Rumor Dynamics in Ethnic Violence

Ravi Bhavnani Michigan State University
Michael G. Findley Brigham Young University
James H. Kuklinski University of Illinois at Urbana-Champaign

Most scholars believe that rumors spark ethnic violence around the world, yet political scientists have been slow to understand the dynamics of these rumors, including whether they emerge and survive as a matter of course or require certain preconditions. Because empirical observation cannot readily answer this question, we use agent-based modeling to explore whether rumor survival depends on selected group follower and leader characteristics. Our first set of experiments assumes one group. They show that violence-promoting rumors are not inevitable, nor are group leaders necessary for their propagation. But rumor survival peaks when leaders espouse extreme beliefs and interact frequently with their followers. Our second set of experiments assumes rival groups whose members interact on a limited basis. Some of the experiments confirm the one-group conclusions. Others reveal that when one group's leaders persist in advocating moderation, rumor propagation remains low in both groups. Finally, when rival leaders try to outbid each other, rumor survival increases markedly in both groups. By way of conclusion, we consider the policy implications of these findings.

Rumors pervade all aspects of life. Often unverified, they can be right or wrong. When people believe rumors to be true, the consequences can be considerable. Financial traders hear about a pending company takeover, and the value of its stock skyrockets or plummets (Rose 1951; Pound and Zeckhauser 1990; Zivney, Bertin, and Torabzadeh 1996). News broadcasters report that two bottles of Tylenol were injected with poison. The news ignites a rumor that all bottles of Tylenol contain poison, with the result that the demand for and price of a substitute, Advil, doubles. A husband hears from friends and associates, who hear from each other, that his wife is having an affair. The rumor proves false, but only after the enraged husband shoots his wife and children and then himself.

All of these consequences pale in comparison to those of racial- and ethnic-centered rumors. Take the infamous 1943 Detroit riot. On June 20 of that year, black and white youths clashed in minor skirmishes on Detroit's Belle Isle (Brown 1943; Capeci and Wilkerson 1991; Langlois 1983; Sugrue 1996). Shortly thereafter, rumors began to fly. Two black youths told a black crowd gathered at a local club that whites had thrown a black mother and her child off the Belle Isle Bridge. In a nearby white neighborhood, another rumor—that blacks had raped and murdered a white woman on the same bridge—moved swiftly among its residents. Within hours, a riot began. The rioting ended 36 hours later, at considerable cost: 34 people, including six Detroit policemen, dead; another 75 policemen injured; more than 1,800 people arrested for looting; and property damage that climbed into the millions of dollars. Neither rumor proved true. Across the Atlantic and more than half a century later, a rumor that Paris police had caused two Muslim youths to be electrocuted in a power substation in Clichy-sous-Bois provoked a level of violence that the city had not experienced in decades (ABC 2005).

Ethnic-centered rumors appear to produce especially dire effects in developing and undeveloped countries. Following the news in Delhi that Indian Prime Minister Indira Gandhi had been killed, extreme rumors—"the Sikhs are celebrating Gandhi's death by distributing sweets," "trainloads of hundreds of dead Hindu bodies just arrived at the Delhi station," "the Sikhs have poisoned the water supply"—proliferated (Tambiah 1996, 237). Although most rumors proved false, rampant and widespread killing of Sikhs followed. In 1992, the Hutu began to hear rumors that the Tutsi were planning to kill their leaders in the Bugesera region of Rwanda. The Tutsi did nothing, but the story persisted. Hutu militias and civilians subsequently massacred hundreds of Tutsi.

The Journal of Politics, Vol. 71, No. 3, July 2009, Pp. 876–892 doi:10.1017/S002238160909077X

ISSN 0022-3816

Hundreds became hundreds of thousands during the 1994 Rwandan genocide, when Hutu extremists killed up to a million ethnic Tutsis and moderate Hutu sympathizers. The catalyst: a then- and still- unsubstantiated rumor among the Hutu that the Tutsi had shot down an airplane carrying Rwandan President Juvenal Habyarimana, who died in the crash (Diamond 2004; Gourevitch 1998; Mamdani 1996).

These examples are not unique. In discussing the precipitants of ethnic riots, Horowitz and Varshney write that "rumors of aggression committed by members of the targeted ethnic groups are nearly universal in events that precede deadly ethnic riots" (2003, 5).[1] Their statement fuses two distinct and equally important observations. First, the presence of rumors seems to be necessary if not sufficient for the rise of ethnic violence.[2] Rumors are, at the least, a prerequisite of ethnic violence, and they likely suffice most of the time. Second, violence-promoting rumors take a remarkably similar form across time and contexts: they are extreme and threaten individuals through their group identities. Indeed, the very same words and themes—ethnic rivals have caused bodily harm to some of the group's own members or have damaged or destroyed its culturally important institutions—appear over and over again.

That such rumors should proximally cause ethnic violence makes good sense. They are, after all, a type of information, and it takes the communication of information to initiate any form of political behavior (Bikhchandani, Hirshleifer, and Welch 1992). Two features of violence-promoting rumors favor them over other information types as a stimulus for action. First, they often lack verifiability when verification is most needed: prior to a decision to initiate violent action against another group. Lacking evidence challenging a rumor's truth or accuracy, group members can easily accept it. Second, rumor initiators play on already-high levels of intergroup animosity to evoke strong emotions. The path from heightened emotions to impulsive, violent behavior is short (Shibutani 1966).

Moreover, the rumor-violence dynamic is self-reinforcing, in that rumor-generated violence has the perverse effect of confirming the rumor's veracity even when it is false. The Hutu kill thousands of Tutsi in response to a rumor that the Tutsi are about to attack their leaders. The attack never occurs, in the eyes of the Hutu, only because they preemptively killed the Tutsi. By implication, the next time the Hutu hear of imminent Tutsi threat, they should again react quickly and strongly. No wonder that Perice, in his intensive anthropological study of Haiti, found "the throng of rumors told over and over" to "calcify into accepted representations of social reality and political life" (1997, 1).

But are there conditions under which violence-promoting rumors emerge and survive and others under which they do not? This question should interest scholars and practitioners striving to reduce if not stop interethnic violence. Controlling rumor propagation should be easier than stopping rumor-generated violence once it begins.

Despite the immeasurable human toll that violence-promoting rumors have extracted, political scientists know little about their dynamics. In fairness, tracking them empirically fringes on the impossible. Unless the group is small, researchers cannot expect to observe all of the simultaneously occurring interactions among the relevant actors. Putting this obstacle aside, even the most conscientious researchers will not see all the failed attempts to propagate rumors. Indeed, it appears that this limitation has led most students of ethnic politics to assume implicitly that rumor emergence and survival occur as a matter of course rather than conditionally. But why make this unnecessary and potentially wrong assumption?

If empirical analysis cannot adequately reveal the dynamics of violence-promoting rumors, what can? For reasons discussed later, and following in the footsteps of economists, we adopt agent-based modeling (ABM) as our analytical tool. A relatively new methodology in political science (but see Axelrod 1997, 2006; Epstein 2002; Miller and Page 2007), ABM is better suited than empirical analysis to determine the conditions under which rumor emergence and survival are most and least likely, given a set of interacting individuals who behave according to a set of rules specified by the researcher. Equally important, ABM lends itself well to counterfactual analysis, i.e., to posing and answering "what if" questions that detailed empirical data, even if available, cannot answer. For example, the researcher can run an experiment to analyze rumor survival among group followers only and then another, identical experiment except for the inclusion of group leaders.

[1] In our review of primary sources and secondary accounts of rumors in India from 1947 to the present, (Brass 1997; Engineer 1995; Horowitz 2001) 27 of the 28 rumors we identified were connected to violent riots.

[2] Tambiah extends the impact of rumors to one of the dominant forms of political violence worldwide, insurgency: "[Rumors] are the currency of mass movements . . . the most potent communicative acts in the spread of insurgency and riot" (1990, 757).

We offer ABM as a complement to, not as a substitute for, empirically based research. Indeed, ABM and empirical observation are well-suited to iteration: existing empirical research guides model specification; the modeling results either do or do not confirm the empirical observations; when confirmatory, the modeling results reveal the mechanisms and dynamics by which the observed outcomes occur; when not, they offer new directions for empirical research, which, when completed, guides the next generation of agent-based models. Like all modes of analysis, ABM requires researchers to select factors carefully and show restraint in the number they choose.

The following section first identifies mechanisms that are essential to the propagation of violence-promoting rumors and then discusses and justifies the potential conditioning factors we included in our model. In a nutshell, we chose characteristics of rank-and-file members (followers) and their leaders that students of ethnic politics generally believe to enhance or inhibit rumor propagation.[3] The group follower characteristics come from selected collective action studies (e.g., Gould 1999; Wood 2003) and the leader characteristics from the "ethnic entrepreneurs" literature (e.g., Lake and Rothchild 1996). Where appropriate, we offer specific hypotheses. Sometimes, however, the complexity precludes hypothesis formulation. Predicting rumor propagation within and across two rival ethnic groups while varying their leaders' behavior, their members' initial beliefs, and the level of interaction and communication across groups, for example, would be folly. ABM has a unique capacity to reveal how different combinations of these factors shape the emergence and survival of violence-promoting rumors. Subsequent sections describe the model and report the experimental results.

The Propagation of Violence-Promoting Rumors

A rumor is a belief that diffuses, usually rapidly, among members of a particular population. The belief is often unverified, and thus it can be factually wrong (Allport and Postman 1947). With respect specifically to interethnic group violence, it tends to come in an extreme form that heightens emotions, instigates panic, intensifies distrust, and justifies brutality. A belief most readily diffuses and becomes a rumor when it affirms the views that one ethnic group's members already hold about the other group's members.

Because ethnic group members live in close proximity to one another, they frequently interact. Word-of-mouth communications, therefore, are instrumental in spreading violence-provoking beliefs. These communications will be especially fast when individuals form into crowds, a condition that often precedes the onslaught of violent behavior. That crowds can become violent almost instantaneously reflects the effectiveness of belief diffusion as a coordinating mechanism: interacting co-ethnics publicly express their otherwise-private beliefs; this public expression immediately increases any particular co-ethnic's certainty about the extent to which others share the belief; seeing that many share the belief, the co-ethnic then also expresses it publicly.

Our use of the conjunction, "emergence and survival," warrants comment. A rumor emerges when a large number of the group's members come to share the belief. It survives when this large number continues to hold the belief long enough for the rumor to influence the group's behavior across repeated encounters.[4] Emergence, then, precedes but does not guarantee survival. We assume that violence-promoting rumors must persist for them to spark interethnic violence.

Mechanisms for Rumor Emergence

Two mechanisms are necessary for belief diffusion and thus rumor emergence: transmission and updating (Banerjee 1993; Kosfeld 2005). Transmission requires that two agents interact with each other, with one conveying a belief to the other. Widespread information transmission does not imply that the information is always accepted. Updating occurs only when the receiver of the transmitted belief replaces his or her existing, and different, belief with it. Many potential recipients judiciously assess the new information before accepting it. The members' probabilities of updating depend on at least two factors: the frequency with which other group members transmit

[3]We do not make broader distinctions between ethnic groups and, say, governments, rebel groups, and terrorist groups. These distinctions require manipulating additional global parameters, such as the number of potential follower transmissions, the threshold beyond which the veracity of a rumor fails to matter, and the returns and costs associated with individual updating. They also require more complicated manipulations than we undertake in this initial, exploratory study.

[4]"Large number" and "long enough" are necessarily vague terms. We define these terms more explicitly in the context of the formal model, where a violence-promoting rumor "survives" if the mean group belief exceeds a certain threshold for a given number of periods during which co-ethnics transmit information among themselves, and in certain cases, among nominal rivals.

the belief to them and the similarity of the transmitted belief to their existing beliefs. Both factors help members to assess the credibility of a transmitted belief.

Ethnic group members interact with each other and, in some cases, with members of opposing groups over time. A given individual might consistently act as a transmitter of a belief and experience more or less success getting others to update their beliefs. The individual might also begin as a receiver and refuse to update or begin as a receiver, update his or her belief, and become a transmitter thereafter. As interactions proceed over time, the overall distributions of people holding the relevant belief usually change.

Mechanisms for Rumor Survival

The survival of violence-promoting rumors occurs when, after the elapse of time, the distribution remains stable at a high level of belief acceptance, that is, when most group members accept and continue to accept the relevant belief. Survival therefore depends on a complex and changing interplay between the distribution of beliefs among group members and an individual member's propensity to accept the emerging belief (Banerjee 1992). As interactions among group members progress, the proportions holding the relevant belief change, which in turn affects the likelihood of any one member accepting or rejecting the belief.

Conformity serves as an especially critical mechanism in this process (Smelser 1962). Members typically do not know the belief's credibility or who will participate in the violence until the episode has ended. Having to accept or reject the belief prior to that time, they often go along with what they view as the majority belief, which they surmise from their interactions with others, and which, in turn, depends on the distribution of beliefs at the moment.

When most members of a group accept and act on a belief, its veracity does not matter for those who adopt it. Those acting in unison share the rewards of collective action and can punish the readily identifiable minority for their nonconformity. In contrast, when multiple beliefs arise from group member interactions, so that no majority of beliefs and behavior exists, then the veracity of individual members' beliefs can determine whether they are awarded or penalized. Over time, group members become more or less prone to accepting a particular belief (or set of beliefs), giving rise to its demise or survival. Group follower and leader behaviors and characteristics, the patterns of interaction amongst followers and leaders, and the nature and extent of communication across groups determine which—demise or survival—prevails.

To summarize our discussion thus far, a violence-promoting rumor emerges within a group when a relevant belief is transmitted by means of interpersonal communication and more and more members of the group accept it; the rumor survives when the changed distribution persists over time. Consider, as an analogy, a community health epidemic. Any one individual's state of health does not comprise an epidemic, but the health deterioration of many individuals living in the community does. Once an epidemic has emerged, it can last for a long or a short time.

Mediating Factors

Our primary concern is the identification of conditions that enhance or inhibit the emergence and survival of violence-promoting rumors. In the case of the health epidemic, for example, frequent interaction among community residents would increase the diffusion of the disease while immunization would slow, if not eliminate, it. There is no immunization equivalent for halting the diffusion of political beliefs. However, scholars who have derived generalizations from available case studies report that certain ethnic group characteristics influence the emergence and survival of violence-promoting rumors.

Rumor propagation requires that beliefs diffuse among ordinary group members, who vary on a variety of characteristics that can independently shape the emergence and survival of rumors. We consider two: the group's initial average belief about the rumor's topic and the distribution of beliefs around that average. Imagine that someone initiates a belief saying that rivals intend to assault its female members within the next few weeks. The average belief among group members might be that such an assault is not likely or, alternatively, that it is eminent. Similarly, the distributions around either of these average beliefs might be small or large. The combination of an initial average belief and the cohesion around it, we hypothesize, shapes rumor emergence and survival:

H1: All else equal, violence-promoting rumors are most likely to emerge and survive when the initial average group belief is extreme and most members hold beliefs close to it.

H2: All else equal, violence-promoting rumors are least likely to emerge and survive when the initial average group belief is moderate and

most members hold beliefs close to it or *when the initial average group belief is moderate and most members' beliefs vary widely around it.*

With respect to ordinary members, the hypotheses predict that both factors—an initially extreme average belief and cohesion around it—enhance rumor propagation, although the initial average belief is more critical.

Many scholars have intimated that rumor propagation requires active and visible group leaders (Brass 1997; Engineer 1995). In their view, ethnic group leaders push certain beliefs to incite violence; take away these leaders, and the rumors and the interethnic violence they generate will fall precipitously. In two of the cases cited earlier, however, the 1943 Detroit riot and the Muslim uprising in Paris, group leaders were notably absent at the outset of the violence. These two cases question the necessity of leaders in rumor propagation. To say they are not necessary does not imply that they lack independent influence, however. Therefore, we hypothesize the following:

H3: All else equal, violence-promoting rumors are more likely to emerge and survive in the presence than in the absence of ethnic group leaders.

This hypothesis speaks only to presence; it does not distinguish among leaders who differ in their behaviors or in their relationships to rank-and-file members.

Anecdotal evidence suggests that most but not all leaders seek to disseminate extreme rather than moderate beliefs because the former, in generating fear, anger, and resentment, more effectively incite violence (Kaufman 2001; Peterson 2002). One might infer, therefore, that extremely stated beliefs also stand a better chance than moderately stated ones of becoming full-fledged rumors. To date, however, this inference remains untested. Scholars relying on empirical observation have often associated the mere presence of leaders' statements with the onset of interethnic violence, without considering whether the extremity of those statements affects rumor emergence and survival (Gagnon 1995).

Group leaders also differ in their connectedness to followers. In some instances, leaders live among their followers and interact daily with them while in others they interact only occasionally. The more frequently leaders interact with their followers, presumably, the greater their chances of propagating rumors among them. The interpersonal interactions can occur through formal associations such as parties, governmental agencies, and unions and business groups; or outside of those contexts through everyday interaction. In any event, increased connectivity implies increased information flow, including the flow of beliefs.

Combining the extremism of leaders' statements and their connectedness to group followers generates the following three hypotheses:

H4: All else equal, violence-promoting rumors are most likely to emerge and survive when ethnic group leaders state extreme beliefs and are highly connected to their followers.

H5: All else equal, violence-promoting rumors are least likely to emerge and survive when ethnic group leaders state moderate beliefs and are highly connected to their followers.

H6: All else equal, violence-promoting rumors are least likely to emerge and survive when ethnic group leaders state extreme or moderate beliefs and are weakly connected to their followers.

Taken together, these three hypotheses posit that leaders most successfully propagate rumors when they state extreme beliefs and are highly connected to their followers, and least successfully propagate rumors when they state moderate beliefs.

ABM also allows us to explore whether and how follower and leader behaviors and characteristics interact to shape rumor dynamics. Particularly important, do group leaders independently affect the emergence and survival of rumors under all follower conditions? Although the preponderance of opinion among students of ethnic violence says yes, Horowitz describes group leader influence as conditional: ''Even those [leaders] with strong motives cannot produce significant violence when masses of people do not share [their] intense preferences'' (2001, 253). If Horowitz is right, then the supposed widespread influence of group leaders in the propagation of rumors could be more illusory than real.

Thus far, we have assumed a single, isolated ethnic group. Although useful for generating initial hypotheses about within-group rumor dynamics, this assumption does not characterize most important real-world situations, where rival ethnic groups exist and rumors emerge and survive, or not, as a function of both within and across group interactions and communications.

Although introducing a second ethnic group increases realism, it also adds complexity. Rumor emergence and survival within each group potentially now depend on the relative sizes of the two groups, the levels of interaction and communication across

them, and the initial characteristics of each group's followers and leaders. The introduction of a second group also raises a set of questions that did not arise in our discussion of a single group: What happens when rumor-related information flows more heavily in one direction than the other? Does it matter if members of one group discount the information they receive from members of the other group, either dismissing it or reacting by changing their existing beliefs in a contrary direction? What happens when one group's leaders propagate extreme rumors and the other group's do not?

The formulation of a priori hypotheses now makes little sense. They would be simplistic speculations that belie the complex, simultaneously occurring across and within group processes by which rumors do or do not emerge and survive. Fortunately, ABM is uniquely designed to study this very type of complexity.

Modeling Rumors

The dynamics underlying rumor propagation include the following: heterogeneous individuals adapting as they interact with others; group characteristics changing as a function of these individual-level interactions; and these changing group characteristics, in turn, shaping the individual-level interactions. Constructed in a ''bottom-up'' manner, beginning with microlevel entities and dynamics, ABMs are well suited to model situations characterized by numerous, adaptive, and heterogeneous agents who interact with other agents and their changing environments.

In an ABM, an agent's state at any point in time can include its beliefs, its social connections (e.g., identities, memberships, and social networks), and a memory of recent events and interactions. In addition, ABMs define agents by their decision-making heuristics and capabilities to act in response to inputs from other agents and a nonagent environment. They can endow agents with a capacity to adapt, thus allowing them to change their heuristics in response to their own experiences. The environment, in turn, changes in response to agent behaviors. But it also follows its own dynamical rules. ABMs, therefore, embody elaborate and interlaced feedback relationships, leading to nonlinear, path-dependent dynamics, where small changes in individual heuristics, characteristics, or interaction patterns generate large changes in collective outcomes, a hallmark of complex adaptive systems.

We do not claim to offer an exhaustive model of rumor transmission. Like all models, ABMs are simplified representations of the entities and processes of interest to the modeler. Whether using game-theoretic models, statistical models, or ABMs, a researcher must have a priori expectations about which factors to include, recognizing, of course, that not all potentially relevant factors can be taken into account. ABMs are especially valuable when the variables and assumptions built into it represent conditions that empirical observation has already identified as crucial. Then they can document the mechanisms that presumably produce those real-world observations.

The Model

Following the earlier definition of rumors as unverified and high-velocity transmissions of information, rumor propagation in the model entails the agents sharing their beliefs with others and also accepting or rejecting others' beliefs. For a rumor to emerge and survive (in modeling terms, to become dominant), large proportions of the actors must accept it. Of course, not all rumors become dominant, and our analysis seeks to identify those conditions under which violence-promoting rumors emerge and survive over time.

Our framework builds upon a tag-tolerance model first specified by Riolo, Cohen, and Axelrod (2001; hereafter RCA). The model is intended to be general and thus not limited to particular substantive domains. In the RCA model, each agent is randomly assigned a tag, $\tau \in [0, 1]$, and a tolerance threshold, $T \in [0, 1]$. Agents are then paired at random such that when paired, an agent A can choose to make a costly donation to the other agent B if B's tag is within A's tolerance threshold, that is, if both tags are sufficiently similar, or $|\tau_A - \tau_B| \leq T_A$. If A donates to B, A pays a cost c and B receives a benefit b. After all agents have participated in all pairings, agents are reproduced on the basis of their scores relative to the scores of others, and the possibility of mutation—a random change to an agent's tag, tolerance, or both—is also introduced. Each run of the RCA model consists of 100 agents and 30,000 generations, and each experimental condition, defined by the combination of selected parameters, is replicated 30 times.

The emergence (or demise) of dominant tag clusters in the RCA framework makes this an appropriate starting point for modeling rumor emergence and survival. Much like the tag clusters in the RCA model, a select set of rumors tends to recur in incident after incident, raising the question of why such sensationalism fails to lose credibility over time. Rumors constitute a problem-solving tool for groups

to cope with uncertainty in that they publicly communicate private beliefs about the state of the world. Rumors heighten emotions, instigate panic, intensify distrust, and justify brutality, all of which serve to increase group leaders' capacities to overcome collective action problems and compel individual members to participate in violence.

Given that group leaders generally know the behavior of their members, they can readily identify and penalize nonconformity in a given episode (Fearon and Laitin 1996). Over time, both leaders and followers learn. Through trial and error, leaders learn which rumors are more successful in generating ethnic violence. They also develop more effective mechanisms for imposing "noncompliance costs" on deviants. In a similar vein, followers learn which rumors are spread by whom and with what associated costs.

In our adaptation of the RCA model, tags represent beliefs, tolerance denotes an agent's receptivity to new beliefs, and donations capture the acceptance of other agents' beliefs. A violence-promoting rumor emerges and survives when a large proportion of agents adopt and retain the same extreme belief over repeated encounters.[5]

Our specific application of the RCA model requires selection of an appropriate interaction topology. The interactions among agents in our model take the form of random pairing with replacement. This choice of interaction topology recognizes that most members of ethnic groups tend to be spatially proximate and well-mixed and that there is a rapid diffusion of beliefs and an exponential rise in agitation immediately before the outbreak of violence. This topology represents the descriptions that Horowitz (2001) and other astute observers have given of how interethnic violence commonly erupts in the real world.

Other interaction topologies exist, most notably star networks and "small worlds" consisting of friends, family, and nearby neighbors. A star network facilitates the rapid diffusion of beliefs from a highly centralized source, but perhaps without much interpersonal reinforcement. In contrast, Centola and Macy (2007; also see Watts 1999) have shown that while new information diffuses relatively slowly through "small world" networks, it strongly shapes beliefs, attitudes, and behavior because of interpersonal reinforcement. Although these and other interaction topologies exist, their inclusion would greatly increase the complexity of an already complex model, which we designed as an initial effort to understand the dynamics of violence-promoting rumors.[6] Nevertheless, we reran selected analyses using two alternative network topologies. The appendix reports the results.

In our specific application of the RCA model, we differentiate between two types of agents: rank-and-file members, or followers, and group leaders. Each follower is endowed with a belief, $\tau \in [0, 1]$, and a threshold for updating this belief, $T \in [0, 1]$; and each acts as a potential transmitter of information with M randomly chosen (with replacement) co-ethnics. When $M = 3$, for example, each follower has three opportunities to transmit information, and, on average, is chosen three times as a potential recipient of new information. When paired with another member of the group j, follower i can choose to accept new information and update her belief if the difference in i's and j's beliefs is within i's tolerance threshold, or $|\tau_{i,t} - \tau_{j,t}| \leq T_i$. When updating, i's posterior beliefs are a convex combination of her prior beliefs and j's beliefs. Thus, if $q = |\tau_{i,t} - \tau_{j,t}|$, then $\tau_{i,t+1} = q(\tau_{i,t}) + (1 - q)(\tau_{j,t})$.

Updating serves as the mechanism by which new information is transmitted between two agents, and it is essential for the spread and emergence of a rumor. We distinguish between extreme ($\tau \rightarrow 1$) and moderate ($\tau \rightarrow 0$) beliefs, where the former promote interethnic violence and the latter promote interethnic cooperation, or at least do not promote violence.

Finally, the distribution of τ determines the levels of belief cohesion among followers. For example, if $\tau = 0.25$ for all followers, then the group is considered highly cohesive or unitary. To be sure, $\tau = 0.25$ is an extreme case, for there is little chance that *all* members hold the same belief. In our model, cohesion is initially distributed as a deviation, σ_τ, surrounding mean group beliefs, $\bar{x}_\tau$. We vary cohesion, where 0.1 is low and 0.025 is high.

To illustrate the logic of transmission and updating in the model as applied to followers, consider the example depicted in Figure 1. A follower—the *transmitter*—randomly picks M partners—or *recipients*—to interact with on a pair-wise basis. Assume that $M = 3$ and the transmitter's belief is $\tau_i = 0.6$. When selected to play, i begins by picking j, who has the values $\tau_j = 0.9$, and $T_j = 0.4$, as a potential recipient. This constitutes a pairing in which recipient j harbors a belief more conducive to interethnic violence than

[5]Rumors that express moderate beliefs can also emerge and survive, although we restrict our focus to extreme-belief rumors (Knapp 1944).

[6]We are indebted to an anonymous referee for pushing us to provide a stronger and more complete rationale for our choice of interaction topology. In subsequent work, we will look more closely at how different topologies affect belief diffusion and rumor propagation.

TABLE 1 Model Parameters

Agent-Level Parameters	
τ_i	Follower belief
τ_{i^*}	Leader belief
T_i	Threshold for updating follower belief
γ_i	Tag denoting group membership
Group-Level Parameters	
$\bar{x}_\tau$	Mean group beliefs
σ_τ	Group cohesion
p	Probability of cross-group interaction
Global Parameters	
N	Number of agents
M	Number of potential follower transmissions
M^*	Number of potential leader transmissions
R	Return
C	Cost
τ_w	True state of the world
δ	Threshold over which veracity of rumor fails to matter
μ_τ	Mutation rate for follower belief
μ_T	Mutation rate for updating follower belief threshold

Note: The probability of cross-group interaction is only applicable to runs in which rival groups exist.

FIGURE 1 Agent Interaction

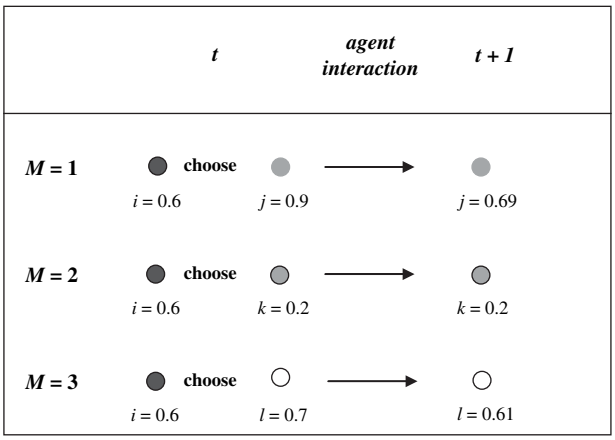

Note: This figure depicts a hypothetical example of agent i's interactions (i has $M=3$ opportunities to interact) with agents j, k, and l. The figure displays changes to the beliefs of agents j, k, and l contingent upon the acceptance of a rumor transmitted by i.

transmitter i; thus there is a possibility for i to exert a moderating influence on j's beliefs. Recall, however, that updating will only occur if the recipient is sufficiently tolerant to accept the transmitter's belief; otherwise the recipient will retain her original belief. Formally, j will accept information passed on by i only when $|\tau_{j,t} - \tau_{i,t}| \leq T_j$. Since this condition holds, j updates her beliefs as follows: $\tau_{j,t+1} = q(\tau_{j,t}) + (1-q)(\tau_{i,t})$, or 0.69, with $q = |\tau_{j,t} - \tau_{i,t}|$. As a result of the interaction with i, j becomes less extreme in her beliefs. Next, i picks k $(\tau_k = 0.2, T_k = 0.3)$, who is moderate but less tolerant of new information than j. Since $|\tau_{k,t} - \tau_{i,t}| \leq T_k$, k does not accept information transmitted by i, and $\tau_{k,t+1}$ remains at 0.2. Finally, i selects l $(\tau_l = 0.7, T_l = 0.2)$. Because the information provided by i falls within l's tolerance threshold, $\tau_{l,t+1}$ falls to 0.61.

This illustration represents the interaction between one transmitter and three potential recipients. More generally, each follower can act as transmitter with up to M potential recipients in each generation. The same group follower can also be selected as a recipient by other members of the group. Because all followers can cause others to update their beliefs as well as update their own beliefs, aggregate beliefs can become more or less extreme.

Leaders i^* differ from followers in that their beliefs are fixed or resistant to change. In addition, leaders vary on two key dimensions: the extremism of their messages and their connectivity among the masses. We model leader extremism as either low $(\tau_{i^*} = 0.1)$ or high $(\tau_{i^*} = 0.9)$ and, similarly, leader connectivity as low or high. Low connectivity gives leaders the same number of opportunities to transmit information as followers $(M^* = 3)$ whereas high connectivity doubles their opportunities to influence followers $(M^* = 6)$.

One generation of a particular run is defined as complete and the true state of the world, τ_w, is revealed when every follower (leader) has had an opportunity to initiate $M(M^*)$, pairings.[7] Let N_τ be the greatest number of followers within a predefined interval who update to a similar belief. We define a threshold δ such that when $N_\tau \geq \delta$, then the veracity of the rumor fails to matter. As the likelihood of collective action increases, the veracity of a rumor (credibility) is no longer relevant; disbelief in the rumor can be effectively suspended while the political, symbolic, and ritualistic value of participation with the group assumes greater salience. In this case, members of the group who acted in unison receive a return R regardless of the true state of the world, whereas dissenters pay a cost, C. Alternatively, when the group fails to act in unison $(N_\tau < \delta)$, then the true state of the world becomes important. In this

[7]For now, we specify the true state of the world, τ_w, stochastically. One extension of this model would entail stipulating the value of τ_w as endogenous.

case, when $\tau_i = \tau_w$, i receives a return, R, or otherwise fails and pays a cost, C.[8]

Next, a pair-wise selection of followers occurs on the basis of relative fitness (the net sum of returns R and costs C for each follower). In particular, we randomly draw a pair of followers, compare their fitness levels, and select the follower with higher fitness to play in the next generation of the model run. In keeping with the RCA model, each selected follower, or offspring in an evolutionary sense, is subject to potential mutation, which might change the offspring's belief, threshold, or both.[9] Each run of the model consists of 100 agents (inclusive of leaders and followers) who interact across 500 generations.

Model Results When There is a Single Ethnic Group

We seek to assess the probability that violence-promoting rumors emerge and survive. In modeling terms, an extreme rumor survives when a sufficient number of followers adopt and retain it over time. This requires, specifically, that the mean belief in a group exceeds 0.8 and remains there for at least 150 generations. To obtain estimates of the percent of all possible rumors that emerge and survive, we replicated each experimental condition, as defined by a combination of selected parameters, 500 times.

We first assume a single ethnic group with no official leaders and vary two member characteristics: their initial average belief and the cohesion of group members' beliefs around that average. We then add group leaders and vary two characteristics that presumably affect their capacities to propagate rumors: the extremism of their messages and their connectedness to group followers. Finally, we add a second group and use the model to predict how various types and levels of group interaction affect rumor emergence and survival, both within and across the two groups.

Rumor Survival in a Single Ethnic Group without Leaders

To provide a basis of comparison throughout our analysis, we constructed two baseline models. Baseline A includes parameters representing the hypothesized pair of least-rumor-optimizing follower characteristics: a moderate average group belief $(\bar{x}_\tau = 0.2)$ with little cohesion around it $(\sigma_\tau = 0.1)$.[10] Baseline B includes parameters representing the hypothesized pair of most-optimizing follower characteristics: an extreme average group belief $(\bar{x}_\tau = 0.8)$ with much cohesion around it $(\sigma_\tau = 0.025)$. Table 2 reports the results of all the single-group experiments.

Rumor survival is almost twice as great in baseline B (23.5%) than in baseline A (12.8%). This first comparison suggests that violence-promoting rumors are more likely to succeed when followers initially hold extreme beliefs. But it also leaves much unanswered.

Experiments 1–4 show that the mean initial belief among followers plays a bigger role than the cohesion of members' beliefs around it. A comparison of experiments 2 and 4 with baseline B underlines this point. The predicted percentage of rumors in experiment 2 (22.9%) is almost identical to that in baseline B, even though experiment 2 assumes low-belief cohesion. In contrast, even though experiment 4 and baseline B both assume high cohesion, experiment 4, with its additional assumption of a moderate initial average belief, predicts fewer successful rumors (15.9%). When the average initial belief is extreme, most followers need not hold beliefs similar to it. This situation can exist, of course, only when those who hold extreme beliefs hold the most extreme beliefs possible.

Rumor Survival in a Single Ethnic Group with Leaders

Now we introduce ethnic group leaders into the model. We hypothesized that leaders most successfully propagate rumors when they state extreme beliefs and

[8]We assume that ethnic groups face collective action problems in retaining the commitment of extremist members during periods of interethnic cooperation and in retaining the commitment of moderate members during periods of interethnic violence. As a result, group leaders utilize a mix of reward (R) and punishment (C) to elicit compliance with their objectives. Rewards might be monetary (land, jobs, housing, or official positions) or symbolic (saving face or elevating one's status or honor in the group). Punishments vary from verbal admonitions, the loss of status, and social ostracism to threats and reprisals against individuals (and family members) ranging in severity from beatings to death. Rewards are arguably more effective in promoting limited compliance and punishment more effective in promoting universal compliance with expected behavior.

[9]The characteristics of leaders remain unchanged across generations, making leaders immune to fitness-based selection.

[10]As we noted earlier when stating hypotheses, experiment 4 in Table 2 represents an equally plausible set of least-optimizing follower conditions. The results indicate that the two experiments produce highly similar results. We chose one as the baseline simply to facilitate discussion.

TABLE 2 Model Results, Single Ethnic Group

	Leader Beliefs τ_{i^*}	Leader Connectivity M^*	Mean Group Beliefs $\bar{x}_\tau$	Group Cohesiveness σ_τ	% of Rumors that Survive
Rumor Survival in a Single Ethnic Group without Leaders					
Baseline A	–	–	0.2	0.1	12.8
Exp 1	–	–	**0.5**	0.1	15.1
Exp 2	–	–	**0.8**	0.1	22.9
Exp 3	–	–	0.2	**0.0625**	14.5
Exp 4	–	–	0.2	**0.025**	15.9
Baseline B	–	–	0.8	0.025	23.5
Rumor Survival in a Single Ethnic Group with Leaders					
Exp 5	0.1	3	0.2	0.1	1.6
Exp 6	**0.9**	3	0.2	0.1	21.7
Exp 7	0.1	**6**	0.2	0.1	0.2
Exp 8	**0.9**	6	0.2	0.1	40.5
Exp 9	0.1	3	0.8	**0.025**	8.6
Exp 10	**0.9**	3	0.8	0.025	20.3
Exp 11	0.1	**6**	0.8	0.025	3.8
Exp 12	**0.9**	6	0.8	0.025	51.8

Note: For all of the model runs examined in the paper, and unless specified, $N = 100$, $\delta = 10$, $M = 3$, $C = 0.1$, $R = 1$, $\mu_\tau = \mu_T = 0.1$, and $\tau_w = rand$.

are well connected to their followers, and least successfully propagate them when they state moderate beliefs and are not well connected to their followers. Our results support both hypotheses.

When rumor-optimizing leader characteristics exist, even in the presence of least optimizing follower characteristics (experiment 8), the percentage of survived rumors (40.5%) is almost twice as great as it is in baseline B, underlining the power of the right combination of leader characteristics. When only one of the two leader characteristics—espousal of extreme beliefs (experiments 6 and 10) *or* high connectivity among followers (experiments 7 and 11)—exists, the rate of rumor emergence and survival never exceeds the baseline B rate. In both of these instances, in other words, group leaders do not add to the rate of rumor propagation that occurs among only followers, most of whom initially hold extreme beliefs.

Even in the presence of group leaders, furthermore, the nature of followers' beliefs affects the rate of rumor emergence and survival. Suppose leaders espouse moderate beliefs and are only weakly connected to their followers. Then, as a comparison of experiments 5 (1.6%) and 9 (8.6%) with experiments 7 (0.2%) and 11 (3.8%) shows, an increase in follower extremism and cohesion independently generates higher levels of rumor survival. In addition, when leaders state extreme beliefs and are highly connected to their followers, an increase in follower extremism and cohesion still increases the rate of rumor emergence and survival (compare experiment 8 [40.5%] with experiment 12 [51.8%]). Indeed, the results in experiment 12 show that a combination of rumor-optimizing follower and leader characteristics produces an especially high rate of rumor propagation.

Using experiment 12 as an example, Figure 2 illustrates that rumor dynamics do not follow a single pattern. At least three distinct trajectories occur: sometimes a rumor never quite emerges during a particular run (plot a); at other times, a dominant rumor rapidly emerges and survives (plot b); and at still other times, a dominant rumor emerges and survives, but only near the end of the run (plot c). Because each experimental condition is run 500 times, the model produces distributions of the various patterns. These distributions show that rumor survival occurs in more than half the replications; the situation where no rumor survives occurs far less often.

So what, in sum, do the ABM results in Table 2 tell us? First, the survival of violence-promoting rumors does not require group leaders. Rumor propagation occurs in their absence, especially when most followers already hold beliefs compatible with the rumors. Second, group leaders can increase the rate of rumor survival above and beyond what their

FIGURE 2 Three Independent Runs of Experiment 12

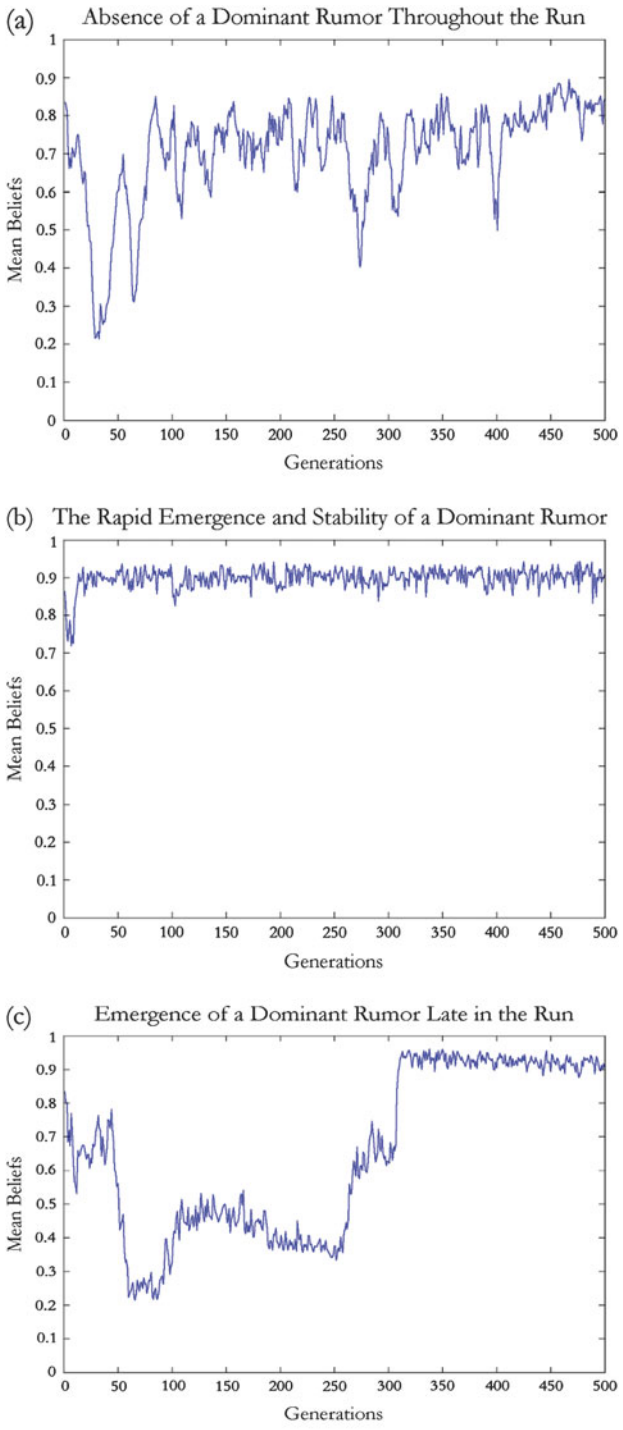

Note: All parameter settings based on experiment 12. Each plot displays mean group beliefs over time.

followers alone produce; but they must take extreme positions and, ideally, be highly interconnected with their followers. When leaders espouse moderate beliefs, they do not succeed in promoting rumors, whether or not they interact frequently with their followers. Finally, and perhaps most significantly, the greatest percentage of violence-promoting rumors emerges and survives when highly connected ethnic group leaders push extreme portrayals *and* most of their followers already hold beliefs compatible with those portrayals. This statement probably describes and explains many real-world situations.

Model Results When There Are Rival Ethnic Groups

We began simply, with a single, isolated ethnic group, to determine how follower and leader characteristics singly and jointly shape rumor dynamics within the group. Members were allowed to interact freely (and randomly) with each other. In most real-world situations, two distinct, usually rival groups exist. Followers and leaders within either group interact freely, but across-group interaction and communication are limited. How, if at all, does the presence of a rival group, and thus a change in the overall social structure of the ethnic population, affect the rumor dynamics we described above?

The introduction of a second, rival ethnic group requires additional nomenclature. Agents now carry an identity, $\gamma_i \in \{a, b\}$, to denote membership in group a or b, respectively.[11] The model also incorporates parameters to reflect the level of interaction and communication among members of the two rival groups. We designate $p^{a,b}$ as the probability that a member of a will initiate an interaction with a member of b, and $p^{b,a}$ as the reverse. Most of the experiments that follow come in pairs, with the model first assuming no interaction between the rival groups' members $\left(p^{a,b} = p^{b,a} = 0\right)$ and then assuming limited interaction, such that members of each group are twice as likely to initiate an interaction with members of their own group as they are to initiate one with members of the other $\left(p^{a,b} = p^{b,a} = 0.5\right)$.

Our research strategy uses and builds on the preceding analyses. We first assume the presence of most-optimizing characteristics in both groups, as represented by the parameters in experiment 12: highly connected leaders propagating extreme statements to followers, most of whom are already inclined to believe the extreme statements. We then systematically change, first, the follower and, second, the leader

[11]Note that agent identities remain unchanged across generations to ensure that fitness-based selection occurs within equally sized groups.

characteristics of one group from most to least optimizing while maintaining the initial characteristics of the other. Table 3 reports the results of all the two-group experiments.

Most Optimizing Characteristics in Both Groups

By assigning parameters identical to those used in experiment 12, experiments 13 and 14 assume rumor-optimizing follower and leader characteristics in both groups. Experiment 13 also assumes the two groups to be completely segregated, while experiment 14 allows for partial across-group integration.

Note, first of all, and as an aside, that the percentage of surviving rumors drops markedly when group size goes from 100 to 50 members. Whereas more than 50% of all rumors survived under optimizing conditions when the group consisted of 100 members, only about a quarter survive under the same conditions when the group consists of 50 members. Although not the focus of this paper, ethnic group size warrants consideration in future studies of the dynamics of violence-promoting rumors.

Second, the rates of rumor propagation in experiments 13 (21.5% and 22.0%) and 14 (26.5% and 26.0%) differ only slightly. When rival ethnic groups begin with the same, rumor-optimizing follower and leader characteristics, it appears not to matter whether members of the two groups are completely segregated or interact and communicate on a limited basis. To the extent that a difference exists, relatively fewer rumors emerge when members of the two groups do not interact.

That rumor propagation declines, or at least does not increase, in the presence of complete segregation contrasts with conventional wisdom, which predicts that segregation of rival ethnic groups will promote especially high rumor emergence and survival in each group. This wisdom, however, consists largely of casual observation (Allport 1954). On the other hand, and in fairness, our epidemiology-like model (Wasserheit and Aral 1996) does not allow one group's leaders to dampen rumor propagation among followers of the other. We expect this to be a rare real-world occurrence, but other researchers might wish to incorporate the mechanism into their models.

Asymmetric Follower Characteristics and Rumor Survival

We next change the follower characteristics in group *b* from most to least rumor optimizing, while maintaining the most optimizing follower characteristics in group *a*, again under both complete and partial segregation conditions. Group leaders' characteristics in both groups remain most optimizing. The results are shown in experiments 15 and 16.

Follower characteristics only slightly affect the emergence and survival of rumors, which can be seen when comparing the percentages of emergent rumors in experiment 15 (22.7% and 21.1%) with those in experiment 13 (21.5% and 22.0%) and the percentages in experiment 16 (25.1% and 24.4%)

TABLE 3 Model Results, Rival Ethnic Groups

	Leader Beliefs $\tau^a_{i^*}, \tau^b_{j^*}$	Leader Connectivity M^*	Mean Follower Beliefs $\bar{x}^a_\tau, \bar{x}^b_\tau$	Mean Cohesiveness $\sigma^a_\tau, \sigma^b_\tau$	p(*a* interacts w/ *b*) $p^{a,b}, p^{b,a}$	% of Rumors that Survive *a*, *b*
Rumor Survival and the Level of Group Segregation						
Exp 13	0.9, 0.9	6	0.8, 0.8	0.025, 0.025	**0, 0**	21.5, 22.0
Exp 14	0.9, 0.9	6	0.8, 0.8	0.025, 0.025	**0.5, 0.5**	26.5, 26.0
Rumor Survival and Asymmetric Follower Characteristics						
Exp 15	0.9, 0.9	6	0.8, **0.2**	0.025, **0.1**	0, 0	22.7, 21.1
Exp 16	0.9, 0.9	6	0.8, **0.2**	0.025, **0.1**	0.5, 0.5	25.1, 24.4
Rumor Survival and Asymmetric Leader Characteristics						
Exp 17	0.9, **0.1**	6	0.8, 0.2	0.025, 0.1	0, 0	21.2, 0.1
Exp 18	0.9, **0.1**	6	0.8, 0.2	0.025, 0.1	0.5, 0.5	7.0, 0.6
Rumor Survival, One-Sided Information Transmission, and Discounting						
Exp 19	0.9, 0.1	6	0.8, 0.2	0.025, 0.1	**0.5, 0**	15.8, 0.8
Exp 20	0.9, 0.1	6	0.8, 0.2	0.025, 0.1	**0, 0.5**	5.6, 0.2
Exp 21	0.9, 0.1	6	0.8, 0.2	0.025, 0.1	**0.5, 0 (d)**	15.4, 0.3

Note: In the rival group experiments, we leave the number of agents N unchanged, dividing the population into two equally sized groups. In experiment 21, "(d)" indicates that members of group *b* discount information transmitted from members of group *a*.

with those in experiment 14 (26.5% and 26.0%). Moreover, although the experiments now assume the least optimizing combination of follower characteristics in one group and the most optimizing in the other, going from complete to partial segregation does not change the results. Moderate interaction $(p^{a,b} = p^{b,a} = 0.5)$ allows group *a* members to transmit more extreme beliefs to group *b* members, but, at the same time, group *b* members are transmitting moderate beliefs to group *a* members. These countervailing dynamics foster relatively equal percentages of dominant rumors in both groups. Finally, these results demonstrate, once again, that leaders who state extreme beliefs and are highly connected to their followers successfully propagate violence-promoting rumors, quite independently of follower characteristics or level of segregation between the two groups. The percentages of rumors that survive in experiments 15 and 16 closely resemble those in experiments 13 and 14, where both follower and leader characteristics are most optimizing.

Asymmetric Follower and Leader Characteristics and Rumor Survival

The experiments in table 2, which assume one group, show that group leaders most successfully propagate violence-promoting rumors when they take extreme positions and are highly connected to their followers. Conversely, when either of these conditions is missing, the emergence and survival of rumors drops, with the bigger drop occurring when leaders espouse moderate rather than extreme beliefs. So what happens, in the case of rival ethnic groups, when one group's leaders possess rumor-optimizing characteristics and the other's not? Recall that, from the preceding experiment, group *a* followers now hold the most, and group *b* followers the least, optimizing characteristics.

Take the completely segregated position first (experiment 17), which requires little comment. More than one in every five possible rumors emerges and survives in group *a*; almost none does in group *b* (21.2% and 0.1%). In other words, rumor propagation in group *a* looks much as it has looked in all of the experiments in Table 3, but, for the first time, rumor propagation in group *b*, whose leaders possess the least optimizing characteristics, comes close to nonexistent. These strong results underline yet again, and perhaps more dramatically than before, the crucial importance of the two leader characteristics we chose to study.

Experiment 18 represents an interesting case in that it allows followers and leaders from one group to interact with those from the other. Again, intragroup interaction occurs at twice the rate of intergroup interaction. Nevertheless, this experiment, unlike experiment 17, affords group *b* members with their largely moderate beliefs an opportunity to counter group *a* leaders' influence on their own followers. The question is, does the percentage of violence-promoting rumors that emerge and survive in group *a* drop?.

The change from complete to partial segregation greatly decreases the percentage of emergent rumors in group *a*, and very slightly increases it in group *b* (7.0% and 0.6%). The significance of this result warrants emphasis: *this is the first time that the presence of optimizing leader characteristics does not ensure relatively high percentages of rumors*. As long as one group's leaders espouse moderate beliefs and there is some communication across groups, the overall percentage of rumors falls precipitously. But now the caveats: first, even two or three rumors can spark intergroup violence in the real world; and, second, the number of real-world situations that meet the parameters specified in experiment 18 might be small. Nevertheless, the experiments offer a basis for reasoned speculation about the nature of rumor emergence and survival when one group's leaders espouse moderate beliefs while the other's leaders espouse extreme ones.

One-Sided Information Transmission, Discounting, and Rumor Survival

The last three experiments in Table 3 go a step further and assume that belief transmission occurs in one direction: from followers and leaders of one group to those of the rival group, but not vice versa. The final of these three experiments allows followers of the receiving group either to dismiss the rumor-based information they receive or to counter it by changing their beliefs in the opposite direction.

Specifically, in experiment 19 we explore how one-sided information transmission from group *a* members to group *b* members $\left(p^{a,b} = 0.5, p^{b,a} = 0\right)$ affects rumor emergence and survival, whereas we explore the reverse condition in experiment 20 $\left(p^{a,b} = 0, p^{b,a} = 0.5\right)$. Conducting both experiments will reveal whether belief transmission from an ethnic group with the rumor-optimizing characteristics to one without differs in its consequences from when transmission goes in the opposite direction. In experiment 21, group *b* members either dismiss the information they receive from group *a* members or

change their beliefs in the opposite direction. When two agents from rival groups hold relatively similar beliefs, then the group *b* member simply dismisses the transmitted belief.[12] If the two nominal rivals differ significantly in their beliefs, however, then belief transmission from group member *a* to group member *b* causes the latter to change her belief in the opposite direction, i.e., to polarize.

Comparing experiment 19 (15.8% and 0.8%) to experiment 18 (7.0% and 0.6%) shows that when group *a* members transmit rumors to those in group *b*, but not vice versa, the percentage of emergent and surviving rumors in group *b* remains largely nonexistent, going from 0.6 to 0.8. Interestingly, the percentage of dominant rumors in group *a* decreases, presumably because group *a* members are now spreading their mostly extreme beliefs among members of both groups, and thus not concentrating on rumor propagation within their own group.

Experiment 20 explores the opposite case, where group *b* members transmit their mostly moderate beliefs to group *a* members but group *a* members do not transmit their mostly extreme beliefs to group *b* members. The results resemble those in experiment 19, albeit with reversed results: there is little change in rumor survival among group *a* members, but a decline among the transmitting group *b* members (5.6% and 0.2%).

Finally, suppose that members of group *a* transmit rumors to group *b* (and not vice versa) but agents in *b* discount the information they receive (experiment 21). Depending on how far the transmitted rumor is from their existing beliefs, they either dismiss the information out of hand or change their beliefs in a direction opposite from the rumor. Then the percentage of rumors that survive, in both groups, is relatively small. In other words, when members of one group who hold mostly moderate beliefs reject extreme rumors at face value, they keep the total number of violence-promoting rumors to a minimum. The following experiment probably better approximates real-world dynamics, however.

[12]In particular, an agent *i* from group *b* will dismiss information passed on by rival *j* from group *a* if $|\tau^b_{i,t} - \tau^a_{j,t}| \leq T_i$ or update her belief based on this information if $|\tau^b_{i,t} - \tau^a_{j,t}| > T_i$, such that if $q' = |\tau^b_{i,t} - \tau^a_{j,t}|$, then:

$$\tau^b_{i,t+1} = \tau^b_{i,t} + (1 - \tau^b_{i,t})(q') \text{if} \tau^b_{i,t} > \tau^a_{j,t}$$

and

$$\tau^b_{i,t+1} = \tau^b_{i,t} - (\tau^b_{i,t} q') \text{if} \tau^b_{i,t} < \tau^a_{j,t}.$$

Note that if $\tau^a_j = \tau^b_i$, then $|\tau^b_i - \tau^a_j| \leq T_i$, and any transmitted information is dismissed.

Rumor Survival in the Presence Outbidding of Rival Leaders

In this final experiment, leaders progressively outbid each other to become more extreme. Parameter settings are identical to those in experiment 13 (Table 3), with one exception: leaders in both groups initially hold moderate beliefs. In each generation, with probability < 0.1, a leader $i^\star$ from group *a* is paired with a leader $j^\star$ from group *b*. The rival leaders compare their respective belief levels and the less extreme agent changes her belief so as to become more extreme than the other. Specifically, leaders update by adopting the extremity of the rival leader and adding to it the absolute value of difference in the two leaders' beliefs at time *t*.[13]

Figure 3 displays the results of a single run of the outbidding experiment. The solid, darker lines represent changes in the mean beliefs of the two groups' leaders. Initially set at approximately 0.5, average beliefs slowly move upwards over the course of 200–300 generations as the two groups' leaders interact. Changes become more acute over time, with leaders of both groups rapidly adopting extreme beliefs toward the end of the run.[14] The thinner, erratic lines, which trace the evolution of group followers' mean beliefs, show how leaders' capacities to propagate rumors among followers simultaneously increase over time. In response to increasingly extreme leader beliefs, followers' beliefs trend mostly upwards. Now and then, followers' beliefs become more extreme than their leaders'.[15]

Discussion

From Detroit to Paris to many parts of Asia, the propagation of extreme rumors has served, by all available accounts, as a catalyst of ethnic violence. Whereas researchers have focused heavily on the

[13]As such,

$$\tau^a_{i^*,t+1} = \tau^b_{j^*,t} + |\tau^a_{i^*,t} - \tau^b_{j^*,t}| \text{ if} \tau^a_{i^*,t} < \tau^b_{j^*,t}$$

and

$$\tau^b_{j^*,t+1} = \tau^a_{i^*,t} + |\tau^a_{i^*,t} - \tau^b_{j^*,t}| \text{ if} \tau^b_{j^*,t} < \tau^a_{i^*,t}$$

Note also that if $\tau^a_{i^*,t} = \tau^b_{j^*,t}$, then neither updates. Leaders in group *a* actually begin at 0.51 whereas those in group *b* begin at 0.5.

[14]We restrict leader extremism from exceeding 0.9 to remain consistent with all other experiments.

[15]In an effort to examine the robustness of our results, we varied three global parameters that were held fixed in our experiments. The results of our sensitivity tests are summarized in the appendix, Table A1.

FIGURE 3 Aggregate Beliefs Under Ethnic Outbidding

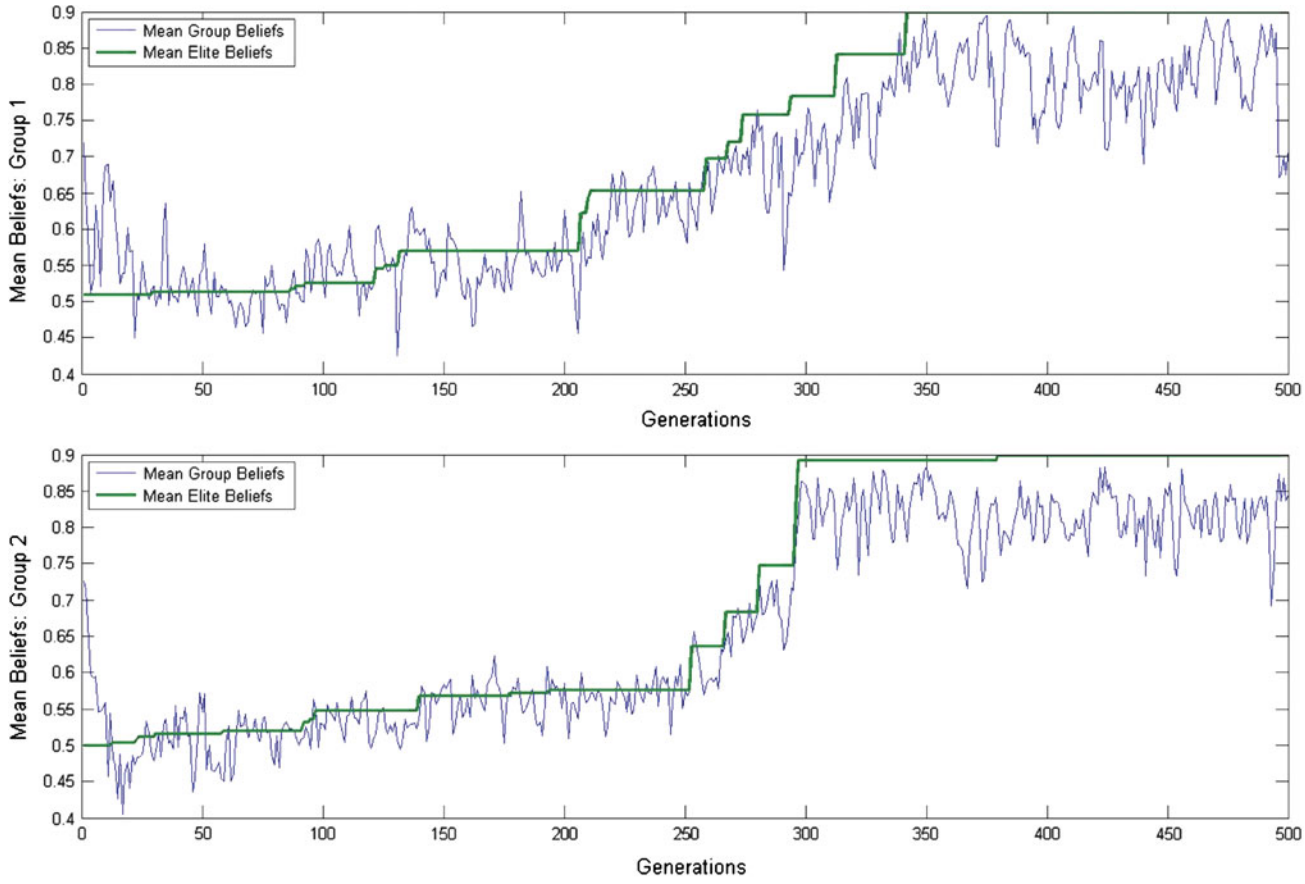

Note: The thinner, erratic lines in the figures plot mean group beliefs and the thicker lines plot mean leader beliefs for each of the two groups over time. Parameter values for this run of the model are as follows: $M^* = 6$; $\bar{x}^a_\tau, \bar{x}^b_\tau = 0.8$; $\sigma^a_\tau, \sigma^b_\tau = 0.025$; $p^{a,b}, p^{b,a} = 0$.

violence, we used ABM to understand the catalyst: do extreme rumors inevitably emerge and survive as a matter of course or, alternatively, does their propagation require certain combinations of group follower and leader characteristics? We began with a single group consisting only of rank-and-file members and varied two group characteristics: the average initial belief of the group as a whole and the clustering of individual members around that average belief. We then added group leaders and varied the two preceding group characteristics as well as two leader characteristics: the extremity of the beliefs they try to disseminate and their connectedness to group members. Finally, we added a rival ethnic group with its own members and leaders and varied their characteristics also, while simultaneously varying the level of segregation between the two groups.

The experiments reveal aspects of rumor dynamics that empirical observation has not shown. One conclusion pervades them all: extreme rumors are not inevitable. Under various suboptimal conditions, they simply never surface. The one-group experiments show that group leaders are not necessary for rumor propagation. Extreme rumors emerge and survive even among leaderless groups, as long as the average initial belief among the members is already extreme. The key term is average; individual members' beliefs need not be closely clustered around this average for rumors to appear.

These same one-group simulations show that group leaders can play a key role in rumor propagation *if* they espouse extreme beliefs and interact frequently with their followers. Of these two factors, the first, disseminating extreme beliefs, is the more essential. The crucial finding, however, is this: group leaders who attempt to propagate extreme statements for purposes of inciting violence fare much better when their followers already hold beliefs that are compatible with the statements than when they do not. In this case, extreme rumors abound. Far fewer

rumors emerge and survive when followers are not initially inclined to accept what their leaders are saying. Our simulations thus support Horowitz's thesis, noted earlier, that *group followers constrain what leaders can do*.

Adding a rival group to the analyses does not change the preceding conclusions. It does produce additional insights. An interesting situation arises when the leaders of one ethnic group advocate extreme, violent behavior while those of the rival group persist in advocating moderation. If, in addition, the rank-and-file members of the two groups are allowed to interact, the rate of rumor propagation within both groups is relatively low. Indeed, this is the one instance where, despite the presence of rumor-optimizing leader characteristics—espousal of extreme beliefs and full integration into the group—the leaders fail to propagate extreme rumors. Allowing rival leaders to try to outbid each other, however, produces a very different end result: a high rate of rumor emergence and survival in both groups. Unfortunately, this latter situation probably describes many real-world situations.

Nevertheless, these findings offer two possible openings by which practitioners might try to stem the current wave of interethnic violence. First, create organizations that not only benefit members of both ethnic groups, but also require the active participation of members from both. For example, micro-aid programs that loan money directly to impoverished citizens have become increasingly prevalent. In countries experiencing interethnic violence, requiring rank-and-file members of the two groups to work together before extending a loan to either could prove especially effective. Presumably these interactions would begin to moderate the beliefs of all participants. Second, identify group leaders who are willing to take moderate positions and provide them incentives to maintain those positions rather than engage in a *quid pro quo* exchange with rival leaders. The incentives could take many forms, from providing monetary assistance to the group to bolstering the status of moderate leaders.

Keep the limitations of our model in mind, however. The experiments explored only part of the model's parameter space. Future analysis could use even finer-grained combinations of the same follower and leader characteristics. In the same vein, researchers could include additional characteristics. The model only distinguishes leaders from followers. This distinction works well with respect to most racial and ethnic groups. However, it is less appropriate to understanding paramilitary organizations such as terrorist groups. The construction of models appropriate to studying such organizations must await another day.

Another useful extension would be to specify interaction topologies—social networks that either assume a particular structural form or that evolve endogenously. Specification of a star network could be used to capture the effects of media reports on rumor survival, whereas a small-world network could capture the phenomenon of strangers being linked by a mutual acquaintance (both of which are discussed in the appendix). We did not distinguish between leader insinuation and appropriation of rumors on the one hand, and decentralized rumor propagation on the other. The ABM could require leaders to assess the costs and benefits of lending credence to existing rumors vis-à-vis those of insinuating new rumors that carry the risk of being traced back to them. Researchers could also use gaming experiments with real subjects to test model predictions about rumor survival. Known as "docking" exercises, human subjects would "play" games similar to those conducted with artificial agents in the model, making it possible to compare or "dock" the ABM and game (Axtell et al. 1996; Duffy 2006).

Given all of these possibilities, one might ask whether our findings are too limited to be meaningful. To that thought, we offer an emphatic "no." Successful scientific endeavors begin modestly and expand in scope and refinement as more and more scholars contribute to them. A more heroic and far less transparent first effort on our part would not have served us or the larger political science community.

Manuscript submitted 18 July 2007
Manuscript accepted for publication 16 November 2008

References

ABC. 2005. "Paris Riots in Perspective: Suburb Realities Highlight the City's Less Glamorous Side." ABC News. http://abcnews.go.com/international/story?id=1280843 (November 30, 2008).

Allport, Gordon. 1954. *The Nature of Prejudice*. Reading, MA: Addison-Wesley.

Allport, Gordon, and Leo Postman. 1947. "An Analysis of Rumor." *Public Opinion Quarterly* 10 (4): 501–17.

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

Axelrod, Robert. 2006. "Agent-Based Modeling as a Bridge between Disciplines." In *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics*, eds. Leigh Tesfatsion and Kenneth Judd. New York: North-Holland/Elsevier, 1565–84.

Axtell, Robert, Robert Axelrod, Joshua Epstein, and Michael Cohen. 1996. "Aligning Simulation Models: A Case Study and Results." *Computational and Mathematical Organization Theory* 1 (2): 123–41.

Banerjee, Abhijit. 1992. "A Simple Model of Herd Behavior." *The Quarterly Journal of Economics* 107 (3): 797–817.

Banerjee, Abhijit. 1993. "The Economics of Rumors." *The Review of Economic Studies* 60 (2): 309–27.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *The Journal of Political Economy* 100 (5): 992–1026.

Brass, Paul. 1997. *Theft of an Idol*. Princeton, NJ: Princeton University Press.

Brown, Earl. 1943. "The Truth about the Detroit Riot." *Harper's Magazine*, November, 488–98.

Capeci, Dominic, and Martha Wilkerson. 1991. *Layered Violence: The Detroit Rioters of 1943*. Jackson: University Press of Mississippi.

Centola, Damon, and Michael Macy. 2007. "Complex Contagions and the Weakness of Long Ties." *American Journal of Sociology* 113 (3): 702–34.

Diamond, Jared. 2004. *Collapse: How Societies Choose to Fail or Succeed*. New York: Penguin.

Duffy, John. 2006. "Agent-Based Models and Human-Subject Experiments." In *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics*, eds. Leigh Tesfatsion and Kenneth Judd. New York: North-Holland/Elsevier, 949–1011.

Engineer, Asghar. 1995. *Lifting the Veil: Communal Violence and Communal Harmony in Contemporary India*. Hyderabad: Sangam Books.

Epstein, Joshua. 2002. "Modeling Civil Violence: An Agent-based Computational Approach." *Proceedings of the National Academy of the Sciences* 99 (3): 7243–50.

Fearon, James, and David Laitin. 1996. "Explaining Interethnic Cooperation." *American Political Science Review* 90 (4): 715–35.

Gagnon, V. P. 1995. "Ethnic Nationalism and International Conflict: The Case of Serbia." *International Security* 19 (3): 130–66.

Gould, R. 1999. "Collective Violence and Group Solidarity: Evidence from a Feuding Society." *American Sociological Review* 64 (3): 356–80.

Gourevitch, Philip. 1998. *We Wish to Inform You that Tomorrow We Will Be Killed With Our Families: Stories From Rwanda*. New York: Picador.

Horowitz, Donald. 2001. *The Deadly Ethnic Riot*. Berkeley: University of California Press.

Horowitz, Donald, and Ashutosh Varshney. 2003. "Lethal Ethnic Riots: Lessons from India and Beyond." *United States Institute of Peace Special Report 101*.

Kaufman, Stuart. 2001. *Modern Hatreds: The Symbolic Politics of Ethnic War*. Ithaca, NY: Cornell University Press.

Knapp, Robert. 1944. "A Psychology of Rumor." *Public Opinion Quarterly* 8 (1): 22–37.

Kosfeld, Michael. 2005. "Rumours and Markets." *Journal of Mathematical Economics* 41: 646–64.

Lake, David, and Donald Rothchild. 1996. "Containing Fear: The Origins and Management of Ethnic Conflict." *International Security* 21 (2): 41–75.

Langlois, Janet. 1983. "The Belle Isle Bridge Incident: Legend Dialectic and Semiotic System in the 1943 Detroit Race Riots." *Journal of American Folklore* 96 (380): 183–99.

Mamdani, Mahmood. 1996. "From Conquest to Consent as the Basis of State Formation: Reflections on Rwanda." *New Left Review* 216: 3–36.

Miller, John, and Scott Page. 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton, NJ: Princeton University Press.

Perice, Glen. 1997. "Rumors and Politics in Haiti." *Anthropological Quarterly* 70 (1): 1–10.

Peterson, Roger. 2002. *Understanding Ethnic Violence: Fear, Hatred, and Resentment in Twentieth Century Eastern Europe*. Cambridge: Cambridge University Press.

Pound, John, and Richard Zeckhauser. 1990. "Clearly Heard on the Street: The Effect of Takeover Rumors on Stock Prices." *Journal of Business* 63 (3): 291–308.

Riolo, Rick, Michael Cohen, and Robert Axelrod. 2001. "Evolution of Cooperation without Reciprocity." *Nature* 414 (6862): 441–43.

Rose, Arnold. 1951. "Rumor in the Stock Market." *Public Opinion Quarterly* 15 (3): 461–86.

Shibutani, Tamotsu. 1966. *Improvised News: A Sociological Study of Rumor*. Indianapolis, IN: Bobbs-Merrill.

Smelser, Neal. 1962. *Theory of Collective Behavior*. New York: Free Press.

Sugrue, Thomas. 1996. *The Origins of the Urban Crisis: Race and Inequality in Post-War Detroit*. Princeton, NJ: Princeton University Press.

Tambiah, Stanley. 1990. "Presidential Address: Reflections on Communal Violence in South Asia." *The Journal of Asian Studies* 49 (4): 741–60.

Tambiah, Stanley. 1996. *Leveling Crowds: Ethnonationalist Conflicts and Collective Violence in South Asia*. Berkeley: University of California Press.

Wasserheit, J. N., and S. O. Aral. 1996. "The Dynamic Topology of Sexually Transmitted Disease Epidemics: Implications for Prevention Strategies." *Journal of Infectious Diseases* 174 (2): 201–13.

Watts, Duncan. 1999. *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton, NJ: Princeton University Press.

Wood, Elizabeth. 2003. *Insurgent Collective Action and Civil War in El Salvador*. Cambridge, MA: Cambridge University Press.

Zivney, Terry, William Bertin, and Khalil Torabzadeh. 1996. "Overreaction to Takeover Speculation." *Quarterly Review of Economics and Finance* 36 (1): 89–115.

Ravi Bhavnani is assistant professor of political science, Michigan State University, East Lansing, MI 48824.

Michael G. Findley is assistant professor of political science, Brigham Young University, Provo, UT 84602.

James H. Kuklinski is professor of political science, University of Illinois at Urbana-Champaign, Urbana, Il 61801.