Influencers, Backfire Effects and the Power of the Periphery

(forthcoming in *Personal Networks,* edited by Mario L. Small, Brea L. Perry, Bernice Pescosolido, and Edward Smith. Cambridge: Cambridge University Press.)

Damon Centola*

*Annenberg School of Communication, School of Engineering and Applied Sciences & Department of Sociology, University of Pennsylvania, 3620 Walnut St., Philadelphia, PA 19104. Contact: dcentola@asc.upenn.edu

Influencers, Backfire Effects and the Power of the Periphery

Abstract:

Katz and Lazarsfeld's *Personal Influence* introduced the world to the impact of networks on the dissemination of mass media. Their "two-step flow" model showed that broadcast signals reached most of the public by being filtered through well-connected people – "opinion leaders" – who were the primary receivers of media messages, and the primary vehicles through which those messages were then disseminated to everyone else. Scientific and industry attention soon shifted to the task of identifying who the opinion leaders were, and how they could be targeted to spread new content. I trace these intellectual developments through to the arrival of social media, which brought greater attention to the idea of "central" players – or "influencers" – in the social network, as the key leverage points for disseminating products, ideas, and political messages. I show how this scientific search for the sources of social influence eventually led to a paradox: the unlikely finding that many social contagions do not spread from the central players to the periphery, but rather from the network periphery to the center. To explain these startling findings, the distinction between simple contagions, like information and viruses, and complex contagions, like social innovations and political movements, shows how the spread of new ideas through social networks depends in counterintuitive ways on the complexity of the contagion and the structure of the social network.

Influencers, Backfire Effects and the Power of the Periphery

How are people convinced to change their minds? What makes them switch to a new political candidate, decide to join a contentious social movement, or become willing to vote in an upcoming election?

In 1944, social scientists Paul Lazarsfeld, Bernard Berelson and Hazel Gaudet published a watershed study that placed social networks into the center of the picture of political influence. The accepted view of the day was that broadcast media dominated the political sphere, reaching out to the electorate and influencing voting behavior. *The People's Choice* showed otherwise. Lazasfeld, Berelson and Gaudet showed that it was people – friends, neighbors and family members – who were primary channels through which people learned about political candidates.

Their real discovery was that mass media still mattered, but surprisingly the way that media signals reached most people was through their social networks. Lazarsfeld, Berelson and Gaudet's new insight was that although media signals were broadcast across the continent, they only landed on a small fraction (perhaps 5%) of individuals. These individuals acted as secondary channels of influence and dissemination for political messages. These rare people were dubbed "opinion leaders" – the human relay stations that political influence traveled through to go from mainstream media to the public.

This insight initiated a new scientific investigation into the topic of opinion leaders, and the question of just how influential they really were for the spread of everything from political opinions to consumer goods. A decade later, the resulting study by Elihu Katz and Paul Lazarsfeld, *Personal Influence*, was an instant classic. It was landmark analysis of who the opinion leaders were, and how they influenced people.

The impact of *Personal Influence* is partly due to the fact that it is a surprisingly subtle and detailed exploration of social influence. Katz and Lazarsfeld's analysis primarily focused on 800 women from Dectur, Illinois and the features of their personalities and relationships that facilitated the flow of information from the mass media through their community. Katz and Lazarsfeld's study revealed that features of social status mattered for influence, but so did the groups that people belonged to, the friends and neighbors they were close to, and their perceived expertise on certain topics. In short, the patterns of influence were complex.

But, one thing was clear. The influence of mass media was not a straight shot from broadcast towers to consumer behavior. More often than not, Katz and Lazarsfeld found that a small group of people gave a disproportionate fraction of their attention to the media. And, those people became the most relevant sources of influence for their peers. Within their social circles, these opinion leaders were the key players in disseminating media messages.

For Katz and Lazarsfeld, opinion leaders are not people like Oprah Winfrey. Rather, Oprah's media empire would be considered a channel of broadcast

communication. Katz and Lazarsfeld's insight was that media signals traveled in steps – from broadcast towers to the opinion leaders, and from opinion leaders to the public. According to this "two-step flow" model, broadcast signals reached most of the public by being filtered through well-connected people in their social networks – their sister-in-law, or their friendly colleague – who were the primary receivers of media messages, and the primary vehicles through which those messages were then disseminated to everyone else.

The two-step flow model was only one part of Katz and Lazarsfeld's investigation. They also developed nuanced, but unfinished accounts of what makes certain people influential, and how that influence can hold sway across contexts, or vary across topics. These questions have since become productive areas of research for other sociologists (Lazarsfeld and Merton 1954, Davis and Greve 1997, Katz 1957, Reagans and McEvily 2003, Centola 2011). But, these are not the insights for which Katz and Lazarsfeld's study became famous. Rather, it was the key idea of the opinion leader that captured people's imagination. The hypothesis that a small number of people pay disproportionate attention to media, and therefore have disproportionate influence in disseminating it, became a focal point for a generation of communication scholars who continued to explore the two-step flow model.

But, it was the *second* step in the two-step model that interested most people. Was it true that a small number of people could influence everyone else? Who are these influential people? How can product advertisers and political campaigns target

them? These questions animated the work of generations of sociologists, political scientists, organizational scholars, marketing professionals, and most recently, network scientists. Many of the subtleties of Katz and Lazarsfeld's original work were eventually refined into a single idea – articulated by Malcolm Gladwell's ominous phrase, "the law of few." The idea is that there are a few special people out there, and, if you can find them, they are the key to disseminating a new idea, product or candidate to everyone else.

The Modern Influencer

Enter social media, and the birth of the modern "influencer."

Today, the notion of the opinion leader is not what it once was. The most significant reason for this is the unprecedented diversity of people's media diets, and the near ubiquity with which people are exposed to niche, almost personalized, media sources. The original question that Katz and Lazarsfeld set out to answer – regarding the influence of mass media on consumer behavior – is less fashionable today, largely because institutions like "the Media" (and "the News") have ceased to exist as they once did. But, Katz and Lazarsfeld's crucial insights into opinion leaders have endured. This is because Katz and Lazarsfeld looked beyond the influence of the media, into the fundamental structure of social relations that controls the spread of personal influence. It thus remains a guiding force in contemporary research on social networks. The question that animates much of today's scientific work on social

influence is a reformulated version of the question that motivated Katz and Lazarsfeld's original investigation: 'How do new ideas, information and innovations become widely accepted by spreading through people's social networks?'

Generations of research building on Katz and Lazarsfeld's ideas have pushed scientific thinking beyond a psychological focus on the individual-level attributes of influencers (like people's personality, charisma or attractiveness) to the structural insight that was nascent within the Katz and Lazarsfeld's original work – the idea that there is a social structure that mediates the process through which social contagions spread from one person to the next. This idea led to the discovery of measurable patterns within social networks – much like Simmel's idea of the "web of group affiliations" (Simmel 1955) – that help to explain how new ideas and behaviors propagate through a society (Granovetter 1973). A central feature of this approach is the insight that a small number of individuals are located at key points in the network structure, which makes them essential for the rapid dissemination of social contagions.

Today, this concept of "structural position" is a core principle of the field of network science. From a networks perspective, the notion of opinion leaders – or "influencers" – boils down to the essential question of which network positions have the most power for spreading new ideas. The classic answer is "centrality" (Newman 2010). The idea is that people at the center of a social network are the individuals who are best positioned to spread social contagions to everyone else. A narrow reading of

Katz and Lazarsfeld, and a charitable reading of Gladwell, unifies all of this research on the problem of personal influence, stretching back over three-quarters of a century, into a single analytical insight: the people at the center of the social network *are* the influencers.

The most popular methods for identifying central individuals in a social network are: (*i*) "degree centrality" (individuals with the most connections); (*ii*) "betweeness centrality" (individuals through which most paths must travel, going from one part of a network to another); and (*iii*) "eigenvector centrality" (individuals whose neighbors are highly connected). In the study of personal networks, these three measures are often collapsed down into a single metric – *viz.*, who is the most highly "connected" individual? That person, or those people are considered to be the influencers.

The Puzzle of the Periphery

Today, this long and important tradition in the study of personal influence has become so refined that it has uncovered a new and unexpected puzzle. This puzzle emerges from the fact that much of the contemporary work on opinion leaders, influencers and network centrality has been crystalized into the core insight that social innovations spread by propagating from the *core* of the network to the *periphery* – that is, it spreads from highly connected people at the "center" of the social network, to the less connected people at the periphery of the network. The puzzle is that this

widely accepted theory of the dynamics of social influence does not match up with the most current data on how social innovations – ranging from new technologies to progressive social movements – actually spread through social networks (State and Adamic 2015, Romero et al. 2011, Steinert-Threkeld 2017, Traag 2017, Sprague and House 2017, Mønsted et al. 2017, Centola 2018, Guilbeault et al. 2018a, Beaman et al. 2016).

The solution to this puzzle reveals the role of personal networks in the flow of social innovations, and the surprising implications for the modern notion of the "influencer." The solution starts with an overlooked distinction between *simple contagions*, like simple information and familiar ideas, which spread from a single contact; and *complex contagions*, like changes in workplace culture, or the adoption of innovative technologies, which typically require reinforcement from several peers before people are willing to adopt them (Centola and Macy 2007, Centola et al. 2007, Centola 2010, Centola 2015). The differences between simple and complex contagions have striking implications for how personal networks control the spread of innovations, the dispersion of ideas, the adoption of inventions, the growth of social movements, the success of political campaigns, and the uptake of new health behaviors (Steinert-Threkeld 2017, Traag 2017, Sprague and House 2017, Guilbeault et al. 2018a).

The classic model of social influence in which the highly connected "influencer" is the source of social change assumes that all contagions are simple.

However, social innovations that are risky or require social coordination are typically complex. For instance, innovative technologies like Facebook and Twitter usually depend on people coordinating with multiple peers before they are willing to adopt (Ugander et al. 2012, Toole et al. 2011). Similarly, successful political campaigns and social movements, like support for marriage equality and the #MeToo movement, are also complex contagions (State and Adamic 2015, Traag 2017, Steiner-Threkeld 2017, Mønsted et al. 2017). People need to believe that a movement is legitimate before they are willing to add their voice to the chorus supporting it. Changes in workplace culture, and the growth of new social norms concerning gender relations within organizations also spread through the dynamics of complex contagion (Kanter 1977, Centola et al. 2018). In fact, the only social contagions that actually spread in the way that the classic model describes – from highly connected influencers at the core of the social network to the modestly connected members network periphery – are familiar ideas, information, and products. Social innovation, however, is a whole other thing.

The Surprising Limits of Influencers for Innovation Diffusion

To see why highly connected people – whom I will refer to hereafter as "influencers" – do not actually play the dominant role in initiating change, we will examine the spread of an innovative technology through the personal work networks of tech industry managers with varying levels of social connectedness.

Consider a C-level manager at a Fortune 500 company who is extremely well-connected. This manager has hundreds of professional contacts, which makes it very easy for her to discover innovative ideas floating around in the vast network that surrounds her. She may be among the first people to find out about an innovative new process-centered management approach, such as Total Quality Management, or "TQM". Although the TQM approach may seem unorthodox to her colleagues and peers, the Fortune 500 manager is nevertheless in an excellent position to be among the leaders in its adoption, responsible for spreading the innovative idea to her large network of connections.

Let's compare this Fortune 500 manager to another manager who is at a small startup company. This manager has many fewer professional connections – perhaps only a few dozen – and thus she has much less access to all of the new management innovations that may be percolating through various parts of the social network. Because the startup manager is less connected, she has much less power to spread any innovative ideas that she discovers. Consequently, the startup manager appears to be much less important than the Fortune 500 manager for spreading an innovative organizational practice like TQM.

However, while highly connected influencers, or "hubs" in the social network, have access to many people in the community –and thus are excellent vehicles for spreading innovations to lots of people –they are also susceptible to peer influences from lots of people. For innovations that are risky or require legitimacy, the influence

from a single adopter is typically weighed against the countervailing influences from all of a person's peer contacts who have not adopted the innovation. When risk or legitimacy are involved, the non-adopters in a person's social network weigh against the decision to adopt. Before adopting a costly or unfamiliar innovation, each manager must also consider the signals coming from all of these countervailing influences.

The Fortune 500 manager has the natural advantage that her vast network of contacts will lead her to discover the innovation early on, before it has become widely adopted. When considering her decision to adoption, her ability to make early discoveries also means that any innovative practices she encounters will not yet be adopted by the vast majority of her contacts. In this light, she needs to consider several things before making the decision to adopt. First, she needs to know if there is enough evidence to support the claim that a new management practice would be successful at her type of company, at her scale, and in her sector. Without seeing peer institutions successfully implement the innovative practice, it would be a gamble to invest in making a large organizational change to an unproven management practice. Second, if she decides to adopt TQM before most of her peers and competitors do, will that decision be seen as savvy or reckless by her more conservative colleagues? That is, will she incur negative reputation effects from adopting an innovation that has not yet been validated by her peer community? Third, and perhaps most importantly, what are the potential reputation effects that she will suffer if the

decision to adopt this innovation is a bad one? Because she is so well-connected, she is a highly visible actor, and thus she is subject to scrutiny from lots of people. If she waits to adopt until the innovation becomes accepted among several of her peer institutions, and many of her contacts have endorsed it, then she cannot be held personally accountable for her decision, given that it had become an industry norm. If the innovation then fails, this failure can be attributed to a collective error rather than to an individual one. By contrast, if she adopts the innovation upon her first exposure to it, as an early mover she will be individually accountable for any failures that come as a result of that decision. For all of these reasons, the Fortune 500 manager may wait until there is substantial social confirmation for the innovative management practice before she will be willing to adopt it.

By contrast, the startup manager is in a different position. She may have all the same concerns as her Fortune 500 colleague, however she has a much smaller network of countervailing influences to consider. Early exposure to an innovative practice from just a few sources of reinforcement will have a much greater impact on her decision-making since this small number of contacts constitutes a much larger fraction of her professional network. For the startup manager, even a modest level of social reinforcement for the innovation may constitute a strong enough social signal for her to consider the innovation seriously. In addition, the startup manager is not subject to the same level of scrutiny as the Fortune 500 manager, nor is she as pressured by the same scale of reputation effects that might serve to inhibit her from

being willing to be an early adopter. For the startup manager, as compared to the Fortune 500 manager, the greater relative impact of a few early adopters, along with her exposure to fewer countervailing influences, make her much more likely to be an early adopter of the innovative practice.

The strategic implication is that innovations may spread far more effectively by percolating through "peripheral" clusters of less connected actors, rather than by trying to spread through highly connected network hubs. For simple contagions, hubs are excellent spreaders because of their vast access to the social network, which makes it very easy for them to become infected with new ideas early on in a spreading process. It also allows them to rapidly spread their infection to their large number of contacts. However, both of these advantages for spreading viral content –that is, early exposure and a large number of contacts –work against the spread of complex contagions.

For complex contagions, early detection of an innovation by hubs also means that most of the hubs' contacts will not have adopted the innovation. Moreover, because a hub is so visible, she is also likely to be very cautious before adopting an unproven innovation. Far from being the initial seed that launches the diffusion of an innovation, a highly connected Fortune 500 manager may be among the last people to adopt an innovative management practice.

Ironically, hubs in the social network may often be a roadblock for spreading particularly novel innovations. In order for "game changing" innovations (that is,

14

innovations that are particularly unfamiliar, unusual, or disruptive) to spread

effectively, they typically need to follow peripheral pathways composed of clustered

networks of interlocking managers at smaller firms who can build a critical mass for

the innovation. The historic growth of innovative technologies such as Facebook,

Twitter and Skype (Ugander et al. 2012, Toole et al. 2011, Karsai et al. 2016, Bakshy et

al. 2009) have all shown that the need for social reinforcement means that diffusion

will be most effective through peripheral channels of overlapping social ties. Once an

innovation starts to catch on and spread through these peripheral channels, it can gain

enough traction in the network to activate a substantial fraction of adopters. Once

enough of the hub's contacts finally adopt, there will be enough reinforcement in the

Fortune 500 manager's network to convince her that this innovation is credible and

that she should adopt it, too. In other words, once an innovation finally takes off,

that's when the highly connected leaders typically adopt it (Centola 2018, 2019).

The fatal mistake often made by entrepreneurs and marketers trying to spread

innovations is to misread the rapid growth in the uptake of an innovation that *causes* a

Fortune 500 manager's adoption, as the *result* of the Fortune 500 manager's adoption.

This fallacy comes from an obvious (but unscientific) observation. This

observation is that the moment of increased growth that marks the sharp acceleration

in the spread of an innovation (often referred to the "elbow" in the growth curve)

happens around the same time that highly connected people start adopting. The

fallacy is that these highly connected adopters are the *cause* of that growth. It is

certainly true that growth rates increase dramatically around the same time that the Fortune 500 managers start adopting. However, the massive acceleration in adoption that coincides with the Fortune 500 manager's adoption is not the *cause* of the innovation's success. Quit the opposite.   The growing success of the innovation is the cause of the highly connected manager's willingness to adopt it. The fallacy is to think that targeting social connectors will jump-start a successful diffusion campaign.

The evidence from research on complex contagions (Toole et al. 2011, Wang et al. 2019, Centola 2018) shows that the most successful strategy for spreading an innovation, particularly and unusual or costly one, is to target peripheral communities of moderately connected individuals, who are much better positioned to grow a groundswell of support. Once this groundswell grows large enough, it can then capture the attention of a Fortune 500 manager. By the time highly connected hubs start adopting an innovative practice, the innovation has typically gained enough support that it has already begun a period of rapid, widespread growth throughout the population.

The more innovative an idea is, the more important this fact about personal networks becomes. For innovations that are unfamiliar or disruptive, success comes from targeting social networks that can provide social reinforcement for a new idea. These social networks can be thought of as "incubators", which can grow expanding pockets of support for an innovation.

Ironically, the reason that incubators in the network periphery can be so effective for sparking successful innovation diffusion is the exact opposite reason that highly connected hubs are so appealing. Hubs are appealing because they provide massive amounts of social exposure for an innovation. We might intuitively expect that increased exposure would lead to increased spreading. This is certainly true for simple contagions. However, for complex contagions, peripheral network channels are successful precisely because they protect potential adopters from being overwhelmed by exposure to too many countervailing influences from the non-adopters in the population. It is this protection that is can be essential for spreading an innovative idea.

Influencer "Backfire" Effects

"The law of the few" is a notion that has become widely accepted partly because there are lots of situations for which it works amazingly well. For instance, many people know the now-infamous (and apocryphal) story of Gaetan Dugas, the highly sexually active flight attendant who claimed to have had over 2,500 sexual partners; and whose promiscuity was thought to have played a significant role in the early spread of HIV/AIDS.

For medical scientists who study the spread of diseases, the notion of highly connected influencers offers an essential insight into how social networks trigger epidemics like HIV/AIDS and COVID-19 – and the pivotal role that highly

17

connected people can play. Of course, this idea goes far beyond epidemiology. It has also helped to animate our imaginations when it comes to thinking about the spread of unusual cultural fads, the best strategies for developing our professional networks, and the social contours that define major historical events.

But, this imagination is precisely the problem. The spread of social innovations – within organizations and within our society – does not actually resemble the spread of a virus. It is the generalization from simple contagions to complex contagions that is problematic. The problem is not simply that this generalization has led to failures. The problem is that it can lead to backfire. Highly connected social stars within personal networks may not only *not help* innovations to succeed, they can inadvertently *undermine* any future attempts at innovation.

Why? Because exposure is not always a good thing. For simple contagions, exposure equals awareness, which translates into adoption. An effective strategy for creating awareness will always be to increase the connectedness of the adopters. This assumption does not generalize to many complex contagions.

For example, let's assume that a highly connected social star has effectively spread the word about an innovation. Everyone knows about it. But, what happens if no one adopts it? Because of the popularity of the influencer who is spreading the word, the awareness campaign may have been *too* successful. Here's why. A successful awareness campaign not only makes people aware of the product, it also makes people aware that *everyone else* is aware of it. An effective word of mouth

campaign includes not just information about the innovation, but also information about the fact that the innovation is well-known.

What does it signal about an innovation if everyone knows about it – and everyone knows that everyone else knows about it – and yet everyone also knows that no one has adopted it? The countervailing influences from a large number of peers – who are known to be aware of an innovation *and* who have also elected not to adopt it – provide evidence for everyone else that the innovation is undesirable. The non-adopters offer an implicit social signal about the illegitimacy of the innovation. Before the awareness campaign succeeded, an innovation may simply have been seen as contentious or unfamiliar. The challenge for spreading this kind of unknown innovation is simply that may not yet be seen as credible or legitimate. Things are different after an awareness campaign succeeds. If awareness of the innovation has been widely diffused, but the innovation itself has still has not caught on, then there is a much bigger problem. The sea of non-adopters *confirms* the innovation's illegitimacy. The problem now is that *it appears that* the innovation has been deliberately and publicly rejected by everyone.

The implicit social signal suggests that the innovation may carry a social stigma. This perception of stigma is created by the success of the influencer's awareness campaign, combined with the utter failure of their adoption campaign. This pinpoints the crucial difference between *awareness* and *acceptance*. For simple contagions, these two are treated as the same thing. For a complex contagion, they can be different,

and that difference can have crucial backfire effects if innovators succeed at creating awareness but fail to create acceptance.

Consider Google's failed attempt to bring Google Glass to market. There was widespread awareness of Glass (we all knew about it), combined with widespread awareness of its lack of acceptance as a product (we all know that it failed). This created a negative social stigma, not just around the product, but also around the company. This stigma became such a significant impediment for the company, that it prevented Google's future attempts to reboot the product line. Google and other companies are now struggling to figure out how to make similar kinds of products grow in the soil that Google's early efforts have inadvertently salted.

The distinction between awareness and acceptance offers a valuable lesson for any diffusion strategy that would attempt to use a "kitchen sink" approach to the spread of innovations – in which all network strategies are attempted simultaneously. One important reason not to use a simple contagion strategy – such as the activation of highly connected influencers – to spread a complex contagion is that successfully creating tremendous awareness for an innovation, while simultaneously failing create corresponding levels of acceptance for it, can backfire. The result is not a *small gain* in adopters. But, a *large loss* in public opinion about the company and the product line. Awareness without acceptance can create a well-known, but socially stigmatized innovation.

The downside of a highly successful awareness campaign is that everyone remembers the product. And, if it fails, everyone also remembers how badly it failed. Future attempts to spread similar innovations are likely to face even greater barriers to adoption than the initial innovation. A better strategy for spreading a novel or unfamiliar innovation is to target the network periphery, where support for the innovation can grow slowly, gaining acceptance in the social network without the unwieldy burden of widespread awareness.

Complex Centrality and the Power of the Periphery

A recent series of studies showing the ineffectiveness of highly connected influencers for spreading change – particularly, for the spread of unusual or contentious ideas on social media – have also highlighted that there may be specific locations in the network periphery that are surprisingly effective for spreading these innovations (Bakshy et al. 2009, Barbera et al. 2015, Steinert-Threkeld 2017). To identify these locations, my student and I tested a series of different models of innovation diffusion on 74 empirically collected large social networks (Guilbeault and Centola 2020a), with the goal of discovering the essential locations for initiating diffusion. Strikingly, the results converged on a new kind of centrality – "complex centrality" – which consistently identified the most influential individuals in the social network for spreading social innovations.

The key to complex centrality is that it pinpoints network clusters that are located at the intersection of "wide bridges" between social groups. When we think of links between communities in a network, we typically imagine these links as a single "bridge tie" from one group to the other. But, a single tie is a very narrow bridge. Only a simple contagion can spread across it. For social innovations to effectively spread across communities, there must be wide bridges – composed of multiple overlapping ties – that can provide social reinforcement for a new idea to propagate from one group to the next (Centola and Macy 2007, Centola 2015). We discovered that the locations in the social network that are most effective for spreading innovations are the social clusters located at the intersection of these wide bridges.

These social clusters (called "bridging groups") are disproportionately influential because they are the most centrally located groups in the entire social network (Centola 2015). Crucially, however, the *individual members* of these social clusters are *not* central to the network. The members of bridging groups are not highly connected influencers, nor are they enterprising information brokers. In fact, individually, the members of these social clusters are indistinguishable from anyone else. They are unlikely even to know that they occupy a special location in the network. But, they are influential nevertheless.

The influence of these social clusters comes from the fact that they sit at the intersection of wide bridges. Collectively, the members of a bridging group have more wide bridges to diverse parts of a social network than any other social cluster in the

population. Bridging groups combine the *strength* of social reinforcement with the *accessibility* of boundary-spanning wide bridges – a rare and effective combination for diffusion (Guilbeault and Centola 2020a).

A recent study of peer influence on Chinese social media showed how effective these network locations can be (Wang et al. 2019, Centola 2019). This study compared the spreading patterns generated by highly connected network influencers (with more than 100,000 connections) to spreading patterns generated by regular users who were clustered amid wide bridges in the network periphery. When media content was a simple contagion, it spread most effectively from the highly connected influencers.

But when media content was politically charged, or pertained to normative topics, then highly connected influencers were only able to spread messages to their direct contacts. Beyond the influencer's immediate social circle, social contagions tended to die a quick death. For this kind of content, social clusters in the network periphery, composed of moderately connected users, significantly increased the spread of the messages – allowing media content to reach as much as five times farther into the population than when it came from highly connected influencers.

Strikingly, these locations in the network periphery are among the *least* effective places for initiating the spread of simple contagions. For viral diffusion, network redundancy and lower connectedness slows down the spreading process. However, these peripheral social clusters are ideal for initiating the spread of a social innovation.

In these network locations, change agents can coordinate with each other and with others to accelerate the growth of a critical mass (Guilbeault and Centola 2020a).

Influencers, Homophily and Bias

One sphere of social life in which influencers can be influential (and dangerously so) is in the spread of ideas and opinions that reinforce existing biases. Online and offline, when communities are organized homophilously, along the lines of shared political, social or cultural beliefs, then ideas that reinforce a community's existing beliefs are often simple contagions. They are easy to understand and easy to spread. Within political echo chambers, highly connected influencers at the center of the conversation can easily spread misinformation that plays to a group's biases (Becker et al. 2017, Becker et al. 2019, Guilbeault et al. 2018b, Guilbeault and Centola 2020b).

By contrast, contentious ideas that challenge a group's biases are complex contagions. These ideas face strong opposition, and thus are not likely to emerge from highly connected individuals facing a sea of countervailing influences. New ideas that challenge the status quo emerge more commonly from the moderately connected network periphery – where everyone's voice is equally heard, and where new ideas can be reinforced among peers and protected from too many countervailing influences (Samuelson and Zeckhauser 1988, Steinert-Threkeld 2017, Centola 2018).

This reveals an important asymmetry in the "influence of the influencer." A highly connected, or authoritative person can effectively spread ideas that reinforce a community's existing biases, but they will not be very effective for spreading beliefs and behaviors that run contrary to people's biases. This asymmetry in the influence of influencers – namely, they're good for spreading simple contagions, but poor at spreading complex contagions – is particularly consequential for disadvantaged communities. The potential danger this poses for online influences has become particularly salient in the discussion about health misinformation among communities that are distrustful of mainstream healthcare.

A generation of research on underserved communities, in particular African-American and Latina women, has found that the members of these communities report disproportionate levels of distrust toward mainstream medical care (Kennedy et al. 2007). And, not without cause – e.g., involuntary sterilization programs in the 1950s and 1960s (Kluchin 2009). As a result, highly connected influencers in these communities can be effective for spreading messages that amplify people's distrust of current preventive health measures (Blankenship et al. 2018), such as birth-control, vaccination and COVID-19 prevention measures. These biases may also make these communities disproportionately susceptible to malicious anti-vaccination campaigns that use highly connected, and well-disguised, social media "bots" to spread misinformation (Broniatowski et al. 2018).

The power of influencers to spread misinformation about topics such as vaccination, birth control and COVID-19 precautions can further exacerbate health inequities, making communities that are already distrustful of mainstream healthcare increasingly vulnerable to suffering negative, but preventable health outcomes. The asymmetry in the network dynamics of simple and complex contagions means that in order to be effective, public health campaigns cannot succeed simply by spreading information (Centola 2018, Centola *forthcoming*). If new preventive health advice challenges a community's existing biases, then successful information-spreading may not only fail, but also lead to *backfire*. Just like Google Glass, if everyone knows about the newly propagated health advice (and everyone knows that everyone knows about it), but everyone also knows that no one in their community has adopted it, then the success of the *informational* campaign may inadvertently manufacture social proof against behavior change. An important direction for future research is to discover new ways that public health dissemination can be approached *not* as a problem of *awareness*, but as one of *acceptance*. Dissemination campaigns that face these challenges can succeed by finding ways to target reinforcing ties among people in the network periphery, which can minimize countervailing influences, and mobilize a challenge to biased expectations about healthcare and health advice.

Conclusion

Katz and Lazersfeld opened the world's eyes to the importance of social networks in the process of media dissemination. Today, media is a different beast. The science of identifying and categorizing media signals has itself become a complex and multi-faceted sub-discipline of communication theory. Nevertheless, the impact of *Personal Influence* remains. Building on Katz and Lazarsfeld's groundbreaking insights, the science of social networks has made remarkable progress mapping the flow of ideas and behaviors through society. As Katz and Lazarsfeld first observed, in the viral spread of information, just as in the viral spread of epidemics, there are highly connected individuals who play a central role in the dissemination process.

These influencers are essential for spreading awareness. But, awareness is a simple contagion. For the spread of social change initiatives and innovative product ideas, the goal is not necessarily to spread awareness. In many cases, the goal is to spread *acceptance*. For these complex contagions, the key value of personal networks is their capacity to spread acceptance through peer reinforcement within the network periphery.

A novel or unusual idea is likely to die a quick death if it is given to a population all at once. It is likely to be overwhelmed by countervailing influences. The more that awareness outpaces acceptance, the greater the backfire may be, making any future diffusion efforts less likely to succeed. By contrast, a campaign that protects an innovation from these countervailing influences, may enable even an usual innovation to spread with surprising effectiveness by targeting the periphery of the

social network. In the periphery, peer reinforcement can allow a new idea to gain

acceptance within diverse pockets of support. Wide bridges and reinforcing ties can

carry the innovation from pocket to pocket, growing broader acceptance for the

innovation.

This strategy of growth through the network periphery does not relegate an

innovation to the margins. Many of the biggest success stories originated this way.

Several once-niche companies, like Facebook, Twitter and Skype, began their run to

fame by providing novel complementary technologies to tech-savvy youth. They all

reached their remarkable levels of global success by following the same pattern of

early growth through the network periphery (Ugander et al. 2012, Toole et al. 2011,

Karsai et al. 2016, Centola 2018). Once an unusual innovation gains a sufficient level

of acceptance in the periphery, it can then trigger a tipping point in the broader

population (Centola et al. 2018). Even an unlikely consumer product – like a photo-

sharing community on the Internet – may become so well-accepted that it overtakes

the vast majority of a population, bringing the "influencers" along with it.

References

Ancona, Deborah G., and David Caldwell. "Bridging the Boundary: External Activity and Performance in Organizational Teams." Administrative Science Quarterly 37 (1992): 634–665.

Bakshy, Etan et al., "Social Influence and the Diffusion of User-Created Content." In Proceedings of the 10th ACM Conference on Electronic Commerce, 325–334. New York: Association of Computing Machinery, 2009;

Beaman, Lori, et al. 2016. "Making networks work for policy: evidence from agricultural technology adoption in Malawi," Impact Evaluation Report 43. New Delhi: International Initiative for Impact Evaluation.

Becker, Joshua et al. 2017. "Network Dynamics of Social Influence in the Wisdom of Crowds." Proceedings of the National Academy of Sciences. 114 (26): E5070-E5076.

Becker, Joshua et al. 2019. "The Wisdom of Partisan Crowds." Proceedings of the National Academy of Sciences. 116 (22) 10717-10722.

Blankenship EB, et al.2018. Sentiment, contents, and retweets: A study of two vaccine-related twitter datasets. Perm J;22:17-138.

Broniatowski, David, et al. (2018) "Weaponized health communication: Twitter bots and Russian trolls amplify the vaccine debate." American Journal of Public Health, 108(10), 1378–1384

Burt, Ronald S. Structural Holes: The Social Structure of Competition. Cambridge, MA: Harvard University Press, 1992.

Centola, Damon and Arnout van de Rijt. 2015. "Choosing Your Network: Social Preferences in an Online Health Community." Social Science & Medicine 125: 19–31.

Centola, Damon and Michael Macy. "Complex Contagions and the Weakness of Long Ties." American Journal of Sociology 113, no. 3 (2007): 702–734

Centola, Damon et al. (2007) Cascade Dynamics of Multiplex Propagation Physica A 374, 449-456;

Centola, Damon et al. (2018). "Experimental evidence for tipping points in social convention." Science, 360(6393), pp. 1116-1119.

Centola, Damon "An Experimental Study of Homophily in the Adoption of Health Behavior." Science 334, no. 6060 (2011): 1269–1272.

Centola, Damon 2018, *How Behavior Spreads*, Princeton, NJ: Princeton University Press.

Centola, Damon (2019). "Influential networks." Nature Human Behavior. 3, 1–2.

Centola, Damon "The Social Origins of Networks and Diffusion." American Journal of Sociology 120, no. 5 (2015): 1295–1338

Centola, Damon "The Spread of Behavior in an Online Social Network Experiment." Science 329, no. 5996 (2010): 1194–1197

Centola, Damon (forthcoming) *Change: How to Make Big Things Happen*, (New York: Little, Brown & Co.).

Davis, Gerald and Henrich R. Greve. "Corporate Elite Networks and Governance Changes in the 1980s." The American Journal of Sociology, Vol. 103, No. 1. (Jul., 1997), pp. 1-37

Gladwell, Malcolm. The Tipping Point: How Little Things Can Make a Big Difference. Boston: Little, Brown, 2000.

Granovetter, Mark. "The Strength of Weak Ties." American Journal of Sociology 78, no. 6 (1973): 1360–1380.

Guilbeault, Douglas et al. 2018b. "Social Learning and Partisan Bias in the Interpretation of Climate Trends." Proceedings of the National Academy of Sciences. 115 (39) 9714-9719.

Guilbeault, Douglas and Damon Centola. "Optimal Seeding Strategies for Social Contagions." Working Paper, Annenberg School for Communication, University of Pennsylvania, Philadelphia, 2020a.

Guilbeault, Douglas and Damon Centola (2020b) "Networked collective intelligence improves dissemination of scientific information regarding smoking risks." PLoS ONE 15(2): e0227813.

Guilbeault, Douglas et al. 2018a. "Complex Contagions: A Decade in Review." In Complex Spreading Phenomena in Social Systems, edited by Yong Yeol Ahn and Sune Lehmann. New York: Springer Nature.

Hansen, Morten T. "The Search-Transfer Problem: The Role of Weak Ties in Sharing Knowledge across Organization Subunits." Administrative Science Quarterly 44, no. 1 (1999): 82–111

Howe, Lauren and Benoît Monin. "Healthier than thou? 'Practicing what you preach' backfires by increasing anticipated devaluation." Journal of Personality and Social Psychology, 2017 May 112(5):735.

Kanter, Rosabeth Moss. 1977. *Men and Women of the Corporation*. New York: Basic Books.

Karsai, Márton, et al. "Local Cascades Induced Global Contagion: How Heterogeneous Thresholds, Exogenous Effects, and Unconcerned Behaviour Govern Online Adoption Spreading." Scientific Reports (2016): 27178. doi.org/10.1038/srep27178.

Katz, Elihu, and Paul Lazarsfeld. Personal Influence. New York: Free Press, 1955.

Katz, Elihu. 1957. "The two-step flow of communication: An up-to-date report on an hypothesis." Public Opinion Quarterly. 21: 61–78.

Kennedy, BR et al. 2007. African Americans and their distrust of the health care system: healthcare for diverse populations. J Cult Divers. Summer;14(2):56-60.

Kluchin, Rebecca (2009) Fit to Be Tied: Sterilization and Reproductive Rights in America, 1950-1980. Rutgers: Rutgers University Press.

Krackhardt, David. "The Ties That Torture: Simmelian Tie Analysis in Organizations." Research in the Sociology of Organizations 16, no. 1 (1999): 183–210.

Lazarsfeld, Paul F., and Robert K. Merton. "Friendship as a Social Process: A Substantive and Methodological Analysis." Freedom and Control in Modern Society 18, no. 1 (1954): 18–66.

Lazarsfeld, Paul, et al. (1944). The People's Choice. Duell, Sloan & Pearce.

Mønsted, Bjarke et al. (2017). "Evidence of complex contagion of information in social media: An experiment using Twitter bots." PLoS ONE 12(9): e0184148.

Newman, Mark. *Networks: An Introduction*. (Oxford University Press, 2010).

Barberá, Pablo et al. "The critical periphery in the growth of social protests." PLoS ONE 10, e0143611 (2015).

Padgett, John F., and Christopher K. Ansell. "Robust Action and the Rise of the Medici, 1400–1434." American Journal of Sociology 98, no. 6 (1993): 1259–1319.

Page, Scott E. The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies. Princeton, NJ: Princeton University Press, 2007.

Reagans, Ray, and Bill McEvily. "Network Structure and Knowledge Transfer: The Effects of Cohesion and Range." Administrative Science Quarterly 48, no. 2 (2003): 240–267.

Romero, Daniel et al. 2011. "Differences in the Mechanics of Information Diffusion across Topics: Idioms, Political Hashtags, and Complex Contagion on Twitter." In Proceedings of the 20th International Conference on World Wide Web, 695–704. New York: Association of Computing Machinery.

Samuelson, William and Richard Zeckhauser. (1988). "Status quo bias in decision making". Journal of Risk and Uncertainty. 1: 7–59.

Simmel, Georg. Conflict and the Web of Group Affiliations. New York: Free Press, 1955.

Sprague, Daniel and Thomas House (2017) "Evidence for complex contagion models of social contagion from observational data." PLoS ONE 12(7): e0180802;

State, Bogdan and Lada Adamic. "The Diffusion of Support in an Online Social Movement: Evidence from the Adoption of Equal-Sign Profile Pictures." In Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work

& Social Computing, 1741–1750. New York: Association of Computing Machinery, 2015.

Steinert-Threlkeld, Zachary C. "Spontaneous Collective Action: Peripheral Mobilization during the Arab Spring." American Political Science Review 111, no. 2 (2017): 379–403.

Toole, Jameson L. et al. "Modeling the Adoption of Innovations in the Presence of Geographic and Media Influences." Plos One 7, no. 1 (2012): e29528.

Traag, Vincent. "Complex Contagion of Campaign Donations." PLoS One 11 no. 4 (2016): e0153539.

Ugander, Johan et al. "Structural Diversity in Social Contagion." Proceedings of the National Academy of Sciences 109, no. 16 2012): 5962–5966.

Vaan, Mathijs de, Balazs Vedres, and David Stark. "Game Changer: The Topology of Creativity." American Journal of Sociology 120, no. 4 (2015): 1144–1194.

Wang, Xioachen et al. 2019. "Anomalous structure and dynamics in news diffusion among heterogeneous individuals." Nature Human Behavior 3, 1–10.

Zhang, Jingwen and Damon Centola. 2019. Annual Review of Sociology. "Social Networks and Health: New Developments in Diffusion, Online and Offline." Vol. 45(1):91-109