


EMMA S. SPIRO AND KATE STARBIRD



Rumors Have Rules

Decades-old research about how and why people share rumors is even more relevant in a world with social media.

Around the time the US government was testing nuclear bombs on Bikini Atoll in the spring of 1954, residents of Bellingham, Washington, inspected their windshields and noticed holes, pits, and other damage. Some blamed vandals, perhaps teenagers with BB guns. Once Bellingham residents reported the pits, people in nearby towns inspected their windshields and found similar damage. Could sand fleas have caused it? Or cosmic rays? As more people examined their windshields and found more pits, a frightening hypothesis emerged: nuclear fallout from government hydrogen bomb testing. Within a week, people around Seattle, 90 miles away, were reporting damage as well.

But the rumors faded almost faster than they began as scientists and local authorities refuted the most prominent theories. What would become known as the “Seattle windshield pitting epidemic” became a textbook example of how rumors propagate: a sort of contagion, spread through social networks, shifting how people perceive patterns and interpret anomalies. Car owners saw dings that they’d previously overlooked and shared observations with others, who repeated the process.

Today this phenomenon would probably be described using the terms *misinformation*, *disinformation*, and perhaps *fake news*. Certainly, communication has changed dramatically since party-line telephone calls and black-and-white television, but scholarship from that era holds critical insights that are essential to the digital era. The study of rumors, which surged around World War II, is still very relevant.

Our team of researchers at the University of Washington have been investigating these issues for more than a decade. Initially, we centered our research on rumors. But we shifted focus in 2016 as public and academic attention lasered onto misinformation and social media manipulation. By 2020, our collaborators in an effort called the Election Integrity Partnership had developed an analytical framework that allowed dozens of students to scour social media platforms in parallel, feeding information to trained researchers for analysis and to authorities and communicators for potential

mitigation. As we worked to build ways to quickly prioritize unsubstantiated claims about election processes and results, we found that the terms *misinformation* and *disinformation* were often cumbersome, confusing, or even inaccurate. But we came full circle in 2022, during a second iteration of our collaboration on election integrity, because the concept of rumors worked easily and consistently to assess the potential for unsubstantiated claims to go viral online.

We are convinced that using rumor as a conceptual framework can enhance understanding of today’s information systems, improve official responses, and help rebuild public trust. In 2022, we created a prioritization tool around the concept of rumors. The idea was to help anticipate rumors that could undermine confidence in the voting process and assess whether a given rumor would go viral. Much of the concept’s utility is that responders can engage with an information cascade before veracity or intent can be determined. It also encourages empathy by acknowledging the agency of people spreading and responding to rumors, inviting serious consideration of how they can contribute to the conversation.

A brief history of rumors

In the scholarly literature, *misinformation* refers to content that is false but not necessarily intentionally so; *disinformation*, which has roots in Soviet propaganda strategies, refers to false or misleading content intentionally seeded or spread to deceive. These terms are useful for certain problematic content and behavior, but they are increasingly politicized and contested. Mislabeling content that turns out to be valid—or potentially valid, like the theory that COVID-19 began with a so-called lab leak—violates public trust, undermines authorities’ credibility, and thwarts progress on consequential issues like strengthening democracy or public health. In contrast, the label of *rumor* does not declare falsity or truth.

Rumoring can be especially valuable when official sources are incentivized to hide information—for example, when an energy company is withholding information about pollution,

or when a government agency is covering up incompetence. Branding expert Jean-Noël Kapferer posits that rumors are best understood not as leading away from truth, but as seeking it out.

Historically, researchers defined rumors as unverified stories or “propositions for belief” that spread from person to person through informal channels. Rumors often emerge during crises and stressful events as people come together to make sense of ambiguous, evolving information, especially when “official” information is delayed. In this light, the numerous rumors that spread in the early days of the pandemic about its origins, impacts, and potential antidotes are unsurprising.

Both the production and spread of rumors are often taken as a natural manifestation of collective behavior with productive informational, social, and psychological roles. For instance, rumors help humans cope with anxiety and uncertainty. A population coming to terms with the risks of nuclear weapons could find a way to voice fears by seeing dings in the windshields of their Ford Thunderbirds and Chevrolet Bel Airs. Recognizing these informational and emotional drivers of rumoring can support more empathetic—and perhaps more effective—interventions.

Recentering rumor

When our research team tracked the use of social media in the 2010 Deepwater Horizon oil spill and the 2013 Boston Marathon bombings, we grounded our work in scholarship of offline rumors. Both events catalyzed numerous rumors: conspiracy theories warning about an impending sea floor collapse, crowd-sleuthed identifications of (innocent) suspects. We uncovered similar patterns in dozens of subsequent events where rumors circulated on social media, including a WestJet plane hijacking (that did not happen), multiple mass shootings in the United States, the downing of Malaysian Airlines flight 17 over Ukraine, and terrorist attacks in Sydney, Australia, in 2014 and Paris, France, in 2015.

Our rumor threat framework draws on much of the foundational social science research in rumoring to create the analytic categories, labeled “Dimensions” in Figure 1. These include underlying conditions (such as uncertainty or trust in official information), features of the rumor (such as novelty and emotional valence), and system effects (such as position in a social network). Both *uncertainty* and *significance* are rooted in the “basic law of rumor” introduced by scholars Gordon W. Allport and Leo Postman in 1946: the strength of a rumor is proportional to its significance to the listener multiplied by the ambiguity of the evidence around it. The condition of *diminished trust* stems from an idea of sociologist Tomatsu Shibutani from 1966, that informal communication surges in the absence of timely official information. The *familiarity/repetition* dimension arises from the “illusory truth effect,” identified in the 1970s, that repetition increases believability. The seemingly contradictory feature of *novelty* tracks to work in 1990 showing that rumors lose value over time.

Assessing these dimensions helps predict which rumors will become viral. Take the 1950s Seattle windshield pitting epidemic: the underlying conditions included *high uncertainty* about both the cause of the windshield damage and the risks of nuclear fallout. The features of the rumor included *high significance* if the nuclear connection were true; *substantial novelty*, as both car ownership and concerns about nuclear weapons had become widespread in the years since World War II; *high emotional valence* pegged to nuclear fears as well as property damage; and *compelling evidence*, since people could see dings on their own cars and photographs of others. The system effects included *participatory potential* as people inspected and discussed their own cars. And, if contemporaneous accounts are true, the rumor declined because of high trust in the authorities who were debunking it.

Consider the COVID-19 pandemic. Every dimension of our framework shows rumor-promoting conditions. A novel virus with uncertain causes, serious consequences, unknown spreading mechanisms, and few remedies scores high in every category. Trust in government and local health officials started out low and declined from there. Impacts include lost lives, jobs, disrupted routines, and isolation—all of which heighten emotions. As for participation, many people all over the world were stuck in their homes and converged on familiar social platforms to share home remedies, hunts for toilet paper, stories of sick loved ones, reactions to lockdowns, and more, often with first-person accounts and video testimonials. Of course, some actors—both organized and disorganized—strategically manipulated information systems to gain attention, sell products, and push political agendas.

Ready for use

We initially developed this framework for research to guide our “rapid response” research. After conversations with local and state election officials who were struggling for guidance about when and how to address false claims about their processes and procedures, we adapted the framework for their perspective. Since then, we have presented it to a small number of local and state election officials for feedback. We aim to develop, deploy, and evaluate trainings based on the framework for 2024.

Though we developed techniques to classify rumors specifically for elections, we see potential for much wider application. Communicators can use the framework to assess vulnerability to rumors and to prepare for outbreaks or other foreseeable crises. Evaluating rumors across multiple dimensions for potential virality can be more useful than deciding whether to apply a label like misinformation or disinformation. It may not be worth drawing attention to a rumor likely to lapse, but it would be valuable to correct a harmful false rumor with high spread potential before it gets started. Such insights can inform how to focus crisis communication, platform moderation, or fact-checking resources.

Our framework may point to other actors and incidents

Figure 1: THE RUMOR THREAT FRAMEWORK HELPS ASSESS WHETHER AN EVENT WILL GENERATE VIRAL RUMORS AND WHICH RUMORS WILL SPREAD

GROUP	DIMENSIONS	DESCRIPTION
INFORMATION AND EVENT CONDITIONS	Uncertainty	As events (infections, train wrecks, elections) unfold, uncertainty powers both generation and spread of rumors. For specific rumors, ambiguous evidence will lead to more spread.
	Diminished Trust	Diminished trust in “official” information providers (government, media, etc.) pushes people toward more informal communication channels, catalyzing rumorizing.
CONTEXTUAL FEATURES	Significance / Impact	The strength of a specific rumor is proportional to its importance in the lives of those spreading it. Events with greater potential impact on people’s lives will catalyze more, and more viral, rumors.
	Familiarity / Repetition	A common set of building blocks underlie many rumors, which may make them resonate with familiarity. This, plus repetition, can enhance plausibility and boost spread.
	Novelty	People spread rumors to inform and entertain. Crises and other emergent events provide novel content that can be assembled into rumors.
	Emotional Valence	Rumors that stimulate strong emotions—including anger, fear, and outrage—will be more likely to spread. Events that stimulate strong emotions may catalyze the creation of viral rumors.
	Compelling-ness of Evidence	Evidence that piques interest and adds tangibility—e.g., first-person accounts, photos, and videos—will catalyze the creation and enhance virality of rumors.
	Participatory Potential	Rumors that allow people to participate—to add evidence or share their experiences and interpretations—are likely to spread further.
SYSTEMS EFFECTS	Position within the Social Network	Social networks shape the spread of rumors. Rumors will spread further if they reach central or high-audience nodes in a network or move from one network to another (e.g., a rumor in an anti-vaccine network jumps into ethnonationalist networks).
	Algorithmic or Network Manipulation	In online environments, rumors can be intentionally seeded or spread for strategic gain. Often those efforts game underlying networks and recommendation systems.

that require further consideration. For instance, online communities that actively engage in conspiracy theorizing are poised to project a common set of ideas onto events’ causes and impacts. And sometimes the sorts of rumors an event will spawn are predictable. For those responding to the toxic train wreck in East Palestine, Ohio, for example, public discourse around oil spills and chemical accidents could reveal what likely rumors and conspiracy theories might appear.

Certainly, this framework cannot capture everything worthy of consideration. Practitioners should examine potential virality in tandem with the potential for harm. There are probably also productive approaches for refocusing the lens around misinformation and disinformation. For example, the term *misinformation* likely remains a better fit for describing false claims spread through low-quality scientific journals, and *propaganda* might better capture concerted efforts to manipulate the masses. *Disinformation* is useful for intentionally misleading and clearly manipulative campaigns.

However, we suspect the power of the rumor concept applies more broadly than misinformation and disinformation. Rather than going straight to the question of what to refute, authorities and analysts would first consider the role that a

rumor is playing within the community—an approach that invites deeper insight.

Making the right calls on information is crucial because these phenomena are now inherently adversarial. Mistakenly assessing intent or accuracy can cause a responder to lose credibility. One overarching benefit of a framework like ours is that journalists, authorities, and researchers can get a handle on ever-shifting flows of ambiguous information without risking a reputation-damaging mistake.

More than that, by inviting a stance that seeks to engage with rumors rather than correct misinformation, the framework could make responses more resistant to bad-faith criticism. It could also allow communicators to acknowledge their own uncertainty, recognize the potential information in communities’ rumors, and help rebuild lost trust. We hope this framework around rumors, and what others might build from it, can support quick, effective responses.

Emma S. Spiro is an associate professor at the Information School at the University of Washington. Kate Starbird is an associate professor in the Department of Human Centered Design & Engineering at the University of Washington.