1 Introduction

1.1 Rock, Paper, Scissors (Lizard, Spock)

In 1996, Barry Sinervo and Curtis Lively, two scientists from the Department of Biology and the Center for the Integrative Study of Animal Behavior at Indiana University, published a paper describing surprising population behaviour concerning the species *Uta stansburiana*, a.k.a. the common side-blotched lizard. In this species of lizard, males occur in three different behavioural types, identified by coloured blotches on their necks. The first type, with an orange throat, aggressively defends large territories containing multiple female lizards. The second type, with a dark blue throat, differs in that it is both less aggressive and prone to defending smaller territories containing fewer female lizards. The third type of male, with a yellow throat, is visually similar to females and – importantly – does not defend any territory at all.

What Sinervo and Lively found was that, over time, these three types of males exhibited an interesting pattern of variation in how frequent each type was in the population. Initially, the orange-throated males increased in number: aggressively defending large swathes of territory enabled them to mate with more female lizards, resulting in their having more offspring than other types. Eventually, though, a tipping point was reached. Defending larger territories meant that a single orange-throated male had to divide his time policing a wider area, and was not always able to prevent the yellow-throated males (which resembled females, remember), from invading their space and mating with the female lizards. This led to the yellow-throated type increasing in number. However, after some time yet another tipping point was reached. The yellow-throated type became vulnerable to the blue-throated type. Why? The fact that the blue-throated type would aggressively defend a small territory mean that they were able to prevent the yellow-throated type from sneaking in. This gave the blue-throats a fitness advantage, causing their type to increase in number. However, once the blue-throated types were populous enough, their less aggressive nature made them vulnerable to the orange-throated type, who would expand their territory into areas previously occupied by the blue-throats. And, then, the cycle would begin again.

Now consider one of the first games young children learn to play: Rock-Paper-Scissors. In this game, each child counts "one–two–three" in unison and then makes the shape of either rock, paper, or scissors with one hand. The rules determining the winner are well known: paper covers rock (so paper wins), but rock breaks scissors (so rock wins), and scissors cut paper (so scissors win). The point to note is that no single choice is the best regardless of what the other person chooses: each choice can either win or lose, depending on what the other person picks. Rock-paper-scissors is thus a game of *strategy*, even if not a very interesting one.[1]

Children have been playing rock-paper-scissors for over two thousand years. The first written description of the game dates from around 1600, when the Chinese author Xie Zhoazhi, writing during the period of the Ming dynasty, stated that the game's origins went as far back as the Han dynasty, which spanned from 206 BC to 220 AD. Back then, instead of rock-paper-scissors the objects of choice were frog-snake-slug, but the game was otherwise the same.

When children play rock-paper-scissors, they understand the rules of the game. Each child knows that whether or not they will win depends on both their choice and the other person's choice. In contrast, the common side-blotched lizard does not understand the structure of their reproductive environment. The lizards do not "choose" their throat colour or their behaviour in any way remotely similar to how children choose rock, paper, or scissors. But the *strategic* aspect to what is going on in both cases is essentially the same. What the example of the lizards demonstrates is the central topic with which this Element is concerned: how the interaction between individual behaviours in an evolutionary setting is such that natural selection *poses*, and then *solves*, problems of *strategy* even though none of the creatures involved are *rational*. This is what gave rise to the field known as *evolutionary game theory*.

The name 'evolutionary game theory' is composed of two parts. The root of the expression, 'game theory' refers to the formal study of problems of strategy, a interdisciplinary study spanning mathematics, economics, computer science, and other disciplines. 'Evolutionary' is an adjective which serves to qualify the particular questions and methods one is interested in when studying those problems of strategy. We'll see in a moment how, exactly, the idea of evolution, an idea fundamental to population biology, became intertwined with the study of strategic problems. But for now the following observation may help: in both cases, the idea of "the best thing to do" makes no sense in the absence of further context. The best way for prey to avoid a predator depends on how the predator pursues the prey. The best way to play chess depends on the skill level of your opponents. And what is even more interesting is that sometimes the "rules" of the game can themselves change, like when a predator invents a new method of pursuit, or when the offside rule is introduced in football. Fans of the TV show *The Big Bang Theory* will be familiar with the extended version of

[1] That said, the fact that the World Rock Paper Scissors Association proudly bills itself as a 'professional organization' for Rock-Paper-Scissors players around the world suggests that some people take the game *very* seriously. See https://wrpsa.com.

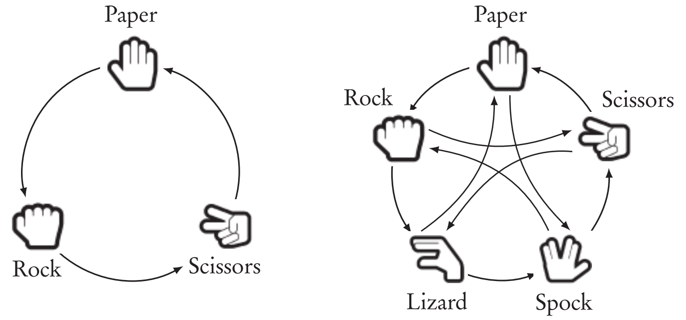

(a) The game of Rock-Paper-Scissors. (b) The game of Rock-Paper-Scissors-Lizard-Spock.

Figure 1 A comparison of the game of Rock-Paper-Scissors (which is found in the animal kingdom) and the extended game of Rock-Paper-Scissors-Lizard-Spock (which is not). An arrow pointing from strategy S_1 to strategy S_2 means that S_1 wins when played against S_2. (Icons courtesy of Font Awesome.)

Rock-Paper-Scissors, originally invented by Sam Kass and Karen Bryla, which adds two additional moves – Lizard and Spock – to the children's game (see Figure 1b). No known species has competitive behaviour which matches the description of this new game, but let's wait and see what evolution produces in the future.

1.2 Game Theory

In 1945, the economist Oskar Morgenstern and the polymath John von Neumann published their seminal book *Theory of Games and Economic Behavior*. The title might strike some as curious: how can a element be both about *games* and *economics*? Surely economics, what Thomas Carlyle called the 'dismal science', is about as far removed from the study of games as possible?

The connection between games and economics derives from an important point about the kinds of choices involved in both. To see this, consider the difference between the kinds of choices a farmer makes when deciding to plant crops, and the kinds of choices a chess player makes when deciding which piece to move. In both cases, there is an optimisation problem. The farmer needs to determine the optimal time to sow the fields, taking into account expectations about future weather. The chess player needs to determine the optimal move to make, given the particular configuration of pieces on the board. But an important difference exists between the two types of optimisation problems. Although we might say that the farmer is trying to 'outwit Nature', that is really just a manner of speaking: Nature does not actually respond to the farmer's actions. Nature does not anticipate that the farmer is going to sow his or her crops at a certain time, and then choose not to rain out of spite (even though it might often feel like that). The weather unfolds in the same way it would have, regardless of what the farmer chooses to do. This is what is known as a problem of *parametric* choice: the farmer is deciding what to do given various parameters, some of which are known for certain, and some of which are either uncertain or unknown altogether. In the case of the chess player, her optimal choice is complicated by the fact that she is interacting with another person. Her opponent responds not only to previous moves as indicated by the position of chess pieces on the board, but to *beliefs* about what is likely to happen in the future. If she sees her opponent make an apparently imprudent move, the thought process that will trigger is readily imagined: ‘Was that a mistake? Or is this an attempt to trick me into making a move whose future consequences I’ve not yet fully considered?’ This is what is known as a problem of *strategic* choice: the choices of each chess player are *interdependent*: what is ‘best’ for one player can only be defined with reference to the choices, plans, and beliefs of the other.

Given that, the connection between the theory of games and economic behaviour should now be clear. Although economic situations are not *games* in the ordinary sense of the term, the multiple interdependencies between buyers and sellers, producers and consumers, and so on, give rise to problems of strategic choice. The best thing for, say, an automobile manufacturer to do depends upon the future demand for their automobiles, which depends on what consumers will want. And what consumers will want depends on what other automobile manufactures may produce, or even on the availability of alternative transportation such as trams or ride-sharing applications that reduce the benefit of owning an automobile. Economic behaviour, on this view, is nothing more than behaviour in a game-theoretic context, given a suitably enlarged conception of what a “game” consists of.

Game theory is the mathematical study of problems of strategic choice. It originated in 1921 with the work of the French mathematician Émile Borel, who analysed the game of poker in order to answer the question of when one should bluff. (As serious poker players know, in order to play poker well you have to bluff.) But, even though his early work was really just at the inception of game theory, Borel had a vision of its potential applications, foreseeing how it could be applied to fields of enquiry far removed from simple parlour games. Despite Borel’s early efforts, the first major theoretical result was due to von Neumann (1928), who proved the influential “Minimax theorem” concerning a special class of two-player games where one player’s gain is another player’s loss. (Such games are known as *zero-sum* games, due to this fact about their payoff structure.) What the Minimax theorem says is that such games always have an optimal course of action for each player such that, if each player follows their respective course of action, they successfully *minimise* the *maximum* loss they might incur. Von Neumann's result was a watershed moment in the development of game theory because it showed that two-player zero-sum games had an effective "solution" guiding the outcome of play. And, given that, the next obvious question to ask was whether the same result, or a similar result, could be shown to hold for other types of games. And, to answer that, we first need to get a bit more precise about the fundamental concepts we have been talking about: what, exactly, we mean by a *game* and a *strategy*.

1.3 Game Theoretic Fundamentals

To begin, let us define a *game* as an interaction between a *finite* and *fixed* number of players, typically denoted by N. This assumption makes sense for many games of strategy like chess (two players), poker (two or more players), and bridge (four players). However, if we consider a "game" such as football, this assumption might seem less appropriate because the number of players on the pitch can vary over time due to penalties, and the *identity* of the players on the pitch can also change over time due to substitutions. In practice this does not present a problem if we conceive of things a bit differently: the total number of players *available* does not change – it corresponds to the complete roster of the team – even though not all players may be *active* at any given time. So the requirement that a game have a finite and fixed number of players raises few practical problems, provided that some flexibility is exercised in the representation.

In a game, each player can choose from one of certain number of *actions*. In the simplest games, like Rock-Paper-Scissors, each player has only a single choice of action and all actions take place simultaneously. But sometimes the actions aren't performed simultaneously, like the game shown in Figure 2a. In that game, each player has a single choice to make, and only if player A chooses A_2 will player B get to perform an action. In more complex games, like Tic-Tac-Toe, the choices the player can make may *take into account the entire history of play up to that point*.[2] A *strategy* is a plan of action that specifies what choice the player will make at every possible decision node they face in the game. Even for a simple game like Tic-Tac-Toe, strategies are vastly complex things. The first player alone has $9 \cdot 7^8 = 51{,}883{,}209$ possible ways to

[2] So, it's not just the state of the board at the time of play that matters, but how the two people *got* to that state. In principle, a strategy could recommend two different moves for the same state of the board, if two different histories of play led to the board looking visually the same.

make their first two moves![3] For a "real" game like chess, a strategy is an almost unbelievably complex object: after the first round of play, there are 400 possible board configurations. After the second round, there are 197,281 possible board configurations. After the third round, there are 119,060,324.

Traditional game theory represents games in two different ways. One representation explicitly tracks every possible way the game can unfold by looking at each possible move available to a player during their turn, and drawing one path for every way the game could be played, from the beginning to the end. At points where a player has a choice to make as to what to do, the path will split according to how many options the player has. This results in a structure known as a *game tree*, because every choice point for a player leads to a "branching" of possibilities (except for the last move of the game), as shown in the simple game of Figure 2a. This is known as the *extensive form representation* of a game. Another representation shows the game as a matrix, with the strategies for each player positioned along one axis and the outcome of the game denoted in the corresponding cell. This is known as the *strategic form representation*, but is also called the *normal form representation*, after von Neumann and Morgenstern, who believed that normally one could adopt this representation without any loss of generality in the subsequent analysis. The strategic form representation is perhaps most natural for games consisting of a single simultaneous move made by each player, like that of Rock-Paper-Scissors shown in Figure 2b, but it turns out that *any* extensive form game can be represented this

[3] This assumes all moves are independent and we do not require consistency for board positions which are strategically equivalent. For example, one strategy for *X* could suggest the following sequence of opening moves:

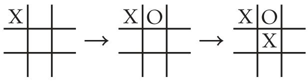

as well as this sequence of opening moves:

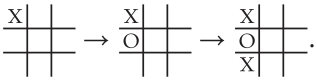

This strategy is peculiar in that it responds differently to *O*'s selection of a center edge position depending on whether it appears on the top edge or on the left edge of the board. Since these two moves, by *O*, are strategically equivalent (they correspond to a reflection of the board along the upper-left to lower-right diagonal, which doesn't fundamentally change the future opportunities available to *O*), the fact that the strategy recommends two strategically different moves suggests a kind of internal inconsistency to the strategy. Nevertheless, strictly speaking, it is permitted according to the given definition of a 'strategy'.

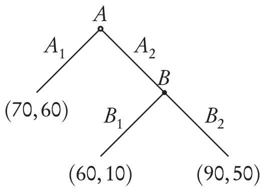

(a) An extensive form representation of a game. Player A moves first, and if she chooses strategy A_2, then player B gets to move. Payoffs are shown with A's payoff listed first.

		Column		
		Rock	Paper	Scissors
	Rock	(0, 0)	(−1, 1)	(1, −1)
Row	Paper	(1, −1)	(0, 0)	(−1, 1)
	Scissors	(−1, 1)	(1, −1)	(0, 0)

(b) The matrix representation of Rock-Paper-Scissors. The payoffs are listed in the form (Row, Column).

Figure 2 Two games illustrating the two different ways they can be represented.

way with a little bit of work.[4] In this Element, most of the time we will consider games shown in the matrix form, but in Section 4 we will consider one extensive form game known as the centipede game.

In game theory, a *solution concept* identifies the strategies which satisfy certain principles of rationality and can arguably be defended as providing a reasonable 'solution' to the game. There are many solution concepts available, but the most important one is the *Nash equilibrium*. A Nash equilibrium occurs when each individual player has settled upon a choice of strategy such that the overall collection of strategies has a certain best-reply property. Suppose that the game under consideration has N players. When each player has settled upon a strategy, the list of strategies used by each player, $\langle S_1, S_2, \ldots S_N \rangle$, is called a *strategy profile*. At a Nash equilibrium, no player has an incentive to change their strategy, provided that everyone else continues to follow the strategy allocated to them by the profile.[5]

---

[4] One simply needs to encode strategies in the right way. If a player has the possibility of making decisions at N nodes in the game tree, then a strategy for that player is a composite object $s_1 s_2 \ldots s_N$ where s_i represents the decision that player will take at node i in the game tree. The strategic form representation can then list all possible strategies for a player down the rows (or across the columns), and the payoff appearing in the corresponding cell will be the payoff realised from the path traced through the game tree by the players' strategies. Notice that this representation of a player's strategy requires specifying what that player would do at nodes in the game tree *which will not be reached* in the actual course of play.

[5] There are two subtleties to bear in mind. First, saying that no player has an incentive to change their strategy does *not* mean that a player would *do worse* if they did. At a Nash equilibrium, a change in strategy may result in a player receiving the same payout that they would have

In a Nash equilibrium, each player adopts a best response to the actions of others. One of the remarkable facts about games, proved by John Nash in 1950 (see Nash, 1950b), is that every game with a finite number of players and a finite number of strategies has at least one Nash equilibrium if we allow players the option of randomising over strategies. More precisely, a Nash equilibrium is guaranteed to exist if we expand the concept of a 'strategy' to include probability distributions over the actions available to players. (You can think of this as a player choosing to create genuine uncertainty in their opponent by selecting an action by flipping a coin, or using a randomisation device; if a player's opponent is *genuinely* uncertain as to what a player will do, then the opponent's best response needs to take into account all the actions a player may take, along with the chance that the player will do that.) In this new framework, what we have been calling a 'strategy', up to now is more properly known as a *pure strategy*; probability distributions over pure strategies are known as *mixed strategies*. Nash's result says that every game with finitely many players and strategies has at least one Nash equilibrium in mixed strategies. One further conceptual shift we must make when speaking of mixed strategies is that we need to talk about a player's *expected* payoff, since there's no guarantee about what outcome will actually result when the game is played. But if we think of a mixed strategy played against another strategy S (pure or mixed) many times, then the average payout over the long run will converge to the expected payout.

In the game of Figure 2a, one Nash equilibrium has player A choosing A_2 followed by player B choosing B_2. This gives a payoff of 90 to A and 50 to B, neither of which could be improved on: had B chosen B_1, he would have received 10 rather than 50, and had A chosen A_1, she would have received 70 rather than 90. Note that, had A chosen A_1, it is true that B would have received 60 rather than 50 – an improvement – but note that this does not result from an alternative choice of strategy by B and so is compatible with the definition of a Nash equilibrium.

In the game of Figure 2b, one can readily check that neither Rock nor Paper nor Scissors can be an equilibrium strategy for either player: if Row picks a strategy and loses, there is always a winning strategy Row could switch to. And, likewise, if Row picks a strategy and wins, then there is always another strategy that Column could switch to and win. (This is the reason underlying the evolutionary cycles in the population behaviour of *Uta stansburiana*.) However, if both players pick a strategy at *random*, with rock, paper, and scissors

received under their original choice. Second, a Nash equilibrium only imposes a requirement for a *single* player contemplating a change. If more than one player were to change strategies at the same time, then anything could happen.

being all equally likely, one can show that no alternative strategy exists – pure or mixed – which one could switch to that would yield a better expected payoff. To see this, first let σ denote the strategy which assigns probability $\frac{1}{3}$ to each pure strategy. For convenience, we will typically adopt the convention of writing a mixed strategy like σ as $\frac{1}{3}\text{Rock} + \frac{1}{3}\text{Scissors} + \frac{1}{3}\text{Paper}$.[6] Now consider the expected payoff when σ plays Rock:

$$\begin{aligned}\pi(\sigma \mid \text{Rock}) &= \tfrac{1}{3}\pi(\text{Rock} \mid \text{Rock}) + \tfrac{1}{3}\pi(\text{Scissors} \mid \text{Rock}) + \tfrac{1}{3}\pi(\text{Paper} \mid \text{Rock}) \\ &= \tfrac{1}{3} \cdot 0 + \tfrac{1}{3} \cdot (-1) + \tfrac{1}{3} \cdot 1 \\ &= 0.\end{aligned}$$

The same is obviously true when σ is played against either Scissors or Paper. From this, it can be easily shown that there is no mixed strategy μ which assigns probability r to Rock, s to Scissors, and p to Paper, where $r + p + s = 1$, such that $\pi(\mu \mid \sigma) > \pi(\sigma \mid \sigma)$.[7] This means that σ is a Nash equilibrium.

This raises an interesting question: why do populations of the common side-blotched lizard not evolve to a state consisting of equal representations of all three types? Since the common side-blotched lizards are essentially playing the game Rock-Paper-Scissors, that would be analogous to the mixed-strategy Nash equilibrium of the game underlying their evolutionary situation. To see this, we can now properly begin our discussion of *evolutionary game theory*. And the first thing we will see is that the solution concept of a Nash

[6] Why this notation? When each player uses a mixed strategy, writing the mixed strategies in this way allows us to compute the expected outcome of the game through a convenient abuse of notation: simply treat the mixed strategies as polynomials and multiply them. For example, suppose Player 1 uses $\sigma = \frac{1}{3}\text{Rock}+\frac{1}{3}\text{Scissors}+\frac{1}{3}\text{Paper}$ and Player 2 uses $\mu = \frac{2}{3}\text{Rock}+\frac{1}{3}\text{Paper}$. Then the expected outcome of σ played against μ is:

$$\begin{aligned}\sigma\mu &= \left(\tfrac{1}{3}\underset{\textbf{Player 1}}{\boxed{\text{Rock}}} + \tfrac{1}{3}\underset{\textbf{Player 1}}{\boxed{\text{Scissors}}} + \tfrac{1}{3}\underset{\textbf{Player 1}}{\boxed{\text{Paper}}}\right)\left(\tfrac{2}{3}\underset{\textbf{Player 2}}{\boxed{\text{Rock}}} + \tfrac{1}{3}\underset{\textbf{Player 2}}{\boxed{\text{Paper}}}\right) \\ &= \tfrac{2}{9}\underset{\textbf{Player 1}}{\boxed{\text{Rock}}}\ \underset{\textbf{Player 2}}{\boxed{\text{Rock}}} + \tfrac{2}{9}\underset{\textbf{Player 1}}{\boxed{\text{Scissors}}}\ \underset{\textbf{Player 2}}{\boxed{\text{Rock}}} + \tfrac{2}{9}\underset{\textbf{Player 1}}{\boxed{\text{Paper}}}\ \underset{\textbf{Player 2}}{\boxed{\text{Rock}}} \\ &\quad + \tfrac{1}{9}\underset{\textbf{Player 1}}{\boxed{\text{Rock}}}\ \underset{\textbf{Player 2}}{\boxed{\text{Paper}}} + \tfrac{1}{9}\underset{\textbf{Player 1}}{\boxed{\text{Scissors}}}\ \underset{\textbf{Player 2}}{\boxed{\text{Paper}}} + \tfrac{1}{9}\underset{\textbf{Player 1}}{\boxed{\text{Paper}}}\ \underset{\textbf{Player 2}}{\boxed{\text{Paper}}}.\end{aligned}$$

Each 'term' of the 'polynomial' represents a possible outcome of play, and the coefficient of each 'term' is the probability that outcome will occur.

[7] Why does $r + p + s = 1$? Because those are the probabilities that a player will use one of the three pure strategies. There are no other strategies available, and a player has to do *something*, so those probabilities must add to one.

	S_1	S_2
S_1	(2, 2)	(1, 1)
S_2	(1, 1)	(1, 1)

Figure 3 The evolutionary drift game.

equilibrium – while a perfectly reasonable solution concept in many contexts – is not the right one to help us understand evolution.

2 Evolutionarily Stable Strategies

2.1 Basic Concepts

Consider the game of Figure 3. Assume that the payoffs listed in each cell of the matrix are the expected number of offspring an individual will have as a result of the interaction, and also assume that we are talking about a species where all offspring are of the same type as their parent.[8] It can be easily seen that the game of Figure 3 has two Nash equilibria in pure strategies: one where both individuals play S_1 and another where both individuals play S_2.

But now suppose that, for historical reasons, the population is in the state where everyone follows the equilibrium strategy S_2. If an S_1-mutant appears, the mutant does not suffer a fitness disadvantage with respect to the rest of the population, because in an $S_1 \bullet\!\!-\!\!\bullet S_2$ interaction the S_1-mutant still receives a payoff of 1, which is exactly what every S_2 individual in the population receives. This means that there is no selection pressure against the S_1-mutant, and so they may persist in the population. If a *second* individual appears following the S_1 strategy (either from an independent mutation or as one of the offspring of the original mutant), the payoff from an $S_1 \bullet\!\!-\!\!\bullet S_1$ interaction is twice that earned by the S_2-type. This gives the individuals following S_1 an explicit fitness advantage over those following S_2, introducing selection pressure *against* the incumbent strategy. Even if, as we are assuming, the S_2-type is the majority of the population, over time we would expect the greater reproductive success of the S_1-type to drive the S_2-type to extinction.[9]

[8] This assumption may raise the question of how applicable evolutionary game theory is to sexually reproducing populations, where the offspring typically feature a blend of traits of their parents. One can introduce refinements to the models which address this, but the additional complexity, at this point, is not worth it.

[9] This intuitive argument highlights the shortcomings of the Nash equilibrium solution concept for capturing the notion of evolutionary stability. Whether or not the S_2-type would actually be driven to extinction depends on details of the underlying evolutionary dynamics. Later, we will show that this happens *always* under a dynamic known as the 'replicator dynamics', but it may not happen if the evolutionary dynamics are modelled using a discrete birth-death process.