Journal of
Peace Research

Journal of Peace Research
2016, Vol. 53(3) 459–471

DOI: 10.1177/0022343316630781
jpr.sagepub.com
SAGE

Interethnic conflict and the potential dangers of cross-group ties

Jennifer M Larson

Department of Politics, New York University

Abstract

Bridging social ties is thought to reduce the likelihood of interethnic violence. This logic has motivated countless development projects and international programs seeking to forge cross-group ties between groups with a conflictual history. However, this article identifies an important mechanism by which certain cross-group ties can make interethnic peace strictly less likely. The results stem from a game-theoretic model which formalizes civil society as a network and relates intergroup cooperation to the particular networks that transmit information from person to person in each group. The model reveals that, first, groups are capable of enforcing cross-group cooperation, even when no cross-group ties are present and the networks within each group are missing links, using peer-enforcement strategies, and their ability to do so depends on the structure of these networks. Second, when attempting to enforce intergroup cooperation, groups with sparse networks may be at risk of a long-lasting series of back-and-forth retaliation that groups with denser networks would avoid. Finally, there exists a mechanism by which some cross-group ties make intergroup cooperation strictly less likely. When interethnic cooperation is enforced by threatening coordinated retaliation for any misbehavior, success depends on expectations about how quickly retaliation can be coordinated and how many will participate in it. Some individuals in a network are in a position to send news to many others quickly; others are not. The latter therefore coordinate retaliation more slowly and would be relatively vulnerable to cross-group defections if they could be identified. Cross-group ties expose the vulnerability and generate incentives to disrupt interethnic peace; cross-group ties between the least embedded individuals in each ethnic group are the most dangerous. Programs seeking to impose ties should avoid exposing this vulnerability without taking steps to mitigate its danger.

Keywords

civil society, cooperation, interethnic conflict, social networks

Introduction

Some neighboring ethnic groups violently conflict with each other; many more coexist peacefully (Fearon & Laitin, 1996). A leading explanation of the difference focuses on the presence or absence of social connections between individuals of different ethnicities, so-called ‘cross-group ties’. Individuals’ relationships have become central to explanations of subnational conflict, so much so that the peace interventions by international organizations frequently have a civil society component that focuses on building social contacts (Belloni, 2001).

This article formalizes the full set of social contacts within and across ethnic groups in order to better isolate their role in intergroup peace. I present a game-theoretic model which explicitly accounts for the constellation of relationships, the social networks, within neighboring ethnic groups. This setup reveals a number of new insights about the role of cross-group ties.

First, this setup shows how the strategic enforcement of intergroup cooperation via a coordinated response (as in Fearon & Laitin, 1996) depends on networks even when no cross-group ties are present. The faster news spreads from person to person within each group, the more easily the groups can coexist in peace.

Corresponding author:
jenn.larson@nyu.edu

Second and counter-intuitively, this setup reveals a mechanism by which cross-group ties may be *antithetical* to interethnic cooperation. Suppose a cross-group tie between i in one group and j in another allows i and j to communicate, and also reveals a minimal amount of personal information to each other: in the model, i knows the size of j's social circle (network neighborhood) and vice versa. If i and j did not share a cross-group tie, they would be strangers and so would not possess this channel of communication and would not know about the other's social circle. It turns out that sometimes the size of one's social circle is better left concealed from the other group. Those with a small social circle are more vulnerable to attack since they would only have enough peers in reach to coordinate a weak retaliation. The ability to identify who these vulnerable individuals are can generate strict incentives to initiate interethnic conflict. Some cross-group ties expose vulnerability.

This second insight not only shows that there is a downside to some cross-group ties, but also that the downside is due to revealed information. The information is not inaccurate or malicious; however, there are some pieces of information which, when learned, actually reduce the likelihood of peace.

Finally, the model reveals a potential source of *protracted* conflict between two groups. If groups try to maintain peace by threatening finite periods of mutual retaliation in response to any misbehavior, groups that spread news too slowly can inadvertently convert a finite period of retaliation to a much longer one. If news spreads slowly, some may learn the two groups have entered a period of retaliation by themselves being retaliated against rather than by being told by a co-ethnic. This poses a problem for coordinating the *end* of the period of retaliation. If, once the retaliation period has ended, some people are still retaliating, then their retaliation triggers more retaliation, which may trigger more again and again due to miscoordination.

These findings suggest that cross-group ties are not *necessarily* beneficial for interethnic peace. An implication is that programs seeking to forge new ties should proceed cautiously. While well-designed programs can indeed promote peace, short-term programs that briefly give individuals in both groups access to information about the other group, especially about the most vulnerable members, run the risk of actually increasing the likelihood of conflict. Understanding the mechanism uncovered in this article that works against peace is essential to designing programs that are on-net beneficial.

Interethnic cooperation

Civil society and conflict

Violent ethnic conflict tends to be longer and deadlier than other forms of civil conflict (Fearon, 2004); however, ethnic tensions rarely erupt into violent conflict despite the prevalence of different ethnicities living in close proximity across the globe (Fearon & Laitin, 1996: 717). When is ethnic pluralism susceptible to the outbreak of violent conflict and when is cooperation possible?

This article joins the growing body of research that turns to civil society for an answer. According to this line of research, the nature of interactions between individuals in their civic lives – day-to-day interactions or encounters in civic organizations – can bear on the likelihood of peace or conflict. This is particularly true of interactions between members of different ethnic groups, which can positively impact the chances for peace between neighboring groups.[1]

Many explanations for the felicitous effect of cross-group interactions in civil society have been posited. They rely on interactions yielding a relationship by which parties to the interaction learn something new and/or have a new source of future information.

In one posited mechanism, the mere act of encountering someone different can change one's feelings toward the out-group. According to the so-called 'contact hypothesis', contact with members of other groups may chip away at prejudice, reduce fear, and foster understanding, and so promote tolerance (for a survey, see Pettigrew & Tropp, 2006).

In another, forging relationships with individuals in other groups also has the potential to change the preferences of one or both parties in ways that improve future interactions. For instance, jointly participating in a civic organization toward a common goal may cultivate shared interests that could reduce friction between groups, or generate mutually beneficial joint endeavors that increase the costs of conflict (e.g. Jha, 2013), which is especially likely in the presence of high levels of specialization among ethnic groups (Jeon, 2014).

According to a third mechanism, interacting with members of another ethnic group may open channels

[1] I refer to these interactions as cross-group ties, though many terms for the presence of cross-group interactions can be found in the literature: cross-cutting as opposed to mutually reinforcing cleavages (Lipset & Rokkan, 1967); brokerage across structural holes as opposed to closure (Burt, 2000); and bridging as opposed to bonding social capital (Putnam, 2001).

of communication between the two groups, which results in having a source of future information in the other group. This source may help prevent interethnic tensions from escalating to interethnic violence: rumors of misconduct or threatening intentions in the other group could ignite conflict, as could political attempts to drum up violence via misinformation (e.g. Wilkinson, 2004). Access to information in the other group that could dispel rumors and offer credible reassurances would reduce the likelihood of violent conflict (Varshney, 2003).

All of the above mechanisms operate at the level of the cross-ethnic tie. An additional mechanism identifies a role for civil society even absent any cross-group ties. Instead, the set of relationships *within* an ethnic group can bear on interethnic violence. This mechanism draws on the logic of community enforcement (Kandori, 1992) and relies on the stylized fact that information relevant to peer sanctioning spreads rapidly among co-ethnics (Miguel & Gugerty, 2005). If ethnic groups threaten to strategically respond to cross-group misbehavior, then the rapid transmission of information through an ethnic group can permit the effective enforcement of interethnic cooperation even without any cross-group ties (Fearon & Laitin, 1996).

The above mechanisms suggest a key insight for how to proceed: both cross-group and within-group ties may matter for interethnic cooperation, potentially in tandem. Any of the cross-group tie explanations could be mediated or magnified by within-group ties (as if, for instance, a cross-group contact shares information that dispels an incendiary rumor, and then within-group ties transmit that information widely). To consider with precision the full set of relationships within and between groups at once, I employ the framework of network analysis (Lin, 1999; Burt, 2000; Lin, Cook & Burt, 2001).

Networks and cooperation

The way groups of actors are connected in a network bears on incentives for cooperation and conflict. In the study of international cooperation and conflict, the relevant actors are states, and the relevant links comprising a network among them may be trade flows (Chyzh, 2016), alliances (Haim, 2016), formal weapons cooperation agreements (Kinne, 2016), nonaggression pacts (Lupu & Poast, 2016), cross-national connections between NGOs (Wilson, Davis & Murdie, 2016), or shared policy preferences (Ward & Dorussen, 2016), as considered elsewhere in this issue. The study of *subnational* cooperation and conflict requires a more micro view: here, relevant actors are taken to be individuals who can spark conflict, and relevant links among them are relationships conducive to sharing information.

The mechanisms by which civil society may bear on conflict among individuals make clear that there is a difference between an interaction, a relationship, and information. Individuals may engage in an interaction with each other to conduct a transaction without having a personal relationship (though personal relationships may make transactions more likely or more civil and frequent interactions may make relationships more likely over time). Individuals in a relationship may have formed the relationship for any number of reasons, but the presence of a relationship opens a potential source of future information: relationships open a channel of communication. The model below separates interactions from relationships, and focuses on the information-spreading role of relationships. The set of relationships that can spread information is represented by a network.

Real groups of people pass news from person to person along networks, and networks among real groups of people tend to be incomplete.[2] News may spread rapidly from person to person within a group, especially in groups we would call ‘tight-knit’. However, unless by ‘tight-knit’ we mean that everyone *directly* and immediately communicates with everyone else, the networks are incomplete and the exact arrangement of the links that *are* present affects how widely and quickly any news spreads. In the model, relevant news concerns behavior in interactions – gossip about who behaved badly.

There is a small but rapidly growing literature that studies cooperation among individuals in a network setup, either by having the network constrain interactions (e.g. Nava & Piccione, 2014) or by having the network constrain information as I do here (e.g. Wolitzky, 2013). To my knowledge, this is the first work to consider cooperation enforced by intergroup punishment in a networks setting, and the first to consider cross-group interactions at all in a networks setting.

One complication that arises in the model is that if peace is enforced through the threat of intergroup retaliatory conflict that lasts some length of time, everyone

[2] Incomplete networks are networks that do not contain every possible link connecting any pair of nodes. In the case of communication networks, this means that some people do not communicate directly with every other person in the group.

must know exactly when the retaliation begins to know when it is due to end. Incomplete networks open the possibility that people disagree about the correct end to retaliation, at least for a while. A similar complication arises in the setting considered in Cho (2011), though with just a single group. The solution there relies on a public randomization device to coordinate a stopping point. Because it can be difficult to identify public randomization devices that real groups use, I consider conditions under which communication networks would allow groups to overcome problems of temporary miscoordination without such a device. I argue that this problem of beliefs out of synch is not a bug of the model but in fact exposes an important source of protracted conflict between real groups.

The model

Overview

The model is a generalization of Fearon & Laitin (1996) which accounts for the networks along which news spreads within and potentially across each of two groups. Because the interest here is in the potential for inter-group conflict, I focus on the spiral equilibrium in which each group threatens large-scale retaliation against the other group if anyone misbehaves in an interaction with a member of the other group.

When news spreads from person to person via links in a network, some players may be out of the loop for some period of time. In the original model, and equivalently, if the networks were complete, when someone in the out-group defects against a member of *i*'s in-group, *i* knows immediately and so can join in the large-scale retaliation right away. When networks are incomplete, it can be the case that someone in the out-group has defected against a member of *i*'s in-group but *i* doesn't know about it yet, so *i* wouldn't know to participate in the retaliation.

In this way, the network structure impacts the extent of cross-group retaliation and, consequently, the expected costs of behaving uncooperatively with a non-coethnic. If a network spreads news of a defection slowly, then the number of those who are informed and can participate in the cross-group retaliation also grows slowly, attenuating the expected costs of retaliation.

Model setup

Begin with two groups, A and B, each containing a set of n players, $A = \{1, \ldots, n\}$, $B = \{n+1, \ldots, 2n\}$, with n even. Players in the model perform two actions: they play games with out-group members, and they share information about games. Players play games with one other randomly selected out-group player at a time.[3] All pairs simultaneously play a single round of the prisoner's dilemma with common payoff matrix

$$\begin{array}{cc} & \begin{array}{cc} C & D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \begin{pmatrix} 1,1 & -\beta,\alpha \\ \alpha,-\beta & 0,0 \end{pmatrix} \end{array}$$

where $\alpha > 1$, $\beta > 0$ and $\frac{\alpha-\beta}{2} < 1$. Players are re-matched every period for an indefinite number of periods. A player's total payoff is then a stream of discounted single round payoffs. Players have common discount factor $\delta < 1$.[4]

Observation and communication occur according to a 'communication network' defined by the pair $(g, A \cup B)$ with $2n \times 2n$ adjacency matrix g where $g_{i,j} = g_{j,i} = 1$ indicates a link between $i \neq j \in A \cup B$. For simplicity, I refer to the network as 'g'. Links in the network are undirected and unweighted.[5] The network g can be thought of as the union of three networks: g^A as the subnetwork induced by A, g^B as the subnetwork induced by B, and g^{AB} as the subnetwork containing the set of ties connecting someone in A to someone in B. That is, g^A

[3] The present interest is in sustaining cross-group cooperation, so for simplicity, all interactions in the model are cross-group. We could also specify a probability with which players are paired with an out-group member as opposed to an in-group member and consider both within- and cross-group cooperation simultaneously. The equilibria of interest here treat cross-group interactions separably. If players had cross-group interactions less frequently, the impact of cross-group punishment would be attenuated and so the conditions would be harder to satisfy, but the comparative statics would all have the same direction as in the present approach.

[4] The prisoner's dilemma is a convenient apparatus for studying cooperation and conflict. It captures the environment in which in any interaction, any individual would like to do something that counts as 'misbehaving' if they could get away with it. This serves as a hard case for cooperation. The main results of the article do not require the PD setup. The downside to cross-group ties persists so long as the strategic environment is any in which the key to preventing misbehavior is disincentivizing it by coordinating a wide-reaching response via a network. The disincentives for misbehaving against someone in more isolated network positions would still be weaker, and revealing this through a cross-group tie would increase the likelihood of conflict.

[5] Undirected and unweighted links are simplifying assumptions. All results hold for directed links; they would just need to be restated with more cumbersome notation. Links are unweighted here following the typical study of communication networks. Even when the relationships themselves are classified as 'strong' and 'weak', both are usually taken to transmit news (see e.g. the job search model of Granovetter, 1973).

contains all of A's within-group ties, g^B contains all of B's within-group ties, and g^{AB} contains all of the cross-group ties. When g^{AB} is the zero matrix, the two groups have no cross-group ties. Let g^A and g^B both be connected networks.[6] The set $N_i = \{j \neq i \in A \cup B : g_{i,j} = 1\}$ is player i's 'neighborhood' and the $j \neq i \in A \cup B : g_{i,j} = 1$ are i's 'neighbors'. The neighborhood of i contains all players who can be reached from i in geodesic distance 1. Generalizing, we can define neighborhoods of size k:

Definition 1 (k-neighborhood): Let $l(i,j)$ be the length of the shortest path from i to j. Define i's **k-neighborhood** in network g, $N_i^k(g)$ to be the set of all j such that the shortest path from i to j is less than or equal to k. That is,

$$N_i^k(g) = \{j \in A : l(i,j) \leq k, i \neq j\}. \qquad (1)$$

That is, i's k-neighborhood is the set of all other players reachable from i in paths of length k or shorter. Information spread through the network is characterized below. Whenever I use the term 'neighborhood' as opposed to 'k-neighborhood', I am referring to $N_i \equiv N_i^1$.

The following term will be useful:

Definition 2 (Incident): An individual $i \in A$ is **incident** to a tie $g_{j,k}$ if $i = j$ or $i = k$. i is **incident to a cross-group tie** if $\exists\, g_{i,j} \in g$ such that $j \in B$. In other words, i is incident to a cross-group tie if i has at least one neighbor in the out-group.

As in Fearon & Laitin (1996), members of a group cannot recognize everyone in an out-group, cannot describe anyone in the out-group, and would not recognize everyone in the out-group if rematched.[7]

In particular, $i \in A$ only recognizes $j \in B$ if $g_{ij}^{AB} = 1$ – when there is a cross-group tie connecting i and j – though i could not describe j well to other members of A who are not also connected to j. A player only knows he's playing 'someone' from the other group unless he is playing his neighbor. This scenario is the hard case for intergroup cooperation, and seems to be a reasonable representation of some ethnic groups' experience (Fearon & Laitin, 1996: 727).[8] More on this below.

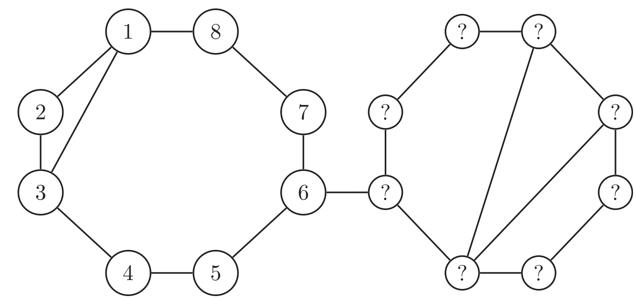

Figure 1. An example communication network from the perspective of player 1

The anonymized shape of the network g is common knowledge among all players in A and B. That is, they could draw the network accurately without labeling any of the nodes. Players know exactly which player is which node in their own group. Players also know which nodes their own cross-group neighbors are, but not which nodes any other out-group members are. So players could add node labels to their drawn networks for every member of their in-group and their personal out-group neighbors (connected through cross-group ties) but no others.[9] Figure 1 shows an example network for group A (here players 1 … 8), group B (players 9 … 16), and the cross-group ties (here just one) from the perspective of player 1. Figure 2 shows the same from the perspective of player 6, and Figure 3 shows the same from the perspective of player 11.

Actions are observable to neighbors in the network g. Observers pass news along to their neighbors. News about rounds spreads truthfully and deterministically through the network between rounds at rate $r = \frac{\text{Degrees spread}}{\text{Rounds played}}$ between rounds. When $r = 1$, news

[6] That is, there exists a finite path between any two players in A and likewise for B.

[7] The relevant feature of a group is it is a set of players who can all recognize each other: they can perfectly identify each other, describe each other, and would recognize each other if rematched. While this applies most naturally to different ethnic groups, in fact the model applies to any pair of groups that can be differentiated in this way.

[8] This depiction of recognizability is the natural extension of no cross-group recognizability in Fearon & Laitin (1996) to the case where there are some cross-group ties.

[9] Individuals can actually know much less about the out-group network for the results to hold: a highly reduced-form view of the out-group network and the out-group neighborhood of their cross-group contacts. The assumption that individuals know more about the network position of their cross-group contacts than other out-group members is key to the results that follow. This assumption seems reasonable – individuals in the out-group with whom i does not share a cross-group tie are effectively strangers. Those with whom i does share a tie have enough of a relationship to share news of day-to-day happenings. This would likely entail knowing something about that contact's social ties, if only a rough impression of the size of his friend group.

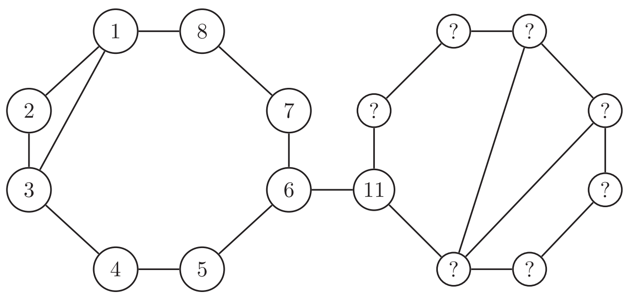

Figure 2. An example communication network from the perspective of player 6

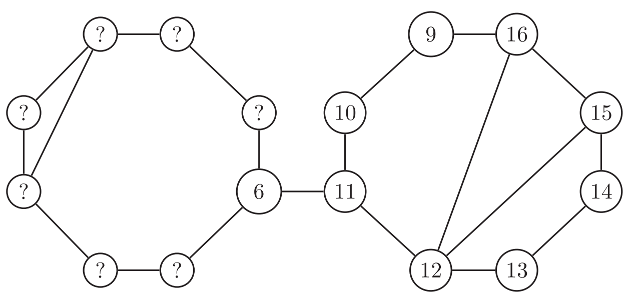

Figure 3. An example communication network from the perspective of player 11

about rounds spreads one degree, to the immediate neighbors, before players are rematched and play again. When $r = 2$, news spreads two degrees, from person to person to person before the next round, and so on. Therefore, at the end of time t, a player has received new relevant information from everyone reachable in paths of length r, that is, from everyone in his r-neighborhood.

The assumptions about recognizability, common knowledge of the network structure, and observability comport with stylized facts about the social structures of real ethnic groups. Ethnic groups are groups of individuals with actual or perceived shared descent and tend to have features that facilitate the rapid spread of information and the issuing of swift sanctions, which can include common language, shared cultural understandings, frequent co-ethnic interactions, and living in close proximity (Horowitz, 1985; Habyarimana et al., 2009). The ability to share news quickly and to sanction within an ethnic group, especially relative to the limited ability to do so across ethnic groups, has become the leading explanation for observed differences in public goods provision between ethnically homogeneous and ethnically diverse areas (Miguel & Gugerty, 2005; Habyarimana et al., 2007). Additionally, to the extent that social ties have a basis in kinship or identification with cultural or familial groupings at the subethnicity level, such as clans and subclans, co-ethnics are more likely to know the network among co-ethnics than among non-coethnics (Larson & Lewis, 2015a).

The infinitely repeated game G proceeds as follows: at the beginning of each round, nature randomly pairs players in group A with players in group B and randomly assigns each pair to one of n possible locations for each pair to play. Players know they are playing 'someone in the other group' at a particular location. Pairs simultaneously play one round of the prisoner's dilemma. News is spread at rate r through the network g, ending the round, and the next round begins.

Equilibrium strategies for cross-group cooperation

Consider a generalization of the spiral equilibrium strategies in Fearon & Laitin (1996), which capture tense behavior between ethnic groups well.[10] Call $t_{i,t}^{spark}$ the first instance of a D played in a cross-group interaction that i knows about by period t. Now we have:

> Definition 3 (Network spiral profile (σ^{NWSPIR})): All players play C in the first round. In subsequent rounds t, all players play C unless they know that anyone has played a D, in which case they play D until $t_{i,t}^{spark} + T$.

Players following σ^{NWSPIR} react to cross-group defections by playing D in all out-group pairings for a finite number of rounds T. Depending on where an individual sits in his group's network, he may be unaware of a D played in some round t. In fact, an individual's first indication that cross-group punishment has been initiated may be when he plays someone in the out-group who knows punishment has been initiated and plays D against him.[11]

[10] These strategies inflict the largest possible finite punishment between groups. Finite punishments are desirable in environments prone to accidents and short tempers so that off-the-equilibrium-path inefficiency is minimized, and can extract maximal cooperation when the action set is binary as it is here.

[11] I follow Fearon & Laitin (1996) and the prisoner's dilemma literature by referring to this phase as 'punishment'. Really, it is one side issuing punishment and the other side reacting defensively since countering D with D is the best response of the group being punished for the T rounds. When a group is the victim of a cross-group defection, that group commits to punishing the other group by playing D; the other group reduces the costs of that punishment by countering with D.

This poses a complication for coordinating the *end* of a punishment phase. A player who has not heard that someone tripped the punishment phase and whose opponent just played D against him in punishment cannot know that the D was played in punishment and instead thinks the D is a first spark of a new punishment. Only when news about the initial spark reaches him will he become aware. For some of the time, members of group A may disagree over when the punishment phase should stop. This disagreement is harmless *so long as everyone agrees by* $t^{spark} + T + 1$. In this round, unless all agree that the punishment triggered in t^{spark} is over, someone will believe they are punishing but instead they will be tripping a new phase of punishment.[12] For certain networks, this miscoordination can persist indefinitely, kicking off a spiral of conflict.[13]

A necessary condition for punishment to remain finite requires a modified definition of network diameter to account for the possibility of cross-group ties. Let $l(i,j)$ be the shortest path between i and j in network g.

Definition 4 (JK-diameter): Let J and K be subsets of nodes of g. The JK-diameter of network g is:

$$diam_{JK}(g) = \max_{i \in J \cup K} \left\{ \min_{j \in J, k \in K} \{l(i,j), l(i,k)\} \right\}. \quad (2)$$

The JK-diameter is the longest shortest path to one of two sources where one source is in one subset of nodes and the other is in another subset. This definition accounts for the fact that there are two sources to any piece of information (both parties to the interaction). Now we can state the following necessary condition for punishment under σ^{NWSPIR} to last only T rounds:

Proposition 1 (Coordinating cross-group punishment): All players always agree on when the last round of punishment should occur under σ^{NWSPIR} only if

$$T \geq \frac{diam_{AB}(g)}{r}. \quad (3)$$

The proof of Proposition 1 is straightforward and can be found in the Online appendix. Satisfying this condition guarantees that by the time someone could act on misinformation and trigger punishment, everyone in both groups will know the correct stopping time.[14] In other words, the punishment period must be sufficiently long relative to the AB-diameter and speed of information spread. This condition suggests that incomplete networks with slow-spreading news can result in a spiral of conflict due to miscoordination.

One interesting consequence of this condition is that some networks will require longer cross-group punishments than others. When news must travel through many intermediaries before reaching everyone (the network has a high AB-diameter), groups must threaten long-lasting retaliations in order for any retaliation to actually remain finite. Retaliations that are designed to be short may actually result in one round of retaliation kicking off another and another.

Full cooperation in equilibrium

Now I turn to the conditions under which the strategy profile σ^{NWSPIR} supports full cross-group cooperation in sequential equilibrium. In order to isolate the consequences of adding cross-group ties, first consider the case in which there are no cross-group ties: $g^{AB} = \emptyset$.

An individual's decision to comply with σ^{NWSPIR} depends on the number of others who know cross-group retaliation is due when cross-group retaliation is in fact due. In any round, a player has to compare his gain from kicking off retaliation to his gain from remaining cooperative. His gain from kicking off retaliation

[12] At first blush, it may seem that if two groups really did enter the retaliation phase, they could tell the difference between it and a new defection. However, consider what the retaliation phase looks like to an individual *i* who learns late. To *i*, relations were warm between the two groups, and now suddenly a member of the out-group behaves uncivilly. If no one in *i*'s group has told him that uncivil encounters are warranted yet, *i* has no reason to believe anything other than that retaliation is due to start now. Even if the member of the out-group who behaved uncivilly broke the rules of the model and told *i* 'I am doing this out of retaliation for what your group did 2 days ago,' this statement would not be credible. Any out-group member would have an incentive to claim they were merely carrying out additional rounds of retaliation any time they were defecting anew to improve the chances of encountering a more profitable player in the future who is still unaware. Even if groups could eventually detect the source of infinitely long conflict, the finite conflict may persist longer than *T*. The general point of the result is that group-wide responses require coordination; some networks facilitate this coordination better than others, and the consequences of poorly facilitated coordination can be protracted conflict.

[13] Note that in the so-called 'spiral equilibrium' of Fearon & Laitin (1996), no actual spiral occurs. One group reacts en masse against the other group, both play *D* for one iteration of the punishment phase, but then punishment ends and cooperation resumes. Here, by accounting for networks, one round of punishment can kick off another and another in a true spiral. This spiral is of concern off the equilibrium path or in reaction to errors.

[14] If there are no cross-group ties, this condition is equivalent to $T \geq \min\{diam(A), diam(B)\}$.

depends on the number of his in-group members who will know about the retaliation in each of the T periods of the retaliation. Part of this number is easy for him to calculate – his neighbors know first, then his neighbors' neighbors, and so on. However, some of his in-group members may learn by being paired with an out-group player who knows and plays D. The number who will learn in this way depends on the number of out-group players who know. But the number of out-group players who know depends on the number who learn through their out-group network *and* the number who learn by being randomly paired with an in-group member who knows and plays D.

The cumbersome recursive expression for the expected number who know each round can be found in the Online appendix. In the condition below, the notation $\#BKNOW_l^i$ means the number of members of group B who can be expected to know about a cross-group defection in $t + l$ that occurred in t when the first defection is committed by i, and likewise for $\#AKNOW_l^i$.

Proposition 2 (Full cooperation under σ^{NWSPIR}): Strategy profile σ^{NWSPIR} is sequentially rational for game G with groups A and B without cross-group ties iff:

$$\alpha - 1 \leq \sum_{l=1}^{T} \delta^l \left[1 - \alpha\left[\frac{n - \#BKNOW_l^i}{n}\right]\right] \quad (4)$$

and

$$(\alpha - 1)\left(\frac{n - \#BKNOW_d^i}{n}\right) + \beta\left(\frac{\#BKNOW_d^i}{n}\right) \geq \sum_{l=1}^{T-d} \delta^l \alpha \left[\frac{\Delta\#BKNOW_l^{-i}}{n}\right] \quad (5)$$

for all $i \in A$ and all values of $1 \leq d < T - 1$; and

$$\alpha - 1 \leq \sum_{l=1}^{T} \delta^l \left[1 - \alpha\left[\frac{n - \#AKNOW_l^i}{n}\right]\right] \quad (6)$$

and

$$(\alpha - 1)\left(\frac{n - \#AKNOW_d^i}{n}\right) + \beta\left(\frac{\#AKNOW_d^i}{n}\right) \geq \sum_{l=1}^{T-d} \delta^l \alpha \left[\frac{\Delta\#AKNOW_l^{-i}}{n}\right] \quad (7)$$

for all $i \in B$ and all values of $1 \leq d < T - 1$

where $\Delta\#BKNOW_l^{-i}$ is the expected difference in the number of people in group B who know in $t + l$ when i participates in punishment in t and when i refrains from participating in punishment in t. These values are all functions of the networks in both B and A, and are defined explicitly in the Online appendix. Likewise, these values depend on whether a player is incident to a cross-group tie, which is key to the results below.

In words, the first and third conditions say that the net gain from defecting must be smaller than the net loss from the punishment phase. The net loss from the punishment phase would be 1 in every round, but if news spreads slowly through the other group, occasionally during the punishment phase a punisher will gain α when playing an out-group member who does not yet know that punishment has begun. The second and fourth conditions say that a player will do his part and punish when he knows that punishment is due as long as his marginal impact on the spread of news through the other group is not so great that he expects to be able to cooperate with someone once more and then punish and earn α in the next round. This second set of conditions rules out an unlikely but possible incentive.

Three comparative statics follow straightforwardly from Proposition 2.

Corollary 1: The longer the cross-group punishment (larger T), the more easily the two groups can sustain cooperation.

The cross-group punishment phase is value-destroying; the longer the destruction of value, the more hesitant any one player is to kick it off.

Corollary 2: The denser both groups' networks are, the more easily the two groups can sustain cooperation.

In the extreme case, if both groups' networks were complete, $\#BKNOW = n$, $\#AKNOW = n$, $\Delta\#BKNOW = \Delta\#AKNOW = 0$, and so all conditions become easier to satisfy. Denser networks ensure that news about the punishment phase reaches everyone immediately, making it maximally painful.[15]

The density of the network ensures that cross-group punishment is massive and swift immediately after a defection, offering a greater deterrent. Additionally, news

[15] General density matters because of the randomness built into becoming informed about retaliation: some are effectively informed by the behavior of a random opponent. Each round, the news spreads along the network and additional messages are seeded at random because some in the out-group are already punishing.

about some individuals' rounds spreads relatively slowly due to their network positions. These individuals have the greatest incentive to defect in the first place.

Corollary 3: Individuals who occupy more peripheral network positions have a greater incentive to deviate from σ^{NWSPIR}.

The punishment phase kicked off by such individuals is slower to cascade through the group, and so the person faces lower expected costs from punishment; these individuals can defect with lower expected cost.

The downside to cross-group ties

Now consider the comparison between two groups with no cross-group ties and two groups with a cross-group tie. The conditions for full cooperation are nearly identical; what differs is what the two individuals incident to the cross-group tie know. While everyone else knows merely the average number of others who know about an out-group defection, someone incident to a cross-group tie would know the exact number of others who would know about an out-group defection against that individual. Now $\#BKNOW_l^{i,j}$ is the number of individuals in B who can be expected to know about a cross-group defection that occurred by i against j in t by $t+l$.

Proposition 3 (Full cooperation under σ^{NWSPIR}): Strategy profile σ^{NWSPIR} is sequentially rational for game G with groups A and B **with** cross-group ties iff:

$$\alpha - 1 \leq \sum_{l=1}^{T} \delta^l \left[1 - \alpha \left[\frac{\#B - \#BKNOW_l^{i,j}}{\#B} \right] \right] \quad (8)$$

and

$$(\alpha - 1)\left(\frac{\#B - \#BKNOW_d^{i,j}}{\#B} \right) + \beta \left(\frac{\#BKNOW_d^{i,j}}{\#B} \right) \geq \sum_{l=1}^{T-d} \delta^l \alpha \left[\frac{\Delta \#BKNOW_l^{-i}}{\#B} \right] \quad (9)$$

for all $i \in A$, $j \in B$, and all values of $1 \leq d < T-1$; and

$$\alpha - 1 \leq \sum_{l=1}^{T} \delta^l \left[1 - \alpha \left[\frac{\#A - \#AKNOW_l^{i,j}}{\#A} \right] \right] \quad (10)$$

and

$$(\alpha - 1)\left(\frac{\#A - \#AKNOW_d^{i,j}}{\#A} \right) + \beta \left(\frac{\#AKNOW_d^{i,j}}{\#A} \right) \geq \sum_{l=1}^{T-d} \delta^l \alpha \left[\frac{\Delta \#AKNOW_l^{-i}}{\#A} \right] \quad (11)$$

for all $i \in B$, $j \in A$, and all values of $1 \leq d < T-1$.

Once again, the exact specification of $\#BKNOW$ can be found in the Online appendix. For simplicity, consider a simplification that represents a cumbersome union of unions with the stand-in '*UNION*'. Now we can write the number of people in group B that i knows will know about a defection occurring in a round between $i \in A$ and $j \in B$ when no cross-group ties are present as:

$$\#BKNOW_l^{i,j} = \#\left[\overline{N}_B^{lr} \cup UNION_l^i \right]. \quad (12)$$

The first term in brackets is the average lr-neighborhood size of group B. When i and j are incident to a cross-group tie, the number of people in group B that i knows will know about a defection occurring in a round between $i \in A$ and $j \in B$ when a cross-group tie is present between i and j is:

$$\#BKNOW_l^{i,j} = \#\left[N_j^{lr} \cup UNION_l^i \right]. \quad (13)$$

Without cross-group ties, i considers the average neighborhood size in group B. With a cross-group tie to which i is incident, i considers the exact neighborhood size of his cross-group contact. This leads to the key result:

Proposition 4 (Dangerous cross-group ties): Player i's incentive to defect against out-group member j is strictly greater when i and j are incident to a cross-group tie and $\#N_j^{lr} < \#\overline{N}_B^{lr}$ for $1 \leq l \leq T$.[16]

In other words, when i's cross-group neighbor j's neighborhood is strictly smaller than the average neighborhood in the out-group, i has a strictly greater incentive to defect on j given this cross-group tie than given its absence. It follows that enforcing cooperation between

[16] This condition is sufficient for a cross-group tie to make full cooperation strictly more difficult, but a slightly weaker version is necessary and sufficient: it must be that the weighted sum in condition (1) is smaller, which does not necessarily require the neighborhoods to be weakly smaller at every value of l.

the two groups is strictly more difficult[17] when the binding group member is incident to a cross-group tie.

We can also establish the relative magnitude of the danger.

Corollary 4 (Particular danger of the poorly connected): The larger the difference between $\#N_j^{lr}$ and $\overline{\#N_B^{lr}}$ for $1 \leq l \leq T$, the more harm the cross-group tie between i and j does to full cooperation.

When a cross-group tie connects $i \in A$ to a member $j \in B$ with smaller-than-average neighborhood size, i faces a greater incentive to defect against j than if there were no such cross-group tie; the more below-average the neighborhood size, the bigger the impact the cross-group tie has on i's incentive.

Discussion

To illustrate the mechanism by which a cross-group tie may be dangerous, consider an example. Suppose there are two neighboring ethnic groups with no cross-group ties, and the two maintain peace by threatening retaliation. If someone in one ethnic group mistreats someone in the other ethnic group, the mistreated ethnic group retaliates for some amount of time, which takes the form of cross-group conflict. The two groups maintain civil relations because they know that an uncooperative act would spark cross-group conflict.

The incentive to behave civilly is strongest when the consequences of cross-group retaliation are greatest. The consequences are higher when more people are involved. People get involved in cross-group retaliation when they hear that they should. Networks that spread news of misbehavior widely and quickly maximize the consequences of retaliation and so increase the incentive to behave cooperatively.

Any individual considering behaving uncivilly toward the other group has to compare what he stands to gain from doing so with the consequences of retaliation that he initiates. First, note that the speed with which news reaches others varies within a network. Some positions within a network are more peripheral – positions from which news reaches the rest of the group more slowly – and some are more central – positions from which news reaches many others quickly. A central individual and a peripheral individual face different consequences from misbehavior; retaliation sparked by the central individual involves many people and exacts high costs quickly. Retaliation sparked by a more peripheral individual escalates more slowly, reducing the costs. Hence, more peripheral individuals within a network are not as strongly deterred from initiating conflict.

Now consider the way the out-group's network figures into each player's estimation of the cost of conflict. With no cross-group ties, no one knows *which* out-group member they are interacting with. They know what the out-group network looks like, but not who sits where in it – all non-coethnics are strangers. The out-group network matters for the cost estimation because the cost of conflict is increasing in the number of out-group members who are involved. This depends in part on how quickly, on average, news from any source in the out-group reaches any other source. If the out-group network has some positions that are more central and some that are more peripheral, and if a player just knows he is playing someone at random, his costs depend on the average over the number of out-group members who would know when he plays any of them.

If, instead, an individual i has a cross-group tie to some j in the out-group, he knows the identity of one of the nodes in the out-group network. He knows exactly how many out-group players would learn news originating with his cross-group neighbor. If his neighbor is more central, this makes him even less likely to behave uncivilly when randomly paired with this opponent. If his neighbor j is more peripheral, though, i may stand to gain more in a defection against j than in a defection against a randomly chosen out-group member.[18] If j is in a particularly vulnerable position, both groups would be better off if no one knew when they were playing the vulnerable j.[19]

[17] We can say that enforcing cooperation is more difficult when the conditions require a large minimum value of the discount factor δ to be satisfied.

[18] In the model, individuals who are incident to cross-group ties know information about the social structure of their contact, which is what potentially allows them to exploit the other individual. Individuals also keep this knowledge to themselves. If instead individuals shared the opportunity for exploitation with their in-group members, the prospects for conflict might be even worse. Likewise, if through a cross-group contact a person learns that not only is his cross-group neighbor a tempting target, but so is one of his co-ethnic friends, then the downsides of cross-group ties would be magnified.

[19] This result may seem surprising given the rough intuition of (Fearon, 1995) – one source of conflict in a rationalist explanation for war is asymmetric information. If both parties engaged in a costly bargain had complete information over resolve and capabilities, a bargaining range could accommodate peace, but asymmetric information can preclude an agreement preferable to war. The main difference here is that agents are not engaged in a costly bargain; they are engaged in a repeated interaction which is cooperative due to the threat of conflict if someone deviates. Cooperation now is worthwhile because punishment in the future is too expensive. The new information is effectively that there is one way to not cooperate now that earns much less expensive future punishment. If neither side had that information, cooperation now would be easier to maintain. Once that information is learned, cooperation is harder to maintain.

The dangers of imposing cross-group ties

The model does not capture every possible source of incentives for conflict or peace, but simply demonstrates the existence of a mechanism by which a cross-group tie can make interethnic peace strictly less likely. The tie conveys information about vulnerability that would otherwise remain unknown. This suggests that cross-group ties are not *necessarily* beneficial on net, which has important implications for empirical interventions that operate in the realm of civil society.

Due to the potential benefits of social contacts that span ethnic groups, development programs and peacekeeping interventions often contain a 'civil society' component with the goal of forging ties across groups (Belloni, 2001). These programs range from large-scale, multiyear efforts to fund and construct community organizations (e.g. Pickering, 2006) to smaller, one-off programs like concerts, soccer camps, and youth programs (e.g. Yablon & Katz, 2001).

The range in program subject, duration, type, and size means there is high variance in ways that ties are promoted. Case study approaches have attempted to evaluate program types and infer conditions under which these programs are most successful at forging relationships and promoting peace. For instance, the programs undertaken in post-conflict Bosnia appeared to be most effective when they offered individuals the chance for frequent, mutually dependent interethnic interactions (Pickering, 2006). The impression is that programs which are unsuccessful may be a waste of resources and have a low marginal impact on peace.

The above model suggests that the worst case scenario for a peace program is not that the marginal impact on peace is zero; the worse case scenario is that the marginal impact is negative.

Ties which connect peripheral individuals *and change nothing else* can generate incentives that undermine peace. Of course reality is more complicated than the model setup. Whether the net consequence of a tie in the real world is negative or not will depend on what else changes – especially details of the strategic environment like the structure of costs or preferences – which will depend on the nature of the intervention.

To explore the plausibility of dangerous cross-group ties, consider two hypothetical interventions. Intervention 1 invites members of both ethnic groups to a year-long education program which meets weekly, after which an organization will be created in which the participants expect to continue to meet indefinitely to perform some mutually beneficial task. Intervention 2 invites members of both ethnic groups to a day-long soccer camp.

Both result in non-coethnics meeting. The difference between the two programs may seem to be that the first program will forge deeper connections and the second one shallower ones, making the first program more likely to succeed and rendering the second less effective. However, there is a scenario in which the soccer camp is not just ineffective but actually reduces peace.

Suppose a member of ethnic group A meets a member of ethnic group B during the day of soccer, and through some casual conversations between play, gleans that the member of B is something of an outcast in his group. His friend group is small, he interacts with the rest of his group infrequently, and so forth. These indicators that the particular member of B is peripheral may make that member of B a tempting target of incivility to the other participant. The ability to identify a particularly isolated member of B may make conflicting with B strictly more profitable.[20]

The same information could be gleaned if the two had interacted via intervention 1, and so the cross-group tie could have negative consequences. However, intervention 1 might result in additional changes which may alter the strategic environment described above, potentially counteracting the downsides identified here. Perhaps the intervention's duration and content actually change the preferences of the participants, or the prospect of more frequent future interactions counteracts the incentive to exploit the out-group contact. Interventions that do not also alter the strategic environment in ways like these may not successfully counteract the incentive to act on the information that someone would be a profitable target.

Moreover, the extent to which an intervention is problematic depends on the recruitment strategy. Corollary 4 suggests that events which recruit as participants the most peripheral individuals in each group exacerbate the negative effects of any added cross-group ties. Cross-group ties to peripheral individuals are those that reveal the most profitable opportunities to take advantage of the out-group. If the intervention seeks to add cross-group ties between the most peripheral, it is seeking to add the potentially most dangerous interethnic ties.

Finally, this reasoning also suggests an additional potential source of selection bias when trying to infer

[20] For the model to apply, interactions between two individuals must have as an option some action that could be taken which would count as a defection, or uncivil, such that each person would find taking that action most profitable if they could get away with it. What exactly this looks like can vary: cheating, hurling insults, physical violence, etc.

the causal benefit of cross-group ties from observational data. Cross-group ties that emerge 'organically', that form between individuals within two groups for some natural reason and not as the product of an intervention, may differ from ties that are imposed (say via an intervention) in the attendant strategic environment. Ties that form naturally between people may be due to genuine fondness or mutual benefit, both of which may make the strategic structure of an interaction between individuals who have formed a relationship on their own different from the structure of an interaction between individuals introduced during an international intervention. If this is the case, cross-group ties that form naturally would be more likely to be helpful to interethnic peace than ties imposed without attempts to alter preferences or other features of the environment. Observing that neighboring ethnic groups have cross-ethnic ties and fight less frequently does not mean that adding cross-ethnic ties through any means would be helpful; to the contrary, some interventions to do so could strictly reduce interethnic peace.

Conclusion

The model formalizes a process by which interethnic cooperation is supported by threats of cross-group retaliation. Interethnic conflict is staved off by the threat of coordinated retaliation: social networks through which news spreads determine how successful this coordination will be. Ethnic groups with particularly dense networks wield a credible threat in response to cross-group misbehavior: news that someone misbehaved will spread rapidly to the rest of the group, allowing many to quickly join in the retaliation. Such groups can best deter conflict.

Sparse networks attenuate the costs of retaliation because coordinating the start of retaliation is slow. Sparse networks are problematic for a second reason as well: coordinating the *end* of a period of retaliation becomes more difficult. In the event that someone has misbehaved, the groups enter a period of costly retaliation that will end after a finite period of time. If someone encounters a retaliator before being told that the retaliation period has begun (because his group's network is sparse and news spreads slowly), he may still be earnestly retaliating once the period is supposed to end. When this happens, his act of retaliation, understood by the other group to be uncalled for, trips a new period of retaliation. This can occur again and again, resulting in protracted conflict.

The model also reveals a mechanism by which some cross-group ties are strictly harmful for interethnic peace. In any network, the more peripheral positions – individuals with few connections to many others – would be more tempting targets of misbehavior. Cross-group ties between individuals in these positions reveal this tempting position and make initiating conflict strictly more likely. Not only does this mean that cross-group ties are not necessarily helpful, but this also suggests that programs aiming to create new cross-group ties should proceed cautiously. Uncareful efforts to promote peace by imposing cross-group ties can do more harm than good, especially if they are aimed at the most peripheral members of both groups.

The model presented here illustrates the potential value of accounting for networks that spread information among groups. Given the surge in network data collection projects in the field (e.g. Conley & Udry, 2010; Barr, Ensminger & Johnson, 2010; Jackson, Rodriguez-Barraquer & Tan, 2012; Banerjee et al., 2014; Larson & Lewis, 2015b), theory relating the precise internal structure of groups to cooperative outcomes may prove to be especially useful.

Online appendix

The Online appendix is available at http://www.prio.org/jpr/datasets.

Acknowledgments

Thanks to the participants of the Networked International Politics conference at the University of St. Gallen for their helpful feedback and suggestions.

References

Banerjee, Abhijit V; Arun Chandrasekhar, Esther Duflo & Matthew O Jackson (2014) Gossip: Identifying central individuals in a social network. Working paper, SSRN 2425379 (http://arxiv.org/pdf/1406.2293.pdf).

Barr, Abigail; Jean Ensminger & Jeffrey C Johnson (2010) Social networks and trust in cross-cultural economic experiments. In: Karen S Cook, Margaret Levi & Russell Hardin (eds) *Whom Can We Trust?* New York: Russell Sage Foundation, 65–90.

Belloni, Roberto (2001) Civil society and peacebuilding in Bosnia and Herzegovina. *Journal of Peace Research* 38(2): 163–180.

Burt, Ronald S (2000) The network structure of social capital. *Research in Organizational Behavior* 22: 345–423.

Cho, Myeonghwan (2011) Public randomization in the repeated prisoner's dilemma game with local interaction. *Economics Letters* 112(3): 280–282.

Chyzh, Olga (2016) Dangerous liaisons: An endogenous model of international trade and human rights. *Journal of Peace Research* 53(3): 409–423.

Conley, Timothy G & Christopher R Udry (2010) Learning about a new technology: Pineapple in Ghana. *American Economic Review* 100(1): 35–69.

Fearon, James D (1995) Rationalist explanations for war. *International Organization* 49(3): 379–414.

Fearon, James D (2004) Why do some civil wars last so much longer than others? *Journal of Peace Research* 41(3): 275–301.

Fearon, James D & David D Laitin (1996) Explaining interethnic cooperation. *American Political Science Review* 90(4): 715–735.

Granovetter, Mark S (1973) The strength of weak ties. *American Journal of Sociology* 78(6): 1360–1380.

Habyarimana, James; Macartan Humphreys, Daniel N Posner & Jeremy M Weinstein (2007) Why does ethnic diversity undermine public goods provision? *American Political Science Review* 101(4): 709–725.

Habyarimana, James; Macartan Humphreys, Daniel N Posner & Jeremy M Weinstein (2009) *Coethnicity: Diversity and the Dilemmas of Collective Action*. New York: Russell Sage Foundation.

Haim, Dotan A (2016) Alliance networks and trade: The effect of indirect political alliances on bilateral trade flows. *Journal of Peace Research* 53(3): 472–490.

Horowitz, Donald L (1985) *Ethnic Groups in Conflict*. Berkeley, CA: University of California Press.

Jackson, Matthew O; Tomas Rodriguez-Barraquer & Xu Tan (2012) Social capital and social quilts: Network patterns of favor exchange. *American Economic Review* 102(5): 1857–1897.

Jeon, Sangick (2014) Interethnic complementarities and decentralized intergroup cooperation. Working paper (http://goo.gl/Ar6ko2).

Jha, Saumitra (2013) Trade, institutions, and ethnic tolerance: Evidence from South Asia. *American Political Science Review* 107(4): 806–832.

Kandori, Michihiro (1992) Social norms and community enforcement. *Review of Economic Studies* 59(1): 63–80.

Kinne, Brandon (2016) Agreeing to arm: Bilateral weapons agreements and the global arms trade. *Journal of Peace Research* 53(3): 359–377.

Larson, Jennifer M & Janet I Lewis (2015a) Rumors, kinship networks and rebel group formation. Working paper (http://goo.gl/Yk0lTr).

Larson, Jennifer M & Janet I Lewis (2015b) Ethnic networks. Working paper (http://goo.gl/I4gEIv).

Lin, Nan (1999) Building a network theory of social capital. *Connections* 22(1): 28–51.

Lin, Nan; Karen S Cook & Ronald S Burt (2001) *Social Capital: Theory and Research*. New Brunswick, NJ: Transaction.

Lipset, Seymour M & Stein Rokkan (1967) *Cleavage Structures, Party Systems, and Voter Alignments: An Introduction*. New York: Free.

Lupu, Yonatan & Paul Poast (2016) Team of former rivals: A multilateral theory of non-aggression pacts. *Journal of Peace Research* 53(3): 344–358.

Miguel, Edward & Mary Kay Gugerty (2005) Ethnic diversity, social sanctions, and public goods in Kenya. *Journal of Public Economics* 89(11–12): 2325–2368.

Nava, Francesco & Michele Piccione (2014) Efficiency in repeated games with local interaction and uncertain local monitoring. *Theoretical Economics* 9(1): 279–312.

Pettigrew, Thomas F & Linda R Tropp (2006) A meta-analytic test of intergroup contact theory. *Journal of Personality and Social Psychology* 90(5): 751–783.

Pickering, Paula M (2006) Generating social capital for bridging ethnic divisions in the Balkans: Case studies of two Bosniak cities. *Ethnic and Racial Studies* 29(1): 79–103.

Putnam, Robert D (2001) *Bowling Alone*. New York: Simon & Schuster.

Varshney, Ashutosh (2003) *Ethnic Conflict and Civic Life: Hindus and Muslims in India*. New Haven, CT: Yale University Press.

Ward, Hugh & Han Dorussen (2016) Standing alongside your friends: Network centrality and providing troops to UN peacekeeping operations. *Journal of Peace Research* 53(3): 392–408.

Wilkinson, Steven I (2004) *Votes and Violence*. Cambridge: Cambridge University Press.

Wilson, Maya; David R Davis & Amanda Murdie (2016) The view from the bottom: Networks of conflict resolution organizations and international peace. *Journal of Peace Research* 53(3): 442–458.

Wolitzky, Alexander (2013) Cooperation with network monitoring. *Review of Economic Studies* 80(1): 395–427.

Yablon, Yaacov B & Yaacov J Katz (2001) Internet-based group relations: A high school peace education project in Israel. *Educational media international* 38(2–3): 175–182.

JENNIFER M LARSON, b. 1984, PhD in Political Science (Harvard University, 2012); Assistant Professor of Politics, New York University (2012–); research interests in political economy, cooperation, and the role of social networks in development and conflict.