Toward an Analytic Framework of Social Influence: Behavioral Diffusion

June 25-26, 2018 Meeting on
Exploring the Development of Analytic Frameworks:
A Pilot Project for the Office of the Director of National Intelligence
National Academy of Sciences Board Room

James A. Kitts
Director, Computational Social Science Institute
Professor, Department of Sociology
University of Massachusetts, Amherst

Yongren Shi
Postdoctoral Fellow
Yale Institute for Network Science
Human Nature Lab
Yale University

Word Count: 11,037

Supported by **National Institutes of Health under award R01HD086259.**
Any opinions, findings and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the NIH.

A GENERAL THEORETICAL FRAMEWORK FOR BEHAVIORAL DIFFUSION

Models of behavioral diffusion typically deal with binary states such as adopting a practice or technology, spreading a rumor, or participating in a protest. If our aim is to promote (or inhibit) diffusion of a particular behavior, we should have a deep understanding of the social context and clear guidance for how to use that understanding to design effective influence campaigns. In this project, **we contribute to the development of a general framework that will allow users to diagnose a situation and identify potent levers for influence** according to the best available research. This framework should then inform the development of a coherent influence (or counter-influence) strategy, including specific tactics and their sequence, timing, and location. To be effective, such a framework should be both general and parsimonious -- integrating strands of research in social psychology, organizational behavior, and network science -- but also concrete enough to translate into empirical interventions. This paper is a first step toward producing such a framework.

I. INFLUENCE AT THE MICRO-LEVEL

We begin with the most basic element of social influence, the *dyad* of Ego and Alter, and focus on the force of influence operating on Ego. Research has identified a variety of micro-level factors that amplify or attenuate the force of influence from Alter to Ego. Some of these are features of the sender(s) of influence, Alter. **First is the *number* of Alters promoting Ego's adoption** -- as a greater number of Alters exerts greater force on Ego. **The number of Alters who oppose Ego's adoption** clearly has the opposite effect. Somewhat less obviously, **the number of other Egos that are targets of influence** also matters because those Egos dilute influence and thus impede adoption. Ceteris paribus, the force of influence on Ego operates as a direct function of the number of proponents, and an inverse function of the number of opponents and fellow targets.

Second, a more nuanced model considers the *weight* of social influence from each Alter, as Ego may be affected by Alter's individual-level power, status, and strength. These weights may apply to influence from each proponent, each opponent, and each bystander or fellow target. See Wejnert (2002) for a discussion of many of these patterns. **Another key input to the weight of influence appears as features of the Alter-Ego dyad: physical distance (Latane et al., 1995), social distance (Suh et al., 2017), tie strength and valence (Macy et al. 2003), Ego's dependence on Alter (Emerson 1964), etc**.

We have proposed **two overall families of factors: the *number* of influencers (proponents, opponents, and bystanders or fellows targets) and the *weight* of influence from each Alter (strength, power, status, relationship or proximity to Ego).** A more elaborate model is offered by Social Impact Theory (Latane 1981), a general framework from social psychology that has been extensively investigated in empirical research and widely used in applied settings. According to SIT, The social force experienced by Ego at a given time is a function of the *number* of Alters influencing Ego, the *importance* of those Alters to Ego (status or power of Alter, strength of the Ego-Alter tie), and the *proximity* of those influence attempts to Ego in space and time (i.e., nearby and recent influence attempts exert the greatest force). In the model, these factors combine multiplicatively in exerting force on Ego. This function is analogous to how light perceived by Ego is a multiplicative function of the number of light bulbs, the intensity of each individual bulb, and their proximity to Ego. This framework is highly flexible, allowing that:

- Influence can be easily weighted by monadic features of sender or target (e.g. status characteristics, power, visibility).
- Influence can be weighted by dyadic features of sender-target pairs (physical distance, social distance, tie strength, affective valence as a positive or negative tie)
- This framework provides a principled way to incorporate temporal dynamics (e.g. attenuation of force on Ego as time passes since Alter's influence attempt).

We take this research program as an inspiration and as a starting point for our preliminary general framework.

A. The Number of Adopters, Opponents, and Targets

The force of influence to adopt a behavior (e.g. to spread a sensational online rumor about a political candidate) that Ego experiences is driven to a large extent by the number of positive influencers (e.g., peers who already support the rumor), the number of counter-influencers (e.g., peers who debunk or disbelieve the rumor), and fellow targets (e.g., peers who are potential adopters but have not yet either supported or opposed).

We begin with the basic relationship between the number of Alters exerting influence to adopt and the force operating on Ego. A simple and prominent approach to understanding how Ego adopting a behavior depends on the number of Alters previously adopting is the *threshold model* (Valente, 1996; Young, 2009). This approach assumes that our actions depend partly on others' actions, but ignores most of what we know about dyadic interpersonal influence to focus on one very specific form of interdependence: Before adopting a new behavior, E may wait for a critical number of others to adopt first. The critical number that would trigger Ego's adoption is called Ego's *threshold*. Threshold models implicitly assume that opponents and bystanders have no effect[1] and that all Alters are weighted identically, so the force of influence operating on Ego can be summarized with a single number, the count of those already adopting. Other determinants of behavior -- e.g., preferences, abilities, constraints -- are implicitly represented as heterogeneity of Egos' thresholds, as our focus is on any marginal contribution of conformity to peers. In other words, an Ego who generally favors adoption and has few constraints to adopt will have a relatively low threshold; i.e., may adopt even if a low number of peers have adopted. An Ego who feels negatively toward adoption or has strong obstacles for adopting will have a relatively high threshold; i.e., may not adopt even if a high number of peers have adopted.

The most basic model assumes that all Egos have the same threshold, *t*. At one extreme, only a single Alter adopting is sufficient to influence Ego to adopt, whereas any further Alters adopting make no difference to Ego. Some researchers refer to the case where *t*=1 as *simple contagion* (illustrated in Fig 1a), where the probability of Ego adopting shifts from 0.0 to 1.0 when exactly one peer adopts. A more nuanced version allows for inertia or resistance, such that the probability of adoption is less than 1.0 even if all Alters have adopted, so Ego's behavior is still a stepfunction but adoption operates probabilistically (Fig 1b). Another variant also allows for innovation or mutation, such that the probability of adoption for Ego is nonzero even if no Alters have adopted (Fig 1c). Lastly, we illustrate

[1] Some researchers assume that the threshold applies to the *proportion* of all actors in the population who have adopted (rather than the number) and this implies that the number of opponents and/or fellow targets may be included in the denominator. But it does so in an extremely restrictive way. Work in social psychology allows that the positive force by the number of adopters may have a functional form that is different from the negative force by the number of bystanders or opponents. These are all conflated in treating the threshold as a proportion.

so-called ***complex contagion*** (Centola and Macy 2007), **the case where Ego will not adopt until multiple peers first adopt (typically, t=2)**. The first Alter may have no effect, as multiple Alters are needed to influence Ego, but once that critical level of support is reached (say, 2 instead of 1), the probability of adoption jumps discontinuously from 0.0 to 1.0 and further Alters have no effect (Fig 1d). Here we highlight that simple and complex contagion are narrow special cases of threshold models, whereas threshold models are a narrow special case of influence models, assuming that influence on Ego operates as a stepfunction of the number of Alters adopting.

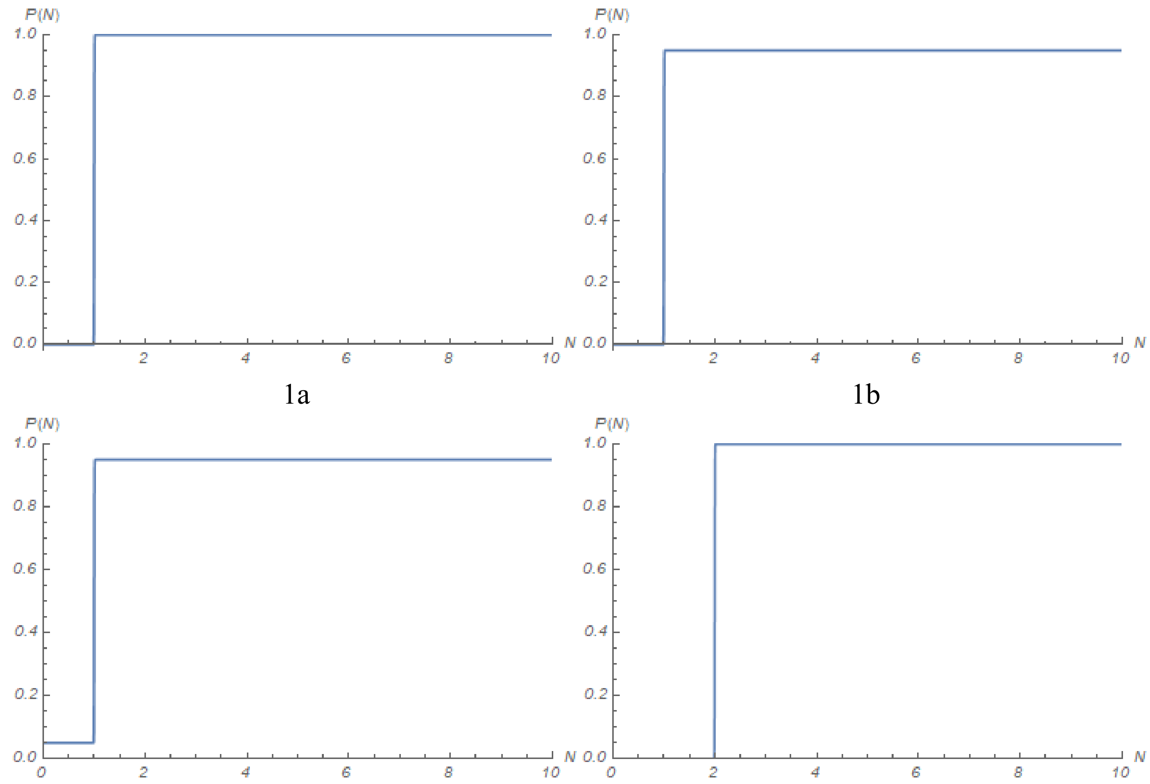

Figure 1. Threshold Models (probability of adoption is a step-function of number of adopters)

Beyond threshold models, more general influence functions may have a range of functional forms, including the stepfunction above, and they can be stochastic (probabilistic, incorporating randomness) as well as deterministic. Using the simplest and most basic threshold models (such as simple and complex contagion) has allowed researchers to focus on other issues such as the diffusion potentials for different network topologies. However, as we seek a general framework of influence we should build a more general influence function from social psychology. In an appendix at the end of this paper, we discuss and compare some alternative functional forms, but until then we will discuss broader and less technical issues.

Here we explore an example of a general framework for social influence from social psychology, *social impact theory* (Latane 1981) which assumes that the pressure on Alter to adopt a behavior is a power function of the number of others that have already adopted, with an exponent in (0,1); that is, each additional Alter adopting increases the pressure on Ego but at a diminishing marginal rate. Decades of experimental and observational research have examined this argument, suggesting that both direct behavior and related precursors to behavior (e.g., attention, memory, attitude) are influenced as a power function of the number of influencers. Importantly, researchers in this tradition have gone beyond the number of Alters (influencers) to also consider the number of Egos (targets). Indeed, research shows that influence on Ego is an *inverse* power function of the number of fellow targets, i.e., with an exponent in (-1,0). Realistically, the size of the population of not-yet-adopters increases resistance for Ego to adopt. Also, this approach can seamlessly integrate the case of counter-influence from *opposing* Alters, a question that is difficult to incorporate in a threshold model. In an empirical experiment (Latane and Wolf 1981), faced with a given number of positive and negative stimuli, research subjects developed evaluations that were power functions of the positive and negative inputs as predicted (i.e., each positive or negative input had a diminishing marginal impact on the impression, but in opposite directions).

Here we have illustrated SIT as a widely used general framework from social psychology that we regard as a starting point. The appeal of this framework is that it subsumes earlier restrictive work on threshold models as a special case, but also may accommodate a great variety of extensions: The framework integrates divergent influence by any number of adopters, opponents, and fellow susceptibles, which can be cumbersome in basic threshold or contagion models. (Whether a power function adequately captures conformity and peer influence processes is an empirical question, and can be addressed by further research applications.) The benefit of this flexibility is that it allows us to make principled use of information that may be available in a real-world influence situation.

B. Weighting Alters (sender strength, dyadic proximity, tie value, recency)

The basic threshold model (including the special cases of simple and complex contagion) and the alternative functional forms discussed in the appendix **all assume that a single force of influence operates from the *number* of Alters who have adopted. The Alters themselves are interchangeable. But in reality we know that some Alters are more important than others in influencing Ego.** Certain features of the influencing Alter may make their behavior more impactful on Ego. For example, some Alters may have social status (Paluck and Shepherd 2012), power (Blau 1964), or reputation (Fischer and Sciarini 2015) that allow them to more easily or effectively influence others in general. There may also be heterogeneity at the dyad level: Some Alters may be connected to Ego by ties of power, dominance, respect, or love, and any of these special bonds may cause Ego to weigh that Alter more strongly in reckoning incoming influence. On the other hand, some Alters may be despised or disrespected and have little or no effect on Ego, or even a paradoxical contrary effect (Liu and Srivastava 2015). So **if Alters vary in their individual-level strength or if ties vary in their strength, this suggests that we should replace the simple *number* of Alters above with a weighted combination of Alters, considering heterogeneity of Alters' strength and of the ties connecting them to Ego.** Taking this information about the importance of individuals or the strength of ties into account will allow for more nuanced influence strategies. **This could inform an influence strategy by focusing attention on the most potent allies and opponents, as well as allowing for strategies that involve strengthening or weakening ties rather than directly targeting behavior.** A framework cannot incorporate this key feature if it considers only the number of contacts adopting (or only the presence of one or two contacts adopting, as in simple or complex contagion).

Modeling the overall impact or force exerted by the weighted combination of Alters can make inconsistent empirical patterns intelligible. For example, empirical research that only considered the number of Alters has suggested that trivial behavior like retweeting seems to operate as 'simple contagion' (only one adopting peer is sufficient to induce Ego to adopt) whereas slightly less trivial behavior like changing a profile photo seems to operate as 'complex contagion' (requiring more than one adopting peer to support it). This pattern could be reframed simply and generally as: More effortful behavioral changes require more force to enact, whereas more trivial behavior requires less force to enact. Ceteris paribus, a greater force may be due to a larger number of adopting Alters or to a greater strength/importance of those Alters to Ego, or both. There is no doubt that one very important Alter (like a spouse or best friend) could induce Ego to change a profile photo, for example. Thus the apparent relationship between the difficulty of adopting a new behavior and complex vs. simple contagion may hold true ceteris paribus, but it is merely one instantiation of the more general statement that more force is needed to enact more effortful behavioral change. A great number of Alters is just one way to increase force, and the other ways are most crucial because the available levers to induce change in the real world may apply to the strength of Alters or the strength of Alter-Ego ties, rather than to the number of Alters.

We have discussed weighting Alters by their individual-level strength (in terms of the force of influence that they convey) and have also discussed weighting dyads by their strength (in terms of the amplification or attenuation of influence passing through the tie). Just as individual Alters or Alter-Ego ties may be more or less strong or emotionally charged, **pressures of influence may reflect the *proximity* of Alters at the time of influence. Proximity here may mean closeness in physical space, as we know actors are influenced most by nearby others even if they are strangers** (Latane et al 1995). Indeed, extensive empirical and simulation research has shown that physical proximity is an important factor leading to behavioral convergence. **Proximity here may also mean closeness in *sociometric* space, i.e., where Ego and Alter occupy nearby positions in a social network (connected directly or by a small number of intermediaries). Lastly, proximity may mean closeness in *sociodemographic* space, i.e., where Ego and Alter have similar identities of race, ethnicity, gender, or age**, and are close in this sense even if they have never met (Suh et al., 2017). Interpersonal influence attenuates as distance in sociodemographic space increases.

This framework also allows us to begin thinking about temporal patterns, which are largely ignored in prominent threshold and contagion models. There is little doubt that an influence attempt or stimulus generally attenuates in its force of influence as time passes; more recent events exert more pressure than more temporally distant events. Analogous to the other forms of distance, Latane and colleagues refer to this as *immediacy* and argue that the force of influence is an inverse power function of the time delay since a stimulus. Empirical research shows strong evidence of the recency effect on adoption of behavior. Using a large-scale dataset from a digital commercial site, Leskovec et al. (2007) found that 35%-40% of books and DVDs were sold within one day after the customers received a recommendation, which is largely based on product similarity and past purchase history. The purchase rates dropped significantly for both DVDs and books in one week after the last recommendation is made.

To recapitulate, a general framework for social influence should allow for senders of influence to be weighted by their importance, strength, or credibility. Dyadic ties should be weighted according to information available about physical, sociometric, and sociodemographic proximity, and the impact of influencing factors should be weighted as an inverse function of time passing since stimulus. A model that incorporates these nuances will more directly inform real-world practice. When an empirical setting offers some of this information but not all of it, the framework lets users know where to 'plug in' the available information, and also guides the users to know what other information would be useful to collect.

II. INFLUENCE AT THE MACRO-LEVEL

The general analytical framework of social influence suggests that the resultant force of influence is determined by the number of influencing Alters (proponents, opponents, and bystanders) weighted at the monadic level by the strength or importance of those influencing sources, weighted at the dyadic level by the physical, sociometric, or sociodemographic proximity of Alter to Ego, and the value of their direct tie. Lastly, we have considered the attenuation of the force of influence as time passes since the stimulus. These conceptual tools allow us to perform a diagnosis and predict where influence will be strongest or weakest at the micro-level. It also may inform interventions to strengthen or weaken influence. That said, **naive social influence interventions based on this micro-analysis may yield unproductive or even counterproductive results because they do not take into account higher-order structure of the population to be influenced. Network topology (structure of relations or contacts above the dyadic level) may have substantial impact, and taking topology into account can allow users to design more effective campaigns**.

Interdisciplinary research assuming the micro-level mechanisms above has derived implications for public opinion and population-level influence dynamics, using computational models of complex adaptive systems (Latané et al 1994; Macy et al 2003), as well as observational and experimental empirical studies (Latané and Bourgeois 1996). These research programs yielded three stylized facts, or emergent social phenomena at the macro-level: *conformity* toward the perceived majority, sociometric *segregation* into cliques or subgroups, and differentiation and *polarization* in the space of ideas and opinons. In the latter case, increasing correlations across a great variety of issues self-organize into attitudes, ideologies, and political party platforms that reduce a massively multidimensional state space into simple higher-order polarization. In this section we discuss these three pervasive population-level or group-level regularities, which prove relevant for planning effective social influence interventions.

A. Conformity / Convergence

The stronger the force of influence on Ego resulting from the micro-processes above, the more Ego will conform to the Alters' behavior. **Intuitively, these generic processes at the dyad level lead to conformity overall, or the reduction in diversity of behaviors and attitudes in the population, also called *convergence.***

The strength of conformity at the micro-level is well known, as is the tendency for populations of actors to adhere to shared norms and to develop common culture. Indeed, formal models of influence on networks notoriously predict an inexorable trend where population members influence one another to converge toward cultural uniformity. In this way, all differences should melt into a homogenized whole. However, in the real-world, we typically see persistent diversity of behavior in organizations, groups, and other local populations.

Many factors that are particularly prevalent in the era of digital social media can modulate the speed and extent to which behavioral convergence occurs. Understanding these factors in light of a general analytical framework of social influence may help to diagnose causes and processes of viral behavioral diffusion online, and identify potential solutions to prevent harmful consequences. In this section, we highlight three such factors.

First, designs and user interfaces of online social media platforms enable quick creation, expression and transmission of content at a negligible cost that are unimaginable in offline social interactions. Features such as Facebook Like and Twitter Retweet allow users to participate in social protests or contagions at a minimal cost (often by a few clicks of mouse), thus easily creating a critical mass at the early stage of diffusion (Lewis et al. 2014; Karpf 2010). The numbers of Likes or Retweets that a user can see on their own personal Facebook or Twitter timelines, or on news content pages provide a proxy of popularity, which eventually can increase sales (Ding et al. 2016), participation (Lewis et al. 2014) and awareness (Rishika et al. 2013). However, sometimes the aggregate numbers create a hyper reality of popularity, greatly exceeding the estimation of popularity by individuals based only on their proximate alters in physical or social spaces. Thousands of likes or retweets in a few hours may look or feel like overwhelming popularity, but this remains miniscule before a population of millions or billions of users. Both the volume and rate of those likes or retweets may create an illusion of overall popularity, even if they are coming from a relatively small but densely interconnected population of fanatics. In other words, **online social platforms create a new media in which individuals who are far apart in offline social or geographic space can readily access, appreciate and amplify each other's opinions and behaviors**.

Second, information technology-enabled viral features, such passive broadcasting (Aral and Walker 2011), can greatly increase behavioral convergence. Though the outcome of global conformity may be realized in models of person-to-person influence, market efforts or political campaigns often use broadcasting channels to quickly reach a large number of individuals at a reasonably low cost (e.g., the Superbowl attracts over 100 million viewers). According to the power function assumed in social impact theory, the broadcasting agent serving as the first influencer can generate a greater impact on the focal person than subsequent interactions, therefore creating an non-negligible force of influence. In an attempt of inferring structural features of viral diffusion, Goel et al. 2016 found strong evidence of broadcasting in producing global diffusion. Though rarely succeeding in their data set of a billion diffusion events on Twitter, products and ideas that did go viral often relied on a mix of broadcast and interpersonal contagion.

Third, the veracity and kind of the information spread through social networks is crucial for the success of social contagion. While the underlying mathematical model of social influence for both true and false news transmission can be the same, the parameters and shape of the influence function are different, and therefore, the outcome of social contagion, measured as the speed, depth, size and breadth, can differ greatly for diffusions of false and true information. In a recent comprehensive study of ~126,000 true, false or mixed stories on Twitter (Vosoughi, Roy and Aral 2018), falsehood (fake news) diffused significantly faster, farther, deeper and more broadly than truth, and the difference is most pronounced in the political domain. Several social and psychological factors contribute to the virality of false news stories. Individual users who shared (retweeted) false news stories tend to have fewer followees (friends) than those sharing true stories, indicating that on average they receive fewer tweets in total and therefore any single news story can have a greater psychological impact on adoption. Users sharing false rumors also tend to be newer users, less engaging on Twitter and not verified. False news stories, compared to true stories, are more likely to contain novel information, and provoke emotions such as surprise or disgust in the content of replies, instead of trust, joy, sadness and anticipation. Both the novelty measured by the content of false rumors, and the emotional response by users strongly suggest the necessity of a parametric specification of the social influence function by carefully measuring behavioral and contextual characteristics.

B. Clustering / Segregation

Given that interpersonal influence occurs locally in direct interaction (situated in networks), it is relevant that social actors are not distributed evenly or arbitrarily in networks; in fact, actors are distributed in clusters containing densely interconnected peers, and those clusters are only sparsely connected with one another. Individuals within these clusters see that their neighbors are each other's neighbors, their coworkers are each other's coworkers, and indeed a high probability that their friends are each other's friends. This clustering impacts almost every aspect of our lives, from marriage choices and job opportunities to disease exposure and behavior adoption.

The most intuitive cause of this clustering is a local phenomenon called *triad closure* -- where a triad is any set of three persons -- which implies that triads containing only two ties (A-B and B-C) will tend to form the third (A-C), thus creating a triangle of three persons who are all tied to one another. In a recent paper investigating such clustering in friendship networks, Goodreau et al (2009) caution that **such clustering can result from multiple intertwined processes, and that triad closure is connected with what they call *sociality* (or heterogeneous tendencies of people to be social, to form ties with others) and with *homophily* (the tendency to form ties with similar others; McPherson et al. 2001).** The problem is illustrated by the figure below, which demonstrates a set of relatively dense but more sparsely interconnected clusters. The cluster of purple friends could be produced by some combination of a force for friends of friends to become friends (triad closure), for purple people to befriend other purple people (homophily), and for purple people just to have more friends (sociality). The three processes work in concert to produce these dense clusters.

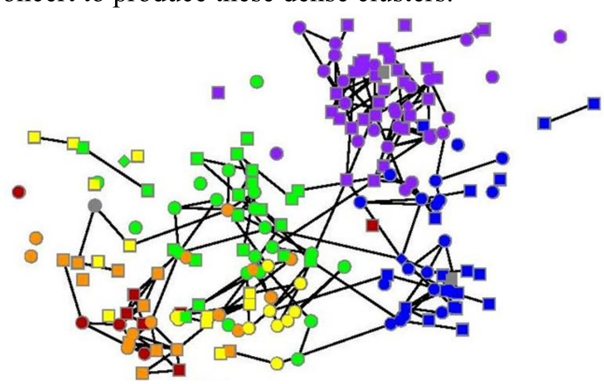

Figure 2. An illustrative friendship network

Of course, purple here is a metaphor for traits of individuals, and may represent race, gender, language, age cohort, or any other relatively fixed and observable trait that could be a basis for tie selection in a self-organizing network. Following the application in Goodreau et al. to adolescent friendship networks, we can take a look at a relatively diverse school population and see patterns exemplified in the photograph below (arbitrarily selected from public domain photos as an example of a high school cafeteria):


Figure 3. A student population in a high school cafeteria

In this typical photo, see that practically all observable interaction connects people of the same race and the same gender. On further examination, interaction partners share similar clothing, similar hairstyles, similar beverages, and similar body types. In all likelihood, they will share similar sociopolitical attitudes, similar career and family trajectories, and similar future health outcomes. This sociodemographic clustering is a lawlike regularity of human populations, resulting predictably from more primitive lawlike regularities of triad closure, homophily, and sociality.

It is worth noting here that **this lawlike regularity of sociometric clustering in networks -- and its deep correspondence to sociodemographic dimensions such as race, gender, age, education, and income --** is not a peculiarity of high schools but **is pervasive across all social contexts, scales, and life stages.**

C. Cultural Differentiation and Polarization

The localized nature of influence in physical and social space, as described in (B) above, leads to a self-limiting of global convergence in (A) above: **If each person continues to interact primarily with others nearby in space, the forces of conformity will be strongest locally, leading to the emergence of clusters of people sharing similar behavior, which we might call subcultures**, like the clusters of clothing styles illustrated in the photograph above. We have already noted that the **pervasive forces toward conformity typically do not lead to homogeneity, but seem to preserve diversity and even result in subcultures that conform internally but are at tension with one another.** More nuanced influence models do generate this pattern. For example, research using computer simulations has demonstrated how localized influence can produce clustered cultures on a 2 dimensional grid (Nowak, Szamrej, and Latane 1990) and in an evolving network (Macy et al 2003), a pattern that also appeared in early online experiments of influence on networks (Latané and Bourgeois 1996) and in observational studies of students living on a university campus (Bowen and Bourgeois 2001; Cullum and Harton 2007).

These points about localized influence in physical space leading to spatial clustering also apply to sociometric spaces (social networks) and sociodemographic spaces (similarity in race, gender, age, etc.). People tend to form social ties with similar others, therefore their behaviors are likely to converge toward those of socially similar Alters. If we draw each social attribute (e.g., gender, age, education, etc) as a dimension in Euclidean space, and locate people as dots whose social attributes correspond to coordinates in social space, we can find that people who are near each other in this space are more likely to share similar behaviors, beliefs, and opinions. Research (McPherson 2004) also found that behavior, beliefs and opinion are not entirely determined by the absolute value of social attributes (e.g., younger people are radical); rather, they can be reinforced and amplified by those similar in social interactions (e.g., a few radical youth share frequent interactions with other youth due to homophily and gradually cause the others to become radical.). These processes of culture/attitude clustering in sociodemographic space combine with the previously described processes of sociodemographic clustering in physical space to yield even more uniform and cohesive clusters of similar actors.

We can apply these insights to the time dimension as well. More recent social interactions convey a greater impact on individual behaviors than earlier interactions. Given repetitive stimuli from many distinct sources in the social network within a short time window, Ego's likelihood of adopting the behavior is greatly increased.[2] Information and communication technologies provide such a platform in which users have the capacity to update and communicate their opinions in real time. And because social networks are locally clustered (e.g., friend of friend is likely my friend too), a piece of information can quickly pass from one Alter to another, thus creating a reinforcing signal in a short period of time such that Ego has no time to experience alternative views. Large social media sites and digital commercial sites can also create echo chambers in which individual users rarely or never encounter views that they may disagree with.

One of the starkest consequences of structural bifurcation is polarization of a population into antagonistic groups that converge on distinct "platforms" of ideology. Not only are individuals densely clustered with others who share the same features and opinions, those opinions tend to become correlated with one another and develop a simple structure like attitudes, ideologies, or culture. This emergence of correlation among previously unrelated issues through network influence processes has been demonstrated in extensive formal models (Latané et al 1994; Macy et al. 2003; DellaPosta et al. 2015) and in observational (DellaPosta et al. 2015; Bishop 2008; Shi et al. 2017) and experimental (Latané and Bourgeois 1996) research.

In application to interaction in online spaces, these lawlike regularities are greatly amplified. For example, social recommender systems (e.g. when Facebook suggests friends) are designed to appeal to a population of users who have well-known *preferences* for homophily and triad closure. Because users want to connect with others who share interests with them and who share clubs or other affiliations with them, in catering to users' preferences the system offers search and partnering mechanisms to allow individuals to find a perfect match of their interests in a great number of social and cultural dimensions (e.g., an ideological leftist who loves cartoon books, watches animation movies, listens to hip hop music, and subscribes to the New York Times). Similarly, the system amplifies a natural tendency toward triad closure by actively recommending friends of Ego's friends as new contacts for Ego.

[2] This raises the possibility that in at least some empirical cases where complex contagion (reinforcement from a larger number of Alters) appears to be decisive, the mechanism may actually be that at the time of decision there has been more *recent* incoming influence. This will be more likely the case with a larger number of influencing alters if their incoming influence or pressure is distributed over time. Strictly speaking, the complex contagion argument could be tested by comparing the effect of p Alters against the effect of p incoming influence stimuli from a single Alter. For example, 5 incoming messages that 5 different friends have bought a book by a given author vs. 5 incoming messages that books by the author were bought by the same friend 5 times, with the same pattern in time. This would rule out the recency effect and isolate the effect of having more than one Alter adopt.

As a consequence of these software features, online arenas typically strengthen the already stark regularity of echo chambers of densely clustered others who are connected to the same people and share the same opinions. **As individuals who find tension and conflict in their personal offline network can move to online discussion rooms to express their views and make new friends with like-minded people, the availability of such spaces exacerbates segmentation and polarization of the overall population.** Deliberation among ideologically similar individuals can not only synchronize their views shared by the "platform" (echo chamber), but it also can lead to extremification and radicalization (Sunstein 1999).

Figure 4 gives another example of correlations of political and cultural dimensions on social media (Twitter; Shi et al. 2017), in which left-wing Twitter users tend to behave similarly to one another, and very different from right-wing Twitter users. Specifically, such partisan users follow (pay attention to) particular profiles of Twitter accounts that are typical for their political persuasion, even for Twitter topical domains that seem substantively apolitical. Each row represents a specific culture domain (such as *Food* or *TV Programs*) containing hundreds of widely followed Twitter accounts within that domain. Red color of the bars indicates a conservative leaning (i.e., ordinary users who follow the accounts in this culture domain are likely to also follow accounts of Republican Congresspeople), and blue color indicates a liberal leaning (followers in this domain tend to follow Democrat Congresspeople). *Internal polarization* is shown as the length of the bars, measuring how left-wing users tend to follow different accounts from the ones right-wing users follow, within that domain. In this way, we see that left-wing individuals tend to follow different domains and different accounts within each domain than do right-wing users, demonstrating politicization and polarization even on seemingly apolitical cultural dimensions. This segmentation and polarization process can also explain emergence of subcultures, including terrorist groups and other extreme organizations.

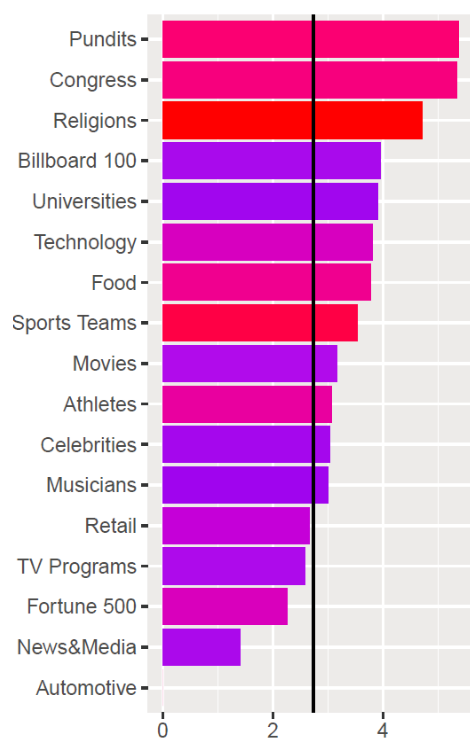

Figure 4. Political alignment (blue color indicates left leaning and red color indicates right leaning) and internal polarization (length of the bars) of cultural dimensions on Twitter based on their co-following patterns. Twitter users who follow right-wing politicians tend to follow accounts in the cultural domains of sports teams, food, and religions. Those who follow left-wing politicians tend to follow in domains of music, celebrities, and universities. Left and right-leaning users also follow different accounts within cultural domains (Billboard100, Universities, Sports teams, etc.).

As we have discussed, social influence between Ego and Alter can be greatly moderated by the social attributes associated with Alter and the relationship between Ego and Alter. **High status, legitimacy, power and credibility of the influencing Alter, and a strong and positive relationship between Ego and Alter can produce a stronger force of influence on Ego. However, the existence of subcultures with discrepant values and attitudes raise serious qualifications for the micro-level model. The status, legitimacy, or credibility that are recognized by the mainstream society may not be perceived in the same way by individuals in subcultures.** Imagine that Ego believes in a conspiracy theory, but Alter tries to debunk it using evidence or appeals to credible sources. If Alter is a professor or government official, this contextual information should strengthen the credibility of the signal from Alter, at least to mainstream audiences. However, if Ego is less educated and belongs to a marginal subculture that distrusts the government, academia, and the mainstream media, then Alter's negative feedback could carry little weight to Ego or even strengthen Ego's conviction of the conspiracy theory. By contrast, if Alter is similarly uneducated and belongs to a fringe militia group, most audiences may disbelieve Alter's perspective but Ego may find it ever more compelling.

III. SUMMARY SCHEMATIC DIAGRAM

We began this paper by identifying desirable features of **a framework for social influence that can guide policy and interventions**. Such a framework **should be general enough to apply to any problem of influence. It needs to be flexible in the sense that any available empirical information can be fed to the framework to yield useful insights; i.e., a great variety of different kinds of data can inform choices. It also needs to be flexible in the sense that it doesn't rely rigidly on a particular kind of information (e.g. social network data) to inform choices.**

The schematic diagram in Figure 5 depicts a web of interconnected constructs. Designing a real-world intervention would involve populating any of the various bins in this diagram with available data. Data for the inputs at the top may or may not be available in a particular application. Similarly, data for the moderating weights (such as on dyadic proximity or social ties among actors, or on the timing of potentially influential events such as messages or observations of Alter's behavior) may or may not be available. The framework will direct our attention to issues that matter from the contextual information we have available.

Formal modeling has proven that the **basic processes described in this figure robustly generate the three outcomes at the bottom -- convergence (reduction in diversity), segregation (clustering of behaviors in physical and network space), and polarization (simplification of the space of attitudes and behaviors by increasing correlations among items). But this is not to imply that the arrows only go one way in the empirical world. For example, segregation/clustering in networks clearly affects the pattern of exposure by Ego to different kinds of actors (e.g. in a neighborhood of**

adopters or a neighborhood of resistors). And the appearance of subcultures greatly affects the differential credibility of different kinds of messages and different kinds of message sources.

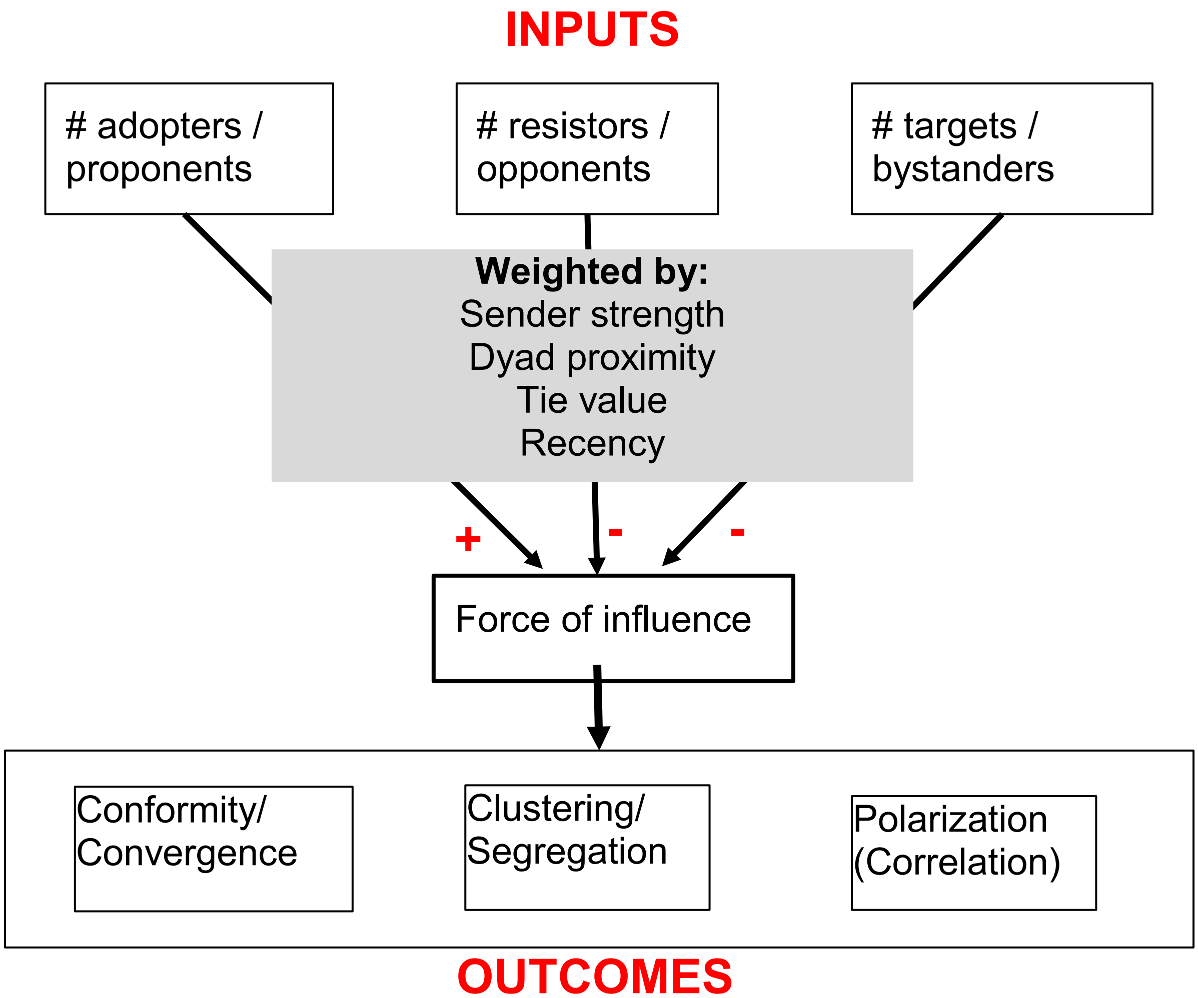

Figure 5. Diagram: Toward a general analytical framework of social influence.

The outcomes at the bottom can be described as nearly ubiquitous lawlike regularities. Being explicitly aware of them will allow them to inform influence strategies. For example, even without direct data on who is a member of which faction or subculture, knowledge that a community contains conflicting ideologies or value systems (as one might observe in online discussion threads) informs us that the community contains relatively segmented networks, and vice versa. And **the presence of segmented networks with corresponding subcultures could inform an intervention relying on diffusion of behavior among peers. A cursory examination of the subcultures could illuminate an influence strategy that is stratified and tailored to those subcultures**.

IV. APPLICATIONS TO PRACTICE: DESIGNING INFLUENCE CAMPAIGNS

Although previous sections may illuminate diagnosis and analysis of real-world situations, and thus inform development of intervention or counter-intervention strategies, this section we will develop some more concrete recommendations for practice. The first part of this paper examined the micro-level regularities that moderate diffusion of behavior on networks (number of influencers/opponents/ bystanders, strength of Alters and Ego-Alter ties, proximity and immediacy of Alter's contact with Ego). This section on implications for practice will be organized according to the macro-level patterns (conformity/convergence, segmentation/segregation, and differentiation/polarization) that are key to intervention strategies, but will be informed by the micro-level theories.

The schematic diagram in Figure 5 provides a visual framework to show how these various factors fit together. If rich contextual information is available in the field, it may be that all of these factors may be populated and elaborated with data. If information is unavailable (say, about dyad proximity, sender strength, or the number of opponents) it can be left out and the framework will operate without it. **If an agent has access to levers to change one or more inputs this framework gives guidance for how those inputs will affect the outcome. It also directs the agent's attention to many less obvious levers that s/he may not even have realized were available. For example, the framework may reveal tactics such as undermining the value of the dyadic ties between opponents and the target, augment the monadic credibility of proponents, or rearrange space to give the targets more ready and visible access to existing adopters. It also recommends gathering information to further populate and inform this framework.**

A. Conformity

Whereas leaders and managers may intuitively seek mass persuasion through public announcements or advertisements, the general framework can provide more pointed recommendations. Understanding the functional form of social influence (or various functional forms that may be applicable in different contexts) can help to design/adjust individual, organizational, and institutional forces that shape the pathway of social influence.

First, consider how Ego's behavior depends on the number of Alters who have adopted, as assumed in threshold models. It seems circular and unhelpful to advise that the secret to getting more people to adopt an innovation is to first get more people to adopt the innovation. However, we can go beyond that: **A key way to influence Ego is to manipulate the *appearance* or *expectation* of Alters' adoptions. Various mechanisms may be used to amplify arrivals and adoptions or dampen departures and rejections,** ultimately giving Ego the impression that a large and growing portion of the population is adopting or is about to adopt. **This can operate as an impression of the magnitude ('everyone is doing it') or the trajectory ('it's taking off, and will be the next big thing')** and these can be notably more effective than appeals to reason or emotions ('it's the right thing to do').

Even when thresholds cannot be directly measured (as is almost always the case), **when we diagram the threshold distribution based on qualitative observation or expert informants (Figure 6), we can then use the diagram to customize an intervention strategy that fosters broad diffusion. That intervention strategy may include which category of individual to target (e.g. innovators, early adopters, late majority, laggards, etc.) and in what order, or even customize particular influence tactics for these different categories.** The framework employing the threshold distribution informs the intervention strategy; for example an agent may target the Early Adopters below, aiming to leverage their momentum to convert the Early Majority (realizing that the Innovators will adopt without being targeted by an intervention). Sufficient momentum could carry through the rest of the population.

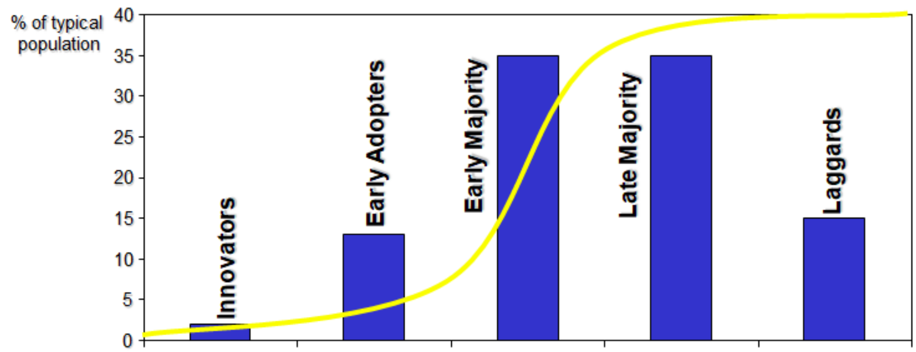

Figure 6. The “S” shape of the diffusion of innovation.

Depending on the shape of the distribution, different strategies may be supported. For a particularly costly or risky innovation, many members may be in the Late Majority or Laggards categories, and interventions may need to approach the distribution beginning with the Innovators and using each category as weight to convert the next. If the distribution is bimodal, with concentrations of actors supporting and opposing the innovation, an intuitive approach would be focus most intervention on the right side, as the left side is likely to adopt anyway and serve as allies, Of course, such a distribution implies that there may be subgroups and subcultures, and knowledge of this would inform structural interventions.

The above assumes a general/global mechanism of influence, where Ego is influenced purely by the number of Alters adopting in the overall population. The framework can give more specific recommendations if there is a more nuanced model of interpersonal influence: **If influence operates locally in *physical* space, an agent can highlight adoption choices of Ego’s nearby peers**. For example, casino jukeboxes only draw the room’s attention to themselves when they yield a large payoff to gamblers, giving everyone in a large room the impression that winning is prevalent and likely. **If influence operates locally in *sociometric* space (e.g. Ego is influenced principally by behavior of friends, family, or colleagues) an influence agent can increase visibility or salience of peers adopting the desired behavior,** so Ego disproportionately notices peers who have adopted. For example, political campaigns will highlight the voting behavior of peers on social media, which can greatly impact Ego’s self-expression, information seeking, and real-world voting behavior (Bond et al 2012). **If influence operates locally in *sociodemographic* space (e.g. Ego is influenced principally by behavior of Alters who are similar to Ego) an agent can highlight adoption choices of others similar to Ego**. For example, marketers will include actors in advertisements who look like the target demographic for the product.

Social media and social networking sites provide ample opportunities for implementing interventions of influence or counter-influence. **We give a few recommendation of practices specifically designed for online platforms. First, to a great extent, Ego’s perception, and eventual adoption of a news item on Facebook (or any other social media site) depends on highly distilled information provided by the platforms such as trending topics, numbers of likes or shares (Chang et al. 2017), or content**

recommended by algorithms. As this aggregate information can be easily manipulated (e.g. by an army of shills or trolls who can induce a cascading momentum for a topic becoming trending), one potential way to counter the harmful trolling behavior is to offer additional breakdown of the dynamics leading to virality. For example, decompose the aggregate viewing counts by demographics; some websites such as YouTube have shown the percentage of viewers from different demographic groups; or show who in the friend list has approved (like and love) or disapproved (angry) the message, a practice already used by social media such as Facebook. Thus, Ego can make more reliable inference of the influencing sources (e.g., are they proximate to Ego in social space? Are they friends of Ego, and what are their reactions?). **Second, the extent to which an opinion (e.g., a piece of news story) is expressed and transmitted is also contingent on the cost associated with the action. Therefore, an increase in the cost of action can discourage careless and hasty behavior and encourage thoughtful, discreet and engaging interactions and conversations.** Cost is broadly defined as the time, attention, and effort that one has to spend on the transmission of a piece of content. Studies show that political retweeting on Twitter (i.e., a click of the retweet button) is more frequent and more prone to direct towards partisan materials than mentioning (i.e., direct message toward others), as mentioning requires an extra effort in composing the message. **Third, social media platforms can also link news stories to external fact check websites, and warn users of potential characteristics and harm of false news content.**

Research on network influence models suggests that we should often target highly connected individuals, individuals who occupy central or brokering positions in networks (see Valente 2012 for a comprehensive review of various network-based targeting techniques). For diffusion of effortful/nontrivial innovations or practices, Ego may require a greater number of Alters to adopt first (i.e., may have a threshold greater than one). **Network threshold models then suggest that an influence agent target subsets of Egos who are also interconnected enough to reinforce each other's adoption** (i.e. exceed their thresholds locally). In this way, activity can be supported/reinforced in local clusters and then diffused to others.

It is important to note that all tactics building on the tendency to conform to group norms must necessarily operate on what Ego *perceives* to be the group norm, which may not be an accurate reporting of others' behavior. If systematic biases operate in how individuals perceive others either through psychological processes (Zou et al. 2009), social biases of communication (Kitts 2003), or difference in expectation of role relationship (parents and children), **this could have strong impacts on collective behavior through the mechanism of conformity.** For instance, in a politically charged environment, **expression of private belief on some social issues may cause emotional conflicts with others, therefore, people would be restrained from disclosing their true beliefs or behaviors *publicly*,** even though a great proportion of those nearby may hold the same beliefs or behaviors *privately*. **This in turn can create an ideal social environment for false norms or false information to spread, as individuals become reluctant to challenge views labelled as ideologically different or opposite, regardless of how absurd they might be.** The general tactic is to build transparent and trustworthy channels of communication, where individuals can have the power of disclosing their private beliefs and have the opportunities to deliberation. **The proper design of online privacy and grouping policies can be crucial in that it determines how much private beliefs a user is willing to share and to whom.** If there is no such function, users on social media sites (e.g., Facebook) would not be willing to share private opinions or the contents that would reveal their private beliefs on their timelines, as everyone in the friend list can see them, and eventually users would abandon the service entirely. On the opposite end, if users disclose their views only through private channels (e.g., Messenger), other users would have trouble to perceive the true norms shared in the community.

B. Segregation

An intervention in practice may aim to have first-order impacts (changing the behavior of target actors who are direct recipients of the intervention) but also higher-order impacts (changing the behavior of actors in the target population who do not receive the intervention directly) and ultimately establish new norms for behavior that take on a life beyond the intervention. They aim to do so by leveraging existing networks within the target population, and a naive application is based on an assumption that those networks are well-mixed, randomly distributed in sociodemographic space, and either uniformly distributed or with a unitary core-periphery structure.

We have seen that real-world social networks -- offline and especially online -- are notoriously *segregated* in sociodemographic space. Imagine an intervention in the student population pictures in the high school cafeteria above. **Even if the intervention is successful in changing the behavior of some number of target individuals, the segmented structure of the population may constrain the outcome.** For example, say that the intervened behavior requires reinforcement by peer adopters, such that intervening on 10 individuals in a well-mixed population would provide sufficient peer reinforcement to establish the innovation and spread it to peers. Now imagine that **the same intervention were to be distributed across 10 relatively disconnected network clusters. As a consequence none of the early adopters has any peers who also adopt, and no individuals are exposed to more than one adopter, so the innovation may fail to spread and may not even survive.** Next imagine that **the intervention is intentionally targeted at a single cluster. In this case, there may be sufficient reinforcement to establish the innovation permanently in a stable subgroup of innovators, but little means for it to spread through the population**. One table of the cafeteria adopts the new behavior, but all other tables ignore or scoff at the innovation. Agents interested in changing behavior in the population may take this into account in designing an intervention on a networked population. **Innovations that are costly, risky, or otherwise difficult to adopt may require a more substantial peer reinforcement, with sufficient attention by the intervention into cohesive subgroups; that said, to keep the intervention from petering out in a small isolated subset of the population, multiple subgroups in relatively distant areas of sociodemographic space may be involved.** Agents may also target subgroups that are relatively connected to other subgroups. For innovations that are simple, safe, and easy to adopt, interventions may be intentionally scattered across subgroups to maximize breadth of impact in the population.

Another approach to intervention is to manipulate the network itself. Institutional forces or organizational practices may intentionally foster cross-cutting ties that bring people together from different social and demographic strata. **Recruiting new members that cross demographic fault-lines (e.g. share features in common with multiple cohesive subgroups) or establishing practices that foster cross-cutting ties among existing members (such as mixer events) will both have the effect of reducing structural bifurcation into disconnected subgroups.**

C. Cultural Differentiation and Polarization

The fact that social proximity leads to behavioral convergence has great implications in monitoring and designing social intervention that controls or eliminate the spread and outbreak of malicious behaviors and opinions. In fact, interventions that aim only at the population level, such as a mass persuasion attempt that fails to take into account the segmentation of the population into distinct subcultures, may yield counterintuitive and even counterproductive results. **Several strategies can potentially bridge the divide between polarized groups.**

Intervention campaigns can create or make use of overarching themes or issues that crosscut different groups as a means to reduce tension between ideologically opposing groups. Such themes or issues include creating/finding common enemies, sports (e.g. Olympics) and entertainment, natural disasters, science projects, etc. Individuals coming from distinct backgrounds or holding contrasting opinions tend to coexist peacefully when facing a common goal. Social integration of individuals in a group is a prior condition of further social blending and cultural integration.

If an agent aims to disrupt an opposing coalition, rather than directly attacking, a more subtle intervention (bypassing reactance motives) would be to simply raise topics or make issues salient that are divisive to that coalition. And we can predict what will be divisive by the patterns of intercorrelation among issues. If a **large number of issues are or have recently become correlated with one another,** such that we can know a person's opinion on issue *j* from knowing her opinion on issue *i*, then **these issues are likely to load on a primary fissure in the ideological space of the group. Raising these issues is likely to drive sociometric segregation and further polarization of ideological space, ultimately fostering a destructive force for the group.**

V. EPILOGUE

Apart from all other issues, we face a problem that is endemic to the era of computational social science: Many contemporary usages of the social network concept envision networks as temporally aggregated time series of interaction events (Kitts and Quintane forthcoming). **Research that involves behavioral diffusion on social networks typically treats those networks as binary (actors are tied or not)** or at most weighted (more or less strong) and continuously active (flow is possible at any time), **even when the ties themselves represent aggregations of communication or exchange acts** observed in continuous time. **For example, a 'network' may be defined as a graph where a given edge is present if the parties shared at least three phone calls in a given month.**

Recent work on diffusion has revealed that such social networks can severely misrepresent interpersonal diffusion processes, because temporal aggregation obscures the sequence and timing of contacts among actors (Moody 2002). For example, if A has contact with B and later B has contact with C, then a temporal aggregation of these events into a network would yield the conclusion that contagion is plausible from C to B, although that is not possible in the underlying sequence of interaction events. **Also, whether contacts occur in sequence or occur concurrently can have strong implications for exposure graphs and resulting dynamics of contagion (Morris et al. 2007).**

Cutting edge work focuses analytically on the temporal and structural dynamics of interaction events, such as interdependencies of event streams of e-mails, IMs, phone calls, or face-to-face meetings. Analyzing the dynamics of interaction (instead of coarsely aggregated 'social ties') both allows us to stay closer to the social processes we are theoretically modeling, and also avoids misrepresenting the substrate for diffusion by obscuring the sequence and dynamics of social contacts. Incorporating these insights into behavioral diffusion research will require fine-grained data, innovative statistical models, and new breeds of theory.

VI. APPENDIX: Continuous influence functions as alternatives to the threshold model

Earlier in this paper we described a class of stepfunction influence models called threshold models, and special cases of those models called simple and complex contagion. An straightforward alternative to a strict stepfunction is a sigmoid function (Van de Rijt, Siegel, and Macy 2009), where the probability of adoption approaches 0.0 when few or no Alters have adopted and approaches 1.0 when many Alters have adopted, but the probability increases smoothly rather than discontinuously in one step (Fig 7a); in the region of ambivalence the probability of adoption is increasingly random.This highly flexible function can approximate a variety of special cases, including non-influence, linear influence, threshold, accelerating or decelerating impact, and forms of stochasticity (exploration, reactance). Adjusting a parameter in the function used to produce Fig 2a allows it to cover the full range from a stepfunction (Fig 1) to a linear function where each Alter has a constant impact on Ego (Fig 7b). Click here for an animation illustrating the full range of shapes.

Much work (including social impact theory) has considered a power function with exponent less than 1.0, widely investigated and supported in experimental research. This function embodies the assumption that the first Alter has the greatest impact and each additional Alter has diminishing marginal impact on Ego (Fig 7c). A parameter of the power function allows it to approximate a stepfunction (Fig 1a) or a linear function (Fig 7b). Click here for an animation illustrating the full range.

We are not trying to adjudicate between continuous and threshold models of social influence, which requires rigorous experimental controls on the type of diffusion content and number of activated alters. Several experimental and observational studies (Centola 2010; Mønsted et al. 2017; Romero, Meeder and Kleinberg 2011) attempting to distinguish between simplex and complex contagions identify compatible patterns of social influence for both continuous and threshold models: 1) influence from multiple neighbors in the social networks greatly increases the likelihood of adoption; 2) the effect of additional number of activated alters is smaller than previous alters. Our goal of presenting multiple theoretical models of influence is to offer a broader analytical framework in which different social elements, many of which are examined in this paper, can be incorporated to reach a more comprehensive understanding of social influence and behavioral diffusion in general.

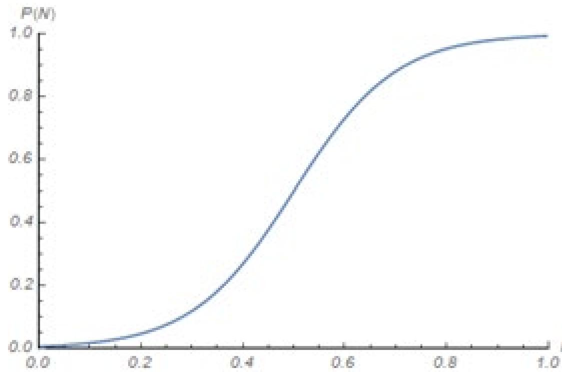

7a

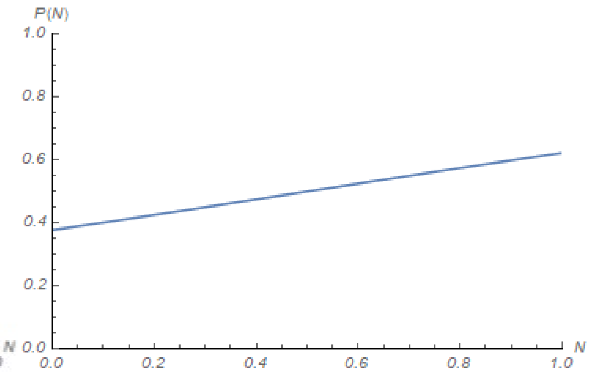

7b

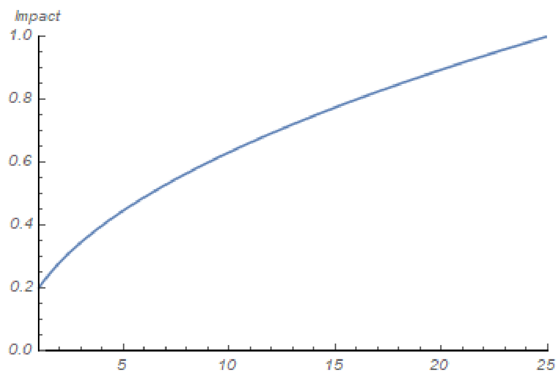

7c

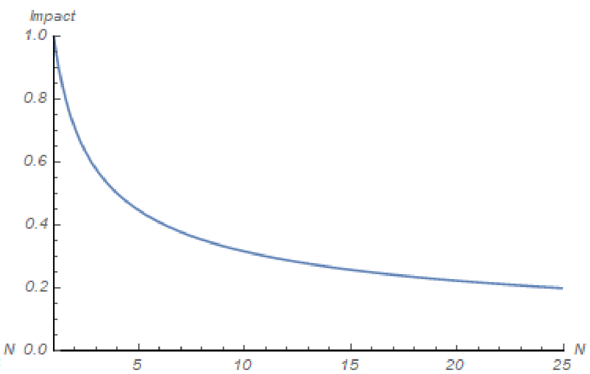

7d

Figure 7. Sigmoid and Power Influence Functions

VII. BIBLIOGRAPHY

Aral,Sinan and Dylan Walker. 2011. "Creating Social Contagion Through Viral Product Design: A Randomized Trial of Peer Influence in Networks." *MANAGEMENT SCIENCE* Vol. 57, No. 9, pp. 1623–1639.

Bishop, Bill. 2008. *The Big Sort: Why the Clustering of Like-Minded America Is Tearing Us Apart*. Boston: Houghton Mifflin.

Blau, Peter. 1964. *Exchange and Power in Social Life*. 1st edition. 1. New York: John Wiley and Sons, Inc.

Bond,Robert, Christopher J. Fariss, Jason J. Jones, Adam D. I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. 2012. "A 61-million-person experiment in social influence and political mobilization." *Nature* 489(7415): . doi:10.1038/nature11421.

Bowen, Anne M. and Martin J. Bourgeois. 2001. "Attitudes toward lesbian, gay, and bisexual college students: The contribution of pluralistic ignorance, dynamic social impact, and contact theories" *Journal of American College Health*, 50:2, 91-96, DOI: 10.1080/07448480109596012.

Centola, D. 2010. "The Spread of Behavior in an Online Social Network Experiment." *Science*, 329(5996), 1194-1197.

Centola, D. & Macy, M. 2007. "Complex contagions and the weakness of long ties." *Am. J. Sociol.* 113, 702–734.

Chang,Jung-Hua, Yu-Qian Zhu, Shan-Huei Wang, and Yi-Jung Li. 2017. "Would you change your mind? An empirical study of social impact theory on Facebook" *Telematics and Informatics* Volume 35, Issue 1, April 2018, Pages 282-292.

DellaPosta, Daniel, Yongren Shi and Michael W. Macy. 2015. "Why Do Liberals Drink Lattes?" *American Journal of Sociology* 120(5):1473-1511.

Ding,C., H.K. Cheng, Y. Duan, and Y. Jin. 2017. The power of the "like" button: the impact of social media on box office. *Decision Support Systems* 94, 77-84.

Emerson, R. (1964). Power-Dependence Relations: Two Experiments. *Sociometry, 27*(3), 282-298.

Feld, Scott L. 1981. "The Focused Organization of Social Ties." *American Journal of Sociology* 86(5):1015–35.

Giddens, Anthony. 1991. *Modernity and Self-Identity: Self and Society in the Late Modern Age*. Stanford, Calif.: Stanford University Press.

Goel, Sharad, Ashton Anderson, Jake Hofman, and Duncan J. Watts. 2016. “The Structural Virality of Online Diffusion" MANAGEMENT SCIENCE Vol. 62, No. 1, January 2016, pp. 180–196.

Goodreau, S.M., Kitts, J.A. & Morris, M. 2009. *Demography* 46: 103.

Hunter, James Davison. 1991. *Culture Wars: The Struggle to Define America*. New York: Basic Books.

Karpf, David. 2010. “Online Political Mobilization from the Advocacy Group’s Perspective: Looking Beyond Clicktivism” Policy & Internet Vol. 2: Iss. 4, Article 2.

Kitts, James A. 2003. “Egocentric Bias or Information Management? Selective Disclosure and the Social Roots of Norm Misperception.” *Social Psychology Quarterly* 66:222–37.

Kitts, James A. and Eric Quintane. Forthcoming. “Rethinking Networks in the Era of Computational Social Science.” *Oxford Handbook of Social Networks*.

Latané, B. 1981. “The psychology of social impact.” *American Psychologist*, 36, 343-356.

Latané, B., A. Nowak, and J. H. Liu. 1994. “Measuring emergent social phenomena: Dynamism, and clustering as order parameters of social systems.” *Behavioral Science*. 39: 1-24.

Latané, B. 1996. “Dynamic social impact: The creation of culture by communication.” *Journal of Communication*, 46, 13-25.

Latané, Bibb, Liu, James H., Nowak, Andrzej, Bonevento, Michael, and Long Zheng. 1995. “Distance Matters: Physical Space and Social Impact.” *Personality and Social Psychology Bulletin.* 21(8): 795 - 805.

Latané, B., & Bourgeois, M. J. 1996. “Experimental evidence for dynamic social impact: The emergence of subcultures in electronic groups.” *Journal of Communication*, 46, 35-47.

Latané, B., & Wolf, S. 1981. “The social impact of majorities and minorities.” *Psychological Review, 88*(5), 438-453.

Leskovec, Jure, Lada A. Adamic, and Bernardo A. Huberman. 2007. “The dynamics of viral marketing.” *ACM Trans. Web* 1, 1, Article 5. DOI=http://dx.doi.org/10.1145/1232722.1232727

Lewis, Kevin, Kurt Gray, and Jens Meierhenrich. 2014. “The Structure of Online Activism.” *Sociological Science* 1: 1-9. DOI: 10.15195/v1.a1

Liu, Christopher C. and Sameer B. Srivastava. 2015. "Pulling Closer and Moving Apart: Interaction, Identity, and Influence in the U.S. Senate, 1973 to 2009." *American Sociological Review*. 80(1): 192-217.

Macy,Michael W., James A. Kitts, and Andreas Flache. 2003. “Polarization in Dynamic Networks: A Hopfield Model of Emergent Structure” *Dynamic Social Network Modeling and Analysis* (2003) by the

National Academy of Sciences, pp. 162-173.

McPherson, Miller. 2004. “A Blau Space Primer: Prolegomenon to an Ecology of Affiliation.” *Industrial and Corporate Change* 13:263–80.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. 2001. “Birds of a Feather: Homophily in Social Networks.” *Annual Review of Sociology* 27:415–44.

Miller, Wesley E., Jr. 1985. “The New Christian Right and Fundamentalist Discontent: The Politics of Lifestyle Concern Hypothesis Revisited.” *Sociological Focus* 18:325–36.

Mønsted, B., Sapieżyński, P., Ferrara, E., & Lehmann, S. 2017. “Evidence of complex contagion of information in social media: An experiment using Twitter bots.” *PloS one*, 12(9), e0184148.

Moody, James. 2002. “The Importance of Relationship Timing for Diffusion.” *Social Forces*. 81, 25-56.

Morris, M, Goodreau, S. and J Moody. 2007. "Sexual networks, concurrency, and STD/HIV." *Sexually Transmitted Diseases.* 4, 109-125.

New York Times. ‘Duck Dynasty’ vs. ‘Modern Family’: 50 Maps of the U.S. Cultural Divide. https://www.nytimes.com/interactive/2016/12/26/upshot/duck-dynasty-vs-modern-family-television-maps.html

Nowak, A., Szamrej, J., & Latané, B. 1990. “From private attitude to public opinion: A dynamic theory of social impact.” *Psychological Review*, 3, 362-376.

Paluck, E. L., & Shepherd, H. (2012). The salience of social referents: A field experiment on collective norms and harassment behavior in a school social network. *Journal of Personality and Social Psychology, 103*(6), 899-915.

Rishika,R., A. Kumar, R. Janakiraman, R. Bezawada, “The effect of customers' social media participation on customer visit frequency and profitability: an empirical investigation,” *Inf. Syst. Res*. 24 (1) (2013) 108–127.

Romero, D. M., Meeder, B., & Kleinberg, J. 2011. “Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter”.WWW '11 Proceedings of the 20th international conference on World wide web 695-704

Shi, Feng, Yongren Shi, Fedor Dokshin, Michael W. Macy and James Evans. 2017. “Millions of Online Book Co-purchases Reveal Partisan Differences in the Consumption of Science.” *Nature Human Behaviour* doi:10.1038/s41562-017- 0079.

Shi, Yongren and Michael W. Macy. 2016. “Measuring Structural Similarity in Large Online Networks” *Social Science Research* 59:97-106.

Shi, Yongren, Kai Mast, Ingmar Weber, Agrippa Kellum and Michael Macy. 2017. “Cultural Fault Lines and Political Polarization.” in *Proceedings of the ACM Conference on Web Science*.

Suh, Chan, Yongren Shi and Matthew Brashears. 2017. “Negligible Connections? The Role of Familiar Others in the Diffusion of Smoking among Adolescents.” *Social Forces.* 96(1):423–448.

Sunstein, Cass R. 1999. The Law of Group Polarization. John M. Olin Program in Law and Economics Working Paper No. 91.

Valente, Thomas W. 1996. “Social Network Thresholds in the Diffusion of Innovations” *Social Networks* 18:69-89.

Valente, Thomas W. 2012. “Network Interventions” Science 337:49-53.

Van de Rijt, Arnout, Siegel, David, and Michael Macy. 2009. “Neighborhood Chance and Neighborhood Change: A Comment on Bruch and Mare.” *American Journal of Sociology*. 114(4): 1166-72, 1176-8.

Vosoughi,Soroush, Deb Roy and Sinan Aral. 2018. “The Spread of True and False News Online.” *Science* 359, 1146–1151.

Wejnert, Barbara. 2002. “Integrating Models of Diffusion of Innovations: A conceptual Framework” Annual Review of Sociology 28:297-326.

Young, H. Peyton. 2009. “Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning” *American Economic Review*, 99:5, 1899–1924.

Zou, X., Tam, K.-P., Morris, M. W., Lee, S.-l., Lau, I. Y.-M., & Chiu, C.-y. (2009). Culture as common sense: Perceived consensus versus personal beliefs as mechanisms of cultural influence. *Journal of Personality and Social Psychology*, 97(4), 579-597.