INSTITUTE FOR THE FUTURE

THE BIOLOGY OF DISINFORMATION

memes, media viruses, and cultural inoculation

DIGITAL INTELLIGENCE FUTURES LAB

ABOUT THE AUTORS

Douglas Rushkoff is a professor of Media Theory and Digital Economics at CUNY/ Queens, is the bestselling author of *Throwing Rocks at the Google Bus*, *Present Shock*, and *Media Virus*, a pioneering 1994 study of memes and popular culture.

David Pescovitz is a research director at Institute for the Future, editor/partner at Boing Boing, and co-producer of the Grammy-winning Voyager Golden Record: 40th

Jake Dunagan is director of the Governance Futures Lab at IFTF. His work centers on social invention—helping people actively participate in the design and implementation of better systems with vision, boldness, and humor.

INSTITUTE FOR THE FUTURE

Institute for the Future (IFTF) is an independent, nonprofit 501(c)(3) strategic research and educational organization celebrating nearly 50 years of forecasting experience. The core of our work is identifying emerging trends and discontinuities that will transform global society and the global marketplace. Our research generates the foresight needed to create insights that lead to action and spans a broad territory of deeply transformative futures, from health and health care to technology, the workplace, learning, and human identity. As an open educational organization, IFTF strives to make all research and materials publicly available under the Creative Commons 4.0 International license.

OUR THANKS TO:

First Draft, a project of the Harvard Kennedy School's Shorenstein Center, formed as a nonprofit coalition with nine Founding Partners in June 2015 to raise awareness and address challenges relating to trust and truth in the digital age.

INSTITUTE FOR THE FUTURE
201 Hamilton Avenue
Palo Alto, CA 94301
www.iftf.com

Table of contents

Introduction

Political campaigns have always relied on values, visions, narratives, and ideologies to win votes. Whether virtuous or cynical, this effort comes down to propaganda: the leverage of social and psychological biases to promote a particular point of view.

Since Woodrow Wilson hired Walter Lippman to create the Creel Commission and win public support for our participation in World War I, social philosophers of all stripes have been debating the merits of manipulating people for political agendas. Lippman called for a "council of experts" to decide what would benefit the public, and an army of public relations specialists to convince them of what was in their own best interests. His protégé, Edward Bernays, took an even more cynical stance, arguing that the public was just too stupid to make informed choices. Elites should figure everything out, and treat the masses like Pavlov treated his dogs.

French legal scholar and social philosopher, Jacques Ellul, understood that propaganda was more than mere psychology. It depended not just on the emotional makeup of the individual, but the social context in which individuals live. In Ellul's words, "Propaganda is a set of methods employed by an organized group that wants to bring about the active or passive participation in its actions of a mass of individuals, psychologically unified through psychological manipulations and incorporated in an organization." Propaganda is meant to change the mindset of the public, so that they not only think but act differently.

Until recently, speeches, TV commercials, talk radio, flyers, and bumper stickers were the primary media through which propaganda could be disseminated. Sometimes, these traditional media campaigns rose to the level of psyops. America's invasion of Guatemala in the 1950's on behalf of United Fruit Company was falsely represented by Ed Bernays' television producers as a liberation of its people, and fake stories of babies being pulled from incubators—the invention of Hill & Knowlton public relations specialists—won public support for America's invasion of Iraq.

Regulation, competition, and a free press are all meant to protect the public from such manipulation, and these safeguards do work to varying extents. But a new breed of high-tech tools for political persuasion have emerged that challenge traditional approaches to public informational health. Politicians, activists, and state actors have attempted to harness media viruses, bots, and computational propaganda to manipulate minds and sway public opinion, often in secrecy, and on a scale unimaginable to previous media watchdogs.

Social media—the new, horizontal landscape in which these lateral and "egalitarian" forms of communication take place—touches not only more than 65 percent of Americans who regularly use it, but also the mainstream media who breathlessly regurgitates its Tweets and updates to the rest of the world. Speed and exclusivity have always been the surest paths to a competitive advantage in journalism, and these abilities are tremendously amplified by digital communications platforms. As a result, today's propagandists understand social media, in particular, as the breeding pool and lifeblood for its most contagious constructs, as well as the delivery system into public opinion at large.

Current investigations are focused on the source of various memetic campaigns, the nations or parties responsible, and the extent to which digital platforms and social networks can be considered accountable for their dissemination. Some groups are going so far as to suggest new, good algorithms to counteract the effects of the bad ones.

Which is the most effective approach to restoring the integrity of public discourse in an age of weaponized memetics: better technological protections, or a more resistant social psyche?

In this research effort, we are focusing instead on the greater shift to a memetic landscape—no matter the origins and sponsorship of particular memes—and the impact of memetic activity on the media, social, and political environment. What does the migration from broadcast propaganda to social and memetic propaganda do to the social organism and its resistance to manipulation? Can technological fixes and government regulations adequately address the problem of propaganda in a computational environment, or must we look at ways to promote a more resilient social fabric? In short, which is the most effective approach to restoring the integrity of public discourse in an age of weaponized memetics: better technological protections, or a more resistant social psyche?

Know Your Meme

In the last few years, the "meme" itself has infected mass culture. Our social feeds spew images with funny block-letter captions, cute cats, and looping animated GIFs of micro-moments culled from YouTube videos—what are commonly known as memes.

Although they've become a slippery topic themselves, for our purposes, memes are really any discrete unit of culture: Madonna, Viagra, lolcats, Occupy, democracy, rampage killings, and, yes, Donald Trump. The easiest way to grasp the difference between a meme and anything else is to ask, does it want me to share it with someone else? If it's an idea or image that seems to ask, "make another one of me" or "pass it on," then it's a meme.

Like its biological analog—the gene, a meme's success is based on its ability to replicate and spread.

For example, as the promulgators of memes on the Internet quickly learned, the image/text format spreads well on social media platforms such as Facebook, where images and video are algorithmically favored over plain text. For example, according to the Columbia Journalism Review, "on Breitbart News's Facebook page, images and videos are overwhelmingly more popular than links. Images made up just 5 percent of Breitbart's total posts in 2016, but they accounted for half of the page's most-shared posts."

"Conservatives often ask, 'How can I help break Establishment Media's stranglehold and get the truth to Americans who need it?'" explained Breitbart News editor-in-chief Alexander Marlow. "The answer is simple and powerful: follow us on Facebook, Twitter, and now Instagram and share and blast the digital 'bullets' we provide for you far and wide. It works … "[1]

Such statements seem to stand in stark contrast to the benign origins of memetics, a biologically inspired model of how information is transferred through culture. The word meme was coined as a nod to the gene, whose

Darwinian struggle leads to a more highly evolved organism. While the scientific community may still deride drawing an equivalence between evolutionary biology and social contagion, the practitioners of propaganda have embraced memetics wholeheartedly. Memetic has been commercialized, and as of the last several election cycles, thoroughly weaponized.

The specter of widespread computational propaganda that leverages memetics through persuasive technologies looms large. Already, artificially intelligent software can evolve false political and social constructs highly targeted to sway specific audiences. Users find themselves in highly individualized, algorithmically determined news and information feeds, intentionally designed to: isolate them from conflicting evidence or opinions, create self-reinforcing feedback loops of confirmation, and untether them from fact-based reality. And these are just early days. If memes and disinformation have been weaponized on social media, it is still in the musket stage. Sam Woolley, director of the Institute for the Future's (IFTF) Digital Intelligence Lab, has concluded that defenders of anything approaching "objective" truth are woefully behind in dealing with computational propaganda. This is the case in both technological responses and neuro-cultural defenses. Moreover, the 2018 and 2020 US election cycles are going to see this kind of cognitive warfare on an unprecedented scale and reach.

But these mechanisms, however powerful, are only as much a threat to human reason as the memetic material they transmit, and the impact of weaponized memetics itself on the social and political landscape. Memes serve as both probes of collective cultural conflicts, and ways of inflaming social divisions. Virulent ideas and imagery only take hold if they effectively trigger a cultural immune response, leading to widespread contagion. This is less a question of technological delivery systems and more a question of human vulnerability. The urgent question we all face is not how to disengage from the modern social media landscape, but rather how do we immunize ourselves against media viruses, fake news, and propaganda? To do this, we need to understand the biology of disinformation.

Our Mediascape

Leading up to the 2016 US presidential election, social media bots drove deceptive campaigns on Twitter and Facebook, creating the appearance of an electorate significantly more engaged and extremist than it really was. Researchers at University of Illinois at Chicago found that bots generated approximately twenty percent of all Tweets.[2]

Propaganda bots are used to share pro-regime, pro-candidate, pro-policy, or anti-opponent messaging. Their purpose is simple: shape faux popularity or the illusion of social proof on platforms like Twitter, Google, or Facebook, gaming the algorithms to surface content in favor of a specific candidate, regime, or policy. This may be intended to promote a particular policy, or to serve a bigger psyops goal, such as increasing divisiveness, exacerbating extremism, or simply eroding public trust in both facts and reasoned debate.

Researchers at the Oxford Internet Institute found that bots generated more than one-third and nearly one-fifth of Twitter activity supporting Donald Trump and Hillary Clinton, respectively, during the first and second Presidential debates. Further, while the social media campaigns of both Clinton and Trump were highly automated, additional research from Oxford indicated that pro-Trump bots outpaced pro-Clinton bots five-to-one from the day of the final debate to Election Day.[3]

Meanwhile, astroturfing bots autonomously retweet, like, or share content with the intent of giving the impression of authentic, grassroots support for a cause or candidate. “In Brexit, we saw one percent of the profiles on Twitter accounting for seventy percent of the traffic in support of the Leave campaign,” determined Woolley, former director of research at the Oxford Internet Institute. “Most of these accounts were highly automated, tweeting hundreds of thousands of times.”[4]

And as tools to build and deploy bot armies proliferate, we must come to understand how the messages they carry are transmitted through our social sphere, and then learn to defend against them. Pervasive social networks offer researchers a massive test-bed to develop and validate theories from memetics, an understanding of cultural information transfer modeled on evolutionary biology.

In the 1976 book *The Selfish Gene*, biologist Richard Dawkins popularized the term “meme” (analogous to “gene”) as a “unit of culture” that can spread and replicate like DNA, from person to person, throughout a culture. Memes start in our brains but travel out from there into the information ecosystem. Until fairly recently in human history, memes traveled by word of mouth. But as technology evolved, so did the way memes moved. “If you want to understand life,” Dawkins wrote, “don’t

think about vibrant, throbbing gels and oozes, think about information technology."[5]

In the early 1980s, memetics emerged as a multidisciplinary field drawing from the likes of microbiology, computer science, sociology, cognitive psychology, linguistics, rhetoric, and media theory. Researchers studied early examples of contagious media forms, such as chain letters, advertising jingles, and watercooler-conversation-inspiring TV shows. Researchers saw how, unlike sexual reproduction, which occurs on a one-to-one basis, memetic reproduction—amplified by the power of media—takes place on a mass scale.

The power of both biological and media viruses reveal less about themselves than they do about their hosts. A virus doesn't make us sick unless we lack an immune system capable of recognizing the shell and then neutralizing the code.

It wasn't until the dawn of networked media, however, that the power of memes to shape society was fully recognized and instrumentalized. Specifically, as top-down, mainstream media ceded control to lateral media from email and YouTube to Facebook and Twitter, new routes for transmission and contagion were opened.

The Internet became the breeding ground for what we dubbed "viral media" in 1993 with the release of the first Web browser. While marketers embraced the idea of creating "viral" messaging for their ads, the original theory of media viruses held that a virus spreads not solely because of some unique trait within itself, but because of the way it interacts with its environment.[6]

A media virus has two parts: a novel media shell (such as the exploitation of a new medium or the breaking of a media standard), and provocative memetic material within it. The virus initially spreads because the media sensationalism makes news. Rodney King's beating is captured on a camcorder, ISIS livestreams a beheading, or a Presidential candidate Tweets conspiracy theories. But the virus only infects us because it exploits a latent yet intrinsic gap in our cultural code. It mines for our cultural vulnerabilities in order to interpolate itself into the greater memetic matrix.

In that sense, the power of both biological and media viruses reveal less about themselves than they do about their hosts. A virus doesn't make us sick unless we lack an immune system capable of recognizing the shell and then neutralizing the code. Until we do that, the virus replicates, and our immune system goes berserk, giving us the fever, chills, congestion, or vomiting—which manifest in culture as media confusion, Twitter wars, protests in the street, sleepless nights, and "homegrown" terror. None of this is spontaneous or unpredictable. It's all just viral memetics in action.

The efficacy of memes involves not just their content, but rather their adjacencies to other symbiotic or hostile memes, the ecosystems and niches in which they exist, and the specific media by which they are transmitted. The medium is indeed the message, or at least a major part of it.

Related Research

Over the past decade, linguistics, behavioral economics, and "fake news" have come to dominate most propaganda efforts as well as most attempts to understand how they're engineered. Those disciplines and methodologies are relevant to memetics research but are often too siloed in their methodologies, or too tightly focused on technologies of transmission instead of human factors of reception.

Linguistics and Discourse Analysis

A recent University of Arizona project employed traditional critical discourse analysis to analyze what the content of Internet memes associated with racism reveal about racial discourse and comedy. Meanwhile, University of Memphis psychologists successfully used existing "theories of working memory, emotion, memory, and psycholinguistics" to increase their ability to predict the virality of particular memes. This approach extends far beyond academia to popular culture.

For example, Me.me is a "meme search engine" that attempts to index and catalog memes as they emerge. In 2017, Me.me conducted a keyword analysis on the memes in their database to identify those words that increased in usage over a twelve-month period. Of course the "winners" were political terms, ranging from MAGA to libertarian to president.

The limits to this approach is that memes can't be reduced to mere words. Memes themselves transcend efforts to frame issues with language or build systems of symbols. This is why teenagers in Russia can launch effective memetic assaults on Americans—with little or no understanding of the linguistic or cultural environment they mean to infect. Memes are better understood as independent actors in a competitive battle of ideas.

News Analysis

Since the 2016 election and subsequent Trump presidency, the "fake news" meme has become viral, itself. This, in turn, has prompted numerous initiatives to model the spread of disinformation disguised as reputable news articles. For example, some researchers are applying learnings from epidemiology to make predictions about how fake news spreads. Mathematicians at Fundação Getúlio Vargas in Brazil and elsewhere suggest that the basic research could lead to automated methods for detecting fake news in news streams. Meanwhile, media outlets like Buzzfeed, have studied data about Facebook user engagement with fake election news stories, to gain insight into how disinformation is affecting the media business. According to the Buzzfeed analysis, "top fake election news stories generated more total engagement on Facebook than top election stories from 19 major news outlets combined."

Facebook and Twitter, themselves, have demonstrated their ability to use algorithms to identify both fake news and posting by automated bots.[7]

Social Network Analysis

Numerous corporations and government agencies are investing significant resources to understand the properties of social networks, as well as how they enhance the spread of particular memetic constructs. For instance, the US Department of Defense's Social Media in Strategic Communication (SMISC) program funds research on algorithms to make sense of the "formation, development and spread of ideas and concepts (memes)" through social media. Certainly, it makes sense to consider how the structures of social media networks affect the diffusion of fake news and memes. But those kinds of maps don't answer the "why" one story, or meme, goes viral while another fizzles. They focus more on the technologies and networks than the humans within them, and are biased toward those understandings of propaganda that treat populations of people as "masses," incapable of exercising full autonomy.

Efforts to inoculate must focus on the effectiveness of memetic suppression by understanding the memes, but also the society, culture, economics, technologies, and other factors that allows particular memes, and memes in general, to thrive.

According to a paper from Indiana University's Center for Complex Networks and Systems Research, the impact that network structure has on a meme's diffusion depends on the type of meme: "While most memes indeed spread like complex contagions, a few viral memes spread across many communities, like diseases." We hope to shed light on the memetic qualities behind those differences in diffusion, and how they trigger specific responses in the people they infect.

Neuro- and Cognitive Science

Why do people believe things that aren't true, even when faced with facts? New efforts in social psychology aim to understand the cognitive biases that lead individuals to hit the "share" button on fake news. Some studies look at confirmation bias, our insatiable appetite for information that confirms what we think is true, while other social psychologists (and journalists) argue that the problem lies in the public's inattention to the credibility of the news source and mistrust of the media.

Experiments from Penn State's Media Effects Research Laboratory showed that readers are "swayed by the source or website that republished or posted the story—in other words, the vehicle that directly delivered them the story," as opposed to the originator of the story, writes lead researcher S. Shyam Sundar. "It's not surprising, then, to hear people say they got their news from 'sources' that don't create and edit news articles… (such as) Facebook and, by proxy, their friends."

Unfortunately, none of these stand-alone approaches are likely to crack the code of memes and help us build collective immune systems more resistant to infection. That's because these existing frameworks tend to account exclusively on the effects of the memes themselves—the "DNA" of the virus and the effectiveness of memetic transmission. We propose that efforts to inoculate must focus on the effectiveness of memetic suppression by understanding the memes, and also the society, culture, economics, technologies, and other factors that allows particular memes, and memes in general, to thrive. What is the "protein soup" in which the virus operates?

Computational Propaganda

Our media ecology is a cacophony of dissenting visions, targeted misinformation, weaponized persuasion, and other attempts to change our minds and our behavior (for helpful, trivial, and sinister reasons). If we ever fooled ourselves into thinking information can be neutral, we can no longer live under that illusion. The crisis of trust in institutions, and the active subversion of the perception and reliability of information sources is one of the defining features of political and social life in the twenty-teens. How we as a society address this situation and the solutions offered will determine the course of the future of democracy.

In this section, we look at the advent of computational propaganda, how it is propagated, and how to protect the body politic from its worst effects.

Sam Woolley has been at the forefront of the emerging field of computational propaganda, and in a recent paper with Phil Howard, defines it as "the assemblage of social media platforms, autonomous agents, and big data tasked with the manipulation of public opinion." We know the usual suspects of social media platforms in the US—mainly Facebook and Twitter, and to a lesser but still significant degree, Instagram, Snapchat, and SMS messaging. Autonomous agents, in the form of software bots, proliferate on Twitter and Facebook, amplifying certain messages through retweets, diluting others by hijacking hashtags and crowding out authentic discourse, and targeting and harassing "enemy" individuals on the platforms. Big data analysis assists in such efforts, using the information trails we leave behind us on the net to sort and reach us more effectively.

For example, in March 2018, it was confirmed that Cambridge Analytica, a data firm that consulted for the Trump campaign, harvested private information from 50 million Facebook users without their permission "to target political advertising and manipulate voters," according to Senate Judiciary Committee member Amy Klobuchar (D-Minn).[8] The revelation led to the US Congress calling Facebook chief Mark Zuckerberg to testify on Capitol Hill about the company's data collection practices.

Leveraging the power of computing, databases, and individualized news feeds, big data and computational media gives propagandists the ability to customize messages to groups or even individuals. Through these tactics, and by utilizing these ubiquitous platforms, marketers, propagandists, and even agents of state-sponsored manipulation have shifted debates, altered public discourse, and influenced the outcomes of elections from the Ukraine to the UK to the US.

To understand and, hopefully, neutralize such efforts, we must explore them not just as technological systems, in *vitro*, but as phenomena that are iterating within the greater cultural matrix—in *vivo*. To this end, let's employ three biological metaphors to examine the biology

of computational propaganda: selection strategies, contagion/inoculation, and symbiosis/endogenous retroviruses.

The conception that a line of code, or any software algorithm, is written and then executed via social networks or the web, is a misleading frame for true understanding of the process.

The use of virus as a metaphor for communication goes back decades, at least as far back as William Burroughs' famous line that "language is a virus from outer space." Following Rushkoff's popularization of the concept of viral media in the 90's, the term has since been an essential part of the common parlance of internet communication, memetics, and information media. But most often, attention has been on the message or code of the "virus," rather than the ecology in which it acts and propagates.

Selection Strategies

Computational propaganda spread through computer networks, at one level, are written (and/or programmed) by people, usually with a specific outcome in mind—even if the outcome is merely to sow doubt, confusion, or chaos in a discourse community. The creators of political bots, especially, try to skew conversations to favor a candidate or ideology, or more often than not, to sully and smear an opposing candidate with innuendo, misinformation, and outright lies.

But the conception that a line of code, or any software algorithm, is written and then executed via social networks or the web, is a misleading frame for true understanding of the process. For example, the way algorithms and bots act in the wild is more akin to "birthing" and nurturing, rather than "executing" a function, or "installing" a program. As computational propaganda researcher Doug Guilbeault notes, "social bots depend greatly on their environment, and develop their personalities by learning from other humans in the system."[9] These bots, acting as "living code" are designed to learn, mimic, and develop in a soup of communicative human and machine feedback loops. As such, even though the process might be called directed evolution, they cannot be controlled like other forms of software code running on a machine. Their behaviors can be forecast by the affordances they inherit by design, yet they are as such still largely unpredictable. Microsoft's Tay.ai experiment was one such notorious tragi-comic example of how bots can be hard to domesticate once unleashed into the wild.[10]

Bots and algorithms are best thought of as bred rather than programmed and, by extension, the memetic code they spread may also best be understood in terms of these more biological reproduction strategies. For viral media, fake news, and social bots alike, a fast-cycle, high-volume, bottom-up approach appears to be most effective. Rapid feedback loops give these evolving forms the opportunity to iterate and improve themselves for maximum effect, all through trial and error, from the bottom up, and at tremendous scale.

Attempts, especially by government agencies and large companies, to engage in top-down, PR-driven "viral" campaigns have backfired in profound ways. For example, New York City's well-meaning #mynypd twitter campaign encouraged people to send pictures of themselves with NYPD officers, but unexpectedly resulted in hundreds of images of police brutality and aggression.[11] While it might have looked good in a boardroom to have a stronger "social media presence," the ecology into which the #mynypd "virus" was born was rampant with anti-police sentiment, cynicism, and anger. The untested virus triggered an immune response.

Evolutionary biologists use the concept of "*r* vs. *K* selection" when analyzing reproduction strategies. Reproduction that is high-volume, frequent, and in which parents provide little nurturing is called an *r* strategy. This is the kind of reproduction of insects, fish, and many other "lower" life forms. *K* strategies are defined by less frequent offspring, a high degree of parental nurturing and care, and longer childhood stages for learning the skills necessary for survival. In a social media ecology that is extremely chaotic, fast-paced, and unpredictable, an *r* strategy, in which thousands or millions of memes or autonomous algorithms are brought forth into the world, and where each individual bot is less important than the overall survival of the "species" of messages, seems to be the more effective approach. Memes following a highly curated and resource-intensive *K* approach might work in certain situations, but the risk associated when they fail is extremely high.

Meme replication seems to be a game of volume and brute force, not specificity of intent and design.

Contagion and Inoculation

Further building on biological metaphors, contagion, and inoculation are often used to describe how viral media spreads and what can be done to resist it. But contagion and inoculation may be best understood less as metaphors than apt descriptors: researchers are delivering compelling evidence that bot networks distribute memes in the same ways that viruses spread, so an epidemiological approach is both enlightening and effective for fighting disinformation contagions.[12]

True, we must be careful in extending this model too far. For example, in behaviors such as exercise, contagious modeling has curious features. In a study of how people are influenced to exercise by peers on global social networks, researchers found that less active runners influenced more active runners, but not the reverse. Furthermore, in this example, men and women influence men, but only women[13] influence women. So memetic contagion is not exactly the same as the biological sort.

Nevertheless, epidemiology and public health frameworks suggest to those who wish to resist the spread of disinformation and manipulative media that they tackle both the code itself, as well as the communicable vulnerabilities that exist in the overall cultural ecology. Viruses work by simulating closely recognizable code and conditions in a host, and then manipulating a micro-environment to create conditions favorable to further spreading of the code. Working from micro-environment to micro-environment, a virus can scale quickly and dominate a system before the host's defenses are able to respond.

Micro-propaganda works in this way as well. And without the need to think about the overall stability and health of the system, a disinformation virus or cancerous meme can single-mindedly focus on reproduction and domination. Rumors run circles around truth, because rumors don't care about the evidence and argumentation needed to confirm something as true. A rumor can spread if it feels right, regardless of its connection to factual truth. (And the more we come to accept the intuitive logic of such memes, the less allegiance we tend to have to objective truth, ourselves.)

Meme replication seems to be a game of volume and brute force, not specificity of intent and design.

And the invocation of "feeling right" also describes why so many viruses and diseases are spread through sex and pleasurable experiences. Media addiction, like sexual or drug addiction, overcomes many of our capacities to resist engaging in destructive behaviors. Psychologist Dannagal Young, acknowledges the limitations we have in consciously resisting the spreading of fake news and disinformation. Satire, humor, and the desire for social acceptance makes "blaming readers for spreading

fake news from a cognitive perspective … somewhat equivalent to blaming a baby for soiling itself. They can't help it."[14] Oliver Vodeb, media scholar and founder of the Memefest Festival of Socially Responsive Communication and Art, recognizing both the hedonistic aspects of viral media as well as the domination of the form by right wing politicians and causes, recommends that the left "think more about pleasure and the pharmacological aspects of media, design, and communication."[15]

So, from a public health perspective, early detection and immunization against infection vectors is key to fighting the spread of disinformation. Containment of micro-propaganda in its early stages is another tactic necessary to success, as well as continued vigilance against secondary outbreaks. Immunization can occur by treating all those that surround an infected person or system, or by trying a broad-based approach that confer some form of herd immunization against the malicious code. Immunization in this case, takes place by priming people to be aware of coercive messages, teaching them to examine sources of messages before sharing them, training them to recognize attempts to hijack cognitive biases such as motivated reasoning and confirmation bias, and slowing down their media sharing practices.

Meme Epidemiology

Viruses aren't all bad, and we have to be careful how we curate our cognitive ecosystems. We have many endogenous retroviruses in our DNA, many of which have helped humans evolve successfully over millions of years, and made us who we are.[16] A biological-ecological approach to disinformation puts the responsibility on us to actively design systems that can grow and evolve, but without destroying their hosts in the process. This is a delicate task, but one that is greatly aided when we look to nature for guidance and grounding—reminding ourselves all the way that any rigorous analysis must accept the relationship between nature and memes—like atoms and bits—is more a metaphorical than actual correspondence.

For a century or more, researchers have been developing mathematical models to analyze and forecast the spread of infectious diseases. As the simulations have increased in resolution, so have their efficacy in forecasting how epidemics may play out. Can these models developed for epidemiology be applied to study and predict the propagation of memes through networks over time and space, both virtual and real?

Since the meme "meme" began in the late 1970s, numerous researchers have explored an epidemiological approach to understanding meme propagation. The common framework is that not only are memes transmitted between people, like human diseases, but they can evolve to become even more viral. Think of it as the natural selection of ideas or thoughts whether the information contained in the meme is objectively true or not. Indeed, SRI International, under a grant from the United State Air Force Research Laboratory and the Intelligence Advanced Research Projects Activity (IARPA), conducts research on "Meme Epidemiology," described as "applying principles and quantitative frameworks from the field of epidemiology" to track meme transmission. According to their project description, they are "exploiting the tools of machine learning, computational linguistics, and information retrieval to predict the emergence and spread of memes as they spur cultural change."[17]

The simplest mathematical model used in both disease epidemiology and some meme epidemiology research is the SIR model where S is the number of susceptible people, I is the number of people infected, and R is the number of people who have recovered from the infection. Historically, it's been applied to the likes of measles, rubella, and mumps where once someone is infected and recovers, they become immune.

Several years ago, mathematicians from the University of New Brunswick in Fredericton, used a modified SIR model to study how memes go viral. In their work, susceptible people were those who hadn't seen the meme, infected people were those actively interested in the meme's content and spread the ideas further, and recovered were those who had seen the meme and lost interest. In this agent-based model, where people are the agents, memes move between groups based on mathematical probabilities. The researchers then used historical data from Google Trends to test their model with humorous memes like the Ultimate Showdown of

Ultimate Destiny animation and “O RLY,” an image of a snowy owl labeled with sarcastic captions. According to their paper, “the successful implementation in the modeling of meme spread as reflected in Internet search data shows that memes may be treated as infectious entities when modeling their propagation over time and across societies.”

They concluded that “the success of the modeling indicates that memes can indeed be considered to be infectious in nature, which opens up whole new frontiers for the likes of education, marketing, and even politics.”[19] While the models may prove out in some cases, it’s extremely difficult to account for all of the variables that may dramatically influence the numbers of susceptible people, how susceptible those individuals may be, or that people’s interests in certain areas of information can change.

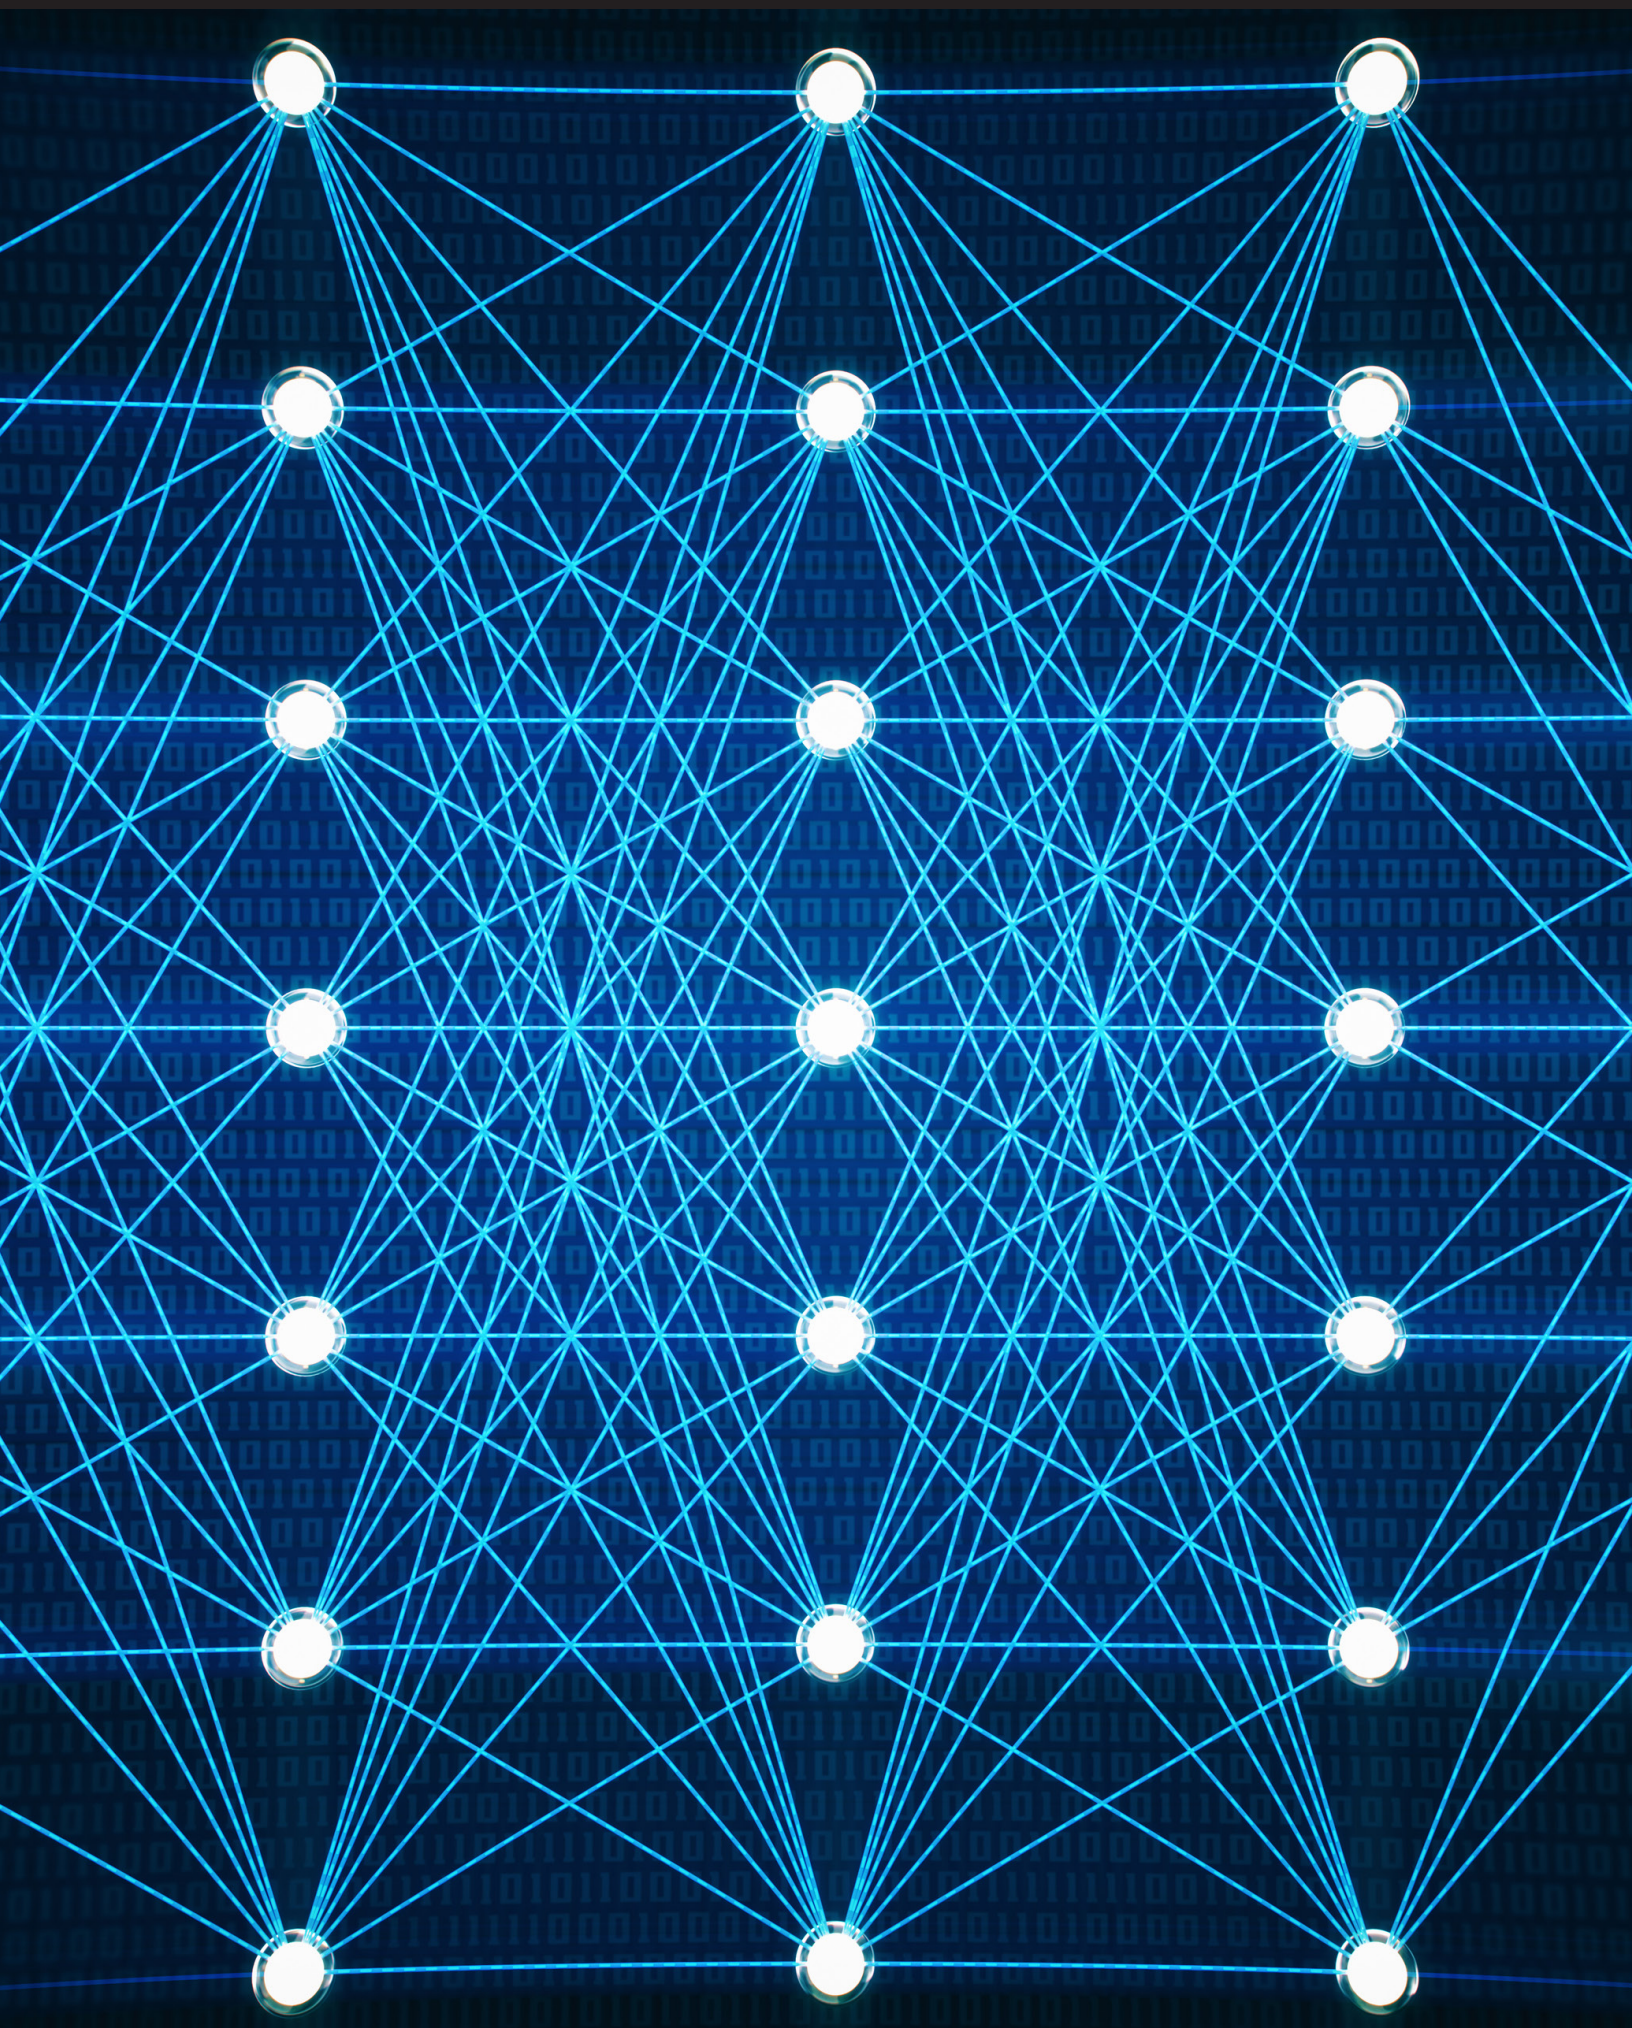

Meme Content vs. Meme Content

Most researchers have favored the study of the content of the memes themselves, for the best clues on how they function and why they spread. While the content of memes certainly matters, there are many other determinative factors in a meme's contagion that are exogenous to the meme itself. The emotional effect on individual recipients, the timing to coincide with major events, the architecture and dynamics of the network, and the media literacy and facility of the target population may be more responsible for memetic contagion than the particular memes themselves.

In 2009, marketing researchers at the University of Pennsylvania's Wharton School analyzed the online sharing of several months of *New York Times* articles. They focused on how the emotions evoked by a piece of content affects its social sharing and report "high-arousal positive (awe) or negative (anger or anxiety) emotions" is more likely to be shared.[20]

In related work, University of Memphis psychologists applied "theories of working memory, emotion, memory, and psycholinguistics" to gain predictive power of "image macro" meme success. Among many other predictor variables, their data concludes that shorter memes (fewer than four words) do well, concrete terms are more memorable than abstract language, and swear words hamper virality.[21]

Computer scientists at the Hebrew University employed an algorithm to predict the spread of certain hashtags on Twitter. They analyzed a dataset of more than 400 million Tweets. Their experimental goal was to see if it's possible to predict the acceptance of a hashtag, as a form of meme, by looking at the characteristics of the tweets containing the hashtag, such as the number of words, the words used, emotional effect, and other content characteristics. Their conclusion is that "there are three main factors to the acceptance of a meme: the meme's content, the meme's context, and the social graph."[22]

According to research from Indiana University's Center for Complex Networks and Systems Research, the content of a meme, its "innate appeal," may contribute far less to its virality than previously thought. After building models that simulate portions of social networks like Twitter, researchers were able to predict the success of a meme based upon how it initially spread across communities. Not only that, but the memes that went viral were no different than those that didn't. The success was due to the structure of the social network.[23]

Does that mean that the content of the memes doesn't matter all? Of course not. According to new research from the Indiana University team and colleagues at

the Shanghai Institute of Technology and Indiana University, we do prefer high-quality digital information. So, the researchers ask, what accounts for "the viral spread of low-quality information, such as the digital misinformation that threatens our democracy?" According to their model, it all comes down to fragmented attention compounded with information overload—even if we are good at discriminating between fact and fiction.[24]

Essentially, we are often too overwhelmed, and our attention too fragmented, to intelligently differentiate. We can only process a small amount of information flowing through our feeds and the competition of ideas is so brutal that we may not even catch a glimpse of the truths among the trash. And it's during that perpetually anxious state of overload when cognitive biases truly kick in.

> Cognitive and behavioral processes for dealing with opinions that challenge one's beliefs may decrease our capability to discriminate between high- and low-quality information. For example, confirmation bias may have evolved as an effective strategy to avoid misinformation, by comparing incoming information with one's own existing beliefs, and adopting it if it is sufficiently concordant. However, in social media, such a bias easily leads to ineffective discrimination; strategies such as accepting new information, if it comes from multiple sources, are not useful because people lack knowledge of the social network structure necessary to determine whether multiple information sources are independent of each other.
>
> Confirmation bias may be reinforced online by our limited capacity to cope with the information overload caused by the messages that flood our screens and our consequent need to quickly discard irrelevant information.[25]

Or, as Indiana University Network Science Institute researcher Giovanni Luca Ciampaglia puts it, the spreading of disinformation really happens in that "last mile between your phone's screen and your eyes."[26]

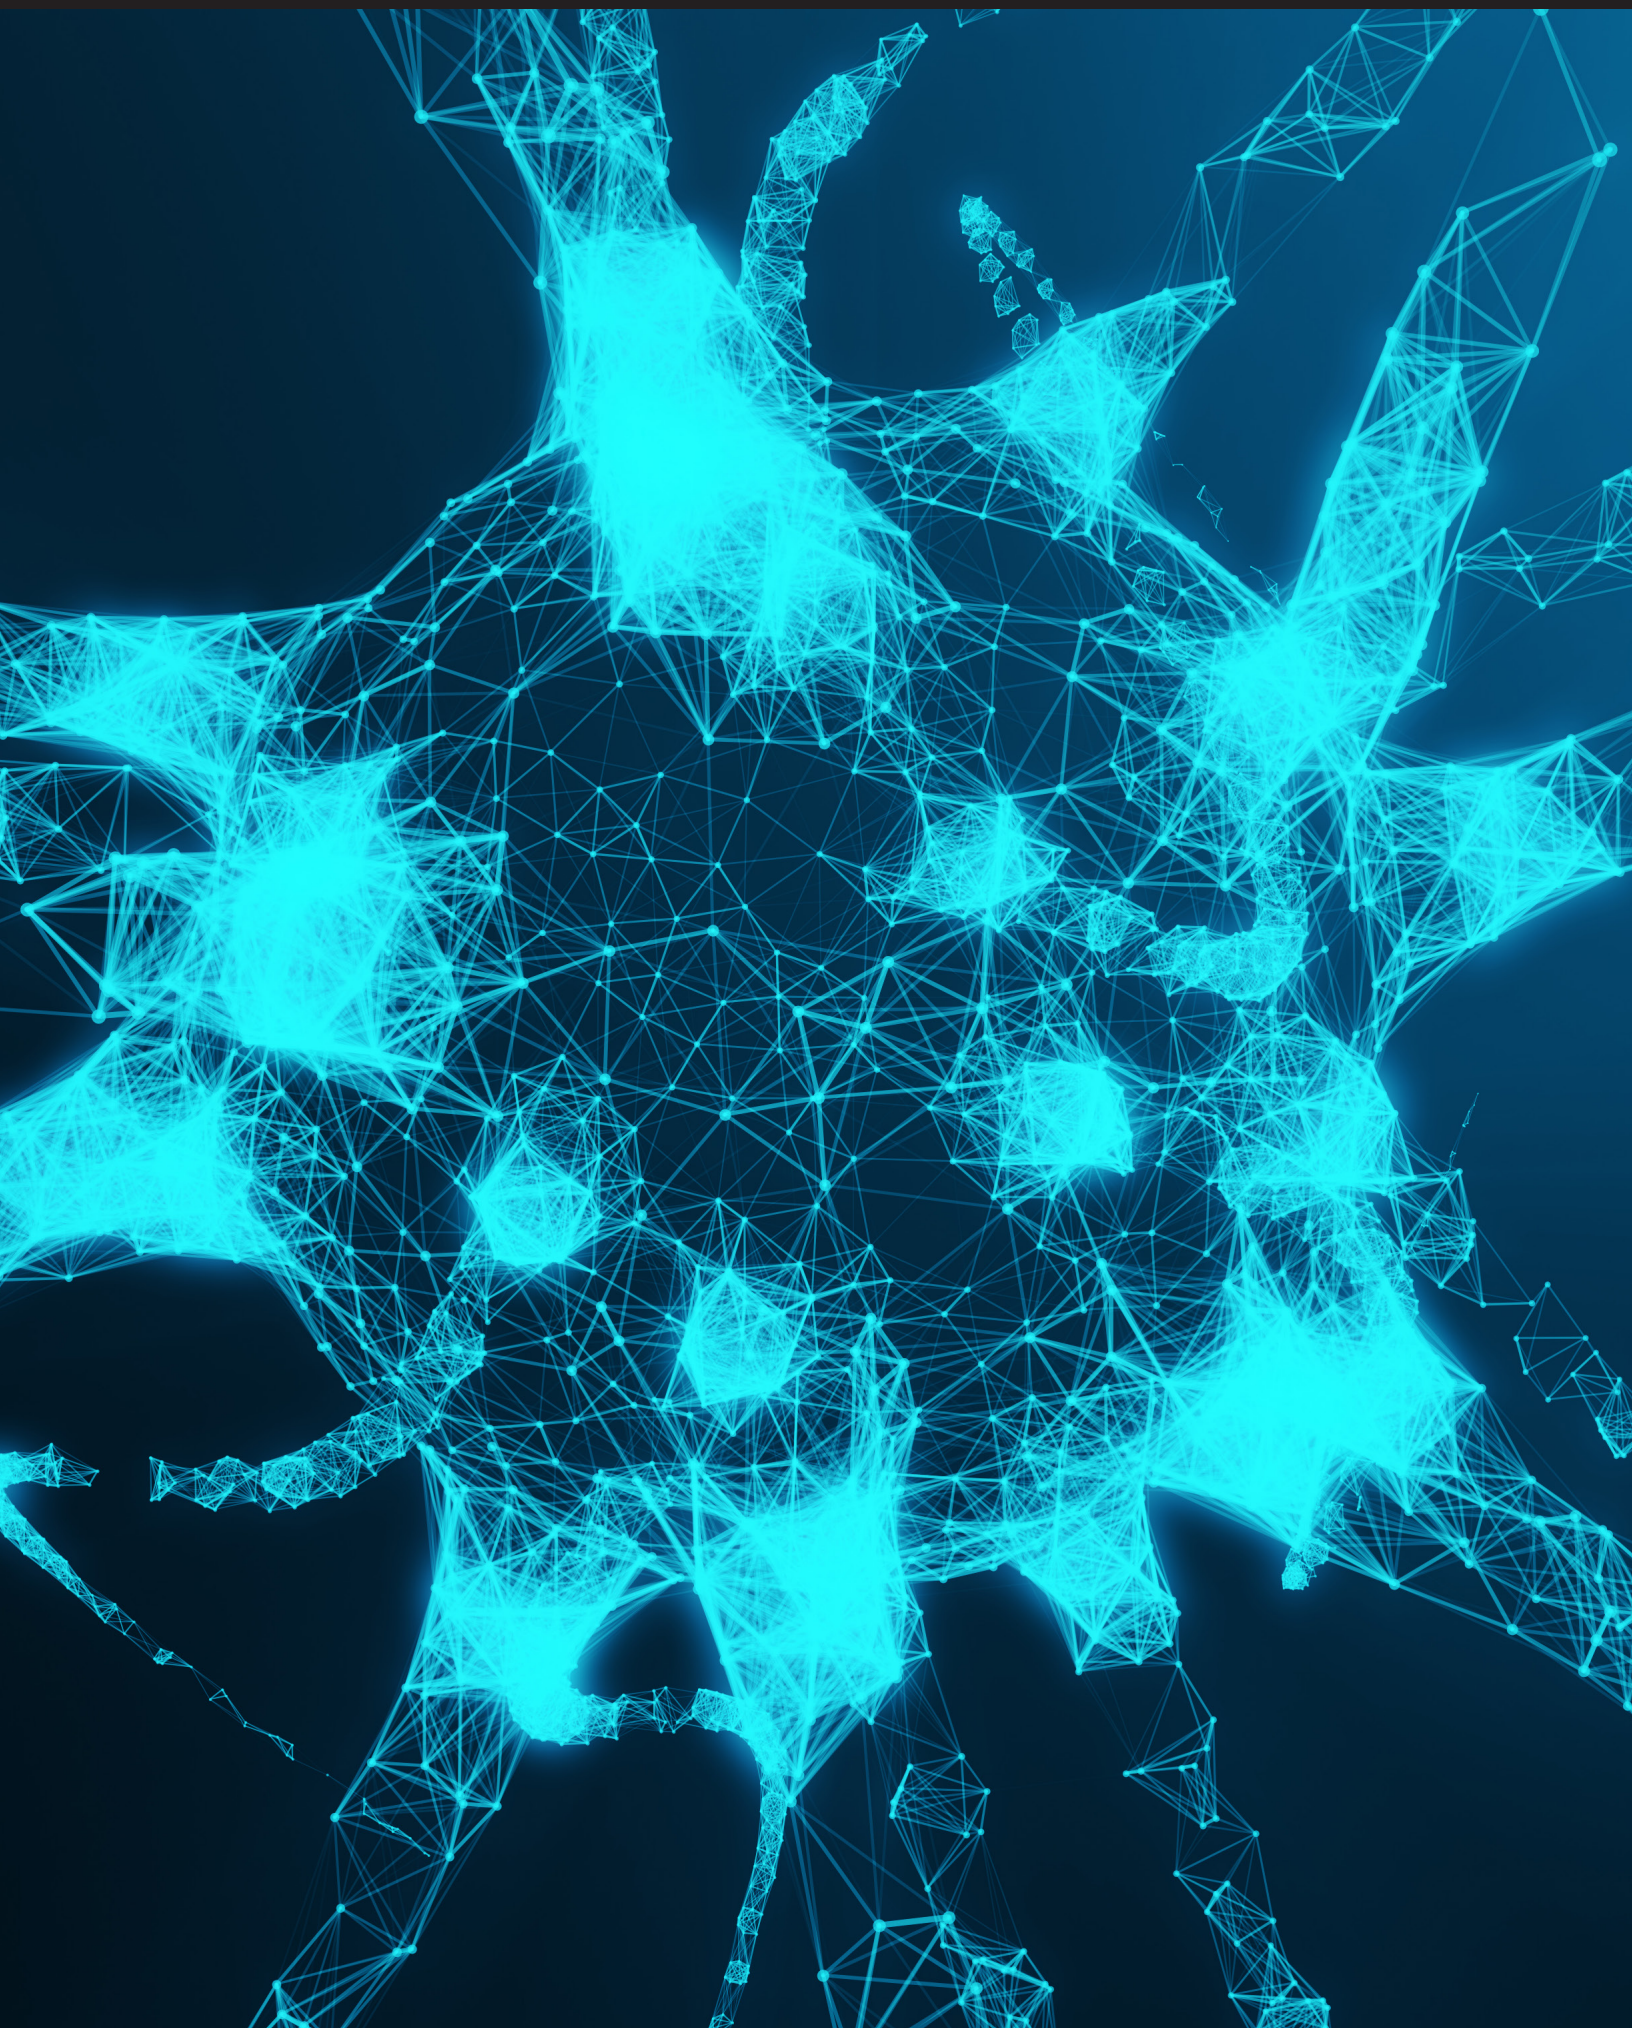

Our Brains on Disinformation

When trying to understand why disinformation works so well, most scientists are drawn to the "source" of our main conscious (and unconscious) information processing system—the brain. New insights from neuroscience and neural imaging seem to flower and bloom every day from research centers around the world. These insights are magnified by a ready press and willing public, ready to soak up some "objective" explanation for why so many of our fellow humans are so easily seduced by fakes or the patently absurd. While one must always tread carefully when claims from neuroscience are popularized, the study of how we process information has yielded powerful windows into the relationship between disinformation and our brains—between memes and human biology. In fact, we could even say that at the level of neurons, information becomes biological.

In this section, we will look at two key aspects of the neuro and cognitive science of disinformation: first, how information sticks in the brain, and second, why stories are shared. These discoveries are helpful for how we might consciously design our media environments and educate our body politic for better outcomes. Right now, however, they are being leveraged effectively by ideologically-driven individuals and entities trying to shift the flow of public opinion in favor of their sponsors or idols.

Sticky Information

In his famous work on cognitive systems, psychologist Daniel Kahneman divides mental functioning into two types. System 1 is the fast, emotional, largely subconscious processing that can react quickly to dangerous situations, instantly recall memorized facts, and guide us through a city street seemingly on "autopilot." System 2 is the slower, logical, contemplative, critical system that helps us to guide our attention to specific things in the environment, analyze a situation or piece of content, and simulate and rehearse interactions.[27]

Disinformation campaigns can lodge ideas and feelings into individuals by using techniques to bypass the slower cognitive system (System 2). This can be done by harnessing cognitive biases such as the availability heuristic—where we favor information that is easily accessed—and flooding data streams with soundbites and examples that enhance a particular agenda. People who can quickly and easily think of emotionally-laden examples, metaphors, and phrases will default to those handy phrases when called upon to defend a position. It is no coincidence that phrases such as "burn the witch," "lock her up," "build the wall," and "drain the swamp" were pushed at every instance by the Trump campaign as well as its human supporters and bot armies both

within the United States and in the server farms of foreign powers. Playing on the availability heuristic, these phrases and memes were sent out in planned tsunamis right before voting in the primaries and in swing states just before the general election. They were used again over the following year, in bot-driven campaigns from #releasethememo to #firemueller.

The Trump campaign and administration's sensibility of provocation, satire, and insulting humor is also backed by neuroscience. Humor and satire, in a flood of information, makes messages stand out, and makes them more memorable. Dannagal Young argues that "humor suspends argument scrutiny of the premise of a given text through various cognitive mechanisms involving processing ability and motivation."[28] Humor is a fantastic Trojan Horse to bypass our critical functions and lodge memes deep into our brains.

Why Stories are Shared

Emily Falk, Director of the University of Pennsylvania's Communication Neuroscience Lab, makes a point that may seem obvious, but is also powerful when connected the way disinformation and fake news is shared. She notes that "people are interested in reading or sharing content that connects to their own experiences, or to their sense of who they want to be. They share things that might improve their relationships, make them look smart, empathic, or cast them in a positive light."[29]

We have a highly socially attuned brain, that tends to analyze situations and behaviors on how they will affect our status, our sense of self, and the quality of our relationships. These social factors play a bigger part in motivating people to share stories than the need to inform one another of facts. The information flow is a byproduct of the social interaction.

Other researchers in Falk's Lab have demonstrated how the expectation of social confirmation and reward influence the likelihood of a person sharing a meme with others. When neural systems associated with "self-related thinking, regions associated with mentalizing— imagining what others might think—and with overall value" were highly activated by stories in the test subjects; those stories had a much higher rate of sharing amongst all users. Media virality, then, could be predicted (and designed) by first confirming that the information causes certain brain areas to be energized when reading or consuming.[30] And this predicted population effect can be seen by looking at a remarkably small number of brains.[31]

Beware Neurocentrism

This report has made an explicit attempt to look at the biology of disinformation by addressing the background ecology and fitness of messages, and not just in the memes or messages themselves. Similarly, one must always be wary of seductive neurocentric explanations for how viral media works. Our brains are situated in our messy bodies, with all the gut bacteria and hormones that influence our thinking. They are also situated in a physical, social, and political environment that plays a constitutive role in how we process and spread information. In fact, gender and education levels have been linked to differing capacities for processing and remembering news information. Recent research has found that "women recognized (accuracy) and recalled (salience) social images better than men, [but] men were more skilled at recognizing, but not recalling, nonsocial images."[32]

Education plays a significant factor as well. From the same study as above, the researchers found that "participants with lower educational levels recognized and recalled fewer images than individuals with higher educational levels. Interactions between demographic variables and time suggest that memory records for social images are more stable than those for nonsocial images."[32]

So, biology is best understood as systems acting within and in concert with other systems. This is true for brains and mental systems.[33]

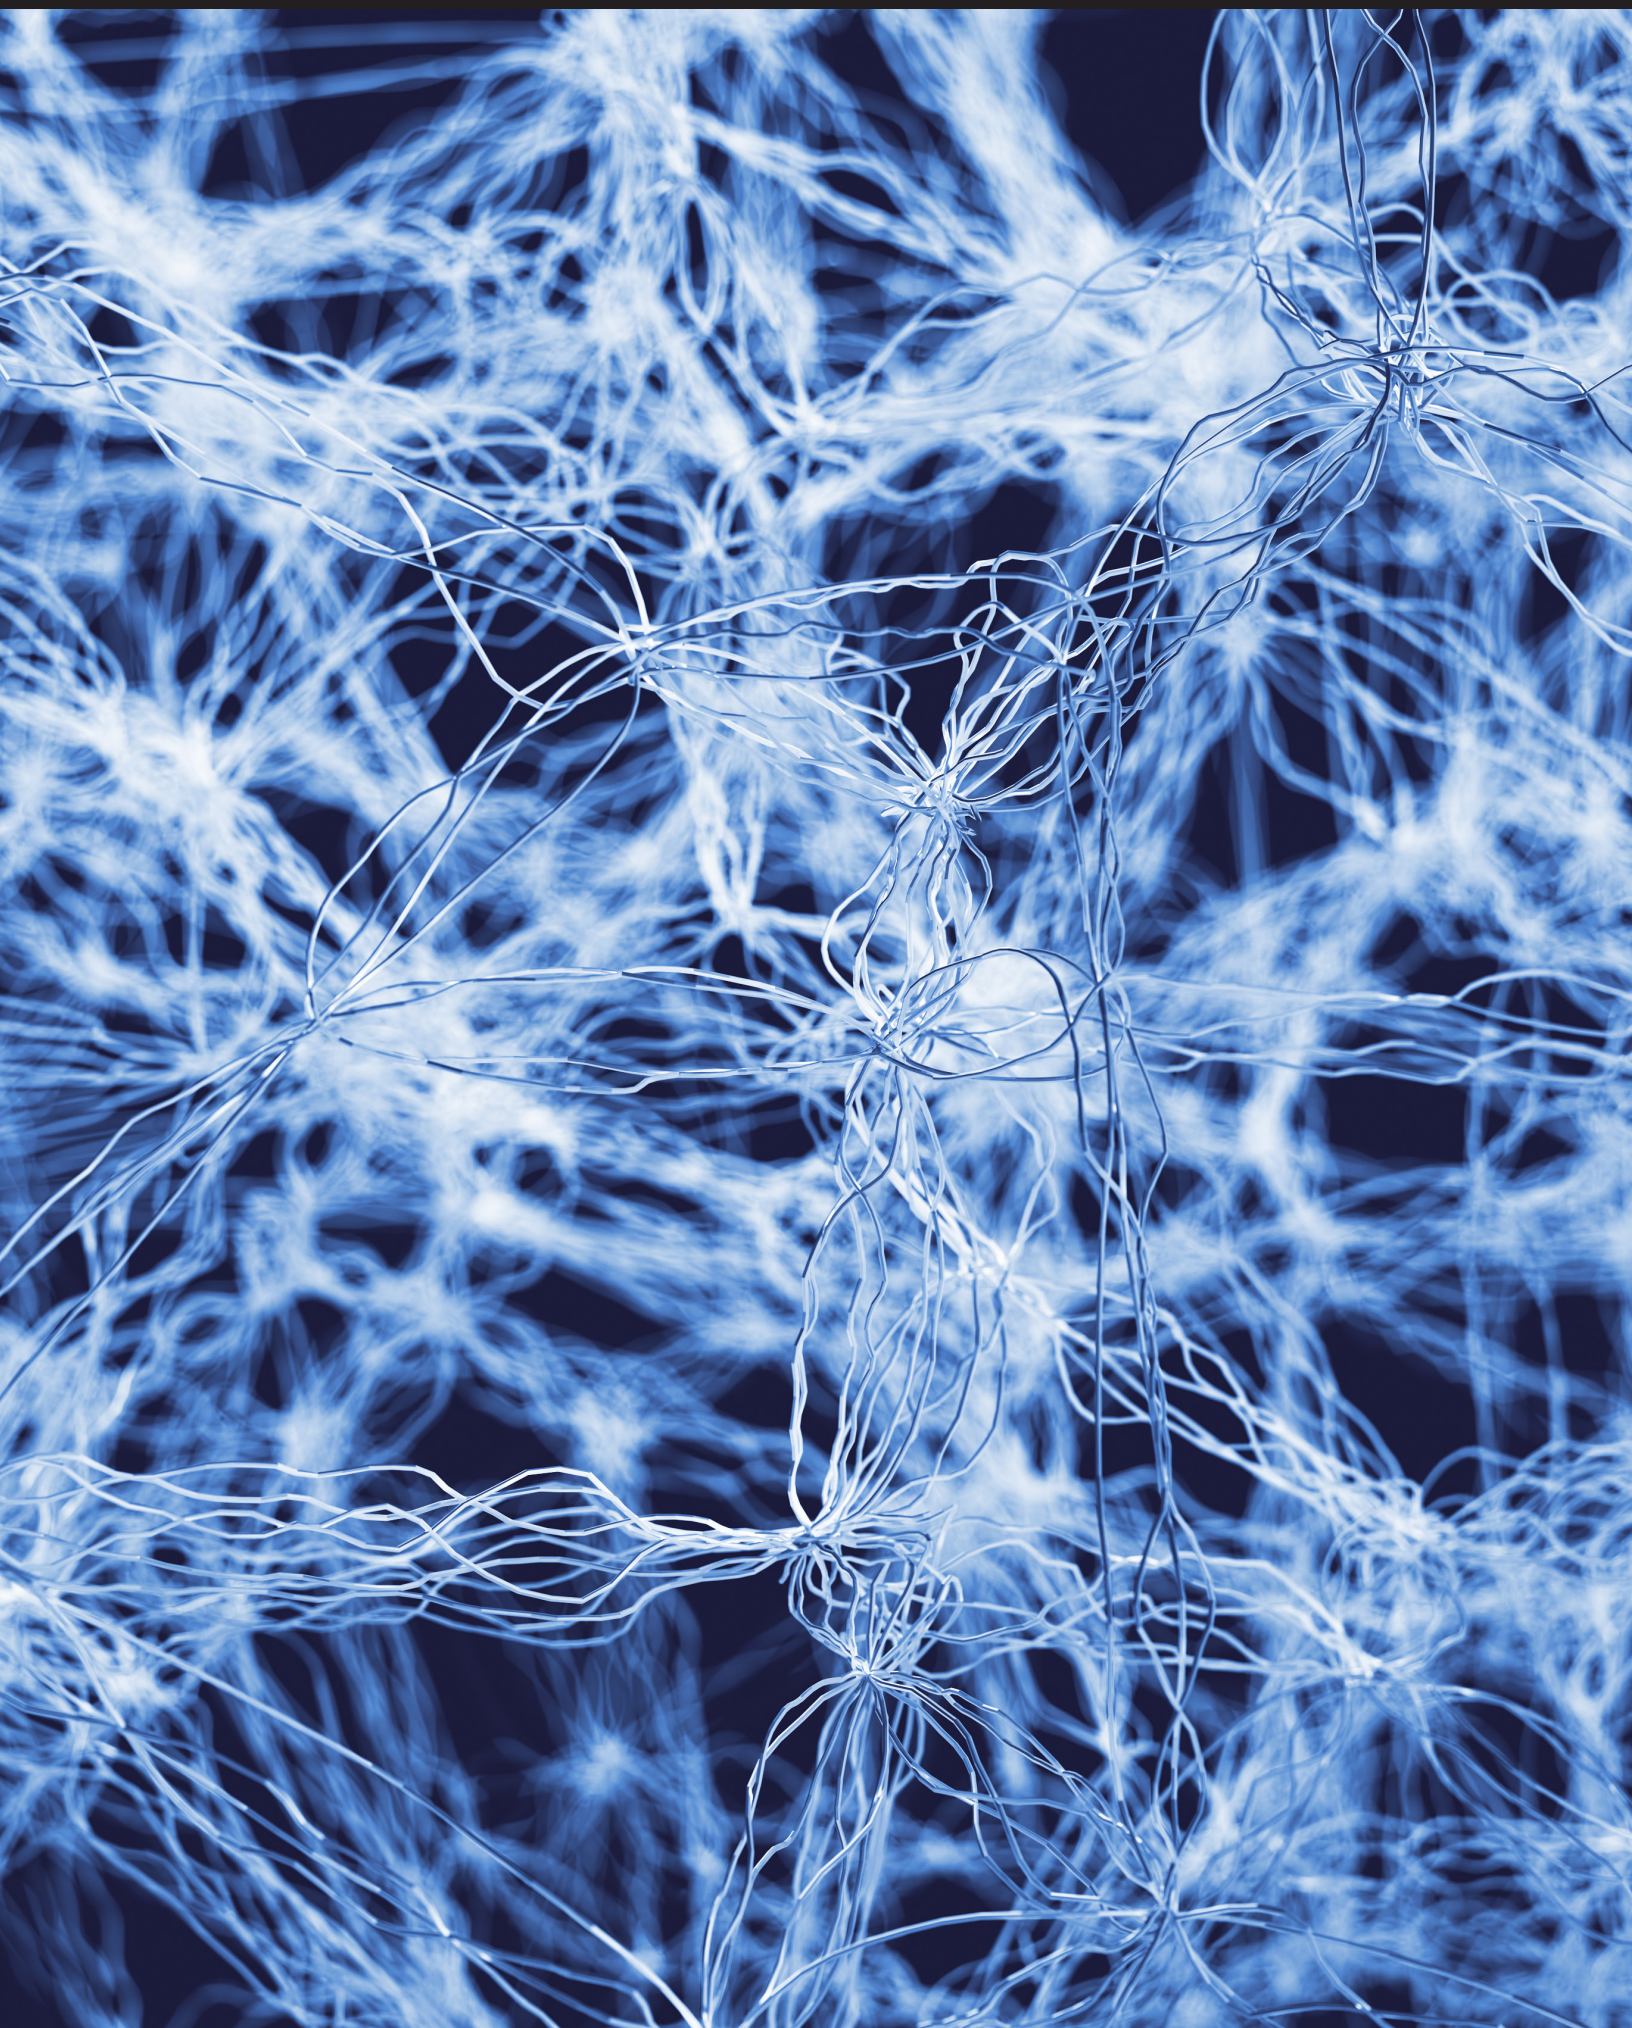

The Ends Don't Justify the Memes

While biology may be the best way of analyzing memetic activity in a human culture, memetic activity itself—particularly when it occurs in digital networks—tends to distance people from the biological and cultural mechanisms that stand a chance of mitigating its effects. The further removed from live, embodied communication we get, the less our interaction is governed by our evolved social skills. It becomes instead, as the many mis-interpreters of Darwin understand the world, a battle for survival of the fittest meme.

At first glance, the horizontal landscape of interactive and social media seemed to promise more lateral communication between peers and less propaganda from above. The elites who owned traditional media outlets would no longer be able to serve as gatekeepers for what the masses read and watched. Anyone with a camcorder, smartphone, web page or social media account would be able to get a message out. And if it was compelling enough, it would be replicated and spread to millions—without the willing cooperation of traditional media organizations.

A videotape of a black man getting beaten by white cops in Los Angeles makes it to cable news before morning. Smaller and tabloid media do not hesitate to broadcast it and, once they do, everyone else must, as well. The original "media virus"[34] is launched, and is so contagious that it leads eventually to full-scale rioting in a dozen American cities.

The term "media virus" was meant to convey the new way ideas could spread in a world with more interactive communications. But the deeper biological underpinnings of this perspective have been forgotten. Like a military spreading chemical weapons to its own troops, we are using a weapon that we do not understand, and at our own collective peril.

A real, biological virus has a novel, never-before-seen protein shell that lets it travel through a person's bloodstream unrecognized. (If the body had identified the virus, it could have sent antibodies to attack it.) The virus then latches onto a cell in the host organism and injects its genetic code inside. The code is, basically, genetic material that wants to get reproduced. So it works its way to the nucleus of the cell, and seeks to interpolate itself into the cell's DNA. It looks for weak spots, then nests in there. The next time the cell reproduces, it replicates the virus's code along with its own.

Then the person carrying the virus begins spreading it to others. If the next person's immune system doesn't recognize the protein shell, then they get infected, too. The virus continues to replicate and spread until, at last, the body learns to reject its code. From then on, the

protein shell will be recognized and attacked—even if it comes back months or years later. Immunity.

A media virus works the same way. It has a novel, unrecognizable shell—but that shell is made of media, not protein. The virus must be packaged sensationally, as part of a unique, rule-breaking use of media that we can't help but spread. A camcorder tape captures police brutality. A voicemail message reveals an actor's abusive relationship or an affair between royals. A TV star posts social media updates on his mental breakdown. An underwear commercial veers too close to child pornography. A rock album is rumored to contain hidden satanic messages. A political candidate's wireless microphone records him making sexist remarks about a female colleague. A woman "live streams" her husband dying of gunshot wounds. A congressman transmits smartphone pictures of his genitals to a minor. A Shakespeare play is reinterpreted as a presidential assassination. A president threatens a nuclear attack in a public, 140-character message typed with his thumbs.

In each case, the story's initial proliferation has more to do with the medium than the message. The viral shell is not just a media phenomenon, but way of grabbing attention and paralyzing a person's critical faculties. What the…? That moment of confusion creates the time and space for infection.

We can't engineer a society through memetics the way a biologist might hope to engineer an organism through genetics. To do so bypasses our higher faculties, our reasoning, and our collective authority.

Once it has been launched, the virus replicates only if its code can successfully challenge our own. That's why the ideas inside the virus—the memes—do matter. They must interpolate into our own confused cultural code, exploiting the issues we haven't adequately addressed as a society, such as racial tension, gender roles, economic inequality, nationalism, or sexual norms.

A fatal car crash on the side of the highway attracts our attention because of the spectacle, but worms its way into our psyche because of our own conflicted relationship with operating such dangerous machinery ourselves, or because of the way it disrupts our ongoing, active denial of our own mortality.

Likewise, a contagious media virus attracts mass attention for its spectacular upending of TV or the net, but then penetrates the cultural psyche by challenging collectively unresolved or repressed anxieties. Surveillance video of a police van running over a black suspect recalls America's shamefully unacknowledged history of slavery and ongoing racism. The social media feed of a neo-Nazi bot in Norway can stimulate simmering resentment of the European Union's dissolution of national identities. Sexual harassment via social media by a sitting president breaks the rules of media decorum, while provoking the animus of a population still resentful of women in the workplace.

The perplexing thing is that it doesn't matter what side of an issue people are on for them to be infected by the meme and provoked to replicate it. "Look what this person said!" is reason enough to spread it. In the contentious social media surrounding elections, the most racist and sexist memes are reposted less by their advocates than their outraged opponents. That's because memes do not compete for dominance by appealing to our intellect, our compassion, or anything to do with our humanity. They compete to trigger our most automatic impulses.

But we can't engineer a society through memetics the way a biologist might hope to engineer an organism through genetics. To do so, bypasses our higher faculties, our reasoning, and our collective authority. It is unethical, and, in the long run, ineffective.

Yes, well-meaning and pro-social counterculture groups from the Situationists to Adbusters and Greenpeace have attempted to spread their messages through the equivalents of viral media. They would cut and paste text and images to subvert the original meanings of advertisements, or the intentions of a corporate logos. It was a form of media aikido, leveraging the tremendous weight and power of an institution against

itself with a single clever twist. With the advent of a new, highly interactive media landscape, Internet viruses seemed like a great way to get people talking about the unresolved issues that needed to be discussed in the light of day. After all, this logic goes, if the meme provokes a response, then it's something that has to be brought up to the surface.

The problem is, the means don't always justify the ends. As we now see, the bottom-up techniques of guerrilla media activists are now in the hands of the world's wealthiest top-down corporations, politicians, propagandists, and everything in between. To them, viral media is no longer about breaking through propaganda and unearthing the truth about inequality or environmental threats. It's simply about generating a response by any means necessary, even if that response is automatic, unthinking, and brutish.

Not that the technique was ever appropriate, even practiced benevolently. The danger with viruses is that they are constructed to bypass the neocortex—the thinking, feeling part of our brain—and go straight to the more primal reptile beneath. The meme for scientifically proven climate change, for example, doesn't provoke the same intensity of cultural response as the meme for "elite conspiracy!"

Logic or truth have nothing to do with it. Memes work by provoking fight-or-flight reactions. And those sorts of responses are highly individualist. They're not pro-social; they're anti-social.

So, for example, a viral assault will not persuade a flood-ravaged town to adopt strategies of mutual aid. It could, on the other hand, help push survivors toward more paranoid styles of self-preservation. Memetic campaigns do not speak to the part of the brain that understands the benefits of tolerance, social connection, or appreciation of difference. They're speaking to the reptile, who only understands predator or prey, fight or flight, and kill or be killed.

Figure and Ground

Memetics was first popularized not by a cultural anthropologist, poet, or media theorist, but by a particularly materialist evolutionary biologist in the 1970's.[35] A devout atheist, Dawkins meant to show how human culture evolves by the same set of rules as any other biological system: competition, mutation, and more competition. Nothing special going on here.

Memes work by provoking fight-or-flight reactions. And those sorts of responses are highly individualist. They're not pro-social; they're anti-social.

It turns out there is something special going on here, and that there are a few things missing from this simplistic explanation of memes and genes, alike. A meme is a great corollary to a gene, for sure, but neither genes nor memes determine everything about an organism or a culture. Surprisingly, DNA is not a static blueprint but acts differently in different situations. It matters which genes we have, but it matters even more how those genes express themselves. That's entirely dependent on the environment, or the protein soup in which those genes are swimming. It's why a locust can be like a tame grasshopper or, in the right conditions, transform into an entirely more gregarious, swarming creature. That's not a sudden mutation within a single lifetime; it is a shift in gene expression that changes the whole organism.[36]

Genes are not solo actors with entirely pre-determined code. They are not selfishly seeking their own replication at all costs. Newer science shows they are almost social in nature, adapting and expressing themselves differently in different environments. Organisms get information from the environment and from one another for how to change. The conditions, the culture, and its connectivity matter as much as the initial code.

If we truly want to understand cultural contagion, we must place equal importance on the memes, the viral shell around those memes, and the ideological soup in which those memes attempt to spread. Early memeticists saw memes as competing against one

another, but that's not quite right. Memes are all attempting to self-replicate by exploiting inconsistencies or weaknesses in our cultural code. They are not attacking one another; they are attacking us humans.

Advertising agencies loved that earlier model, because it meant all they had to do was work on crafting the best meme for it to "go viral." But that's not how it actually works, and why most of those campaigns failed miserably. A meme can only go viral if it is unleashing a repressed cultural agenda. The potential has to be there, already. The Trump viral shell was his reality show persona and its unique migration to real world politics. But the memes within the Trump virus replicated—at least in part—because there was already a widespread but unexpressed white nationalist rage in America.

It's not the meme that matters, but the culture's ability to muster an effective immune response against it.

Human societies must come to recognize the importance of developing a healthy cultural immune response to an onslaught of hostile memes. The technologies through which they are being transmitted are changing so rapidly that it would be impossible to recognize their new forms—their shells—in advance. We must instead build our collective immune system by strengthening our organic coherence—our resistance to socially destructive memes.

This is particularly difficult when the enemies of democracy and their unwitting allies (the communications directors of political campaigns) are busy upscaling memetic warfare with each of social media's latest tricks, from predictive algorithms to artificial intelligence. In addition to artificially amplifying the "scale" of memes that may not have gained any organic traction on their own, these algorithms and bots are designed to engage with us individually, disconnect us from one another, neutralize our defense mechanisms, and program our behaviors as if we were computers. Television advertisers may have normalized the idea that consumers can be experimented on like lab rats, but social media takes it to an entirely new level. At least TV happened in public. TV ads were expensive, proving that there was a big company behind the product willing to invest in its success. And TV stations censored ads they found offensive. Social media manipulates us individually, one private screen at a time. Messages may cost pennies or nothing at all, and they're sold and placed by bots with no regard to their content. When media is programmed to atomize us and the messaging is engineered to provoke our most competitive, reptilian sensibilities, it's much harder to muster a collective defense.

The powers working to disrupt democratic process through memetic warfare understand this well. Contrary to popular accounts, they invest in propaganda from all sides of the political spectrum.[37] The particular memes they propagate through social media are less important than the immune reactions they hope to provoke. Memetic warfare, regardless of the content, discourages cooperation, consensus, or empathy. The reptile brain it triggers doesn't engage in those prosocial behaviors. Instead, in an environment of hostile memes and isolated by social media, human beings become more entrenched in their positions and driven by a fear for their personal survival. Worst of all, since these platforms appear so interactive and democratic, we experience this degradation of our social processes as a form of personal empowerment. To be truly social starts to feel like a restraint—like the yoke of political correctness, or a compromising tolerance of those whose very existence weakens our stock.

This may not have been the intent of social media, or any of the communications technologies that came before it. The internet doesn't have to be used against a person's critical faculties any more than language has to be used to lie or text be used to inventory slaves. But each extension of our social reality into a new medium requires that we make a conscious effort to bring our humanity along with us.

A Way Forward

To combat hostile memes, we can either attack the memes themselves with antagonistic memes, or strengthen our cultural immune response to their codons.

The former approach is more straightforward, though it risks further weaponizing the media culture. The technique—currently practiced by a Hungarian memetic agency called Darwin38—involves analyzing the landscape of memes around a particular idea in order to understand the various memetic neighborhoods, how to position or re-position meme, and which other memes may complement or degrade its virality.

While such an approach may be appropriate in a crisis, the problem is that it increases the amount of weaponized memetics in play at any particular time. The enemy memes may be weakened, but so, too, is the community of humans under attack. This simply makes them less capable of fighting off viral infection the next time. As when treating a person with antibiotics, the only one that gets stronger in the long term is the bacteria.

Protecting the public from destructive memes requires either external regulation of the platforms, or self-policing of platform content. Regulation is problematic for a number of reasons. Since the current social media platforms enjoy near-monopolies in their particular messaging types, they would be likely to encourage regulations that only cement their power in the landscape. Any serious regulations of their content or distribution would also raise concerns of reducing American competitiveness against China and other nations.

Until now, social media companies have been depending on users to flag suspicious or malicious content—a self-policing that has clearly failed. Now, they are offering to upgrade their policing capability by using AI's and machine learning. But given that the CEOs of social media companies got into the current mess by depending on technologies they didn't fully understand, we must be guarded about their ability to find solutions using technologies whose ramifications are even less predictable.

A less dramatic, but ultimately more powerful approach is to strengthen the cultural immune response of the society under attack. This could mean educating people about the facts around a particular issue or bringing very controversial but memetically potent issues into the light of day. A society having an open, honest conversation about race, guilt, and fear of change is less vulnerable to a memetic attack invoking white supremacy than a society still afraid to have that painful conversation.

Bringing repressed issues up and out into the light of day reduces the potential difference—the voltage—between the expressed and unexpressed cultural

agendas of that moment. The hostile memes will either not be able to locate confused code in which to nest, or will fail to produce a rapid acceleration of reproduction if they do.

The downside to such strategies, of course, is whose curriculum is used to educate the public about a particular issue? Town halls and other public forums are great for airing grievances, but at some point the conversation will have to turn to real history, real facts, or real science. Whose real is accepted? We end up back in the highly criticized situation envisioned by the father of public relations, Walter Lippman, his "council of experts" informing government officials of the appropriate action, and an army of public relations specialists engineering public consent.

In the currently militarized socio-political environment, any efforts at education would be interpreted as partisan at best, and elitist an untrustworthy at worst.

The longest-term strategy to defend against memetic attack, and ultimately the most effective one in our opinion, is to strengthen the social and cultural resiliency of the population under attack. Human beings have evolved complex and adaptive strategies for social cohesion. Our neurology is primed to establish rapport with other humans, to utilize reciprocal altruism, and to work toward common goals. Such social relationships require real-world, organic calibration to take effect. The establishment of rapport, for example, depends on eye contact, synchronized respiration, and recognition of subtle changes in vocal timbre.

In virtual spaces, these mechanisms cease to function. In fact, when human beings fail to establish "social resonance" through digital media, they tend to blame not the low fidelity of the medium, but the trustworthiness of the other party.[39] We repeat: the inability to establish organic social bonds through digital media increases our suspicion of one another, not the medium through which we are failing to connect.

This creates the perfect preconditions for memetic attack. The people, newscasters, friends, and experts we encounter through digital media are not trusted. The bots, algorithms, images and ideas to which we are exposed, on the other hand, are accepted at face value.

The only surefire safeguard against this state of vulnerability is to reaffirm the live, local, social, organic relationships between the people in the target population. This means challenging the value of time spent socializing and entertaining themselves on digital platforms, and giving people enough minutes of non-connected, social experiences each day to anchor live human-to-human connection as the primary form of social engagement.

People with some live experience of local politics, mutual aid, and environmental maintenance will be more resistant to the memetic constructions of the synthetic ideological landscape. They will be more likely to blame low fidelity on technology than one another, and less likely to accept the false, anti-social premises of angry, sensationalist memes. The less alienated a population is from one another, the harder it is to turn them against one another through polarizing memetics.

Of course, we recognize the tremendous challenge this poses to an economy committed to the growth of digital platforms for its sustainability. But re-establishing organic human relationships and a local social fabric requires activation and leverage of real biological mechanisms, not the further application of abstract biological metaphors.

Memetics best serves memes themselves. Memetic countermeasures only further weaponize the environment. Upgraded algorithmic filtering of dangerous memes can only result in a technological arms race between memetic engineers and platform censors. As collateral damage, it adds black boxes to the space of public discourse, magnifying distrust while doing nothing to arm the human beings against memetic provocation. It is also unlikely that publicly traded companies dependent on the perception of the size and centrality of their platforms would voluntarily restrict the activity of the bots and algorithms generating such a high volume of activity.

While a turnkey, technological solution to the problem of memetic attack is certainly an appealing prospect, the strengthening of local, human social relationships is the surest countermeasure to ideological warfare that depends on alienation for its survival.

Author biographies

Douglas Rushkoff | Research Fellow, IFTF

Douglas Rushkoff is a writer, documentarian, and lecturer whose work focuses on human autonomy in a digital age. He is the author of fifteen bestselling books on media, technology, and society, including *Program or Be Programmed, Present Shock,* and *Throwing Rocks at the Google Bus.* He has made such award-winning PBS Frontline documentaries as *Generation Like, Merchants of Cool,* and *The Persuaders,* and is the author of graphic novels including *Testament* and *Aleister & Adolf.* Rushkoff is the recipient of the Marshall McLuhan Award for his book *Coercion,* The Jacques Ellul Award for his documentary *The Merchants of Cool,* and the Neil Postman Award for Career Achievement in Public Intellectual Activity. Named one of the world's ten most influential intellectuals by MIT, he is responsible for originating such concepts as "viral media," "social currency," and "digital natives." Today, Dr. Rushkoff serves as Professor of Media Theory and Digital Economics at CUNY/Queens, where he recently founded the Laboratory for Digital Humanism and hosts its TeamHuman podcast.

David Pescovitz | Research Director, IFTF

David Pescovitz is co-editor/partner at the influential tech/culture Web site Boing Boing. He's also the co-founder of Ozma Records, a music label focused on the intersection of science and art to spark the imagination. The label's first release is the Voyager Golden Record: 40th Anniversary Edition, a lavish vinyl box set documenting the iconic phonograph record that NASA launched into space in 1977. Pescovitz was also the founding editor-at-large for *MAKE:,* the DIY technology magazine, and co-wrote the book *Reality Check* (HardWired, 1996), based on his long-running futurist column in *Wired* magazine. He has written for *Scientific American, Popular Science,* The *New York Times,* The *Washington Post, Salon,* and *New Scientist,* among many other publications. From 2000 to 2007 Pescovitz was the first ever writer-in-residence at UC Berkeley's College of Engineering. He holds a Bachelor of Fine Arts in Electronic Media from the University of Cincinnati and a Master's in Journalism from UC Berkeley.

Jake Dunagan | Research Director, IFTF

Jake Dunagan's focus is on the examination of social invention and the redesign of systems in light of current and emerging tools and knowledge. He has conducted research, written, and lectured on a wide range of topics, including strategic foresight, governance design, comprehensive cognitive design, intellectual property, judicial foresight, kids' technologies, Internet law and politics, mobility, and the future of learning and work. Dunagan joined IFTF in 2008 and brought with him the alternative futures methodology developed at the Manoa School of Futures Studies. Jake holds a BA in visual anthropology from Auburn University, an MA in visual anthropology from Temple University, and a PhD from the Manoa School of Futures Studies at the University of Hawaii with an emphasis in neuropolitics, governance design, alternative futures, and communication of foresight. He is currently an adjunct professor in the MBA in Design Strategy program at the California College of the Arts, where he teaches strategic foresight, tactical media, and social invention.

Endnotes

1. Renner, N. (January 30, 2017). Memes trup6mp articles on Breitbart's Facebook page. *Columbia Journalism Review.* Retrieved from: http://bit.ly/2pCf3Rs
2. UIC journal shows 'bots' sway political discourse, could impact election. (November 4, 2016). *UIC Today.* Retrieved from: http://bit.ly/2iGmzt6
3. Pro-Trump highly automated accounts 'colonised' pro-Clinton Twitter campaign. (November 17, 2016). *University of Oxford News.* Retrieved from: http://bit.ly/2y9o3lZ
4. Trabulsi, A. The Next Front Line in the Fight for Hearts and Minds. (March 4, 2014). *Stratfor Worldview*. Retrieved from: http://bit.ly/2iIhsly
5. Dawkins, R. (1986). *The Blind Watchmaker.* New York: Norton.
6. Rushkoff, D. (1994). *Media Virus!: Hidden Agendas in Popular Culture.* New York: Ballantine Books.
7. Wagner, K. (April 14, 2017). Facebook is deleting a bunch of spam accounts in its effort to fight fake news. *Recode.* Retrieved from: http://bit.ly/2oLxru9; Confessore, N et al. (January 27, 2018). The follower factory. *New York Times.* Retrieved from: http://nyti.ms/2rJ8YZM
8. Rosenberg, M et al. (March 3, 2018). How Trump Consultants Exploited the Facebook Data of Millions. Retrieved from: https://nyti.ms/2HH74vA
9. Guilbeault, D. (September 17, 2017). Personal interview.
10. West, J. (April 2, 2016). Microsoft's disastrous Tay experiment shows the hidden dangers of AI. *Quartz.* Retrieved from: http://bit.ly/2mPetCI
11. Russo, M. (April 23, 2014). NYPD Social Media Outreach Backfires When Twitter Answers #myNYPD Campaign. *NBC 4 New York.* Retrieved from: http://bit.ly/2zGPvfj
12. Mønsted, B. et al. (2017). Evidence of Complex Contagion of Information in Social Media: An Experiment Using Twitter Bots. *arXiv.* https://arxiv.org/abs/1703.06027
13. Aral, S., Nicolaides, C. (2017). Exercise contagion in a global social network. *Nature Communications 8.* doi:10.1038/ncomms14753. Retrieved from: http://go.nature.com/2oIvUEI
14. Akpan, N. (December 5, 2015). The very real consequences of fake news stories and why your brain can't ignore them. *PBS News Hour.* Retrieved from: http://to.pbs.org/2hfldF4
15. Lovink, G. (June 22, 2017). Interview with Oliver Vodeb (Memefest) on the Addictive Power of Memes Today. I*nstitute of Network Cultures.* Retrieved from: http://bit.ly/2jdbC6d
16. Arnold, C. (Sept 28, 2016). The Viruses That Made Us Human. *Nova Next.* Retrieved from: http://to.pbs.org/2dgcXTh
17. Meme Epidemiology. *SRI International.* Retrieved from: http://bit.ly/2m6PU4T
18. Wang, L., Wood, B. C. (2011). An epidemiological approach to model the viral propagation of memes. *Applied Mathematical Modeling,* 35 (11). Retrieved from: http://bit.ly/2jeFeQJ
19. Ibid
20. Berger, J., Milkman, K. L. (2012). What Makes Online Content Viral? *Journal of Marketing Research* 49 (2). Retrieved from: http://bit.ly/2yqOHY9
21. Shubeck, K.T., Huette, S. (2015). Predicting Meme Success with Linguistic Features in a Multilayer Backpropagation Network. *Proceedings of the 37th Annual Conference of the Cognitive Science Society.* Retrieved from: http://bit.ly/2AnGSCu
22. He, S., Zheng, X., & Zeng, D. (2016). A model-free scheme for meme ranking in social media. *Decision Support Systems,* 81. doi:10.1016/j.dss.2015.10.002. Retrieved from: http://bit.ly/2hTYSOl
23. Weng, L. Menczer, F., Ahn, Y. (2014). Predicting Successful Memes using Network and Community Structure. Proceedings of 8th AAAI Intl. Conf. on Weblogs and social media. Retrieved from: http://bit.ly/2zsDqrd
24. Qiu, X., Oliveira, D.F.M., Shirazi, A.S., Flammini, A., Menczer, F. (2017). Limited individual attention and online virality of low-quality information. *Nature Human Behavior,* 1. doi:10.1038/s41562-017-0132. Retrieved from: http://go.nature.com/2AuS7Kn
25. ibid
26. Ciampaglia, G. (September 11, 2017). Personal interview.
27. Kahneman, D. (2011). *Thinking, fast and slow.* New York: Farrar, Straus and Giroux.
28. Young, D. G. (2008). The Privileged Role of the Late-Night Joke: Exploring Humor's Role in Disrupting Argument Scrutiny. *Media Psychology,* 11. doi:10.1080/15213260701837073. Retrieved from:http://bit.ly/2g1IADX

29. How Your Brain Makes Articles Go Viral. (February 27, 2017). *Communication Neuroscience Lab, Annenberg School for Communication News, University of Pennsylvania.* Retrieved from: http://bit.ly/2yJZTTi

30. Schola, C., et al. (2017). A neural model of valuation and information virality. *PNAS 114, 11*. Retrieved from: http://www.pnas.org/content/114/11/2881

31. Genevsky, A., Yoon, C., Knutson, B. (2017). When Brain Beats Behavior: Neuroforecasting Crowdfunding Outcomes. *Journal of Neuroscience,* 37 (36). doi:10.1523/JNEUROSCI.1633-16.2017. Retrieved from: http://bit.ly/2hl7Yqe

32. Ingeborg van Driel, I., Grabe, M. E., Bas, O., Kleemans, M. (2016). Demographic variation in how the social brain processes news messages. *Politics and the Life Sciences* 35 (1). Retrieved from: http://bit.ly/2zsINXb

33. Ingeborg van Driel, I., Grabe, M. E., Bas, O., Kleemans, M. (2016). Demographic variation in how the social brain processes news messages. *Politics and the Life Sciences* 35 (1). Retrieved from: http://bit.ly/2zsINXb

34. Rushkoff, D. (1994). *Media virus!: Hidden agendas in popular culture.* New York: Ballantine Books.

35. Dawkins, R. (1989). *The selfish gene.* Oxford: Oxford University Press.

36. Dobbs, D. (December 3, 2013). Die, selfish gene, die. *Aeon.* Retrieved from: http://bit.ly/2znbkQF

37. Isaac, M., Shane, S. (October 2, 2017). Facebook's Russia-Linked Ads Came in Many Disguises. *The New York Times.* Retrieved from: http://nyti.ms/2g4eVIj

38. http://www.darwins.hu

39. Softky, W., Benford, C. (2017). Sensory Metrics of Neuromechanical Trust. *Neural Computation* 29 (9). Retrieved from: http://bit.ly/2zoCLtd