## Project Overview

Online misinformation is an urgent and growing concern because of its potential to influence individual decision-making and behavior. Misinformation can come in many forms—e.g., the deliberate reporting of events that didn't happen (i.e., disinformation), anecdotes or rumors promoted in the place of credible evidence, or inadvertent signals about the status of scientific knowledge. Once adopted, it persists because it is hard to debunk [34]; attempts to do so can even cause it to become more deeply entrenched [21, 54, 66]. These complex, yet prevalent, dynamics of misinformation have led psychologist Steven Lewandowsky to argue that the proliferation of misinformation online has produced an alternative epistemic space, now shared by millions, in which conventional standards of evidence no longer apply and facts are a matter of opinion [52].

Misinformation is not directly responsible for harmful behaviors or bad decisions on the part of people that are exposed to it though. Rather, misinformation becomes dangerous when it is woven into complex, socially maintained narratives, which consist of many interconnected and self-reinforcing stories. Narratives spread easily amongst online groups, and people use them to interpret events in the world and make decisions about who to affiliate with and how to act [3, 11, 24, 43, 70, 78]. This research project focuses on anti-vaccination narratives and the misinformation which has been implicated in the resurgence of anti-vaccine attitudes [42, 43, 71, 79]. This is an urgently needed area of research because the value of vaccination depends on the principle of herd immunity, and small drops in immunization rates could have large impacts on public health, with large costs and potential casualties well beyond those who choose not to vaccinate [17],

The proposed effort will improve our understanding of how misinformation becomes woven into anti-vaccination narratives online, how technology influences this process, and how design might be used to alter it. Researchers have wrestled with how to address anti-vaccine attitudes for many years [18], and significant research has gone into understanding the psychology of vaccine hesitancy in individuals [6]. This proposal will focus upon socially constructed online narratives to address the gap between our understanding of individual psychology and the broader dynamics of many psychologies interacting online. Prior research has begun to examine such narratives by studying the diffusion of digital artifacts like hashtags [e.g., 3, 93]. Such studies can tell us, for instance, that some kinds of narratives are viral, resilient, or immiscible, but not why. The proposed project is unique because it will examine how interdependencies among networks of information, embedded within networks of stories, embedded within narratives, interact with social technology to generate observed narrative dynamics.

To pursue this research, I propose a series of crowd-based experiments to investigate how people in online networks work together to create, modify, and protect narratives. These experiments build on existing research [e.g., 32, 33] and leverage an experimental platform I have developed for this purpose. Narratives are the primary means by which anti-vaccination concerns are spread online [43, 55, 75] and as such an imperative focus in developing a substantive understanding of the interplay between misinformation and anti-vaccination attitudes, specifically, and the dynamics of misinformation online, more generally.

### Vaccination Hesitancy and the Internet

Anti-vaccination sentiments have existed as long as vaccines themselves, but vaccination became widely accepted during the 20$^{th}$ century, ultimately leading to the eradication of diseases like smallpox and polio [18]. In 1998, Andrew Wakefield, along with 12 co-authors, published a study linking the MMR vaccine to autism in the prestigious medical journal *The Lancet*. In 2004, after multiple epidemiological studies found no such connection and revelations that Wakefield had accepted money from legal firms seeking evidence against vaccine manufacturers, Wakefield and 10 of his co-authors retracted their names from the study. In 2010, the UKs General Medical Council found that Wakefield acted 'dishonestly and irresponsibly' and revoked his right to practice medicine in the UK. *The Lancet* soon after took the uncommon step of retroactively retracting his article [18].

Despite the fact that his work has been soundly rejected, Wakefield went on in 2016 to direct the movie "Vaxxed: From Cover-Up to Catastrophe," promoting conspiracy theories about efforts to cover up the link between the MMR vaccine and autism. He has become something of a cult hero in online anti-vaccination forums [43], leading many to pin the current and continuing concerns about autism and vaccines squarely on his shoulders [11]. Accordingly, a sizable body of research literature (reviewed below) focusing on why people continue to believe misinformation about vaccines even after it has been retracted, and how best to correct it [e.g., 31, 64, 68], has since developed.

Yet, medical historian Elena Conis points out that concerns about autism predate Wakefield's article in *The Lancet*, and suggests that Wakefield merely gave voice to one storyline in a complex and evolving narrative [11]. The narrative that encompasses vaccine hesitancy comprises numerous issues, including parental fears about "too many" vaccines "too soon," concerns about the use of additives like mercury in vaccines, shifting attitudes towards drug manufacturers, and an over-zealous medical establishment [11]. These themes are easy to identify in online anti-vaccination content—online posts frequently voice concerns about other toxins and contaminants in vaccines, civil liberties, and the profit motives of drug companies [42, 43, 62]. Moreover, public health communication efforts do little to address this complex set of concerns. For instance, a study of people's reactions to pro-vaccine messages drawn from materials produced by the CDC revealed that people found these messages to be one-sided, manipulative, simplistic, and perhaps motivated by 'Big Pharma' and hence not credible [57].

The anti-vaccine narrative might thus be best understood as a system of many stories, woven not just of misinformation, but also themes of power, privilege, and agency. Misinformation does proliferate in online anti-vaccination discussions, but vaccine-hesitant individuals should not be viewed as "dupes" or victims of unfortunate cognitive biases. Instead, it may be more useful to consider vaccine-hesitancy as an aspect of an emergent online culture [38, 47, 72, 77]. My proposed research embraces this perspective.

**Intellectual Merit of the Proposed Research**

This research will advance our understanding of the complex interactions among technologically mediated social systems and information. From these insights, scholars and designers will be able to clarify how different design features interact with the production of narratives in the context of misinformation and, concomitantly, how sociotechnical strategies should be designed to reduce the resilience and spread of vaccine hesitancy. Most importantly, this research will also generate a set of insights about how the design of online networks can influence the correction of misinformation.

Specifically, this research will improve current models of information contagion to account for the fact that individual pieces of information to which individuals are exposed depend on one another as well as the background knowledge, beliefs, and attitudes of receivers. This project will also develop an experimental paradigm for exploring the dynamics of information in online social networks. While various crowd-based studies using simulated online platforms have been explored previously [9, 45, 58], my proposed methods offer a novel approach for understanding how people seek, exchange, and integrate information in complex online knowledge environments.

The studies proposed here are representative of a large class of such studies that might be run other researchers. The project will generate new datasets that illustrate the coordinated diffusion of narrative information, and these will be made available on the project website. The project will also refine my software platform for running crowd-based experiments, and the refined software will be documented and released for public download. Overall, the proposed effort will help demonstrate the potential of these studies and help disseminate these experimental methods more broadly.

**Broader Impacts**

Widespread vaccination against infectious diseases is widely considered to be one of the greatest achievements of modern public health programs. According to the Centers for Disease Control, vaccines save millions of lives and billions of dollars annually [8], and their success has led to strong public support for childhood vaccination [25]. At the same time, there remain significant misconceptions about

vaccination, and it is feared that the ease with which rumors and misinformation spread in social media is a significant contributing factor [6, 42, 43, 71, 79].

Others have documented the presence of anti-vaccination views online [19, 43, 67, 71, 79] and the impact that different interventions have on the vaccination attitudes of individuals [6, 37, 68]. Yet, there have been no empirically tested recommendations for how to address persistent anti-vaccination attitudes in technologically mediated social contexts. My proposed research project will do precisely this: it will generate realistic sociotechnical interventions that account for the complex and dynamic interplay between online interaction and the proliferation of vaccine hesitancy.

More broadly, the proposed research will contribute new understandings and strategies for coping with digital misinformation that has begun to be integrated into our broader national discourse. There is a dire need for this research. In a recent Pew Report [1], roughly half of the 1,116 experts surveyed voiced pessimistic responses to the following question:

*In the next 10 years, will trusted methods emerge to block false narratives and allow the most accurate information to prevail in the overall information ecosystem? Or will the quality and veracity of information online deteriorate due to the spread of unreliable, sometimes even dangerous, socially destabilizing ideas?*

This project will offer concrete recommendations for how to design social technologies that help limit the impact of online misinformation. The products of this research will be widely disseminated via a public project website, presented not only to academic conferences, but purposively disseminated in popular and social media. The project will generate concrete, empirically based strategies that policy-makers and technology companies can consider as we strive to insulate our democracy from the threat of rampant digital misinformation.

## Stories, Narrative, and Culture

The imperative to investigate narratives in relation to online misinformation dynamics requires a parallel move to clarify the terms *narrative* and *story* in relation to one another. *Narrative* refers to a system of stories [28], and a *story*, according to Labov and Waletsky [50], is a sequence of events told in a specific order to make a rhetorical point. Most commonly, the sequence of events that unfold in a story details the causally connected, intentional actions of characters [81]. In this sense, a narrative is the abstract expression of many different, interconnected stories.

An important property of both narratives and stories is *coherence*. Coherence is how well the information in a story "hangs together" [24]. This concept of hanging together has been investigated relative to story credibility and persuasiveness [13, 29, 30]. It has also been researched relative to causal structure [7, 84] by introducing a content analytical approach for transforming a story into a graph of causal relations [84, 85]. In so doing, Trabasso and colleagues found that a story's structure explains several kinds of variance in how people processed the information it contained: the total number of causal links in and out of a story element (roughly, a noun phrase) correlates with its recall [85]; stories that have linear causal structures (fewer peripheral causal chains) can better be recalled than those that do not; and elements not on the main causal trunk of a narrative tend to be forgotten more easily [83–85].

Other structural factors also play a role in impressions of coherence and story processing. People recall stories as a series of episodes, each centered on the goal-directed actions of characters [81]. Stories that follow this pattern and present clear causal connectives between episodes are easier to understand and remember [35, 56, 80, 81]. These findings hold across cultures, leading some researchers to suggest that the structure of a well-formed story is innate and universal [35, 46, 80]. When stories do not follow these patterns, people recall them with distortions at points where the presumed ideal structure is violated [56] and reorganize the event-based information so that it follows this ideal [2, 36].

Some scholars have looked at how such structural biases might accumulate over time to shape stories as they are passed among groups [56]. In one recently completed (not yet published) study we examined the structure of health-related stories influences the decay of the information they contain as they are passed

from person to person. To accomplish this, we develop four stories, each containing some health information (simple information about metabolic processes underlying a health concern, such as obesity or osteoporosis) and a recommended health behavior (e.g. eating five servings of fruits vegetables), so that the health information justified the health behavior. We then manipulated the structure of each story in various ways to alter the coherence of health information the story. Using a software framework I developed, we used Amazon's Mechanical Turk (AMT) service to run a serial large serial transmission study (colloquially the 'telephone game'), generating over 4000 stories along transmission chains of up to four people.

We found that health information decayed more rapidly than any other kind of information along a transmission chain, but that improving its structural coherence (i.e., placing it on the main trunk of the story) reduced this decay. Somewhat surprisingly though, we also found that when recalled, health information tended to displace the more story-like elements, such as the main character and health behavior itself. In other words, more complex metabolic information was both hard to retain and did not seem to fit well within a story-like form.

As pertains to the current study, these findings offer one explanation for why technical information about the mechanisms of a vaccine or the importance of herd immunity might never outlast a story about a child's health problems following vaccination online. Moreover, it also highlights the need for a deeper understanding of how social interaction may function to filter and shape the kinds of stories that are likely to be re-shared online as we begin to consider how best to design messages for such contexts.

The narratives that people have already adopted also play an important role in the new stories they are able to accept. Acknowledgement of this entanglement leads to the concept of *narrative coherence* [24]. Unlike story coherence, narrative coherence has not been the focus of targeted empirical research. Still, scholars across a range of disciplines have identified different organizing principles that constrain it. For instance, the term *cultural narrative* is used in sociology [69, 70], cultural studies [74] and psychology [46] to describe narratives that convincingly embed culturally-specific norms of behavior and identity into them. These narratives typically contain abstracted plot lines and characters that appear and connect across multiple stories. As with culture in general, cultural narratives are often nested within hierarchical social groupings (e.g., nation, city, family).

Narratives create expectations about why events take place or what comes next; people use these expectations and assumptions to make sense of otherwise disconnected events and fill in the gaps when information is incomplete [69]. At the same time, narratives people have already adopted can constrain what they are willing, or perhaps able, to believe [69]. Finally narratives are social objects. Sociologist Margaret Somers writes, "all of us come to be who we are (however ephemeral, multiple, and changing) by being located or locating ourselves (usually unconsciously) in social narratives rarely of our own making" [78]. As cultural creatures, individuals affiliate with social groups, and these social groups adopt certain cultural narratives, which, in turn, become an integral part of the identity of the social group and its individual members [78].

In this sense, narratives may be understood as systems that reify and perpetuate culture, and the internet is a particularly rich environment for their development [23]. The pervasiveness and similarity of shared stories in online anti-vaccination discussions are one of the hallmarks of such online cultures [43, 55, 62, 75]. This beckons consideration of Douglas and Wildavsky's cultural theory of risk [16], which posits that individuals form opinions of risk that are drawn from the cluster of beliefs and values shared by others in their social context. If we accept this, the challenge of addressing vaccine-hesitancy online is one of how best to reshape an emerging online culture.

Misinformation is a pervasive component of the stories that circulate in online anti-vaccination cultures [42]. Considering the individual, psychological aspects of misinformation and attempts to correct it may help us build towards sociotechnical interventions to dislodge it.

## The Psychology of Misinformation

It is well understood that people do not integrate information to form narratives and beliefs as purely rational information processors. Rather, cognitive limitations lead individuals to rely on a set of heuristic shortcuts [40]. One important cognitive bias evident in decision making is the *confirmation bias*, wherein people tend to seek confirmatory evidence for their prior beliefs and reject disconfirmatory evidence [48]. People's likelihood of adopting information is also influenced by their social connections, generally referred to as *social influence* [92]. These factors can have a large impact on how people process information, evaluate risk, and update misinformation.

Cognitive bias is well illuminated in the case of the *continuing influence* of misinformation [20, 21, 34]. The continuing influence effect describes the phenomenon in which misinformation that is corrected continues to have an effect on the types of inferences people make. For example, various authors have blamed the continuing influence effect for ongoing concerns that the MMR vaccine causes autism even after the research making this connection was retracted [18].

The experimental paradigm used to test the continuing influence effect involves three steps: in an *encoding step*, individuals are exposed to a series of news items that either contain or imply misinformation; in an *intervention step*, subjects are exposed to retractions, rebuttals and / or warnings about the possibility of misinformation; and in an *evaluation step*, subjects are tested (typically after a time lag of roughly a week [82]) via survey items and open-ended interview questions to determine whether or not misinformation still affects their information processing. These studies have consistently found that despite corrections, misinformation persists and continues to inform how people make future inferences [21, 53, 63, 65, 66].

Story coherence, as discussed above, may be an important factor in the continuing influence effect. Highly coherent stories are hard to disrupt, even when they contain misinformation, and so retractions may not have much impact [13, 29, 30, 53]. This may be because people seek to fill in gaps in the causal structure of a story, and a retraction that deletes, but does not replace a previously established causal connection is likely to regress [13, 29, 30, 53]. As evidence of this, Nyhan and Reifler [63] illustrated how repairing the coherence of an encoded story following a retraction (e.g., by providing alternative causal explanations) reduces the continuing influence effect.

Narrative coherence might further exacerbate the continuing influence effect. It has been found that in some cases, retractions can cause an individual become more deeply committed to the original misinformation [53, 66]. Lewandowsky et al. [53] hypothesize that this can happen if retractions challenge an individual's *worldview*. Lewandowsky uses the term worldview to refer to a person's ideology, which is not precisely the same as either culture or narrative, but is heavily informed by them. From this perspective, a rebuttal or retraction that does too much violence to the coherence of an individual's narrative might cause them to develop defenses against future attacks. Whether or not there is a specific degree of incoherence that trigger such a worldview-backfire effect, and how other factors might mitigate it are current areas of study [12, 86].

In the case of social influence, the perceived credibility of the source of information plays a large role in people's willingness to accept it. Credibility includes both the expertise and trustworthiness of the source, but research suggests that the latter plays a far larger role [27]. For instance, trust can increase with the simple repetition of an unknown name [53]. Online, social endorsements such as 'likes' or crowd-based ratings also play an important role in signifying credibility [60, 61], as do signals that suggest that there are similarities between the message sender and recipient (homophily) [87].

Another dimension of social influence is the innate need for social acceptance, which can lead people to yield to conformance pressures, and/or engage in social interactions intended to advance one's status [10, 90]. Ultimately, numerous factors operate together such that the probability a given individual will adopt information correlates with number of their social contacts that already have it. This has led to the development of the *threshold model of social contagion*, which suggests that the diffusion of information

across a social network depends on the topology of the network and the distribution of a social influence threshold parameter across that network [14, 26, 88]. Research has demonstrated that social influence and contagion can have strong effects on false memory and beliefs in misinformation [22, 59, 91]. However, it is important to note that research on social contagions has tended to focus on immutable objects like hashtags and hyperlinks [51, 73, 76]. As illustrated by my serial transmission study, complex knowledge objects like stories are likely to change as they are shared, complicating the current threshold contagion model.

Drawing together a cultural theory of risk [16], the confirmation bias, and social influence theory, Kahan has proposed a *cultural cognition* theory of risk [37–39]. In this theory, cognitive biases are explained as stemming from individual's desires to protect and promote a worldview (as with Lewandowky's work, a worldview is a kind of ideology). Drawing on Douglas' [15] typology of worldview, Kahan found that people are in general likely to reject messages that are inconsistent with their worldview, but that this effect is substantially reduced when the source of the message is perceived to be aligned with the receiver's worldview [37]. This may be understood as a particular form of homophily that is connected to both a cultural theory of risk and the confirmation bias.

My proposed research builds from this point, and considers the particular dynamics of story and narrative construction in online social environments. These considerations are necessary as we move towards strategies for dealing with misinformation that might be realistically deployed in such contexts.

## Research Objectives and Study Design

The objective of this research is to advance the understanding of the complex interactions among narratives, technologically mediated social systems, and misinformation. As a basis for my investigation, I combine prior work that examines narrative, in particular as part of an online culture, and the psychology of belief in misinformation. This entails understanding the dynamic processes though which misinformation becomes embedded in and sustains the stories and narratives that arise in online networks. I pursue this investigation via three research questions, which are organized in terms of increasing breadth:

1) What is the relationship between coherence and how quickly a story saturates a social network?
2) How does coherence interact with signifiers of credibility to guide how people select among competing sets of information?
3) How do network diversity and the content of corrective messages influence attempts to correct misinformation embedded in socially shared stories?

Each of the three research questions correlates to a separate experimental study. The cumulative order of these studies aims to contrast baseline collective performance with increasingly complex factors. The first focuses on establishing the relationship between coherence and how quickly a story saturates a social network; the second considers how coherence interacts with signifiers of credibility to guide how people select among competing sets of information; the third and final considers the impact of homophily and content to evaluate the efficacy of different kinds of retractions. Each study includes a vaccine-relevant narrative and will manipulate the alignment of the experimental population with each story's perspective, using a previously validated scale such as the one described in [31]. To perform each of these studies, I propose to use my bespoke software platform, the Story Loom, described below.

### Experimental Platform: The Story Loom

To gain further insights into how interdependencies among pieces of information influence their uptake in online networks, I have designed and tested a platform called The Story Loom. The Story Loom is a synchronous, web-based game that connects players in an online network. The game is designed to work with (but does not depend upon) Amazon's Mechanical Turk service, and scales to hundreds of players who can be arranged in a network that is structured to the experimenter's specification.

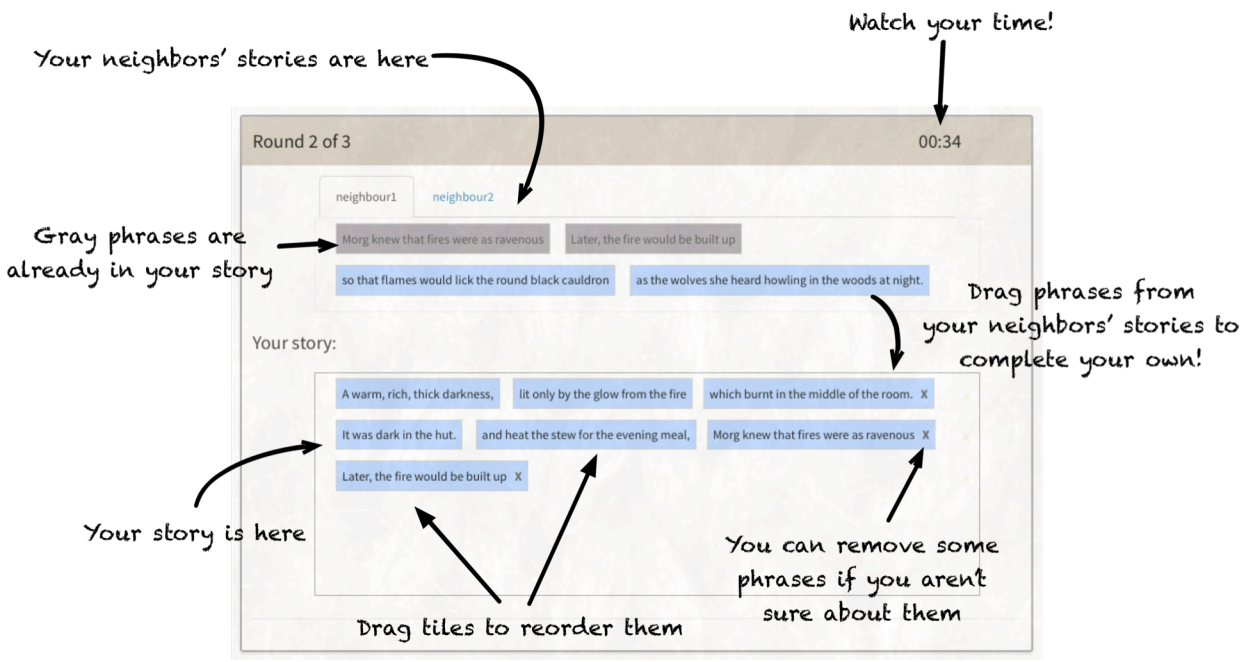


**Figure 1: The user's workspace in the Story Loom game. Image taken from the training component of the platform.**

The goal of the game for each player is to combine disparate information elements to form the "best" story they can. The information elements may be phrases in a story, or brief news snippets, or any other text-based material the experimenter chooses. When the game starts, each player is provided with a small set of information elements, randomly drawn from the pool of all possible information elements. The player's interface provides him or her with a workspace in which to reorder their elements, as well as a view of their neighbors' workspaces (Figure 1). Players can copy elements they don't already have from their neighbors. The game proceeds in time-limited rounds, and in each round, a player tries to improve their story. Players are rewarded for stories that have more elements, but penalized when elements are out of order. The game thus sets up a situation where the diffusion of information across a network is constrained by people's perceptions of how coherent each element is with their current, partially formed story.

To recruit participants for the platform so that studies can be run in a timely manner, it is necessary to cultivate a large pool of users who I can draw upon. To achieve this, I've developed a training platform that collects demographics information, provides initial screening (e.g., for reading ability, or in this case, vaccination views), and then assigns a "qualification" (a technical affordance of the AMT platform) that gives them access to the game, and provides me with a way to contact these individuals. For example, to run pilot tests, I first cultivated a large pool of subjects (approximately 500 MTurk workers). With this population, I was able to recruit 50-100 users for any given study in a matter of minutes.

I have run several pilot studies with the platform. In one, I manipulated the causal structure of a story in order to reduce its structural coherence. I ran the study with two groups of fifty participants, each arranged in a lattice network butusing a different version of the story. The group with the modified story took roughly twice as long to attain 80% saturation (80% of the participants with a complete, correct story) as the group with the original story, suggesting that the reduction in coherence (as defined by Trabasso and colleagues [84]) can impair the speed with which is saturates a network.

In another pilot, I examined the speed with which each element in a story moved across two different groups of fifty people arranged in different network configurations—a lattice network and a small world network. Although the overall speed of saturation varied (the small world network was faster), the relative speed with which individual story elements diffused across the network was highly correlated ($R=.65$),

suggesting that the role an information element plays in an story (and not the particular group of people, or network configuration) plays a large role in how quickly it will move across a network and be adopted by individuals.

My initial pilot work demonstrates the feasibility of the framework, and suggests that the structure and composition of stories indeed play a significant role in how people, arranged in networks, process them. The Story Loom enables a wide range of novel experiments for exploring the interplay between impressions of story coherence and other factors that are know to influence information spread. For example, with small modifications to the user interface, it is possible for me to examine how different types of design interventions (e.g., signals about source credibility) change how people interact with stories. This is a key component of my proposed study design, described below.

For the proposed project, I will recruit an initial pool of approximately 5000 users, split across those with pro- and anti-vaccine views. Several studies [31, 68] have published validated scales for vaccination preferences, and I will reuse one of these. I will also collect information about each participant's worldview, using the scale used in [37], and other factors that might have important bearing on the performance of individuals, such as political leaning, numeracy, and reading ability. These will be reserved as control factors in statistical modeling.

**Study 1: How does coherence influence the development of stories in social networks?**

The threshold contagion model [26] has been highly influential in describing how information is transmitted across online social networks, leading to extensive modeling of the spread of many types of information, digital artifacts, and behaviors. However, this model cannot account for how interdependencies among different information elements influence the diffusion process. For example, adopting a technique for preventing aliens from reading one's thoughts presumably depends upon one's belief that they are trying to do so. Improving our understanding of how such path dependencies influence the diffusion of information may open the door to new strategies for combating misinformation.

With Study 1, I will operationalize these interdependencies via a coherence function and use this function to develop a model of information diffusion using empirical data generated via a series of Story Loom studies. To clarify, consider a social network wherein all knowledge can be represented as a *k*-length vector of a set of bits, where each bit is a distinct piece of information (e.g., a proposition). Every individual *i* is associated with a specific binary vector $\mathbf{s}^i$, where any element equal to 1 indicates that agent *i* has adopted it (i.e., believes it to be true). The set **S** is the set of all possible knowledge states; that is $\mathbf{S} = \{s \in \{0,1\}^k\}$. Thus, we can define the coherence function *coh* as a function that maps each vector **s** and a new state for information element *l* to the real range (0,1):

$$\begin{aligned} &coh: X \to (0,1) \\ &X = \{(\mathbf{s}, s_l) | \forall \mathbf{s} \in \mathbf{S}; 0 < l \leq k\} \end{aligned} \tag{1}$$

The coherence function may be interpreted as a conditional probability table that determines the likelihood any given individual will adopt (or reject) any information element *l* given that individual's current state of knowledge ***s***.

With coherence defined in this manner, one can develop a model of information diffusion wherein the probability that an individual *i* in a social network will adopt a given piece of information *l* at time *t* is a function of the coherence of *l* with the individual's existing knowledge, the number of *i*'s neighbors that have that knowledge, and the individual's social influence threshold $\tau^i$:

$$P\left(s_{l,t}^i = 1\right) = f\left(\sum_{j=0}^{\infty} w_{ij}\left[s_{k,t-1}^j = 1\right], \tau^i, coh\left(\mathbf{s}^i, s_k\right)\right) \tag{2}$$

Equation (2) is a modified version of a generalized threshold contagion model. Neither the behavior of *coh* nor its precise relationship with $\tau^i$ are currently known. I will use data from a series of Story Loom

trials to determine these unknowns. Considering a Story Loom experiment with a story that has $k$ elements, the state of any participant in a round of a game is exactly the vector $\mathbf{s}^i$ in the preceding model. Participants can adopt new information when they drag a new information element into their workspace, thus transforming an element in $\mathbf{s}$ from 0 to 1. The function *coh* is thus fully specified by $k * 2^k$ parameters (entries in the conditional probability table described above). Following prior work [41], I will further assume that $\tau$ is drawn from a normal distribution (truncated at zero) and anticipate additional parameters that dictate the interaction between $\tau$, the overall degree of social influence (quantified as the summation in equation (2)), and *coh*.

This is a generative model that is fully specified by a multivariate probability distribution. The model generates the sequence of adoptions for any given individual and the overall diffusion pattern given a network of individuals. It is possible to infer the parameters of the model using a Markov chain Monte Carlo (MCMC) algorithm such as a Gibbs sampler. This approach is a well-understood computational procedure, perhaps best known for its application to probabilistic topic models (e.g., Latent Dirichlet Allocation [4]). By applying this approach to multiple stories, I will use the parameters for each story-specific coherence function to examine how different variables influence perceptions of coherence, as described below.

*Experiment Design*: I will develop a set of six stories. Two of the stories will address vaccination issues: one pro-vaccination and the other anti-vaccination. Another pair of stories will address medical issues (e.g., the treatment of a condition with a new drug) and will not include any information about vaccination, but will otherwise have parallel elements. That is, if the anti-vaccination story includes "Big Pharma" as an element, the health story will also include "Big Pharma." The causal and episodic organization of these stories will also be identical to the vaccination stories. The content of the third pair of stories will not be related to the initial stories—e.g., they may be simple folk tales, without a health message—but they will have the same episodic and causal structure as the initial pair of stories. This design will provide me with an indication of how much variance can be accounted for by causal structures alone, and conversely, how much impact vaccination views have on perceptions of coherence. In some sense, the experiment will enable me to compare the relative impacts of story and narrative coherence.

For each story, I anticipate running approximately 10 groups of 50 participants balanced across two conditions: groups with pro-vaccination views and groups with anti-vaccination views. I will determine the precise number of groups to use in advance by simulating the experiment using equation (2), evaluating the perplexity of the induced model (an information theoretic measure of probabilistic fit) and conducting a power analysis of the overall experimental framework.

*Analysis*: The experiment will generate numerous types of information, including: how information elements spread in the context of a larger story, how easily people are able to infer the correct ordering of that information, the impact of network topology on the virality of information, the potential for information overload, among others. However, to answer my orienting study question, I will focus on two types of analysis: refining the model offered in equation (2) to deduce the structure of *coh* in each story (using the probabilistic techniques described above), and analyzing the element probabilities in *coh* in relation to a several factors.

Refining equation (2) is a matter of testing different assumptions about how $\tau$ interacts with *coh*. I will test at least two different assumptions. One assumption is that $\tau$ and *coh* are linearly independent, and their relative contributions are identical across information elements for each individual. The alternative assumption is that these elements are not independent—for example, as the coherence of an element increases the impact of $\tau$ may be reduced. To select between these models, I will use perplexity (a measure of probabilistic fit) to evaluate how well the different model choices fit the data.

Having arrived at a concrete model, I will examine the different coherence functions using multilevel regression modeling to understand how different factors influence the coherence of different information elements. In this case, the coherence of an element is an aggregate of its conditional probabilities. Prior research suggests several factors, but I anticipate that the analysis will not be purely deductive. Among

1. **Birth Announcement**: Mark and Wendy Smith welcome baby Joseph; 7 lbs. 8 oz.
2. **CDC Website**: Infants can be infected with Hepatitis B at birth and should be administered their first dose of the vaccine unless they are experiencing medical issues.
3. **St. Mary's Hospital, Medical Records for Joseph Smith**: Difficulty breathing, diagnosed with Meconium Aspiration Syndrome. Admitted to NICU.
4. **St. Mary's Hospital, Medical Records for Joseph Smith:** 8 days old, post-natal visit. Minor lung congestion. Advised follow-up ped visit if conditions do not improve. Administered Hepatitis B vaccine.
5. **Mark and Wendy Smith's blog**: "Joseph received his hepatitis B shot on his eighth day of life. Hours later, he was crying inconsolably, refusing to eat, and taking on seizure-like postures."

6a. **[NON-COHERENT] Ridgewood Diagnostics, Allergy Panel results for Joseph Smith**: Highly allergic to ampicillin, gentamicin.

6b. **[COHERENT] Ridgewood Diagnostics: Allergy Panel results for Joseph Smith:** No allergies detected.

7. **Ridgewood Pediatric Center: Dr. Fatulah's handwritten notes:** Suspected lung infection. Administered ampicillin.
8. **Mark and Wendy Smith's blog:** "Later, when we were at Children's Hospital, nearly a dozen specialists saw Joseph's. All said about the same thing, 'Something insulted his system.'"

9a. **[NON-COHERENT] Mark and Wendy Smith's blog:** When Mark and I suggested it was an allergic reaction, we were told that this was quite possible."

9b. **[COHERENT] Mark and Wendy Smith's blog:** When Mark and I suggested it was the hepatitis B vaccine, we were continually told that this was not possible."

10. **Children's Hospital: Medical Records for Joseph Smith:** Systemic inflammation; low platelet count; blood in lungs. Admitted to ICU.
11. **Children's Hospital: Medical Records for Joseph Smith**: 24 days old. Deceased 10:42 PM.

**Figure 2: Sample materials for the anti-vaccination story in RQ2**

the factors to be considered are: the alignment of the test population with the story (in the case of vaccine narratives), the structural coherence of each element (following Trabasso et al.'s work [84]), the role that element plays in an idealized story grammar [e.g., 81], and the identity of the story itself. This analysis will allow me to quantify the relative influence of structural coherence, story role (the role a story element plays in a story), and narrative coherence (i.e., compatibility of the story with pre-existing attitudes), yielding a model that explains how the structure of a story interacts with people's pre-existing attitudes to influence it spread within an online network.

**Study 2: How do different signals about information credibility influence creation of stories?**

The perceived credibility of an information source can have a large impact on whether or not people are willing to adopt that information. I will examine how credibility signals interact with social influence and narrative coherence to influence the development of stories in social networks. Study 2 places design factors that influence perceived credibility in competition with coherence to examine how groups select from competing sets of information when constructing stories. Note that in this study, the coherence of an information element will be determined according to its alignment with the background narrative of each individual (see Figure 2 for an example of coherent information items).

To facilitate these experiments, I will introduce two design changes to the Story Loom platform. The first will provide participants with direct information about the trustworthiness of the source of each information element. Different media types—e.g., a blog, a scientific journal, or news article—will be used, but (following Kahan [37]) the trustworthiness of each source will be manipulated by providing information about each source's worldview. The second design change will provide a "social endorsement" signal, indicating how many others in the network have adopted a piece of information. This will be manipulated as part of the experiment.

*Experiment Design*: I will develop two stories: one for and the other against vaccination. For each story, I will construct a set of incompatible information items (referred to here as the target items) that render the story more or less coherent with its vaccine stance. Figure 2 provides an example of a possible anti-vaccination story (drawn from a viral online story described in [75]). During the experiment, all elements will be made available in the network, and participants will be told that some elements might not fit.

The experiment will manipulate source credibility in relation to the coherence of the target elements. Source trustworthiness will vary in three ways: *not present, coherent*, and *anti-coherent*. In the *not present* condition, the trustworthiness signal will be absent. In the *coherent* condition, the trustworthiness of the source will correlate with the coherence of the target item. That is, more coherent target items will come from more trustworthy sources. In the *anti-coherent* condition, coherence will be inversely correlated with trustworthiness. In other words, less coherent target items will be marked as coming from trustworthy sources.

Social endorsements will be manipulated in a similar fashion: in a *no manipulation* condition, social endorsements will reflect the actual number of people that have adopted the information. In the *coherent* condition, social endorsements will be manipulated so that the signaled rate of adoption correlates with the degree of coherence of the target item. In the *anti-coherent* condition social endorsements will be manipulated to inversely correlate with the degree of coherence in the target item.

The proposed manipulations yield a total of nine conditions for each story. As in Study 1, the exact number of replications will be subject to power analysis, but I anticipate running roughly 10 replications with 50 people in each group, balanced across pro- and anti-vaccination groups.

*Analysis*: To analyze the data, I will follow the approach outlined in Study 1 to establish coherence estimates for information elements. Here, it will be necessary to update the coherence function to include transitions that allow for the *rejection* of information items. Then I will use regression modeling to evaluate the impact of the manipulated factors on the aggregate coherence probabilities of the target items. I hypothesize that the coherence probabilities will be greatest for coherent items that are aligned with the participants' vaccination stance. For example, in Figure 3, I anticipate that anti-vaccination groups will adopt item 6b more quickly than 6a. A more significant question is the degree to which the two credibility indicators interact with the coherent items. The results of this study will develop deeper insights on how different technology-based signals about credibility moderate the development of stories when competing, less coherent information is present.

**Study 3: How do network diversity and the content of corrective messages influence attempts to correct misinformation embedded in socially shared stories?**

Moving towards realistic interventions in online misinformation, it is important to understand how online social environments interact with the different kinds of corrections that are presented. Study 3 examines the impact of perceptions of homophily and the relative breadth of a correction, where breadth concerns how many different storylines in a background narrative that a correction impacts.

Homophily may impact the spread and uptake of misinformation in online anti-vaccination groups in two ways. First, a high degree of homophily in an anti-vaccination group makes it is unlikely that corrections that will spread. On the other hand, if a correction indeed comes from a source that is aligned with one's worldview (hence a homophilous tie) it may not be so impeded [37]. How these factors interact in diverse networks is not known.

The breadth of the correction with respect to a background narrative has not yet been explored, but prior work suggests that this is a potentially important variable. Studies have suggested that the coherence of a correction with a person's worldview may be a critical consideration [37, 53], and while a worldview is not a narrative, but a background narrative as described here embodies a worldview. I hypothesize that correcting information in a manner that impairs just one storyline in a narrative will leave that narrative in an incoherent state. Consistent with Lewandowsky's worldview-backfire effect [53], an individual will subsequently repair their narrative by forcefully rejecting the correction. If, however, corrections address

a broader spectrum of stories within a person's background narrative, the narrative may be altered but left in a more coherent state, avoiding the need for repair and hence the backfire effect. This logic is consistent with Nyhan and Rieffler's finding's regarding the delivery of causal repairs along with corrections to misinformation [63].

I will use the Story Loom platform to examine the impact of homophily, the breadth of a correction, and their interaction, in the context of corrections made to an anti-vaccine story. I will distribute a set of information items and ask individuals in a group to gather all of the relevant information needed to answer a probe (e.g., "What was the most likely cause of the illness?"). As in study two, a subset of the information elements provided will be highly coherent with an anti-vaccine perspective, but will connect to different storylines (referred to here as the target elements). Other elements that are not related to the story will also be distributed, and each person will be asked to create their own story by selecting just the relevant information, adopting information from their neighbors as they see fit. Following this, each participant will be asked to record his or her answer to the probe question.

Once this task is complete, I will present the group with additional information that retracts (or refutes) one or more (subject to the experimental condition, described below) target elements in the initial story. Participants will be given an opportunity to incorporate this information into their previous set of items, and will be given an opportunity to update their answer to the probe question.

After a period of time has elapsed (1 week, following prior work [82]), I will expose the original group to a new set of information that builds on the story that was initially presented. Participants will be told that some of the information is inaccurate and asked to construct the most plausible story from a subset of the information provided. The information will support two coherent interpretations: one will contain an inference drawn from the initially presented story and the other will contain an inference drawn from the "corrected" story.

*Experiment Design*: The envisioned experiment will manipulate two dimensions: the content of the intervention and perceptions of homophily. To manipulate content, the story I develop will include three items that correspond to major themes identified in online anti-vaccination discussions [42, 43]. The first element will suggest unethical practices on the part of a drug manufacturer, the second element will raise concerns about ingredients in a vaccine, and the third will raise concerns about interactions between multiple vaccines given in quick succession. The content manipulation will vary the number of items that are retracted (1, 2 or 3 items), and participants will have a choice in whether or not to adopt the intervention.

To manipulate perceptions of network diversity, I will first modify the Story Loom interface so that participants are provided with information about their neighbors' vaccine-related opinions (a homophily signal). I will then recruit even numbers of individuals with pro-vaccine and anti-vaccine attitudes and arrange them so that each person has the same number of neighbors that are ideologically aligned as are not. I will manipulate perceptions of homophily across three conditions: in the *anti-vaccination* condition, all neighbors will appear to hold anti-vaccination attitudes; in the *pro-vaccination* condition, all neighbors will appear to hold pro-vaccination attitudes; in the *actual* condition, signaled attitudes will correspond to the true attitudes of the participants. As with the prior questions, it will be necessary to perform a power analysis, but I anticipate running approximately ten replications in each unique condition.

*Analysis:* I will perform two levels of analysis in Study 3. The first will use statistical methods to consider the diffusion of corrections in each condition. More concretely, I will analyze the number of individuals who receive the corrected information at the end of the correction phase as a function of the manipulated factors. In the second level of analysis, I will focus on the continuing influence of misinformation in each condition. That is, I will use an analysis of variance (ANOVA) to evaluate the impact of each experimental condition on the likelihood that people make the inference that is based on the corrected information. The results of Study 3 will shed light on how perceived homophily influences the spread of ideologically aligned (or mis-aligned) corrections to misinformation.

**Summary**

The proposed series of studies will build new knowledge about how a variety of factors influence how pieces of information (factual or not) become woven into stories that in turn connect with background narratives maintained by online networks. In particular, it will highlight how perceived trustworthiness, social endorsements, and homophily can influence the choices that people make about which kinds of information they are willing to accept. These latter findings offer a concrete basis to begin thinking about how to design social technologies that remain permeable to good information.

However, this series of experiments offers yet another potential avenue for analysis that I have not surfaced, but intend to take advantage of. This analysis focuses on the longitudinal process via which online networks of people connecting the bits of information they encounter online to construct knowledge as embodied by stories that are connected through narratives. This would be a significant advance beyond current models of information diffusion—such as with the study of opinion dynamics [5, 89]—that are used to model collective decision processes.

For example, the *coh* function developed in Studies 1 and 2 can be understood as a probabilistic transition matrix that produces a dynamic system similar to Kauffman's NK networks [44] and Kosko's fuzzy cognitive maps [49]. Although a rarefied model of human knowledge, these transition matrices capture a likely progression of knowledge states in groups under a variety of different conditions. Examining the rules that govern these dynamic systems can yield new insights about the ways groups process information, and perhaps the kinds of knowledge structures that are likely to evolve under different conditions. In other words, the data produced by my proposed studies will become a rich resource for developing deeper insights about collective intelligence, and how technology might impact it.

**Risks and Challenges**

This is an ambitious project, and each study will face somewhat different risks and challenges. However, in each study, the most significant risk is that of a non-finding. For example, in study 1, it is possible that I may not be able to identify any reliable correlates of element coherence among the variables I collect. However, my pilot results and prior work with story decay suggest that this is an unlikely outcome. I have already uncovered some preliminary evidence that causal structures explain some of the variance in how different story elements interact. If a given study does not result in the anticipated findings, this will not have any impact on the other studies, as the individual studies do not build on one another.

I have previously addressed the technical challenges that might otherwise impact all of the studies. My platform has been tested with groups up to 200 people, and I have encountered no difficulty in recruiting from the AMT population. One possible concern is that Amazon might discontinue its AMT service. This seems unlikely, but it is worth noting that the StoryLoom does not explicitly depend on AMT. The platform maintains it's own database with login credentials for each user, and runs on a standalone web-server. Any other recruitment tool will suffice, such as CrowdFlower, CrowdFactory, Qualtrics, etc. I have budgeted to pay workers $11.25 an hour which is a fair wage across a range of platforms.

A more speculative risk is that in generating public media attention, the project itself becomes a target of anti-vaccination groups, potentially reducing the impact of my results, or more plausibly, leading the project to become labeled as an interventionist social engineering effort. I do not believe the impact of this to be that serious, but I believe it is important to be as transparent as possible in the research effort, and moreover to focus attention on *misinformation* in online anti-vaccination groups, rather than vaccine hesitancy itself. The latter is a choice that individuals are free to make, but they should make such choices with the best possible information.

## Results from Prior NSF Support

Although I have not been the recipient of any prior NSF support, I have a great deal of experience managing large and complex research projects. While an industry scientist, I served as PI or acting PI for a range of Small Business Innovation Research (SBIR) grants funded by DARPA, the Air Force (AFRL), and the Army (ARDEC). I was also the Chief Architect of the MIT Climate Colab, which is a web

**Table 1: Project plan. Each column reflects an academic term (Summer, Spring, Fall), with an anticipated start date of 6/19**

| Task | Y1 | | | Y2 | | | Y3 | | |
|---|---|---|---|---|---|---|---|---|---|
| 1. **Study 1: Story coherence and Information Contagion** | | | | | | | | | |
| 1.1. Materials preparation | | | | | | | | | |
| 1.2. Recruitment | | | | | | | | | |
| 1.3. Project Web Site | | | | | | | | | |
| 1.4. Study | | | | | | | | | |
| 1.5. Analysis & Reporting | | | | | | | | | |
| 2. **Study 2: Social signifiers and narrative coherence** | | | | | | | | | |
| 2.1. Materials preparation | | | | | | | | | |
| 2.2. Pilot testing | | | | | | | | | |
| 2.3. Platform updates | | | | | | | | | |
| 2.4. Study | | | | | | | | | |
| 2.5. Analysis & Reporting | | | | | | | | | |
| 3. **Study 3: The continuing influence effect in online network** | | | | | | | | | |
| 3.1. Materials preparation | | | | | | | | | |
| 3.2. Pilot testing | | | | | | | | | |
| 3.3. Study | | | | | | | | | |
| 3.4. Credibility validation | | | | | | | | | |
| 3.5. Analysis & Reporting | | | | | | | | | |

platform for crowdsourcing climate change solutions. I am currently PI on a project funded by the Michigan Department of Education ($275,000) to develop a mobile platform to address food access issues in Flint, MI., as well as an industry-funded project ($50,000) to develop an SMS-based platform to provide social support for low income HIV/AIDS patients.

## Personnel & Project Plan

I will serve as the lead on all aspects of the project, and work with one graduate research assistant to develop the materials for and execute these studies. My research assistant will be responsible for organizing the experiment materials, updating the Story Loom codebase as necessary, overseeing the studies, and preparing results for the public website. The research assistant will have an authorship role all of the papers generated by this research, and I will mentor her or him in all aspects of the project.

Although the proposed studies are large and complex, I have substantial experience running such studies. As discussed previously, the Story Loom platform simplifies the process of recruiting large pools of users, and I will train my research assistant to use the platform and interact with the workers on AMT. The Story Loom does not require high-powered computational hardware, and may be run using my existing server infrastructure.

The overall project plan is structured so that the three studies are run sequentially, one in each year of the experiment. Table 2 provides an overview of the sequencing of activities. In general, the overall flow of the project is straightforward and nearly identical in each year. The bulk of the preparatory work will be done each summer, each study run in the fall, and the spring devoted to data analysis and paper writing.

In each year, significant attention will be devoted to analysis and reporting. For the duration of the project and foreseeable future all products, including datasets, papers, and software will be made available on a project website. I will also host the source code in a publicly available source code repository like GitHub using a permissive open source license, and make the datasets available on an

open source data repository like ICPSR. I have budgeted for two conferences a year for my research assistant and myself in order to disseminate this to academic audiences. In addition, I will maintain a weblog on the project website, to be updated in the Spring of each year in order to provide a less technical overview of the project, and I will work with my college's highly effective communication team to disseminate this work through popular media and to non-academic audiences.

**References**


[1] Anderson, J. and Rainie, L. 2017. *The Future of Truth and Misinformation Online*. Pew Research Center: Internet, Science & Tech.

[2] Bartlett, F. 1932. Remembering: An Experimental and Social Study. *Cambridge: Cambridge University*. (1932).

[3] Bessi, A., Coletto, M., Davidescu, G.A., Scala, A., Caldarelli, G. and Quattrociocchi, W. 2015. Science vs Conspiracy: Collective Narratives in the Age of Misinformation. *PLoS ONE*. 10, 2 (Feb. 2015), e0118093.

[4] Blei, D.M., Ng, A.Y. and Jordan, M.I. 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.* 3, (Mar. 2003), 993–1022.

[5] Bravo-Marquez, F., Gayo-Avello, D., Mendoza, M. and Poblete, B. 2012. Opinion Dynamics of Elections in Twitter. (Oct. 2012), 32–39.

[6] Brewer, N.T., Chapman, G.B., Rothman, A.J., Leask, J. and Kempe, A. 2017. Increasing Vaccination: Putting Psychological Science Into Action. *Psychological Science in the Public Interest*. 18, 3 (Dec. 2017), 149–207. DOI:https://doi.org/10.1177/1529100618760521.

[7] van den Broek, P. 1994. Comprehension and memory of narrative texts: Inferences and coherence. *Handbook of psycholinguistics.* Academic Press. 539–588.

[8] Centers for Disease Control and Prevention (CDC) 1999. Impact of vaccines universally recommended for children--United States, 1990-1998. *MMWR. Morbidity and mortality weekly report*. 48, 12 (Apr. 1999), 243–248.

[9] Centola, D. 2010. The Spread of Behavior in an Online Social Network Experiment. *Science*. 329, 5996 (2010), 1194–1197. DOI:https://doi.org/10.1126/science.1185231.

[10] Cialdini, R.B. and Trost, M.R. 1998. Social influence: Social norms, conformity and compliance. *The handbook of social psychology, Vols. 1 and 2 (4th ed.)*. D.T. Gilbert, S.T. Fiske, and G. Lindzey, eds. McGraw-Hill. 151–192.

[11] Conis, E. 2017. Vaccines, Pesticides, and Narratives of Exposure and Evidence. *Canadian Bulletin of Medical History*. 34, 2 (Sep. 2017), 297–326. DOI:https://doi.org/10.3138/cbmh.190-21122016.

[12] Cook, J. and Lewandowsky, S. 2016. Rational Irrationality: Modeling Climate Change Belief Polarization Using Bayesian Networks. *Topics in Cognitive Science*. 8, 1 (Jan. 2016), 160–179. DOI:https://doi.org/10.1111/tops.12186.

[13] Dieckmann, N.F., Slovic, P. and Peters, E.M. 2009. The Use of Narrative Evidence and Explicit Likelihood by Decisionmakers Varying in Numeracy. *Risk Analysis*. 29, 10 (Oct. 2009), 1473–1488. DOI:https://doi.org/10.1111/j.1539-6924.2009.01279.x.

[14] Dodds, P.S. and Watts, D.J. 2004. Universal Behavior in a Generalized Model of Contagion. *Physical Review Letters*. 92, 21 (May 2004), 218701. DOI:https://doi.org/10.1103/PhysRevLett.92.218701.

[15] Douglas, M. 1996. *Natural Symbols: Explorations in Cosmology*. Psychology Press.

[16] Douglas, M. and Wildavsky, A. 1983. *Risk and Culture: An Essay on the Selection of Technological and Environmental Dangers*. University of California Press.

[17] Dubé, E., Laberge, C., Guay, M., Bramadat, P., Roy, R. and Bettinger, J.A. 2013. Vaccine hesitancy. *Human Vaccines & Immunotherapeutics*. 9, 8 (Aug. 2013), 1763–1773. DOI:https://doi.org/10.4161/hv.24657.

[18] Dubé, E., Vivion, M. and MacDonald, N.E. 2015. Vaccine hesitancy, vaccine refusal and the anti-vaccine movement: influence, impact and implications. *Expert Review of Vaccines*. 14, 1 (Jan. 2015), 99–117. DOI:https://doi.org/10.1586/14760584.2015.964212.

[19] Dunn, A.G., Surian, D., Leask, J., Dey, A., Mandl, K.D. and Coiera, E. 2017. Mapping information exposure on social media to explain differences in HPV vaccine coverage in the United States. *Vaccine*. 35, 23 (May 2017), 3033–3040. DOI:https://doi.org/10.1016/j.vaccine.2017.04.060.

[20] Ecker, U.K., Swire, B. and Lewandowsky, S. 2014. Correcting misinformation—A challenge for education and cognitive science. *Processing inaccurate information: Theoretical and applied perspectives from cognitive science and the educational sciences*. D.N. Rapp and J.L.G. Braasch, eds. The MIT Press. 13–38.

[21] Ecker, U.K.H., Lewandowsky, S., Fenton, O. and Martin, K. 2014. Do people keep believing because they want to? Preexisting attitudes and the continued influence of misinformation. *Memory & Cognition*. 42, 2 (Feb. 2014), 292–304. DOI:https://doi.org/10.3758/s13421-013-0358-x.

[22] Edelson, M., Sharot, T., Dolan, R.J. and Dudai, Y. 2011. Following the Crowd: Brain Substrates of Long-Term Memory Conformity. *Science*. 333, 6038 (Jul. 2011), 108–111. DOI:https://doi.org/10.1126/science.1203557.

[23] Fernback, J. 2003. Legends on the net: an examination of computer-mediated communication as a locus of oral culture. *New Media & Society*. 5, 1 (Mar. 2003), 29–45. DOI:https://doi.org/10.1177/1461444803005001902.

[24] Fisher, W.R. 1989. *Human Communication as Narration: Toward a Philosophy of Reason, Value, and Action*. University of South Carolina Press.

[25] Gellin, B.G., Maibach, E.W., Marcuse, E.K. and Committee‖, for the N.N. for I.I.S. 2000. Do Parents Understand Immunizations? A National Telephone Survey. *Pediatrics*. 106, 5 (Nov. 2000), 1097–1102. DOI:https://doi.org/10.1542/peds.106.5.1097.

[26] Granovetter, M. 1978. Threshold Models of Collective Behavior. *American Journal of Sociology*. 83, 6 (May 1978), 1420–1443. DOI:https://doi.org/10.1086/226707.

[27] Guillory, J.J. and Geraci, L. 2013. Correcting erroneous inferences in memory: The role of source credibility. *Journal of Applied Research in Memory and Cognition*. 2, 4 (Dec. 2013), 201–209. DOI:https://doi.org/10.1016/j.jarmac.2013.10.001.

[28] Halverson, J., Corman, S. and Goodall, H.L. 2011. *Master Narratives of Islamist Extremism*. Springer.

[29] HARRIS, S. 2001. Fragmented Narratives and Multiple Tellers: Witness and Defendant Accounts in Trials. *Discourse Studies*. 3, 1 (Feb. 2001), 53–74. DOI:https://doi.org/10.1177/1461445601003001003.

[30] Hastie, R. and Pennington, N. 2000. Explanation-Based Decision Making. *Judgment and Decision Making: An Interdisciplinary Reader*. Cambridge University Press. 212–228.

[31] Horne, Z., Powell, D., Hummel, J.E. and Holyoak, K.J. 2015. Countering antivaccination attitudes. *Proceedings of the National Academy of Sciences*. 112, 33 (Aug. 2015), 10321–10324. DOI:https://doi.org/10.1073/pnas.1504019112.

[32] Introne, J., Gokce Yildirim, I., Iandoli, L., DeCook, J. and Elzeini, S. 2018. How People Weave Online Information Into Pseudoknowledge. *Social Media + Society*. 4, 3 (Jul. 2018). DOI:https://doi.org/10.1177/2056305118785639.

[33] Introne, J., Iandoli, L., DeCook, J., Yildirim, I.G. and Elzeini, S. 2017. The Collaborative Construction and Evolution of Pseudoknowledge in Online Conversations. *Proceedings of*

*the 8th 2017 International Conference on Social Media & Society* (Toronto, Canada, Jul. 2017).

[34] Johnson, H.M. and Seifert, C.M. 1994. Sources of the continued influence effect: When misinformation in memory affects later inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 20, 6 (1994), 1420–1436. DOI:https://doi.org/10.1037/0278-7393.20.6.1420.

[35] Johnson, N. and Mandler, J. 1980. A tale of two structures: Underlying and surface forms in stories. *Poetics*. 9, 1–3 (Jun. 1980), 51–86. DOI:https://doi.org/10.1016/0304-422X(80)90012-1.

[36] Johnson, R.E. 1970. Recall of prose as a function of the structural importance of the linguistic units. *Journal of Verbal Learning and Verbal Behavior*. 9, 1 (Feb. 1970), 12–20. DOI:https://doi.org/10.1016/S0022-5371(70)80003-2.

[37] Kahan, D.M., Braman, D., Cohen, G.L., Gastil, J. and Slovic, P. 2010. Who Fears the HPV Vaccine, Who Doesn't, and Why? An Experimental Study of the Mechanisms of Cultural Cognition. *Law and Human Behavior*. 34, 6 (Dec. 2010), 501–516. DOI:https://doi.org/10.1007/s10979-009-9201-0.

[38] Kahan, D.M., Jenkins-Smith, H. and Braman, D. 2011. Cultural cognition of scientific consensus. *Journal of Risk Research*. 14, 2 (Feb. 2011), 147–174. DOI:https://doi.org/10.1080/13669877.2010.511246.

[39] Kahan, D.M., Slovic, P., Braman, D. and Gastil, J. 2006. Fear of Democracy: A Cultural Evaluation of Sunstein on Risk. *Harvard Law Review*. 119, 4 (2006), 1071–1109.

[40] Kahneman, D., Slovic, P. and Tversky, A. 1982. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press.

[41] Karampourniotis, P.D., Sreenivasan, S., Szymanski, B.K. and Korniss, G. 2015. The Impact of Heterogeneous Thresholds on Social Contagion with Multiple Initiators. *PLOS ONE*. 10, 11 (Nov. 2015), e0143020. DOI:https://doi.org/10.1371/journal.pone.0143020.

[42] Kata, A. 2010. A postmodern Pandora's box: Anti-vaccination misinformation on the Internet. *Vaccine*. 28, 7 (Feb. 2010), 1709–1716. DOI:https://doi.org/10.1016/j.vaccine.2009.12.022.

[43] Kata, A. 2012. Anti-vaccine activists, Web 2.0, and the postmodern paradigm – An overview of tactics and tropes used online by the anti-vaccination movement. *Vaccine*. 30, 25 (May 2012), 3778–3789. DOI:https://doi.org/10.1016/j.vaccine.2011.11.112.

[44] Kauffman, S.A. 1993. *The origins of order: Self organization and selection in evolution*. Oxford University Press, USA.

[45] Kearns, M., Suri, S. and Montfort, N. 2006. An Experimental Study of the Coloring Problem on Human Subject Networks. *Science*. 313, 5788 (Aug. 2006), 824–827. DOI:https://doi.org/10.1126/science.1127207.

[46] Kintsch, W. and Greene, E. 1978. The role of culture-specific schemata in the comprehension and recall of stories. *Discourse Processes*. 1, 1 (Jan. 1978), 1–13. DOI:https://doi.org/10.1080/01638537809544425.

[47] Kitta, A. and Goldberg, D.S. 2017. The significance of folklore for vaccine policy: discarding the deficit model. *Critical Public Health*. 27, 4 (Aug. 2017), 506–514. DOI:https://doi.org/10.1080/09581596.2016.1235259.

[48] Klayman, J. and Ha, Y. 1987. Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*. 94, 2 (1987), 211–228. DOI:https://doi.org/10.1037/0033-295X.94.2.211.

[49] Kosko, B. 1986. Fuzzy cognitive maps. *International Journal of Man-Machine Studies*. 24, 1 (1986), 65–75.

[50] Labov, W. and Waletzky, J. 2003. Narrative analysis: Oral versions of personal experience. *Essays on the Verbal and Visual Arts*. J. Helm, ed. University of Washington Press Seattle. 12–44.

[51] Lerman, K. and Ghosh, R. 2010. Information Contagion: An Empirical Study of the Spread of News on Digg and Twitter Social Networks. *ICWSM*. 10, (2010), 90–97.

[52] Lewandowsky, S., Ecker, U.K.H. and Cook, J. 2017. Beyond Misinformation: Understanding and Coping with the "Post-Truth" Era. *Journal of Applied Research in Memory and Cognition*. 6, 4 (Dec. 2017), 353–369. DOI:https://doi.org/10.1016/j.jarmac.2017.07.008.

[53] Lewandowsky, S., Ecker, U.K.H., Seifert, C.M., Schwarz, N. and Cook, J. 2012. Misinformation and Its Correction Continued Influence and Successful Debiasing. *Psychological Science in the Public Interest*. 13, 3 (Dec. 2012), 106–131. DOI:https://doi.org/10.1177/1529100612451018.

[54] Lewandowsky, S., Gignac, G.E. and Oberauer, K. 2013. The Role of Conspiracist Ideation and Worldviews in Predicting Rejection of Science. *PLOS ONE*. 8, 10 (Oct. 2013), e75637. DOI:https://doi.org/10.1371/journal.pone.0075637.

[55] Ma, J. and Stahl, L. 2017. A multimodal critical discourse analysis of anti-vaccination information on Facebook. *Library & Information Science Research*. 39, 4 (Oct. 2017), 303–310. DOI:https://doi.org/10.1016/j.lisr.2017.11.005.

[56] Mandler, J. and Johnson, N. 1977. Remembrance of things parsed: Story structure and recall. *Cognitive Psychology*. 9, 1 (Jan. 1977), 111–151. DOI:https://doi.org/10.1016/0010-0285(77)90006-8.

[57] Masaryk, R. and Hatoková, M. 2017. Qualitative inquiry into reasons why vaccination messages fail. *Journal of Health Psychology*. 22, 14 (Dec. 2017), 1880–1888. DOI:https://doi.org/10.1177/1359105316656770.

[58] Mason, W. and Watts, D.J. 2012. Collaborative learning in networks. *Proceedings of the National Academy of Sciences*. 109, 3 (Jan. 2012), 764–769. DOI:https://doi.org/10.1073/pnas.1110069108.

[59] Meade, M.L. and Roediger, H.L. 2002. Explorations in the social contagion of memory. *Memory & Cognition*. 30, 7 (Oct. 2002), 995–1009. DOI:https://doi.org/10.3758/BF03194318.

[60] Messing, S. and Westwood, S.J. 2014. Selective Exposure in the Age of Social Media: Endorsements Trump Partisan Source Affiliation When Selecting News Online. *Communication Research*. 41, 8 (Dec. 2014), 1042–1063. DOI:https://doi.org/10.1177/0093650212466406.

[61] Metzger, M.J., Flanagin, A.J. and Medders, R.B. 2010. Social and Heuristic Approaches to Credibility Evaluation Online. *Journal of Communication*. 60, 3 (Sep. 2010), 413–439. DOI:https://doi.org/10.1111/j.1460-2466.2010.01488.x.

[62] Nicholson, M.S. and Leask, J. 2012. Lessons from an online debate about measles–mumps–rubella (MMR) immunization. *Vaccine*. 30, 25 (May 2012), 3806–3812. DOI:https://doi.org/10.1016/j.vaccine.2011.10.072.

[63] Nyhan, B. and Reifler, J. 2015. Displacing Misinformation about Events: An Experimental Test of Causal Corrections. *Journal of Experimental Political Science*. 2, 1 (Apr. 2015), 81–93. DOI:https://doi.org/10.1017/XPS.2014.22.

[64] Nyhan, B. and Reifler, J. 2015. Does correcting myths about the flu vaccine work? An experimental evaluation of the effects of corrective information. *Vaccine*. 33, 3 (Jan. 2015), 459–464. DOI:https://doi.org/10.1016/j.vaccine.2014.11.017.

[65] Nyhan, B. and Reifler, J. 2012. Misinformation and Fact-checking. *Research Findings*. (2012).

[66] Nyhan, B. and Reifler, J. 2010. When Corrections Fail: The Persistence of Political Misperceptions. *Political Behavior*. 32, 2 (2010), 303–330.

[67] Nyhan, B., Reifler, J. and Richey, S. 2012. The Role of Social Networks in Influenza Vaccine Attitudes and Intentions Among College Students in the Southeastern United States. *Journal of Adolescent Health*. 51, 3 (Sep. 2012), 302–304. DOI:https://doi.org/10.1016/j.jadohealth.2012.02.014.

[68] Nyhan, B., Reifler, J., Richey, S. and Freed, G.L. 2014. Effective messages in vaccine promotion: a randomized trial. *Pediatrics*. 133, 4 (2014), 1–8.

[69] Polletta, F. 1998. Contending Stories: Narrative in Social Movements. *Qualitative Sociology*. 21, 4 (Dec. 1998), 419–446. DOI:https://doi.org/10.1023/A:1023332410633.

[70] Polletta, F., Chen, P.C.B., Gardner, B.G. and Motes, A. 2011. The Sociology of Storytelling. *Annual Review of Sociology*. 37, 1 (Jul. 2011), 109–130. DOI:https://doi.org/10.1146/annurev-soc-081309-150106.

[71] Radzikowski, J., Stefanidis, A., Jacobsen, K.H., Croitoru, A., Crooks, A. and Delamater, P.L. 2016. The Measles Vaccination Narrative in Twitter: A Quantitative Analysis. *JMIR Public Health and Surveillance*. 2, 1 (Jan. 2016). DOI:https://doi.org/10.2196/publichealth.5059.

[72] Reich, J.A. 2014. Neoliberal Mothering and Vaccine Refusal: Imagined Gated Communities and the Privilege of Choice. *Gender & Society*. 28, 5 (Oct. 2014), 679–704. DOI:https://doi.org/10.1177/0891243214532711.

[73] Romero, D.M., Meeder, B. and Kleinberg, J. 2011. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. *Proceedings of the 20th international conference on World wide web* (New York, NY, USA, 2011), 695–704.

[74] Sclater, S. 2017. *The Uses of Narrative: Explorations in Sociology, Psychology and Cultural Studies*. Routledge.

[75] Shelby, A. and Ernst, K. 2013. Story and science. *Human Vaccines & Immunotherapeutics*. 9, 8 (Aug. 2013), 1795–1801. DOI:https://doi.org/10.4161/hv.24828.

[76] Shifman, L. 2012. An anatomy of a YouTube meme. *New Media & Society*. 14, 2 (Mar. 2012), 187–203. DOI:https://doi.org/10.1177/1461444811412160.

[77] Sobo, E.J. 2015. Social Cultivation of Vaccine Refusal and Delay among Waldorf (Steiner) School Parents. *Medical Anthropology Quarterly*. 29, 3 (Sep. 2015), 381–399. DOI:https://doi.org/10.1111/maq.12214.

[78] Somers, M.R. 1994. The narrative constitution of identity: A relational and network approach. *Theory and society*. 23, 5 (1994), 605–649.

[79] Song, M.Y.-J. and Gruzd, A. 2017. Examining Sentiments and Popularity of Pro- and Anti-Vaccination Videos on YouTube. *Proceedings of the 8th International Conference on Social Media & Society* (New York, NY, USA, 2017), 17:1–17:8.

[80] Stein, N. and Albro, E. 1997. Building Complexity and Coherence: Children's Use of Goal-Structured Knowledge in Telling Stories. *Narrative Development: Six Approaches*. M.G.W. Bamberg, ed. Psychology Press.

[81] Stein, N. and Glenn, C. 1979. An analysis of story comprehension in elementary school children. *New directions in discourse processing*. 2, (1979), 53–120.

[82] Swire, B., Ecker, U.K.H. and Lewandowsky, S. 2017. The role of familiarity in correcting inaccurate information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 43, 12 (Dec. 2017), 1948–1961. DOI:https://doi.org/10.1037/xlm0000422.

[83] Trabasso, T. and van den Broek, P. 1985. Causal thinking and the representation of narrative events. *Journal of Memory and Language*. 24, 5 (Oct. 1985), 612–630. DOI:https://doi.org/10.1016/0749-596X(85)90049-X.

[84] Trabasso, T., Secco, T. and van den Broek, P. 1982. Causal cohesion and story coherence. *Learning and comprehension of text*. H. Mandl, N. Stein, and P. van den Broek, eds. Lawrence Erlbaum Associates. 83–112.

[85] Trabasso, T. and Sperry, L.L. 1985. Causal relatedness and importance of story events. *Journal of Memory and Language*. 24, 5 (Oct. 1985), 595–611. DOI:https://doi.org/10.1016/0749-596X(85)90048-8.

[86] Trevors, G.J., Muis, K.R., Pekrun, R., Sinatra, G.M. and Winne, P.H. 2016. Identity and Epistemic Emotions During Knowledge Revision: A Potential Account for the Backfire Effect. *Discourse Processes*. 53, 5–6 (Jul. 2016), 339–370. DOI:https://doi.org/10.1080/0163853X.2015.1136507.

[87] Wang, Z., Walther, J.B., Pingree, S. and Hawkins, R.P. 2008. Health Information, Credibility, Homophily, and Influence via the Internet: Web Sites Versus Discussion Groups. *Health Communication*. 23, 4 (Aug. 2008), 358–368. DOI:https://doi.org/10.1080/10410230802229738.

[88] Watts, D.J. 2002. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*. 99, 9 (Apr. 2002), 5766–5771. DOI:https://doi.org/10.1073/pnas.082090499.

[89] Watts, D.J. and Dodds, P.S. 2007. Influentials, Networks, and Public Opinion Formation. *Journal of Consumer Research*. 34, 4 (Dec. 2007), 441–458. DOI:https://doi.org/10.1086/518527.

[90] Wood, W. 2000. Attitude Change: Persuasion and Social Influence. *Annual Review of Psychology*. 51, 1 (2000), 539–570. DOI:https://doi.org/10.1146/annurev.psych.51.1.539.

[91] Wright, K. 2000. Computer-mediated social support, older adults, and coping. *Journal of Communication*. 50, 3 (Sep. 2000), 100–118. DOI:https://doi.org/10.1111/j.1460-2466.2000.tb02855.x.

[92] Zimbardo, P.G. and Leippe, M.R. 1991. *The psychology of attitude change and social influence.* Mcgraw-Hill Book Company.

[93] Zollo, F., Bessi, A., Del Vicario, M., Scala, A., Caldarelli, G., Shekhtman, L., Havlin, S. and Quattrociocchi, W. 2017. Debunking in a world of tribes. *PLOS ONE*. 12, 7 (Jul. 2017), e0181821. DOI:https://doi.org/10.1371/journal.pone.0181821.