Chapter 7

The Contagion of Prosocial Behavior and the Emergence of Voluntary-Contribution Communities

Milena Tsvetkova and Michael Macy

Abstract Every day, millions of people write online restaurant reviews, leave product ratings, provide answers to an unknown user's question, or contribute lines of code to open-source software, all without any direct reward or recognition. People help strangers offline as well, as when people anonymously donate blood or stop to help a stranded motorist, but these behaviors are relatively rare compared to the pervasiveness of online communities based on user-generated content. Why are mutual-help communities far more common online than in traditional offline settings that are not mediated by the Internet? We address this puzzle in two steps. We begin with empirical evidence from an online experiment that tests two mechanisms for the contagion of helping behavior: "generalized reciprocity" and "third-party imitation". We then use an empirically-calibrated agent-based model to show how these mechanisms interact with the rivalness of contributions, that is, the extent to which the benefit from a contribution is limited to just one beneficiary (as when helping a stranded motorist) or benefits many people at once (as when contributing a product review online). The results suggest that the non-rivalness of most user-generated content provides a plausible explanation for the rapid diffusion of helping behavior in online communities

Milena Tsvetkova
Department of Sociology, Cornell University, 372 Uris Hall Ithaca, New York 14853-7601 e-mail: mvt9@cornell.edu

Michael Macy
Department of Sociology, Department of Information Science, Social Dynamics Laboratory, Cornell University, 372 Uris Hall Ithaca, New York 14853-7601 e-mail: m.macy@cornell.edu

7.1 Introduction

The health regime we follow [1], the music we listen to [2], the new technologies we adopt [3], the news stories we read [4], and even the likelihood that we vote in an election [5] are all to a large degree influenced by our friends and peers. Many human behaviors spread through social contact, including some that are often assumed to be acquired independently, such as obesity and fertility [6].

Prosocial behavior has also been shown to be contagious. Fowler and Christakis [7] found experimental evidence that if you help someone, you not only increase the likelihood that they help others, but that those they help will also help others, and so on, out to three steps. Suri and Watts [8] and Jordan et al. [9] similarly found that generous behavior was contagious at least in direct interaction. These groundbreaking studies have provoked new questions. What are the mechanisms through which prosocial behavior spreads among strangers? How do these mechanisms affect the contagion dynamics? Can they lead to the emergence of cooperation in an initially non-cooperating population?

7.1.1 The puzzle of online generosity

The puzzle of contagious generosity is compounded further by the emergence of online communities with user-generated content, from open source software development to advice forums to Wikipedia [10]. Why are mutual-help communities far more common online than in traditional offline settings that are not mediated by the internet?

We address this puzzle using an empirically calibrated agent based model. The results suggest that the answer may lie in the differences in the rivalness of online and offline public goods involving anonymous contribution. Many offline public goods – like blood donation, charities, and giving up one's seat – are rivalrous, meaning that the contribution transfers resources from the giver to a particular receiver. In contrast, many online public goods, especially in communities based on user-generated content, are non-rival – everyone in the community can benefit from a given contribution. The difference is not limited to the effect of non-rival incentives on the independent probability of contribution by a member of the community. Computer simulation shows that this "within individual" difference is amplified by the "between individual" effects of the contagion dynamics. More precisely, we identify two mechanisms of contagion – "generalized reciprocity" and "third-party imitation" – and show how these mechanisms interact with differences between rival and non-rival contributions to explain the spread of helping behavior in online communities.

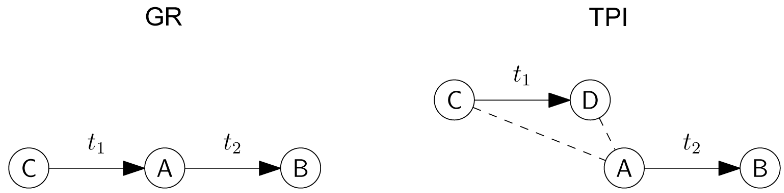

Fig. 7.1: Two mechanisms for the contagion of prosocial behavior.(GR) Generalized reciprocity: A helps B because C has helped A. (TPI) Third-party influence: A helps B because A has observed C help D.

7.1.2 Outline of a theory of prosocial contagion

Previous research has suggested that there are two distinct mechanisms for the contagion of prosocial behavior among strangers: generalized reciprocity and third-party influence. Generalized Reciprocity (GR) refers to cases in which those who *benefit* from a stranger's prosocial behavior behave more prosocially towards another in the future. As diagrammed in Fig. 7.1, A helps B because C has helped A [11, 12]. Third-party influence (TPI) refers to cases in which those who *observe* prosocial behavior by strangers behave more prosocially towards a stranger: A helps B because A has seen C help D. GR characterizes "pay it forward" behavior triggered by a normative or affective response to being helped [13], while TPI characterizes social learning through imitation of others' behavior.

GR and TPI also differ in the pattern of transmission. GR transmits the contagion from person to person through direct contact and hence its contagious effect is constrained to the chain of those who were previously helped. In contrast, TPI has the potential to broadcast the contagion from one person to any number of observers. The interaction of the two mechanisms could generate a powerful self-reinforcing dynamic that dramatically increases the rate of prosocial behavior in an initially uncooperative population.

In this chapter, we summarize an online experiment that distinguished between the behavioral effects of the two contagion mechanisms [14] and use an agent-based model to investigate the contagion dynamics and the population-level outcomes that they entail. The results show that receiving help can increase the willingness to be generous towards others, but observing help can have the opposite effect, particularly among those who have not received help. We use a threshold model with dynamic interaction structure and adaptive behavior to simulate a population of agents with this behavior. The computational experiments indicate that the agents can self-organize in communities based on voluntary contributions in two possible ways. On the one hand, when contributions are rival, a handful of altruists can lead to the emergence of small clusters of contributors as long as agents observe contribution beneficiaries in a relatively large radius (for example, via gossip) and unsatisfied agents are not too mobile. On the other hand, when contributions are non-rival, communities are much more likely to emerge and the level of contributions is higher when agents observe contributors rather than recipients. These two pathways roughly correspond to offline and online interactions. They offer explanation for the fact that cultures of kindness are rare for anonymous face-to-face interactions but common on the Web, for example, in the form of communities based on user generated content.

7.2 Testing individual mechanisms

Causal mechanisms are notoriously difficult to observe in natural settings, and controlled diffusion experiments with large groups are highly impractical in traditional laboratory settings. To test the two contagion mechanisms, we therefore designed and conducted a large behavioral experiment online. The experiment used anonymity to isolate the effects of GR and TPI from other cooperation-inducing mechanisms, including direct and indirect reciprocity, as well as peer pressure based on reputation effects. To isolate GR from TPI, we manipulated the extent to which participants received and observed help.

7.2.1 Online experiment

The study was designed as a sequential two-player investment/gift-exchange game in groups of 150 with random partner selection. In the game, a participant could choose to return part of their payment so that another anonymous participant could benefit.

We first recruited a pool of potential participants by posting a task on the online crowdsourcing platform Amazon Mechanical Turk (AMT). The task invited AMT users to sign up for a study that offered the chance to earn up to $14-21 for doing the same $2-3 ten-minute task multiple times. The AMT users were informed that they could only participate in the task and earn the promised amount if they were randomly selected from the pool of potential participants. Participants were eligible to be selected multiple times but there was no guarantee that they would be selected even once. If selected, the participant was to receive an e-mail notification with further instructions.

The email invitation informed recipients that they were randomly chosen to participate in the game, which they had to complete within 24 hours. Participants were then directed to our website, where they read a description of the game and made a single decision about whether to donate money to benefit a stranger. The game description explained to each participant that they would be paid the amount promised in the original solicitation, which included a “base” payment plus a “bonus” payment. Participants were also told that they were part of a group of 150 AMT users and that only members of this group who received an invitation could actually participate and receive the promised payment. The instructions further informed participants that the study had allocated a limited number of invitations to be distributed to randomly selected participants (“seeds”). The seeds were invited by the experimenters to participate. In addition to these invitations created by the experimenters, each participant who received and accepted an invitation had the option to create a new invitation and allow one more person to participate. However, in order to create a new invitation, the participant had to be willing to donate his or her bonus, even though this would reduce the participant’s earnings. If the participant chose to donate his or her bonus, a recipient of the new invitation (the “invitee”) would then be randomly selected from the other 149 AMT users in the group. The instructions explained further that when a participant donated his or her bonus, we supplemented the bonus amount so that the next invited participant received the same base payment and bonus and had the same options: to keep his or her bonus or donate it and create a new invitation for one more participant.

All participants knew that the person who receives the donated invitation would not know the identity of the participant who made the donation. Thus, anyone receiving a donated invitation was unable to directly reciprocate or to pass along a favorable reputation. We referred to participants by their AMT worker ID, randomly anonymized in a way that precluded the possibility to identify the same individual and be influenced by reputation.

The experiment involved five manipulations: whether the participant received a donated invitation created by another participant (i.e. being a “link”), the number of times the participant was invited to play the game (ranging from one to six), whether the participant was able to observe donated invitations, the number of donated invitations the participant observed (ranging from zero to 223), and the payment the participant received ($2 base rate and $1 bonus or $1 base rate and $1 bonus).

The observation and payment manipulations were crossed to define four between-individual treatment groups to which participants were randomly assigned. The number of invitations received and observed varied within individual. Further, some participants were only selected as seeds, others were only selected as invitees, and still others were selected as invitees after having been previously selected as seeds.

7.2.2 *Results*

After removing data from participants who did not demonstrate an adequate understanding of the instructions, we were left with 518 individuals and 1,070 observations. We used random-intercepts logistic regression models of observations nested in individuals to estimate the change in the odds of donating under the different manipulations. The models allow us to adjust for the non-independence of repeated measures and control for the effect of payment level and two other potential confounders, the time elapsed between subsequent interactions and the number of previous interactions. To better isolate the mechanisms, the models pool data only form the relevant treatment conditions: we tested GR in the no-observation condition only, we tested TPI for seeds only, and we tested the interaction of GR and TPI in the observation condition only.

Consistent with GR, participants were more likely to be generous towards a stranger after experiencing generosity (Table 1A). However, the effect is limited to the first receipt of generosity as the critical event in triggering GR. The odds of donating do not continue to increase with receiving additional donated invitations.

Consistent with TPI, there was a statistically significant increase in the odds of donating among the seeds who had observed between 0 and 75 donated invitations, compared to those who had not observed any (Table 1B). However, the level of donation among those who observed more than 75 invitations was not significantly greater than the baseline level. In other words, similarly to GR, the effect of TPI appears to be concave, with most of the effect evident at relatively low levels of observed donation and little subsequent change.

Less intuitively, the effect from observing widespread generosity is significantly different for those who have recently benefited from generosity compared to those who have not. When observing more than 75 donated invitations, the odds of donating decrease for seeds but do not change for invitees (Table 1C). This difference between seeds and invitees is statistically significant (χ^2 (1 df) $= 3.88$, $p = 0.049$ for observing $76-150$; χ^2 (1 df) $= 5.55$, $p = 0.019$ for observing 151+) and suggests the possibility that seeds succumb to a "free-riding" effect from which invitees are immune due to having been recipients of generosity. Free riding represents the temptation to refrain from contributions, especially when one becomes aware that others are already contributing. The behavior is common in collective-action situations [15] and is also known as social loafing [16] and as the "bystander effect" or "diffusion of responsibility" [17].

In sum, the experimental results show that receiving and observing generosity can significantly increase the likelihood to be generous towards a stranger. However, the willingness to contribute can be offset by lower perceived need when the level of helping is sufficiently high. This "bystander effect" is especially evident among those who have not themselves benefited from generosity. In other words, the norm to "be generous if that is what others are doing" weakens when the level of helping behavior is high, unless it interacts with the normative obligation to "pay it forward."

The implications of the effects of the two contagion mechanisms for the dynamics of helping cascades are not intuitively obvious. We therefore incorporated the empirical findings in an agent-based model to investigate the macro-level effects of GR and TPI.

7.3 Extrapolating to population outcomes

Our model is a threshold model of collective behavior. Such models have been previously used to study the emergence of collective action and the resolution of social dilemmas [18, 19, 20]. In this literature, a threshold is the critical number or proportion of contributors at which an individual becomes willing to contribute to a collective action or to join a collective behavior. Depending on the distribution of individual thresholds, cascades are possible in which each additional participant triggers participation by others. It has been established that the emergence of widespread participation critically depends on the composition of the population, and in particular, the existence of a critical mass of altruists, or unconditional contributors.

We model diffusion through the dynamics of selection and influence by relaxing two common assumptions in existing threshold models: fixed interaction structure and fixed individual interests in contributing. Our model assumes that agents both move in space (similarly to Ref. [21]) and adapt their behavior (similarly to Ref. [19, 22, 23, 24]). By combining dynamic interaction structure with adaptive behavior, our model is similar to evolutionary-game models on cooperation [25, 26, 27, 28, 29, 30]. In these models, agents choose an action or a strategy in the Prisoner's Dilemma and play it against each of their interaction neighbors. The agents update their behavior by imitating successful neighbors and find more beneficial interaction partners by moving on a spatial grid or rewiring their interaction network. In our model, agents play a gift game with a different number of their neighbors, depending on the rivalness of the exchanged gifts. Influence occurs not because agents imitate others but because they condition their behavior on others' behavior.

7.3.1 Simulation model

7.3.1.1 Assumptions

The model assumes that agents are heterogeneous with respect to their natural proclivity to condition their contributions on others' behavior and their own outcomes. These proclivities are exogenously predetermined and remain fixed throughout social interactions. In addition to generalized reciprocity, third-party influence, and free riding, the model assumes two other behavioral mechanisms: unconditional altruism and aspiration. Unconditional altruism captures the extent to which individuals are willing to help strangers regardless of others' behavior or their own outcomes. Aspiration is the expectation about the extent to which one should benefit from others' contributions. Aspiration is the benchmark against which the agent evaluates outcomes as satisfactory [31]. If outcomes are unsatisfactory, the agent can decide to move to a different community (similarly to Ref. [21]). We set the aspiration as $\theta_A \sim Uniform(0, 0.5)$.

Following previous research [24], agents are assigned a level of unconditional altruism that is randomly drawn from a beta distribution: $\theta_{UA} \sim Beta(\alpha, \beta)$. The model fixes $\alpha = 5$ and $\beta = 5$. The resulting distribution lacks a critical mass of altruists because the majority of individuals have values close to 0.5. This distribution matches the empirical distribution of behavioral types in the general population, characterized by few unconditional altruists (about 13%) and a majority of conditional contributors (50-63%; [32, 33]). Nevertheless, previous analytical work on deterministic threshold models in fixed populations has shown this type of distribution not to favor the emergence of high-levels of contribution [18, 24]. Compared to these earlier models, we start from a lower level of unconditional altruism that is more empirically plausible.

The model assumes that generalized reciprocity $GR \sim Uniform(0, 1)$ and third-party influence $TPI \sim Uniform(0, 1)$. The higher the value of GR (TPI) the more the agent's contribution behavior is sensitive to benefits received (observed). For consistency, the free-riding value is always at least as large as the unconditional-altruism value: $\theta_{FR} = \theta_{UA} + FR(1 - \theta_{UA})$, where $FR \sim Uniform(0, 1)$. The higher the value of FR, the lower the observed level of contribution at which the agent refrains from contributing in order to free-ride on others' effort.

The model also assumes that the interaction structure is a square lattice that wraps into a torus. This structure is characterized by a high average clustering, long average path-lengths, and regularity in network positions. The structure is a poor representation for persistent social relations such as friendships and business contacts. However, it is a suitable heuristic for interactions between strangers in geographical space. Further, we assume that an agent's interaction neighborhood does not entirely coincide with the agent's observation neighborhood. In both cases, the neighborhood is a Moore neighborhood (a square with the focal agent in the center) but the radius of the neighborhood can vary. A larger interaction neighborhood corresponds to a larger community size while a larger observation neighborhood corresponds to a higher degree of gossip or centralized broadcasting.

7.3.1.2 Behavioral rules

The five behavioral mechanisms come together in two separate threshold functions that determine whether agents contribute to a neighbor (or multiple neighbors) from

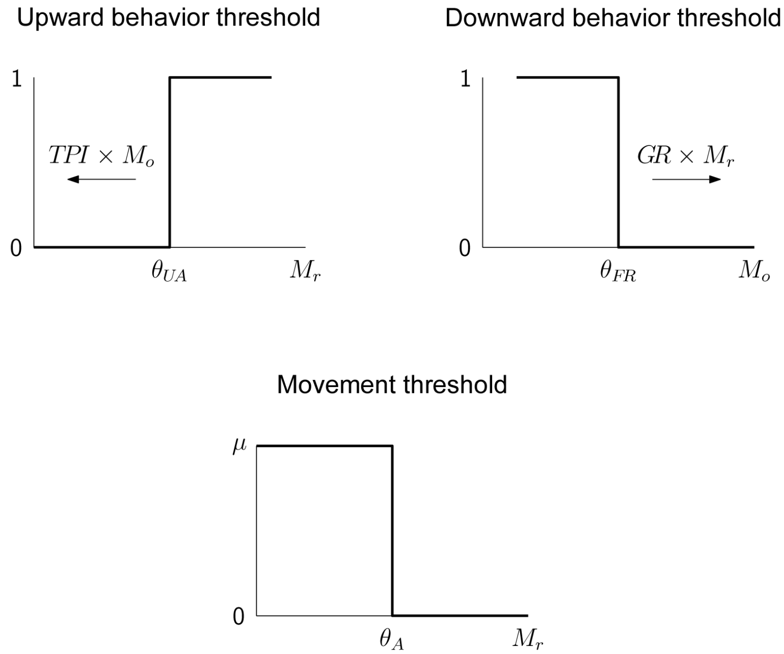

Fig. 7.2: Three thresholds in the simulation model. The upward behavior threshold depends on unconditional altruism (θ_{UA}) but can decrease due to third-party influence ($TPI \times M_o$). The downward behavior threshold depends on the proclivity to free ride (θ_{FR}) but can increase due to generalized reciprocity ($GR \times M_r$). The movement threshold depends on the aspiration (θ_A). The agent makes a contribution to the benefit of a random neighbor(s) within her interaction neighborhood if the contributions she remembers receiving (M_r) match or surpass her upward threshold but the contributions she remembers observing (M_o) do not exceed her downward threshold. The agent moves to a new empty site within her observation neighborhood if the contributions she remembers receiving (M_r) fall below her aspiration.

their interaction neighborhood and whether agents move to a new location in their observation neighborhood.

The contribution threshold models the combined effect from receiving and observing others' contributions on one's likelihood to contribute. As in the empirical results, benefiting from others' contributions increases one's likelihood to contribute and decreases one's likelihood to free-ride, while observing others' contributions could increase both one's likelihood to contribute and one's likelihood to free-ride. Following previous models of non-monotonic threshold functions [22, 23, 24], the function is characterized by two thresholds: an upward threshold $\theta_{0\to1}$ and a downward threshold $\theta_{1\to0}$. The agent contributes as long as the number of received and observed contributions is within these two thresholds. The upward threshold is predetermined by the agent's unconditional altruism but decreases if the agent experiences third-party influence. The downward threshold is anchored by the agents' proclivity to free ride but increases if the agent succumbs to generalized reciprocity (Fig. 7.2). More specifically:

$$\begin{aligned}\theta_{0\to 1}(t) &= \theta_{UA} - TPI \times M_o(t) \times \theta_{UA},\\ \theta_{1\to 0}(t) &= \theta_{FR} - GR \times M_r(t) \times (1-\theta_{FR}),\end{aligned} \tag{7.1}$$

where $M_r(t)$ is the number of contributions the agent remembers receiving and $M_o(t)$ is the proportion of contributions the agent remembers observing in her observation neighborhood. The agent makes a contribution to the benefit of a random neighbor(s) within her interaction neighborhood if the contributions she remembers receiving match or surpass her upward threshold but the contributions she remembers observing do not exceed her downward threshold:

- **Behavior Rule 1:** Contribute if $M_r(t) \geq \theta_{0\to 1}$ and $M_o(t) < \theta_{1\to 0}(t)$.

Similarly, the agent moves with probability μ (mobility) to a random empty site within her observation neighborhood if the contributions she remembers receiving do not match her aspiration:

- **Behavior Rule 2:** Move with probability μ if $M_r < \theta_A$.

Thus, agents who are satisfied with their outcomes tend to stick to the community they have found but unhappy agents tend to move to communities with higher levels of contribution. $M_r(t)$ and $M_o(t)$ are simply the number of contributions the agent received and the proportion of local contributions the agent observed in the previous m time periods, where m is the length of memory. More formally, $M_r(t) = \frac{\sum_{t-m}^{t-1} r_t}{m}$ and $M_o(t) = \frac{\sum_{t-m}^{t-1} o_t n_t^{-1}}{m}$, where r_t is the number of times the agent benefited from a contribution at time t, o_t is the number of contributions the agent observed at time t, and n_t is the size of the agent's neighborhood at time t. For the model, $m = 5$ was chosen because this value produced high variability in the results. Increasing constricts the conditions for emergence of contributions since more random events become necessary in an agent's neighborhood in order to convert that agent into a contributor.

Updating is synchronous for both the decision to contribute and to move. At each time period, agents are drawn in random order to decide whether to contribute, given the contributions they observed and the amount of contributions they received up until the last period. Once all agents have had the chance to update their behavior, the agents decide whether to move, given the amount of contributions they have received until the end of the current period. Thus, the model assumes that agents observe and receive contributions within each time period and then decide whether to contribute (Behavior Rule 1) and whether to leave a community (Behavior Rule 2). Since threshold models have been shown not to be robust to noise [34], the model assumes that there is a small probability $\varepsilon = 10^{-3}$ that an agent's contribution or movement decision is reversed.

7.3.1.3 Parameter space

To preclude sensitivity to initial conditions and synchronous updating, the model used behavioral and movement noise, the simulations were run for a sizable agent population, and the results were averaged over multiple repetitions. The fixed parameters in the model (the shape and the range of the distributions and the length of memory) were chosen with the goal to keep them as simple as possible while producing the highest variation in results along the variable parameters.

The computational experiments were run for a population of 1000 agents on a torus (40% occupied locations). The experiments investigated the average contribution level (i.e. the proportion of contributors) for two different levels of rivalness: we assume that rival contributions benefit one recipient, while non-rival contributions benefit 3 recipients. The effects of four parameters are explored:

- The mobility $\mu \in [0, 0.05, 0.5]$. This is the probability to move if the agent is unhappy with what she receives from the current community. This parameter represents community turnover. (Turnover could also be adjusted by varying the average aspiration θ_A.)
- The radius of the interaction neighborhood $\in [1,2,3,4,5,7,10,15]$. Since the model uses Moore neighborhoods, this is equivalent to a maximum of $[8,24,48,80,120,224,440,960]$ neighbors for each agent. This parameter corresponds to community size.
- The radius of the observation neighborhood $\in [0,1,2,3,4,5,6,10,15]$. Since the model uses Moore neighborhoods, this is equivalent to a maximum of $[0,8,24,48,80,120,224,440,960]$ neighbors for each agent. This parameter is related to gossip and centralized broadcasting.
- The observation targets $\in [recipients, contributors]$. Agents observe either the proportion of contributors or the proportion of beneficiaries within their observation neighborhood.

The simulations were run for 5000 periods which was sufficient for convergence to an equilibrium. The equilibrium proportion of contributors was then estimated by averaging the proportion of contributors over the last 1000 periods.The resulting equilibrium proportion of contributors was then averaged over 25 replications for each parameter combination.

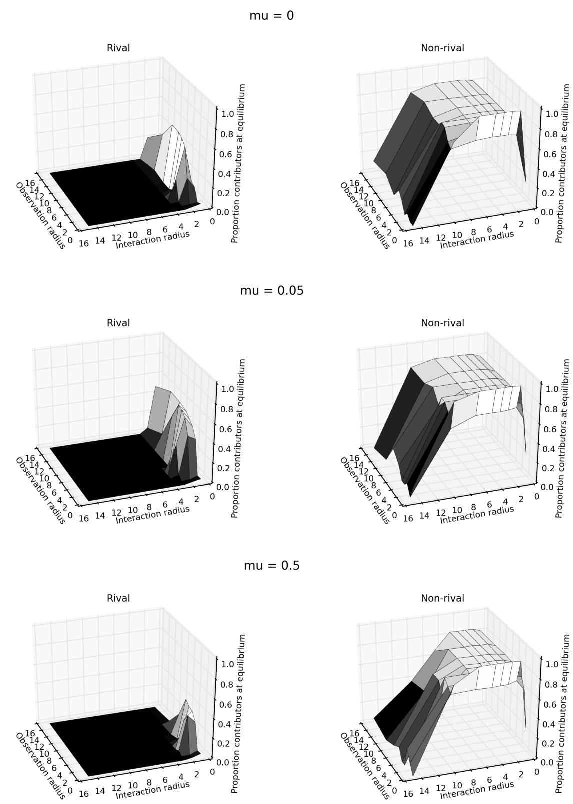

Fig. 7.3: The equilibrium proportion of contributors by observation radius and interaction radius for contributors as observation target and mobility $\mu = 0$ (top), $\mu = 0.05$ (middle), and $\mu = 0.5$ (bottom). Results are shown for rival (left) and non-rival (right) contributions.

7.3.2 Results

Fig. 7.3 shows that for non-rival contributions, the equilibrium level of contributing is visibly higher than for rival contributions. Further, for non-rival contributions, the conditions for the emergence of contribution-based communities are significantly less restricted.

When the exchanged contributions are non-rival, the global level of contribution is high over a large range of interaction radii. Widespread contribution fails to emerge only when the interaction radius and/or the observation radius are extremely large. This implies that non-rival exchange allows for relatively large contribution-based communities. For relatively large communities (interaction radius > 1), observed

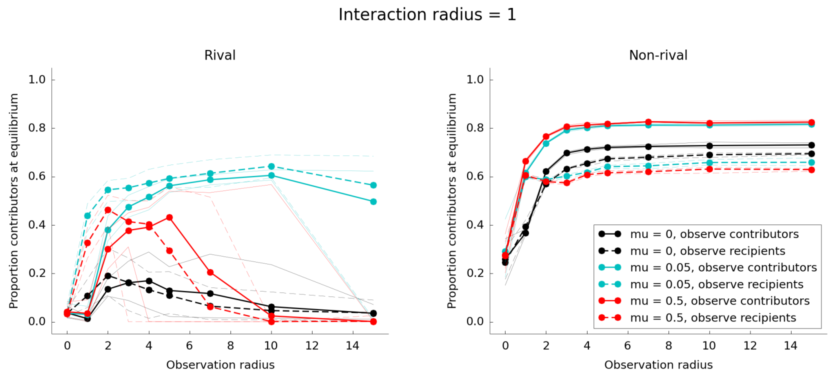

Fig. 7.4: The equilibrium proportion of contributors by observation radius when interaction is constrained to immediate neighbors only (interaction radius $= 1$). Line colors show levels of mobility and line types differentiate the observation target. Results are shown for rival (left) and non-rival (right) contributions. The thick lines show the proportion averaged over 25 replications for that particular parameter combination. The thin lines show the minimum and the maximum proportions achieved in the replications.

contribution has little effect, and 100% contribution is possible even when there is no observation (observation radius $= 0$). Overall, observing contributors has a greater effect than observing recipients (right column in Fig. 7.3 and 7.4). Community turnover does not affect outcomes except when the communities are small (interaction radius $= 1$) or when observation is widespread in large communities. In the first case, some mobility is better than no mobility (Fig. 7.4, right) and in the second case, too much mobility is bad (right column in Fig. 7.3).

When the exchanged contributions are rival, only small communities can have high levels of contribution (optimal interaction radius $\sim 2-3$; left column in Fig. 7.3). Further, observation is crucial for the emergence of contribution communities: the level of contribution is zero when there is no observation. As the observation radius increases, the level of contribution radically increases initially but eventually starts decreasing slowly (left in Fig. 7.4 and 7.5). The optimal observation radius is between 2 and 5, depending on the target of observation. Compared to observing contributors, observing recipients requires a smaller observation radius to achieve the maximum level of contribution. Finally, the effect of mobility is non-monotonic: low mobility ($\mu = 0.05$) is better than no mobility ($\mu = 0$) or too much mobility ($\mu = 0.5$).

Fig. 7.6 identifies the reason for differences between rival and non-rival contributions. Non-rivalness implies that a larger number of individuals can benefit from a single contribution, as when a user is given advice that benefits many others in an online community. This leads to the easy formation of multiple small communities in

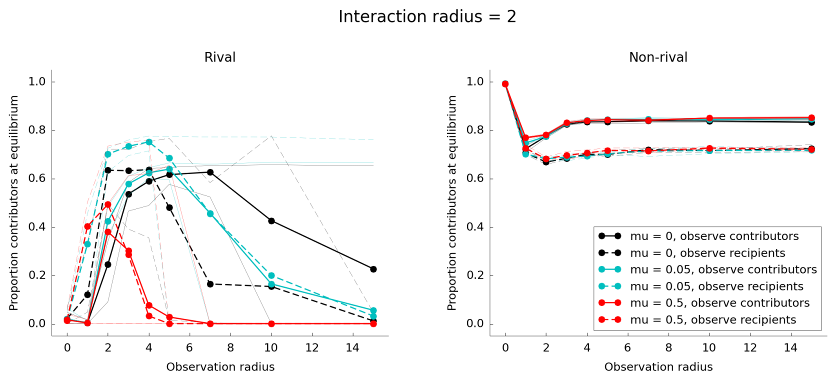

Fig. 7.5: The emergence of contribution by observation radius when interaction is constrained to immediate neighbors and neighbors of neighbors (interaction radius = 2). Line colors show levels of mobility and line types differentiate the observation target. Results are shown for rival (left) and non-rival (right) contributions. The thick lines show the proportion averaged over 25 replications for that particular parameter combination. The thin lines show the minimum and the maximum proportions achieved in the replications.

which contributors benefit and hence continue contributing, despite free-riders who benefit enough to hang around the periphery of the clusters. When contributions are rival and only one individual can benefit from each contribution, contribution-based communities are much less likely to emerge and persist. If they do, this usually happens around a core of unconditional altruists (agents with low θ_{UA} and θ_{FR} high) who form a critical mass. These agents (the blue agents in Fig. 7.6, left column) continue contributing regardless of what others around them do. When outcome-based mobility is relatively low, the agents remain in the neighborhood long enough to have a chance to benefit from a contribution or to observe many others benefiting. (If they were observing contributors instead of recipients, they would have only observed the altruist or the few altruists that started contributing, not the many neighbors who benefit). As a result, a few clusters form around the handful of altruists in the population but the contagion does not spread to agents in other corners of the space.

The differences in the macro-outcomes between rival and non-rival contributions result from the structure of interactions and not from the differences in effect size. Assuming that the GR and TPI effects for non-rival contributions are weaker than the GR and TPI effects for rival contributions does not significantly affect the emergence of non-rival contributions.

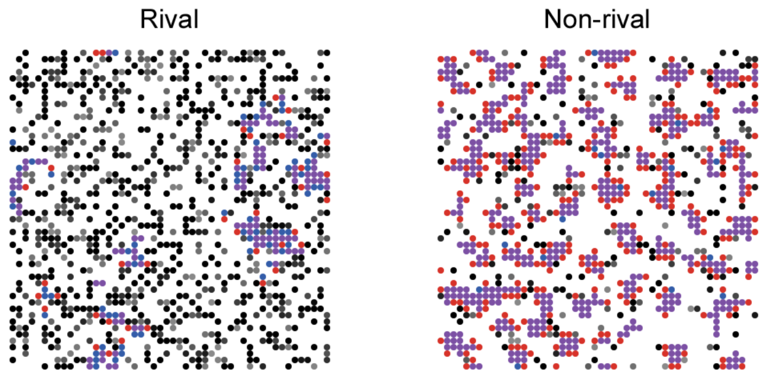

Fig. 7.6: The emergence of contribution communities for rival (left) and non-rival (right) contributions for interaction radius = 1, observation radius = 5, observing recipients, and mobility. Agents in blue contribute but do not benefit, agents in red benefit but do not contribute, and agents in purple both contribute and benefit.

7.4 Discussion

Selfless acts of kindness and anonymous voluntary donations can be puzzling, even though they are not uncommon. In daily life, people donate blood, contribute money to charity, hold the door open for the person behind, or vacate a subway seat for an elderly passenger. In the online world, users review services, rank products, or answer strangers' questions on forums. Why do communities vary in the level of member contributions? This study suggests that the answer could lie in the contagion of prosocial behavior. We first presented empirical evidence from an online experiment for the existence and interaction of two distinct mechanisms of contagion – generalized reciprocity and third-party influence. We then implemented these mechanisms in an agent-based model to investigate the conditions under which they lead to high levels of contributions at the population level.

The empirical results showed that receiving and observing helping behavior can increase the likelihood to help a stranger. However, the willingness to contribute can be offset by lower perceived need when the level of helping is sufficiently high, particularly among those who have not themselves been helped.

We implemented these findings in a threshold model with dynamic interaction structure and adaptive behavior. The computational experiments suggested two alternative pathways to the emergence of contribution-based communities. It is useful to think of these two pathways in the context of rival face-to-face interactions on the one hand and non-rival online contributions on the other hand. In face-to-face interactions, acts of generosity are rival if the benefit is limited to the intended recipient, as happens when holding the door open or vacating one's seat for a stranger. The simulation results show that these contributions can emerge and spread in small and stable communities, that is, communities that are tightly knit and have little turnover. In such communities, hearing about or seeing other people who benefit from the kindness of strangers increases contributions. As a result, a relatively small number of persistent altruists can trigger the spread of helping behavior. In this situation, gossip and newspaper reports about anonymous acts of generosity play an important role. For example, in an office environment, a single active anonymous altruist could trigger a chain of generosity so long as there is sufficient gossip about the level of charitable behavior such that observers come to believe that generosity is normative and conform to this "office culture."

In comparison, non-rival contributions, such as writing a product review on the Web or answering a question in an online forum, are much more likely to emerge and spread across a wider range of conditions, including in much larger groups with high turnover. For example, small esoteric-interest groups and large general-topic online portals could be equally successful user-generated content communities. In such communities, hearing about or seeing other people who contribute sustains high levels of contribution, while awareness of the number of beneficiaries decreases contribution (perhaps due to the belief that there is little need for additional sacrifice).

However, a disclaimer is in order. The chapter provides a plausible explanation for the emergence and persistence of voluntary contribution-based communities but is mainly intended to address the emergence of contribution communities among anonymous individuals. Undoubtedly, once a community forms and anonymity diminishes, cooperation-inducing mechanisms based on social sanctions (for example, reputation systems or long-term-membership privileges) become more prominent and more effective.

7.5 Acknowledgements

This work was supported by grants SES-1260348 and SES-1303526 from the National Science Foundation. The first author thanks Scott E. Page, John H. Miller, and participants at the Graduate Workshop in Computational Social Science at the Santa Fe Institute for valuable comments and suggestions.

References

[1] Centola, Damon. 2010. "The spread of behavior in an online social network experiment." *Science* 329(5996):1194–97.

[2] Salganik, Matthew J., Peter Sheridan Dodds, and Duncan J. Watts. 2006. "Experimental study of inequality and unpredictability in an artificial cultural market." *Science* 311(5762):854–56.

[3] Rogers, Everett M. 2003. *Diffusion of Innovations*. New York: Free Press.

[4] Muchnik, Lev, Sinan Aral, and Sean J. Taylor. 2013. "Social influence bias: a randomized experiment." *Science* 341(6146):647–51.

[5] Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D. I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. 2012. "A 61-million-person experiment in social influence and political mobilization." *Nature* 489(7415):295–98.

[6] Christakis, Nicholas A., and James H. Fowler. 2009. *Connected: The Surprising Power of Our Social Networks and How They Shape Our Lives*. New York: Little, Brown and Company.

[7] Fowler, James H., and Nicholas A. Christakis. 2010. Cooperative behavior cascades in human social networks. Proceedings of the National Academy of Sciences 107(12), 5334–5338.

[8] Suri, Siddharth, and Duncan J. Watts. 2011. Cooperation and contagion in web-based, networked public goods experiments. *PLoS ONE* 6(3), e16836.

[9] Jordan, Jillian J., David G. Rand, Samuel Arbesman, James H. Fowler, and Nicholas A. Christakis. 2013. "Contagion of cooperation in static and fluid social networks." *PLoS ONE* 8(6):e66199.

[10] Kollock, Peter. 1999. "The economies of online cooperation: gifts and public goods in Cyberspace." Pp. 220–42 in *Communities in Cyberspace*, edited by Marc A. Smith and Peter Kollock. London: Routledge.

[11] Pfeiffer, Thomas, Claudia Rutte, Timothy Killingback, Michael Taborsky, and Sebastian Bonhoeffer. 2005. "Evolution of cooperation by generalized reciprocity." *Proceedings of the Royal Society B* 272:1115–1120.

[12] Stanca, Luca. 2009. "Measuring indirect reciprocity: Whose back do we scratch?" *Journal of Economic Psychology* 30(2):190–202.

[13] Bartlett, Monica Y., and David DeSteno. 2006. "Gratitude and prosocial behavior." *Psychological Science* 17(4):319–325.

[14] Tsvetkova, Milena, and Michael W. Macy. 2014. "The social contagion of generosity." *PLoS ONE* 9(2):e87275.

[15] Oliver, Pamela, Gerald Marwell, and Ruy Teixeira. 1985. "A theory of the critical mass. I. Interdependence, group heterogeneity, and the production of collective action." *American Journal of Sociology* 91(3):522–56.

[16] Karau, Steven J. and Kipling D. Williams. 1993. "Social loafing: a meta-analytic review and theoretical integration." *Journal of Personality and Social Psychology* 65 (4): 681–706.

[17] Darley, John M., and Bibb Latané. 1968. "Bystander intervention in emergencies: diffusion of responsibility." *Journal of Personality and Social Psychology* 8: 377-83.

[18] Granovetter, Mark. 1978. "Threshold models of collective behavior." *American Journal of Sociology* 83(6):1420–43.

[19] Macy, Michael W. 1991. "Chains of cooperation: threshold effects in collective action." *American Sociological Review* 56(6):730–47.

[20] Oliver, Pamela E. 1993. "Formal models of collective action." *Annual Review of Sociology* 19:271–300.

[21] Schelling, Thomas C. 1971. "Dynamic Models of Segregation." *The Journal of Mathematical Sociology* 1(2):143.

[22] Granovetter, Mark, and Roland Soong. 1983. "Threshold models of diffusion and collective behavior." *The Journal of Mathematical Sociology* 9(3):165–179.

[23] Granovetter, Mark, and Roland Soong. 1986. "Threshold models of interpersonal effects in consumer demand." *Journal of Economic Behavior & Organization* 7(1):83–99.

[24] Lopez-Pintado, Dunia, and Duncan J. Watts. 2008. "Social influence, binary decisions and collective dynamics." *Rationality and Society* 20(4):399–443.

[25] Eguluz, Victor M., Martin G. Zimmermann, Maxi San Miguel, and Camilo J. Cela-Conde. 2005. "Cooperation and the emergence of role differentiation in the dynamics of social networks." *American Journal of Sociology* 110(4):977–1008.

[26] Biely, Christoly, Klaus Dragosits, and Stefan Thurner. 2007. "The Prisoner's Dilemma on co-evolving networks under perfect rationality." *Physica D: Nonlinear Phenomena* 228(1):40–48.

[27] Hanaki, Nobuyuki, Alexander Peterhansl, Peter S. Dodds, and Duncan J. Watts. 2007. "Cooperation in evolving social networks." *Management Science* 53(7):1036–50.

[28] Helbing, Dirk and Wenjian Yu. 2009. "The outbreak of cooperation among success-driven individuals under noisy conditions." *Proceedings of the National Academy of Sciences* 106(10):3680–85.

[29] Meloni, S., A. Buscarino, L. Fortuna, M. Frasca, J. Gomez-Gardenes, V. Latora, and Y. Moreno. 2009. "Effects of mobility in a population of Prisoner's Dilemma players." *Physical Review E* 79(6):067101.

[30] Fehl, Katrin, Daniel J. van der Post, and Dirk Semmann. 2011. “Co-evolution of behaviour and social network structure promotes human cooperation.” *Ecology Letters* 14(6):546–51.

[31] Macy, Michael W., and Andreas Flache. 2002. “Learning dynamics in social dilemmas.” *Proceedings of the National Academy of Sciences of the United States of America* 99(3): 7229–7236.

[32] Fischbacher, Urs, Simon Gachter, and Ernst Fehr. 2001. Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters* 71(3):397–404.

[33] Kurzban, Robert, and Daniel Houser. 2005. Experiments investigating cooperative types in humans: a complement to evolutionary theory and simulations.” *Proceedings of the National Academy of Sciences of the United States of America* 102(5):1803–7.

[34] Macy, Michael W., and Milena Tsvetkova. 2013. “The signal importance of noise.” *Sociological Methods & Research*, doi 0049124113508093.

Manipulation	A)GR	B)TPI	C)GRxTPI
Invitee (receives a donated invitation)	7.006 (0.030)*		0.327 (0.262)
Has previously received donated invitations	0.712 (0.686)		1.021 (0.982)
Seeds			
Observes 0-75		11.414* (0.043)	(baseline)
Observes 75-150		1.341 (0.787)	0.136 (0.101)
Observes 151+		0.219 (0.280)	0.015* (0.022)
Invitees			
Observes 76-150			19.907* (0.041)
Observes 151+			89.948* (0.026)
High Payment	64.103** (0.007)	2.532 (0.300)	3.235 (0.295)
Time waited (in hours)	0.972* (0.023)	0.992 (0.577)	0.976 (0.075)
Previous participations	0.690 (0.379)	0.784 (0.622)	0.454 (0.171)
Baseline odds	4.305 (0.181)	5.323 (0.100)	268.707*** (0.000)
Number of observations	516	371	554
Number of participants	252	277	266
Wald χ^2	5 df, 11.93* (0.036)	6 df, 6.66 (0.354)	8 df, 11.98 (0.214)

Two-sided tests:* $p<0.05$, ** $p<0.01$, ***$p<0.001$

Table 7.1: Odds ratios for donating across treatments. The table reports odds ratios and p values (in brackets) from random-intercept logistic regression models for A) seeds and invitees in the no-observation treatment by number of donated invitations received; B) seeds in the observation and no-observation treatments by number of donated invitations observed; and C) seeds and invitees in the observation treatment by number of donated invitations observed by invitees compared to seeds. Results show that receiving and observing donations initially increases the willingness to help others, and that invitees are less susceptible to a subsequent decline in helping.