

Markets, Hierarchies, and Networks: An Agent-Based Organizational Ecology

Author(s): Danielle F. Jung and David A. Lake

Source: *American Journal of Political Science*, Vol. 55, No. 4 (October • 2011), pp. 972-990

Published by: Midwest Political Science Association

Stable URL: <https://www.jstor.org/stable/23025131>

Accessed: 31-07-2018 17:22 UTC

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



JSTOR

Midwest Political Science Association is collaborating with JSTOR to digitize, preserve and extend access to *American Journal of Political Science*

Markets, Hierarchies, and Networks: An Agent-Based Organizational Ecology

Danielle F. Jung University of California, San Diego
David A. Lake University of California, San Diego

Markets, hierarchies, and networks are widely understood to be the three primary forms of social organization. In this article, we study the choice between these forms in a general, agent-based model (ABM) of cooperation. The organizational ecology is the product, an emergent property, of the set of choices made by agents contingent on their individual attributes and beliefs about the population of agents. This is one of the first attempts to theorize explicitly the choice between different organizational forms, especially networks and hierarchies, and certainly the first to do so in an ABM. The insights of the model are applied to current research on transnational networks, social capital, and the sources of hierarchy and especially autocracy.

Markets, hierarchies, and networks are widely understood to be the three primary forms of social organization (see Powell 1990; Ronfeldt 1996). Economic transactions occur in arm's-length markets between anomic buyers and sellers, hierarchically ordered corporations (Williamson 1975, 1985), or networks of co-ethnic traders (Greif 2006; Rauch and Trindade 2002) or, in Japan, of production *keiretsu*. Insurgencies are carried out by ad hoc "groups of guys" who come together to plot and carry attacks (Sageman 2008), by militias with commanders, ranks, and insignias, or by networks of loosely connected cells (Arquilla and Ronfeldt 2002). Social movements can arise in spontaneous, uncoordinated protests, as top-down, disciplined organizations or through tightly linked policy entrepreneurs who mobilize followers (Tarrow 2005). Countries cooperate with one another under anarchy, in supranational organizations, and in small groups of deeply interdependent or "networked" states (Kahler 2009a). Similarly, in the late medieval period, newly emergent sovereign territorial states, the epitome of hierarchy, decisively beat out trading networks like the Hanseatic League to become the primary units of international political life (Spruyt

1994). Likewise, legal contracting may be displacing social networks based on personal relationships and private knowledge as the predominant form of organization in American society (Putnam 2000). These same organizational forms recur at all levels of social interaction, but we understand little about the conditions that give rise to one form or another or why actors choose between alternative organizations.

In this article, we study the choice between markets, hierarchies, and networks in a general, agent-based model (ABM) of cooperation. Agents choose to cooperate or not in a market, hierarchy, or network as a function of their individual attributes and their beliefs about the attributes of the other agents with whom they may interact. The set of organizations in a population—the organizational ecology—is an emergent property of the choices made by agents. Unlike closed-form formal models that focus on a small number of actors in a well-defined strategic setting, our ABM shifts attention to the attributes of the *population* of interacting agents.

We theorize explicitly the choice between all three organizational forms. Markets have in the past been compared separately to hierarchies and networks, but few

Danielle F. Jung is a Ph.D. candidate in the Department of Political Science, University of California, San Diego, 9500 Gilman Dr. #0521, La Jolla, CA 92093–0521 (dfjung@ucsd.edu). David A. Lake is the Jerri-Ann and Gary E. Jacobs Professor of Social Sciences and Distinguished Professor in the Department of Political Science, University of California, San Diego, 9500 Gilman Dr. #0521, La Jolla, CA 92093–0521 (dlake@ucsd.edu). Jung and Lake participated equally developing the architecture of the model and writing this article; Jung was entirely responsible for implementing the model and the simulations reported here. We are grateful for the programming assistance of Daniel Han and Robert Chen and helpful comments from Robert Axelrod, James Fowler, Peter Gourevitch, Emilie Hafner-Burton, Robert Keohane, Dylan Lake, Woody Powell, David Ronfeldt, Nick Weller, Wendy Wong, David Zaring, and two anonymous reviewers. This research was funded by the California Institute for Telecommunications and Information Technology (CalIT2) at the University of California, San Diego. Lake also gratefully acknowledges the support of the Center for Advanced Study in the Behavioral Sciences.

American Journal of Political Science, Vol. 55, No. 4, October 2011, Pp. 972–990

©2011, Midwest Political Science Association

DOI: 10.1111/j.1540-5907.2011.00536.x

have compared hierarchies and networks directly to each other.¹ Past ABMs of cooperation have not included institutional or organizational features while computational models of organizations have focused on intra-organizational attributes.² This is the first ABM to study the choice between organizational forms in a general model of cooperation.

The ABM reveals theoretical limitations and inconsistencies in existing theories, highlights interaction effects and population dynamics—thereby offering new explanations for phase shifts and other phenomena absent from purely verbal and even formal models—and generates new theoretical insights. Networks, for instance, both provide information to agents and permit agents to select specific agents with whom to interact. We find that the informational benefits of networks are somewhat limited, valuable only to actors with contingent strategies, such as tit-for-tat within a Prisoner's Dilemma, who can learn from the knowledge of others. In a static population with no exogenous changes in the attributes of agents, the informational benefits of networks quickly decline as agents develop their own histories of play with others. When permitted to choose other agents with whom to interact within their networks, however, all types of agents may join, yet this permits not only cooperation but also exploitation by “bullies” and offers a far less sanguine view of networks than is common. Conversely, hierarchy is preferred in relatively “nastier” populations with larger numbers of uncooperative and opportunistic agents. Paradoxically, it is the most cooperative agents who first join hierarchies to reap the benefits of centralized enforcement when the population turns nasty. In many cases, population dynamics create phase shifts in which agents flip from one organization to another as a result of small changes in their environments, captured by varying different parameters in our model.

After explaining the ABM in some detail, we apply it to the literatures on transnational networks and international governance, networks and social capital, and the sources of hierarchy. These applications are not tests of our model; indeed, our ABM is intended to be general and is not designed to capture issues specific to these literatures. Nonetheless, we believe the model clarifies

propositions central to each literature and raises important new questions for research.

Markets, Hierarchies, and Networks as Organizations

We focus on the generic problem of cooperation among self-seeking actors choosing between different organizational forms. By using simple, ideal type representations we aim to identify broad principles of organizational ecology that can be applied to an array of cooperation problems. For each organization, we distill the form to its essence as characterized in the existing literature. There are, no doubt, many hybrid forms in the real world, but to keep the analysis simple we focus only on ideal types of markets, hierarchies, and networks.

The problem of cooperation is characterized here as a repeated two-player Prisoner's Dilemma (PD) game (see Figure 1).³ As Axelrod (1984) and others have shown, such a model captures the essential features of a broad class of cooperation problems. To model hierarchy appropriately, however, we modify the standard PD setup slightly. Specifically, we permit agents to have individual preferences (p_i) defined by ideal points along a finite continuum. Our intuition is that mutual cooperation does not mean the same thing or carry the same value for all pairs of political actors, especially under hierarchy. Cooperation with an actor who shares one's preferences is different from cooperation with an actor with preferences distant from one's own. Assuming that cooperation occurs at the median of their ideal points, two “left” actors, for instance, gain greater utility from cooperating with one another than might one “left” and one “right” actor. If cooperation means working together to promote a political cause, two left actors will pursue a policy closer to their preferences than would a left and right actor, for whom the median would be further from their ideal points. To anticipate a technical point below, when actors both cooperate, we subtract the weighted difference between their ideal points from the payoffs from mutual cooperation ($k_{ij} = w(|p_i - p_j|/2)$). In all cases, any weight on preferences greater than zero makes cooperation less likely as it reduces its value relative to other possible outcomes. As the weight on preferences increases, agents

¹ On markets and hierarchies, see Williamson (1975), and on markets and networks, see Rauch and Hamilton (2001). For comparisons of networks and hierarchies, see Powell (1990) and Kahler and Lake (2009).

² On agent-based models of cooperation and politics, see Axelrod (1984, 1997), Cederman (1997), Epstein (2007), Kollman, Miller, and Page (2003), and Miller and Page (2007). On computational models of organizations, see Prietula, Carley, and Gasser (1998), Ilgen and Hulin (2000), and Chang and Harrington (2006).

³ We characterize this as a repeated PD because the same interaction structure recurs throughout the population. Agents may, and often do, interact iteratively, just not always in subsequent rounds. While interactions are modeled in dyads, as will be explained in more detail below, they are done so based on beliefs about the *population* rather than the individual.

FIGURE 1 The Modified Prisoner's Dilemma Game

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

$k = k_{ij}$ for markets and networks, k_{ih} for hierarchies

who might otherwise choose to cooperate will now defect.⁴ The primary implication of this amendment to the standard PD game is that actors with more similar preferences will be more likely to cooperate than agents with more dissimilar preferences. In hierarchy, by contrast, agents cooperate at the hierarchy's ideal point (p_h) and payoffs for cooperation are adjusted by the difference not between their individual preferences but between each agent's ideal point and that assigned for the hierarchy as a whole ($k_{ih} = w|p_i - p_h|$). A key attribute of hierarchy is the ability of a third party—typically the ruler, leader, or boss—to command legitimately certain actions between the members of the organization (see below). By assuming that cooperation occurs at the hierarchy's ideal point, we capture, in part, the notion of command or authority that is central to hierarchy. The intent here is to model not just vertical interactions between the leader and the members of the hierarchy, but cooperation between members that is “commanded” and centrally enforced (see below). In the final substantive section below, we vary the hierarchy's ideal point relative to the mean in society to reflect variations in regime type.

Although a common term, the concept of market lacks a fixed analytical definition. Once referring only to a site for trading, since the early twentieth 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, conceived as networks of linked firms of factors of production.⁵ As an organization, according to Powell, markets are “the paradigm of individually self-interested, noncooperative, *unconstrained* social interaction” (1990, 302; italics added) that engages 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 and generalizes it to a greater range of interactions.⁶

The essential feature of markets represented in our model is that risky exchanges occur among strangers. We model this in terms of random encounters between two individuals who play a round of a PD. In framing the institution this way, we do not include features of markets such as trade associations, third-party recommendations, or public enforcement of contracts. While we acknowledge this conception neglects some common features of economic markets, we also believe it applies more generally to markets as generic organizations. The PD is analogous to many other risky situations and allows us to have a baseline “state of nature” against which we can compare other organizational forms.

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 networks here in two ways. First,

⁵ On economic and sociological views of markets, see Swedberg (2003).

⁶ 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, again, 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.

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). On edges that add information to facilitate coordination, see Enemark et al. (2011) and McCubbins, Paturi, and Weller (2009). We treat networks more as governance structures, but focus on the information flows and reciprocity between agents common to both

⁴ All results reported in this article, except those on hierarchy in Figures 4 and 5, have been replicated with the weight on preferences set to zero ($w = 0$), deleting preferences from the analysis (i.e., the standard PD). The results on markets and networks remain nearly identical with no substantive implications. We do not focus extensively on preferences here except in the final section on hierarchy, but they are included in the parameter sweeps in the appendix.

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) and what is j 's ideal point (p_j). With this information, agent i can then decide whether to cooperate or defect with j . Thus, networks provide information that supplements what 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 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 (see Podolny and Page 1998, 59; Powell 1990, 303). Only agents who possess a contingent strategy (defined below) will ever choose to join a network to gain information about others, and having joined they will play reciprocally.

Second, networks also permit agents to intentionally select other agents with whom to interact, a more direct form of reciprocity. As explained below, 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 within networks 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.

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. 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.

Thus, the essential features of networks captured in our model are information sharing and some ability to

select one's partner (selective affinity). Proponents of networks may find this conceptualization too simple and lacking in what they believe are core features. In part, such criticisms follow from the absence of any consensus on networks as organizations. Even Powell (1990), which opened the systematic comparison of organizational forms, largely defines networks by what they are not (i.e., markets and hierarchies) and less by what they are in any positive sense. Nonetheless, information sharing and selective affinity appear common to nearly all treatments, and we focus on these attributes here.⁸

Third-party enforcement stands at the core of all definitions of hierarchy. In our model, agents within the hierarchy cooperate with one another at the hierarchy's ideal point, subject to punishments for (random) defection.⁹ If an agent defects, it receives the temptation (T) payoff less the punishment, while the other receives the sucker's payoff (S).¹⁰ We treat both the probability of cooperation within the hierarchy (q) and the magnitude of the punishment (v) as exogenous. 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

⁸ This conception of networks does not allow permanent structures to be created, but each individual agent creates its own network of cooperators that reflects its recent behavior as well as that within the population. While not establishing permanent structures, this representation allows networks to remain fluid, often a key characteristic of this organization.

⁹ 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).

¹⁰ 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.

approaches. On the different conceptions and uses of networks in political science, see Kahler (2009a).

organizational forms.¹¹ This representation allows us to investigate how the probability of defection and levels of punishment affect the expected utility of cooperation under hierarchy. We include a variable tax on members joining a hierarchy (τ), subtracted from the expected utility of joining the hierarchy.

Agents in the hierarchy who interact with agents outside the same hierarchy play as in the market. In a firm, 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 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, the rule of law represented in cooperation at the hierarchy's ideal point and centralized punishment for defection does not apply "extraterritorially" or beyond the members of the hierarchy.

In our model, agents join only one organization and select at random and play only one other agent in each round of the game (although they may be selected multiple times by other agents, particularly under selective affinity). 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 and serve on a community organization's board with that same person. In our model, such complex relationships are simply treated as separate rounds of the game, and the conditions that lead one interaction to take place in a network and another to occur in a hierarchy are studied as variables. This analytic move simplifies but does not, we believe, unduly distort more complex relationships.

Similarly, agents in the model choose freely each round to join the organization that promises the high-

est 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, however, especially 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 nonvoluntary 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 his will—including governing participation—on reluctant others.¹² Conversely, the exit of all agents from a hierarchy approximates the loss of popular support for a political regime.

Another important assumption of the model is that all actors must play the agent to whom they are assigned or, under selective affinity, whom they choose in any given round. In other words, agents cannot "opt out" of an interaction. Although this restriction has important consequences as cooperators cannot choose to form closed groups that exclude defectors, creating a "not play" option in addition to cooperation or defection changes the game itself (even if the PD is retained as a subgame). For reasons of both simplicity and comparability, we restrict our analysis here to the standard PD game and overlay organizations on this structure.

Finally, as endogenous products of the choices of many independent agents, organizations are created anew each round of play. Which agents constitute the market, the hierarchy, or a network is established by their choices, which may differ by round. Social scientists often treat organizations as sticky or long-lived, whereas individuals are variable and short-lived. As our interest is in the origins and survival of organizational forms, the assumption of a static population of agents seems to us to be a reasonable simplification. To the extent that organizations change the pattern of cooperation, this will inevitably feed back upon the population in some dynamic evolutionary process. We intend to study selection and evolution in the future. But understanding how individuals choose one organization

¹¹ 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.

¹² On the collective nature of authority and hierarchy, see Lake (2009a, chap. 1).

over another at any moment in time is a prerequisite to modeling more complex dynamic processes.

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 inevitably carries some cost in understanding specific organizations and individual choices.¹³ 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.

Modeling Organizational Ecologies in an Agent-Based Framework

We describe the ABM here in its three stages: initialization, learning, and organizational choice. The model, along with the expected utility equations for each organization and full parameter scans for each variable, is detailed in the online appendix. There is no convention for evaluating ABMs. The parameter sweeps in the appendix are intended to allow the reader to assess the robustness of the simulations. In these scans, we run each of the user-defined parameters from a minimum to maximum under different population mixes, holding all other parameters at their default values, and track strategy types over organizations to describe changes in the organizational ecology. 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 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 levels such that reasonable changes in any single parameter are likely to lead at least some agents to alter their organizational choices. Different default values, of course, might change the organizational ecologies that emerge from the model. This said, the sweeps in the appendix nonetheless permit the reader to judge parameters of interest from the same baseline used below.

¹³ For example, Coleman (1990) places extended third-party reciprocity/enforcement as central to networks, but this conceptualization is not universally shared. For the sake of generality, we do not model interactions and decisions in a way that will capture this hybrid, though this may limit the scope of our results slightly.

Initialization

The model begins with specification of 24 user-defined parameters (see Table 1). Payoffs for the various outcomes are set: T, R, P, and S.¹⁴ The user defines the population of actors, defined by the distribution of strategy types, and their preferences. Following Hirshleifer and Coll (1988), we focus on three basic strategies: all cooperate (ALLC), all defect (ALLD), and tit-for-tat (TFT). ALLC and TFT are *nice* strategies that begin by cooperating with new agents, while ALLD is a *nasty* strategy.¹⁵ Below, we refer to nice and nasty populations as defined by the relative proportions of these two sets of agents. Preferences (p_i) are defined over a $[0,1]$ space and randomly assigned from a normal distribution.¹⁶ The weight on preferences (w) can also be varied.

The organizational parameters, known to all agents, 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 polled [a 3×3 ($\alpha = 3, l = 3$) network is illustrated in the appendix]. 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 polls. 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.¹⁷ Selective affinity is defined by the

¹⁴ We set the default cardinal payoffs in the PD game as in Axelrod (1984) for purposes of comparability. All other default parameters were then set relative to these default payoffs. In defining the strategy space, we build off of the now accepted space defined by Axelrod and others, including Cohen et al. (2001).

¹⁵ Although this is a limited set of strategies, our intent is to see how extremely intuitive strategies, on either extreme (ALLC and ALLD), as well as how contingent strategies, particularly TFT, are able to fare in an organizational environment.

¹⁶ We have no priors about different distributions of populations. For intuitive and analytical ease, we chose a normal distribution. We anticipate that the role of preferences will be amplified in deeply divided/bimodal societies, but in an effort to make this model as generalizable as possible we feel this is an appropriate assumption.

¹⁷ 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. 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.

TABLE 1 User-Defined Parameters and Default Values

Parameter	Symbol	Description	Default Value
General			
Increments		Number of times the simulation is run incrementing a parameter	21
Repetitions		Number of times the identical simulation is repeated with different random seeds	1000
Rounds		Number of rounds of play	20
Mean for ideal point		Distribution of actors' policy preferences in population	0.5
Weight on ideal	w	Weight on policy preferences	1.0
Learning rounds		Set as either number of rounds or population convergence to within a proportion of the true population mean	5 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
Ideal point	p_h	Ideal point of the hierarchy	0.5
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
Selective Affinity			
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

probability (η) an agent gets to select an agent from its affinity memory (m_a) with whom to interact, with one minus this probability being the rate at which that agent

Conversely, without this restriction on network memory, the network would return "too much" information in early rounds and become obsolete almost immediately.

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 its ideal point (p_h), 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). 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.

Learning

Agents begin the simulation without any knowledge of the distribution of the other agents’ strategies or ideal points. In the learning phase, agents are randomly paired with other agents with whom they play a round of the game according to their fixed strategy type with the specified payoffs. Agents develop beliefs about two parameters from their interactions with other agents. First, they learn about the distribution of other strategy types. Observing their own payoffs, they then back out whether the other agent cooperated or defected, store this action in their memory, 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 others’ actions, limited to cooperation or defection, not their underlying types. This is equivalent to not being able to observe an individual’s intent or strategy, 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. Second, when they cooperate with other agents, agents also learn about the distribution of preferences in the population and whether their own preferences are relatively extreme or moderate. Again, knowing only their own preference, agents who cooperate with one another examine their payoffs and back out the ideal point of the other agent, store this in memory, and then update their beliefs about the mean ideal point in the population (\hat{P}). 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 its course of play, meaning that even agents with the same strategy type and similar or identical ideal points will make different organizational choices in the next stage. This introduces heterogeneity of agents even within a fixed population of only three basic strategy types.¹⁸ Agents who believe the population is nas-

tier 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 concludes, the main simulation of interest begins and continues for a fixed number of rounds. 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.

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 to join the network (ϕ). The utility for entering a hierarchy depends on the proportion of the population in the hierarchy (θ), weighed against the likelihood of cooperation within the hierarchy (q), the punishment for defection (v), the tax (τ), and the ideal point of the hierarchy (p_h).

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 noncontingent 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. In this way, we assume that all agents

¹⁸ We could assign these beliefs and allow them to be updated over the course of play. We favor an endogenous learning mechanism, however, as it allows us to vary the accuracy of these beliefs by the number of rounds of play during the learning phase.

are sincere in their reporting and are known to be so by all other agents.¹⁹ 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 hierarchy, the agent will cooperate at the rate the hierarchy enforces (q). If the agent defects ($1-q$), it will be punished at the defined level (v). If a player is matched with a player outside of its hierarchy, it will play as if it were interacting in the market.

Following play, real payoffs are calculated as a function of the outcome of play, adjusted for the players' ideal points (k) if the outcome was cooperative, punishments, and fees prescribed by their organizations. Actual payoffs can differ from expected payoffs, but are on average the same.

We are primarily interested here in the organizations selected overall and by specific strategy types under varying parameters, and the real payoffs of the agents. Our strategy is to simulate organization choice and payoffs under varying conditions by incrementing selected parameter values over some range; this is roughly equivalent to comparative static predictions in closed form models. 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, unless noted otherwise, we replicate the simulation 1,000 times for each increment of the parameter and report the average of the results.²⁰

¹⁹ 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. 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. See Downs, Rocke, and Siverson (1986). 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.

²⁰ Despite the number of replications, the results reported below in graphs are not always smooth functions due to the discrete increments of the parameters. Each run of the model creates a

Illustrations of the Model

Like others, our ABM “is a way of doing thought experiments” that, because of complex interactions, may have nonobvious conclusions (Axelrod 1997, 4). We illustrate the potential of the ABM to provide new insights into organizations and cooperation by briefly summarizing simulations that capture core features of three disparate literatures in political science. Our model reveals theoretical limitations and inconsistencies in existing theories. In the case of transnational networks, for instance, we focus on the information value of networks and find that, contrary to much of the existing literature, networks rapidly decline in use. Selective affinity causes networks to be robust, on the other hand, suggesting it is not information but the opportunity to select partners which sustains networks in everyday life. By highlighting interaction effects and population dynamics, the ABM also offers new explanations for phenomena absent from purely verbal and even two-player, closed-form formal models. In the case of social capital, we show the value of population models in explaining phase shifts in behavior now unexplained in the literature. Finally, the model generates new theoretical insights. Again, in the social capital literature, we demonstrate how hierarchy is a viable alternative to social networks and may more accurately characterize modern American society than market interactions. Similarly, in a simple depiction of the emergence of political hierarchy, we not only derive the core logic of Hobbes's *Leviathan* from the model, but also show how hierarchy can emerge even when the ruler has preferences that are extreme or distant from the mean of society. This produces important insights into the nature of autocratic rule.

Transnational Networks and International Governance

Transgovernmental networks (TGNs) are, Slaughter (2004, 8–11) claims, the solution to the governance dilemma created by a need for global institutions and a continuing fear of centralization.²¹ According to Slaughter, TGNs have become prominent in coordinating

large number of observations. We are limited by the number of increments for each variable versus the number of replications we can perform. In the appendix, we present Figure A2 with 95 percent confidence intervals plotted around the result. As this figure demonstrates, additional replications are unlikely to produce significantly different results. Since confidence intervals make already complex graphs harder to read, we do not include them in other figures.

²¹ See also Raustiala (2002) and Eilstrup-Sangiovanni (2009). For a related view emphasizing transparency, often through networks, see

central banking, corporate regulation, the international legal system, and more. Such networks are, in her view, not only a building wave but also an effective solution to the absence of hierarchical, authoritative institutions in world politics. Although TGNs do many things, in Slaughter's view, primary among them are creating "incentives to establish a good reputation and avoid a bad one" and exchanging "regular information about their own activities and . . . best practices" (2004, 3).

Transnational economic networks (TENs) are also seen as key to economic growth and governance. In his study of the Maghribi traders, entrepreneurs active in long-distance exchange around the Mediterranean in the eleventh to the fourteenth centuries, Greif (2006, 59) finds two attributes were central to their success: linking each agency transaction to all future agency transactions with other merchants in the network—in a word, reciprocity—and information sharing on agents among the merchants.²² Similarly, as Spruyt notes, one of the key tasks of the Hanseatic League, a medieval network of city-states engaged in international trade, was "to facilitate the exchange of information between merchants" (1994, 123). These same traits are key to the efficiency of Japan's corporate networks in the modern era (Lincoln and Gerlach 2004). In their emphasis on information sharing, TENs are essentially similar to TGNs.

Even though our ABM is not identical to any specific network in these different literatures, it captures the essence of networks as governance structures in its focus on information sharing and selective affinity. Demonstrating much of the promise of transnational networks, our model nonetheless suggests that the conditions under which networks will be preferred to markets and hierarchies are contingent in ways not yet appreciated by the current literature. We focus on two key limitations of networks not because we dispute their benefits but because this is where the ABM reveals further theorizing is most necessary.

Networks as a source of information only quickly become obsolete with time (rounds of the game). In our characterization, agents acquire information about the strategy type and ideal point of another agent directly

through interactions or indirectly through the network of agents with whom they have cooperated in the past. Networks are valued for the information about other agents they can provide. As agents acquire knowledge of other agents through their own interactions or the network, the value of the network declines. At an extreme, after an agent has interacted with or acquired knowledge through the network about every other agent in the population, the network can return no new information of value to that agent; if there is any cost to belonging to a network, agents will then choose some other organizational form (see Figure 2a). Paradoxically, the larger the network relative to the population—making it more beneficial and attractive in early rounds of the game—the more quickly it becomes obsolete (not shown). Transnational networks may be initially useful in coordinating diverse actors, but all else held constant their utility declines over time as the actors become more familiar with one another.

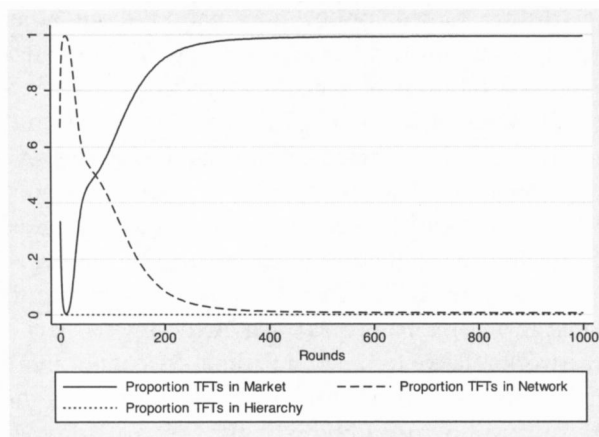
Similarly, the larger the population, the less likely networks are to be selected by agents (see Figure 2b). It might seem that larger populations favor networks as it takes more rounds of the game for agents to acquire direct knowledge of other agents and, therefore, networks are more valuable. Yet, for networks of a given size, larger populations also mean that the network is less likely to return information useful to the agent about the agent with whom it is randomly paired.²³ In very large populations, "small" networks are of little value and, therefore, will not be chosen by agents. This suggests that networks may develop among, say, the functional ministers of relatively small groups of countries, such as the G8, but not among broader groups like the G77 or all UN members. Likewise, networks may function effectively among small groups of traders, like the Maghribi or Hansa, but not among all traders in a region.

In contrast to the informational benefits of networks, selective affinity produces robust networks that persist indefinitely (see Figure 2c). With a small chance of selecting a partner, agents leave the network due to the declining value of information relatively quickly (solid line), but with higher rates of selective affinity agents join and stay in the network for the dyadic cooperation it sustains (dashed lines). As might be expected, ALLC and TFT agents join the network in hopes of getting to select a partner with whom they have cooperated in the past, an effect that does not diminish as learning occurs. Selective affinity

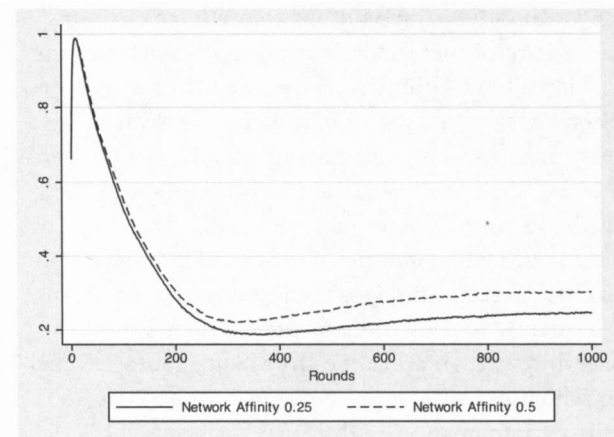
Florini (2003). Transnational advocacy networks (Keck and Sikkink 1998) are related but different in that the information being shared is not about other members of the network but the states that are the targets of political change.

²² Another well-known model, Milgrom, North, and Weingast's (1990) "law merchant," possesses attributes of networks, in its use of strategies of reciprocal punishment, and hierarchy, in its centralized dissemination of information. Their analysis becomes a true network only if the law merchant is depicted as a central node to which all other nodes are directly linked.

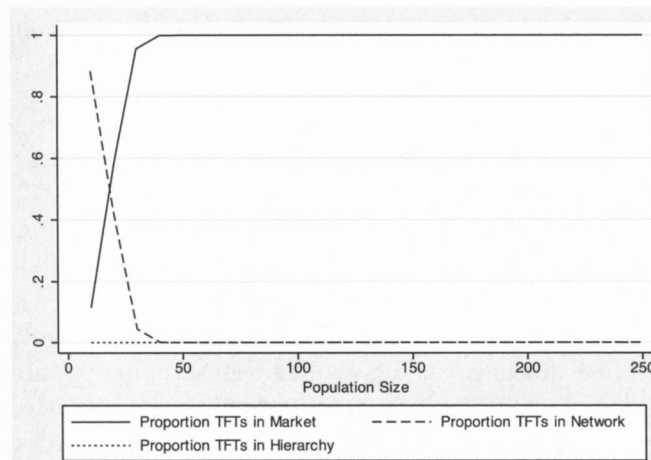
²³ The probability of a network returning a useful reply is $(\frac{m}{n-1})(\sum_{y=1}^l \beta \alpha^y)$. As n increases, the probability falls. In all cases, allowing redundancy reduces the probability of a useful response. In the model, we do not adjust the expected utility of networks for redundant responses.

FIGURE 2 Declining Network and Population Size

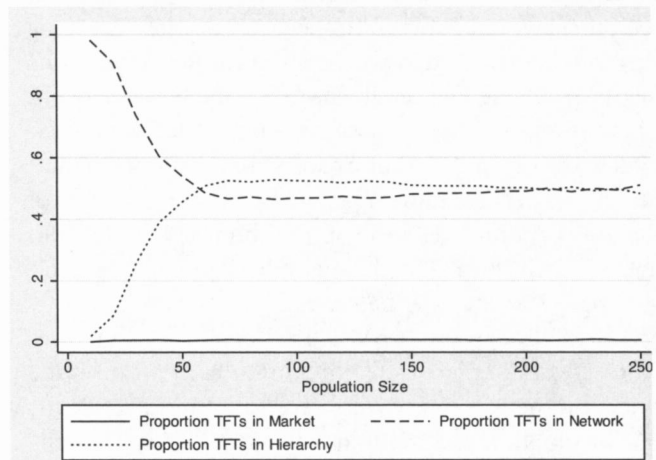
Panel a.



Panel c.



Panel b.



Panel d.

Panel a. Network decay in a world without selective affinity. TFTs quickly leave the network. The population used for this simulation—a contingent, but predominantly nice one composed of 90 TFT and 10 ALLD agents—is the most likely to use the network for information.

Panel b. Population size and TFT organizational choice. As population size increases, TFTs leave the network for the market. This population is 40% cooperative types (30% TFTs, 10% ALLCs) and 60% ALLDs.

Panel c. Network decay in a world with selective affinity. In a population identical to that in Figure 2a, affinity in the network is enabled, first at 25%, then at 50%. The higher the rate of selective affinity, the larger the proportion of TFT agents who remain in the network even after the information value has “worn off.”

Panel d. Population size and TFT organizational choice with selective affinity. As population size increases, TFTs leave the network for the hierarchy. Population is identical to that in Figure 2b. At all population sizes, TFTs who go into the hierarchy believe the world is significantly nastier than those who go into the market after leaving the network ($t = 1,900$; $df = 16,000,000$; $t < 0.0001$).

also offsets the size effect just noted (see Figure 2d). This confirms the relatively optimistic view in much of the literature that networks are indeed an effective facilitator of cooperation. It also suggests that observed networks that endure for long periods are more likely founded on gains that arise from selecting one's partners than from information.

Paradoxically, however, selective affinity also carries a “dark side.” With selective affinity, nasty agents will also choose to join a network in anticipation of “suckering” or

exploiting agents who cooperated with them in past plays of the game (not shown).²⁴ These nasty agents become essentially schoolyard bullies who identify and repeatedly exploit a victim, especially ALLC types who cannot retaliate in future rounds. These victims eventually update their beliefs and perceive the world as nastier than

²⁴ This dark side of networks, of course, is a product of the inability of victims to decline to interact with the agent who has selected them. If agents could opt out of interactions (see above), this effect would go away and only the cooperative effect would endure.

it really is and subsequently escape to the hierarchy for protection. Examples abound in transnational terrorist and criminal networks that intimidate locals into providing resources and intelligence (see Kahler 2009b; Kenney 2009). This dark side of networks is not anticipated in the theoretical literature. Thus, selective affinity both sustains networks and drives some nice strategy types who would otherwise remain in the network into the hierarchy.

The supposed benefits of transnational networks should be treated with caution. The question is not whether networks substitute for alternative forms of governance but, rather, what are the ranges of conditions under which networks will be selected. As the declining benefits of information within networks suggest, these conditions may be more restrictive than they first appear. In addition, networks may not only facilitate cooperation but also create opportunities for exploitation. More attention must be paid to the details of any specific network before assessing its overall effects on cooperation.

Networks and Social Capital

In 2000, Robert Putnam published a path-breaking study on the decline of social capital and civic engagement in the United States. For Putnam, “the core idea of social capital theory is that social networks have value” (2000, 19). As he elaborates, “social capital refers to connections among individuals—social networks and the norms of reciprocity and trustworthiness that arise from them.” If social capital is at its core a social network, as Putnam indicates, our ABM may shed light on this sea change in American society.²⁵ Indeed, although our model is designed to capture the general effects of organizational forms on cooperation and not specifically to represent Putnam’s theory, it nonetheless has important implications for understanding the decline in social capital and alternatives to social networks.

We can model the decline of social capital within our ABM in four ways, each of which captures slightly different dimensions of Putnam’s analysis. In the online appendix, we vary each representation separately. In Figure 3, we manipulate simultaneously all four representations, increasing population size, the cost of joining the network, and the proportion of nasty players in the

population while decreasing the rate of selective affinity. This captures the decline of social capital in all forms and along all dimensions.

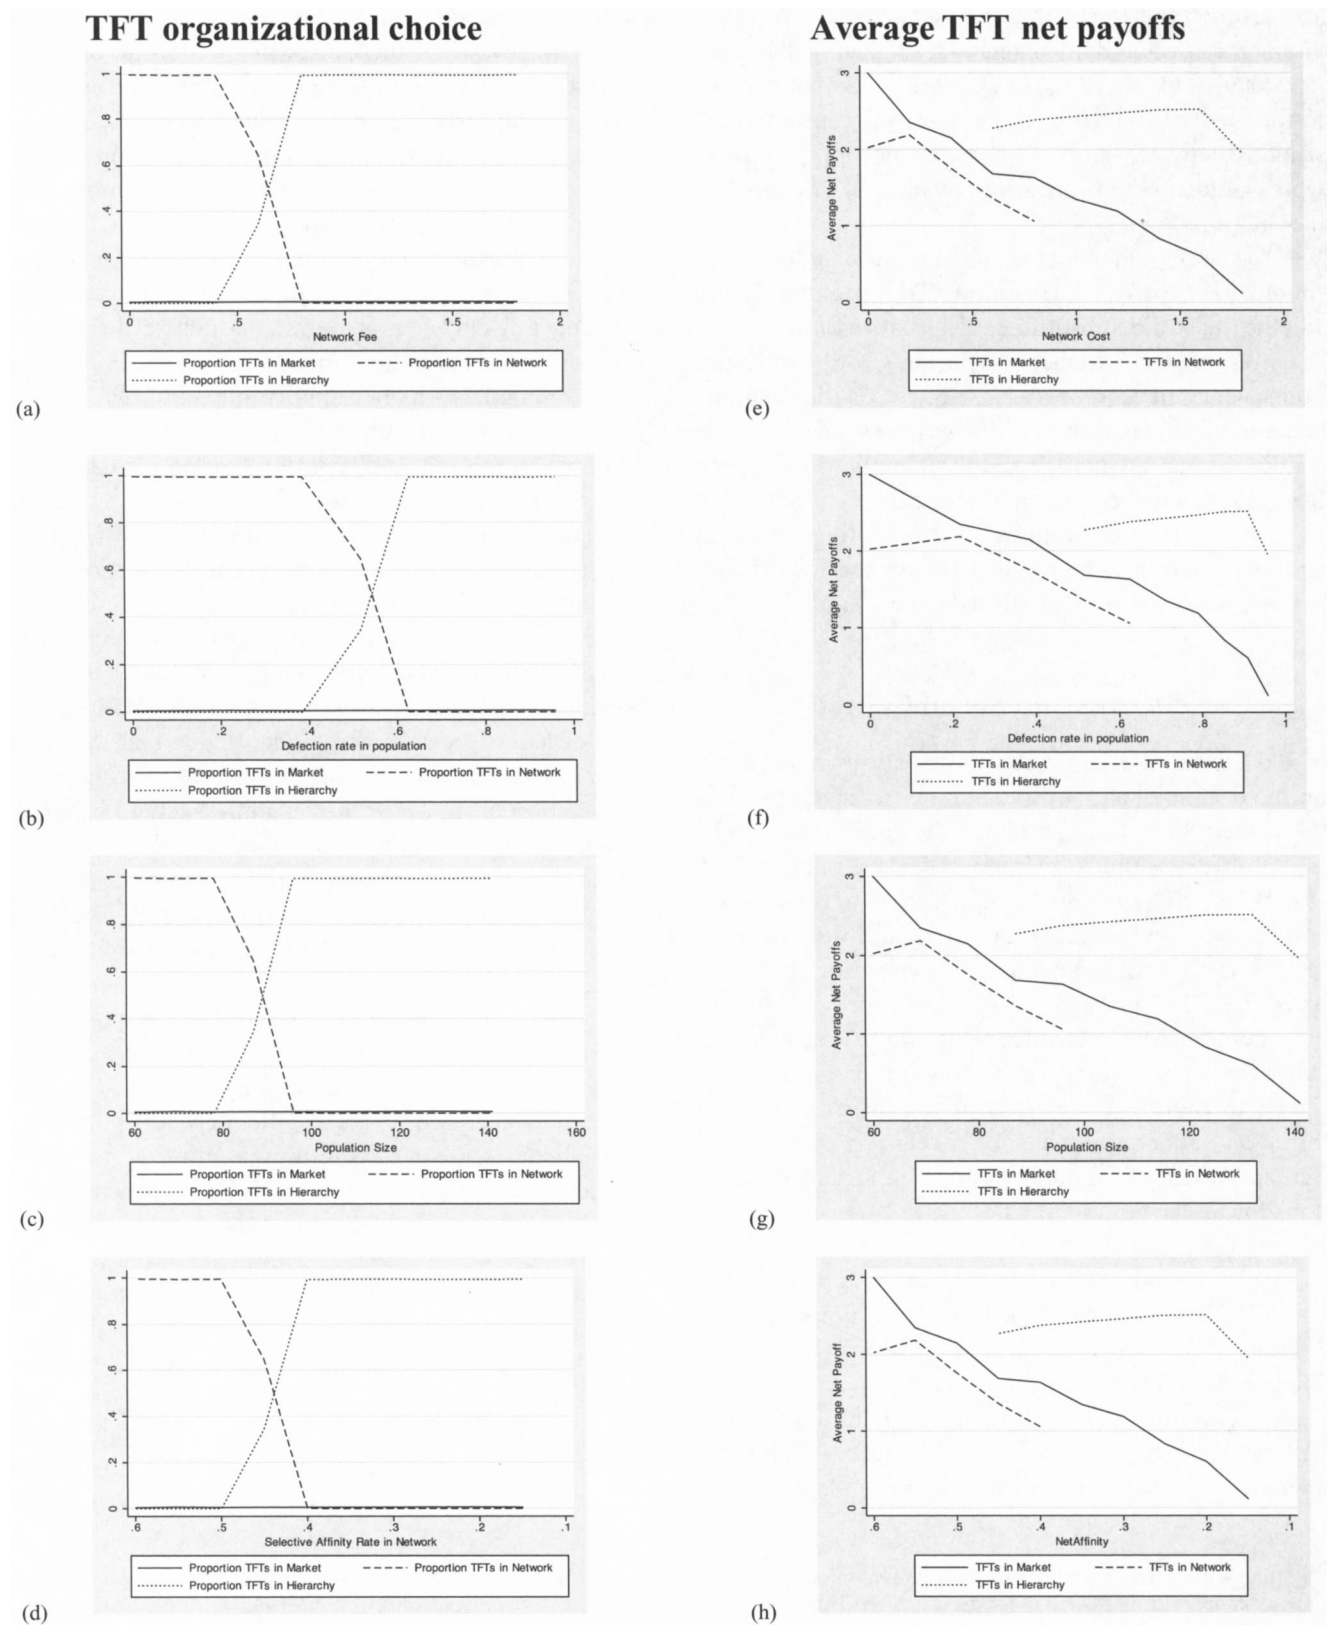
First, several of the causes of the decline of social capital identified by Putnam can be represented as an increase in the costs of joining a network (see Figure 3a). Specifically, the pressures of time, money, and suburbanization, and the pull of electronic entertainment, all of which Putnam cites, can be understood as increasing the opportunity costs of networking. As Americans work longer hours to earn more money while commuting longer distances and face more attractive alternatives for their shrinking leisure time, the effort spent building social capital has a higher opportunity cost. Our cost of joining a network captures this opportunity cost directly. Although the result is straightforward and predictable, it is consistent with Putnam’s description of change in American society over the last decades. As the cost of joining a social network increases, agents of all types leave the network and join the hierarchy.²⁶

Second, the decline of social capital, or the perception of decline, can be represented by increasing the proportion of nasty strategy types in the population (see Figure 3b). As social capital erodes, individuals perceive others as less trustworthy and less likely to reciprocate cooperation or, in our terms, as more likely to be nasty strategy types. Although in our model the beliefs of agents will eventually converge on the true distribution of strategy types in the population, for any given agent its beliefs are the product of its “lived” experience of interacting with other agents. This is, we believe, a close analog to the perceptions of individuals about the changing social world they inhabit. As the proportion of nasty strategy types in the population increases, TFTs leave the network and join the hierarchy to protect themselves through centralized enforcement. Less intuitively, nicer agents leave the network first (even before TFTs) and nastier agents, who benefit from the ability to exploit others both in the market and via selective affinity, leave the network last (see Figure 4a). This again shows the dark side of networks and is not predicted by Putnam or others.

Third, the increasing opportunities for interaction with others through both an increasingly integrated national market and declining transportation and communication costs, implicit in Putnam, erode the utility of social capital. While it may be possible to know

²⁵ Putnam’s empirical claim that social capital has declined in the United States is contested. For an example, see Paxton (1999). We sidestep this debate and, for purposes of this analysis, accept Putnam’s general description as correct. Our interest is more on the theoretical side in clarifying whether his conclusions about the effects of declining social capital actually follow from this theory.

²⁶ We present only the multidimensional simulations here. The effects here are quite distinct from the single-dimension, comparative static results presented in the appendix. As we increase all four parameters concurrently, hierarchy becomes the dominant organizational form at lower values of each parameter and displaces markets as the likely alternative.

FIGURE 3 Multidimensional Simulation of Declining Social Capital

The figures above illustrate round 10 of each of four representations of the decline of social capital (left panel) and the welfare effects (right panel). Each figure displays the effects of varying one parameter while the other three parameters are varied as well.

everyone within a small community—or at least know someone who knows someone who knows the relevant individual—this is increasingly difficult to maintain as individuals are pulled by opportunities outside that community. We can represent this increasing opportunity structure, as above, as an increase in the population of agents (see Figures 3c, 2b/d). We see again the same pattern of a growing population leading TFTs to leave the network and join the hierarchy.

Finally, a decline in social capital may also be associated with a reduced ability to select partners for interaction within the network. As with population, in a larger and more integrated society the ability to select particular agents with whom you have interacted in the past may erode as opportunities expand. As selective affinity declines, agents again leave the network more rapidly (Figures 3d, 2c).

All four ways of modeling the decline of social capital point in the same direction: contingent strategy types begin in the network and then move to the hierarchy.²⁷ Agreeing with the general pattern Putnam outlines, the model nonetheless helps to resolve a key tension in his analysis. To explain why social capital has declined in the United States, Putnam examines, as we have seen, the effects of longer working hours, suburbanization, electronic media, and other causes. Putnam is restrained in his conclusions on these possible causal variables, however, because all appear to be gradual and incremental changes but the decline in civic engagement is sharp and dramatic. Our representations of social capital, on the other hand, have nonmonotonic effects on network membership that produce a phase shift brought on by very small changes in the relevant parameters. These effects follow from population dynamics that interact with agent attributes to magnify the effect of changes on network joining. Putnam is trapped by his implicit assumption of monotonic effects. Even in our simple depiction of a social system, the interactions are sufficient to create phase shifts in network membership.

Although resolving this key empirical puzzle, our analysis also suggests that Putnam's alternatives to social networks are drawn too narrowly. Although he is correct to see markets as an alternative to networks, hierarchy is also an option, and increasingly so as network costs, the defection rate in the population, and population size increase and selective affinity decreases. Indeed, if the

population is sufficiently nasty (see Figure 3b), nearly all TFTs will leave the network for the benefits of centralized enforcement in the hierarchy and remain there. This may be what we are witnessing in the United States today. Accepting Putnam's description of the decline of social capital, we see individuals insulating themselves from opportunism by turning to the centralized, legal enforcement mechanisms of the state. Rather than relying on a personal relationship with a local business owner, for example, bankers today depend upon standardized credit reports, contracts, and legal penalties for breaches. As is frequently observed, the United States has become a significantly more litigious society. One way to interpret this is that networks are being displaced by various forms of hierarchy. If such interactive processes as represented in Figure 3 are at work in the United States today, this may explain why and how "small town" America has given way so dramatically to a legalized form of enforcement over the last generation.

What then are the welfare implications of these changing organizational ecologies? Putnam clearly expects a world of markets or hierarchy to produce less welfare for individuals and society than a world of social networks. It is not just nostalgia that leads him to highlight the virtues of social capital, but a fear that markets or, by our extension, hierarchy will leave all less well off than in the past. In all of our representations, however, a similar pattern emerges in which payoffs in hierarchy are higher than in markets and payoffs in markets are generally higher than in networks (see Figures 3e–h), although the small number of agents selecting markets in this simulation cause payoffs for market interactions to be especially unstable. The welfare benefits of hierarchy are clearly inconsistent with Putnam's expectations about the effects of declining social capital. In Section IV of the book, Putnam describes a variety of ways in which the welfare of Americans has declined as social capital has decayed. Our analysis suggests that agents may be better off under hierarchy in a world with less social capital than in networks in a world with greater social capital. Participating in a network is costly. Hours spent in a bowling league cultivating social ties and trust are hours not spent doing something else—including time with one's family or possibly acquiring greater human capital. Moreover, in worlds with very little social capital, agents are driven into the hierarchy where they are then subject to punishment for defecting on other members of the hierarchy, thereby creating a virtuous circle that leaves members better off, paradoxically, than they are in worlds with more social capital. The alternatives to social networks are not only an anomic market of declining of cooperation, but also a civil society in which mutually

²⁷ In the single-parameter representations, included in the online appendix, TFTs typically move from the network to the market and then to the hierarchy at higher values. The effect of changing all four parameters simultaneously is to make the world significantly more "harsh" and to drive agents directly into the hierarchy, bypassing the intermediate stage of the market.

beneficial cooperation is enforced by the threat of centralized punishment.

Hobbesian Hierarchy, or Why Don't Men Rebel?

Hierarchy is common in social life. It has been explained as an innate characteristic of individuals or societies, a function of initial social inequalities, a form of socially constructed power relations or an institutional solution to collective action and contracting dilemmas (for a review, see Lake 2009b). Our ABM suggests an approach in which hierarchy is an emergent property of the choices of many egoistic actors. As already indicated in the discussion of social capital, given a sufficiently nasty population, agents join a hierarchy and submit to its possible punishments in order to secure the benefits of cooperation it facilitates. By enforcing cooperation between agents, hierarchy improves their expected utility such that they chose to subordinate themselves to third-party rule. Our ABM is, in some ways, a computational representation of Thomas Hobbes's classic argument for *Leviathan*. Nonetheless, the model has several surprising implications.

First, counterintuitively, as a population becomes nastier, it is the nicer types of agents who join the hierarchy first, and the nastiest types who join last (see Figure 4a). A naïve expectation might posit that the ALLD agents would join the hierarchy first, as this is the only way they can escape mutual defection with one another and informed TFTs.²⁸ However, there is another, countervailing process occurring simultaneously. Left to otherwise fend for themselves in the market, ALLC types are increasingly exploited by ALLDs as the latter increase as a proportion of the population. ALLCs join the hierarchy not because they are uncooperative players but precisely because they no longer have sufficient opportunities to interact with other cooperative agents. TFT types draw upon the information in the network and then their own knowledge of other agents to protect themselves from being suckered by ALLD types. Less vulnerable to exploitation, TFTs "hold out" until the population gets even nastier but eventually join the hierarchy as well. In contrast to the naïve expectation, ALLDs are the last type of agent to join the hierarchy because they benefit from exploiting others in the market. In the end, for at least the most optimistic ALLDs in our simulations, the expected benefits of defecting on the ALLCs outweigh the gains they would otherwise anticipate from cooperating under hierarchy.

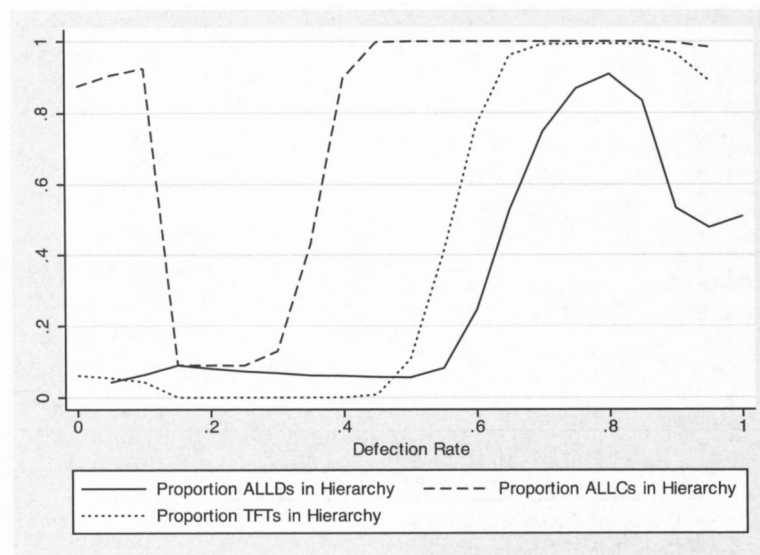
²⁸ Orbell and Dawes's (1991) work on the cognitive miser might also produce this intuition, but because defectors will expect defection.

This pattern is magnified the higher the rate of selective affinity. Paradoxically, in an increasingly nasty world, net payoffs increase on average for ALLC and ALLD agents and remain relatively constant for TFTs (see Figure 4b). Interestingly, even the ALLD agents achieve their highest average payoffs when the majority of them enter the hierarchy. This is, again, consistent with Hobbes's view that individuals subordinate themselves to the Leviathan to escape the state of nature and improve their welfare.


Second, the ABM also explains why hierarchies can be stable over long periods. By design in the ABM and by analogy to the real world, within a hierarchy agents do not learn anything about the strategy types or ideal points of other agents in the hierarchy. If both are in the hierarchy and agent j cooperates with agent i , i cannot learn whether j cooperated because it "wanted to" or did so only under threat of punishment. Having joined the hierarchy because it believed the population was sufficiently nasty, i then has fewer opportunities to revise its beliefs. Perversely, given these fewer opportunities, agent i will actually develop more skewed beliefs that lead it, over subsequent rounds, to believe the population is nastier than it really is, reinforcing its initial choice of hierarchy. Unable to learn from others in the hierarchy, agent i nonetheless continues to interact with randomly paired others in the market and, as we see in Figure 4, these others are likely to be disproportionately nasty, leading i to update its beliefs with increasing bias. In this way, agents and, in the real world, individuals get locked into hierarchy and become complicit in the perpetuation of their own subordination.

The emergence of hierarchy may be most counterintuitive when the hierarchy is autocratic, or when the hierarchy has an ideal point that is "extreme" within the population and, by analogy, cannot stay in power simply because he reflects broadly shared preferences. It is on this point that our modification to the standard PD game (see Figure 1 and discussion above) becomes perhaps most important. Autocracy is one of the great, unexplored frontiers of political science. Although there are many insightful and informative case studies, they have largely failed to cumulate into a theory that explains when autocracies are likely to arise and why they persist. General theories of autocracy, to the extent that they exist, fall into at least one of three approaches. The first treats autocracy as a default condition with analysis focusing on the fragile nature of democracy and the determinants of successful democratic transitions (see Przeworski 1991; Przeworski et al. 2000). A second approach, exemplified by the selectorate model of Bueno de Mesquita et al. (2003), focuses on the means by which a ruler satisfies a minimum winning coalition. How the selectorate

Figure 4 Proportion of Different Strategy Types Joining the Hierarchy as the Population Becomes Increasingly Nasty



Panel a.

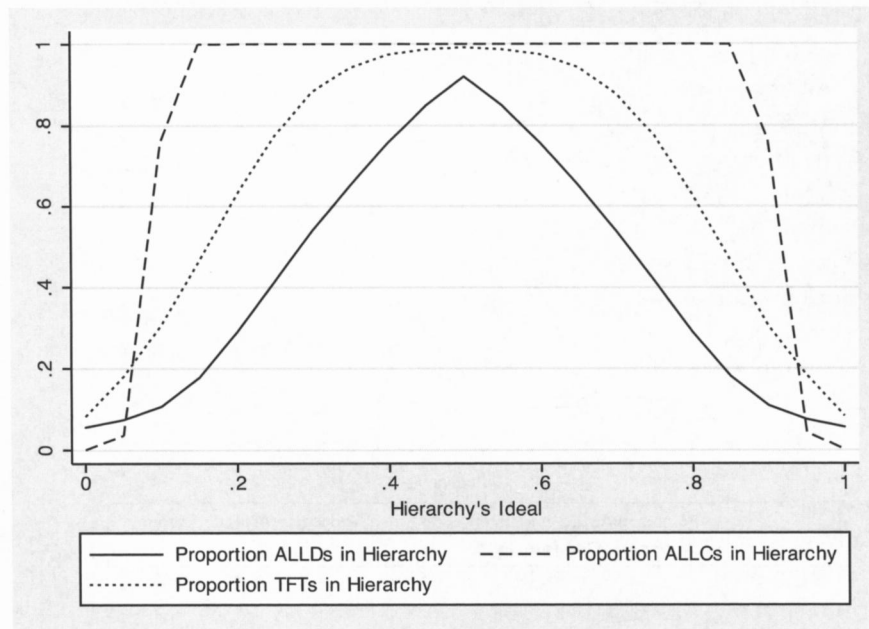


Panel b.

Panel a. Proportion of each strategy type choosing hierarchy. ALLCs are the first to join the hierarchy as the proportion of nasty strategy types increases in the population, followed by TFTs and then the most pessimistic ALLDs. After all of the “nice” agents are in the hierarchy, some ALLD agents return to the market to try to sucker the ALLCs and TFTs.

Panel b. Net payoffs by strategy type in hierarchy. Average net payoffs for all agents increase as the proportion of agents in the hierarchy begins to rise. The decline in payoffs in the nastiest populations is due to the mutual defection of ALLDs that have moved from the hierarchy back to the market.

FIGURE 5 Proportion of Each Strategy Type Joining the Hierarchy as the Hierarchy's Ideal Point Varies



This is a relatively nasty population of 70 ALLDs, 10 ALLCs, and 20 TFTs, corresponding roughly to the point where all TFT agents enter the hierarchy in Figure 4. Ideal points of agents are normally distributed with a mean of 0.5. As the hierarchy's ideal point moves toward either extreme (closer to zero or closer to one), fewer agents join the hierarchy. Importantly, however, except for very extreme values, agents still join the hierarchy for the cooperation it facilitates.

succeeds in deterring challenges or even revolution “from below” is left implicit.²⁹ A third approach posits that the masses who might otherwise rebel are repressed by the coercive power of the state (see Wintrobe 1998). In this view, successful autocrats divide and conquer the subject population to thwart collective action, promote false ideological and normative appeals to persuade individuals that others support the government, and repress dissidents who might otherwise rally the masses to stand up to the regime. All these approaches agree, however, that autocrats do not rely on popular support and reflect the political preferences of a smaller group within the population. We represent different regime types in our ABM by varying the hierarchy's ideal point. The more “extreme” the hierarchy's preference relative to the population, the more “autocratic” the ruler is likely to be.

Our model, in turn, suggests that autocracy can emerge due to the cooperation it facilitates even when levels of distrust within the population are high. In Figure 5, agents and especially the ALLC strategy types join

the hierarchy unless its ideal point is very far from the median. In this model, individuals choose not to exit the hierarchy or “rebel”—and indeed, voluntarily subordinate themselves to a hierarchy even with extreme preferences—because the coercive power of the state is believed to be the only mechanism for ensuring cooperation in a sufficiently nasty population. This implies autocracy is most likely when many individuals believe others will exploit them or do not trust one another to cooperate in market or even networked interactions. In other words, hostile environments in which agents are sincerely nasty or believed to be nasty are likely to be organized as autocracies. This further implies that autocratic hierarchs drive wedges between individuals and groups not to suppress collective action, as traditionally understood in the notion of divide and conquer, but to exacerbate the lack of trust otherwise necessary for self-enforcing cooperation in markets or networks.

Conclusion

Like all theories and models, ABMs are only as useful as the empirically supported, nonobvious propositions

²⁹ For an effort to deal with this problem within the same model, see Bueno de Mesquita and Smith (2009). For another variant, see Acemoglu and Robinson (2006).

they generate. In this article, we limit our empirical applications to the established work of others. The obvious propositions generated by the model largely serve to validate our ideal types of markets, hierarchies, and networks and, equally, our implementation. The nonobvious propositions show the promise of the ABM and, especially, the value in studying population dynamics. Striking in our view is the declining utility of networks over time and in large populations. To our minds, the biggest “surprise” of the model is that as the population becomes nastier, agents of all types are more likely to enter the hierarchy and, furthermore, nicer types will enter the hierarchy before nastier types. Possibly obvious once stated, this was certainly not a proposition that we anticipated before developing the basic architecture of the model. Indeed, it was not until we ran the model and saw this consistent pattern that we understood the exploitation that occurs in the market and explains this result.

Building off a relatively simple conceptualization of cooperation has produced new and, we think, important insights into the conditions under which networks are preferred organizations. These insights, we believe, reveal assumptions about networks left implicit in existing literatures on networks and social capital. This same conceptualization offers a fresh if disturbing perspective on the emergence of hierarchy in nasty populations and provides a new explanation for the persistence of autocracy.

A key but also unexpected finding of the model is that different organizational forms will often coexist across a broad range of parameter values. That is, different agents (even of the same strategy type) will join markets, hierarchies, or networks at sufficient rates to sustain multiple forms of organization simultaneously. Indeed, it is only under relatively extreme values of the parameters that one organizational form ever triumphs over the others. This suggests that research ought to shift, first, from assessing the superiority of markets, hierarchies, and networks to determining superiority for whom, when, and why and, second, from organizations to organizational ecologies so as to understand how different forms complement, compete, and survive in different populations and environments.

References

Acemoglu, Daron, and James Robinson. 2006. *Economic Origins of Dictatorship and Democracy*. New York: Cambridge University Press.

Arquilla, John, and David Ronfeldt. 2002. *Networks and Networks: The Future of Terror, Crime, and Militancy*. Santa Monica, CA: RAND.

Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books.

Axelrod, Robert. 1997. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton, NJ: Princeton University Press.

Bueno de Mesquita, Bruce, and Alastair Smith. 2009. “Political Survival and Endogenous Institutional Change.” *Comparative Political Studies* 42(2): 167–97.

Bueno de Mesquita, Bruce, Alastair Smith, Randolph M. Siverson, and James D. Morrow. 2003. *The Logic of Political Survival*. Cambridge, MA: MIT Press.

Cederman, Lars-Erik. 1997. *Emergent Actors in World Politics: How States and Nations Develop and Dissolve*. Princeton, NJ: Princeton University Press.

Chang, Myong-Hun, and Joseph E. Harrington Jr. 2006. “Agent-Based Models of Organizations.” In *Handbook of Computational Economics: Vol. 2. Agent-Based Computational Economics*, ed. Leigh Tesfatsion and Kenneth L. Judd. New York: Elsevier, 1273–1337.

Cohen, Michael D., Rick L. Riolo, and Robert Axelrod. 2001. “The Role of Social Structure in the Maintenance of Cooperative Regimes.” *Rationality and Society* 13(1): 5–32.

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

Downs, George W., David M. Rocke, and Randolph M. Siverson. 1986. “Arms Races and Cooperation.” In *Cooperation under Anarchy*, ed. K. Oye. Princeton, NJ: Princeton University Press, 118–46.

Eilstrup-Sangiovanni, Mette. 2009. “Varieties of Cooperation: Government Networks in International Security.” In *Networked Politics: Agency, Power, and Governance*, ed. M. Kahler. Ithaca, NY: Cornell University Press, 194–227.

Enemark, Daniel, Mathew McCubbins, Ramamohan Paturi, and Nicholas Weller. 2011. “Does More Connectivity Help Groups to Solve Social Problems?” *Proceedings of the ACM Conference of Electronic Commerce 2011*, 21–27.

Epstein, Joshua M. 2007. *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton, NJ: Princeton University Press.

Florini, Ann. 2003. *The Coming Democracy: New Rules for Running a New World*. Washington, DC: Island Press.

Greif, Avner. 2006. *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*. New York: Cambridge 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): 367–98.

Ilgen, Daniel R., and Charles L. Hulin, eds. 2000. *Computational Modeling of Behavior in Organizations: The Third Scientific Discipline*. Washington, DC: American Psychological Association.

Jackson, Matthew O. 2008. *Social and Economic Networks*. Princeton, NJ: Princeton University Press.

Kahler, Miles, ed. 2009a. *Networked Politics: Agency, Power, and Governance*. Ithaca, NY: Cornell University Press.

Kahler, Miles. 2009b. “Collective Action and Clandestine Networks: The Case of al Qaeda.” In *Networked Politics: Agency, Power, and Governance*, ed. M. Kahler. Ithaca, NY: Cornell University Press, 103–24.

- Kahler, Miles, and David A. Lake. 2009. "Economic Integration and Global Governance: Why So Little Supranationalism?" In *The Politics of Global Regulation*, ed. N. Woods and W. Mattli. Princeton, NJ: Princeton University Press, 242–75.
- Keck, Margaret E., and Kathryn Sikkink. 1998. *Activists Beyond Borders: Advocacy Networks in International Politics*. Ithaca, NY: Cornell University Press.
- Kenney, Michael. 2009. "Turning to the 'Dark Side': Coordination, Exchange, and Learning in Criminal Networks." In *Networked Politics: Agency, Power, and Governance*, ed. M. Kahler. Ithaca, NY: Cornell University Press, 79–102.
- Kollman, Ken, John H. Miller, and Scott E. Page, eds. 2003. *Computational Models in Political Economy*. Cambridge, MA: MIT Press.
- Lincoln, James R., and Michael L. Gerlach. 2004. *Japan's Network Economy: Structure, Persistence, and Change*. New York: Cambridge University Press.
- Lake, David A. 2009a. *Hierarchy in International Relations*. Ithaca, NY: Cornell University Press.
- Lake, David A. 2009b. "Hobbesian Hierarchy: The Political Economy of Political Organization." *Annual Review of Political Science* 12: 263–83.
- McCubbins, Mathew, Ramamohan Paturi, and Nicholas Weller. 2009. "Connected Coordination: Network Structure and Group Coordination." *American Politics Research* 37(5): 899–920.
- Milgrom, Paul R., Douglass C. North, and Barry R. Weingast. 1990. "The Role of Institutions in the Revival of Trade: The Medieval Law Merchant, Private Judges, and the Champagne Fairs." *Economics and Politics* 2(1): 1–23.
- Miller, John H., and Scott E. Page. 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton, NJ: Princeton University Press.
- Nowak, Martin A., and Karl Sigmund. 2005. "Evolution of Indirect Reciprocity." *Nature* 437: 1291–98.
- Orbell, John, and Robyn M. Dawes. 1991. "A 'Cognitive Miser' Theory of Cooperators' Advantage." *American Political Science Review* 85(2): 515–28.
- Paxton, Pamela. 1999. "Is Social Capital Declining in the United States? A Multiple Indicator Assessment." *American Journal of Sociology* 105(1): 88–127.
- 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.
- Prietula, Michael J., Kathleen M. Carley, and Les Gasser, eds. 1998. *Simulating Organizations: Computational Models of Institutions and Groups*. Cambridge, MA: MIT Press.
- Przeworski, Adam. 1991. *Democracy and the Market: Political and Economic Reforms in Eastern Europe and Latin America*. New York: Cambridge University Press.
- Przeworski, Adam, Michael E. Alvarez, Jose Antonio Cheibub, and Fernando Limongi. 2000. *Democracy and Development: Political Institutions and Well-Being in the World*. New York: Cambridge University Press.
- Putnam, Robert D. 2000. *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon & Schuster.
- Rauch, James E., and Gary G. Hamilton. 2001. "Networks and Markets: Concepts for Bridging Disciplines." In *Networks and Markets*, ed. J. E. Rauch and A. Casella. New York: Russell Sage Foundation, 1–29.
- Rauch, James E., and Victor Trindade. 2002. "Ethnic Chinese Networks in International Trade." *Review of Economics and Statistics* 84(1): 116–30.
- Raustiala, Kal. 2002. "The Architecture of International Cooperation: Transgovernmental Networks and the Future of International Law." *Virginia Journal of International Law* 43 (1): 1–92.
- Ronfeldt, David. 1996. *Tribes, Institutions, Markets, and Networks: A Framework about Social Evolution*. RAND Paper P-7969. Santa Monica, CA.
- Sageman, Marc. 2008. *Leaderless Jihad*. Philadelphia: University of Pennsylvania Press.
- Slaughter, Anne-Marie. 2004. *A New World Order*. Princeton, NJ: Princeton University Press.
- Spruyt, Hendrik. 1994. *The Sovereign State and Its Competitors: An Analysis of Systems Change*. Princeton, NJ: Princeton University Press.
- Swedberg, Richard. 2003. "Economic and Sociological Approaches to Markets." In *Principles of Economic Sociology*, ed. Richard Swedberg. Princeton, NJ: Princeton University Press, 104–30.
- Tarrow, Sidney. 2005. *The New Transnational Activism*. New York: Cambridge University Press.
- 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.
- Wintrobe, Ronald. 1998. *The Political Economy of Dictatorship*. New York: Cambridge University Press.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

- Appendix A presents additional information about the simulations presented in the article.
- Appendix B is a supplement that describes in greater detail the architecture of the model and provides additional parameter scans to assess the robustness of the conclusions.

Please note: Wiley-Blackwell is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.