Forthcoming in Computational and Mathematical Organization Theory, 1999

Private Games Are Too Dangerous

© Ronald S. Burt
Graduate School of Business, University of Chicago, Chicago, IL 60637,
and INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France
email: ron.burt@gsb.uchicago.edu; http://gsbwww.uchicago.edu/fac/ronald.burt

Given the difficulty of observing interpersonal relations as they develop within an organization, I use iterated prisoner's dilemma games to simulate their development. The goal is to understand how trust could develop as a function of private games, that is, as a function of interaction sequences between two people independent of their relationships with other people. My baseline is Axelrod’s results with TIT for TAT showing that cooperation can emerge as the dominant form of interaction even in a society of selfish individuals without central authority. I replicate Axelrod’s results, then show that the results only occur in a rare social context — maximum density networks. Where people form less dense networks by withdrawing from unproductive relationships, as is typical in organizations, the competitive advantage shifts from TIT for TAT to abusive strategies. A devious PUSHY strategy wins in moderate to high density networks. A blatantly HOSTILE strategy wins in less dense networks. Abusive players do well in sparse networks because their abuse is lucrative in the initial exchanges of a relationship — before the other person knows to withdraw. Wise players avoiding the abusive players leaves the abusive players free to concentrate on naive players (con men thrive in big cities). The implication is that what keeps abusive players at bay are friends and acquaintances warning managers away from people known to exploit their colleagues. I reinforce the point with illustrative survey data to conclude that private games are not only too dangerous, but also too rare and too slow to be the foundation for trust within organizations. The results are an evidential call for the sociological intuition that trust and distrust cannot be understood independent of the network context in which they are produced.

Keywords: trust, organization, social network, density, prisoner’s dilemma,

1. Introduction

There are two ways to think about the probability of trust within a relationship. The simpler is to think in terms of private games — that is, in terms of interaction sequences between two people independent of their relationships with other people. Trust emerges between the two people as a function of their history and future with one another. Two prominent examples of this thinking in sociology are Homans' (1961) analysis of social behavior, and Blau's (1964) analysis of social exchange (see Blau, 1994:152-172, for Blau's contemporary view, esp. pp. 156-158, explaining his continued focus on dyadic exchange). Blau (1964:112-113) argues that trust develops because social exchange involves unspecified obligations for which no binding contract can be written. When you exchange sensitive information with someone, for example, trust is implicit in the risk you now face that the other person might leak the information. Putting aside Blau's moral obligation aspect of exchange to focus on parameters of cost-benefit calculation (cf. Ekeh, 1974:175), Coleman (1990: Chap. 5) captured trust more concretely for his systems of two-party exchange: Trust is committing to an exchange before you know how the other person will behave. You anticipate cooperation from the other person, but you commit to the exchange before you know how the other person will behave. This is trust, pure and simple. Anticipated cooperation is a narrow segment in the spectrum of concepts spanned by richer images such as Barber's (1983) distinctions between trust as moral order, competence, and obligation. However, anticipated cooperation is much of the trust essential to people in organizations. The issue isn't moral. It is flexible cooperation (e.g., Macauley, 1963; Uzzi, 1996).

Viewed as anticipated cooperation, trust is twice created by repeated interaction; from the past and from the future. From the past, repeated experience with a person is improved knowledge of the person. Cooperation in today's game is a signal of future cooperation. Across repeated games with cooperative outcomes, you build confidence in the other person's tendency to cooperate. At minimum, the cumulative process can be cast as a statistical decision problem in which you become more certain of the other person across repeated samples of his or her behavior. The repetition of cooperative exchange promotes trust. More generally, the cumulative process involves escalation. From tentative initial exchanges, you move to familiarity, and from there to more significant exchanges. The gradual expansion of exchanges promotes the trust necessary for them. Whatever the cumulative process, past cooperation is a basis for future cooperation (cf. Zucker, 1986, on process-based trust; Larson, 1992; on the importance of the long term for trust between firms; Lawler and Yoon, 1996, for laboratory evidence; Stinchcombe, 1990:164-165, on the information advantages of current suppliers for building trust; Gulati 1995; Gulati and Garguilo, Forthcoming, for empirical evidence). Further, the history of cooperation is an investment that would be lost if either party behaved so as to erode the relationship — another factor making it easier for each party to trust the other to cooperate. Blau (1968:454) summarized the process as follows: ". . . social exchange relations evolve in a slow process, starting with minor transactions in which little trust is required because little risk is involved and in which both partners can prove their trustworthiness, enabling them to expand their relation and engage in major transactions. Thus, the process of social exchange leads to the trust required for it in a self-governing fashion." Where sociologists explain trust emerging from past exchanges (e.g., Coleman, 1990; Granovetter, 1992), economists look to the incentives of future exchanges (e.g., Kreps, 1990; Gibbons, 1992:88ff). The expectation that violations of trust will be punished in the future leads players to cooperate even if defection would be more profitable in a single play of the game. The information contained in past experience and the potential for future interactions are inextricably linked. A player's willingness to forego short-term gains is based on the expectation that his or her current behavior will be used to predict his or her behavior in the future.

The alternative is to think more broadly in terms of public games — interaction sequences in which two people display their behavior to one another, but are mindful of their reputation with friends and acquaintances observing the game. Relationships play out in a social context of friends, acquaintances, and enemies who are third parties to the relationship and transform what was private into public. Three network arguments describe how third parties to a relationship can be expected to affect trust within the relationship. There is a balance argument in which people are presumed to have a preference for consistency between relationships adjacent in a network (e.g., Heider, 1958; Davis, 1970), a reputation argument in which strong indirect connections through third parties are a medium through which social norms of cooperation can be enforced (e.g., Granovetter, 1985, 1992; Coleman, 1988, 1990; Kreps, 1990, Greif, 1993; Putnam, 1993), and a gossip argument in which etiquette-biased information amplifies predispositions toward trust or distrust (e.g., Burt and Knez, 1995; Burt, 1999a, 1999b; Sobel, 1999). The gossip argument reproduces the first two in predicting that third parties increase the probability of trust within strong relationships, and contradicts them in predicting that third parties also increase the probability of distrust within weak relationships. Evidence consistent with the gossip argument is beginning to appear as scholars test for network effects on negative relationships (e.g., Burt and Knez 1995, on distrust associated with third parties within organizations; Labianca, Brass, and Gray 1998, on perceptions of conflict associated with third parties; Gulati and Westphal, Forthcoming, on third-party interlocks amplifying positive and negative predispositions to alliance between organizations; Burt, 1999c, on three generic research findings supporting the gossip argument).

Whatever the cross-sectional and historical evidence of association between trust and third parties, the evidence needed to establish third-party effects in an authoritative way is a study of interpersonal networks over time to see how variation in the structure of third parties around a relationship affects its development. The problem is that such observation typically involves (a) substantial cost within an organization of even moderate size, and (b) the distortions associated with intrusive observers. The problem must be more serious toward the top of the organization where much of trust is an implicit understanding and so not easily calibrated over time. Systematic evidence from the field has been limited to cross-sectional data showing that trust and distrust are more likely in relations embedded in third-party ties (Gulati and Westphal, Forthcoming; Gulati and Garguilo, Forthcoming, are rare exceptions that use panels of archival data on relations between organizations). How relations develop as a function of third parties remains unknown, so causal inferences are accordingly suspect (Burt and Knez, 1995: 284-285).

Where observation over time is difficult, simulations can be a productive complement to cross-sectional data. I use iterated prisoner's dilemma games between eight game strategies to simulate the development of relationships among kinds of managers within an organization. My baseline is Axelrod’s (1980a, 1980b, 1984) widely-known use of such simulations to argue that cooperation can emerge as the dominant form of interaction even in a society of selfish individuals without central authority.

Two points are established with the simulations. I first replicate Axelrod's tournament results on cooperation. TIT for TAT is the most successful kind of player. The TIT for TAT strategy is to begin with cooperation, then reciprocate whatever move your partner made in the prior game.

Second, I show that TIT for TAT's success is limited to dense networks. The simulated network of Axelrod's tournament was dense in two ways rare in organizations: repeated games with every other person, and no exit from a bad relationship. This is a network characteristic of families more than organizations. I make the simulation more relevant to organizations by shifting to a sparse network in which relations rarely have a chance to develop beyond initial exchanges and players can withdraw from a bad relationship. The result is a change in the quality of relationships and the relative success of player strategies. Relations in the sparse network are on average more rewarding and more pleasant because players can withdraw from bad relationships. The irony is that the higher quality relations serve most to benefit nasty players, players who exploit cooperative partners. In anything less than maximum-density conditions, the advantage shifts from TIT for TAT players to the cautiously exploitative PUSHY players. Just before the network disintegrates to a condition more sparse than dense, the advantage quickly shifts to the more aggressively exploitative HOSTILE players, who dominate all populations more sparse than dense. Nasty players do well in sparse networks because they do well in the initial games of a relationship, and few relations in a sparse network develop beyond the initial games.

I use the term density as it is usually used in network analysis, but let me be clear about it at the outset because density is a key variable in what follows. A dense network is one in which each person has a strong (positive or negative) relationship with every other person. At the other extreme, a sparse network is one in which relations are often weak or absent. When I allow in the simulation for people to play only few games with specific others, and withdraw from unproductive relationships, I transform what were dense networks into sparse networks. There are other ways in which a sparse network could occur, but in this paper I only study the consequences of infrequent interaction and withdrawal. It is useful to keep separate the condition of a network (e.g., dense versus sparse) from the many processes by which the network could have come to its current condition.

2. Private Games in Dense Networks

To ensure that the substantive content of the simulations is clear, let me take a moment to describe the prisoner's dilemma game. Hardin (1990:364) provides a delightful glimpse into the game's origin: "This game was discovered or invented — it is not clear which is the more apt term here — by Merrill Flood and Melvin Dresher, two early game theorists who were trying to test bargaining theories with experimental games. Oddly, two of the games with which Flood experimented before the prisoner's dilemma involved simple exchanges — of old cars for money. He seems not to have seen that his prisoner's dilemma game was a simplification and generalization of such exchanges. Unfortunately, this association got lost in the later naming of the game by A. W. Tucker, who saw in the game a perverse analog of American criminal justice, in which prosecutors extract confessions on the promise of reduced sentences." The analog is illustrated in Figure 1 with an anecdote often used, in this or similar form, to introduce the game.[1]

—— Figure 1 About Here ——

You are a prisoner deciding whether to confess. The best outcome for you (ego) and your partner (alter) occurs if you cooperate with one another and don't confess, but cooperation is not the optimum strategy for you personally. If your partner confesses, you are better off confessing (five years in prison versus ten if you don't confess). If your partner doesn't confess, you are still better off confessing (no years in prison versus five if you don't confess). Regardless of your partner's behavior, you personally are better off confessing. Trust is the essential tension of the game. You have to decide to cooperate before you know how your partner will reciprocate. Can you trust your partner to keep his mouth shut? In the absence of trust, the optimum choice is to defect.[2]

---

[1]Hardin's note here is a citation to Flood (1958), and a citation to an earlier Rand Corporation research memorandum RM-789-1 (20 June 1952), which is, to his knowledge the first instance of the appearance of the prisoner's dilemma. Hardin also reports, from private communication with Merrill Flood, that Flood gives Tucker credit for giving the game its current name (which must have been an informal communication because I have yet to see a Tucker publication credited as the source).

[2]There is a voluminous literature on behavior in prisoner's dilemma games as a function of players (e.g., sex, occupation, nationality, personality), the magnitude of the temptation to defect, and the relation between the player (e.g., previous experience, opportunities to communicate). Rapoport, Guyer, and Gordon (1976) provide a useful review of results, Rapoport (1987) is a succinct overview of issues, Roth (1995:26-30) provides behavioral economics review, Oskamp (1971) focuses on results with experiment subjects playing against

The optimum strategy is less obvious across repeated games. Today's game has implications for tomorrow’s game. With respect to the prisoner anecdote, there will be future games in prison or on the street with the partner betrayed today (not to mention the future disdain of peers — no one likes a snitch).

Axelrod took a behavioral approach to determining the optimum strategy. Whatever the initial value of the approach, it is productive for the ease with which it extends to more complex network settings. Experts were invited to submit strategies for a computer tournament to determine the optimum strategy (Axelrod, 1980a:7): "In a computer tournament, each entrant writes a program which embodies a decision rule to select the cooperative or noncooperative choice on each move. The program has available to it the history of the game so far, and may use this history in making a choice. Because the participants are recruited primarily from those who have written on game theory and especially the prisoner's dilemma, the entrants are assured that their decision rule will be facing rules of other experts. Such recruitment also guarantees that the state of the art is represented in the tournament." There are obvious questions about who would enter such a tournament, and there were few submissions to the initial tournament.[3] However, the results of the first tournament were sufficiently intriguing (Axelrod, 1980a) to draw 63 contestants for a second tournament (Axelrod, 1980b).

2.1 NETWORK MODEL

The 63 contestants were not completely different. Qualities of some strategies appear in other strategies. I wrote a computer program to simulate the behavior of eight strategies, or kinds of players, that span significant distinctions in Axelrod's tournament. Game outcomes

---

programmed players, and Lomborg (1996:278-283) provides a contemporary review of the sociology of iterated prisoner's dilemma games.

[3] There are also questions about the validity of simulated play. In his behavioral economics review, Roth (1995:29) calls attention to inconsistency between the success of cooperation in Axelrod's simulations and the negative results in some experiments with human subjects, and "suspects" that the inconsistencies reflect the learning that occurs with human subjects (see Macy, 1991:818-820, for negative results in simulations that allow for ego learning across repeated play; cf., Roth and Erev, 1995). I am not concerned with the validity of simulated play for two reasons: First, sociologists are well beyond the issue of simulation versus experiments with human subjects because published research goes straight to the trust phenomenon in vivo with field data on senior managers (as cited in the text). Second, I study the simulations as a social fact, not as reality. Axelrod's

are proportional to the rewards in Axelrod's tournament and Figure 1 (10 points for mutual defection, 30 points for mutual cooperation, 50 points to the player who defects against a cooperating partner, and 0 points to the player who cooperates with a partner who defects).

Table 1 is a summary of the simulations. Cell (A,B) is the average outcome for player A in games with player B (cf. Axelrod, 1980a:11, 1980b:387, 1984:194, 197-199). There are eighty, not eight, players in the table. Table 1 is a summary of relations in an eighty-person network where each relation is based on games transacted at random within the network. The table begins with each relation z_{ij} from row i to column j equal to zero. Pick one of the rows i at random to play a game. Pick column j at random to be the other player. The two players make moves according to their respective strategies. The results are added to the appropriate two cells of the table: If i and j both cooperate, add 30 points to relations z_{ij} and z_{ji}. If i and j both defect, add 10 points to relations z_{ij} and z_{ji}. If i cooperates but j defects, then add 50 points to relation z_{ji}, nothing to z_{ij}. Draw a second random pair of players for game two. Make each player's move. Add the results to the appropriate two relations in the table. Draw a third random pair of players, and so on. Continue for 800,000 games. Each relation is the cumulative result of an average of 250 repeated games. To get the average outcome for player i in a game with player j, divide the summed pay-offs in relation z_{ij} by the number of games summed in the relation. Cell (A,B) in Table 1 is the average of the 100 relations between the ten players of row kind A with the ten players of column kind B. The standard deviation of average outcomes across the 100 relations in each cell is reported in parentheses. The number in brackets is the average number of games played to define each of the relations (slightly more than 250 because the 10 self-relations z_{ii} in each diagonal cell are excluded). In network analysis terms, Table 1 is a density table of the simulated relations between the eighty players.

—— Table 1 About Here ——

Note that these results do not describe an N-way game; they describe two-person games in a system of N players. In an N-way game, each player is engaged in a game with the N-1 other players. As N increases, monitoring the other players is more difficult, and

---

results are widely cited as evidence of trust emerging from private games. They need to be put into perspective.

cooperation is less likely because so many players have to cooperate at the same time. This is the central point of Hechter's (1987) argument for the importance of control mechanisms to enforce cooperation. In fact, Knez and Camerer (1994) describe experiments in which adding one person to a two-person coordination game significantly erodes cooperation. The network summarized in Table 1 is limited to two-person games. It is the network implicit in Axelrod's simulations. Players do not monitor all N-1 other players, and there is no indirect monitoring through third parties. The only games known to a player are the games in which he or she is a participant.[4]

2.2 NONRESPONSIVE PLAYERS

The first three kinds of players in Table 1 are nonresponsive in the sense that they give no attention to the behavior of their partner. HAPPY cooperates. HOSTILE defects. ERRATIC is equally likely to cooperate or defect. Always cooperating, HAPPY players always earn 30 points in games with one another. Always defecting, HOSTILE players always earn 10 points from one another. HOSTILE does well against HAPPY because HAPPY continues to cooperate despite HOSTILE's defections. The outcome of these games are 50 points to HOSTILE, 0 to HAPPY. HOSTILE also does well against random play. When ERRATIC cooperates with HOSTILE he earns nothing and when he defects he earns 10 points. Half his moves are cooperation, so his average profit from a game with HOSTILE is only 5 points. On the other side, HOSTILE earns 50 points when ERRATIC cooperates and 10 points when he defects, for an average profit of 30 points

---

I replicate his results, then show their implausibility in the sparse networks that characterize organizations.

[4] Marwell and Oliver (1993:Chap. 5) illustrate the logical next step in their simulations of exchange network solutions to the collective goods problem. Their strategy is to have every player be an entrepreneur monitoring the feasibility of becoming an "organizer." Each player only knows the other players with whom he or she has direct contact. Contact is established at random between players (constrained to a lognormal distribution of network sizes within the system). Players calculate whether they could write a contract providing a collective good paid with contributions from contacts such that the cost of each contact's contribution is less than the benefit each would receive from the collective good. Marwell and Oliver compare systems restricting games to different social structures of pairing players to see how the structures affect the probability of the collective good. In addition to significant results on covariate effects of interest and resource heterogeneity, their central result is that centralized systems are more likely to provide themselves with the collective good because there is more likely to be a player with a large network who can bring together the most rewarding set of contributors. The players in Table 1 are less sophisticated than Marwell and Oliver's in the sense that they do not transfer experience across relationships. A next step for the simulations in this paper

HOSTILE is the most successful of the three nonresponsive players. Axelrod presents two kinds of evidence for such conclusions; expected outcomes and ecological survival. In terms of expected outcomes, HOSTILE is most successful. These games equally distributed across the three nonresponsive players yield an average profit to HOSTILE of 29.9 (divide by three the row sum from Table 1 of 50.0, plus 10.0, plus 29.6). The other two nonresponsive players earn less — 14.9 for HAPPY, and 22.5 for ERRATIC.

2.3 NICE PLAYERS

The next three players in Table 1 are what Axelrod (1980b:389) termed "nice" players. Nice players respond to their partner, but are never the first to defect. TIT for TAT begins with cooperation, then reciprocates the other person's last move — cooperates in response to cooperation, defects in response to defection. This was Anatol Rapoport's submission, and won the tournament (Axelrod, 1980b:389-390). OTHER CHEEK begins with cooperation, then reciprocates, but turns the other cheek in response to a defection, allowing the other person two defections before retaliating. This tit-for-two-tats strategy was John Maynard Smith's submission to the Axelrod tournament. RIGHTEOUS cooperates until the other person is discovered to be dishonest, then never again cooperates with that person. This strategy consistent with the legal philosophy of the Puritan Massachusetts Bay Colony was James W. Friedman's submission to the Axelrod tournament (cf. Gibbons, 1992:91ff, on the "trigger" strategy).

RIGHTEOUS players are the most successful among these six kinds of players, but no one kind of player is dominant. The players in descending order of expected profit are RIGHTEOUS (26.6 mean outcome), TIT for TAT (25.4), OTHER CHEEK (24.7), HAPPY (22.5), ERRATIC (21.1), and HOSTILE (20.1). HOSTILE is now the least successful player because his constant defection denies him the cooperation benefits shared by the nice players. Further, nice players reciprocate HOSTILE's defection so he can't exploit them like he exploits HAPPY. The primary difference among the nice players is their relative success against random play. RIGHTEOUS does best against ERRATIC. As soon as RIGHTEOUS

---

would be to allow players to learn how to behave in a relationship from their experience in other relations and discovers that ERRATIC defects against a cooperative partner, it never cooperates again. Thereafter, RIGHTEOUS earns at least 10 points per game with ERRATIC.

2.4 NASTY PLAYERS

Nasty players defect against a cooperating partner. HOSTILE players are the pure form; they defect in all games with all players. ERRATIC players are a weak form in the sense that they do sometimes defect against a cooperating partner. The last two players in Table 1 are strategic forms. They respond to their partner's behavior, but they look for opportunities to get away with defecting against a cooperating partner.

SNEAKY looks for a little extra from a tit-for-tat strategy. The strategy is to play the TIT for TAT strategy, but occasionally defect against a cooperating partner. If the other player defects, SNEAKY defects. If the other cooperates, SNEAKY cooperates at random 90% of the time. In other words, SNEAKY's deceit can go as high as defecting in 10% of his moves against a cooperating partner. This was Johann Joss's submission to the Axelrod tournament (JOSS in Axelrod, 1980a:14, 23; 1984).

PUSHY tries to exploit, but backs away if threatened. PUSHY defects on the first move to test the other player. If the other player cooperates, then PUSHY cooperates on the second move and defects on alternate subsequent moves (4, 6, 8, etc.). If the other defects, PUSHY apologizes by cooperating on the next move and playing TIT for TAT for the rest of the game. In other words, PUSHY's deceit can go as high as 50% against a cooperating partner. This was David Gladstein's submission to the Axelrod tournament (TESTER in Axelrod, 1980b:391; 1984).

The nasty players reveal important differences between nice players. SNEAKY reveals the importance of forgiveness (Axelrod, 1980a:14-18). RIGHTEOUS never forgives, TIT for TAT forgives after retaliating with a punishment, and OTHER CHEEK is the most forgiving in the sense that he turns the other cheek and waits for a second defection before retaliating. The result is that SNEAKY drags RIGHTEOUS and TIT for TAT into an unproductive equilibrium state of mutual retaliation while OTHER CHEEK does well. The

---

from successful players elsewhere in the network (see footnote 9).

average outcome of games between SNEAKY and TIT for TAT are 11.7 points to TIT for TAT, 11.9 points to SNEAKY. The expected outcome varies as much as it does because it depends on the timing of SNEAKY's initial defection. When SNEAKY does defect, TIT for TAT retaliates, then SNEAKY retaliates in return, and the two players are thereafter in an equilibrium state of mutual retaliation. In contrast, when SNEAKY defects against OTHER CHEEK, there is no immediate retaliation and SNEAKY is unlikely to defect twice in a row, so OTHER CHEEK usually forgives the one defection with SNEAKY's subsequent cooperation. Two results: (a) SNEAKY players get more from their games with OTHER CHEEK because defections earn them 50 points and give OTHER CHEEK nothing (29.4 mean outcome for SNEAKY versus 25.1 mean outcome for OTHER CHEEK). (b) OTHER CHEEK makes money while the other nice players get runner-up awards (25.1 mean outcome for OTHER CHEEK versus 11.7 and 10.9 mean outcomes for TIT for TAT and RIGHTEOUS). It pays to forgive.

It also pays to punish. PUSHY reveals the importance of provocability (Axelrod, 1980b:391-393). RIGHTEOUS and TIT for TAT retaliate immediately against a defection, OTHER CHEEK waits. The result is that PUSHY drags RIGHTEOUS into an unproductive equilibrium of mutual retaliation, exploits OTHER CHEEK, and does well with TIT for TAT. On the first point, RIGHTEOUS never trusts a PUSHY player after PUSHY's opening defection, so they both end up with minimum profit (10.3 points to RIGHTEOUS in Table 1, 10.1 to PUSHY, slightly higher for RIGHTEOUS because PUSHY apologizes with two cooperative moves before going into permanent defection). Second, the largest gains are against OTHER CHEEK. PUSHY never defects twice in a row, so OTHER CHEEK never punishes his defections. Games alternate between mutual cooperation (30 points to both players) and exploitation (50 points to PUSHY, zero to OTHER CHEEK), giving PUSHY the advantage on average (40.0 points to PUSHY and 15.0 to OTHER CHEEK in Table 1). Third, TIT for TAT does well against PUSHY because he punishes PUSHY immediately for his opening defection, after which PUSHY cooperates with TIT for TAT. The average outcome of their games is 30.0 points to both PUSHY and TIT for TAT. This is also why PUSHY does so well against players of his own kind. After earning only 10 points with their initial defections against one another, PUSHY players cooperate and earn 30 points per game (29.8 mean across two 10-point and 249 30-point games).

Having SNEAKY and PUSHY players in the population affects the relative success of the nice players. TIT for TAT is the most successful player, and ultimately dominant, but not by much. Average outcomes are reported at the top of Table 1. These are the average relation for each row kind of player with all players (an average of 790 relations from each row kind of player). The four most successful players are TIT for TAT (24.3 mean outcome), PUSHY (24.3), OTHER CHEEK (23.6), and now the least successful of the nice players, RIGHTEOUS (22.6 mean outcome).

2.5 ECOLOGICAL RESULTS

The simulation also replicates Axelrod's ecological results. Allow two reasonable baseline assumptions: players breed in proportion to the expected outcome of their games with one another (i.e., successful players are more numerous in the next generation), and players die in proportion to their numbers (death affects all players equally).

Then, the results in Table 1 can be used to project the mix of players in the population across successive generations. Where P_1 and P_2 are the proportions of HAPPY and HOSTILE players in a population, then P_1P_2 is the proportion of random games between HOSTILE and HAPPY, and 50.0 is the expected outcome for the HOSTILE player. Therefore, since players breed in proportion to their expected outcomes, expected player i births, B_i, are:

$$B_i \text{ in next generation} = \Sigma_j \, P_i P_j z_{ij}, \tag{1}$$

where z_{ij} is the average outcome in Table 1 for player i in a game with player j (e.g., z_{21} is 50.0, the average outcome for a HOSTILE player against HAPPY). Let players die in proportion to their numbers, say a 25% death rate to speed the ecological process to equilibrium. Player i proportion in the next generation is the current number minus deaths, $(1-.25)N_i$, plus births, B_i, quantity divided by the new population total:

$$P_i \text{ in next generation} = (.75N_i + B_i) / \Sigma_j (.75N_j + B_j). \quad (2)$$

For example, consider a population of 10 HAPPY, 10 HOSTILE, and 10 ERRATIC players. The first generation will produce five new HAPPY players, (30.0+14.8)/9, ten new HOSTILE players, and seven and a half new ERRATIC players. A death rate of 25% will claim two and half players of each kind in the first generation. Therefore, from proportions of .33, .33 and .33 in the first generation, the population shifts to proportions of .28, .39, and .33 in the second generation. The second generation contains more HOSTILE players and fewer HAPPY players because games between these two kinds of players favor HOSTILE.

—— Figure 2 About Here ——

Figure 2 shows what happens when all eight kinds of players in Table 1 interact across generations. The two most successful strategies, TIT for TAT and PUSHY, quickly replace the other players. There are slightly more PUSHY players at that point because they do better at exploiting OTHER CHEEK and HAPPY players. Once PUSHY's victims are dead, however, all games are among PUSHY and TIT for TAT players. PUSHY players are at a slight disadvantage that ultimately destroys them. TIT for TAT players do slightly better with one another because they always cooperate. The opening move in which PUSHY tests the other player lowers their gains from one another (29.8 between PUSHY players versus 30.0 between TIT for TAT). That slight disadvantage multiplies through time to eliminate the PUSHY players. The population is about equally composed of PUSHY and TIT for TAT players in the 200th generation, then 44% PUSHY by the 500th generation, 36% by the one thousandth generation, and so on until the last PUSHY player dies in generation 74,296. The dominance of TIT for TAT isn't affected by the initial numbers of players, but the number of generations required to reach dominance is affected. The smaller the (non-zero) proportion of TIT for TAT players in the initial population, the more generations that pass before TIT for TAT players replace all others.[5]

---

[5]Initial numbers matter for other mixtures of players (a less circulated result from Axelrod's tournament, 1980b:401-403; 1984:40, 48). If the population begins with equal numbers of only the first six kinds of players in Table 1, it evolves to an equilibrium of multiple players. After ninety generations, the HOSTILE and ERRATIC players are all dead, and the population is stable at 16% HAPPY, 27% TIT for TAT, 23% OTHER CHEEK, and 34% RIGHTEOUS. There is no ecological pressure on the four finalists after the ERRATIC players disappear. All four finalists cooperate with one another and so breed in proportion to their existing numbers. The RIGHTEOUS players become most numerous because they profit most from games with

3. From Voice to Exit

The simulated exchange network summarized in Table 1, and the kindred density tables in analyses such as Axelrod's tournament, are maximum-density networks. Every person plays many repeated games with every other person in the network.

Networks are rarely this dense in organizations. Managers do have contacts with whom they meet frequently, or trust because of previous frequent meetings. But more often than not, life in a complex organization involves work with new contacts and people met only rarely — senior managers, peers in other functions, colleagues in other firms. Play in a sparse network involves relations defined by fewer games, and choice between alternative partners. Instead of many repeated games with every person, players complete introductory games with many people, and limit their numerous repeated games to a few preferred people.

The cooperation implications of a sparse network are not clear. The effect of few games is well known. Defection is the optimum prisoner's dilemma move against a partner you won't meet again. The implications of choice between alternative partners is less obvious. Tullock (1985: 1076) noted on the surprising lack of attention to a withdrawal option in prisoner's dilemma games; "All of the previous articles that I have seen are situations in which the individual does not have a choice of partners. Either he is stuck with his colleague, like the two prisoners in the initial story, or the state of nature is such that he cannot change partners." Tullock (1985:1080-1081) concluded that introducing the threat of withdrawal will generate cooperation:

ERRATIC players. The two key variables are: (a) the relative advantage of nice players in games with ERRATIC players — strong to weak in Table 1 are RIGHTEOUS, then TIT for TAT, then OTHER CHEEK, and (b) the relative numbers of nice players in the initial population. On the first point, more ERRATIC players in the initial population means a higher proportion of RIGHTEOUS players at equilibrium. On the second point, a kind of nice player that begins with a numerical disadvantage might not expand fast enough to overtake the other nice players before all of the available ERRATIC players are consumed. For example, a population composed of twenty OTHER CHEEK players to ten each of the other five kinds evolves to an equilibrium in which OTHER CHEEK is most numerous (21% TIT for TAT, 28% RIGHTEOUS, 37% OTHER CHEEK). A population initially composed of twenty TIT for TAT players to ten each of the others moves to an equilibrium in which TIT for TAT is most numerous (42% TIT for TAT, 26% RIGHTEOUS, 19% OTHER CHEEK). Population composition notwithstanding, I am interested in a composition as in Table 1 that spans the space of Axelrod’s contestants and replicates his results, because my goal is to show why the results do not generalize to the sparse networks typical within organizations.

> "...the main theme of this discussion has been that the prisoner's dilemma, strictly speaking, occurs only in a rather narrow area. Where there are a number of potential players available, the dilemma is proportionately weakened. Indeed, here we have what amounts to a mapping of the usual economic distinction between monopoly on the one side and competition on the other. If there is only one person with whom I can play the game, both he and I very likely will decide not to cooperate. As the number of people playing increases, the prospect that either he or I can get another partner, if we find our current partner objectionable, exerts steadily increasing pressure to always play cooperatively. . . . Where the market is broad and there are many alternatives, you had better cooperate. If you choose the noncooperative solution, you may find you have no one to noncooperate with."

Another view is to think in terms of Hirschman's (1970) contrast between voice and exit. In dense networks, players respond to uncooperative partners with voice. The responsive players in Table 1 are petulant in the sense that they punish uncooperative partners by being uncooperative in return. This is their only communication option since they are obliged to continue the relationship whether it is productive or not. In contrast, players in sparse networks can respond to uncooperative partners with exit. They are not obliged to waste time on uncooperative partners. They can punish by avoiding uncooperative partners. The exit option, however, has the potential to be a private solution to a collective problem. Ego rejecting an uncooperative partner is likely to do better with a different partner, but what does that leave the uncooperative partner free to do to other players? What it does is free nasty players to focus on the most cooperative players, and so earn maximum gains from each round of completed games.

4. Private Games in Sparse Networks

Table 2 summarizes game play in a sparse network. As in Table 1, each cell in Table 2 is the average outcome for row players in games with the column players. The average game outcome for a kind of player is reported at the top of the table.

Table 2 differs from Table 1 in two ways. First, fewer games define relations (16,000 games were played to construct Table 2, versus 800,000 to construct Table 1). The average relation in Table 2 is defined by five games, but almost a fifth of the relations are defined by only one or two games, and 14 pairs of players never meet one another — versus the 253 game average, and 201 game minimum in Table 1. In contrast to the relations in Table 1 defined by a long history of repeated games, Table 2 describes relations based primarily on the initial games in a relationship.[6]

—— Table 2 About Here ——

Second, and this is the more consequential difference from Table 1, players in a sparse network can choose between alternative partners, so they can punish uncooperative partners by rejecting them in future games.[7] Retaliatory defection in Table 1 is replaced by rejection in Table 2. For example, TIT for TAT in Table 1 retaliates by defecting when the other

---

[6]Drawing pairs of players at random is now a stronger assumption than in Table 1. Drawing players at random for Table 1 imposes independence on the timing and number of games. However, differences between relations in their number of games is trivial because all relations are defined by a large number of games. Table 2 is more complex. Pairs of players vary in the extent to which they develop a relationship. It is now a strong assumption to say that the number of games is independent between relations. Available work on network dynamics is of little use here because the work remains largely abstract and methodological, guided by little or no substantive research on network dynamics (e.g., Burt, 1979, Marsden, 1983, Zeggelink, 1994, Stokman and Zeggelink, 1996 for illustrative models and literature review). Turning to common sense, there is reason to believe at minimum that the initial games of adjacent relations are not independent. The factors that result in me first interacting with you, lead me to meet the people with whom you interact. When someone meets you in the hall, for example, they also meet the people you meet in the hall. This implies that the probability of games between persons i and j should increase with the extent to which i and j play games with the same people (instead of being independent of games with other people). Looking for information on a functional form for the contingency, I studied network data on manager relationships in three study populations (Figure 7). A manager cites two people j and k especially close with each other. How does the manager's frequency of meeting person j change as meetings with person k become more frequent? If either person is within the manager's work group, frequency is contingent (11.6 chi-square, 4 d.f., P = .02). Managers meet more often the people that their immediate co-workers meet often. Beyond the immediate workgroup, however, frequency is independent in adjacent relationships, and more cited colleagues are beyond the immediate work group, so frequency in the aggregate is independent in adjacent relationships (5.1 chi-square, 4 d.f., P = .27). Managers are involved in such a diversity of events that the frequency with which they meet one key colleague is independent of the frequency with which they meet other key colleagues close to the first. Therefore, it is not unreasonable to draw players at random for Table 2. A next step would be to draw players in sparse networks as a function of their ties to other players (see footnote 9 on learning, footnote 11 on allowing players to select productive relationships).

[7]By more consequential, I mean that larger differences are created. Here are correlations between the corresponding 64 cells in four density tables:

1.00	.70	.35	.37	dense, no withdrawal (Table 1)
	1.00	.66	.69	sparse, no withdrawal
		1.00	.87	dense, allowing withdrawal
			1.00	sparse, allowing withdrawal (Table 2)

If players were not allowed to reject partners in Table 2, then average outcomes in the table would be correlated .70 with the averages in Table 1. The outcomes would be almost as correlated, .69, with the outcomes in Table 2. If the dense relations in Table 1 were defined by games in which players could reject partners, average outcomes would be correlated .87 with the outcomes reported in Table 2. Limiting the number of games person defects. In Table 2, TIT for TAT rejects the defecting person as a partner the next time he is assigned to a game with the person. TIT for TAT picks someone else to be their partner.[8] TIT for TAT in Table 1 forgives after one retaliatory defection. TIT for TAT in Table 2 forgives after one retaliatory rejection. The other responsive players in Table 2 similarly correspond to players in Table 1. OTHER CHEEK rejects a partner after the other defects in two successive games, then forgives. RIGHTEOUS rejects a partner after the other defects once, and continues to reject them thereafter. SNEAKY and PUSHY use TIT for TAT's retaliation rule.[9]

---

defining relations has an effect (Table 2 no rejection is more correlated than Table 1 with the last two tables), but allowing rejections creates the largest differences from Table 1 (.35 and .37 are the lowest correlations).

[8]In a population of fixed size, exit doesn't guarantee an acceptable partner. How many partners can a player reject before having to play with the next partner drawn? I constructed Table 2 allowing different numbers of rejections: none, one, two, three, and so on, allowing up to nine rejections before a player had to play with the tenth partner drawn. There is a .676 correlation between corresponding cells in the tables allowing none versus one rejection, a .985 correlation between cells in the tables allowing one versus two rejections. The correlation increases to .993 for two versus three rejections, .998 for three versus four rejections, .999 and higher between tables allowing more rejections. I also constructed dense networks allowing different numbers of rejections. The correlation between corresponding cells in the aggregate table are .508 for none versus one rejection (Table 1 describes the network allowing no rejections), .994 for one versus two, .992 for two versus three, .997 for three versus four, .999 and higher between tables allowing more rejections. Limiting rejections to a specific number has different implications for different kinds of players in different networks, but in eighty-person dense and sparse networks, I get similar aggregate structures with limits of five percent or more of the network (four person limit is five percent of the eighty-person network). I use a ten percent limit for Table 2. Players can reject up to eight random partners before they have to play with the ninth.

[9]Limiting players to exit in place of defection excludes most of the learning I expect in a sparse network. Given a choice between alternative partners, players should learn to avoid hostile partners and seek out cooperative partners. I do not model this because I want to hold the strategies in Table 1 constant to see what happens to their relative success in a sparse network. The next step is to see if the same outcomes result when players can learn their neighbor's strategy where it helps distinguish hostile from cooperative partners. Macy (1991:834ff) offers a promising lead for such work with his simulations of a sparse network in which players select future partners on the basis of prior games (also Boone and Macy, 1998, with undergraduates). Lomborg (1996) offers results on players learning to replace their strategy with more productive strategies (also Macy and Skvoretz, 1998, on strategies diffusing through game play). Lomborg's two central modifications to the Axelrod simulations are that ego learns (a random sample of players compare themselves to a randomly selected other player and switches to the other player's strategy if the other player is earning higher payoff) and a veil of noise obscures ego's vision of alter's moves (random error in perceiving alter's move; alter is paid according to his actual move, but ego responds to the misperceived move). He shows that learning makes it possible for cooperation to emerge and be sustained even with substantial amounts of misunderstanding due to noise. A particularly nice point is that learning generates higher levels of cooperation than would occur with TIT FOR TAT playing against itself because a TIT FOR TAT player misperceiving the other player's move triggers a continuing exchange of defections (Lomborg, 1996:289-290; also see the next footnote on noisy networks).

4.1 RESULT 1: BETTER RELATIONSHIPS

I draw three conclusions from the simulation. First, the sparse-network relations are of higher quality. This is so in two ways, both due to players being able to choose between alternative partners.

Relations in the sparse network are more rewarding. Not everyone gains more, but the average game in Table 2 yields a higher outcome than the average game in Table 1. Players earn 27.0 points on average in the 6,320 games defining the sparse network summarized in Table 2. In the dense network summarized in Table 1, the average outcome is significantly lower (21.6 points; 23.7 t-test, $P < .001$).

Relations in the sparse network are also more pleasant, in the sense that there are fewer games with uncooperative players. Players have numerous repeated games with every other player in a dense network. The only way to punish is to be uncooperative, so many relations degenerate into a permanent cycle of retaliation. This happens in organizations, but not often. Vindictive behavior is a self-indulgence I associate with managers past the productive segment of their careers. If efforts to build a cooperative relationship with someone are met with abuse, managers don't spend time punishing the person. They avoid her. They shift to more productive relationships.

—— Figure 3 About Here ——

The more pleasant aspect of relations within the sparse network is illustrated in Figure 3. Bars show the proportion of games in which the other person defects. Dark bars describe games in dense networks (Table 1). Defections are high because of the obligatory repeated games with uncooperative partners in dense networks. Almost half of the games in Table 1 are with an uncooperative partner (47%). Defection is most often against HOSTILE and SNEAKY players. HOSTILE players elicit defections from all responsive players, so they have the highest proportion of games with uncooperative partners (81%). SNEAKY players often get caught making their occasional defection, so they have the next highest proportion (67%). White bars describe games in the sparse network (Table 2). Less than a quarter of the games are with an uncooperative partner (23%). Almost every kind of player benefits. Of the nice players, RIGHTEOUS players benefit most. Their Puritanical standard of never forgiving embroils them in mutual retaliation relations in the Table 1 dense network. In the sparse network, they can stay away from the partners they despise. Their 44% of games with uncooperative partners in the dense network drops to 10% in the sparse network. Even HOSTILE players have more positive games in a sparse network. They face uncooperative partners in only 30% of their games, down from 81% in the dense network. Their behavior is no less aggressive, but other players now avoid them. The result is that HOSTILE players in a sparse network are more often in games with naive partners.

4.2 Result 2: Nasty Player Advantage

The irony is that the more rewarding and pleasant relations in a sparse network most benefit the nasty players. Nasty players do well in sparse networks because they do well in the initial games of a relationship, and few relations in a sparse network develop beyond the initial games. Given relations based on preliminary games, players can go either of two ways: (a) treat the initial games as a honeymoon in which you are cooperative to encourage trust within the relationship, or (b) treat the initial games as a window of opportunity to exploit your partner. In private games, the player who takes the latter course dominates.[10]

The nasty player advantage is illustrated in Figure 4. The vertical axes are cumulative average outcomes in repeated games between two kinds of players. The top graph shows what happens between HOSTILE and TIT for TAT. HOSTILE earns 50 points in the opening game by defecting against TIT for TAT's cooperation. TIT for TAT earns nothing. On the second move, TIT for TAT is uncooperative in retaliation and the HOSTILE remains uncooperative. They each earn 10 points. The average outcome after the second game is 30 points to HOSTILE (average of 50 and 10), and five points to TIT for TAT (average of zero

[10]The TIT for TAT strategy has been criticized in recent simulations for being too quick to punish when moves cannot be interpreted clearly. Add a random function to the simulation so that ego occasionally misperceives alter's move. If alter cooperates, but ego believes that the move was a defection, then ego playing a TIT for TAT strategy punishes with a defection on the next move, which triggers a subsequent defection by alter playing a TIT for TAT strategy, and so locks the two players into a perpetual exchange of reprisals. Strategies more forgiving than TIT for TAT can avoid the exchange of reprisals, and so dominate TIT for TAT where noise is an issue (see Bendor, Kramer, and Stout, 1991; Bendor, Kramer, and Swistak, 1996; Kollak, 1993, 1996). An alternative is to have players learn to switch to more productive game strategies as in Lomborg's (1996) simulations discussed in the preceding footnote. A sparse network probably involves a greater amount of noise than a dense network, because there is less experience with alter to rely upon, but the generosity that is a solution to noise is not a solution to the problems of low density. In a sparse network, strategies more forgiving than TIT for TAT will be eaten alive. The nice strategy that fares best is

and 10). Additional games move their average outcomes to the 10 points expected in a continuing relation of mutual retaliation (cells 2,4 and 4,2 in Table 1). Games between PUSHY and RIGHTEOUS in Figure 4 illustrates the same point with slightly more complex players. PUSHY earns 50 points in the first game by defecting against a cooperative RIGHTEOUS, who earns nothing. RIGHTEOUS never trusts PUSHY again, defecting in every subsequent game. PUSHY cooperates on games two and three, then retaliates with defection. Additional games move their average outcomes to the 10 points expected in mutual retaliation (cells 6,8 and 8,6 in Table 1).

—— Figure 4 About Here ——

What is different in a sparse network is that the Figure 4 relations between nice and nasty players never develop beyond the first game. The nice player exits after the first exploitative move. When a HOSTILE player defects in the first game, for example, TIT for TAT refuses the next assigned game with that player. Note the game frequencies in Table 2. The first number in brackets is the average number of games played between row and column players. The second number in brackets is the average frequency of a row player rejecting a game with a column player. The [4, 3] in the second column of row four means that each TIT for TAT player, on average, completes four games with each HOSTILE player, and refuses the player three times as a partner.

But after rejecting HOSTILE, TIT for TAT forgives. TIT for TAT is open for a new first game with HOSTILE — who again exploits TIT for TAT. The games are infrequent, but the outcomes are consistently no points for TIT for TAT, 50 points for HOSTILE (cells 4,2 and 2,4 respectively in Table 2).

RIGHTEOUS is the optimum nice strategy in a sparse network because it is never exploited twice by the same partner (26.7 row average at the top of Table 2, versus 24.3 and 22.6 for TIT for TAT and OTHER CHEEK). HOSTILE's initial defection is never forgiven. The [1, 6] frequency in the second column of row six of Table 2 means that HOSTILE plays only one game with RIGHTEOUS, after which RIGHTEOUS refuses all future games (an average of six times). Similarly, each RIGHTEOUS player completes only one game with

---

RIGHTEOUS, which never forgives a defection. As for learning, the strategies most likely to be learned under a criterion of who is earning the largest payoffs are abusive strategies, HOSTILE and PUSHY.

each PUSHY player ([1, 5] in row six, column eight). Responding to PUSHY's initial defection, RIGHTEOUS rejects all subsequent games.

But RIGHTEOUS is the best of a bad lot. Paired against nasty players in a sparse network, nice players lose. The most successful are HOSTILE (37.0 row average at the top of Table 2), ERRATIC (30.2), and PUSHY (29.4).

—— Figure 5 About Here ——

The survival graphs in Figure 5 provide more dramatic illustration. The left-hand graph describes what happens in a population that remains sparse across generations; something like the city-street society of strangers. Nice players die quickly (bold solid lines). The two runner-up nasty players then disappear (bold dashed lines), leaving a population of HOSTILE players. If there is even one HOSTILE player in the initial population, his advantage is sufficient to eventually replace the other seven kinds of players. And the other players disappear quickly. In dense-network Figure 2, more than 100 generations pass before TIT for TAT and PUSHY jointly dominate the dense network, and thousands more pass before the population is entirely TIT for TAT players. In sparse-network Figure 5, almost everyone is a HOSTILE player (99%) within 50 generations.[11]

The graph to the right of Figure 5 describes what happens in a mixture of sparse and dense networks.[12] Nice players have the advantage in established relations (Table 1). Nasty

---

[11]The results would be different if players could "lock-in" productive relationships. Stanley, Ashlock, and Tesfatsion (1994) show in theory that allowing exit in what I here describe as a sparse network can lead to cooperative players dominating the network if they can lock-in productive relationships in the sense that once they find a cooperative partner they can stay in that relationship for subsequent games. Boone and Macy (1998) present empirical results with undergraduates in a computer laboratory card game that suggest behavioral qualifications to the theoretical conclusion, but the theoretical point remains that each cooperative player should eventually all pair with another cooperative player, leaving nasty players to games with one another, which means that nasty players disappear in an evolutionary graph such as Figure 6. Locking-in productive relationships makes sense in a marriage market, but what managers stay working with the same contacts across their career? Successful managers are continually meeting new people and forming new work relationships, with new relations seeming to emerge almost independently beyond the immediate work group (footnote 6). Networks within organizations are obviously not random, but neither are they are they a marriage market in which successful managers can withdraw into relationships with a few friends. The "lock-in" solution is not a realistic solution, for managers, to the competitive advantage of nasty players in sparse networks (reputation is; see Rapoport, Diekmann, and Franzen's, 1995, demonstration with students playing against all other students in a laboratory network that mutual cooperation and mutual defection relations tend to "lock-in" when students know their partner's history of behavior with other students in the group).

[12]Generations are projected for the left-hand graph in Figure 5 as described in the text by Eqs. (1) and (2). Player i proportion in the next generation for Figure 5 is the current number, minus deaths, plus births, quantity divided by the new population total (cf. Eq. 2): $(.75N_i + B_i)/NTOT$, where NTOT is the new total of players in

players have the advantage with new acquaintances (Table 2). Like a confidence man in the big city, nasty players can thrive if sufficient new players enter the population. The established relations sustaining nice players means that it takes longer for the nice players to disappear, but they do. The last TIT for TAT player dies in generation 224. HOSTILE doesn't do as well because exploiting partners in dense networks results in low-yield relations of mutual retaliation (Figure 4). PUSHY is more careful about defecting against cooperating partners, and so dominates the mixed sparse and dense population (80% by 100th generation in Figure 5, 99% by the 400th generation).

4.3 Result 3: Mixtures of Dense and Sparse

The mixture graph to the right in Figure 5 is one specific mixture of dense and sparse. New generations are 25% new players because I use a 25% death rate. I use the outcomes in Table 2 to describe sparse-network games with new players, so the mix in Figure 5 is 75% dense to 25% sparse. In a mixture of 75% dense, the graph to the right in Figure 5 shows that PUSHY players are dominant.

In different mixtures of dense and sparse, other players are dominant. The results in Figure 6 tie together the results at specific alternative mixtures of dense and sparse.[13] The

---

the population, Σ_j (.75N_j + B_j), and births are a function of player i success in games with each other kind of player j (cf. Eq. 1):

$$B_i \text{ in next generation} = \Sigma_j P_i P_j Z_{ij} = \Sigma_j P_i [(.75N_j + B_j) Z_{ij}]/NTOT,$$

where Z_{ij} is cell (i,j) in Table 2. This is what happens when little knowledge is carried from one game to the next. In relations with more history, these results are less reasonable because players have to be too forgiving. For example, a TIT for TAT player in Table 2 plays four games with each HOSTILE player, rejecting the player once between each game. The cycle with HOSTILE is play game, get exploited, punish by refusing to take the person as a partner in the next game, then forgive. Four times with each HOSTILE player. TIT for TAT has the forgiveness of a saint. Generations are projected for the right-hand graph in Figure 5 allowing a mix of dense and sparse networks. Births are computed with the following equation:

$$B_i \text{ in next generation} = \Sigma_j P_i [.75N_j Z_{ij(1)} + B_j Z_{ij(2)}]/NTOT.$$

Players breed in proportion to their success in dense games with experienced players and their success in sparse games with new players (where $Z_{ij(1)}$ and $Z_{ij(2)}$ are cells (i,j) in Tables 1 and 2 respectively).

[13] Generations are projected for each population with weighted averages of Table 1 and Table 2 and the populations all begin with equal numbers of the eight kinds of players. I tried linking the tables to new and old players as in the preceding footnote, then manipulating the death rate as a way to mix dense and sparse. This turned out to be a poor strategy because more changes than the mix between dense and sparse. A 0% death rate in the preceding footnote, for example, doesn't generate the graph in Figure 2 for survival in dense networks. It generates a population extremely slow to change because older players never leave. Generations are projected for Figure 6 by simply averaging the outcome matrices in Table 1 and Table 2:

horizontal axis is a continuum from 0% dense networks to 100% dense. The graphs shows how dominance shifts from HOSTILE, to PUSHY, to TIT for TAT players as density increases. Maximum density is to the right. Survival in this population is described by the graph in Figure 2 and the outcomes in Table 1. TIT for TAT is ultimately dominant, but at 47% is still less than PUSHY in the 100th generation (indicated by the gap above the end of the bold line at the right in Figure 6). Minimum density is to the left in the graph. HOSTILE players are dominant. Survival is described by the graph to the left in Figure 5 and the outcomes in Table 2. The .75 density population in Figure 6 lies between the extremes. This is the population described by the graph to the right in Figure 5. PUSHY players are ultimately dominant, but are only 80% of the population in Figure 5 by the 100th generation, which is the generation reported in Figure 6.

TIT for TAT Requires Maximum Density

Three points are illustrated in Figure 6. First, TIT for TAT is dominant only in the maximum density conditions — the conditions in Axelrod's tournament, and kindred analyses such as the simulations summarized in Table 1. Give players a choice between alternative partners for their relationships, which is the usual situation for managers in an organization, and TIT for TAT loses to PUSHY. The lower the density, the faster TIT for TAT loses to PUSHY. This is indicated in Figure 6 by the rapid decline in percent TIT for TAT at lower levels of density levels below 1.0 (bold line); 46% TIT for TAT at .99 density, drops to 40% at .95 density, and 34% at .90 density.

—— Figure 6 About Here ——

Not indicated in Figure 6 is the fact that TIT for TAT players do not survive in populations that are anything less than maximum density. Their disappearance requires more generations at higher levels of density. The specific number of generations is trivial because that changes if the death rate is changed, but the increasing relative numbers of generations is

---

$$B_i \text{ in next generation} = \Sigma_j P_i P_j [(\text{Density}) Z_{ij(1)} + (1\text{-Density}) Z_{ij(2)}].$$

where density is the population network density on the horizontal axis of Figure 6 and P_i is defined as in the preceding footnote. Generations are projected for each of 101 populations distinguished by density increments of .01 from zero to one inclusive. The proportional numbers of HOSTILE, PUSHY, and TIT for TAT players are reported in Figure 6.

stable.[14] At .8 density, TIT for TAT declines to one percent of the population after 567 generations. At .9 density, they decline to one percent after 1,495 generations. At .95 density, they last for 6,750 generations. At .96 density, they last for 32,073 generations. The advantage then shifts. TIT for TAT players begin to hold a minority equilibrium position against a PUSHY majority. At .97 density, they decline to 22.5% of the population after seventy thousand generations, then remain at that proportion in subsequent generations. At .98 density, they reach an equilibrium of 47.6% after sixty thousand generations. At .99 density, they reach an equilibrium of 73.4% after twenty-five thousand generations. Only at 1.0 density, as in Table 1 and Figure 2, do TIT for TAT players expand to eventually replace all others.

PUSHY Dominates Populations More Dense than Sparse

Second, the cautious exploitation by PUSHY players is the dominant strategy in mixed populations more dense than sparse. PUSHY players speed to dominance faster as the dense network disintegrates to sparse — reaching their peak at .7 density (bold dashed line in Figure 6). At .7 density, PUSHY is 81% of the population by the 100th generation (reported in Figure 6), and 99% after 300 generations. Then, quickly, the advantage shifts to more aggressively exploitative players.

HOSTILE Is Dominant in Most Mixtures of Dense and Sparse

Third, HOSTILE is the dominant strategy in predominantly sparse populations (dotted line in Figure 6). The advantage shifts from PUSHY to HOSTILE within a narrow range of density (where the dotted and dashed lines cross in Figure 6). The critical range is .6 to .5 density.

---

[14] The specific numbers of generations in this paragraph are trivial because they change with the death rate, but their relative numbers are significant. The three results in the paragraph are: (1) TIT for TAT eventually disappears in a population below .97 density, (2) TIT for TAT goes to an equilibrium proportion at densities .97 to .99, and (3) at 1.0 density, TIT for TAT eventually replaces all other players. If I replace the 25% death rate in the text with the maximum, a 100% death rate, I get the same results but in fewer generations. At .9 density, for example, TIT for TAT players decline to one percent of the population after 375 generations (versus 1,495 in the text). But in populations with a 100% death rate, it is still true that: (1) TIT for TAT disappears in all populations less than .97 density, (2) TIT for TAT goes to an equilibrium of 22.5% at .97 density, 47.6% at .98 density, and 73.4% at .99 density, and (3) at 1.0 density, TIT for TAT eventually replaces all other players.

Here are the numbers plotted in Figure 6 (population network density, proportion HOSTILE, and proportion PUSHY in 100th generation):

.60	.20	.65
.59	.31	.57
.58	.46	.44
.57	.65	.29
.56	.81	.16
.55	.90	.08

At .60 density, PUSHY players are dominant. They are 65% of the population in the 100th generation, 99% after 250 generations. A few decimal points later, at .55 density, the advantage has shifted to HOSTILE. They are 90% of the population in the 100th generation, 99% after 130 generations. The dotted line in Figure 6 shows that the advantage stays with HOSTILE players at all lower levels of network density.

5. Conclusion and Discussion

To determine the value of understanding trust as a function of public rather than private games, relationships can be studied over time to see how variation in the structure of third parties around a relationship affects its development. The problem is that such observation is expensive and intrusive. Therefore, systematic empirical evidence from the field has been limited to cross-sectional data showing that trust and distrust are more likely in relations embedded in strong third-party ties. How trust and distrust developed as a function of the third parties remains unknown, and causal inferences are accordingly suspect.

Where observation over time is difficult, simulations can be a productive complement to cross-sectional data. I have used iterated prisoner's dilemma games between eight game strategies to simulate the development of relationships among kinds of managers within an organization. My baseline has been Axelrod’s widely-known use of such simulations to argue that cooperation can emerge as the dominant form of interaction even in a society of selfish individuals without central authority.

Two points have been established. First, the simulations replicate Axelrod's results on cooperation. TIT for TAT is the most successful kind of player (Table 1, Figure 2). The TIT for TAT strategy is to begin with cooperation, then reciprocate whatever move your partner made in the prior game. TIT for TAT players are successful because they are: *nice* (they never defect before the other person), *forgiving* (punished partners are presumed rehabilitated), and *easily provoked* (immediate retaliation). As Axelrod summarized the virtues of a TIT for TAT strategy (1980b:403): "Its niceness prevents it from getting into unnecessary trouble. Its provocability discourages the other side from persisting whenever defection is tried. And its forgiveness helps restore mutual cooperation."

Second, TIT for TAT's success is limited to dense networks. The simulated network of Axelrod's tournament was dense in two ways rare in organizations: repeated games with every other person, and no exit from a bad relationship. This is a network characteristic of families more than organizations. I made the simulation more relevant to organizations by shifting to a sparse network in which relations rarely have a chance to develop beyond initial exchanges and players can withdraw from a bad relationship. The result is a change in the quality of relationships and the relative success of player strategies (Table 2, Figure 3, and Figure 6). Relations in the sparse network are on average more rewarding and more pleasant because players can withdraw from bad relationships. The irony is that the higher quality relations serve most to benefit nasty players, players who exploit cooperative partners. In anything less than maximum-density conditions, the advantage shifts from TIT for TAT players to the cautiously exploitative PUSHY players. Just before the network disintegrates to a condition more sparse than dense, the advantage quickly shifts to the more aggressively exploitative HOSTILE players, who dominate all populations more sparse than dense. Nasty players do well in sparse networks because they do well in the initial games of a relationship, and few relations in a sparse network develop beyond the initial games.

The results in sparse networks are less surprising than they are an important reminder. The effect of fewer games is well known. Defection is the optimum prisoner's dilemma move against a partner you won't meet again. Sparse networks are a merely social context in which exploiting cooperative partners can be the optimum strategy.

My conclusion is that trust is unlikely within an organizations based on private games. The rewards of cooperation depend on being able to avoid unproductive relations before they begin. The key to cooperation in a sparse network is not the threat that partners will withdraw from bad relationships. Exploitative players in a large network of private games can always find new partners. The key is the threat that potential partners will avoid unproductive relationships. What keeps nasty players at bay are the friends and colleagues who warn managers away from people known to exploit their partners.[15]

—— Figure 7 About Here ——

Two findings from illustrative survey data reinforce the point. First, private games are rare in the sense that manager relationships are typically embedded in relations with third parties. The amount of embedding varies between relationships, but all tend to involve some third-party presence. For example, a probability sample of 284 senior managers in a large American electronic components and computer manufacturer surveyed in 1989 cited 3,015 colleagues they deemed to be key contacts for information, personal advice, and political support within the organization (Burt, 1992, 1999c). Two thirds of the 3,015 relationships were embedded in relations with one or more mutual friends close to the manager and the colleague (65%), and every cited colleague had some third-party connection with the manager citing them. Similar network questions were given in 1996 to a representative sample of 317 staff officers in two large American financial services firms (Burt, Jannotta, and Mahoney, 1998; Burt, 1999c). The 317 staff officers cited 3,324 key colleague contacts, of whom almost two-thirds were embedded in relations with mutual friends (64%), and again, every cited colleague had some third-party connection with the manager citing them. The same network questions were given in 1997 to a representative sample of 60 senior managers in a large French chemicals firm (Burt, Hogarth, and Michaud, 1998). The 60 managers cited 656 key colleague contacts, of whom slightly more than two-thirds were embedded in relations with mutual friends (68%), and only nine (2%) were with colleagues who had no third-party connection with the manager citing them. These are only illustrative

[15] My purpose in this paper has been to simulate a baseline model of trust in a sparse network of private games to support the conclusion that trust is unlikely without making the games in some way public so that new partners can signal their trustworthiness to one another. My preference at this point is to return to field data on managers to learn how trust is associated with third parties in public games (e.g., Burt and Knez, 1995; Burt, 1999c). Alternatively, one could continue with the network simulation by adding assumptions about how trust might be facilitated by embedding games in a broader social structure (Willer and Skvoretz, 1997, outline directions for such work). For example, Macy and Skvoretz (1998) induce cooperation in a sparse network of simulated prisoner's dilemma games by adding assumptions to define a neighborhood around each player, define certain players as "like-minded," assume that trust is more likely between "like-minded neighbors," and define probabilities with which players interact with neighbors rather than strangers outside the neighborhood. Trust emerges within neighborhoods, then expands between neighborhoods.

data, but they agree on the tendency for manager relationships to be embedded in relations with third parties. Nothing is private at the top of organizations. What begins in private, is unlikely to stay private.

Second, the probability of trust increases with the age of a relationship as expected in private games, but the change involves several years. Figure 7 shows the probability of a colleague being cited for trust or distrust as a function of relationship age in each of the three study populations in the preceding paragraph. The substantive content of the trust and distrust indicators is beyond the scope of this paper (see Burt and Knez, 1995). Most of the cited relationships in the three study populations (4,831 of 6,995) were cited for neither trust nor distrust. Trust is indicated by a manager citing a colleague as someone with whom s/he would discuss potentially damaging personal information, and the increasing bold lines in Figure 7 show that older relationships are more likely to be cited for trust. Distrust is indicated by a manager citing someone as his or her most difficult colleague and the decreasing dashed lines show that recent relationships are more likely to be cited for distrust. The sign of the association between trust and time is not the point here. The point is that several years pass before the association occurs. In the electronics and computer firm at the bottom of Figure 7, trust and distrust were equally likely for the first five years of a relationship. Relations more than five years old are significantly more likely to be cited for trust, but there is no correlation with age for the first five years and no correlation with age after five years. Colleagues seem to be sorted into two groups after five years (a period about equal to two job assignments in this firm); those you trust and with whom you maintain relations, versus others you allow to drift away. The same pattern appears in the chemicals firm except that the watershed number of years is three, after which trust is more likely (and colleagues are especially likely to be cited for trust in relations that have lasted for a decade). The same pattern appears in tendency for staff officers to cite colleagues for distrust except that the watershed number of years is four, after which distrust is much less likely. Only the staff officers show continuous increase in the probability of trust as a relationship ages. Regardless of trust being significantly more likely in relationships that have lasted for three (chemicals firm), four (financial services firms), or five years (electronics and computer firm), the point here is that years is too long a gestation period. Managers typically have to decide over the course of a few days whether to trust a colleague. There is no time to accumulate a history with the colleague. The expedient course is to ask mutual friends and acquaintances about their history with the colleague, whereupon the game is public.

Thus, private games are not only too dangerous, they are too rare and too slow to be the foundation for trust within organizations. Managers are protected from exploitation by the information third parties provide, they more often than not have mutual friends to whom they can turn as third parties, and they can get that information quickly. It seems safe to conclude that trust and distrust cannot be understood independent of the network context in which they are produced.

Acknowledgment

This material was the subject of a 1994 session of James Coleman's Mathematical Sociology Seminar at the University of Chicago. I am grateful to Kazuo Yamaguchi and James Coleman for comments that guided my revision of the initial manuscript.

References

Axelrod, Robert, (1980a), “Effective choice in the prisoner's dilemma,“ Journal of Conflict Resolution, 24, 3-25.

Axelrod, Robert, (1980b), “More effective choice in the prisoner's dilemma,“ Journal of Conflict Resolution, 24, 379-403.

Axelrod, Robert, (1984), The Evolution of Cooperation, New York, Basic Books.

Barber, Bernard, (1983), The Logic and Limits of Trust, New Brunswick, NJ: Rutgers University Press.

Bendor, Jonathan, Roderick Kramer, and Suzanne Stout, (1991), “Cooperation in a noisy prisoner's dilemma,“ Journal of Conflict Resolution, 35, 691-719.

Bendor, Jonathan, Roderick Kramer, and Piotr Swistak, (1996), “Cooperation under uncertainty: what is new, what is true, and what is important,“ American Sociological Review, 61, 333-338.

Blau, Peter M, (1964), Exchange and Power in Social Life, New York: John Wiley.

Blau, Peter M, (1968), “Interaction: social exchange,” Pp. 452-458 in Vol. 7 of The International Encylopedia of the Social Sciences, New York: Free Press and Macmillan.

Blau, Peter M, (1994), Structural Contexts of Opportunities, Chicago, IL: University of Chicago Press.

Boone, R. Thomas, and Michael W. Macy, (1998), "Dependence, selectivity, and cooperation," Paper presented at the annual meetings of the American Sociological Association.

Burt, Ronald S., (1979), "Relational equilibrium in a social topology," Journal of Mathematical Sociology, 6, 211-252.

Burt, Ronald S., (1992), Structural Holes, Cambridge, MA: Harvard University Press.

Burt, Ronald S., (1999a), "Entrepreneurs, Distrust, and Third Parties," In Leigh Thompson, David Messick, and John Levine (Eds.) Shared Cognition in Organizations, Mahwah, NJ: Lawrence Erlbaum Associates.

Burt, Ronald S., (1999b), "Trust and gossip on the path to equilibrium," In Alessandra Casella and James E. Rauch (Eds) Network and Market Models of the Economy, New York: Russell Sage Foundation.

Burt, Ronald S., (1999c), "Evidence of density, distrust, and character assassination." Graduate School of Business, University of Chicago.

Burt, Ronald S., Joseph E. Jannotta, and James T. Mahoney, (1998), "Personality correlates of structural holes," Social Networks, 20, 63-87.

Burt, Ronald S., Robin M. Hogarth, and Claude Michaud, (1998), "The social capital of French and American managers," Graduate School of Business, University of Chicago.

Burt, Ronald S. and Marc Knez, (1995), "Kinds of third-party effects on trust," Rationality and Society, 7, 255-292.

Coleman, James S., (1988), "Social capital in the creation of human capital," American Journal of Sociology, 94, S95-S120.

Coleman, James S., (1990), Foundations of Social Theory, Cambridge, MA: Harvard University Press.

Davis, James A., (1970), "Clustering and hierarchy in interpersonal relations: testing two graph theoretical models on 742 sociograms," American Sociological Review, 35, 843-852.

Ekeh, Peter P., (1974), Social Exchange Theory, Cambridge, MA: Harvard University Press.

Flood, Merrill, (1958), "Some experimental games,“ Management Science, 5, 5-26.

Gibbons, Robert, (1992), Game Theory for Applied Economists, Princeton, NJ: Princeton University Press.

Granovetter, Mark S., (1985), "Economic action, social structure, and embeddedness," American Journal of Sociology, 91, 481-510.

Granovetter, Mark S., (1992), "Problems of explanation in economic sociology," In Nitin Nohria and Robert G. Eccles (Eds.) Networks and Organization, Boston, MA: Harvard Business School Press.

Greif, Avner, (1993), "Contract enforceability and economic institutions in early trade: the Maghribi traders coalition." American Economic Review, 83, 525-548.

Gulati, Ranjay, (1995), "Familiarity breeds trust? The implications of repeated ties for contractual choice in alliances," Academy of Management Journal, 38, 85-112.

Gulati, Ranjay and Martin Gargiulo, (Forthcoming), "Where do interorganizational networks come from?" American Journal of Sociology.

Gulati, Ranjay and James D. Westphal, (Forthcoming), “The dark side of network embeddedness: an examination of the content of board interlocks, third-party ties, and the formation of interfirm joint ventures.” Administrative Science Quarterly.

Hardin, Russell, (1990), “The social evolution of cooperation,“ In Karen S. Cook and Margaret Levi (Eds.) The Limits of Rationality, Chicago, IL: University of Chicago Press.

Hechter, Michael, (1987), Principles of Group Solidarity, Berkeley, CA: University of California Press.

Heider, Fritz, (1958), The Psychology of Interpersonal Relations, New York: John Wiley.

Hirschman, Albert O., (1970), Exit, Voice, and Loyalty, Cambridge, MA: Harvard University Press.

Homans, George C., (1961), Social Behavior, New York: Harcourt, Brace & World.

Knez, Marc and Colin Camerer, (1994), “Creating expectational assets in the laboratory: coordination in weakest-link games,“ Strategic Management Journal, 15, 101-119.

Kollock, Peter, (1993), “An eye for an eye leaves everyone blind: cooperation and accounting systems,“ American Sociological Review, 58, 768-786.

Kollock, Peter, (1996), “The logic and practice of generosity,“ American Sociological Review, 61, 341-346.

Kreps, David M., (1990), “Corporate culture and economic theory,” In James E. Alt and Kenneth A. Shepsle (Eds.) Perspectives on Positive Political Economy, New York: Cambridge University Press.

Labianca, Giuseppe, Daniel J. Brass and Barbara Gray, (1998), “Social networks and perceptions of intergroup conflict: the role of negative relationships and third parties,” Academy of Management Journal, 41, 55-67.

Lawler, Edward J. and Jeongkoo Yoon, (1996), “Commitment in exchange relations: test of a theory of relational cohesion,” American Sociological Review, 61, 89-108.

Larson, Andrea, (1992), “Network dyads in entrepreneurial settings: a study of the governance of exchange relationships,” Administrative Science Quarterly, 37, 76-104.

Lomborg, Bjørn, (1996), “The evolution of social structure in the iterated prisoner's dilemma,“ American Sociological Review, 61, 278-307.

Macaulay, Stewart, (1963), “Non-contractual relations in business: a preliminary study,” American Sociological Review, 28, 55-67.

Macy, Michael W., (1991), “Learning to cooperate: stochastic and tacit collusion in social exchange,“ American Journal of Sociology, 97, 808-843.

Macy, Michael W. and John Skvoretz, (1998) “Trust and cooperation between strangers,” American Sociological Review, 63, (in press).

Marwell, Gerald and Pamela Oliver, (1993), The Critical Mass in Collective Action, New York: Cambridge University Press.

Marsden, Peter V., (1983), “Restricted access in networks and models of power,” American Journal of Sociology, 88, 686-717.

Oskamp, Stuart, (1971), “Effects of programmed strategies on cooperation in the prisoner's dilemma and other mixed-motive games,“ Journal of Conflict Resolution, 15, 225-229.

Putnam, Robert D., (1993), Making Democracy Work, Princeton, NJ: Princeton University Press.

Rapoport, Anatol, (1987), "Prisoner's dilemma," In John Eatwell, Murray Milgate, and Peter Newman (Eds.) Game Theory, New York: Macmillan Press.

Rapoport, Anatol, M. Guyer, and D. Gordon, (1976), The 2 x 2 Game, Ann Arbor, MI: University of Michigan Press.

Rapoport, Anatol, Andreas Diekmann, and Axel Franzen, (1995) "Experiments with social traps IV: reputation effects in the evolution of cooperation," Rationality and Society, 7, 431-441.

Roth, Alvin E., (1995), "Introduction to experimental economics," In John H. Kagel and Alvin E. Roth (Eds.) The Handbook of Experimental Economics, Princeton, NJ: Princeton University Press.

Roth, Alvin E. and Ido Erev, (1995), "Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term," Games and Economic Behavior, 8, 164-212.

Sobel, Joel, (1999), "Another view of trust and gossip," In Alessandra Casella and James E. Rauch (Eds) Network and Market Models of the Economy, New York: Russell Sage Foundation.

Stanley, E., D. Ashlock, and L. Tesfatsion, (1994), "Iterated prisoner's dilemma with choice and refusal of partners," In Christopher G. Langton (Ed.), Artificial Life III, Vol. 7, New York: Addison-Wesley.

Stinchcombe, Arthur L., (1990), Information and Organizations, Berkeley, CA: University of California Press.

Stokman, Frans N. and Evelien P. H. Zeggelink, (1996), "Self-organizing friendship networks," In W. B. G. Liebrand and D. M. Messick (Eds.) Frontiers in Social Dilemmas Research, Berlin: Springer-Verslag.

Tullock, Gordon, (1985), "Adam Smith and the prisoner's dilemma," Quarterly Journal of Economics, 99, 1073-1081.

Uzzi, Brian, (1996), "The sources and consequences of embeddedness for the economic performance of organizations: the network effect," American Sociological Review, 61, 674-698.

Willer, David, and John Skvoretz, (1997) "Games and structures," Rationality and Society, 9, 5-35.

Zeggelink, Evelien P. H., (1994), "Dynamics of structure: an individual oriented approach," Social Networks, 16, 295-333.

Zucker, Lynne G., (1986), "Production of trust: institutional sources of economic structure, 1840-1920," Research in Organizational Behavior, 8, 53-111.

Ronald S. Burt is the Hobart W. Williams professor of sociology and strategy at the University of Chicago and the Shell Professor of Human Resources and Organisational Development at INSEAD. His work is on the social structure of competitive advantage in careers, organizations, and markets.

Figure 1.
The Prisoner's Dilemma.

		Alter's Move	
		Cooperate	**Defect**
Ego's Move	**Cooperate**	E gets 3, A gets 3 Reward for mutual cooperation	E gets 0, A gets 5 Sucker's payoff and temptation to defect
	Defect	E gets 5, A gets 0 Temptation to defect and sucker's payoff	E gets 1, A gets 1 Punishment for mutual defection

The prosecutor has two suspects in jail. The suspects are questioned separately and cannot communicate with one another. There is insufficient evidence to convict the suspects of their crime, but the prosecutor has sufficient evidence to convict them both of a lesser offense. If neither suspect confesses, then neither will be convicted of their serious crime but both will be convicted of the lesser offense. If both confess, then both will be convicted of the serious crime. Playing the two suspects against one another, the prosecutor tells each that he will go free if he is the first to testify against the other. The informer will go free. The other person will get the book thrown at him.

Figure 2.
Survival in Private Games.

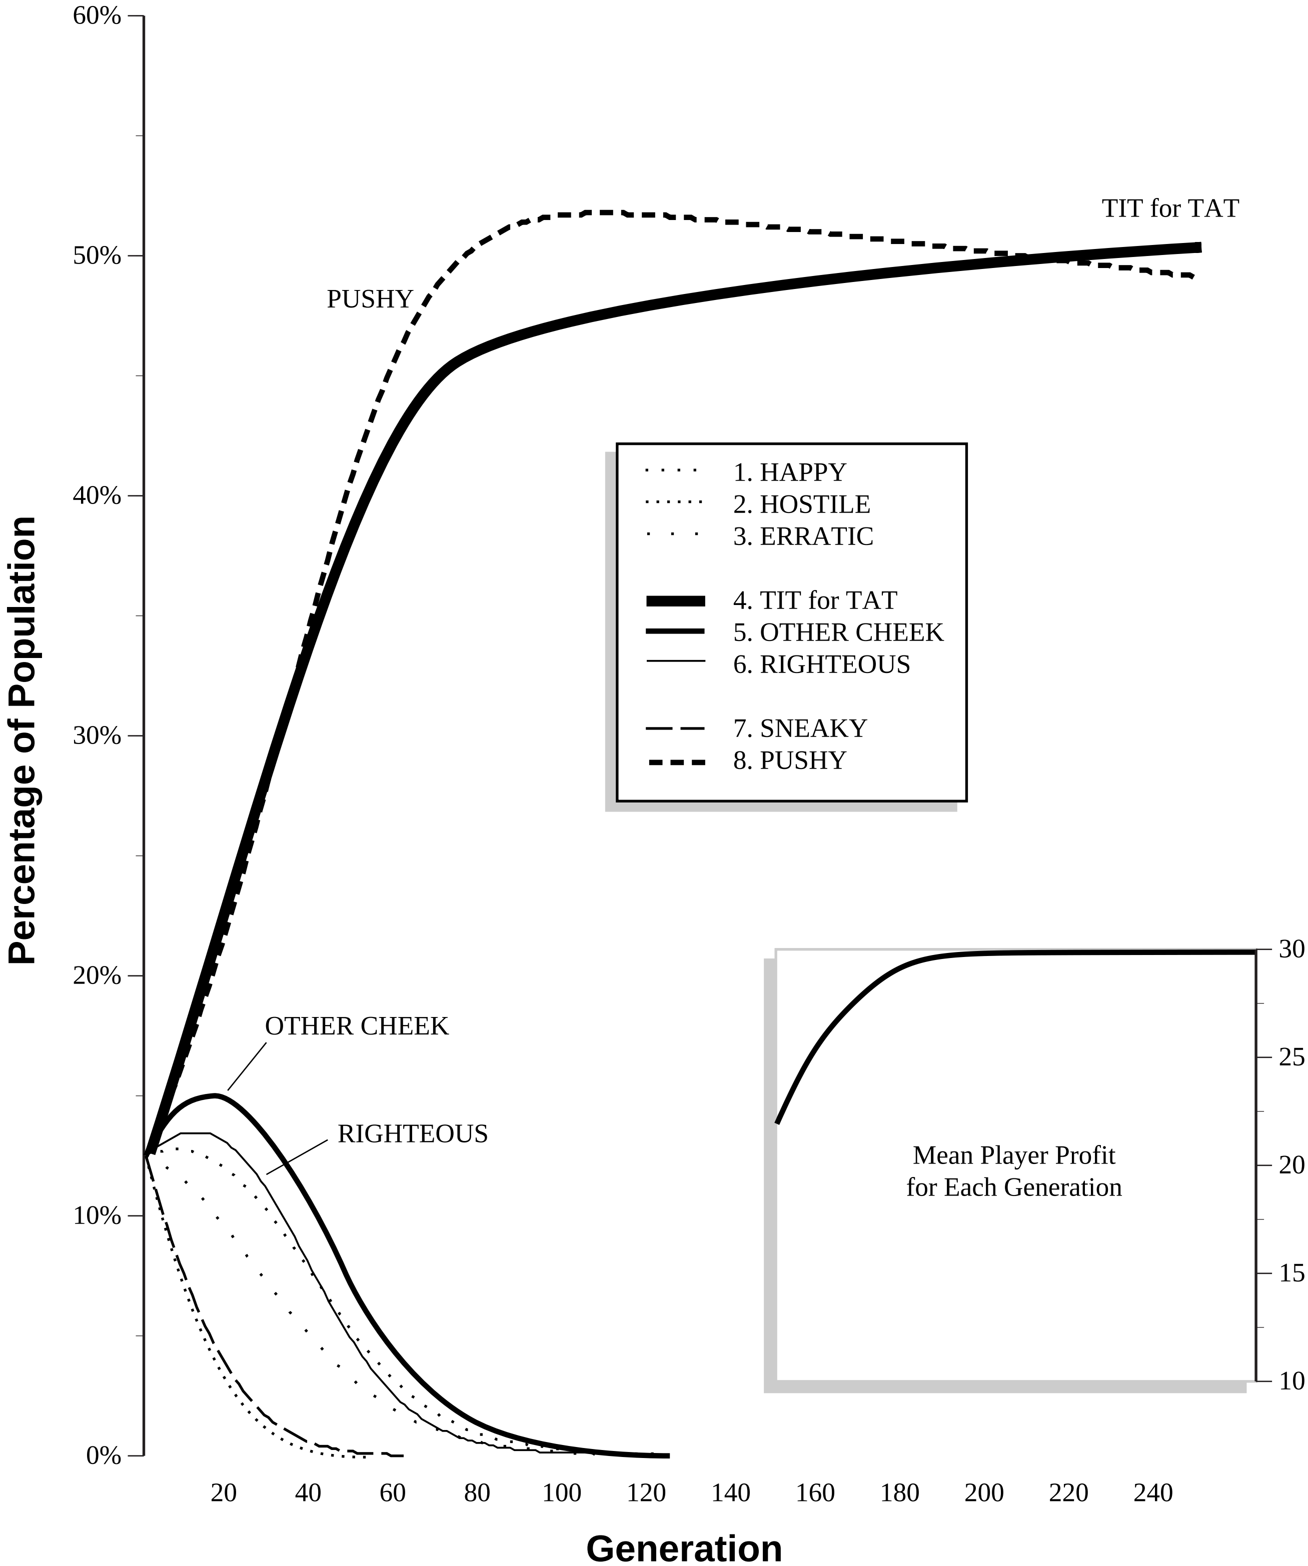

Figure 3.
Games with Uncooperative Partners in Dense and Sparse Networks.

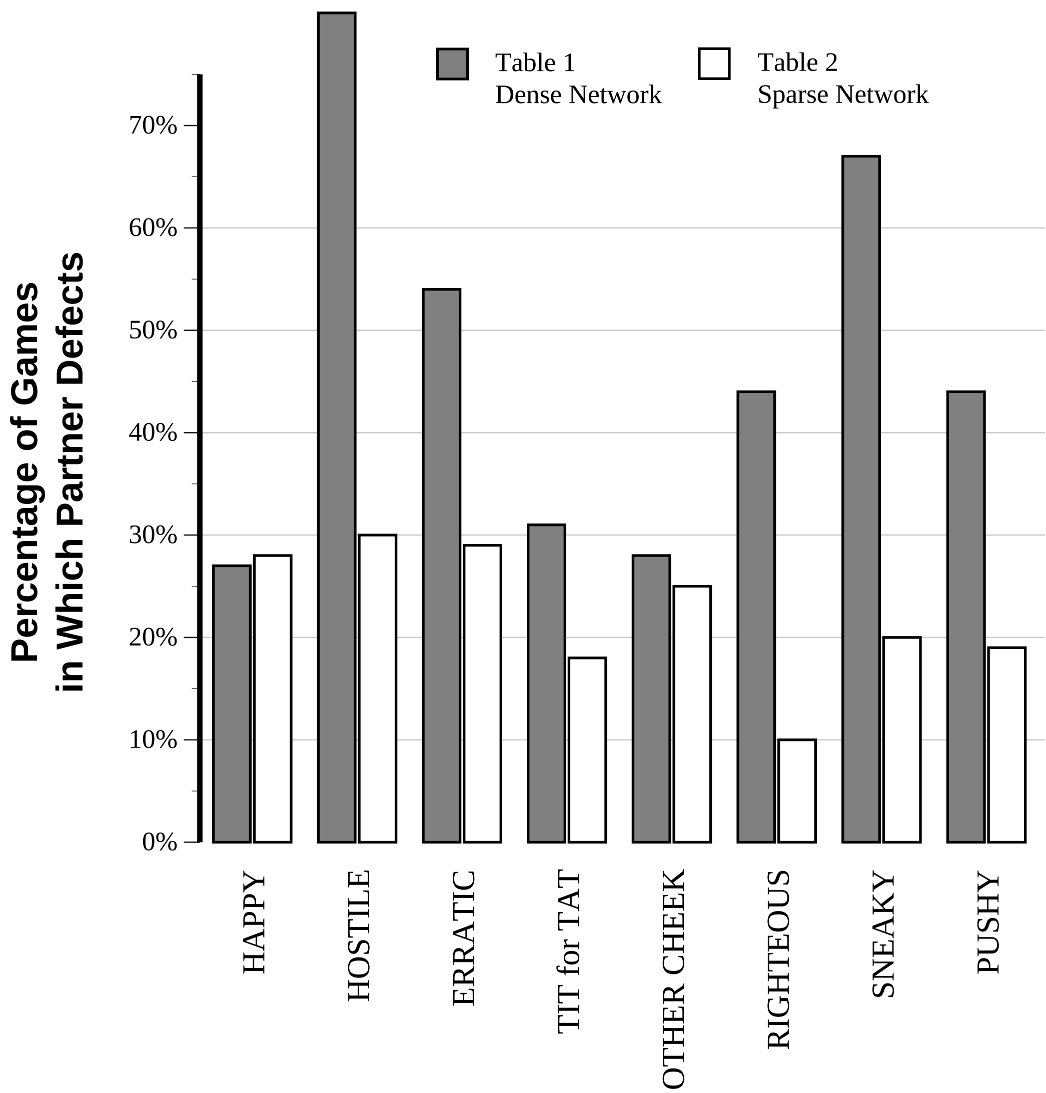

Figure 4.
Nasty Players Do Well in the Initial Games of a Relationship.

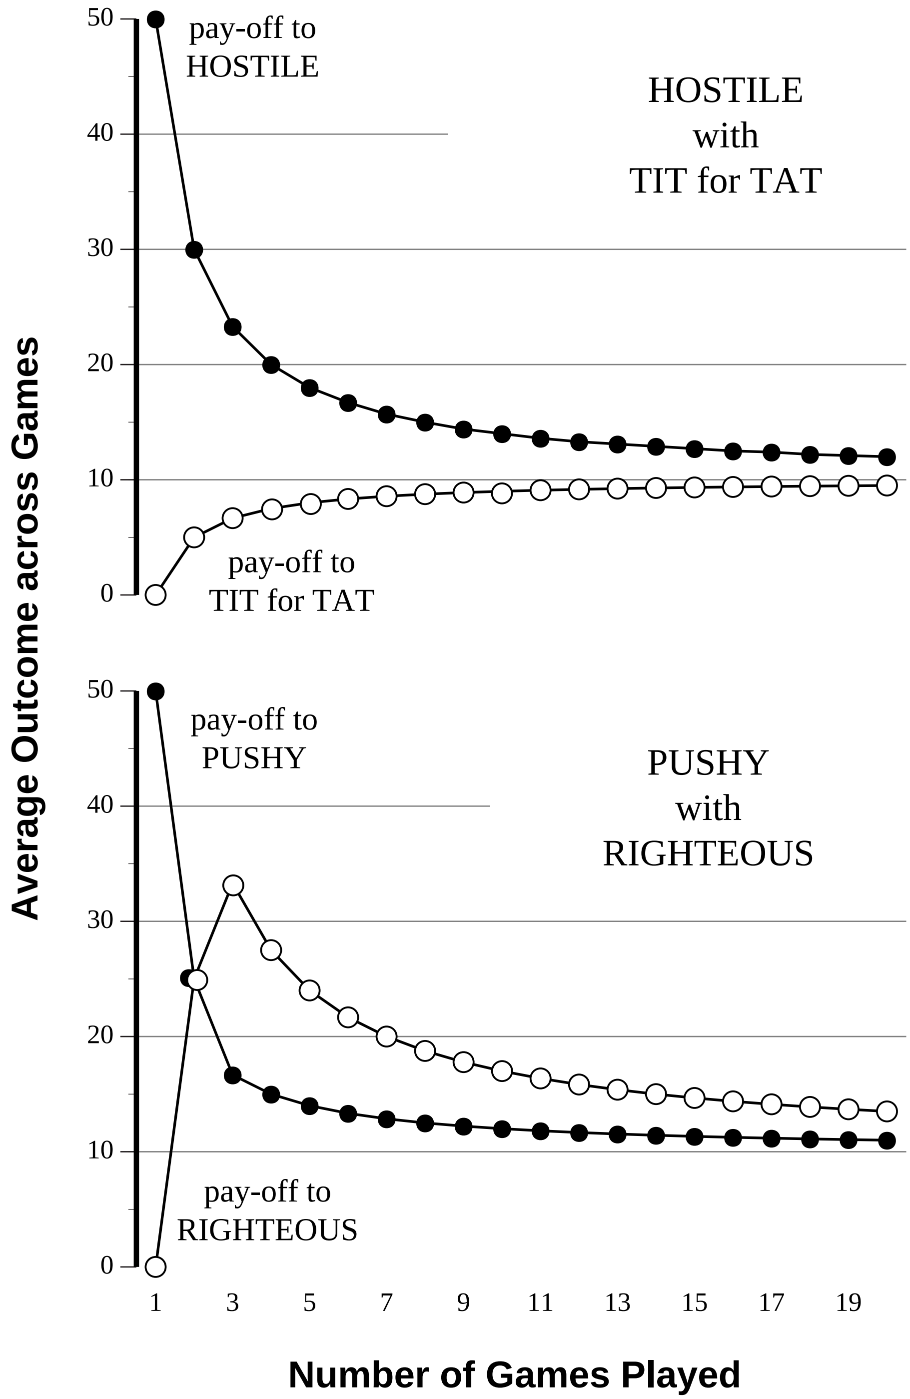

Figure 5.
Survival in a Sparse Network.

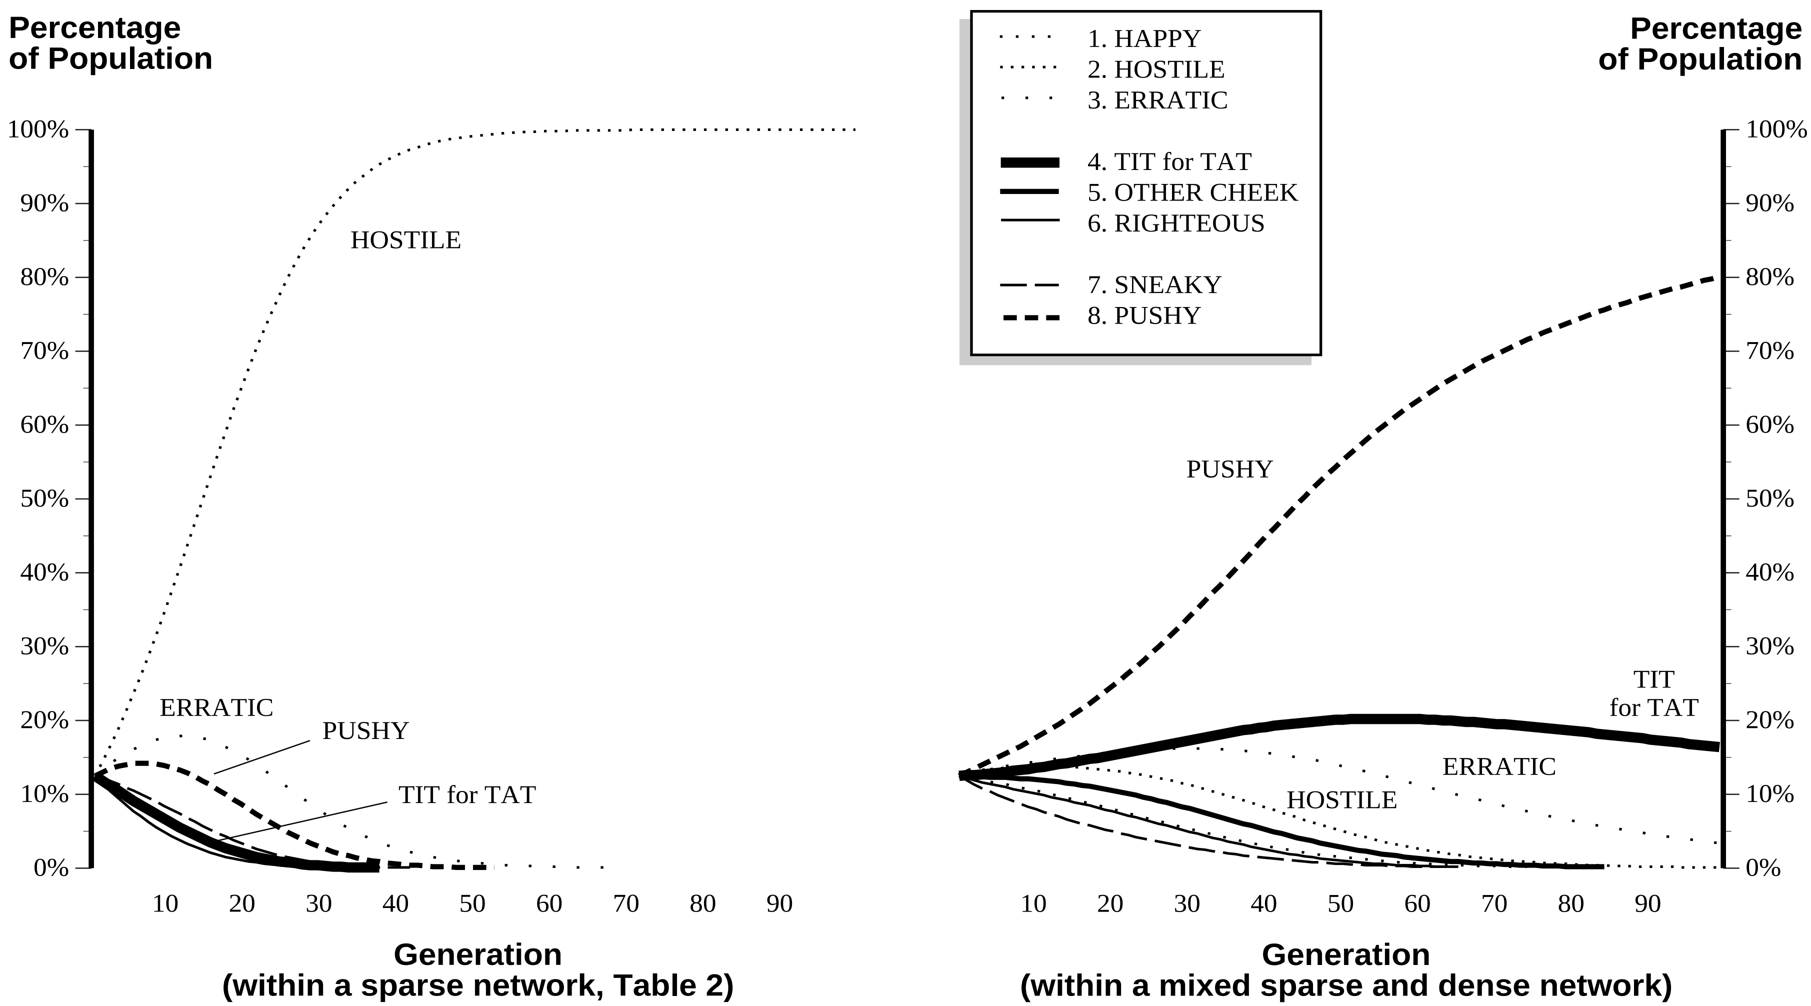

Figure 6.
Flash-Point Shift in a Sparse Network between PUSHY and HOSTILE.

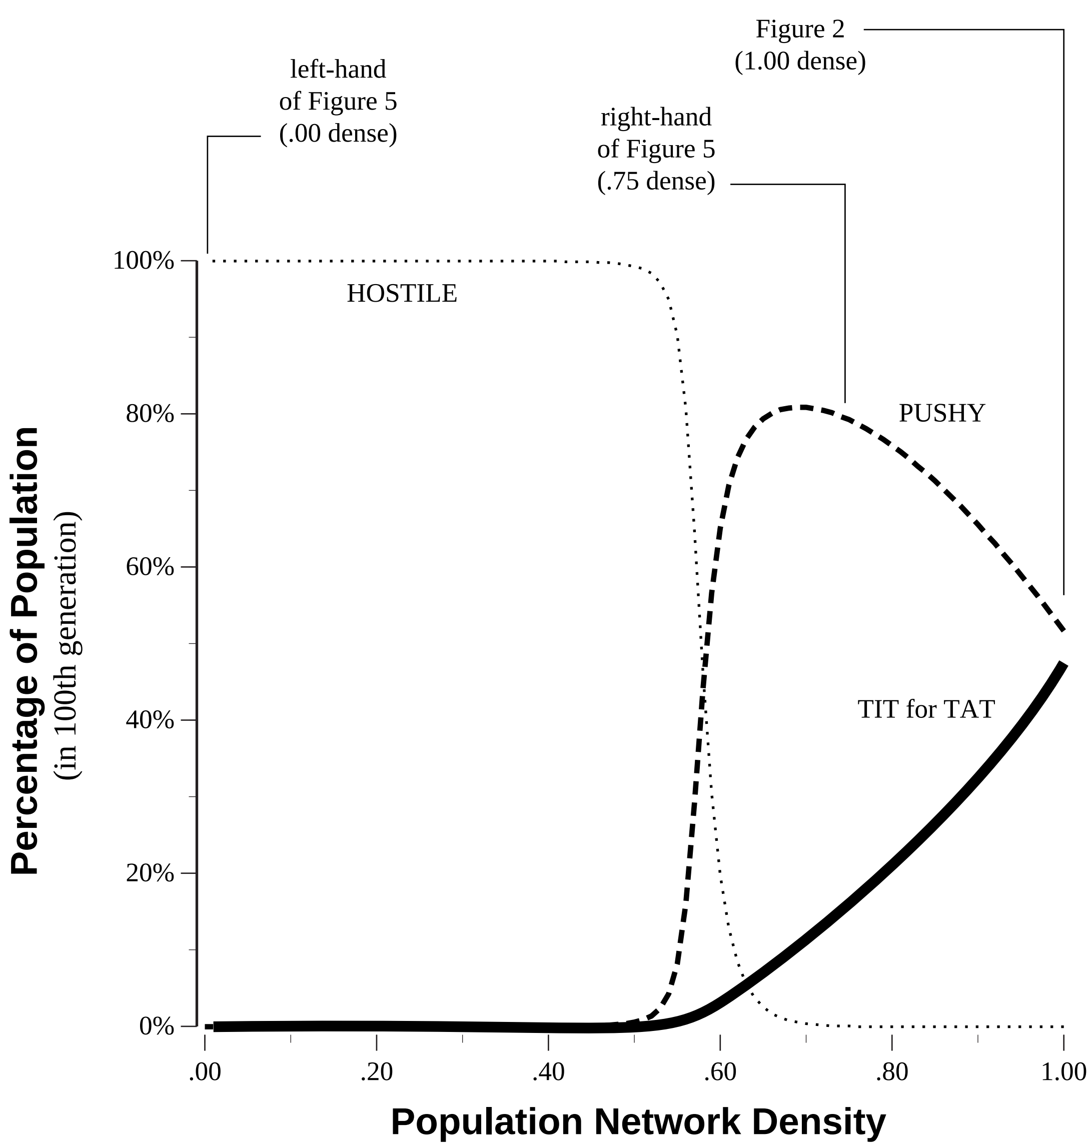

Figure 7
Probability of Trust Increases with Time.

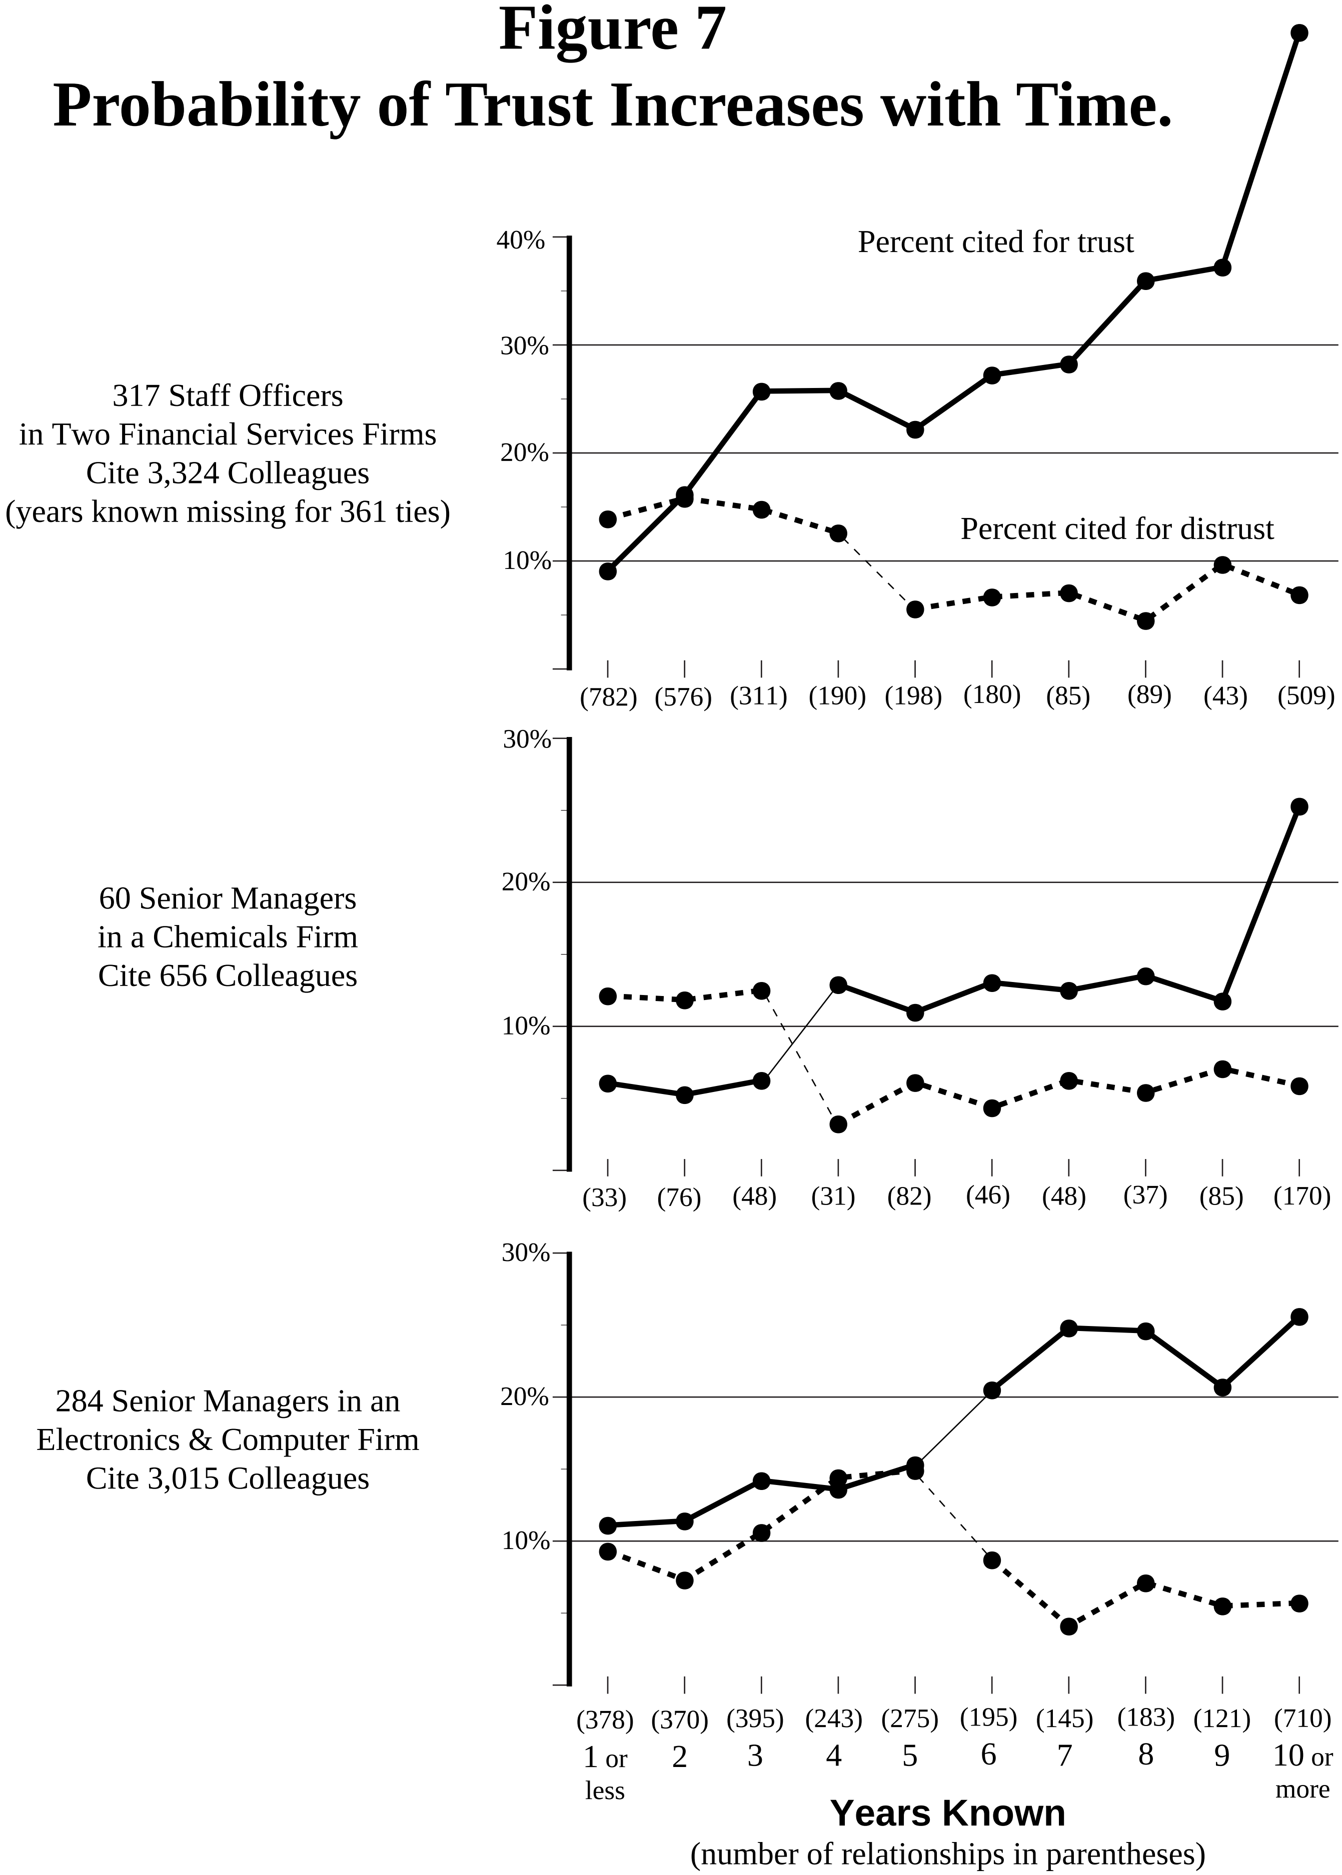

Table 1.
Average Rewards from Private Games.

	1	2	3	4	5	6	7	8
Row Averages:	22.1	17.6	21.2	24.3	23.6	22.6	17.6	24.3
1. HAPPY	**30.0**	**0.0**	**14.8**	**30.0**	**30.0**	**30.0**	**27.1**	**15.0**
(always cooperate)	(0.0)	(0.0)	(0.8)	(0.0)	(0.0)	(0.0)	(0.6)	(0.0)
	[256]	[253]	[254]	[255]	[252]	[252]	[257]	[255]
2. HOSTILE	**50.0**	**10.0**	**29.6**	**10.2**	**10.3**	**10.2**	**10.2**	**10.2**
(never cooperate)	(0.0)	(0.0)	(1.3)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)
	[253]	[253]	[254]	[254]	[253]	[253]	[255]	[252]
3. ERRATIC	**40.2**	**5.1**	**22.3**	**22.3**	**31.2**	**5.4**	**20.7**	**22.3**
(random)	(0.6)	(0.3)	(1.3)	(0.6)	(0.9)	(0.3)	(0.7)	(0.7)
	[254]	[254]	[254]	[250]	[252]	[254]	[253]	[254]
4. TIT for TAT	**30.0**	**10.0**	**22.2**	**30.0**	**30.0**	**30.0**	**11.7**	**30.0**
(reciprocate)	(0.0)	(0.0)	(0.6)	(0.0)	(0.0)	(0.0)	(1.3)	(0.0)
	[255]	[254]	[250]	[251]	[253]	[252]	[248]	[255]
5. OTHER CHEEK	**30.0**	**9.9**	**18.5**	**30.0**	**30.0**	**30.0**	**25.1**	**15.0**
(tit for two tats)	(0.0)	(0.0)	(0.6)	(0.0)	(0.0)	(0.0)	(4.2)	(0.0)
	[252]	[253]	[252]	[253]	[255]	[253]	[252]	[249]
6. RIGHTEOUS	**30.0**	**10.0**	**29.5**	**30.0**	**30.0**	**30.0**	**10.9**	**10.3**
(never forgive)	(0.0)	(0.0)	(1.1)	(0.0)	(0.0)	(0.0)	(0.8)	(0.0)
	[252]	[253]	[254]	[252]	[253]	[253]	[254]	[254]
7. SNEAKY	**32.0**	**10.0**	**23.0**	**11.9**	**29.4**	**10.9**	**11.0**	**12.2**
(up to 10% deceit)	(0.4)	(0.0)	(0.8)	(1.3)	(5.3)	(0.8)	(0.5)	(1.6)
	[257]	[255]	[253]	[248]	[252]	[254]	[257]	[253]
8. PUSHY	**40.0**	**10.0**	**22.3**	**30.0**	**40.0**	**10.1**	**12.0**	**29.8**
(up to 50% deceit)	(0.0)	(0.0)	(0.8)	(0.0)	(0.0)	(0.0)	(1.6)	(0.0)
	[255]	[252]	[254]	[255]	[249]	[254]	[253]	[251]

NOTE — Each cell represents 100 relations between ten players of the row kind and ten players of the column kind (except diagonal cells which exclude the ten self-relations of players with themselves). The relation from player i to player j is the average pay-off for player i in repeated games with player j. Average outcomes at the top of the table are the average of 790 relations from each kind of player. The three numbers in cell (A,B) of the table are: the average of the 100 row A relations with column B players, the standard deviation of the relations (in parentheses), and the average number of games played to define each relation [in brackets].

Table 2.
Game Rewards in a Sparse Network.

	1	2	3	4	5	6	7	8
Row Averages:	21.6	37.0	30.2	24.3	22.6	26.7	25.6	29.4
1. HAPPY	**30.0**	**0.0**	**13.8**	**30.0**	**30.0**	**30.0**	**28.2**	**16.3**
(always cooperate)	(0.0)	(0.0)	(6.6)	(0.0)	(0.0)	(0.0)	(3.2)	(2.9)
	[5, -]	[7, -]	[6, -]	[5, -]	[5, -]	[6, -]	[6, -]	[6, -]
2. HOSTILE	**50.0**	**10.0**	**30.8**	**50.0**	**50.0**	**50.0**	**47.7**	**39.0**
(never cooperate)	(0.0)	(0.0)	(9.3)	(0.0)	(0.0)	(0.0)	(4.8)	(3.8)
	[7, -]	[8, -]	[7, -]	[4, -]	[5, -]	[1, -]	[4, -]	[4, -]
3. ERRATIC	**40.9**	**4.8**	**22.3**	**41.5**	**41.4**	**43.7**	**37.8**	**31.5**
(random)	(4.4)	(2.3)	(7.3)	(4.7)	(5.1)	(6.3)	(6.7)	(4.9)
	[6, -]	[7, -]	[7, -]	[4, -]	[5, -]	[2, -]	[5, -]	[5, -]
4. TIT for TAT	**30.0**	**0.0**	**12.7**	**30.0**	**30.0**	**30.0**	**26.1**	**21.9**
(reciprocate)	(0.0)	(0.0)	(7.1)	(0.0)	(0.0)	(0.0)	(3.7)	(4.7)
	[5, -]	[4, 3]	[4, 2]	[6, -]	[6, -]	[7, -]	[5, 1]	[5, 1]
5. OTHER CHEEK	**30.0**	**0.0**	**12.9**	**30.0**	**30.0**	**30.0**	**27.6**	**16.4**
(tit for two tats)	(0.0)	(0.0)	(7.6)	(0.0)	(0.0)	(0.0)	(3.5)	(2.4)
	[5, -]	[5, 2]	[5, 1]	[6, -]	[5, -]	[6, -]	[6, 0]	[6, -]
6. RIGHTEOUS	**30.0**	**0.0**	**9.4**	**30.0**	**30.0**	**30.0**	**25.9**	**0.0**
(never forgive)	(0.0)	(0.0)	(9.5)	(0.0)	(0.0)	(0.0)	(5.2)	(0.0)
	[6, -]	[1, 6]	[2, 6]	[7, -]	[6, -]	[8, -]	[6, 1]	[1, 5]
7. SNEAKY	**31.2**	**0.6**	**15.7**	**32.6**	**31.6**	**32.7**	**28.7**	**23.2**
(up to 10% deceit)	(2.2)	(1.2)	(8.1)	(2.5)	(2.3)	(3.5)	(3.6)	(6.2)
	[6, -]	[4, 3]	[5, 2]	[5, 0]	[6, -]	[6, 0]	[4, 1]	[5, 1]
8. PUSHY	**39.2**	**2.7**	**18.8**	**35.4**	**39.0**	**50.0**	**33.0**	**26.8**
(up to 50% deceit)	(1.9)	(0.9)	(7.9)	(3.2)	(1.6)	(0.0)	(6.0)	(1.5)
	[6, -]	[4, 3]	[5, 1]	[5, -]	[6, -]	[1, 1]	[5, 0]	[7, -]

NOTE — Each cell represents 100 relations between ten players of the row kind and ten players of the column kind (except diagonal cells which exclude the ten self-relations of players with themselves). The relation from player i to player j is the average pay-off for player i in repeated games with player j. Average outcomes at the top of the table are the average of 790 relations from each kind of player. The results describe the 6,306 relations in which one or more games were played (six cells contain one pair of disconnected players, four cells contain two pairs). The numbers in cell (A,B) of the table are: the average of row A relations with column B players, the standard deviation of the relations (in parentheses), the average number of games played to define each relation [first number in brackets], and the average number of times that each row player rejected each column player as a partner [second number in brackets, "-" means never].