

Organizations and the Evolution of Cooperation

Danielle Jung and David A. Lake

Department of Political Science
University of California, San Diego
La Jolla, CA 92093-0521

dfjung@ucsd.edu

dlake@ucsd.edu

Draft 1.9.2, August 25, 2010

Comments welcome; please do not cite without permission.

Prepared for the Annual Meeting of the American Political Science Association,
Washington, DC, September 1-5, 2010

Organizations and the Evolution of Cooperation¹

The Prisoner's Dilemma (PD) is a metaphor for the difficulties of achieving cooperation in social life. Despite the mutual gains from cooperation, each actor in the one-shot game has a dominant strategy of defecting, leading both players into a Pareto suboptimal outcome. Even if they recognize the gains from mutual cooperation, their best strategy is to seek to exploit others or to defend themselves against exploitation by their opponent. In this world, as political philosopher Thomas Hobbes (1651/1962) originally envisioned it, life is "solitary, poor, nasty, brutish, and short."

This bleak view of the prospects for cooperation was transformed in the 1980s by new understandings of the PD.² Most important, Axelrod and colleagues demonstrated that within a repeated PD (RPD) tit-for-tat (TFT) not only produced cooperation but performed better than other strategies, was stable in that it could not be "invaded" or displaced by competing strategies, and that if permitted to evolve optimal strategies would emerge that were quite similar to TFT (Axelrod 1984; 1997). Cooperation was not only possible, many inferred, but likely. Most important, TFT was a "nice" strategy that began by cooperating with others, implying that nice guys finish first. Emergent cooperation, in turn, was discovered between opposing armies on the World War I battlefield (Axelrod 1984, Chapter 4) and even between cells in the growth of cancer (Axelrod et al. 2006). The theoretical discovery of the possibility of "spontaneous order"

¹ Jung and Lake shared equally in developing the architecture of the model and writing this paper. Jung is responsible for implementing the model and the simulations.

² In addition to the work of Axelrod discussed here, the "folk theorem" demonstrated that in repeated PD any outcome could be an equilibrium and that, given a sufficiently small discount factor, sustained cooperation was possible. Flood (1958, 16) reports a personal communication from John Nash explaining that two trigger strategies are in equilibrium in the infinite version of the repeated PD game. The general result is known as the folk theorem. See Fudenberg and Maskin (1986).

affirmed long held liberal or progressive beliefs on the potential for human progress (Taylor 1976; Sugden 1989; Hardin 1999).

In both its PD and RPD variants the game is fixed and exogenous, limited to the two moves of cooperation or defection. Social actors, however, also build organizations to transform the game and, thus, the possible strategies and outcomes.³ Even without changing the underlying incentives embodied in the PD payoffs, actors by intention or even through “dumb” selection in an evolutionary environment can create organizations that facilitate cooperation. Two ideal type organizations are common in the literature: 1) networks that provide information to actors and permit them to select their own partners and 2) hierarchies that enforce cooperation through third-party punishments. These two ideal types, and their many variants, are ubiquitous in social life, found everywhere from the animal kingdom (Boehm 1999) to the economic marketplace (Williamson 1975; 1985; Rauch and Hamilton 2001), to global politics among nation-states (Kahler 2009; Lake 2009). As Hobbes himself argued, organizations – in his case, the third party enforcer of a Leviathan – can create a civil society that is superior to the state-of-nature.

In this paper, we present results from an agent-based model (ABM) of the RPD in which agents interact in a “market,” equivalent to the non-organizational interactions within the standard RPD, or have the option of interacting within networks or hierarchies. Agents may, for some variable cost, join a network of varying characteristics that provides information about other agents and allows them to select for repeat play from among those agents they have played in the past. Agents may also, for a variable tax, join a hierarchy in which defection is punished by an exogenous third-party. We focus at

³ On the more general category of institutional solutions to the PD, see Axelrod and Keohane (1986).

present on the evolution of strategy types within a population of agents and endogenous organizations.⁴ This is, to our knowledge, the first attempt to study systematically the effects of alternative organizations within a RPD framework.

We find that organizations affect the evolution of cooperation in significant ways. Given that organizations are a staple of social and political theories, this is perhaps not a surprise. More important is when and how organizations affect cooperation, conditions on which our ABM sheds new light. In relatively nice worlds, defined by a large initial share of the population comprised of TFT and always cooperative (ALLC) agents, TFTs join networks and thereby insulate themselves against the initial and often devastating defections of nasty players (agents that begin by defecting, here ALLD). This allows TFTs to survive in greater numbers in initial rounds of the game and, in turn, to sustain their share of the total population over multiple rounds of play. In relatively nasty worlds, TFTs enter hierarchies and cooperate sufficiently often to also preserve their share of the population. In moderate worlds, on the other hand, TFTs are extremely vulnerable and typically decline as a share of the population. The information advantages are not large enough to justify joining networks and the population of nasty agents is not threatening enough to drive TFTs into the hierarchy. As a consequence, TFTs are defenseless and easily exploited. These middling worlds are, surprisingly, the most dangerous for TFTs. Organizations also improve the welfare of all agents – both nasty and nice – but only in relatively nice worlds. The benefits of civil society are contingent on the characteristics of the initial population in which it emerges.

⁴ Future work will include research on the evolution of organizations within the RPD and introduce mutation at the level of individual agents as a form of uncertainty.

This paper begins with a brief review of the literature on computational models of cooperation. We then lay out our theory of organizations and how they are modeled within the ABM (section II) and describe the ABM in detail (section III). In section IV, we then present the results of our model both for non-organizational and organizational worlds and compare the results.

I. Evolution of Cooperation

Axelrod's *The Evolution of Cooperation* (1984), reproducing material in two earlier articles (Axelrod 1980; 1980), stimulated a new research program using agent-based modeling, which we build on here, to investigate the success of different strategies in the RPD, to which we seek to add. Using a round-robin computer tournament of 64 strategies submitted by others, Axelrod found that TFT scored better than its opponents and was a best-reply to itself, such that, once established, it could not be invaded or displaced by other strategies. This proved theoretically the possibility of cooperation within the RPD in the absence of central enforcement. Using a replicator dynamic (see below), Axelrod also found that TFT and other reciprocal cooperating strategies were evolutionarily superior and would, over time, grow as a proportion of the total population of strategies. This is because TFT does well against itself and other cooperative strategies, whereas defecting strategies lag as they spread and their victims die out. Axelrod later used a genetic algorithm, developed in evolutionary biology, to simulate agent learning and/or selective fitness (Axelrod 1987). The algorithm employs mutation and cross-over (similar to sexual reproduction) to evolve increasingly efficient strategies,

as defined by higher payoffs. In most simulations, the algorithm generated strategies that resemble TFT.

This research has generated a large literature and many criticisms, most of which were anticipated by Axelrod in his original presentation (see Binmore 1994; 1998; Hoffmann 2000). We do not attempt to summarize this voluminous literature here. Most generally, however, the evolutionary superiority of TFT is now understood to rest on a number of consequential modeling assumptions made by Axelrod. Three such assumptions are particularly important for our ABM and discussion below.

First, the composition of the initial population strongly conditions the prospects for cooperation. Axelrod used a round-robin tournament in which each strategy was pitted once each round against all other strategies, its twin, and a strategy that randomly cooperated or defected. The population of agents, in turn, was a product of the strategies submitted for the competition. Subsequent work, most importantly by Hirshleifer and Coll (1988), finds that a sufficient proportion of “nice” strategies in the population relative to “nasty” strategies is necessary for cooperation to spread. Since TFT can never “win” an interaction by successfully suckering an opponent (T,S), can at best “tie” by getting mutual cooperation (R,R), and can after the first interaction with any agent do no worse than mutual defection (P,P), it will never outperform the pure all defect strategy (ALLD) except in a population with sufficiently many other conditional cooperators and sufficiently few all cooperate (ALLC) agents. All of our results are presented with varying initial populations.

Second, results are sensitive to the way agents are matched for interaction. As just noted, Axelrod employed a round-robin where each strategy played all others each round.

Other models have matched agents on identity, giving each agent an option of not interacting with a given opponent (Vanberg and Congleton 1992; Stanley et al. 1994), or on spatial location wherein agents interact more intensively with their neighbors, permitting clusters of cooperation to emerge in otherwise hostile environments (Nowak and May 1993). As explained below, our model randomly matches agents each round (with the exception of selective affinity). This is generally expected to promote cooperation (Hoffmann 2000, 3.14). At the same time, this random matching of agents each round means that it takes much longer for TFTs to “learn” the strategy types of their opponents, leaving them vulnerable to elimination. This greatly weakens the evolutionary fitness of TFT and, as we shall see, forms an important incentive for such agents to join networks early in the game.

Third, the selection mechanism also matters. In his round-robin model, Axelrod (1984, 217, fn 217) used a selection model based on “the weighted average of the (cardinal) scores of a given rule with all other rules, where the weights are the number of the other rules which exist in the current generation.” In other words, the proportion of each strategy type in the next generation is the product of that strategy’s average score and its proportion in the current generation. Alternatively, Hirshleifer and Coll (1988) use an elimination contest in which defeated players are removed from the population (and the population shrinks round-by-round). Because ALLDs can successfully sucker TFTs in their first matchups, and TFTs can never do better than ALLDs in repeated play, elimination contests are biased against TFT except in relatively nice populations. As will be explained in more detail below, we model selection here as an elimination contest with replacement in which the lowest performing agents are deleted from the population and

replaced by “offspring” of the highest performing agents. These offspring are identical to their parents in all ways except their accumulated memories, rendering them naïve and vulnerable when they first enter the competition. This provides an important incentive for such new agents to join networks or hierarchies throughout the simulation.

II. Markets, Hierarchies, and Networks as Organizations

Following the extant literature, the problem of cooperation is characterized here as a two-player RPD game (see Figure 1). In organizational terms, the RPD can be thought of as a market.⁵ Although a common term, the concept of market lacks a fixed analytical definition. Once referring only to a site for trading, since the early 20th century economists have tended to use market as a synonym for exchange, and to focus on variations in market structure, including the numbers of buyers and sellers, the information available to each, and so on. Sociologists focus more on production markets (of firms or factors of production), conceived as networks of linked agents.⁶ As an organization, according to Powell (1990, 302, italics added), markets are “the paradigm of individually self-interested, noncooperative, *unconstrained* social interaction” that engage strictly anomic agents who can form only self-enforcing agreements and know only their own past interactions with each other. This view of markets as an organization strips the concept of its focus on the exchange of goods but generalizes it to a greater

⁵ Markets might also be described as anarchy, in contradistinction to hierarchy. We defer to the larger literature, however, in the use of the term “market” in this way. We have some sympathy for critics of this thin conception of markets but follow common practice. Moreover, although modern economic exchange is often embedded in larger social organizations, such as the state, this simply implies that those “markets” are actually a hybrid form of social organization. Economic exchange does occur within markets that approximate our ideal type, commonly when price accurately reflects all known information (i.e., in “competitive” markets), the costs of third party adjudication are high relative to the value of the good, contracts are incomplete, or the goods are illicit.

⁶ On economic and sociological views of markets, see Swedberg 2003, Chapter 5.

range of interactions and is fully captured in a RPD game. Without attaching any significant meaning to the term, we will, for convenience, refer to the default RPD as a market. In a static, non-evolutionary model with sufficiently nice populations – composed of agents that at least begin by cooperating with other agents -- all agents tend to interact within the market. TFTs and ALLCs (agents that cooperate in every round regardless of the opponent’s strategy) do not need the network or hierarchy for protection, while the relatively small number of ALLDs (agents that defect in every round regardless of the opponent’s strategy) prefer to remain in the market in expectation of suckering nice strategy types (See Jung and Lake 2009).

Figure 1 About Here

In the canonical definition, networks as organizations are characterized by “voluntary, reciprocal, and horizontal patterns of communication and exchange” (Keck and Sikkink 1998, 8; Podolny and Page 1998, 59).⁷ Accordingly, we model two characteristics of networks. First, networks are mechanisms for acquiring information on agents from other agents with whom an agent has cooperated in the past. Intuitively, networks allow one agent, say i , to ask a defined number of agents with whom i has previously cooperated if they have played agent j , and if so what j did (cooperate or defect). With this information, agent i can then decide whether to cooperate or defect with j . Thus, networks provide information that supplements the information i may have acquired through its own past interactions with j . The primary effect of information from the network is to prevent agents from being suckered in the first round of play with any

⁷ Defined as a set of nodes (agents) and edges (interactions), almost any set of actors can be described as a network. Social network theory, in turn, has developed a host of tools and concepts for measuring and describing the structure of such networks (see Jackson 2008). We treat networks more as governance structures, but focus on the information flows and reciprocity between agents common to both approaches. On the different conceptions and uses of networks in political science, see Kahler (2009).

new agent. Information sharing can be understood as a form of indirect reciprocity (see Nowak and Sigmund 2005). Often treated as a defining attribute of networks, this first form of reciprocity is an emergent property of the agents who tend to select themselves into networks (Powell 1990, 303; Podolny and Page 1998, 59). Only agents that possess a contingent strategy such as TFT will ever choose to join a network to gain information about others, and having joined they will play reciprocally. As the population becomes increasingly nasty – composed of agents that begin by defecting – and the value of knowing the strategy type of one’s possible partner increases, TFTs will increasingly join the network in a non-evolutionary model. As the population becomes too nasty, however, TFTs leave the network and simply assume that their partners are likely to be defectors. In addition, this informational advantage of networks is a short-lived attraction. As agents play other agents and accumulate their own histories of interaction the benefits of acquiring information from others diminishes. This effect is stronger the smaller the population and weaker the larger the population (see Jung and Lake 2009).

Second, in some conceptions, networks permit agents to engage in a more direct form of reciprocity by intentionally selecting specific agents with whom to interact, often on a repeated basis. As already noted, in the core model agents are randomly paired in any given round of the game. Yet, in the real world, agents do not necessarily interact with a uniform probability. We implement this second type of reciprocity by a variable rate of selective affinity (η) in which nature permits an agent to select for play another agent it has interacted with in the past. With selective affinity, agents of all strategy types may choose to join the network. In a non-evolutionary model with a sufficiently nasty population, TFTs and ALLCs value being able to select an agent they have cooperated

with successfully in the past. ALLDs, in turn, value being able to select an agent they have suckered successfully in the past – an effect similar to that of a schoolyard bully who picks on the same victim day after day. Thus, selective affinity both supports greater cooperation among nice strategy types and, somewhat counter-intuitively, facilitates the exploitation of nice types by nasty types. The ability to engage in this more direct form of reciprocity is not necessarily supportive of greater cooperation (Jung and Lake 2009).

Participating in a network is always costly, however, represented in the model as a variable fee (ϕ) subtracted from the agent's payoffs no matter the outcome of the interaction. This fee is intended to capture the transaction costs of networking, variously interpreted as the opportunity costs of providing information, engaging in activities intended to develop social capital, and sending costly signals of commitment to the group necessary to establish trust or reputation. In a straightforward way, the greater the network fee the less likely agents of all types are to join the network. An agent may join a network and gain information about or select its partner even if that other agent chooses a market or hierarchy during its turn of the game. In such a case, the networked agent plays with the information acquired from past cooperators, but the other agent plays using only its private knowledge.

Third party enforcement stands at the core of all definitions of hierarchy. In our model, agents within the hierarchy cooperate with one another subject to punishments for (random) defection.⁸ If an agent defects, it receives the temptation (T) payoff minus the

⁸ That is, agents who join a hierarchy, regardless of their strategy type (see below), play a mixed strategy in which they cooperate with other agents in the hierarchy with some exogenous but commonly known probability (q).

punishment, while the other receives the sucker's payoff (S).⁹ We treat both the probability of cooperation within the hierarchy (q) and the magnitude of the punishment (v) as exogenous, but variable within the model. Our intuitive analogy is to agents working in a corporation and tasked to cooperate with their fellow employees, but cooperation within the firm is contingent on factors beyond the agent's control – including the state of the macroeconomy, fickle consumer tastes, a capricious boss, and so on. Some portion of the time, the agent's best efforts to cooperate may nonetheless appear to be a defection for which it is punished. This intuition extends to families, clans, religious orders, and more hierarchies in which individuals are mandated to cooperate (uphold contracts) with one another and are punished by a central enforcer if they defect. It also extends to states --both democratic and autocratic, local and national -- in which law regulates the behavior of individuals in relations with one another (cooperate, observe contracts, follow established conventions, etc.) under threat of (imperfect) monitoring and sanctioning. Although random defection at an exogenously defined probability is somewhat crude, some such mechanism is necessary to prevent hierarchy from dominating all other organizational forms.¹⁰ This representation allows us to investigate how the probability of defection and levels of punishment affect the expected utility of cooperation under hierarchy.

⁹ When both agents in a hierarchy defect simultaneously they each receive the DD payoff minus the punishment. With our default settings in the ABM, mutual defection is typically rare but remains a possibility.

¹⁰ Ideally, one might want to endogenize defection by strategy type and levels of punishment necessary to sustain cooperation. But if so, the punishment could always be set at a level to induce cooperation by the least cooperative agent, and mutual cooperation would always occur. If agents always cooperate in hierarchy, this form will always dominate other organizational forms, which is neither true in the real world nor theoretically interesting. One might also prefer that exogenous shocks be allowed to occur in markets and networks as well. This is reasonable. In this case, we can easily interpret the exogenous probability of defection as the *difference* between exogenous shocks in markets or networks versus hierarchies.

Again, in a relatively straightforward way, the higher the probability of cooperation within the hierarchy and the smaller the punishment, the more likely all agents are to join the hierarchy. More importantly, in a non-evolutionary model, the nastier the population the more likely agents of all strategy types are to join the hierarchy. Somewhat counter-intuitively, however, nicer agents (ALLCs) will join first, followed by contingent but nice strategy types (TFTs), and then nasty strategy types (ALLDs). Nicer agents flee to the hierarchy in the hopes of avoiding exploitation by nasty types in the market or network. Although it might seem that nasty agents would benefit most from being able to avoid the sub-optimal outcome of mutual defection, they actually benefit from being able to sucker nice types in the market or network and, thus, remain outside the hierarchy longer than other types of agents (Jung and Lake 2009).

We include a variable tax on members joining a hierarchy (τ), subtracted from the expected utility of joining the hierarchy. As with network fees, the greater this tax the less likely agents will be to join the hierarchy. Agents in the hierarchy who interact with agents outside the same hierarchy play as in the market. In a firm, for instance, some portion of any individual's daily interactions are with other employees of the same organization (e.g., as part of a team producing a new widget), but many others are with other actors outside the corporation (e.g., other firms – the local grocer, friends and families). Similarly, individuals governed by one authority, such as a state with a distinct set of laws, may interact both with one another and more or less frequently with foreigners in a second state with different laws. Cooperation is mandated and subject to centralized enforcement only with other members of one's own hierarchy or, in this case,

state. In other words, centralized punishment for defection does not apply extraterritorially or beyond the members of the same hierarchy.

In our model, agents choose freely each round to join the organization that promises the highest expected payoffs to the game, given updated beliefs. For most social organizations this is a reasonable approximation. Individuals choose whether to ask associates about the reliability and political views of potential partners and to work for one corporation or participate in one civic association rather than another. Participation in other social organizations, especially some hierarchies like the state, is less purely voluntaristic. Individuals are “born” into a state, though they may choose to immigrate at more or less cost. Young boys may be forced to join militias, and can escape only at greater or lesser personal risk. Such presumed or forced memberships are admittedly not captured well in our model. One must be careful in generalizing our results to non-voluntary organizations. Even here, however, the model helps identify conditions under which individuals and, in turn, the population (or significant portions of a population) would choose to subordinate themselves to a hierarchy and, in so doing, collectively empower the hierarch to enforce cooperation on reluctant others.¹¹ Conversely, the exit of all agents from a hierarchy approximates the loss of popular support for a political regime. Finally, as endogenous products of the choices of many independent agents, organizations are created anew each round of play. Which agents constitute a hierarchy or network is established by their choices, which may differ by round.

Our ideal types and the model in general cannot capture all aspects of all interactions in all real world social organizations. We emphasize generality, but this

¹¹ On the collective nature of authority and hierarchy, see Lake (2009a, Chapter 1).

inevitably carries some cost in understanding specific organizations and individual choices. Nonetheless, given the basic character of markets, hierarchies, and networks, their ubiquitous presence in the real world, and their similar treatment across very different academic literatures, we believe the model – even or perhaps especially in its highly simplified form -- has broad applicability.

III. Modeling Organizational Ecologies in an Agent-Based Framework

We describe the ABM here in its four stages: initialization, learning, organizational choice and, selection and evolution. A non-evolutionary version of the model that is identical in all particulars except for the selection mechanism described below, along with the expected utility equations for each organization and full parameter scans for each variable, are detailed in Jung and Lake (2009) and its online appendix.¹²

Initialization

The model begins with the specification of 25 user-defined parameters. These parameters and their default values, used in all the simulations presented below unless otherwise specified, are listed in Table 1. The default values for the parameters are admittedly arbitrary but are calibrated to make all organizational forms somewhat likely in any given simulation. By setting certain parameters higher or lower than our defaults, it would be trivial to simulate worlds in which either markets, hierarchies or networks always predominate or never arise. Instead, our defaults are set relative to one another at

¹² One significant difference between that paper and this is that we turn the weight on preferences to zero in all simulations here, thus deleting the effects of political preferences. This makes the results in this paper more directly comparable to the rest of the literature on cooperation within the RPD.

levels such that reasonable changes in any single parameter are likely to lead at least some agents to alter their organizational choices.

Table 1 About Here

Payoffs for the various outcomes are set: T, R, P, and S. We set the default cardinal payoffs in the RPD game as in Axelrod (1984) for purposes of comparability. All other default parameters are then set relative to these default payoffs. The user defines the population of actors, specifically the distribution of strategy types. As implied above, and following Hirshleifer and Coll (1988), we focus on three basic strategies: ALLC, ALLD, and TFT. As also implied above, we refer to nice and nasty populations as defined by the relative proportions of ALLC and TFT agents, on the one hand, and ALLD, on the other.

The organizational parameters are also set at this stage. Networks are defined by their width (α), the number of other agents each agent can directly ask about the agent it has been randomly paired with, and their depth (l), the number of levels of agents that are polled. Although each agent has a potentially infinite memory of its own interactions with each other agent in the population, the network is limited to a fixed memory (m_n) defined by the number of previous rounds over which it can poll. That is, if memory is set at five, any agent can poll only those agents with whom it has cooperated in the last five rounds whether they have interacted with the other agent with whom it has been randomly paired in the current round. The longer the memory (the larger is m_n) for the network, the more useful information it returns to the agent.¹³ Selective affinity is defined by the probability

¹³ In allowing an agent's own memory of past play and the network's "memory" to differ, we are essentially assuming that an individual's memory of others lasts longer than that individual's social interactions. This seems reasonable. Those of us who hold grudges and have only fleeting friendships

an agent gets to select an agent (η) from its memory (of length m_a) with whom to interact, with one minus this probability being the rate at which that agent will be randomly paired with another agent as in the base model ($1-\eta$). The fee for joining the network (ϕ) is also set.

A hierarchy is defined by the probability that any agent will cooperate with other agents in the hierarchy (q), the penalty that is imposed on agents for defecting on other agents in the hierarchy (v), and the tax assessed on members (τ). These parameters are common knowledge. Since the expected utility for joining the hierarchy is contingent on the number of other agents in the hierarchy (θ), in the first round of organizational play the user sets an “advertised” number of agents in the hierarchy, which need not be the same as the actual number of agents who join. In subsequent rounds, agents know the actual number of agents who joined the hierarchy in the previous round.

Evolution parameters are also defined by the user at this stage. The number of agents selected for replacement is set. As mentioned above, these low performing agents are replaced by the highest performing agents in the population. The agents selected out become offspring of the higher performing “parents.” The method of establishing the selection is also set, either using cumulative average payoffs (the method used in this paper), or by selecting on the most recent z rounds. The characteristics passed from parent to offspring are also set. In all situations, offspring will enter the simulation with no memory of any past interactions. They also inherit their parent’s strategy type (subject

typically remember others who have treated us badly in the past longer than we engage in sets of social relationships. This assumption is consequential only for the transient nature of networks discussed below. If agent memory were limited to the same as the network memory, networks would remain more robust over more rounds of the game. Conversely, without this restriction on network memory, the network would return “too much” information in early rounds and become obsolete almost immediately.

to some defined boundary of mutation) and their parent's belief about the nastiness of the population.¹⁴

Learning

Agents begin the simulation without any knowledge of the distribution of the other agents' strategies. In both the learning and organizational phases, agents are randomly paired with other agents with whom they then play a round of the game according to their strategy type with payoffs as specified (except in selective affinity, see below). Each agent is always an "A" agent that is then randomly paired with a "B" agent. In any particular round, an agent may not be selected as a B player or may be selected multiple times by different agents, but will on average serve as B once each round. For purposes of selection, payoffs are averaged over all interactions each round.¹⁵

Agents learn about the distribution of other strategy types from their interactions with other agents. Observing their own payoffs, they then back out whether the other agent cooperated or defected, store this action in memory by agent, and update a running estimate of the proportion of cooperators and defectors in the population (β_i). From this, agents learn whether the environment is relatively nice or nasty. Importantly, agents observe only the other's actions, limited to cooperation or defection, and not their underlying types. This is equivalent to not being able to observe an individual's intent or

¹⁴ For this paper, mutation on both characteristics is set to zero, meaning strategy type and belief about the nastiness of the population are passed directly from parent to child in the current form. In later iterations of this project we anticipate introducing uncertainty into the population using these characteristics.

¹⁵ In the real world, individuals may participate in many different social organizations nearly simultaneously, sometimes with the same partners. One might, for example, gain information from a neighbor about a new job opening, a form of networking, and serve on a community organization's board with that same person. Our model does not fully capture such complex relationships, although because of the pairing of each A actor with a random B actor agents may actually interact multiple times in the same round of the game. Similarly, agents may interact by sharing information multiple times in multiple networks.

character, only what he or she actually does. Thus, each agent assigns and then subsequently updates for each agent it plays a single running probability of cooperation. In this phase of the simulation, agents are restricted to the knowledge they accumulate about other agents through direct play. Each agent develops unique beliefs over the course of play, meaning that even agents with the same strategy type will make different organizational choices in the next stage. This introduces heterogeneity of agents even within a fixed set of strategy types. Agents who believe the population is nastier than it really is are *pessimists* and agents who believe the population is nicer than in actuality are *optimists*.

Organizational Choice and Play

Once the learning period is concluded, the main simulation of interest begins and continues for a fixed number of rounds. In this phase, a round is defined by two actions: the organizational choice of each agent for that round and the actual play in that round. Agents begin each round by calculating their expected utility for joining each type of organization and select the one they calculate will yield the highest return. Agents continue to update their beliefs about the distribution of strategy types throughout play (except in hierarchy, see below). As these beliefs change, the same agents will make different organizational choices. Over many rounds, the beliefs of agents will tend to converge to one another, reducing heterogeneity in the population, and to the “true” population parameters.

The expected utility for market interactions is the same as an agent would get in play during the learning phase described above. Agents can choose to pay the cost to join the network (ϕ) of a known selective affinity (η), affinity memory (m_a), width (α) and

depth (l) of agents with whom she has a history of cooperation in the last number of rounds as defined by memory (m_n). The expected utility from the network is essentially the likelihood that the player receives information about its current partner that changes its behavior plus the likelihood it does not and the likelihood that the agent gets to select its partner from memory, less the fee imposed to join the network (ϕ). The utility for entering a hierarchy will depend on the proportion of the population in the hierarchy the player will join (θ), weighed against the likelihood of cooperation within the hierarchy (q), the punishment for defection (v), and the tax (τ).

After agents choose the organization they will join for that round, the next stage is actual play within each organization. If a player selects the market it plays its fixed strategy. For non-contingent strategy types (ALLC and ALLD), information from the network is irrelevant, since they play the same move regardless of the type of other agent. Without selective affinity, such agents never choose to join the network even at zero cost. Since only contingent strategy types (TFT) can potentially benefit from information on other agents, only these agents will consider joining the network in the absence of selective affinity. If an agent selects the network, it will query the specified past cooperators about the agent with whom it has been randomly paired and be given a number $[0,1]$ representing the probability of cooperation to expect from that partner. If that agent believes the other agent is likely to cooperate (the probability is ≥ 0.5), it will cooperate, otherwise the agent defects. The information returned from the network is treated as equivalent to the agent's own beliefs about the randomly paired agent acquired through direct play. That is, if agent i has no past play with agent j , and it receives a signal from the network that j cooperates with a probability of 0.7, it will update its belief

about j 's type to 0.7. Similarly, if i believes on the basis of a single past interaction that j cooperates 1.0 and it receives a signal from the network that j has cooperated with five networked agents at a rate of 0.7, it revises its belief about j to 0.75—weighting its own experience equally with those received from the network. In this way, we assume that all agents are sincere in their reporting and are known to be so by all other agents.¹⁶ If an agent joins the network and is given by nature the opportunity to select its own partner (η), it chooses the agent within memory (m_a) with whom it earned its highest payoff in previous rounds. If the agent chooses to join the hierarchy, its play depends on whether or not it is matched with another player in the hierarchy. If the two players belong to the same hierarchy, the agent will cooperate at the rate that the hierarchy enforces (q). If the agent defects ($1-q$), it will be punished at the defined level (v). Within the hierarchy, players do not update their beliefs about the other agent's strategy type. As cooperation within the hierarchy is induced, agents cannot tell if their partners cooperated because of their strategy type or from fear of punishment. As a result, agents cannot learn anything from the behavior of the other agent. In a world in which all agents join the hierarchy, this decision becomes irrevocable as the beliefs that led an agent to join the hierarchy in one period carry over into subsequent periods. If a player is matched with a player outside of its hierarchy, it will play as if it were interacting in the market and update its beliefs accordingly. Since nasty agents tend to remain in the market even when nice

¹⁶ This is an important assumption. If agents lie or even communicate poorly (e.g., perform the kinds of minor distortions familiar to children from the “telephone game”), networks may actually harm rather than increase utility by causing contingent players to engage in bouts of mutual punishment. This is equivalent to the problem of noise within the RPD. See Hoffman (2000). In this version of the model, we do not discard or discount redundant responses from the network. Intuitively, in real interactions we often do not know exactly where a friend of a friend received their information about some other actor. Given that the strategy types we examine here are pure, this assumption has no consequence for any of our results.

agents enter the hierarchy, nice agents will tend to form more pessimistic beliefs over time in incompletely hierarchical worlds.

Following play, real payoffs are calculated as a function of the outcome of play, punishments, and fees prescribed by their organizations. Actual payoffs can differ from expected payoffs, but are on average the same.

Selection and Evolution

As previewed in the literature review, our selection mechanism is an elimination contest with replacement in which the lowest performing agents are deleted from the population and replaced by offspring of highest performing agents. Once actual payoffs have been assigned, the lowest achieving agents are selected out of the population and the highest achieving agents beget new agents who inherit their strategy type and beliefs, but not their memories. The selection rate can vary, but is set by default at 10 percent, meaning that in a population of 100 agents, the 10 agents with the lowest actual payoffs are eliminated at the end of each round and replaced by the same number of agents that are identical in all ways except their histories of play to the 10 agents with the highest actual payoffs. The greater the selection rate, the greater the evolutionary pressure on the population and the faster the population will evolve.

IV. Results and Discussion

We are primarily interested here in the evolutionary success of particular strategies and organizations. Our strategy is to simulate organization choice and payoffs under varying populations of agents in which we create relatively nice worlds of many TFT and ALLC agents and relatively nasty worlds of more ALLD agents. Because

several parameters are randomly assigned according to specified distributions in the initialization phase, and agents are randomly paired at each round of play in both the learning and organizational phases (unless in selective affinity), no two simulations will be identical. For the results below, we replicate the simulation 10 times and report the average of the results. With high rates of random variation across rounds of the game, we also smooth the results depicted in our payoff graphs.

The Evolution of Cooperation without Organizations

We begin by simulating our model without organizations. This is important because of the several differences in the basic setup of the ABM. Figure 2 demonstrates that population dynamics are, indeed, sensitive to the initial population. The primary result here is that in our elimination contest with replacement, TFT is not an evolutionarily stable strategy in relatively nice worlds. In Figure 2, we have drawn positive and negatively sloped 45 degree lines. Any portion of the ALLD curve above the positively sloped line indicates that ALLDs formed a greater share of the population in round 1000 than in the first, and any portion of the curve below the 45 degree line indicates that they were a smaller share of the population. As seen in Figure 2, ALLD expands beyond its initial share of the population rapidly in comparatively nice populations, and then returns to its initial share in comparatively nastier populations. Similarly, any portion of the TFT curve above the negatively sloped line indicates that TFT grew as a share of the population, and any portion of the curve below the line means that it shrank as a share of the population. As can be seen, TFT fares poorly in relatively nice populations and retains its initial size in very nasty populations.

Figure 2 About Here

The dynamics of selection can be seen more clearly in representative worlds in which we depict the composition of the population and average payoffs by strategy type. In a relatively nice world of 10 percent ALLD agents, displayed in Figure 3, ALLC agents are quickly eliminated by ALLDs, who successfully sucker their opponents and then grow as a proportion of the population. In early rounds, TFTs have not yet played many ALLD agents and are also suckered on their first interactions, reflected in their plummeting numbers and extremely low payoffs. Once they have gained experience with the ALLD players, however, TFTs stabilize themselves as a proportion of the population and match payoffs, on average, with ALLDs. TFTs either cooperate with one another, earning a payoff of three, defect with known ALLDs, receiving a payoff of one, or cooperate on the first interaction and get suckered by ALLDs, thereby getting a payoff of zero. The naïve TFTs who get suckered are eliminated from the population via selection and are typically replaced by offspring of the TFTs who successfully cooperated with one another. These new TFTs, with no experience of play, then get suckered and eliminated again by ALLDs, creating a revolving “death door” at the bottom of the payoff rankings. With a relatively low selection rate of ten percent in these simulations, naïve TFTs who have not yet played against many of the ALLDs in the population accumulate more rapidly than they are eliminated and replaced. As experience builds among the remaining TFTs, who face a nearly constant population of ALLDs, the population eventually stabilizes. Overall, average payoffs for both ALLDs and TFTs converge at a relatively low level of one, the payoff for mutual defection.

Figure 3 About Here

Interestingly, virtually the same dynamic and stable population emerges in a

relatively nasty world of 75 percent ALLD agents (Figure 4). ALLCs are, again, eliminated almost immediately and replaced by new ALLDs. Without as many ALLC and TFTs to sucker, ALLDs do not enjoy the initially high payoffs they earned in the nicer world. Heavily suckered by ALLDs, TFTs begin with low average payoffs that eventually converge to one as they learn defend themselves. With a still significant number of other TFTs in the population with whom to cooperate, and a relatively low selection rate, successfully cooperating TFTs (on average, about 20 percent of the interactions) beget enough new TFTs to maintain a relatively constant population and even grow slightly as a proportion of the total population. Once again, the naïve offspring are then suckered by the ALLDs they are likely to encounter.

Figure 4 About Here

These results confirm that TFT does well against itself and holds its own over sufficient rounds against nasty strategies. Once TFT has learned about its partners, it does no worse than ALLD and, thus, sustains itself as a share of the population. Only after it has gained experience, however, is TFT evolutionarily stable. As indicated by Figure 2, however, TFTs do not generally expand their share of the population in any population.

As noted in the literature review, our random matching design means that it takes much longer for agents to acquire information on other agents than does a round-robin tournament in which the “niceness” or “nastiness” of all strategy types are known by all agents after the first round. This design makes TFT highly vulnerable to exploitation in early rounds of the game when the agents are still learning the strategy types of their opponents. This feature interacts with the initial population and, although not shown here, the size of the population to determine the evolutionary course of the simulation (the

larger the population, the longer it takes to learn about all other agents). Although we do not demonstrate this point here, the evolutionary course will also be greatly exacerbated by the severity of the selection mechanism. With the relatively low rate of selection used here, TFTs can maintain their proportion of the population by cooperating enough with one another to beget replacements for all the naïve TFTs that are eliminated each round. A lower rate of selection allows TFTs to survive at a greater proportion of the population in nicer worlds, but does not fundamentally alter the dynamics just discussed. A high enough rate of selection, on the other hand, eliminates TFTs in early rounds before they have had an opportunity to learn the strategy types of other agents and leads to a population collapse for the strategy type. With fewer others of their type to cooperate with, TFTs are no longer among the highest performing agents and cannot produce offspring fast enough to replace themselves. The population of TFTs, thus, quickly erodes.

The Evolution of Cooperation with Organizations

Introducing organizations substantially alters the results described above. Most dramatically, as summarized in Figure 5a, the ability of agents to join a network or hierarchy permits TFTs to retain their proportion of the population more often, including in relatively nice worlds in which they were severely disadvantaged without organizations. With organizations, the curve for TFTs as a share of the population line is often above the negatively-sloped 45 degree line in very nice populations and on the line in very nasty populations (compare with Figure 2). This is especially true in early rounds. Figure 5b depicts population shares in round 10 of the game. Here, TFTs are able to

retain their initial share even where without organizations they were already under assault by ALLDs (see Figure 3). This trend is sustained through round 50 of the game, illustrated in Figure 5c. As Figure 5a shows, however, there is a lot of variability in relatively nice worlds as ALLDs and TFTs battle for survival over longer runs. Perhaps most surprisingly, ALLCs enter the hierarchy and not only survive in relatively nice worlds but actually expand slightly as a proportion of the population in moderate worlds (up through an initial population of about 40 percent ALLDs). Organizations transform the results of the game quite dramatically to the benefit of nice strategies.

Figure 5 About Here

As shall be explained in more detail below, in relatively nice worlds, TFTs enter the network and use its informational mechanism to protect themselves against being suckered in the early rounds of the game. Correspondingly, without the ability to sucker TFTs early on, ALLDs do not expand as a proportion of the population (see Figure 6, which is comparable to Figure 3 without organizations). In moderate worlds, on the other hand, TFTs begin in the network, as in nicer worlds, but are not able to sustain their numbers against the progressively growing population of ALLDs in the market; as the population of TFTs declines, they move somewhat into the hierarchy but still cannot maintain ground when paired with the ALLDs in interactions that occur outside the hierarchy (see Figure 7). Once on a “losing” track despite using the network, modest numbers in the hierarchy do not allow the TFTs to recover. Most dramatically, in comparatively nasty worlds, TFTs move immediately into the hierarchy and are able to successfully defend themselves against the large number of ALLDs they face (see Figure

8). By joining the hierarchy, even in worlds of more than 70 percent ALLD agents, TFTs are able to maintain their initial shares of the population (see Figure 5).

Figures 6-8 About Here

Organizational choice varies by strategy type and initial population. In relatively nice worlds, ALLCs use the network in early rounds for the benefits of selective affinity (see Figure 9a). As they become more confident of finding other cooperating agents, they avoid the fees and taxes imposed by networks and hierarchies and mostly move into the market. More pessimistic ALLCs, who believe the world is nastier than it really is, go into the hierarchy for the protection it affords even in comparatively nice worlds. As the population becomes progressively nastier, ALLCs still begin in the network but move increasingly into the hierarchy (Figure 9b) until, in very nasty worlds, they are completely in the hierarchy (Figure 9c). Hierarchy allows some ALLCs to survive and grow as a proportion of the population in moderate worlds (see Figure 5a). In the very nastiest of worlds though, ALLCs are quickly eliminated. In worlds without organizations, ALLCs are always eliminated immediately by predatory ALLDs. In worlds with organizations, ALLCs are able to successfully cooperate under the shadow of third party enforcement with each other and other strategy types to survive at a significant rate up to some initial population limit. Hirshleifer and Coll (1988, 395) find that ALLCs are parasitic on TFTs, surviving only because they can cooperate sufficiently with the latter. In our simulations, ALLCs cannot survive without organizations. By entering the hierarchy, however, ALLCs not only survive but aid TFTs in their non-hierarchical interactions.

Figure 9 About Here

In nice worlds, ALLDs begin in disproportionately in the network to take advantage of selective affinity and its opportunity to exploit nice strategy types (Figure 10a). As TFTs acquire information and become harder to bully, ALLDs then move mostly into the market, where they can continue to exploit ALLCs and naïve TFTs, and into the hierarchy. As the world becomes nastier, ALLDs go into the hierarchy at high rates and remain there in higher proportions (Figures 10b and 10c). However, even in very nasty worlds a majority continue to remain in the market in order to exploit others.

Figure 10 About Here

In relatively nice and moderate worlds, TFTs begin in the network both for the benefits of selective affinity, which permits cooperation with known others, but more importantly for the information it provides about other agents they have not yet played (Figures 10a and b). This information benefit, however, eventually wanes as TFTs themselves come to play directly an ever larger fraction of the population. This effect is greater the nastier the population. TFTs then move into the market or hierarchy, with optimists joining the former and pessimists joining the latter. As the world becomes nastier, TFT move increasingly into the hierarchy (Figures 10b and c). Where networks provide valuable information to TFTs in early rounds of any game, and allow them to avoid being suckered as often by ALLDs, they move disproportionately into the hierarchy as the game evolves in all but the nicest worlds. By joining organizations, TFTs avoid being decimated by ALLDs in early-to-middle rounds and survive in comparatively greater numbers in relatively nice and nasty worlds.

Figure 11 About Here

Counterintuitively, moderate populations are the most dangerous for TFTs in worlds with organizations. TFTs exhibit significantly greater evolutionary fitness in nice worlds, where they can use the network for information to protect themselves against initial suckerings, and in nasty worlds, where they use the hierarchy to sustain cooperation despite the large number of ALLDs they confront. TFTs are extremely vulnerable in moderate worlds a) in which information is not enough to boost the payoffs of surviving agents or protect naïve TFTs from exploitation and b) which are not sufficiently threatening to drive them into the hierarchy. Relatively defenseless, TFTs are eliminated in disproportionate numbers in these middling worlds. As indicated by the stability of the negative population trend for TFTs (and the inverse for ALLDs) in the middle of Figure 5a, this is a robust pattern.

The welfare implications of organizations may be equally important and interesting. In worlds without organizations (see prior section), average payoffs for all actors tend to converge to one regardless of the initial population (Figures 3b and 4b). Even in nice worlds, TFTs are sufficiently exploited by ALLDs that, on average, all strategy types are receiving the payoffs for mutual defection (P). This is the Hobbesian state of nature. In worlds with organizations, this same result occurs in comparatively nasty initial populations (compare Figures 4b and 8b). In nice worlds with organizations, however, all agents on average do significantly better than without organizations. With TFTs surviving initial defections by ALLDs and sustaining themselves as a large share of the population, TFTs cooperate sufficiently with one another that net payoffs (after network fees and hierarchy taxes) are close to 2.5, more than twice as high as without organizations (compare Figures 3b and 6b). Paradoxically, even ALLDs do better with

organizations, achieving net payoffs of approximately 1.5, fifty percent higher than without organizations. In moderate worlds, TFTs do even less well than in non-organizational worlds, averaging below one, while ALLDs benefit from the opportunity to exploit the TFTs and earn average payoffs of nearly 1.5. If a world of organizations is a Hobbesian civil society, the benefits of that society are achieved only under some initial conditions – a relatively nice population – and are not equally shared by all. Civil society does not produce higher welfare all by itself but only in conjunction with relatively nice populations.

Conclusions

The rise of spontaneous order is fascinating. That cooperative actors can emerge and prove resilient against competition from uncooperative others is heartening. Axelrod and ABMs of the RPD have helped scholars understand this process. But spontaneous order is also limited. That even small changes in our ABM – different matching and tournament structures – can affect the prospects for cooperation so drastically suggests just how fragile spontaneous cooperation can be.

To compensate for that fragility, humans -- and perhaps other organisms as well -- have developed organizations to support cooperation and improve welfare. This is virtually a truism, if still an important one. Our ABM begins to explore the conditions under which organizations will be formed. Equally, if not more important, it allows us to begin to identify which agents will prefer which type of organization under what conditions and, finally, with what affects on welfare. Our results suggest few universal rules. Rather, our simulations show that similar agents in different populations will often

make different organizational choices, and that the same organizations in different populations will have different effects on cooperation. Population, choices, and organizations interact in complex and contingent ways.

Nonetheless, our model does imply some general propositions. Nice agents in relatively nice worlds benefit from acquiring information from social networks in order to identify who they can trust to cooperate and who is likely to defect. This ability to identify cooperative and uncooperative partners prevents contingent but nice strategy types, like TFTs, from being exploited in first interactions and, as a result, to sustain their shares of the total population over time. The ability to survive in initial interactions has long term ecological consequences.

In turn, nice agents in nasty worlds enter the hierarchy to protect themselves under third party enforcement. Hierarchy does not necessarily constrain the “bad guys,” as is often supposed. In our simulations, nasty strategy types play mostly in the market in all conditions in hopes of suckering nice strategy types. Rather, hierarchy is preferred by nice types as a way of ensuring sufficient cooperation with other cooperative partners. As worlds become increasingly nasty, it is the most cooperative types that join the hierarchy first (ALLCs), following by nice contingent types (TFTs). By entering the hierarchy and exposing themselves to the possibility of punishment, they raise their own welfare and ensure their reproduction. Although hierarchy does not allow nice strategy types to take over the world in our simulations, it does allow them to continue to thrive even in hostile populations.

Table 1. User defined parameters and default values.

Parameter	Symbol	Description	Default Value
<i>General</i>			
Increments		Number of times the simulation is run incrementing a parameter	20
Repetitions		Number of times the identical simulation is repeated with different random seeds	10
Rounds		Number of rounds of play	1000
Learning rounds		Set as either number of rounds or population convergence to within a proportion of the true population mean	10 rounds
Agents (Total)			100
All Cooperate		Number of actors of type always cooperate	
All Defect		Number of actors of type always defect	
TFT		Number of actors playing tit-for tat strategy	
Payoffs			
R	R	Payoff for CC outcome	3
S	S	Payoff for CD outcome	0
T	T	Payoff for DC outcome	5
P	P	Payoff for DD outcome	1
Hierarchy			
Initial size	θ	Proportion of the population in hierarchy. In first round of play, this variable is set exogenously; after the first round, this variable is endogenous and defined as the number of players in the previous round.	10
Penalty	V	Penalty for defection within the hierarchy	0.5
Probability of Cooperation	Q	Rate at which the agents cooperate with other agents in the hierarchy	0.95
Tax	τ	Tax assessed on members of the hierarchy	0.3
Network			

Cost	ϕ	Fee for joining the network	0.3
Width	α	Number of past cooperative partners each agent i can ask for information about agent j	3
Depth	L	Number of levels agent i can survey	3
Memory	m_n	How many past moves each agent remembers within the network	5
Network Affinity	η	Probability of network players being able to pick their partner	0.1
Affinity Memory	m_a	How far back affinity players can look into their memory	5
Evolution			
Number of Agents replaced		Number of lowest performing agents replaced by the same number of highest performing agents	10
Ranking Mechanism		Agents are ranked either on cumulative average payoffs or on their average payoffs for the last z rounds (also specified by user)	Cumulative
Mutation on Strategy		The range around the parent agent's strategy type from which the child's strategy type is drawn	0.0
Mutation on PC belief (β_i)		The range around the parent agent's belief about the proportion of cooperators and defectors in the population (β_i)	0.0

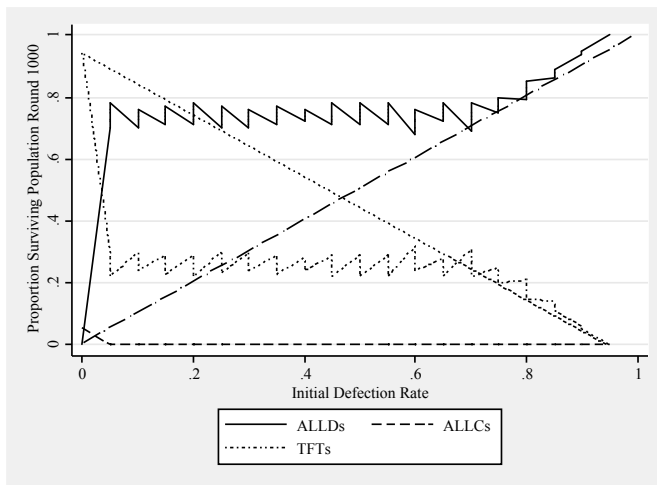
Figure 1. The prisoner's dilemma game

		Agent j	
		C	D
Agent i	C	R, R	S, T
	D	T, S	P, P

$$T > R > P > S$$

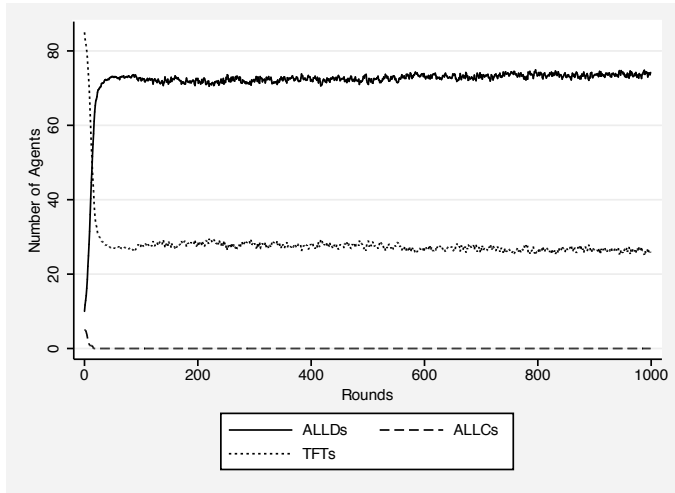
$$2R > T + S$$

Figure 2. Proportion of Each Strategy Type Surviving Until Round 1000, Given Initial Populations, Without Organizations

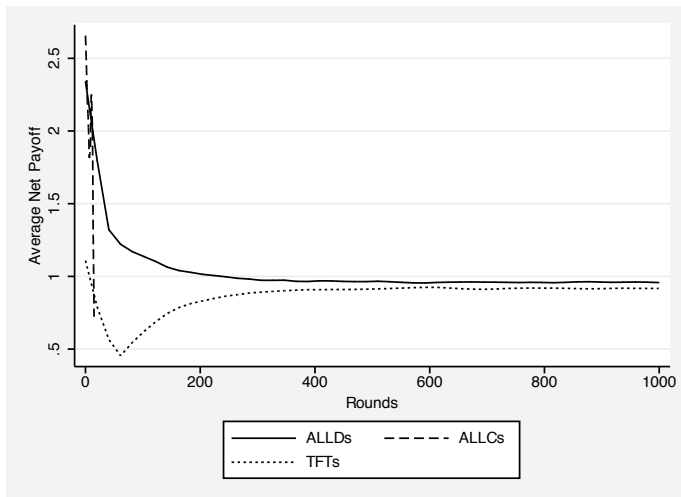


Initial population becomes “nastier” along the horizontal axis. Surviving members of each strategy type, given initial proportions, are depicted on the vertical axis. The 45 degree lines indicate constant initial and ending proportions of ALLD and TFT agents. TFT performs relatively poorly in nicer populations, falling far below the negatively-sloped 45 degree line, and ALLD performs relatively well in nicer populations, rising above the positively-sloped 45 degree line. The small number of ALLCs are eliminated even in relatively nice worlds. [Seed 436834; Population begins with 5 ALLCs and 95 TFTs, 5 TFTs are subtracted and 5 ALLDs are added with each of 20 iterations.]

Figure 3. A Relatively Nice World without Organizations

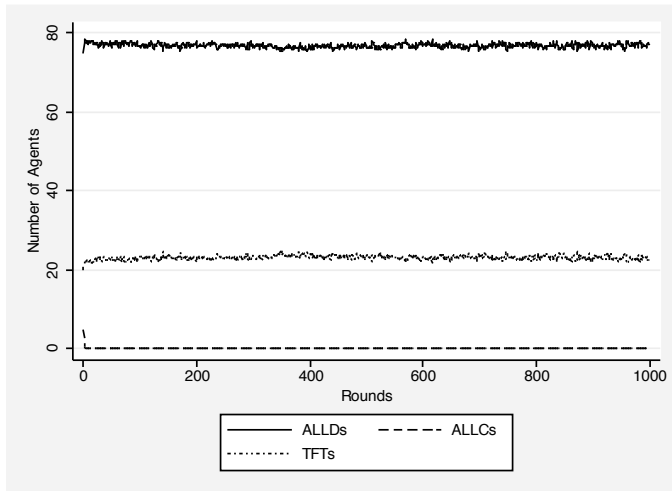


Panel a. Number of agents by strategy type surviving by round (100 total, initial population of 85 TFT, 5 ALLC, 10 ALLD). [Seed 436834]

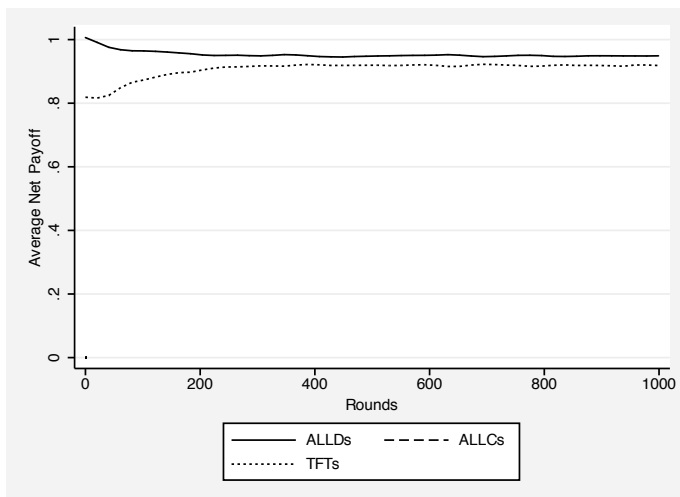


Panel b. Smoothed payoffs by strategy type by round.

Figure 4. A Nasty World without Organizations

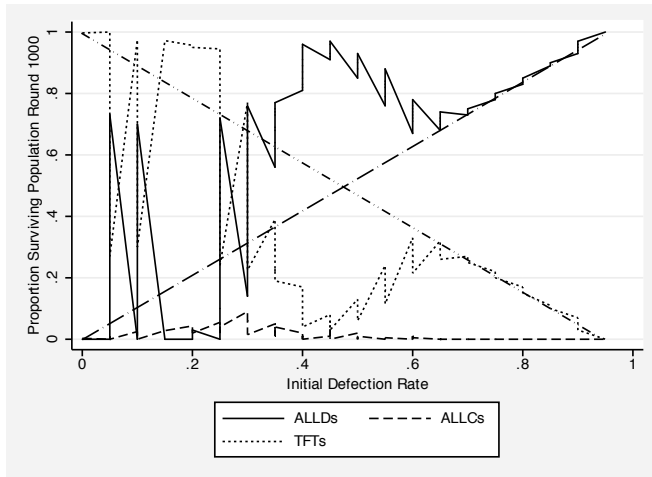


Panel a. Number of agents by strategy type surviving by round (100 total, initial population of 20 TFT, 5 ALLC, 75 ALLD). [Seed 436834]

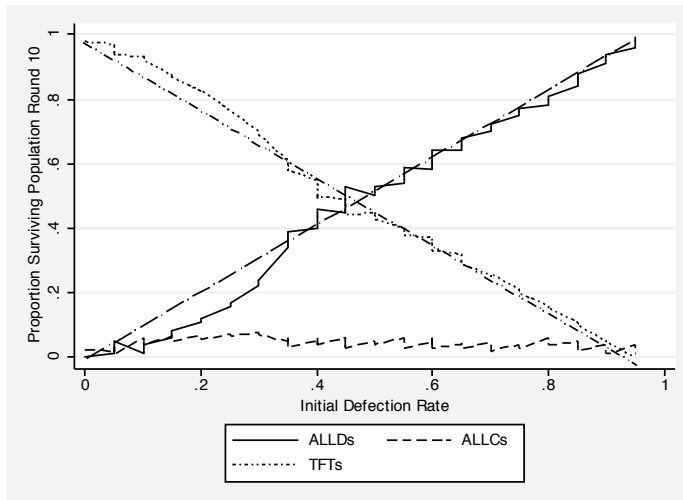


Panel b. Smoothed payoffs by strategy type by round.

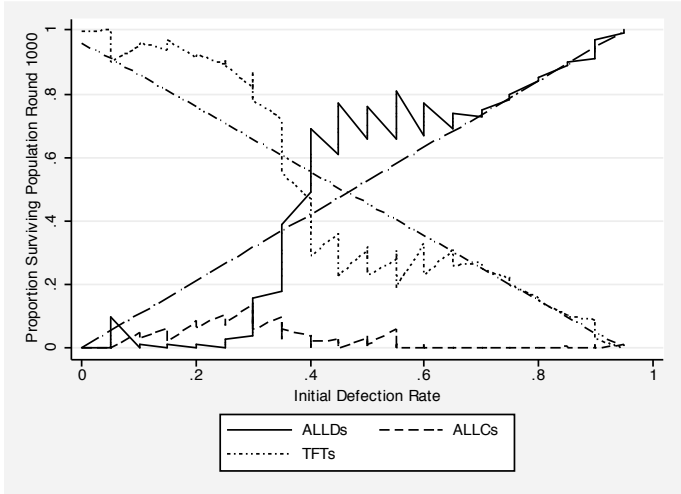
Figure 5. Proportion of Each Strategy Surviving Until Round N, Given Initial Population, with Organizations



Panel a. 1000 Rounds. This figure is comparable to Figure 2, but includes the effects of organizations. TFTs do comparatively well in relatively nice and nasty populations, although there is considerable instability in nicer worlds. TFTs fare poorly in moderate worlds. The small number of ALLCs survive and even grow as a share of the population in nice and moderate worlds. [Seed 545554; Initial population begins with 5 ALLCs and 95 TFTs; 5 TFTs are subtracted and 5 ALLDs added with each of the 20 iterations.]

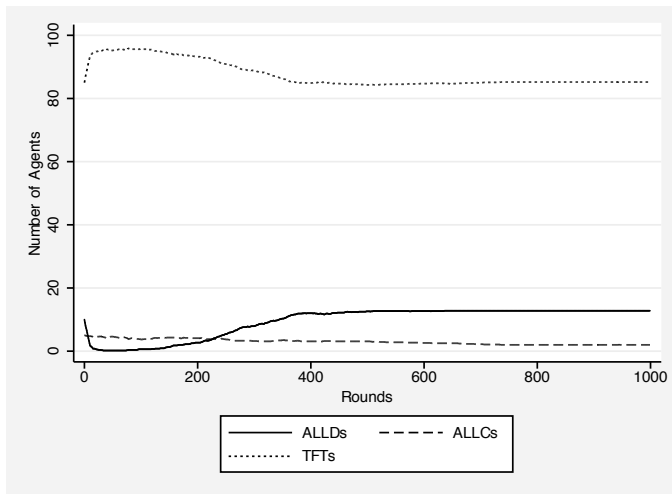


Panel b. Ten Rounds. In very early rounds, agents of all types are able to resist invasion and maintain their initial proportions.

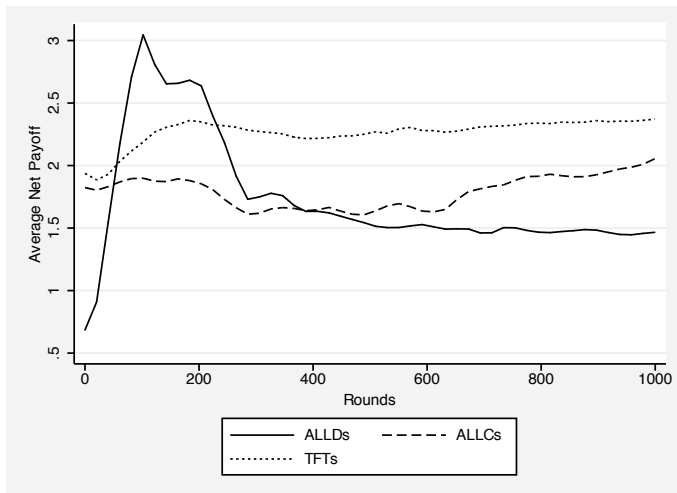


Panel c. 50 Rounds. In slightly later rounds, nice types are able to grow as a proportion of the population in nice to moderate populations. In moderate to nasty populations, TFTs are able to maintain their proportion of the population.

Figure 6. A Relatively Nice World with Organizations

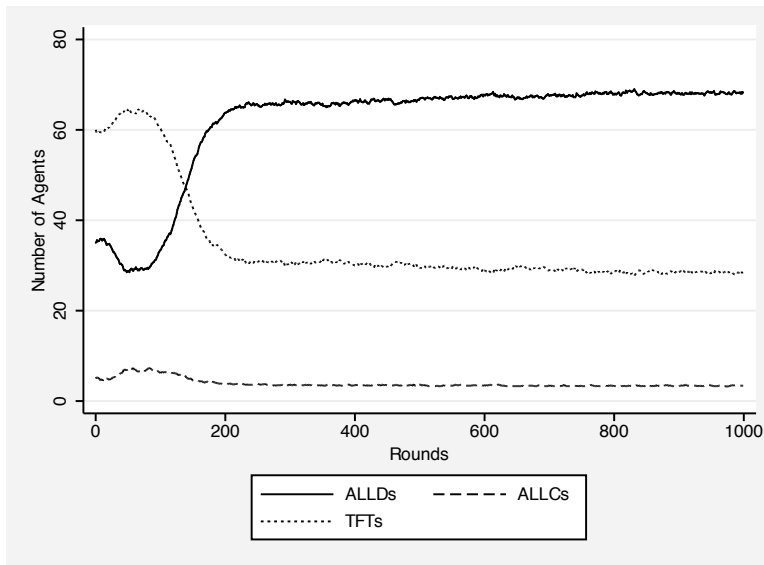


Panel a. This simulation begins with the same initial population as Figure 3. TFTs survive at higher rates early in the game and maintain this success throughout. [Seed 545554]

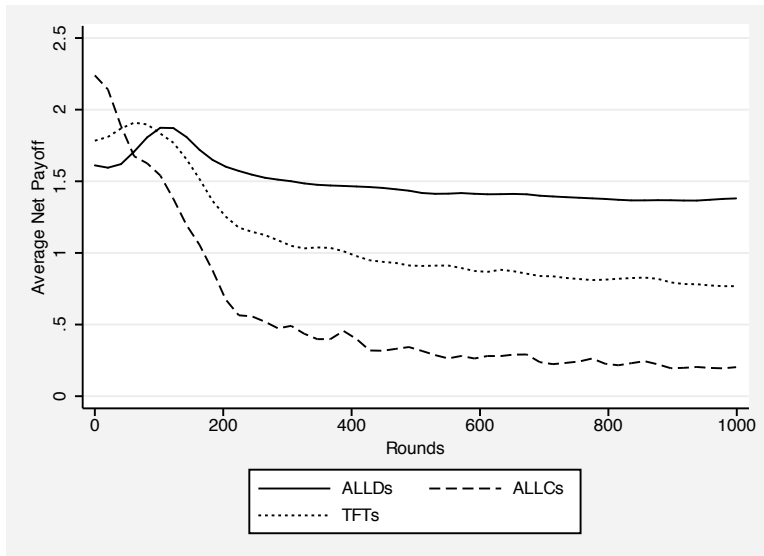


Panel b. Smoothed payoffs by strategy type

Figure 7. A Moderate World with Organizations

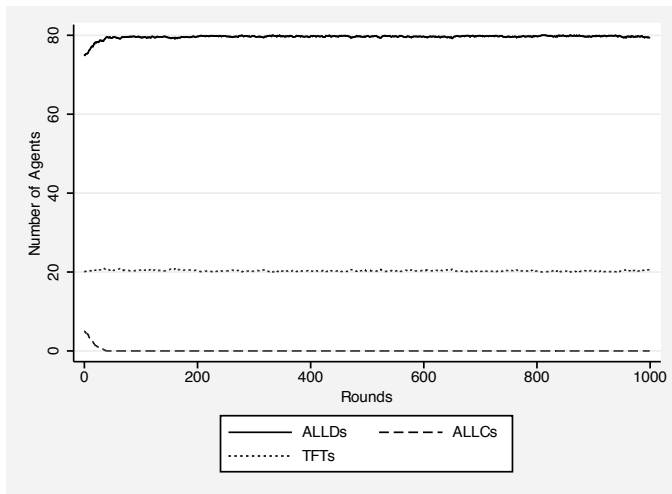


Panel a. With an initial population of 60 TFTs, 5 ALLCs and 35 ALLDs, TFTs decline and ALLDs increase in early rounds, stabilizing after approximately round 200. [Seed 545554]

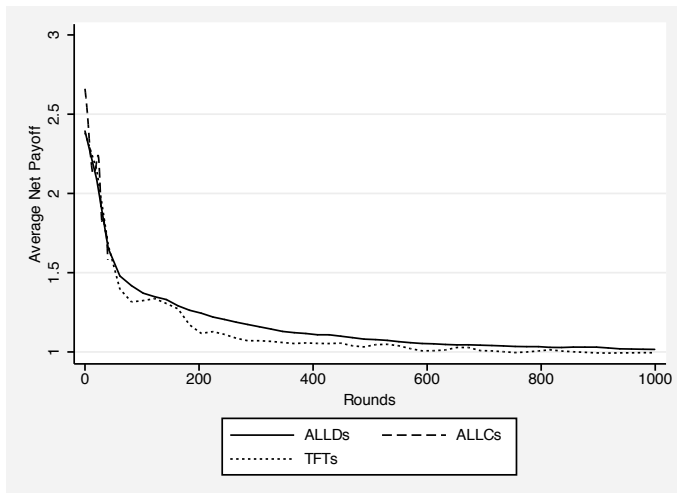


Panel b. Smoothed payoffs by strategy type

Figure 8. A Relatively Nasty World with Organizations

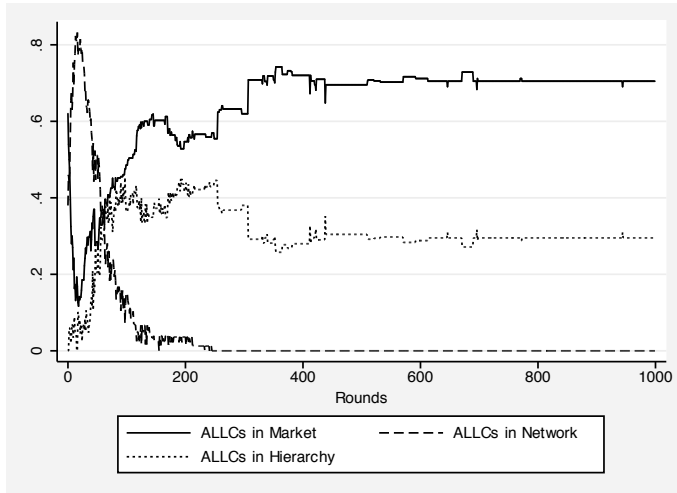


Panel A. This simulation begins with the same initial population as Figure 4 (75 ALLDs, 20 TFTs and 5 ALLCs). ALLDs increase slightly at the expense of ALLCs. TFTs retain their initial population share. [Seed 545554]

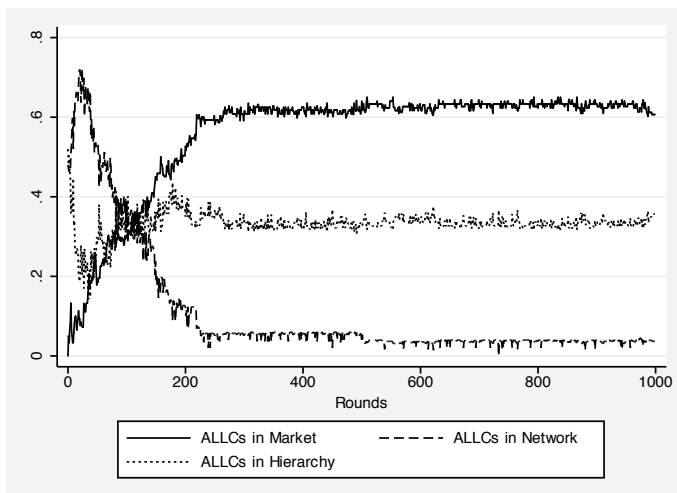


Panel b. Smoothed payoffs by strategy type.

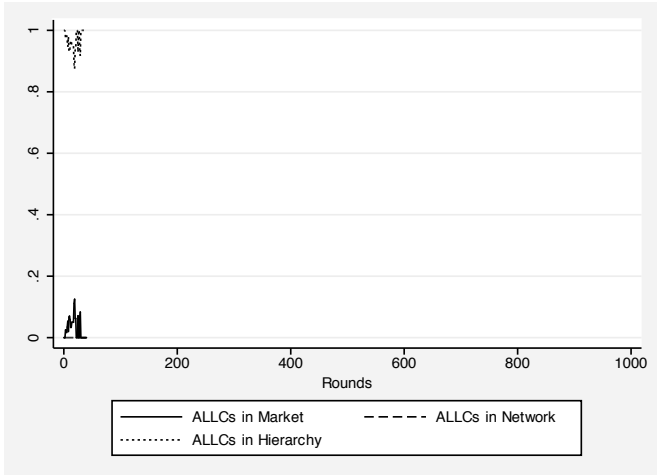
Figure 9. ALLC with Organizations



Panel a. In a nice world (5ALLCs, 90 TFTs, 5ALLDs) ALLCs start in the network and move into the market and hierarchy.

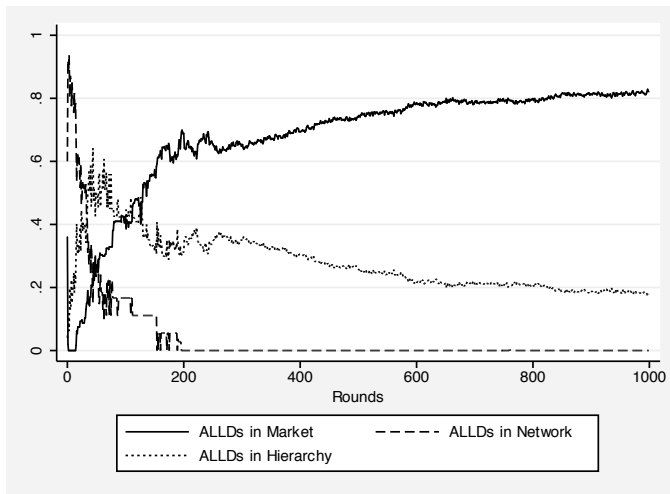


Panel b. In a moderate world (35 ALLDs, 5 ALLCs and 60 TFTs) ALLCs start in the network and and hierarchy and move to the market and hierarchy.

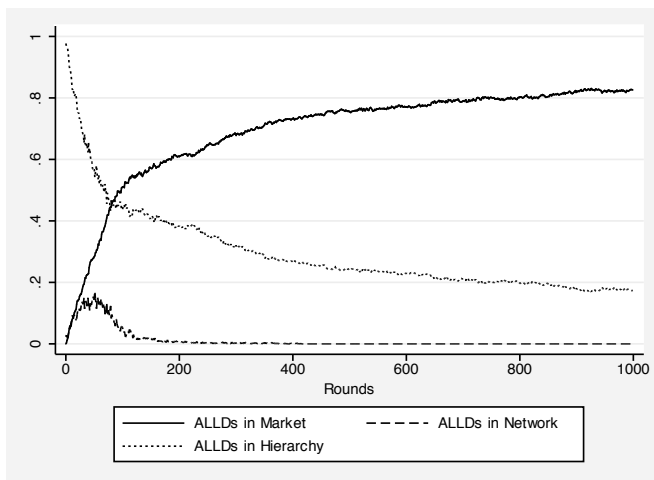


Panel c. In a nasty world (80 ALLDs, 5 ALLCs and 15 TFTs) ALLCS start almost entirely in the hierarchy but quickly are eliminated.

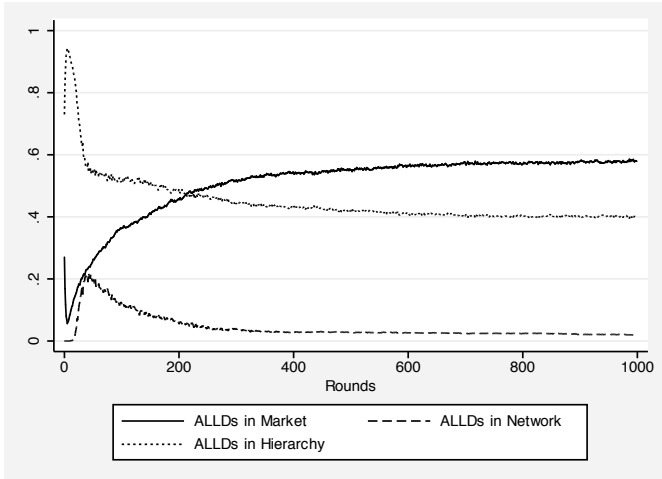
Figure 10. ALLD with Organizations



Panel a. In a Nice world (5 ALLDs, 5 ALLCs, and 90 TFTs) the ALLDs start in the network and hierarchy and move out of the network into the market and hierarchy.

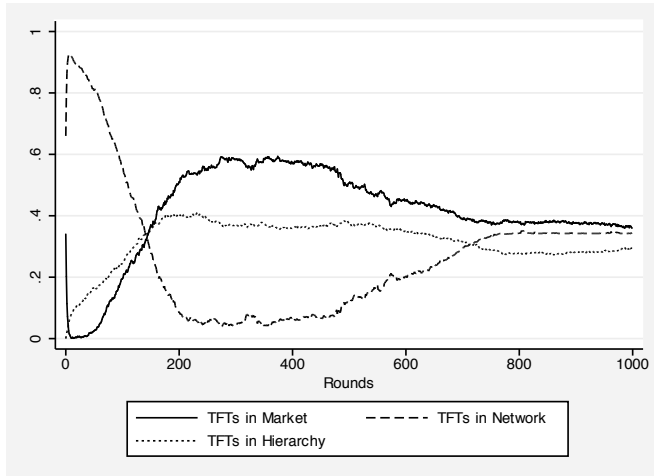


Panel b. In a moderate world (35 ALLDs, 5 ALLCs and 60 TFTs) ALLDS begin in the hierarchy and network and move into the market.

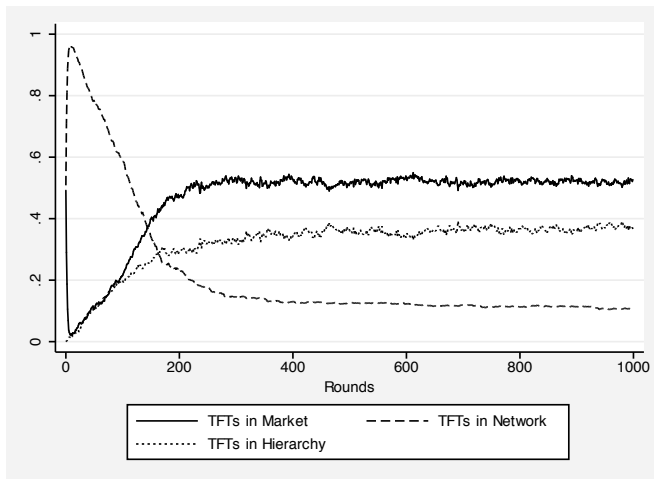


Panel c. In a nasty world (80 ALLDs, 5 ALLCs, and 15 TFTs) the ALLDs begin in the hierarchy and network and move into the market and hierarchy.

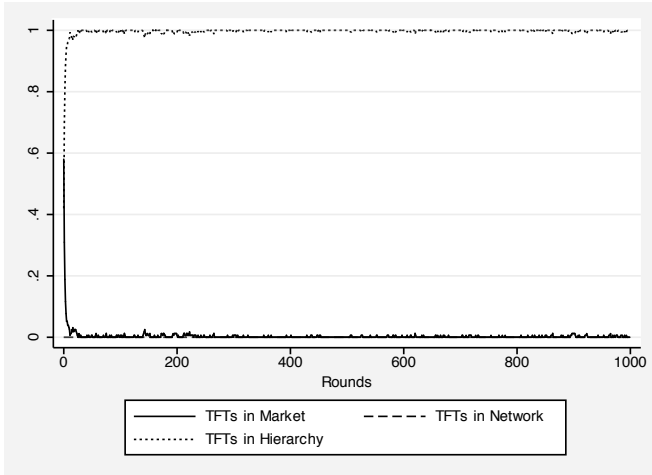
Figure 11. TFT with Organizations



Panel a. In a nice world (5 ALLCs, 5 ALLDs and 90 TFTS) the TFTs move from the network to the market, network and hierarchy.



Panel b. In a moderate world (5 ALLCs, 35 ALLDs, and 60 TFTs) TFTs move from the network to the market and hierarchy.



Panel c. In a nasty world (80 ALLDs, 5 ALLCs and 15 TFTs) the TFTs are universally in the hierarchy.

References

- Axelrod, Robert (1980). "Effective Choice in the Prisoner's Dilemma." *Journal of Conflict Resolution* **24**(1): 3-25.
- Axelrod, Robert (1980). "More Effective Choice in the Prisoner's Dilemma." *Journal of Conflict Resolution* **24**(3): 379-403.
- Axelrod, Robert (1984). *The Evolution of Cooperation*. New York, Basic Books.
- Axelrod, Robert (1987). The Evolution of Strategies in the Iterated Prisoner's Dilemma. *Genetic Algorithms and Simulated Annealing*. L. Davis. Los Altos, CA, Morgan Kaufman: 32-41.
- Axelrod, Robert (1997). *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton, NJ, Princeton University Press.
- Axelrod, Robert, et al. (2006). "Evolution of Cooperation among Tumor Cells." *Proceedings of the National Academy of Sciences* **103**(36): 12474-13479.
- Axelrod, Robert and Robert O. Keohane (1986). Achieving Cooperation under Anarchy: Strategies and Institutions. *Cooperation Under Anarchy*. K. A. Oye. Princeton, NJ, Princeton University Press: 226-254.
- Binmore, K. G. (1994). *Playing Fair*. Cambridge, MA, MIT Press.
- Binmore, K. G. (1998). "Review: The Complexity of Cooperation." *Journal of Artificial Societies and Social Simulation* **1**(1).
- Boehm, Christopher (1999). *Hierarchy in the Forest: The Evolution of Egalitarian Behavior*. Cambridge, MA, Harvard University Press.
- Flood, M. M. (1958). "Some Experimental Games." *Management Science* **5**: 5-26.
- Fudenberg, Drew and Eric Maskin (1986). "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." *Econometrica* **54**(3): 533-554.
- Hardin, Russell (1999). *Liberalism, Constitutionalism, and Democracy*. New York, Oxford University Press.
- Hirshleifer, Jack and Juan Carlos Martinez Coll (1988). "What Strategies Can Support the Evolutionary Emergence of Cooperation?" *Journal of Conflict Resolution* **32**(2): 367-398.
- Hobbes, Thomas (1651/1962). *Leviathan*. London, Collier.
- Hoffmann, Robert (2000). "Twenty Years On: The Evolution of Cooperation Revisited." *Journal of Artificial Societies and Social Simulation* **3**(2): 1-15.
- Jackson, Matthew O. (2008). *Social and Economic Networks*. Princeton, NJ, Princeton University Press.
- Jung, Danielle and David A. Lake (2009). Markets, Hierarchies, and Networks: An Agent-Based Organizational Ecology. *Annual Meeting of the American Political Science Association*. Boston, MA.
- Kahler, Miles, Ed. (2009). *Networked Politics: Agency, Power, and Governance*. Ithaca, NY, Cornell University Press.

- Keck, Margaret E. and Kathryn Sikkink (1998). *Activists Beyond Borders: Advocacy Networks in International Politics*. Ithaca, NY, Cornell University Press.
- Lake, David A. (2009). *Hierarchy in International Relations*. Ithaca, NY, Cornell University Press.
- Nowak, Martin A. and Robert M. May (1993). "The Spatial Dilemmas of Evolution." *International Journal of Bifurcation and Chaos* **3**(1): 35-78.
- Nowak, Martin A. and Karl Sigmund (2005). "The Evolution of Indirect Reciprocity." *Nature* **437**(October 27): 1291-1298.
- Podolny, Joel M. and Karen L. Page (1998). "Network Forms of Organization." *Annual Review of Sociology* **24**: 57-76.
- Powell, Walter W. (1990). "Neither Market nor Hierarchy: Network Forms of Organization." *Research in Organizational Behavior* **12**: 295-336.
- Rauch, James E. and James G. Hamilton (2001). Networks and Markets: Concepts for Bridging Disciplines. *Networks and Markets*. J. E. Rauch and A. Casella. New York, Russell Sage Foundation: 1-29.
- Stanley, E. A., et al. (1994). Iterated Prisoner's Dilemma with Choice and Refusal of Partners. *Artificial Life III*. C. G. Langton. Reading, MA, Addison-Wesley.
- Sugden, Robert (1989). "Spontaneous Order." *Journal of Economic Perspectives* **3**(4): 85-97.
- Taylor, Michael (1976). *Anarchy and Cooperation*. New York, Wiley.
- Vanberg, Viktor J. and Roger D. Congleton (1992). "Rationality, Morality, and Exist." *American Political Science Review* **86**(2): 418-431.
- Williamson, Oliver E (1975). *Markets and Hierarchies: Analysis and Antitrust Implications*. New York, Free Press.
- Williamson, Oliver E (1985). *The Economic Institutions of Capitalism: Firms, Markets, and Relational Contracting*. New York, Free Press.