

6-2-2011

QUIT-for-TAT and the Endogenous Structure of Cooperation in Voluntary Dilemmas

John T. Scholz
FSU, jscholz@fsu.edu

T K. Ahn

Follow this and additional works at: http://opensiuc.lib.siu.edu/pnconfs_2011

Recommended Citation

Scholz, John T. and Ahn, T K., "QUIT-for-TAT and the Endogenous Structure of Cooperation in Voluntary Dilemmas" (2011). 2011. Paper 3.
http://opensiuc.lib.siu.edu/pnconfs_2011/3

This Article is brought to you for free and open access by the Conference Proceedings at OpenSIUC. It has been accepted for inclusion in 2011 by an authorized administrator of OpenSIUC. For more information, please contact opensiuc@lib.siu.edu.

QUIT-for-TAT and the Endogenous Structure of Cooperation in Voluntary Dilemmas

T.K. Ahn, Seoul National University

John T. Scholz, Florida State University

ABSTRACT

The ability to select and reject partners creates a powerful means of supporting cooperation when a common set of actors faces repeated possibilities for playing the prisoner's dilemmas with each other, a common situation that we refer to as a voluntary dilemma. The cooperative quit-for-tat (QFT) strategy that maintains all relationships with mutually cooperative partners but quits any relationship after a defection can maintain cooperation in voluntary dilemmas by joining together and excluding nasty and exploitative strategies. We develop the implications of the QFT model for the dynamics and structure of mutual cooperation, and test these implications in an experimental voluntary dilemma. The results confirm that the simple QFT model accounts for observed dynamics and structure of cooperative relationships, and that high-scoring subjects follow strategies that resemble QFT. We discuss the relative importance of niceness, forgiveness, and optimistic search in accounting for the success of QFT strategies, and note that the observed clustering of cooperators in this experimental setting is an artifact rather than a necessary support for cooperation.

How do government agencies learn to cooperate with others on mutually-advantageous joint projects in newly-emergent policy arenas? How do legislative sponsors or academic coauthors find trustworthy collaborators? How can individuals, firms and agencies gain the advantages of specialization required for collaborative high-value exchanges while avoiding the associated risks imposed by opportunistic partners? These situations involve *voluntary dilemmas* in which the opportunities for high-value exchanges are constrained by problems of moral hazard and adverse selection.

We investigate these questions in the context of laboratory experiments in which subjects propose who they want to play with in each period, and may play multiple prisoner's dilemma games in each period whenever both subjects propose to each other. Subjects may choose different strategies in games with different partners. The result can lead to repeated mutual cooperation among consenting partners, but it can also lead to no collaboration when either partner refuses to play or to a repeated "market for lemons" (Akerlof 1970) in which safer low value exchanges (mutual defection) drive out riskier high value exchanges, thereby reducing the potential gains from exchange.

We analyze a self-organizing solution to this problem in which high-value traders independently follow strategies that create communities of cooperation and banish low value traders to "Nash-ville," a ghetto of defectors. The critical requirement for this solution is that high-value traders play "Quit-for -Tat" (QFT), an extension of the well-known "Tit-for-Tat" strategy of the prisoners' dilemma that has been applied to the realm of involuntary dilemmas. QFT players exchange high-value goods, or cooperate, in any new exchange and continue to trade and cooperate until the partner defects by exchanging an inferior good. QFT then ceases to

trade with that partner, and seeks others instead. When multiple players use the QFT strategy a dense, overlapping network of cooperation can endogenously evolve.

We provide experimental evidence that this trial and error strategy explains the evolution of cooperative communities formed by the most successful subjects in the repeated voluntary dilemma game. We find that the QFT partner selection process yields the dense, overlapping networks of reciprocation among cooperators associated with the social capital perspective (Coleman 1987, 1988, Putnam 1993). However, in voluntary dilemmas it is the partner selection process that incidentally produces dense networks of cooperators, not the dense networks that produce cooperation. The ability of QFT to support Pareto superior solutions by selecting appropriate partners is another instance of the general evolutionary selection process noted for other types of games (Skyrms and Pemantle 2000).

We first argue that the kinds of risky high-value exchanges involved when government agencies seek collaborators, when authors seek coauthors, or when firms seek alliances is best represented by the voluntary dilemma in which players can choose and reject partners from a pool of potential players. We then consider the simple solution provided by QFT, and use this perspective to analyze the results of a voluntary dilemma experiment. The results confirm the power of QFT both to allow subjects to maintain cooperative payoffs in the experiment and to predict the observed patterns of exchange shaped by the QFT search strategy. We discuss the implications of this and related research for self-organizing collaboration in the absence of institutional support.

Cooperation in Voluntary Dilemmas

Our analysis of cooperation in voluntary dilemmas explores the general propositions that exchange can enhance an actor's welfare, that greater gains generally require riskier exchanges,

and that opportunism associated with risky exchanges may require institutional support to make risky exchanges possible (North 1990, Williamson 1985). Although the project was initially designed to investigate cooperation among government agencies, the generic problem affects all forms of risky exchange. In the policy field, the problem arises whenever agencies attempt to negotiate agreements to mitigate negative externalities and enhance positive externalities imposed by one agency's policies on others (Steinacker 2010). For example, decisions by a water supply authority to pump more groundwater can reduce stream flows, increasing the concentration of contaminants to levels above thresholds set by an environmental protection agency; two agencies that "exchange" tradeoffs in policies they control individually can maximize their joint benefits, but at the risk of losses if one of the partners proves unreliable (Scholz, Berardo and Kile 2008).

Smaller projects that take advantage of each partner's existing specialized authority, resources and knowledge can provide some joint gain with relatively little risk whenever partners together can create a better project or policy than either could develop on their own, so low-risk exchanges are generally unproblematic. However, more ambitious projects to tackle externalities with greater potential gains will generally require the development of greater technical and administrative capabilities and their dedicated deployment to a unique situation. Once these resources are committed and the agency is locked in to the project, other partners can potentially maximize their profits by shirking on their commitments to take advantage of the agency's investments that are locked in to the project.

Arend and Seale (2005) model this alliance process among businesses as an iterated prisoners' dilemma (IPD) with an exit option. Each partner makes a series of decisions about the level of resources to commit to the activity that are irreversible in the short run, but can be

discontinued contingent on the behavior of other partners. In each period a partner decides between a cooperative but risky high level investment and a shirking but risk-free low level of investment. Both levels are selected to maximize joint output if all partners invested at the same level, with the shirking level set by the Nash equilibrium. As in any IPD, mutual cooperation produces higher expected payoffs than mutual shirking, but shirkers gain the highest overall payoffs and cooperators the lowest if partners choose different strategies. Although decisions at each stage must be made with no clear assurance about the partners' levels of investment, the history of the partnership and the threat of termination provide some support for mutual cooperation. In order to represent the policy example, the *voluntary dilemma* model extends the alliance model by allowing actors to select possible alliance partners among a fixed set of actors involved in mutually interacting policy arenas (cf. Seale, Arend and Phelan 2006).

Experimental Voluntary Dilemmas

To provide a more precise example of the voluntary dilemma to be explored, we next describe the experiment that will be used to test the empirical implications of the QFT model.¹ The experimental voluntary dilemma was designed in part to reflect the conditions in a relatively new policy arena (cf. Scholz, Bererdo and Kile 2008) in which opportunities for both low and

¹ The experiment assesses the relative effectiveness of three different reputational mechanisms that provide information about other players, as described fully and analyzed in Ahn, Esarey and Scholz (2009). They find that a condition in which cooperators could share information about other cooperators outperforms other information-sharing mechanisms, and they analyze differences in the provision and utilization of reputational information to account for the result. We now analyze the results to test the potential role of QFT in explaining the relatively high levels of cooperation across all conditions.

high risk collaboration abound among a small set of subjects who have not yet engaged in collaboration. Subjects have limited resources for collaboration and initially have no knowledge about the behavior of other subjects. To simplify the analysis, the experiment is limited to 2-person exchanges.

The experiment involves 14 subjects who choose partners and play two-person prisoner's dilemmas for 20 periods. In each period, subjects can propose a partnership to any number of other subjects, and an alliance is formed whenever both subjects propose to each other. The partners then play an iterated prisoners' dilemma (IPD) as long as both partners continue to propose to play with each other, which provides either partner with the option to exit the alliance simply by not proposing in the next period. To reflect the many opportunities for exchange in relatively new policy arenas, even low-risk trades (mutual defection in the PD) provide positive payoffs. To represent the limited capability of partners to expand their exchange alliances in these policy arenas, marginal costs for maintaining each exchange increase as new alliances are added. Thus the number of partners sought by a subject will be limited, with the optimal number dependent on the payoffs per exchange. As we will see, the increasing cost limitation plays a critical role in shaping the emergent patterns of high-value exchange.

The experiment allows 14 subjects in each session to simultaneously make proposals for exchanges with any of the other subjects without knowing the proposals of the other subjects. If two subjects propose each other, an exchange takes place in the current period. Each player then chooses between providing a low-value and a high-value investment in the project in the current period. If both choose the high value investment, both receive payoffs of 75 Experimental Currency Units (ECUs where 400 ECUs=\$1). If both choose the low value investment, both receive 25 ECUs, representing the positive value of low-value exchanges even at the shirker's

equilibrium. If one invests the high-value while the other invests the low value, the one providing the low value gains the highest payoff of 100 while the other gains nothing from the exchange. Since the marginal cost of exchange² increases with the number of exchanges, payoffs for mutual low-value exchanges reach their optimal level at four for a net payoff of $(4 \times 25) - 39.6 = 60.4$ ECUs in a given period. High-value exchanges support up to ten exchanges worth $(10 \times 75) - 356.4 = 393.6$ ECUs before the additional exchange would cost more than it would return in gross payoffs.

This procedure is repeated for a total of 20 periods of exchange. In each period subjects repeat the process of proposing to specific partners and of exchanging when mutual proposals are made. Thus each partner can unilaterally terminate any exchange at the end of any period by simply not proposing to the same subject. They can try to maintain the relationship by proposing to the same player, but the other player must also propose to the subject for the relationship to be maintained.

The experiment contained three different reputational mechanisms). The baseline condition provided no information apart from the subject's own experience. The local condition allowed subjects to ask their exchange partners for recommendations, while the central condition allowed subjects to post recommendations on a central information board where other subjects could request the recommendations for any other subject. Four sessions were run for each of three conditions. The 168 subjects were recruited from social science courses at ----, and earned performance-based payoffs averaging \$20-\$25 for sessions that lasted from 60 to 90 minutes. An earlier analysis (Ahn, Esarey and Scholz 2009) found that the local condition outperforms other information-sharing mechanisms because of the coupling of the information provision dilemma

² This is modeled as a cost function $c = 4.4(n - 1)^2$ for $n > 0$ and $c = 0$ if $n = 0$, where n is a non-negative integer indicating the number of exchanges the player has in a period.

with the original exchange dilemma. We now analyze the results to test the extent to which the QFT model predicts the observed patterns of cooperation.

The Quit-For-Tat Model of Cooperation

The QFT model extends the logic of the well-known “tit-for-tat” (TFT) strategy in IPD to explore the ability of a “quit-for-tat” strategy to foster cooperation in voluntary dilemmas. Like TFT, the QFT strategy is both nice and vengeful in the new context (Bendor and Swistak 1997). It is nice because it always cooperates until it meets a defector, and it is vengeful because it exits any relationship immediately if the partner defects and never proposes to play with that partner in the future. The pure QFT strategy thus resembles the “grim trigger” strategy (cooperate until your partner defects, and then never cooperate again), but we later consider more forgiving alternatives willing to accept more than one defection before exiting.

Nice, vengeful strategies have the advantage of requiring the lowest clustering proportion in evolutionary games based on IPD and other similar games (Bendor and Swistak 1997). That is, strategies like TFT and grim trigger can guarantee survival (under proportionate fitness selection rules) as long as they comprise a threshold level of a fully mixed infinite population, whereas nasty strategies like always defect (AllD) require greater proportions before they can guarantee survival. Any clustering mechanism that increases the ability of nice, vengeful strategies to interact more frequently with each other will increase the proportion of mutual cooperation payoffs to mutual defection, thus increasing their survival advantage over AllD (e.g., Axelrod 1981). In smaller populations, the advantage of TFT translates into a higher probability that TFT rather than AllD will take over the population (Nowak 2006). In noisy evolutionary processes, this evolutionary advantage ensures that small populations will spend more time

dominated by TFT than by ALLD, with the ratio of domination determined by the ratio of payoffs. (Fudenberg and Imhof 2006)

Nice, vengeful strategies can also support cooperation in the amended IPD when an exit option is added. Hirschliefer and Rasmusen (1989) demonstrate that “out-for-tat”, a simple extension of the TFT strategy that always cooperates and simply exits from the relationship whenever the partner defects, can support mutual cooperation in IPD with exit for both 2-person games and n-person extensions. They assume a fixed value for the exit option, but other studies have explored the advantages of “out-for-tat” for various contexts that can alter the value of the exit option.³ For example, Orbell and Dawes (1991) provide experimental evidence that the ability to opt out can create a *cooperators’ advantage* even in one-shot games. They find that potential defectors are pessimists who think that others are likely to defect, and are thus considerably more likely to opt out whenever the payoff for not playing exceeds the expected Nash equilibrium payoff from mutual defection. Cooperators, on the other hand, are optimistic in expecting others to cooperate, and therefore prefer to play whenever the exit option is lower than that for mutual cooperation. Optimistic cooperators thus choose to play at much higher frequencies, leading to higher payoffs for those who do choose to play.

The cooperator’s advantage and the value of exit may change, however, when exit from one relationship leads to the beginning of another series of games played with a random draw from the available pool, as in the voluntary dilemma. When exit leads to a random draw of a new partner, for example, nice strategies can be exploited by “roving defectors” who defect in

³ Various exit options are explored in laboratory experiments (Biele and Reiskamp 2008, Boone and Macy 1999, Coricelli, Fehr and Fellner 2004, Seale, Arend and Phelan 2006), Axelrod-type tournaments of strategies (Hayashi and Yamaguchi 1998), and agent-based models of the evolution of cooperation (Hanaki et al 2007, Hruschka and Henrich 2006, Takahishi 2000).

the first round and then quit (Boone and Macy 1999, Dugatkin and Wilson 1991). Exit in this case allows roving defectors to maximally exploit nice strategies by exiting before retaliatory punishment can be imposed.

“Quit-for-Tat”, the Cooperative Community, and Nash-ville

The QFT model of voluntary dilemmas extends this literature on cooperation by exploring the potential ability of cooperators who can choose their partners to *endogenously* create their own cooperative community that protects against roving defectors by banishing them to Nash-ville, a community that shares the Nash-dominant mutual defection outcome. While most studies have explored how clustering influences strategy choice or the emergent frequency of cooperative strategies, the QFT model analyzes how partner selection can endogenously evolve clusters capable of sustaining mutual cooperation.

Network analysis provides a useful foundation for analyzing the emergent structures of cooperation in voluntary dilemmas settings. Each player in the voluntary dilemma is represented as a node, and each game played between players is represented as a link between the two players. An ego network consists of a given player (ego) and all partners (alters) linked to ego by a game. Consider the emergent structure for a simple voluntary dilemma model in which fixed populations of QFT confront the voluntary dilemma equivalent of the IPD “all defect” strategy that always defects and never quits (DNQ). A QFT player will continue to terminate partnerships with DNQ and search for other QFT players until the optimal number of mutual cooperation exchanges is reached. DNQ never quits, but can never maintain relationships with QFT, so it eventually gets paired with the optimal number of mutual defectors.

Pairwise Stability, QFT, and the Structure of Cooperation

One means of analyzing the structure of cooperation is based on the concept of a *pairwise stable* equilibrium for choice process in which individual pairs are free to change links at a given moment, but cannot coordinate with other pairs (cf. Snijder's 2001 actor-oriented model). Pairwise stability requires that no two players can both be made better off by either dropping a current link between them or adding a new link (Jackson 2008). Any pairwise stable equilibrium structure of links in the voluntary dilemma would segregate QFT and DNQ in separate populations, since QFT always prefers to break ties with DNQ. At least one pairwise stable equilibrium exists in the experimental setting in which QFT and DNQ form two separate rings, with each QFT player linked to (up to) the nearest five neighbors in both directions in their ring and each DNQ player linked to the nearest two neighbors in both directions.

Whenever the population of QFT is smaller than the optimal number of partners, only a fully-linked cluster containing all QFT would be pairwise stable. In our experiment with a total population of 14, for example, all QFT players would seek to link to all others whenever there are 11 or fewer QFT players in the population. In this case, pairwise stability in the QFT model would predict a fully-linked cluster of all QFT players with no links to the DNQ players.

For larger populations facing the same payoffs, no structure will be pairwise stable as long as there are two QFT players with fewer than ten links each with other QFT players or two DNQ with fewer than four DNQ links. For either strategy, the two similar strategies would both be better off linking with each other, so by definition the structure would not be pairwise stable. Thus the cooperator's advantage is enormous in any pairwise stable structure as long as there are enough QFT players since QFT receives the higher mutual cooperation payoff from up to ten players minus the cost of ten links while DNQ receives the lower mutual defection payoff from only four players. The payoffs in the experiment provide a big cooperators' advantage, since

QFT will outscore DNQ as long as there is at least 1 other QFT player in the experiment (Min QFT=75, Max DNQ=60.4).

In sum, the basic QFT model applied to our experimental setting implies that:

H1: cooperators (QFT) and defectors (DNQ) will form segregated clusters.,

H2: average payoffs will be considerably higher among cooperator populations than among defectors.

The Dynamics of Partner Search

Given sufficient time, it is plausible to expect a fully-connected cluster of QFT players segregated from the DNQ population to emerge. On the other hand, most plausible dynamic processes of partner selection in a large population would be unlikely to exactly generate the unusual double ring outcome noted above as one potential pairwise stable equilibrium. There are most likely many interesting pairwise stable structures of cooperative relationships implicit in the QFT model that are worth exploring, and the expected frequency of these structures would depend on the nature of the partner search process defined in an extended model.

The potential complexity of this dynamic process requires us to limit our inquiry to the specific conditions in our experiment and to simple search processes. Specifically, we consider only the case of a new exchange arena in which no exchange relationships exist and nothing is known initially about any potential partners. Consider an introspective partner selection extension of QFT and DNQ based on Orbell and Dawes (1991). If introspective choice strategy myopically assumes that other strategies will do what the strategy itself intends to do, the more optimistic QFT will make considerably more proposals in the early rounds than will the pessimistic DNQ since QFT seeks 10 mutual cooperation links compared to DNQ's four mutual defection links. In each succeeding round both strategies will make fewer proposals as they gain

permanent partners sharing their strategy and thus seek fewer new partners. The process will continue until some pairwise stable structure is reached, at which time no one has an incentive to make additional proposals. Even this simple process provides one testable hypothesis: QFT strategies with more aggressive early search will most quickly establish the optimal number of mutual cooperation partnerships, thereby earning the higher payoffs for a longer period of time.

The optimal search strategy for QFT will be influenced by the relative frequency of QFT and other strategies in the population as well as by the information available about other strategies. For example, aggressive search provides an additional advantage for QFT when roving defectors are present in addition to DNQ. Unlike DNQ, roving defectors prefer seeking new suckers over preserving relationships of mutual defection. As QFT and DNQ strategies respectively increase the number of stable relationships over time, they reduce the proposals made to new partners. Since roving bandits continually seek new partners, they increasingly dominate the pool making proposals to new partners. In each succeeding period, QFT faces a higher likelihood of encountering the sucker's payoff. At the same time, both the likelihood and the remaining value of gaining another cooperative partner decreases, particularly since each additional partner increases the costs of all remaining partnerships. Risk-averse QFT may stop searching for new cooperative partners in this environment even though the number of current partners falls short of what is theoretically optimal. Thus QFT strategies that aggressively search in the more favorable early rounds enjoy both a higher likelihood of encountering QFT and a greater remaining value for the established relationships.

At some point in this process, the expected loss of proposing to a nasty new partner is likely to exceed the expected gain from finding a nice new partner, and the optimal search strategy for QFT would stop seeking new partners. Thus search costs in nasty environments may

stop partner acquisition far short of long-term optimal number of QFT partners. To defend against the threat of roving defectors, QFT strategies in the cooperator community in effect lower the community's gates after the first several periods of selection, relegating an increasing concentration of nasty strategies to deal with each other in Nash-ville.

Search costs are affected not only by the frequency and type of strategies, but also by the mechanism through which subjects can learn about potential partners. For example, the local experimental condition allows players to ask current partners for recommendations about new potential partners. This information mechanism would allow QFT strategies already involved in mutual cooperation to share recommendations, thereby considerably reducing search costs for new QFT partners. As a result of lower search costs, we would expect QFT strategies in the local condition to more readily overcome the roving defector problem, establish more mutual cooperation, and outperform their counterparts in conditions with less informational support.⁴

Furthermore, effective sharing of information would strongly encourage clustering if each QFT player recommends all known QFT partners to each other; all QFT partners of player A in period one would propose to all others and establish mutual cooperation with them by period two (at least up to the optimum mutual cooperation limits). Even in large populations where the likelihood that player A's QFT partners would by chance discover each other is small, the local information mechanism can create the fully-linked cooperative clusters that are assumed to support cooperation in the social capital literature (Coleman 1988; Scholz and Berardo 2010). But note that in the QFT model the clustering is a product of the information mechanism, an artifact that by definition does not alter the propensity of QFT to cooperate.

⁴ Although a central information condition could potentially provide an equally valuable support for QFT, the central mechanism implemented in the experiment was not successful due to subjects' inability to solve the information provision dilemma involving the entire population (Ahn, Esarey and Scholz 2009).

In sum, the empirically-testable implications of the introspective QFT model and the dynamics of search in the experimental setting include

H3: Cooperators will select more new partners than will defectors in early rounds, and will select fewer new partners as the number of mutually cooperative links grows;

H4: Optimistic cooperators who aggressively seek new partners in early rounds will outscore pessimistic cooperators who will need to seek cooperators in an increasingly nasty environment; and

H5: Cooperators in the local experimental condition will exhibit tighter clustering and will outperform cooperators in the other conditions.

Forgiveness in Noisy and Nasty Environments

The most critical aspect of QFT is the niceness that preserves mutual cooperation with all other nice strategies. Like its TFT namesake, however, QFT may perform poorly in noisy and nasty environments (Bendor and Mookherjee 1987) in which more forgiving QFT strategies would perform better against a wider variety of both nice and nasty strategies. If strategies sometimes make mistakes in noisy environments, for example, a single mistake by a QFT partner would eliminate that partner from all future interactions. Then QFT would have to take the chance of trying to locate a new QFT partner, and in small populations QFT would eventually end up with no partners at all. A forgiving QFT that would only break with a partner after two defections would mitigate this problem. In addition, forgiving QFT could potentially convert into cooperators the suspicious strategies that defect first but respond to prior-round cooperation with cooperation. This conversion would add to the pool of potential cooperators that would not be available to unforgiving QFT.

On the other hand, forgiveness in this case comes at the expense of vengefulness; forgiving QFT is exploited twice rather than just one before breaking links with strategies like DNQ. This exposure to exploitation would provide an unnecessary cost if QFT partners are relatively easy to locate, so the advantage of forgiveness for QFT thus depends on conditions in which QFT players are either prone to making mistakes or are difficult to locate. In addition, forgiveness imposes a greater disadvantage on QFT in the early stages of new exchange arenas, since the added time spent with a defector postpones search in an increasingly nasty pool of available partners. Finally, forgiveness does little good in extremely nasty environments unless coupled with the vengeance of TFT—that is, if it maintains links with defectors but defects until the partner cooperates. In very nasty environments TFT would at least be able to gain the profit from four mutual defection links that DNQ earns, while an unforgiving QFT suffers the suckers payoff from continuous search.

In short, the role of forgiveness is not as clear as the role for niceness in the success of QFT strategies, particularly in nasty environments. The testable implications for our experiment are

H6: Observed cooperative strategies are likely to vary more in forgiveness than in niceness; and

H7. QFT without forgiveness will outperform more forgiving QFT as long as the environment is sufficiently nice.

Testing the Structural Implications of the QFT Model

The dynamic process of cooperators clustering together over time and banishing defectors from their cluster (H1) can be seen in the period 17 snapshot of cooperation and defection choices in Figure 1. Each of the four sociograms represents a different session of the

local condition—the same patterns emerge in the other two conditions as well, although they are not quite as clear. Circles represent each of the 14 subjects in each session, with the size of the circle representing total earnings over all periods. Solid grey arrows represent a choice of the subject at the base of the arrow to cooperate with the subject indicated by the arrowhead; dotted lines represent the choice to defect. No arrow is present when players do not play. The circles have been placed to emphasize the two groupings that have emerged in each session, with the tightly-connected circle composing the community of cooperators and a more dispersed peripheral group composing the Nash-ville of defectors. The clustering of cooperators and banishment of defectors to Nash-ville arises from the stability of mutual cooperation ties and instability of ties with defectors, as can be readily observed in a dynamic view of relationships over the 20 periods in the online appendix (appended for reviewers as Appendix 1).

H1: Cooperators and Defectors Form Separate Populations

To confirm the tendency to form two separate populations and identify members of each for comparison purposes, we used the TABU clustering algorithm in UCINET (Borgatti et al. 2002) to partition each session into two groups that maximize the similarity of linkage patterns or associations within each of the two groups based on the sum of game links observed in the ten periods (8-17) prior to period 18, when endgame effects become evident. This timeframe allows us to observe the pattern of established exchanges after the early-round sorting occurs and before the endgame effects. Note that decisions to cooperate or defect play no role in defining the partition that depends only on game links. Significant partitions were found for ten of the twelve sessions, with partitions being most significant in the local condition and least significant in the central condition sessions, where two sessions did not produce significant partitions.

H2: QFT Cooperators Outscore Defecting Nash-ville Subjects

Differences in cooperation and payoff levels between the two clusters are as striking as suggested by the QFT model, confirming both that the two clusters correspond to cooperators and defector populations and that cooperators outperform defectors (H2). Figure 2 compares the proportion of cooperation choices across all games (including games both within and between communities) in each period for members of the partitions that correspond to the cooperator community and Nash-ville in each experimental condition. The patterns for mutual cooperation and for average payoffs are very similar and support the same conclusions, so are not displayed.⁵ The cooperative communities score dramatically higher than their Nash-ville counterparts on all measures, with cooperators in the local condition approaching full cooperation and the baseline and central conditions ranking respectively lower. The commonly-observed endgame drop in cooperation in the last few rounds is evident in all partitions; since subjects knew that the game lasted 20 periods, the expected decrease in cooperation indicates that cooperation is conditional on expected future payoffs and not due to unconditional altruism even for subjects in the cooperators cluster.

Figure 2 provides evidence that the higher average performance of the local condition that was reported and analyzed in Ahn, Esarey and Scholz (2009) is almost fully accounted for by the enhanced performance of the cooperator community, since they outscore cooperators in both other information conditions and there is relatively little difference across conditions in the

⁵ Across all conditions, high value exchanges (mutual cooperation) occur ten times more frequently (2069 vs. 206) than low value exchanges within the cooperator community, while low value exchanges (mutual defection) occur over three times more frequently (1391 vs. 412) than high value exchanges within Nash-ville. Support for high value exchange is particularly striking in the local condition, with the percentage of all high value exchanges (mutual cooperation) in the session that occur within the cooperator community being 82.7% (1203 of 1454 games), 75% (501 of 668), and 50.2% (365 of 727) for the local, baseline, and central conditions, respectively. The proportion of low value exchanges (mutual defection) in the session that occur within Nash-ville is more similar across conditions at 66.8% (533 of 821), 64.9% (368 of 551), and 69.5% (490 of 705) respectively.

more uncooperative Nash-ville. By allowing cooperators to recommend each other, the partner-based information exchange appears to accelerate and stabilize the selection process of QFT strategies that creates the cooperative community and bans defectors to Nash-ville (H5). The information differences across conditions appear to have little effect on members of Nash-ville, where there are no such systematic differences in cooperation and in earnings.

H3. Cooperators aggressively seek ties in early rounds, but not in later rounds.

The difference in search for new partners between cooperator and defector communities is illustrated clearly in the comparison of the highest scoring cooperators cluster (session 3) and lowest scoring Nash-ville (session 2) in the local condition provided in Figure 3—the more general case will be considered later. The horizontal axis represents the periods, while the vertical axis represents the average number of each type of proposal made in each period per member. A triangle indicates the number of proposals made in the period to new partners not linked in the previous period, while a circle indicates proposals made to partners linked in the previous period who cooperated and an “x” indicates proposals to partners who defected. The solid line indicates the average number of links that were established in each period as a result of these proposals. Note that the total number of proposals generally exceeds the number of links because links are only formed when the proposed partner has simultaneously proposed to the subject.

In confirmation of H4, the cooperators cluster in Figure 3 reflects the introspective QFT strategy that aggressively proposes to many new partners (triangles) in the first few rounds, but increasingly limits proposals exclusively to cooperating partners (circles) as the number of cooperating partners increase. Cooperators seldom propose to defecting partners (“x”), and practically stop searching for new links after establishing six to seven stable cooperative

exchanges. At this point the increasing marginal costs of links makes the expected value of a new link less tempting, since several periods of mutual cooperation with a seventh or eighth partner would be required to compensate for a single period of becoming a sucker.⁶ As in the previous figure, endgame effects appear to affect the stability of cooperative links in the later periods.

Subjects in the Nash-ville cluster in Figure 3 initially propose to fewer new partners than the cooperative cluster does, as suggested by the introspective DNQ strategy (H3). However, they do not follow the expectation that over time DNQ would increase proposals to defecting partners to the optimum four low-value exchanges and consequently reduce proposals to new partners. If anything, proposals to new partners increase over time, while proposals to prior defectors remain below two until the final periods. The continuing high number of new proposals and low number of proposals to defectors in this worst-performing Nash-ville cluster suggests the instability of mutual defection relationships favored by DNQ,⁷ and represents instead a pattern more consistent with a roving defector strategy that prefers to seek new suckers rather than to stick with known defectors.

The combined patterns in Figure 3 illustrates a tendency for new proposals in later rounds to come increasingly from the less cooperative players in Nash-ville and very few from the cooperative community, in which case the proportion of proposers who are nasty will increase

⁶ The marginal cost of an eighth partner is 61.6 given the link cost function $4.4(n-1)^2$. A player with 7 mutually cooperative links proposes to an eighth in a session and gets the sucker's payoff of 0 incurs a pure cost of 61.6. Even when an eighth partner turns out to be a cooperator, the net benefit is $75-61.6=13.4$. Thus, five periods of mutual cooperation with a cooperative eighth partner are necessary to compensate for a single defection encountered while searching for an eighth cooperative partner.)

⁷ The average .75 probability (reported in Table 1) that a Nash-ville subject will propose to a defecting partner would correspond to slightly more than an even chance that a low-value exchange will last more than one round, which is much lower than would be expected from DNQ.

over time.⁸ This supports the contention that more optimistic or aggressive search in early periods will improve the performance of QFT strategies in this small, fixed population (H4), as will be further examined. In addition, the increasing nastiness in available partners over time further clarifies why even the most successful cooperators cluster portrayed in Figure 3 achieves only seven enduring relationships of mutual cooperation, far short of the optimal level of ten. As the environment becomes increasingly nasty, the expected cost of encountering a nasty strategy more rapidly outweighs the future gains from a successful search.

Testing the Strategic Dimensions of QFT: Niceness, Search, and Forgiveness

The analysis to this point has compared behavior and performance of the two communities to predictions of the QFT model. We next focus on the observed strategies of individual community members to see whether strategies of subjects in the cooperators community actually resemble QFT and to test the relative effectiveness of QFT's niceness, forgiveness, and search in earning high payoffs in the experiment.

Representing Strategy Types We categorize observed strategies by extending a simple representation from evolutionary game theory (e.g., Nowak 2006) that allows us to identify the critical QFT dimensions of niceness, forgiveness, and search.⁹ This class of one-period strategies specifies the probability that the strategy will cooperate after each possible state reached in the previous period, which in IPD includes the probability of cooperation if the other

⁸ Overall, the two highest scoring groups decrease their probability of proposing to new partners from about .5 in the first 3 periods to around .1 in periods 15-17, and the probability of cooperating with a new partner for the two lowest-scoring groups decreases from around .3 to less than .05 in the corresponding periods.

⁹ Note that TFT, forgiving TFT, grim, allC and allD can all be represented by extreme values in a continuous representation of strategies (e.g., Nowak 2006). Including both ego and partner's choice to create five possible previous period outcomes (No Game, CC, CD, DC, and DD) turned out not to produce different results than the simpler model, so we base our analysis on the three category classification of previous outcomes based on the partner's choice.

player cooperated, defected, or did not play in the previous period. For voluntary dilemmas in which proposals are also made, we extend the specification to include the probability of proposing to a player who cooperated, or defected, or did not play in the previous round.

Table 1 reports the average probabilities for each element of the strategy observed during the first five periods of play for the categories of subjects identified in each row. The last columns in the table also report average total earnings over the full 20 periods in ECUs and the observed number of mutual cooperation outcomes in period five. These columns provide a comparison of performance for the observed strategies associated with each subject type.

The first six columns report probabilities for each strategic dimension, with the first three columns reporting the probabilities associated with *niceness*. Column 1 reports the average probability of cooperating (C) with new players after no play in the previous period, column 2 the probability of proposing to a player who cooperated (*c*) in the previous period, and column 3 the probability of cooperating with a player who cooperated in the previous period if a game takes place (total cooperation choices divided by the sum of cooperation and defect choices). The next two columns report the *forgiveness* of the subject in the period after a partner defects (*d*), with Column 4 reporting the probability of proposing and column 5 reporting the probability of cooperation if the game is played. Finally, column 6 reports the *search* dimension of the strategy in terms of the probability of proposing after no play; that is, the subject's total proposals to new players in the current period over the number of players not linked to the subject in the previous period.

The first three rows in Table 1 provide the baseline expectations of the three previously-considered strategies for comparison with observed probabilities for different classes of subjects reported in the remaining rows. QFT is as nice as possible (probability of cooperation =1)

because it always cooperates with new players and always proposes and cooperates with partners who cooperated in the previous period. QFT is completely unforgiving ($p=0$) because it never proposes to a defecting partner and would not cooperate if somehow a game were played. Finally, QFT is aggressive in search, proposing to a relatively high number of new partners in comparison with the DNQ strategy, although what would be interpreted as a high probability would depend on the number of other players, the frequency of QFT-like strategies, and the information available about potential partners.

DNQ is nasty, never cooperating ($p=0$) with new players or previous cooperators, although it would have no reason to not propose again to cooperators ($p=1$). It is forgiving in terms of proposing to previous defectors, but it will never cooperate with previously cooperating partners. And it has a lower probability of proposing to new players in comparison to QFT since it seeks only four stable mutual defection links while QFT seeks ten mutual cooperation links. The roving defector (RD) strategy differs from DNQ only in aggressively seeking new links and never proposing to defectors, both of which allow RD to continually seek new suckers.

To classify subject strategies, each row in Table 1 reports the observed probabilities in the first five periods for the subject types identified on the left. Averaging all choices over five periods provides a sufficiently large sample for measuring the proportions, based on the assumptions that a given subject employs the same conditional strategy (not the same choice) with all partners and that the strategy is stable over the entire period.

Do Cooperators Play QFT? Consider first the differences in niceness between subjects who become members of the cooperator community and those relegated to Nash-ville, which are presented separately for each information condition. The biggest difference between the two communities corresponds to the difference in niceness between QFT and DNQ for the two

strategic elements where differences are expected. Table 1 indicates that the cooperators community cooperates about twice as frequently as the Nashville counterparts both with new partners (62% vs 28% averaged across conditions) and with existing partners who cooperated in the last period (87% vs. 46%). The differences in niceness are critical for achieving more mutual cooperation in period 5 (averaging 2.26 vs 0.54), which in part accounts for the higher average earnings (4101 vs 2694 ECUs) for cooperators. Since the product of all niceness probabilities determine the likelihood that a chance encounter with QFT will lead to a second period of mutual cooperation, this likelihood is about five times greater in the cooperators community ($0.62 \times 0.93 \times 0.87 = .50$) than in Nashville ($0.28 \times 0.86 \times 0.46 = 0.11$).

Once mutual cooperation is established, it continues into the next round with a probability of .85 across all subjects and all periods, so the biggest difference accounting for more mutual cooperation payoffs in the cooperators community is the difference in commencing the relationship in the early periods.¹⁰ Ironically, mutual defection continues with a probability of .5, which means it has only a .25, .12, and .06 chance of surviving a consecutive third, fourth and fifth period. Thus low risk, low payoff exchanges are actually less stable than high risk, high payoff exchanges, at least when both possibilities are clearly presented to subjects!

Table 1 again shows that cooperators in the local condition outperform cooperators in the other conditions both in terms of mutual cooperation and total earnings. The local condition scores highest in all niceness elements, but since mutual cooperation is almost certain to continue in all conditions,¹¹ the difference in performance must reflect differences in search efficiency. Information exchanged in the local conditions leads to greater observed niceness toward new

¹⁰ Continuation probabilities are listed for each condition in online appendix 2 (attached as separate document for reviewers.)

¹¹ If the same ratio of niceness to earnings from the local condition were applied to the other conditions, the number of mutual cooperation links would be 40% higher (2.92) in baseline and 50% higher (1.5) in central conditions.

partners in Table 1 (.73 for local versus .66 for baseline and .47 for central), presumably reflecting nice responses to new partners recommended by a cooperative partner. In addition, sincere recommendations presumably increase the likelihood that the partner is also nice, and hence doubly magnify the likelihood of rapidly establishing mutual cooperation for QFT strategies. Neither of these effects of recommendations would be relevant for nasty strategies, suggesting why recommendations enhance only the cooperators advantage in this experimental setting.

The differences between cooperators and Nash-ville members in forgiveness and search are both smaller and less consistent with the QFT and DNQ strategies. Nash-ville members are closer to DNQ in the forgiveness dimension where differences are expected, with Nash-ville subjects considerably more willing to propose a link to defectors (75% vs 59%). But the cooperator community members appear much more willing to propose to defectors than expected in the pure QFT model, as suggested in H6. Finally, the expected differences in search observed between the highest and lowest earning clusters in Figure 3 are not observed on average across communities or conditions. We therefore turn next to analyzing differences in strategies between high and low performing subjects to better understand the roles of forgiveness and search.

Do Top Performers Play QFT? To address this issue, the next rows in Table 1 first analyze the strategies of the highest-earning subjects in each condition, and then the strategies of the top and bottom quartile of subjects as determined by total earnings. The strategies of the top performers in all three conditions strongly confirm the advantages of QFT, since they earn almost double the average payoffs (6670 vs 3420) by employing the niceness, vengeance, and aggressive search of unforgiving QFT. Given their niceness and willingness to propose to new

partners, they quickly establish five to seven mutually cooperative relationships by the fifth period, still short of the optimal 10 for reasons discussed previously.

Differences between the top and bottom quartile earners in the final two rows reinforce the critical role of niceness, since top performers differ in niceness from bottom performers by even greater margins than do the cooperators community and Nash-ville. In this comparison the likelihood that a chance encounter with a pure QFT would develop and maintain a mutually cooperative relationship in the following period is even higher for top quartile earners ($.8 \times .95 \times .96 = .73$) and lower for bottom quartile earners ($.24 \times .79 \times .41 = .08$), which again helps explain the greater number of mutually cooperative relationships achieved by high quartile earners by the fifth period (3.64 vs 0.26) and hence the greater earnings over the full 20 periods (5132 vs 2123).

Although top quartile earners manifest the advantages of QFT's niceness, like the cooperator community they also diverge from QFT in both search and forgiveness (H6). As in the cooperative community, top quartile earners are only slightly more aggressive than bottom quartile earners in searching for new partners (.43 vs .38). They are considerably less willing to propose to defectors (.46 to .65), although this is still far more willing than pure QFT. In addition, they are actually more willing than low quartile earners to cooperate with defectors (.32 vs .20) when a game does take place.¹² In sum, the comparison of high and low performers confirms again that niceness is most clearly related to high performance (H2), and supports the contention that observed forgiveness is more likely than observed niceness to deviate from the pure QFT strategy (H6). The strategies of top individual performers in each condition support the hypotheses that aggressive search (H4) and minimal forgiveness (H7) are associated with

¹² In short, they are more forgiving even than a pure TFT strategy which would retain links with defectors, but would never cooperate until the partner cooperated first.

higher performance, but the broader comparisons of the two communities and of top and bottom quartile earners provides more tenuous support for the importance of forgiveness and search.

Are Forgiveness and Search Important for QFT? The regression results reported in Table 2 provide another means of testing the role of niceness, forgiveness and search on a strategy's performance. The regressions use measures of strategy based on behavior in the first five periods to predict *total earnings* for the full 20 periods, so the analysis specifically tests the impact of early period strategies on total earnings. *Niceness*, *forgiveness* and *search* each represent the average across all relevant elements as indicated in Table 1. OLS estimates are reported since total earnings is distributed normally except for somewhat fatter tails, but robust standard errors clustered by session are reported to account for the potential for correlated errors for subjects within each of the 12 session.

To control for the known differences in performance across experimental conditions, both estimates in Table 2 include dummy variables for the baseline and central condition to account for the lower performance in those conditions in comparison to the omitted local condition, which is represented by the constant. An interaction of condition with the QFT category is included to test for differences in effectiveness across these conditions. Finally, we account for individual differences in early environments because subjects who by chance encounter a higher proportion of nicer strategies in the early rounds are expected to do better regardless of their strategy type. *Proportion nice partners* measures the proportion of new partners in the first three rounds who cooperate initially.

Model 1: High niceness, high search, and low forgiveness outscore other strategies. Attempts to directly test the impact of pure QFT and its optimistic and forgiving variants are hampered by the difficulty of determining an appropriate cutoff point for the QFT category,

since few subjects are fully nice and unforgiving. Estimates in the first column in Table 2 therefore test the independent impact of each dimension of strategy, and include the interaction of niceness with both search and forgiveness to represent QFT's additional gain from combining these elements. The QFT model would be supported with positive coefficients for niceness (H2) and for its interaction with search (H4) and negative coefficients for the interaction with forgiveness (H7).

The coefficients for niceness, search and forgiveness indicate significant positive impacts of each dimension. Since all strategy variables represent proportions with the same potential [0,1] range, the larger coefficient for niceness confirms a considerably greater impact of changes in niceness on performance, with search a distant second and changes in forgiveness having the least impact. More importantly for testing H7, the interaction with forgiveness is actually negative and significant, confirming that vengeance rather than forgiveness increases performance for nice strategies like QFT. Forgiveness improves performance for the nastiest strategies like DNQ who are forgiving only in terms of proposing to defectors, but increasingly hurts strategies as they become increasing nicer. The interaction with search is positive as predicted in H4, although it is not significantly greater than the positive impact of search even for nasty strategies.

The interactions are perhaps easiest to see in Figure 4, which plots the change in estimated performance as forgiveness increases for subjects with the four combinations of low (0) and high (1) levels of both niceness and search.¹³ The highest performance is achieved by the nicest subjects with the highest search (solid blue line) and lowest forgiveness (leftmost point) that most closely resemble pure QFT. DNQ-like nasty strategies with low search (dashed purple

¹³ The extreme conditions are portrayed in order to compare pure QFT with DNQ, even though a value of zero is not plausible for search since it would lead to no games. Since the equation is linear, any intermediate values can be readily imagined by shifting the lines between the extremes.

line) but high forgiveness (rightmost point) earn less than 50% of QFT (3.6 vs 7.8), and even with the highest search (dashed green line) earns at best less than 60% of QFT (4.6 vs 7.8). For nasty subjects, increases in forgiveness and search improve performance. For nice subjects, on the other hand, forgiveness actually lowers scores. In sum, the highest search and lowest forgiveness are both critical for QFT's performance; a complete shift to the lowest search would decrease estimated performance to 71% of QFT's earnings (5.6 vs 7.8), while a shift to complete forgiveness would decrease estimated performance to 85% (6.6 vs 7.8).

Model 2: High Search and Low Forgiveness are required for QFT performance. A more direct test of the importance of search and forgiveness for QFT is attempted in the second model in Table 2, which reports results for a QFT5050 dummy that equals one whenever niceness $>.5$ and forgiveness $<.5$. This generous cutoff identifies 46 subjects with QFT-like strategies. Less generous cutoffs (e.g. above 75th percentile in both niceness and forgiveness) yielded too few subjects to meaningfully test the role of forgiveness and search within the QFT component. The significant interaction terms in Model 2 confirm that higher search and lower forgiveness significantly improve performance for subjects in the QFT5050 category, although QFT5050 itself did not significantly outperform other strategies. The range of niceness and forgiveness included in QFT5050 are apparently too broad to capture the full advantages of QFT in comparison with other strategies comparable in independent dimensions of niceness, search and forgiveness. Only when combined with higher search and lower forgiveness do QFT5050 subjects outperform others, underscoring the importance of combining niceness, unforforgiveness, and optimistic search as QFT does.

Selection, Learning and the Structure of Cooperation

Returning to our policy example, the experimental results suggest that risky high-value trades can emerge and be sustained to the extent that a sufficient number of stakeholders who have developed nice, vengeful, QFT-like strategies elsewhere transfer these strategies to new voluntary dilemmas. Cooperation emerges most rapidly when strategies aggressively seek appropriate partners, particularly in institutional settings allowing partners to share recommendations about other potential partners.

Our analysis finds that mutual cooperation in sustaining high risk, high value exchanges occurs primarily within clusters of cooperators, that subjects in these clusters pursue strategies with niceness similar to QFT, and that in general the best-performing strategies are those that combine niceness with the aggressive search and unforgiving vengeance of the QFT model. Thus the simple QFT model predicts experimental outcomes remarkably well, at least for the newly-forming exchange arenas with predefined trade options and a relatively restricted population represented in the experiment.

Note that the clustering of cooperation in the QFT model is dependent only on the ability of nice strategies to find and maintain cooperative relationships while rejecting defectors; selection creates clusters of QFT, while clustering does not (by definition) enhance cooperation. Furthermore, clustering will decrease with the ratio of available QFT strategies to the optimum number of partners sought; in sufficiently large populations a QFT player may find the optimum number of partners with no clustering, so no partners would know the other partners. Of course, if partners can share recommendations as in the local experimental condition, clustering is likely to emerge even in large populations. In short, the structure of cooperative exchanges will be determined by the nature of the selection process. A dynamic network model (Jackson 2009)

that specifies the information available about potential partners, the selection strategies available to players, and the distribution of QFT and other strategy types could be developed to analyze patterns of pairwise stable outcomes that could then be compared with observed patterns of high-value exchange under the specified conditions.

In newly-emerging voluntary dilemma exchange arenas, the experimental results suggest that the levels and patterns of high value, risky exchanges will be most affected by selection effects among preexisting strategy types, at least in the short run represented in the experiment. To the extent that potential partners transfer at least some QFT-like strategies learned in other arenas, the emergence of a cooperative community and Nash-ville is likely to be observed, particularly in smaller exchange arenas.

The results challenge the relevance rather than the truth of the established argument that clusters of exchange relationships positively influence the likelihood of cooperation in involuntary dilemmas (Coleman 1988, Granovetter 1985). Once cooperative clusters emerge in the experiment, they could potentially encourage greater niceness among Nash-ville subjects and eventually absorb reformed DNQ players into the higher-value exchange communities. Indeed, we can find instances of initially nasty subjects who successfully join the cooperators community.

Unfortunately, the increasingly-nasty environment produced by the QFT selection process makes this transformation very difficult within the short time frame represented by the experiment. By the time that nasty subjects realize the instability even of mutual defection and adapt a nicer approach they are increasingly likely to encounter exploitation, sometimes even by those in the cooperators community already maintaining a full set of mutually cooperative exchanges. In richer informational contexts and longer time frames, perhaps all low-value

traders will eventually understand and seek the advantages of high-value trade within the cooperators community. Even then, the problem of coordinating on higher paying QFT strategies imposes the very difficult problem of achieving pareto superior outcomes in weak-link games (e.g., Brandts, Cooper and Fatas, 2007). Perhaps the most critical remaining challenge in understanding cooperation in voluntary dilemmas is to explain the conditions under which these Nash-ville subjects can transform themselves in order to enhance the benefits of exchange within new cooperative communities.

References

- Ahn, T.K, Justin Esarey and John T. Scholz. 2009. "Reputation and Cooperation in Voluntary Exchanges: Comparing Central and Local Mechanisms." *Journal of Politics* 71(2): 398-413.
- Akerlof, George. 1970. "The Market for Lemons: Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 84(3): 488-500.
- Arend, Richard and Darryl A. Seale. 2005. "Modeling Alliance Activity: An Iterated Prisoners' Dilemma with Exit Option." *Strategic Management Journal* 26: 1057-1074.
- Axelrod, Robert. 1981. "The Emergence of Cooperation among Egoists." *American Political Science Review* 75: 306-318.
- Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books.
- Bendor, Jon and D. Mookherjee. 1987. "Institutional structure and the logic of ongoing collective action." *American Political Science Review* 81(1):129-154.
- Bendor, Jonathan and Piotr Swistak. 1997. "The Evolutionary Stability of Cooperation." *American Political Science Review* 91:290-307.
- Biele, Guido, Jo"rg Rieskamp, and Uwe Czienskowski. 2008. "Explaining cooperation in groups: Testing models of reciprocity and learning." *Organizational Behavior and Human Decision Processes* 106: 89-105.
- Boone, R. Thomas and Michael W. Macy. 1999. "Unlocking the Doors of the Prisoners' Dilemma: Dependence, Selectivity and Cooperation." *Social Psychology Quarterly*. 62(1): 32-52.
- Borgatti, S.P., M.G. Everett, and L.C. Freeman. 2002. *UCINET for Windows: Software for Social Network Analysis*. Harvard: Analytic Technologies.
- Brandts, Jordi, David J Cooper and Enrique Fatas. 2007. "Leadership and overcoming coordination failure with asymmetric costs." *Experimental Economics* 10(3): 269-284.
- Coleman, James. 1987. "Norms as Social Capital." in Gerard Radnitzky and Peter Bernholz (eds), *Economic Imperialism*, pp. 133-155. New York: Paragon.
- Coleman, James. 1988. "Social Capital in the Creation of Human Capital." *American Journal of Sociology* 94: S95-S120.
- Coricelli, Giorgio, Dietmar Fehr, Gerlinde Fellner. 2004. "Partner Selection in Public Goods Experiments." *The Journal of Conflict Resolution* 48(3): 356-378.
- Dugatkin, Lee Alan and David Sloan Wilson. 1991. "Rover: a strategy for exploiting cooperators in a patchy environment." *American Naturalist* 138: 687-701.

- Feiock, Richard and John T. Scholz. Eds. 2009. *Self-Organizing Federalism: Collaborative Mechanisms to Mitigate Institutional Collective Action*. New York: Cambridge University Press.
- Fudenberg, Drew and Lorens A. Imhof. 2006. "Imitation Processes with Small Mutations." *Journal of Economic Theory* 131: 251-262.
- Goyal, Sanjeev. 2007. *Connections: An Introduction to the Economics of Networks*. Princeton University Press.
- Gulati, Ranjay and Martin Gargiulo. 1999. "Where Do Interorganizational Networks Come From?" *The American Journal of Sociology* 104(5): 1439-1493.
- Hanaki, Nobuyuki, Alexander Peterhansl, Peter S. Dodds and Duncan J. Watts. 2007. "Cooperation in Evolving Social Networks." *Management Science* 53(7):1036-1050.
- Hayashi, Nahoko and Tosio Yamagishi. 1998. "Selective Play: Choosing Partners in an Uncertain World." *Journal of Personality and Social Psychology Review* 2(4): 276-289.
- Hirshleifer, D., & Rasmusen, E. 1989. Cooperation in a Repeated Prisoners' Dilemma with Ostracism. *Journal of Economic Behavior and Organization* 12: 87-106.
- Hruschka, Daniel J. and Joseph Henrich. 2006. "Friendship, cliquishness, and the emergence of cooperation." *Journal of Theoretical Biology*. 239(1): 1-15.
- Jackson, Matthew O. 2008. *Economic and Social Networks*. Princeton: Princeton University Press.
- Kollock, Peter. 1994. "The Emergence of Exchange Structures: An Experimental Study of Uncertainty, Commitment, and Trust." *The American Journal of Sociology* 100(2): 313-345.
- Miller, Gary. 1992. *Managerial Dilemmas: The Political Economy of Hierarchy*. New York: Cambridge University Press.
- Nowak, Martin A. 2006. *Evolutionary Dynamics: Exploring the Equations of Life*. Cambridge, MA: Belknap/Harvard University Press.
- Orbell, John and Robin M. Dawes. 1991. "A Cognitive Miser Theory of Cooperators' Advantage." *American Political Science Review* 85: 515-528
- Putnam, Robert (with Robert Leonardi and Raffaella Nanetti). 1993. *Making Democracy Work*. Princeton, NJ: Princeton University Press.

- Scholz, John T. and Cheng-Lung Wang. 2009. "Learning to Cooperate: Learning Networks and the Problem of Altruism." *American Journal of Political Science* 53(3): 572-587.
- Scholz, John T. Ramiro Berardo and Brad Kile. 2008. "Do Networks Solve Collective Action Problems? Credibility, Search, and Collaboration." *Journal of Politics* 70(2):393-406.
- Seale Darryl A., Richard J. Arend and Steven Phelan. 2006. "Modeling alliance activity: Opportunity cost effects and manipulations in an iterated prisoners' dilemma with exit option." *Organizational Behavior and Human Decision Processes* 100:60-75.
- Skyrms, Brian, and Robin Pemantle. 2000. "A Dynamic Model of Social Network Formation." *PNAS* 97(16): 9340-9346.
- Takahashi, Nobuyuki. 2000. "The Emergence of Generalized Exchange." *The American Journal of Sociology* 105(4): 1105-1134.
- Williamson, Oliver. E. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. Free Press: New York, NY.
- Williamson, Oliver. E. 1985. *The Economic Institutions of Capitalism*. Free Press: New York, NY.

Table 1. Cooperative Community and High Earners Play Quit-for-Tat

Type of Subject	DIMENSIONS OF QFT (probabilities)						EARNING AND MUTUAL COOPERATION	
	Niceness		Forgiveness		Search	Total Earning *	CC in period 5**	
	C after no play	Propose after <i>c</i>	C after <i>c</i>	Propose after <i>d</i>	C after <i>d</i>			Propose after no play
-----	-----	-----	-----	-----	-----	-----	-----	
EXPECTATIONS FOR STRATEGIES								
QFT	1	1	1	0	0	high		10
DNQ	0	1	0	1	0	low		0
RD	0	1	0	0	0	high		0
COOPERATOR COMMUNITY								
Baseline	0.66	0.93	0.90	0.53	0.34	0.43	3969	2.12
Local	0.73	0.96	0.92	0.63	0.29	0.44	5116	3.47
Central	0.47	0.90	0.79	0.61	0.19	0.41	3219	1.19
NASHVILLE								
Baseline	0.30	0.86	0.55	0.76	0.11	0.44	2757	0.74
Local	0.30	0.81	0.36	0.73	0.14	0.45	2935	0.46
Central	0.25	0.90	0.47	0.77	0.16	0.42	2389	0.41
HIGHEST EARNER IN								
Baseline	1	1	1	0	N/A	0.70	6934	7
Local	0.82	0.95	1	0.17	0	0.56	6631	5
Central	1	1	1	0	N/A	0.60	6443	6
TOP QUARTILE EARNERS***								
	0.80	0.95	0.96	0.46	0.32	0.43	5132	3.64
BOTTOM QUARTILE EARNERS								
	0.24	0.79	0.41	0.65	0.20	0.38	2123	0.26

Notes: Probabilities to propose a link and to cooperate are calculated for period 1 to 5, total earnings are for the entire 20 periods, and mutual cooperation links are the average per subject in period 5. Probability to cooperate is conditional on the game being played, so includes only decisions in the set of games played in the period being evaluated. Thus probability to cooperate after no link includes total number of cooperate decisions divided by the total number of games played with alters who did not play with subject in the previous period.

* Treatment averages: 3298(baseline), 4182(local) and 2789(central)

** Treatment averages: 1.6 (baseline), 2.2(local) and 0.8(central)

*** Quartile earners for each treatment (thus, the top and bottom 14 subjects in each treatment)

Table 2
Impact of Niceness, Forgiveness and Search on Experimental Payoffs
(Dependent variable: total payoffs, in 1000 ECU units)

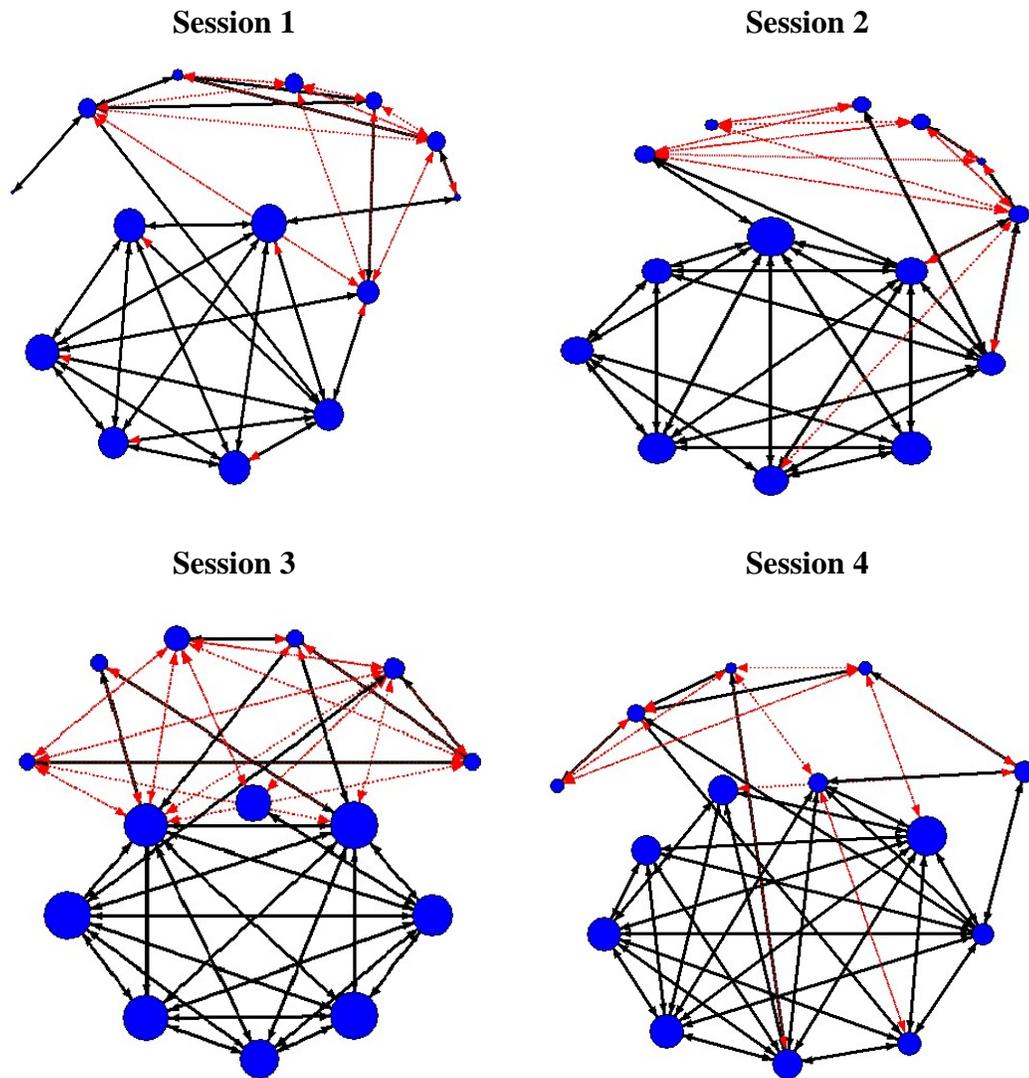
	Strategic Dimensions			QFT5050	
	Coefficients	Robust Standard Error		Coefficients	Robust Standard Error
<i>Strategic Dimensions</i>					
Niceness	2.44	0.46		1.78	0.34
Search	0.98	0.43		1.23	0.32
Forgiveness	0.50	0.22		0.08	0.29
<i>Interactions</i>					
Niceness x Search	1.28	0.95		--	
Niceness x Forgiveness	-1.69	0.46		--	
QFT5050	--			0.10	0.43
QFT5050 x Search	--			2.12	0.60
QFT5050 x Forgive	--			-2.19	0.54
<i>Environment</i>					
Proportion Nice Partners	1.16	0.32		1.16	0.26
Baseline	-0.54	0.16		-0.53	0.13
Central	-0.68	0.35		-0.71	0.30
QFT5050x Baseline	--			-0.19	0.49
QFT5050 x Central	--			0.21	0.28
Constant	1.46	0.38		1.66	0.28

Notes: Regression analyses are clustered by experimental session for 165 subjects. Dependent variable is total payoff for session in 1000 ECUs, where 400 ECUs = \$1(check). Bold entries are significant at $p < .05$

Figure 1

**The Cooperators' Community and Nash-Ville:
Cooperate and Defect Choices in Period 17**

(Cooperate=solid black, Defect=dotted red)



Notes: These sociograms are produced with the Netdraw program from UCINET. Circles represent the 14 subjects in each session, with the diameter representing the subject's total earnings. Arrows represent decisions to cooperate (solid black) or to defect (dotted red), with arrowheads pointing away from decisionmaker toward the target of the choice. The circles have been placed manually to emphasize the clustering of the cooperators' community.

Figure 2. Comparison of Proportion of Cooperation by Condition and Community

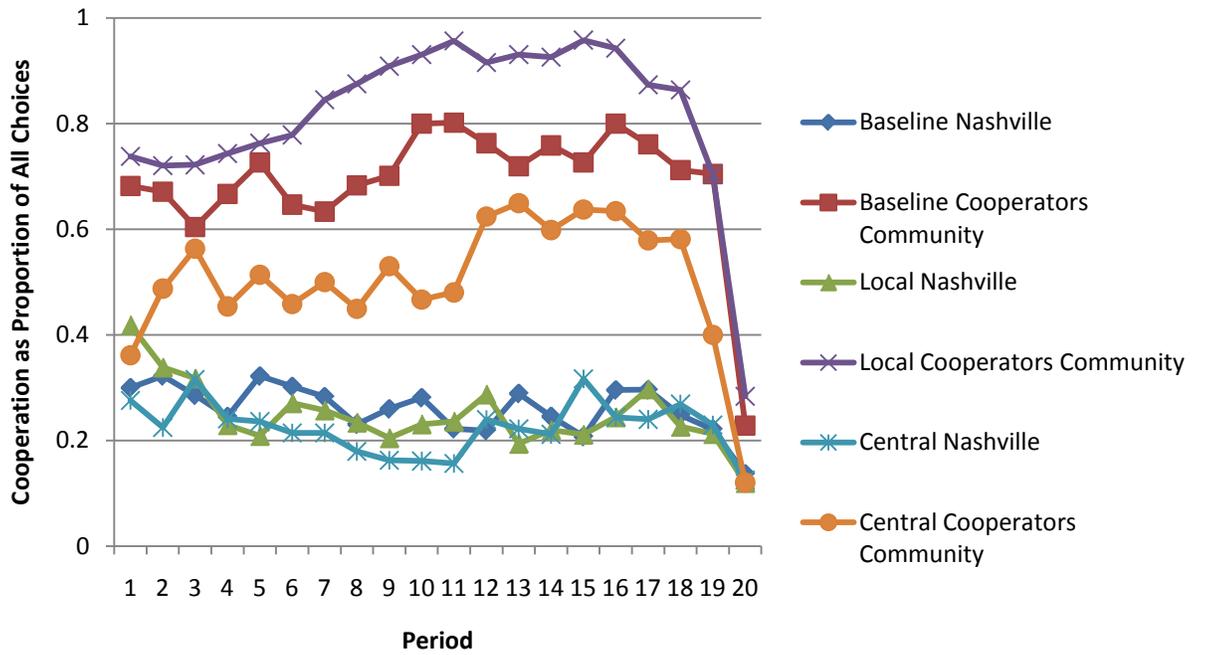
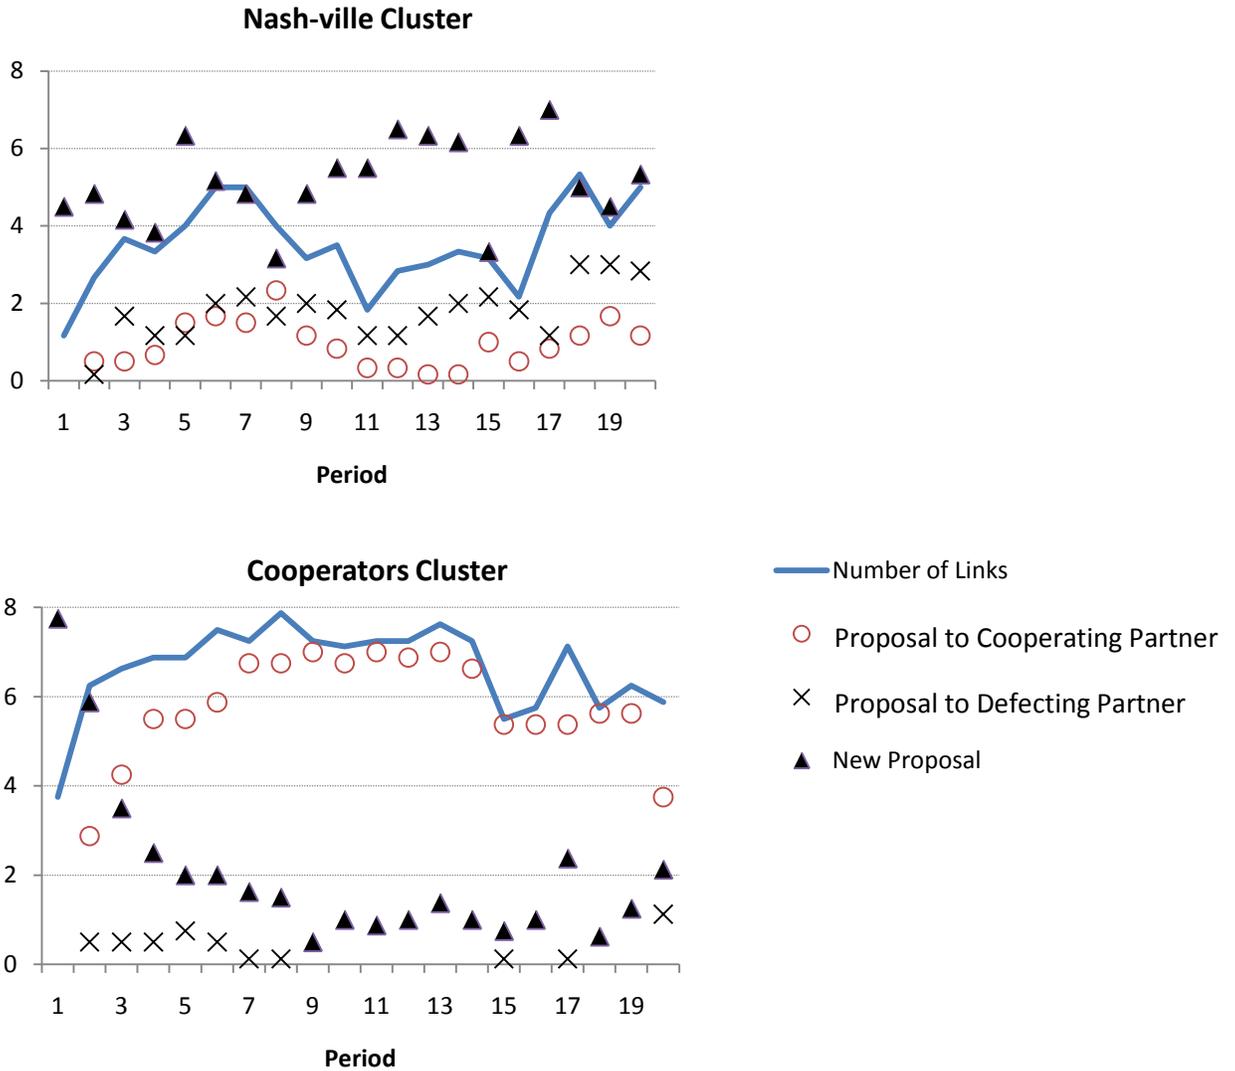


Figure 3

Comparison of Proposal Patterns for Nash-ville and Cooperators Cluster.

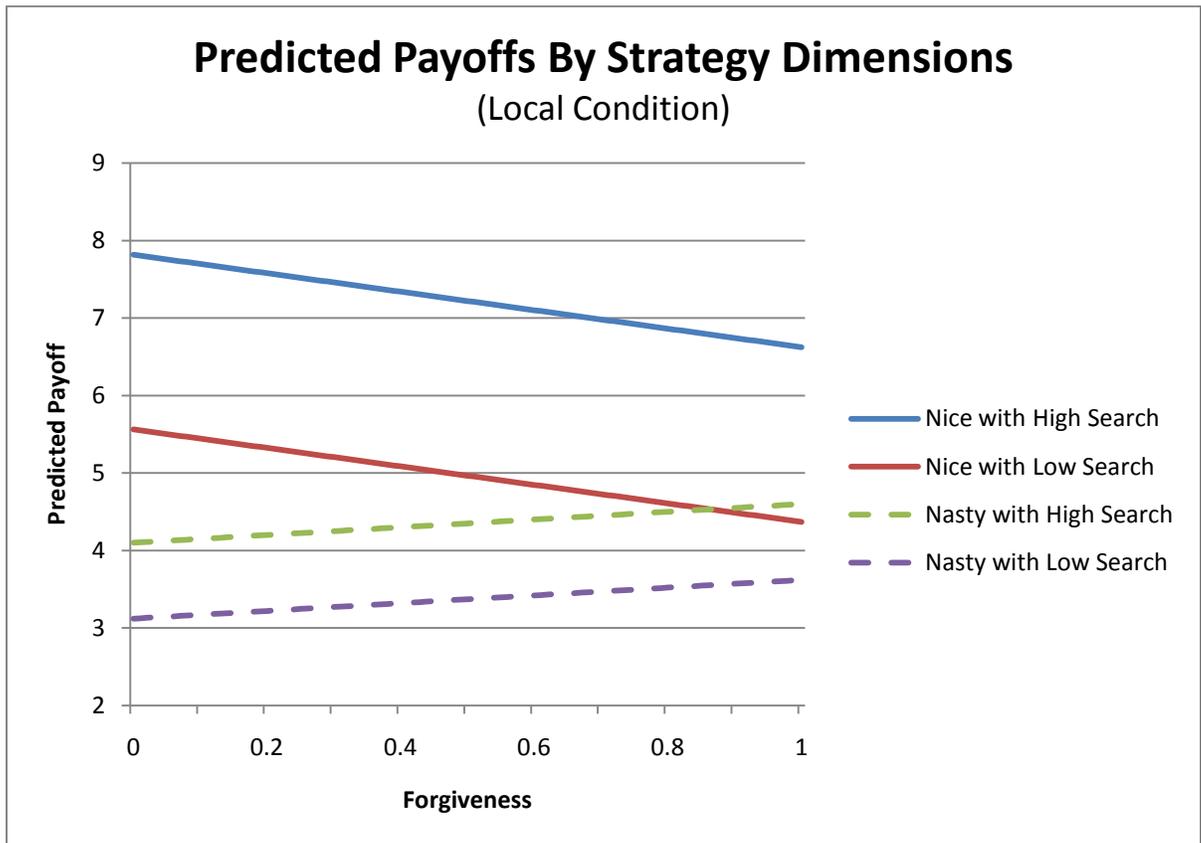


Note:

The numbers of proposals per member of the cluster are indicated by triangles for proposals to new partners, circles for proposals to partners who cooperated and “x” for partners who defected in the previous period. The resultant links per member are indicated by the line.

As explained in the text, each session is divided into a cooperators cluster and Nash-ville using TABU optimization procedure in UCINET, returning a total of 8 clusters for each treatment. The highest scoring cooperators cluster belongs to session 3 and has 8 subjects. The lowest scoring Nash-ville has 6 members and belongs to session 2.

Figure 4



Note: Predictions are calculated for the Local Condition based on coefficients in the first column of Table 2, with Niceness = 1 for Nice and 0 for nasty strategies and search = 1 for high search and 0 for low search, and Proportion Nice Partners = .5068, the mean value across all sessions