How Behaviour Spreads: The Science of Complex Contagions by Damon Centola - summary

January 2022

This document is a summary of Damon Centola's book *How Behaviour Spreads: The Science of Complex Contagions*.

Chapter 1: Introduction

A common intuition about how behaviour spreads is based on the notion of 'virality'. In this conception, what matters are the 'weak ties' that connect huge numbers of people - who themselves typically have meaningful relationships with only Dunbar-sized groups - into 'small-world' networks. This is embodied by the idea of 'six degrees of separation', which has become a kind of received wisdom for how *everything* spreads. For spreading actual viruses, or simple information (memes, job openings) to remote parts of a network, leveraging weak ties can indeed be an optimal strategy (Figure 1), since transmission can be effectively instantaneous via simple mechanisms (infection, social media shares/likes), and is often involuntary.

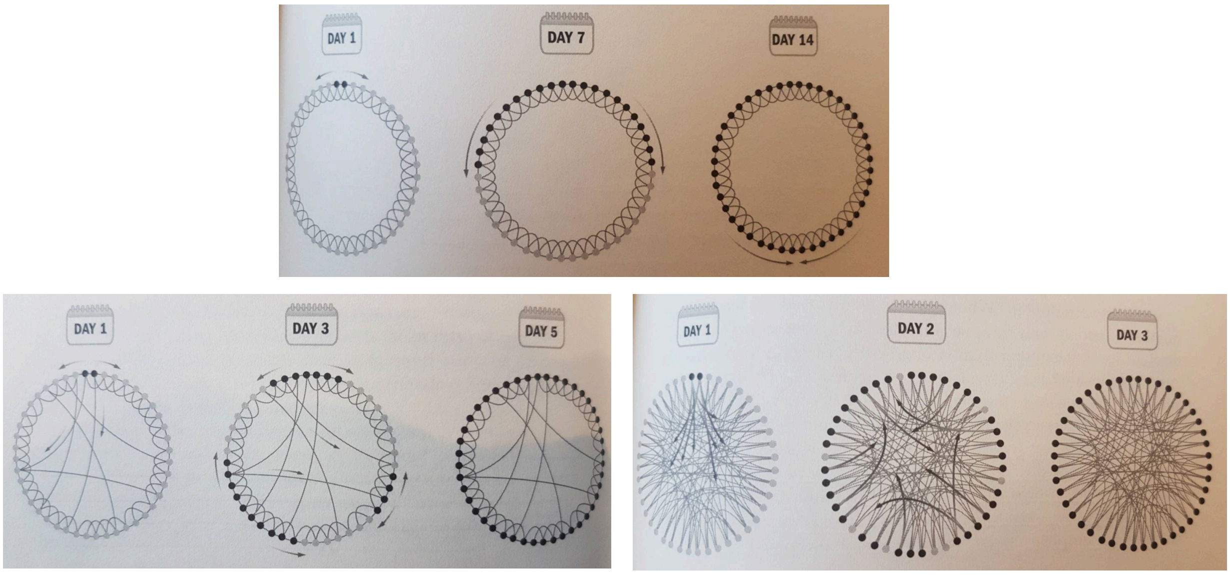

Figure 1: a simple contagion spreads slowly through a 'large-world' network (top) where people are only connected to their immediate neighbours, but much more rapidly when there are even a few ties to distant parts on the network (bottom left). In a highly connected 'small-world' network (bottom right), complete exposure is almost instantaneous. This illustrates the dynamics of a 'viral' contagion.

Given the prevalence of this intuitive model, anyone wishing to diffuse a new behaviour or technological innovation may assume that they should focus on weak ties, as opposed to what appear to be the redundant, dense, overlapping connections that are characteristic of e.g. friendship groups. However, there is plenty of empirical evidence that desirable innovations often fail to diffuse as hoped, and Centola demonstrates (through computational models, experiments, and case studies) that for behaviours that meet some minimal threshold for complexity - which can be as low as giving some contact details to an online health forum sign-up form - the viral model of diffusion is misleading, and strategies based on it can be actively counterproductive.

For these ‘complex’ behaviours, it is precisely the mutually-reinforcing ‘strong ties’ found in clustered social groups that lead to adoption and maintenance (Figures 2 and 3). The contrast is encapsulated by the spread of a virus versus the spread of attitudes to vaccination.

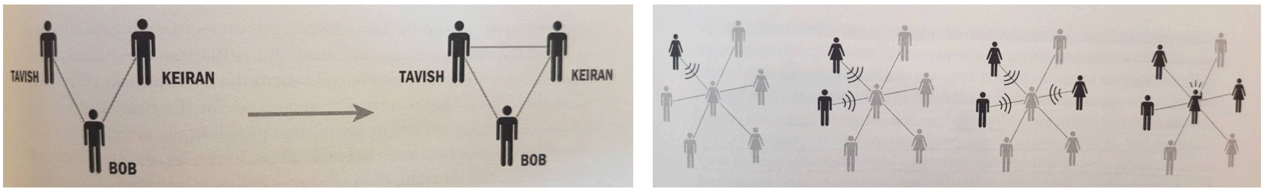

Figure 2: a ‘triadic closure’ (left) as the archetype of the kind of mutually-reinforcing social structure necessary to spread a ‘complex’ behavioural contagion, and an example of adoption of a complex behaviour following reinforcing exposures (right).

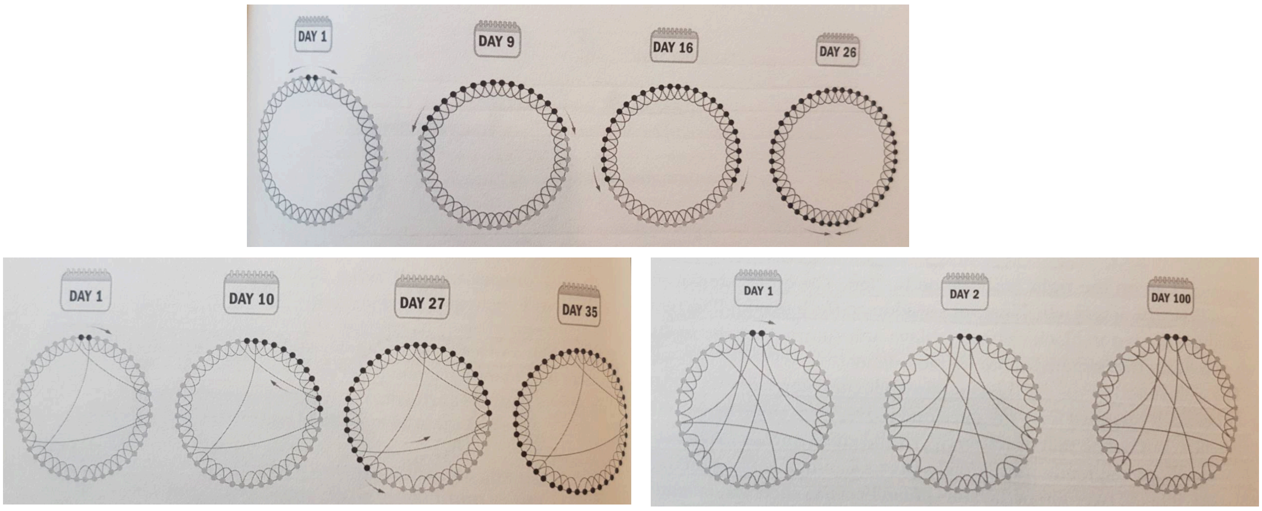

Figure 3: although it moves slowly compared to the simple contagion in Figure 1, a complex contagion successfully spreads through a ‘large-world’ network (top) faster than it does when there are even a few ties to distant parts on the network (bottom left). Adding only a few more weak ties (bottom right) results in the complex contagion completely failing to diffuse (note that this network is substantially less connected than a small world). This illustrates the dynamics of a complex contagion.

In terms of framing and scope, Centola assumes that whatever is intended for diffusion cannot be readily improved, instead focusing on the social network factors that ‘can transform a failed technology into a successful innovation’. Furthermore, he looks primarily at the potential of *online* technologies to influence outcomes.

A final important framing (that is implicit in the network approach, but which should still be emphasised) is that the book investigates collective rather than individual factors in achieving change. Although tailoring the messages people are exposed to can be effective, influencing social norms and structures can be far more so. In the same way that schooling behaviour in fish could never be predicted from studying a single fish in isolation, many behavioural diffusion phenomena only make sense when complex interdependent social relations are understood.

Specific applications discussed include how a novel technology can be incubated in populations in which an alternative is already established (Chapter 6), the use of *wide bridges* (exemplifying a distinction between the *structural* and *relational* importance of ties) for encouraging diffusion across organisational boundaries (Chapter 7), how to design networks to increase diffusion amongst strangers (Chapter 8), and how to control the kinds of behaviour that spread using incentives for social support or comparison (Chapter 9).

Centola concludes by arguing that his basic approach of investigating how 'imperceptible changes in the structure of social relationships produce significant differences in collective outcomes' yields results that give grounds for optimism; whilst people frequently resist even beneficial behavioural changes, they are not incorrigible. Diffusion can succeed if the right strategies are followed.

Part I - Theory

Chapter 2: Understanding Diffusion

The essential problem that is addressed in the book is that whilst some things diffuse despite our best efforts to stop them, others that we wish to see spread as widely as possible fail to do so. Resistance to an innovation can arise because it is unfamiliar, risky, inconvenient, or otherwise costly. Whilst in some cases the innovation can be redesigned to make it more 'contagious', this is not always possible or necessarily effective.

Another dimension that influences uptake is the culture of the would-be adopters. Opposition is to be expected from individuals directly threatened with a loss of power or privilege, but a change in familiar social arrangements can also be perceived as undesirable by those who stand to benefit. The chapters of the book consider how to identify and target relevant social pathways to boost the spread of innovative behaviours.

Over several decades, developments in how sociologists studied social networks (particularly Mark Granovetter's work on 'the strength of weak ties') were boosted by new methods in network and data science, often originating in the mathematical and physical sciences. The ability to study entire populations led to a growing understanding of how different social structures, such as *small-world* and *scale-free* networks, could influence the flows of behaviour across populations.

A common assumption when studying such flows has been that they are analogous to the spread of a pathogen; once an individual is 'activated', they remain so. In reality, all sorts of behavioural changes do not consist of one-off interventions, but rather require on-going maintenance.

A further source of confusion has been the double meaning of ‘strong’ and ‘weak’ in Granovetter’s usage; there are two dimensions of social ties that sometimes get conflated. The first dimension is *relational* and operates at the dyadic level; how powerful is the interpersonal bond of mutual influence between two people? The second is *structural* and operates at the population level; how effective is a tie at facilitating diffusion, cohesion and integration by linking distant parts of a social world together?

The small-worlds insight is based on the idea that despite their limited ‘bandwidth’, relationally weak ties can be structurally influential in their role as shortcuts across a network, which can certainly be true in the case of ‘simple’ contagions. The bulk of the book is dedicated to modifying network models of diffusion to take these distinctions into account, and showing that this is not true for ‘complex’ contagions.

These can only be effectively spread through strong relational ties, which tend to cluster together in friendship groups that have one manifestation as sets of ‘closed triads’ (whereby friends know each others’ friends - see Figure 2). By the same token, long-range structural ties tend to be relationally weak. As a result, relational strength translates into network structure, and vice versa.

From an individual’s perspective, strong relational ties are by definition the most familiar, trusted, and salient. From an ‘objective’ bird’s eye view, however, these appear as dense clusters of overlapping, redundant triangles, a tangle through which a behaviour must slowly struggle, ‘wasting signals’, to reach other parts of the network. In contrast, weak relational ties stand out as structurally important, branching randomly across the population such that the ‘degrees of separation’ between any two individuals remains small, even as a population becomes vast.

Where affect is unimportant, these weak-tie pathways really can be shortcuts for diffusion. However, empirical studies of how behaviour (such as social movements) spreads show that the common intuition that has developed around this understanding is misleading. Collective action actually spreads spatially through clustered networks, more like the wavefront after a pebble is dropped into a pond than a sudden viral explosion leaping across long-range ties. This raises the question of how behaviour can propagate itself effectively despite its limited reliance on what would appear to be the most efficient channels.

Chapter 3: The Theory of Complex Contagions

Centola points out that the puzzle could be explained by a counterfactual; just because behaviour is observed to spread through spatially clustered social networks, doesn’t mean that that is necessarily the best way for it to do so. A behaviour could spread even more rapidly in another network structure. In other words, perhaps behaviours spread through clusters because those are the only pathways that are available to them, leading to the counterfactual hypothesis that *if there had been* more weak ties in the network, diffusion would have been more efficient.

Computational experiments are used to refute the counterfactual hypothesis, and support the proposed alternative; that the *type* of contagion that is being propagated is the determining factor for the importance of weak ties. As outlined above, weak ties are critical for spreading 'simple' contagions which can be transmitted by a single contact, but do not boost 'complex' contagions which require multiple contacts. In fact, as networks are rewired to enhance the spread of simple contagions, they make complex contagions less viable. The 'disease model' of diffusion is therefore inappropriate to describe the latter. The distinction between the two types can be understood in terms of the social mechanisms that create resistance to adoption.

Simple contagions are defined as those for which transmission can be achieved by a single source, such as diseases and information (multiple exposures to the same 'infected' individual also fall into this category). Transmission of complex contagions requires reinforcement from multiple points of contact; they typically involve some cost to the (deliberate) adopter, impose a risk, or the outcome depends on what others are doing (complementarity). This uncertainty means that social confirmation plays the deciding factor in whether or not diffusion of meaningful behaviour change occurs. Centola states that there are at least four social mechanisms that explain the need for contact with multiple sources of activation:

1. *Strategic complementarity*: awareness rarely translates directly into adoption. Network effects mean that the economic impact of a decision depends on how many others have made similar choices, with a critical mass often required to make the efforts of newcomers worthwhile for themselves.

2. *Credibility:* most people will not adopt an innovation without confirmatory evidence from trusted peers that it works.

3. *Legitimacy*: widespread adoption increases the expectation of potential adopters that their decision will be met with social approval, and decreases the reputational risk if things turn out wrong; there is safety in numbers.

4. *Emotional contagion*: socially-concentrated gatherings can reinforce and amplify shared emotional states such as excitement.

The rest of the chapter systematically tests the two hypotheses introduced as competing explanations for why real-world behaviours are frequently observed not to conform with the viral model of diffusion.

In summary, a complex contagion that crosses a weak tie arrives alone, without the necessary reinforcement for adoption. However, weak ties do not only not provide useful diffusion pathways; they can actively hurt adoption if they are added at the expense of eroding the 'wide bridges' of multiple ties between clustered neighbourhoods that are in fact the useful conduits. Wide bridges explain the tendency for complex behaviours to spread in a wave-like spatial manner; in relational terms, they connect individuals who are physically proximate, but they are also strong in structural terms (the key factor in aiding diffusion). The result is that spatial networks composed of clusters of strong ties provide highly efficient routes for diffusion across large and diverse populations (Figure 4).

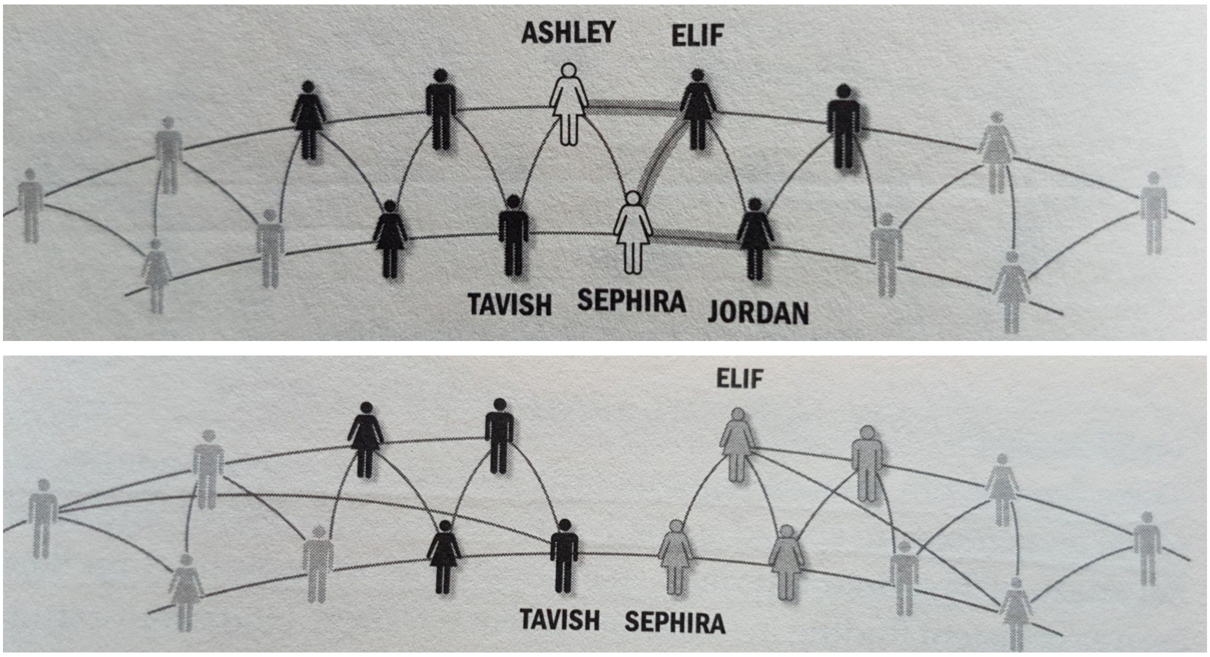

Figure 4: Zooming in on the large world network in Figure 3, Ashley and Sephira are the shared friends of Tavish and Elif, and form a wide bridge between their neighbourhoods that a complex contagion can propagate across (top). Rewiring some of the ties into weak links to make it more like a small world network decreases the bridge's width such that a behavioural contagion cannot spread (although information still can) (bottom).

This finding becomes even more robust as greater realism - such as skewed degree distributions (corresponding to hubs), status differences, variations in adoption thresholds, and variation in tie strength - is added to the network model. Whilst hubs may be ideal for spreading simple contagions, if adoption thresholds are fractional (i.e. people are adversely influenced by non-adopters), then clustering of peripheral sets of nodes connected by chains of wide bridges might be the only way to overcome the countervailing influences preventing a central hub from adopting (and thereby acting as a blockage to diffusion).

Furthermore, fractional thresholds also mean that adding new weak ties (rather than rewiring them from strong ties, the obvious solution to the erosion of wide bridges) doesn't work; because these lone ties cannot transmit the behaviour, people are simply exposed to more non-adopters and it becomes harder for them to meet their fractional threshold, even from their wide bridges. An additional finding is that the greater the status inequality within a network, the more important wide bridges become.

Finally, relative weighting of tie strength was introduced, reflecting Granovetter's original formulation of strong and weak ties. Wide bridges became not just structurally necessary, but relationally efficient for spreading complex contagions. Now, rewiring to add weak ties means that not only are there fewer overlapping neighbourhood ties, but also that those that do exist are less effective at providing social reinforcement. Both of these factors prevent the diffusion of complex contagions.

Chapter 4: A Social Experiment on the Internet

Traditional observational methods are frequently incapable of establishing causality in the spread of behaviour; large populations (due to scope conditions of the small-world model), complete time-series knowledge of adoption behaviour, complete network data (in both cases to exclude confounding factors on the hypothesised driver of diffusion), and replication (since the 'unit of observation' is a population, not an individual) are all necessary to test any given theory. However, online experiments make it possible to not merely watch the spreading process in settings that meet these criteria, but also to investigate the influence of network structure on propagation rates. They offer the control and theoretical precision associated with small-group experiments combined with the scale, observational precision, and natural setting provided by data science.

Centola reports on an experiment designed to show that even among an engaged and easy-to-reach group, the actual usage of helpful resources on improving health would be influenced by the structure of the social network he created. The results are summarised as follows:

- Even the minor 'cost' of registering to use a health forum was enough to prevent many participants who received an initial invite link from a 'health buddy' from joining, and they only did so after receiving further invites. This illustrates the low threshold to qualify a contagion as 'complex'.

- In all trials, the behaviour spread extremely easily through clustered social networks, spilled over into proximate neighbourhoods as they saturated and eventually reached the entire population. Clustering boosted not only the reach, but also the speed of diffusion.

- Awareness of the health forum spread rapidly through the random networks, but was not followed by adoption. In contrast, messages spread locally in the clustered networks, meaning that although initial exposures were much lower, redundancy drove uptake (Figure 5).

- Credibility and complementarity are posited as the driving mechanisms of uptake at the individual level. Both are largely determined by the number of other adopters; as more people participate, the forum can be expected to host a greater number of more diverse materials, recommendations, and people to interact with. These factors explain why multiple invitations made people more likely to pay the 'cost' of signing up; a 67% increase was observed after one additional invite, followed by a further 32% increase for a third, with no evidence of saturation or backfire effects.

- Additional invites were also strongly associated with maintenance of the behaviour. Only 12% of members who joined after one invite ever came back, with none coming back more than twice. In comparison, 34% of those who received two invites came back, rising slowly to 45% after five. Although this is only a correlation, Centola suggests that the findings support a structural interpretation in terms of the reinforcement provided by the clustered network, noting that participants who were most resistant to joining were also the most committed once they had done so. This implies that the same factors that make a behaviour complex also make it ‘sticky’; complementarity can make a technology hard to abandon once it has been adopted. Complexity can therefore be interpreted not as a limitation to diffusion, but rather a necessary hurdle to achieve lasting change.

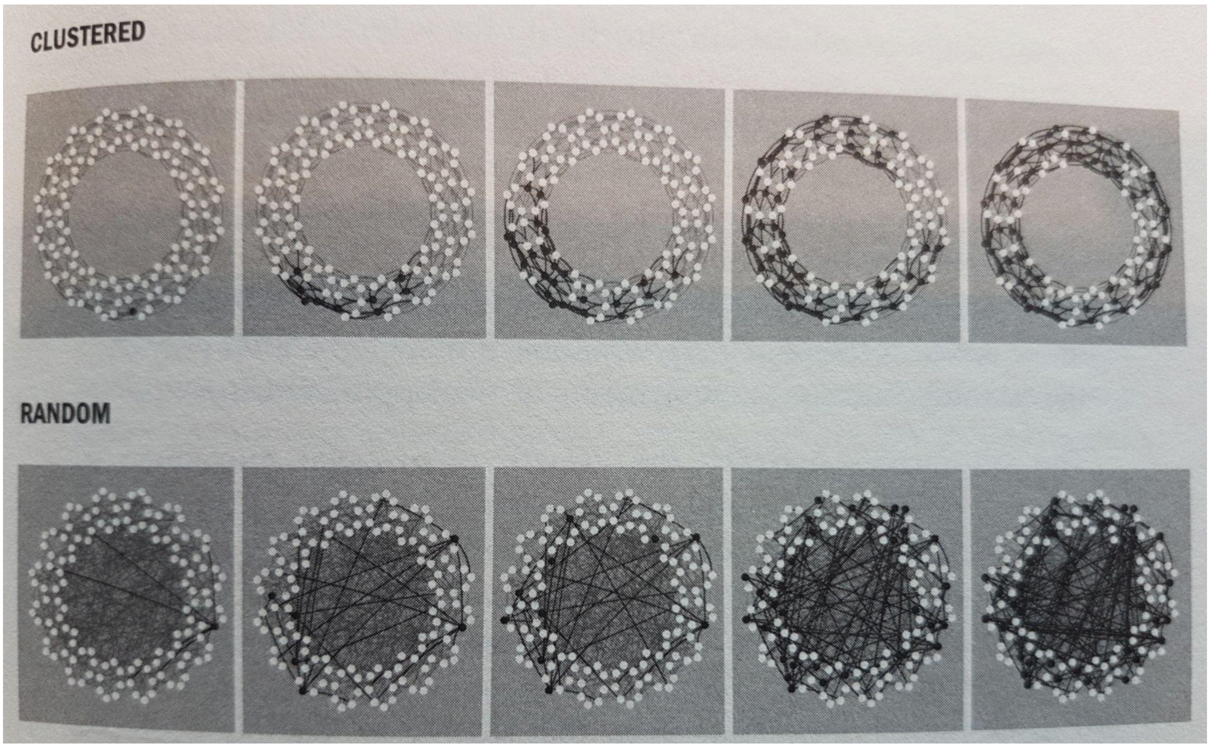

Figure 5: the spread of adoption of an online health forum - a complex contagion - in a clustered (top) and random (bottom) social network. Adoption propagated further and faster in the clustered network, even though all parts of the random network were exposed to the idea much sooner.

In summary, redundancy rather than reach is the key factor in diffusion and maintenance. Centola suggests that because the costs associated with the behaviour in this experiment are trivial compared to major lifestyle changes (for example), the results are likely to be conservative. Furthermore, known reinforcing factors for diffusion (such as strong interpersonal affect) that were by design not included in the study are expected to further amplify the importance of network clustering.

Part II - Applications

Chapter 5: Complex Contagions in Other Contexts

There has been much recent work looking at the factors that determine complexity in online diffusion. New media technologies either have a financial cost or, if they are free, require high complementarity to be worth joining; unilateral action will likely result in wasted time, effort, or the opportunity costs of social exclusion. Facebook, Twitter, and Skype all spread as complex contagions, and although these sites also make use of elements that strengthen the social diffusion of simple contagions within their networks (such as ‘share’ and ‘retweet’ buttons), intra-platform complex diffusion can also be observed. On Twitter, how hashtags spread depends on their subject matter; conventional idioms diffuse as simple contagions, whilst political hashtags involve a degree of risk and so diffuse in a complex manner.

The distinction is much more important than the specific platform, and highly central individuals who are effective at spreading idioms are far less so at diffusing political messages, which propagate best through densely-connected people at the periphery. Furthermore, perceived legitimacy is boosted when diverse members of a user’s network adopt a behaviour (indicating greater normative acceptance), such as the ‘equality’ banner on Facebook.

In terms of social movements, if joining has low risk/cost, then the primary problem is one of access and a word-of-mouth campaign can be effective for boosting exposure and hence recruitment. In contrast, for high risk/cost collective action the problem is resistance, which is best overcome by close-knit networks that provide trust and reinforcement.

As an interesting aside, in *The Origins of Totalitarianism* Hannah Arendt observed that citizens in repressive regimes are allowed to live in a web of weak-tie interactions, but actively prevented - in part through the fostering of a culture of suspicion - from forming strong, cohesive associations. The result is a society through which regime propaganda can easily spread, but in which the dense networks that can support high-risk opposition cannot form.

Chapter 6: Diffusing Innovations that Face Opposition

Strategies for overcoming the barriers to complex contagions and ensuring they reach the maximum number of people include clustered seeding of health interventions, and using incubator neighbourhoods that allow new technologies to gain a foothold against an established competitor. In both cases, minimising early adopters’ exposure to the rest of the (resistant) network is critical.

In the first computational study reported, the aim was to find the most effective seeding strategy for a health intervention that faced entrenched opposition, and required ongoing maintenance. The ‘viral’ strategy of random seeding to maximise network exposure to the intervention initially showed some success, but the lone seeds were generally incapable of persuading their neighbours to adopt and, unsupported themselves, were soon worn down by social inertia. In contrast, treating all the members of a clustered neighbourhood resulted in self-sustaining spread and lock-in of the behaviour as a new social norm (Figures 6 and 7).

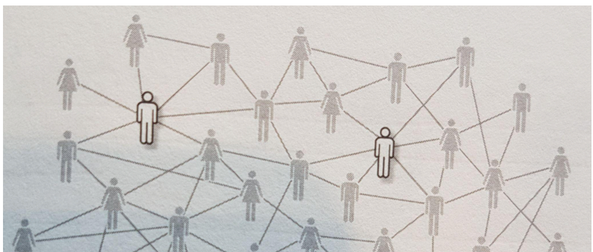

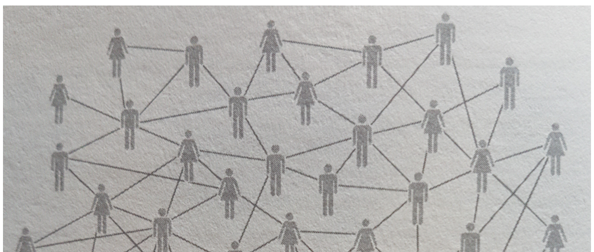

Figure 6: random seeding (left) that maximises network exposure to the health intervention, and clustered seeding (right) that creates mutually-reinforcing signals.

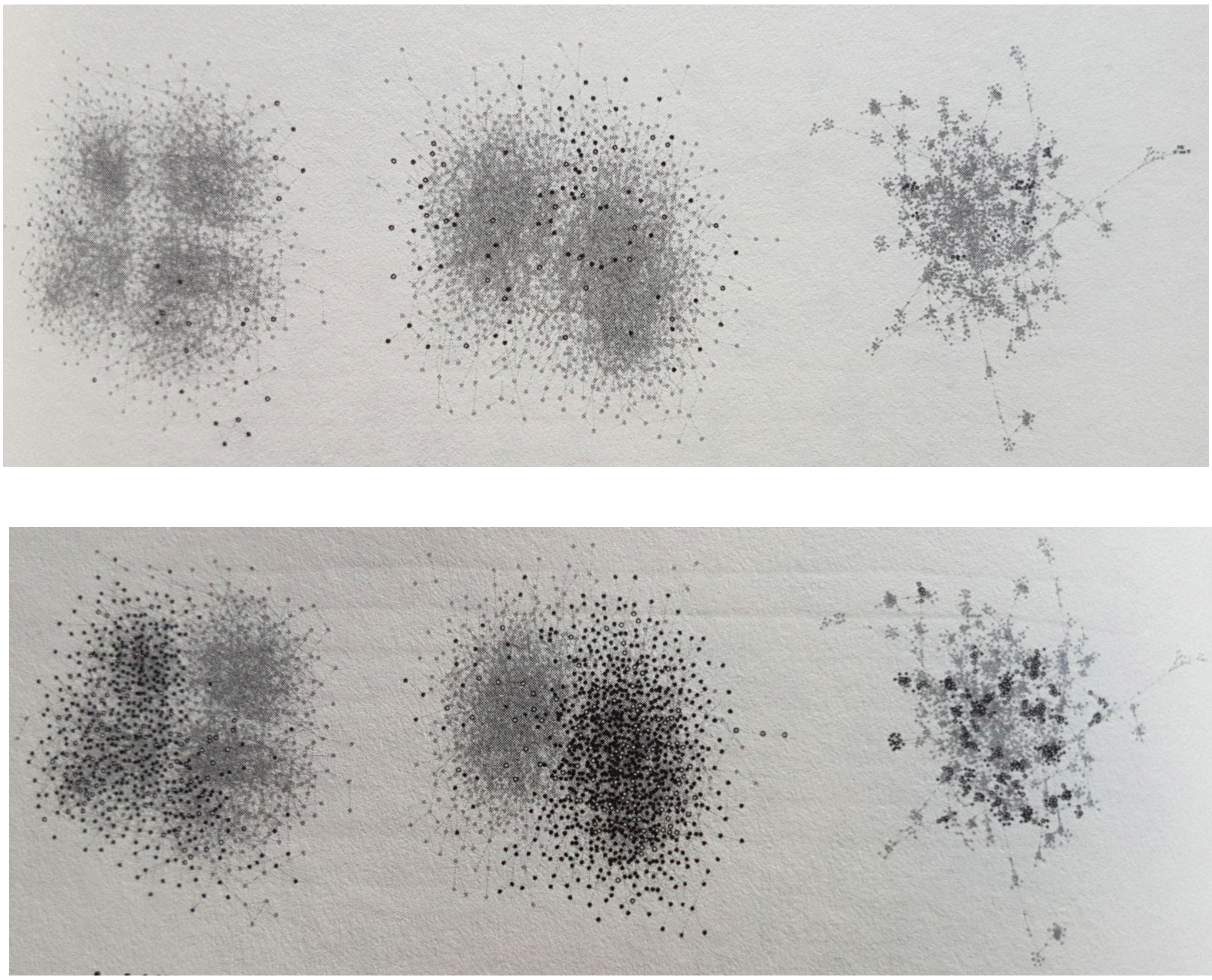

Figure 7: final adoption for random seeding (top) was extremely low, as isolated signals could not overcome social inertia, and in fact the seeds themselves abandoned the treatment after a while due to lack of support. Final adoption for clustered seeding (bottom) was far higher - the intervention became 'locked in' as a stable social norm.

The second study looked at whether the structure of a social network could determine whether an innovative technology could displace an established alternative. The important difference with the previous experiment is that here, both technologies were freely available and known to all individuals from the start. Whereas the first study looked at exogenous seeding strategies, this one looked at how the network structure could be used to initiate endogenous diffusion.

The results showed that clustered networks have the key advantage that ‘dyadic’ coordination is embedded within a more robust ‘triadic’ structure; two people experimenting with a new technology could easily persuade their mutual contact to do so, rapidly solidifying neighbourhood adoption. The challenger technology could thereby gain an initial foothold (even in the face of an entrenched alternative), from which it could spread via wide bridges to other neighbourhoods, which in turn reinforce each other and prevent recidivism.

The resulting stability provides insurance against accidents and individual defections that would easily derail coordination across long ties. Clustered neighbourhoods can therefore be used as social incubators that protect the early adopters of a contested innovation from countervailing influences from the rest of the population.

Interestingly, Centola notes that the spread of cooperative behaviour in a population of defectors resembles the diffusion of a contested innovation. Clusters of altruists providing each other mutual benefits can outcompete nearby defectors, and eventually flip the population into a new state. Their capacity to do so is dependent on wide interlocking bridges, and so encouraging a critical mass of these is an effective leverage point for social innovators.

Chapter 7: Diffusing Change in Organisations

Network theory as applied to organisational performance resulted in the concept of ‘structural holes’, which are gaps between diverse clusters that prevent them from sharing information. Narrow bridges may solve this problem, but the previous results suggest that they will not aid the diffusion of any behaviour that requires social reinforcement. In this chapter the focus is on the requirement for means of coordination within an organisation as the source of complexity, rather than resistance.

Centola notes that there may be a tension between the interests of someone playing the role of a knowledge broker (a weak tie between two clusters) and the interests of the organisation as a whole. Controlling the flow of information provides an opportunity to behave strategically, and because the broker can exploit both groups for personal gain, they actually have an incentive to prevent the formation of additional linkages that would aid the diffusion of complex behaviours.

A computational experiment in which innovations were seeded in a single neighbourhood, each embedded in a different network structure, showed that diffusion as a function of brokerage potential follows an inverted U-shape; with either low brokerage (a snowflake-like network) or high brokerage (clusters joined by a single individual), adoption failed. The only case for which the behaviour was integrated across the vast majority of the population was where clustering in the seed neighbourhood was supplemented by wide bridges between neighbourhoods (Figure 8).

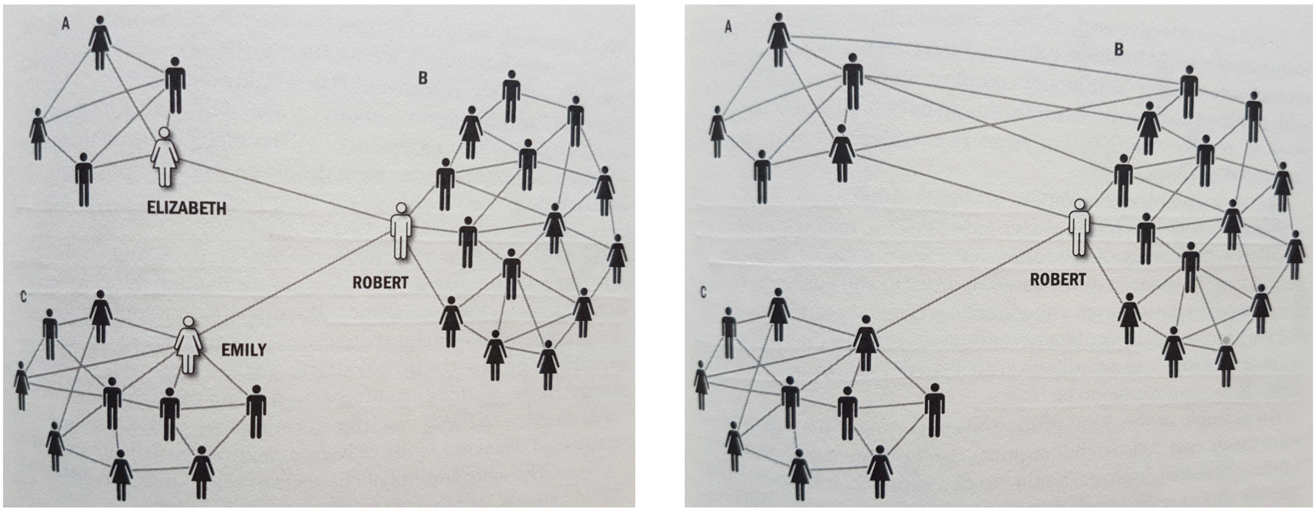

Figure 8: weak ties between groups (left) allow ‘brokers’ to share information, but not coordinate complex behaviour change. Wide bridges (right) can support the propagation of complex contagions.

Wide bridges offer three advantages over brokers: brokerage ties are weak and potentially untrustworthy (whereas reputational effects constrain the actions of bridge members); divergence between interests and goals of an innovating and receiving group may limit the ability of a broker to persuade the latter to adopt (whereas wide-bridge members of a receiving group can confirm a behaviour’s value with the innovating group, and work together to integrate it appropriately); and finally, wide-bridge channels are far more robust in the face of attrition due to staff turnover.

The same arguments also apply to linkages across organisations. In summary, brokers can be essential for creating new pathways for information diffusion, but their privileged structural position can hinder the spread of innovative practices, which require wide bridges to support their complexity.

These bridges can be induced in face-to-face settings by strategically influencing the social network, i.e. the opportunities people have to meet and interact with each other. This requires an understanding of people’s organisational identities, which are profiles of their social contexts. These profiles are often easier to construct than a complete picture of their social network, frequently consisting of information that is already known or easy to discover, such as where and who they work with. With a collection of these profiles, it becomes possible to explore how the corresponding structure of network ties can be redesigned by altering people’s affiliations, and hence the organisation’s capacity for innovation.

The single most important factor for the network structure is whether organisational identities are *focused* or *expansive*. In a highly focused setting, people in a given organisational subunit interact almost exclusively with others in the same unit; given knowledge of which one people belong to, it is extremely simple to predict their social contexts. Such organisations tend to be hierarchical with rigid boundaries between clustered cliques. In contrast, in an expansive setting, knowing one aspect of a person’s profile tells you almost nothing about what other social groups they may belong to; people are free to form diverse ties, almost at random, across these flat organisations.

As might be expected from previous chapters, expansive organisations may be very effective at spreading information, but individuals may struggle to coordinate to effect a change in organisational norms, whereas focused organisations tend to exhibit fragmented but entrenched behaviours, with little integration across groups. *Balanced* identities represent a compromise of approximately 70% focused and 30% expansive, corresponding to a network of overlapping groups connected by wide bridges. Centola suggests four practical implications of these ideas:

1. *Measuring social networks*: it is easier to collect data on people's identities than on their social networks, and these can give a good impression of the overall social network structure.

2. *Changing social networks*: people's organisational identities (and therefore patterns of social interaction) can be changed by altering their working groups, locations, and other affiliations.

3. *Initiating critical mass*: organisational networks can be designed to provide the necessary structure for accelerating change. Achieving critical mass to implement a change is more important than just getting everyone to agree that it is a good idea in principle. Coordinating this critical mass will be difficult in an expansive setting. In a focused setting, it may be easy enough in one cluster, but the challenge will be spreading it to other clusters. Balanced organisations make this much easier.

4. *Making adaptive organisations*: furthermore, whilst balanced organisational identities are capable of generating innovative ideas, they also automatically filter them, since an idea must gain support from multiple overlapping groups in order to take on momentum.

Part III - Social Design

Introduction to Part III

This part of the book considers the more general implications of how diffusion can be influenced through social identities. It starts by noting the vast implications of the findings on complex diffusions; as the world becomes more connected (more like a small-world network), the proportion of weak ties in people's lives increases, meaning that simple contagions will increasingly be selected for over complex contagions. This could result in a 'cultural amnesia', with society finding it ever harder to adopt behaviours that contribute to the common good.

Survival for beneficial innovations may therefore depend more on the strategies that are used to propagate them than on their intrinsic merit. Resorting to making the innovations simpler - in the hope that they might spread virally - might merely precipitate an arms race to the bottom, with the scope for complex behaviours eliminated altogether. However, Centola argues that people will always seek out spaces that provide them with the social capital to achieve beneficial changes, and the lessons from previous chapters can aid in designing them. The next two chapters demonstrate empirically how this can be done.

Chapter 8: Designing Social Networks for Diffusion

Centola presents an illustrative example of building a website to encourage health-related behaviours. He opens by noting that the single most important insight from previous chapters is that wide bridges of social reinforcement are necessary to spread complex contagions. However, other features (in addition to network topology) have yet to be considered; *homophily* - the similarity of social contacts to each other - is one influential factor that does not matter for simple contagions.

Peer networks and settings can be designed that make relevant similarities (which are largely determined by the social context) as salient as possible, at the expense of factors that would otherwise inhibit influence. The result is that informal relationships with effective strangers in an online health forum can sometimes be more influential than the formal face-to-face patient physician relationship, as similarities of experience (in living with a condition, finding resources, choosing treatments etc.) provide hard-to-find empathy, information, and a sense of social agency.

Two observations in particular suggest that a strategy for social planners can be deduced: the first is that people are highly choosy about who they interact with, and the second is that these selections often result in influential relationships. The same factors determine both how people choose their network ties, and who can change their behaviour. If the traits that characterise attraction can be identified, they can also be used to channel influence. At least for online settings, it is possible to narrow the focus to a very specific context to test the 'selection and influence' strategy for turning strangers into relevant sources of influence.

In the online health setting, it turned out that just three factors (age, gender, and BMI) were important for homophilous social selection. This finding was then used to construct a social network on the basis of these preferences, with the resulting social reinforcement driving in a 200% increase in the uptake of a dieting tool (with the greatest effect in the least healthy members of the community, who had been least likely to use it in a non-homophilous network). Note that both the homophilous and non-homophilous networks had a clustered structure (Figure 9). This effect recalls findings from behavioural economics on choice architecture, and confirms that identifying selection preferences and using them to match people into influential relationships can be an effective strategy.

Figure 9: individuals with varying traits (indicated by shade) randomly distributed (left) and homophilously arranged (right). Both networks had the same structure that offered mutually-reinforcing signals, but only the second provided reinforcement from socially-relevant peers, leading to much greater rates of behaviour change.

Chapter 9: Creating Social Contexts for Behaviour Change

This chapter extends the previous finding to explore how the relational context (particularly people's expectations) of social ties can also be used to influence behaviour. Considering social settings as institutions, the rules, rewards, and punishments can shape expectations of behaviour, and hence perceptions of relevance. In online settings, it is possible to design such narrow contexts that very few features come to define who people are, how they behave, and what their motives for interaction are. This presents an opportunity to increase interpersonal relevance and meaning simply by omitting details on the status characteristics that inevitably affect offline encounters.

A key limitation up to this point has been that in almost all the diffusion settings considered, the behaviour has been seeded exogenously, as if a pebble were dropped into a pond. Although a useful approximation, particularly where interventions are one-off, it does not necessarily work well when behaviours emerge endogenously. For example, the decision to start dieting may be a behavioural contagion, but the constituent elements of actually dieting (meal timing, size, frequency, location, content etc.) are complexities - with social contexts that mediate and coordinate them - that should not be overlooked. A more appropriate metaphor is therefore the interference patterns that emerge as raindrops splash into the pond; rather than a single behaviour spreading through the population, interdependent interactions give rise to a change in collective behaviours.

An example based on a fitness programme created from the ground-up - with every aspect designed to enable detection of minimal online 'health-buddy' network influences on offline exercise behaviour - to study these effects is described. Health buddy connections were chosen on the basis of social reinforcement and social relevance described in earlier chapters. The results showed that the rules of the constructed online social worlds directly controlled, with predictable and sizeable effects, offline activity levels. Four settings were created:

1. Control: monetary rewards based on individual activity, but no social interaction. This provides the null hypothesis for this experiment; participants will be motivated solely by monetary reward, with social incentives from anonymous online health buddies irrelevant to activity levels.

2. ‘Social comparison’: individual rewards and no social interaction apart from visible comparison metrics. The hypothesis here is that reinforcement from anonymous peers will increase aspirations and hence activity levels.

3. ‘Social support’: team rewards and with social interaction features based on what participants thought would help them (a chat tool for sharing advice, encouragement, and offline contact details with team mates). The hypothesis is that collective incentives and communication technologies will allow participants to form affective ties, increasing social influence and forming norms for greater activity.

4. ‘Group comparison’: team rewards with visible comparison metrics for other teams. The hypothesis is that competition with out-groups will result in inter-team comparison effects (analogous to those seen in the social comparison setting) that motivate in-group performance.

The health-buddy networks had a significant effect on physical activity, and so the null hypothesis could be rejected. The social comparison and group comparison settings resulted in significant outperformance of the control setting baseline, with the latter giving a 90% increase (note that the collective incentives did not result in free-rider effects or a reduction in individuals’ engagement, contrary to the ‘economist’s prediction’ that the null hypothesis was also derived from). Contrary to expectations, the social support setting produced the worst outcomes, with participants 17% less active than those in the control setting; they would have been better off without their health buddies.

Clearly, the relational context of network ties matters. These results suggest some pitfalls and solutions for social planners:

- In a competitive setting, the most salient individuals for comparison are the high-performers, against which participants benchmark their own aspirations. In a supportive setting, the expectation of encouragement can mean that the least-engaged become the most salient, and the same benchmarking process results in a downward spiral of activity reminiscent of the social inertia/countervailing influence effects seen in Chapter 6. Although many studies claim to show that supportive social media settings can help people exercise, diet, quit smoking etc. the programmes they report on typically include some level of social comparison that is not taken into account.

- Social support is an intuitive explanation for online influence, and it is plausible that it would have been cited by members of the group comparison setting as the reason for their success (even though it is explained entirely by inter-group competition). Although support can be valuable in affecting expectations and thereby creating relevance, this shouldn’t be assumed.

- Sometimes less is more; in the best-performing setting (group comparison), participants could not communicate but were simply exposed to how they stacked up against competing groups, whereas the worst-performing group had access to chat tools with real-time connectivity. The right kind of social capital is much more important than sophisticated social technologies.

- Furthermore, the fact that people actively requested the tools that led to disastrous outcomes suggests that using self-reported preferences can be dangerous for social design.

- Overall, however, the results are encouraging in that they demonstrate that small changes in relation contexts (in particular, how social relevance is created) can have a transformative effect on the dynamics of group behaviour in online settings. These findings are cumulative with those of previous chapters, and successful diffusion relies on all of the elements identified.

Chapter 10: Conclusion

The results here show that the diffusion of novel complex behaviours can get started and become fixed as a stable norm in a diverse range of contexts, with minimal assumptions about the settings and individuals involved. The core requirement is simply that individuals are enmeshed in social networks that offer them relevant sources of reinforcement.

These findings are only comprehensible in light of a distinction between the structural and relational strength of ties, factors that have often been conflated in studies of contagions across networks. This is understandable, since affectively potent ties also tend to connect structurally proximate individuals in clusters, but drawing the distinction is necessary in order to effectively apply these concepts to specific contexts. Otherwise, there is a high risk that intuitive strategies based on the viral model of diffusion will backfire, with the risk increasing as the desired behaviour becomes more complex.

Part I of the book describes and experimentally tests new theoretical results on network diffusion, demonstrating that network clustering improves the speed and reach of complex behavioural contagions. Reinforcement from multiple contacts is also found to be important for fixing a new behaviour.

Practical applications are considered in Part II, with clustering of early adopters (limiting their exposure to countervailing influences from non-adopters) found to be essential for propagating an innovative technology in the face of an entrenched alternative. Organisational identities can also be used (by altering the social contexts people encounter) to enhance the spread of complex behaviours. The overarching conclusion to this point is that the presence and cultivation of wide bridges of social reinforcement are essential for complex behavioural diffusion.

Part III then looks at how the relational aspects of network ties - in addition to the structural factors just described - can be used to further boost the effectiveness of online policy initiatives. Empathy is identified as a key determinant of diffusion patterns, and it can be induced even between near-strangers without a shared past or future by emphasising homophilous aspects of their relationships. The resulting social capital influences kinds of collective behaviour that can subsequently diffuse.

Social inertia is identified as a potential obstacle in all of these settings. Strategies to avoid a failed diffusion effort include favouring peripheral networks over hubs (unless they can definitely be recruited), intentionally seeding interventions in clustered neighbourhoods (rather than at random), and designing settings in which social comparison results in a focus on individuals exemplifying the desired behaviour.

Reinforcement and relevance are found to be key factors in all contexts considered. However, some conclusions are more nuanced and require further research; high homogeneity may be required to induce people to participate in ritualistic activities, whereas signals from diverse constituents may increase the perceived legitimacy of political movements. In general, contexts that support the creation of the necessary social capital can be included in the design of institutions (broadly understood).