies somewhere between 0 and 1. At that point, though, the population reaches a payoff minimum and branching occurs. The final population consists of a large number of unconditionally lazy individuals who sit back, crank

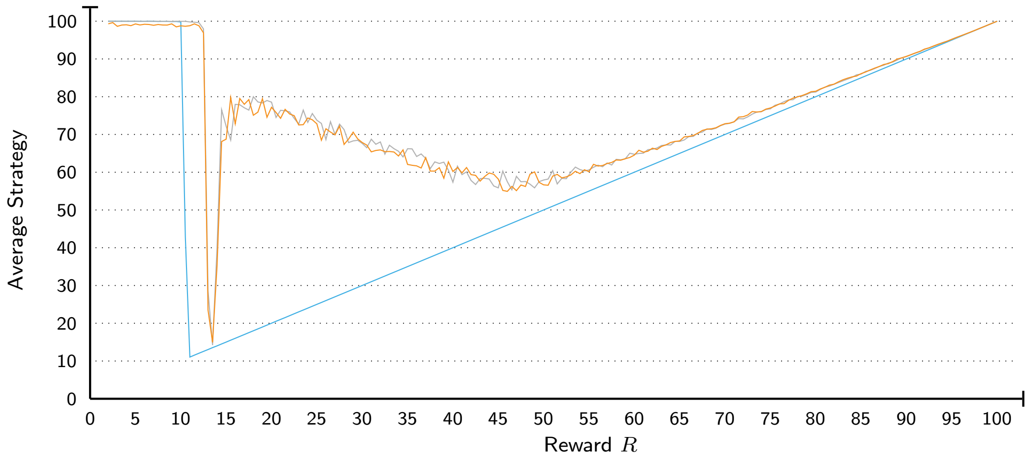

FIGURE 1. A plot of the average strategy versus reward/punishment R for populations of size $n = 1000$ with well-mixed (blue), regular graph of degree 4 (gray), and Erdős-Renyi random graph of average degree 4 (orange) structures.

the heat, and listen to the radio as the rest of the group shovels away the impasse. The adaptive dynamics for this game gives a very simple mechanism that reproduces the Pareto Principle, sometimes called the 20–80 Rule [DHK04]:

Twenty percent of the people do eighty percent of the work.

2. The Adaptive Program: Smoothing

Many of the classic social dilemmas in the economic literature include a step function in their payoff. This category of game includes the traveler's dilemma and the minimum effort coordination game. The Nash equilibria of such games often propose cold, calculated and sometimes counter-intuitive behaviors.

We believe that ineffectiveness of the Nash equilibrium in these cases stems from two facts: first, the classical analysis is the same regardless of the personal cost or punishment involved. To wit, human subjects playing the traveler's dilemma will bid liberally when the potential punishment is low and play conservatively when the punishment is high. Second, people are unlikely to calculate risk in absolute terms. Instead, individuals use a smoothed, expected risk. In this section, we will use adaptive dynamics to analyze a broad class of smoothed symmetric two-person games with continuous strategies.

Consider a game with payoff function π of the form

$$\pi(x,y) = \begin{cases} m(x,y) & x < y \\ \frac{1}{2}[m(x,y) + M(x,y)] & x = y \\ M(x,y) & x > y. \end{cases}$$

for affine functions $m(x,y)$ and $M(x,y)$. We may write $\pi(x,y)$ more succinctly with the aid of the Heaviside step function $H(t)$ as

$$\pi(x,y) = \underbrace{m(x,y)H(y-x)}_{x<y} + \underbrace{M(x,y)H(x-y)}_{x>y}.$$

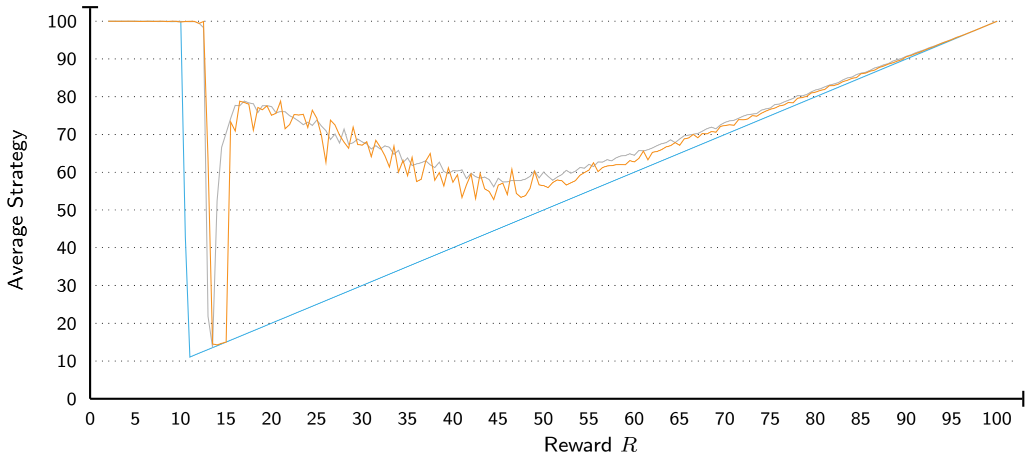

FIGURE 2. A plot of the average strategy versus reward/punishment R for populations of size $n = 1000$ with well-mixed (blue), a 32×32 rectangular lattice (gray), and scale-free graph of average degree 4 (orange) structures.

The step function is non-differentiable at $t = 0$. We will approximate the step function with any smooth real-valued function $\theta(t)$ such that

$$\lim_{t\to-\infty} \theta(t) = 0, \lim_{t\to\infty} \theta(t) = 1, \text{ and } \theta(0) = \tfrac{1}{2}.$$

The selection gradient for the smoothed version of the game is

$$D(x) = \theta'(0)[(M(x,x) - m(x,x)] + \frac{1}{2}\frac{\partial}{\partial y}\Big[M(y,x) + m(y,x)\Big]\Big|_{y=x}.$$

Almost miraculously, the selection gradient depends on the choice of smoothing θ and its derivative only at a single point. The first term in the selection gradient measures how much a player would stand to gain (or lose) by nonconforming weighted by a perceived risk factor. The risk factor $\theta'(0)$ describes how related one situation $(x < y)$ is with the other $(x > y)$. It is an assessment of uncertainty. The second term is the derivative of success to behavior that conforms to the norm of the resident population.

Travelers Dilemma. In the traveler's dilemma, the payoff functions are $m(x,y) = x + R$ and $M(x,y) = y - R$. The adaptive dynamics for this game is therefore

$$\dot{x} = \frac{1}{2} - 2R\theta'(0).$$

The resulting selection gradient is nowhere vanishing. Therefore there are no evolutionarily singular strategies to investigate in this game. Under imitative learning, the population should either increase to the maximum or decrease minimum available bids, depending on whether the sign of the gradient is positive or negative. This analysis is in accord with experimental results. Figures 1 and 2 show the average strategy of populations evolved playing the traveler's dilemma with a varying penalty/reward parameter. When the penalty is low, bids are high; when the penalty is high, bids are low.

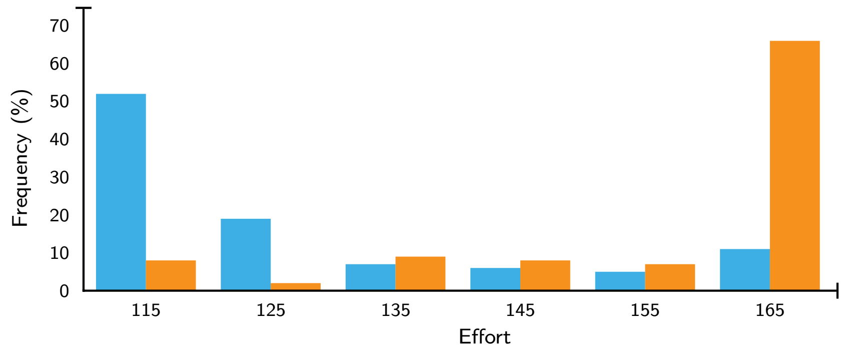

FIGURE 3. Effort choice frequencies for a minimum effort coordination game with costs $k = 0.9$ (blue) and $k = 0.1$ (orange). (*Adapted from* [**GH01**].)

Minimum Coordination Games. Here the payoff functions are given by $m(x, y) = x - kx$ and $M(x, y) = y - kx$. The smoothing θ cancels out entirely, leaving the surprising adaptive dynamics

$$\dot{x} = \frac{1}{2} - k.$$

This time the adaptive framework singles out one of the infinite otherwise indistinguishable Nash equilibria for fixed k. When k is greater than one half, the population evolves to the maximum strategy value 1. If k is less than one half, the population moves toward 0. However, when k is precisely equal to one half, the model is again entirely silent and all strategies are singular. A small perturbation, though, will send the population hurdling toward one of the extreme values.

Agent-based simulations are in good accord with the predictions of the adaptive dynamical model, even for modest population sizes of 500. These simulations reproduce human behavior. When the cost of effort is high, the risk of over-shooting is high and the game dynamics favor low level of effort. Similarly, when the cost is low, human players. Figure 3 sketches the inverse relationship between cost and effort predicted by the framework in real, human trials. Figures 4–6 show the average strategy in populations evolved by playing the minimum effort coordination game with varying costs. Again, the low costs give rise to high, coordinated levels of effort. High costs drive down efforts to zero.

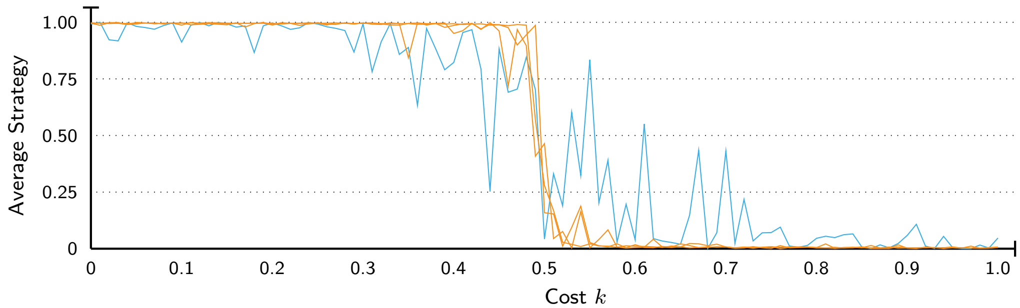

Figure 4. A plot of average strategy versus cost k for well-mixed populations of size $n = 100$ (blue), 500, 1000, and 5000 (orange) playing the minimum effort coordination game.

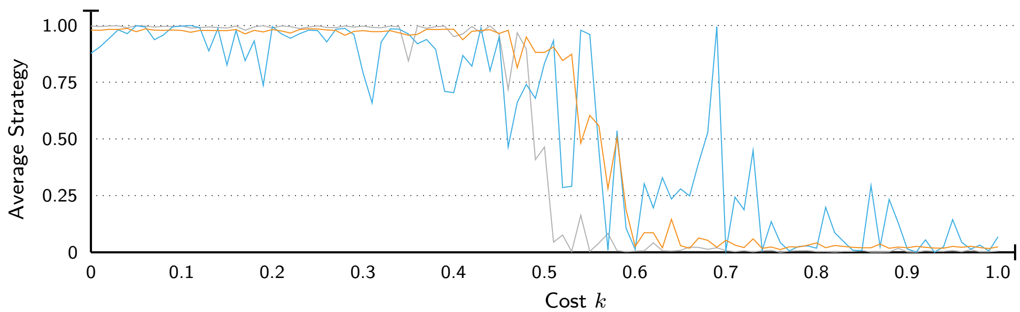

Figure 5. A plot of the average strategy versus cost k for populations of size 500 with well-mixed (gray), scale-free graph with average degree 4 (blue), and classical Erdős-Renyi random graph with average degree 4 (orange) structures.

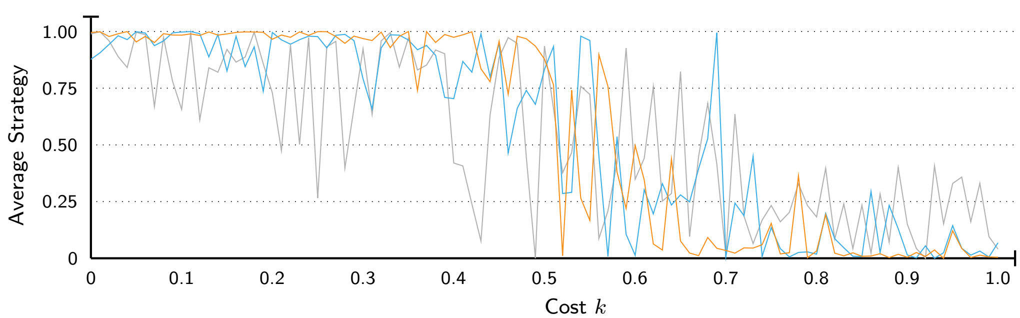

Figure 6. Scale-free populations of average degree 4 and sizes $n = 100$ (gray), 500 (blue), and 1000 (orange).

Bibliography

[All94] L.V. Allis, *Searching for solutions in games and artificial intelligence*, Ph.D. thesis, Universiteit Maastricht, 1994.

[Bas94] K. Basu, *The traveler's dilemma: Paradoxes of rationality in game theory*, Papers and Proceedings of the Hundred and Sixth Annual Meeting of the American Economic Association, vol. 84, The American Ecnomic Review, American Economic Association, May 1994, pp. 391–395.

[Bas07] ———, *The traveler's dilemma*, Scientific American Magazine **296** (2007), no. 6, 90–95.

[BCN05] Tilman Becker, Michael Carter, and Jörg Naeve, *Experts playing the traveler's dilemma*, Diskussionsbeiträge aus dem Institut für Volkswirtschaftslehre, Universität Hohenheim **252** (2005).

[CGGH99] C. Monica Capra, Jacob K. Goeree, Rosario Gomez, and Charles A. Holt, *Anomalous behvior in a traveler's dilemma?*, American Economic Review **89** (1999), no. 3, 678–690.

[Cre03] R. Cressman, *Evolutionary dynamics and extensive form games*, MIT Press, 2003.

[Dal73] H.E. Daly, *Toward a steady-state economy*, WH Freeman & Company, 1973.

[Daw72] R.M. Dawes, *Fundamentals of attitude measurement*, Wiley New York, 1972.

[Dev89] R.L. Devaney, *An introduction to chaotic dynamical systems*, 2nd ed., Perseus Press, 1989.

[DGP06] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou, *The complexity of computing a nash equilibrium*, STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing (New York, NY, USA), ACM, 2006, pp. 71–78.

[DHK04] M. Doebeli, C. Hauert, and T. Killingback, *The evolutionary origin of cooperators and defectors*, 2004, pp. 859–862.

[GH01] J.K. Goeree and C.A. Holt, *Ten little treasures of game theory and ten intuitive contradictions*, American Economic Review (2001), 1402–1422.

[Gib92] R. Gibbons, *Game theory for applied economists*, Princeton University Press Princeton, NJ, 1992.

[GS01] G. Grimmett and D. Stirzaker, *Probability and random processes*, Oxford University Press, USA, 2001.

[GZ89] I. Gilboa and E. Zemel, *Nash and correlated equilibria: Some complexity considerations*, Games and Economic Behavior **1** (1989), no. 1, 80–93.

[HS98] J. Hofbauer and K. Sigmund, *Evolutionary games and population dynamics*, Cambridge Univ Press, June 1998.

[Mac03] D.J.C. MacKay, *Information theory, inference, and learning algorithms*, Cambridge Univ Pr, 2003.

[Mea08] D. Meadows, *Thinking in systems: A primer*, Chelsea Green Publishing Company, December 2008.

[Nis07] N. Nisan, *Algorithmic game theory*, Cambridge Univ Pr, 2007.

[Now06] M.A. Nowak, *Evolutionary dynamics: Exploring the Equations of Life*, Belknap Press of Harvard University Press, September 2006.

[ON06] Hisashi Ohtsuki and Martin A. Nowak, *The replicator equation on graphs*, Journal of Theoretical Biology **243** (2006), 86–97.

[ON08] ———, *Evolutionary stability on graphs*, Journal of Theoretical Biology **251** (2008), 698–707.

[Pap07] C.H. Papadimitriou, *The complexity of finding nash equilibria*, Algorithmic Game Theory (2007), 29–51.

[Pou93] W. Poundstone, *Prisoner's dilemma*, Doubleday New York, NY, USA, 1993.

[San08] W.H. Sandholm, *Population games and evolutionary dynamics*, Unpublished manuscript, University of Wisconsin, June 2008.

[SF07] G. Szabó and G. Fath, *Evolutionary games on graphs*, Physics Reports **446** (2007), no. 4-6, 97–216.

[Wut93] R. Wuthnow, *Acts of compassion: Caring for others and helping ourselves*, Princeton Univ Pr, 1993.