"All Together Now"

Linking the Public Goods Game and Prisoner's Dilemma For Robustness Against Free-Riders

Kathleen Pax Donahue

March 21, 2016

Senior Thesis submitted in partial fulfillment of the honors requirements for the degree of Bachelor of Arts in Mathematics.

Abstract

Experimental studies have demonstrated that humans engage in free-riding behavior much less frequently than would be expected theoretically. This thesis proposes that such high willingness to cooperate in the Public Goods Game (PGG) can be explained by linking individual's behavior in groups with pairwise interactions in a Prisoner's Dilemma (PD). Using an evolutionary game theoretic framework, I derive analytical conditions under which cooperation in the PD maintains cooperation in the multi-player PGG. Certain probabilistic strategies that are civic-minded, but also punish selfish players, are robust to invasion from other strategies, including free-riders. Furthermore, I use computer simulations to demonstrate that these "Civic Partner" strategies can outperform selfish strategies in populations where many strategies are present.

Acknowledgments

This thesis was a fascinating exploration of a topic I had previously never considered. I would like to thank Christian Hilbe and Oliver Hauser for suggesting this project, for helping me work through many iterations of proofs and code, and for connecting me with hugely useful resources on this topic (some of which they wrote!). I would not have gotten a quarter as far without their guidance. I would also like to thank Professor Martin Nowak for his feedback, especially for helping to simplify the project at a particularly tricky point in the fall. Professor Clifford Taubes gave invaluable suggestions for the Perron-Frobenius section, as well as overall considerations for thesis structure. Julie Estrada, my writing center tutor, read many terrible drafts of this thesis and helped make them much better. Jack Donahue (who happens to be my father) bravely read an early draft and helped improve many portions. He also suggested this thesis's title. Sarah Chapin (who happens to be my girlfriend) helped by providing editing, support, and grilled cheese. Finally, I would like to thank everyone who has helped me along my math education and anyone I've cornered to talk about game theory in the past year. All mistakes are my own, but there would be many more without the the above people.

Contents

Chapter 1

Introduction

1.1 A Motivating Example

In 2015, “Martin Shkreli” suddenly became a widely-reviled household name. The reason was his drug company, Turing Pharmaceuticals, whose strategy consisted of buying licenses for out-of-patent drugs. Because of the difficulty of getting another license for a generic drug approved, frequently Turing Pharmaceuticals remained the only provider, enabling it to dramatically raise prices. The drug that precipitated Shkreli’s rise to notoriety was Pyrimethamine, a generic anti-malarial and antiparasitic frequently used by AIDS patients. After Turing Pharmaceuticals bought the last licence, it raised the prices by 5,556%. After the change, in the United States a monthly dose cost $75,000, while in Canada (where Turing’s monopoly didn’t exist) the monthly dose was about $60. Shkreli’s actions brought broad condemnation and even prompted the FDA to change its rules to prevent such actions in the future[1]. Intriguingly, other drug companies with very high profit margins generally have not incited such widespread disgust. One unique aspect of Shkreli’s situation is that he did nothing to create the drug he is selling, instead piggy-backing off the medical advances of others. While his strategy is individually profitable, if everyone behaved like him by selling existing drugs rather than developing new ones, there would be no new drugs at all, leaving everyone worse off.

Shkreli is a colorful example of the common problem of **free riders**, those who benefit off of a resource that they do not contribute to. A more common example is the “tax gap”, the difference between the amount of taxes due and the amount actually paid to the government. In 2006, that gap in the United States was $385 billion dollars. Those who do not pay their taxes, unless caught and punished, continue to enjoy the benefits of public schooling, paved roads and clean water that their tax-paying neighbors do. In the absence of enforcement, each individual’s logical choice would be to not pay their taxes. However, a world in which no one paid their taxes is one where each person must clean their own drinking water, pave their own roads and build their own schools, likely at much greater expense. This is in part why society exists: cooperation within communities can make daily life easier than it would be for individuals acting alone. As Justice Oliver Wendell Holmes said, “Taxes are the price we pay for civilized society”[2].

From the game-theoretist point of view, Shkreli’s case raises two questions:

1. Why do such “free-riders” exist?

[1] [6]
[2] [12]

2. Why aren't there more of them?

It is easy to address the first question by showing mathematically the benefits of such a strategy (which will be done in later chapters). However, the second question is more difficult. Most people are not Martin Shkreli. They pay their taxes, help neighbors in need, and donate to charity. Why is there a difference between the "logical" profit-maximizing strategy and the one most people choose to follow? Is it possible to create a theoretical framework under which free riders aren't a problem?

This thesis seeks to answer this question by abstracting the idea of human interactions into game theory, with strategies, people's framework of interpersonal relations, represented numerically. Specifically, this thesis revolves around two simple and widely-studied games: the Prisoner's Dilemma and the Public Goods Games. Together, these games create a model of human interactions that I will use to show the existence of a family of strategies that cooperate with each other, are generous, and prevent the rise of free riders. Before reaching the novel portion of this investigation, this thesis will thoroughly investigate the history of evolutionary game theory, summarizing related papers and introducing core concepts in game theory and evolutionary dynamics that will be applied to the problem at hand.

1.2 Theoretical Framework

This section presents an introduction to the two theoretical games that have been the central focus of game theory for decades. Later sections will expand upon them. Note that the "players" can in general refer to anything that takes independent action, from people to countries to computer programs.

1.2.1 Public Goods Game

In the **Public Goods Game** (PGG), there can be n players, for any positive value of n. Each of them chooses whether or not to cooperate and pay a cost c_G into a common pot of money. If m players cooperate, then the total in the pot is $m * c_G$. This amount is multiplied by a factor r with $1 < r < n$ and divided evenly among the n participants, regardless of if they cooperated or not. Thus, for each player there are two scenarios: they cooperated or they didn't. The end [benefit gained] - [cost paid] is called a player's **payoff**. The payoff of players in the PGG goes as follows:

1. Player cooperated. Payoff: $r * \frac{m * c_G}{n} - c_G$.
2. Player defected. Payoff: $r * \frac{m * c_G}{n}$.

Because $c_G > 0$, the payoff for each individual player is generally higher when they defect. However, if all the players defect, the payoff is identically 0. The analogous counterpart to the Shkreli example would be drug research: each person can choose to add to the pool of medical knowledge which helps all people in the community. People who don't contribute, like Shkreli, benefit from the history of medical advances anyway, but if nobody contributed, there would be no medical knowledge to share.

A more canonical example of the PGG is as follows: Imagine a forest that nearby villagers periodically log. The health of the forest is best with limited logging and periodic maintenance so it can produce taller trees and continue to produce them for years to come. Yet during any year and for any single villager, action that would benefit them the most would be to do no help with maintenance and simply log the entire forest, despite the fact that this would leave the other villagers without wood and even hurt the greedy villager the next year when the forest is all gone. The inherent tension between cooperation and defection is what makes the PGG so interesting to study.

1.2.2 Prisoner's Dilemma

The **Prisoner's Dilemma** (PD) involves two players who I will call A and B. In one "round" of the Prisoner's Dilemma, each player has a choice: they can cooperate or defect. If they cooperate, they pay a cost c_d to give their partner a benefit b. Both c_d and b are positive and here I will assume $b > c_d$. In general, the payoff of a strategy A is denoted π_A. Since both A and B are given the same choices, there are 4 possible outcomes.

1. A and B both cooperate. Payoffs: $\pi_A = \pi_B = b - c_d$
2. A cooperates, B does not. Payoffs: $\pi_A = -c_d$. $\pi_B = b$.
3. A defects, B cooperates. Payoffs: $\pi_A = b$. $\pi_B = -c_d$.
4. A and B both defect. Payoffs: $\pi_A = \pi_B = 0$.

The dilemma in the "Prisoner's Dilemma" comes from the difficulty of maintaining a stable strategy that is maximally beneficial to both players. Situation 1 has a high payoff for both A and B, but each player could get a higher payoff by defecting and moving to Situation 2 or 3. However, if both strategies defect, they end up at Situation 4, which is bad for both of them. This dilemma will be examined in further mathematical detail in later chapters.

The Prisoner's Dilemma comes up in many real-life situations, such as cartels and nuclear war. Its name comes from one example involving two prisoners who have been arrested for conspiracy in a crime. If player A confesses and rats on B, he will get off free while B will get a sentence of 10 years, and vice versa. If both confess, they will both end up with 9 year sentences, and if neither confesses they will only each have a 1 year sentence. Thus, the optimal choice for both is for neither to confess, ending up with 1 year sentences. However, each might be tempted to rat out their partner, leading to a worse overall situation.

A common and fascinating extension of this set-up is the **Iterated Prisoner's Dilemma**, where the PD multiple rounds are played with the same partners. In the memory-one version of the game, each player does not remember the entire history of their interactions, but only what was done in the last round. Based on that information, the player determines what their action in this round will be. This action might be deterministic, such as "Always defect if my partner defected against me in the last round, never defect if it didn't" or probabilistic, such as "Cooperate with probability 95% if my partner cooperated with me in the last round and with probability 50% if it did not". These can be represented as vectors of numbers. For example, the deterministic one would be written $[1, 0]$ to represent a 100% chance of cooperation in the first case (partner cooperated) and 0 otherwise. The second could be written $[.95, .5]$. A strategy of this form can more generally be denoted

$$S_A = [p_c, p_d] \tag{1.1}$$

1.2.3 Combining the Games

The scenario this thesis will examine is how the math of the strategies change when both games are combined. The inspiration for this model is a small town. All members of the community need to contribute to the public good, perhaps through maintaining the town green. They also interact with each other on an individual basis, such as offering to help to shovel a front walk. In the linked version of the game, a person's level of contribution to the public good is public knowledge. A logical strategy in this situation would be for people to refuse to cooperate individually (in the PD) with the neighbors who did not contribute to the public good (were free-riders in the PGG). Such a strategy might encourage those free-riders to cooperate more in the PGG.

The theoretical model is as follows. In the population, there are n players divided into $n/2$ pairs. Each pair plays a memory-one Iterated Prisoner's Dilemma as before. However, before playing each round of the Prisoner's Dilemma, as a whole community they play a round of a Public Goods Game. Each player chooses to cooperate or not in the PGG based on their partner's behavior in the PD. In the **unlinked** version of the game, a player's probability of cooperation in the PD depends on its partner's actions in the last round of PD only. In the **linked** game, a player's probability of cooperation in the PD depends on its partners actions in the PD and PGG.

The focus of this thesis will *not* be in studying real-life small-towns, but rather examining whether in the theoretical, mathematical model involving PD and PGG, there exists some strategies that can maintain high cooperation while punishing free-riders.

1.3 Literature Review

Given how widely both the Prisoner's Dilemma and Public Goods Game have been studied, it's not surprising that the idea of combining the two has occurred to other researchers. This section discusses those papers, noting both their strengths and weaknesses.

1.3.1 Experiments with People

These first two papers study how humans play the PD and PGG in controlled laboratory settings. They can be useful in quantifying the choices people make in interacting with each other.

Milinski et al. "Reputation helps solve the tragedy of the commons"

The first paper[3] was based off an experiment done with students who, via computer, played series of PD or PGG games with each other for money. There were four different patterns in the order of how the PGG and PD games were played. The participants were not told in advance of the schedule of games unless otherwise noted. The distinctions between the four schedules are presented in Table 1.1.

Figure 1.1 shows how cooperation levels varied with the schedule the players were assigned to. Schedule 1 and 2 show higher cooperation in the PGG game and PD in the first 16 games. However, in the last 4 rounds, when it was announced that only PGG games would be played as in Schedule 2, cooperation dropped dramatically as compared to Schedule 1, where the only-PGG

[3] [4]

Schedule ID	Rounds 1-8	Rounds 9-16	Rounds 17-20
1	Alternate between PGG and PD		Only PGG, which was NOT announced in advance
2	Alternate between PGG and PD		Only PGG, which WAS announced in advance
3	PGG only	PD only	Only PGG, which was NOT announced in advance
4	PGG only	PD only	Only PGG, which WAS announced in advance

Table 1.1: This table shows the difference between the four schedules of games in the Milinski paper.

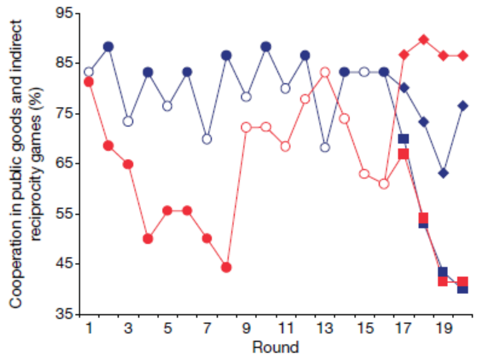

Figure 1.1: Figure from Milinski et al paper showing differing levels of cooperation among different schedules of playing PGG versus PD. Schedules 1 and 2 are in blue and Schedules 3 and 4 are in red. Filled circles denote PGG games and empty circles denote PD games. For the last four rounds, squares denote PGG games that were announced in advance and diamonds represent PGG games that were not announced in advance.

announcement was not made. For schedules 3 and 4, back-to-back rounds of PGG cooperation declined, but cooperation did not decline in back-to-back rounds of PD. Again, if the last four rounds were announced in advance to be only PGG (as in Schedule 4), cooperation dropped, but if they were not announced (Schedule 3), cooperation levels remained high.

Analyzing the first 16 rounds of the experiment, the authors believe that alternating between the PD and PGG game is what allowed higher cooperation to be maintained in Schedules 1 and 2, rather than decline as in Schedules 3 and 4. The evidence for this interpretation is that players are less likely to cooperate in the PD with a player that defected in the previous PGG. Figure 1.2 shows that players had about a 20% probability of defecting in the PD against players that had cooperated in the last round of the PGG and about a 60% probability of defecting against those that had defected in the last round of the PGG. Players that defect in one round of PGG could be convinced to change their behavior by a poor payoff in the PD. Analyzing the last 4 rounds, it seems more important whether players *believe* that PD games will return than whether they actually do. The difference between cooperation levels in Schedule 1 and Schedule 2 shown in Fig. 1.1 shows that where there was still a possibility of PD rounds returning, cooperation was higher than those where the threat of retaliation in PD rounds was known to be gone.

Why is greater cooperation important? The second graph in Fig. 1.2 shows that players in schedules with alternating PD and PGG not only had higher cooperation levels, but also had higher payoff, leading to a greater benefit to all players.

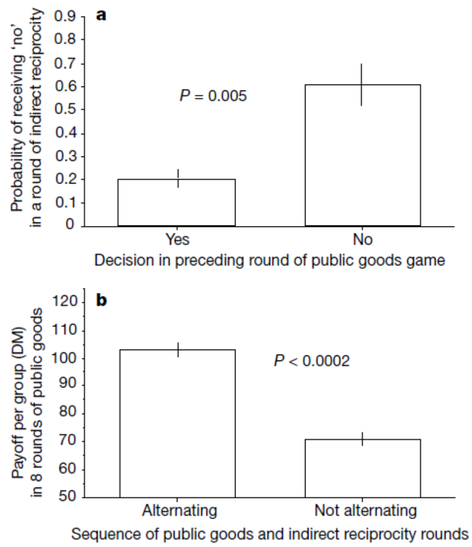

Figure 1.2: Figure from Milinski et al paper showing evidence for the theory that players punish in the PD those players that defect in the PGG.

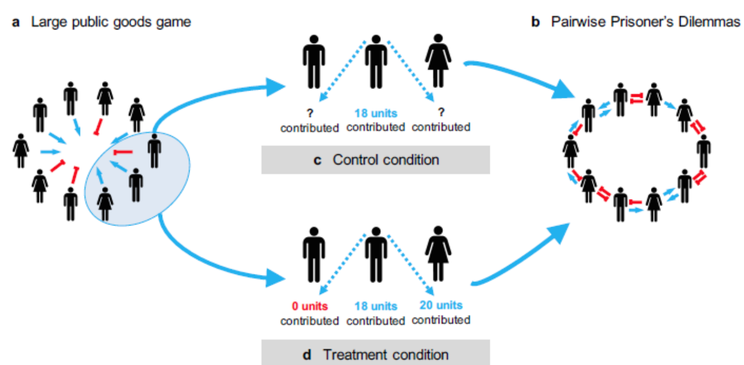

Figure 1.3: Figure from Hauser et al. showing the experimental set-up. In the control condition, the central figure does not know his partner's actions in the PGG, but in the treatment, he is able to see that one contributed much more than the other.

Hauser et al. "Preserving the Global Commons"

The other example[4] involving humans that I will examine comes from the dissertation of one of my advisers, Oliver Hauser. Similarly to the Milinski example, "players" were people interacting via computers. They were in communities of mean size 39, much higher than the 6-person groups in the Milinski trial. In Hauser's set-up, each person was given 20 "Monetary Units", MUs, to start. They could then donate however much they wanted to the central pool in the PGG. The amount was doubled and divided evenly among all players, regardless of if they cooperated or not. Then, each player was arranged in a ring and could played a Prisoner's Dilemma game with each of its two neighbors. In one version of the game, called the control or **unlinked** game, players did not know their partner's history of cooperation in the PGG. In the treatment or **linked** game, they did know their PGG history. Figure 1.3 has a visual representation of the different game scenarios.

[4] [8]

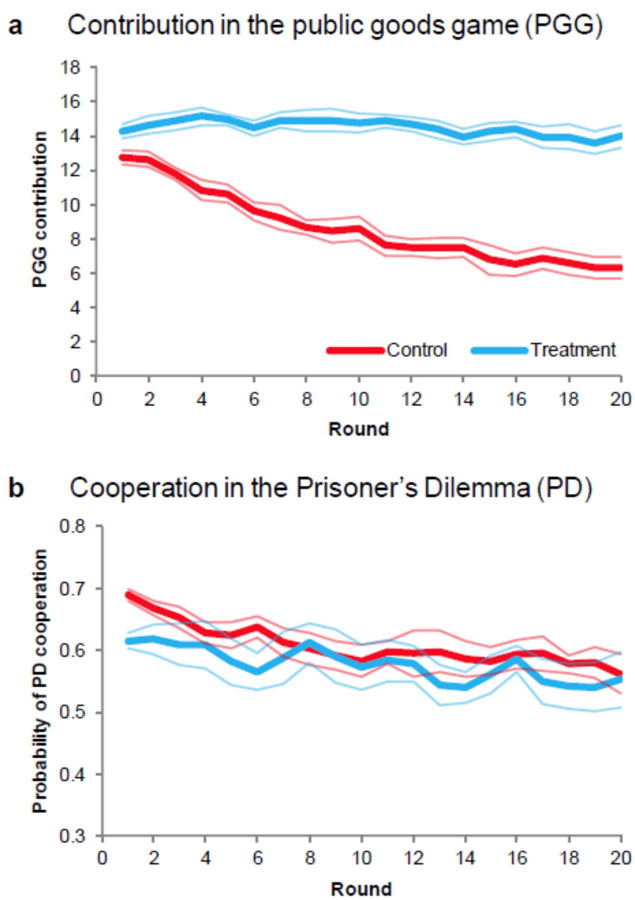

Figure 1.4: Figure from Hauser et al. showing the experimental results. The PGG y axis is out of a total 20 possible Monetary Units that could be donated. Note that under the treatment, a much higher level of PGG cooperation was maintained.

Experimental results are shown in Fig. 1.4. Compared to the control group, linking the games caused a much higher level of cooperation in the PGG to be maintained over time, while it had no significant impact on the cooperation levels in the PD. Hauser found that players were statistically less likely to cooperate with players that had made low contributions to the PGG, punishing them in the exact same way as seen in the Milinski paper.

Human Subject Studies Summarized

These studies provided a powerful illustration of how actual humans have existing strategies for how to play these very theoretical games. The downside is that it is very hard to see exactly which strategy each person is playing, especially if a person's strategy changes over time. Modeling a human's complete decision-making process would be very difficult, almost impossible. In order to make study more manageable, simplified models of how humans interact, I turn to other papers.

1.3.2 Theoretical Investigations: Computer Simulations and Proofs

The next two papers also investigate linking the PGG and PD, but instead of human players, their participants are computers making decisions based off pre-programmed strategies. The human-based papers showed that strategies exist with high PGG cooperation and propensity to punish those in the PD who defect in the PGG. These theoretical papers will investigate potential explicit strategies that give the same results.

Panchanathan and Boyd "Indirect reciprocity can stabilize cooperation without the second-order free rider problem"

Panchanathan and Boyd's[5] structure has n players that first play a PGG, then multiple rounds of a "mutual aid game" similar to the PD. This game repeats with probability w, and when it ends, it has reproduction of strategies based on payoff and death. **Reproduction** and **death** always are matched, so the overall population remains constant. If player A dies and player B reproduces, then the strategy B follows is copied to player A, so in the future player A will follow the strategy of player B. Any player's probability of reproducing is usually proportional to its payoff, so more successful strategies spread throughout a population and less-successful ones die off. An analogy to a human population could be the spread of good ideas: if people notice one person is doing particularly well in life, they may abandon their previous strategy in life and instead copy someone elses.

In Panchanathan and Boyds set-up specifically, in the mutual aid game a player is selected at random to be "needy". Players have the option of helping it or not. Players can be in good or bad standing, which is known to other players. Bad standing is earned permanently by refusing to cooperate in the PGG and earned temporarily by refusing to help a needy player of good reputation. However, if bad reputation is gained temporarily, the player can regain good standing by helping any needy player in the future.

As opposed to Milinski's paper, which allowed each human participant to play any "strategy" they chose, Panchanathan strictly regulates which strategies are allowed to exist. There are only three:

[5] [9]

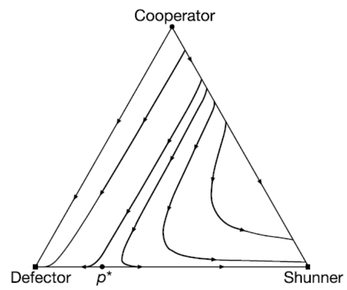

Figure 1 Evolutionary dynamics of the Shunner, Defector and Cooperator strategies, plotted in trilinear coordinates. The arrowheads depict evolutionary trajectories. Populations composed of all Shunners and all Defectors are both evolutionarily stable. When the initial frequency of Shunners is too low, the population evolves towards the asocial Defector equilibrium. When it is higher, the population reaches a stable equilibrium in which collective action is enforced by exclusion from the benefits of indirect reciprocity. The model parameters for this figure are: $B = 10$, $C = 5$, $b = 2$, $c = 1$, $n = 100$, $w = 0.95$, $W_0 = 100$ and $e = 0.05$.

Figure 1.5: Figure from Panchanathan et al paper showing how populations evolve to contain more or fewer of each strategy type. Shunners and Defectors are evolutionarily stable because slight perturbations from their vertex end up on lines with arrows pointing back to the original vertex. By contrast, a slight perturbation from the Cooperator vertex sends it away from the original vertex and towards Defectors or Shunners.

1. Defectors do not cooperate in the PGG or the mutual aid game, except by error with probability e.

2. Cooperators cooperate in the PGG and help all needy players in the mutual aid game regardless of standing. However, with probability e, they will mistakenly not help players of good standing.

3. Shunners cooperate in the PGG game and then try to help all players of good standing in the mutual aid game, but again mistakenly do not help with probability e.

The goal of this paper was to show that strategies that punish those that don't help others (like the Shunners) are evolutionarily stable. Roughly speaking, a strategy is evolutionarily stable if a homogenous population of that strategy cannot be invaded by a single mutant. It is a desirable property for a strategy because it shows a robustness to competing strategies. (This topic will be explored in detail in later sections). A useful figure is shown in Fig. 1.5. Each vertex of the triangle shows a population composed entirely of players with that strategy. For example, the top vertex is when all players are Cooperators. A point in the very center would have equal numbers of each strategy. Arrows indicate paths the population follows over time. For example, the all-cooperator initial population gradually gains more defectors through death and reproduction until it is completely Defectors. A population with some cooperators, some Shunners and some Defectors either becomes a population of completely Defectors or completely Shunners based on the initial ratios. Because the population of Shunners is evolutionarily stable, this shows the stability of a strategy that punishes anti-social strategies, as desired.

However, their paper doesn't end there. The authors were also interested in seeing how exactly Shunners could invade and if there were situations in which they could not. They introduced a fourth strategy: Reciprocators. A Reciprocator does not cooperate in the PGG and does not care about whether other players cooperated in the PGG or not. However, it does care if they cooperate in the mutual aid game and will only and always help those who have helped those in need in the mutual aid game (subject to a probability e error as before). This is essentially a Shunner that ignores the PGG. Additionally, it could be analogous to a strategy in the "unlinked" game, where players ignore the PGG when playing the PD.

In this set-up, they removed Cooperators and Defectors. Both Shunners and Reciprocators are stable. However, Shunners do less well when the number of mutual aid games is high. The authors suggest this is because Shunners never help Reciprocators in the mutual aid game: they failed to cooperate in the PGG, so are still in bad standing according to Shunners. The Reciprocators don't care about PGG cooperation, but when Shunners fail to help Reciprocators, they fall into the Reciprocator's definition of bad standing and so will not be helped by other Reciprocators when they are needy. This is more important when there are many mutual aid games because probabilistically, more Shunners will have refused to help Reciprocators by the end of a long game.

This paper is helpful in some respects. It shows strategies (Shunners) that punish "antisocial" strategies can be evolutionarily stable against strategies that always cooperate (Cooperators) or always defect (Defectors). However, they aren't always successful, as is shown by how Shunners in some situations lose out against Reciprocators, a type of free-rider. However, this paper has one major drawback: it only considers 4 strategies. It's unlikely that these four simple strategies cover the entirety of strategies that players can use in life, or even in relatively simple combinations of theoretical games. Restricting the scope this severely impacts the usefulness of this paper's conclusions.

Hilbe et al. "Partners or rivals?"

The last paper[6] I will examine was written by one of my advisers for this thesis. Hilbe et al studied Iterated Prisoner's Dilemma (IPD), not PGG. Rather than trying to prove specific theorems about when strategies were stable, it focused on categorizing and understanding the wide array of strategies that can emerge. The strategy space for memory-1 strategies, discussed in the vast majority of the paper, was determined by 1 vector: $(p_{CC}, p_{CD}, p_{DC}, p_{DD}; p_0)$. Each number represents the probability of a player cooperating a specific situation. For example, p_{CD} is the probability of a player cooperating if it cooperated and its partner defected in the last round. p_0 is the probability that it cooperates in the first round, when there is no history.

This strategy space is continuous and 5-dimensional, which makes it substantially more complicated than Panchananthan's 4-strategy space. Instead of simulation, Hilbe et al focuses on dividing the strategy space into strategies with different properties. For the sake of brevity, I will skip the proofs and only list several definitions. For example, a **competitive** strategy is a strategy A such that a player against it with any strategy B gets a payoff than is less than or equal to A's payoff. A **partner** strategy A is one that is never the first to defect and additionally has the property that for any co-player strategy B, if the payoff to a partner strategy A is less than $b - c_d$, the payoff to B is also less than $b - c_d$. This particular bound is important because it is the level of payoff for two players both cooperating in the PD. Therefore, it is impossible to do better than mutual cooperation against a partner strategy. Both of these are potentially useful properties for a strategy to have and will be useful definitions to bear in mind going forward.

Hilbe et al uses the idea of a **zero-determinant strategy**, a strategy for player A such that

[6] [1]

there exist constants α, β, γ, that if π_A=payoff of strategy A and π_B=payoff of strategy B, then

$$\alpha * \pi_A + \beta * \pi_B + \gamma = 0$$

or in other words, there is a linear relationship between the payoff of A and the payoff of B. Dividing by a constant and rearranging allows this to be more conveniently written as

$$\pi_B - \kappa = \chi(\pi_A - \kappa)$$

Certain strategy types for A that are zero-determinant cause this equation to have different properties. **Equalizer** strategies have $\chi = 0$, forcing the payoff of B to be κ regardless of π_A and regardless of Bs strategy. **Extortion** strategies have $\kappa = 0$ (the payoff in the case both players defect) and $\chi \in (0, 1)$, so π_B is always some fraction of π_A. Note that when two extortion strategies meet each other, though, they both end up getting the same payoff: 0. A **generous** strategy has $\kappa = b - c_d$ and $\chi \in (0, 1)$. A generous strategy ensures its payoff and its partner's payoff vary together: what helps on also helps the other. It also forces π_B to be constrained to vary within a narrower range than π_A [7]

While these specific results are less relevant for this thesis because they deal only with IPD, the formulation is extremely useful. Much of this thesis will similarly deal with very large strategy spaces that need to be carefully studied, rather than Panchathan's severely restricted 4-strategy space. However, in contrast to Hilbe et al's paper, this thesis will endeavor to build upon this work by not only identifying strategies, but also showing via computer simulation many additional properties certain strategies have. Later, some strategies in the novel portion of this thesis will be shown to be partner strategies. In a certain limit, they will also be equalizers.

[7] Note that Hilbe et al.'s definition of "generous" is very specific and different from how the term will be used in other contexts. Other portions of this thesis will just use it to refer to a strategy that cooperates sometimes even when its partner had defected against it.

Chapter 2

Simple Games

The previous chapter introduced the basic concepts of the Prisoner's Dilemma and Public Goods Game, as well as recent studies of how they might be combined. This chapter will present an overview of the key concepts in evolutionary dynamics and game theory, especially as they relate to the concept of linking PGG and PD games, which will be explored in depth in later chapters. Much of this section is drawn from a textbook by one of my advisers, Professor Martin Nowak[1].

2.1 Nash Equilibrium

Imagine all the potential actions that could be taken in a game as a physical space called $\mathcal{A}$. $\mathcal{A}$ is n dimensional, where n is the number of players in the game. A point $p \in \mathcal{A}$ is a vector describing each player's actions. An example for a point where $n = 3$ might be "Player 1 cooperates, Player 2 cooperates, Player 3 defects". A **Nash Equilibrium** is a point in $\mathcal{A}$ with the special property that no player can change their action (holding all other players' actions constant) and get a higher payoff. However, it might be able to switch its action and get the same payoff. It is a *strict* Nash Equilibrium if no player can switch its action and even get the same payoff. Any switch will give it a strictly worse payoff. Essentially, a Nash equilibrium is a local maximum in the graph of payoff as a function of player's actions. Let's consider an example of a game with two players and the following payoffs:

1. If both players cooperate, they each get payoff of 3.
2. If player 1 cooperates and player 2 defects, player 1 gets payoff 0 and player 2 gets payoff 5.
3. If player 1 defects and player 2 cooperates, player 1 gets payoff 5 and player 2 gets payoff 0.
4. If both players defect, they each get payoff 1.

[1] [5]

	Player 2 Cooperates	**Player 2 Defects**
Player 1 Cooperates	(3,3)	(0,5)
Player 1 Cooperates	(5,0)	(1,1)

Table 2.1: An alternate presentation of the payoff in a PD in a matrix format. For each entry in the table, the first value in the vector is Player 1's payoff and the second is Player 2's payoff.

An alternate presentation of this information is shown in Table 2.1. Point 1 isn't a Nash equilibrium: either player 1 or player 2 could defect and get a higher payoff ($5 > 3$). Points 2 and 3 similarly can't be Nash equilibrium: the cooperating player could defect and get a higher payoff ($1 > 0$). Point 4 (defect) is a Nash equilibrium: if they are both defecting, neither can increase their payoff by cooperating. However, note that if they both cooperated, they could move to Point 1 with a higher overall payoff, 3. The Nash Equilibrium is the best each player can do given the current situation, but better solutions might be possible through broader coordination. A Nash equilibrium is a property of a point in strategy space, but can also be used to refer to a strategy, that when played by all players, causes them to land in the Nash equilibrium. Here, that strategy is "defect".

Consider the concept of a Nash Equilibrium in the PGG. Each of the n players has the option of cooperating (contributing to the pot of money to be multiplied and divided) or defecting and contributing nothing. If a player switches from defection to cooperation, then their net change for payoff is

$$\pi_{cooperate} - \pi_{defect} = (\frac{r * m * c_G}{n} - c_G) - \frac{r * (m-1) * c_G}{n} = \frac{r * c_G}{n} - c_G = c_G(\frac{r}{n} - 1) \tag{2.1}$$

It is a better choice to cooperate than defect when $\frac{r}{n} - 1 > 0$ or $r > n$. This makes intuitive sense: if the player's contribution is multiplied by a larger amount than the number of players it is divided among, it always makes sense to contribute regardless of the other players' choices. For example, a PGG with 3 players and $c_G = 1$ where $r = 10,000$, switching from defection to cooperation would increase each cooperating player's payoff by $1 * (\frac{10000}{3} - 1) \gg 0$. Thus, if $r > n$, the Nash Equilibrium will be cooperation (all players cooperating), but if $r < n$, the Nash Equilibrium will be defection. If $r = n$ then each player gets exactly the same payoff regardless of if they cooperate or defect, so every single combination of actions is a Nash Equilibrium (though none are strict Nash Equilibrium). The commonly-accepted definition of PGG has $r < n$ and so the Nash equilibrium is universal defection.

One interesting phenomenon is that humans rarely express this complete defection. As shown in Fig. 1.1, the Milinski et al paper found that the lowest cooperation dropped to is about 35%, which is much higher than the 0% that would be the logical Nash Equilibrium. One question this thesis seeks to answer is why some innate propensity for irrational cooperation in the PGG could be evolutionarily helpful. If theoretical calculations fail to find a strategy that logically explains this behavior, perhaps the models of interactions are more complicated than previously thought. On the other hand, if theoretical calculations find a strategy that would explain human behavior, we will grow closer to understanding how people make decisions in daily life.

2.2 Evolutionarily Stable Strategy (ESS)

2.2.1 Population with Two Strategies

A Nash Equilibrium isn't the only way to consider which strategies are the "best". The ESS was dreamed up by John Maynard Smith independently of the concept of the Nash Equilibrium. Its set-up considers a population of players of strategy A and some small amount of "mutants" of strategy B. Suppose that the fraction of players with strategy B is $\epsilon \ll 1$, so the remainder of the $1 - \epsilon$ are playing strategy A. Strategy A is **evolutionarily stable** if players using this strategy have a higher payoff than any mutant strategy B, and so can resist invasion by it. **Invasion** occurs when rare strategy has a higher payoff on average than a common strategy. Because players with high

	... **playing A**	... **playing B**
Payoff of A ...	(a,a)	(b,d)
Payoff of B ...	(d,b)	(c,c)

Table 2.2: Table showing payoff with notation as used in ESS set-up. For each entry, the first item in the parenthesis gives the payoff of the strategy in the row when playing the strategy in the column.

payoff are more likely to reproduce, B then increases in number and displaces the more common ("resident") strategy.

Assume that each player is repeatedly playing games with randomly-selected members of the general population. In the limit of a large population, each player of strategy A has a $1-\epsilon$ probability of being paired with a player of strategy A and an ϵ probability of being paired with a player of strategy B. Similarly, each player of strategy B has a $1-\epsilon$ chance of being paired with a player of strategy A and an ϵ probability of being paired with a player of strategy B. Strategy A is ESS if and only if it has a higher average payoff than the players of strategy B. Now, I'll show an example of what ESS implies in practice. Assume that when A plays itself, it has a net payoff of a and when B plays itself it has a payoff of c. Additionally, assume that when A plays B, A has a payoff of b and B has a payoff of d. This information is represented in Table 2.2. Then, saying A is ESS means that [2]

$$a(1-\epsilon)+b\epsilon > d(1-\epsilon)+c\epsilon \tag{2.2}$$

In the limit of $\epsilon \ll 1$, this requires $a \geq d$. Intuitively, this means that since strategy B is rare, most of the time strategy A will be playing itself (payoff a) and strategy B will be playing strategy A (payoff d). Therefore, having $a > d$ in the limit of small ϵ guarantees that strategy A will have a higher payoff than strategy B. In the case that $a < d$, B will have a higher comparative payoff and be able to invade. If $a = d$, then this term drops out. In this case, A is ESS if $b > c$: if A gets a higher payoff when it plays B than B gets when it plays others of strategy B.

This thesis will use a related concept called **robust to invasion by lone mutant** or occasionally **robust to invasion**. This set-up is as follows: There are $n-1$ players of strategy A and 1 of strategy B. Here, $n-2$ of the A players are paired with each other and one will be paired with the strategy B player. When strategy A players play each other, they get payoff a, when A plays B, A gets payoff b and B gets payoff d. Note that as there is only one player of strategy B, the payoff it gets when playing itself (c) is irrelevant in this setup. Then, A is robust to invasion by a lone mutant if

$$\frac{(n-2)*a+b}{n-1} \geq d \tag{2.3}$$

In the limit of large n (equivalent to small ϵ), this simples to $a \geq d$, similarly to before.

2.2.2 Population with Many Strategies

Of course, the particular set-up of mainly A with a few B isn't the only way a strategy can be invaded. For example, it's possible that B fails to invade A, but B loses heavily to a third strategy C, and the temporary presence of B allows C to invade A. The presence of more than two strategies is much more complicated. This thesis will focus less on the math of this particular situation, but it is useful to know the rough idea because occasionally in computer simulations more than two

[2] [5] pg. 53

strategies will be present at one time[3].

First, a bit of notation. S_i stands for strategy i and $E(S_i|S_j)$ is the payoff S_i gets when playing against S_j. In the many-player game, a strategy S_k is a strict Nash Equilibrium if

$$E(S_k|S_k) > E(S_i|S_k) \quad \forall i \neq k$$

In words, this means that S_k gets a higher payoff from playing itself than any other strategy gets from playing S_k. It's a Nash Equilibrium because none of the players playing S_k can do any better by switching their strategy. Turning the $>$ to $\geq$ tells is that S_k is a non-strict Nash Equilibrium if

$$E(S_k|S_k) \geq E(S_i|S_k) \quad \forall i$$

Similarly, this means that none of the players with S_k can do any better (at most, equally well) if they switch to a different strategy. Similarly, a strategy S_k is ESS if

$$E(S_k|S_k) > E(S_i|S_k) \quad \forall i \neq k$$

or

$$E(S_k|S_k) = E(S_i|S_k) \quad \text{and} \quad E(S_k|S_i) > E(S_i|S_i) \quad \forall i \neq k$$

The top condition is exactly the condition for Nash equilibrium, and the bottom equation means that if a strategy S_i does as well against S_k as S_k does against itself, then S_k must do better against S_i than S_i does while playing itself.

Similarly, we can weaken the last $>$ to $\geq$, which tells us that S_k is **robust against invasion by selection** if

$$E(S_k|S_k) > E(S_i|S_k) \quad \forall i \neq k$$

or

$$E(S_k|S_k) = E(S_i|S_k) \quad \text{and} \quad E(S_k|S_i) \geq E(S_i|S_i) \quad \forall i \neq k$$

This means that either S_i is a strict Nash equilibrium or, if S_i does as well against S_k as S_k does against itself, then S_k does *at least* as well against S_i as S_i does against itself, a slightly weaker condition than ESS. This is a generalization of the "robust to invasion" definition previously given for the two-player case. Finally, a strategy is called **unbeatable** if

$$E(S_k|S_k) > E(S_i|S_k) \quad \text{and} \quad E(S_k|S_i) > E(S_i|S_i) \quad \forall i \neq k$$

In words, this means that S_k gets a higher payoff when playing against itself than any other strategy gets by playing against it, and S_k gets a higher payoff against any other strategy than that strategy gets even playing itself. This criteria is quite strict and strategies that satisfy it aren't very common.

In the following chapter of this thesis, I will define a strategy that is robust to invasion by three specific strategies. I will later find that this simple requirement is sufficient to show that strategy is robust to invasion from all strategies in the linked strategy space, and also that it is stable against invasion by selection against a subset of 7 strategies.

[3] [5] pg. 54

2.3 Interesting Strategies in the Prisoner's Dilemma

One main focus of this thesis will be examining a particular class of strategies with certain nice properties. This is only the latest in a long history of research identifying and describing certain strategies with special properties, a history that I will review here.

2.3.1 Tit for Tat

While it is easy to theorize about why certain strategies perform better than others, it is difficult to find empirical evidence. In 1980, Robert Axelrod, a political scientist, proposed an interesting experiment that could perhaps shed light on the issue of what strategies are the most successful.

He invited researchers from across disciplines to participate in a tournament of Iterated Prisoner's Dilemma. Each entrant could submit a strategy for their player based off the entire history of actions by both players in that game to date. Each strategy would play each other and would get a certain number of points for each round based on the Prisoner's Dilemma payoff diagram. In the end, the strategy with the highest number of points won. In total, 14 strategies were submitted.

Writing in "The Evolution of Cooperation"[4] in 1984, Axelrod recalls the original tournment and one strategy in particular called Tit for Tat (TFT). TFT cooperates on the first round and then for each subsequent rounds does exactly what its partner did in the previous round. At a time t, it can be written

$$\begin{cases} \text{cooperate} & t = 1 \\ P_{t-1} & t > 1 \end{cases}$$

where P_{t-1} represents the action the partner of the TFT player took on the previous round.TFT is a memory-one strategy, meaning it only remembers the last round that occurred. Actually, it only remembers what its partner did in the last round, not even its own behavior! Some strategies entered in Axelrod's tournament used the entire history of actions in the game for both players to try to optimize for a higher payoff, so TFT was a very simple strategy indeed.

At the time of the tournament, TFT was well known, both in games with human participants and in computer simulations. TFT got first and second place in two preliminary tournaments whose results were sent to the participants of the actual tournament. Many entrants took TFT and tried to improve upon it. Despite this, Axelrod notes, "the striking fact is that *none* of the more complex programs submitted was able to perform as well as as the original, simple TIT FOR TAT". In fact, in the actual tournament, TFT won first place. In fact, in a second tournament held in later years with 62 participants, all of whom knew about the success of TFT, TFT again won.

What accounts for TFT's success? Axelrod believes that there are four properties of a strategy that indicate its success: being nice, forgiving, retalitory, and clear. I will discuss the first two now and the latter two later in the chapter. **Nice** means that the strategy is never the first to defect. TFT only defects in response to defection, so meets this requirement. In both tournaments, almost all of the top strategies were nice. However, being nice wasn't sufficient for success. Another property Axelrod noticed was the **forgiveness** of a strategy: "its propensity to cooperate in the

[4] [10]

moves after a strategy has defected"[5]. In this thesis, I use the word "generous" interchangeably to mean "having high forgiveness". TFT is a very forgiving strategy: after punishing its partner once, TFT is willing to cooperate again if the partner indicates goodwill by cooperating first. One strategy that was nice but not forgiving in the first tournament was called Grim. Once a player defected against it once, Grim would never cooperate with it again, even if that player repeatedly cooperated. Grim did the worst out of all of the nice strategies in the first tournament.

When the results of the first tournament were published, many players studied TFT and tried to do better than it. Axelrod writes that there were two potential lessons to be drawn from the first tournament:, "Lesson one: Be nice and forgiving. Lessson two: If others are going to be nice and forgiving, it pays to try to take advantage of them"[6]. Naturally, people who decided to base their strategies around Lesson 1 were key targets to be exploited by those who followed Lesson 2. However, similarly to the results of the first tournament, those strategies that were exploitative tended not to do very well. The best strategies tended to be variants of TFT that controlled for one weakness of TFT: its poor showing against RANDOM, a strategy that took actions randomly. They did better by cooperating *less*. Apparently, though, it did not occur to many researchers to cooperate *more* than TFT.

In advertisements for the second tournament, there were instructions for how to write up your strategy. Amusingly, the sample strategy including with the advertisement would have beaten all strategies in the second tournament, including TFT! This prodigious strategy is "Tit for Two Tats": it cooperated first, and only stopped cooperating with its partner after two defections in a row. It is even more forgiving than TFT. The success of Tit-for-Two-Tats shows that TFT isn't some magical unbeatable strategy. It also shows that there is room for successful strategies that are more forgiving. However, there is likely a limit to how forgiving a successful strategy can be: AllC, the strategy that always cooperates, would likely not be the most successful option.

AllC's failure could be because it lacks other properties of a good strategy. Axelrod defines a strategy that is **retalitory** as one that "immediately defects after an 'uncalled for' defection by the other"[7], leaving open the definition of "uncalled for". TFT is realitory: it responds immediately to both defection and cooperation. However, AllC is not retaliatory: it continutes cooperating no matter how frequently its partner defects. In the second tournament, Axelrod notes that some strategies focused excessively on ferreting out whether their partner strategy was retaliatory or not in order to exploit more "easy-going" strategies like AllC. Axelrod defines one more property of a successful strategy: its **clarity**[8], which is how easy it is for other strategies to recognize its properties and the best way for them to interact with it. With TFT, an intelligently-designed strategy will recognize that cooperation is the best strategy.

However, that last caveat identifies one of the weaknesses of the Axelrod tournament. TFT was successful, but against a very limited set of strategies: in total, fewer than 100, all specially designed, not randomly generated. If we consider real life, many strategies organisms might develop will likely be not intelligently created, but rather contain elements of randomness. A more robust test would have TFT play a very large number of randomly-generated strategies for a very long

[5] [10] pg. 36
[6] [10] pg. 47
[7] [10] pg. 44
[8] [10] pg. 54

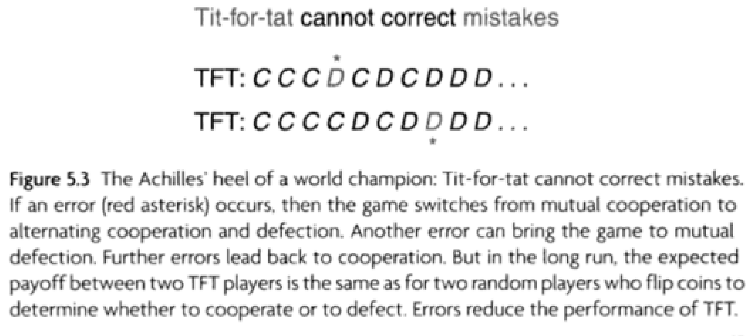

Figure 2.1: Illustration of how TFT playing itself can slip into mutual defection. Courtesy of [5].

time. At the time of the tournaments, computational power was likely a limiting factor. However, later studies would explore this issue.

2.3.2 Generous Tit for Tat

In 1992, Martin Nowak and Karl Sigmund published a paper called "Tit for tat in heterogeneous populations"[9]. It built off Axelrod's results, particularly the weakness of TFT in the presence of "errors". In this paper, an **error** would be when a player intended to cooperate but instead defected or vice versa. Errors are a useful concept because they reflect the randomness of real life: sometimes people will mistakenly believe their partner defected when they cooperated and vice versa. TFT noticeably did much worse in the presence of even a small error rate, though still better than the competing submitted strategies. The reason why it did worse is that it would mistakenly defect against cooperating strategies (thus impinging its property of being nice), occasionally bringing down retaliatory pressure. Even when playing itself, a single error could lead it into a cycle of players cooperating and defecting in turn (reducing the player's payoff). A second error would lead to an unending cycle of defection. This pattern is shown in Figure 2.1. With such poor performance, it seemed possible that another strategy could do better than TFT in the presence of errors.

Nowak and Sigmund introduced a new strategy called "Generous Tit for Tat" (GTFT). Similarly to the TFT, it cooperated on the opening round and every time when its partner had cooperated in the last round. However, while the TFT defected when its partner had defected on the last round, the GTFT cooperated with a certain probability p even in the case that its partner defected in the last round. Because of this, GTFT is more forgiving (more generous) than TFT. Note that GTFT, unlike TFT, is probabilistic, with an element of randomness even if all relevant information is known. They defined this probability to be

$$\min\{1 - \frac{T-R}{R-S}, \frac{R-P}{T-P}\} \tag{2.4}$$

where R is the payoff strategies get if they both cooperate, P is the payoff if they both defect, and if one defects and the other cooperates, the defector gets payoff T and the cooperator gets payoff S. Translated into my notation,

$$R = b - c_d \quad P = 0 \quad T = b \quad S = -c_d$$

[9] [7]

Therefore, the first term equals:

$$1 - \frac{(b) - (b - c_d)}{(b - c_d) - (-c_d)} = 1 - \frac{c_d}{b} = \frac{b - c_d}{b} \tag{2.5}$$

The second works out to

$$\frac{(b - c_d) - 0}{b - 0} = \frac{b - c_d}{b} \tag{2.6}$$

which is the exact same equation. This equation will show up later in the novel portion of my thesis in calculating an analogous probability of cooperation.

Nowak and Sigmund note that if a population of randomly-generated strategies is initiated, it will tend towards the rise of AllD-like strategies and away from more nice and forgiving strategies. This is because those with lower cooperation probabilities will exploit those with higher cooperation probabilities. However, in the case where one of the initial strategies is TFT, the population first goes towards AllD and later TFT grows to dominate. In the end, though, TFT gives way to GTFT, which has a higher payoff against itself in the presence of errors.

This concludes the literature and evolutionary game theory review portion of this thesis. Future chapters will focus solely on the novel work I produced related to this topic. However, I will continue to tie back my results to previous findings. This provides context for how my new findings fit into puzzles generations of mathematicians have been exploring.

Chapter 3

Analytical Portions

Previous chapters introduced the main theoretical frameworks of evolutionary game theory. This chapter will go through the analytical derivations related to the central question of linking the Public Goods Game and Prisoner's Dilemma. My central focus is two ways these games can be combined, shown in Table 3.1. The particular question at hand is whether there exists some family of strategies the cooperate frequently, are generous (forgiving), and can't be invaded by free-riders. Additionally, I will investigate, if a strategy of this form exists, if there are differences in whether it exists for the two versions of the game (linked and unlinked).

3.1 Theoretical Framework

This section explain the techniques used and key definitions I will focus on in later portions of this chapter.

3.1.1 Linked vs. Unlinked Games

In this thesis, the most interesting case is the linked one because it allows us to examine how information from the PGG and PD can both inform actions in the PD. However, the unlinked game is essential to this analysis because it acts as a control case. In the unlinked game, the PD is played as it would be in the standard Iterated Prisoner's Dilemma. This allows me to compare my findings with existing IPD research. I keep the PGG in the unlinked game, even though it does not affect the PD, because this keeps the structure of the unlinked game as close as possible to the linked game. Building the unlinked game so it mimics the linked game as closely as possible enables me to see how one small change (cooperation in the PD depending on actions in the PD *and* PGG) impacts my results.

As mentioned before, in the linked game players can have different cooperation probabilities in

	Unlinked	**Linked**
Probability of cooperation in the PGG depends on partner's actions in last round of	PD	PD
Probability of cooperation in the PD depends on partner's actions in last round of	PD	PGG and PD

Table 3.1: Table showing the various ways the two games can be combined in this thesis.

the PD based on their partner's actions in the PGG. Specifically, a particular strategy could have a lower probability of cooperation in the PD with players that defect in the PGG, effectively punishing free-riders. However, this distinction is impossible in the unlinked game because a player can base its cooperation probability only on its partner's actions in the PD. Based on this, I would expect that free-rider strategies have a higher payoff in general in the unlinked game than the linked game. If reproduction is proportional to payoff, this would lead to a greater frequency of free-rider strategies.

Definition 1. *Notation. The following notation is used for denoting the strategy space of each player, which is a probability of cooperation in various situations. Consider the terms*

$$P(\text{ cooperate in PGG}|\text{Partner cooperated in PGG but not PD})$$

$$P(\text{ cooperate in PD}|\text{Partner cooperated in PGG and PD})$$

In order to distinguish between the two terms and shorten the length, let's write the two above terms like this:

$$P_{Cd} \quad p_{Cc} \tag{3.1}$$

Here, capital P denotes the probability of cooperation in the PGG, while lower case p is the probability of cooperation in the PD. C is the event of partner's cooperation in the PGG, while c and d are cooperation and defection respectively in the PD. With this notation, there are 8 possible situations to specify:

$$\{P_{Cc}, P_{Cd}, P_{Dc}, P_{Dd}|p_{Cc}, p_{Cd}, p_{Dc}, p_{Dd}\} \tag{3.2}$$

In this thesis, the most complicated situation I will consider will be a subset of this space above and will be called the **Linked Game**

$$\{P_c, P_d|p_{Cc}, p_{Cd}, p_{Dc}, p_{Dd}\} \tag{3.3}$$

Here, note that cooperation in the PGG depends only on the partner's action in the PD, not the PGG. I will also consider the **Unlinked Game**

$$\{P_c, P_d|p_c, p_d\} \tag{3.4}$$

where cooperation in the PGG and PD both only depend on the PD. Note that the Unlinked game is a subset of the Linked Game space. In this space, it can be written as

$$\{P_c, P_d|p_c, p_d, p_c, p_d\} \tag{3.5}$$

3.1.2 Transition Matrix

Many of the strategies studied in this thesis are probabilistic: for any set of circumstances, there is a given probability of that player cooperating. It is frequently useful to analyze this type of situation by using a **transition matrix**. Transition matrices relate to two matched players. Each entry in the matrix corresponds to the probability of going from one particular point in strategy space to another point. Transition matrices come from the study of **Markov chains**. A Markov chain describes a process of probabilistic transition between various states X_i. Markov chains are also memory-1, meaning that which state you end up in at time t depends only which state you are in at time $t-1$. Your location at time $1, 2, \ldots t-2$ are all irrelevant. This is exactly the set-up of the game I am analyzing: each strategy bases its decision to cooperate or not only on its partner's actions in the previous round of the game, but not any other previous actions. This formulation greatly simplifies the analysis and allows for powerful statements about properties of the game. In the remaining portions of this thesis, I will rely on the results of Markov Chain theory to bolster my results.

Let's explore the specific set-up of this thesis. As described in the "Notation" section, six numbers are required to describe a strategy in the linked space. A particular strategy can be written:

$$S_l = \{P_c, P_d | p_{Cc}, p_{Cd}, p_{Dc}, p_{Dd}\} \tag{3.6}$$

In each round, there are 16 possible states: for each of the two players, they have two options for states in the PGG and two options for states in the PD, and $(2 \cdot 2)^2 = 16$. I could represent the situation where the first player cooperates in the PGG and PD and the second player cooperates in the PGG but not PD, for example, as

$$CC|CD \tag{3.7}$$

where the first entries correspond to action in the PGG and the second two correspond to actions in the PD. In the transition matrix T, there is one row that corresponds to starting in state $CC|CD$. Its entries give the probability of going from this state to any other. One entry in the row could represent the probability of the transition:

$$CC|CD \quad \rightarrow \quad CC|CC \tag{3.8}$$

I can calculate this probability by hand, assuming we have player 1 and player 2 with the following given strategies:

$$P_1 = \{P_c, P_d | p_{Cc}, p_{Cd}, p_{Dc}, p_{Dd}\} \quad P_2 = \{\tilde{P}_c, \tilde{P}_d | \tilde{p}_{Cc}, \tilde{p}_{Cd}, \tilde{p}_{Dc}, \tilde{p}_{Dd}\} \tag{3.9}$$

1. The first relevant term is the probability player 1 cooperates in the PGG given that in the last round of the PD, its partner defected. This is P_d.
2. The second relevant term is the probability player 2 cooperates in the PGG given that its partner cooperated with it in the last round PD. This is $\tilde{P}_c$.
3. The third relevant term is the probability that player 1 cooperates in the PD given that its partner defected in the last round of PD and cooperated in the last round of PGG. This is p_{Cd}.
4. The fourth relevant term is the probability player 2 cooperates in the PD given that its partner cooperated in the last round of PD and PGG. This is $\tilde{p}_{Cc}$.

Taken together, this gives us:

$$P(CC|CD \rightarrow CC|CC) = P_d \tilde{P}_c p_{Cd} \tilde{p}_{Cc} \tag{3.10}$$

Next, let's consider

$$P(CC|CD \rightarrow CC|DC) \tag{3.11}$$

I don't need to completely re-do my reasoning. Note that all of the items above are the same except the third. Before I was trying to find the probability player 1 *cooperates* in the PD given that its partner defected in the last round of PD. Here, I want to find the probability player 1 *defects* given the same situation. I can very simply write

$$P(CC|CD \to CC|DC) = P_d \tilde{P}_c (1 - p_{Cd}) \tilde{p}_{Cc} \tag{3.12}$$

In a similar way, once I know the the four probabilities associated with each state, then the probability of transition to each other state just involves various combinations of those probabilities and their complements.

3.2 Investigation of Civic Partner Strategies

There is one particular type of strategy I am especially interested in. This is the family of civic partner strategies.

Definition 2. *A strategy is a* **civic partner** *if it fits the following definition:*

1. *If two civic partner strategies are playing each other, they stay in state $CC|CC$ with probability 1.*
2. *In a population of n strategies, if $n-1$ of them are civic partner strategies and the one mutant is AllD, CD or DC, the payoff of the mutant cannot be higher than the average payoff of the civic partners.*
3. *No matter what its partner does, a civic partner cooperates with some probability strictly greater than 0. This means none of the entries in its vector of cooperation probabilities can be equal to 0.*

The second property can be rephrased as saying that a civic partner is **robust to invasion by AllD, CD and DC**.

As a reminder, the overarching goal of this thesis is to study how linking the PGG and PD games affects the ability of free riders to have high payoff. This definition was created to enable us to study a particular type of strategy that cooperates with high probability among its own type (property 1) and be generous (property 3) but does not allow less-generous strategies (like AllD, DC and CD) to invade by getting a higher payoff (property 2). The second property might seem a bit narrow because it only considers three mutant strategies. However, in the computer-aided portion of the thesis, I will show by simulation that a strategy that fulfills property 2 also has a higher average payoff against a single invader of *any* strategy, making a civic partner much more generally powerful. There can be multiple civic partner strategies that each cooperate with each other with probability 1, so sometimes I might refer to the "family of civic partner strategies".

The rest of this chapter will study the linked and unlinked strategy space, showing that civic partner strategies exist in the linked game but not the unlinked game. This supports the key premise of this thesis: **linking the games enables a strategy to exist that is generous and prevents the invasion of free-riders**.

3.3 Linked Civic Partner Investigation

Let's examine the criteria for a civic partner in sequence. The results of this section will give us the following theorem:

Theorem 1. *In the linked game, civic partner strategies exist and are of the form:*

$$\{1, P_d|1, p_{Cd}, p_{Dc}, p_{Dd}\} \tag{3.13}$$

where

$$0 < p_{Dd} < \frac{\frac{N-2}{N-1}b - c_G\frac{N-2+P_d}{N-1} - c_d\frac{N-2}{N-1}}{b + \frac{c_d}{N-1}} \quad 0 < P_d < 1 \tag{3.14}$$

$$0 < p_{Dc} < \frac{b - c_G + \frac{cd}{N-1}}{b + \frac{c_d}{N-1}} \quad 0 < p_{Cd} < \frac{b * \frac{N-2}{N-1} + \frac{c_G}{N-1}(1 - P_d) - c_d\frac{N-2}{N-1}}{b + \frac{c_d}{N-1}} \tag{3.15}$$

In the limit of large N, these bounds go to

$$0 < p_{Dd} < \frac{b - c_G - c_d}{b} \quad 0 < P_d < 1 \quad 0 < p_{Dc} < \frac{b - c_G}{b} \quad 0 < p_{Cd} < \frac{b - c_d}{b} \tag{3.16}$$

Note that this formulation looks very similar to Nowak's definition of a Generous Tit-For-Tat strategy. The cooperation probability in the case of defection again looks like $\frac{cost}{benefit}$. This could be because a civic partner strategy has very similar properties to a GTFT. Both punish defection by lowering their probability of cooperating, but still try to remain as generous as they can. Note that if the PGG game is dropped (c_G set to 0), the probability bounds are exactly what they are in the GTFT.

Additionally, this family of strategies has 3 of Axelrod's properties. They are nice (never defecting when they start in mutual cooperation), forgiving (cooperating with probability > 0) and retaliatory (responding to defection by lowering cooperation probability).

This proof has two parts correspondingly to the two main requirements.

Lemma 1. First Property *If two civic partner strategies are playing each other, they stay in state $CC|CC$ with probability 1.*

Proof. Let's consider some general strategy playing itself, apply the requirement, and see if a strategy that fits the requirement exists in the linked space.

$$P_1 = P_2 = \{P_c, P_d|p_{Cc}, p_{Cd}, p_{Dc}, p_{Dd}\} \tag{3.17}$$

Then I want:

$$P(CC|CC \rightarrow CC|CC) = 1 \tag{3.18}$$

Writing out this probability:

1. First item is the probability player 1 cooperates in the PGG, given its partner cooperated with it in the PD. This is P_c.

2. Second item is the probability player 2 cooperates in the PGG, given its partner cooperated with it in the PD. This is also P_c.

3. Third item is the probability player 1 cooperates in the PD, given its partner cooperated with it in the PD and PGG. This is p_{Cc}.

4. Fourth item is the probability player 2 cooperates in the PD, given its partner cooperated with it in the PD and PGG. This is also p_{Cc}.

Thus, I need

$$P_c^2 p_{Cc}^2 = 1 \quad \Rightarrow \quad P_c = p_{Cc} = 1 \quad \Rightarrow \quad L_{CP} = \{1, P_d | 1, p_{Cd}, p_{Dc}, p_{Dd}\} \tag{3.19}$$

where L_{CP} denotes a strategy that is a candidate civic partner. □

This gives us a specific subset of strategies in the linked game that have this first civic partner property, so the civic partner strategies must be a subset of this subset.

The second property is that a lone mutant of the form AllD, CD or DC will not be able to have a higher average payoff than the resident civic partner.

Lemma 2. *There exists a family of strategies in the linked game that, when playing each other, stay in state CC|CC with probability 1 and are robust to invasion by AllD, DC and CD.*

Proof. This is more tricky. I'll first try considering the case of AllD. In this scenario, I have N total players. $N-1$ of them are my (as of yet undefined) civic partner strategy while the last is AllD. The candidate civic partner strategies cooperate with each other in the PGG and PD with probability 1. The payoff for the $N-2$ civic partner players that are in **homogeneous pairs** is (pairs of two identical strategies playing each other)

$$\left(\frac{N-2}{N} * r * c_G + \frac{1}{N} * r * P_d * c_G - c_G\right) + (b - c_d) \tag{3.20}$$

The first part is the payoff of the PGG: each player in the homogeneous pairs contributes and the candidate civic partner contributes with probability P_d. Each player in the homogeneous pairs has cost c_G. In the PD, each candidate civic partner cooperates with probability 1 and so has cost c_d and benefit b.

Next, I'll calculate the payoff of the civic partner player paired with the AllD player. The PGG payoff is almost identical, with the only change being that it only pays a cost c_G with probability P_d when it cooperates with the defector. The PD is different. It never receives a benefit b because its partner never cooperates. With probability p_{Dd}, it pays a cost c_d to cooperate. This works out to:

$$\left(\frac{N-2}{N} * r * c_G + \frac{1}{N} * r * P_d * c_G - P_d * c_G\right) + (-p_{Dd} * c_d) \tag{3.21}$$

Finally, the payoff of the AllD player is below. It receives the same benefit from the PGG as all the players, but never pays a cost. It receives a benefit b in the PD with probability p_{Dd} but never pays any cost.

$$\left(\frac{N-2}{N} * r * c_G + \frac{1}{N} * r * P_d * c_G\right) + (b * p_{Dd}) \tag{3.22}$$

In order for my desired property to hold, I must have

$$\text{payoff of AllD} < \frac{(N-2)\text{payoff of homogenous civic partner} + \text{payoff of last civic partner}}{N-1} \tag{3.23}$$

Note that each player gains exactly the same benefit from the PGG, so this can cancel from both sides. First, let's write out the righthand side. This is

$$(-\frac{N-2}{N-1} * c_G * r - \frac{1}{N-1} * P_d * c_G) + (\frac{N-2}{N-1} * b - c_d * \frac{N-2}{N-1} - c_d \frac{p_{Dd}}{N-1}) \tag{3.24}$$

Note again that as $N \to \infty$, this simplifies to $r * c_G - c_G + b - c_d$, the payoff of a population of entirely civic partner players. Therefore, in the limit as $N \to \infty$, the full equation simplifies to

$$(r * c_G) + (b * p_{Dd}) < r * c_G - c_G + b - c_d \tag{3.25}$$

$$p_{Dd} < \frac{b - c_G - c_d}{b} \tag{3.26}$$

In the small N case, this is written as:

$$p_{Dd} < \frac{\frac{N-2}{N-1}b - c_G \frac{N-2+P_d}{N-1} - c_d \frac{N-2}{N-1}}{b + \frac{c_d}{N-1}} \tag{3.27}$$

Next, I need to investigate p_{Cd} and p_{Dc}, two terms I have not determined as of yet. I can do this by repeating the procedure above. I know that a civic partner strategy must have a higher average payoff than any other strategy, including (00|1111), the strategy that always cooperates in the PD but never in the PGG. Let's name this DC to represent the first portion (defection) and the second portion (cooperation). Note that because DC always cooperates in the PD, the civic partner it is paired with always cooperates in the PGG and cooperates in the PD with probability p_{Dc}. In this setup, the payoff of each player in the homogeneous civic partner pairs is

$$(\frac{N-1}{N} * r * c_G - c_G) + (b - c_d) \tag{3.28}$$

The payoff of the civic partner paired with DC is

$$(\frac{N-1}{N} * r * c_G - c_G) + (b - c_d * p_{Dc}) \tag{3.29}$$

DC's average payoff is

$$(\frac{N-1}{N} * r * c_G) + (p_{Dc} b - c_d) \tag{3.30}$$

I set up the inequality equation as before, requiring that the average payoff of the civic partner is higher than the payoff of DC. The right hand side of the equation is

$$\frac{N-1}{N} * r * c_G - c_G + b - \frac{N-2+p_{Dc}}{N-1} c_d \tag{3.31}$$

Cancelling out parts of the PGG that are the same, I find my inequality becomes:

$$p_{Dc} * b - c_d < -c_G + b - \frac{N-2+p_{Dc}}{N-1} * c_d \tag{3.32}$$

which simplifies out to

$$p_{Dc} < \frac{b - c_G + \frac{c_d}{N-1}}{b + \frac{c_d}{N-1}} \tag{3.33}$$

In the case where $N \to \infty$, this simplifies to

$$p_{Dc} < \frac{b - c_G}{b} \tag{3.34}$$

As compared with the bound for p_{Dd}, this term has dropped the c_d in the numerator. This is because p_{Dc} considers the case where the player's partner has defected in the PGG only. Defection increases its partner's payoff by c_G, the money it would have spent if it cooperated in the PD. If the goal of the candidate civic partner, then, is to limit the payoff of strategies that defect, then the probability of cooperating with them must be lower when the player gets a very high payoff from defection, so when c_G is large. This is exactly what this probability gives us.

Finally, let's consider how a civic partner would fare against a player with strategy (11|0000), CD, who always cooperates in the PGG but never in the PD. The civic partner in the heterogenous pair would cooperate in the PGG with probability P_d (previously calculated) and in the PD with probability p_{Cd}. The civic partner players in the homogeneous pairs would have payoff

$$(\frac{N-1}{N} * r * c_G + \frac{1}{N} * r * c_G * P_d - c_G) + (b - c_d) \tag{3.35}$$

The civic partner in the heterogeneous pair would have payoff

$$(\frac{N-1}{N} * r * c_G + \frac{1}{N} * r * c_G * P_d - c_G * P_d) + (-c_d * p_{Cd}) \tag{3.36}$$

Finally, the CD strategy would have payoff

$$(\frac{N-1}{N} * r * c_G + \frac{1}{N} * r * c_G * P_d - c_G) + (b * p_{Cd}) \tag{3.37}$$

I again want a average civic partner payoff to be higher than the payoff in the CD player. This simplifies to

$$-c_G + b * p_{Cd} < \frac{-c_G(N - 2 + P_d) + (N-2) * b - c_d * (N - 2 + p_{Cd})}{N-1} \tag{3.38}$$

This gives us a bound of

$$p_{Cd} < \frac{b * \frac{N-2}{N-1} + \frac{c_G}{N-1}(1 - P_d) - c_d \frac{N-2}{N-1}}{b + \frac{c_d}{N-1}} \tag{3.39}$$

In the limit of $N \to \infty$, this simplifies to

$$p_{Cd} < \frac{b - c_d}{b} \tag{3.40}$$

Similarly to the last case discussed, this only involves one of the costs, here the one relating to the benefit that the partner got by defecting in the PD. □

The final requirement of the civic partner definition requires that no entry in the vector of probabilities is equal to 0. This is possible if all of the upper bounds I just calculated, p_{Cd} etc are greater than 0. In most experimental cases they are because it is common practice to assume the benefit b is greater than the costs c_d or c_g. Thus, these lemmas show that the linked space contains a family of civic partner strategies: all strategies whose cooperation probabilities satisfy the requirements of the civic partner. Next, let's move on to considering the unlinked case.

3.4 Unlinked Civic Partner Investigation

This section investigates whether a civic partner strategy exists in the unlinked strategy space and concludes that it does not. There are strategies that stay in $CC|CC$ with probability 1 and are robust to invasion from AllD, which I will call a **weak civic partner** strategy. However, these are not civic partner strategies.

Theorem 2. *The family of weak civic partner strategies for the unlinked game in the limit as $N \to \infty$ looks like*

$$U_{wcp} = \{1, P_d|1, p_d\} \quad 0 < p_d < \frac{b - c_d - c_g}{b} \quad 0 < P_d < 1 \tag{3.41}$$

In the small N case, this family looks like

$$U_{wcp} = \{1, P_d|1, p_d\} \quad 0 < p_d < \frac{(N-2)(b-c_d) - c_g(N-2+P_d)}{(N-1)b + c_d} \quad 0 < P_d < 1 \tag{3.42}$$

Lemma 3. *There exists a family of strategies in the unlinked game that with probability 1 continues mutual cooperation ($CC|CC$).*

Proof. This is exactly the same analysis as for the linked case. If I have,

$$P_1 = P_2 = U_{wcp} = \{P_c, P_d|p_c, p_d\} \tag{3.43}$$

then I want

$$P(CC|CC \to CC|CC) = 1 \tag{3.44}$$

where player 1 and player 2 both have the same strategy. Writing out this probability in the identical way as done in the linked game tells me that I require

$$P_c^2 p_c^2 = 1 \quad \to \quad P_c = p_c = 1 \tag{3.45}$$

□

Lemma 4. *In the unlinked game, there exists a subset of the family of strategies stay in $CC|CC$ with probability 1 and is also robust to invasion from AllD.*

Proof. First, let's consider the payoff of the homogeneous weak civic partner pairs. In the PGG, they would each contribute c_g. This total sum of $(N-1)c_g$ would be multiplied to become $r(N-1)c_g$. The weak civic partner in the heterogeneous pairs cooperates in the PGG with probability P_d. In the PD, each player would pay c_d to give its partner b, which it would also receive in turn. So the overall payoff is

$$(r * c_g - c_g) + (b - c_d) \tag{3.46}$$

Next, let's consider a situation where the population of almost all pairs has civic partners, but one pair has a civic partner and an AllD player. The civic partner payoff would go as follows: in the PD, with probability p_d it will pay a cost c_d. It will never receive any benefit. In the PGG, it will receive a payout $\frac{N-2}{N} r * c_g$ always, and with probability P_d it will pay a cost c_g and receive an additional $\frac{1}{N} r * c_g$ from its own contribution. Its payoff is

$$(\frac{N-2}{N} r * c_g + P_d \frac{1}{N} r * c_g - P_d c_g) + (-p_d * c_d) \tag{3.47}$$

Finally, let's consider the case of the payoff of the AllD player. With probability p_d, it receives a benefit b in the PD. It receives a benefit of $\frac{N-2}{N} r * c_g + \frac{1}{N} * r * c_G$ in PGG. It has no costs. Therefore, its overall payoff is

$$(\frac{N-2}{N} r * c_g + P_d \frac{1}{N} r * c_g) + (p_d * b) \tag{3.48}$$

If I want average payoff of the civic partner strategies to be higher than that of the non- civic partner, I would want

$$\text{payoff of AlllD} < \frac{(N-2)\text{payoff of WCP player pairs} + \text{payoff of WCP player paired with AllD}}{N-1} \tag{3.49}$$

Cancelling the PGG benefit, which is identical for all players, gives us

$$p_d * b < -c_g \frac{(N-2) + P_d}{N-1} + \frac{N-2}{N-1}(b - c_d) - \frac{p_d}{N-1} c_d \tag{3.50}$$

$$p_d < \frac{b * \frac{N-2}{N-1} - c_d * \frac{N-2}{N-1} - c_G \frac{N-2+P_d}{N-1}}{b + \frac{cd}{N-1}} \tag{3.51}$$

In the limit where $N \to \infty$, then $(N-1)/(N-2) \approx 1$, and since $P_d \in [0,1]$, $(N-2+P_d)/N \approx 1$ so I can write

$$p_d < \frac{b - c_d - c_g}{b} \tag{3.52}$$

This makes sense: as $N \to \infty$, the average payoff increasingly looks like the payoff of all the candidate civic partner players not in the AllD pair. Additionally, note if we set $c_g = 0$ then this reduces down to the case where the PGG is irrelevant. This is the exact Generous Tit-for-Tat situation in the repeated Prisoner's Dilemma case (without including the PGG game) that Nowak proposed[1]. □

Finally, in order for the weak civic partner defined here to be forgiving, I require that the probability of cooperation in any situation be strictly greater than 0. This is usually possible, because for most formulations $b > c_d + c_g$.

Having defined a weak civic partner strategy, what does this tell us about whether the unlinked strategy space has a civic partner strategy?

Lemma 5. *A weak civic partner strategy as defined above is not a civic partner strategy.*

Proof. In order to show that the weak civic partner defined above isn't a civic partner, I will use the fact that the strategy space of the unlinked game is a subset of the strategy space of the linked game. This was explained in the introduction where it was shown how each unlinked strategy can be written as a strategy in the linked space. Therefore, if a civic partner strategy exists in the unlinked strategy space, it should also be a member of the civic partner strategy family in the linked space. Additionally, it should also be a weak civic partner, based on the definitions of civic partner and weak civic partner. A weak civic partner strategy is of the form

$$\{1, P_d | 1, p_d\} \quad \to \quad \{1, P_d | 1, p_d, 1, p_d\} \tag{3.53}$$

[1] [7]

where $P_d, p_d < 1$ and where the arrow indicates translation into the linked space. As proven in the previous section, the civic partner is of the form

$$\{1, P_d | 1, p_{Cd}, p_{Dc}, p_{Dd}\} \tag{3.54}$$

where $P_d, p_{Cd}, p_{Dc}, p_{Dd} < 1$. If the weak civic partner is a civic partner, it must fit the form of the above equation. However, this is impossible: the fifth entry in the civic partner is $p_{Dc} < 1$ and the fifth entry in the weak civic partner strategy is exactly 1. Therefore, the unlinked game contains a weak civic partner strategy family but no civic partner strategies. □

Chapter 4

Computer-Aided Analysis

Certain mathematical questions can be solved analytically, as was done in previous sections. However, in studying populations of strategies, I often wish to investigate systems that are too complicated to study by hand because they involve too many players or involve randomness in complicated ways. This chapter contains the results of experiments done using Matlab code I wrote for this thesis. Each portion will grow progressively more complicated, including more sources of randomness and thus better approximating the messy, complicated world we live in. A summary of the experiments in order of increasing complexity goes as:

1. Calculating the payoff of resident civic partner players against a randomly-generated mutant. Randomness: selection of mutant strategy.

2. Ability of resident civic partner to defend against invading mutant, and ability of civic partner to invade a resident mutant. Additional randomness: reproduction and death proportional to payoff.

3. In a population of 7 named strategies, ability of the civic partner strategies to become the resident and/or maintain high prevalence in population. Additional randomness: interactions of multiple strategies at a time and a mutation rate during reproduction.

In every case, the unlinked case has been retained for comparison.

4.1 Theoretical Background

Because this section involves computer simulations, some portions of the set-up had to be translated from their analytical formulation. Fortunately, there are commonly-accepted ways to do corresponding work in analytical and computer methods, many related to the study of Markov Chains. This portion of the chapter will explain them.

4.1.1 Finding Payoff Via Eigenvectors

Let's consider the transition matrix examined in the previous chapter, denoted T. Suppose two players each start in some state, like $CC|DC$. I can represent this as a 16 by 1 vector v with a 1 in the location corresponding to the current state and 0s all elsewhere. The term

$$v \cdot T$$

is again a 16 by 1 vector. Here, the entries represent probabilistic distributions over each combination of cooperate/defect situations. For example, a 0.23 in the row corresponding to $CC|CC$ means it has a 23% chance of moving to that particular combination. Some entries may be 0, but the entire vector will sum to 1 to because there is probability 1 of existing in some state. Next, the quantity

$$(v \cdot T) \cdot T$$

is the probability distribution of particular scenarios after two games have been played between those players since they started with their definite history, like $CC|DC$. I could continue this process infinitely many times. There is a certain type of vector, v', that has the property

$$v' \cdot T = v'$$

Mathematically, it is an eigenvector with eigenvalue 1. In this context, this means that playing another round doesn't change the distribution of probabilities. It is a **steady state**. In order to calculate the payoff between two players after infinite rounds of the games, I follow three simple steps[1].

1. Calculate the payoff for each player under each situation, ie $CC|DC$.
2. Multiply the payoff for each player against the probability of each situation occuring in the long-term steady state vector v'.
3. For each player, sum up their total expected payoff across the multiplied values calculated above. Use these two numbers for comparisons between the two strategies and reproduction, which is proportional to payoff.

Let's call this type of strategy the **matrix-driven method**. This method has been frequently used by other mathematicians in analyzing Markov Chains. However, it's not immediately clear why this method is any better than the following:

1. Two players are matched and they play a series of rounds of games with each other.
2. For each round, players consider their partner's history of actions. If their strategy space is probablistic, then they decide whether to cooperate or not based off of a randomly-generated number. For example, if a partner has a 70% chance of cooperating in a particular situation, they cooperate if a random number between 0 and 1 is less than 0.7 and defect if it is greater than .3.
3. After a certain number of rounds, but at least 1, reproduction can be done proportional to payoff.

This **many-rounds method** procedure seems much more logical and is consistent with how previous portions of this thesis have been talking about how two strategies interact. However, it has two key drawbacks. The first is that the many-rounds method is *much* slower to run on a computer. How much slower depends on the exact method by which you run the simulation. Later portions of this chapter will discuss simulations done with 100,000 rounds of reproduction each involving 20 players. Using the matrix method, such a simulation can be done a laptop computer within 6 minutes, but using the second method could easily take an hour, making experimentation long and

[1]This explanation is drawn from Sigmund's "Calculus of Selfishness"[3].

cumbersome. The second reason involves precision and repeatability. The many-rounds method involves randomly generated numbers. Randomness can be useful, but it also adds uncertainty into our calculations. For example, consider the strategies

$$S_1 = [\epsilon, \epsilon|\epsilon, \epsilon, \epsilon, \epsilon] \quad S_2 = [1-\epsilon, 1-\epsilon|1-\epsilon, 1-\epsilon, 1-\epsilon, 1-\epsilon] \tag{4.1}$$

for some small ϵ. For any particular situation, S_2 is more likely to cooperate than S_1. I would expect that if we played the two strategies against each other, S_1 would spend most of its time defecting and S_2 would spend most of its time cooperating. However, let's take the other approach and consider the probability that over N rounds, the players behave very uncharacteristically: S_1 spends all of its time cooperating or S_2 spends all of its time defecting. The probability of the first situation situation happening is

$$P(S_1 \text{ cooperates N times out of N trials}) = P(S_2 \text{ defects N times out of N trials}) = \epsilon^N \tag{4.2}$$

Alarmingly, this probability is strictly greater than 0! Such a result would give us a severely skewed interpretation of the true way the strategies usually interact. For some common values, say $\epsilon = .01$ and $N = 5$, this type of situation might happen with probability 10^{-5}. Over the course of 100,000 separate payoff calculations between players, I would expect such an event to occur on average, once. I could lower this probability by increasing N (thereby increasing our computational time!) but never eliminate it completely. Additionally, the exact situation posited in the above equation isn't the only "unusual" situation that could result. S_1 cooperating $N-1$ times and defecting once or S_2 defecting $N-1$ times and cooperating once. Either of these could occur with probability $(N-1)*\epsilon^{N-1}(1-\epsilon)$, and would also be an unfortunate situation because it is far from the "typical behavior" I might expect from looking at their strategies. By contrast, the matrix-driven method given both player's strategies spits out the same vector each time describing the average payoff. Essentially, the matrix-driven method gives the same result as for the many-rounds method if it were done infinitely long, but it's repeatable and very speedy. Because of these benefits, the remainder of this thesis will use the matrix-driven method for payoff calculation.

4.1.2 Perron-Frebonious

However, the matrix-driven method isn't without drawbacks. Most notably, the algorithm provided in the previous section assumed that the transition matrix T has exactly one eigenvector with eigenvalue 1. For any general transition matrix T, this isn't necessarily true. For example, consider the situation where

$$S_1 = S_2 = [1, 0|1, 0, 1, 0] \tag{4.3}$$

Each strategy is a type of Tit-for-Tat. Logically considering the strategies, it becomes obvious that there are 3 potential long-term situations. If both players start out cooperating, they will continue cooperating forever. If they both start out defecting, they will defect forever. If one starts out cooperating and one defecting, they will alternate between cooperation and defection forever. Indeed, the matrix procedure correctly identifies 3 eigenvectors with eigenvalue 1, but the correct probability distribution can't be found by simply averaging over these three because the third starting position (one cooperating, one defecting) would occur twice as frequently as the others for random initial conditions. It's difficult to write a standard procedure to identify and weight each of these situations correctly.

However, let's examine the situation if we change each entry slightly from begin either 0 or 1, writing instead

$$S_1 = S_2 = [1-\epsilon, \epsilon | 1-\epsilon, \epsilon, 1-\epsilon, \epsilon] \tag{4.4}$$

Then, the matrix-driven procedure gives us one matrix with approximately .25 in entries corresponding to each of the situations where both players are defecting, where both are cooperating, where player 1 cooperates while player 2 defects, and where player 2 cooperates while player 1 defects. Essentially, it has weighted each of the 4 separate situations correctly for us. The price was the addition of an error rate, which can be made arbitrarily small. I get this result because of a theorem called Perron-Frebeonius.

4.1.3 Perron-Febronious: Theory and Proof

This proof is drawn from the lecture notes of one of my advisers, Professor Clifford Taubes[2]. Perron-Febronious theory studies transition matrices T for Markov chains. There are several useful results it gives, but the one I am most interested in is:

Theorem 3. *If a transition matrix T has all entries strictly greater than 0 and each row sums to 1, then the kernel of $T - I$ (where I is the identity matrix) is one-dimensional. Additionally, there is a unique vector in this kernel whose entries are positive and sum to 1.*

This theorem, if true, tells us that for some transition matrix with positive entries everywhere, a single steady-state (eigenvalue of 1) vector of probabilities (positive, sums to 1) exists. This is exactly what I want for my transition matrix: one vector that describes the steady-state situation for the two players will always exist! This will be proved through three lemmas.

Lemma 6. *Any transition matrix T with non-negative entries has at least one eigenvector with eigenvalue 1.*

Proof. Proving this lemma is equivalent to proving that the kernel of $T - I$ is non-trivial because

$$a(T - I) = 0 \quad \Rightarrow \quad aT = aI = a \tag{4.5}$$

If a matrix's kernel is non-trivial, then because of rank-nullity, the dimension of the image is smaller than the dimension of the side of the matrix, n. This means the map that T represents cannot be injective and thus cannot be inverted. Because the map can't be inverted, the matrix T is not invertible, so its determinant must be 0. Proving that $det(T - I) = 0$ would tell me that the kernel is non-trivial. To do this, let's note that $det(T - I) = det(T^T - I)$. This comes from the fact that $det(A) = det(A^T)$ and $(T - I)^T = T^T - I$. My problem has been reduced to proving that $det(T^T - I)$ is zero, equivalent to showing it has some non-zero element in its kernel. Let's consider w, the vector with 1 in all of its entries. I have that

$$w(T^T - I) = wT^T - wI = w - w = 0 \tag{4.6}$$

This is because of how I defined T: the rows sum of 1, so T^T's columns sum to 1. The operation wT^T is summing along the columns, so each entry of the output will be exactly equal to 1, as desired. Because the kernel of $T^T - I$ is non-trivial, the kernel of $T - I$ is non-trivial, and there

[2] [11]

exists some v such that $v(T-I)=0$ and so $vT=v$, an eigenvector with eigenvalue 1. This lemma could be summarized as

$$ker(T-I) \text{ non-trivial} \Leftrightarrow det(T-I)=0 \Leftrightarrow det(T^T-I)=0 \Leftrightarrow ker(T^T-I) \text{ non-trivial} \tag{4.7}$$

□

Lemma 7. *If each entry in the transition matrix T is strictly positive, then every non-zero vector in the kernel of $T-I$ has either all positive entries or all negative ones.*

Proof. This is a proof by contradiction. Assume there's some v in the kernel of $T-I$ with $k<n$ entries ≤ 0 and all the others >0. Because the indexing of matrix T is arbitrary, I can rearrange them to have the first k entries of $v \leq 0$ and all following entries ≥ 0. Then, if I have $vT=v$, the output of vT must equal v. Looking at the first k entries of vT and v gives me these k equations:

$$v_1 = v_1 * T_{1,1} + v_2 * T_{2,1} + \ldots v_n * T_{n,1} \quad \ldots \quad v_k = v_1 * T_{1,k} + \ldots + v_n * T_{n,k} \tag{4.8}$$

If I add these equations together, I get

$$(v_1 + \ldots + v_k) = v_1 * (T_{1,1} + \ldots + T_{1,k}) + \ldots v_n * (T_{n,1} + \ldots + T_{n,k}) \tag{4.9}$$

If I subtract to gather all the $v_1, \ldots v_k$ together on the left hand side, I get

$$v_1(1-(T_{1,1}+\ldots+T_{1,k}))+\ldots+v_k(1-(T_{k,1}+\ldots+T_{k,k})) = v_{k+1}(T_{k+1,1}+\ldots T_{k+1,k})+\ldots v_n(T_{n,1}+\ldots T_{n,k}) \tag{4.10}$$

The $T_{i,1}+\ldots+T_{i,k}$ terms are summing across the rows of the matrix T, so they are positive, and for $k<n$ sum to something strictly less than 1. Therefore, the right hand side is strictly positive (because each v_i term is strictly positive by hypothesis) and the lefthand side is negative or 0 (because the v_i terms on the lefthand side are 0 or negative by hypothesis). This is a contradiction, so it must be impossible that some entries of v are positive and some negative. □

Lemma 8. *If all the entries of T strictly positive, then the kernel of $T-I$ (ie, the eigenvectors) is one-dimensional, and there is a unique vector in the kernel whose entries are positive and sum to 1.*

Proof. Again, let's do a proof by contradiction. Assume there exist two seperate u', v' where one is not the multiple of the other, but both of them are in the kernel of $T-I$. By the previous lemma, I know the entries of u', v' are both either all positive or all negative. We can normalize them by dividing each by their sum to obtain u, v that are both all positive and sum to 1. If $u \neq v$ then there is some first entry where they differ. WLOG, let's assume $u_1 < v_1$. Because both u and v sum to 1, there must be some later i such that $u_i > v_i$. WLOG, let's assume $u_2 > v_2$. I know that $u-v$ is in the kernel of $T-I$ because

$$(u-v)(T-I) = uT - vT - uI + vI = u(T-I) - v(T-I) = 0 \tag{4.11}$$

However, from the previous lemma I know that all vectors in the kernel of $T-I$ must have all positive or all negative entries. I again arrive at a contradiction, so all vectors v in the kernel of $T-I$ must be multiples of each other. Given this, I can take any vector in the kernel and divide it by the sum of its entries to get a vector v^* that is positive and whose entries sum to 1. No other vector in the kernel can have this property because it must be a multiple of v^* and therefore cannot sum to 1. □

These three lemmas taken together form the entirety of the proof of the theorem.

4.1.4 Fitness

Previous portions of this thesis have frequently relied on payoff as a method of comparing the strength of two strategies. However, reproduction generally is done not involving not payoff, but **fitness**. In this thesis, the fitness of a player, $F(S_1)$ is a function of their payoff, π_i, in a particular situation. The equation is

$$F(S_1) = e^{\beta * \pi_1} \quad \beta \geq 0 \tag{4.12}$$

The probability of reproduction comes from normalizing this factor:

$$P(P_1 \text{ is selected for reproduction }) = \frac{e^{\beta * \pi_1}}{\sum_{i=1}^{N} e^{\beta * \pi_i}} \tag{4.13}$$

β is a scaling factor that denotes the strength of selection. Larger β means that a the same difference in payoff translates into a greater difference in probability of being selected for reproduction. Note that for $\beta = 0$ a player's fitness does not depend on their payoff, so reproduction is at random. A negative β value means that higher payoff indicates lower likelihood of being selected for reproduction. For almost all situations I will consider $\beta > 0$ only. To see how β works, consider a simplified version with two players. One has payoff π_1 and the other has payoff $\pi_1 + a$. The probability of the first being selected for reproduction is

$$\frac{e^{\beta * \pi_1}}{e^{\beta * \pi_1} + e^{\beta * (\pi_1 + a)}} = \frac{e^{\beta * \pi_1}}{e^{\beta * \pi_1}(1 + e^{\beta * a})} = \frac{1}{1 + e^{\beta * a}} \tag{4.14}$$

The probability of the player with the lower payoff getting selected decreases exponentially with increased a and increased β. A higher β value can be more desirable because it speeds up evolution and reduces the chance of a sub-par strategy reproducing, but it also makes the procedure more determinist, perhaps detrimentally so. For most portions of this thesis, a β value of 1 is used in reproduction, and death is done by random selection.

4.1.5 Mutation Rate

Mutation occurs during reproduction. Usually, a dying player's strategy is replaced by the reproducing player's strategy. With mutation, the dying player's strategy is replaced by a randomly-selected strategy. Mutation adds to the complexity of a simulation, but it also adds to the realism of the the model because mutation is common in areas like biological reproduction.

Additionally, mutation is useful because it ensures strategies are introduced in a variety of scenarios throughout the game. Perhaps some strategy S_1 fails to do well in the beginning of a simulation because another strategy S_2 is dominant, and S_1 does particularly badly against S_2. However, during a later portion of the simulation maybe S_3 is more frequent, but S_3 does badly against S_1. Without re-introducing S_1 to the population that would be impossible to realize this. A mutation rate describes the probability of any given round of death and reproduction involving mutation. Unless otherwise noted, the most frequently used mutation rates are either 0 or 2.5%, which means that 2.5% of the time, when reproduction occurs, a new mutant is generated.

4.2 Experiments Involving Computers and their Results

This section describes the actual experiments I ran. They were all done in Matlab using code written entirely by me.

4.2.1 Mutant Payoff Calculation

The analytical portion of this thesis designed a civic partner with respect to three possible invading mutants: AllD, DC and CD. However, I mentioned that designing the civic partner so it is resistant to invasion from these three strategies is sufficient for the resident civic partner to have on average a higher payoff than any other invading mutant. In this portion I will support my analytical solution by actually checking the average payoff of civic partners against many, many randomly-generated invading mutants. Because each number detailing the strategy can be any real number between 0 and 1, the strategy space for the linked game is continuous and 6 dimensional, making it extremely difficult to explore analytically. Instead, I used a method similar to that used in "Stochastic evolutionary dynamics of direct reciprocity" by Lorens Imhof and Martin Nowak[3].

The set-up is as follows: in total, there are n players in the population, $n-1$ of which are civic partners. Of these, $n-2$ of these are in homogeneous pairs (both participants are civic partner). One civic partner is paired with a mutant whose strategy is randomly generated. The program calculates the steady-state payoff of the homogeneous civic partner, the civic partner in the heterogeneous pair, and the mutant. It stores this information. While Imhof and Nowak's process included reproduction and death, my method differs. In the interest of exploring the strategy space as quickly as possible, my method immediately discards the mutant, generates another one, and calculates the payoff for that. This methodology allows me to rapidly explore a large number of mutants and see if there are mutants that have a higher payoff than the average civic partner payoff.

As discussed before, in order to calculate the steady-state probability distribution, it is necessary that the transition matrix doesn't have any 0s. This requires the addition of an error rate ϵ to be added to probabilities of cooperation that are equal to 0 and subtracted from probabilities of cooperation that are equal to 0 for each player's strategy. Repeated experimentation (omitted here, for space considerations) shows that varying ϵ does not dramatically change the results. Unless otherwise noted, all graphs made below were done with $c_G = c_d = 1$, $b = 5$, $r = 3$, $\epsilon = 10^{-5}$ and $P_d = 1$. Usually the tests were done with around 10,000 randomly generated mutants.

Civic Partner Strategies (in the Linked Game)

The beginning of this chapter (and the previous chapter calculating the linked civic partner formula) gave a strict upper bound on civic partner cooperation probabilities, but no lower ones. These upper bound values are denoted by P_d, p_{Cd} etc. This section investigated civic partner strategies of the form

$$\{1, P_d * (1-\lambda)|1, p_{Cd}(P'_d) * (1-\lambda), p_{Dc} * (1-\lambda), p_{Dd}(P'_d) * (1-\lambda)\} \tag{4.15}$$

where $\lambda \in [0,1]$ and p_{Cd} and p_{Dd} have been written to explicitly note the upper bound's dependence on $P_d(1-\lambda)$, denoted P'_d. Lower levels of λ lead to higher levels of cooperation with less-cooperative strategies, which could be considered generosity or forgiveness. At higher levels of λ, the civic partner begins to approximate TFT, which here I interpret as $[1, 0, |1, 0, 0, 0]$.

Figure 4.1 shows the results of trials with three different types of civic partners: one that is the most generous ($\lambda = 0$) and two others that are slightly less generous ($\lambda = .2$) and much less generous ($\lambda = .7$). Note that the $\lambda = 0$ strategy is not technically a civic partner strategy because it corresponds to probabilities of cooperation exactly at the upper limit. However, it is included as

[3] [2]

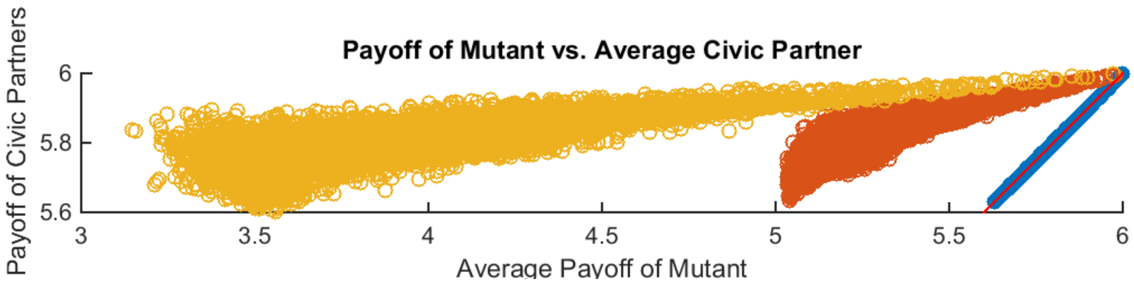

Figure 4.1: Each dot (of which there are 10,000 in each color) corresponds to the ordered pair {payoff of mutant, payoff of average civic partner} for a randomly-generated mutant. The blue dots correspond to trials with civic partners that are the most generous. The orange dots are civic partners that are a bit less generous, and the yellow dots are related to trials with civic partners that are the least generous. A red line gives the $y = x$ axis. Above this line, the civic partners have higher payoff, and below it, the mutant has a higher payoff.

reference case. This figure shows that the $\lambda = 0$ strategy has exactly the same payoff as the mutant for all 10,000 mutants. The less generous strategies do better than the mutant for every mutant tried. This proves a key fact about civic partners: defining them to be robust to AllD, DC and CD was sufficient to make them robust to *all* strategies in the linked space! This is a very useful result - it shows that by building a civic partner to be robust to invasion by only three specific mutants (CD, DC and AllD), we automatically get that it is resistant to invasion from the entire space of mutants!

Additionally, all of the shapes in the figure slope upwards. This indicates that a mutant in general can only get a higher payoff if the civic partner also gets a higher payoff. In other words, the civic partner doesn't allow one side or the other to gain through exploitation.

The experiment described above was done with $n = 20$ total players in the population. Because of this, I used the formula for calculating the civic partner in the small n case. Next, I was interested in seeing how the properties of the civic partner varied with different population size. Having more civic partners and fewer mutants means that the average payoff of all the civic partners is more determined by the civic partners in homogeneous pairs (playing other civic partners). Because of this, the civic partner playing the mutant can afford to have lower payoff and still have the overall civic partner payoff be reasonably high.

Figure 4.2 shows a trial with three different types of civic partner strategies in populations ranging from small ($n = 10$) to medium ($n = 20$) to large ($n = 500$). All of these civic partners were done with $\lambda = .2$. Note that increasing the population size causes the average payoff of the civic partners to move up and away from the $y = x$ line, so it does even better than the mutant. In the limit of large n, the average civic partner payoff is exactly the same no matter what mutant is present.

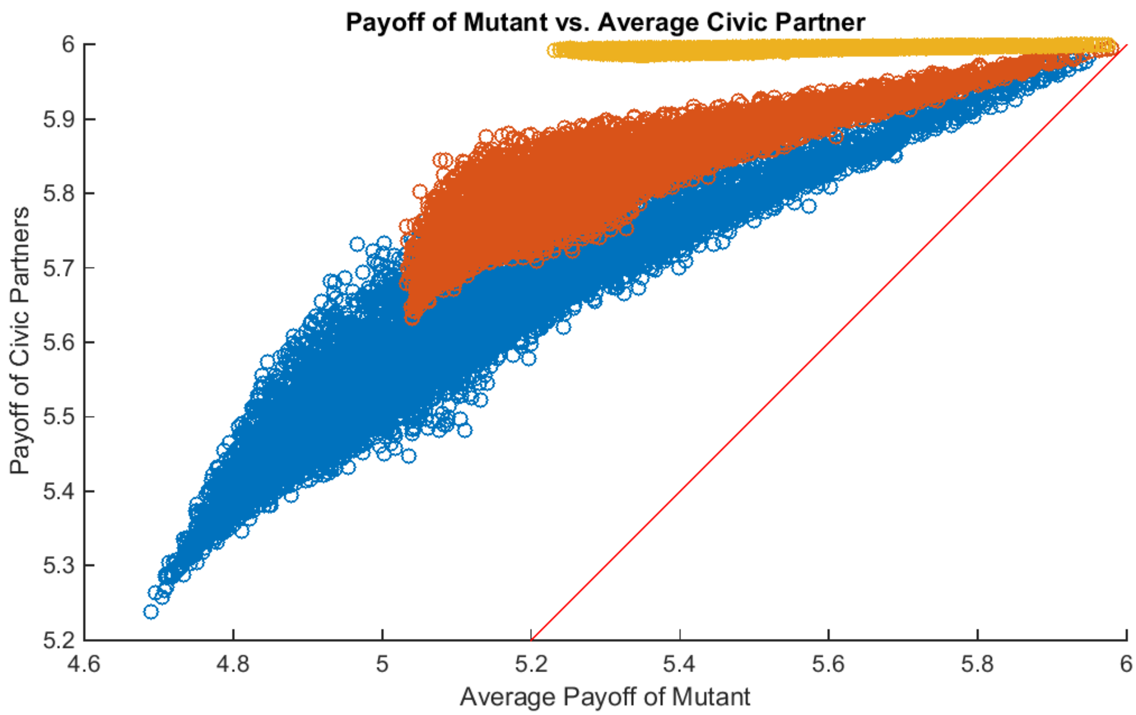

Figure 4.2: Each dot (of which there are 10,000 in each color) corresponds to the ordered pair {payoff of mutant, payoff of average civic partner} for a randomly-generated mutant. The blue dots correspond to civic partners in a population of $n = 10$ total players, the orange to civic partners in a population of 20 players, and the yellow to civic partners in a population of 500 total players. The red line is the $y = x$ line. Points above the line indicate that the civic partners had a higher payoff than the mutant, and points below it indicate the opposite occurred.

The next question I investigated was how exactly the difference between civic partner and mutant payoff was kept to within such a narrow range even while tested across tens of thousands of mutants. The cooperation level in the PD and PGG between homogeneous civic partner pairs stays constant regardless of the mutant, so instead I looked at the relationship between the payoff of the civic partner paired with the mutant (afterwards called the "Partner") and the payoff of the mutant itself.

Figure 4.3 shows graphs with those payoff charts. The top chart shows how decreasing λ (making the civic partner more generous) enables the mutant to more frequently get a higher payoff than the civic partner, as shown by how more of the dots are below the $y = x$ line. Additionally, it makes the relationship more linear. The second chart shows how varying the population size makes the relationship between the payoffs even more constrained.

The civic partner's strategy here has links to Hilbe et al's "Partners or rivals? Strategies for the iterated prisoners dilemma". In it, the authors define a strategy A to be a "partner" strategy if the following condition is met: if the strategy A player gets a payoff less than it would under full

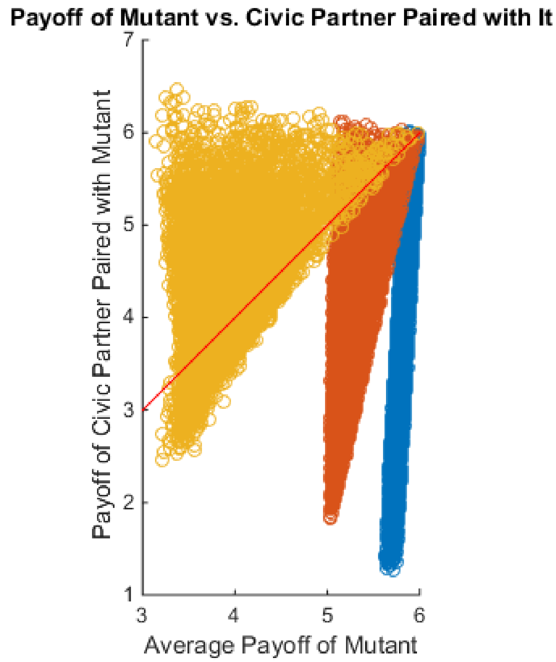

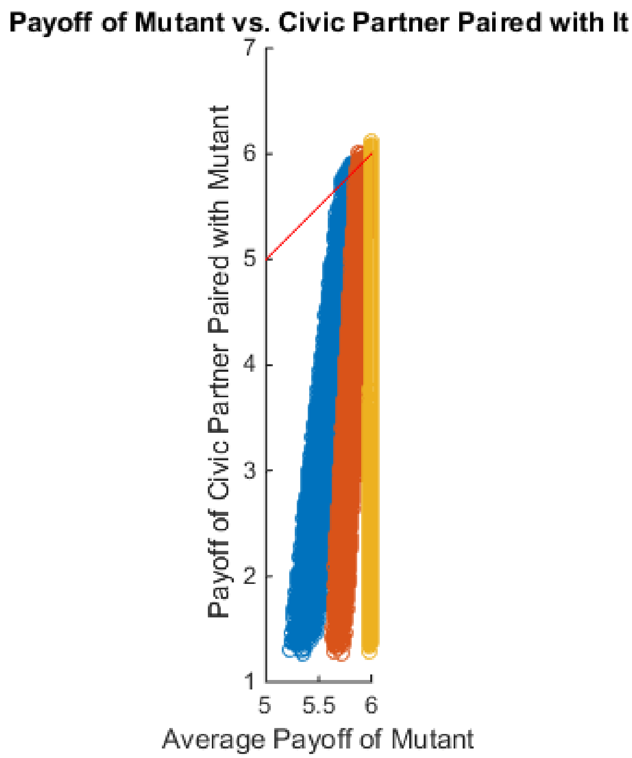

Figure 4.3: Each dot (of which there are 10,000 in each color) corresponds to the ordered pair {payoff of mutant, payoff of civic partner paired with mutant} for a randomly-generated mutant. The top chart shows varying λ values with a fixed population size $n = 20$: yellow has $\lambda = .7$, orange has $\lambda = .2$ and blue has $\lambda = 0$. As the civic partner gets more generous, the payoff of the civic partner paired with the mutant varies more and the payoff of the mutant varies much less. The bottom chart shows variation over population size with a fixed $\lambda = 0$. Recall that the civic partner definition varies based on n, the population size. Blue relates to population size $n = 10$, orange to $n = 20$ and yellow to $n = 500$. As the population size increases, the payoff of the mutant is constrained to vary even less and the payoff of the civic partner paired with it varies more.

cooperation by both partners, then A's partner also must get a lower payoff than it could get under full cooperation. In this case, the payoff for full cooperation would be 2 (payoff for full cooperation in PGG) + 4 (payoff for full cooperation under PD), which is 6. Note that both graphs in Fig. 4.3 show that when the civic partner's payoff is below 6, the mutant's payoff is also below 6, implying that the civic partner is a partner strategy. Additionally, in the limit as $\lambda \to 0$ and $n \to \infty$, the payoff of the mutant is held constant while the payoff of the mutant is allowed to vary. This fits Hilbe's definition of an equalizer strategy, one that holds its partner's payoff constant while its own payoff varies. Knowing that the civic partner satisfies these properties is very useful, because it means I could potentially draw on Hilbe's results to gain further insights into the properties of my civic partner family of strategies.

Weak Civic Partner Strategy (in the Unlinked Game)

As a control case, I ran a simulation where the weak civic partner (in the unlinked strategy space) was the resident and measured its ability to fend off invading mutants, again with a λ value moderating its generosity. The general formula for the unlinked weak civic partner looks like

$$\{1, 1-\lambda | 1, p_d * (1-\lambda)\}$$

because cooperation in the PGG and PD both depend solely on the partner's actions in the PD, so it can be described with 4 entries. Translated into the linked space, this can be written:

$$\{1, 1-\lambda | 1, p_d * (1-\lambda), 1, 1-\lambda\}$$

The key difference between the civic partner and the weak civic partner is that the former penalizes defection in the PGG in the PD, while the latter maintains unwavering cooperation with partners who defect in the PGG but cooperate in the PD. Additionally, the weak civic partner will continue to cooperate in the PGG when playing with this type of partner, effectively subsidizing their refusal to contribute to the public good.

In the previous analytical chapter, I showed that the weak civic partner is not a civic partner. In order to give a fuller picture of how it fails to compete, I ran some of the same experiments as done above, but with a weak civic partner substituted in. Figure 4.4 shows that over a wide range of population sizes and λ values, a weak civic partner still is not robust against invasion by many mutants. This can be seen by the fact that a large number of points are below the $y = x$ axis, indicating that the mutant had a higher payoff than the weak civic partner. Compare these charts to Figures 4.1 and 4.2, where the average payoff of the civic partner is always at or above the $y = x$ line, indicating a higher or at least equal payoff. This analysis shows that, even though the weak civic partner meets some of the criteria for being a civic partner, it failing to meet *all* of the criteria causes it to perform much worse against a wide variety of mutants.

4.2.2 Ability of Civic Partner to Defend Against Invading Mutants

The last section showed that in a population of mostly civic partners and one mutant, the civic partners had a higher average payoff for all mutants tested, while the weak civic partner family did not have this property. The next section builds on the last section by seeing how higher payoff translates into fitness and the likelihood that that strategy will succeed in taking over the entire population. Again, I slowly introduce more randomness with each investigation, here by adding in reproduction proportional to fitness.

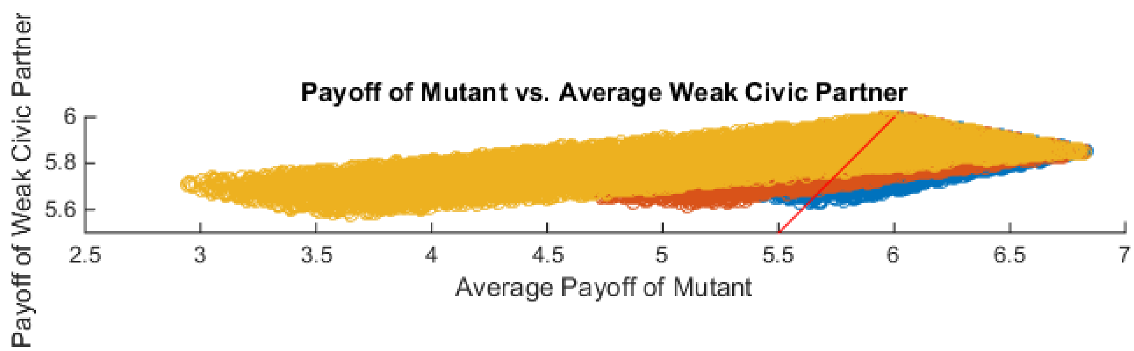

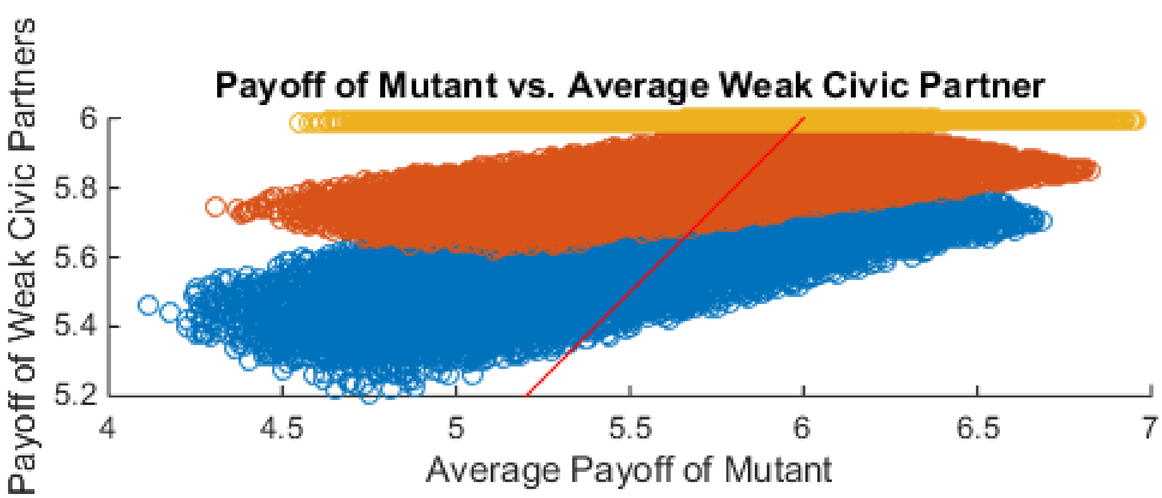

Figure 4.4: Each dot (of which there are 10,000 in each color) corresponds to the ordered pair {payoff of mutant, average payoff of weak civic partners } for a randomly-generated mutant. Above the red $y = x$ line, the weak civic partner paired with the mutant has higher payoff, and below the line, the mutant has a higher payoff. The top chart shows varying λ values with a fixed population size $n = 20$: yellow has $\lambda = .7$, orange has $\lambda = .2$ and blue has $\lambda = 0$. No matter how what the λ value, there are still mutants that do better than the average weak civic partner.The bottom chart shows variation over population size with a fixed $\lambda = 0$. Recall that the civic partner definition varies based on n, the population size. Blue relates to population size $n = 10$, orange to $n = 20$ and yellow to $n = 500$. Again, no matter what the population size, there are mutants that have higher payoff than the average weak civic partner payoff.

Linked

One trial is set up as follows: $n-1$ players play one strategy (the resident) and the last is a randomly-generated mutant, as before. However, this payoff is translated into fitness through an exponential function: fitness= $e^{\beta * \text{payoff}}$ where β determines the strength of selection: higher values means players with the highest payoff are even more likely to be selected. Reproduction is done proportional to fitness, and death is done at random. For this experiment, mutation rate is kept at 0. Whenever one strategy, the resident or mutant, has **"won"** by being the first strategy played by every player in the population, the trial was declared over. Alternatively, if the trial was still running after a pre-determined number of reproductive cycles, the trial was ended anyways and data saved on the proportion of each strategy present. A **reproductive cycle** is one round of gameplay and then birth/death. The maximum trial length was selected so almost all runs ended well before the maximum. In this field, one cycle of reproduction and birth is referred to as **"1/nth of a generation"** because one **generation** is commonly understood to be n reproductive cycles long, where n is the number of players. This formulation is useful because it allows for comparisons of trial length across different population sizes. However, I find discussing reproductive cycles useful as well, so I will mention both. Here, each trial is done for a maximum of $750/20 \approx 38$ generations or 750 reproductive cycles.

Higher payoff implies higher fitness levels, so a civic partner strategy should be able to resist invasion. Again, weak civic partners in the unlinked game are retained for comparison. Here, the added randomness comes from the process of death and reproduction. Because reproduction is done probabilistically, it is unlikely but possible that a player with a very low payoff will be chosen to reproduce.

Figures showing results are in Fig. 4.5. Note that $\lambda = 0$ is the most generous level and the upper bound on cooperation. It becomes immediately apparent that increased λ causes increased frequency of civic partner wins and decreased frequency of mutant wins, along with shorter average times until civic partner wins. However, this general pattern fails right around where $\lambda = 1$, or TFT. The culprit is decreased homogeneous cooperation between these TFT strategies. TFT cooperates with itself 100% of the time in the absence of errors, but as this computer program introduces a small error ($10^{-3}\%$), the long-term cooperation rates of TFT drop dramatically, causing its inability to maintain a high payoff.

For reference, I ran the same simulations but with other well-known strategies as the resident instead of civic partner. AllD wins 99.46±0.21% and AllC wins 39.80±1.372%. When AllD wins, it does so in 24.13±.80 reproducive cycles, AllC 22.53±1.51. Comparing the civic partner results to the AllD and AllC results again shows the power of the civic partner family. A population of civic partner when faced with a lone mutant is able to win as frequently as AllD and within approximately the same speed, though it is much more generous!

For future experiments, I used selected variants of civic partner, $\lambda = .2$ and $\lambda = .7$. The former was selected because it is a relatively low value of λ that achieves most of the benefits of higher λ. The latter was chosen as a comparison that is less generous, potentially with a greater ability to compete with other strategies like AllD.

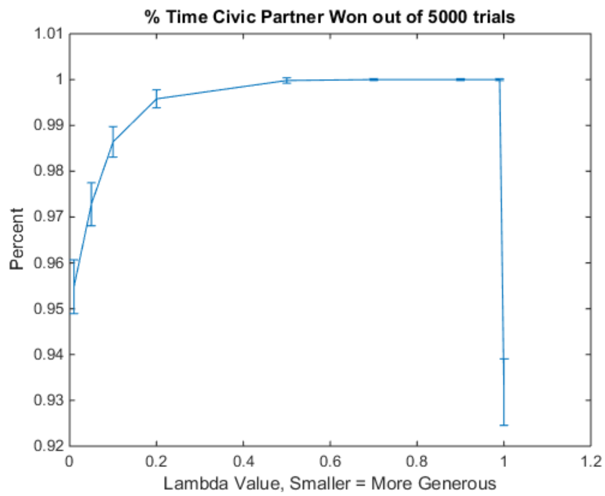

(a)

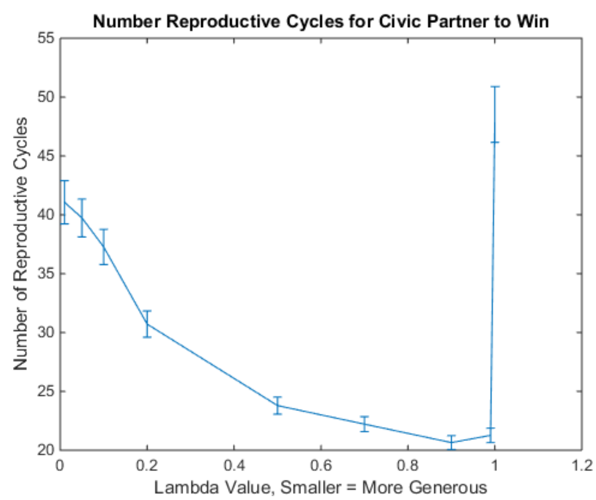

(b)

Figure 4.5: Charts showing how frequently civic partner won for varying levels of λ and the average time to civic partner win. Both plots have 95% confidence intervals. Done with beta value for birth equal to 1 and beta for death equal to 0 (random). Lower λ values (to the left of the x axis) are more generous and higher λ values (to the right) are less generous. A strategy "won" if it was the first strategy to be played by all the players in the population.

Unlinked

The simulation was run analogously to the linked case. For $\lambda = 0$, the weak civic partner won 87.14±0.94% of the time. For $\lambda = .2$ and .7, it won 91.66±0.78%, and 95.98±0.57% of the time respectively. First of all, the weak civic partner wins much more frequently than AllC, which is impressive given how generous the weak civic partner family is. However, a weak civic partner wins less frequently than a civic partner for identical λ values, showing again that the civic partner family is more successful in this type of situation.

4.2.3 Ability of Civic Partner to Invade Homogeneous Population

The last two sections explored properties of civic partner and weak civic partner strategies that were intimately related with the properties they have by definition: when the civic partner strategies are resident and faced with a single mutant, they should have a higher payoff on average than the mutant. This coming section marks a pivot in my analysis as I shift to exploring what other possible properties the civic partner strategy family might have. In some cases, these properties are far-removed from its original definition. The first example of a property I will explore is the reverse of what the civic partner was designed to do: when faced with a population made up of $n-1$ other strategies, how well is a civic partner able to *invade* this population? How does this compare with the ability of a weak civic partner to invade? How does the λ value impact any of this?

Linked

This section investigates whether two different civic partner strategies ($\lambda = .2,\ \lambda = .7$) were able to invade a population of either AllD or TFT. All of the parameters are kept constant from the previous trial. A civic partner ($\lambda = .2, \lambda = .7$) is relatively unable to invade AllD, succeeding 0.68±0.25% and 7.32±0.74% of the time, respectively. However, it is frequently able to invade TFT, succeeding 89.54±1.37% and 88.94±1.39% of the time, respectively.

These results accord with previous results found by Nowak on the Generous Tit-for-Tat strategy. For small numbers of invaders, GTFT is unable to invade AllD. However, given TFT's inability to maintain high levels of mutual cooperation in the presence of errors, it is not surprising that civic partner would be more successful at invading in a situation, as here, with a non-zero error rate.

Unlinked

Analogously, I ran simulations to see if the unlinked civic partner would be able to invade a population of either AllD or TFT. The results were very similar to that of the linked game for AllD: It was able to invade AllD $0.44 \pm 0.196\%$ and 7.02±0.79% of the time for $\lambda = .2$ and .7 respectively. However, it was able to invade TFT (spelled $[1, 0, 1, 0, 1, 0]$ in the unlinked situation) 72.55±1.96% and 72.89±1.96% of the time. As compared with the (linked) civic partner, the weak civic partner is less able to invade TFT, but similarly the choice between various λ values did not seem hugely important.

4.2.4 Multiple Strategies

This last trial is the most complicated one and puts the civic partner through its greatest test. It adds two new sources of randomness. The first is a greater diversity in the number of strategies allowed at any given time: 7. (Previously, only 2 different strategies were present at any given time). They are

1. AllD
2. TFT, in the linked game $[1, 0, |1, 0, 0, 0]$, in the unlinked game $[1, 0|1, 0] = [1, 0|1, 0, 1, 0]$.
3. Civic partner either linked or unlinked depending on the game, $\lambda = .2$.
4. Civic partner either linked or unlinked depending on the game, $\lambda = .7$.
5. The strategy that always cooperates in the PD but never in the PGG: $[0, 0|1, 1, 1, 1]$, referred to as **DC**.
6. The strategy that always cooperates in the PGG but never in the PD: $[1, 1|0, 0, 0, 0]$, referred to as **CD**.
7. AllC

These were chosen because they are all commonly-used strategies relevant to the study of the civic partner. AllD, DC and CD were the original mutant comparisons for deriving the upper probability limits for the civic partner calculation. TFT and AllC are both strategies that cooperate with other strategies of the same type 100% of the time (without errors), but are not as robust to other types of strategies to errors. Before the trial began, each of the $n = 20$ players was randomly selected to have one of the above strategies, with each selection independent and with the strategies equally likely to be chosen.

The second area of randomness comes during mutation rates, which were nonzero but small (2.5%). During death and reproduction, if mutation occurred, the dying strategy could only mutate to one of the 7 above strategies. This experiment was aiming to investigate whether there were differences in the ability of civic partner to dominate in the linked or unlinked game when playing a variety of opponents. The initial hypothesis was that the linked civic partner would do better because it would be able to distinguish between players that contributed in both the PGG and PD and players that only cooperated in the PD. DC is a prime example of a "spoiler" strategy the unlinked weak civic partner would be unable to distinguish from AllC and might cooperate excessively with DC.

Linked

Before beginning trials, I calculated the payoff each strategy had while playing each other, shown in Table 4.1. Note that payoff from the PGG is dependant on the entire population's actions, not just the pair's, so this is excluded. Assuming each strategy encounters each other strategy with equal frequency, averaging along the rows provides an estimate of each player's average payoff (again, excluding benefit from PGG). This is shown in Table 4.2. Players tend to encounter each other with the same frequency around the beginning of the game, before reproduction and death has begun to alter the frequency with which strategies appear in the population. The civic partners $\lambda = .7$ does better on average than the AllD player, and the $\lambda = 0.2$ does not much worse than AllD.

	AllD	TFT	λ =0.2	λ=0 .7	DC	CD	AllC
AllD	0	0	2.22	0.84	5	0	5
TFT	0	0	3	3	4	0	3
λ= 0.2	-1.24	3	3	3	3.36	-1.4	3
λ=0.7	-0.47	3	3	3	3.76	-0.53	3
DC	-1	-1	2.21	0.2	4	-1	4
CD	-1	-1	2.01	0.14	4	-1	4
AllC	-2	3	3	3	3	-2	3

Table 4.1: Payoff for strategies in the linked case playing each other. Each entry represents the payoff the player along the rows gets while playing the strategy in that column. For example, DC gets a payoff of 2.21 when playing civic partner $\lambda = .2$.

	Avg Payoff
AllD	1.87
TFT	1.86
$\lambda = 0.2$	1.82
$\lambda = 0.7$	2.11
DC	1.06
CD	1.02
AllC	1.57

Table 4.2: Average payoff to each strategy in the linked game, excluding benefit from PGG and assuming each strategy interacts with each other with equal frequency.

Next, I ran simulations. Because of the higher mutation rate, I increased the maximum number of reproductive cycles to 1500 (75 generations), but again ended a trial when one strategy had taken over the entire population. There were 5000 independent trials done in total. Table 4.3 shows these results. Note that AllD is the most frequent winner, but civic partner (both varieties) do much better than any other strategy, with the less-generous one doing better. This seems to fit with what I was expecting. The strategies that have the highest average payoff tend to win the most frequently. The one exception is TFT and this is likely because it has a very low payoff with other TFT players, so whenever TFT becomes common, it is easy for another strategy like AllD to invade.

Unlinked

As a comparison, I ran the exact same set of experiments with the unlinked game. Table 4.4 shows payoff levels for strategies in the unlinked space, and Table 4.5 shows their average payoff. Note that DC has a much higher average payoff, coming mainly from its increased payoff when paired

	AllD	TFT	$\lambda = 0.2$	$\lambda = 0.7$	DC	CD	AllC
Percent time wins	46%	2%	16%	25%	3%	6%	1%
sd above	0.7%	0.2%	0.5%	0.6%	0.2%	0.4%	0.2%
mean length to win	149	230	301	276	261	181	392
sd above	1.6	2.0	2.7	2.7	2.9	2.0	3.1

Table 4.3: Results from evolutionary trial for linked game. Length is given in reproductive cycles.

	AllD	TFT	$\lambda = 0.2$	$\lambda = 0.7$	DC	CD	AllC
AllD	0	0	2.24	0.83	5	0	5
TFT	0	1.5	3	3	3	0	3
$\lambda = 0.7$	-1.25	3	3	3	3	-1.25	3
$\lambda = 0.7$	-0.47	3	3	3	3	-0.47	3
DC	-1	4	4	4	4	-1	4
CD	-1	-1	1.24	-0.17	4	-1	4
AllC	-2	3	3	3	3	-2	3

Table 4.4: Relative payoffs (excluding PGG benefit) to players in the unlinked game. Note how, as compared with the linked TFT, the unlinked TFT is more generous and therefore has a higher payoff when paired with itself.

	Avg Payoff
AllD	1.87
TFT	1.93
$\lambda = 0.2$	1.79
$\lambda = 0.7$	2.01
DC	2.57
CD	0.87
AllC	1.57

Table 4.5: Average payoff for each strategy (excluding payoff from PGG), assuming they meet with equal frequency.

with with TFT and the weak civic partner strategies. Table 4.6 confirms my initial hypothesis: in the linked game, the weak civic partner strategies do much worse, and the DC strategy does much better! Because the weak civic partner strategies in the unlinked game are unable to punish DC for its anti-social defection in the PGG, DC is able to exploit them and dominate. This portion of the experiment shows that linking the games enables the civic partner strategies to do much better and DC to do much worse, as desired!

4.2.5 Long-term Cycles of Multiple Strategies: Frequency and Cooperation levels

The section studied when certain strategies "won" a simulation by occupying all players in the population. However, with a non-zero mutation rate the simulation doesn't need to end at this point. In fact, a more interesting area of analysis can come from studying a population over many cycles of reproduction and death, on the order of 100,000, to see if certain patterns emerge. These trials were run for 5000 generations or 100,000 reproductive events. This number was chosen by selecting the length that seems to show several patterns of strategies dominating and being invaded,

	AllD	TFT	$\lambda = .2$	$\lambda = .7$	DC	CD	AllC
Percent of time strategy won	65%	2%	3%	5%	18%	7%	0.%
sd above	0.7%	0.2%	0.3%	0.3%	0.6%	0.4%	0.1%
mean length to win	151	239	268	248	184	176	305
sd above	1.3	1.3	2.5	2.4	1.8	1.4	3.7

Table 4.6: Percent of the time each strategy dominated and the length of time it took each strategy to win in the unlinked game.

	AllD	**TFT**	$\lambda = .2$	$\lambda = .7$	**DC**	**CD**	**AllC**
mean frequency	29%	3%	22%	40%	2%	2%	3%
sd	2.2%	0.1%	1.7%	2.6%	0.3%	0.2%	0.1%

Table 4.7: The frequency of the seven named strategies in a 5000-generation trial (100,000 reproductive events with 20 players in the population).

which enables me to show how the strategies interact with each other.

Linked

Results for the linked trial are shown in Table 4.7. The standard deviations are fairly large, which reflects the high underlying level of variation in the 10 trials that were done to produce this table. However, some high-level interpretations can still be made, especially in comparison with Table 4.3, which displays how frequently each strategy was the first one to win (have all players in the population be using that strategy). This comparison seems to show that AllD wins more frequently, but that the the less-generous civic partner ($\lambda = .7$) actually is present more frequently in the long term. The more generous civic partner strategy, $\lambda.2$ doesn't do much worse than the AllD player in terms of frequency in the long-term. Notably, all other strategies do much worse than those three main leaders.

The same strategies can behave in very different ways, depending on what other strategies they are paired with. For this reason, it's useful to examine the mean levels of cooperation in the PD and PGG. I calculated this by keeping track of the average cooperation rate for each player in each pair (calculated using the matrix method explained previously) and averaging over the entire population present at one time. If you average over all reproductive cycles for all 10 sample trials, the mean level of cooperation in the PD was 0.67 with a standard deviation of 0.025. For PGG, its 0.67 with a standard deviation of 0.0232. These numbers are all between 0 and 1: 0 would represent no player cooperating ever and 1 would represent all players always cooperating.

Additionally, I'd like to explore the dynamics of strategies. One key question we'd like to be able to answer is: "If strategy x is currently being used by many players in the population, how did this come about, and how will it stop?" To try and answer that question, I studied, for each reproductive cycle, what the strategy was that occupied the majority of the players, and what was the last strategy to occupy the majority of strategies (be **dominant**). I also calculated how many separate instances a strategy invaded (became dominant) and for how many total reproductive cycles that strategy was dominant. I calculated these values for each of the 10 trials of 5000 generations (100,000 reproductive events).

Table 4.8 contains information about how frequently various strategies invaded. Interestingly, AllD invaded relatively rarely, but each period of dominance lasted for a long time. The two civic partner strategies had the highest number of invasions, and while shorter than AllD's average, they were longer than any other strategies. Table 4.9 shows how frequently each strategy was the last dominant strategy before the uprise of a new strategy. The two civic partner strategies invade each other very frequently. TFT occasionally invades and is invaded by the civic partner strategies. When AllC invades, it does so almost always after one of the civic partner strategies.

As a visual demonstration, Figure 4.6 contains close-up views of brief interactions between

	AllD	TFT	$\lambda = .2$	$\lambda = .7$	DC	CD	AllC
Dominance Frequency	47	75	177	207	16	16	94
Total Dominance Length	29951	910	22484	41563	1520	1046	1890
Average Dominance Length	644	12	132	207	90	66	20

Table 4.8: Information on how frequently a strategy was dominant (occupied more players than any other strategy) for the linked trials.

	AllD	TFT	$\lambda = .2$	$\lambda = .7$	DC	CD	AllC
AllD	0 %	27	17	20	9	11	16
TFT	22	0	24	42	4	2	6
$\lambda = .2$	5	11	0	65	1	2	17
$\lambda = .7$	8	17	54	0	2	1	19
DC	9	12	20	17	0	10	32
CD	21	6	18	7	11	0	38
AllC	2	6	36	50	3	4	0

Table 4.9: What percentage of the time each strategy was the last dominant strategy before another one became dominant, averaged over 10 trials (linked). For example, the 27 in the upper left means that for an average of 27% of the time, TFT was the last dominant strategy before AllD became dominant. The % sign is dropped everywhere for neatness except the upper left entry to remind the reader that the entries are percentages. Rows that correspond to strategies with fewer instances of dominance, such as AllD, DC and CD, have greater standard deviations and so are less likely to provide useful information.

the seven strategies.

Unlinked

As done before, this same experiment was repeated in the unlinked strategy space for comparison. Table 4.10 shows the frequency of each of the seven strategies. The first useful comparison is to Table 4.7, which contains the same values but for the linked game. Relative to the linked results, in the unlinked ones AllD is much more prevalent, and the civic partner strategies are much less prevalent. Additionally, the strategy DC is much more common. This makes sense: the unlinked weak civic partners don't punish defection in the PGG, so they cooperate with this strategy just as frequently as they would with AllC. Because they base their cooperation choices in the PGG off the behavior of their partner in the PD, a weak civic partner paired with a DC will cooperate in the PGG with probability 1. However, DC will not cooperate in the PGG, so DC will be allowed to free-ride without punishment. One potential reason why DC isn't more common is that it does badly against AllD, which is a very common strategy in the unlinked situation. Finally, the standard deviations in the unlinked table are lower, perhaps indicating that unlinked trials tend to follow more similar patterns.

In terms of cooperation level, for the PD the mean cooperation rate was 0.35 with standard deviation 0.0081. The PGG mean cooperation rate was 0.21 with standard deviation 0.0082. The unlinked games all have much lower levels of cooperation than the linked game, as would be expected by its much higher frequency of AllD. The PGG cooperation is lower than PD, which makes sense given the prevelance of DC.

Again, I calculated how frequently each strategy was dominant. Table 4.11 shows frequency

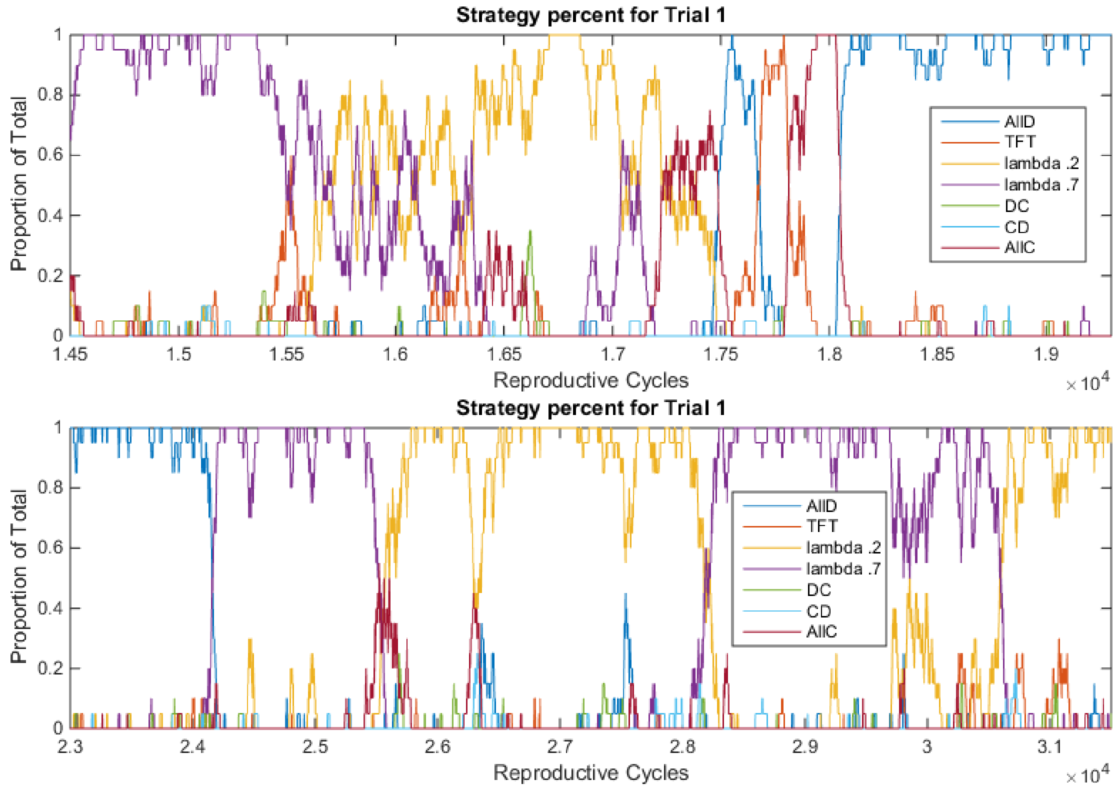

Figure 4.6: Examples from a linked trial showing the interactions of the seven strategies. Here, the two civic partner strategies frequently invade each other, which makes sense because all civic partners cooperate with other civic partners with probability 1. The top figure shows that AllC is able to invade a civic partner, likely because it also has a high cooperation rate. However, AllC is not retaliatory (able to punish players that defect against it). This could be the reason why AllD quickly invades it. The bottom chart shows how the two civic partner strategies can trade off dominance for extended periods of time.

	AllD	**TFT**	$\lambda = .2$	$\lambda = .7$	**DC**	**CD**	**AllC**
mean frequency	59%	5%	5%	10.%	17%	2%	2%
sd	0.8%	0.3%	0.7%	0.7%	0.7%	0.2%	0.1%

Table 4.10: The frequency of the seven named strategies in a 5000-generation trial (100,000 reproductive events with 20 players in the population).

	AllD	TFT	$\lambda = .2$	$\lambda = .7$	DC	CD	AllC
Dominance Frequency	104	117	61	102	95	23	37
Total Dominance Length	61261	4148	4428	9827	17191	1694	910
Average Dominance Length	600.	36	71	100	183	78	26

Table 4.11: Information on how frequently a strategy was dominant (occupied more players than any other strategy) for the unlinked trials.

	AllD	TFT	$\lambda = .2$	$\lambda = .7$	DC	CD	AllC
AllD	0%	21	4	9	56	7	4
TFT	39	0	18	27	3	2	11
$\lambda = .2$	10	46	0	31	8	1	5
$\lambda = .7$	22	37	18	0	12	2	10
DC	25	12	15	33	0	11	4
CD	17	4	0	3	63	0	12
AllC	3	47	10	31	5	4	0

Table 4.12: How frequently each strategy was the last dominant strategy before another one became dominant, averaged over 10 trials (unlinked). For example, the 22 in the upper left means that out of all the times AllD became dominant, 21% of the time TFT was the last dominant strategy. The % sign is dropped everywhere for neatness except the upper left entry to remind the reader that the entries are percentages. Rows that correspond to strategies with fewer instances of dominance, such as CD and AllC, have greater errors and so are less likely to provide useful information.

information. When compared to the linked game in Figure 4.8, I see that AllD again has the largest average dominance length, but has a much higher invasion frequency count. DC has both more invasions and longer ones, while the weak civic partners have fewer invasions and shorter ones. Table 4.12 shows how frequently each strategy rose to dominance after the dominance of another strategy. AllD invades very frequently from DC and TFT invades very frequently from AllD and the weak civic partner strategies. DC invades frequently from AllD, the weak civic partner strategies and TFT. These results indicate that DC's increased strength (as compared to the linked case) comes from the weakness of TFT and the weak civic partner strategies. AllD's increased strength comes from the weakness of DC as compared with AllD.

Figure 4.7 shows example snap-shots from interactions of the seven cycles in the unlinked simulations.

4.3 Conclusion of Computer-Aided Analysis

This chapter's goal was to build upon the analytically-derived definition of the civic partner family and investigate other properties it might have. The strength of using a computer-based method is that it allows me to explore situations that would have taken infeasibly long if done by hand, such as calculating the payoff of a civic partner against 10,000 randomly-generated strategies. The weakness of using a computer-based method is that I can never actually "prove" a result as I would in an analytical setting, only show that a property seems to be true when tried under many conditions.

Nevertheless, this chapter gave me some powerful results. I found that a civic partner strategy seems to be robust against invasion for every mutant strategy that was tested. Less-generous civic

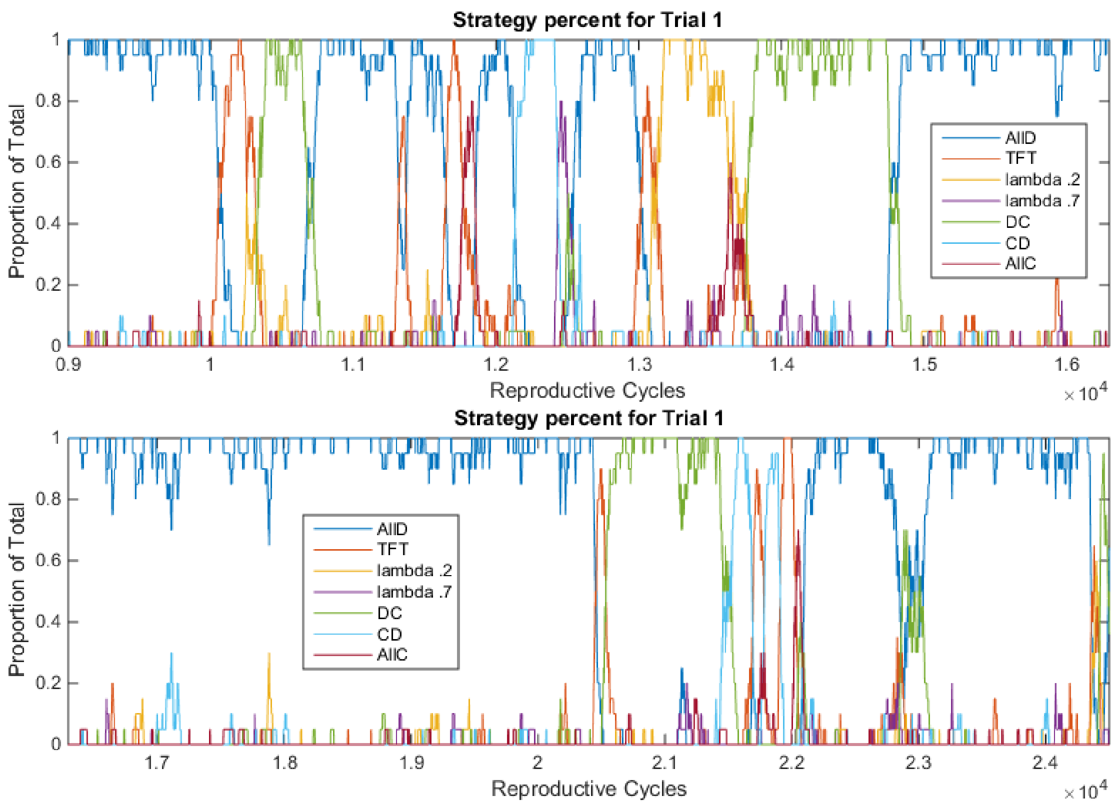

Figure 4.7: Examples from an unlinked trial showing the interactions of the seven strategies. The top chart shows how AllD is invaded by TFT, which is invaded by DC, subsequently invaded by AllD again. A weak civic partner at one point becomes dominant, but it is eventually invaded by DC. The bottom chart shows how after TFT invades AllD, DC quickly invades TFT again, allowing a flurry of invasions until AllD rises to dominance again.

partners were even more robust against invasion because they had an even higher average payoff. This higher payoff translated into an increased likelihood of being able to repel invading mutants when death and reproduction were added to the experiment. By contrast, weak civic partners had a lower average payoff than many mutants. This translated into a reduced ability to defend against invading mutants when death and reproduction was added.

I found that the civic partner happened to have additional beneficial properties that it was not specifically designed to have. Firstly, it was more able to invade a homogeneous population of TFT than a weak civic partner was. Secondly, it performed better when put in a mixture of seven named strategies. As compared with a weak civic partner, a civic partner was more frequently able to win (be the first strategy to occupy all players). When the simulation was run for longer and overall frequency was measured, the civic partner strategies are present more frequently in their population than the weak civic partners are present in their unlinked population. In addition, the weak civic partners allow free riders like DC to appear very frequently and gain high payoff, while the civic partners do not allow DC to do as well.

These results can be made more intuitive by anthropomorphizing the strategies. For example, a civic partner strategy can be considered to be a person who is kind to their neighbors and contributes to the common good. However, they are less kind to those who to not contribute to the common good. DC is a person who is kind to their neighbors, but does not contribute to the common good. In the linked situation, a civic partner punishes the free-rider by being less generous in daily person-to-person interactions. Because people can change their strategies, free-riders become less common and civic partners become more common. In the unlinked game, though, no one knows how much their neighbors are contributing to the common good. This is why the weak civic partner is unable to punish free-riders like DC in person-to-person interactions. Because contribution to the public good is costly (such as taxes or helping maintain town greens), free-riders do better and become more frequent.

Given that situation seen more commonly in human interactions is a low-prevalence of free-riders, it could make sense that human society is better reflected by the linked game. This would imply that humans generally consider contributions to the public good when they interact with other humans, punishing those that are free-riders.

Chapter 5

Conclusion

This section will summarize the results in the previous analytical and computational chapters. Additionally, it will discuss how my findings relate to the original motivation for this project and describe future areas of research, including extensions I have already begun work on.

5.1 Review of Motivation

I began my thesis by discussing the case of Martin Shkreli, whose company's business model was to dramatically raise the prices on generic drugs and make a fortune without advancing medical knowledge or improving patient welfare. This case is interesting partially because of how rare it is: most people do not choose such anti-social free-rider methods to make their living. In fact, psychological studies of human participants playing classical games such as the Public Goods Game showed that humans usually cooperate more frequently than would make sense mathematically. For the PGG, the Nash Equilibrium is universal defection, meaning that each player optimizing their own payoff should choose to defect. However, controlled experiments have shown the lowest cooperation ever dips is about 30%.

Other studies suggested a possible explanation for why humans participate more in the PGG. Humans seem to cooperate much more frequently in the PGG when it is paired with the Prisoners Dilemma. This could perhaps be because this linkage makes knowledge of a player's cooperation status public, enabling detection of free-riders. However, is it a "smart" strategy to punish this free-riders? Investigating this question is at the core of my thesis.

5.2 Analysis of Analytical and Computational Findings

Because human psychology is difficult to analyze computationally, I abstracted a person's behavior into a 6-dimensional vector indicating a player's willingness to cooperate under a variety of circumstances. This abstraction is a very standard way of approaching game-theoretical problems. I reviewed the history of analyzing repeated games, such as Axelrod's tournaments, and described a few famous strategies like Tit-for-Tat and Generous Tit-for-Tat. Additionally, I described some properties various strategies can have, such as being robust to invasion.

Having established a solid foundation, I next began with the novel portion of my thesis. I defined a type of strategy I aimed to investigate, one that with probability 1 continued mutual cooperation, was robust to invasion by AllD, DC and CD, and always had a greater than 0 probability of cooperation. I named this family of strategies "civic partners" and wished to investigate whether any strategies that meet this criteria actually exist. In particular, I wanted to see if they exist in 1) a linked game, where players know their partners PGG cooperation status and may change their likelihood to cooperate in the PD based on this information and 2) the unlinked game, where in the PD players know only their partner's actions in the PD. I found that a family of civic partners did exist in the linked game, and I completely described the subset of the strategy space they occupy. In the unlinked space, I proved that no civic partners exists. However, a family of weak civic partners, which fulfilled all the properties of the civic partner except robustness against AllD and DC, does exist. I retained this family in future sections as a reference to the civic partners in the linked strategy space.

Following the analytical portion, I set out to further investigate the properties of the civic partners as compared to the weak civic partners. I showed via simulation that civic partners are robust to invasion not only by AllD, CD and DC, but also by any strategy in the continuous 6-dimensional space of strategies under study. Less-generous civic partners had substantially higher payoff than a lone mutant, but even the most generous civic partner never did worse than a mutant. I found that the civic partner paired with the lone mutant is able to force a relationship between its payoff and the mutant, ensuring that the mutant can never get a higher payoff than mutual cooperation unless the civic partner also got a higher payoff than the bound. The strategy space of a civic partner depends on the population size, and as the population increases, civic partners emerge that are more generous and are able to maintain a constant payoff no matter what invading mutant appears. Using terminology from previously-published research papers, I showed that the civic partner is a "partner" strategy, and in the limit of increasing generosity and increasing population, it is also an "equalizer" strategy. The weak civic partner (in the unlinked game) does not have any of these properties and frequently allows invading mutants to have higher payoffs than the average weak civic partner payoff. When reproduction is added, the civic partner is very frequently able to eliminate the mutant, and less-generous civic partners are able to do this more frequently and more quickly. By contrast, the weak civic partner eliminates a lone mutant less frequently.

Finally, I ran simulations with up to 7 named strategies present in the population at a time: AllD, TFT, two variations of civic partners, DC, CD, and AllC. As compared to an identical set-up with weak civic partners in the unlinked space, the linked civic partner wins (is the first strategy to be employed by the entire population) more frequently. Conversely, in the unlinked game, the weak civic partner wins less frequently against AllD. DC, the free-rider strategy, wins more frequently. When this set-up is run for longer and with mutation, the linked game shows that the civic partners are more prevalent and the AllD and DC players less prevalent than in the unlinked game.

5.3 Implications

The analytical and computational results support my initial hypothesis. The strategy space of a linked game includes a family of strategies that have several desirable properties: high cooperation probabilities, yet robust to invasion and much more able to do well in a mixed population of other named strategies. By contrast, the unlinked game contains no such civic partner, and the next best thing, the weak civic partner, which it does have, lacks many desirable properties.

These results suggest that linking the PGG and PD may be instrumental in ensuring that free-riders aren't able to proliferate in a population. An analogy to a human population might be a mindset that maintains high cooperation and is generous with those who defect, always maintaining the possibility of future cooperation, and yet is not so generous that it invites exploitation by anti-social free-riders. Showing that this strategy is able to do well in a mixed population is analogous to showing that people with this mindset are able to survive and prosper in a community where other people have different approaches to cooperation, even where free-riders appear in large numbers.

If you replace "people" with "communities", my results could be seen as describing how communities with high levels of generosity towards other populations, but not so high as to be taken advantage of, have managed to thrive while other types of communities have failed to take advantages of opportunities for mutual cooperation. If many participants that have a firm but forgiving view towards other interactions with other participants, this could deter other members from free-riding (sharing in the spoils of cooperation without contributing itself). A very rough example could be seen in international trade. Here, the PGG is international sanctions on a rouge nation. For any individual country, violating the sanction would bring the benefits of near-monopolistic trade with the sanctioned nation. However, it is best for the entire community (we assume) if sanctions are maintained. Any country that violates the sanctions can be punished by revoking nation-to-nation trade agreements with it (the PD). The analysis in my thesis could suggest that because the PGG and PD are linked where trade patterns are known, nations tend to follow civic partner strategies. Analyzing exactly how applicable my mathematical simulations are to the messy, unpredictable world of international relations could be a thesis of its own, but introducing this simple example gives a sense of why the questions addressed in my thesis are relevant to a non-mathematical situation.

5.4 Areas of Future Study

My results certainly leave open options for further research in psychology and neuroscience, as well as community-building and international relations. However, being a mathematician, I will focus on areas of mathematical advance. I plan to continue work on this project post-thesis and am already investigating the next step of the analysis.

This portion would allow a population to be initially seeded with strategies from anywhere within the six-dimensional strategy space, not just the 7 named strategies discussed previously. The options are technically infinite, but reproduction according to fitness should enable a civic-partner-like strategy to naturally emerge from the chaos, further bolstering the civic partners strength and relevancy to evolutionary problems. Because the strategy space is so much larger, more advanced analysis is needed to ensure the results are clear and robust, but initial results are promising. I look forward to better understanding the complicated, fascinating topic of multi-player, multi-game scenarios.

Chapter 6

Glossery

Definition 3. *Games*

- **PGG**, *the Public Goods Game, is a multi-player game where each player has the opportunity to contribute a cost* c_g *or not to a central pot of money, which is then divided by some* $r > 1$ *and divided evenly among all players, regardless of if they contributed or not.*

- **PD**, *the Prisoner's Dilemma, is a two-person game. Each player has the choice of paying a cost* c_d *to give its partner a benefit b. When the same two players repeatedly play each other, this is called the Iterated Prisoner's Dilemma, or* **IPD**.

Definition 4. *Game-Specific Notation.*

- **r** *is the multiplication factor in the PGG game. It must be* ≥ 1.

- **b** *is the benefit the receiving player gets in the PD. It must be* ≥ 0 *and can be assumed to be* > 0 *unless otherwise noted.*

- **N** *is the population size. It must be even so there are* $N/2$ *pairs.*

- $\mathbf{c_g}$ *is the cost of cooperation in the PGG game. It must be* ≥ 0 *and can be assumed to be* > 0 *unless otherwise noted.*

- $\mathbf{c_d}$ *is the cost of cooperation in the PD game. It must be* ≥ 0 *and can be assumed to be* > 0 *unless otherwise noted.*

Definition 5. *The* **payoff** *of a player for a given round is its benefit for a round minus its cost paid out. Its* **fitness** *is a function of payoff that determines how likely it is to survive and pass on its strategy. The fitness strategy can vary, but it must preserve order: if player A has a greater payoff than player B, it must also have a higher fitness.*

Definition 6. *A* **linked game** *is one where the cooperation status of your partner is known to you in the PD. A point in the* **linked strategy space** *is a strategy that may have different cooperation probabilities in the PD based on their partner's cooperation decision in the PGG. The* **unlinked game** *is where cooperation in the PGG depends on cooperation in the PD and cooperation in the PD depends only on cooperation in the PD. A point in the* **unlinked strategy space** *is a strategy that does not differentiate between partner's that cooperated in the PGG and those that did not in PD cooperation rates.*

Definition 7. *A* **reproductive cycle** *is a period that includes calculating payoff for all players, calculating fitness and then 1 reproductive event and 1 death. This is also referred to as a* $1/nth$ *generation. A* **generation** *is when* n *reproductive events have occurred, where* n *is the number of players in the population.*

Definition 8. **AllD** *is a strategy that always defects and* **AllC** *is a strategy that always cooperates.* **Tit for Tat** *is a strategy that does back to its partner what its partner did to it last round. A* **Generous Tit for Tat** *is a Tit for Tat strategy that cooperates with probability* p *even when its partner defected in the last round.*

Definition 9. *From Hilbe et al's paper: A strategy is a* **partner** *strategy if it is never the first to defect, and additionally has the property that if its payoff is less than the payoff for mutual cooperation, the payoff to its partner is also less than the payoff for mutual cooperation. An* **equalizer** *strategy, also from Hilbe et al, is one that forces its partner to have a constant payoff. A strategy is* **evolutionarily stable** *if it has an average higher payoff than a rare mutant. A strategy is a with the property that no player can unilaterally change its strategy and get a higher payoff.*

Definition 10. *When a player is selected for* **death***, its strategy is removed and replaced by either the reproducing player's strategy or a randomly-selected strategy through mutation. At the same time that a player dies, another is selected (proportional to payoff) to* **reproduce** *and its strategy is copied into the strategy of the player who dies. When* **mutation** *occurs, the dying players strategy is not replaced with the reproducing one, but instead a randomly-selected mutant. Mutation occurs with some probability called the mutation rate.*

Bibliography

[1] Christian Hilbe, Arne Traulsen, and Karl Sigmund. Partners or rivals? strategies for the iterated prisoner's dilemma. *Games and Economic Behavior*, 92:41–52, 2015.

[2] Lorens A. Imhof and Martin A. Nowak. Stochastic evolutionary dynamics of direct reciprocity. *Proceedings of the Royal Society*, 277:463–468, 201.

[3] Karl Sigmund. *The Calculus of Selfishness*. Princeton University Press, 2010.

[4] Manfred Milinski, Dirk Semmenn and Hans-Jurgen Krambeck. Reputation helps solve the "tragedy of the commons". *Nature*, 415:424–426, 2002.

[5] Martin A. Nowak. *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard University Press, 2006.

[6] Beth Mole. The FDA quietly changed the rules to block the next Martin Shkreli, 2016.

[7] Martin A. Nowak and Karl Sigmund. Tit for tat in heterogenous populations. *Nature*, 355:250–254, 1992.

[8] David G. Rand Oliver P. Hauser, Achim Hendriks and Martin Nowak. Preserving the global commons. 2016.

[9] Karthik Panchanathan and Robert Boyd. Indirect reciprocity can stabilize cooperation without the second-order free-rider problem. *Nature*, 432:499–503, 2004.

[10] Robert Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.

[11] C.H. Taubes. Lecture notes on probability, statistics and linear algebra. 2016.

[12] University of Groningen. Justice Oliver Wendell Holmes, 2012.