The BOISI CENTER for

RELIGION and AMERICAN PUBLIC LIFE

Graduate Symposium on Religion and Politics

2019-2020

DEEP STORIES:

NARRATIVE'S ROLE IN AMERICAN

RELIGION AND POLITICS

Reading Packet 2

Boisi Center 2018-2019 Graduate Symposium on Religion and Politics

DEEP STORIES:

NARRATIVE'S ROLE IN AMERICAN RELIGION AND POLITICS

Table of Contents

The Grim Conclusions of the Largest-Ever Study of Fake News

Falsehoods almost always beat out the truth on Twitter, penetrating further, faster, and deeper into the social network than accurate information.

ROBINSON MEYER   **MARCH 8, 2018**

KRISTA KENNELL / STONE / CATWALKER / SHUTTERSTOCK / THE ATLANTIC

"Falsehood flies, and the Truth comes limping after it," Jonathan Swift once wrote.

It was hyperbole three centuries ago. But it is a factual description of social media, according to an ambitious and first-of-its-kind study published Thursday in *Science*.

The massive new study analyzes every major contested news story in English across the span of Twitter's existence—some 126,000 stories, tweeted by 3 million users, over more than 10 years—and finds that the truth simply cannot compete with hoax and rumor. By every common metric, falsehood consistently dominates the truth on Twitter, the study finds: Fake news and false rumors reach more people, penetrate deeper into the social network, and spread much faster than accurate stories.

"It seems to be pretty clear [from our study] that false information outperforms true information," said Soroush Vosoughi, a data scientist at MIT who has studied fake news since 2013 and who led this study. "And that is not just because of bots. It might have something to do with human nature."

The study has already prompted alarm from social scientists. "We must redesign our information ecosystem in the 21st century," write a group of 16 political scientists and legal scholars in an essay also published Thursday in *Science*. They call for a new drive of interdisciplinary research "to reduce the spread of fake news and to address the

underlying pathologies it has revealed."

"How can we create a news ecosystem … that values and promotes truth?" they ask.

The new study suggests that it will not be easy. Though Vosoughi and his colleagues only focus on Twitter—the study was conducted using exclusive data that the company made available to MIT—their work has implications for Facebook, YouTube, and every major social network. Any platform that regularly amplifies engaging or provocative content runs the risk of amplifying fake news along with it.

Though the study is written in the clinical language of statistics, it offers a methodical indictment of the accuracy of information that spreads on these platforms. A false story is much more likely to go viral than a real story, the authors find. A false story reaches 1,500 people six times quicker, on average, than a true story does. And while false stories outperform the truth on every subject—including business, terrorism and war, science and technology, and entertainment—fake news about politics regularly does best.

Twitter users seem almost to *prefer* sharing falsehoods. Even when the researchers controlled for every difference between the accounts originating rumors—like whether that person had more followers or was verified—falsehoods were still 70 percent more likely to get retweeted than accurate news.

And blame for this problem cannot be laid with our robotic brethren. From 2006 to 2016, Twitter bots amplified true stories as much as they

amplified false ones, the study found. Fake news prospers, the authors write, "because humans, not robots, are more likely to spread it."

Political scientists and social-media researchers largely praised the study, saying it gave the broadest and most rigorous look so far into the scale of the fake-news problem on social networks, though some disputed its findings about bots and questioned its definition of news.

"This is a really interesting and impressive study, and the results around how demonstrably untrue assertions spread faster and wider than demonstrable true ones do, within the sample, seem very robust, consistent, and well supported," said Rasmus Kleis Nielsen, a professor of political communication at the University of Oxford, in an email.

"I think it's very careful, important work," Brendan Nyhan, a professor of government at Dartmouth College, told me. "It's excellent research of the sort that we need more of."

"In short, I don't think there's any reason to doubt the study's results," said Rebekah Tromble, a professor of political science at Leiden University in the Netherlands, in an email.

What makes this study different? In the past, researchers have looked into the problem of falsehoods spreading online. They've often focused on rumors around singular events, like the speculation that preceded the discovery of the Higgs boson in 2012 or the rumors that followed the Haiti earthquake in 2010.

This new paper takes a far grander scale, looking at nearly the entire

lifespan of Twitter: every piece of controversial news that propagated on the service from September 2006 to December 2016. But to do that, Vosoughi and his colleagues had to answer a more preliminary question first: *What is truth? And how do we know?*

————————

It's a question that can have life-or-death consequences.

"[Fake news] has become a white-hot political and, really, cultural topic, but the trigger for us was personal events that hit Boston five years ago," said Deb Roy, a media scientist at MIT and one of the authors of the new study.

On April 15, 2013, two bombs exploded near the route of the Boston Marathon, killing three people and injuring hundreds more. Almost immediately, wild conspiracy theories about the bombings took over Twitter and other social-media platforms. The mess of information only grew more intense on April 19, when the governor of Massachusetts asked millions of people to remain in their homes as police conducted a huge manhunt.

"I was on lockdown with my wife and kids in our house in Belmont for two days, and Soroush was on lockdown in Cambridge," Roy told me. Stuck inside, Twitter became their lifeline to the outside world. "We heard a lot of things that were not true, and we heard a lot of things that did turn out to be true" using the service, he said.

The ordeal soon ended. But when the two men reunited on campus,

they agreed it seemed seemed silly for Vosoughi—then a Ph.D. student focused on social media—to research anything but what they had just lived through. Roy, his adviser, blessed the project.

He made a truth machine: an algorithm that could sort through torrents of tweets and pull out the facts most likely to be accurate from them. It focused on three attributes of a given tweet: the properties of its author (were they verified?), the kind of language it used (was it sophisticated?), and how a given tweet propagated through the network.

"The model that Soroush developed was able to predict accuracy with a far-above-chance performance," said Roy. He earned his Ph.D. in 2015.

After that, the two men—and Sinan Aral, a professor of management at MIT—turned to examining how falsehoods move across Twitter as a whole. But they were back not only at the "what is truth?" question, but its more pertinent twin: How does *the computer* know what truth is?

They opted to turn to the ultimate arbiter of fact online: the third-party fact-checking sites. By scraping and analyzing six different fact-checking sites—including *Snopes*, *Politifact*, and FactCheck.org—they generated a list of tens of thousands of online rumors that had spread between 2006 and 2016 on Twitter. Then they searched Twitter for these rumors, using a proprietary search engine owned by the social network called Gnip.

Ultimately, they found about 126,000 tweets, which, together, had

been retweeted more than 4.5 million times. Some linked to "fake" stories hosted on other websites. Some started rumors themselves, either in the text of a tweet or in an attached image. (The team used a special program that could search for words contained within static tweet images.) And some contained true information or linked to it elsewhere.

Then they ran a series of analyses, comparing the popularity of the fake rumors with the popularity of the real news. What they found astounded them.

Speaking from MIT this week, Vosoughi gave me an example: There are lots of ways for a tweet to get 10,000 retweets, he said. If a celebrity sends Tweet A, and they have a couple million followers, maybe 10,000 people will see Tweet A in their timeline and decide to retweet it. Tweet A was broadcast, creating a big but shallow pattern.

Meanwhile, someone without many followers sends Tweet B. It goes out to their 20 followers—but one of those people sees it, and retweets it, and then one of *their* followers sees it and retweets it too, on and on until tens of thousands of people have seen and shared Tweet B.

Tweet A and Tweet B both have the same size audience, but Tweet B has more "depth," to use Vosoughi's term. It chained together retweets, going viral in a way that Tweet A never did. "It could reach 1,000 retweets, but it has a very different shape," he said.

Here's the thing: Fake news dominates *according to both metrics*. It consistently reaches a larger audience, *and* it tunnels much deeper into

social networks than real news does. The authors found that accurate news wasn't able to chain together more than 10 retweets. Fake news could put together a retweet chain 19 links long—and do it 10 times as fast as accurate news put together its measly 10 retweets.

These results proved robust even when they were checked by humans, not bots. Separate from the main inquiry, a group of undergraduate students fact-checked a random selection of roughly 13,000 English-language tweets from the same period. They found that false information outperformed true information in ways "nearly identical" to the main data set, according to the study.

What does this look like in real life? Take two examples from the last presidential election. In August 2015, a rumor circulated on social media that Donald Trump had let a sick child use his plane to get urgent medical care. *Snopes* confirmed almost all of the tale as true. But according to the team's estimates, only about 1,300 people shared or retweeted the story.

In February 2016, a rumor developed that Trump's elderly cousin had recently died and that he had opposed the magnate's presidential bid in his obituary. "As a proud bearer of the Trump name, I implore you all, please don't let that walking mucus bag become president," the obituary reportedly said. But *Snopes* could not find evidence of the cousin, or his obituary, and rejected the story as false.

Nonetheless, roughly 38,000 Twitter users shared the story. And it put together a retweet chain three times as long as the sick-child story

managed.

A false story alleging the boxer Floyd Mayweather had worn a Muslim head scarf to a Trump rally also reached an audience more than 10 times the size of the sick-child story.

Why does falsehood do so well? The MIT team settled on two hypotheses.

First, fake news seems to be more "novel" than real news. Falsehoods are often notably different from the all the tweets that have appeared in a user's timeline 60 days prior to their retweeting them, the team found.

Second, fake news evokes much more emotion than the average tweet. The researchers created a database of the words that Twitter users used to reply to the 126,000 contested tweets, then analyzed it with a state-of-the-art sentiment-analysis tool. Fake tweets tended to elicit words associated with surprise and disgust, while accurate tweets summoned words associated with sadness and trust, they found.

———

The team wanted to answer one more question: Were Twitter bots helping to spread misinformation?

After using two different bot-detection algorithms on their sample of 3 million Twitter users, they found that the automated bots were spreading false news—but they were retweeting it at the same rate that they retweeted accurate information.[9]

"The massive differences in how true and false news spreads on Twitter cannot be explained by the presence of bots," Aral told me.

But some political scientists cautioned that this should not be used to disprove the role of Russian bots in seeding disinformation recently. An "army" of Russian-associated bots helped amplify divisive rhetoric after the school shooting in Parkland, Florida, _The New York Times_ has reported.

"It can both be the case that (1) over the whole 10-year data set, bots don't favor false propaganda and (2) in a recent subset of cases, botnets have been strategically deployed to spread the reach of false propaganda claims," said Dave Karpf, a political scientist at George Washington University, in an email.

"My guess is that the paper is going to get picked up as 'scientific proof that bots don't really matter!' And this paper does indeed show that, if we're looking at the full life span of Twitter. But the real bots debate assumes that their usage has recently escalated because strategic actors have poured resources into their use. This paper doesn't refute that assumption," he said.

Vosoughi agrees that his paper does not determine whether the use of botnets changed around the 2016 election. "We did not study the change in the role of bots across time," he told me in an email. "This is an interesting question and one that we will probably look at in future work."

Some political scientists also questioned the study's definition of

"news." By turning to the fact-checking sites, the study blurs together a wide range of false information: outright lies, urban legends, hoaxes, spoofs, falsehoods, *and* "fake news." It does not just look at fake news by itself—that is, articles or videos that look like news content, and which appear to have gone through a journalistic process, but which are actually made up.

Therefore, the study may undercount "non-contested news": accurate news that is widely understood to be true. For many years, the most retweeted post in Twitter's history celebrated Obama's re-election as president. But as his victory was not a widely disputed fact, *Snopes* and other fact-checking sites never confirmed it.

The study also elides *content* and *news.* "All our audience research suggests a vast majority of users see news as clearly distinct from content more broadly," Nielsen, the Oxford professor, said in an email. "Saying that untrue content, including rumors, spread faster than true statements on Twitter is a bit different from saying false news and true news spread at different rates."

But many researchers told me that simply understanding *why* false rumors travel so far, so fast, was as important as knowing that they do so in the first place.

"The key takeaway is really that content that *arouses strong emotions* spreads further, faster, more deeply, and more broadly on Twitter," said Tromble, the political scientist, in an email. "This particular finding is consistent with research in a number of different areas, including

psychology and communication studies. It's also relatively intuitive."

"False information online is often really novel and frequently negative," said Nyhan, the Dartmouth professor. "We know those are two features of information generally that grab our attention as human beings and that cause us to want to share that information with others—we're attentive to novel threats and especially attentive to negative threats."

"It's all too easy to create both when you're not bound by the limitations of reality. So people can exploit the interaction of human psychology and the design of these networks in powerful ways," he added.

He lauded Twitter for making its data available to researchers and called on other major platforms, like Facebook, to do the same. "In terms of research, the platforms are the whole ballgame. We have so much to learn but we're so constrained in what we can study without platform partnership and collaboration," he said.

"These companies now exercise a great deal of power and influence over the news that people get in our democracy. The amount of power that platforms now hold means they have to face a great deal of scrutiny and transparency," he said. "We can study Twitter all day, but only about 12 percent of Americans are on it. It's important for journalists and academics, but it's not how most people get their news."

In a statement, Twitter said that it was hoping to expand its work with outside experts. In a series of tweets last week, Jack Dorsey, the company's CEO, said the company[1] hoped to "increase the collective

health, openness, and civility of public conversation, and to hold ourselves publicly accountable toward progress."

Facebook did not respond to a request for comment.

But Tromble, the political-science professor, said that the findings would likely apply to Facebook, too. "Earlier this year, Facebook announced that it would restructure its News Feed to favor 'meaningful interaction,'" she told me.

"It became clear that they would gauge 'meaningful interaction' based on the number of comments and replies to comments a post receives. But, as this study shows, that only further incentivizes creating posts full of disinformation and other content likely to garner strong emotional reactions," she added.

"Putting my conservative scientist hat on, I'm not comfortable saying how this applies to other social networks. We only studied Twitter here," said Aral, one of the researchers. "But my intuition is that these findings are broadly applicable to social-media platforms in general. You could run this exact same study if you worked with Facebook's data."

———————

Yet these do not encompass the most depressing finding of the study. When they began their research, the MIT team expected that users who shared the most fake news would basically be crowd-pleasers. They assumed they would find a group of people who obsessively use Twitter

in a partisan or sensationalist way, accumulating more fans and followers than their more fact-based peers.

In fact, the team found that the opposite is true. Users who share accurate information have more followers, and send more tweets, than fake-news sharers. These fact-guided users have also been on Twitter for longer, and they are more likely to be verified. In short, the most trustworthy users can boast every obvious structural advantage that Twitter, either as a company or a community, can bestow on its best users.

The truth has a running start, in other words—but inaccuracies, somehow, still win the race. "Falsehood diffused further and faster than the truth *despite* these differences [between accounts], not because of them," write the authors.

This finding should dispirit every user who turns to social media to find or distribute accurate information. It suggests that no matter how adroitly people *plan* to use Twitter—no matter how meticulously they curate their feed or follow reliable sources—they can still get snookered by a falsehood in the heat of the moment.

It suggests—to me, at least, a Twitter user since 2007, and someone who got his start in journalism because of the social network—that social-media platforms do not encourage the kind of behavior that anchors a democratic government. On platforms where every user is at once a reader, a writer, and a publisher, falsehoods are too seductive not to succeed: The thrill of novelty is too alluring, the titillation of disgust too difficult to transcend. After a long and aggravating day, even the

most staid user might find themselves lunging for the politically advantageous rumor. Amid an anxious election season, even the most public-minded user might subvert their higher interest to win an argument.

It is unclear which interventions, if any, could reverse this tendency toward falsehood. "We don't know enough to say what works and what doesn't," Aral told me. There is little evidence that people change their opinion because they see a fact-checking site reject one of their beliefs, for instance. Labeling fake news as such, on a social network or search engine, may do little to deter it as well.

In short, social media seems to systematically amplify falsehood at the expense of the truth, and no one—neither experts nor politicians nor tech companies—knows how to reverse that trend. It is a dangerous moment for any system of government premised on a common public reality.

———

We want to hear what you think about this article. Submit a letter to the editor or write to letters@theatlantic.com.

SOCIAL SCIENCE

The spread of true and false news online

Soroush Vosoughi,[1] Deb Roy,[1] Sinan Aral[2]*

We investigated the differential diffusion of all of the verified true and false news stories distributed on Twitter from 2006 to 2017. The data comprise ~126,000 stories tweeted by ~3 million people more than 4.5 million times. We classified news as true or false using information from six independent fact-checking organizations that exhibited 95 to 98% agreement on the classifications. Falsehood diffused significantly farther, faster, deeper, and more broadly than the truth in all categories of information, and the effects were more pronounced for false political news than for false news about terrorism, natural disasters, science, urban legends, or financial information. We found that false news was more novel than true news, which suggests that people were more likely to share novel information. Whereas false stories inspired fear, disgust, and surprise in replies, true stories inspired anticipation, sadness, joy, and trust. Contrary to conventional wisdom, robots accelerated the spread of true and false news at the same rate, implying that false news spreads more than the truth because humans, not robots, are more likely to spread it.

Foundational theories of decision-making (*1*–*3*), cooperation (*4*), communication (*5*), and markets (*6*) all view some conceptualization of truth or accuracy as central to the functioning of nearly every human endeavor. Yet, both true and false information spreads rapidly through online media. Defining what is true and false has become a common political strategy, replacing debates based on a mutually agreed on set of facts. Our economies are not immune to the spread of falsity either. False rumors have affected stock prices and the motivation for large-scale investments, for example, wiping out $130 billion in stock value after a false tweet claimed that Barack Obama was injured in an explosion (*7*). Indeed, our responses to everything from natural disasters (*8*, *9*) to terrorist attacks (*10*) have been disrupted by the spread of false news online.

New social technologies, which facilitate rapid information sharing and large-scale information cascades, can enable the spread of misinformation (i.e., information that is inaccurate or misleading). But although more and more of our access to information and news is guided by these new technologies (*11*), we know little about their contribution to the spread of falsity online. Though considerable attention has been paid to anecdotal analyses of the spread of false news by the media (*12*), there are few large-scale empirical investigations of the diffusion of misinformation or its social origins. Studies of the spread of misinformation are currently limited to analyses of small, ad hoc samples that ignore two of the most important scientific questions: How do truth and falsity diffuse differently, and what factors of human judgment explain these differences?

Current work analyzes the spread of single rumors, like the discovery of the Higgs boson (*13*) or the Haitian earthquake of 2010 (*14*), and multiple rumors from a single disaster event, like the Boston Marathon bombing of 2013 (*10*), or it develops theoretical models of rumor diffusion (*15*), methods for rumor detection (*16*), credibility evaluation (*17*, *18*), or interventions to curtail the spread of rumors (*19*). But almost no studies comprehensively evaluate differences in the spread of truth and falsity across topics or examine why false news may spread differently than the truth. For example, although Del Vicario *et al.* (*20*) and Bessi *et al.* (*21*) studied the spread of scientific and conspiracy-theory stories, they did not evaluate their veracity. Scientific and conspiracy-theory stories can both be either true or false, and they differ on stylistic dimensions that are important to their spread but orthogonal to their veracity. To understand the spread of false news, it is necessary to examine diffusion after differentiating true and false scientific stories and true and false conspiracy-theory stories and controlling for the topical and stylistic differences between the categories themselves. The only study to date that segments rumors by veracity is that of Friggeri *et al.* (*19*), who analyzed ~4000 rumors spreading on Facebook and focused more on how fact checking affects rumor propagation than on how falsity diffuses differently than the truth (*22*).

In our current political climate and in the academic literature, a fluid terminology has arisen around "fake news," foreign interventions in U.S. politics through social media, and our understanding of what constitutes news, fake news, false news, rumors, rumor cascades, and other related terms. Although, at one time, it may have been appropriate to think of fake news as referring to the veracity of a news story, we now believe that this phrase has been irredeemably polarized in our current political and media climate. As politicians have implemented a political strategy of labeling news sources that do not support their positions as unreliable or fake news, whereas sources that support their positions are labeled reliable or not fake, the term has lost all connection to the actual veracity of the information presented, rendering it meaningless for use in academic classification. We have therefore explicitly avoided the term fake news throughout this paper and instead use the more objectively verifiable terms "true" or "false" news. Although the terms fake news and misinformation also imply a willful distortion of the truth, we do not make any claims about the intent of the purveyors of the information in our analyses. We instead focus our attention on veracity and stories that have been verified as true or false.

We also purposefully adopt a broad definition of the term news. Rather than defining what constitutes news on the basis of the institutional source of the assertions in a story, we refer to any asserted claim made on Twitter as news (we defend this decision in the supplementary materials section on "reliable sources," section S1.2). We define news as any story or claim with an assertion in it and a rumor as the social phenomena of a news story or claim spreading or diffusing through the Twitter network. That is, rumors are inherently social and involve the sharing of claims between people. News, on the other hand, is an assertion with claims, whether it is shared or not.

A rumor cascade begins on Twitter when a user makes an assertion about a topic in a tweet, which could include written text, photos, or links to articles online. Others then propagate the rumor by retweeting it. A rumor's diffusion process can be characterized as having one or more cascades, which we define as instances of a rumor-spreading pattern that exhibit an unbroken retweet chain with a common, singular origin. For example, an individual could start a rumor cascade by tweeting a story or claim with an assertion in it, and another individual could independently start a second cascade of the same rumor (pertaining to the same story or claim) that is completely independent of the first cascade, except that it pertains to the same story or claim. If they remain independent, they represent two cascades of the same rumor. Cascades can be as small as size one (meaning no one retweeted the original tweet). The number of cascades that make up a rumor is equal to the number of times the story or claim was independently tweeted by a user (not retweeted). So, if a rumor "A" is tweeted by 10 people separately, but not retweeted, it would have 10 cascades, each of size one. Conversely, if a second rumor "B" is independently tweeted by two people and each of those two tweets is retweeted 100 times, the rumor would consist of two cascades, each of size 100.

Here we investigate the differential diffusion of true, false, and mixed (partially true, partially false) news stories using a comprehensive data set of all of the fact-checked rumor cascades that spread on Twitter from its inception in 2006 to 2017. The data include ~126,000 rumor cascades spread by ~3 million people more than 4.5 million times. We sampled all rumor cascades investigated by six independent fact-checking organizations

[1]Massachusetts Institute of Technology (MIT), the Media Lab, E14-526, 75 Amherst Street, Cambridge, MA 02142, USA. [2]MIT, E62-364, 100 Main Street, Cambridge, MA 02142, USA.
*Corresponding author. Email: sinan@mit.edu

16

(snopes.com, politifact.com, factcheck.org, truthor-fiction.com, hoax-slayer.com, and urbanlegends. about.com) by parsing the title, body, and verdict (true, false, or mixed) of each rumor investigation reported on their websites and automatically collecting the cascades corresponding to those rumors on Twitter. The result was a sample of rumor cascades whose veracity had been agreed on by these organizations between 95 and 98% of the time. We cataloged the diffusion of the rumor cascades by collecting all English-language replies to tweets that contained a link to any of the aforementioned websites from 2006 to 2017 and used optical character recognition to extract text from images where needed. For each reply tweet, we extracted the original tweet being replied to and all the retweets of the original tweet. Each retweet cascade represents a rumor propagating on Twitter that has been verified as true or false by the fact-checking organizations (see the supplementary materials for more details on cascade construction). We then quantified the cascades'

depth (the number of retweet hops from the origin tweet over time, where a hop is a retweet by a new unique user), size (the number of users involved in the cascade over time), maximum breadth (the maximum number of users involved in the cascade at any depth), and structural virality (23) (a measure that interpolates between content spread through a single, large broadcast and that which spreads through multiple generations, with any one individual directly responsible for only a fraction of the total spread) (see the supplementary materials for more detail on the measurement of rumor diffusion).

As a rumor is retweeted, the depth, size, maximum breadth, and structural virality of the cascade increase (Fig. 1A). A greater fraction of false rumors experienced between 1 and 1000 cascades, whereas a greater fraction of true rumors experienced more than 1000 cascades (Fig. 1B); this was also true for rumors based on political news (Fig. 1D). The total number of false rumors peaked at the end of both 2013 and 2015 and again at the

end of 2016, corresponding to the last U.S. presidential election (Fig. 1C). The data also show clear increases in the total number of false political rumors during the 2012 and 2016 U.S. presidential elections (Fig. 1E) and a spike in rumors that contained partially true and partially false information during the Russian annexation of Crimea in 2014 (Fig. 1E). Politics was the largest rumor category in our data, with ~45,000 cascades, followed by urban legends, business, terrorism, science, entertainment, and natural disasters (Fig. 1F).

When we analyzed the diffusion dynamics of true and false rumors, we found that falsehood diffused significantly farther, faster, deeper, and more broadly than the truth in all categories of information [Kolmogorov-Smirnov (K-S) tests are reported in tables S3 to S10]. A significantly greater fraction of false cascades than true cascades exceeded a depth of 10, and the top 0.01% of false cascades diffused eight hops deeper into the Twittersphere than the truth, diffusing to depths
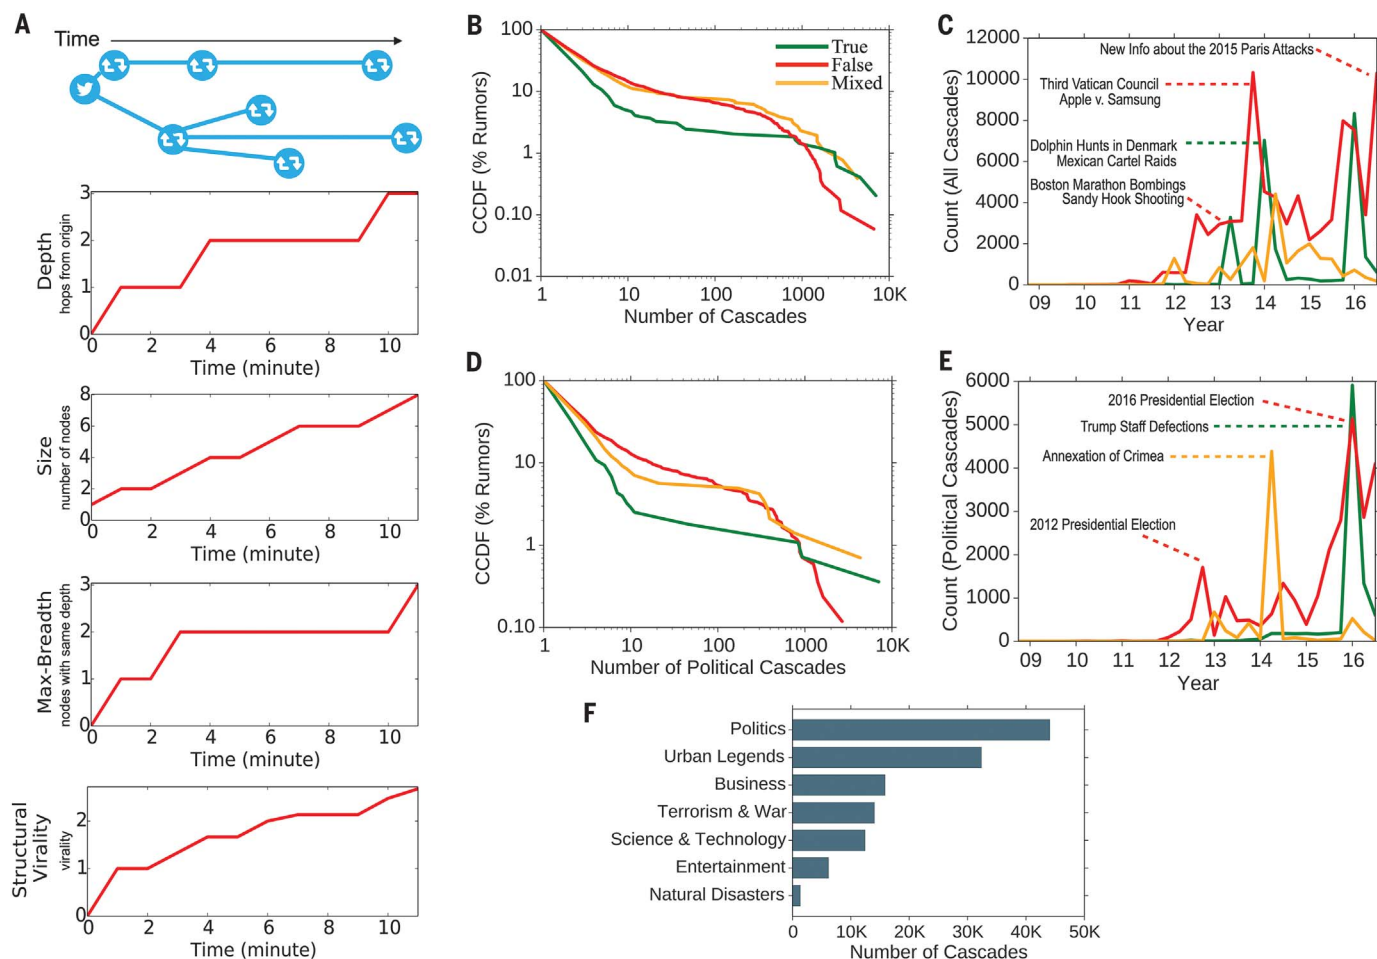
Fig. 1. Rumor cascades. (**A**) An example rumor cascade collected by our method as well as its depth, size, maximum breadth, and structural virality over time. "Nodes" are users. (**B**) The complementary cumulative distribution functions (CCDFs) of true, false, and mixed (partially true and partially false) cascades, measuring the fraction of rumors that exhibit a given number of cascades. (**C**) Quarterly counts of all true, false, and mixed rumor cascades that diffused on Twitter between 2006 and 2017, annotated with example rumors in each category. (**D**) The CCDFs of true, false, and mixed political cascades. (**E**) Quarterly counts of all true, false, and mixed political rumor cascades that diffused on Twitter between 2006 and 2017, annotated with example rumors in each category. (**F**) A histogram of the total number of rumor cascades in our data across the seven most frequent topical categories.

17
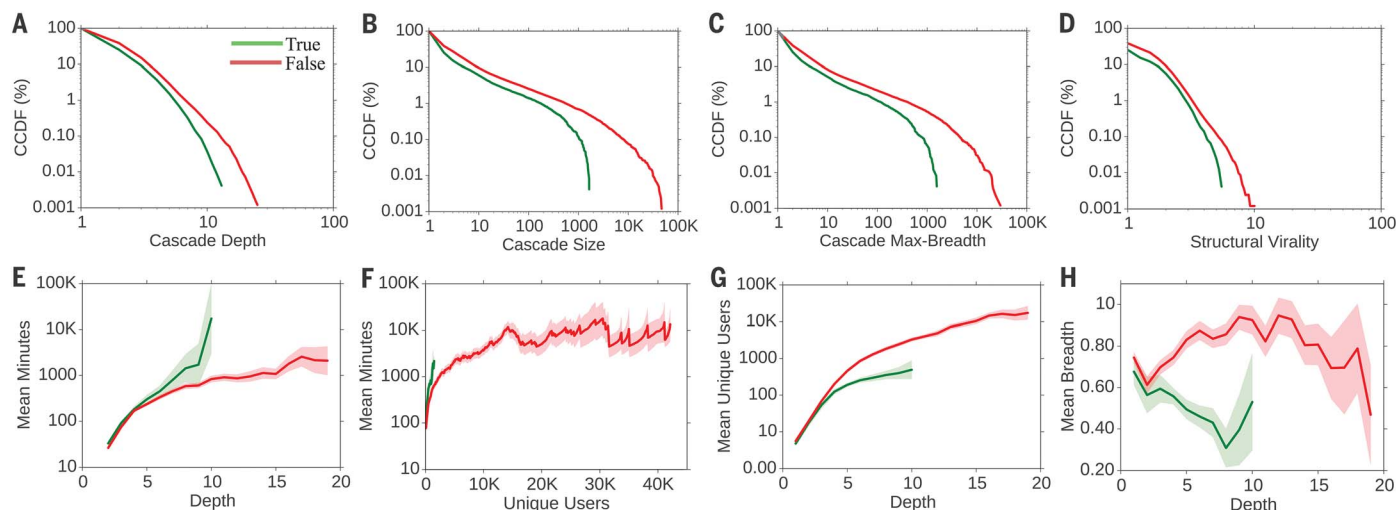
Fig. 2. Complementary cumulative distribution functions (CCDFs) of true and false rumor cascades. (**A**) Depth. (**B**) Size. (**C**) Maximum breadth. (**D**) Structural virality. (**E** and **F**) The number of minutes it takes for true and false rumor cascades to reach any (E) depth and (F) number of unique Twitter users. (**G**) The number of unique Twitter users reached at every depth and (**H**) the mean breadth of true and false rumor cascades at every depth. In (H), plot is lognormal. Standard errors were clustered at the rumor level (i.e., cascades belonging to the same rumor were clustered together; see supplementary materials for additional details).
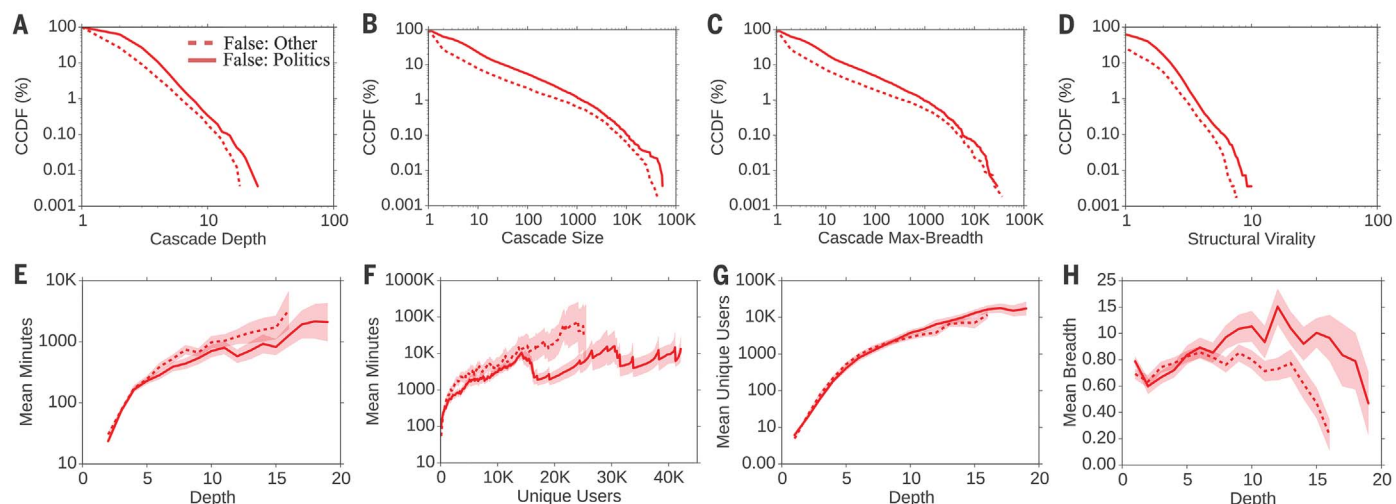
Fig. 3. Complementary cumulative distribution functions (CCDFs) of false political and other types of rumor cascades. (**A**) Depth. (**B**) Size. (**C**) Maximum breadth. (**D**) Structural virality. (**E** and **F**) The number of minutes it takes for false political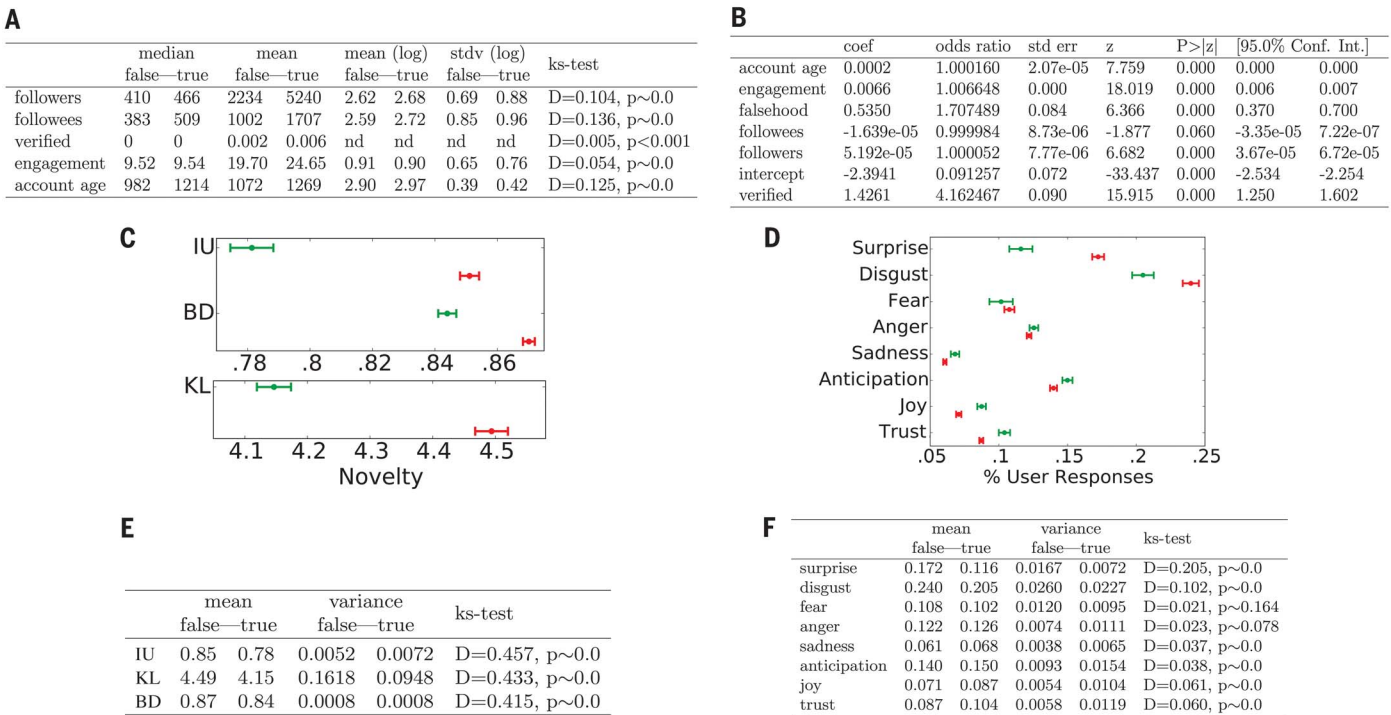 and other false news cascades to reach any (E) depth and (F) number of unique Twitter users. (**G**) The number of unique Twitter users reached at every depth and (**H**) the mean breadth of these false rumor cascades at every depth. In (H), plot is lognormal. Standard errors were clustered at the rumor level.

greater than 19 hops from the origin tweet (Fig. 2A). Falsehood also reached far more people than the truth. Whereas the truth rarely diffused to more than 1000 people, the top 1% of false-news cascades routinely diffused to between 1000 and 100,000 people (Fig. 2B). Falsehood reached more people at every depth of a cascade than the truth, meaning that many more people retweeted false-hood than they did the truth (Fig. 2C). The spread of falsehood was aided by its virality, meaning that falsehood did not simply spread through broadcast dynamics but rather through peer-to-peer diffusion characterized by a viral branching process (Fig. 2D).

It took the truth about six times as long as falsehood to reach 1500 people (Fig. 2F) and 20 times as long as falsehood to reach a cascade depth of 10 (Fig. 2E). As the truth never diffused beyond a depth of 10, we saw that falsehood reached a depth of 19 nearly 10 times faster than the truth reached a depth of 10 (Fig. 2E). Falsehood also diffused significantly more broadly (Fig. 2H) and was retweeted by more unique users than the truth at every cascade depth (Fig. 2G).

False political news (Fig. 1D) traveled deeper (Fig. 3A) and more broadly (Fig. 3C), reached more people (Fig. 3B), and was more viral than any other category of false information (Fig. 3D). False po-

litical news also diffused deeper more quickly (Fig. 3E) and reached more than 20,000 people nearly three times faster than all other types of false news reached 10,000 people (Fig. 3F). Although the other categories of false news reached about the same number of unique users at depths between 1 and 10, false political news routinely reached the most unique users at depths greater than 10 (Fig. 3G). Although all other categories of false news traveled slightly more broadly at shallower depths, false political news traveled more broadly at greater depths, indicating that more-popular false political news items exhibited broader and more-accelerated diffusion dynamics

18

A

	median		mean		mean (log)		stdv (log)		ks-test
	false	true	false	true	false	true	false	true	
followers	410	466	2234	5240	2.62	2.68	0.69	0.88	D=0.104, p∼0.0
followees	383	509	1002	1707	2.59	2.72	0.85	0.96	D=0.136, p∼0.0
verified	0	0	0.002	0.006	nd	nd	nd	nd	D=0.005, p<0.001
engagement	9.52	9.54	19.70	24.65	0.91	0.90	0.65	0.76	D=0.054, p∼0.0
account age	982	1214	1072	1269	2.90	2.97	0.39	0.42	D=0.125, p∼0.0

B

	coef	odds ratio	std err	z	P>|z|	[95.0% Conf. Int.]	
account age	0.0002	1.000160	2.07e-05	7.759	0.000	0.000	0.000
engagement	0.0066	1.006648	0.000	18.019	0.000	0.006	0.007
falsehood	0.5350	1.707489	0.084	6.366	0.000	0.370	0.700
followees	-1.639e-05	0.999984	8.73e-06	-1.877	0.060	-3.35e-05	7.22e-07
followers	5.192e-05	1.000052	7.77e-06	6.682	0.000	3.67e-05	6.72e-05
intercept	-2.3941	0.091257	0.072	-33.437	0.000	-2.534	-2.254
verified	1.4261	4.162467	0.090	15.915	0.000	1.250	1.602

C (IU, BD, KL plots — Novelty axis: .78, .8, .82, .84, .86 and 4.1, 4.2, 4.3, 4.4, 4.5)

D (Surprise, Disgust, Fear, Anger, Sadness, Anticipation, Joy, Trust — % User Responses axis: .05, .1, .15, .2, .25)

E

	mean		variance		ks-test
	false	true	false	true	
IU	0.85	0.78	0.0052	0.0072	D=0.457, p∼0.0
KL	4.49	4.15	0.1618	0.0948	D=0.433, p∼0.0
BD	0.87	0.84	0.0008	0.0008	D=0.415, p∼0.0

F

	mean		variance		ks-test
	false	true	false	true	
surprise	0.172	0.116	0.0167	0.0072	D=0.205, p∼0.0
disgust	0.240	0.205	0.0260	0.0227	D=0.102, p∼0.0
fear	0.108	0.102	0.0120	0.0095	D=0.021, p∼0.164
anger	0.122	0.126	0.0074	0.0111	D=0.023, p∼0.078
sadness	0.061	0.068	0.0038	0.0065	D=0.037, p∼0.0
anticipation	0.140	0.150	0.0093	0.0154	D=0.038, p∼0.0
joy	0.071	0.087	0.0054	0.0104	D=0.061, p∼0.0
trust	0.087	0.104	0.0058	0.0119	D=0.060, p∼0.0

Fig. 4. Models estimating correlates of news diffusion, the novelty of true and false news, and the emotional content of replies to news.
(**A**) Descriptive statistics on users who participated in true and false rumor cascades as well as K-S tests of the differences in the distributions of these measures across true and false rumor cascades. (**B**) Results of a logistic regression model estimating users' likelihood of retweeting a rumor as a function of variables shown at the left. coeff, logit coefficient; z, z score. (**C**) Differences in the information uniqueness (IU), scaled Bhattacharyya distance (BD), and K-L divergence (KL) of true (green) and false (red) rumor tweets compared to the corpus of prior tweets the user was exposed to in the 60 days before retweeting the rumor tweet. (**D**) The emotional content of replies to true (green) and false (red) rumor tweets across seven dimensions categorized by the NRC. (**E**) Mean and variance of the IU, KL, and BD of true and false rumor tweets compared to the corpus of prior tweets the user has seen in the 60 days before seeing the rumor tweet as well as K-S tests of their differences across true and false rumors. (**F**) Mean and variance of the emotional content of replies to true and false rumor tweets across seven dimensions categorized by the NRC as well as K-S tests of their differences across true and false rumors. All standard errors are clustered at the rumor level, and all models are estimated with cluster-robust standard errors at the rumor level.

(Fig. 3H). Analysis of all news categories showed that news about politics, urban legends, and science spread to the most people, whereas news about politics and urban legends spread the fastest and were the most viral in terms of their structural virality (see fig. S11 for detailed comparisons across all topics).

One might suspect that structural elements of the network or individual characteristics of the users involved in the cascades explain why falsity travels with greater velocity than the truth. Perhaps those who spread falsity "followed" more people, had more followers, tweeted more often, were more often "verified" users, or had been on Twitter longer. But when we compared users involved in true and false rumor cascades, we found that the opposite was true in every case. Users who spread false news had significantly fewer followers (K-S test = 0.104, P ∼ 0.0), followed significantly fewer people (K-S test = 0.136, P ∼ 0.0), were significantly less active on Twitter (K-S test = 0.054, P ∼ 0.0), were verified significantly less often (K-S test = 0.004, P < 0.001), and had been on Twitter for significantly less time (K-S test = 0.125, P ∼ 0.0) (Fig. 4A). Falsehood

diffused farther and faster than the truth despite these differences, not because of them.

When we estimated a model of the likelihood of retweeting, we found that falsehoods were 70% more likely to be retweeted than the truth (Wald chi-square test, P ∼ 0.0), even when controlling for the account age, activity level, and number of followers and followees of the original tweeter, as well as whether the original tweeter was a verified user (Fig. 4B). Because user characteristics and network structure could not explain the differential diffusion of truth and falsity, we sought alternative explanations for the differences in their diffusion dynamics.

One alternative explanation emerges from information theory and Bayesian decision theory. Novelty attracts human attention (24), contributes to productive decision-making (25), and encourages information sharing (26) because novelty updates our understanding of the world. When information is novel, it is not only surprising, but also more valuable, both from an information theoretic perspective [in that it provides the greatest aid to decision-making (25)] and from a social perspective [in that it conveys so-

cial status on one that is "in the know" or has access to unique "inside" information (26)]. We therefore tested whether falsity was more novel than the truth and whether Twitter users were more likely to retweet information that was more novel.

To assess novelty, we randomly selected ∼5000 users who propagated true and false rumors and extracted a random sample of ∼25,000 tweets that they were exposed to in the 60 days prior to their decision to retweet a rumor. We then specified a latent Dirichlet Allocation Topic model (27), with 200 topics and trained on 10 million English-language tweets, to calculate the information distance between the rumor tweets and all the prior tweets that users were exposed to before retweeting the rumor tweets. This generated a probability distribution over the 200 topics for each tweet in our data set. We then measured how novel the information in the true and false rumors was by comparing the topic distributions of the rumor tweets with the topic distributions of the tweets to which users were exposed in the 60 days before their retweet. We found that false rumors were significantly more

19

novel than the truth across all novelty metrics, displaying significantly higher information uniqueness (K-S test = 0.457, $P \sim 0.0$) (28), Kullback-Leibler (K-L) divergence (K-S test = 0.433, $P \sim 0.0$) (29), and Bhattacharyya distance (K-S test = 0.415, $P \sim 0.0$) (which is similar to the Hellinger distance) (30). The last two metrics measure differences between probability distributions representing the topical content of the incoming tweet and the corpus of previous tweets to which users were exposed.

Although false rumors were measurably more novel than true rumors, users may not have perceived them as such. We therefore assessed users' perceptions of the information contained in true and false rumors by comparing the emotional content of replies to true and false rumors. We categorized the emotion in the replies by using the leading lexicon curated by the National Research Council Canada (NRC), which provides a comprehensive list of ~140,000 English words and their associations with eight emotions based on Plutchik's (31) work on basic emotion—anger, fear, anticipation, trust, surprise, sadness, joy, and disgust (32)—and a list of ~32,000 Twitter hashtags and their weighted associations with the same emotions (33). We removed stop words and URLs from the reply tweets and calculated the fraction of words in the tweets that related to each of the eight emotions, creating a vector of emotion weights for each reply that summed to one across the emotions. We found that false rumors inspired replies expressing greater surprise (K-S test = 0.205, $P \sim 0.0$), corroborating the novelty hypothesis, and greater disgust (K-S test = 0.102, $P \sim 0.0$), whereas the truth inspired replies that expressed greater sadness (K-S test = 0.037, $P \sim 0.0$), anticipation (K-S test = 0.038, $P \sim 0.0$), joy (K-S test = 0.061, $P \sim 0.0$), and trust (K-S test = 0.060, $P \sim 0.0$) (Fig. 4, D and F). The emotions expressed in reply to falsehoods may illuminate additional factors, beyond novelty, that inspire people to share false news. Although we cannot claim that novelty causes retweets or that novelty is the only reason why false news is retweeted more often, we do find that false news is more novel and that novel information is more likely to be retweeted.

Numerous diagnostic statistics and manipulation checks validated our results and confirmed their robustness. First, as there were multiple cascades for every true and false rumor, the variance of and error terms associated with cascades corresponding to the same rumor will be correlated. We therefore specified cluster-robust standard errors and calculated all variance statistics clustered at the rumor level. We tested the robustness of our findings to this specification by comparing analyses with and without clustered errors and found that, although clustering reduced the precision of our estimates as expected, the directions, magnitudes, and significance of our results did not change, and chi-square ($P \sim 0.0$) and deviance (d) goodness-of-fit tests ($d = 3.4649 \times 10^{-6}$, $P \sim 1.0$) indicate that the models are well specified (see supplementary materials for more detail).

Second, a selection bias may arise from the restriction of our sample to tweets fact checked by the six organizations we relied on. Fact checking may select certain types of rumors or draw additional attention to them. To validate the robustness of our analysis to this selection and the generalizability of our results to all true and false rumor cascades, we independently verified a second sample of rumor cascades that were not verified by any fact-checking organization. These rumors were fact checked by three undergraduate students at Massachusetts Institute of Technology (MIT) and Wellesley College. We trained the students to detect and investigate rumors with our automated rumor-detection algorithm running on 3 million English-language tweets from 2016 (34). The undergraduate annotators investigated the veracity of the detected rumors using simple search queries on the web. We asked them to label the rumors as true, false, or mixed on the basis of their research and to discard all rumors previously investigated by one of the fact-checking organizations. The annotators, who worked independently and were not aware of one another, agreed on the veracity of 90% of the 13,240 rumor cascades that they investigated and achieved a Fleiss' kappa of 0.88. When we compared the diffusion dynamics of the true and false rumors that the annotators agreed on, we found results nearly identical to those estimated with our main data set (see fig. S17). False rumors in the robustness data set had greater depth (K-S test = 0.139, $P \sim 0.0$), size (K-S test = 0.131, $P \sim 0.0$), maximum breadth (K-S test = 0.139, $P \sim 0.0$), structural virality (K-S test = 0.066, $P \sim 0.0$), and speed (fig. S17) and a greater number of unique users at each depth (fig. S17). When we broadened the analysis to include majority-rule labeling, rather than unanimity, we again found the same results (see supplementary materials for results using majority-rule labeling).

Third, although the differential diffusion of truth and falsity is interesting with or without robot, or bot, activity, one may worry that our conclusions about human judgment may be biased by the presence of bots in our analysis. We therefore used a sophisticated bot-detection algorithm (35) to identify and remove all bots before running the analysis. When we added bot traffic back into the analysis, we found that none of our main conclusions changed—false news still spread farther, faster, deeper, and more broadly than the truth in all categories of information. The results remained the same when we removed all tweet cascades started by bots, including human retweets of original bot tweets (see supplementary materials, section S8.3) and when we used a second, independent bot-detection algorithm (see supplementary materials, section S8.3.5) and varied the algorithm's sensitivity threshold to verify the robustness of our analysis (see supplementary materials, section S8.3.4). Although the inclusion of bots, as measured by the two state-of-the-art bot-detection algorithms we used in our analysis, accelerated the spread of both true and false news, it affected their spread roughly equally. This suggests that false news spreads farther, faster, deeper, and more broadly than the truth because humans, not robots, are more likely to spread it.

Finally, more research on the behavioral explanations of differences in the diffusion of true and false news is clearly warranted. In particular, more robust identification of the factors of human judgment that drive the spread of true and false news online requires more direct interaction with users through interviews, surveys, lab experiments, and even neuroimaging. We encourage these and other approaches to the investigation of the factors of human judgment that drive the spread of true and false news in future work.

False news can drive the misallocation of resources during terror attacks and natural disasters, the misalignment of business investments, and misinformed elections. Unfortunately, although the amount of false news online is clearly increasing (Fig. 1, C and E), the scientific understanding of how and why false news spreads is currently based on ad hoc rather than large-scale systematic analyses. Our analysis of all the verified true and false rumors that spread on Twitter confirms that false news spreads more pervasively than the truth online. It also overturns conventional wisdom about how false news spreads. Though one might expect network structure and individual characteristics of spreaders to favor and promote false news, the opposite is true. The greater likelihood of people to retweet falsity more than the truth is what drives the spread of false news, despite network and individual factors that favor the truth. Furthermore, although recent testimony before congressional committees on misinformation in the United States has focused on the role of bots in spreading false news (36), we conclude that human behavior contributes more to the differential spread of falsity and truth than automated robots do. This implies that misinformation-containment policies should also emphasize behavioral interventions, like labeling and incentives to dissuade the spread of misinformation, rather than focusing exclusively on curtailing bots. Understanding how false news spreads is the first step toward containing it. We hope our work inspires more large-scale research into the causes and consequences of the spread of false news as well as its potential cures.

REFERENCES AND NOTES

1. L. J. Savage, *J. Am. Stat. Assoc.* **46**, 55–67 (1951).
2. H. A. Simon, *The New Science of Management Decision* (Harper & Brothers Publishers, New York, 1960).
3. R. Wedgwood, *Noûs* **36**, 267–297 (2002).
4. E. Fehr, U. Fischbacher, *Nature* **425**, 785–791 (2003).
5. C. E. Shannon, *Bell Syst. Tech. J.* **27**, 379–423 (1948).
6. S. Bikhchandani, D. Hirshleifer, I. Welch, *J. Polit. Econ.* **100**, 992–1026 (1992).
7. K. Rapoza, "Can 'fake news' impact the stock market?" *Forbes*, 26 February 2017; www.forbes.com/sites/kenrapoza/2017/02/26/can-fake-news-impact-the-stock-market/.
8. M. Mendoza, B. Poblete, C. Castillo, in *Proceedings of the First Workshop on Social Media Analytics* (Association for Computing Machinery, ACM, 2010), pp. 71–79.
9. A. Gupta, H. Lamba, P. Kumaraguru, A. Joshi, in *Proceedings of the 22nd International Conference on World Wide Web* (ACM, 2010), pp. 729–736.

10. K. Starbird, J. Maddock, M. Orand, P. Achterman, R. M. Mason, in *iConference 2014 Proceedings* (iSchools, 2014).
11. J. Gottfried, E. Shearer, "News use across social media platforms," Pew Research Center, 26 May 2016; www.journalism.org/2016/05/26/news-use-across-social-media-platforms-2016/.
12. C. Silverman, "This analysis shows how viral fake election news stories outperformed real news on Facebook," *BuzzFeed News*, 16 November 2016; www.buzzfeed.com/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook/.
13. M. De Domenico, A. Lima, P. Mougel, M. Musolesi, *Sci. Rep.* **3**, 2980 (2013).
14. O. Oh, K. H. Kwon, H. R. Rao, in *Proceedings of the International Conference on Information Systems* (International Conference on Information Systems, ICIS, paper 231, 2010).
15. M. Tambuscio, G. Ruffo, A. Flammini, F. Menczer, in *Proceedings of the 24th International Conference on World Wide Web* (ACM, 2015), pp. 977–982.
16. Z. Zhao, P. Resnick, Q. Mei, in *Proceedings of the 24th International Conference on World Wide Web* (ACM, 2015), pp. 1395–1405.
17. M. Gupta, P. Zhao, J. Han, in *Proceedings of the 2012 Society for Industrial and Applied Mathematics International Conference on Data Mining* (Society for Industrial and Applied Mathematics, SIAM, 2012), pp. 153–164.
18. G. L. Ciampaglia *et al.*, *PLOS ONE* **10**, e0128193 (2015).
19. A. Friggeri, L. A. Adamic, D. Eckles, J. Cheng, in *Proceedings of the International Conference on Weblogs and Social Media* (Association for the Advancement of Artificial Intelligence, AAAI, 2014)
20. M. Del Vicario *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **113**, 554–559 (2016).
21. A. Bessi *et al.*, *PLOS ONE* **10**, e0118093 (2015).
22. Friggeri *et al.* (*19*) do evaluate two metrics of diffusion: depth, which shows little difference between true and false rumors, and shares per rumor, which is higher for true rumors than it is for false rumors. Although these results are important, they are not definitive owing to the smaller sample size of the study; the early timing of the sample, which misses the rise of false news after 2013; and the fact that more shares per rumor do not necessarily equate to deeper, broader, or more rapid diffusion.
23. S. Goel, A. Anderson, J. Hofman, D. J. Watts, *Manage. Sci.* **62**, 180–196 (2015).
24. L. Itti, P. Baldi, *Vision Res.* **49**, 1295–1306 (2009).
25. S. Aral, M. Van Alstyne, *Am. J. Sociol.* **117**, 90–171 (2011).
26. J. Berger, K. L. Milkman, *J. Mark. Res.* **49**, 192–205 (2012).
27. D. M. Blei, A. Y. Ng, M. I. Jordan, *J. Mach. Learn. Res.* **3**, 993–1022 (2003).
28. S. Aral, P. Dhillon, "Unpacking novelty: The anatomy of vision advantages," Working paper, MIT–Sloan School of Management, Cambridge, MA, 22 June 2016; https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2388254.
29. T. M. Cover, J. A. Thomas, *Elements of Information Theory* (Wiley, 2012).
30. T. Kailath, *IEEE Trans. Commun. Technol.* **15**, 52–60 (1967).
31. R. Plutchik, *Am. Sci.* **89**, 344–350 (2001).
32. S. M. Mohammad, P. D. Turney, *Comput. Intell.* **29**, 436–465 (2013).
33. S. M. Mohammad, S. Kiritchenko, *Comput. Intell.* **31**, 301–326 (2015).
34. S. Vosoughi, D. Roy, in *Proceedings of the 10th International AAAI Conference on Weblogs and Social Media* (AAAI, 2016), pp. 707–710.
35. C. A. Davis, O. Varol, E. Ferrara, A. Flammini, F. Menczer, in *Proceedings of the 25th International Conference Companion on World Wide Web* (ACM, 2016), pp. 273–274.
36. For example, this is an argument made in recent testimony by Clint Watts—Robert A. Fox Fellow at the Foreign Policy Research Institute and Senior Fellow at the Center for Cyber and Homeland Security at George Washington University—given during the U.S. Senate Select Committee on Intelligence hearing on "Disinformation: A Primer in Russian Active Measures and Influence Campaigns" on 30 March 2017; www.intelligence.senate.gov/sites/default/files/documents/os-cwatts-033017.pdf.

21

When Narrative Matters More Than Fact

A teacher argues that helping students analyze the stories they care so much about is more effective than pushing pure fact-checking.

ASHLEY LAMB-SINCLAIR **JANUARY 9, 2017**

When I was in high school, one of my history teachers was also the football coach. "Coach Mac," we called him. For a right-brained creative like me, history was often a toss up. There were certain parts of the curriculum that I loved, but I loathed (and was generally inept at) memorizing dates and obscure facts. But Coach Mac taught us history

through football plays and storytelling. Through a series of Xs, Os, and arrows detailing their paths, Coach Mac told stories of Roman invasions, the Crusades, Genghis Khan, and the rise of Stalin. I sat in the front row, took copious notes, and was a star student every day in that class.

Because of Coach Mac, I became a history minor in college. And yet, if you asked me dates and details of these events Coach Mac and my college professors taught me, I could not tell you any of them without the aid of Google. The truth is, history stole my heart not because of the facts, but because of the stories.

Joseph Campbell famously said that there are only two stories in the whole world: Hero takes a journey and stranger comes to town. As an English teacher, I enjoy telling my students this nugget of wisdom and challenging them to defy it. They never can because, although stories are powerful, they are also simple. There are certain constructs, rhythms, and traits to a well-crafted story. Stories, at their heart, are either about heroes on a journey or strangers coming into a new setting.

For many Americans, Donald Trump is a hero on a journey; for others, he is a villainous stranger who has come to town. No one knows how the story will play out, but to deny that the country is in the midst of a fascinating <u>rising action</u>, to use a literary term, is to admit that you're not paying attention.

Like many educators, I am appalled at the wealth of <u>fake news that</u>

floats around social media and the power it has over young people who do not necessarily have the skills to interpret it. Many adults are worried about how to best teach strategies for interpreting fake news, and many of those strategies seem to surround the idea of fact checking. And although classrooms like mine should place a strong focus on helping students navigate the evolving world of the internet and social media, to be critical consumers of media, and to develop a general desire to seek facts above fiction, to concentrate solely on fact checking is a naive approach to the problem.

Just as it was for me so many years ago in Coach Mac's class, narrative, both fiction and nonfiction, will always be more alluring than a collection of facts—for better or worse—because narrative is rooted in the human experience. People want to connect with characters, want to see a plot develop to its end, and want to engage in the fascinating layers of conflict.

Explaining to someone, however accurately, that Donald Trump didn't help save 2,100 jobs with the Carrier deal, but rather 850, and that he may have actually had very little to do with it, or that the deal may have negative implications for the economy and job growth down the road, means virtually nothing to someone who has lost a job and gotten it back. To this person, there is a clear narrative that resonates: Trump is the hero. Telling someone whose only image and interaction ever with a woman wearing a hijab is through negative stereotypes on social media that five of the last 12 Nobel Peace Prize winners were Muslim means little to someone whose mind has generalized such a character as the villain.

Facts (or the lack thereof) mean very little to people caught up in storylines. The best way to teach true understanding is not by teaching students facts (although that is still a valuable lesson); it is to teach them to analyze, as one does with elements of narrative.

When I was growing up in a small town, my only contact with Latinos was with two men who worked with me in a local restaurant. I was a waitress, and they worked in the kitchen. These men were a little more flirty with me than is probably appropriate for grown men to be toward a 17-year-old girl. They sometimes made me uncomfortable, and because of that, I began to develop a perception about all Latinos, based solely on these two men. The story I crafted from this experience, regardless of the facts, was that Latino men were inappropriately flirtatious toward women. Facts and statistics would have meant very little to changing the story in my mind about what it meant to be a Latino man. My facts were wrong, but my story was what mattered to me. If I saw anyone I perceived as a Latino man in public, and I was alone, I would feel myself become anxious. How could facts in a moment of fear for a 17-year-old girl make any difference? It would be equivalent to telling someone who is afraid of flying that more deaths occur by car than plane.

But what did change the story for me was moving to Southern California. There, I joined a sorority with mostly Latina women who became my new "sisters." I worked on campus as a telemarketer, surrounded by Latinos who became my family. Because I was living away from home, one of my new friends invited me over for dinner most Sundays, and on those days, her father cooked the best carne

asada I have ever eaten and welcomed me kindly into their home—creating a new image of Latino men for me. I went on dates with several Latino men who treated me with kindness and respect. Ultimately, truth is subjective. What was true for me at 17 was not true for me years later. The narrative I had crafted as a teenager suddenly seemed ridiculous, and not because someone presented me with facts, but because I understood much more of the story. I had analyzed various characters and could now understand how false my perception had been.

Obviously, it is unlikely that every single small-town young person like me will have the experiences I had. But it is possible to emulate such experiences in the classroom.

Now is the time for teachers to teach students not only to be critical thinkers who question the validity of facts, but also to analyze narratives. That is what Coach Mac did in his classroom through his football plays. When a certain Roman general marked as an X on his chalkboard acted in a way that developed the plot of the story, Coach Mac would ask the class, "Why do you think he did that?" We didn't have Google then, but even if we had, he wasn't asking us to simply look up the facts; he was asking us to analyze what had taken place thus far, how X had behaved up to that point, and what the possibilities were for X's next actions.

When living in California, the storyline of Latinos-as-villains no longer made sense when I analyzed as Coach Mac had taught me. And even if students can't go to California as I had, a teacher can still expose

students to various types of characters and plotlines from many perspectives, both fictional and real. Teachers can—and do—ask the same types of questions of those narratives that Coach Mac asked of me and my peers of historical ones.

It is a human endeavor to create and tell stories. From the caves of Lascaux to oral-storytelling traditions around the world, humans have sought ways to share truth as they see it, to develop narratives in ways that makes sense for each individual. Young people use social media to tell stories and share their perception of truth, and it is also on these platforms that they seek truth.

I was lucky enough to move to a new place and experience other cultures that changed my perception of some people, but others don't take such leaps. Young people have opportunities to use the global space that is social media to broaden their perceptions and be critical analysts of false narratives. It is up to adults to teach students these skills.

Adults can teach students about unreliable narrators, about character motivation, about the need of any good storyteller to create conflicts and obstacles. Just as I explained recently to the students in my creative-writing class who are writing 10-minute plays, a good storyteller should plant minor obstacles in the beginning of the story that will indicate what the climax will be. So, as critics of stories, students might have noticed, as I did, that Donald Trump planted seeds of a treacherous media and rigged elections early on as minor

obstacles in his story, so that as his story progressed those conflicts and the people who enacted them became more and more like the villains, while he <u>became more and more the hero</u>. Because I am a storyteller, I could see the plot unfolding. I want the same skills for my students because facts aren't enough when it is time to understand the difference between a hero and a villain.

———————

We want to hear what you think about this article. <u>Submit a letter</u> to the editor or write to letters@theatlantic.com.

Life's Stories

How you arrange the plot points of your life into a narrative can shape who you are—and is a fundamental part of being human.

Story by Julie Beck

AUGUST 10, 2015 | **HEALTH**

In Paul Murray's novel *Skippy Dies,* there's a point where the main character, Howard, has an existential crisis. "'It's just not how I expected my life would be,'" he says.

"'What did you expect?'" a friend responds.

"Howard ponders this. 'I suppose—this sounds stupid, but I suppose I thought there'd be more of a *narrative arc.*'"

But it's not stupid at all. Though perhaps the facts of someone's life, presented end to end, wouldn't much resemble a narrative to the outside observer, the way people choose to tell the stories of their lives, to others and—crucially—to themselves, almost always does have a narrative arc. In telling the story of how you became who you are, and of who you're on your way to becoming, the story itself becomes a part of who you are.

"Life stories do not simply *reflect* personality. They *are* personality, or more accurately, they are important *parts of* personality, along with other parts, like dispositional traits, goals, and values," writes Dan McAdams, a professor of psychology at Northwestern University, along with Erika

Manczak, in a chapter for the *APA Handbook of Personality and Social Psychology.*

In the realm of narrative psychology, a person's life story is not a Wikipedia biography of the facts and events of a life, but rather the way a person integrates those facts and events internally—picks them apart and weaves them back together to make meaning. This narrative becomes a form of identity, in which the things someone chooses to include in the story, and the way she tells it, can both reflect and shape who she is. A life story doesn't just say what happened, it says why it was important, what it means for who the person is, for who they'll become, and for what happens next.

"Sometimes in cases of extreme autism, people don't construct a narrative structure for their lives," says Jonathan Adler, an assistant professor of psychology at Olin College of Engineering, "but the default mode of human cognition is a narrative mode."

When people tell others about themselves, they kind of have to do it in a narrative way—that's just how humans communicate. But when people think about their lives to themselves, is it always in a narrative way, with a plot that leads from one point to another? There's an old adage that everyone has a book inside of them. (Christopher Hitchens once said that inside is "exactly where I think it should, in most cases, remain.") Is there anyone out there with a life

story that's not a story at all, but some other kind of more disjointed, avant-garde representation of their existence?

"This is an almost impossible question to address from a scientific approach," says Monisha Pasupathi, a professor of developmental psychology at the University of Utah. Even if we are, as the writer Jonathan Gottschall put it, "storytelling animals," what does that mean from one person to the next? Not only are there individual differences in how people think of their stories, there's huge variation in the degree to which they engage in narrative storytelling in the first place.
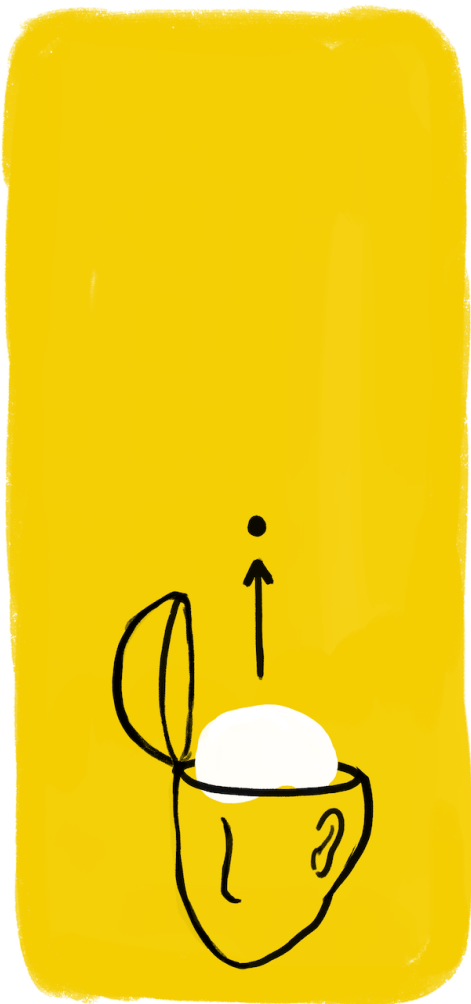
"Some people write in their diaries and are very introspective, and some people are not at all," says Kate McLean, an associate professor of psychology at Western Washington University. Journal-keeping, though a way of documenting the life story, doesn't always make for a tightly-wound narrative. A writer I interviewed several months ago —Sarah Manguso—has kept a diary for 25 years, and still told me, "Narrative is not a mode that has ever come easily to me."

Nevertheless, the researchers I spoke with were all convinced that even if it's not 100 percent universal to see life as a story, it's at least extremely common.

"I think normal, healthy adults have in common that they can all produce a life story," Pasupathi says. "They can all put one together … In order to have relationships, we've all had to tell little pieces of our story. And so it's hard to be a human being and have relationships without having some version of a life story floating around."

But life rarely follows the logical progression that most stories—good stories—do, where the clues come together, guns left on mantles go off at the appropriate moments, the climax comes in the third act. So narrative seems like an incongruous framing method for life's chaos, until you remember where stories came from in the first place. Ultimately, the only material we've ever had to make stories

out of is our own imagination, and life itself.

Storytelling, then—fictional or nonfictional, realistic or embellished with dragons—is a way of making sense of the world around us.

"Stories don't have to be really simple, like fairy tales. They can be complicated. It can be like James Joyce out there."

"Life is incredibly complex, there are lots of things going on in our environment and in our lives at all times, and in order to hold onto our experience, we need to make meaning out of it," Adler says. "The way we do that is by structuring our lives into stories."

It's hardly a simple undertaking. People contain multitudes, and by multitudes, I mean libraries. Someone might have an overarching narrative for her whole life, and different narratives for different realms of her life—career, romance, family, faith. She might have narratives within each realm that intersect, diverge, or contradict each other, all of them filled with the micro-stories of specific events. And to truly make a life story, she'll need to do what researchers call "autobiographical reasoning" about the events—"identifying lessons learned or insights gained in life experiences,

marking development or growth through sequences of scenes, and showing how specific life episodes illustrate enduring truths about the self," McAdams and Manczak write.

"Stories don't have to be really simple, like fairy-tale-type narratives," McAdams says. "They can be complicated. It can be like James Joyce out there."

If you really like James Joyce, it might be a lot like James Joyce. People take the stories that surround them—fictional tales, news articles, apocryphal family anecdotes—then identify with them and borrow from them while fashioning their own self-conceptions. It's a Möbius strip: Stories are life, life is stories.

* * *

People aren't writing their life stories from birth, though. The ability to create a life narrative takes a little while to come online—the development process gives priority to things like walking, talking, and object permanence. Young children can tell stories about isolated events, with guidance, and much of adolescence is dedicated to learning "what goes in a story… and what makes a good story in the first place," Pasupathi says. "I don't know how much time you've spent around little kids, but they really don't understand that. I have a child who can really take an hour to tell you about *Minecraft*." Through friends, family, and fiction, children

learn what others consider to be good storytelling—and that being able to spin a good yarn has social value.

It's in the late teens and early years of adulthood that story construction really picks up—because by then people have developed some of the cognitive tools they need to create a coherent life story. These include causal coherence—the ability to describe how one event led to another—and thematic coherence—the ability to identify overarching values and motifs that recur throughout the story. In a study analyzing the life stories of 8-, 12-, 16-, and 20-year-olds, these kinds of coherence were found to increase with age. As the life story enters its last chapters, it may become more set in stone. In one study by McLean, older adults had more thematic coherence, and told more stories about stability, while young adults tended to tell more stories about change.

McAdams conceives of this development as the layering of three aspects of the self. Pretty much from birth, people are "actors." They have personality traits, they interact with the world, they have roles to play—daughter, sister, the neighbor's new baby that cries all night and keeps you up. When they get old enough to have goals, they become "agents," too—still playing their roles and interacting with the world, but making decisions with the hopes of producing desired outcomes. And the final layer is "author," when people begin to bundle ideas about the future with experiences from the past and present to form a narrative self.

This developmental trajectory could also explain why people enjoy different types of fictional stories at different ages. "When you're a kid, it's mostly about plot," McAdams says. "This happens and this happens. You're not tuned into the idea that a character develops." Thus, perhaps, the appeal of cartoon characters who never get older.

Recently, McAdams says, his book club read *Ethan Frome* by Edith Wharton. "I read it in high school and hated it," he says. "All I could remember about it was that this sled hits a tree. And we read it recently in the club, and whoa, is it fabulous. A sled does hit the tree, there's no doubt that is a big scene, but how it changes these people's lives and the tragedy of this whole thing, it's completely lost on 18-year-olds. Things are lost on 8-year-olds that a 40-year-old picks up, and things that an 8-year-old found compelling and interesting will just bore a 40-year-old to tears sometimes."

And like personal taste in books or movies, the stories we tell ourselves about ourselves are influenced by more than just, well, ourselves. The way people recount experiences to others seems to shape the way they end up remembering those events. According to Pasupathi's research, this happens in a couple of ways. One is that people tailor the stories they tell to their audiences and the context. (For example, I tell the story of the time I crashed my mom's car much differently now, to friends, than the way I told it to my mom at the time. Much less crying.)

When people drop the cheesy pick-up line "What's your story?", like a man who nicks his carotid artery while shaving, they've accidentally hit upon something vital.

The other is that the act of telling is a rehearsal of the story, Pasupathi says. "And rehearsal strengthens connections between some pieces of information in your mind and diminishes connections between others. So the things I tell you become more accessible to me and more memorable to me. Those can be pretty lasting effects." So when people drop the cheesy pick-up line "What's your story?" at a bar, like a man who nicks his carotid artery while shaving, they've accidentally hit upon something vital.

But just as there are consequences to telling, there are consequences to not telling. If someone is afraid of how people might react to a story, and they keep it to themselves, they'll likely miss out on the enrichment that comes with a back-and-forth conversation. A listener "may give you other things to think about, or may acknowledge that this thing you thought was really bad is actually not a big deal, so you get this richer and more elaborated memory," Pasupathi says. If you don't tell, "your memory for that event may be less flexible and give you less chance for growth." This is basically

the premise of talk therapy.

And all of this doesn't even account for all the conversations you plan to have, or elaborately imagine having and never have. The path from outside to inside and back out is winding, dark, and full of switchbacks.

* * *

Once certain stories get embedded into the culture, they become master narratives—blueprints for people to follow when structuring their own stories, for better or worse. One such blueprint is your standard "go to school, graduate, get a job, get married, have kids."

That can be a helpful script in that it gives children a sense of the arc of a life, and shows them examples of tentpole events that could happen. But the downsides of standard narratives have been well-documented—they stigmatize anyone who doesn't follow them to a T, and provide unrealistic expectations of happiness for those who do. If this approach were a blueprint for an Ikea desk instead of a life, almost everyone trying to follow it would end up with something wobbly and misshapen, with a few leftover bolts you find under the couch, boding ill for the structural integrity of the thing you built.

"I think that's particularly pernicious frame for people who become parents," Pasupathi says. "That's a narrative where

the pinnacle is to get married and have kids and then everything will be sort of flatly happy from then on."

And these scripts evolve as culture evolves. For example, in centuries past, stories of being possessed by demons might not have been out of place, but it's unlikely most people would describe their actions in those terms nowadays.

Other common narrative structures seen in many cultures today are redemption sequences and contamination sequences. A redemption story starts off bad and ends better —"That horrible vacation ultimately brought us closer as a family"—while a contamination story does the opposite —"The cruise was amazing until we all got food poisoning." Having redemption themes in one's life story is generally associated with greater well-being, while contamination themes tend to coincide with poorer mental health.

Many people have some smaller stories of each type sprinkled throughout their greater life story, though a person's disposition, culture, and environment can influence which they gravitate to. People can also see the larger arc of their lives as redemptive or contaminated, and redemption in particular is a popular, and particularly American, narrative. "Evolving from the Puritans to Ralph Waldo Emerson to Oprah Winfrey… Americans have sought to author their lives as redemptive tales of atonement,

emancipation, recovery, self-fulfillment, and upward social mobility," McAdams writes in an <u>overview of life story research</u>. "The stories speak of heroic individual protagonists —the chosen people—whose manifest destiny is to make a positive difference in a dangerous world, even when the world does not wish to be redeemed."

The redemption story is American optimism—things will get better!—and American exceptionalism—I can make things better!—and it's in the water, in the air, and in our heads. This is actually a good thing a lot of the time. <u>Studies have shown</u> that finding a positive meaning in negative events is linked to a more complex sense of self and greater life satisfaction. And even controlling for general optimism, McAdams and his colleagues found that having more redemption sequences in a life story was still associated with higher well-being.

The trouble comes when redemption isn't possible. The redemptive American tale is one of privilege, and for those who can't control their circumstances, and have little reason to believe things will get better, it can be an illogical and unattainable choice. There are things that happen to people that cannot be redeemed.

It can be hard to share a story when it amounts to: "This happened, and it was terrible. The end." In research McLean did, in which she asked people who'd had near-death

experiences to tell their stories to others. "The people who told these unresolved stories had really negative responses," she says. If there wasn't some kind of uplifting, redemptive end to the story (beyond just the fact that they survived), "The listeners did not like that."

"The redemptive story is really valued in America, because for a lot of people it's a great way to tell stories, but for people who just can't do that, who can't redeem their traumas for whatever reason, they're sort of in a double bind," she continues. "They both have this crappy story that's hanging on, but they also can't tell it and get acceptance or validation from people."

In cases like this, for people who have gone through a lot of trauma, it might be better for them not to autobiographically reason about it at all.

There are things that happen to people that cannot be redeemed.

"The first time I ever found this association, of reasoning associated with poor mental health, I thought that I had analyzed my data incorrectly," McLean says. But after other researchers replicated her findings, she got more confident that something was going on. She thinks that people may repress traumatic events in a way that, while not ideal, is still

"healthy enough."

"The typical idea is that you can repress something but it's going to come back and bite you if you don't deal with it," she says. "But that's still under the assumption that people have the resources to deal with it."

In one study, McLean and her colleagues interviewed adolescents attending a high school for vulnerable students. One subject, Josie, the 17-year-old daughter of a single mother, suffered from drug and alcohol abuse, bipolar disorder, rape, and a suicide attempt. She told the researchers that her self-defining memory was that her mother had promised not to have more children, and then broke that promise.

"I'm the only person that I can rely on in my life because I've tried to rely on other people and I either get stabbed in the back or hurt, so I really know that I can only trust myself and rely on myself," Josie said when recounting this memory.

"That's pretty intensive reasoning," McLean says. "So that's meaningful in understanding who you are, but it doesn't really give you a positive view of who you are. It may be true in the moment, but it's not something that propels someone towards growth."

It's possible to over-reason about good things in your life as well. "There's been some experimental research that shows

that when people are asked to reflect on positive experiences, it makes them feel worse, because you're like 'Oh, why *did* I marry that person?'" McLean says. "Wisdom and maturity and cognitive complexity are all things that we value, but they don't necessarily make you happy."

* * *

Though sometimes autobiographical reasoning can lead to dark thoughts, other times it can help people find meaning. And while you may be able to avoid reasoning about a certain event, it would be pretty hard to leave all the pages of a life story unwritten.

"I think the act of framing our lives as a narrative is neither positive nor negative, it just is," Adler says. "That said, there are better and worse ways of doing that narrative process for our mental health."

In his research, Adler has noticed two themes in people's stories that tend to correlate with better well-being: agency, or feeling like you are in control of your life, and communion, or feeling like you have good relationships in your life. The connection is "a little fuzzier" with communion, Adler says—there's a strong relationship between communion and well-being at the same moment; it's less clear if feeling communion now predicts well-being later.

But agency sure does. It makes sense, since feelings of helplessness and hopelessness are classic symptoms of depression, that feeling in control would be good for mental

health. Adler did <u>a longitudinal study</u> of 47 adults undergoing therapy, having them write personal narratives and complete mental health assessments over the course of 12 therapy sessions. What he found was not only that themes of agency in participants' stories increased over time, and that mental health increased, and that the two were related, but that increased agency actually appeared in stories *before* people's mental health improved.

"It's sort of like people put out a new version of themselves and lived their way into it," Adler says.

(There's something about the narrative form, specifically— while expressing thoughts and feelings about negative events seems to help people's well-being, <u>one study</u> found that writing them in a narrative form helped more than just listing them.)

But, he continues, "I'm not like Mr. Agency, agency at all costs. I don't believe that. If you have stage 4 cancer, agency may be good for you, but is it a rational choice? And I do think [redemption] is good in the long term, but in the throes of really struggling with illness, I don't know that it actually helps people."

But I wondered: Though agency may be good for you, does seeing yourself as a strong protagonist come at a cost to the other characters in your story? Are there implications for empathy if we see other people as bit players instead of

protagonists in their own right?

"That's actually kind of an interesting empirical idea," Pasupathi says. "I don't know that anybody's looking at that."

As Adler's work shows, people need to see themselves as actors to a certain degree. And Pasupathi's work shows that other people play a big role in shaping life stories. The question, perhaps, is how much people recognize that their agency is not absolute.

According to one study, highly generative people—that is, people who are caring and committed to helping future generations—often tell stories about others who helped them in the past. McAdams suggests that narcissists are probably more likely to do the opposite—"People [who] are really good at talking about themselves and pushing their own narrative, but they're not willing to listen to yours."

"If our stories are about us as triumphant agents going through life and overcoming, and they underplay the role of other people and the role of institutional support in helping us do those things, we are likely to be less good at recognizing how other people's lives are constrained by institutions and other people," Pasupathi says. "I think that has real implications for how we think about inequity in our society. The more the whole world is designed to work for you, the less you are aware that it is working for you."

It's a dizzying problem: People use stories to make sense of life, but how much do those stories reflect life's realities? Even allowing for the fact that people are capable of complex Joyce-ian storytelling, biases, personality differences, or emotions can lead different people to see the same event differently. And considering how susceptible humans are to false memories, who's to say that the plot points in someone's life story really happened, or happened the way she thought they did, or really caused the effects she saw from them?

"Any creation of a narrative is a bit of a lie."

Pasupathi's not convinced that it matters that much whether life stories are perfectly accurate. A lot of false memory research has to do with eyewitness testimony, where it matters a whole lot whether a person is telling a story precisely as it happened. But for narrative-psychology researchers, "What really matters isn't so much whether it's true in the forensic sense, in the legal sense," she says. "What really matters is whether people are making something meaningful and coherent out of what happened. Any creation of a narrative is a bit of a lie. And some lies have enough truth."

Organizing the past into a narrative isn't just a way to understand the self, but also to attempt to predict the future. Which is interesting, because the storytelling device that seems most incompatible with the realities of actual life is foreshadowing. Metaphors, sure. As college literature class discussion sections taught me, you can see anything as a metaphor if you try hard enough. Motifs, definitely. Even if you're living your life as randomly as possible, enough things will happen that, like monkeys with typewriters, patterns will start to emerge.

But no matter how hard you try, no matter how badly you want to, there is no way to truly know the future, and the world isn't *really* organizing itself to give you hints. If you're prone to overthinking, and playing out every possible scenario in your head in advance, you can see foreshadowing in everything. The look your partner gives you means a fight is on the horizon, that compliment from your boss means you're on track for a promotion, all the little things you've forgotten over the years mean you're definitely going to get dementia when you're old.

"Actual life is full of false clues and signposts that lead nowhere," E.M. Forster once wrote. These become obvious in the keeping of a diary: "Imagine a biography that includes not just a narrative but also all the events that failed to foreshadow," Manguso writes in *Ongoingness,* the book about her 25-year diary. "Most of what the diary includes foreshadows nothing."

So what to do, then, with all the things that don't fit tidily? There is evidence that finding some "unity" in your narrative identity is better, psychologically, than not finding it. And it probably is easier to just drop those things as you pull patterns from the chaos, though it may take some readjusting.

But Pasupathi rejects that. "I would want to see people do a good job of not trying to leave stuff out because they can't make it fit," she says. "We're not trying to make pieces of your life go away."

And so even with the dead ends and wrong turns, people can't stop themselves. "We try to predict the future all the time," Pasupathi says. She speculates that the reason there's foreshadowing in fiction in the first place is because of this human tendency. The uncertainty of the future makes people uncomfortable, and stories are a way to deal with that.

"The future is never a direct replica of the past," Adler says. "So we need to be able to take pieces of things that have happened to us and reconfigure them into possible futures." For example, through experience, one learns that "We need to talk" rarely foreshadows anything good. (Life has its own clichés.)

There's been some brain research supporting this link between the past and the future, showing that the same regions of the brain are activated when people are asked to remember something and when they're asked to imagine an event that hasn't happened yet. On the flip side, a patient with severe amnesia also had trouble imagining the future.

Similarly, the way someone imagines his future seems to affect the way he sees his past, at the same time as his past informs what he expects for the future.

"If you're planning to be a doctor, and you're a 25-year-old starting medical school, and you have expectations about what the next five to 10 years are going to be like, you've probably construed a narrative from your past that helps you understand how you got to this point," McAdams says. "Then, say, you get into med school and you hate it and you drop out, you probably at the same time are going to change your past. You rewrite the history."

A life story is written in chalk, not ink, and it can be changed. "You're both the narrator and the main character of your story," Adler says. "That can sometimes be a revelation —'Oh, I'm not just living out this story, I am actually in charge of this story.'"

Whether it's with the help of therapy, in the midst of an identity crisis, when you've been chasing a roadrunner of foreshadowing towards a tunnel that turns out to be painted on a wall, or slowly, methodically, day by day—like with all stories, there's power in rewriting.

"The past is always up for grabs," McAdams says.

weaknesses of the polygraph method illustrate the inherent problems facing mechanised lie detectors.[13] That said, the pace at which AI is evolving suggests that these problems may be overcome sooner than we would suppose.

FACTS ARE NOT ENOUGH

In the meantime, *Homo sapiens* must do combat with Post-Truth. In his book on the perils of statistics, the psychologist Daniel Levitin insists that the due diligence required of today's citizens is part of 'an implicit bargain we've all made'. The trivial tasks of research and information retrieval that used to consume days can now be accomplished in seconds on a smartphone or a tablet.

'We've saved incalculable numbers of hours of trips to libraries and far-flung archives, of hunting through thick books for the one passage that will answer our question,' writes Levitin. 'The implicit bargain that we all need to make explicit is that we will use just *some* of that time we saved in information acquisition to perform proper information verification.'

He commends as a tool kit the methods devised by Thomas Bayes (1701–61), the English statistician and philosopher, whereby the probability of a proposition's truth is determined by the incremental accumulation of evidence.[14] The more a doctor learns of our symptoms,

the more accurately he can diagnose our condition. The more verifiable evidence we have of President Trump's connections in Russia, the more we can say with confidence about the probity of his relationship with its government. We need to recover the patience to apply this technique.

This is a fair demand to make. But to stand a chance of success, such strategies must be advanced in the world as it is, rather than the world as it once was. In particular, as the 'backfire effect' illustrates, it is a mistake to imagine that Post-Truth will crumble under the weight of freshly verified information repeated relentlessly and ubiquitously.

Indeed, it is a common error to confuse data with truth: the former informs the latter, but they are not the same thing. In the Brexit referendum, the Remain camp's greatest mistake was to assume that torrents of statistics would win the day. In the Vietnam War, what Viktor Mayer-Schönberger and Kenneth Cukier call the 'dictatorship of data' had a disastrous impact upon US strategy. Robert McNamara, Secretary of Defense under Presidents Kennedy and Johnson, displayed an almost religious faith in the power of statistics to guide public policy. As a consequence, the daily 'body count' – the number of enemy dead – became the crucial 'data point'.

Privately, and later publicly, the generals believed this statistic to be a useless measure of success in such a complex military and political context. Its centrality to

the debate also invited fabrication: officers in the field were alleged to have inflated the figures as a matter of course. The truth of battle could not be captured in a spreadsheet or a set of graphs – any more than the case for Britain's continued membership of the EU could be reduced to a series of statistics.[15]

In the right circumstances, a lie may be defeated by the skilful deployment of facts. But Post-Truth is, first and foremost, an emotional phenomenon. It concerns our attitude to truth, rather than truth itself.

From this, it should be clear that the counter-attack has to be emotionally intelligent as well as rigorously rational. In their account of science denial, Sara and Jack Gorman insist that academics and researchers must raise their game accordingly in the public sphere. Just as the opponents of vaccination have deployed the techniques of celebrity, so 'scientifically credible charismatic leaders' are needed to counter their claims.[16] At a time when Professor Brian Cox can sell out stadiums, the British astronaut Tim Peake is mobbed by fans and Professor Stephen Hawking is a cultural icon, this is not an unreasonable expectation. As the Gormans put it:

We propose that scientists not only become fluent in the kind of information we confront on the Internet but also that they join the conversation in a much more active way. We think that scientific and medical societies in particular have much to gain from formalizing a broad,

far-reaching online and social media strategy. Yes, all scientific and medical societies these days have websites and send out e-newsletters and digitize their journals. But how many of them have a truly active Twitter or Facebook account providing up-to-the-minute coverage for the general public about important scientific and medical issues?

Rapid rebuttal is only the start, though. As they acknowledge:

Discussion should center on not only how to make the material accessible but also how to present it in a manner that will discourage irrational response ... Instead of just telling people the percent risk, scientists need to understand what these percentages actually mean to people and, most important, frame them in a way that they will be most convincing and acceptable to the nonscientists. A small amount of training in cognitive psychology and behavioral economics should make scientists more aware of the biases and heuristics people use to interpret scientific information and teach them how to communicate around these psychological processes ...

In a Post-Truth world, in other words, it is not enough to make an intellectual case. In many (perhaps most) contexts, facts need to be communicated in a way that recognises emotional as well as rational imperatives. For

doctors talking to patients, for instance, one such model is 'Motivation Interviewing', a clinical method developed in addiction treatment that goes beyond the simple transference of information to explore the subject's motives, anxieties and ambivalence to encourage behavioural change. At a time when most British family doctors spend an average of seven to eight minutes per patient, this is an ambitious proposal.[17] But something like it will be necessary to deal definitively with the propaganda onslaughts of – for instance – the anti-vaccine lobby.

It is important, too, to understand the multiplicity of methods available to those propagating lies. In China, state-sponsored commenters – millions of them – fabricate about 448 million social media posts a year. But, as one study has shown,

the Chinese regime's strategy is to avoid arguing with skeptics of the party and the government, and to not even discuss controversial issues. We infer that the goal of this massive secretive operation is instead to regularly distract the public and change the subject, as most of these posts involve cheerleading for China, the revolutionary history of the Communist Party or other symbols of the regime.[18]

If distraction can be the enemy of truth, it follows that its protectors must engage in the battle for attention. It is not enough to issue a press release, appear on a news

channel or tweet a correction. The *means* of correction have to match the prevailing culture. A viral podcast, a demonstration or an online petition may do more to banish a falsehood than a straightforward assertion of fact. This is a slippery slope, of course: an endless battle of distraction and counter-distraction would do nothing for democratic discourse. Veracity must never be compromised by theatricality. But it is naive to think that the battle against Post-Truth will be won with sole recourse to the routine techniques of verification.

TRUMP THE NARRATIVE

Progress is sequential: which is to say, those who hope for change, or to combat a pernicious social trend, must adapt with steely discipline to the circumstances in which they find themselves. This is less obvious than it seems. There is a powerful instinct simply to reinstate that which is lost or jeopardised, to reassert the status quo ante. But, as I said at the beginning of this chapter, those who wish to defend Enlightenment values in this transformed context – hectic mobility, technological revolution, emotional ferment – must operate within its parameters. All else is delusion.

In his book *The Myth Gap*, Alex Evans argues that 'we need new myths that speak about who we are and the world we inhabit'.[19] It has, of course, become commonplace to use the word 'myth' as a synonym for 'familiar

A THEORY OF RELIGIOUS RHETORIC IN AMERICAN CAMPAIGNS

> *Beyond all differences of race or creed, we are one country, mourning together and facing danger together. Deep in the American character, there is honor, and it is stronger than cynicism. And many have discovered again that even in tragedy—especially in tragedy—God is near. In a single instant, we realized that this will be a decisive decade in the history of liberty, that we've been called to a unique role in human events.*
>
> —President George W. Bush, 2002 State of the Union Address

During the 2004 presidential election, voters chose between candidates advocating starkly different approaches to a myriad of issues of national consequence. The United States was entangled in two costly wars and was still feeling the effects of the September 11, 2001, terrorist attacks. Domestically, President George W. Bush and the Congress had just passed major prescription drug reform, enacted controversial tax cuts, and legislated dramatic changes to American education. Yet in the aftermath of the 2004 presidential election, many political observers roundly concluded that Bush's reelection was not due to any of these factors but was largely the product of Americans' concern with *moral values*.

In the months before the election, political observers were already predicting that existing religious and moral cleavages might decide the day. One *New York Times* headline read "Battle Cry of Faithful Pits Believers against the Rest" (Kirkpatrick, October 31, 2004, 24). The *Chicago Tribune* reported that "this presidential campaign had become one of the most spiritually saturated in memory with people of faith bombarded with entreaties from Republicans and Democrats" (Anderson, November 4, 2004, C1). The significant role of religion in the election gained considerable support from the Election Day exit polling, illustrating a substantial "God gap" between religious and secular

voters. George W. Bush received 64 percent of the vote among those attending religious services more than once a week, whereas Kerry received 62 percent support among those never attending services. Moreover, fully 22 percent of the voting public responded that "moral values" were the most important issue facing the nation, a group of voters that swung decidedly toward the Republican Party (however, see Hillygus and Shields 2005). Asked to interpret this statistic in a *Meet the Press* interview shortly following the election, Karl Rove characterized these voters as a group of Americans most concerned about a certain "coarseness of our culture." Voters, Rove argued, saw in President Bush the "vision and values and ideas that they supported" (*Meet the Press* 2004).

This was not the story of the 2008 election, however. In that election, voters gave comparatively little weight to religious or moral considerations. The day before the election, Stephen Prothero, a religion scholar at Boston University, editorialized that "much of the energy that Democrats and Republicans alike have pumped into the religion question seems to have dissipated. Voters tomorrow will be thinking more about the economy, health care and war than about the social and sexual issues that preoccupied 'values voters' in the 2004 election" (Prothero, November 3, 2008, 15A). Exit polls were consistent with this assertion. The Republican support among those attending church more than once a week had been reduced from 64 percent in 2004 to 55 percent.[1] The Democratic candidate, Barack Obama, actually won the vote among those attending church "monthly," a demographic group that John Kerry had failed to capture four years earlier. A Pew study published immediately after the election concluded that, although sizable religion gaps persisted, "Among nearly every religious group, the Democratic candidate received equal or higher levels of support compared with the 2004 Democratic nominee, John Kerry" (Pew Research Center for the People and the Press 2008). And, whereas *religion* and *values* were the hot topics on *Meet the Press* following the 2004 election, these words were not even mentioned on the 2008 post-election roundtable of the program.

Why was religion the story of the 2004 election but not in 2008? The difference in electoral dynamics is puzzling for two reasons. First, in both elections there were other, deeply salient, competing issues that may have distracted voters. Whereas the economy and the first viable African

American candidacy may have diverted attention away from religion in 2008, the economy was also highly salient in 2004, along with terrorism, two major wars, and tax cuts.[2] But despite these strong similarities in salient secular issues, religious cleavages decided the day in 2004 but not in 2008. Second, there are plenty of reasons to suspect that religion should have been even *more* important in the 2008 election. For example, 2008 had Sarah Palin, a vice presidential candidate who was in part selected to bring moral and cultural issues to the forefront. Moreover, 2008 was witness to one of the most intense religious campaign issues in recent memory when Obama's former pastor, Reverend Jeremiah Wright, made controversial remarks at the intersection of religion, race, and politics, arguing that "The government gives [African Americans] the drugs, builds bigger prisons, passes a three-strike law and then wants us to sing 'God Bless America.' No, no, no, God damn America, that's in the Bible for killing innocent people" (quoted in Murphy 2008). Given all this, it is surprising that religion factored more prominently in the public consciousness in 2004 than in 2008.[3]

We know relatively little about why the influence of religion waxes and wanes from one election to the next, raising ambiguities that operate at multiple levels. At one level, the intermingling of religion and electoral politics raises important normative questions about the nature of political representation in a country characterized as having a "wall of separation" between church and state. When religion is a factor in an election, how should leaders deliver representation to their constituents? Moreover, given the tremendous religious diversity of Americans, is a genuinely inclusive religious representational style even possible?

At another level, religion plays an ambiguous role in U.S. elections because of the varied forms that public religious expression can take. Historically, religious political rhetoric can be roughly classified into two genres. *Culture war* religious expression generally focuses on deep-seated religious differences in American society and the intractable political conflicts produced by these divisions (Hunter 1991; Evans and Nunn 2005).[4] *Civil religion* appeals, on the other hand, are nondenominational declarations of spiritualized American national identity. Civil religion appeals generally stress points of spiritual commonality among all Americans and posit a transcendent religious ethos that permeates American institutions and culture (Bellah 1967, 1975). Despite volumes of research on

these rhetorical genres, we know relatively little about how civil religion and culture war messages are actually received by the public at large. When candidates deploy religious messages, do divisions emerge, or are religious appeals a cultural glue uniting Americans across diverse backgrounds?

In this book, I am principally concerned with the ambiguous role of religious expression and how it comes to shape the American politics. Grappling with this issue will not only explain the role of religion in the 2004 and 2008 presidential elections—it will also help make sense of the meaning of religious and cultural divisions in a country premised on church-state separation. The way in which political elites use religious rhetoric in the public sphere determines the exact role that religion plays in American elections, political culture, and the representative dynamics of the country. At the heart of the argument is the observation that nearly all forms of religious rhetoric can be understood in terms of how they express themselves along two key dimensions: *emotion* and *identity*. Religious rhetoric gains its unique political command because it is well equipped to resonate with individuals' emotions and identities—two factors that, not coincidentally, are central to political persuasion.

The extent to which an election takes on a religious character depends on how successfully elites use religious language to activate emotions and identities. This analysis is not limited to 2004 and 2008. Indeed, I contend that religious rhetoric is an evolving genre that has its roots in early Puritan sermonizing and Revolutionary pamphleteering. The success of the genre is due to its use having been so congruent with basic psychological persuasive properties and its being flexible enough to fit with the religious sensibilities of an incredibly varied religious constituency. The religious character of American politics—both now and throughout history—depends on how well religious political rhetoric activates the emotions and identities of a diverse and deeply religious public. Moreover, the activation of religious considerations has consequences that extend far beyond electoral outcomes. The religious nature of political debate and discussion ultimately shapes the nature of political representation and the contours of the political community. As we will see, American civil religion has a special place in this story. Civil religion rhetoric can simultaneously be an electorally powerful persuasive tool and point of shared religious identification, and also a source of alienation from the political process.

THE PERSUASIVE APPEAL OF AMERICAN CIVIL RELIGION

In this chapter, I elaborate a theoretic framework for how religious rhetoric intensifies the emotions and identities of the American public and what the consequences of these rhetorical choices are. Doing so requires knitting together several diverse strands of scholarship, encompassing research on religion, voting behavior, political communication, and social psychology. The basic argument, illustrated in figure 1.1, has several moving parts. The religious character of American politics is shaped by a confluence of three factors: the religious makeup of the U.S. electorate, the psychological basis of persuasion, and the political demands imposed by competitive winner-take-all elections. Religious rhetoric (particularly the civil religion genre) is uniquely adept at satisfying the demands imposed by each of the factors, enabling candidates to form deep connections with voters in competitive electoral environments. Moreover, when specific forms of religious rhetoric are strategically deployed by candidates, the consequences extend far beyond the realm of any single electoral contest. Precisely *how* candidates locate themselves at the intersection of these factors not only determines their electoral success but also has implications for defining the boundaries of
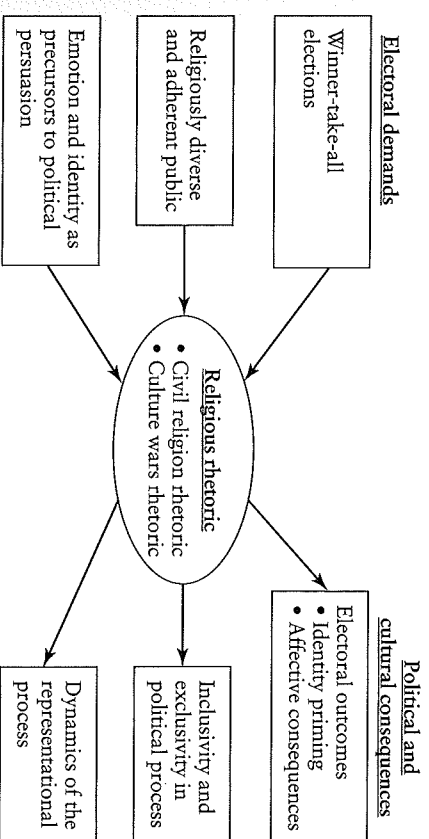

Electoral demands

- Religiously diverse and adherent public
- Winner-take-all elections
- Emotion and identity as precursors to political persuasion

Religious rhetoric
- Civil religion rhetoric
- Culture wars rhetoric

Political and cultural consequences

- Electoral outcomes
  - Identity priming
  - Affective consequences
- Inclusivity and exclusivity in political process
- Dynamics of the representational process

FIGURE 1.1 Causes and consequences of religious political rhetoric
How religious rhetoric connects electoral demands to political and cultural consequences. The dominant forms of religious political expression are the results of a confluence of political, religious, and psychological factors. In turn, these forms of expression have consequences on electoral behavior and American political culture.

the American political community and for shaping the dynamics of the representative process.

Psychological, Religious, and Political Factors in American Elections

As figure 1.1 illustrates, candidates making their case before the American public must deal with multiple crosscutting pressures, the first of which is political. Most U.S. races are winner-take-all, meaning that to hold office a candidate must win a plurality of the votes (or a majority of electoral votes at the presidential level). Unlike many other electoral systems, seats are not allocated for second place. Accordingly, candidate rhetoric must appeal to an audience that holds a diverse array of religious beliefs. The United States is unique among world democracies in this regard, having both high levels of religious adherence and no single dominant sect (Greeley 1972; Pew Research Center for the People and the Press 2007; Putnam and Campbell 2010). Thus, candidates cannot afford to ignore religion, nor can they afford to privilege a particular faith tradition.

These political and religious pressures have an interesting point of intersection with what is known about the psychological basis of political persuasion. As previously noted, political psychologists have identified two factors—identity and emotion—that play a central role in how voters think about political candidates. While many factors can awaken emotions and identities in the public, religious appeals are particularly well suited to this task and are at the same time capable of effectively satisfying the competing political and religious pressures incumbent on candidates.

By *identity*, I refer here to "social identities," or individuals' awareness of objective group membership and the sense of attachment they get from belonging (Tajfel 1981; Conover 1984, 1988). Although "being religious" need not imply social identity as a matter of definition, scholars have recognized that religion does play important identity-relevant functions (Emmons and Paloutzian 2003).[5] Social identities have been found to have numerous political implications, the most consequential and widely replicated being ingroup favoritism, even when group attachment is fairly minimal (Huddy 2001). That is to say, even when group boundaries are arbitrarily assigned, individuals still tend to demonstrate a persistent bias toward their own group. And, because religious group attachment is far from arbitrary, religious social identities should engender substantial favoritism toward those with whom the individual shares group membership.

Candidates commonly use political rhetoric to prime identities. *Priming* refers to "provok[ing] opinion or behavior change not because in-

dividuals alter their beliefs or evaluations of objects, but because they alter the relative weight they give to various considerations that make up the ultimate evaluation" (Mendelsohn 1996, 113). In any given election, voters have numerous competing considerations, from issues to images to social group memberships (Valentino 1999; Druckman and Holmes 2004; Druckman, Jacobs, and Ostermeier 2004; Jackson 2005). By rhetorically emphasizing social identities (or any number of other considerations), candidates make these group attachments a salient basis of political evaluation. In this way, expressions of religious identity can create deep feelings of group favoritism between candidates and the public.

Of course, candidates can make direct appeals to denominational subgroups (such as Catholics and Baptists) even though, with no single dominant religious denomination in the United States, subgroup appeals have a somewhat limited audience. What is more likely is that religious identity priming operates by engendering a sense of *civil religion identity*. The concept of American civil religion asserts that a broad religious identity unites virtually the entire nation. In this way, civil religion appeals should theoretically serve as a solution to the challenge of appealing to a religious constituency that is both committed and diverse. With the possible exception of appeals to national identity, no other group-based appeal (e.g., to race, gender, or class) has the potential to codify political support around such a broad (yet salient) group. Thus, when Bush said (see quotation at the outset of this chapter) that "God is near. In a single instant, we realized that this will be a decisive decade in the history of liberty," that we've been called to a unique role in human events," he was asserting a leadership role in an overarching spiritual community.

Scholars emphasize different points regarding precisely how public expressions of civil religion should be characterized and where it gains its cultural and political significance. For example, Martin Marty (1987) identifies both "priestly" civil religion, which is primarily concerned with legitimizing state practices, and "prophetic" civil religion, which seeks to guide the nation to meet certain ethical benchmarks (see also Wald and Calhoun-Brown 2006). Whereas Marty ascribes an ongoing sociological significance to both forms of civil religion in American politics, others have suggested that the cultural force of civil religion has declined since a peak in the 1950s (Ahlstrom 1972; Marty 1987). Other scholarly debates revolve around whether civil religion sits in tension or in harmony with religious pluralism. Some argue that civil religion can unite common elements of different religious traditions; others contend that it carries with it an implicitly sectarian

impulse (Lambert 2008, 26–27; see also Mead 1974). Taking up this latter point, Herbert Richardson concludes that, in the end, "civil religion always tends to generate the very situation it seeks to prevent" (1974, 165).

I address these debates in this book. The theory of identity priming, however, suggests a somewhat different starting point for grappling with language of civil religion and its consequences for American politics and culture. Specifically, understanding religious identity priming requires that we draw a clear connection between rhetoric in the public square and its influence on the attitudes and opinions of the electorate. What considerations are primed by these nondenominational appeals? Moreover, how do these appeals shape political behavior and political life? Civil religion appeals are quite common in political rhetoric, and evidence suggests that they should have broad appeal to religiously diverse constituencies. Indeed, the perseverance of the genre throughout American political history is a testament to the degree to which it has mass appeal (Bellah 1967). If civil religion references tended to fall flat, it is unlikely that candidates would still be making references to America as a "City on a Hill" and asserting a shared spiritual bond (Wimberley 1980). It should be noted that just because the genre does not appeal to a readily definable group (such as Methodists) does not mean that it cannot activate group identity. Michael Billig (2003) argues that adept rhetoricians often attempt to foster a sense of shared identity with the members of their target audience by linking them with cherished national values. Even the "banal use of political clichés" and strategic deployment of pronouns like "we" can cement ingroup allegiances and commitments to group values (Billig 2003, 238). Because these values (in the American case) are often religious or quasi-religious, it makes sense that civil religion appeals are amenable to the activation of group identity. Empirical evidence has documented the political significance of these broad identity appeals (Gaertner and Dovidio 2000). John Transue (2007), for example, has found that national identity can supersede the effects of subgroup attachment. In short, public figures' use of language that yokes together religion and country has a dramatic impact on the "self-image" of the public and ultimately "what it means to be an American" (Domke and Coe 2010).

In stark contrast, culture wars appeals drive a wedge into the American public, asserting that there are exactly two religious groups in American politics and that they are locked in an intractable political conflict over the moral standing of the nation.[6] The potential target membership of culture wars appeals is smaller than civil religion appeals, making it

unlikely that this rhetorical style will be deployed with the same regularity as civil religion rhetoric, despite fears by many scholars that a culture war is on the rise (Hunter 1991). Nevertheless, the stakes are always high in culture wars rhetoric, suggesting that, although they might not mobilize support in a pluralistic state as effectively as civil religion rhetoric, self-identification with orthodox or progressive camps will still be deeply felt and politically salient.

Culture wars and civil religion appeals thus implicate different understandings of religious identity. Each also carries a specific emotional tenor, ranging from enthusiasm to anxiety to anger. Understanding the tenor of religious rhetoric is important because emotions are known to have significant consequences on political judgment (Marcus, Neuman, and MacKuen 2000) and because political messages are, at least in part, responsible for bringing about these emotions in message recipients (Brader 2006).

Emotions work in two principle ways. They can work directly, by transferring the emotive content of a stimulus onto a message recipient. Simply put, if a candidate puts you in a positive mood, you will like her more, and the opposite is also true (Ladd and Lenz 2008). Emotions can also work indirectly by altering the decision-making process (Marcus, Neuman, and MacKuen 2000; see also Schwarz, Bless, and Bohner 1991). Psychologists argue that emotions play an evolutionarily adaptive role, and accordingly, different discrete affective states have arisen to meet specific situational demands (Bodenhausen, Sheppard, and Kramer 1994; Nabi 1999; Lerner and Keltner 2000; Marcus, Neuman, and MacKuen 2000; DeSteno, Petty, et al. 2004; Brader, 2006). For example, enthusiasm cues tend to be consistent with heuristic processing or with reliance on stable preexisting political affiliations (Brader 2006). Negative emotions often do the opposite. Norbert Schwarz concludes that, "In a nutshell, we usually feel bad when things go wrong and feel good when we face no particular problems. Hence, negative affective states may signal that the current situation is problematic and may hence elicit a processing style that pays close attention to the specifics of the apparently problematic situation" (2000, 434). Not all negative moods lead to systematic processing, however. For example, Galen Bodenhausen, Lori Sheppard, and Geoffrey Kramer (1994) find evidence that, whereas sad individuals tend to engage in effortful processing as a means to alleviate the sad situation, angry individuals tend to engage in heuristic processing due to reduced cognitive capacity and reduced motivation for thoughtful analysis.

Likewise, David DeSteno, Nilanjana Dasgupta, and colleagues find that "anger, because of its basic association with intergroup competition and conflict, evok[es] a psychological readiness to evaluate outgroups negatively vis-à-vis ingroups, thus creating an automatic prejudice against the outgroup from thin air" (2004, 323).

Anxiety has also been theorized as a distinct emotional state associated with a distinct processing style.[7] Ted Brader argues that "fear ads" used in political campaigns elicit anxiety in individuals, causing them to "place less weight on prior convictions and more weight on contemporary evaluations" (2006, 182), a conclusion that is consistent with the Affective Intelligence model of George Marcus, W. Russell Neuman, and Michael MacKuen (2000). For example, enthusiastic individuals tend to rely on heuristic judgments in their reasoning about candidates, whereas anxious individuals tend to engage in a deeper, more effortful information search (Brader 2006). In short, by making voters anxious, angry, enthusiastic, sad, or calm, candidates may be activating a number of psychological processes that are important to how voters think about candidates.

More important, the civil religion and culture wars genres are each closely identified with a highly emotive communication style. Culture wars rhetoric, for example, regularly uses anxiety and fear to characterize competing worldviews. The civil religion tradition, in contrast, is characterized by its hopefulness and optimism about the future of America, as well as a lament about U.S. moral shortcomings (Murphy 2009). Thus, given this close association between the dominant forms of religious political communication and emotion, on the one hand, and the importance of emotion in political persuasion, on the other, it is important to investigate how religious rhetoric influences the public mood of the electorate (Rahn, Kroeger, and Kite 1996).

The Consequences of Religious Rhetoric

It is clear that political, religious, and psychological factors combine to make religious rhetoric an optimal strategy for ambitious public elites, but in this book I am not concerned just with why various modes of religious communication are employed but also with the consequences of religious rhetoric in the public sphere. As figure 1.1 illustrates, three consequences of religious rhetoric are of particular interest: its impact on electoral outcomes, the contours of political community, and the dynamics of political representation.

First and foremost, religious appeals are common because they work. In this book, I provide evidence that religious rhetoric is used to activat[e] religious identities as a basis of candidate evaluation and to elevate th[e] emotive tone of campaigns. But the connection between rhetoric and vo[t]ing behavior is complex and sometimes counterintuitive. For examp[le,] although it is often suggested that Bush garnered favor among religiou[s] constituencies because of his stance on issues such as same-sex marriag[e] in chapters 5 and 6 I provide evidence that, instead, Bush's success in 200[4] had to do with his effective use of civil religion appeals, causing much o[f] the electorate to evaluate him on his ability to provide moral leadershi[p] for the nation. These patterns of voting behavior were ultimately lodge[d] in a sense of shared religious identity with Bush—not any particula[r] affinity with his stance on the so-called cultural issues. In this way, scho[l]ars and pundits bemoaning the rise of a culture war have missed a critica[l] component of religious rhetoric. What made Bush effective with religiou[s] audiences (at least through the 2004 campaign) was his ability to solidif[y] the support of a rather diverse group of believers rather than the use o[f] religious rhetoric calling attention to cultural differences.

If citizens are voting for candidates based on a shared religious iden[t]ity and emotional arousal, this raises a corresponding set of question[s] about how religious rhetoric influences American political culture. Tha[t] is, if religious identity is part and parcel of political identity, then religio[n] may be playing a role in who is and who is not included in particula[r] visions of the American political community. Religious rhetoric sets th[e] tone for political debate and discussion, and whether it is conducted in a spirit of cohesion or competition, optimism or anger. Thus, candidates' religious discourse is at the heart of questions of inclusion and exclusion in the political community.

Perhaps even more critically, insofar as religious communication influences the voting behavior of the mass public, it is also influencing the representational activities of leaders, guiding how they govern based on the substantive and symbolic demands of various religious constituencies. In a country premised on an ostensible wall of separation between church and state, it is important to explore what consequences free religious expression has on promoting a religious mandate for officeholders, once elected. If religious identities are able to sway elections, how are officeholders to govern a constituency that is religiously diverse?

Understanding the dynamics between would-be representatives and the governed ultimately requires a deeper theoretical elaboration of the

recipients of religious messages—the voters. Rather than treating the religious vote as monolithic and static, it is more appropriate to think of the numerous and varied religious identities that could potentially be activated by candidate rhetoric. Although it is common to think about direct appeals to different denominations or to the religiously progressive and orthodox, we need to take stock of an additional politically salient religious identity, informed by the importance of civil religion rhetoric in American politics. Specifically, I articulate the concept of a civil religion identity, arguing that identification with this nondenominational American spiritual community is key to understanding the dynamics of political representation in America. Civil religion identifiers hold a deep sense of attachment to an explicitly spiritualized understanding of America. For civil religion identifiers, the United States—vested with a sacred sense of purpose in the world order—is as much a religious community as a political entity. As the psychological analog to civil religion rhetoric, civil religion identity provides the key to understanding how commonplace religious rhetoric can turn the attention of voters to religious evaluative criteria. Many Americans strongly identify with the basic tenets of American civil religion, and in fact for many a latent attachment to this quasi-religious identity is a foundational component of their political DNA. The evidence indicates that civil religion voters are not motivated by cultural issues such as abortion and same-sex marriage, issues which often dominate political debate.[8] Instead, civil religion voters are looking for a prototypical leader who offers a representational style consistent with the tenets of American civil religion. In this sense, religious representation may not be a mandate for policy change but, rather, an endorsement of leadership offering a nondenominational spiritualized sense of the place of America in the world order.

This model of religious rhetoric has consequences, not just for our understanding of religion but for how we think about political campaigns and political behavior more generally. Scholars have long questioned whether campaigns can substantively change Americans' attitudes, especially when the factors that often influence electoral decision making (such as party identification) tend to be stable from election to election. The model I present here suggests that stability and change are not mutually exclusive. Stable predispositions such as religiosity might not fluctuate much in the American public, but campaign rhetoric can certainly play a role in activating these predispositions and making them germane to the task of evaluating candidates. Moreover, campaigns do not merely

mobilize fixed constituencies. Instead, who is and is not a member of a particular constituency is itself the product of rhetorical forces. Instead of speaking just about mobilizing constituencies, we also need to think of campaigns as activating different parts of the individual, changing the American consciousness about group membership and group values, and making different identities politically salient. Even though there is no formal membership in the civil religion community, civil religion is something that is very real to many Americans. Civil religion exists as an identity that is itself actively reified through political rhetoric and made salient to the electoral process.

OVERVIEW OF THE CHAPTERS

In this book, I establish a number of conclusions about the impact of religious rhetoric on American politics, ranging from broad historical judgments about the place of civil religion in American culture to narrowly tailored arguments about the psychology of religiosity and vote choice. The key premise is that these macro and micro approaches are mutually reinforcing. It is impossible to understand the success of religious appeals across history without thinking critically about the psychological processes that undergirded the persuasive efforts of historical figures, and it is impossible to understand present-day patterns of religious voting without placing the campaign rhetoric in its proper historic context.

Accordingly, the evidence presented in this book unfolds at several levels of analysis, moving from a historical examination of religious rhetoric, to a content analysis of religious rhetoric in contemporary campaigns, and finally to an examination of individuals' voting behavior using surveys and experiments. Although the scope of the book changes, the evidence consistently indicates that religious rhetoric is effective primarily because of its ability to induce a sense of *shared identity* and *emotions* in message recipients. In chapters 2 to 4, I focus on identity-laden and emotive cues in rhetoric, addressing exactly how public figures craft religious expression and what this rhetorical nuance says about American political culture. In chapters 5 and 6, I shift the scope of the argument from identity cues in rhetoric to religious identity as experienced by individuals and from emotive language to emotions in the mass public. Ultimately, the argument is consistent across levels of analysis. Emotion and identity are both important components of the terms and tenor of

political debate, and they are important parts of how individuals understand the political world.

Chapter 2 begins by addressing emotion and identity from a historical vantage, examining the use of religious rhetoric in American politics from early Puritan political communities through the twentieth century. By examining the evolution of religious rhetoric over time, I provide insight into how religious rhetoric is constitutive of American political culture and how it is used politically across contexts. My principal argument here is that emotion and identity have been central elements in religious rhetoric throughout American history (although how they have been invoked has evolved over time). They are neither fleeting elements nor elements emerging only in contemporary politics, and they have consistently played a prominent role in shaping American political discourse. This finding has important consequences for the book as a whole, suggesting that, even amid substantial contextual variation, identity and emotion provide considerable insight into American political culture and political preference formation. Religious rhetoric is not joined at the hip to any one political issue or ideological outlook. It is a flexible genre that has been appropriated to fit numerous political causes.

Chapter 3 builds on this historical analysis by exploring the invocation of religious identities in modern campaigns. In this chapter, I present evidence that religious rhetoric is rarely concerned with taking stances on issues or rationalizing a complex policy agenda. Rather, religious rhetoric is principally in the business of building a sense of shared identity between citizens and candidates. Three types of identity references are common: *subgroup* references to specific denominations and faith traditions; *civil religion* appeals, aimed at engendering a spiritualized sense of national identity; and *culture wars* identities, which seek to make cultural fault lines salient. These identities have important consequences for how the boundaries of the American political community are characterized. A content analysis of over 1,300 speeches by Republican and Democratic presidential hopefuls from 1980 to 2008 indicates that candidates make frequent reference to a nondenominational spiritual identity that permeates the American experience. As we will see, the language of civil religion identity is far more common than the language of culture wars, indicating that American religious identity—at least in the realm of political discourse—is more a source of unity than a source of division. But even though these civil religion appeals typically downplay religious pluralism in the interest of emphasizing what faiths and peoples have in

common, they also often marginalize and alienate key segments of the American electorate. Ultimately, then, the rhetoric of religious identity—although central to understanding American politics—neither sows the seeds of a large-scale cultural battle nor constitutes a panacea for social cohesion.

In addition to identity, emotion is a key element of American religious political rhetoric. Chapter 4 begins by developing a strategy for identifying emotive elements in speeches. Using the same rhetoric database as in chapter 3, I conclude that candidates adopt specific emotive frames to make identity-based appeals. Religious rhetoric tends to be exceptionally optimistic—far more positive and hopeful than secular campaign speech. I also present evidence that is generally inconsistent with claims that there is a growing divisiveness in religious rhetoric; there is little evidence to suggest that presidential campaign religious rhetoric is being used as a tool to leverage large-scale cultural rifts. I do find, however, that there are significant partisan dimensions in the emotive characteristics of religious rhetoric, a conclusion that follows from the nature of existing religious divisions in the electorate.

Chapter 5 argues that religious rhetoric is one significant cause of the relationship between religion and politics in the mass public and that this process can be understood by examining both the specific qualities of religious messages and how varied message types interact with different religious predispositions. I use statistical tools to merge the religious rhetoric variables from chapters 3 and 4 with survey data collected in presidential elections from 1980 to 2004. The results indicate that how candidates craft the identity and emotive elements of religious rhetoric influences how they are viewed by the electorate. Specifically, different rhetorical nuances tend to activate different dimensions of individuals' religious orientations. For example, when a candidate speaks in the language of American civil religion, the religiously committed become more favorably predisposed toward that candidate, but the religiously orthodox actually tend to lower their opinion of the candidate. These findings are consistent with the hypothesis that identity priming and emotion are a basis for political persuasion and provide strong evidence that candidates' religious rhetoric is, in part, responsible for the relationship between voters' religious and political attitudes.

Chapter 6 uses experiments to add a layer of confirmatory evidence to the findings from chapter 5, and using the concept of civil religion identity, it extends the conclusions drawn in chapter 5 to provide a deeper understanding of political representation. Civil religion identity

has broad adherence in the American public, and political rhetoric often primes civil religion identity as a basis of candidate evaluation. Voters who strongly identify with the American civil religion are attracted to candidates who invoke this genre and see them as prototypical group leaders who share their moral outlook (although not necessarily their substantive policy outlook). The representational consequences of this identity, however, are crosscutting. Although theoretically civil religion identity is nondenominational, in practice its membership is essentially limited to Christians. Non-Christians find themselves unrepresented by candidates who embrace the tenets of American civil religion. The consequences of civil religion rhetoric are thus mixed—although it provides a form of meaningful leadership for many Americans, it leaves others feeling excluded from the political system.

Using corroborating evidence from multiple methodological approaches, I conclude that religious rhetoric is a central force responsible shaping the contours of American political culture. Religious rhetoric is also electorally consequential and culturally significant, with important implications for how we interpret American political representation. Even though its use has changed over time, it has been remarkably consistent in its ability to stir the emotions of the mass public and to engender a sense of shared spiritualized identity.

2

RELIGIOUS RHETORIC IN AMERICAN POLITICAL HISTORY

Let us resolve tonight that young Americans will always see those Potomac lights; that they will always find there a city of hope in a country that is free. And let us resolve they will say of our day and of our generation that we did keep faith with our God, that we did act "worthy of ourselves;" that we did protect and pass on lovingly that shining city on a hill.

—President Ronald Reagan, 1980

Religious political rhetoric can overwhelm citizens with an array of different emotions, leading individuals to identify with a broad and varied range of groups and interests. We know very little, however, about exactly which group identities and emotions religious rhetoric is bringing to the surface. Although it seems likely that many voters will have some sort of emotional response to a passage such as the Reagan statement quoted here, it is unclear precisely what kind of affective punch this statement will have on a religiously diverse public. Moreover, it is unclear exactly which religious identity Reagan is calling forth as a standard of political evaluation. Phrases such as "our God" implies an appeal to a religious group; however, the boundaries of this group are not easily identifiable. Thus, although there are strong theoretical (and intuitive) reasons to suspect that Reagan's city on a hill speech mattered politically, there is little research documenting how religious communication translates into political opinion.

Part of the reason this connection is so elusive is that religious identity is so complex. If candidates were only making simple appeals to denominational subgroups, connecting candidate rhetoric to voter attitudes would be a straightforward task. But there is nothing straightforward about religious appeals in American politics. Indeed, it is virtually impossible to make sense of the substance and the tenor of an appeal such as Reagan's without first understanding how nuanced rhetorical constructions

Although the U.S. Constitution erects a wall of separation between church and state, voters regularly bring religious standards to bear on political figures. The dynamics of this complex representational process depends on the religious standards in question, the voters' own religious identity, and the nature of candidate rhetoric. Although many have interpreted religious voting patterns as placing a substantive issue-based mandate on the officeholder, most religious voters are actually more concerned with a symbolic representational style. Indeed, the most common style of religious rhetoric—American civil religion—tends to direct many voters' attention to the religious image, not any particular religious policy platform. It does so by promoting the view that, in American politics, religious evaluative criteria are social acceptable and normatively desirable.

On one hand, these findings are an effective counterexample to the view that a rhetorical cultural war is driving cultural cleavages in the United States. Although self-identified religious fundamentalists may be overwhelmingly concerned with abortion and same-sex marriage, *fundamentalist* does not describe most of the American electorate. Thus, consistent with Geoffrey Layman and John Green (2005), we find that an issue-driven culture war is real but is being "waged by limited religious troops on narrow policy fronts" (83). On the other hand, the analysis also reveals deep tensions with respect to the way religious political rhetoric is normally expressed. Candidates' characterizations of America as a blessed nation resonate with many; but they leave many others unhappy with these candidates and less likely to weigh religious evaluative criteria. Thus, although civil religion is properly considered a broad superordinate identity, it is far from an identity with universal appeal.

7

THE RHETORICAL CONSTRUCTION
OF RELIGIOUS CONSTITUENCIES

With a good conscience our only sure reward, with history the final judge of our deeds, let us go forth to lead the land we love, asking His blessing and His help, but knowing that here on earth God's work must truly be our own.

—President John F Kennedy, 1961 Inaugural Address

Religious rhetoric is a defining feature of the American political campaign. Although the contours of the genre have changed over time, it contains two enduring elements that make it well suited to be a highly persuasive tool given the unique American religious landscape. Specifically, the genre is defined by the rhetorical expression of politically salient collective identities and the use of highly emotive rhetorical cues. By rhetorically leveraging emotions and identity, political elites have thus used—and will continue to use—this genre to their electoral advantage. The evidence is clear that, for vast segments of the American public, religious rhetoric is a desirable, if not necessary, component of a candidate's public self-presentation. Nevertheless, the prevalence of this genre in the public sphere produces crosscutting effects on the ability of candidates to deliver adequate representation to all constituents and on the meaning of an inclusive and tolerant democracy.

In this chapter, I explore the interconnections among religious persuasion, representation, and culture. The evidence suggests that to fully understand the politics of religious appeals we need to fundamentally retheorize the nature of religious constituencies. Religious rhetoric should be thought of less in terms of appealing to stable preexisting religious groups and more in terms of rhetorically activating latent religious identities. Of particular importance is the activation of an American civil religion identity. Even though the tenets of civil religion have a broad adherence in the American public, civil religion is also quite exclusive, putting forward an explicitly religious conception of American national

identity. Moreover, a candidate's invoking American civil religion recasts the relationship between the representative and voter, bringing symbolic representational demands to bear on the political process. In the end, *separation of church and state* may ultimately be a misnomer in American politics, in that religious rhetoric is responsible for actively creating religious constituencies that can drive election results.

RELIGIOUS RHETORIC AND POLITICAL PERSUASION

It is no accident that religious rhetoric is such a robust feature of American election campaigns. In American politics, religious rhetoric provides a unique solution to the convergence of three challenges faced by candidates. That is, in a winner-take-all electoral system, candidates need to develop a rhetorical style with broad appeal to a religiously diverse constituency that leverages the psychological underpinnings of persuasion. Religious rhetoric, as it has evolved across American history, sits comfortably at the intersection of all these forces.

Of particular importance is the civil religion tradition, which is used to activate a spiritualized sense of collective identity in the American public. As demonstrated in chapter 6, civil religion finds broad identification in the American public, and candidate rhetoric routinely makes this identity electorally salient. It should not be lost on us that civil religion identity gains its power from unique religious makeup of America. In a marketplace of competing religious traditions, civil religion identity attempts to unify an otherwise diverse set of religious affiliations and orientations. Along these lines, however, civil religion rhetoric is not the only choice for candidates. In different electoral environments, candidates have sought to make denominational identities salient and even to activate a schism between orthodox and progressive religious factions.

The key observation here is that religious constituencies are not permanent facts in American elections, defined by bright lines and intractable group allegiances. Rather, individuals have numerous and crosscutting religious identities, and religious political rhetoric works by strategically making these identities politically salient. Scholars of electoral behavior have long understood that individuals hold an array of competing considerations on matters of foreign and domestic policy, and which considerations are brought to the forefront of political evaluation has much to do with how skilled politicians make their case. The process of religious

identity priming follows a similar course. To fully understand the role of religion in voting, we should not ask just how campaigns activate religious groups in the electorate but also how campaigns activate different religious identities in the individual.

This has important implications for understanding and interpreting elections. It cautions us not to perceive the religious vote as monolithic but, rather, as a diverse group with multiple interests and desires. It also informs our understanding of candidate strategy. Previous research has found substantial evidence that candidates craft their rhetoric to strategically prime the issues on which they will be favorably evaluated (Jacobs and Shapiro 1994). The present research adds to this understanding of campaign dynamics by showing that candidates also actively construct a common group identity with voters in the electorate. This social group identity need not be formed around any particular issue in the way that farmers might unite around agriculture subsidy policies or union members might unite around changes in labor law. Identity itself can be grounds for persuasion—it need not have a substantive basis in political issues.

This ultimately may be the best explanation for the difference in the religious vote from 2004 to 2008. George W. Bush made significant gains with religiously committed voters during the course of the 2004 election, accentuating the already sizable religion gap enjoyed by Republican candidates. Bush's rhetoric deserves the credit for this. His religious self-presentation was not sectarian but, rather, cast American greatness in religious terms, thus appealing to the identities of a broad array of voters. Although John McCain used religious rhetoric, he did not do so in a manner consistent with the activation of religious identity. In contrast, Obama deftly primed civil religion identity, lamenting the loss of U.S. status in the world order while, at the same time, offering a promise of American greatness. Even though McCain still enjoyed an advantage among religious voters, Obama's rhetorical style was probably responsible for closing the large gap that had been present four years earlier.

FROM PERSUASION TO REPRESENTATION

After the election is over, the active yoking of religion to politics in political campaigns leaves an indelible imprint on the American social fabric. When candidates use religious rhetoric to actively promote civil religion

identity as a basis of vote choice, the effect is that many voters understand political leadership in symbolic terms. Given the right rhetorical cues, electoral behavior becomes more about a candidate's image of moral character and less about substantive issues. This is consistent with Pitkin's (1967) understanding of symbolic representation, in which political elites are conceived of as active symbol-makers rather than passive agents who simply stand for their constituents. Pitkin has serious concerns about this representational style, writing that symbolic representation is not "merely ritual activity. Rather, it is a kind of activity to foster belief, loyalty, satisfaction with their leaders, among the people.... Since there can be no rational justification of the symbolic representative's position as leader, the emphasis (as with symbols) must fall on the nonrational or emotive elements in belief, and on leadership techniques which exploit such elements" (1967, 107).

Given how subtle and ecumenical religious rhetoric often is in practice, "exploit" may be too strong of a word to describe most candidates' use of rhetorical style. Moreover, although religious rhetoric can shift vote choice, the evidence is clear that, even in religiously charged environments, many other factors (e.g., party identification) remain important. Nevertheless, Pitkin's analysis directs our attention to the normatively problematic aspects of a rhetorical style steeped in religious language. Voters may rush to the polling place unaware that their preferred candidate may not ultimately stand for (or even care about) their substantive interests. Religious rhetoric can thus create representational disjointedness, whereby purely symbolic behavior is taken to imply a substantive mandate.

On top of this representational challenge for those who do identify with civil religion is the question of representation for those who do not. As Murray Edelman notes, "Signs evoke an intense response only for those already taking the roles that make them sensitive to the cues that are given off" (1964, 122). The case of civil religion rhetoric is interesting in that it clearly does not stand for everyone but purports to do exactly that. This rhetorical exclusion has observable consequences. Many Americans simply do not feel represented by candidates who employ the language of civil religion identity. This represents a major challenge, especially considering the growing religious diversity of Americans' religious faiths. As a result, civil religion rhetoric has changed over time to become more inclusive, although exclusion is still a part of the genre. One representational challenge for political elites wielding the language of religious

identity is whether the genre can expand once again to offer an even more inclusive definition of the American civil religion.

All this paints a fairly negative portrayal of the representational thrust of American civil religion; however, this portrayal needs to be tempered by our recalling that, theoretically, civil religion exists as a solution to the complex representational challenge posed by a religious constituency that is both diverse and highly committed. For all its potential representational drawbacks, civil religion still stands as a clear alternative to the rhetoric of culture wars, which seeks to actively deepen religious differences. Political observers should take some comfort in the fact that the vast majority of religious rhetoric is not preoccupied with fostering deep societal divisions; in fact, the vast majority of religious language seeks to assert a point of shared collective identity. Moreover, as many have noted, civil religion also plays a role in providing basic political legitimacy for American institutions (Wald and Calhoun Brown 2006). It has helped to rhetorically construct a movement identity for diverse constituencies, as was the case with the Populist movement (Williams and Alexander 1994). And it has been used to direct national attention toward moral shortcomings and to urge reform, as Obama's campaign rhetoric frequently did.

In sum, civil religion plays a complex representation role, engendering a collective identity in a diverse public while, at the same time, directing the electorate toward largely image-based standards of political evaluation. Multiple layers of evidence indicate that, for all its purported attributes, civil religion rhetoric creates serious representational challenges in a pluralistic religious society. Ultimately, perhaps the genre defies an assessment painted with a broad brush. Civil religion rhetoric plays an important representational role, but in doing so it produces serious normative challenges that must be acknowledged as well.

FROM PERSUASION TO AMERICAN POLITICAL CULTURE

In appealing to voters, presidential candidates typically invoke one of three religious identities: subgroup religion references to specific denominations and faith traditions; references to American civil religion, which involves spiritualized, yet banal expressions of American national identity; and divisive references to a culture war in American political life. The expression of these religious identities has important consequences

for how faith and American national identity are understood in the public sphere. The concept of an American civil religion, for example, is often theorized to play a positive role in the maintenance of democratic institutions. Jean-Jacques Rousseau famously theorized that a "civil profession of faith" legitimizes democratic institutions (1762). In the American case, Tocqueville asserted a variant of the Rousseauian argument. Even though Americans, Alexis de Tocqueville argues, are divided into numerous sects, "they all see their religion in the same light" (1840, 449). The unity and the general moral wherewithal provided by religion is indispensable to American democracy: "Every religion . . . imposes on each man some obligations toward mankind, to be performed in common with the rest of mankind, and so it draws him away, from time to time, from thinking about himself. . . . Thus religious people are naturally strong just at the point where democratic peoples are weak" (Tocqueville 1840, 445).

If Tocqueville and Rousseau are correct, civil religion is indispensable to American democracy. At the same time, the evidence in this book raises serious doubts about whether the American civil religion casts a wide enough net to generate an inclusive and tolerant political culture. If it does not, civil religion may have a corrosive effect on religious freedom. As Rousseau (1762) himself recognized, civil religion is closely interwined with intolerance for those who do not ascribe to its basic tenets. In chapter 3, I have shown that civil religion rhetoric has potentially divisive undercurrents. Of its adherents, it requires not only a belief in God but a belief in a very specific conception of God. This detailed attention to God image is not just a theoretical exercise. Evidence presented in chapter 6 suggests that these rhetorical nuances have observable consequences for the American public, making some feel suspicious of and distant from political candidates. In this way, civil religion rhetoric, although politically persuasive, does not live up to its billing as a source of political cohesion.

Even though the cultural impact of civil religion is potentially problematic, it is important not to oversell the negatives. Consider Kennedy's remarks quoted at the outset of this chapter. Kennedy evokes this strong statement of American civil religion from the vantage of a religious outgroup facing substantial religious prejudice. Months earlier in his campaign, Kennedy worked to assure voters that "I am not the Catholic candidate for president. I am the Democratic Party's candidate for president, who happens also to be a Catholic" (September 12, 1960, Houston, TX, Annenberg/Pew Archive of Presidential Campaign Discourse). Rhetori-

cally, Kennedy supplanted fears that he would deliver descriptive representation that privileged Catholics with a symbolic alignment of wills. In this sense, the idea that leaders are active symbol-makers can be seen as potentially working to overcome virulent religious discrimination, insofar as the leader reconstructs the symbol in the spirit of civil inclusivity. Even though American civil religion can marginalizes some voters, history suggests that the genre is adaptive enough to rhetorically accommodate a growing diversity of citizens. Ultimately, the future of religious rhetoric will be assessed on these terms—on how well it manages to carve out a rhetorical space that accommodates diversity and promotes tolerance while still managing to offer meaning and vision.

70