Simulating International Cooperation under Uncertainty

THE EFFECTS OF SYMMETRIC AND ASYMMETRIC NOISE

CURTIS S. SIGNORINO
Harvard University

The repeated prisioner's dilemma is representative of a broad range of situations in international security and trade. This article examines the effects of asymmetric noise on the emergence and maintenance of cooperation under such conditions. The results show that positive and negative asymmetric noise have very different effects on strategy performance. For forgiving strategies, positive noise provides a stimulus out of perpetual defection or unsynchronized retaliations, but also opens them to exploitation. For provocable strategies, negative noise triggers unsynchronized retaliations or perpetual defection, although this may be tempered by generosity and contrition. The effects of neutral noise reflects the signature of each asymmetric noise type. Of the strategies examined, contrite tit-for-tat (CTFT) is generally one of the best performers in both homogeneous and heterogeneous systems. Moreover, one generally sees the evolutionary models moving from heterogeneous bilateral interaction to cooperative norms of behavior, often including or even dominated by CTFT.

The repeated prisoner's dilemma (RPD) is representative of a broad range of situations in international relations, both in security and in international political economy. Of particular interest to international relations theorists is the evolution of cooperation among states under such conditions. It is often theoretically informative to construct hypothetical but plausible international contexts to examine whether cooperation is possible in such contexts and, if cooperation is possible, also to examine what characteristics of state behavior are necessary for the emergence and maintenance of that cooperation. Along these lines, this article examines how international cooperation may emerge and best be maintained in a world characterized by a "noisy" RPD.

AUTHOR'S NOTE: I would like to thank James Alt, Jonathan Bendor, Bruce Bueno de Mesquita, Marc Busch, Robert Keohane, Gary King, Lisa Martin, Edward Schwartz, and the participants of the Harvard Government Department's Rational Choice Discussion Group for their valuable comments. I would also like to thank the Center for International Affairs, Harvard University, for its support during the writing of this article.

JOURNAL OF CONFLICT RESOLUTION, Vol. 40 No. 1, March 1996 152-205

Studying cooperation requires that one distinguish it from other international conditions. Keohane (1984, 51-4) identified three distinct situations: harmony, cooperation, and discord. Under conditions of harmony, each actor's policies automatically facilitate the attainment of the other's goals. In contrast, conditions of discord arise when an actor's goals are hindered or rendered unattainable because of the policies of another actor and when that actor refuses to modify those policies. As Oye (1986, 6) stated, if there is harmony, there is no need for cooperation, and if deadlock exists, there is no possibility for cooperation. In seeking to maximize their individual security, nations are tempted to act deceptively and to cheat. As a result, discord generally reigns over harmony as the initial result of independent action (Keohane 1984, 54). However, somewhere in between exist conditions wherein cooperation may emerge, wherein actors have both conflicting and complementary interests, and wherein mutually beneficial outcomes do not follow automatically but neither are rendered unattainable.

The prisoner's dilemma is a simple but powerful model often used to analyze the evolution of cooperation under anarchy. It represents an important problem concerning cooperation: myopic pursuit of self-interest can be disastrous (or at the very least inefficient), whereas cooperation can lead to a mutually more beneficial outcome. The story motivating the prisoner's dilemma is that two prisoners are being interrogated individually for a crime in which they took part. The decision each must make is whether to cooperate with each other by not confessing to the authorities or to defect by "ratting" on the other. The incentives to cooperate or defect are such that a prisoner receives the reward payoff R if both cooperate, the punishment payoff P if both defect, the temptation payoff T if one defects when the other cooperates, and the sucker's payoff S if one cooperates when the other defects. Because the incentives are structured such that $T > R > P > S$, in a single-shot game, each's myopic self-interest leads to an equilibrium of mutual defection. As implied earlier, mutual defection is not pareto optimal, and both players would have done better had they both cooperated. However, the situation is not necessarily so disappointing under repeated play. If the actors believe their interaction will continue infinitely into the future or if they believe their interaction will occur over a finite period of unknown length, but with a high probability of future interaction, then cooperation may emerge.

Although it has been suggested that games of stag, chicken, and deadlock predominate in world politics (Axelrod and Keohane 1986, 231), prisoner's dilemma situations are not uncommon. In arms-reduction treaty compliance and nuclear nonproliferation, superpowers must decide whether to cooperate or defect in their production of additional weapons or establishment of nuclear arms programs (Downs, Rocke, and Siverson 1986, 128-32; see also Downs and Rocke 1990). Axelrod (1984, 73) described the "live-and-let-live" system in World War I trench warfare in which, because of the static nature and repeatedness of the interactions, both sides developed a system of cooperation in which each would aim so as to miss the other—but would retaliate if the other side aimed to hit. In international trade, two industrial nations may erect trade barriers to each other's exports. Both nations have an incentive to retain their trade barriers, but in doing so, their outcome is worse than if they had cooperated (Axelrod 1984, 7). Milgrom, North, and Weingast (1990) showed how the availability of information via the law merchant in the Champagne Fairs trade system turned single interactions by traders into what amounted to repeated interactions, which, in turn, led to cooperative trading.[1] Conybeare (1986, 152-8) argued that the Anglo-Hanse trade wars from 1300 to 1700 involved prisoner's dilemma incentives for the two parties. And Pahre (1994) took examples surrounding the Marshall Plan to show how multilateral cooperation may obtain in an iterated prisoner's dilemma. Thus, although they may not be the predominant form of interaction, prisoner's dilemmas do exist in world politics. Moreover, with the prisoner's dilemma model, both political-economic and international security issues may be studied using the same analytical framework (Axelrod and Keohane 1986, 231).

Over the last decade, research in cooperation theory has increased greatly. In one of the seminal works in cooperation theory, Axelrod (1984) showed that a surprisingly simple strategy based on immediate reciprocity outperformed all the other strategies submitted for his tournament. This strategy, tit-for-tat (TFT), starts by cooperating and thereafter plays whatever the opponent played the last move. According to Axelrod, TFT includes the four properties that make a decision rule successful: avoiding unnecessary conflict by cooperating when the other player cooperates, retaliating for unprovoked defection, forgiveness after retaliating, and clarity of behavior, allowing the other player to adapt to one's pattern of behavior (Axelrod 1984, 20).

However, Axelrod (1984) noted that his analysis omitted verbal communication, direct influences of third parties, problems with implementing a chosen move, and uncertainty about what the other player did on the previous move. More recent work has shown that the latter two have dramatic effects on the performance of TFT. For example, Molander (1985); Downs, Rocke, and Siverson (1985); and Bendor (1987, 1993) all analytically examined TFT in a RPD with noise and found that TFT's performance deteriorates dramatically at even very low noise levels. Molander showed that performance can be improved by adding unconditional generosity to TFT. This is done by randomly playing cooperate with probability p and TFT with probability $(1 - p)$. Because the generosity is not deterministic, Molander suggested this avoids the pitfalls that make strategies like tit-for-two-tats highly exploit-

1. For an interesting analysis of how noise affects the law merchant institution, see Schwartz (1993).

able. Bendor, Kramer, and Stout (1991) conducted a computer tournament in a noisy RPD and similarly found that TFT performs poorly, whereas generous strategies are effective. More recently, Wu and Axelrod (1995) conducted a tournament with 1% noise and found that generosity and contrition help restore cooperation, contingent on the strategy being faced.

This study differs from the preceding research in two main ways. First, all prior research incorporating noisy prisoner's dilemmas has done so using only neutral noise—that is, disturbances that are symmetric and centered about zero. In a prisoner's dilemma, an assumption of neutral noise implies there is an equal probability that an intended cooperation by one party is "realized" by the other party as a defection and that an intended defection is realized as a cooperation. In this study, I analyze the effects of both symmetric and asymmetric noise on the performance of the strategies employed here.[2] I argue that understanding the effects of different noise distributions is especially important when examining international cooperation, because many situations of human interaction are more representative of biased noise environments (especially negatively biased ones) than neutrally noisy environments.[3] In a negatively biased noise environment, the probability of an intended cooperation being realized as a defection is higher than the probability of an intended defection being realized as a cooperation. For example, the lack of control of bureaucratic administration or the requirements for congressional approval might cause the failed implementation of an executive's intended foreign policy action—but will generally not generate a cooperative action out of an executive's intended "defection." In the dismantling of weapons systems or in abiding by treaties, nations have difficulty verifying commitments and may view their partner's actions with an inherent negative bias; witness, for example, the intense perceptual bias of the United States and the former Soviet Union in dealing with each other. In regional conflicts, during a mutually beneficial cease-fire, lack of control over factions in one's army may result in shelling that is unintended by the regional leaders, causing the two sides to fall back into retaliations. In international trade, produce and other goods may be damaged in transit. Finally, although the results presented in this study may have application across a number of disciplines, my focus here is on how symmetric and asymmetric noise affect the prospects for cooperation in international politics. Assuming that policy makers operate in a boundedly rational world of fairly heterogeneous modes

2. With the exception of Bendor (1993), the literature generally has not differentiated between implementation and perception errors, combining them into a definition of noise that refers to any disturbance that causes the perceived move to differ from the intended move. Recognizing this, I still take the more general approach, leaving the refinement of analyses of asymmetric perception error and asymmetric implementation error for future research.

3. Throughout this article, I use the terms *biased noise environment* and *asymmetric noise environment* interchangeably.

TABLE 1
Prisoner's Dilemma Payoff Matrix

		Player 2: *Cooperate*	Player 2: *Defect*
Player 1	Cooperate	R, R	S, T
	Defect	T, S	P, P

NOTE: R = reward payoff; S = sucker's payoff; T = temptation payoff; and P = punishment payoff.

of bilateral interaction, I examine whether a robust heuristic exists for the symmetrically and asymmetrically noisy environments modeled here and whether incentives exist for states to evolve from heterogeneous modes of behavior to international institutions facilitating cooperation.

By way of a road map, the following section describes the game structure, strategies, and methods used to conduct the analysis. The third section presents the results of the simulation-based analysis for the different noise environments. As one will see, positive and negative asymmetric noise have very different effects on the strategies' performances in homogeneous systems and in heterogeneous systems, and on the evolution of populations employing those strategies. The effects of neutral noise reflect the signature effects of each of the two asymmetric noise types. The fourth section extends these results to broader questions of world politics. Of the sample of strategies examined here, contrite tit-for-tat (CTFT) proves to be very robust to the effects of the different noise types—both in the heterogeneous system and as an institution. Under the conditions of this study, actors will almost always fare better in a CTFT institution than by employing any other strategy in a heterogeneous system. Thus there are incentives for the formation of institutions—and, in fact, one sees the boundedly rational actors of the evolutionary models moving from heterogeneous modes of bilateral interaction to cooperative norms of behavior. Finally, the fifth section concludes.

GAME STRUCTURE, STRATEGIES, AND METHOD

The game theoretic model employed here is the RPD. The single-stage game payoffs are illustrated in Table 1 and are structured ordinally such that $T > R > P > S$, with the additional constraint of $R > (T + S)/2$.[4] As in Axelrod

4. This constraint ensures that the payoff for mutually cooperating is greater than that from alternatingly (or randomly with $p = .5$) exploiting and being exploited.

(1984), I use $T = 5$, $R = 3$, $P = 1$, and $S = 0$. In this study, I simply assume the discount factor is 1. I do this primarily because I am interested in looking at the emergence and maintenance of cooperation when the shadow of the future looms large for the actors involved. My objective is not to conduct a parametric analysis of the discount factor's effect. Those previous simulation-based studies on cooperation that implement a discount factor less than 1 do so to preclude entrants in the tournaments from exploiting knowledge of when the last play is. No such strategies are implemented here.

STRATEGIES

Because the universe of possible strategies is infinite, selecting a sample of strategies to play in a tournament is no easy matter. Yet the selection process should not be ad hoc. In this case, three general guidelines are used. First, the sample of strategies employed here consists of simple (or ideal or heuristic) strategies one might expect to see in world politics, and variations thereof that one would hypothesize might perform better under different conditions of noise. The strategies are kept as simple as possible because there are no theoretical or empirical reasons for assuming that decision makers use highly complex strategies over long periods of time (or over a large number of interactions) and because simplicity allows one to assess more directly the effects of specific characteristics of the strategies. As will be shown, even with the simple strategies used here, complex patterns of interaction emerge. Second, because of the success of TFT in Axelrod (1984) and because of the extensive research on it since then, most of the strategies are modifications of TFT along substantively interesting dimensions.[5] Third, a number of strategies were implemented to determine how a parametric change in a strategy's characteristics affects its performance. Having identified these general criteria, I turn now to the strategies themselves.

The concept of reciprocity is most clearly embodied by TFT, which is considered a baseline strategy for this study. Using TFT, a player will cooperate on the first move and thereafter play what it perceives the opponent played the previous move. As the baseline strategy, an important question is whether this strategy is ever observed in international politics. There is some evidence to suggest that it is. The World War I live-and-let-live system described in Axelrod (1984, 73) is one example of TFT behavior. King (1989, 203) used a seemingly unrelated Poisson regression model estimator to identify TFT behavior between the United States and the Soviet Union. Goldstein (1991) found evidence of reciprocity between the United States and the Soviet Union during the period from 1950 to 1980. Finally, in examining reciprocity between South Africa and 36 other countries, Van Wyk

5. In fact, 12 of the 17 overall strategies are TFT based (including TFT itself).

and Radloff (1993) found significant levels of reciprocity among dyads in which the number of actions targeted by the dyad partners at each other are roughly equivalent.

Another TFT-based strategy is weighted tit-for-tat (WTFT*n*), which incorporates memory and weighting of an opponent's past moves. On the first move, WTFT*n* cooperates. Thereafter, like the reciprocity strategy developed in Bendor (1987), a player's response y_t is based on a weighted average of the opponent's past moves x_{t-i} such that $y_t = \text{round}(\alpha_1 x_{t-1} + \ldots + \alpha_n x_{t-n})$, where $\sum_{i=1}^{n} \alpha_i = 1$ and $0 \leq \alpha_i \leq 1$. Thus, depending on the memory depth, WTFT*n* may consider an opponent's defection to be a mistake (and therefore not retaliate) if the opponent's reputation is cooperative. On the other hand, if an opponent's history is entirely noncooperative, then depending on the weighting and memory depth, a WTFT*n* strategy might require an opponent to cooperate multiple times before it will resume cooperation. Here, I use a fairly simple weighting scheme for three instantiations of this strategy:[6] WTFT3 has a memory depth of 3 and is weighted (3/6, 2/6, 1/6) for past moves, WTFT4 has a memory depth of 4 and is weighted (4/10, 3/10, 2/10, 1/10), and WTFT5 has a memory depth of 5 and is weighted (5/15, 4/15, 3/15, 2/15, 1/15). In these cases, the weights are ordered such that the most recent move is weighted the highest.

CTFT (Sugden 1986; Boyd 1989) is also TFT based and is very similar to the strategy played by traders in the Champagne Fairs under the law merchant (Milgrom, North, and Weingast 1991). A player begins by cooperating on the first move. Play in subsequent moves is then dependent on the good standing of the player and the opponent: if a player is in bad standing or if the opponent is in good standing, the player will cooperate; otherwise, the player defects. A player is denoted as in bad standing if it unprovokedly defects when the other player is in good standing. Thus, if CTFT unprovokedly defects, then it will take the sucker's payoff on the next move as an act of "contrition." In turn, if CTFT's opponent unprovokedly defects, then it will require the opponent to cooperate on the next move while it defects. If the opponent does not make the act of contrition, CTFT will consider the opponent still in bad standing—and still requiring of contrition.[7]

Three exploiting TFT-based strategies are also included. Cheating tit-for-tat (CHTFT*n*) plays TFT (including cooperating on the first move) but defects after its opponent has cooperated *n* times in a row since either player last defected (or, prior to the first defection, since the start of the game). Three

6. TFT is equivalent to WTFT1 with memory depth of 1 and a weight of 1, so one could argue that there are actually four instantiations of WTFT.

7. One difference between CTFT and the Law Merchant is that if a CTFT player spuriously cooperates when the opponent is in bad standing, then the CTFT player is not considered in bad standing.

instantiations of this strategy are implemented: CHTFT2 defects after its opponent has cooperated twice in a row since either last defected, CHTFT3 defects after its opponent has cooperated three times in a row since either last defected, and CHTFT4 defects after its opponent has cooperated four times in a row since either last defected. The second exploiting strategy, randomly defecting tit-for-tat (RDTFT*p*), plays TFT (including cooperating on the first move) except that when its opponent cooperates, on the next turn RDTFT*p* cooperates with probability p and defects with probability $(1 - p)$. Three versions of this strategy are implemented: RDTFT99, RDTFT95, and RDTFT90 cooperate with probabilities .99, .95, and .90, respectively, after their opponents have cooperated. Another exploiting strategy implemented here is suspicious tit-for-tat (STFT), which is identical to TFT except that instead of cooperating on the first move, it defects.

The remaining strategies are nonreciprocating. The first, GRIM*n*, cooperates until its opponent defects a total of *n* times, defecting thereafter. Three instantiations of GRIM are implemented: GRIM1, GRIM2, and GRIM3 cooperate until their opponents defect a total of one, two, and three times, respectively, defecting thereafter. Thus, GRIM1 is an example of the severe situation in which a player will not cooperate after the opponent has defected once, whereas GRIM3 embodies the phrase "three strikes and you're out."

Finally, always defect (ALLD) and random (RAND) do just what their names imply: ALLD always defects, and RAND cooperates and defects with probability .5. Substantively, the former may represent a nation that is antagonistic or simply never cooperates, which, depending on the scenario, may be different actions. Although the latter could be interpreted as a nation intentionally playing a mixed strategy, it could also represent a nation that is without strategy. A situation such as this might occur if a nation's decision makers have no set policy concerning certain issues or when a nation is politically unstable, undergoing frequent changes in decision makers, who in turn implement different policies. Aside from their possible substantive interpretations, these last three classes of strategies also allow one to analyze particular characteristics that the TFT-based strategies do not share—and, similarly, how the reciprocating strategies react to those characteristics.

The above strategies can be classified according to a number of characteristics. Table 2 presents such a stylization of the strategies, identifying whether a given strategy is nice, nasty, provocable, retaliatory, generous, forgiving, contrite, or exploiting. These traits are defined with respect to the strategy's long-term behavior in a noiseless environment and are given as follows:

Nice (*N*): A strategy that is nice is never the first to intentionally defect (Axelrod 1984, 33). As Table 2 shows, only certain strategies played here are nice: TFT, WTFT*n*, CTFT, and GRIM*n*.

TABLE 2
Classification of Strategies

	Subclass								
Strategy	N	A	P	R	G	F	C	E	*Superclass*
TFT	x		x	x		x			$N\bar{A}PR\bar{G}F\bar{C}\bar{E}$
WTFT3	x		x		x	x			$N\bar{A}P\bar{R}GF\bar{C}\bar{E}$
WTFT4	x		x		x	x			$N\bar{A}P\bar{R}GF\bar{C}\bar{E}$
WTFT5	x		x		x	x			$N\bar{A}P\bar{R}GF\bar{C}\bar{E}$
CTFT	x		x	x		x	x		$N\bar{A}PR\bar{G}FC\bar{E}$
STFT		x	x	x		x		x	$\bar{N}APR\bar{G}F\bar{C}E$
CHTFT2			x	x		x		x	$\bar{N}\bar{A}PR\bar{G}F\bar{C}E$
CHTFT3			x	x		x		x	$\bar{N}\bar{A}PR\bar{G}F\bar{C}E$
CHTFT4			x	x		x		x	$\bar{N}\bar{A}PR\bar{G}F\bar{C}E$
RDTFT99			x	x		x		x	$\bar{N}\bar{A}PR\bar{G}F\bar{C}E$
RDTFT95			x	x		x		x	$\bar{N}\bar{A}PR\bar{G}F\bar{C}E$
RDTFT90			x	x		x		x	$\bar{N}\bar{A}PR\bar{G}F\bar{C}E$
GRIM1	x		x	x					$N\bar{A}PR\bar{G}\bar{F}\bar{C}\bar{E}$
GRIM2	x		x						$N\bar{A}P\bar{R}\bar{G}\bar{F}\bar{C}\bar{E}$
GRIM3	x		x						$N\bar{A}P\bar{R}\bar{G}\bar{F}\bar{C}\bar{E}$
ALLD		x						x	$\bar{N}A\bar{P}\bar{R}\bar{G}\bar{F}\bar{C}E$
RAND						x		x	$\bar{N}\bar{A}\bar{P}\bar{R}\bar{G}F\bar{C}E$

NOTE: N = nice; A = nasty; P = provocable; R = retaliatory; G = generous; F = forgiving; C = contrite; E = exploiting; $\bar{X}$ = not X.

Nasty (*A*): A strategy that is nasty is never the first to cooperate (Bendor 1993, 731). Here, only STFT and ALLD are nasty strategies.[8]

Provocable (*P*): If a strategy is provocable by another player's defection, there must be a strictly positive probability of it responding with a defection at some later move (Axelrod 1984, 62). Here, all but ALLD and RAND are provocable.[9]

Retaliatory (*R*): A strategy that is retaliatory is one that immediately defects after the unprovoked defection of another strategy (Axelrod 1984, 44). Being provocable and being retaliatory are related; the difference is that the latter requires certainty of an immediate response, whereas the former does not (Axelrod 1984, 218 n. 6). Thus being retaliatory implies provocability but not vice versa. So, for example, whereas all but ALLD and RAND are provocable, only TFT, CTFT, STFT, CHTFT*n*, RDTFT*p*, and GRIM1 are retaliatory.

Generous (*G*): A generous strategy is one that returns more cooperation in the long term than it receives (Bendor, Kramer, and Stout 1991, 696-701; Molander

8. Note that a strategy may be both nonnice and nonnasty—for example, RAND has a .5 probability of cooperating or defecting on the first move.

9. Although there is a "strictly positive probability" of ALLD defecting after the defection of its opponent—in fact, exactly equal to 1—ALLD's defection is not a response to the opponent's defection. Because my definition of provocability connotes conditional behavior, I categorize neither ALLD nor RAND as being provocable.

1985). Thus strategies like WTFT*n* are generous because they regularly allow for periodic defections against them without retaliating. On the other hand, GRIM2 and GRIM3 are not considered generous, because, although they allow for one or two defections before severely retaliating, in the limit the excess cooperation given goes to zero.

Forgiving (*F*): The extent to which a strategy is forgiving is determined by its propensity to cooperate after the other player has defected (Axelrod 1984, 36). Although all reciprocating strategies (such as the TFT-based strategies) are by definition forgiving, not all forgiving strategies reciprocate. For example, RAND does not reciprocate, yet it is defined as being forgiving, because there is some probability—specifically, $p = .5$—that it will cooperate after its opponent has defected.

Contrite (*C*): Contrition is defined as the propensity to accept punishment after unprovokedly defecting. Only one strategy here has that quality: CTFT.[10]

Exploiting (*E*): The characteristic of being exploiting is defined as in Bendor (1993, 715): *A* exploits *B* in a given period if *A* intentionally (or unprovokedly) defects while *B* cooperates. Thus the exploiting strategies in this study are CHTFT*n*, RDTFT*p*, ALLD, and RAND.

As will be shown in the following section, it is often useful to refer to a class of strategies rather than to list each individual strategy. This provides for economy of description, and it also allows one to examine whether certain characteristics are associated with higher or lower levels of performance. The component subclasses in Table 2 can be assembled to form superclasses of varying length. Table 2 shows the overall superclass for each strategy. Using this notation, the WTFT*n* strategies are all $NPGF\bar{E}$—nice, provocable, generous, forgiving, and nonexploiting. The bar over a letter indicates that the strategies in that class specifically do not have that characteristic; multiple letters together, such as that just given, indicate that the strategies have all the characteristics described by the letters. Generally, characteristic requirements are restrictive in the sense that the more the characteristic requirements, the fewer the number of strategies there will be in a sample that have those qualities. Because this notation will be used henceforth, I provide the following list of commonly appearing classes and their included strategies:

$NPGF\bar{E}$: WTFT3, WTFT4, WTFT5
$NPF\bar{E}$: TFT, WTFT3, WTFT4, WTFT5, CTFT
$NP\bar{E}$: TFT, WTFT3, WTFT4, WTFT5, CTFT, GRIM1, GRIM2, GRIM3
$\bar{N}E$: STFT, CHTFT2, CHTFT3, CHTFT4, RDTFT99, RDTFT95, RDTFT90, ALLD, RAND

10. The act of contrition may be interpreted in a number of different ways. It may require a formal apology, for which there is some cost—perhaps in terms of international reputation or status. Or it may take the form of financial compensation, as with the fee demanded by Milgrom, North, and Weingast's (1990) law merchant. Here, it is simply modeled as accepting the sucker's payoff after unprovokedly defecting.

PF: TFT, WTFT3, WTFT4, WTFT5, CTFT, STFT, CHTFT2, CHTFT3, CHTFT4 RDTFT99, RDTFT95, RDTFT90
$\bar{F}$: GRIM1, GRIM2, GRIM3, ALLD

METHOD OF ANALYSIS

As previously mentioned, both symmetric and asymmetric errors are implemented in this study. Because the moves are binary (i.e., cooperate or defect), noise is modeled by allowing for some probability p that a player's implemented move is not what it intended. In the case of symmetric or neutral noise, p_{ε_0} represents the probability that an intended cooperation will be implemented as a defection and that an intended defection will be implemented as a cooperation. For asymmetric noise, there are two cases: positive and negative.[11] For positive noise, p_{ε_+} represents the probability that an intended defection will actually be implemented as cooperation; for negative noise, p_{ε_-} represents the probability that an intended cooperative move will be implemented as a defection.[12]

For each nontrivial noise environment (neutral, positive, and negative), four types of analyses are conducted. First, each of the strategies is assessed for how it performs against itself under noise levels that range from $p = 0$ to $p = .5$ in increments of .01. Each dyadic simulation plays for 200 iterations, is replicated 50 times, and is then averaged to obtain an estimate of a strategy's performance against itself. This score is then divided by 200 to obtain the average payoff.

What does this first type of analysis show? At its most basic level, it shows how two decision makers with identical strategies perform against each other under various noise conditions. However, it can be extended beyond this. For example, it also represents how a population with identical strategies (i.e., homogeneous in terms of strategies) will fare when the population members interact dyadically. Moreover, the strategies implemented here represent *rule-based behavior* (Mailath 1992, 264-5). Thus the results of this analysis also give one insight into how different norms or institutions—qua rules—are affected by various distributions of noise.

In the second type of analysis, each strategy is matched dyadically against every other strategy in a tournament, again under noise levels that range from $p = 0$ to $p = .5$ in increments of .01. As before, each dyadic simulation plays for 200 iterations, is replicated 50 times, and is averaged to obtain the

11. The biased (or asymmetric) noise distributions modeled here are completely asymmetric. Thus, in the positively biased environment, the probability of a negatively directed error is 0, as is the probability of a positively directed error in the negatively biased environment.

12. Although the cooperation can often be used for exploitative purposes, I use the terms *positive* and *negative* in accordance with the normative connotations of traditional cooperation theory.

estimates of its performance against the others. Once all the dyadic simulations have been run for a given value of p, each strategy's average payoff is calculated by summing its performance scores against all of the strategies and then dividing by 1,400, the total number of moves a strategy made against all 17 strategies.

Whereas the first type of analysis examines how noise affects the performance of strategies in homogeneous populations, here the analysis investigates the effects of noise in heterogeneous populations of strategies. Although the results of the heterogeneous systems are specific to the sample of strategies analyzed, they often highlight how the general strategy characteristics affect a strategy's performance relative to other strategies and relative to the noise environment. This information is also important for three other reasons. First, one can compare the payoffs of an individual strategy in homogeneous versus particular heterogeneous systems. The performance of some strategies may degrade when moving from a homogeneous to a heterogeneous system, whereas the payoffs of others may actually increase. Second, if a particular strategy in a homogeneous system performs better than any of the strategies in the heterogeneous system, then gains may be obtained by institutionalizing the rule-based behavior of the strategy. In addition to this, the tournament information is important for a third reason—namely, it provides the payoffs used in analyses of population dynamics, which comprise the two remaining types of analyses conducted in this article.

The population dynamics of the strategies are examined using ecology and territory models—both of which fall under the rubric of "evolutionary models." As others have noted elsewhere, evolutionary models are inherently models of *bounded rationality* (Mailath 1992, 264; Morrow 1994, 308-11). An assumption of these models is that actors do not optimize over all possibilities of action but are rather unsophisticated, instead adopting or slightly modifying the successful strategies they see around them. Evolutionary models generally incorporate some mechanism for the "reproduction" or the spreading of successful strategies through a population, sometimes also incorporating mutation for the generation of new strategies. In the two evolutionary models employed in this article, I do not incorporate mutation. Rather, the focus is on how the population evolves in terms of the distribution of the strategies already presented.

Ecology models assume a system with a changing distribution of strategies, in which the proliferation of a strategy is proportional to how well it does in the population. The better performing strategies are increasingly adopted by members of the population, whereas the less successful strategies become less common (Axelrod 1984, 51). Calculation of the ecological population dynamics is conducted as in Axelrod (1984). The process is started using the current proportion of each strategy in the population. Because there are 17 strategies, the starting proportions are 1/17. The proportions of subsequent generations are calculated as follows. First, the weighted average of each strategy's scores against the other strategies is calculated, where the weights are the current proportions of the strategies. Then, the next generation's proportions are calculated as a weighted average of the current proportions, where the weights are the weighted performance scores just calculated. Thus a strategy's proportion is determined not just by its scores against other strategies but also by the proportions of those strategies currently in the population. This process is conducted until the population dynamics stabilize.

Territory analysis also provides information on population dynamics, but under slightly different conditions. Here, population members are assumed to be utility maximizing, but boundedly so. In this model, population members have fixed locations in the territory and have four neighbors: one above, one below, one to the right, and one to the left.[13] Although a population member is assumed to be utility maximizing, the set of options available to her is limited to her own strategy, plus the strategies of her four neighbors. The purpose of the territory analysis in this study is to examine the converged distributions of strategies in the population when starting under randomized configurations of strategies and under varying noise levels. Here, the territory consists of a 25 × 25 matrix of players. For each type of noise, the level of noise is varied from 0 to .5. At each noise level, the players in the 25 × 25 matrix are first randomly assigned strategies. Then, in each generation, each player is given a score for how its strategy performs on average against the strategies of its four neighbors. Each player then examines how well its neighbors performed on average against their neighbors, and if any neighbor performed better than it did, it will adopt the strategy of the highest performing neighbor.[14] This process continues until the territory converges to a steady-state population, at which point the proportion of each strategy in the territory is recorded.[15] This process of randomly assigning strategies to players and iterating until convergence is conducted 100 times for each noise level. For the territory results in the following sections, I report at each noise level the average proportion of the population over the 100 trials.[16]

13. The concepts of location and neighbors can be abstracted in a number of ways. See Axelrod (1984, 158-60) for more on territory analysis.

14. To ensure that the edge locations are not biased, the territory "wraps around" at the edges. For example, given a 25 × 25 matrix where (i,j) represents the location of row i and column j, the neighbor "below" (25,1) is (1,1), and the neighbor on its "left" is (25,25).

15. Sometimes, small locales of a territory would exhibit (generally two-generation) cycling between (generally two) strategies. When territories converged to this type of stable cycling, I took the first phase of the cycle as the converged territory. Because the whole process was conducted 100 times for each noise level, the effects of this should have averaged out.

16. This process uses *synchronous* updating of the territory—that is, every player in the population updates its strategy simultaneously each generation. Based on the suggestion of an anonymous referee, I also investigated how *asynchronous* updating (see Huberman and Glance

As with the first two analyses, the results of the evolutionary models can be interpreted in a number of different ways. First, one can examine how the distribution of strategies evolves among the boundedly rational members of the population—that is, how the type and level of noise affect the distribution of strategies in the steady-state population. Second, if the population converges on a single strategy or a single class of strategies, then one has conditions under which the boundedly rational members of this system converge on certain rules of behavior. In this way, one has a story for how bounded rationality can lead to the formation of certain norms or institutions.

SIMULATION RESULTS

In this section, I examine how the strategies perform under the various noise environments. In addition, I also analyze the noiseless environment, because it serves as a baseline for the nontrivial environments. Within each subsection, I address the four areas of analysis: how strategies perform against themselves, how the strategies perform in a tournament, the population dynamics of strategies in an ecology, and the population dynamics of strategies in an initially randomized territory.

NO NOISE

As a baseline, one might ask how this sample of strategies performs under noiseless conditions. Table 3 presents the ranking and scores of the strategies against themselves and in a tournament. For the homogeneous system, the top eight strategies are all tied in score, receiving the reward payoff. The reason for this is because they are all $N\bar{E}$—nice and nonexploiting. Thus these strategies perpetually cooperate with themselves in a noiseless environment. In contrast, the worst performing strategies are all $\bar{N}E$; ALLD received the lowest average payoff among the strategies—that being the punishment payoff.

In the heterogeneous system, the top eight strategies (all scoring 3 on average) are those that are $NP\bar{E}$, whereas those that perform the worst are $\bar{N}E$.

1993) affects the resulting distribution of strategies over the noise levels. In contrast to synchronous updating, asynchronous updating consists of updating the territory population by repeatedly randomly choosing individual players to update with their four neighbors. A total of 1,871 asynchronous updates were required of each generation so that the probability of an individual player not being chosen for updating during that generation was less than .05. This process of asynchronous updating over each generation was then iterated until the territory distribution stabilized. For each noise level, this was conducted 100 times and the average taken over the 100 trials. Surprisingly, conducting the above territory analyses with asynchronous updating yielded almost identical results as with synchronous updating. Because of that, I present only the synchronous results in this article.

TABLE 3
Strategy Scores (and Rankings) in Noiseless Environment

	Scores	
Strategy	*Homogeneous System*	*Heterogeneous System*
WTFT5	3.00 (1)	2.56 (1)
WTFT3	3.00 (1)	2.54 (2)
WTFT4	3.00 (1)	2.53 (3)
TFT	3.00 (1)	2.50 (4)
CTFT	3.00 (1)	2.50 (4)
GRIM3	3.00 (1)	2.34 (5)
GRIM2	3.00 (1)	2.31 (6)
RDTFT99	2.11 (3)	2.19 (7)
GRIM1	3.00 (1)	2.13 (8)
STFT	1.00 (9)	2.09 (9)
CHTFT2	1.02 (8)	2.08 (10)
CHTFT3	1.03 (7)	2.06 (11)
CHTFT4	1.04 (6)	2.06 (11)
RAND	2.25 (2)	1.86 (12)
RDTFT95	1.26 (4)	1.58 (13)
RDTFT90	1.13 (5)	1.35 (14)
ALLD	1.00 (9)	1.14 (15)

In contrast to Axelrod's (1984) tournaments, TFT is not the top performer in this tournament of strategies, although it ties for fourth, and its average performance (2.50) is very close to that of the top performer WTFT5 (2.56). Examining the average performance of the strategies in the overall tournament shows that all of the top seven scoring strategies (WTFT5, WTFT3, WTFT4, CTFT, TFT, GRIM3, and GRIM2 [ranked 1, 2, 3, 4, 4, 5, and 6, respectively]) are nice, provocable, and nonexploiting ($NP\bar{E}$); the top five strategies are also forgiving ($NPF\bar{E}$); and the top three strategies are generous as well ($NPGF\bar{E}$). Being provocable (and especially retaliatory) tends to keep a strategy from being exploited, whereas the qualities of generosity and forgiveness allow it to maintain and resume cooperation.

Because the top-ranked strategies in the noiseless tournament are nice—and thus will always cooperate with each other—their relative rankings are determined primarily by how they perform against the exploiting strategies. In general, the generous strategies perform best against the exploiters. However, CHTFT*n* strategies perform extremely well against WTFT*n* strategies precisely because they are so generous. In fact, the highest scores in the tournament are obtained by the CHTFT*n* strategies against the WTFT*n* strategies, because CHTFT*n* can sneak in defections against which WTFT*n* will not retaliate if CHTFT*n*'s reputation is cooperative. However, the

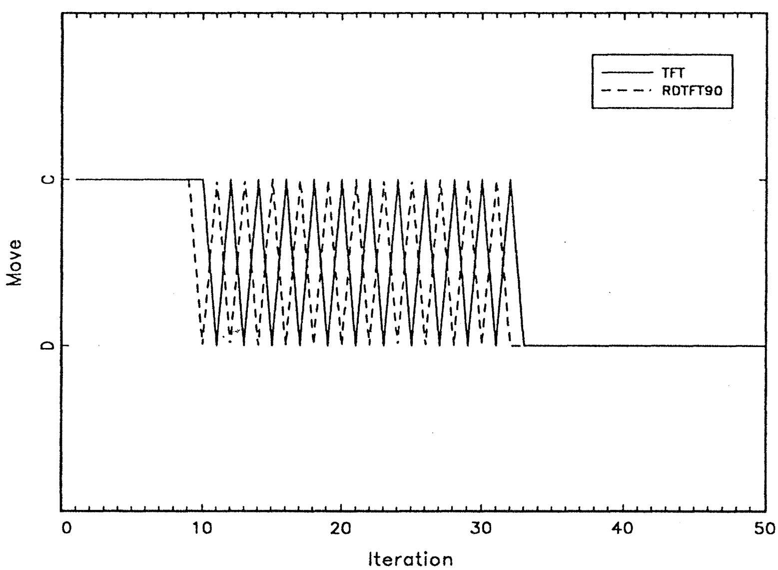

Figure 1: The Onset of Unsynchronized Retaliations and the Resulting Perpetual Defection between TFT and RDTFT90
NOTE: C = cooperation; D = defection.

CHTFT*n* strategies perform poorly against themselves and against the remaining strategies, resulting in poor overall performance.

The lowest average scoring strategies (ALLD, RDTFT90, and RDTFT95 [ranked 15, 14, and 13, respectively]) are both nonnice and exploiting ($\bar{N}E$). ALLD scores poorly because it incurs permanent defection from the provocable strategies (which include 15 of the 17 strategies). The RDTFT*p* strategies tend to score poorly because they trigger an unsynchronized sequence of retaliations by both parties, similar to the *echoes* described by Axelrod (1984, 37). Figure 1 shows an example of this process between TFT and RDTFT90. Here, RDTFT90 randomly defects at iteration 10, causing TFT to retaliate in iteration 11. However, RDTFT90, which cooperated in iteration 11 because TFT cooperated in iteration 10, then retaliates against TFT's defection. TFT retaliates in turn, and this process continues until RDTFT90 randomly defects again, at which point both TFT and RDTFT90 are synchronized in their defections and both defect for the remainder of iterations, receiving the punishment payoff.

Interestingly, although the CHTFT*n* strategies suffer from this problem, their performance is not harmed to the same extent. Figure 2 shows why.

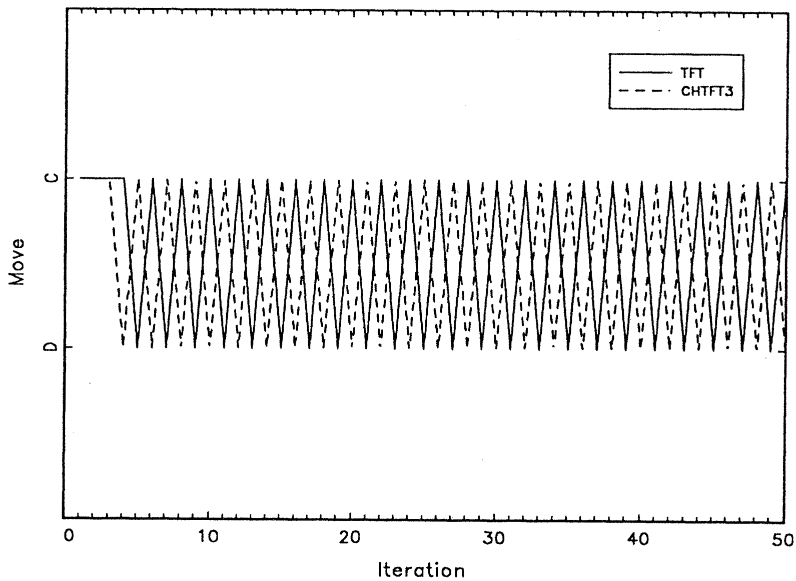

Figure 2: The Emergence of Perpetual Unsynchronized Retaliations between TFT and CHTFT3
NOTE: C = cooperation; D = defection.

Here, CHTFT3 sneaks in a defection after the third cooperation by TFT. TFT retaliates against this, and the unsynchronized retaliations begin. However, whereas RDTFT90 randomly defected again, resynchronizing the process into permanent defection, CHTFT3 does not do so, resulting in both strategies alternatingly receiving the temptation and sucker payoffs until they stop interacting. Although alternating *T* and *S* is less than *R*, it is greater than *P*, so CHTFT3 scores higher than RDTFT90 on average against a strategy like TFT. This generalization holds for the two classes of exploiting strategies when the RDTFT*p* probability of defection is high enough.[17]

Although Table 3 gives the results of the strategies' performances for heterogeneous and homogeneous ecologies, it is difficult to assess how these strategies fare together unless one analyzes them in the context of population dynamics. As mentioned in the previous section, two ways of doing this are by analyzing ecological systems and territorial systems of the strategies. The mechanism behind both of these is that better performing strategies prolifer-

17. Note that RDTFT99 is ranked 7 versus 10, 11, and 11 of CHTFT2, CHTFT3, and CHTFT4, respectively.

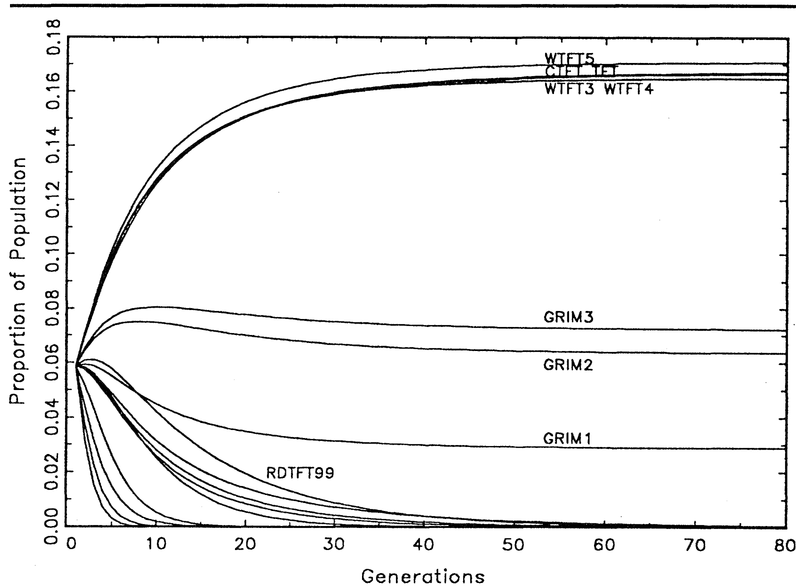

Figure 3: Population Dynamics of Strategies in Noiseless Ecology

ate, whereas those that do not perform well either become obsolete and die out or are reduced to a negligible level of representation.

Figure 3 shows the population dynamics of the strategies in an ecology with no noise; each line represents the proportion of that strategy in the population from generation 0, at which point all strategies start out as 1/17 of the population, until the system stabilizes. Not surprisingly, the top seven scoring strategies (WTFT5, WTFT3, WTFT4, CTFT, TFT, GRIM3, and GRIM2) are all present to varying degrees in the system when it stabilizes. Because each of these strategies is nice and because there is no noise in the system, each of these strategies will indefinitely cooperate with each other, resulting in the reward payoff. In fact, it is extremely difficult to differentiate the performance of CTFT, TFT, and the WTFT*n* strategies. Although the seventh-ranked RDTFT99 initially performs better than eighth-ranked GRIM1, the strategies against which RDTFT99 scores well begin to disappear quickly. By the ninth generation, RDTFT99 drops below GRIM1, eventually disappearing, whereas GRIM1 stabilizes at a low (but not negligible) level of representation in the system.

In the territorial system, each strategy does not meet each other strategy every generation. Rather, each location's strategy interacts only with the strategies of its four neighbors. If a location's average performance against

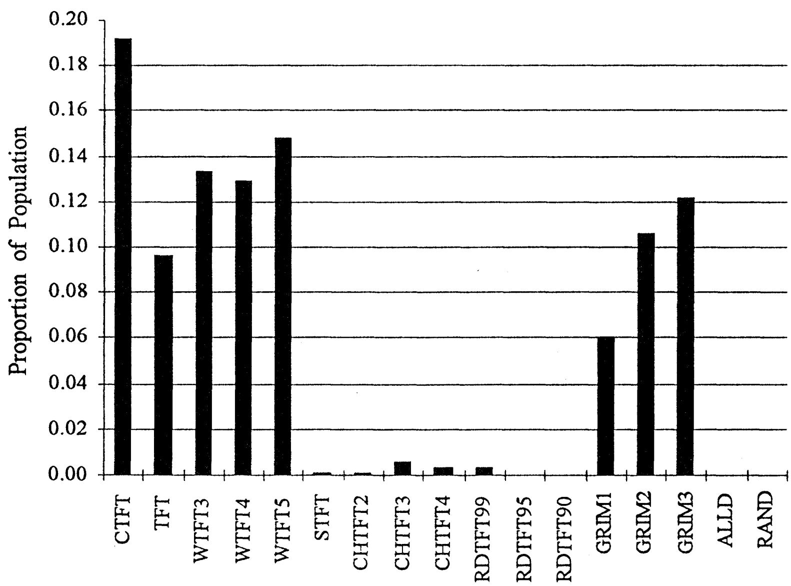

Figure 4: Average Proportion of Population in Noiseless Territory

these four neighbors is less than the average performance of any of its neighbors, then it adopts the strategy of its highest scoring neighbor. Figure 4 displays the resulting average distribution of strategies for a randomly populated 25 × 25 territory under noiseless conditions. One sees that there is a plurality of strategies present in the territory; no one strategy truly dominates. Noticeably underrepresented (and in some cases unrepresented) are the exploiting strategies. Additionally, although no single strategy dominates, a class of strategies does: $NP\bar{E}$. I noted earlier that the $NP\bar{E}$ strategies cooperate indefinitely with each other and tend to perform better against the exploiters than the exploiters do against each other. After the exploiters disappear from the system, each $NP\bar{E}$ strategy is indistinguishable from the other, resulting in a plurality of strategies from this class.

NEUTRAL NOISE

Having examined the case of a noiseless environment, I now examine the effects of neutral noise on the performance of the strategies interacting in the homogeneous populations, in the heterogeneous populations, and subsequently, in the ecologies and territories. Here, the probability p_{ε_0} that an intended cooperation will result in a defection (and vice versa) is varied from

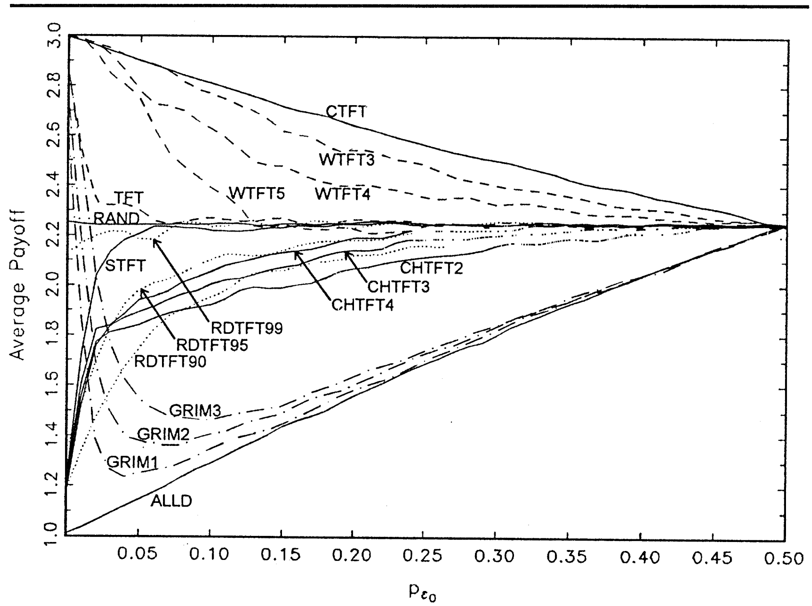

Figure 5: Average Payoff for Each Strategy in a Homogeneous System under Neutral Noise

0 to .5 in increments of .01. Figure 5 displays the average payoff of each strategy when played against itself under conditions of neutral noise.[18] Note that payoffs of 1 and 3 correspond to defecting and cooperating on average, respectively. Loosely speaking, one sees here that CTFT is an upper bound, ALLD is a lower bound, and RAND is a limit toward which all strategies tend as the noise increases. As the earlier studies mentioned in the first section show, TFT drastically drops in performance with even the smallest amount of neutral noise; with greater than 5% noise in the system, one might as well be flipping a coin.

It is interesting to examine how the groups of these strategies are affected by the noise. Take, for example, CTFT, WTFT*n*, and TFT. All of these strategies perform at least as well as RAND—and most do much better overall. All five of the strategies in this group are *NPF$\bar{E}$*: each starts out being nice, each is provocable, each is forgiving and will resume cooperation when the other

18. For presentation, the data in the homogeneous and heterogeneous payoff graphs have been smoothed using the formula $Y_t^* = .25y_{t-1} + .5y_t + .25y_{t+1}$. Because the variations in the raw data are relatively small, this represents only a very mild smoothing of the curves, and no significant information is lost in the process. Moreover, the raw data are used for all subsequent ecological and territorial analyses. The raw data are available on request from the author.

player cooperates, and none attempt to exploit the other player. Of these strategies, TFT does so poorly because, when playing against itself, it falls into unsynchronized retaliations once a spurious defection occurs. It fails to do as well as the other four strategies because it lacks the additional qualities they have: generosity and contrition.

The WTFT*n* strategies all add generosity to the baseline TFT characteristics. Therefore, they will allow spurious defections on the opponent's part if the opponent's history is cooperative. This tends to maintain cooperation when any negative noise is present in the environment, because noise-induced defection by one player will be less likely to provoke retaliation by the other. Because the WTFT*n* strategies are playing only against themselves in the homogeneous systems, there is no problem with this generosity being exploited.

In contrast, CTFT implements contrition, which attacks the pathology of unsynchronized retaliations from a different angle. Whereas the generosity of WTFT*n* reduces the onset of feuds because of unprovoked defections on the part of the *other* player, the contrition of CTFT ensures this pathological behavior will not ensue because of its *own* spurious defections. With CTFT playing in a homogeneous system, one should never see unsynchronized retaliations, because each player will offer contrition for any mistaken defection, and cooperation will then follow.

A dividing line of sorts appears to be the strategy RAND. Because RAND randomly mixes cooperation and defection with equal probability, it is not surprising that its payoff against itself—$(T + R + P + S)/4 = 2.25$—remains constant across all neutral noise levels.[19] Moreover, as neutral noise approaches 50%, strategy simply does not matter, as the realized outcomes become indistinguishable from RAND's.

The strategies whose average performance against themselves falls below RAND's are either nonnice and exploiting (e.g., STFT, CHTFT*n*, RDTFT*p*, and ALLD) or not forgiving (e.g., GRIM*n* and ALLD). The worst performing of these, ALLD, falls into both categories. In contrast to the $NPF\bar{E}$ strategies, the $\bar{N}E$ and $\bar{F}$ strategies are generally helped by the neutral noise—with the exception of the GRIM*n* strategies, which initially plummet drastically in payoff. The positive element of the neutral noise helps the TFT-based exploiters to break out of permanent defections or unsynchronized retaliations. For example, a dyad of nasty STFTs, which will perpetually defect against each other under noiseless conditions, can accidentally be shocked into unsynchronized retaliation by the spurious cooperation of one player and then even subsequently shocked into cooperation by spurious cooperation from

19. I provide a model in the appendix for the expected payoff of mixed, stationary strategies in environments with either symmetric or asymmetric noise. Payoffs for RAND and ALLD in noisy, homogeneous systems can be calculated using this model.

the other player. Although the neutral noise does contain negative error, which will stimulate the interaction in the opposite direction, the overall payoff from supergames containing even brief periods of unsynchronized retaliation will be better than the payoff from those consisting of perpetual defection.[20]

The dynamic for the unforgiving strategies is a bit different. It is interesting that, of the sample of strategies examined here, the GRIM*n* are the only strategies in the homogeneous system whose payoffs display a nonmonotonic behavior as neutral noise increases.[21] With even very small amounts of neutral (more specifically, negative) error in the system, the GRIM*n* strategies become almost indistinguishable from ALLD. However, as the amount of noise increases, the $\bar{F}$ strategies benefit because of the positive element of the noise and the fact that occasionally obtaining equal amounts of T and S yields a better overall payoff than perpetual defection.

Parametrically, the qualities that differentiate the success between groups also determine the ranking of strategies within a group. For example, WTFT3 is more generous than WTFT4, which is more generous than WTFT5. And one sees their payoffs positively related to this generosity. Similarly, as strategies become more exploiting, their performance decreases. Thus CHTFTn_1 performs worse against itself than does CHTFTn_2 when $n_1 < n_2$, and RDTFTp_1 fares worse against itself than does RDTFTp_2 when $p_1 < p_2$. Finally, GRIMn_1 does worse against itself than does GRIMn_2 when $n_1 < n_2$.

Whereas Figure 5 displays the performance of strategies in homogeneous systems, Figure 6 shows how each strategy fares in the heterogeneous sample. Here, one sees a similar convergence of the strategies' performance to $(T + R + P + S)/4 = 2.25$ at $p_{\varepsilon_0} = .50$. *However, an interesting effect occurs in the region* $p_{\varepsilon_0} \approx .26$, dividing the results into two regions. In the region $p_{\varepsilon_0} < .26$, *the* $NPF\bar{E}$ strategies are the best performers, the TFT-based exploiters are mediocre, and the unforgiving ALLD and GRIM*n* strategies fare poorly. The strategies within each parametric group retain primarily the same relative ranking as before. In contrast, the results are almost reversed in the region $p_{\varepsilon_0} > .26$. Here, the top performers are the $\bar{F}$ strategies, whereas the top two performers in the previous region (CTFT and WTFT3) become the lowest performers. Finally, comparing the heterogeneous results with the homogeneous results, one sees that in moving from the homogeneous system to this system of strategies, the performance of the exploiting strategies is not affected greatly. However, the $NPF\bar{E}$ strategies do not perform nearly as well

20. Bendor (1993, 731-2) identified that for TFT and STFT, neutral noise makes the reciprocative characteristics decisive over whether the strategies are nice or nasty. Thus, in symmetrically noisy systems, TFT and STFT will reach the same limiting distribution of outcomes. One sees this in the simulation results by the relatively quick convergence of TFT and STFT to the same payoff—that of RAND.

21. This nonmonotonic pattern of the GRIM*n* payoffs is also identified in Bendor (1993, 725-6).

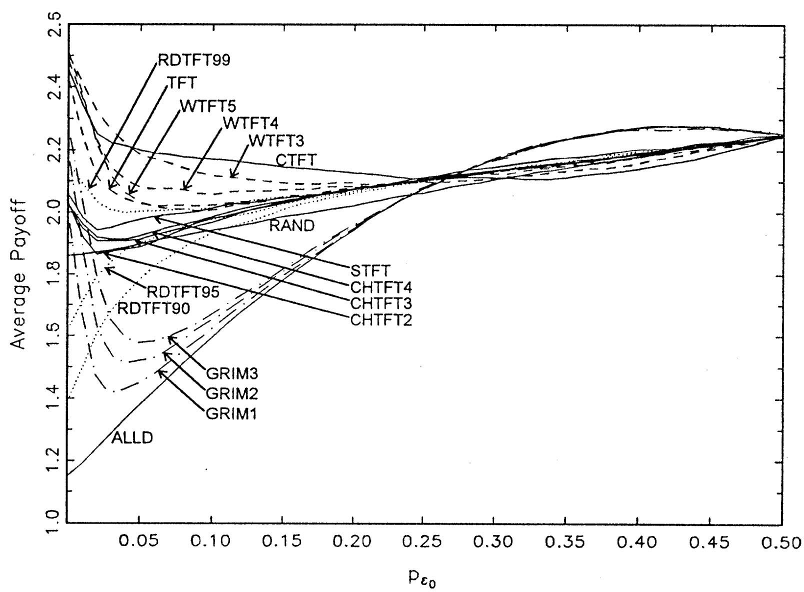

Figure 6: Average Payoff for Each Strategy in the Heterogeneous System under Neutral Noise

overall and degrade at a quicker rate (in terms of noise levels), whereas the $\bar{F}$ strategies improve at a quicker rate.

A natural question is, Why the switch in ranking around $p_{\varepsilon_0} \approx .26$? The answer lies in how the *NPF$\bar{E}$* strategies interact with other *NPF$\bar{E}$* strategies and in how they interact with the $\bar{F}$ strategies. To see this, consider a hypothetical tournament between only TFT and ALLD. TFT is an ideal type of the *NPF$\bar{E}$* strategies, and in environments with any negative noise, ALLD is an ideal type of the $\bar{F}$ strategies. In this simplified tournament, TFT's average tournament payoff would be the average of its payoff against itself and against ALLD. Similarly, ALLD's average tournament payoff would be the average of its payoff against TFT and against itself.

Consider first the average payoff of ALLD against itself. This is the same as ALLD's average payoff in the homogeneous system, which is a line from 1.00 at $p_{\varepsilon_0} = 0$ to 2.25 at $p_{\varepsilon_0} = .5$. Next, consider the average payoff of TFT when played against ALLD. It is not difficult to see that TFT's average payoff against ALLD will always be less than ALLD's average payoff against itself, except at $p_{\varepsilon_0} = 0$ and $p_{\varepsilon_0} = .5$, where they are equal. This is because TFT will have a greater probability of being exploited by ALLD than ALLD does against itself. At the same time, TFT cannot garner large enough gains from spurious mutual cooperation or from spurious exploitation of ALLD to offset the exploitation by ALLD. Nevertheless, ALLD's average payoff against itself is never larger than TFT's payoff against ALLD by more than approximately .2.[22] Because the average payoffs of the two strategies against ALLD never cross and because they are fairly close over the entire noise range, the interaction of these two dyads cannot explain the crossover identified in Figure 6; rather, they simply shift the resulting tournament payoff curves downward in the averaging.

Instead, one must look to how the two strategies interact with TFT. Consider the average payoff of TFT against itself. Again, this is just the average payoff of TFT in a homogeneous system. As noted in Figure 5, as the neutral noise increases, TFT's payoff in a homogeneous system quickly converges to that of simply flipping a coin. In contrast, ALLD's payoff against TFT starts low—a payoff of 1 at $p_{\varepsilon_0} = 0$—but rises as the neutral noise increases. ALLD's payoff against TFT surpasses TFT's payoff against itself for the same reason that TFT can never do better against ALLD than ALLD can against itself—namely, the proclivity of TFT to be exploited by ALLD when positive noise exists in the environment. Whereas ALLD naturally defects—and thus will only be exploited because of a spurious cooperation on its part—TFT will be exploited because of both a spurious cooperation on its own part and a spurious cooperation on ALLD's part, which will generally lead to TFT cooperating on the next move and ALLD resuming its defection. Although ALLD also eventually converges to the random tournament payoff of 2.25, for moderate to high levels of neutral noise it is able to capture gains from exploiting TFT that push its payoff beyond TFT's average payoff when played against itself. This simple example extends to the $NPF\bar{E}$ and $\bar{F}$ strategies in the larger tournament, and it is this dynamic that explains the crossover in Figure 6.

Given these tournament results, I now examine the population dynamics of the strategies at various levels of neutral noise. Figure 7 displays the population dynamics of an ecology with 1% noise. Here, the $NPF\bar{E}$ strategies are all present during the transition period,[23] but the noncontrite strategies disappear after 800 generations, leaving CTFT dominating the system.[24] This

22. This analysis was conducted after the simulations using a Markov model of the transition probabilities between states. Results are available on request from the author.

23. In analyzing system dynamics, one often observes a transition period before the system reaches a steady state.

24. The robustness of CTFT under low levels of noise is further corroborated by Wu and Axelrod's (1995, 187) ecology analysis under 1% noise, which shows CTFT becoming the predominant strategy in the population. The significance of this is that, between their analysis and this one, CTFT is shown to be highly robust under low levels of neutral noise in two different samples of strategies.

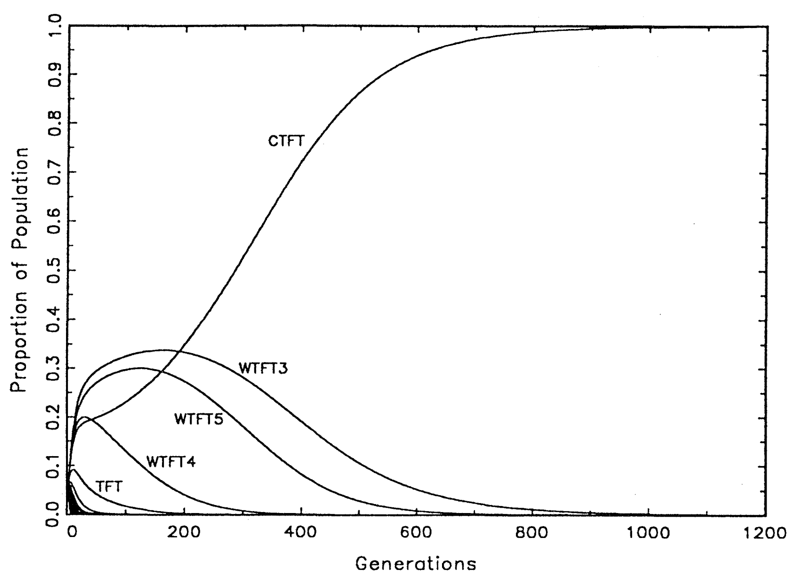

Figure 7: Population Dynamics of Ecology with Neutral Noise, p_{ε_0} = .01

hegemony of CTFT continues for neutral noise levels up to 20% but decreases, becoming only part of the transition period, as the error level approaches 26%. For values of p_{ε_0} near .26, the transition region becomes more prominent. Figure 8 shows that for $p_{\varepsilon_0} = .23$, a mix of cooperation and exploitation results: after maintaining a large presence in a fairly long transition period, CTFT is reduced to a negligible level, whereas the exploiters CHTFT2 and CHTFT4 increase and remain stably present in the system. With greater than 26% noise in the system, the outlook is—quite literally—grim. Figure 9 shows that for a noise level of 27%, the system in steady state is composed of $\bar{F}$ strategies. Interestingly, Figure 10 shows that for 35% noise, the system is highly unstable, oscillating in system dominance by each of the three GRIM*n* strategies. However, I suspect this may be due more to the strategies' scores all converging as the noise increases and to the stochastic simulation error, given that the GRIM*n* scores are so close. For example, at $p_{\varepsilon_0} = .35$, the average payoffs of the GRIM*n* strategies against all other strategies are all very similar. Additionally, their payoffs against each other are all very close. Thus they appear to cycle in population predominance, when, in fact, they all perform nearly identically.

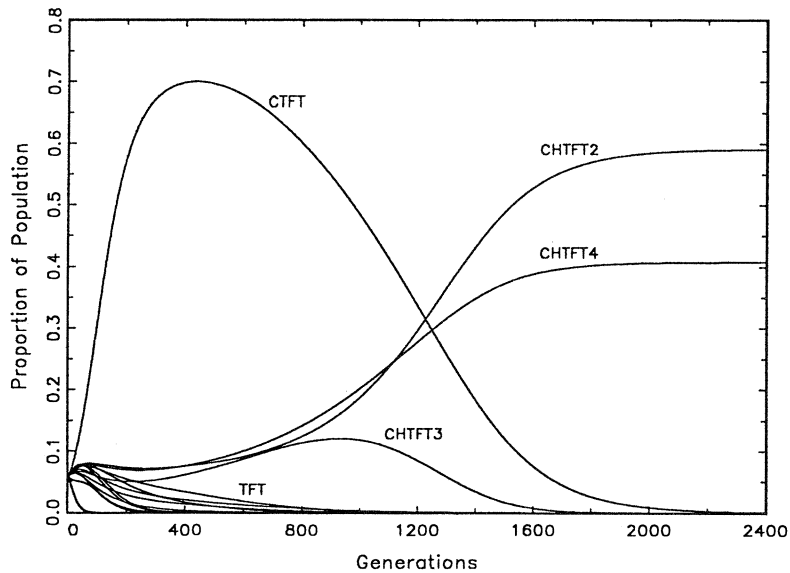

Figure 8: Population Dynamics of Ecology with Neutral Noise, $p_{\varepsilon_0} = .23$

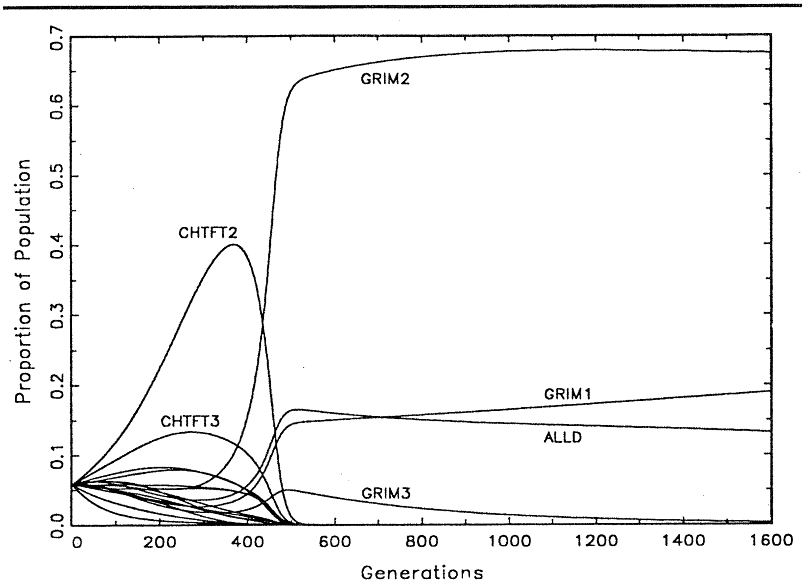

Figure 9: Population Dynamics of Ecology with Neutral Noise, $p_{\varepsilon_0} = .27$

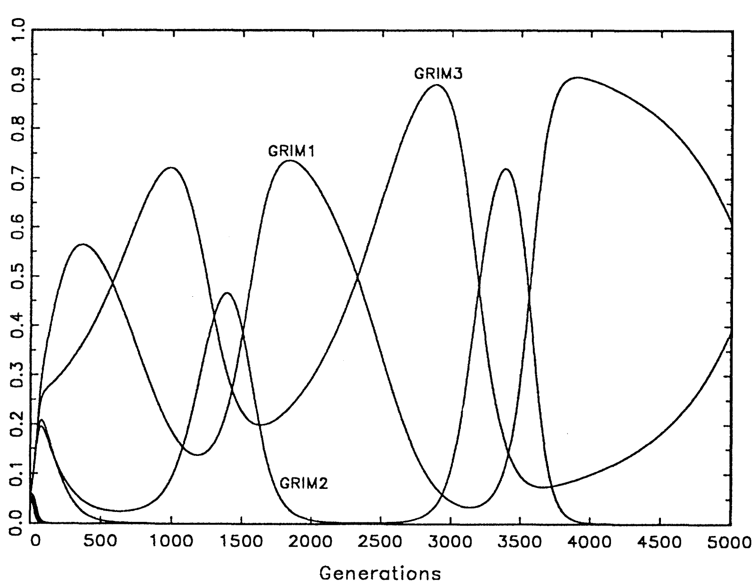

Figure 10: Population Dynamics of Ecology with Neutral Noise, $p_{\varepsilon_0} = .35$

Turning now to the population dynamics of the strategies in a randomized territory with neutral noise, Figure 11 plots for each noise level the proportion of each strategy represented in the territory after convergence. As in Figure 4, one sees that for zero to very low levels of noise, the $NP\bar{E}$ strategies are all represented. As the noise increases, the GRIM*n* strategies fade quickly from the converged population, with the *NPFE* strategies still appearing until about 5% noise. However, from approximately 5% to 26% noise, CTFT almost entirely dominates the territory. After that, CTFT quickly plummets to almost no representation. The region from 30% to 50% noise is quite turbulent in terms of strategy representation.[25] However, in this region, the $\bar{F}$ strategies tend to dominate, just as they did in the same region of the tournament scores and in the ecology dynamics. Thus, whereas extremely minor amounts of neutral noise allow for a plurality of $NP\bar{E}$ strategies, small to moderate amounts of neutral noise result in a hegemony of CTFT, until high levels of noise again allow for a plurality of strategies, but this time dominated by the unforgiving strategies.

25. Another indicator of this shift is that the average number of generations to convergence is 12.8 for the region $0 \leq p_{\varepsilon 0} \leq .3$, but jumps to 110.6 for the region $.31 \leq p_{\varepsilon 0} \leq .5$.

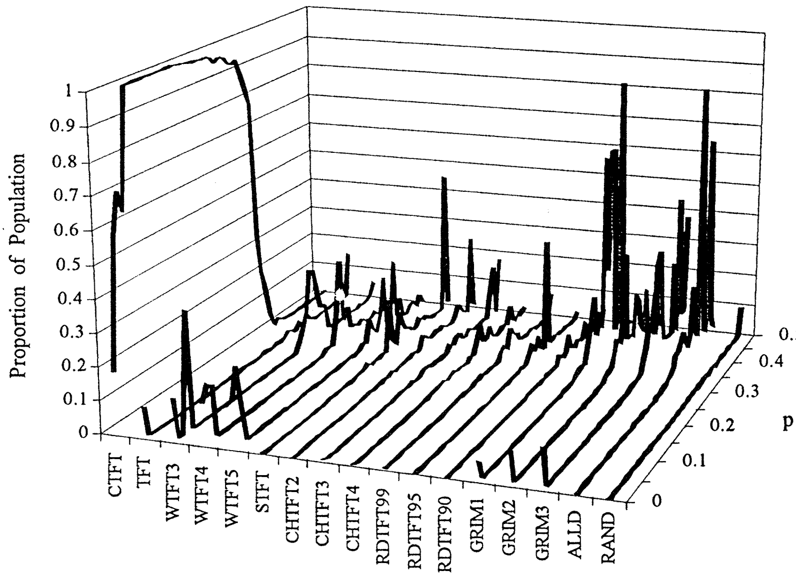

Figure 11: Population Dynamics of Territory with Neutral Noise

POSITIVE NOISE

I now turn to asymmetric noise, starting first with positive noise—that is, where there is some probability p_{ε_+} that an intended defection will be implemented as a cooperation. As before, I examine the effects of positive noise on homogeneous and heterogeneous systems, and also in the subsequent ecology and territory models.

Figure 12 displays the average payoffs of strategies when played against themselves under conditions of positive error. Here, any $N\bar{E}$ strategy will cooperate every time with itself, receiving the reward payoff. Therefore, adding positive noise to the environment does not provide $N\bar{E}$ strategies anything over what they already obtain in a noiseless environment. However, the $\bar{N}E$ strategies do benefit from positive error, as was the case in the homogeneous environment under neutral noise. TFT-based exploiters benefit for the same reason mentioned earlier: the positive noise allows them to break out of both permanent defections and unsynchronized retaliations. In contrast to the neutral noise environment, here there is no negative element to the noise, so there is no stimulus back toward defection outside of the strategy itself. Using the model in the appendix, one sees that RAND varies from 2.25 at $p_{\varepsilon_+} = 0$ to 2.68 at $p_{\varepsilon_+} = .50$. Similarly, ALLD varies linearly from 1.00 at

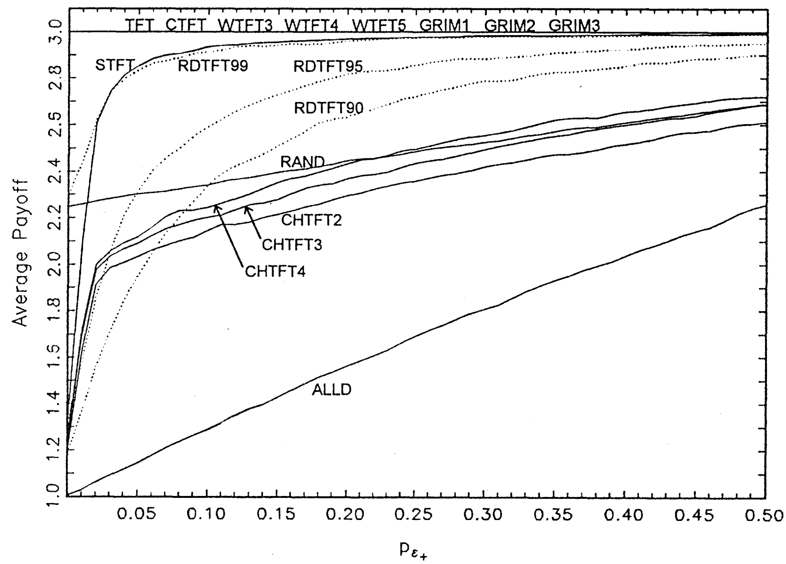

Figure 12: Average Payoff for Each Strategy in a Homogeneous System under Positive Noise

$p_{\varepsilon_+} = 0$ to 2.25 at $p_{\varepsilon_+} = .50$. Finally, the within-group rankings of parametric strategies stay the same as in the environment with neutral noise—namely, the more a strategy exploits, the worse it does against itself.

Figure 13 displays the average payoffs of each strategy played in a heterogeneous system with positive noise. In general, relative to the noiseless tournament, all strategies benefit from the positive noise, even those whose performance degrades in moving from a homogeneous system to a heterogeneous tournament. Some, like the *NPFĒ* strategies, benefit only marginally from the positive noise, whereas others, like RAND and (especially) ALLD, reap large benefits. As in the neutral noise environment, one sees different classes of strategies performing better than others in different regions of noise. For example, the *NPFĒ* strategies fare the best in the region $0 \leq p_{\varepsilon_+} \leq .22$, whereas the GRIM*n* strategies overtake them in the region $.22 \leq p_{\varepsilon_+} \leq .30$ and surpass them thereafter. Finally, in terms of within-parametric class rankings, one sees some deviations from the neutral noise case. Although the GRIM*n* and RDTFT*p* strategies hold to the same ranking, the CHTFT*n* strategies reverse: here, CHTFT2 performs better than CHTFT3, which performs better than CHTFT4. Finally, the WTFT*n* strategies are so close in performance that it is difficult to determine the rankings.

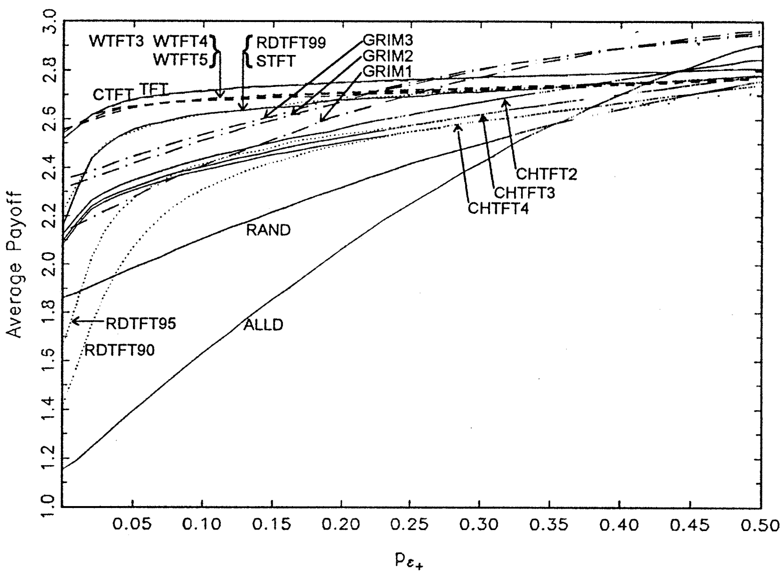

Figure 13: Average Payoff for Each Strategy in the Heterogeneous System under Positive Noise

What accounts for this shift from the $NPF\bar{E}$ to the $NP\bar{E}$ strategies as the positive noise increases? It was noted that because both the $NPF\bar{E}$ and $NP\bar{F}\bar{E}$ strategies are nice, they will always cooperate under conditions of solely positive noise (as well as under noiseless conditions). Because the $NPF\bar{E}$ and $NP\bar{F}\bar{E}$ strategies will achieve exactly the same payoffs when playing against each other, the determining factor in their relative performance is how each group performs against the $\bar{N}E$ strategies. Specifically, forgiveness is the characteristic that determines whether the $NPF\bar{E}$ or $NP\bar{F}\bar{E}$ strategies perform better against the $\bar{N}E$ strategies over the range of positive noise.

To see this, consider the interactions of WTFT*n* and GRIM*n* versus CHTFT*n*. For low levels of positive noise, WTFT*n* versus CHTFT*n* will never enter into states of permanent defection but, rather, will alternate between periods of cooperation and unsynchronized retaliation, with less cooperation than unsynchronized retaliation. In contrast, GRIM*n* versus CHTFT*n* will result in GRIM*n* quickly being triggered into permanent defection. Once this occurs, positive noise may infrequently cause CHTFT*n* to cooperate while GRIM*n* is defecting, but the situation immediately returns to permanent defection. Similarly, GRIM*n* may infrequently accidentally cooperate. This will cause CHTFT*n* to cooperate on the next move, but

GRIM*n* will return to defection, and the system will then return to a state of mutual defection. At low levels of noise, the frequent state of unsynchronized retaliation, peppered with periods of mutual cooperation, will yield higher payoffs to WTFT*n* than GRIM*n* will receive from a general state of mutual defection that is interspersed with infrequent temptation payoffs. Thus, for low levels of noise, the $NPF\bar{E}$ strategies (all of which are provocable here) reap higher average payoffs against the $\bar{N}E$ strategies because they allow for unsynchronized retaliations and mutual cooperation, in contrast to the $NP\bar{F}\bar{E}$ strategies, which on average do not receive enough temptation payoffs to overcome the payoffs from mutual defection.

However, at high levels of positive noise, the situation is reversed. Here, WTFT*n* and CHTFT*n* will again never enter into mutual defection, but the payoffs will, as the positive noise increases, tend toward mutual cooperation rather than unsynchronized retaliation. These seem to be high payoffs, but the GRIM*n* strategies seem to do better yet. In the case of very high levels of positive noise, GRIM*n* versus CHTFT*n* will generally not have extended periods of mutual defection. Instead, there will be more frequent occurrences of one or both strategies accidentally cooperating. Because CHTFT*n* is forgiving, whereas GRIM*n* is not, CHTFT*n* will be exploited by GRIM*n*.[26] In the case in which only CHTFT*n* cooperates, GRIM*n* will receive the temptation payoff, and both will attempt to mutually defect the next period. In the case in which only GRIM*n* cooperates, CHTFT*n* will cooperate in the next period, but GRIM*n* will defect. And, in the case in which both accidentally cooperate, CHTFT*n* will cooperate in the next period, whereas GRIM*n* will defect. Thus, as the positive noise level increases, GRIM*n* receives fewer mutual defection payoffs, more unsynchronized retaliation payoffs, and because it is unforgiving, even more temptation payoffs. Even though the average unsynchronized retaliation payoff is by construction less than the mutual cooperation payoff, as the level of noise becomes high, the frequent temptation payoffs plus the unsynchronized retaliation payoffs of GRIM*n* average to more than mutual cooperation payoff of WTFT*n*.

Turning to the population dynamics, one sees in Figure 14 that an ecology with 1% positive noise does not differ much from a noiseless ecology: steady state produces an ecology consisting of $NP\bar{E}$ strategies, with the $NPF\bar{E}$ strategies having the highest levels of representation. This basic trend continues in ecologies through 10% and 20% positive noise, with the GRIM*n* strategies rising in proportion as the level of noise increases. Figure 15 shows

26. The reader may feel some tension, as I do, that the characteristic of exploitation appears to be context specific. My original specification of GRIM*n* as nonexploiting was done with respect to a noiseless environment. However, here the success of GRIM*n* at high positive noise levels is due mainly to the fact that it is unforgiving, which allows it to exploit forgiving strategies once it has been triggered. Although I do not attempt it here, this would indicate the need for more precise—perhaps context-specific—classifications of the strategies.

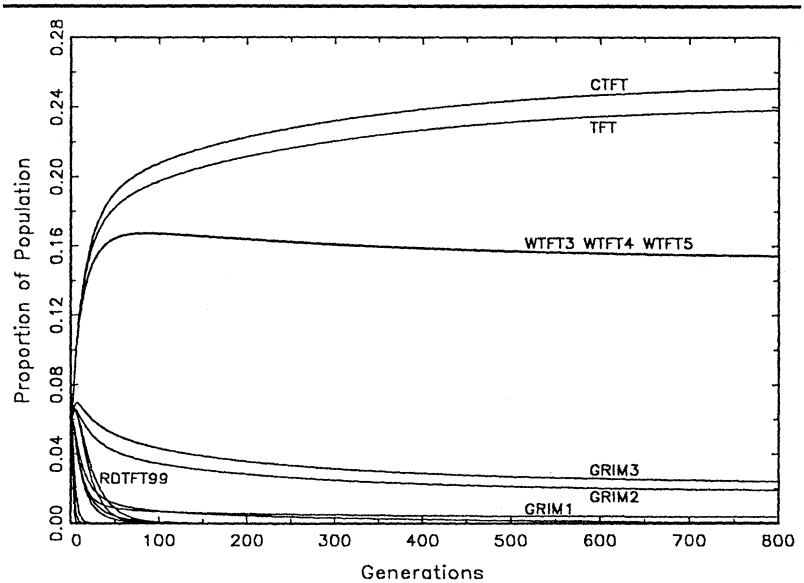

Figure 14: Population Dynamics of Ecology with Positive Noise, $p_{\varepsilon_+} = .01$

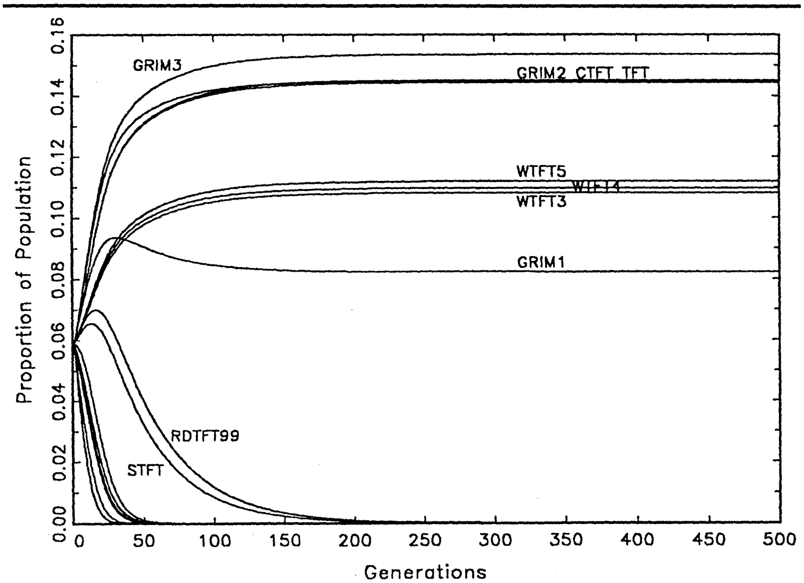

Figure 15: Population Dynamics of Ecology with Positive Noise, $p_{\varepsilon_+} = .30$

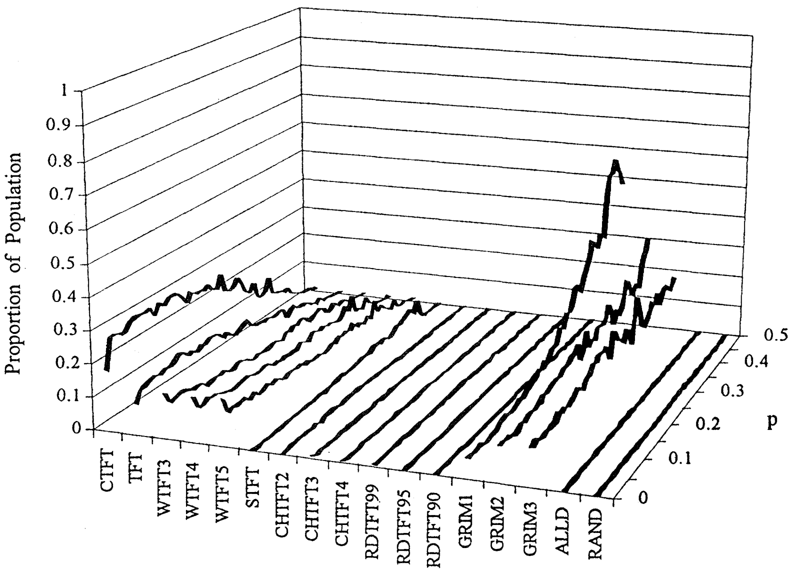

Figure 16: Population Dynamics of Territory with Positive Noise

that at 30% noise, the $NPF\bar{E}$ strategies are still well represented in the population, but now GRIM3 is the largest, and GRIM2 is tied for second with CTFT and TFT. By no means does any single strategy dominate the population. However, noticeably absent in the steady-state regions are any exploiters. Finally, note that the time to steady state is much shorter here than in the neutral noise ecology, which was characterized by large transition periods and wild fluctuations at high noise levels.

In the population dynamics of the randomized territory under positive noise, one sees results similar to those in the ecology. Over all noise levels, the territory converges to $NP\bar{E}$ strategies. The $NP\bar{E}$ strategies dominate because they will always cooperate with each other under conditions of positive noise. This payoff is higher than any the $\bar{N}E$ strategies will achieve, allowing the $NP\bar{E}$ strategies to spread throughout the territory. Once the $\bar{N}E$ strategies have disappeared, all the nice strategies will appear exactly alike, allowing a plurality of the $NP\bar{E}$ strategies to exist in the converged territory. However, just because a plurality of $NP\bar{E}$ strategies exists across all noise levels does not preclude any differentiation based on subclasses. In fact, one sees that different subclasses of $NP\bar{E}$ strategies fare better over different ranges of noise. For example, at low levels of noise, the $NPF\bar{E}$ strategies have the highest representation in the population. As the positive noise increases, the $NPF\bar{E}$ strategies slowly decrease in proportion, whereas the $NP\bar{F}\bar{E}$ (i.e., GRIMn) strategies increase. At high levels of noise, the $NP\bar{F}\bar{E}$ strategies dominate, with GRIM1 comprising up to 50% of the population.[27]

The reason one sees an increase in $NP\bar{F}\bar{E}$ strategies in the converged population as the level of noise increases is the same as that described for the shift in the heterogeneous system. At lower levels of noise, $NPF\bar{E}$ strategies fare better against the $\bar{N}E$ strategies and, therefore, spread faster in the population than do the $NP\bar{F}\bar{E}$ strategies. When the $\bar{N}E$ strategies have died out, the $NPF\bar{E}$ and $NP\bar{F}\bar{E}$ strategies are left perpetually cooperating with each other, stabilizing the population. Conversely, at higher levels of positive noise, it is the $NP\bar{F}\bar{E}$ strategies that fare better against the $\bar{N}E$ strategies, resulting in a higher proportion of $NP\bar{F}\bar{E}$ strategies at convergence.

NEGATIVE NOISE

The second type of asymmetric noise, negative noise, causes an intended cooperation to be implemented as a defection. Playing strategies against themselves in this environment yields a number of interesting results. As Figure 17 shows, with the exception of RAND, the best performing strategies are $NP\bar{E}$. Among these, $NPF\bar{E}$ strategies fare better, and CTFT fares best by far. However, most strategies are severely affected by the negative noise, and the punishment payoff is a lower bound toward which most strategies tend. Although the worst performing strategies are the exploiters, all generally do poorly. At only a 5% level of negative noise, most of the strategies are reduced to an average payoff of less than 1.4. Nasty strategies such as ALLD and STFT will maintain perpetual defection against themselves. At low levels of noise, the GRIMn strategies quickly become undifferentiable from ALLD. The $PF\bar{C}$ strategies succumb to unsynchronized retaliations and perpetual mutual defections; for these strategies, there is no counteracting positive stimulus to shock them from perpetual defection into unsynchronized retaliations and then into cooperation. Because of their generosity, the WTFTn strategies fare better than others, but at approximately 7% error, they are surpassed in performance by RAND and at 15% error are also reduced to an average payoff of less than 1.4.

Of the strategies, CTFT is clearly the most robust in a homogeneous environment, decreasing seemingly linearly in average payoff from 3 at $p_{\varepsilon_-} = 0$ to a little over 2 at $p_{\varepsilon_-} = .5$, whereas most other strategies are already below

27. The positive noise environment also does not experience the shift in generations to convergence that the neutral noise environment did. Here, the average generations to convergence is 18.1 for the region $0 \leq p_{\varepsilon_+} \leq .2$, rising only slightly to 23.1 for the region $.21 \leq p_{\varepsilon_+} \leq .5$.

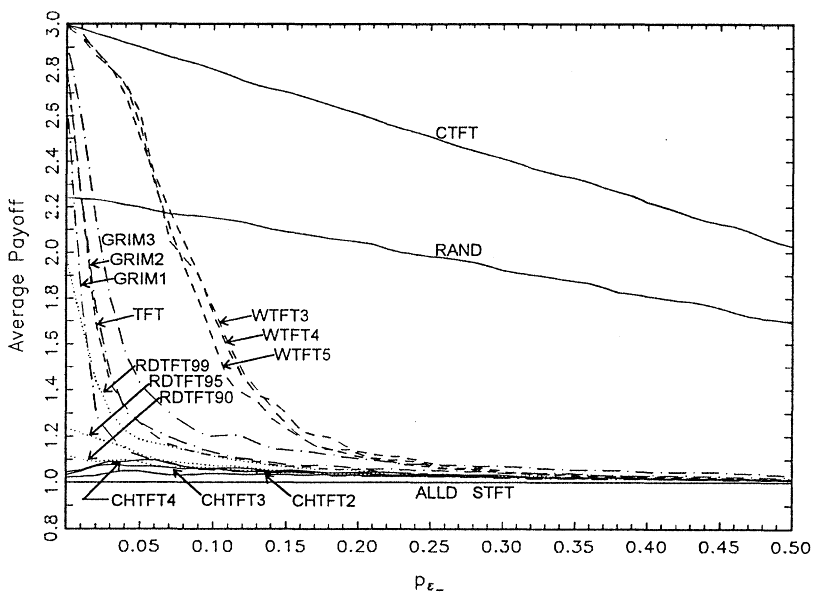

Figure 17: Average Payoff for Each Strategy in a Homogeneous System under Negative Noise

1.4 with only 5% noise in the system.[28] CTFT's high performance is due to its willingness when it is not in good standing to take the sucker's payoff to resume cooperation. In effect, when two players in an environment of negative noise both employ CTFT, no state of permanent defection will ever result, nor will they ever enter into unsynchronized retaliations. Rather, if uncertainty causes a spurious defection, the defecting player will take the sucker's payoff in the next iteration to put itself back in good standing so that cooperation may resume. Thus, even at high levels of negative noise, the two players will tend to alternate receiving the temptation and sucker payoffs, rather than permanently defecting. The point $p_{\varepsilon\_} = .25$ is noteworthy for CTFT playing against itself, because the average payoff at this point is approximately 2.5—what each player would receive by alternatingly receiving the temptation and sucker payoffs. In the region $p_{\varepsilon\_} > .25$, the additional noise decreases the average payoff below that obtained by alternating the tempta-

28. Again, it is robust in the sense that noise distribution does not affect its ranking as a top-performing homogeneous system relative to the other homogeneous systems examined. However, no claims are presented here as to whether a CTFT system is robust in the sense that it cannot be invaded by another strategy.

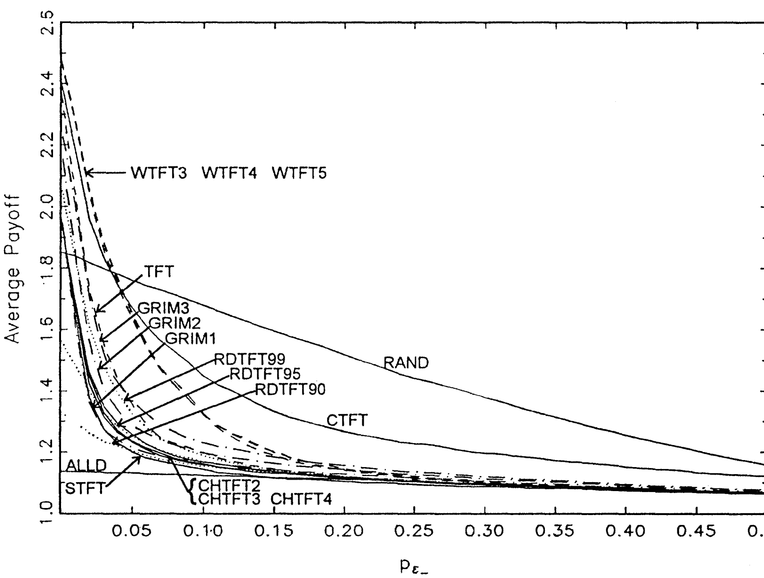

Figure 18: Average Payoff for Each Strategy in the Heterogeneous System under Negative Noise

tion and sucker payoffs, but not nearly to the same extent as for the other strategies.

Interestingly, in homogeneous play, the other robust strategy is RAND. Although it obviously does not operate on the same principle as CTFT, it follows a similar almost linear degradation in average payoff, starting at 2.25 for $p_{\varepsilon\_} = 0$, decreasing to 1.98 for $p_{\varepsilon\_} = .25$, and then to 1.68 for $p_{\varepsilon\_} = .5$. RAND's payoff is not affected to the same extent as the provocable strategies' or ALLD's payoffs since it is not provocable, but it is still forgiving—and, therefore, will not degenerate into perpetual mutual defection. Thus for the homogeneous systems over a wide range of negative noise levels, flipping a coin beats every other strategy here except for CTFT. In terms of the rankings of parametrically grouped strategies, the patterns are generally the same here as before.

As Figure 18 shows, negative noise also severely degrades the performance of strategies in the heterogeneous system. With only 5% negative error in the system, the average payoff of most strategies is below 1.4, heading quickly toward the punishment payoff. As with the homogeneous examples, the *NPF$\bar{E}$* strategies and RAND are also the top performers in the heterogeneous system. Similarly, the exploiters other than RAND fare poorly. For low levels of noise, the WTFT*n* strategies receive the highest average payoff. However, they are surpassed by RAND at about 4% error and CTFT at 5% error. None of the *NPF$\bar{E}$* strategies performs nearly as well as they do in the homogeneous system. With 10% negative noise, the WTFT*n* strategies are reduced to an average payoff of barely more than 1.3, whereas CTFT is reduced to the same with only 15% noise in the system. In moving from a homogeneous system, CTFT loses that which allows it to perform so well in the homogeneous system: contrition on the part of both players, ensuring that neither unsynchronized retaliations nor perpetual defection will occur. Thus, in the heterogeneous system, CTFT is not able to protect its overall payoff from the pathology resulting from the spurious defections of others.

Surprisingly, for all but very low levels of noise, RAND turns out to be the best performer in the tournament under negative noise. Even more so than in the homogeneous system, it appears that flipping a coin becomes a better strategy than being either strategically "hard-nosed" or forgiving and generous. The reasons for this can be seen by examining the dynamics between the *PF* strategies and the dynamics between the *PF* strategies and RAND. As was shown in the homogeneous system under negative noise, the *PF$\bar{C}$* strategies will quickly degrade into unsynchronized retaliations and perpetual mutual defections when playing against themselves. This dynamic also tends to occur when *PF* strategies play each other in the heterogeneous system—including when CTFT plays *PF$\bar{C}$* strategies.

In contrast, whereas RAND does not perform spectacularly against itself, its payoff is not bad relative to the homogeneous payoffs of the *PF$\bar{C}$* strategies. In the heterogeneous system, RAND's payoffs against the *PF* strategies (and vice versa) are also generally better than those of the *PF* strategies against the *PF$\bar{C}$* strategies. In the sample of strategies examined here, RAND's average performance in the heterogeneous system is then the average of mostly mediocre scores—the results of playing itself and mostly *PF* strategies. In contrast, the average performance of the *PF* strategies is the average of generally very low scores—the results of playing mostly other *PF$\bar{C}$* strategies. Thus Figure 18 shows RAND performing better than all other strategies over most of the negative noise range.

The population dynamics of an ecology with 1% negative noise are shown in Figure 19. Interestingly, the dynamics in Figure 19 are almost identical to those in the ecology with 1% neutral noise (see Figure 7). Here, as in the neutral noise ecology, one sees a hegemony of CTFT, with initial representations of other *NPF$\bar{E}$* strategies during a transition period. Figure 20 shows that as one moves to 10% negative error, RAND dominates during a transition period until the strategies against which it performs well die out. As RAND declines, CTFT rises, eventually dominating the system in a stable manner. This dynamic of a dominating but transitory RAND giving way to

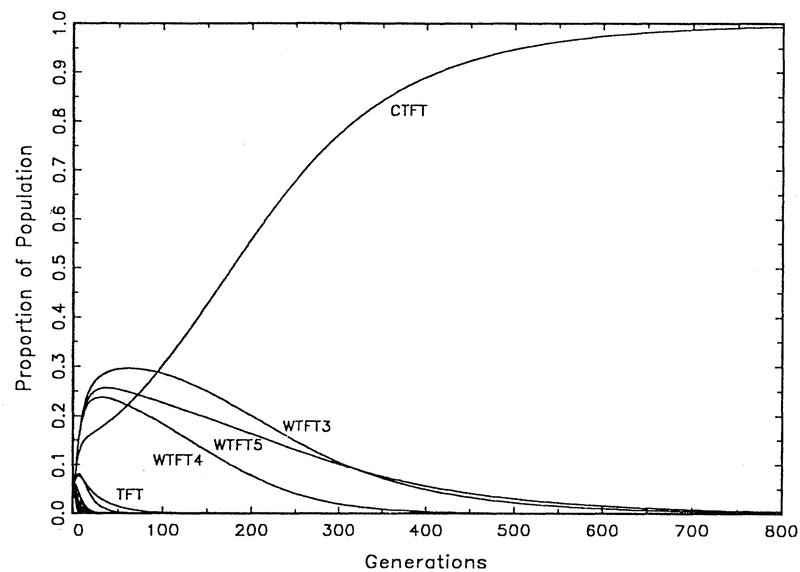

Figure 19: Population Dynamics of Ecology with Negative Noise, $p_{\varepsilon\_}$ = .01

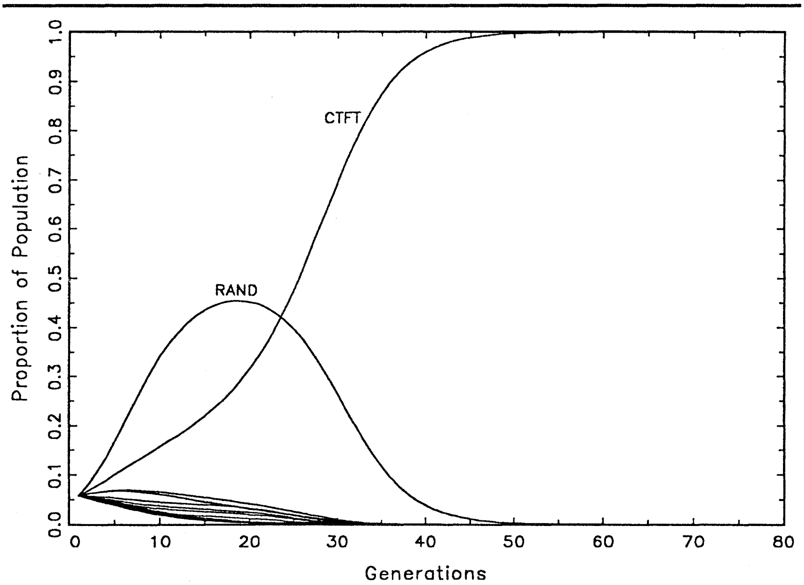

Figure 20: Population Dynamics of Ecology with Negative Noise, $p_{\varepsilon\_}$ = .10

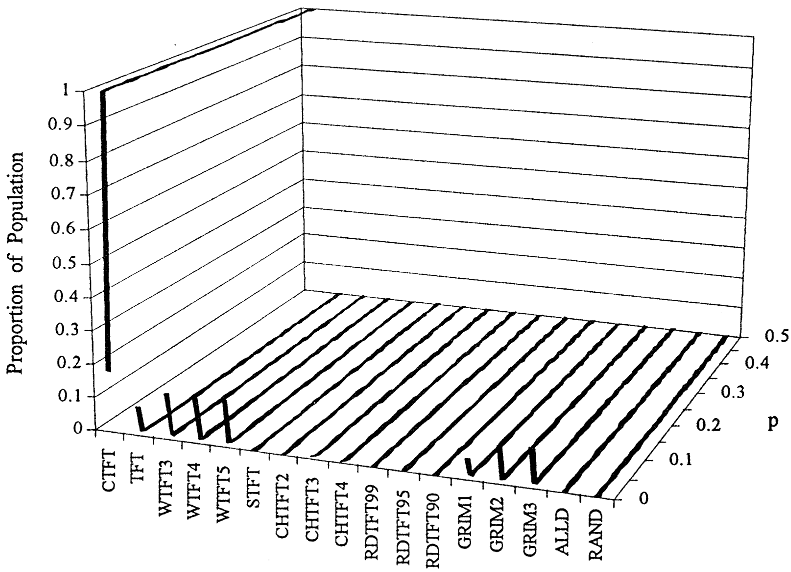

Figure 21: Population Dynamics of Territory with Negative Noise

a steady state of CTFT continues up through 30% negative noise. That CTFT should so dominate the ecology is somewhat surprising given that RAND outperformed CTFT in the heterogeneous system. Another interesting aspect of these ecologies is the extremely short time it takes for them to stabilize. In the neutral and positive environments, it was not uncommon for the system to take 500 to 2,000 generations to stabilize, whereas here, for moderate to high levels of noise, it takes less than 80 generations.

The territory dynamics in Figure 21 further emphasize the overwhelming hegemony of CTFT. Here, one sees that once even a small amount of negative noise is introduced, CTFT completely dominates the territory—and this holds all the way up to a 50% level of noise.[29] Given the tournament results, this again comes as somewhat of a surprise, because RAND was relatively more robust to negatively biased noise in the heterogeneous system.

One can see why this should be the case, however, by looking back on the analysis of the homogeneous and heterogeneous systems. In the heterogeneous system, all *PF* strategies do much worse than RAND for a wide range of noise levels. At these noise levels, because RAND performs better relative to these strategies, it initially grows, whereas the $PF\bar{C}$ strategies decrease in proportion. However, as the $PF\bar{C}$ strategies begin to die out, their deleterious

29. The time to convergence is also rather quick: for a large region of p_{ε_-}, it is between 9 and 10 generations, and it never requires more than 14 generations.

TABLE 4
Summary of Strategies' Performances
in Homogeneous and Heterogeneous Systems

	Homogeneous Population	*Heterogeneous Population*
No noise	• Best: $N\bar{E}$ • Worst: $\bar{N}AE$ (then $\bar{N}E$)	• Best: $NPGF\bar{E}$ (then $NPF\bar{E}$, then $NP\bar{E}$) • Worst: $\bar{N}E$
Positive noise	• Best: $N\bar{E}$ • Worst: $\bar{N}A\bar{F}E$ (then $\bar{N}E$) • As p increases, payoffs tend toward the reward payoff	• Best: $NPF\bar{E}$ for $0 < p \leq .22$, GRIMn for $.30 < p < .50$ • Worst: $\bar{N}A\bar{F}E$ until very high levels of p • As p increases, payoffs tend toward the reward payoff
Negative noise	• Best: $NPFC\bar{E}$ • Worst: $\bar{N}AE$ (then $\bar{N}E$, excluding RAND) • Except for CTFT and RAND, all quickly degrade • As p increases, most payoffs quickly approach the punishment payoff	• Best: $NPF\bar{E}$ for $0 < p < .04$, RAND for $.04 \leq p < .50$ • Worst: $\bar{N}AE$ (then $\bar{N}E$, excluding RAND) • Except for RAND, all quickly degrade, including CTFT • As p increases, most payoffs quickly approach the punishment payoff
Neutral noise	• Best: $NPFC\bar{E}$ (then $NPGF\bar{E}$) • Worst: $A\bar{F}$ (then $\bar{F}$) • As p increases, all payoffs approach RAND payoff of $(T + R + P + S)/4$	• Best: $NPF\bar{E}$ for $0 < p \leq .26$, $\bar{F}$ for $.26 < p < .50$ • Worst: $\bar{F}$ for $0 < p \leq .26$, CTFT for $.26 < p < .50$ • As p approaches .5, all payoffs approach the RAND payoff of $(T + R + P + S)/4$ at $p = .5$

NOTE: Numerical values are based on simulations and are, therefore, approximate. Class: N = nice; A = nasty; P = provocable; G = generous; F = forgiving; C = contrite; E = exploiting; $\bar{X}$ = not X. Payoffs: T = temptation payoff; R = reward payoff; P = punishment payoff; S = sucker's payoff.

effect on CTFT's payoff also decreases. Because CTFT performs so well against itself, it then quickly grows in the population; because its homogeneous performance is much better than RAND's, it overtakes RAND and eventually dominates the system.

DISCUSSION

Tables 4 and 5 summarize the results of the previous section, displaying the major findings for the homogeneous systems, heterogeneous systems,

TABLE 5
Summary of Population Distribution after Convergence in Ecology and Territory Analyses

	Ecology	Territory
No noise	• Plurality of $NP\bar{E}$	• Plurality of $NP\bar{E}$
	• Largest proportion are $NPF\bar{E}$	
Positive noise	• Plurality of $NP\bar{E}$	• Plurality of $NP\bar{E}$
	• $NPF\bar{E}$ have largest representation at lower p. • However, $NP\bar{F}\bar{E}$ increase in proportion as p increases	• As p increases, $NPF\bar{E}$ decrease in proportion, whereas $NP\bar{F}\bar{E}$ increase
Negative noise	• Hegemony of CTFT	• Hegemony of CTFT for all but very small levels of p
	• As p increases up to .30, RAND becomes an increasingly larger part of the transition period but always eventually disappears, leaving only CTFT	
	• System stabilizes very quickly	
Neutral noise	• Hegemony of CTFT for $0 < p < .20$	• Plurality of $NP\bar{E}$ for very small p
	• For p near .25, CTFT begins to appear only in transition period, whereas $\bar{N}E$ appear in steady state	• Hegemony of CTFT for $.04 < p \leq .25$
	• Hegemony of $\bar{F}$ for $.27 \leq p \leq .35$	• Plurality of $\bar{N}E$ and $\bar{F}$ for $.25 < p < .35$
	• Oscillating system of GRIMn at $p = .35$	• Hegemony of $\bar{F}$ for $.35 \leq p \leq .50$
		• Time to convergence increases greatly as p increases

NOTE: Numerical values are based on simulations and are, therefore, approximate. Class: N = nice; P = provocable; F = forgiving; E = exploiting; $\bar{X}$ = not X.

ecology models, and territory models. Additionally, Table 6 summarizes the general effects of the noise environments on the dynamics of the strategies examined here. As the results of the preceding section show, the distribution of noise makes a big difference in the dynamics of how cooperation emerges and is maintained. Moreover, depending on the sample of strategies, certain

TABLE 6
General Effects of Noise Environments on Strategy Dynamics

Environment	*Effects*
No noise	Nice dyads will perpetually cooperate Nasty dyads will perpetually defect Exploitative strategies will trigger unsynchronized retaliations or perpetual mutual defections with provocable strategies
Positive noise	Nice dyads will perpetually cooperate, regardless of the level of positive stimulus For forgiving dyads, positive noise provides a stimulus out of perpetual mutual defection and unsynchronized retaliations Opens forgiving strategies to additional exploitation by unforgiving strategies Separates nasty strategies
Negative noise	Nasty dyads will perpetually defect, regardless of the level of negative stimulus Contrition prevents pathologies because of error on one's own part Generosity helps alleviate pathologies because of error on the part of others but opens one to increased exploitation Dyads with at least one provocable, noncontrite strategy degenerate into unsynchronized retaliations and perpetual mutual defection Forgiving but unprovocable randomized strategies (e.g., RAND) will not fall prey to noise pathologies but will also not perform as well as a contrite dyad Separates nice strategies
Neutral noise	Positive component Provides a stimulus out of perpetual mutual defection and unsynchronized retaliations among forgiving strategies Opens forgiving strategies to additional exploitation by unforgiving strategies Separates nasty strategies Negative component Can trigger unsynchronized retaliations and perpetual mutual defection among provocable strategies Contrition prevents pathologies because of negative error on one's own part Generosity helps alleviate pathologies because of negative error on the part of others but opens one to increased exploitation Separates nice strategies

strategy characteristics will be more important in the development and maintenance of cooperative behavior in different noise environments.

In the positively biased noise environment, the positive noise does not provide $N\bar{E}$ dyads anything beyond what a noiseless environment already allows. $N\bar{E}$ strategies will perpetually cooperate with each other, so the stimulus from defection to cooperation gives no additional utility for the strategies in these dyads. However, positive noise can benefit both $N\bar{E}$

strategies and $\bar{N}E$ strategies in heterogeneous dyads. For an $NPF\bar{E}$ strategy playing an $\bar{N}PFE$ strategy, the positive noise can provide a stimulus out of mutual defections and unsynchronized retaliations. At moderate to high levels of positive noise, *$NP\bar{F}\bar{E}$ strategies benefit when playing against $\bar{N}PFE$* strategies because the forgiveness on the part of the $\bar{N}PFE$ strategies allow them to be exploited at a greater rate than the $NP\bar{F}\bar{E}$ strategies. Overall, the exploiters are helped the most by positive noise, but relative only to a noiseless environment, not relative to $N\bar{E}$ strategies. Thus in the evolutionary models with positively biased noise, the populations converge to a plurality of $N\bar{E}$ strategies.

In the negatively biased noise environment, everyone generally loses the more noise is increased—the exceptions being the nasty strategies, which perpetually defect in a noiseless environment and are, therefore, unaffected by increases in negative noise. Especially susceptible are dyads with at least one provocable, noncontrite strategy. These will degenerate into unsynchronized retaliations and extended mutual defection. However, if to err is human and forgiveness divine, then generosity and contrition may be even more divine, for these are two ways of alleviating the pathologies of negatively directed error. Generosity dampens the pathologies arising from error on the part of the other player. Contrition prevents the pathologies from arising because of error on one's own part. Because great generosity is required to achieve the same effect as contrition and because generosity opens one to increased exploitation,[30] contrition appears to be the more efficient means of alleviating the potential for unsynchronized retaliations and extended mutual defection.[31] In the analysis here, a CTFT dyad has been shown to be highly robust in the face of negative noise, primarily because each side accounts for its own negative error—thus allowing cooperation to be reestablished. Because of this robustness under negative error, one sees the evolutionary models invariably converge to complete dominance by CTFT.

When combined, the two types of noise do not simply cancel each other out. Rather, in the neutral noise environment, one sees the effects of each noise type, similar to how each type affects the strategies in its individual asymmetric noise environment. For example, the positive component provides a stimulus out of perpetual mutual defection and unsynchronized retaliations among forgiving strategies, whereas the negative component can trigger unsynchronized retaliations and perpetual mutual defection among provocable strategies. The payoffs of $\bar{N}E$ strategies are helped by the addition

30. Turning the other cheek may result not in a resumption of cooperation but simply in the other cheek being slapped.

31. This, however, begs the question of whether there are costs associated with obtaining information under a contrite strategy—to ensure that one can know when one has accidentally defected. If so, then these costs will cut against the efficiency of a contrite strategy versus a generous one.

of neutral noise (because of the positive noise component), whereas those of $N\bar{E}$ strategies are degraded (because of the negative component). Although one again sees that a CTFT dyad is most resistant to the pathologies of the negative component, the evolutionary models show a convergence of the population to CTFT for low to moderate levels of neutral noise, but giving way to a plurality of generally unforgiving strategies for higher levels of noise.

All of this suggests one should be careful in how one models noise—that one's assumptions of noise distribution should be based on the substantive issue under examination, so that the prescriptions derived from the model may have greater external validity. Empirically, one may not expect many situations in international politics to include much (if any) positively directed noise. Assumptions based on entirely asymmetric negative noise or noise that includes both negative and positive components, but is heavily skewed toward the negative side, may be more appropriate in these cases.

A ROBUST HEURISTIC FOR A BOUNDEDLY RATIONAL WORLD

Given that different distributions of noise affect the performance of strategies in different ways, can anything be said about which strategy might be a robust heuristic across all noise environments in this artificial, boundedly rational world? Or does each noise environment require a different strategy for success? As Tables 4 and 5 show, CTFT is relatively robust across all noise environments and for all the analyses run here. For example, in a homogeneous population, it fares as well as any other $N\bar{E}$ strategy in noiseless and positive noise environments—and is very robust in the negative and neutral noise environments. In the heterogeneous population, CTFT is robust in noiseless environments and at low noise levels in positive and neutral noise environments. However, it performs almost as poorly as the other strategies under solely negative noise—and it is outperformed by RAND at all but small probabilities of negative error. Although this latter result might weaken one's confidence in CTFT as a robust strategy, the ecology and territory models support the notion that CTFT is, in fact, robust. As Table 5 displays, for noiseless and positive noise conditions, CTFT is among the plurality of strategies present in the system; for the negative noise environment, CTFT maintains a hegemonic proportion of the population across all levels of noise; and, when neutral noise is present, CTFT dominates for low to moderate levels of noise. Thus CTFT appears to perform well under varying types and levels of noise and, under the assumptions of the evolutionary models, generally either coexists among a plurality of similarly nice and nonexploiting strategies or spreads throughout the population completely.

There are a number of reasons for CTFT's success. Because it is nice, it will always cooperate with itself and other nice strategies under noiseless and positive noise conditions. It is conditionally retaliatory, so it is resistant to exploitation. It is forgiving, so it will resume cooperating once its opponent cooperates. It is nonexploiting, so it will never intentionally start a sequence of unsynchronized retaliations or mutual defection. Finally, an act of contrition will be made for any accidental defection on its part (e.g., under negative and neutral noise environments), generally resulting in a return to cooperation when playing against strategies that are both provocable and forgiving.

Although contrition allows CTFT to protect itself against the pathologies associated with its own accidental defections, CTFT cannot protect itself against the unsynchronized retaliations that ensue when its opponent accidentally defects. In fact, CTFT's retaliatory character is partly responsible for this dynamic. To mitigate the pathological effects of accidental defections on the part of others, one would need to relax CTFT's retaliatory character to one that is only provocable and, in doing so, also add generosity. Based on the simulation results and the analysis of which characteristics foster cooperation under different types of noise, one might hypothesize that a $NPGFC\bar{E}$ strategy—for example, a generous CTFT strategy—would be even more robust to the effects of noise than CTFT. However, the performance of such a strategy—and even the manner in which one were to implement generosity—would be ecology dependent. In the sample of strategies examined here, an $NPGFC\bar{E}$ strategy would probably do very well. On the other hand, in an ecology of ALLD's, an $NPGFC\bar{E}$ strategy would most likely fare worse than CTFT, because generosity opens one to exploitation.

INCENTIVES FOR INTERNATIONAL INSTITUTIONS

In the second section, I noted that the strategies may be viewed as rule-based behavior and that norms or institutions of bilateral interaction may be represented as homogeneous systems of this rule-based behavior. Thus defined, for an institution to exist in a population, it is not required that every actor's strategy be identical. Rather, if there is some behavioral characteristic that is common to all members, then one can say that that behavior is a norm or has become institutionalized. This allows for different degrees of institutionalization of behavior in populations. For example, the $N\bar{E}$ institution is fairly weak, because there are a number of different strategies that share these characteristics. Whereas actors with $N\bar{E}$ norms will behave similarly under certain conditions, their behavior under other conditions will diverge—sometimes drastically. In contrast, a population composed of members with $NPGF\bar{E}$ behavior would have a higher degree of institutionalization. One

TABLE 7
Comparative Statics of Institutional and Noise Incentives

	Institutional Incentives, Holding Noise Constant	*Noise Incentives, Holding Institutions Constant*
No noise	Move to $N\bar{E}$ institution[a]	$N\bar{E}$ institutions: no incentive to increase any type of noise; however, indifferent about positive noise $\bar{N}E$ institutions: increase positive noise
Positive noise	Move to $N\bar{E}$ institution[a]	$N\bar{E}$ institutions: indifferent about positive noise $\bar{N}E$ institutions: increase positive noise
Negative noise	Move to $NPFC\bar{E}$ institution[b]	A institutions: indifferent about negative noise $\bar{A}$ institutions: decrease negative noise
Neutral noise	Move to $NPFC\bar{E}$ institution[b]	$NPF\bar{E}$ institutions: decrease neutral noise $\bar{N}E$ or $\bar{F}$ institutions: increase neutral noise

NOTE: N = nice; A = nasty; P = provocable; F = forgiving; C = contrite; E = exploiting; $\bar{X}$ = not X.
a. May require $NPF\bar{E}$ characteristics to deal with exploiters during institutional formation and maintenance.
b. $NPGFC\bar{E}$ haracteristics may be helpful during early (very heterogeneous) phase of institutional formation.

composed of members that all employed an identical strategy could be considered the highest form of institutionalization examined in this study.

When viewed from this perspective, a number of interesting questions arise. First, does each noise environment examined here display a clear incentive to choose one institution over another, and is there a robust institution that fares well in all environments? Second, given the different institutions, what are the incentives for increasing or decreasing the types of noise in each noise environment? Third, based on the homogeneous and heterogeneous system results, are there incentives to move from heterogeneous systems of interaction to a common form of bilateral interaction? Finally, does one see the boundedly rational actors of the evolutionary models moving from heterogeneous to homogeneous behavior?

To answer the first two questions, I take the results of the homogeneous systems and examine the incentives to move to a given institution at a given noise type and level. Similarly, for each of the institutions in a given noise environment, I examine the incentives for increasing or decreasing the level of that noise type. Table 7 provides summary comparative statics of these incentives. The second column shows the institutional incentive, holding level of noise constant, whereas the third column shows the noise incentive, holding institutions constant.

Concerning the question of institutional incentives, one sees that homogeneous systems of $N\bar{E}$ behavior reap the greatest rewards in noiseless and positive noise environments. They do so because the members of $N\bar{E}$ institutions will perpetually cooperate with each other, whereas the $\bar{N}E$ institutions examined here will often give rise to unsynchronized retaliations or periods of mutual defection—and any gains from exploitation will be offset by the losses due to noncooperation. Yet, saying that there is an incentive to move to *any* $N\bar{E}$ institution is very different than saying that one *can* move to an $N\bar{E}$ institution in any particular system of strategies, even those presented here. For example, a heterogeneous system during the institutional formation phase may include exploiters against whom the $N\bar{E}$ strategies must protect themselves—here, by being provocable. Similarly, once the institution has been established, the possibility for exploiters to arise may require provocability to be retained, even though it is not necessary for cooperative interaction between those within the institution itself.

When negative noise is present, the situation becomes more complicated than the noiseless or positive noise environments, because the $N\bar{E}$ norms are embedded in different strategies, most of which also include provocable or retaliatory behavior. Here, forgiveness is required so that there is some chance of cooperation resuming after an accidental defection. However, even with forgiveness, members of such a system will still fall prey to the unsynchronized retaliations and mutual defections that arise from negatively directed spurious defections. Under neutral noise, norms of generosity and contrition help compensate for this, with CTFT having the highest payoff as an institution. However, when negative noise alone is present, contrition appears to be the only deterrence to these pathologies—although generosity may be helpful during the formation phase of the institution.

Looking at each of the rows in the second column of Table 7, one notices that there is a common institution that meets each of the noise environments' incentives: CTFT. Under noiseless and positive noise, members of a CTFT institution will perpetually cooperate with each other—and will be resistant to exploitation during institution formation and maintenance. Under negative and neutral noise, members of a CTFT institution will always apologize for accidental defections on their own part, removing the possibility of entering into unsynchronized retaliations or permanent defection—and thus providing the opportunity to renew cooperation.

Turning to the third column of Table 7, I now examine the incentives for the level of noise in each environment while holding institutions constant. As before, large classes of institutions react similarly to the noise levels, so one may group them in the analysis. Here, one sees that in a noiseless environment, the $N\bar{E}$ institutions have no incentive to increase any type of noise, whether positive, neutral, or negative, but they are indifferent to all positive noise for the reasons previously mentioned. However, the $\bar{N}E$ institutions have an incentive to increase positive noise. The incentives for the positively biased noise environment are similar to those of the noiseless environment. Here, the $N\bar{E}$ institutions are indifferent to the positive noise, whereas the $\bar{N}E$ institutions benefit from increasing the level of noise. In contrast, in a negative noise environment, almost all institutions have an incentive to decrease the level of noise—the exceptions being nasty institutions, which are indifferent to increases of negative noise. In the neutral noise environment, $NPF\bar{E}$ institutions have an incentive to decrease neutral noise, whereas $\bar{N}E$ and $\bar{F}$ institutions have an incentive to increase neutral noise. The former suffer more from the deleterious effects of the negative noise component. The latter benefit more from the positive noise component. Finally, holding the institutions constant and comparing the incentives for different noise types—as opposed to levels within types—one sees that there is generally an incentive to move from a negatively biased environment to a neutral noise environment. For $N\bar{E}$ institutions, there is then an incentive to move from a neutral noise environment to either a positively biased environment or a noiseless environment, whereas for $\bar{N}E$ institutions, there is an incentive to move to a positively biased environment.

It has been shown that CTFT is a robust institution across all noise environments. Yet, turning to the third question, one might still ask, Do states have an incentive to move from a system with heterogeneous modes of interaction to one with institutionalized interaction—for example, CTFT? Put differently, when will a state agree to an international institution? The answer is that it depends on whether the gains from exploitation are greater than the gains from cooperation, and that depends on the sample of strategies being used by the other nations in the international system. One can think of a number of cases in which a state would prefer a heterogeneous system to a homogeneous system of any kind. For example, for a wide range of noise distributions, a member playing ALLD in a system in which everyone else unconditionally cooperates would not prefer to move to any norm of interaction—including CTFT. Thus just as statements concerning an "optimal" strategy are conditional on the sample of strategies to which it is compared, so too are statements about the incentives to move from heterogeneous to homogeneous rules of behavior. Having noted that, conditional on the homogeneous and heterogeneous analyses conducted in this study, there appear to be incentives under all noise distributions to move from heterogeneous systems to a CTFT institution: there is almost no case of a strategy faring better in the heterogeneous system than CTFT does in the homogeneous system.[32]

32. The only exceptions to this are under neutral noise for $p_{\varepsilon_+} = .49$ and $p_{\varepsilon_+} = .50$, in which a number of strategies obtain higher average scores in the tournament than CTFT does against

Does one see the actors of the evolutionary models moving from heterogeneous to homogeneous modes of interaction? Generally, yes. Although Table 5 implies that the ecologies and territories converge to a plurality of strategies under noiseless and positive noise environments, the populations, in fact, converge on behavior that is nice and nonexploiting. Provocability is required during the evolution of a population to protect against exploitation. However, once the population has converged on the norm of niceness and nonexploitation, its members will henceforth perpetually cooperate, and all other characteristics become irrelevant to the maintenance of this cooperation.[33] Moreover, Table 5 shows that the populations under negative noise converge only to CTFT behavior at all but the smallest levels of noise. Finally, in the neutral noise environment, the population converges to a CTFT norm for low to moderate amounts of noise but does not appear to converge on any single mode of interaction for higher levels of noise. Thus, except for the neutral environment at higher levels of noise, one generally sees the boundedly rational actors of the evolutionary models moving from heterogeneous to common modes of interaction.

This analysis provides one explanation for how norms might arise out of systems of heterogeneous behavior: egoistic but boundedly rational individuals adopt the behavior of those around them that appears to benefit them the most. Over time, the population converges on certain modes of behavior. Were one to see all evolutionary models converge on an institution at every noise type and level, one would suspect that the convergence arose by construction—that is, simply because of the assumptions of the model. That the neutral noise environment shows signs of nonconvergence during high levels of noise actually bolsters the argument that the type and level of noise can significantly affect which institutions are formed.

As a final note, the various homogeneous system analyses indicate that under certain conditions, it may be quite difficult to distinguish empirically the extent of institutionalization in a population and the underlying strategies being employed. For example, one saw in the noiseless and positive noise environments that not only are homogeneous populations of $N\bar{E}$ strategies observationally equivalent, but they are also observationally equivalent to any heterogeneous population composed solely of $N\bar{E}$ strategies. Therefore, to determine the degree of institutionalization and to distinguish the possibly

itself. However, the differences are very small—most likely attributable to stochastic error. Additionally, it should be noted that in the positive noise environment, any $N\bar{E}$ strategy is just as good an institution as CTFT.

33. However, as previously indicated, if one allows for even a very small proportion of the population to be exploitative after convergence (e.g., as in the ecology model in which strategies never really die out) or if one allows for mutations in strategies or for invasions, then certain characteristics (e.g., provocability) not required for the maintenance of cooperation within the

very different strategies of a given institution, one would somehow need to perturb the system in a way that would allow one to separate strategies. Neutral noise appears to do this naturally for the strategies examined here: the negative noise component separates the $N\bar{E}$ strategies, which are observationally equivalent under noiseless and positively biased conditions; the positive noise component separates the nasty strategies, which are observationally equivalent in the negative noise environment.[34] Having noted this, one should not forget that noise can also degrade one's ability to distinguish strategies, especially as the level of noise increases. As Bendor (1993, 731-2) showed, neutral noise causes TFT and STFT to converge to the same limiting distribution of outcomes. Thus the ability of noise to separate strategies is tempered by its ability to obscure.

CONCLUDING REMARKS

In the context of an international system—in which there is no common government to enforce agreements made by nations, in which nations have both conflicting and complementary interests, and in which mutually beneficial outcomes do not follow automatically but neither are rendered unattainable—the RPD is a simple but powerful tool for analyzing the evolution of cooperation. In this article, I have examined how international cooperation may emerge and best be maintained in a world characterized by an RPD, but with symmetric or asymmetric noise.

The major finding is that the noise distribution greatly affects the payoffs of strategies in homogeneous systems and heterogeneous systems and that it affects the evolution of populations employing those strategies. Interestingly, for a policy maker in the boundedly rational world modeled here, CTFT identifies itself as a robust heuristic across all noise environments, performing well in the heterogeneous systems and in the evolutionary models. Additionally, CTFT is also a robust institution across all noise environments, and actors will almost always fare better in a CTFT institution than by employing any other strategy here in a heterogeneous system. Thus there are incentives for the formation of institutions, and in fact, one sees the boundedly rational members of the evolutionary models moving from heterogeneous modes of bilateral interaction to cooperative norms of behavior.

This study has not addressed a number of other topics that may provide interesting avenues of future research. First, the interaction modeled here has

34. In fact, the symmetry of the noise is not important here—just that both noise components are represented, allowing them to differentiate strategies that are not observationally equivalent given the appropriately directed stimulus. Moreover, only a small amount of noise is required for separation to occur.

been purely bilateral. The effects of noise on multilateral punishment strategies and the benefits of multilateral information sharing would provide insight that is perhaps more directly applicable to many situations in trade and security today. Second, a refinement of the analyses would be to model polychotomous or continuous cooperation-defection levels. Third, although CTFT was found to be a robust heuristic in heterogeneous systems and a robust institution compared with other homogeneous systems, this finding may be an artifact of the sample of strategies examined here. Thus, before recommending CTFT as a policy option, one should examine the extent to which it remains evolutionarily stable against small invasions of strategies.[35] Fourth, I noted that certain characteristics of strategies—as defined in the second section—are not consistent across noise distributions. For example, certain strategies that are nonexploiting under noiseless and positive noise conditions become exploitating under negative and neutral noise. Thus a typology is called for that either is consistent across all noise distributions or recognizes that the strategy characteristics are context specific.

Finally, throughout this study, the underlying assumption has been that domestic politics do not affect the prisoner's dilemma payoff structure at the international level. Given these assumptions, CTFT was found to be very robust in heterogeneous systems and as an institution, and it was hypothesized that a generous CTFT may be even more robust given the sample of strategies examined here. However, one might well ask how likely a nation is to implement a strategy whereby it must (1) accept an unprovoked defection from another nation or (2) make some act of contrition for an accidental defection on its own part. If one assumes that a nation's leader is dependent on some domestic interest group for the maintenance of office, then the extent to which that group views retaliation favorably and contrition unfavorably will affect whether a generous and contrite strategy will be implemented by the policy maker. Additionally, the domestic audience may assign different costs to each issue area. An institution like the Law Merchant, in which the act of contrition is a fee, may be more domestically acceptable than one that hurts national feelings of pride or perceptions of security. Thus a nation composed of very proud (or patriotic) citizens, who believe in a firm and quick response to any defection against them, will not favor a generous and contrite strategy. In this case, the irony is that the strategy that is most robust under different noise conditions and against different strategies—and that is most likely to foster cooperation—may be the least likely to be implemented.

35. In fact, Boyd (1989) showed that CTFT is evolutionarily stable over a range of neutral noise given a sufficiently high shadow of the future. However, I am not aware of any similar studies being conducted for the asymmetric noise environments.

APPENDIX
Expected Payoff for Mixed Strategies in a Noisy Environment

Consider the game in Table 1A between two players, each of whom may make one of two possible moves (C and D), resulting in a payoff W, X, Y, or Z, where player 1's payoffs are denoted by a subscript 1 and player 2's payoffs are denoted by a subscript 2. Assuming the players use mixed strategies, let c_1 be the probability that player 1 plays C_1, $(1 - c_1)$ be the probability that player 1 plays D_1, c_2 be the probability that player 2 plays C_2, and $(1 - c_2)$ be the probability that player 2 plays D_2. In addition, assume there is noise in the system, such that if player 1 chooses D_1, there is probability p_1 that the move will either mistakenly result in or be perceived by player 2 as C_1. Similarly, let q_1 be the probability that an intended move of C_1 results in D_1, p_2 be the probability that an intended move of D_2 results in C_2, and q_2 be the probability that an intended move of C_2 results in D_2.

Given this, the expected payoffs for each player, EV_i, are

TABLE 1A
Payoff Matrix for 2 × 2 Game

		Player 2	
		C_2	D_2
Player 1	C_1	W_1, W_2	X_1, X_2
	D_1	Y_1, Y_2	Z_1, Z_2

$$EV_1 = W_1\Pr(C_1,C_2) + X_1\Pr(C_1,D_2) + Y_1\Pr(D_1,C_2) + Z_1\Pr(D_1,D_2) \quad (1)$$

and

$$EV_2 = W_2\Pr(C_1,C_2) + X_2\Pr(C_1,D_2) + Y_2\Pr(D_1,C_2) + Z_2\Pr(D_1,D_2), \quad (2)$$

where $\Pr(i,j)$ is the probability of outcome i,j. Assuming each player's move is independent of the other's, the $\Pr(i,j)$ are defined as follows:

$$\begin{aligned}\Pr(C_1,C_2) = {} & (1 - c_1)p_1(1 - c_2)p_2 + (1 - c_1)p_1c_2(1 - q_2) + \\ & c_1(1 - q_1)(1 - c_2)p_2 + c_1(1 - q_1)c_2(1 - q_2) = \\ & (-c_1 - p_1 + c_1p_1 + c_1q_1)(-c_2 - p_2 + c_2p_2 + c_2q_2),\end{aligned} \quad (3)$$

(continued)

APPENDIX Continued

$$\begin{aligned}\Pr(C_1,D_2) &= (1-c_1)p_1(1-c_2)(1-p_2) + (1-c_1)p_1c_2q_2 \\ &\quad + c_1(1-q_1)(1-c_2)(1-p_2) + c_1(1-q_1)c_2q_2 \\ &= -(-c_1 - p_1 + c_1p_1 + c_1q_1)(1 - c_2 - p_2 + c_2p_2 + c_2q_2),\end{aligned} \tag{4}$$

$$\begin{aligned}\Pr(D_1,C_2) &= (1-c_1)(1-p_1)(1-c_2)p_2 \\ &\quad + (1-c_1)(1-p_1)c_2(1-q_2) + c_1q_1(1-c_2)p_2 + c_1q_1c_2(1-q_2) \\ &= -(-c_2 - p_2 + c_2p_2 + c_2q_2)(1 - c_1 - p_1 + c_1p_1 + c_1q_1),\end{aligned} \tag{5}$$

$$\begin{aligned}\Pr(D_1,D_2) &= (1-c_1)(1-p_1)(1-c_2)(1-p_2) \\ &\quad + (1-c_1)(1-p_1)c_2q_2 + c_1q_1(1-c_2)(1-p_2) + c_1q_1c_2q_2 \\ &= (1 - c_1 - p_1 + c_1p_1 + c_1q_1)(1 - c_2 - p_2 + c_2p_2 + c_2q_2).\end{aligned} \tag{6}$$

It can be easily shown that $\Pr(C_1,C_2) + \Pr(C_1,D_2) + \Pr(D_1,C_2) + \Pr(D_1,D_2) = 1$.

The significance of this model is that it allows not just for asymmetric payoffs between players but also for asymmetric noise. Furthermore, the noise may be not only asymmetric for each player (i.e., $p_1 \neq q_1$ and $p_2 \neq q_2$) but also asymmetric between players (i.e., $p_1 \neq p_2$ and $q_1 \neq q_2$).

REFERENCES

Axelrod, Robert. 1984. *The evolution of cooperation.* New York: Basic Books.

Axelrod, Robert, and Robert Keohane. 1986. Achieving cooperation under anarchy: Strategies and institutions. In *Cooperation under anarchy,* edited by K. Oye, 226-54. Princeton, NJ: Princeton University Press.

Bendor, Jonathan. 1987. In good times and bad: Reciprocity in an uncertain world. *American Journal of Political Science* 31:531-8.

———. 1993. Uncertainty and the evolution of cooperation. *Journal of Conflict Resolution* 37:709-34.

Bendor, Jonathan, R. Kramer, and S. Stout. 1991. When in doubt . . . : Cooperation in a noisy prisoner's dilemma. *Journal of Conflict Resolution* 35:691-719.

Boyd, R. 1989. Mistakes allow evolutionary stability in the repeated prisoner's dilemma game. *Journal of Theoretical Biology* 136:47-56.

Conybeare, John. 1986. Trade wars: A comparative study of Anglo-Hanse, Franco-Italian, and Hawley-Smoot conflicts. In *Cooperation under anarchy,* edited by K. Oye, 147-72. Princeton, NJ: Princeton University Press.

Downs, George W., and David M. Rocke. 1990. *Tacit bargaining, arms races, and arms control.* Ann Arbor: University of Michigan Press.

Downs, George, D. Rocke, and R. Siverson. 1986. Arms races and cooperation. In *Cooperation under anarchy,* edited by K. Oye, 118-46. Princeton, NJ: Princeton University Press.

Goldstein, Joshua S. 1991. Reciprocity in superpower relations: An empirical analysis. *International Studies Quarterly* 35:195-209.

Huberman, Bernardo A., and Natalie S. Glance. 1993. Evolutionary games and computer simulations. *Proceedings of the National Academy of Science* 90:7716-8.

Keohane, Robert. 1984. *After hegemony*. Princeton, NJ: Princeton University Press.

King, Gary. 1989. *Unifying political methodology*. Cambridge, UK: Cambridge University Press.

Mailath, George J. 1992. Introduction: Symposium on evolutionary game theory. *Journal of Economic Theory* 57:259-77.

Milgrom, P., D. North, and B. Weingast. 1990. The role of institutions in the revival of trade. *Economics and Politics* 2 (March): 1-23.

Molander, P. 1985. The optimal level of generosity in a selfish, uncertain world. *Journal of Conflict Resolution* 29:611-8.

Morrow, James D. 1994. *Game theory for political scientists*. Princeton, NJ: Princeton University Press.

Oye, Kenneth. 1986. Explaining cooperation under anarchy. In *Cooperation under anarchy*, edited by K. Oye, 1-24. Princeton, NJ: Princeton University Press.

Pahre, Robert. 1994. Multilateral cooperation in an iterated prisoner's dilemma. *Journal of Conflict Resolution* 38:326-52.

Schwartz, Edward P. 1993. Trade in a world most uncertain: Error and incentives in the system of the Medieval Law Merchant. Harvard University. Mimeographed.

Sugden, R. 1986. *The economics of rights, co-operation, and welfare*. Oxford: Basil Blackwell.

Van Wyk, Koos, and Sarah Radloff. 1993. Symmetry and reciprocity in South Africa's foreign policy. *Journal of Conflict Resolution* 37:381-96.

Wu, Jianzhong, and Robert Axelrod. 1995. How to cope with noise in the iterated prisoner's dilemma. *Journal of Conflict Resolution* 39:183-7.