

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank) 2. REPORT DATE 1990 3. REPORT TYPE AND DATES COVERED ~~XXXXX~~/Dissertation

4. TITLE AND SUBTITLE Essays on Cooperation and Competition 5. FUNDING NUMBERS

6. AUTHOR(S) Bruce George Linster AD-A224 480

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AFIT Student at: University of Michigan 8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/CI/CIA -90-018D

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFIT/CI Wright-Patterson AFB OH 45433 10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for Public Release IAW AFR 190-1 Distribution Unlimited ERNEST A. HAYGOOD, 1st Lt, USAF Executive Officer, Civilian Institution Programs 12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

DTIC ELECTE AUGO 1 1990 S B D

14. SUBJECT TERMS 15. NUMBER OF PAGES 156

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED 18. SECURITY CLASSIFICATION OF THIS PAGE 19. SECURITY CLASSIFICATION OF ABSTRACT 20. LIMITATION OF ABSTRACT

ABSTRACT

ESSAYS ON COOPERATION AND COMPETITION

by
Bruce George Linster

Co-Chairs: Theodore C. Bergstrom, Kenneth G. Binmore

Understanding the basic concepts of cooperation and competition is fundamental to understanding economic and social behavior. These essays explore two somewhat different areas in which cooperation and competition play a role.

This dissertation explores how cooperative behavior evolves and is sustained in situations which can be modeled with the Prisoners' Dilemma. This is accomplished through a replication of Robert Axelrod's famous Prisoners' Dilemma tournament with the payoffs calculated to take the infinite nature of the game into account and computer simulations which analyze the stability these results in the presence of mutation. We can then see what characteristics the successful strategies have in various situations.

The rent-seeking games originally modeled by Gordon Tullock are then investigated. Two modifications to the existing literature are explored. First, these games are modified to be played sequentially. Then, the players' valuations for the prize in these games are modified to be vectors. This allows players to have different preferences over who wins the prize. The results of this study indicate total rent-seeking expenditure depends on which player goes first and their relative valuations. This work also explains why some players may choose not to participate in these contests. If prizes are public goods, it is shown that if some players share interests with some of the others, the more they have in common with those players the less likely any of them are to win. The results here have applications in political, international, and military competition.

90 07 31 078

ESSAYS ON COOPERATION AND COMPETITION

by
Bruce George Linster

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
1990

Doctoral Committee:

Professor Theodore C. Bergstrom, Co-Chair
Professor Kenneth G. Binmore, Co-Chair
Assistant Professor Elazar Berkovitch
Professor Carl P. Simon



Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

© Bruce George Linster 1990
All Rights Reserved

For Debbie
my partner in all endeavors

ACKNOWLEDGEMENTS

My sincere thanks goes to my Dissertation Committee, especially Ken Binmore and Ted Bergstrom. Their guidance and recommendations helped make what seemed an impossible task a reality. In addition to helping me as a committee member, Carl Simon has been a valued friend and advisor, and I am deeply indebted to him. I also owe a great deal to John Miller, Hal Varian, and Larry Samuelson for sharing their immense insight with me. Finally, without the help of Mary Braun I never would have made it through what seems to be the most complex of bureaucratic mazes.

I offer my love and deepest thanks to Debbie, Greg, and Andrea. Without them to put things in perspective and keep me going when things got tough, this would have been one of those things I'd wished I'd done.

PREFACE

This dissertation began as a number of individually motivated studies into two different areas. The first area of inquiry was how cooperation evolves and is sustained in the infinitely repeated Prisoners' Dilemma, and the other was how the results of game theoretic rent-seeking models are altered if we allow players to act sequentially or have preferences which vary over who wins the prize.

Any study of cooperation must at some time touch on the works of Robert Axelrod. His provocative 1984 book *The Evolution of Cooperation* is the seed which grew into many academic papers, and it still evokes controversy over his methodology and results. I begin this collection of essays by replicating part of his work to capture the effect of the infinite nature of the game he described. The results of the simulation reveal there is indeed a difference in how the strategies do when the game's infinite time horizon is taken into account.

In the second essay I consider the evolution of cooperation in the presence of trembles or perturbations. I find that how the trembles are applied affects the evolutionary results significantly. I explore the effects of repeated perturbations to the evolutionary process in the form of mutation. I then study evolutionary stability in the infinitely repeated Prisoners' Dilemma in an environment where the strategies must be implemented by two state Moore machines or finite automata. Again, I explore the effects of mutation and find TIT-FOR-TAT, clearly the best strategy in Axelrod's tournament, lacks qualities which are necessary to be evolutionarily successful under certain circumstances.

In the last essay I examine models originally studied by Gordon Tullock to analyze political competition. I alter the model in two ways. First, I analyze the model as a sequential game. I allow someone to go first in a game which is similar to a lottery. Next, I allow for a degree of publicness in the prize. These models have applications in the study of international or military competition, as well as in analyzing what are generally referred to as rent-seeking expenditures.

The total effect of these four essays, I hope, is to provide some insight into cooperation and competition. Although these words are part of our everyday vocabulary, they are interesting social phenomena which can be studied and at least partially understood using game theory.

TABLE OF CONTENTS

DEDICATION ii
ACKNOWLEDGEMENTS iii
PREFACE iv
LIST OF FIGURES viii
LIST OF TABLES x

CHAPTER

I. AUTOMATA PLAY THE PRISONERS' DILEMMA: ANOTHER LOOK AT THE EVOLUTION OF COOPERATION 1

Introduction

The Prisoners' Dilemma

The Evolution of Cooperation

Nachbar's Criticism

Evolution in an Approximately Infinite RPD

Summary

Bibliography

II. THE EVOLUTION OF COOPERATION IN A WORLD WITH DISTURBANCES 34

Introduction

Evolutionary Stability: Basic Concepts

Stochastic Stability in Evolutionary Games

	Simulations of Stochastic Evolutionary Processes	
	Summary	
	Bibliography	
III .	EVOLUTIONARY STABILITY IN THE REPEATED PRISONERS' DILEMMA: AN EMPIRICAL ANALYSIS	67
	Introduction	
	Cooperation and the Infinite RPD	
	The Model	
	Simulation Results	
	Discussion	
	Summary	
	Bibliography	
IV .	NEW DIMENSIONS IN RENT-SEEKING	108
	Introduction	
	Sequential Rent-seeking	
	Rent-seeking When the Prize is a Public Good	
	Conclusions	
	Bibliography	

LIST OF FIGURES

Figure

1.1 A Prisoners' Dilemma	3
1.2 Equilibrium Sets in the Infinite RPD	7
1.3 Ecological Success of the Strategies in the Finite RPD	21
1.4 Ecological Success of the Strategies in the Approximately Infinite RPD	22
1.5 Ecologically Successful Strategies in the Finite RPD from Random Initial Distributions	25
1.6 Ecologically Successful Strategies in the Approximately Infinite RPD from Random Initial Distributions	26
2.1 A Game with no ESS	41
2.2 A Game with Multiple ESSs	42
2.3 The Prisoners' Dilemma	45
2.4 Moore Machines for GRIM and ALWAYS DEFECT	52
2.5 Evolution with Mutant Strategies $\delta = .7, \mu = .2$	53
2.6 Evolution with Mutant Strategies $\delta = .7, \mu = .1$	54
2.7 Evolution with Mutant Strategies $\delta = .7, \mu = .01$	54
2.8 Evolution with Mutant Strategies $\delta = .55, \mu = .2$	55
2.9 Evolution with Mutant Strategies $\delta = .55, \mu = .1$	55
2.10 Evolution with Mutant Strategies $\delta = .55, \mu = .01$	56
2.11 Evolution with Different Mutation Rates $\delta = .7, \mu = .1$	59
2.12 Evolution with Different Mutation Rates $\delta = .7, \mu = .01$	59
2.13 Evolution with Constant Mutation Rates $\delta = .7, \mu = .02$	61
2.14 Evolution with Constant Mutation Rates $\delta = .7, \mu = .002$	62
3.1 The Prisoners' Dilemma	70
3.2 Equilibrium Sets in the Infinite RPD	75
3.3 Some Moore Machines	79
3.4 The Stylized GRIM Machine	81

3.5	Mutation of UTIL	85
3.6	Possible Automata	87
3.7	Evolutionary Success without Mutation	89
3.8	Evolutionary Success with Uniform Mutation ($\mu = 0.0001$)	90
3.9	Evolution with Stylized Mutation ($\mu = 0.0001$)	92
3.10	Evolution with Mutants Who Enter as Groups (Uniform Mutation)	94
3.11	Evolution with Mutants Who Enter as Groups (Stylized Mutation)	95
3.12	Evolution with Costly States ($c = 0.05$)	97
3.13	Evolution with Costly States ($c = 0.1$)	98
3.14	Evolution with Costly Transitions ($p = 0.01$)	99
3.15	Evolution with Costly Transitions ($p = 0.1$)	99
4.1	Cournot Equilibrium in the Rent-seeking Game	114
4.2	An Interior Stackelberg Equilibrium in the Rent-seeking Game	118
4.3	A Corner Stackelberg Equilibrium ($v_1/2 \geq v_2$)	119
4.4	Cournot and Stackelberg Equilibrium with Equal Valuations	120
4.5	Total Rent-seeking with Complete and Incomplete Information $v_h = 2, v_l = 1, v_2 = 1.5$	124

LIST OF TABLES

Table

1.1	Results of the Replicated RPDs	19
4.1	Equilibrium Contributions in the Street Light Game	140
4.2	Equilibrium Contributions in the Auto Maker Game	141

CHAPTER I

AUTOMATA PLAY THE PRISONERS' DILEMMA: ANOTHER LOOK AT THE EVOLUTION OF COOPERATION

Introduction

The study of how cooperation can be sustained in repeated play of the Prisoners' Dilemma has been an important topic in game theory and economics. This essay presents the results of a variation on the work reported by Robert Axelrod in his influential 1984 book, *The Evolution of Cooperation*, which was based on the results of a tournament he organized where participants submitted computer programs to play a repeated Prisoners' Dilemma (RPD). Although Axelrod considered primarily political science applications, many economists found his work interesting. Perhaps the most provocative part of his work was his use of evolutionary dynamics to demonstrate the robustness of his results. This part of Axelrod's work provides the motivation for this essay.

A number of authors have explored the relationship between evolutionary dynamical systems and game theory. Josef Hofbauer and Karl Sigmund (1988) provide an excellent survey of this literature, but perhaps the best known of the theoretical work in this field is by John Maynard Smith (1982). Axelrod's evolutionary simulation, however, still provides the most well-known application of these ideas.

Recently, some economists have studied the relationship between equilibrium concepts and the results of evolutionary dynamical systems. A partial list of the

most recent contributors in this area include Larry Samuelson (1989), John Nachbar (1989a), and Dean Foster and Peyton Young (1988). Larry Samuelson shows under very general conditions that reasonable evolutionary dynamical processes lead to either rationalizable or perfect equilibria.¹ Specifically, he proves if the dynamical system employed in this simulation converges, it converges to a perfect equilibrium. John Nachbar has studied convergence of evolutionary selection dynamics under somewhat more restrictive conditions. His work in this area forms the basis of his critique of Professor Axelrod's simulation. Dean Foster and Peyton Young study the behavior of dynamical systems which are subject to stochastic influences. The random effects they have in mind can be thought of as perturbations to the payoff matrix. In the language of biology, we can think of these as fluctuations around some mean reproduction rate or the possibility of encountering mutations in the play of the game. Foster and Young examine what happens as the perturbations vanish and characterize the set of stochastically stable vectors.

The primary purpose of this essay is to report the results of a simulation which is to a large degree a replication of Axelrod's work. Overall, this analysis supports Axelrod's conclusions, and it directly answers part of Nachbar's (1989b) criticism of the results. That is, I show if the game is modeled differently, defection at every iteration need not emerge as the outcome of the evolutionary process. Then I simulate the play of a variation of the RPD game analyzed by Axelrod and examine the robustness of the results. I begin by discussing the Prisoners' Dilemma and some of the arguments used to justify cooperation in both the finitely and infinitely repeated games. Then I describe Axelrod's Prisoners' Dilemma tournament as well as his conclusions. I then review Nachbar's criticism of Axelrod's results and methodology. Finally, I discuss the results of my replication of Axelrod's evolutionary simulation.

¹ For an enlightening discussion on this and other Nash equilibrium refinements see Binmore (1987).

The Prisoners' Dilemma

The version of the Prisoners' Dilemma I employ in this essay is identical to the stage game used by Axelrod in his tournament. There are many possible variants of the Prisoners' Dilemma, but the specific game used for his tournament is represented in figure 1.1.

Figure 1.1 — A Prisoners' Dilemma

		Player II	
		C	D
Player I	C	3, 3	0, 5
	D	5, 0	1, 1

The Prisoners' Dilemma is a common way to represent situations in which either player can benefit by noncooperation regardless of what his opponent does. The story which goes along with this payoff matrix is well known and involves two unseemly people trying to maximize their own well being after they are captured and are being interrogated for some unspeakable misdeed. Either culprit can reduce his sentence by squealing on the other scofflaw regardless of what the other does. Yet, if they were to cooperate with each other and remain tight-lipped they would both be better off.²

It is easy to see the only reasonable outcome in a single play of the game is for both players to defect (D,D) yielding a payoff of one to each player. If both players could be induced to cooperate they would realize a payoff of three each. However, mutual cooperation cannot be an equilibrium in a single play of the game because either player can improve his lot by defecting and receiving a higher payoff. This unique equilibrium in a single play of the game has the unpleasant quality of being the only pure strategy outcome which is not Pareto efficient. Also, the mutually

² For other interesting applications of this type of game see Luce and Raiffa (1957) and Schelling (1960).

defecting equilibrium is the utility minimizing outcome in the sense that the sum of the payoffs is smallest. Not only is (D,D) the only equilibrium outcome in a single play of the game, any finite number of repetitions of the stage game has only one subgame perfect equilibrium outcome which has both players defecting at every turn. This fact can easily be seen using a backward induction argument. Consider the last play of the stage game. Certainly, we expect both players will defect on the last turn. Now consider the penultimate round. We can see neither player has any reason to cooperate since they will both defect in the last round. We continue all the way back to the first round, and we can see both players must defect at every stage. This result seems to be at odds with the body of evidence reported from the experimental literature as well as real world observations.

There is a strand of the game theory literature which attempts to explain why cooperation is often observed. Some of the more well-known papers in the theoretical literature justifying cooperation in the finite RPD have been written by Roy Radner (1986); Abraham Neyman (1985); and David Kreps, Paul Milgrom, John Roberts, and Robert Wilson (1982).

Radner attempts to justify cooperation using uncertainty about which trigger strategy one's opponent will use. He demonstrates uncertainty of this type can make cooperation at the beginning of a finite RPD an optimal strategy. That is, if the players are Bayesian in their approach to the game, the chance one player may cooperate for a while justifies cooperation by the other player. One criticism of this approach is the resulting strategies are not in equilibrium. If the degrees of uncertainty are common knowledge, each player can compute the other's strategy and defect one turn earlier. Radner also uses his notion of ϵ -equilibrium, or near optimality, to justify cooperation in the finitely repeated game. He shows if players are satisfied to get nearly the maximum payoff possible given the other player's strategy, cooperation can be supported in this game. Radner's third bounded rationality

argument to justify cooperation in these situations considers only strategies which can be implemented by finite automata, or Moore machines, which are bounded in size. His argument is similar to that of Neyman who shows how cooperation can arise because the players must use up states of their automata to keep track of what the other player is doing. If the number of states in an automaton is between two and $T - 1$, where T is the number of times the game will be repeated, cooperation can be supported at every iteration. For machines which have a very large number of states relative to the number of iterations, we can come arbitrarily close to the cooperative outcome. The criticisms to this approach are (1) the bound on the complexity of the implementing machines is arbitrary and exogenously imposed on the players and (2) it doesn't seem to represent how players choose their strategies. The first criticism is obvious. The second can be summarized by the argument that a player in a one hundred iteration Prisoners' Dilemma may very well cooperate at the beginning, but it is almost certainly not because he cannot count high enough.

Kreps et al. show if it is common knowledge both that one player is rational and that there is some uncertainty whether the other player is rational (always defects) or plays TIT-FOR-TAT (TFT)³, then there is a unique sequential equilibrium in which cooperation takes place until nearly the end of the game. This argument is based on the idea that since one player is known to be unsure about the state of the world, the other player can profitably pretend to be irrational at the beginning of the game. His opponent can then profit by pretending to be fooled even though he is almost sure his opponent is rational. The above arguments for cooperation in the finite RPD rely on either bounded rationality or incomplete information. Also, "Folk Theorem" type results have been discovered for finite games, but they do not apply to the Prisoners' Dilemma.⁴

³ TFT is the strategy where the player cooperates on the first round and then takes the action chosen by his opponent in the previous round.

⁴ See Benoit and Krishna (1986) and Fudenberg and Maskin (1986).

If we consider only the class of finite RPD games characterized by complete information for all of the players, the only perfect equilibrium has the unique outcome of defection at every stage of the game. However, infinitely repeated games have an abundance of perfect equilibria. The "Folk Theorem" of repeated games allows any feasible individually rational payoff vector to be a perfect equilibrium outcome of the infinite RPD if the discount factor is large enough.⁵ In other words, there is a discontinuity at infinity because any arbitrarily long finite game has a unique perfect equilibrium, but an infinite RPD has an infinite number of them. Much of the theoretical work in this area has been an attempt to reduce the size of the set of equilibrium outcomes in the infinite RPD.

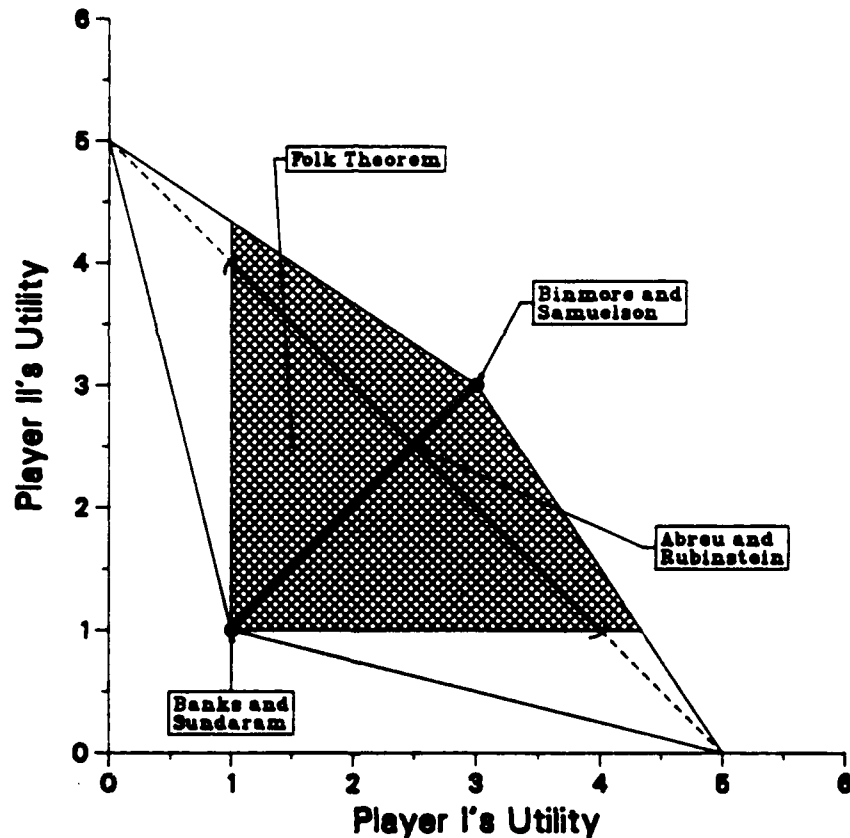
Ariel Rubinstein (1987) and Dilip Abreu and Rubinstein (1988) employed a model where the metagame strategies are implemented by finite automata, and the complexity, or number of states, of the machine is endogenized by making the metagame payoffs depend positively on stage game payoffs and negatively on complexity. By this I mean if two machines yield the same stage game payoffs, the machine with fewer states yields higher metagame payoffs. One model they employed had complexity enter the metagame payoffs lexicographically. Abreu and Rubinstein were able to reduce the set of equilibrium outcomes to the rational payoff vectors on the main and alternate diagonals of the set of feasible outcomes which provide each player with more than his security level.

Ken Binmore and Larry Samuelson (1989) have done interesting work recently using the model developed by Abreu and Rubinstein. They show if we consider the same utility functions used by Abreu and Rubinstein, any evolutionarily stable outcome must have both players cooperating in all but the first round. In other words, Abreu and Rubinstein refine the set of possible equilibrium outcomes by considering complexity of the implementing machine, and Binmore and Samuelson

⁵ See Aumann (1981).

further refine this set to one outcome with an evolutionary stability argument. This result depends crucially on the definition of complexity we choose. Jeffrey Banks and Rangarajan Sundaram (1989) have shown if we use preferences which are lexicographic in complexity and use the number of transitions in the Moore machine as the measure of complexity, the only Nash equilibrium machine defects always.

Figure 1.2 — Equilibrium Sets in the Infinite RPD



These ideas are summarized in figure 1.2. The equilibrium outcomes allowed by the “Folk Theorem” are all those vectors in the shaded region. The equilibrium outcome set from the Abreu Rubinstein model are the rational points on the main and alternate diagonals. Binmore and Samuelson reduced the set to the point (3,3), and Banks and Sundaram’s model yields the unique equilibrium outcome (1,1).

It is interesting to note even though TIT-FOR-TAT plays a central role in Axelrod's theory of cooperation, it is not a subgame perfect equilibrium strategy in the infinite RPD. This fact is easily verified by looking at what happens if the players find themselves at a node in the game tree where one player has defected while the other cooperated. It is clear TFT is not an optimal response for the player who cooperated because the players get themselves "out of sync" and begin a pattern which has the players taking turns defecting while the other cooperates.

The Evolution of Cooperation

Robert Axelrod's theory of cooperation described in *The Evolution of Cooperation* was based on the results of two rounds of a computer Prisoners' Dilemma tournament. Axelrod took the results of the tournament to indicate there are some characteristics common to successful strategies, and these attributes tend to generate cooperative outcomes in the Repeated Prisoners' Dilemma. These characteristics are niceness⁶, provokability, forgiveness, and clarity.⁷ In this section I will describe Axelrod's tournament and subsequent tests of how robust his results are.

Axelrod's tournament consisted of two rounds. The decision rules, or machines, could use the opponent's last move, the sum of payoffs to each player, the iteration number, and a randomly generated set of digits to make its choice for the next move. The first round had fourteen entries which along with the RANDOM strategy⁸ played against each other in an RPD of two hundred iterations. The most successful program submitted in the first round of the tournament was TIT-FOR-TAT. After the results of the first round were announced, a second round of the tournament was held.⁹ This round was designed to eliminate end of game effects by announcing that after each play of the stage game the RPD would continue with probability

⁶ Axelrod uses the term "nice" to describe strategies which are not the first to defect.

⁷ See Axelrod and Dion (1988) for a discussion of these characteristics.

⁸ This strategy chooses between "Cooperate" and "Defect" with equal probability.

⁹ For more information on the first round of the Prisoners' Dilemma tournament see Axelrod (1978).

0.99654. This probability was selected so the expected median game length would be 200 repetitions. This makes the game strategically equivalent to an infinite RPD with the payoffs calculated using a discount factor of 0.99654. In the tournament, however, Axelrod computed the payoffs as the mean of the payoffs from five games whose lengths were randomly determined beforehand.¹⁰ These randomly selected lengths were sixty-three, seventy-seven, 151, 156, and 308 iterations. In this round sixty-two decision rules¹¹ were submitted and played against each other and the RANDOM strategy. TFT proved again the most successful strategy.¹² No one was able to improve on TFT as a strategy in the RPD even though it was reported as the best performer in the first round.

Axelrod then used two different methods to examine how robust TFT's success was in a wide variety of environments. First he constructed a series of hypothetical tournaments, each having a different distribution of the types of programs participating. He reported TFT won five of the six major variants of the tournament and came in second in the sixth. Another, and to me more interesting, test of the results' robustness was to construct a sequence of hypothetical rounds of the tournament employing an evolutionary process widely used in game theoretical and evolutionary biology models.

The evolutionary process can best be understood if we imagine players in an infinite population who are matched randomly and implement strategies with automata which may choose their actions stochastically. That is, a player chooses one strategy and plays it. He cannot randomize between different metagame strategies, but an individual automaton may randomize between "Cooperate" and "Defect" in the repeated game. After playing the RPD, each player is told his payoff and

¹⁰ The metagame payoffs were calculated simply as the undiscounted sum of the payoffs in the stage games.

¹¹ Actually only sixty-one distinct strategies were submitted because two were identical. However, the tournament was run as if there were sixty-two different strategies.

¹² For the complete results see Axelrod (1984).

some information about the payoffs of the other players. Using this information, he may revise his strategy. To see how this works, imagine a game in which there is a strategy space S with n pure strategies, $S = \{s_1, s_2, \dots, s_n\}$. The RPD is then played at time t , and each strategy earns a payoff depending on the strategy with which it is matched. These payoffs can be represented in a $n \times n$ matrix $A(t)$ with elements $a_{ij}(t)$ being the payoff to a player who plays s_i if he is matched against a player who plays s_j at time t . The payoff matrix is indexed for time in this case because the game was designed so the actual number of stage games played by a generation, and hence payoffs, will depend on when the RPD takes place. Then we have

$$A(t) = \begin{pmatrix} a_{11}(t) & a_{12}(t) & \dots & a_{1n}(t) \\ a_{21}(t) & a_{22}(t) & \dots & a_{2n}(t) \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1}(t) & a_{n2}(t) & \dots & a_{nn}(t) \end{pmatrix}.$$

Also, let $\mathbf{p}(t) = (p_1(t), p_2(t), \dots, p_n(t))^T$ be the vector of proportions of each type of player in the population at time $t \in \{0, 1, 2, \dots\}$. Here T indicates the transpose of a matrix. We can then define the expected payoff to a player of strategy s_i at time t as $(A(t)\mathbf{p}(t))_i$. This is the i th element of the vector $A(t)\mathbf{p}(t)$. Also, the expected payoff for a member of the population at time t is $\mathbf{p}(t)^T A(t)\mathbf{p}(t)$. The process we are considering has the proportion of strategy i at time $t + 1$ equal to its proportion at time t multiplied by the ratio of its own expected payoff to the expected payoff of all players. Then we have the following dynamic process:

$$p_i(t + 1) = p_i(t) \left(\frac{(A(t)\mathbf{p}(t))_i}{\mathbf{p}(t)^T A(t)\mathbf{p}(t)} \right).$$

It is easy to see whether a strategy's proportion in the population gets larger or smaller depends on whether or not it is doing better or worse than average. The proportion of a strategy in the next generation depends not only on its own performance, but also on its proportion in the current population. This captures the idea that a strategy must be both successful and observed by others to be copied.

The dynamic process described above is based on the notion some programs or rules would be so unsuccessful they would probably not be used in future rounds, while the superior strategies would be imitated. We can also think of the payoffs from playing the RPD as fitness in the biological sense. That is, strategies with higher payoffs are able to reproduce (asexually) more prolifically than strategies with lower payoffs. In either case, we would expect to see the best strategies flourish and the worst strategies die out. The term "evolutionary process" is used loosely here since the process is not truly evolutionary, but ecological, because no new strategies can be added to the original sixty-three. This process allows us to evaluate how well a strategy will do when the less effective ones cease being important, and only the best ones remain to play each other.

In the evolutionary process Axelrod simulated, $A(t)$ is a 63×63 matrix. He chose $\mathbf{p}(0)$ such that $p_i(0) = 1/63$ for $i = 1, 2, \dots, 63$. Professor Axelrod used the matrix of payoffs he obtained from averaging the results of the five games whose lengths were determined prior to the second round of the tournament for the $A(t)$ matrix. This was held constant for all generations. After simulating one thousand generations of the above dynamical process, Axelrod found TFT was again the best rule in the sense that it was the strategy with the largest proportion of the population. In this paper I take another look at how the robustness of these results can be evaluated in this dynamic framework. Specifically, I reevaluated the results he reported by using a payoff matrix derived from a different procedure; however, I used the same dynamic process to simulate the future rounds of the RPD.

An important point to note in the description of the evolutionary simulation performed by Axelrod that is he used the same matrix of payoffs based on a finite number of plays of the repeated Prisoners' Dilemma for each simulated generation of strategies. This point leads to Nachbar's criticism of the results. He argues in a finitely repeated Prisoners' Dilemma, mutual defection is the only Nash equilibrium

play. Also, the game is dominance solvable to an equilibrium in which players defect on every turn. I will show for the situation Axelrod was describing, the result is not as clear cut as Nachbar would have us believe. I demonstrate what happens if the payoff matrix is derived taking the infinite nature of the RPD into account.

Nachbar's Criticism

John Nachbar (1989b) proved if the evolutionary dynamic process we are discussing converges, it must be to a symmetric Nash equilibrium.¹³ He actually proved a somewhat more general result, but his criticism of Axelrod's simulation rests on this fact. His objection to Axelrod's results is that they are merely a reflection of the strategy set used. Nachbar (1989b) showed for any finite RPD the resulting limit of the dynamical process will never include TIT-FOR-TAT if the strategy space is sufficiently rich. Moreover, if "always defect" (DD) is included in the strategy set, it will be the unique solution after deleting weakly dominated strategies, and defection will be the only equilibrium play, provided a sequence of bridging strategies exist.¹⁴ Hence, if the evolutionary process I described is convergent, it must converge to a point where defection occurs at every move. Any other result must come from rounding error which approximates a very small proportion to zero, terminating the evolutionary simulation too early, or the choice of possible strategies.

Based on the Axelrod's methodology, Nachbar's criticism is well founded. Cooperation may have survived the evolutionary process only because there is no completely defecting Nash equilibrium. Although the 63×63 payoff matrix does not, strictly speaking, reflect a finite Prisoners' Dilemma game, the validity of Nachbar's argument is not affected. It will become clear as I describe Nachbar's argument we

¹³ This is also proved by others. For example, see also Samuelson (1989).

¹⁴ A sequence of bridging strategies in the case of a finite RPD iterated T times is a sequence of strategies, $\{\sigma_1, \sigma_2, \dots, \sigma_{T-1}\}$, where σ_i is the strategy which plays TFT until iteration i and then defects on iteration i and every one afterward.

will obtain defection at every stage in every Nash Equilibrium if we expand the strategy space appropriately.

To understand Nachbar's point, consider the version of the Prisoners' Dilemma described in figure 1.1 as the stage game. Now, I will examine, as Nachbar (1989b) did, a finitely repeated version of this game with only six repetitions of the stage game. We limit our attention to a small subset of all possible strategies. This is necessary because even such a simple game as this has 9.2×10^{18} possible non-randomizing strategies. He considered only TFT, DD, and all possible time bomb strategies. These are strategies which play TFT until some predetermined time and then begin defecting. The strategies are numbered in the following way:

- 1) TFT,
- 2) TFT until stage 6 then DD,
- 3) TFT until stage 5 then DD,
- 4) TFT until stage 4 then DD,
- 5) TFT until stage 3 then DD,
- 6) TFT until stage 2 then DD,
- 7) DD.

Now look at the payoff matrix for this finitely repeated game.

$$A = \begin{pmatrix} 18 & 15 & 13 & 11 & 9 & 7 & 5 \\ 20 & 16 & 13 & 11 & 9 & 7 & 5 \\ 18 & 18 & 14 & 11 & 9 & 7 & 5 \\ 16 & 16 & 16 & 12 & 9 & 7 & 5 \\ 14 & 14 & 14 & 14 & 10 & 7 & 5 \\ 12 & 12 & 12 & 12 & 12 & 8 & 5 \\ 10 & 10 & 10 & 10 & 10 & 10 & 6 \end{pmatrix}$$

It is clear from the above payoff matrix the only pure strategy symmetric Nash equilibrium has DD played by both players. This is also the only rationalizable strategy because it is the unique remaining strategy after successively eliminating weakly dominated strategies. Nachbar showed the only limit point of this evolutionary process has defection at every stage when applied to this game. He also

showed, through a computer simulation, if we were to terminate the simulation early, it could appear we had reached convergence when in fact we had not.

Nachbar asserts this argument will hold for Axelrod's game. That is, if the strategy space were sufficiently rich, the evolutionary simulation would yield the result that defection occurs at every round. To see how this argument would work, imagine first we introduced a strategy which played TFT for the first 307 iterations in Axelrod's tournament and then defected. Clearly this does better than TFT or any other nice strategy. We can also add a strategy which plays TFT for the first 306 moves and then defects on move 307, and so on.

Nachbar's argument correctly reveals if all strategies are possible, defection at every stage would eventually emerge as the limit of the dynamic process. A note of caution is in order here so there is no misunderstanding. It is not true that if all possible strategies are included the result of the evolutionary process will necessarily have DD played by all players. The only Nash equilibrium play has defection at every move; however, there is a continuum of equilibria with this property. As an example, consider a two-stage RPD as Nachbar did. Suppose the strategy "defect then play TFT" (DTFT) is one of the strategies included in the strategy space and was assigned a positive weight initially in the evolutionary simulation. This strategy will always have a positive weight (even in the limit) for the dynamic process under consideration. The intuition behind this is not difficult to see. Here we have to assume convergence because I am not aware of any result which assures convergence in cases like this. If the evolutionary process converges, we know it must converge to a symmetric Nash equilibrium. There will exist in this game an equilibrium which consists of a mixture of DD and DTFT. In the limit the strategies which do not defect in both stages against "always defect" will go to extinction. However, since "always defect" is only better than DTFT when other strategies are present, DTFT will survive in strictly positive proportion. In

other words, although "always defect" weakly dominates DTFT, "always defect" does better only so long as other strategies have a positive weight, but in the limit they will not.¹⁵

Nachbar's argument is not valid for an infinite RPD. I have already noted any feasible individually rational outcome can be justified as a perfect equilibrium outcome in infinitely repeated games if the discount factor is large enough. It is of special interest here that TFT is an equilibrium strategy in the infinite RPD as long as the discount factor is sufficiently close to unity.¹⁶ Axelrod and Douglas Dion (1989) refer to this as the shadow of the future being sufficiently large. In the infinite RPD, the best a player can do against TFT without being nice is to either defect at every stage or alternate defection and cooperation. In this particular stage game, alternating "Cooperate" and "Defect" yields a higher payoff against TFT than always defecting. This means any discount factor δ with $1 > \delta \geq \frac{2}{3}$ will make TFT an equilibrium strategy in the infinite RPD. Note also cooperation can be supported by the GRIM strategy¹⁷ in the infinite RPD.

Evolution in an Approximately Infinite RPD

Nachbar's criticism of Axelrod's work appears to be aimed more at his methodology than the results. This section addresses that methodology. Here I report the results of a replication of Axelrod's simulation which is immune from part of Nachbar's criticism.

The theoretical part of Professor Axelrod's work dealt with cooperation in infinitely repeated games. However, Axelrod's modeling choice for the evolutionary simulation makes his game essentially a finite RPD, subject to the criticisms leveled

¹⁵ See Nachbar (1989b).

¹⁶ A proof of this proposition is provided in Axelrod (1984), 207-8.

¹⁷ The GRIM strategy begins by cooperating and subsequently cooperates as long as its opponent reciprocates. It punishes either its own defection or the defection of its opponent with permanent defection thereafter. This is very similar to Friedman's strategy in the second round of Axelrod's RPD tournament. His strategy cooperates until the opponent defects, then it defects always.

by Nachbar. If the evolutionary simulation were based on the results of an infinite RPD then it would not be at all surprising to find cooperation flourishing.

Perhaps the easiest way to model strategies in an infinite game is use finite automata or Moore machines. A Moore machine is described by a four-tuple $\langle Q, q_0, \lambda, \mu \rangle$, where $Q = \{q_0, q_1, q_2, \dots, q_m\}$ is a finite set of states, $q_0 \in Q$ is the initial state, $\lambda : Q \rightarrow \{C, D\}$ is the output function which maps the state into a strategic choice, and $\mu : Q \times \{C, D\} \rightarrow Q$ is the transition function which maps the current state and the opponent's choice into a state (not necessarily different). Moore machines not only allow us to model strategies in repeated games concisely and conveniently, they let us analytically calculate payoffs for infinite games. When two finite automata play each other they must eventually enter a cycle. Therefore, the sequence of payoffs can be represented by an infinite sequence which can be summed with discounted stage game payoffs. Alternatively, the metagame payoff can be calculated as the limit of the mean of the payoffs in each stage game. This modeling choice was not feasible here because 22 of Axelrod's strategies have randomization as a possibility. Also, some strategies count the moves to compute some summary statistic. Clearly a strategy which counts all stages in an infinitely repeated game cannot be modeled using a finite number of states. Although Moore machines have frequently been used to study cooperation in infinitely repeated games, their use is not appropriate here.

Since calculating the infinite RPD payoffs analytically is not possible, I decided to use computer simulations to analyze the infinite RPD. Again, there are a number of possible ways to simulate the metagame. One possibility is to carry Axelrod's method further and actually determine the length of each repeated game randomly. This alternative is not workable because of the time required. Another variant of this idea is to determine beforehand a number of payoff matrices which correspond to games of different lengths. Then we can apply these as a discrete approximation

to the distribution implied by a termination probability of 0.00346. Performing these simulations would not be a viable option, either, because of the time required. This particular method is also open to another criticism. It would fail to capture the notion that over a sufficiently long time period the players would learn the length of possible RPDs and evolve so as to maximize fitness in that environment. Finally, the stage game could be repeated enough times with discounted payoffs so the total payoff after stopping the simulation will approximate the sum of the payoffs in the infinite game. Perhaps the best way to think of this is as a numerical approximation to a game which is strategically equivalent to Axelrod's proposed tournament.

I decided to use the approach which approximates an infinite game. I simulated five thousand iterations of the stage game so the payoffs would be accurate to four decimal places. Using a discount factor of 0.99654, this means if one player were able to obtain the maximum payoff in every stage game from move 5001 on, his payoff is under reported by 4.9×10^{-6} percent or 4.286×10^{-5} units. In order to compensate for the randomness in some of the strategies, I simulated the infinite RPD five times and averaged the payoffs for the five games. These average scores were then used to form the payoff matrix for the evolutionary simulation.

To perform this analysis I first obtained the FORTRAN code originally used by Axelrod and wrote Turbo Pascal programs to duplicate his tournament strategies. I simulated the RPD on a Zenith 248 computer, and the compiler generated the random numbers. First I replicated Axelrod's tournament. I also replicated one thousand generations of the ecological dynamic process, using the results of the replication of Axelrod's finite RPD tournament for the payoff matrix. I followed Professor Axelrod's original procedures as closely as possible, but there is no practical way to be certain the programs I used do exactly the same thing the programs did in Axelrod's tournament. However, it is clear I have sixty-three programs which closely approximate those in his tournament. My purpose here was not to find an

error in Axelrod's work, but to provide a basis for comparison when I derive results for the approximately infinite game. Therefore, any comparisons I make will be between my finite game replication and the approximately infinite game.

Next, I completed the numerical approximation of an infinitely repeated game with payoffs computed as the discounted sum of the stage game payoffs, with the discount factor 0.99654. I then used the payoff matrix derived from this exercise to simulate one thousand generations of the evolutionary process from an initial distribution where each strategy had an equal weight. In addition, I performed one hundred simulations of the one thousand generation ecological dynamical process with randomly chosen starting points in Δ^{62} , the unit simplex in R^{63} . I used the results of both the finite replication and the approximately infinite game. I will discuss the results of these simulations shortly.

The results of this replication of Axelrod's finite RPD tournament and ecological process are similar to his. The significant exception is strategy #2 (submitted by D. Champion).¹⁸ This strategy is subject to extreme random fluctuations because the author uses a variable to count the number of times the opponent cooperates. However, this variable was never initialized. Therefore, if it was arbitrarily assigned a very large (in absolute terms) negative value initially, the program should have behaved very much like TFT.¹⁹ I initialized the counting variable to zero every time the strategy began playing an RPD. The results of this replication seem remarkably similar to Axelrod's despite the amount of randomization.²⁰

¹⁸ The strategies are numbered according to their order of finish in Axelrod's original tournament.

¹⁹ In Axelrod's tournament the difference in average score between this strategy and TFT was 0.85.

²⁰ For a comparison, see Axelrod (1984), 51.

Table 1.1 — Results of the Replicated RPDs

<i>Submitter's Name</i>	<i>Infinite Replication</i>	<i>Strategy Number</i>	<i>Finite Replication</i>
Anatol Rapoport	1	1	1
Otto Boruhsen	2	3	2
T. Nicolaus Tideman & P. Chieruzzi	3	9	3
Rob Cave (R)	4	4	6
William Adams (R)	5	5	8
Herb Weiner	6	7	4
Francois Leyvraz (R)	7	12	12
Danny C. Champion (R)	8	2	11
Graham Eatherly (R)	9	14	13
Charles Kluepfel (R)	10	10	5
Jim Graaskamp & Ken Katzen	11	6	7
Abraham Getzler (R)	12	11	10
Paul D. Harrington (R)	13	8	9
Paul E. Black (R)	14	15	14
Brian Yamachi	15	17	19
Richard Hufford	16	16	16
D. Ambuelh & K. Hickey	17	26	20
John Maynard Smith	18	24	23
Jonathan Pinkley	19	30	24
Ray Mikkelson	20	20	21
Glenn Rowsam	21	21	17
Edward White, Jr. (R)	22	13	15
John W. Colbert	23	18	26
Tom Almy	24	25	33
Scott Appold (R)	25	22	18
Gail Grisell	26	23	22
Rudy Nydegger	27	31	28
Bernard Grofman	28	28	25
Craig Feathers (R)	29	27	30
Stanley F. Quayle	30	38	31
Nelson Weiderman	31	34	29
Leslie Downing	32	40	35
Martyn Jones	33	43	42
Steve Newman	34	42	37
Roger Falk & James Langsted	35	33	39
Robert Adams	36	35	38
Robyn M. Dawes & Mark Batell	37	36	32
Johann Joss (R)	38	29	48
George Zimmerman	39	41	41
E. E. H. Shurmann (R)	40	44	40
Robert Pebly (R)	41	32	36
George Lefevre	42	37	34
David Gladstein	43	46	44
Fred Mauk (R)	44	19	27
Henry Nussbacher (R)	45	45	45

Table 1.1 — Results of the Replicated RPDs (cont.)

<i>Submitter's Name</i>	<i>Infinite Replication</i>	<i>Strategy Number</i>	<i>Finite Replication</i>
R. D. Anderson	46	39	43
Michael F. McGurrin	47	50	49
Mark Batell	48	47	47
David A. Smith (R)	49	48	46
James W. Friedman	50	52	50
Howard R. Hollander (R)	51	51	52
Robert Leyland (R)	52	49	51
Ric Smoody (R)	53	54	54
W. H. Robertson	54	58	53
Gene Snodgrass	55	56	57
George Hufford	56	53	55
Scott Feld (R)	57	55	56
George Duisman	58	57	58
Harold Rabbie	59	59	59
James E. Hill	60	60	60
Edward Friedland	61	61	61
RANDOM (R)	62	62	62
Roger Hotz (R)	63	63	63

(R) indicates randomizing strategy.

The results of the approximately infinite RPD do not appear to be significantly different than those of the finite RPD if we look only at the order of finish in the tournament. I have summarized the results of Professor Axelrod's original tournament as well as those of my replications in table 1.1. This table doesn't reveal any major differences in the order in which the strategies finished in the finite replication or the approximately infinite game. Switching from a finite game with undiscounted payoffs to an approximately infinite game with discounted payoffs doesn't seem to have caused any major change in the strategies' relative success.

The results of the evolutionary simulation in the case of the approximately infinite game, however, reveal a substantial difference in the relative evolutionary success of strategies in both tournaments. Figures 1.3 and 1.4 depict the approximate dynamic path of the eighteen most successful strategies in the ecological

simulations. Although TIT-FOR-TAT was the most successful strategy after simulating the ecological dynamic process with the payoffs from both the approximately infinite RPD and the finite RPD, some inspection indicates its relative superiority is diminished in the simulation using the approximately infinite game results.

Figure 1.3 — Ecological Success of the Strategies in the Finite RPD





Figure 1.4 — Ecological Success of the Strategies in the Approximately Infinite RPD



If we consider the first fifteen decision rules (from Axelrod's tournament) many of the strategies which were most successful in the finite game were less successful in the approximately infinite game, and some of the strategies which were less successful in the finite game tended to be more successful in the infinite game. Specifically, decision rules 1, 3, 4, 6, 7, 8, 9, 10, and 14 did worse in the approximately infinite game than they did in the replication of Axelrod's tournament. However, strategies 2, 5, 11, 12, 13, and 15 were more successful in the approximately infinite game.

The difference in relative success of the strategies decreased. After one thou-

sand generations, TIT-FOR-TAT's proportion of the population was almost 40 percent larger using the payoff matrix from the finite RPD than it was after the same number of generations using the payoff matrix from the infinite RPD. The first fifteen strategies had a mean proportion of the population of 0.064326 and a variance of 0.002368 using the payoffs from the finite replication. When the simulation is accomplished using the payoffs from the approximately infinite game, the same strategies have a mean proportion of 0.060784 and a variance of 0.001152. Also, these first fifteen strategies accounted for about 96.5 percent of the population in the replication with the finite RPD payoff matrix, but they accounted for only 91.2 percent after simulating the ecological dynamic process with the approximately infinite RPD payoffs. This means the first fifteen strategies are more evenly distributed after 1000 generations using the approximately infinite game payoffs than they are after the simulation with the finite RPD results. When we moved from the finite to the infinite RPD the results changed so after one thousand generations we have less of the population represented by a few strategies, and the difference in the level of success of those strategies decreases.

The difference in the distribution of strategies after one thousand generations is interesting, but its explanation is not entirely clear. One possibility is, of course, the random element. There is no way to separate the stochastic effects. However, in my attempts to mirror Axelrod's simulation as closely as possible, I performed the experiment many times, and what I report here is representative of what emerged. These results are robust in the sense that the chance of getting very different results seems very small. A more likely explanation is the strategies perform better in the game I simulated than they did in the game Axelrod played. I suggest it is not surprising the strategies perform better on average in this game because they were designed to play in a game which is strategically equivalent to the one I simulated.

It is also interesting to note cooperation in general did better in the infinitely

repeated game in a certain sense. The only mean²¹ strategy to survive Axelrod's evolutionary simulation and the finite replication was Paul Harrington's, #8. However, in the approximately infinite simulation it became insignificant by the eight hundredth generation. Again, this result is robust. This is not surprising in the infinite game. The rewards of cooperation overshadow the possible gains from probing for weakness against strategies which retaliate. Once the exploitable strategies have been eliminated, defecting strategies will not do well.²²

Since Harrington's strategy was clearly the most successful, or least unsuccessful, of the mean strategies, a brief description of how it acts is in order. This strategy analyzes the opponent's play of the game and attempts to exploit weakness in the opposing strategy. This rule plays cooperatively for the first thirty-six iterations against a nice strategy. then it defects without provocation. If its opponent makes its first defection on the same move, this strategy assumes it is playing itself unless the opponent defects again. If it thinks it is playing a strategy identical to itself, it cooperates. However, if it is not playing itself, it attempts to take advantage of the opponent. It decides randomly when it should probe the other strategy for weakness. If the opponent appears to be a consistent defector, Harrington's strategy will respond with continual defection. This strategy is interesting because it attempts to identify its twin, and it tracks the opponent's responses to its action over the course of the game. Other strategies monitored the opposing player's actions during the game, but this one was the only mean strategy to have any discernable degree of success.


As another test of the robustness of the evolutionary superiority of cooperation in this environment I simulated the evolutionary process with different initial population distributions. Specifically, I simulated one thousand generations of the ecological dynamic process one hundred times, with randomly chosen initial distri-

²¹ A "mean" strategy is one that is not nice.

²² See Axelrod (1982), 52.

butions. The resulting populations were again very cooperative.

Figure 1.5 — Ecologically Successful Strategies in the Finite RPD from Random Initial Distributions



When I used the payoff matrix from the finite replication, only seven different strategies came in first place. In fact, three of the strategies (1, 3, and 9) accounted for 73 percent of the ecological success in the sense that one of these three strategies came in first in seventy-three of the one hundred simulations. Not surprisingly, TIT-FOR-TAT came in first most frequently, but it only came in first 26 percent of the time. The second best strategy in this exercise, which was submitted by O. Borufsen, finished first 23 percent of the time. In Axelrod's tests for robustness TFT finished first in five of six tournaments which had different distributions of strategies. This result differs significantly from the results obtained when we started

with equal weights on all strategies, and suggests TFT is not as evolutionarily fit as Axelrod's evidence indicates. In all cases, the simulations seemed to be converging to a cooperative equilibrium. These findings are summarized in figure 1.5, where I show which strategies finished first and how often they did so when starting from randomly determined initial distributions.

I also performed similar simulations, starting from the same randomly chosen initial distributions, using the payoff matrix from the approximately infinite game. The set of strategies which finished first here contains the set of those which finished first using the results of the finite game. However, strategies 5, 12, and 14 also finished first in this simulation. In fact, these three rules accounted for 9 percent of the first place finishes. Again, convergence was always to a cooperative equilibrium. These findings are summarized in figure 1.6.




Figure 1.6 — Ecologically Successful Strategies in the Approximately Infinite RPD from Random Initial Distributions

Next I will examine the successful strategies other than TIT-FOR-TAT. I will look specifically at the nine other strategies which finished first in one of the simulations which started from a random initial distribution using the payoff matrix from either the finite or approximately infinite games.

1. Borufsen's strategy (#3) uses TFT as its main rule. However, it maintains statistics on the play of the game so it can take advantage of specific irrational opponents. For example, if it detects its opponent is random or defective, it defects on 25 consecutive moves. This program also checks to see how the opponent responds to his defection. However, like all the other successful strategies, this one will never be the first to defect.

2. The third most successful strategy was Tideman and Chieruzzi's (#9). This strategy also has characteristics which are similar to TFT. It is nice, but less provokable than TFT. The authors of the rule use an expression which depends on the difference in scores between the two players, the change in the difference in the players' scores, and the number of defections the strategy has detected to determine whether or not to retaliate. After the opponent defects, this strategy will continue defecting until it has either closed the gap in the scores sufficiently or defected ten consecutive times. After ten consecutive defections, this strategy returns to cooperating as long as the opponent does not switch back from cooperate to defect and has not defected too often in the past.

3. Weiner's strategy (#7) is a variation of TFT. It can be thought of as TFT with special forgiveness because it ignores a defection if it has been more than twenty moves since the last one and the last defection was followed by cooperation. Also, this program always defects if the opponent has defected at least five times in the last twelve moves.

4. Graaskamp and Katzen's strategy (#6) can best be thought of as checking its own score against certain milestones during the game. As long as its score is high

enough at predetermined moves it will play TFT. However, if its score falls below the minimum acceptable at a checkpoint, it enters an absorbing state which defects at every turn.

5. Cave's strategy (#4) is similar in spirit to #6 because it evaluates the play of the game at predetermined points. In addition, if the opponent has defected less than eighteen times, this rule defects in response to defection by his opponent with a probability of $1/2$. However, after the opponent's eighteenth and subsequent defections, this strategy defects with certainty. Also, if the opponent is overly defective according to certain rules, this strategy gives up and defects until the opponent's percentage of cooperative moves increases to acceptable levels.

6. The last of the strategies to finish first in the simulations with the finite game payoff matrix is Kluepfel's (#10). This strategy maintains a history of the last three moves. It cooperates until the opponent defects. Then it begins randomizing depending on the three iteration history.

The next three strategies finished first only when the infinite game payoffs were used.

7. Leyvraz's strategy (#12) is also similar to TFT; however, its retaliation rules are slightly different. This strategy keeps track of the opponent's last three moves. If he defected the last two times, this rule defects with probability .75. If the opponent defected two moves ago but not on the last move, this rule defects with certainty. Finally, if the last move was the opponent's only defection in the last three iterations, this strategy defects with probability .5. Otherwise, it will cooperate.

8. William Adams's strategy (#5) is again similar in spirit to TFT, but it is less provokable. It starts with a threshold of four defections. Once the threshold is crossed, it defects and then adjusts the threshold by cutting it in half. It continues calculating the threshold after it is less than one because it then becomes the probability this rule cooperates after a defection.

9. Eatherly's strategy (#14) is a very simple strategy. This rule calculates the proportion of defections in all previous moves and uses this as the probability it will retaliate against a defection.

Overall, the results of this work reinforce Professor Axelrod's theory of cooperation. That is, these simulations indicate cooperation can flourish and can be supported through swift and sure retaliation to a defection. TFT proved to be superior to all other strategies in these replications, just as it was in Axelrod's original simulation. However, it seems to be substantially less superior in this environment than it appeared to be in Axelrod's report. Although TIT-FOR-TAT was more successful than any other strategy in these simulations, it is most unlikely this evolutionary process is converging to virtually all TIT-FOR-TAT as conjectured by Axelrod.²³ In fact, if there is convergence here, it can be to any of a continuum of symmetric Nash equilibrium.

The simulation of the dynamical process using the payoff matrix from the approximately infinite RPD is immune from at least part of Nachbar's criticism. In the simulations I have reported here, no argument can be made TFT is subject to successive elimination of weakly dominated strategies. However, a purely defecting equilibrium still has no chance of success in this game because there are no purely defecting strategies. Therefore, we must be careful to refer to TFT's success as being conditional on the submitted strategies.

There still may be reasons why cooperation cannot be sustained in this situation, but Nachbar's criticism does not negate the robustness of Axelrod's theory in this RPD which is modeled only slightly differently. There does remain, however, the possibility a defecting equilibrium outcome may prevail if we allow defecting strategies. In subsequent essays I will discuss this further, but for now I will note both cooperation and defection are possible equilibrium outcomes in the infinite

²³ See Axelrod (1984), 55.

RPD with discounted payoffs.

Summary

Robert Axelrod's evolutionary simulation of future rounds of his RPD tournament has provided a basis for our understanding of cooperation, and how it can evolve in a population. However, for what I believe to be methodological rather than sound theoretical reasons, his study was open to the criticism a purely defecting equilibrium would emerge if the strategy space was sufficiently rich. I demonstrated if the game was modeled differently to represent an infinitely repeated game, it is not necessarily the case a defecting equilibrium would emerge.

The infinitely repeated version of Axelrod's simulation provided some additional insight into the relative evolutionary fitness of the strategies used by Axelrod. TIT-FOR-TAT proved to be less evolutionarily superior to the other decision rules in the approximately infinite game than it was in the finite replication. By examining the evolutionary process from random initial population distributions, I obtained results which also tend to support Axelrod's theory and findings. However, TIT-FOR-TAT did not emerge as exceptionally dominant when starting from these initial distributions. In all cases, the simulated future generations revealed the robust nature of cooperation in this tournament. All of these findings are environment specific because there is no possibility of defection emerging here. Although Nachbar's criticism does not fully apply to the simulation I performed using the results of an approximately infinite game, it is still true a defecting equilibrium had no chance because it was never introduced into the game. It is clear if we introduced the "always defect" strategy into this environment, it would not fare well. However, it is equally clear if we added a sufficiently large number of defecting strategies or if we start close enough to the point where every player defects, the evolutionary process would converge to a purely defecting equilibrium.

The notion of evolutionary stability also has no force here if the strategy set is sufficiently rich. I will expand on this idea in subsequent essays. For now, however, it must suffice to note evolutionary success of cooperation is possible when the game is modeled to reflect the infinite nature of the RPD when each stage game ends with a positive probability less than unity.

Bibliography

- Abreu, Dilip and Ariel Rubinstein. "The Structure of Nash Equilibrium in Repeated Games with Finite Automata." *Econometrica* 56 (November 1988):1259-1281.
- Aumann, Robert J. "Survey of Repeated Games." In *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*. Bibliographisches Institute Mannheim. Wein, Zurich. 1981.
- Axelrod, Robert. "A Report on the Computer Tournament for the Iterated Prisoner's Dilemma." Institute for Public Policy Studies Discussion Paper No. 128. Institute for Public Policy Studies, The University of Michigan, 1978.
- Axelrod, Robert. *The Evolution of Cooperation*. New York: Basic Books, 1984.
- Axelrod, Robert and Douglas Dion. "The Further Evolution of Cooperation." *Science* 242 (9 December 1988):1385-1390.
- Banks, Jeffrey and Rangarajan Sundaram. "Repeated Games, Finite Automata, and Complexity." University of Rochester Working Paper No. 183, March 1989.
- Benoit, Jean-Pierre and Vijay Krishna. "Finitely Repeated Games." *Econometrica* 53 (July 1985):905-922.
- Binmore, Kenneth G. "Modeling Rational Players, Part I." *Economics and Philosophy* 3 (1987):179-214.
- Binmore, Kenneth G. and Larry Samuelson. "Notes on Evolutionarily Stable Strategies in Repeated Games Played by Finite Automata." unpublished manuscript, 1989.
- Foster, Dean and Peyton Young. "Stochastic Evolutionary Game Dynamics." unpublished manuscript, 20 September 1988.
- Fudenberg, Drew and Eric Maskin. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." *Econometrica* 54 (May 1986): 533-554.
- Hofbauer, Josef and Karl Sigmund. *The Theory of Evolution and Dynamical Systems*. Cambridge: Cambridge University Press, 1988.

- Kreps, David, Paul Milgrom, John Roberts, and Robert Wilson. "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma." *Journal of Economic Theory* 27 (1982):245-252.
- Luce, R. Duncan and Howard Raiffa. *Games and Decisions*. New York: John Wiley & Sons, 1957.
- Maynard Smith, John. *Evolution and the Theory of Games*. Cambridge: Cambridge University Press, 1982.
- Nachbar, John. "Evolutionary Selection Dynamics in Games: Convergence and Limit Properties." unpublished manuscript, 1989.
- Nachbar, John. "The Evolution of Cooperation Revisited." unpublished manuscript, 1989.
- Neyman, Abraham. "Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoner's Dilemma." *Economic Letters* 19 (1985):227-229.
- Radner, Roy. "Can Bounded Rationality Resolve the Prisoner's Dilemma?" *Essays in Honor of Gerrard Debreu*, New York: North-Holland, 1986.
- Rubinstein, Ariel. "The Complexity of Strategies and the Resolution of Conflict: An Introduction." unpublished manuscript, 1987.
- Samuelson, Larry. "Evolutionary Foundations of Solution Concepts for Finite, Two-player, Normal-form Games." unpublished manuscript, 1989.
- Schelling, Thomas C. *The Strategy of Conflict*. Cambridge: Harvard University Press, 1960.

CHAPTER II

THE EVOLUTION OF COOPERATION IN A WORLD WITH DISTURBANCES

Introduction

In the first essay I looked at the ecological dynamics of a very special population. I showed the evolutionary success of cooperative strategies was possible in Robert Axelrod's tournament if the game was properly modeled. I also simulated an approximately infinitely Repeated Prisoners' Dilemma (RPD), as well as a dynamic ecological process, to determine the effect on the results with the game modeled differently. Mutual cooperation was the most evolutionarily fit outcome in every case. However, the results did not preclude the possibility a purely defecting equilibrium could still emerge in another environment. In Axelrod's environment purely defecting strategies never had a chance to succeed because none were ever introduced. We still have not explored what conditions lead an evolutionary process to result in a cooperating or a defecting population. In this essay, I look at how perturbations influence the evolution of a population playing an infinite RPD. I will consider perturbations, or trembles, in the play of the game as well as perturbations in the payoff matrix.

Axelrod's (1984) analysis of the success of cooperation in the repeated Prisoners' Dilemma (RPD) is the most famous of many studies on the subject. However,

many authors have studied RPDs more generally in an evolutionary framework. Here I review some of these works. These papers offer a wide variety of models, and the results vary similarly. John Maynard Smith and G. R. Price (1973) are among those credited with introducing evolutionary game theory, but perhaps the most well known of the works on the subject is Maynard Smith's 1982 book *Evolution and the Theory of Games*. Since its introduction many people in a broad range of disciplines have applied evolutionary game theory to the study of natural and social science problems. It has frequently been used to study coordination games, common interest games, and the Prisoners' Dilemma. My purpose here is to provide a brief review of some of those studies, especially those analyses which deal with the Prisoners' Dilemma.

In his 1984 book *The Evolution of Cooperation* Axelrod used an evolutionary argument to justify cooperation in the his RPD tournament. The notion of evolutionary stability he used is less restrictive than the one I apply here. For now, I will define an evolutionarily stable strategy (ESS) as one which cannot be invaded. This definition is somewhat less rigorous than the one I use in then next section. An ESS must be a best reply to itself. In addition, if a strategy is not a unique best reply to itself, an alternate best reply must do worse against itself than the indigenous strategy does against the alternate best reply. This is called the stability condition which Axelrod eliminated. He called his version of this idea collective stability.¹ This is nothing more than requiring a strategy be part of a symmetric Nash equilibrium. In choosing this definition, Axelrod disallowed the possibility an alternate best reply strategy which earns an equal payoff against both the indigenous and invading strategies would be able to successfully infiltrate the native population. The idea nice² populations can be invaded has been used frequently to

¹ See Axelrod (1984) p. 217.

² Axelrod (1982) defined nice strategies to be those which never defect first.

demonstrate the evolutionary instability of strategies like TIT-FOR-TAT (TFT)³ and is the basis for the computer simulations which I discuss later.

Recently, Robert Boyd and Jeffrey Lorberbaum (1987) and Joseph Farrell and Roger Ware (1989) show an ESS does not exist in the infinite RPD in the sense of Maynard Smith.⁴ Specifically, part of Boyd and Lorberbaum's contribution in this area was their proof no pure strategy ESS is possible in the infinite RPD if every strategy can appear through mutation. That is, no pure strategy is immune from invasion in the infinite RPD. To gain some intuition, consider any candidate ESS. If we allow arbitrary strategies to appear through mutation, there will be strategies which play the same as the indigenous strategy along the equilibrium path but differently off it. These actions off the equilibrium path may allow it to be invaded. Farrell and Ware extended this and showed there is no ESS in finitely mixed strategies. Hence, unless we alter the notion of evolutionary stability we are doomed to failure in trying to find a solution to the infinite RPD. Here, then, we find one of the weaknesses of the ESS concept; there is a nonexistence problem. Reinhard Selten's concept, the limit ESS, is the most significant attempt to find a satisfactory way to eliminate the nonexistence problem. However, as we shall see, there are also nonexistence problems with the limit ESS.

Finding perfect equilibrium outcomes in infinitely repeated games is not a problem of scarcity but one of surplus. The "Folk Theorem" of repeated games assures any feasible and individually rational outcome is a possible perfect equilibrium outcome.⁵ Many economists and game theorists argue that when multiple equilibria exist, the solution to the game has the players achieving an efficient outcome. Using the terminology of Harsanyi (1977) and Harsanyi and Selten (1988),

³ This is the strategy which cooperates in the first period and subsequently chooses the action played by the opponent on the previous move.

⁴ This is what Selten calls a normal form ESS.

⁵ See Auman (1981).

the players should be expected to choose this outcome based on payoff dominance⁶ if possible. This intuition is especially appealing for repeated games since there is recurrent interaction between the players. The argument for selecting payoff dominant outcomes is also very strong when some sort of coordinating device is present. For example, if a payoff dominant equilibrium outcome exists and we allow preplay communication between players, we should expect the equilibrium they agree upon to produce that outcome. Even in the absence of such a coordinating device, many game theorists and economists argue we should expect to see a payoff dominant equilibrium outcome. In the infinite RPD, of course, this means we should expect to see a cooperative outcome.

Drew Fudenberg and Eric Maskin (1990) have shown cooperation in infinite RPD games is evolutionarily stable in an environment where only strategies of finite complexity can be used, payoffs are calculated as the limit of the mean stage game payoffs, and noise exists. In their model, "noise" refers to some chance an action may be misperceived by a player's opponent. John Miller (1987) performed interesting simulations of the evolution of automata which were modeled as a string of digits. The impact of noise on the model was one of the factors he examined. He simulated the evolution of a population playing a finite RPD in which there were strictly positive probabilities of errors in the transmission of information about what a player's opponent did on the previous move.⁷ His results are difficult to characterize in a sentence or two; however, the methodology he used is interesting and the results indicate the effects of noise on a model like this are not negligible.

Ken Binmore and Larry Samuelson (1989) proved a result similar to Fudenberg and Maskin's using a model in which metagame strategies are implemented by finite automata and metagame payoffs, or profits, increase in stage game payoffs and

⁶ A payoff dominant equilibrium outcome, in this case, is one in which both players obtain a higher payoff than they would get in any other equilibrium outcome. Obviously these do not always exist.

⁷ Specifically, he considered error rates of 1 percent and 5 percent.

decrease lexicographically in complexity. The complexity measure in their model is the number of states in the implementing machine, and the stage game payoff is computed as the limit of the mean. Their theorem is general in that it is true for a large class of games. They show any equilibrium strategy which yields less than a utilitarian payoff⁸ can be invaded by a strategy which does as well, in stage game payoffs, as a native does if it plays a native but does better against itself than it does against a native. They rely on Dilip Abreu and Ariel Rubinstein's (1988) result that any payoff vector on the main diagonal of the set of feasible and individually rational payoffs is achievable as the outcome of a pair of equilibrium machines. The payoff maximizing equilibrium machine is the only one which can be invasion-proof because it is the minimum complexity equilibrium machine which results in the utilitarian outcome. In their model, the invading machines are able to signal their opponents. If they are not playing their own type, it makes no difference in the stage game payoffs because they are computed as the limit of the mean, and in the limit they get the same payoff as the native strategy. The fact they earn higher payoffs against machines like themselves insures them a higher average payoff than the native machine gets.

Arthur Robson (1989) has studied evolution in coordination games and the Prisoners' Dilemma. In his model a mutant is introduced which destroys any ESS which has a lower payoff than another. He shows the efficient outcome can be temporarily attained in the Prisoners' Dilemma using a signal he calls the "secret handshake." Robson claims such a model makes the evolution of cooperation unavoidable. Cooperation cannot be permanent in this model because the same possibilities which allow those with the secret handshake to invade a defecting population will allow other strategies to give the signal and then exploit the invading strategy.

⁸ This is an outcome in which the sum of payoffs to both players is maximized.

In a recent paper, Yong-Gwan Kim (1989) showed there does not exist an ESS in infinite RPDs without perturbations. Additionally, he proved a "Folk Theorem" which states we can find a limit ESS which is arbitrarily close to any point on the main diagonal of the convex hull of feasible payoffs.

In the next section I discuss evolutionary stability and the RPD. Specifically, I will review Selten's ESS concepts. After that I will discuss the introduction of chance elements to the dynamic process. This can be distinguished from the perturbations in Selten's limit ESS which are probabilities a player will make an error in the play of the stage game. The random elements in the stochastic dynamic model are continuous perturbations to the dynamic process. We can think of these as variations in the payoffs or disturbances caused by mutations from the failure of strategies to breed true. In a very simple framework, we can characterize when we should expect to find defecting populations and when we should expect to find cooperating populations. Then I describe the results of computer simulations where strategies mutate in an intuitively appealing way based on the idea metagame strategies are implemented by finite automata, or Moore machines. Finally, I summarize the findings in this essay.

Evolutionary Stability: Basic Concepts

The purpose of this section is to review the evolutionary stability concepts which have been applied to games. Most of the work in this field is due to Selten. He developed the concept of limit ESS which is a modification of the normal form ESS introduced by John Maynard Smith.⁹ There are two somewhat contradictory problems with the notion of evolutionary stability. I already identified the first one, nonexistence. The second problem is it may fail to select the most intuitively appealing, or "best," equilibrium when multiple equilibria exist. In this section, I discuss Selten's attempts to solve the nonexistence problem with the limit ESS.

⁹ For a concise explanation of Maynard Smith's concept see Maynard Smith (1982).

The ESS concept is based on the idea that the evolutionary success from playing a certain strategy depends not only on which strategy a given player chooses but also on the characteristics of the population in which it plays. Loosely speaking, an evolutionarily stable strategy is one which can not be invaded. The requirements for a strategy to be an ESS are more stringent than those of a Nash equilibrium. A symmetric Nash equilibrium strategy is neutrally stable in the sense that a population which plays this strategy will have no opponents which are strictly more successful than the native strategy. This is obvious from the definition of Nash equilibrium. For a strategy to be an ESS, it must be immune from invasion from any feasible mixed strategy.

I begin by defining evolutionary stability in the context of a two-player symmetric normal form game. By convention, a two-player normal form game G is defined by a four-tuple $G = \langle S_1, S_2, u_1, u_2 \rangle$ where S_i is Player i 's strategy space, $S = S_1 \times S_2$, and $u_i : S \rightarrow R$ is Player i 's payoff function. In this game the two players simultaneously choose $s \in S_1$ and $t \in S_2$, and receive payoffs $u_1(s, t)$ and $u_2(s, t)$. A strategy for Player i is defined as $\sigma_i \in \Delta(S_i)$ where $\Delta(S_i)$ is the set of all probability measures on S_i . More formally, if there are k pure strategies, then $\Delta(S_i) = \{\sigma_i \in R^k : \sigma_i(s) \geq 0, \sum_{s \in S_i} \sigma_i(s) = 1\}$. A game is symmetric if $S_1 = S_2$ and $u_1(s, t) = u_2(t, s)$. In other words, the role of each player is irrelevant in a symmetric game. Except where necessary for clarity, I will abuse notation slightly and use $S = S_1 = S_2$ and $u = u_1 = u_2$.

Using the above notation, a strategy σ is an evolutionarily stable strategy of a symmetric game if and only if:

$$u(\sigma, \sigma) \geq u(\sigma', \sigma) \tag{1}$$

and

$$u(\sigma, \sigma) = u(\sigma', \sigma) \implies u(\sigma, \sigma') > u(\sigma', \sigma'). \tag{2}$$

for all $\sigma' \neq \sigma, \sigma' \in \Delta(S)$. In words, these conditions require any strategy be a best reply against itself, and if there are alternate best replies, the ESS must do strictly better against an alternate best reply than the alternate best reply does against itself. The first requirement above means any ESS is a Nash equilibrium. However, the converse is not true because of the second condition, which is often called the stability condition.

The above definition is actually a characterization of a more basic requirement which holds in pairwise random matching models.¹⁰ Assuming von Neumann-Morgenstern utility functions, the following must hold for a sufficiently small proportion of invaders, $\epsilon > 0$:

$$(1 - \epsilon)u(\sigma, \sigma) + \epsilon u(\sigma, \sigma') > (1 - \epsilon)u(\sigma', \sigma) + \epsilon u(\sigma', \sigma').$$

This requires the expected payoff from playing σ in a population with proportion ϵ of σ' to be strictly greater than the expected utility of those playing σ' .

Figure 2.1 — A Game with no ESS

		Player II		
		A	B	C
Player I	A	0, 0	4, -5	-5, 4
	B	-5, 4	0, 0	4, -5
	C	4, -5	-5, 4	0, 0

To see some of the problems which can arise, consider the game in figure 2.1. This is a version of the children's game rock-scissors-paper. This game has a unique trembling hand perfect¹¹ symmetric Nash equilibrium $(1/3, 1/3, 1/3)$ and no ESS.¹² The equilibrium is easily verified. However, this equilibrium cannot be an ESS

¹⁰ See Maynard Smith (1982) for details.

¹¹ See Binmore (1987) for an excellent discussion of this and other Nash equilibrium refinements.

¹² This game is adapted from one in Foster and Young (1989).

because condition (2) fails. Since the unique equilibrium is a completely mixed strategy, it must be true that condition (1) holds with equality for all $\sigma' \in \Delta(S)$. The equilibrium strategy will have a payoff of $-1/3$ when matched with any other mixed strategy. However, any other strategy will get more than $-1/3$ when it plays itself. Hence, it fails condition (2) above.

As an example of the second problem with the ESS, consider the common interest game presented in figure 2.2. This game has three Nash equilibria, all of which are perfect. They are $(1, 0)$, $(0, 1)$, and $(1/3, 2/3)$.¹³ Evolutionary stability eliminates only the mixed strategy equilibrium in this case. The more intuitively appealing of these is $(1, 0)$ because it yields the payoff dominant outcome. However, evolutionary stability fails to eliminate the $(0, 1)$ equilibrium.

Figure 2.2 — A Game with Multiple ESSs

		Player II	
		A	C
Player I	A	2, 2	0, 0
	C	0, 0	1, 1

Next we shall see the limit ESS fails to help in either of these cases. That is, since any ESS is also a limit ESS, it provides no ability to discriminate between multiple ESSs. Application of the limit ESS concept also fails to provide any help in the rock-scissors-paper game.

In order to describe Selten's more general limit ESS concept, we must first introduce the idea of perturbed games. I do this following Larry Samuelson's (1989) exposition modified for the fact that I will deal only with symmetric games in this essay. For any game G , we can define a perturbed game \hat{G} by replacing the strategy set $\Delta(S)$ with $\{\sigma \in R^k : \sigma(s) \geq \eta(s) \geq 0, \sum_{s \in S} \sigma(s) = 1\}$. The function

¹³ The notation here is $(p, 1 - p)$ where p is the probability assigned to strategy A.

$\eta : S \rightarrow [0, 1]$ assigns a minimum probability of a tremble to some of the strategies in S . We can then define a strategy σ^* to be a limit ESS of a game G if, for every $\epsilon > 0$, there is at least one perturbed game $\hat{G}(\epsilon)$ with an ESS, σ' , such that

$$\max_{s_i \in S} \{\eta(s_i)\} < \epsilon \quad \text{and} \quad \left[\sum_{s_i \in S} (\sigma^*(s_i) - \sigma'(s_i))^2 \right]^{1/2} < \epsilon.$$

The conditions above require the trembles to be less than ϵ in magnitude and the equilibrium in the perturbed game to be close enough, in the Euclidean metric, to the candidate limit ESS. It is clear any ESS is also a limit ESS because we can always set $\eta(s_i)$ to zero for all $s_i \in S$.

As Samuelson points out, the trembles in the definition of the limit ESS have the effect of breaking ties in the payoffs to strategies which yield identical payoffs without the trembles. There may be a weak best reply to an equilibrium strategy which keeps it from satisfying the conditions for evolutionary stability. These trembles have the effect of turning weak best replies into inferior replies, hence potentially enlarging the set of evolutionarily stable strategies.

Now we can see the limit ESS concept fails to be of any help in the two games I described. In the case of the rock-scissors-paper game the intuition is clear. Since $(1/3, 1/3, 1/3)$ is a completely mixed strategy Nash equilibrium, it must be true if one player employs the equilibrium strategy the other player is indifferent among the strategies he can play. We have already seen $(1/3, 1/3, 1/3)$ is not an ESS. Also, any perturbation away from $(1/3, 1/3, 1/3)$ will make another strategy a strict best reply. For example, suppose we choose the trembles so as to make A slightly more likely for Player I. This will make C a unique best reply to the perturbed strategy. Hence, the perturbed strategy cannot be an ESS.

Now consider the second game, the common interest game. As I stated earlier, $(1, 0)$ is the more intuitively plausible result of an evolutionary process. However, the limit ESS concept is again of no help in distinguishing between the two pure

strategy equilibria. The significance of this is even more plain if we substitute 10^{1000} for 2. That is, for any arbitrarily high payoff for (1, 0) relative to the payoff for (0, 1), both A and C are limit ESSs.

The problems I have described with the ESS and limit ESS are not trivial. Certainly, any theory of evolutionary games must have as one of its primary goals the ability to find the game's solution. In other words, given a game and a dynamical process, we would expect a reasonable theory of this class of games to provide us with an idea of what we can expect as the result of the evolutionary process. In the next section, I describe a theory of evolutionary processes which avoids the two problems I described here.

Stochastic Stability in Evolutionary Games

In the last section we looked at how disturbances in the form of trembles in the play of the game could be used in the evolutionary stability framework. We found by allowing these trembles we could expand the set of ESSs to help reduce the nonexistence problem. However, we also found the set of limit ESSs contains the set of normal form ESSs. That is, any normal form ESS must also be a limit ESS. The converse obviously fails. Now I will describe a theory of evolutionary game dynamics developed by Dean Foster and Peyton Young (1988) in which stochastic disturbances are explicitly modeled. The basic difference between these two notions is that in the case of the evolutionary stability we are looking at stability against a one-time disturbance in the system. In other words, if the system is displaced by a single disturbance some arbitrarily small distance from an equilibrium and returns to the equilibrium, that equilibrium is an ESS. On the other hand, the idea behind stochastic stability is that in the presence of continuously applied disturbances the system will select a set of states near which it will stay.

Foster and Young are the only social scientists to have studied evolution in the presence of this type of stochastic influences. W. G. S. Hines (1982) examined

the effects of strategy mutations in a different context. He looked at changes in diversity of the population due to mutation while the population is in strategy equilibrium, which he defined as a state when the average strategy remains constant over time. He shows under certain conditions, including no deterministic trend and independence of the mutation rates, the population will become increasingly diverse over time. The idea Foster and Young developed is quite different. However, it is not a refinement of either of the ESS concepts I described in the last section since not all stochastically stable strategies are ESSs, and not all ESSs are stochastically stable.

Figure 2.3 — The Prisoners' Dilemma

		Player II	
		C	D
Player I	C	R, R	S, T
	D	T, S	P, P

In order to discuss the notion of stochastically stable strategies, I introduce an extremely simple version of the infinite RPD. The version of the Prisoners' Dilemma I will use as the stage game is represented in figure 2.3 with $T > R > P > S$ and $R > (T + S)/2$.

I will begin by considering only two possible metagame strategies in the infinitely repeated version of this stage game. I will consider one cooperative strategy, the GRIM strategy¹⁴, and the "always defect" (DD) strategy. DD is always a subgame perfect equilibrium strategy in the RPD because D is the dominant strategy in the stage game. GRIM is also a subgame perfect Nash equilibrium strategy in the infinite RPD if the common discount factor, δ , is at least $(T - R)/(T - P)$. The values I will use in the payoff matrix are $T = 5$, $R = 3$, $P = 1$, and $S = 0$.¹⁵ This

¹⁴ The GRIM strategy begins by cooperating, and if either player defects it defects continually afterwards.

¹⁵ These are the same values used in Axelrod's tournament and in the previous essay.

means if $\delta \geq 1/2$ both strategies form symmetric subgame perfect equilibria. I will limit the discussion to those cases.

I will begin with a discount factor of $\delta = 0.9$. This yields the following payoff matrix for the infinitely repeated game.

$$A = \begin{pmatrix} 30 & 9 \\ 14 & 10 \end{pmatrix}$$

In the payoff matrix, A , an element a_{ij} represents the payoff to a player who employs metagame strategy i when he meets a player who uses strategy j . Here strategy 1 is GRIM and strategy 2 is DD. As we can see, neither of the ESS concepts I described earlier can help us decide which of the metagame strategies is more evolutionarily fit. In this simple game both pure strategies are symmetric Nash equilibria. There is also a mixed strategy symmetric Nash equilibrium, $(\frac{1}{17}, \frac{16}{17})$ which is not an ESS. Applying the normal form and limit ESS concepts here eliminates the mixed strategy equilibrium but does nothing to help us select between the cooperative and defecting equilibria. Both pure strategies are limit ESSs. However, as we shall see, Foster and Young's stochastic stability concept will always select at least one of the two strategies in this simple example.¹⁶

Suppose the payoffs in the above matrix are continuously subjected to random perturbations. We can think of these disturbances as fluctuations in fitness rates due to chance or mutational effects. The difference between this concept and evolutionary stability is if a strategy is evolutionarily stable it is stable against a single small perturbation. However, if these disturbances are continuously taking place, evolutionary stability doesn't insure a population will be secure against invasion by mutants. We can apply the notion of stochastic stability to show as the variance of the perturbations gets close to zero the population will almost always be arbitrarily close to the cooperative equilibrium in this game.

¹⁶ As we shall see, both strategies are stochastically stable in this particular game if and only if $\delta = 0.6$.

Foster and Young demonstrate this idea using the continuous-time dynamical evolutionary process

$$\dot{p}_i(t) = p_i(t) \left[(A\mathbf{p}(t))_i - \mathbf{p}(t)^T A\mathbf{p}(t) \right].$$

Here, $\mathbf{p}(t) = (p_1(t), p_2(t))^T$ is the vector of proportions of the population playing each strategy at time t , and $p_i(t)$ is the proportion of the population playing strategy i at time t . T indicates the transpose of a matrix. The vector $\mathbf{p}(t)$ can be thought of the state of the system. In pairwise random matching models with an infinite population, we can think of a state of the system as a mixed strategy against which each participant plays. This dynamic process has essentially the same dynamic behavior as the discrete time model used by Axelrod (1984). Although generalization from the continuous time process to the discrete time process is not straightforward, the two models yield similar dynamics since we consider only symmetric situations.¹⁷

Now we will add a random element to the process. Suppose the dynamic process can be represented by the stochastic differential equation

$$dp_i(t) = p_i(t) \left[(A\mathbf{p}(t))_i - \mathbf{p}(t)^T A\mathbf{p}(t) \right] dt + \omega \Gamma(\mathbf{p}(t)) dB(t).$$

In the above equation $B(t)$ represents a perfectly random perturbation or white noise process. $B(t)$ is normally distributed with mean zero and unit rate covariance matrix. In Foster and Young's most general formulation, the covariance matrix, Γ , is dependent on the state of the dynamical system and bounded away from zero. Also, in order to assure the process stays within the unit simplex, we assume the boundaries are reflecting in the sense described by Karlin and Taylor (1975). Foster and Young analyze the asymptotic behavior of this stochastic dynamic process as $\omega \rightarrow 0$. They find the dynamical system will select among the different metagame strategies as the noise term gets arbitrarily small.

¹⁷ See Nachbar (1988).

For analytic tractability, we assume $\Gamma(p) = 1$ in this simple example. Then, if we abuse notation slightly and denote the proportion of the population playing the cooperative strategy at time t as $p(t)$, we have

$$dp(t) = p(t) \left[30p(t) + 9(1 - p(t)) - [30p(t)^2 + 9p(t)(1 - p(t)) + 14p(t)(1 - p(t)) + 10(1 - p(t))^2] \right] + \omega dB(t)$$

$$dp(t) = p(t) \left[-17p(t)^2 + 18p(t) - 1 \right] dt + \omega dB(t).$$

Foster and Young use a result from Freidlin and Wenzel (1984) to show minimizing an appropriately derived potential function yields stochastically stable solutions, and at least one of these always exists. To avoid unnecessary technicalities, I will loosely define the stochastically stable set as the smallest closed set such that as $t \rightarrow \infty$ and $\omega \rightarrow 0$, it is almost certain the system is arbitrarily close to the stochastically stable set regardless of the initial state.

For the case at hand the appropriate potential function is

$$V(p) = - \int_0^p x \left[(A \cdot (x, 1 - x)^T)_1 - (x, 1 - x) \cdot A \cdot (x, 1 - x)^T \right] dx$$

$$= \int_0^p x(17x^2 - 18x + 1) dx.$$

The above potential function has a unique global minimum at $p = 1$. The point $p = 0$ (all players playing "always defect") is a local minimum, and the mixed strategy equilibrium corresponds to a local maximum.

It is easy to see both $p = 1$ (all GRIM) and $p = 0$ are locally stable. However, $p = 1$ is a deeper trough. It is not the steepness of the potential function near the trough which matters in determining stochastic stability but its depth.¹⁸ Foster and Young indicate as $\omega \rightarrow 0$ the dynamic process will spend increasingly more

¹⁸ See Freidlin and Wenzel (1984).

time near $p = 1$, and the population will almost surely be all cooperators in the limit.

Foster and Young proved two other important properties of stochastically stable sets. First, they always exist. If we use stochastic stability as a solution concept, we always have a nonempty solution set, even if there is no limit ESS. Second, even if a game has a unique ESS, it needn't coincide with the stochastically stable set. It is clear, then, stochastic stability is not a refinement of the ESS concept.

If we consider the two games we discussed in the previous section, we can see they both have nonempty, stochastically stable strategy sets. The stochastically stable set of the rock-scissors-paper game is the set consisting of the three corners in the unit simplex in R^3 . This means as $\omega \rightarrow 0$, the probability the state of the system will be arbitrarily close to one of the corners is one.¹⁹ The stochastically stable set in the common interest game is the ESS (1, 0). This solution is intuitively plausible and yields the payoff dominant equilibrium outcome.

This discussion leaves us wondering under what conditions in the simple infinite RPD will cooperation emerge as a stochastically stable solution. In other words, will cooperation always emerge if we have only the two simple strategies available and they are both ESSs? We saw in this very simple framework with a common discount factor of .9 the population will tend toward the cooperative strategy. Is this generally true?

The answers to the above questions are completely intuitive. If the discount factor is sufficiently small, the stochastically stable solution will have defect always as the only strategy in the stochastically stable set. To find the critical discount factor (the value for δ where "always defect" becomes a stochastically stable solution) we must look at a more general payoff matrix. If the infinite RPD has the payoff structure I described earlier and a common discount factor δ , the payoff matrix

¹⁹ For a more detailed discussion see Foster and Young (1988).

from using the GRIM and DD strategies is the following:

$$A = \begin{pmatrix} \frac{R}{1-\delta} & S + \frac{\delta P}{1-\delta} \\ T + \frac{\delta P}{1-\delta} & \frac{P}{1-\delta} \end{pmatrix}$$

After some algebra, we can see if $\delta \leq 1 - (R - P)/(T - S)$ "always defect" will be in the stochastically stable set. If the inequality is strict, it will be the unique element of the set. For this to be of interest it must be the case $\delta \geq (T - R)/(T - P)$ so there are two ESSs in this simple game. In the example above where $T = 5$, $R = 3$, $P = 1$, and $S = 0$ we find "always defect" is stochastically stable if $\delta \leq 3/5$.

We have shown in a very simple environment either defection or cooperation can be stable in a stochastic setting. If the discount factor, or shadow of the future as it is described in Axelrod and Dion (1988), is sufficiently large, cooperation should emerge. However, there are at least two problems standing in the way of this being particularly interesting. First, the stochastic effects of any reasonable mutation scheme will almost certainly have a covariance matrix which is meaningfully state dependent. For example, if we think of these strategies mutating from one to the other at different rates, the variance will depend on the proportions of the various strategies in the population. Second, we are looking at only two of an infinite number of possible strategies. In the next section, I will address these problems. Specifically, I will simulate an evolutionary dynamic process with mutation. Then I will allow a simple mutation scheme and simulate the evolutionary process with multiple mutants.


Simulations of Stochastic Evolutionary Processes

There are two reasons for performing simulations instead of examining this problem analytically. The first is tractability. Any sensible mutation scheme will prove to be analytically intractable. The covariance matrix which would reasonably represent the dynamic process I have in mind would make solving the problem

extremely difficult, if it could be solved for a closed-form solution at all. The second problem is adding more strategies further increases the analytic complexity. To make the exposition as clear as possible, I will explore the different concepts of stability I discussed earlier using computer simulations which shed some light on the problem. These simulations will also help us answer the questions with which we ended the last section.

I begin by discussing the mutation scheme I have in mind. The two strategies I discussed in the simple version of the infinite RPD, the GRIM strategy and "always defect," can easily be represented by finite automata, or Moore machines. A Moore machine is described by a four-tuple $\langle Q, q_0, \lambda, \mu \rangle$, where $Q = \{q_0, q_1, q_2, \dots, q_m\}$ is a finite set of states, $q_0 \in Q$ is the initial state, $\lambda : Q \rightarrow \{C, D\}$ is the output function which maps the state into a strategic choice, and $\mu : Q \times \{C, D\} \rightarrow Q$ is the transition function which maps the current state and the opponent's choice into a state (not necessarily different). Moore machines not only allow us to model strategies in repeated games concisely and conveniently, they let us analytically calculate payoffs for infinite games. When two finite automata play each other, they must eventually enter a cycle. Hence, the sequence of payoffs can be represented by an infinite sequence which can be summed if we discount the stage game payoffs. Figure 2.4 provides representations of the two automata which implement the two metagame strategies I described in the last section. Here the circles represent the different states, and the letter inside the circle indicates the action taken in that state. The initial state is the one on the left, and the arrows indicate transitions. For example, GRIM transitions from the cooperating state to the defecting state if the opponent defects. Otherwise, it stays in the cooperating state.

Figure 2.4 — Moore Machines for GRIM and ALWAYS DEFECT



The mutation scheme I use in the simulations is based on the fact that the two automata above have two easily distinguishable parts. They both have the DD strategy in common because the second state of GRIM is simply DD. If we think of these strategies as being made up of two fundamental building blocks, we can visualize these strategies as changing readily, one into the other. One building block is a state which cooperates as long as the opponent cooperates but moves to the next state if the opponent defects, and the other is an absorbing state which defects regardless of what his opponent does. I assume it is just as easy to gain a part of a strategy as it is to lose one. I begin by assuming a strategy can have no more than two states. I hypothesize a mutation rate of μ . This means with probability μ a GRIM strategy will change into DD, or a DD strategy will mutate to GRIM.

Using these two strategies, I simulated one thousand generations of the evolutionary process used by Axelrod in his tournament²⁰ with mutation. The simulation was done in discrete time with a population size of one hundred.²¹ The procedure I used simulates random matching with no memory. Each player was assigned a fitness based on how well he did against the other strategy and other players like himself. The proportion in the next generation, before mutation, is the proportion

²⁰ For details, see Axelrod (1982).

²¹ The simulations were accomplished on a Zenith 248 computer with programs written in Turbo Pascal. Some algorithms were taken from John Miller's (1990) program ECOSIM which graphically analyzes ecological dynamics.

in the current generation times the ratio of the individual strategy's fitness to the average fitness in the population. After the evolutionary process takes place, a strategy changes into the other with probability μ . The initial population for all simulations was half GRIM and half DD. However, regardless of the initial distribution, the same qualitative results are observed.

The simulations support the hypothesis that when the discount factor is small (below .6) DD will be the result of the evolutionary dynamic process as the mutation rate, μ , gets smaller, and if μ is large, GRIM will be the evolutionary result. Figures 2.5 - 2.10 reveal as μ decreases the populations select one strategy, DD for $\delta = .55$ and GRIM for $\delta = .7$. This is more clear when we simulate the dynamic process with a mutation rate of .01.




Figure 2.5 — Evolution with Mutant Strategies

$$\delta = .7, \mu = .2$$

Figure 2.6 — Evolution with Mutant Strategies

$$\delta = .7, \mu = .1$$





Figure 2.7 — Evolution with Mutant Strategies

$$\delta = .7, \mu = .01$$

Figure 2.8 — Evolution with Mutant Strategies

$$\delta = .55, \mu = .2$$






Figure 2.9 — Evolution with Mutant Strategies

$$\delta = .55, \mu = .1$$

Figure 2.10 — Evolution with Mutant Strategies

$$\delta = .55, \mu = .01$$



These simulations suggest Foster and Young's notion of stochastic stability holds in this dynamic process. Although these simulations were accomplished in discrete time, these results suggest the validity of their model for continuous time processes. Although there is only one mutant in each generation on average when $\mu = .01$, the population selects one equilibrium outcome. This indicates the importance of repeated perturbations in this model. A small perturbation to the system would not cause the population to move away from either ESS, but the repeated perturbations cause such a movement to take place.

The intuition is not difficult to see. As the discount factor decreases, DD becomes a less inferior reply to GRIM. That is, the difference between what a GRIM player gets against himself and what a DD player gets against GRIM diminishes. Therefore, as δ gets smaller a DD player will persist in the population long enough for another perturbation to take place. On the other hand, a GRIM player doesn't

do as well relative to a DD player because the future is less important. If the discount factor is near one, though, the difference in payoffs to GRIM and DD players when they play a GRIM player is large enough so defectors die out quickly if almost all the population plays GRIM. If most of the population defects, a GRIM player can persevere until the next perturbation because the future isn't discounted much. Also, if other cooperators appear, the GRIM strategy does a much better against them than DD does, so it can increase its proportion of the population.

Now we will examine what happens in a population when we add feasible mutant strategies. The idea of feasible mutants is important here. As I pointed out earlier, no strategy is immune from invasion in the infinite RPD. However, we should have some idea of what feasible strategies are in a given environment. Charles Darwin (1859) first suggested invading mutants should be readily derived from the native population. Boyd and Lorberbaum (1987) suggest a way to destroy TIT-FOR-TAT as an ESS in this game which relies on invasion by the right proportions of Suspicious TFT (STFT)²² and TIT-FOR-TWO-TATS (TF2T).²³ However, if those two strategies are not possible in the environment, we should consider the strategy stable against invasion.

Gregory Pollock (1989) proposed a heuristic for deciding what strategies are possible in a given environment. Pollock analyzed the playing of the infinite RPD in a viscous lattice where players interact only with their neighbors. He chose to describe strategies in terms of niceness and provokability. For example, GRIM is nice and maximally provokable, and DD is not nice and maximally provokable. He allowed the properties to be altered to introduce new strategies. The way I chose to model the mutant strategies in the simulations is similar in spirit to Pollock's mutation heuristic. Specifically, I allow possible mutants to be constructed of the

²² Suspicious TIT-FOR-TAT is a strategy which begins with defection, and then plays TFT.

²³ TIT-FOR-TWO-TATS is the strategy which cooperates in the first two periods, and then defects if the opponent defects twice in a row. It is the same as TFT except it requires two defections in a row to trigger retaliation.

two building blocks I discussed earlier. The possible mutants are, in addition to GRIM and DD, what I call 2-GRIM, 3-GRIM, ..., n -GRIM. That is, a strategy will transition to the absorbing defection state after n defections by his opponent. However, it should not be plausible to have 5-GRIM mutate to DD as easily as GRIM does. I assume the probability a building block is added or dropped is μ . Therefore, the probability an n -GRIM machine mutates into a DD machine is μ^n . More generally, an m -GRIM machine will change into an n -GRIM machine with the probability $\mu^{|m-n|}$.

We can describe the probabilities of mutation with a matrix M . An element m_{ij} is the probability a j machine changes into an i machine in one period. I chose to use six strategies in the simulations, so we have the following mutation matrix where α_j is defined as $\sum_{i \neq j} m_{ij}$, or the sum of the off-diagonal elements in the j th column.

$$M = \begin{pmatrix} 1 - \alpha_1 & \mu & \mu^2 & \mu^3 & \mu^4 & \mu^5 \\ \mu & 1 - \alpha_2 & \mu & \mu^2 & \mu^3 & \mu^4 \\ \mu^2 & \mu & 1 - \alpha_3 & \mu & \mu^2 & \mu^3 \\ \mu^3 & \mu^2 & \mu & 1 - \alpha_4 & \mu & \mu^2 \\ \mu^4 & \mu^3 & \mu^2 & \mu & 1 - \alpha_5 & \mu \\ \mu^5 & \mu^4 & \mu^3 & \mu^2 & \mu & 1 - \alpha_6 \end{pmatrix}$$

The results of the simulations are summarized in figures 2.11 - 2.12. The graphs indicate the possibility of mutation makes cooperation unsustainable in this stochastic environment. This is true even if we use a discount factor in the range so DD will not be stochastically stable. We can see even this simple mutation scheme can destroy GRIM as a unique stochastically stable strategy. If the six mutant strategies I chose didn't eliminate the cooperative ESS, I could have allowed more strategies. It is easy to imagine strategies mutating into increasingly less provokable strategies. The results of the simulation suggest for any discount factor less than one, some finitely less provokable strategy will allow DD to invade the population.

Figure 2.11 — Evolution with Different Mutation Rates

$$\delta = .7, \mu = .1$$





Figure 2.12 — Evolution with Different Mutation Rates

$$\delta = .7, \mu = .01$$

These simulations suggest the population will tend to cycle over time between cooperative and defecting outcomes if δ is large enough. The same forces which cause the population to select GRIM in the two strategy evolutionary process will cause the population to choose GRIM initially when we allow more mutant strategies. However, over time the less provokable strategies drift into the population. This is because they do as well against the GRIM strategy as it does against itself. As the population becomes less provokable, it becomes susceptible to being invaded by DD. After the population is overtaken by DD the same forces which worked initially are present again. Hence, the population will cycle over time. This argument can be seen in figure 2.12. We can see the cooperating strategy dominated the population for a while, and then the less provokable strategies allowed DD to invade again, only to be overcome by the GRIM strategy. If δ is small enough, there are no forces moving the population toward cooperation.

These graphs suggest the validity of Foster and Young's ideas in this process. The mutation scheme is an important characteristic of this model. If we allow the population to mutate at a constant rate among all strategies, Foster and Young's theory indicates eventually these less provokable strategies would make DD the unique stochastically stable strategy.²⁴ The correctness of this argument is suggested in figures 2.13 and 2.14. These figures show the results of two simulations of five thousand generations each, starting from a distribution with equal proportions of DD and GRIM. The mutation scheme used for these simulations puts the same probability on mutating from one strategy to any of the others. I used mutation rates of 0.02 and 0.002 to keep these simulations comparable to those of figures 2.11 and 2.12. Because of differences in the mutation schemes, mutation rates which are about one-fifth of those used in the earlier simulations were needed to maintain the expected number of mutations per generation approximately constant for the

²⁴ See Foster and Young (1988) for an example of this.

different simulations.

The difference between the results from these mutation schemes is obvious. If we have constant mutation rates across strategies, only DD will be stochastically stable, and cycling will not take place. Both mutation schemes appear to eliminate GRIM as the unique stochastically stable strategy. If mutation rates are constant across strategies, DD becomes the unique stochastically stable strategy. Foster and Young showed this, and it can be seen in the results of the simulations I accomplished. However, if mutations occur in the way suggested by the mutation scheme with differential rates, it appears the population will cycle between being nearly all DD and almost all cooperative strategies.

Figure 2.13 — Evolution with Constant Mutation Rates

$$\delta = .7, \mu = .02$$


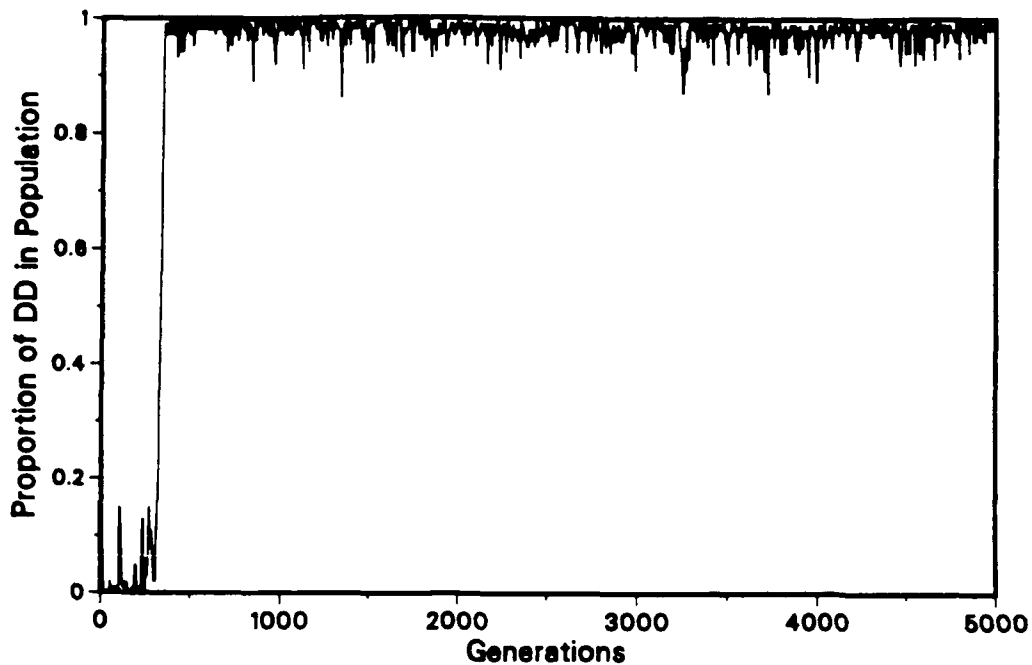


Figure 2.14 — Evolution with Constant Mutation Rates

$$\delta = .7, \mu = .002$$



These results are caused by the perturbations adding upon one another. If all of the strategies were available and we started with a population of all GRIM players, one small disturbance would not change the equilibrium play of the game. The population could change in composition, but a small perturbation would not cause the population to change to a defecting one. The cycling is the result of repeated perturbations.

Now consider Axelrod's simulation. It is obvious adding a single defecting strategy would not have affected the outcome in any significant way. However, if we allow a mutation scheme similar to the one I use here, these simulations suggest the outcome of an infinitely long evolutionary process may have ended up differently.

Summary

The simulations I performed and report here suggest cooperation may not be the outcome of an evolutionary process subject to continuous perturbations. In this essay I discussed the notion of evolutionary stability in an environment subject to disturbances. Much work has been done on the subject, and it both supports and denies the possibility cooperative outcomes can survive evolutionary processes.

As I pointed out, it has been proved no ESS exists in the infinite RPD. However, a "Folk Theorem" for the infinite RPD indicates an infinite number of limit ESSs are possible. The limit ESS is an embellishment on the ESS concept which allows us to hypothesize a player may make a tremble in the play of the game. This is not unlike Selten's trembling hand perfect equilibrium refinement to the Nash equilibrium.²⁵

Foster and Young have examined the stability of populations in an environment where the dynamical process is continuously subjected to perturbations. They show the population will generally select a set of strategies and will almost surely be arbitrarily close to that set as the variance of the disturbance goes to zero. The computation of stochastically stable sets is extremely cumbersome with distribution functions which are meaningfully state dependent or have a large number of possible strategies. Therefore, I performed a number of simulations using a simple mutation scheme.

The principal result which comes out of these simulations is that if we employ a reasonable mutation scheme in this particular evolutionary process, the population may not tend monotonically toward cooperation. In fact, in these simulations the population cycled between being cooperative and defective.

I chose GRIM and DD as the initial strategies because they are both subgame perfect in the infinite RPD. However, if we think of them as being implemented

²⁵ For an excellent comparison, see Samuelson (1989).

by Moore machines and imagine mutation as adding and subtracting parts of those machines, we can see neither defection nor cooperation can survive indefinitely if the shadow of the future is large enough. If the discount factor is small enough, though, DD will be stochastically stable, even in the presence of these mutants.

Bibliography

- Abreu, Dilip and Ariel Rubinstein. "The Structure of Nash Equilibrium in Repeated Games with Finite Automata." *Econometrica* 56 (November 1988):1259-1281.
- Auman, Robert J. "Survey of Repeated Games." *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*. Bibliographisches Institute:Manheim, 1981, pp.11-42.
- Axelrod, Robert. *The Evolution of Cooperation*. New York: Basic Books, 1984.
- Axelrod, Robert and Douglas Dion. "The Further Evolution of Cooperation." *Science* 242 (December 9, 1988):1385-1390.
- Binmore, Kenneth G. "Modeling Rational Players, Part I." *Economics and Philosophy* 3 (1987):179-214.
- Binmore, Kenneth G. and Larry Samuelson. "Notes on Evolutionarily Stable Strategies in Repeated Games Played by Finite Automata." unpublished manuscript, 1989.
- Boyd, Robert and Jeffrey Lorberbaum. "No Pure Strategy is Evolutionarily Stable in the Repeated Prisoner's Dilemma Game," *Nature* 327 (1987):58-59.
- Darwin, Charles. *On the Origin of Species*. London: John Murray, 1859.
- Farrell, Joseph and Roger Ware. "Evolutionary Stability in the Repeated Prisoners' Dilemma," *Theoretical Population Biology* 36 (1988):161-166.
- Foster, Dean and Peyton Young. "Stochastic Evolutionary Game Dynamics." unpublished manuscript, September 20, 1988.
- Fudenberg, Drew and Eric Maskin. "Evolution and Cooperation in Noisy Repeated Games," Harvard Institute for Economic Research Discussion Paper 1469, 1990.
- Harsanyi, John. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge: Cambridge University Press, 1977.
- Harsanyi, John and Reinhard Selten. *A General Theory of Equilibrium Selection in Games*, Cambridge: MIT Press, 1989.
- Hines, W. G. S. "Strategy Stability in Complex Populations," *Journal of Applied Probability* 19 (1982):19:204-205.

- Freidlin, M. and A. Wenzel. *Random Perturbations of Dynamical Systems*, New York: Academic Press, 1975.
- Kim, Yong-Gwan. "Evolutionarily Stable Strategies in the Repeated Prisoner's Dilemma," unpublished manuscript, October 1989.
- Maynard Smith, John. *Evolution and the Theory of Games*. Cambridge: Cambridge University Press, 1982.
- Maynard Smith, John and G. R. Price. "The Logic of Animal Conflict," *Nature* 246 (1973):15-18.
- Miller, John. "The Evolution of Automata in the Repeated Prisoner's Dilemma," unpublished manuscript, 1987.
- Nachbar, John. "'Evolutionary' Selection Dynamics in Games: Convergence and Limit Properties." unpublished manuscript, 1989.
- Pollock, Gregory. "Evolutionary Stability of Reciprocity in a Viscous Lattice," unpublished manuscript, 1989.
- Samuelson, Larry. "Evolutionary Stability in Asymmetric Games," Pennsylvania State University Department of Economics Working Paper, 1989.
- Selten, Reinhard. "Evolutionary Stability in Extensive Two-Person Games," *Mathematical Social Sciences* 5 (1983): 269-363.

CHAPTER III

EVOLUTIONARY STABILITY IN THE REPEATED PRISONERS' DILEMMA: AN EMPIRICAL ANALYSIS

Introduction

A large literature which attempts to explain how cooperative outcomes can be supported in the repeated Prisoners' Dilemma (RPD) has emerged in recent years. An interesting branch of this literature has analyzed the infinite RPD as a game in which two "metaplayers" choose a strategy which is implemented by a finite automaton, or Moore machine. This line of research has been especially interesting because it allows us to capture the notion of "bounded" rationality in the players. Also, the use of these machines as devices to implement strategies permits us to quantify the notion of strategic complexity. We can then apply the idea that complexity is costly to the analysis, and the results have proven very interesting.

So far, the work in this area has applied the concepts of Nash equilibrium and evolutionary stability in determining what reasonable outcomes are in these games. However, it has frequently been pointed out these equilibrium concepts may be either too restrictive or not restrictive enough to be useful in analyzing the infinitely repeated Prisoners' Dilemma. That is, if we use the Nash equilibrium as the appropriate equilibrium concept, the set of equilibrium outcomes is infinitely

large. On the other hand, if we require the equilibrium to be evolutionarily stable, the set of such equilibrium outcomes is frequently empty. It has been suggested the requirement for evolutionary stability may be too stringent a criterion.

The purpose of this essay is to examine what happens in these games in an evolutionary framework under various conditions. The results of the simulations I performed and report here indicate certain combinations of strategies may exist which, although not evolutionarily stable in the sense of Maynard Smith (1982), prove to be invasion-proof against permissible mutant strategies. This is akin to saying a set of possible strategies may exist which cannot be invaded, but neither the individual strategies nor the mixed strategy represented by the population is evolutionarily stable. We can imagine the mix of strategies changing as various mutants attempt to invade the population but the set of strategies remaining the same as the one with which we started. In fact, it may well be the case that this set of strategies will not last forever, but it may last for a very long time. I formalize this idea later in the essay. In order to test the robustness of the hypothesis, a number of different mutation schemes were simulated to examine what sort of outcome we should expect to find in a population similar to the one I explore.

I begin by reviewing the applicable literature and discussing the philosophy underlying this exercise. Then I describe the set of automata I consider as well as the various evolutionary models I simulated. As always, the results we obtain depend upon the assumptions of the model. The assumptions in this case include the set of permissible strategy implementing machines and the rules for how these strategies evolve over time. This line of research differs from other attempts to capture the results of an evolutionary story in two important ways. First, since the strategies I consider are all those which can be implemented by an automaton with not more than two states, there is no predisposition, either intentional or unintentional, toward a given outcome. The claim can be made the strategies which

were entered in Robert Axelrod's (1984) now famous Prisoners' Dilemma tournament were in some sense biased toward the TIT-FOR-TAT (TFT) type strategies because Axelrod identified TFT as the most successful strategy in a preliminary round of the tournament. The simulations I performed are not predisposed toward a particular outcome except through the evolutionary processes I use, and these are made explicit so we can either accept them or reject them on more reasonable grounds. Second, I implement different schemes for mutation which seem plausible in a sense I will explain later. Using Axelrod's tournament as an example again, we can imagine mutation entering his dynamic ecological process. Such mutant strategies could easily effect the final population in significant ways. However, what a reasonable mutant looks like in Axelrod's framework is not clear. In this essay, I formalize how the strategies are modeled so we can then evaluate the results based on the underlying assumptions.

Cooperation and the Infinite RPD

The Prisoners' Dilemma is a well known bimatrix game which has been used to study such economic problems as oligopolistic collusion, international trade, and public goods provision. The players in the game must simultaneously choose between "Cooperate" (C) and "Defect" (D). The payoffs I will use in this analysis are described in figure 3.1. Regardless of what one's opponent does, a player does better by defecting. This means (D,D) is the only Nash equilibrium in the one-shot game. The dilemma is that if the players could be induced to cooperate they would both do better than if they receive the equilibrium payoff. In fact, any finitely repeated version of the Prisoners' Dilemma has defection at every stage as the only perfect equilibrium outcome. It is not until we consider the infinitely repeated game that cooperation can be sustained. The "Folk Theorem" of repeated games tells us any individually rational payoff vector can be the outcome of a perfect equilibrium of

an infinitely repeated game when the payoffs are calculated as “the limit of the mean.”¹

Figure 3.1 — The Prisoners’ Dilemma

		Player II	
		C	D
Player I	C	3, 3	0, 4
	D	4, 0	1, 1

John Maynard Smith and G. R. Price (1973) are among those credited with the introduction of evolutionary game theory, but perhaps the most well known of the theoretical works on the subject is Maynard Smith’s 1982 book *Evolution and the Theory of Games*. Axelrod’s (1984) analysis of the success of cooperation in the repeated Prisoner’s Dilemma (RPD) is probably the most famous of many studies of the game in an evolutionary framework. Evolutionary game theory has frequently been used to study *coordination games*, *common interest games*, and the Prisoners’ Dilemma. Since its introduction many people in a broad range of disciplines have applied it to the study of natural and social science problems. Here I provide a brief review of some of those studies, especially those analyses which deal with the Prisoner’s Dilemma.

Robert Axelrod, in his 1984 book *The Evolution of Cooperation*, suggested cooperation is evolutionarily stable in the RPD. However, the definition of evolutionary stability he employed differs from that used by Maynard Smith (1982). Specifically, Axelrod applied a concept he called “collective stability” which requires only a strategy be a Nash equilibrium when it plays itself. In choosing this definition, Axelrod disallowed the possibility an alternate best reply strategy, which earns an equal payoff against both the indigenous and invading strategies, would be able to successfully infiltrate the native population.

¹ See Robert Aumann (1981) for a proof and further discussion.

This idea is embodied in the formal definition of an evolutionarily stable strategy (ESS). A strategy σ , which can be either a pure or mixed strategy, is an ESS of a symmetric game if and only if:

$$u(\sigma, \sigma) \geq u(\sigma', \sigma)$$

and

$$u(\sigma, \sigma) = u(\sigma', \sigma) \implies u(\sigma, \sigma') > u(\sigma', \sigma').$$

for all $\sigma' \neq \sigma$, where $u(\sigma_1, \sigma_2)$ represents the payoff to a player if he plays strategy σ_1 and his opponent chooses strategy σ_2 . In words, these conditions require an ESS be a best reply to itself, and if there are alternate best replies, the ESS must do strictly better against an alternate best reply than the alternate best reply does against itself. The first requirement above means any ESS is a Nash equilibrium. However, the converse is not true because of the second condition, which is often called the stability condition. Axelrod eliminated this stability condition as part of his definition of evolutionary stability.² Therefore, it is in this weaker sense TFT is stable. There were a number of possible strategies which met this requirement but did not do well in Axelrod's tournament.

The above definition of an ESS is actually a characterization of a more basic requirement which holds in pairwise random matching models.³ Assuming von Neumann-Morgenstern utility functions, the following must hold for some sufficiently small proportion of invaders, $\epsilon > 0$:

$$(1 - \epsilon)u(\sigma, \sigma) + \epsilon u(\sigma, \sigma') > (1 - \epsilon)u(\sigma', \sigma) + \epsilon u(\sigma', \sigma').$$

This requires the expected payoff from playing σ in a population with proportion ϵ of σ' is strictly greater than the expected utility of those playing σ' . Loosely speaking,

² See Axelrod (1984) p. 217.

³ See Maynard Smith (1982) for details.

then, an evolutionarily stable strategy (ESS) is one which cannot be invaded by an arbitrarily small proportion of potential invading strategies.

Robert Boyd and Jeffrey Lorberbaum (1987) have shown no pure strategy is evolutionarily stable in this more restricted sense in the infinite RPD. Specifically, they show a population of TFT players can be invaded by the appropriate mixture of Suspicious TIT-FOR-TAT, which is the strategy which defects on the first move and then reciprocates the opponent's move, and TIT-FOR-TWO-TATS (TF2T), which reciprocates defection only after two consecutive defections by the opponent. Moreover, they show if every strategy is possible through mutation, no pure strategy is immune to invasion by some mixture of strategies. However, they suggest every strategy will not be possible through mutation, so in reality cooperation based on reciprocity may indeed flourish. Joseph Farrell and Roger Ware (1989) extended the work of Boyd and Lorberbaum to show for any evolutionarily stable mixture of strategies in the infinite RPD, every finite history must occur with positive probability. The implication of this result is no finite mixture of strategies is evolutionarily stable in the infinite RPD. They use this result as evidence the ESS concept is too stringent a criterion to be useful.

More recently, Yong-Gwan Kim (1989) showed an ESS in the infinite RPD does not exist unless we allow perturbations. Then he applied Reinhard Selten's (1983) concept of "limit ESS." This concept expands the set of evolutionarily stable outcomes by allowing perturbations which may put some minimum probability on other strategies in the game. In this way, ties which arise may be broken to allow for more equilibria.⁴ Kim was able to prove a "Folk Theorem" of sorts which states we can find a limit ESS outcome which is arbitrarily close to any convex combination of the payoffs from two purely defecting strategies and two purely cooperating strategies.

⁴ The idea is similar to Selten's concept of perfect equilibrium. For an excellent discussion, see Samuelson (1989).

If we limit the set of possible strategies to those which can be modeled as finite automata, or Moore machines, we can include measures of complexity in the players' utility functions. This has led to interesting results and has only recently been used to make evolutionary arguments. It was Robert Aumann (1981) who first suggested modeling strategies as machines with a finite number of states as a way to handle bounded rationality. Abraham Neyman (1985) showed if we model strategies this way and limit the number of states exogenously, we can find equilibrium machines which cooperate at every stage of the game if the number of states is less than the number of iterations of the game. The intuition here is easy to see. Suppose the Prisoners' Dilemma will be repeated one hundred times. As long as the machines which implement the strategies have fewer than one hundred states, cooperation can be maintained by the GRIM strategy.⁵ In order to beat the GRIM strategy, an opponent must be able to identify the last stage of the game so it can defect. Since it requires one hundred states to count that high, any machine with fewer than one hundred states cannot improve on the payoff to playing GRIM. Roy Radner (1986) used a similar argument to show how cooperation can be maintained using strategies which are implemented by finite automata which are similarly bounded in size.

Dilip Abreu and Ariel Rubinstein (1988) have studied models where complexity, which in their model is defined to be the number of states in the automaton, is endogenously determined. The metagame strategies are implemented by finite automata, and the complexity of the machine enters the players' decisions by making the metagame payoffs depend positively on stage game payoffs and negatively on complexity. By this I mean if two machines yield the same stage game payoffs, the machine with fewer states yields higher metagame payoffs. One model they analyzed had complexity enter the metagame payoffs lexicographically. Abreu and

⁵ The GRIM strategy begins by cooperating, but if either player defects it defects forever.

Rubinstein were able to reduce the set of equilibrium outcomes to the rational payoff vectors on the main and alternate diagonals of the set of feasible outcomes which provide each player with more than his security level. The interpretation Abreu and Rubinstein put on their model is that of a rational decision maker who must choose someone to play the game for him. However, the player who implements the strategy can only carry out simple instructions. We can think of these decision rules as "rules of thumb" which evolve over time. Another way to think of this is to assign some evolutionary process the roll of the metaplayer. In other words, some evolutionary process selects the strategies. It is this interpretation which I will explore.

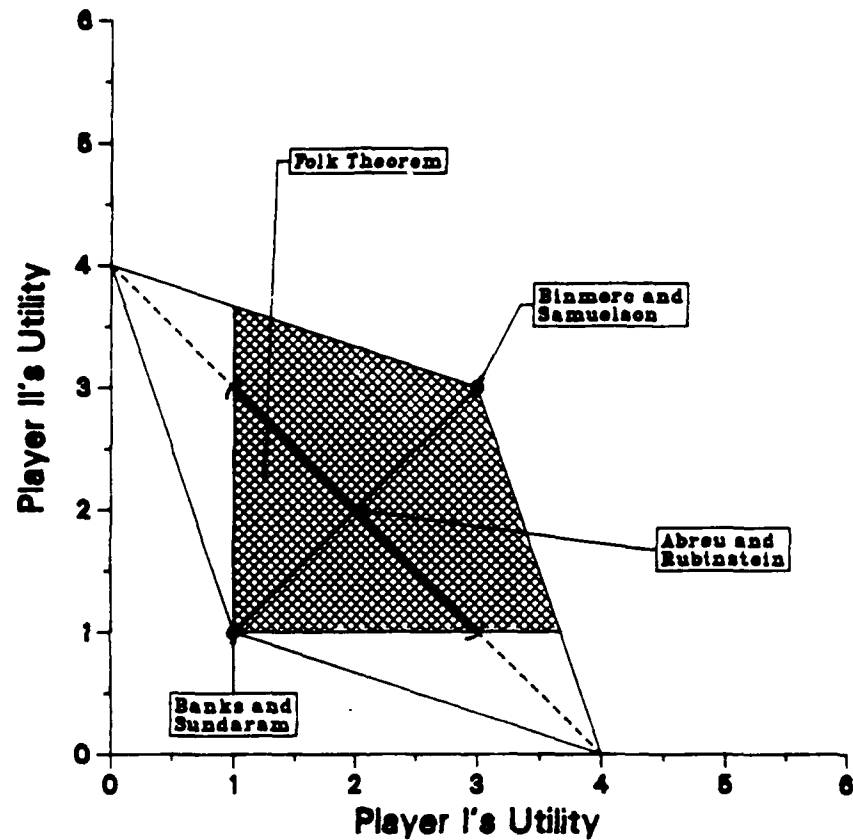
Ken Binmore and Larry Samuelson (1989) have done interesting work recently using the model developed by Abreu and Rubinstein. They show if we consider the same type of utility functions used by Abreu and Rubinstein, any evolutionarily stable strategy has both players earning the cooperative, or utilitarian,⁶ payoff. In other words, Abreu and Rubinstein refine the set of equilibrium payoffs by considering complexity of the implementing machine, and Binmore and Samuelson further refine this set to one outcome with an evolutionary stability argument. This result depends crucially on the definition of complexity we choose. Jeffrey Banks and Rangarajan Sundaram (1989) have shown if we use preferences which are lexicographic in complexity and use the number of transitions in the Moore machine as the measure of complexity, the only Nash equilibrium machine defects always.

These ideas are summarized in figure 3.2. The equilibrium payoffs allowed by the "Folk Theorem" are all those vectors in the shaded region. The equilibrium payoff set from the Abreu/Rubinstein model contains the rational points on the main and alternate diagonals which assure each player a payoff greater than one. Binmore and Samuelson's evolutionary model reduced the set to the unique point

⁶ By this they mean the sum of the payoffs is maximized.

(3,3), and the Banks/Sundaram model yields the one equilibrium outcome in which both players always defect and receive the payoff vector (1,1).

Figure 3.2 — Equilibrium Sets in the Infinite RPD



Selecting the “right” equilibrium is frequently the result of some *ad hoc* decision rule. For example, many economists and game theorists argue when multiple equilibria exist, the solution to the game should have the players achieving an efficient outcome. Using the terminology of Harsanyi (1977) and Harsanyi and Selten (1988), the players should choose this outcome based on payoff dominance⁷ if possible. This intuition is especially appealing for repeated games since there is recurrent interaction between the players. The argument for selecting payoff dominant equilibrium outcomes is especially strong when some sort of coordinating device is present. For

⁷ A payoff dominant equilibrium outcome is one in which both players obtain a higher payoff than they would get in any other equilibrium outcome. Obviously these do not always exist.

example, if we allow preplay communication between players, we should expect the equilibrium they agree upon will be payoff dominant if possible. Even in the absence of such a coordinating device, many game theorists and economists expect to see the payoff dominant equilibrium outcome. In the infinite RPD, of course, this means we should expect to see a cooperative outcome. However, if such selection criteria are to be meaningful, they should be based on first principles and not chosen because they give nice results.

This area of game theory has just recently become the target of significant research. The results so far have proven interesting but not conclusive. The work I described by Binmore and Samuelson is the first research in this area. They prove any symmetric equilibrium strategy which yields less than the cooperative payoff can be invaded by a strategy which does as well, in stage game payoffs, as a native does if it plays a native, but it does better against itself than it does against a native. *The invading machines are able to send a signal to their opponents.* If they are not playing their own type, it makes no difference in the metagame payoffs because payoffs are computed as the limit of the mean, and in the limit they get the same payoff as the native strategy. However, they earn greater average payoffs than the native strategy because they do better when playing strategies like themselves. The idea behind signalling to opponents is of ancient vintage. Arthur Robson (1989) speaks of a "secret handshake" which identifies members of a certain group to each other.

Drew Fudenberg and Eric Maskin (1990) have shown, independently of Binmore and Samuelson, cooperation in RPD games is evolutionarily stable in an environment where only strategies of finite complexity can be used, metagame payoffs are calculated as the limit of the mean and noise exists. In their model, the term "noise" refers to some chance an action may be misperceived by a player's opponent. John Miller (1987) performed interesting simulations of the evolution of automata which

were modeled as a string of digits. The impact of noise on the model was one of the factors he examined. He simulated the evolution of a population playing a finite RPD in which there were strictly positive probabilities of errors in the transmission of information about what a player's opponent did on the previous move.⁸ His results are difficult to characterize in a sentence or two; however, the methodology he used is interesting and the results indicate the effects of noise on a model like this are not negligible.

There is no definitive theoretical result explaining why a certain outcome should emerge from some evolutionary process. We are caught between arguments like those of Farrell and Ware, which indicate no strategy or mix of strategies is evolutionarily stable in the infinite RPD, and those which identify one particular strategy as the only reasonable result of an evolutionary process. With this in mind, I intend to explore the evolution of strategies which can be implemented by two state Moore Machines in the infinite RPD. Axelrod's work has been the standard in this area. However, as I pointed out earlier, his results may have been part of a self-fulfilling prophecy in which TFT was identified as the most successful strategy. Another problem with Axelrod's work is it is purely ecological in nature. That is, the next generation's population was determined only by the ecological dynamic process, and new strategies were not allowed to enter. In other words, the strategies were assumed to be perfect in their ability to breed true. In the next section, I will describe the basics of the simulations I performed. I describe how I modeled the strategies, how I chose to model possible mutations, and how I attempted to capture the idea of strategic complexity.

⁸ Specifically, he considered error rates of 1 percent and 5 percent.

The Model

In this essay, I consider only two-player games. The underlying game is the version of the Prisoners' Dilemma described in figure 3.1. More formally, I will consider the game G which is specified by the four-tuple $\langle S_1, S_2, \pi_1, \pi_2 \rangle$ where $S_i = \{C, D\}$ is the strategy set and $\pi_i : S_1 \times S_2 \rightarrow R$ is the payoff function. These elements are summarized in the bimatrix form of the game.

The infinitely repeated game $G^\infty = \langle R_1, R_2, P_1, P_2 \rangle$ is constructed with G as the stage game. Each player's strategy space, R_i , is the set of functions which map any history of play into S_i . That is, $R_i = \{f : h_t \rightarrow S_i\}$ where h_t is the history of the game up to and including time t . The payoff functions P_1 and P_2 are defined as the limit of the mean of the stage game payoffs. This is always defined because we are considering only strategies which can be implemented using finite automata. Since the machines must eventually enter a cycle, we know this limit exists. Therefore, we have

$$P_i(r_1, r_2) = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \pi_i(r_1(h_t), r_2(h_t))$$

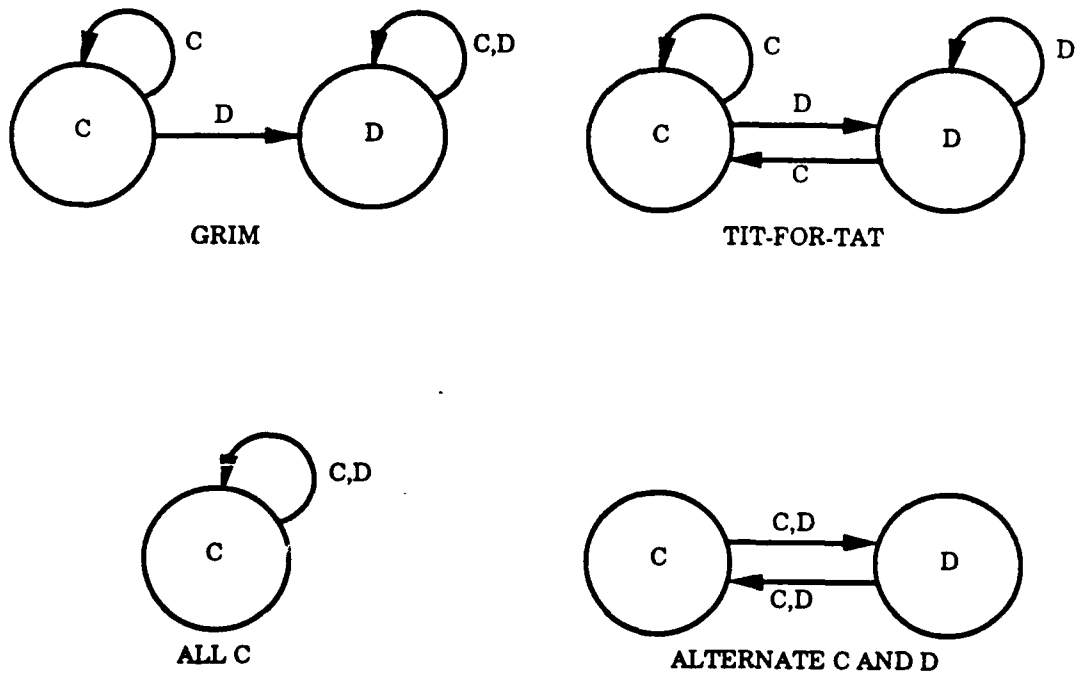
where $r_i \in R_i$.

I also use Abreu and Rubinstein's definition of an automaton selection game. Here $G^\#$ is the automaton selection game defined as $\langle A_1, A_2, U_1, U_2 \rangle$. The strategy space A_i is the set of Moore machines which I will describe more completely later. The functions U_i are the true utility, or profit, functions of the players. At various times during this essay their precise form will vary. However, these functions basically adjust the payoff functions P_i for complexity. In this automaton selection game we can think of automaton $a \in A_i$ as actually being Player i in G^∞ . In a sense, this automaton is Player i 's agent who follows the simple instructions represented by the Moore machine.

Now I will define more formally the kind of machine I have in mind as implementing these simple instructions. A Moore machine is described by a four-tuple

$\langle Q, q_0, \lambda, \mu \rangle$, where $Q = \{q_0, q_1, q_2, \dots, q_m\}$ is a finite set of states, $q_0 \in Q$ is the initial state, $\lambda : Q \rightarrow \{C, D\}$ is the output function which maps the state into a strategic choice, and $\mu : Q \times \{C, D\} \rightarrow Q$ is the transition function which maps the current state and the opponent's choice into a state (not necessarily different). Moore machines not only allow us to model strategies in repeated games concisely and conveniently, they let us analytically calculate payoffs for infinite games.

Figure 3.3 — Some Moore Machines




These machines can be represented quite easily as demonstrated in figure 3.3. Each circle in a machine represents a possible state of the machine, and the state on the left is the initial state. I represent the action taken when the machine is in that state by the letter inside the circle. The arrows indicate how the machines transition after a play of the stage game. The strategies which are "nice" by Axelrod's definition are those which begin by cooperating and continue cooperating as long as the opponent does. An example of a nice strategy is GRIM or TFT.

Determining a reasonable mutation scheme is not without problems. One alternative is to assume it is possible for one strategy to mutate to any other strategy with equal likelihood. I call this uniform mutation. However, this is not a very satisfying way to proceed. If we are thinking of these machines as decision rules evolving over time, we must develop some idea of "closeness" in the strategies and only allow strategies to mutate to those which are close. Another idea we should try to capture is that simplifying mutations are more likely than complicating mutations. I will use this idea later as a way to pick up the costs of complexity.

In order to capture these notions of "closeness" and simplification, I developed a mutation scheme based on how these strategies would appear if they were machines in the physical sense and subject to breakdowns. We can think of these machines as consisting of transmitters and the appropriate wires hooked up to receivers to control which signal is transmitted. The simplest machines are those which need only transmit either "cooperate" or "defect." These correspond to the single state machines which always defect or always cooperate. Two state machines are made up of the two transmitters which send the signals corresponding to cooperation and defection. I chose to model these strategies so the machines attempt to move to the next state, or change the signal it sends, after taking an action. However, the transition may be blocked for one of two reasons. First, there may be no circuit from one state to the next. In this case a machine will continually send the same response because it can reach no other state. Another way the transition can be inhibited is by perceiving an appropriate action from the opponent to prevent the transition. We can think of these inhibitors as switches which open to keep the state from changing. For example, consider the GRIM strategy which I have modeled in the machine in figure 3.4. In this case, the machine will begin cooperating and attempt to switch to the defecting state unless the machine receives the signal the opponent cooperated. If GRIM fails to detect cooperation by the opponent, it switches to the

defecting state where it defects forever because there is no wiring available to get back to the cooperating state.

Figure 3.4 — The Stylized GRIM Machine



This stylized model of how strategies are implemented offers an opportunity to consider what sensible mutations look like. For example, we can assume a reasonable mutation involves a wire in the machine breaking. We can imagine a strategy being misinterpreted by the player and some transition being missed. If the mutant machine continues to function as well or better in the environment than the machine from which it evolved, we can expect the mutant machine will be copied.

In this essay I will consider “breakdowns” of the machines, or possible mutations in which a wire breaks. Also, I assume it is possible for the signal to get reversed in the machine so the automaton sends C when it should send D and vice versa. This is not the only way to model these strategies as mechanical or electrical devices. However, this is one way which is relatively simple and yields reasonable mutation patterns. I will not examine increasing complexity through mutation. This model of the strategies allows us one opportunity to capture the cost of complexity. The more complex a machine is, the more likely it is to break down into other simpler machines. Therefore, if complexity adds nothing to the payoffs, it will not endure in the population.

Another way I try to capture the costs of complexity is to model parts of the machines as having some cost. This can be applied to the evolutionary dynamic process in one of two ways which closely follow the notions of complexity described by Abreu/Rubinstein and Banks/Sundaram. To capture the idea states are costly, I adjust the payoffs so the single state machines receive a premium in the play of the game. It turns out the size of the premium affects the result in an evolutionary framework. Also, I impose a cost for maintaining each of these stylized wires/transitions in the machines. Again, the payoff matrix is modified to reflect this idea.

The model of Abreu and Rubinstein does not lend itself to evolutionary simulations. Operationalizing the idea of lexicographic preferences is not straightforward in an evolutionary context. In the model I have in mind, the proportion of the population which plays a particular strategy at time $t + 1$ depends on how well that strategy does at time t relative to the average member of the population. Lexicographic preferences do not allow such simple averaging of payoffs, but we can handle the idea preferences depend on complexity by imposing costs on the metaplayers.

The evolutionary process I have in mind for the simulations has basically two parts. The first part attempts to capture the idea successful strategies or rules of thumb will tend to be copied. The second part reflects the idea that over time these strategies will not be perfectly duplicated, or will not breed true in biological terms. First I will describe the replication dynamic, and then I will describe how the strategies mutate.

To see how the replication dynamic process works, imagine a game in which there is a strategy space A with n strategies or automata, $A = \{a_1, a_2, \dots, a_n\}$. The RPD is then played at time t , and each strategy earns a payoff depending on the strategy with which it is matched. These payoffs can be represented in a $n \times n$ matrix B with elements b_{ij} being the payoff to a player who uses machine a_i if he

is matched against a player who uses a_j . Then we have

$$B = \begin{pmatrix} b_{11} & b_{12} & \dots & b_{1n} \\ b_{21} & b_{22} & \dots & b_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ b_{n1} & b_{n2} & \dots & b_{nn} \end{pmatrix}.$$

Also, let $\mathbf{p}(t) = (p_1(t), p_2(t), \dots, p_n(t))^T$ be the vector of proportions of each type of player in the population at time $t \in \{0, 1, 2, \dots\}$. Here T indicates the transpose of a matrix. We can then define the expected payoff to a player of strategy a_i at time t as $(B\mathbf{p}(t))_i$. This is the i th element of the vector $B\mathbf{p}(t)$. Also, the expected payoff for a member of the population at time t is $\mathbf{p}(t)^T B\mathbf{p}(t)$. The process we are considering has the proportion of strategy i at time $t + 1$ equal to its proportion at time t multiplied by the ratio of its own expected payoff to the expected payoff of all players. Then we have the following dynamic process:

$$p_i(t + 1) = p_i(t) \left(\frac{(B\mathbf{p}(t))_i}{\mathbf{p}(t)^T B\mathbf{p}(t)} \right).$$

It is easy to see whether a strategy's proportion in the population gets larger or smaller depends on whether or not it is doing better or worse than average. The proportion of a strategy in the next generation depends not only on its own relative performance, but also on its proportion in the current population. This captures the idea a strategy must be both successful and observed by other strategies to be copied.

The dynamic process described above is based on the idea some machines or rules of thumb would be so unsuccessful they would probably not be used in the future, while the superior machines would be imitated. We can also think of the payoffs from playing the RPD as fitness in the biological sense. That is, strategies with higher payoffs are able to reproduce (asexually) more successfully than strategies with lower payoffs. In either case, we would expect to see the best strategies flourish and the worst strategies die out. This process allows us

to evaluate how well a strategy will do when the less effective ones cease being important, and only the best strategies remain to play each other. The dynamic process I used simulates evolution with an infinitely large population and random matching. Essentially, when a player of a certain strategy type plays the game, he is playing against a mixed strategy which is represented by the different proportions in the population.


In order to capture the notion strategies mutate in the evolutionary process, I used a matrix M to represent the rate at which strategies change into each other. A typical element $m_{i,j}$ is the probability a strategy of type j mutates into a type i player. Therefore, after the reproduction dynamic process described above has taken place, the strategies undergo a mutation process which is described by M . Then, we can find the proportions of the population, $\mathbf{p}^*(t + 1)$, which play the infinite RPD in period $t + 1$ with the equation

$$\mathbf{p}^*(t + 1) = M\mathbf{p}(t + 1)$$

where $\mathbf{p}(t + 1)$ is the vector of proportions which are the result of the dynamic replication process described above.

Some very simple examples may help see what the mutation matrix does. For example, if we assume one strategy can change into any other strategy with equal probability, μ , then the matrix M has form $m_{i,j} = \mu$ for $i \neq j$ and $m_{i,j} = 1 - (n - 1)\mu$ if $i = j$ and there are n total strategies. On the other hand, if no mutation takes place the above is true with $\mu = 0$. The only restrictions we need to place on this matrix is the columns must add to one and each element must be in the interval $[0, 1]$. These restrictions are completely reasonable and merely require an individual either stay the same or change. We do not allow one machine of type a to turn into some other number of type b machines.

Figure 3.5 — Mutation of UTIL



With the stylized machines we can derive a mutation matrix. Rather than list the entire matrix, I will just show how one well known strategy mutates. Consider the strategy which implements the cooperative outcome in the Abreu and Rubinstein automaton selection game. A diagram of the stylized machine is depicted in figure 3.5. It is easy to see what happens as each of the imagined wires breaks. For example, if the wire breaks which attempts to change actions from D to C (this represents the transition from the defecting state to the cooperating state in the diagram of the Moore machine) the machine will only be able to send the D signal. If, however, the inhibitor wire on the top of the diagram were to malfunction, this machine would mutate into AC. If the signals got reversed, this CC machine would become cc. We need not go any further to see the CC machine can mutate into DD, AC, CA, CD, or cc if only one of the wires breaks. For simplicity I assume

only one wire can break in any generation. We can think of this as approximating independent mutations where the mutation rates are so small the probability of two mutations in a given generation is negligible.

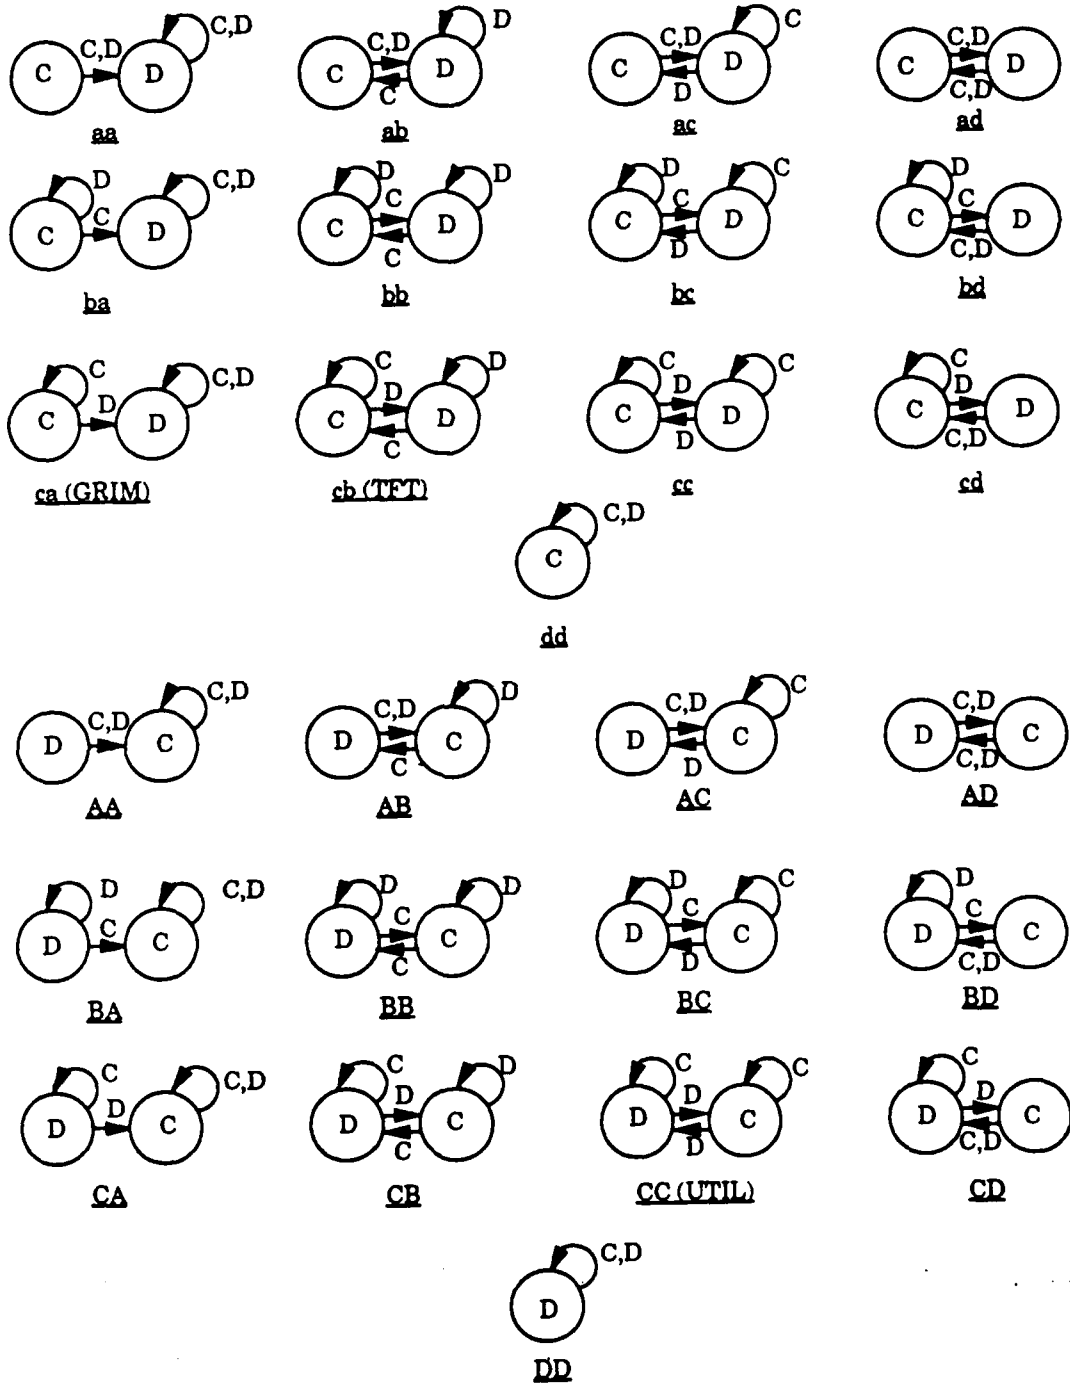
Finally, before actually discussing the results of the simulations, I will describe the twenty-six strategies I used. These strategies can be represented by all of the possible one and two state Moore machines. These simple strategies can be thought of as decision rules which are very easily implemented. It is not correct to say these machines are restricted to a memory of only one period. In a sense, the GRIM strategy has a memory of infinite length. However, the strategies I describe here are unable to allow for strategies which need to count to any number greater than one. It is interesting to note the most successful strategy in Axelrod's RPD, TFT, tournament is represented here. Nearly all the strategies which enjoyed some degree of success in the tournament used TFT as the main rule.⁹ All of the strategies are described in figure 3.6.

Eight of the possible twenty-six strategies form a Nash equilibrium when they play against an identical machine. These equilibrium strategies include all of the "nice" strategies except "always cooperate." Using the notation here, these are ca, cb, cc, and cd. In addition, the utility maximizing equilibrium strategy from the Abreu and Rubinstein model,¹⁰ CC, is an equilibrium strategy when it plays itself. The three other strategies which form equilibria when they play themselves are the "always defect" strategy (DD), aa, and AC. The aa machine cooperates on the first move and then begins defecting forever. The AC machine begins by defecting and then switches to a cooperating state where it stays until the opponent defects. An opponent's defection will move this machine back to the initial state. Once in the initial state again, it repeats the same sequence of play by defecting once and then moving to the cooperating state until it detects another defection.

⁹ For more discussion on Axelrod's tournament, see Linster (1990).

¹⁰ I follow Binmore and Samuelson and call this strategy "UTIL."

Figure 3.6 — Possible Automata



Simulation Results

In this section I will briefly describe the results of the simulations I performed. I elaborate on the results only briefly in this section. I save most of the discussion for the next section where I analyze which strategies do well under the various mutation schemes and draw conclusions about what characteristics the successful strategies have. These will not be in strict agreement with the characteristics identified by Axelrod. There are some circumstances in which "niceness"¹¹ appears to be a successful characteristic, but under other circumstances it may not be.

The simulations were accomplished on a Zenith 248 computer with programs written in Turbo Pascal. Some of the algorithms for the programs were taken from John Miller's program ECOLSIM which graphically analyzes ecological dynamics with mutation. In the evolutionary dynamic simulations, each machine was assigned a fitness based on how well it does against the other strategies and other machines like itself. The proportion in the next generation, before mutation, is the proportion in the current generation times the ratio of the individual strategy's fitness to the average fitness in the population. After the evolutionary process takes place, a strategy changes into another with a probability determined by the appropriate mutation scheme.

Evolution Without Mutation


I performed one hundred simulations of one thousand generations without any mutation as a control to determine which, if any, particular strategies are the likely results of a strictly ecological process. These simulations began from randomly selected starting points in the unit simplex in R^{26} . One thousand generations of the evolutionary dynamic process described above were then simulated. These simulations were repeated one hundred times. The significance of the ecological

¹¹ Axelrod (1984) referred to strategies which are not the first to defect as "nice."

simulations is that once a strategy nearly dies out, there are no perturbations or trembles to increase its proportions again. This ecological simulation is very similar to what Axelrod did with the results of his RPD tournament.

I graphically describe the average of the proportions of the population for the final generation in each of the simulations in figure 3.7. It is clear the cooperative outcomes proved the most evolutionarily fit in this environment. In fact, we can see the only strategies to survive this dynamic process are the "nice" ones, and all of them survived.¹² It is interesting that the strategy which is most exploitable in some sense, "always cooperate," survives in this ecological process. This is because it does well against those other strategies which do well. Specifically, this strategy does well against the other "nice" strategies, and those "nice" strategies drive the "mean" strategies to near extinction quickly enough so the nicest and infinitely forgiving strategy, ALL C, can survive.

Figure 3.7 — Evolutionary Success without Mutation




¹² By "surviving" I mean they were represented in the population by a proportion greater than 10^{-10} .

Evolution with Uniform Mutation

I then performed similar simulations using uniform mutation with the same one hundred random starting points I used in the previous simulation. I chose a mutation rate of 0.0001. I present the data in figure 3.8 in the same format I used for the evolution without mutation. The same set of strategies that did well in the last simulation are successful in this simulation.¹³ The two strategies which appear to do significantly differently are TFT and ALL C. TIT-FOR-TAT does not do as well in the presence of mutation, and ALL C performs better. The intuition for why TFT does not do as well seems relatively clear. TFT is not as aggressive in exploiting poor strategies as GRIM. Hence, when these poor strategies appear, GRIM does relatively better than TFT. Hence, TFT's proportion of the population diminishes over time.

Figure 3.8 — Evolutionary Success with Uniform Mutation

($\mu = 0.0001$)



¹³ Here I refer to those strategies which survived in a proportion at least three times the mutation rate as doing well.

Note in order for a strategy which is a large proportion of the population to keep from becoming smaller, it must do well relative to the others because of the mutation process. The larger the proportion of a given strategy there is in the population, the better that strategy must do to keep from shrinking in its proportion. If all strategies did equally well, this mutation scheme would equilibrate when all strategies are in equal proportion. However, the periodic introduction of poor strategies allows GRIM to do better than TFT. The improvement of ALL C seems to be the result of the mutation scheme. Although ALL C doesn't do very well against "mean" strategies, there are not many of them around. Since ALL C is a small proportion of the population, it is a net gainer when the mutation process works. Additionally, since most of the population is nice it does not do very badly in terms of payoffs, so over time it grows in size until it attains the levels found in this simulation. We can think of strategies like GRIM as "enforcers" in this game. That is, GRIM keeps the proportions of the "mean" strategies low so ALL C can survive.

Evolution with Stylized Mutation


The work of Abreu and Rubinstein hypothesizes less complex strategies are better than more complex strategies. The notion I attempt to capture in this simulation is that more complex strategies will tend to break down more often than less complex strategies. Hence, unless the extra complexity results in increased stage game payoffs, the more complex strategies will tend to be replaced by less complex strategies through evolution. We can think of these rules of thumb mutating to simpler rules. If making a rule less complex does not reduce its payoffs in the stage games, we should expect the simpler rule to be copied more often. Again, which rules are successful will depend on the environment.

I simulated one hundred evolutionary processes of one thousand generations once again, and I started at the same random starting points as in the previous

two simulations with the same mutation rate of 0.0001. I summarize the results in figure 3.9. Again, the GRIM strategy does much better than any of the others. The intuition is slightly different here. Not only does the GRIM strategy exploit irrational strategies, but it is one of the least complex two state machines in terms of the model I chose because it only needs to perceive one deviation and perform one transition. Hence, it will increase in proportion as a result of mutation from other successful strategies as well as its own strategic fitness.

Figure 3.9 — Evolution with Stylized Mutation

($\mu = 0.0001$)



Mutants Who Enter as Groups (Uniform Mutation)

Next, I attempted to capture the possibility strategies can grow in their own small communities and then attempt to invade the population all at once. The biological story which goes with this simulation has a certain group of animals

which are separated from the rest of their species when the land they are on is cut off from the main land mass by water. The animals live in isolation and develop their own characteristics. At some point, a land bridge forms, and this relatively large group of mutants enters the population. I attempt to capture the essence of this thought experiment by allowing one percent of the population to become a mutant type and attempt to invade the population. Then I allow the population dynamics to settle down again. Once the mutant strategy's influence is absorbed by the population, another band of mutants enters. This process continues.


I performed this simulation one time for thirty thousand generations. I chose to use the midpoint of the unit simplex in R^{26} as the initial distribution. That is, I began with every strategy as $1/26$ of the population. Also, I chose to regard a particular strategy as being extinct once its proportion of the population falls below 10^{-5} . This allows us to more easily identify when the population stabilizes. In this exercise, I allow all of the strategies to enter with equal probability. That is, I assume uniform mutation. I show the final population in figure 3.10. It is interesting to note after the population stabilizes from the initial starting point, it changes very little. Specifically, after initial stabilization and before mutants attempt to invade the population the nice strategies are in these proportions: ca, 0.575; cb, 0.152; cc, 0.164; cd, 0.105; and dd, 0.005. Here again, the "nice" strategies do very well.

We cannot say these strategies are immune from invasion in this model because there is a very small probability the "always cooperate" strategy will be the invading mutant for a sufficient number of consecutive generations so the population can be successfully invaded by "always defect." Since this probability, although extremely small, is greater than zero, the event will take place if the process continues long enough. However, we can say this set of strategies is nearly invasion-proof in the above sense. This idea is akin to the concept formalized by Binmore and Samuelson (1990) which they call a Payoff-equivalent, Polymorphous Modified ESS (PPMESS)

without the complexity criteria. The idea behind this concept is certain groups of strategies may exist which cannot be successfully invaded. The individual proportions of each strategy will change as various types of potential invaders appear, but the population cannot be invaded.

It is somewhat surprising the ALL C strategy survives in such a large proportion. The reason for this is clear. Although the ALL C strategy can be exploited, it does not exist in large enough proportions to allow a "mean" strategy to invade. It is not difficult to see if only the GRIM, ALL C, and ALL D strategies were possible, ALL D would not be able to invade a "nice" population until ALL C accounted for at least 2/3 of the population. However, this is extremely unlikely because as ALL D appears through mutation it keeps the proportion of ALL C low.


**Figure 3.10 — Evolution with Mutants Who Enter as Groups
(Uniform Mutation)**



Evolution with Mutants Who Enter as Groups (Stylized Mutation)

In this simulation I attempted to capture the same elements as in the last simulation except I use the mutation scheme from our stylized strategy implementing machines. In this example the strategies which attempt to invade in a group must come from the set of mutants which are possible from one of the strategies in the population. The probability a strategy mutates in this simulation is its proportion of the population. Then any possible mutant from that strategy is equally likely. Here we require those who are "stranded on the island" and then become the invading mutants be reasonably similar to the population from which they came. This simulation was performed like the previous one. The final population is described in figure 3.11.

**Figure 3.11 — Evolution with Mutants Who Enter as Groups
(Stylized Mutation)**



The success of GRIM and ALL C is evident here. These two strategies, which

appeared successful in the previous simulations, are the only two survivors of this process. The intuition here again relates to their relative simplicity. Here we seem to have captured the idea simple strategies will be more evolutionarily fit. Although GRIM is more complex than ALL C, the fact that ALL D enters periodically as a mutant invader keeps ALL C in small enough proportions so ALL D cannot successfully invade the population. The same sort of symbiotic relationship described by Binmore and Samuelson is present again. I will discuss the reason why the "stable" group of strategies in this model are different than those in Binmore and Samuelson's model in the next section.

Evolution with Costly States


In this section it may be more accurate to think of the automata as representing costly rationality in addition to bounded rationality. We can think of this as increased strategic complexity requiring more time and resources. Rather than deal with different types of mutation schemes, I will allow mutation in these simulations from our stylized model of the strategies with a mutation rate of 0.0001. Here I want to focus primarily on the effects of costly complexity on the evolutionary dynamic process.

In order to capture the idea that states are costly, I supplement the payoffs of the one state machines (ALL C and ALL D) by 0.05 and 0.1. I simulated five thousand generations of this dynamic process which began from the center of the unit simplex in R^{26} . I present the results in figures 3.12 and 3.13. These diagrams show how the population changes over time. I have chosen to represent only the GRIM, ALL C, and ALL D strategies because their dynamic behavior accounts for all of the interesting phenomena.

Again, the GRIM and ALL C strategies play an important role in the evolutionary process. When having states is less costly, we can see the ALL D strategy

cannot successfully invade the population on a permanent basis. However, when ALL C gets very large, which will occur through both mutation and the supplemented payoffs, the ALL D strategy nearly invades but then dies out before it can complete the invasion. However, if the cost of maintaining a state is sufficiently large, ALL D can successfully invade. Since it is one of the simplest machines, there will not be any mutation which can threaten its existence. Note, however, this result is dependent on specific values for complexity costs.


Figure 3.12 — Evolution with Costly States ($c = 0.05$)



We can see in figures 3.12 and 3.13 if states are sufficiently costly, ALL D will be the evolutionary outcome. The population will initially tend to being all "nice." Since the mean strategies are in very small proportions in the population, the ALL C strategy does very well because of the higher payoff. When it becomes a sufficiently large proportion, it is invaded by ALL D. If the supplemental payoff



is not great enough, ALL D will enjoy a brief period of success, but will again be displaced by "nice" strategies, especially GRIM, and the cycle starts again.

Figure 3.13 — Evolution with Costly States ($c = 0.1$)



Evolution with Costly Transitions

In these simulations I attempt to capture the idea that maintaining transitions is costly. That is, any time a decision must be made by the person implementing these rules, there is an associated cost. The major difference in this simulation relative to the previous ones is that in these simulations GRIM is more successful than any of the other nice two state machines. The reason is GRIM is less complex than some of the other strategies. That is, it is an easier rule to apply than rules like TFT, but clearly it is more complex than ALL C. Again, GRIM is able to take advantage of any of the poor strategies and reap the benefits of cooperation.

Figure 3.14 — Evolution with Costly Transitions ($p = 0.01$)Figure 3.15 — Evolution with Costly Transitions ($p = 0.1$)

The evolutionary results are depicted in figures 3.14 and 3.15. Here I use penalties for complexity of 0.01 and 0.1 per "wire" in the stylized machines. We can see as long as the penalty is not too great, the population will tend to cycle. It begins by evolving to nearly all GRIM. However, since ALL C earns a higher payoff because it is less complex, it eventually becomes a very large portion of the population. When the proportion of ALL C in the population is large enough, ALL D's proportion grows rapidly. Then, unless the penalty is too large, GRIM emerges as the dominant strategy, and the cycle begins once more. If the penalty is large enough, though, ALL D will eventually dominate the population.

Discussion

These simulations suggest a number of things. The most conspicuous result is the evolutionary success of the GRIM strategy. This is surprising since it has not appeared as the result of any previous evolutionary simulation of which I am aware. The argument against the GRIM strategy being successful is that it is not forgiving. This is not a problem in these simulations because two state Moore machines cannot probe for weakness in the strategies periodically. However, there is something to be learned from this result. One possible reason for GRIM's success is it can exploit poor strategies. It is not difficult to see TFT can never do strictly better than it's opponent because it merely mirrors it's opponent. GRIM can exploit irrational play better than TFT. When GRIM and TFT play the other strategies, GRIM is usually better at taking advantage of bad play. Of the twenty-four other strategies, GRIM earns more than TFT against fifteen of them. On the other hand TFT does better than GRIM against only two of them.

The results of these simulations fall far short of suggesting GRIM is evolutionarily superior to all other strategies in the RPD, but they do show what I feel is a profound weakness in the TFT strategy. The strategy which just copies it's

opponent cannot take advantage of poor play which will be present through such processes as mutation. These simulations do, however, suggest the most successful strategies will be those who are able to take advantage of irrational play in an evolutionary situation where mutants enter over time.

It is obvious the GRIM strategy doesn't exploit all irrational play because the ALL C strategy survived in positive proportions in all of the simulations without complexity costs. The reason for this appears to be the GRIM strategy cannot identify the ALL C strategy when they play each other. Binmore and Samuelson (1990) suggested the possibility irrational strategies and rational strategies could exist together as the consequence of an evolutionary dynamic process. In these simulations there may be, in the long run, a substantial number of irrational players. However, the fact they look exactly like the rational strategy GRIM when playing him keeps them from being exploited. Also, any strategy which attempts to exploit these irrational strategies is forced to take the noncooperative payoff in the game most of the time. Therefore, the "nice" strategy GRIM protects, in a sense, the irrational players. I refer to it as an "enforcer" strategy which keeps the "mean" strategies from invading the population. If for some reason the number of ALL C players were to increase, ALL D could invade the population. In the simulations I performed that was not a likely event. Generally, whenever the proportion of ALL C grew relative to the other strategies, the ALL D and other mean strategies would force it to become smaller. However, because nearly all of the strategies in the population were nice, ALL C survived in small proportions while the "mean" strategies earned very small payoffs relative to the cooperative strategies and, hence, died off rapidly.

The "nice" strategies in this simulation have a property which is similar in spirit to Binmore and Samuelson's Payoff-equivalent Polymorphous MESS (PPMESS). They form something akin to a PPMESS in the simulations I performed. It is

clear this set of strategies is not a PPMESS in the model analyzed by Binmore and Samuelson. The reason for this gets to the heart of the differences between the models analyzed by both Abreu/Rubinstein and Binmore/Samuelson and the evolutionary processes I simulated here.

The simulations I performed have a number of different strategies entering over time either simultaneously or as a single strategy type in a group. In other words, my simulations have the populations being repeatedly perturbed. Dean Foster and Peyton Young (1987) analyzed this type process analytically in a very simple environment. However, most studies in this area analyze the stability of populations against a single perturbation. Here, though, we look at how population mixtures fare while being repeatedly attacked by possible mutant strategies.

A discussion of Binmore and Samuelson's results will help clarify the issue. They suggest a mixture of the cc, cd, CC, and CA machines form a PPMESS. That is, an appropriate mixture of these strategies cannot be successfully invaded by a mutant strategy. This is true because they earn equal payoffs against each other and have equal complexity. There is a symbiotic relationship among these strategies which protects them from invasion. However, this is true only if we consider these perturbations as one time events. Certainly this mix of strategies is immune to a small proportion of invading mutants. Although the exact mix of strategies may change after a mutant invasion is thwarted, the same set of strategies will survive. However, in repeated simulations I found these machines are not immune to invasion when mutants repeatedly attempt to enter the population because the GRIM strategy will be able to invade after some number of unsuccessful tries. When a GRIM machine plays the Abreu/Rubinstein utility maximizing equilibrium machine (UTIL) it earns higher payoffs against UTIL than UTIL earns against it. Hence, we have a situation where cc, cd, and CA attempt to invade the population over time through mutation and are successful because they earn the cooperative payoff when

they play themselves and UTIL. However, after these strategies enter the population, GRIM can successfully invade if it appears often enough. Each time GRIM attempts to invade, the proportion of UTIL remaining decreases relative to cc and cd (both "nice" strategies) because GRIM earns the cooperative payoff against them and does better against UTIL than UTIL does against GRIM. Also, GRIM earns the maximum possible payoff against CA. Therefore, after enough attempts GRIM will successfully invade the population.

The ideas of perfect equilibria and trembles are important to these simulations. Reinhard Selten (1983) and others have used the idea of trembles in the play of a game to select among multiple equilibria. He also applies this idea to evolutionary stability to expand the set of evolutionarily viable strategies.¹⁴ However, applying the idea of trembles to the play of the automaton selection game is not straightforward. We cannot allow a machine to "misplay" during an infinite RPD because there appears to be no sensible way to define such a tremble. If a machine were to change strategies with positive probability during the game, the change in payoffs could be discontinuous because we calculate the payoffs as the limit of the mean. To see this, suppose two GRIM strategies were playing the RPD, and with some arbitrarily small positive probability one changes to ALL D during the game. We know the change will occur in finite time. Hence, the payoffs will switch to the defecting payoff for both machines. Such a discontinuity makes this sort of tremble unusable. Instead, I consider trembles in the evolutionary process in these simulations. That is, what happens if the strategies do not breed true. Another way of looking at the problem is to imagine what happens if there are perturbations to the payoff matrix.

These trembles point out another difference between these simulations and the work by Abreu/Rubinstein and Binmore/Samuelson. In this analysis, there is no

¹⁴ See Selten (1983).

such thing as an unused state. That is, because mutant strategies can appear over time, the idea of unused states means nothing here. This goes back to the fact that in other evolutionary models only one strategy or a mix of strategies enters at a time. Or, perhaps more precisely, we only analyze how a strategy or mix of strategies does against potential mix of invaders. Without the aid of computer simulations it is very difficult to capture the interactions between strategies as well as the impact of a certain mutation scheme.

It is also interesting to note if the penalty for complexity is sufficiently small, the nice strategies will remain immune to successful invasion. It is not until the cost of complexity gets large that it affects the qualitative results of the evolutionary simulation. We can think of complexity as being the cost of monitoring your opponent and responding accordingly. If the cost of monitoring actions increases enough, the defecting outcome will prevail.

Finally, I will comment on the obvious success of the GRIM strategy in these simulations. We must keep in mind we considered only two state machines. This is a severe limitation on the complexity we allow. In fact, a strategy essentially identical to GRIM was submitted to Axelrod's RPD tournament and finished fifty-second out of sixty-three strategies. However, we can say the success of GRIM in this tournament captures something which has frequently been overlooked in this sort of model. In order for a strategy to survive over time, it should be able to exploit possible bad strategies. The idea TFT should be in some sense the "best" strategy in an infinite RPD has the significant failing that TFT doesn't do strictly better than any strategy it meets. It seems reasonable to expect a strategy which succeeds over time should be able to take advantage of bad players who appear.

Summary

This essay reports the results of a number of computer simulations of infinite RPDs. The results are somewhat surprising because the GRIM strategy is clearly the most successful of the cooperative strategies. I attribute this to the fact GRIM is able to both exploit poor strategies and earn the cooperative payoff against other "nice" strategies. The results do not suggest GRIM would be the "best" strategy in all environments.

Although this work appears to be to the contrary, it supports various results of Binmore and Samuelson's recent work. That is, the idea of a PPMESS seems valid in evolutionary models. The differences in our results stem from the fact that the strategies in my simulations are subject to repeated perturbations, and they analyze a game which is more tranquil since they look at what strategies can invade a population if the intruders appear only once either one at a time or in certain mixes. The reason for performing these simulations is to capture what happens when many things are happening at once. Analyzing these problems analytically is intractable because of the complex interrelationships among the strategies and the dynamic process. We are able to see how all of these forces affect the final population under these specific conditions through the use of computer simulations.

Bibliography

- Abreu, Dilip and Ariel Rubinstein. "The Structure of Nash Equilibrium in Repeated Games with Finite Automata." *Econometrica* 56 (November 1988):1259-1281.
- Aumann, Robert J. "Survey of Repeated Games." *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*. Bibliographisches Institute:Manheim, 1981, pp.11-42.
- Axelrod, Robert. *The Evolution of Cooperation*. New York: Basic Books, 1984.
- Banks, Jeffrey and Rangarajan Sundaram. "Repeated Games, Finite Automata, and Complexity." University of Rochester Working Paper No. 183. March 1989.
- Binmore, Kenneth G. and Larry Samuelson. "Notes on Evolutionarily Stable Strategies in Repeated Games Played by Finite Automata." unpublished manuscript, 1989.
- Binmore, Kenneth G. and Larry Samuelson. "Evolutionary Stability in Repeated Games Played by Finite Automata." unpublished manuscript, 1990.
- Boyd, Robert and Jeffrey Lorberbaum. "No Pure Strategy is Evolutionarily Stable in the Repeated Prisoner's Dilemma Game," *Nature* 327 (1987):58-59.
- Farrell, Joseph and Roger Ware. "Evolutionary Stability in the Repeated Prisoner's Dilemma," *Theoretical Population Biology* 36 (1989):161-166.
- Foster, Dean and Peyton Young. "Stochastic Evolutionary Game Dynamics." unpublished manuscript, September 20, 1988.
- Fudenberg, Drew and Eric Maskin. "Evolution and Cooperation in Noisy Repeated Games," Harvard Institute for Economic Research Discussion Paper 1469, 1990.
- Harsanyi, John. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge: Cambridge University Press, 1977.
- Harsanyi, John and Reinhard Selten. *A General Theory of Equilibrium Selection in Games*, Cambridge: MIT Press, 1989.
- Kim, Yong-Gwan. "Evolutionarily Stable Strategies in the Repeated Prisoner's Dilemma," unpublished manuscript, October 1989.

- Linster, Bruce G. "Automata Play the Prisoners' Dilemma: Another Look at the Evolution of Cooperation," unpublished manuscript, 1990.
- Maynard Smith, John. *Evolution and the Theory of Games*. Cambridge: Cambridge University Press, 1982.
- Maynard Smith, John and G. R. Price. "The Logic of Animal Conflict," *Nature* 246 (1973):15-18.
- Miller, John. "The Evolution of Automata in the Repeated Prisoner's Dilemma," unpublished manuscript, 1987.
- Neyman, Abraham. "Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoner's Dilemma." *Economic Letters* 19 (1985):227-229.
- Radner, Roy. "Can Bounded Rationality Resolve the Prisoner's Dilemma?" *Essays in Honor of Gerrard Debreu*, New York: North-Holland, 1986.
- Robson, Arthur J. "Efficiency in Evolutionary Games: Darwin, Nash, and the Secret Handshake," University of Western Ontario Department of Economics Working Paper, 1989.
- Samuelson, Larry. "Evolutionary Stability in Asymmetric Games," Pennsylvania State University Department of Economics Working Paper, 1989.
- Selten, Reinhard. "Evolutionary Stability in Extensive Two-Person Games." *Mathematical Social Sciences* 5 (1983): 269-363.

CHAPTER IV

NEW DIMENSIONS IN RENT-SEEKING

Introduction

When something of value is awarded to the winner of a contest, there is frequently wasteful rent-seeking expenditure. This is true of such prizes as monopoly rights, tariffs, quotas, government contracts, or favorable legislation. Gordon Tullock's (1967) insight into this problem was that the social loss associated with rent-seeking exceeds the deadweight loss identified by A. C. Harberger (1954). Tullock argued the inefficiencies caused by tariffs or regulation also inflict a social cost. More importantly for the purposes this essay, he pointed out there will be expenditures made in attempts to realize the economic rents. From society's point of view, significant resources may be wasted in attempts to form a monopoly or obtain quota rights from the government. Examples of this type of activity are such things as political lobbying, bribes, or even studying for employment tests.¹ These expenditures generally exceed any transfer from individuals to businesses. The original authors on this subject, most notably Anne Krueger (1974), Richard Posner (1975), and Gary Becker (1968), formulated models in which rent-seeking activity completely dissipates the rent. In the case of monopoly and regulation, the economic rent is

¹ For an excellent introduction to this problem see Buchanan (1980).

any monopoly profit which could be generated. However, the loss to society in these models is the dissipated rent plus any deadweight loss.

Tullock (1980) developed a game theoretic approach to this problem in which the prize is awarded to the winner of a lottery-like contest. The probability an individual wins this contest depends on his contribution and the total contributions of all players. He showed total rent-seeking expenditures can be less than, equal to, or greater than the value of the prize. The degree of rent dissipation in his model depends on the number of players and the parameter values for a probability function.

Many extensions of Tullock's basic model have been analyzed. William Corcoran (1983), for example, considered a long run setting for Tullock's model and found rents will be completely dissipated if free entry is allowed. Richard Higgins, William Shugart, and Robert Tollison (1985) examined the situation where the prize goes to the player who appears to be trying hardest, but effort is observed imperfectly. Their analysis shows rent-seeking activity will occasionally exceed and occasionally be less than the value of the prize, but on average rents will be fully dissipated. William P. Rogerson (1982) examined rent-seeking activity in the context of a monopoly where firms face differential start up costs or monopoly rights are periodically reassigned but the current monopolist has an advantage. He finds rents will be less than completely dissipated in this situation. Ayre Hillman and Eliakim Katz (1984) employ a model with risk averse rent-seekers and large prizes to reach similar results. Corcoran and Gordon Karels (1985) extend the analysis of Corcoran (1983) by allowing different types of long-run competitive response. Ayre Hillman and John Riley (1987) allow players to value the prize differently, but each player is indifferent to who wins if it is someone else. Their model is similar in spirit to this essay.

In addition to the above authors, significant contributions have been made to

this literature by international economists. The most notable of these are Jagdish Bhagwati (1980, 1982), Bhagwati and T. N. Srinivasan (1982), and Bhagwati, R. A. Brecher, and T. Hatta (1985). This branch of the literature deals primarily with rent-seeking, or as they are called in these works Directly Unproductive Profit-seeking (DUP), activities in an international trade context. An interesting result from this branch of the literature shows these activities may be welfare enhancing as a "second best" solution. This essay is similar in spirit to the works of the "public choice" economists like Tullock and others rather than these studies.

In this essay I extend Tullock's model in new directions. In the next section, I examine the rent-seeking game played sequentially. That is, I allow one player to go first in contributing toward the prize. We can think of these contributions as lobbying expenditures or bribes to legislators or bureaucrats for political favors. After analyzing a sequential rent-seeking game, I examine what happens in the Tullock model if we allow players to have different preferences over who else wins the prize. The basic model I use is due to Ted Bergstrom and Hal Varian (1987). They were the first to suggest this type of framework for analyzing military alliances and arms races. Finally, I summarize the results.

Sequential Rent-seeking

Suppose there are two individuals, each trying to influence the award of a prize. We can think of the prize as monopoly or quota rights, or anything else of value which can be allocated through the political process. This is the type of model first analyzed by Tullock (1980). The most closely related work in the rent-seeking literature is that of Hillman and Riley (1987) where they examined a game in which players compete for a prize each values differently. Hillman and Riley found, roughly speaking, rents would be less than fully dissipated in imperfectly discriminating contests² with different valuations among the players. However, throughout this

² The term "imperfectly discriminating" contests refers to those contests where the player making

literature the analyses have examined the Nash equilibrium assuming the players move simultaneously. I will refer to these simultaneous move games as Cournot games because they are similar in nature to Cournot duopoly games. These rent-seeking models have not been analyzed as Stackelberg games, or games in which individuals choose contributions sequentially. In this game, I assume one player can commit himself to a contribution. Then, the other player makes her decision knowing the first player's choice. This section is similar in spirit to Hal Varian's (1989) analysis of public goods provision when agents act sequentially.

Looking at rent-seeking models as a sequential game has a significant impact on the results. I will use Hillman and Riley's analysis as a benchmark against which I can compare the Stackelberg equilibrium. The difference between the equilibria of the simultaneous and sequential games depends on which agent goes first, as well as the relative difference between the players' valuations of the prize. I will examine how individual and total equilibrium outlays in the sequential game compare with those in the simultaneous game.

The Simultaneous Move Game

Before looking at the sequential contributions game, I will briefly describe the simultaneous move game. This is a simplified version of part of Hillman and Riley's work. In this game there are two players: Player I and Player II. They compete for something of value, and the prize goes to Player i with a probability depending on both players' contributions. The strategy space for both players is $S_i = \{x_i | x_i \in R_+\}$ for $i \in \{1, 2\}$. Finally, to complete the specification of the game I describe each player's payoff function:

$$U_i(x_1, x_2) = v_i \cdot p_i(x_1, x_2) - x_i, \quad i \in \{1, 2\}.$$

the largest contribution doesn't necessarily win.

In this expression, v_i is Player i 's valuation for the prize. I always assume the valuations to be strictly positive, and they are common knowledge among the players. I relax the common knowledge assumption later. Also, $p_i(x_1, x_2)$ is the probability Player i wins the contest when the contributions from Players I and II are x_1 and x_2 respectively. In the model I consider, the probability function is defined as follows:

$$p_i(x_1, x_2) = \frac{x_i}{x_1 + x_2}, \quad i \in \{1, 2\}.$$

There are more general formulations which have been analyzed in simultaneous move games. For example, Tullock considered the case with the probability Player i wins as $p_i(x_1, x_2) = x_i^r / (x_1^r + x_2^r)$. Bergstrom and Varian have generalized this to be the following for some function $g_i(x_i)$:

$$p_i(x_1, x_2) = \frac{e^{g_i(x_i)}}{e^{g_1(x_1)} + e^{g_2(x_2)}}, \quad i \in \{1, 2\}.$$

It is easy to see Tullock's formulation is a special case of Bergstrom and Varian's model where $g_i(x_i) = r \ln x_i$. Of course, these are easily adapted to more players in the obvious way. I will analyze the special case of Bergstrom and Varian's formulation where $g_i(x_i) = \ln x_i$. The contest is similar to a lottery, and making contributions is akin to purchasing lottery tickets. The other probability functions have similar interpretations.

In Tullock's formulation, the parameter r captures information about the rate of change in the marginal cost of influencing the outcome of the contest. A lower value of r indicates the marginal cost curve rises more steeply. Put another way, a small r value means a marginal contribution will have a less significant effect on the outcome than it would if r were greater.³ For some intuition, imagine an r value very close to zero. This means each player has a probability of about 1/2 of winning regardless of his contribution, so the marginal cost is very high. Hence,

³ For a complete discussion and some examples see Tullock (1980).

we should expect to see little rent-seeking activity. As r gets larger, an increase in contributions has a greater impact on a player's probability of winning.

Solving for the Cournot/Nash equilibrium is accomplished by first differentiating both players' payoff functions to obtain the first order utility maximizing conditions. Then, we can simultaneously solve them. The first order conditions are

$$\frac{\partial U_1(x_1, x_2)}{\partial x_1} = \frac{v_1 x_2}{s^2} - 1 = 0, \quad (1)$$

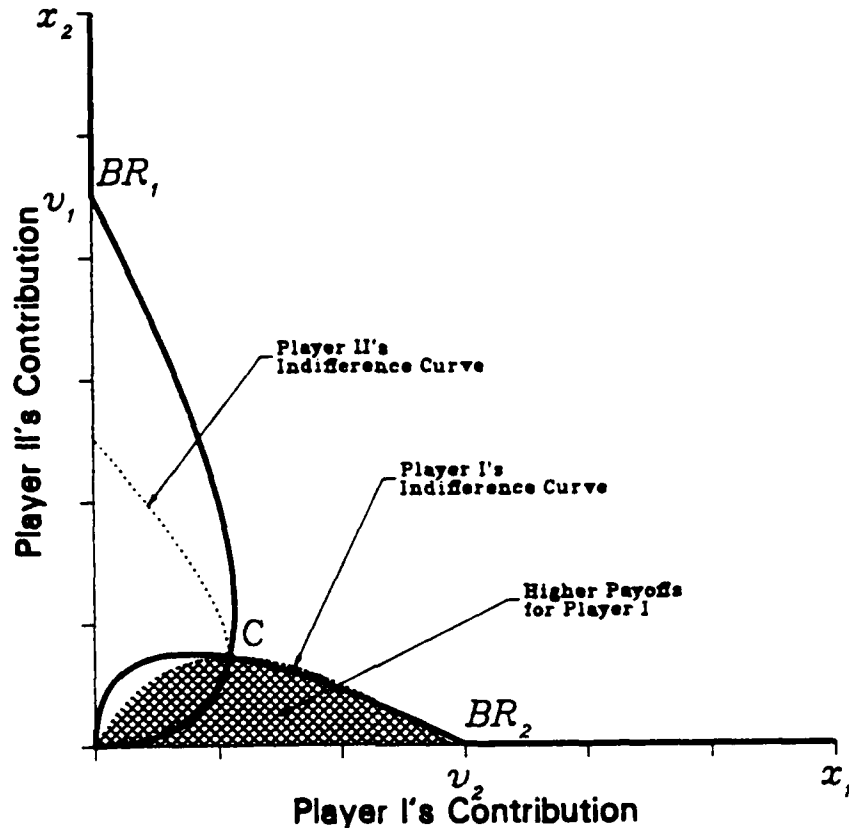
$$\frac{\partial U_2(x_1, x_2)}{\partial x_2} = \frac{v_2 x_1}{s^2} - 1 = 0, \quad (2)$$

where $s \equiv x_1 + x_2$. Solving for s yields equilibrium total outlays of $s^* = v_1 v_2 / (v_1 + v_2)$. In general, Hillman and Riley showed the equilibrium total outlay is $(n-1)\hat{v}/n$ where n is the number of contestants, and \hat{v} is the harmonic mean⁴ of the valuations. The equilibrium individual outlays in the contest are $x_1^* = v_1^2 v_2 / (v_1 + v_2)^2$ and $x_2^* = v_2^2 v_1 / (v_1 + v_2)^2$. An interesting aspect of this game is that in equilibrium the ratio of outlays to valuations, x_i^* / v_i , is the same for both players. That is, both contestants expend the same proportion of their valuations in equilibrium. This also means the equilibrium "odds" Player I wins, is the ratio of valuations, or $p_1 / p_2 = v_1 / v_2$. Finally, the equilibrium expected payoff to Player i is $v_i^3 / (v_1 + v_2)^2$.

Before examining the Stackelberg version of the game it will be helpful to see how the problem looks when solved graphically. I do this in figure 4.1. In this diagram I show what each player's best reply function and equal payoff curves look like. It is important to note each player's upper contour, or "better than," set lies below his equal payoff curve. The Cournot/Nash equilibrium outcome is the point where the best reply functions intersect. That is, if the players make the contributions indicated at that intersection neither player can unilaterally increase his utility.

⁴ The harmonic mean, \hat{y} , of a set of data $\{y_1, y_2, \dots, y_n\}$ is defined as follows: $\hat{y} = \frac{n}{\sum_{i=1}^n \frac{1}{y_i}}$.

Figure 4.1 — Cournot Equilibrium in the Rent-seeking Game



Player i 's best reply is found by solving his first order utility maximization condition and solving for x_i . Hence, Player i 's best reply is (for $i \neq j$)

$$BR_i(x_j) = \begin{cases} (v_i x_j)^{1/2} - x_j, & \text{if } x_j \in (0, v_i]; \\ 0, & \text{if } x_j > v_i. \end{cases}$$

The two best reply curves intersect at the point $(v_1^2 v_2 / (v_1 + v_2)^2, v_1 v_2^2 / (v_1 + v_2)^2)$. Notice in the simultaneous move game both players always contribute in equilibrium. We will see this need not be true for the sequential game. Another fact about the best replies which will become more important later is that $BR_i(x_j)$ is maximized at $x_j = v_i/4$. In order to analyze the sequential game, I will return to this sort of graph a number of times.⁵

⁵ See Congleton (1980) for more discussion of the best reply functions.

The Stackelberg Game

In order to change the model I have been considering into a Stackelberg game I must redefine one player's strategy space and describe the order of play. Player I will always be the first to move. His strategy space is unchanged, but the strategy space for Player II is now $S_2 = \{f | f : R_+ \rightarrow R_+\}$. In words, Player I's strategy space is the set of nonnegative real numbers, and Player II's strategy space is the set of all functions which map a nonnegative real number into a nonnegative real number. The utility functions remain unchanged.

The sequence of moves in the game is the following: Player I moves first by selecting a contribution; then Player II chooses a contribution level after observing Player I's choice. In order to eliminate Nash equilibrium strategies where Player II makes incredible threats, I will look only at subgame perfect⁶ Nash equilibria in the sequential game. This eliminates equilibria where Player II makes a threat she would not want to carry out if Player I strays from the equilibrium path. That is, I will not allow equilibria which depend on II's unbelievable threats. The requirement the equilibrium be subgame perfect means regardless what Player I contributes, Player II chooses her outlay to maximize her utility conditional on what Player I does. An interesting and simple extension to the sequential game is to add a third player, say a government official, to the model. We can allow him to decide who gets the opportunity to make the first offer. I will discuss this shortly.

I will now describe the unique subgame perfect equilibrium for this game. First, consider what Player II will do when it is her turn. She will choose an outlay using her best reply function. We find this by differentiating her utility function with respect to her expenditure. We then set the derivative equal to zero and solve for x_2 as a function of x_1 . If we solve equation (2) for $x_2(x_1)$ we get $x_2(x_1) = (v_2 x_1)^{1/2} - x_1$. However, Player II will not participate if Player I contributes more than v_2 because

⁶ For an enlightening discussion of this and other Nash equilibrium refinements see Binmore (1988).

doing so would assure her a negative payoff. She can always obtain a payoff of zero with certainty by choosing not to participate. This means her equilibrium strategy is the following: $x_2^*(x_1) = \max\{(v_2 x_1)^{1/2} - x_1, 0\}$. In words, Player II will choose x_2 to maximize $v_2 x_2 / (x_1 + x_2) - x_2$, unless doing so assures her a negative expected payoff. Now, substituting $x_2^*(x_1)$ for x_2 in Player I's utility function, we can solve for I's optimal choice by differentiating his utility function with respect to x_1 and setting the result equal to zero. This is simplified because we know $x_1 + x_2^*(x_1) = \max\{(v_2 x_1)^{1/2}, x_1\}$. Substituting this into Player I's utility function we get

$$U_1(x_1, x_2(x_1)) = \frac{v_1 x_1}{\max\{(v_2 x_1)^{1/2}, x_1\}} - x_1.$$

Maximizing Player I's utility function yields $x_1^* = \min\{v_1^2/4v_2, v_2\}$. This means if Player II chooses her outlay from her best reply function, Player I's best strategy is to choose x_1 to maximize $v_1 x_1 / (v_2 x_1)^{1/2} - x_1$ unless doing so yields an outlay greater than v_2 . Player I will never contribute more than Player II's valuation. If he did, Player II would never rationally choose to participate in the contest, so any contribution larger than v_2 would be wasteful.

With these equilibrium strategies we can determine the equilibrium outlays for both players. First, suppose we have an interior solution. Substitution and some algebra establish the Stackelberg equilibrium outlays are $(\frac{v_1^2}{4v_2}, \frac{v_1}{2} - \frac{v_1^2}{4v_2})$. An interesting property of the sequential rent-seeking game is that it is possible to have only one player make a strictly positive contribution in equilibrium. This will occur if Player II's valuation is less than half of Player I's valuation.⁷ The Stackelberg equilibrium outlays for a corner solution are $(v_2, 0)$. Therefore, the sum of equilibrium outlays is $\min\{v_1/2, v_2\}$. This leads us to the first proposition.

Proposition 1: Suppose the players' valuations for the prize are different. If the player with the larger (resp. smaller) valuation goes first, the sum of the resulting

⁷ This can easily be seen by computing the equilibrium outlays if $v_1/2 \geq v_2$.

Stackelberg equilibrium outlays will be larger (resp. smaller) than those in the Cournot game.

To prove this, I must consider two cases of an interior solution and the corner solution. Recall $v_1 v_2 / (v_1 + v_2)$ is the sum of the equilibrium contributions in the simultaneous move game. First I look at the interior solutions when $v_1 > v_2$. In other words, the player with the larger valuation goes first. To prove the result I suppose it is not true. That is, suppose $v_1/2 \leq v_1 v_2 / (v_1 + v_2)$. After dividing through by v_1 we have $1/2 \leq v_2 / (v_1 + v_2)$ which is a contradiction since $v_1 > v_2$ by hypothesis. Hence, the proposition holds in this case. The proof for $v_1 < v_2$ is the same argument with the weak inequalities reversed. If we are at a corner solution, $v_1/2 \geq v_2$. Therefore, all I have remaining to show is $v_2 > v_1 v_2 / (v_1 + v_2)$. I can again assume the proposition is false and derive the contradiction $1 \leq v_1 / (v_1 + v_2)$ since $v_2 > 0$ by assumption.

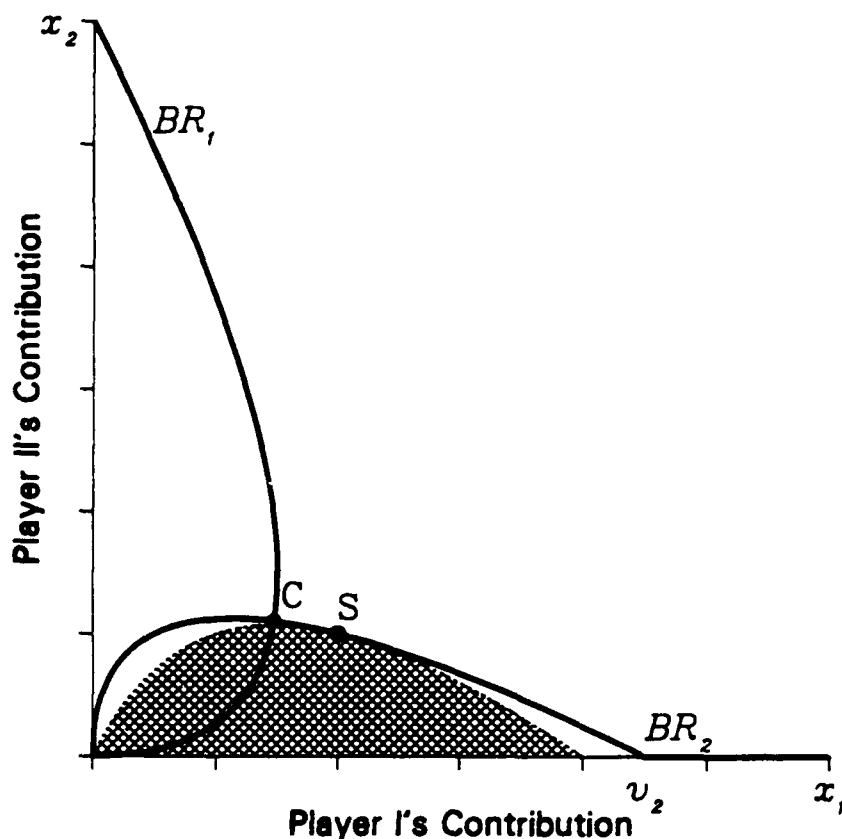
Again, it is of some value to solve for the equilibrium outlays graphically. I begin with the interior equilibrium. Consider figure 4.2, which assumes $v_1 > v_2 > v_1/2$. I labeled the Cournot equilibrium by C and the Stackelberg equilibrium by S . Since Player I's "better than set" lies below his equal payoff curves, it is easy to see Player I is as well off at S as he can possibly be. It is easy to verify the tangency of the two curves. The slope of BR_2 can be easily derived and is $(dx_2(x_1)/dx_1)_{BR_2} = \frac{1}{2}(v_2/x_1)^{1/2} - 1$. If we calculate the equilibrium expected payoff for Player I, we find it is $v_1^2/4v_2$. The slope of the equal payoff curve for $U_1 = v_1^2/(4v_2)$ can be shown to be

$$\left(\frac{dx_2(x_1)}{dx_1} \right)_{U_1 = \frac{v_1^2}{4v_2}} = \frac{\frac{v_1^2}{4v_2} (v_1 - \frac{v_1^2}{4v_2} - 2x_1) - x_1^2}{(\frac{v_1^2}{4v_2} + x_1)^2}$$

Substituting the equilibrium outlays into both expressions we can see the slopes are the same and equal to $\frac{v_2}{v_1} - 1$. Hence, the curves are tangent. We can also see the sum of the payoffs at S is greater than the sum of the payoffs at C . This is verified

by noting S is to the right of the line passing through C with a slope of -1 .

Figure 4.2 — An Interior Stackelberg Equilibrium in the Rent-seeking Game




Now I will graphically describe the corner solution, which occurs when $v_1/2 \geq v_2$. I have drawn the best reply curves in figure 4.3 for the limiting case where $v_1/2 = v_2$. BR_2 is not differentiable at $x_1 = v_2$; however, as $x_1 \rightarrow v_2$, the slope of BR_2 approaches $-1/2$ from above. By this I mean the slope of BR_2 is greater than $-1/2$ for $x_1 < v_2$. If Player I chooses an outlay of v_2 , he assures himself of a payoff of $v_1 - v_2$. The slope of Player I's equal payoff curve for $U_1 = v_1 - v_2$. It is

$$\left(\frac{dx_2(x_1)}{dx_1} \right)_{U_1=v_1-v_2} = \frac{(v_1 - v_2)(v_1 - 2x_1 - v_1 + v_2) - x_1^2}{(v_1 - v_2 + x_1)^2}$$

If we evaluate the above derivative at $x_1 = v_2$ we get the slope of the equal payoff curve to be $-v_2/v_1$. If $v_1/2 = v_2$ as we assumed, the slope is $-1/2$. If $v_1/2 > v_2$,

the slope of the equal payoff curve at $(v_2, 0)$ is greater (smaller in absolute value) than the slope of BR_2 if $x_1 < v_2$ and smaller (greater in absolute value) for $x_1 > v_2$. Hence, the best Player I can do is achieve the payoff $v_1 - v_2$.


Figure 4.3 — A Corner Stackelberg Equilibrium ($v_1/2 \geq v_2$)



Now we can consider adding another player. Suppose we allow a third player, say a government official, who gets to choose who goes first. If the new player benefits by these rent-seeking expenditures, he will want the player with the higher valuation to go first. On the other hand, if a benevolent government official wants to minimize the rent-seeking expenditure he will have the lower valuation player go first. In his 1980 essay, Gordon Tullock identified the fundamental problem to be addressed in these models as finding ways to reduce the rent-seeking expenditures. This proposition lends some insight to that problem if the game is played sequentially.

Proposition 2: In the Stackelberg game the first player is not always strictly better off than he would be in the Cournot game. Certainly, he is no worse off by going first.

Figure 4.4 — Cournot and Stackelberg Equilibrium with Equal Valuations



In Stackelberg duopoly games a firm is generally strictly better off being the leader even if both duopolists are identical in every respect except the order of play. This is not true in this rent-seeking game. To prove this proposition consider a sequential rent-seeking game where players have identical valuations, which I normalize to one. The Stackelberg equilibrium outlays are $(1/4, 1/4)$ which are the same as the Cournot equilibrium outlays. The reason for the difference is the best reply curves and equal payoff curves are shaped differently than they are in the

standard duopoly models. When both players' valuations are equal, the equilibrium outlays in the Cournot and Stackelberg games occur at the maximal points on the best reply curves. I have drawn the appropriate equal payoff curves in figure 4.4.

It is easy to see the indifference curves which yield at least the Cournot equilibrium outcome are everywhere inside the best reply curves except at the equilibrium point. By definition, the slope of a player's indifference curve is zero everywhere along his best reply curve, and the slope of the best reply curve is zero at $(1/4, 1/4)$ because it is at a maximum. Therefore, each player's indifference curve is tangent to the other player's best reply curve at the intersection of the best reply curves. Hence, it makes no difference who goes first. The Stackelberg leader is not able to do any better than the Cournot equilibrium. However, if the valuations are different, then the leader can improve his payoff over that of the Cournot game.

Proposition 3: In the sequential game the equilibrium ratio of outlays for an interior solution (also, the "odds" Player I wins) is $x_1^*/x_2^* = v_1/(2v_2 - v_1)$. This is greater than the equilibrium ratio of the Cournot game if and only if $v_1 > v_2$. If we are at a corner equilibrium, this ratio is not defined because $x_2^* = 0$, and Player I wins with certainty.

The proposition is easily verified by taking the ratio of equilibrium outlays and simplifying. This means if the person with the lower (resp. higher) valuation goes first, he is less (resp. more) likely to win. However, the leader will have a higher expected payoff than he would in the Cournot game. It may seem counter intuitive at first that the leader may be less likely to win in the sequential game than he would be in the simultaneous game. However, by contributing less when he has the first choice he increases his expected payoff as long as his opponent plays the equilibrium strategy.

An implication of this proposition is the high valuation player is more likely to win whether he goes first or second. Hence, in some sense, we can say the outcome

from a sequential rent-seeking game in which the high valuation player goes second is socially superior to the outcome of the simultaneous game if the players value the prize differently. This is because there is less rent-seeking expenditure and the prize is more likely to go to the player who values it more highly.

Incomplete Information

Another interesting question involves incomplete information. Suppose Player I is unsure of Player II's type. In this case, Player I will choose an expenditure level to maximize his expected payoff. It is not very interesting to allow uncertainty about Player I's type because Player II merely responds to Player I's choice to make herself as well off as possible based on her type.

For simplicity I will assume Player II can only have one of two possible valuations; $v_2 \in \{v_l, v_h\}$ where $v_h > v_l$. Also, I assume the probability Player II is the v_h type is q , and the probability she is the v_l type is $(1 - q)$. Applying what we found in the last section, we can see Player II's equilibrium strategy is now simply

$$x_2^*(x_1) = \begin{cases} \max\{0, (v_h x_1)^{1/2} - x_1\}, & \text{if } v_2 = v_h; \\ \max\{0, (v_l x_1)^{1/2} - x_1\}, & \text{if } v_2 = v_l. \end{cases}$$

The significant difference comes in the problem Player I must solve. He must choose x_1 to maximize the following expression:

$$E[U_1(x_1)] = q \cdot v_1 \Pr(\text{I wins} | v_2 = v_h) + (1 - q) \cdot v_1 \Pr(\text{I wins} | v_2 = v_l) - x_1.$$

We can solve this problem for Player I's optimal outlay, x_1^* , as we did earlier. The twist here is x_1^* depends on the valuations in a more complicated way. I begin by calculating x_1^* for the situation where the possible values of v_2 are such that Player II contributes whether the realization is v_h or v_l . After substituting $(v_h x_1)^{1/2}$ and $(v_l x_1)^{1/2}$ for $(x_1 + x_2)$, this means Player I maximizes the following expression:

$$E[U(x_1)] = q \cdot \frac{v_1 x_1}{(v_h x_1)^{1/2}} + (1 - q) \cdot \frac{v_1 x_1}{(v_l x_1)^{1/2}} - x_1.$$

The first order condition for maximizing this expression is

$$\frac{v_1 x_1^{-1/2}}{2} [q v_h^{-1/2} + (1-q) v_l^{-1/2}] = 1.$$

Solving for Player I's optimal outlay we find if the possible values of v_2 are such that either kind of Player II participates we get $x_1^* = v_1^2/4 \cdot [q v_h^{-1/2} + (1-q) v_l^{-1/2}]^2$. Keeping this in mind, we know if $x_1^* \geq v_l$ a low valuation type Player II will choose not to participate in the contest when it is her turn. This occurs if $x_1^* = v_1^2/4 \cdot [q v_h^{-1/2} + (1-q) v_l^{-1/2}]^2 \geq v_l$, or $v_1 \geq 2v_l^{1/2} [q v_h^{-1/2} + (1-q) v_l^{-1/2}]^{-1}$. Therefore, if $v_1 \geq 2v_l^{1/2} [q v_h^{-1/2} + (1-q) v_l^{-1/2}]^{-1}$, Player II's best reply will be zero in equilibrium if she is the low valuation type. Then, Player I will maximize the following expression:

$$E[U(x_1)] = q \cdot \frac{v_1 x_1}{(v_h x_1)^{1/2}} + (1-q) \cdot v_1 - x_1$$

subject to: $x_1 \geq v_l$.

This maximization problem is easily solved using the method of Lagrange. When $x_1 > v_l$, the first order maximization condition is $v_1 q (v_h x_1)^{-1/2} / 2 = 1$. Solving for Player I's optimal expenditure we get $x_1^* = \max\{v_l, q^2 v_1^2 / (4v_h)\}$. Again, if $x_1^* = q^2 v_1^2 / (4v_h) \geq v_h$, Player II will offer the best reply of zero. This will be true when $v_1 \geq 2v_h/q$. If this condition holds, Player I will choose v_h as his outlay.

Now I can summarize the equilibrium strategies.

If $v_1 < 2v_l^{1/2} [q v_h^{-1/2} + (1-q) v_l^{-1/2}]^{-1}$, then

$$x_1^* = \frac{v_1^2}{4} [q v_h^{-1/2} + (1-q) v_l^{-1/2}]^2.$$


If $2v_l^{1/2} [q v_h^{-1/2} + (1-q) v_l^{-1/2}]^{-1} \leq v_1 < 2v_h/q$, then

$$x_1^* = \max\left\{v_l, \frac{q^2 v_1^2}{4v_h}\right\}.$$

If $v_1 \geq 2v_h/q$, then $x_1^* = v_h$.

Figure 4.5 — Total Rent-seeking with Complete and Incomplete Information

$$v_h = 2, v_l = 1, v_2 = 1.5$$



I have provided figure 4.5 to show how uncertainty affects the total contributions in this game with incomplete information. In the case I depicted, the parameter values are $q = .5$, $v_l = 1$, and $v_h = 2$. I compare that to the case of complete information with $v_2 = 1.5$. We can see as long as Player I's contribution is less than v_l , the uncertainty increases the total contributions. However, if the values of v_l are in the range where only the high valuation type Player II participates it is possible the total contributions for the case of complete information exceed the total contributions for the incomplete information case.

The equilibrium in the incomplete information game is intuitively appealing. If q is small, Player I's equilibrium choice is close to what it would be with perfect

information. The same is true for q close to one. Also, the comparative statics are completely intuitive. We can see the equilibrium outlay, x_1^* , is nondecreasing in v_1 and q , and nonincreasing in both v_h and v_l . In the region where v_h and v_l type players contribute, x_1^* strictly increases as they do. Obviously, this model can be extended to include distributions with more possible types.

Summary

The above analysis has been concerned with the case where one agent can make the first move by committing to an outlay and allowing the second agent to respond. This is an interesting problem because we often see rent-seeking activity occur sequentially in the world. For example, one interest group will choose to be the first to endorse a candidate or lobby an agency for the program it wants. There is frequently an advantage to moving first, and this model captures it. Here it makes sense to be the first to commit to a choice. Also, if valuations differ between players, it makes a difference in the total rent-seeking expenditure in which order the players move. If a Congressman or bureaucrat were looking to maximize the bribes he can receive, he would want to offer the high valuation contestant the opportunity to go first. If we are looking to minimize the amount of political rent-seeking, the low valuation contestant should go first.

Another interesting feature of this model is that it also explains why both players don't always participate in a contest. In the case of perfect information we saw if the high valuation player goes first and has a valuation at least twice the lower valuation, the second player will choose not to compete. In other words, the high valuation leader will preempt the lower valuation follower in equilibrium. In the simultaneous move game, both players compete in equilibrium regardless of the relative size of the valuations. In this game it is rational at times to make a contribution equal to your opponent's valuation of the prize to ensure he does

not compete. I also showed the higher valuation player is more likely to win in sequential games than in the Cournot game. This is true regardless which player goes first.

Finally, I showed how this analysis can be extended to games of incomplete information. If Player I is unsure of what Player II's value of the prize is, he will still choose a level of expenditure to maximize his expected payoff. The calculations are straightforward.

Rent-seeking When the Prize is a Public Good

The work on rent-seeking so far has ignored any "publicness" in the prize. In this section I look at what happens in a Tullock type model where the players move simultaneously and the prize is to some extent a public good. Another way of thinking of this is to say a person is not indifferent to who wins the prize if he doesn't. This allows us to consider a wide range of problems. For example, if international political or military competition can be modeled this way, we can think of this as forming the basis for a theory of alliances. This can also represent a model of political lobbying where a lobbyist is not indifferent to whose favorite law passes if his doesn't. In order to capture this idea I represent a player's valuation of the prize by a vector instead of a scalar. Each player's valuation of the prize is a vector $\mathbf{v}_i = (v_{i1}, v_{i2}, \dots, v_{in})^T$, where v_{ij} is the value to person i if person j wins the prize, and T indicates the transpose of a matrix. In this section I will consider only imperfectly discriminating contests where the probability person i wins the prize is $p_i(\mathbf{x}) = x_i^r / s_r$, where $s_r \equiv \sum_{j=1}^n x_j^r$, $\mathbf{x} = (x_1, x_2, \dots, x_n)^T$ and x_i is person i 's outlay to influence the contest in his favor. I begin by looking at the special case of $r = 1$. Later I will look at changes in r through numerical examples. We can denote the vector of probabilities by $\mathbf{p}(\mathbf{x}) = (p_1, p_2, \dots, p_n)^T$. Finally, we can

express person i 's expected payoff for the contest as

$$U_i(\mathbf{x}) = \mathbf{v}_i \cdot \mathbf{p}(\mathbf{x}) - x_i.$$

This special case is a nice starting point for this discussion. The model considered by Hillman and Riley can be represented in this formulation with $v_{ij} = 0$ for $i \neq j$. Similarly, Tullock's original model is a special case with $v_{ij} = 0$ for $i \neq j$ and $v_{11} = v_{22} = \dots = v_{nn}$.

In order to see how this model is applied, I will begin by showing how simply Hillman and Riley's result, $\sum_{i=1}^n x_i = [(n-1)/n]\hat{v}$ where \hat{v} is the harmonic mean of the valuations, can be derived using this model. V is the $n \times n$ matrix created by stacking \mathbf{v}_1^T on \mathbf{v}_2^T , and so on. A non-cooperative Nash equilibrium solution can be derived by simultaneously solving the n first order conditions to the system of equations below:

$$\begin{pmatrix} v_{11} & 0 & \dots & 0 \\ 0 & v_{22} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & v_{nn} \end{pmatrix} \begin{pmatrix} \frac{x_1}{s} \\ \frac{x_2}{s} \\ \vdots \\ \frac{x_n}{s} \end{pmatrix} - \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} = \begin{pmatrix} U_1 \\ U_2 \\ \vdots \\ U_n \end{pmatrix}.$$

The first order maximization condition from Player i 's utility function is $v_{ii}(s - x_i)/s^2 = 1$. We can summarize all n first order conditions in the following matrix expression:

$$\begin{pmatrix} v_{11} & 0 & \dots & 0 \\ 0 & v_{22} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & v_{nn} \end{pmatrix} \begin{pmatrix} \frac{s-x_1}{s^2} \\ \frac{s-x_2}{s^2} \\ \vdots \\ \frac{s-x_n}{s^2} \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix}.$$

If we premultiply both sides of the above expression by V^{-1} we can simultaneously solve the first order conditions above. Finding V^{-1} is trivial in this case, and we get

$$\begin{pmatrix} s - x_1 \\ s - x_2 \\ \vdots \\ s - x_n \end{pmatrix} = s^2 \begin{pmatrix} \frac{1}{v_{11}} & 0 & 0 & \dots & 0 \\ 0 & \frac{1}{v_{22}} & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \frac{1}{v_{nn}} \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix}.$$

Summing the n equations we get

$$(n-1)s = s^2 n / \hat{v}$$

which yields Hillman and Riley's result, $s = (n-1)\hat{v}/n$ where $\hat{v} = \frac{n}{\sum_{i=1}^n 1/v_i}$ is the harmonic mean of the valuations. However, this formulation can handle more complicated problems.

An Example which Depends on Distance

As motivation for further analysis, consider the following problem. Suppose three people live in a neighborhood without a street light. They are identical in every respect except where they live. A street light is to be placed somewhere on their street, and each of the residents wants it in his favorite location. However, each person is not indifferent to the location of the street light if he doesn't win the contest. I assume the players' valuations are as follows ($0 \leq \gamma < 1$):

$$\mathbf{v}_1 = \begin{pmatrix} 1 \\ \gamma \\ \gamma^2 \end{pmatrix}; \mathbf{v}_2 = \begin{pmatrix} \gamma \\ 1 \\ \gamma \end{pmatrix}; \mathbf{v}_3 = \begin{pmatrix} \gamma^2 \\ \gamma \\ 1 \end{pmatrix}.$$

In words, each person values the street light equally if it is placed in his favorite place. Each also prefers having it located at his next-door-neighbor's favorite place to having it placed farther away.

The system we want to find the Nash equilibrium for is

$$\begin{pmatrix} 1 & \gamma & \gamma^2 \\ \gamma & 1 & \gamma \\ \gamma^2 & \gamma & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} - \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} U_1 \\ U_2 \\ U_3 \end{pmatrix}.$$

In general, the first order conditions for a problem like this can be summarized by $(J_{n \times n} - V)\mathbf{x} = s^2 \mathbf{1}_n$, where $J_{n \times n}$ is the $n \times n$ matrix which has v_{ii} as every element in the i th row and $\mathbf{1}_n$ is the n column vector of ones. In this case we have the following first order conditions:

$$\begin{pmatrix} 0 & 1-\gamma & 1-\gamma^2 \\ 1-\gamma & 0 & 1-\gamma \\ 1-\gamma^2 & 1-\gamma & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}.$$

Inverting the matrix $(J_{n \times n} - V)$ yields the following:

$$(J_{n \times n} - V)^{-1} = \frac{1}{2(1 - \gamma^2)} \begin{pmatrix} -1 & 1 + \gamma & 1 \\ 1 + \gamma & -(1 + \gamma)^2 & 1 + \gamma \\ 1 & 1 + \gamma & -1 \end{pmatrix}.$$

Performing the appropriate matrix multiplications we get the following:

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \frac{s^2}{2(1 - \gamma^2)} \begin{pmatrix} -1 & 1 + \gamma & 1 \\ 1 + \gamma & -(1 + \gamma)^2 & 1 + \gamma \\ 1 & 1 + \gamma & -1 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}.$$

Now, solving for the x_i s we get

$$x_1^* = x_3^* = \frac{(s^*)^2}{2(1 - \gamma)}, \quad x_2^* = \frac{(s^*)^2}{2}.$$

If we add the x_i^* s and solve for s^* we get $x_1^* = x_3^* = 2(1 - \gamma)/(3 - \gamma)^2$, $x_2^* = 2(1 - \gamma)^2/(3 - \gamma)^2$, and $s^* = (2 - 2\gamma)/(3 - \gamma)$.

Notice how this equilibrium compares with the case where players have scalar valuations. If the street light were a purely private good, the equilibrium level of expenditure would be $2/9$ for each player, and the total expenditure would be $2/3$. This corresponds to this problem with $\gamma = 0$. A simple comparative statics exercise shows as γ increases from zero to one, s decreases. That is, the more "publicness" there is to the street light, the less people will pay to get it in their yard in equilibrium. Also, each player will individually be willing to pay less as γ increases. Another interesting property of this equilibrium is that if $\gamma > 0$, Player I and Player III will always pay more than Player II in equilibrium. To see this, note $x_1^*/x_2^* = x_3^*/x_2^* = 1/(1 - \gamma) > 1$.

In a utilitarian sense, the optimal placement of the street light is in II's yard because the total value to the neighborhood is $1 + 2\gamma$ while if it is placed in either of the other yards the total value of the street light is $1 + \gamma + \gamma^2$. However, the socially optimal placement is now the least likely. In this case, then, the public nature of the prize changes the equilibrium outcome in two conflicting ways. First, there is less rent-seeking expenditure if the prize is a public good than if it were a

purely private good. Second, the socially optimal outcome is less likely when the prize is public in nature. This problem can easily be modified to allow for dissimilar valuations, different forms for the disutility of being farther away from the street light, and more players, but the same qualitative results will continue to hold.

An Example with Common Interests

The above example is certainly not the only kind of problem we can analyze using this framework. As another example, suppose there are three players again, and we can begin by thinking of them as automobile makers. Each has a legislative proposal it favors. Let Player I be General Motors, Player II be Ford, and Player III be Toyota. General Motors has a favorite proposal which it values at one. However, it would rather see Ford's proposal accepted than Toyota's. Ford has a favorite proposal with value one to itself, too, but it would rather see GM's proposal accepted than Toyota's. Toyota also values its own proposal at one, but it cares only that its proposal is accepted and views the other two as equally undesirable. I assume GM and Ford value the proposals made by each other equally, so the problem can be set up in the following way:

$$\begin{pmatrix} 1 & \gamma & 0 \\ \gamma & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} - \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} U_1 \\ U_2 \\ U_3 \end{pmatrix}.$$

The first order conditions give the following system of equations:

$$\begin{pmatrix} 0 & 1-\gamma & 1 \\ 1-\gamma & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}.$$

Again, we get the following expression after some simple algebra:

$$\begin{pmatrix} x_1^* \\ x_2^* \\ x_3^* \end{pmatrix} = \frac{s^2}{2(1-\gamma)} \begin{pmatrix} -1 & 1 & 1-\gamma \\ 1 & -1 & 1-\gamma \\ 1-\gamma & 1-\gamma & -(1-\gamma)^2 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}.$$

If we sum the x_i^* s and solve for s^* , we get $s^* = 2/(3+\gamma)$ with $x_1^* = x_2^* = 2/(3+\gamma)^2$ and $x_3^* = 2(1+\gamma)/(3+\gamma)^2$. Here too, we have the result that an increase in γ

results in less total rent-seeking expenditures and less expenditures by GM and Ford. However, Toyota will now pay more as γ increases. That is, if GM and Ford become more in agreement over the legislation, in equilibrium their individual contributions decrease, but Toyota's increases. However, the increase in Toyota's contribution is less than the combined decrease in GM's and Ford's contributions. This shows as Ford and GM become closer to being in perfect agreement over which proposal is best, it becomes costlier for Toyota to compete. Here again, the fact that the prize is a public good means less rent-seeking expenditure. It also means the agent for whom success is less socially desirable will spend more and become more likely to win in equilibrium. This again will lead to a decrease in social welfare.

The problem I described above can be given other interpretations. Suppose instead of calling our agents GM, Ford, and Toyota, we call them country A, country B, and country C respectively. Then we have a model of military or political competition. Here p_i has the interpretation of being the probability country i wins the contest, or it can be thought of as the proportion of disputed land country i can obtain. At first glance this example poses a puzzle of sorts because countries A and B are less likely to win any competition with country C as their interests become more closely aligned. Each of A and B has a probability of winning the competition of $1/(3 + \gamma)$, and the country C has a probability of winning of $(1 + \gamma)/(3 + \gamma)$. This is an example which supports a conjecture by Olson and Zeckhauser (1968) in their initiatory work on the economic theory of alliances.

“... a decline in the amity, unity, and community of interest among allies need not necessarily reduce the effectiveness of an alliance because the decline in these alliance ‘virtues’ produces a greater ratio of private to social benefits.”

This idea is captured explicitly in the example above. As γ increases the “community of interest” between A and B increases accordingly, but their effectiveness, measured by the probability either wins the competition, decreases. This result is

general in the sense that the probability one interest group wins declines as their interests become more closely aligned in this noncooperative game. The intuition is easy to see. As the prize's "publicness" increases, countries with common interests can "free ride" on each other's contribution. Modeling the competition in this way picks up elements not captured in Olson and Zeckhauser's model. Specifically, here the level of the other players' contributions is not exogenously given but endogenously determined by the expenditures of both groups.

Generalizations

Next, I examine what happens as the number of players increases in this type of model. To do this, I will add players with interests similar to either those of Players I and II or those of Player III from the last example. Suppose there are n_1 agents who have interests in common with Players I or II and n_2 players who share the same concerns as Player III. We can think of the player set as consisting of two partitions. The first partition consists of the set $\{1, 2, \dots, n_1\}$ and the other partition is the set $\{n_1 + 1, n_1 + 2, \dots, n_1 + n_2\}$.

The preferences of individuals in the two groups can be characterized by the set of players with whom they share interests and the amount of publicness there is inside the group. For the first partition, the value of the prize to each player is one if he wins it and γ if another player in the partition wins it. A member of the second partition values winning the prize herself at one and attaches a value of δ to the event someone else in her partition wins it. The utility functions which the players want to maximize are summarized in the following, using matrices for

notational simplicity:

$$\begin{pmatrix} 1 & \gamma & \dots & \gamma & 0 & \dots & 0 & 0 \\ \gamma & 1 & \dots & \gamma & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \gamma & \gamma & \dots & 1 & 0 & \dots & 0 & 0 \\ 0 & 0 & \dots & 0 & 1 & \delta & \dots & \delta \\ 0 & 0 & \dots & 0 & \delta & 1 & \dots & \delta \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 & \delta & \delta & \dots & 1 \end{pmatrix} \begin{pmatrix} \frac{x_1}{s} \\ \vdots \\ \frac{x_{n_1}}{s} \\ \frac{x_{n_1+1}}{s} \\ \vdots \\ \frac{x_{n_1+n_2}}{s} \end{pmatrix} - \begin{pmatrix} x_1 \\ \vdots \\ x_{n_1} \\ x_{n_1+1} \\ \vdots \\ x_{n_1+n_2} \end{pmatrix} = \begin{pmatrix} U_1 \\ \vdots \\ U_{n_1} \\ U_{n_1+1} \\ \vdots \\ U_{n_1+n_2} \end{pmatrix}.$$

The first order conditions are derived in the same way as in the previous examples.

$$\begin{pmatrix} 0 & 1-\gamma & \dots & 1-\gamma & 1 & 1 & \dots & 1 \\ 1-\gamma & 0 & \dots & 1-\gamma & 1 & 1 & \dots & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1-\gamma & 1-\gamma & \dots & 0 & 1 & 1 & \dots & 1 \\ 1 & 1 & \dots & 1 & 0 & 1-\delta & \dots & 1-\delta \\ 1 & 1 & \dots & 1 & 1-\delta & 0 & \dots & 1-\delta \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \dots & 1 & 1-\delta & 1-\delta & \dots & 0 \end{pmatrix} \begin{pmatrix} \frac{x_1}{s^2} \\ \vdots \\ \frac{x_{n_1}}{s^2} \\ \frac{x_{n_1+1}}{s^2} \\ \vdots \\ \frac{x_{n_1+n_2}}{s^2} \end{pmatrix} = \begin{pmatrix} 1 \\ \vdots \\ 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix}$$

A matrix of the form $(J_{(n_1+n_2) \times (n_1+n_2)} - V)$ can be inverted for any $\gamma, \delta \neq 1$ and $n_1, n_2 \geq 1$. I will use the notation I_n for the $n \times n$ identity matrix. We can then solve for the vector of equilibrium individual contributions, \mathbf{x}^* .

$$\mathbf{x}^* = (s^*)^2 (J_{(n_1+n_2) \times (n_1+n_2)} - V)^{-1} \cdot \mathbf{1}_{(n_1+n_2)}$$

In the above expression we have

$$(J_{(n_1+n_2) \times (n_1+n_2)} - V)^{-1} = \begin{pmatrix} W_{11} & W_{12} \\ W_{21} & W_{22} \end{pmatrix}$$

with W_{ij} defined below. Here $\mathbf{1}_{n \times m}$ is an $n \times m$ matrix of ones.

$$W_{11} = -\frac{1}{(1-\gamma)} \left[I_{n_1} - \frac{(1-\gamma)(1-\delta)(n_2-1) - n_2}{(n_1-1)(n_2-1)(1-\gamma)(1-\delta) - n_1 n_2} \cdot J_{n_1 \times n_1} \right]$$

$$W_{ij} = \frac{-1}{(n_1-1)(n_2-1)(1-\gamma)(1-\delta) - n_1 n_2} \cdot \mathbf{1}_{n_i \times n_j}, \quad i \neq j$$

$$W_{22} = -\frac{1}{(1-\delta)} \left[I_{n_2} - \frac{(1-\gamma)(1-\delta)(n_1-1) - n_1}{(n_1-1)(n_2-1)(1-\gamma)(1-\delta) - n_1 n_2} \cdot J_{n_2 \times n_2} \right]$$

From the above equation we can derive expressions for equilibrium total contributions, s^* , and the equilibrium individual contributions. I will use x_1^* to denote the equilibrium contributions of the first n_1 players and x_2^* for the second group. We have

$$s^* = \frac{n_1 n_2 - (n_1 - 1)(n_2 - 1)(1 - \gamma)(1 - \delta)}{n_1 [n_2 - (1 - \delta)(n_2 - 1)] + n_2 [n_1 - (1 - \gamma)(n_1 - 1)]}; \quad (3)$$

$$x_1^* = (s^*)^2 \left[\frac{n_2 - (1 - \delta)(n_2 - 1)}{n_1 n_2 - (n_1 - 1)(n_2 - 1)(1 - \gamma)(1 - \delta)} \right]; \quad (4)$$

$$x_2^* = (s^*)^2 \left[\frac{n_1 - (1 - \gamma)(n_1 - 1)}{n_1 n_2 - (n_1 - 1)(n_2 - 1)(1 - \gamma)(1 - \delta)} \right]. \quad (5)$$

These expressions can be solved for x_1^* and x_2^* in terms of the parameters of the game. It is interesting to note everyone will choose to contribute in equilibrium. This is not generally true if we allow "own-valuations" (v_{ii}) to differ among players. It is easy to see Olson and Zeckhauser's conjecture holds true from the above equilibrium outlays.

With these in mind, we can derive some simple comparative statics results. As either of the publicness parameters increases the sum of the expenditures will decrease. That is, $\partial s^*/\gamma, \partial s^*/\delta < 0$. Also, the representative contribution of a member of an interest group will decrease as its own publicness parameter increases. More formally, we have $\partial x_1^*/\gamma, \partial x_2^*/\delta < 0$. Perhaps the easiest way to see what happens as we allow parameters to vary is to look at the ratio $n_1 x_1^*/n_2 x_2^*$, which is also the ratio of the probability the first group wins to the probability the second group wins. It is easy to see this ratio increases as δ increases, and it decreases as γ increases. Also, note as the number of players in a group increases, that group's probability of winning the contest increases. Even though having more players in a group with a common interest allows more opportunity for free riding among the participants, having more players with your interests helps your chance of winning as long as the players' interests are not identical (γ or $\delta \neq 1$).

Now I will analyze a game where players have different "own-valuations" for the prize. I will show how the equilibrium strategies change in this model if we allow the prize to have a different own-value for each player. To keep the model simple, I will analyze a three player contest where a player in an interest group values the prize won by another in his group as a constant proportion of his "own-valuation." This can be summarized in the following valuation matrix

$$V = \begin{pmatrix} v_{11} & \gamma v_{11} & 0 \\ \gamma v_{22} & v_{22} & 0 \\ 0 & 0 & v_{33} \end{pmatrix}.$$

Determining the first order conditions for this problem is slightly different than in previous examples because we allow $v_{ii} \neq v_{jj}$. However, the same general rule continues to hold. Hence, the first order conditions the players simultaneously solve are

$$\begin{pmatrix} 0 & v_{11}(1-\gamma) & v_{11} \\ v_{22}(1-\gamma) & 0 & v_{22} \\ v_{33} & v_{33} & 0 \end{pmatrix} \begin{pmatrix} \frac{x_1}{s^2} \\ \frac{x_2}{s^2} \\ \frac{x_3}{s^2} \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}.$$

Solving the above simultaneous equations yields the following equilibrium outlays if all players participate:

$$x_1^* = \frac{2v_{11}v_{22}v_{33} \cdot [v_{11}v_{33} - v_{22}v_{33} + v_{11}v_{22}(1-\gamma)]}{[v_{11}v_{22} + v_{22}v_{33} + v_{11}v_{33} + v_{11}v_{22}\gamma]^2(1-\gamma)} \quad (6)$$

$$x_2^* = \frac{2v_{11}v_{22}v_{33} \cdot [v_{22}v_{33} - v_{11}v_{33} + v_{11}v_{22}(1-\gamma)]}{[v_{11}v_{22} + v_{22}v_{33} + v_{11}v_{33} + v_{11}v_{22}\gamma]^2(1-\gamma)} \quad (7)$$

$$x_3^* = \frac{2v_{11}v_{22}v_{33} \cdot [v_{22}v_{33}(1-\gamma) + v_{11}v_{33}(1-\gamma) - v_{11}v_{22}(1-\gamma)^2]}{[v_{11}v_{22} + v_{22}v_{33} + v_{11}v_{33} + v_{11}v_{22}\gamma]^2(1-\gamma)} \quad (8)$$

$$s^* = \frac{2v_{11}v_{22}v_{33}}{[v_{11}v_{22} + v_{22}v_{33} + v_{11}v_{33} + v_{11}v_{22}\gamma]} \quad (9)$$

It is important to note the above outlays are in equilibrium only when all players participate. Next I will discuss the conditions for participation by the players. Certainly if the numerator in any of the above expressions is negative, a player will choose not to play. I will examine the conditions for a player to want to enter the game. It turns out the same conditions which assure strictly positive contributions

in equations (6) - (8) have a very nice interpretation. Those conditions are the answer to the question, "If two of the players are already active in this game, when will the third player want to make a marginal contribution rather than not participate?" We also want to find out if it makes any difference in which order the players join the game in determining the final set of players, and under what conditions will all three agents participate in this game. I answer these questions next.

Suppose first, Players I and II are playing the game against each other, and Player III must decide whether or not to play. Certainly, she will not want to contribute if her marginal utility would be negative.⁸ That is, III will want to contribute if $\partial U_3(x_1, x_2, x_3)/\partial x_3 > 0$. In this example, the condition for participation is

$$\frac{\partial U_3(x_1, x_2, x_3)}{\partial x_3} = \frac{(x_1 + x_2)v_{33}}{s^2} - 1 = \frac{(s - x_3)v_{33}}{s^2} - 1 > 0.$$

Since we are evaluating this derivative at $x_3 = 0$, s is the sum of the equilibrium outlays for Players I and II if they play the game by themselves, which in this case is $s = v_{11}v_{22}(1 - \gamma)/(v_{11} + v_{22})$. Hence, we have the result that Player III will want to participate if

$$v_{33} > \frac{v_{11}v_{22}(1 - \gamma)}{v_{11} + v_{22}}.$$

Consider the above expression with $\gamma = 0$. This is exactly what Hillman and Riley proved for the scalar valuation case. In other words, Player III will choose to participate if his valuation of the prize is greater than one half the harmonic mean of the valuations of the first two players.

If we allow publicness in the prize, the condition for another player's participation is modified somewhat. However, we get the same result as Hillman and Riley if we think of each person's valuation as the incremental valuation to him if he wins

⁸ See Hillman and Riley (1987) for a discussion of participation in this game with scalar valuations.

the prize over what he gets if someone else wins it. Note the same condition which produces a positive outlay for Player III in equation (8) is the same condition which gives him positive utility on the margin. It is also interesting to note as γ increases, the minimum valuation for Player III's active participation decreases. The intuition is Players I and II will compete less vigorously as γ increases, so a lower valuation Player III can improve his payoff by joining the contest.

If we consider the entry conditions for one of the first two players, the analysis is only slightly less straightforward because of the public nature of the prize. Suppose Players I and III are playing the game. Now we ask ourselves when II would want to play. Using the same type of analysis as we used above, we know Player II will want to contribute if $\partial U(x_1, x_2, x_3)/\partial x_2 > 0$. We know from the first order conditions this requirement is

$$\frac{\partial U(x_1, x_2, x_3)}{\partial x_2} = \frac{v_{22}(1-\gamma)x_1}{s^2} + \frac{v_{22}x_3}{s^2} - 1 > 0.$$

This can be simplified somewhat by multiplying through by s . This gives us

$$\frac{v_{22}(1-\gamma)x_1}{s} + \frac{v_{22}x_3}{s} > s.$$

Rearranging this we have

$$\frac{v_{22}(s-x_2)}{s} - \frac{\gamma v_{22}x_1}{s} > s.$$

Players I and III will be playing their equilibrium strategies for a two player game where Player I values the prize at v_{11} and III assigns a value of v_{33} to the prize. Recall the equilibrium contributions in the two player game are $x_1 = v_{11}^2 v_{33}/(v_{11} + v_{33})^2$, $x_3 = v_{11} v_{33}^2/(v_{11} + v_{33})^2$, and $s = v_{11} v_{33}/(v_{11} + v_{33})$. Now we can substitute and evaluate the derivative at $x_2 = 0$. This yields

$$v_{22} \left(1 - \frac{\gamma v_{11}}{(v_{11} + v_{33})} \right) > \frac{v_{11} v_{33}}{(v_{11} + v_{33})}$$

as the condition for Player II's participation. This means II will want to make a contribution if

$$v_{22} > \frac{v_{11}v_{33}}{v_{11}(1-\gamma) + v_{33}}. \quad (10)$$

Again this is the same condition which generates a positive contribution in equation (7). Analyzing the situation with Player I inactive and Players II and III participating is symmetric. If $\gamma = 0$, this is again the same result Hillman and Riley obtained. If we think about condition (10) above, we can see Player II will have to value the prize more in order to participate if v_{11} , v_{33} , or γ increase. The intuition is Player I will make a larger contribution in the two player game if v_{11} increases, and Player II can "free ride" on it. If γ increases, Player II will want to contribute less because he gains more utility from Player I's outlay. Also, an increase in v_{33} will make a contribution by Player II less profitable so he will have to value winning the prize more before he chooses to be active in the game.

Now I can describe the equilibrium strategies in this game. First, however, I will introduce the notation that the minimum valuation for Player k to actively participate in the game is \bar{v}_{kk} . That is,

$$\bar{v}_{ii} = \frac{v_{jj}v_{33}}{v_{jj}(1-\gamma) + v_{33}} \quad \text{for } i, j \in \{1, 2\} \quad \text{and } i \neq j;$$

$$\bar{v}_{33} = \frac{v_{11}v_{22}(1-\gamma)}{v_{11} + v_{22}}.$$

The player's equilibrium strategies are the following:

$$\begin{aligned}
 x_1^* &= \begin{cases} \frac{2v_{11}v_{22}v_{33} \cdot [v_{11}v_{33} - v_{22}v_{33} + v_{11}v_{22}(1-\gamma)]}{[v_{11}v_{22} + v_{22}v_{33} + v_{11}v_{33} + v_{11}v_{22}\gamma]^2(1-\gamma)} & \text{if } v_{ii} > \bar{v}_{ii} \text{ for } i \in \{1, 2, 3\}; \\ \frac{v_{11}^2 v_{33}}{(v_{11} + v_{33})^2} & \text{if } v_{22} \leq \bar{v}_{22}; \\ \frac{v_{11}^2 v_{22}(1-\gamma)}{(v_{11} + v_{22})^2} & \text{if } v_{33} \leq \bar{v}_{33}; \\ 0 & \text{if } v_{11} \leq \bar{v}_{11}. \end{cases} \\
 x_2^* &= \begin{cases} \frac{2v_{11}v_{22}v_{33} \cdot [v_{22}v_{33} - v_{11}v_{33} + v_{11}v_{22}(1-\gamma)]}{[v_{11}v_{22} + v_{22}v_{33} + v_{11}v_{33} + v_{11}v_{22}\gamma]^2(1-\gamma)} & \text{if } v_{ii} > \bar{v}_{ii} \text{ for } i \in \{1, 2, 3\}; \\ \frac{v_{22}^2 v_{33}}{(v_{22} + v_{33})^2} & \text{if } v_{11} \leq \bar{v}_{11}; \\ \frac{v_{22}^2 v_{11}(1-\gamma)}{(v_{11} + v_{22})^2} & \text{if } v_{33} \leq \bar{v}_{33}; \\ 0 & \text{if } v_{22} \leq \bar{v}_{22}. \end{cases} \\
 x_3^* &= \begin{cases} \frac{2v_{11}v_{22}v_{33} \cdot [v_{22}v_{33}(1-\gamma) + v_{11}v_{33}(1-\gamma) - v_{11}v_{22}(1-\gamma)^2]}{[v_{11}v_{22} + v_{22}v_{33} + v_{11}v_{33} + v_{11}v_{22}\gamma]^2(1-\gamma)} & \text{if } v_{ii} > \bar{v}_{ii} \text{ for } i \in \{1, 2, 3\}; \\ \frac{v_{33}^2 v_{22}}{(v_{22} + v_{33})^2} & \text{if } v_{11} \leq \bar{v}_{11}; \\ \frac{v_{33}^2 v_{11}}{(v_{11} + v_{33})^2} & \text{if } v_{22} \leq \bar{v}_{22}; \\ 0 & \text{if } v_{33} \leq \bar{v}_{33}. \end{cases}
 \end{aligned}$$

Note at least two players will be active in this game because it is a simultaneous move game.⁹ The above equilibrium strategies yield the answers we were looking for. In this model it makes no difference in which order the players join in the competition. The conditions for players switching from being inactive to active participants are identical to the conditions for them contributing more than zero if they are already in the game. The required condition for all three players being active is obvious from the equilibrium strategies.

If we assume all players are active participants and the set of active players doesn't change for a small change in a parameter value, we can make some observations about the comparative statics of this game. It should not be surprising to see as γ increases, x_1^* , x_2^* , and s^* decrease and x_3^* increases. Again, it will prove illuminating if we consider the ratio

$$\frac{x_1^* + x_2^*}{x_3^*} = \frac{2v_{11}v_{22}}{v_{33}(v_{11} + v_{22}) - v_{11}v_{22}(1-\gamma)}.$$

It is clear the probability either Player I or Player II wins decreases as γ increases or as v_{33} increases. An increase in either v_{11} or v_{22} will increase this ratio. These

⁹ For a more detailed explanation, see Hillman and Riley (1987).

results are not unexpected and verify the results obtained in the examples at the beginning of this section hold true generally in these games.

Changes in the Probability Function

Now I will briefly look at the effects of changing the probability function. Specifically, I will consider what happens as r changes using Tullock's formulation. I will look only at the street light and automobile maker examples.

Recall according to this setup we have

$$p_i(x_1, x_2, x_3) = \frac{x_i^r}{\sum_{j=1}^3 x_j^r}.$$

If we denote $x_1^r + x_2^r + x_3^r = s_r$, we can restate the street light problem as

$$\begin{pmatrix} 1 & \gamma & \gamma^2 \\ \gamma & 1 & \gamma \\ \gamma^2 & \gamma & 1 \end{pmatrix} \begin{pmatrix} \frac{x_1^r}{s_r} \\ \frac{x_2^r}{s_r} \\ \frac{x_3^r}{s_r} \end{pmatrix} - \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} U_1 \\ U_2 \\ U_3 \end{pmatrix}.$$

The above equations yield the following first order conditions:

$$\begin{pmatrix} 0 & rx_1^{r-1}(1-\gamma) & rx_1^{r-1}(1-\gamma^2) \\ rx_2^{r-1}(1-\gamma) & 0 & rx_2^{r-1}(1-\gamma) \\ rx_3^{r-1}(1-\gamma^2) & rx_3^{r-1}(1-\gamma) & 0 \end{pmatrix} \begin{pmatrix} \frac{x_1^r}{s_r} \\ \frac{x_2^r}{s_r} \\ \frac{x_3^r}{s_r} \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}.$$

This problem cannot be solved analytically for a closed-form solution. However, I solved it numerically for certain parameter values, and I present equilibrium contributions for various r values and levels of publicness in table 4.1. The first value in each cell is the x_1^* or x_3^* , and the second value is x_2^* .

Table 4.1 — Equilibrium Contributions in the Street Light Game

$r \backslash \gamma$	0	.25	.5	.75
.5	.111, .111	.095, .081	.072, .052	.041, .025
1	.222, .222	.198, .149	.16, .08	.099, .025
5	1.11, 1.11	.848, .915	.542, .625	.248, .304

Now we can look at the auto maker or three country alliance problem. The utility functions in this problem are summarized by the following:

$$\begin{pmatrix} 1 & \gamma & 0 \\ \gamma & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \frac{x_1}{s_r} \\ \frac{x_2}{s_r} \\ \frac{x_3}{s_r} \end{pmatrix} - \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} U_1 \\ U_2 \\ U_3 \end{pmatrix}.$$

These yield the following first order conditions:

$$\begin{pmatrix} 0 & rx_1^{r-1}(1-\gamma) & rx_1^{r-1} \\ rx_2^{r-1}(1-\gamma) & 0 & rx_2^{r-1} \\ rx_3^{r-1} & rx_3^{r-1} & 0 \end{pmatrix} \begin{pmatrix} \frac{x_1}{s_r} \\ \frac{x_2}{s_r} \\ \frac{x_3}{s_r} \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}.$$

This problem cannot be solved analytically either. Once again, I present equilibrium contributions which were calculated numerically for various r values and levels of publicness. In table 4.2, the first value in each cell is the equilibrium contribution by Players I and II, and the second value is the equilibrium contribution of Player III.

Table 4.2 — Equilibrium Contributions in the Auto Maker Game

$r \backslash \gamma$	0	.25	.5	.75
.5	.111, .111	.095, .148	.082, .184	.071, .218
1	.222, .222	.189, .237	.163, .245	.142, .249
5	1.11, 1.11	.941, .99	.816, .885	.711, .795

Tables 4.1 and 4.2 do not yield surprising results. As r increases there is more rent-seeking expenditure which is the same result Tullock obtained. It is interesting to note as r increases there is a point where the person in the middle house contributes more than each of the players on the ends in the street light problem. Also, we can see as the degree of publicness in the prize increases there will be less rent-seeking. An interesting result from table 4.1 is total expenditures may be less than the value of the prize to a single player if γ is close to one, even for

large values of r . This contrasts with Tullock's results when the prize is a private good. I should note, however, if the sum of the three contributions from tables 4.1 and 4.2 exceed the common valuation, these are not equilibrium outlays. In these instances no equilibrium exists in this game.¹⁰

Summary

The framework developed here can handle a large number of problems in which there is publicness in the prize over which competition arises. We have been able to analyze specific situations in this section because of the symmetry of the problems. However, this framework is much more general. If we can identify the valuation matrix we will generally be able to solve for the equilibrium contributions which result from the non-cooperative game if we use Tullock's probability function with $r = 1$.

It is clear adding "publicness" to the rent-seeking problem means less expenditure for the prize. In the cases I have examined, the prize has a common value. This analysis serves two basic purposes. First, it shows if we consider the public nature of the prizes in rent-seeking games there may be less socially wasteful expenditure than Tullock first argued. Second, it provides a useful framework for analyzing political and military competition.

An excellent example of where this sort of analysis can be useful is in analyzing expenditures among countries with common goals. This provides evidence in support of Olson and Zeckhauser's paradoxical conclusion which we described in this section. The fundamental difference between this model and what has been done previously in the theory of alliances is the models of Olson and Zeckhauser and others are partial equilibrium in nature.¹¹ A country's demand for defense or security expenditures is derived from an exogenously given threat in the other models.

¹⁰ See Tullock (1980) for the details of this problem and his results.

¹¹ Some of the more important papers in this literature are Sandler (1977), Sandler and Forbes (1980), Murdoch and Sandler (1982, 1984).

Here we begin to address the problem which arises when we try to see how enemy expenditures are derived. We have shown changes in valuations or publicness alter the equilibrium outlays for all players. As far as I know, no other models have been developed which address variations both inside and outside the interest groups.

This model of competitive and cooperative rent-seeking behavior can easily be used as a model of political activity. Now, instead of nations, the players are lobbying organizations. This model provides insight into how players choose the level of expenditure for rent-seeking activities. If the benefits are public goods, we can expect to see little expenditure. However, if the benefits are specific to a player, we should see much more rent-seeking behavior in equilibrium.

Conclusions

In this essay I have considered two basic extensions to the rent-seeking models originated by Tullock and refined by others. I analyzed a rent-seeking game which I assumed was played sequentially, and I allowed players to have valuations described by vectors in an imperfectly discriminating game. The last topic extended work primarily done by Hillman and Riley.

In the sequential game we found the equilibrium total expenditures can be greater or less than those in the simultaneous move game depending on which player goes first. This has implications for the holder of the contest. We also found, contrary to what we see in duopoly theory, it may not be strictly beneficial to go first in these contests. If both players have the same valuation for the prize, it makes no difference who goes first. I also showed how the game can be extended to situations of incomplete information.

In the third section, I described a model where there is some degree of publicness in the prize. In other words, a player is not indifferent as to who wins the prize if he doesn't. I looked at specific formulations which could be handled analytically. In this framework we can verify Olson and Zeckhauser's conjecture that

a decrease in the coincidence of interests may not decrease the effectiveness of an alliance. In other words, as the "publicness" parameter increased for one common interest group in these models, that interest group became less likely to win. I also described some numerically generated results based on different valuations and probability functions.

This essay examined rent-seeking behavior in situations which haven't been considered explicitly before. We can see the amount of rent-seeking expenditure depends on the structure of the game. The framework developed here can be used to analyze many different problems of this type.

Bibliography

- Bergstrom, Theodore and Hal R. Varian. "Competitive and Cooperative Rent-seeking." An unpublished research proposal, The University of Michigan, 1987.
- Bhagwati, Jagdish. "Lobbying and Welfare." *Journal of Public Economics* 14 (1980):355-363.
- Bhagwati, Jagdish. "Directly-unproductive Profit-seeking (DUP) Activities." *Journal of Political Economy* 90 (1982):988-1002.
- Bhagwati, J. N. and T. N. Srinivasan. "The Welfare Consequences of Directly-unproductive Profit-seeking (DUP) Lobbying Activities: Prices versus Quantity Distortions." *Journal of International Economics* 13 (1982):33-44.
- Bhagwati, J. N., R. A. Brecher and T. Hatta. "The Generalized Theory of Transfers and Welfare: Exogenous (Policy-imposed) and Endogenous (Transfer-induced) Distortions." *Quarterly Journal of Economics* 100 (1985):697-714.
- Binmore, Kenneth G. "Modeling Rational Players, Part I." *Economics and Philosophy* 3 (1987):179-214.
- Buchanan, James. "Rent Seeking and Profit Seeking" In *Toward a Theory of the Rent-seeking Society*, ed. James Buchanan, Robert Tollison, and Gordon Tullock, 3-15. College Station: Texas A & M University Press, 1980.
- Buchanan, James, Robert Tollison and Gordon Tullock, eds. *Toward a Theory of the Rent-seeking Society*. College Station: Texas A & M University Press, 1980.
- Congleton, Roger. "Competitive Process, Competitive Waste, and Institutions." In *Toward a Theory of the Rent-seeking Society*, ed. James Buchanan, Robert Tollison, and Gordon Tullock, 153-179. College Station: Texas A & M University Press, 1980.
- Corcoran, William J. "Long-run Equilibrium and Total Expenditures in Rent-seeking." *Public Choice* 43 (1984):89-94.
- Corcoran, William J. and Gordon V. Karels. "Rent-seeking Behavior in the Long-run." *Public Choice* 46 (1985):227-245.
- Harberger, Arnold. "Monopoly and Resource Allocation." *American Economic Review* 44 (May 1954):77-87.

- Higgins, Richard, William Shugart, and Robert Tollison. "Free Entry and Efficient Rent-seeking." *Public Choice* 46 (1985):246-258.
- Hillman, Ayre and Eliakim Katz. "Risk Averse Rent Seekers and the Social Cost of Monopoly." *Economic Journal* 94 (1984):104-110.
- Hillman, Ayre and John Riley. "Politically Contestable Rents and Transfers." Paper presented at the World Bank Conference on Political Economy: Theory and Policy Implications, Washington D. C., June 17-19, 1987.
- Krueger, Anne O. "The Political Economy of the Rent-seeking Society." *American Economic Review* 64 (1974):291-303.
- Murdoch, James C. and Todd Sandler. "A Theoretical and Empirical Analysis of NATO." *Journal of Conflict Resolution* 26 (June 1982):237-263.
- Murdoch, James C. and Todd Sandler. "Complementarity, Free Riding, and the Military Expenditures of NATO Allies." *Journal of Public Economics* 25 (1984):83-101.
- Olson, Mancur and Richard Zeckhauser. "An Economic Theory of Alliances." *Review of Economics and Statistics* 48 (1968):266-279.
- Rogerson, William P. "The Social Cost of Monopoly and Regulation: A Game Theoretic Analysis." *Bell Journal of Economics* 13 (1982):391-401.
- Posner, Richard. "The Social Cost of Monopoly and Regulation." *Journal of Political Economy* 83 (1975):83:807-827, Reprinted in *Toward a Theory of the Rent-seeking Society*, ed. James Buchanan, Robert Tollison, and Gordon Tullock. College Station: Texas A & M University Press, 1980.
- Sandler, Todd. "Impurity of Defense: An Application to the Economics of Alliances." *Kyklos* 30 (1977):443-460.
- Sandler, Todd and John F. Forbes. "Burden Sharing, Strategy, and the Design of NATO." *Economic Inquiry* 18 (1980):425-444.
- Tullock, Gordon. "The Welfare Costs of Tariffs, Monopolies, and Theft." *Western Economic Journal* 5 (1967):224-232, Reprinted in *Toward a Theory of the Rent-seeking Society*, ed. James Buchanan, Robert Tollison, and Gordon Tullock. College Station: Texas A & M University Press, 1980.
- Tullock, Gordon. "Efficient Rent-seeking." In *Toward a Theory of the Rent-seeking Society*, ed. James Buchanan, Robert Tollison, and Gordon Tullock, 97-112. College Station: Texas A & M University Press, 1980.
- Varian, Hal. "Sequential Provision of Public Goods." Center for Research on Economic and Social Theory (CREST) Working Paper, The University of Michigan (Nov 1989).