

'Change: How to Make Big Things Happen' by Damon Centola

February 2022, summary by Tom Woodroof

This document is a summary of Damon Centola's book '*Change: How to Make Big Things Happen*', with concepts selected according to their relevance to MCS' strategy for implementing the Credit Commons. Implications for different contexts are to be described elsewhere - this document is essentially a highly condensed version of the book itself.

Introduction

The standard ideas behind 'viral' growth are that well-connected 'influencers' play a dominant role in the spread of an innovation, with the 'stickiness' (an analogue with infectiousness) of the innovation itself the other key factor in social marketing. This works for spreading information or simple ideas, but not beliefs and behaviours, which are much harder to influence. However, there are strategies that can be used to spread these complex contagions through social networks, online and offline.

Social change is about norms rather than information. Social networks should be viewed as prisms through which people understand ideas, beliefs, and behaviours, not as mere pipes for transmitting them. This network bias can reinforce people's existing biases, but can also be used to trigger enthusiasm for an innovation or social movement.

Part I - Pervasive Myths that Prevent Change

Chapter 1: The Myth of the Influencer: the (Un)Popularity Paradox

The ability of social stars (or 'influencers') to spread change is one of the most enduring and misleading myths in social science. The huge spread of Twitter is often attributed to Oprah Winfrey's adoption in 2009, but in fact this was the result of its success, not the reason. In fact, its growth had already peaked when Oprah sent her first tweet, and although it kept on growing, it did so at a slower rate. The important question is not 'how did Twitter recruit Oprah', but rather 'how did Twitter get so big that Oprah herself stood to gain from using it?' It turns out that the social periphery, not the social stars, are where companies and campaigns can grow from obscurity into household names.

We are typically moved to adopt an innovation that requires overcoming a 'coordination problem' (i.e. something that only makes sense if other people also adopt it, such as using a handshake versus a high-five as a greeting) by the fraction of people in our network who have already started using it, not by the absolute number. This means that the more connected people are, the less likely they are to adopt a new idea or behaviour, as it is much

harder to convince them that it is legitimate; non-adopters act as 'countervailing influences', and a typical leader in a central network position is much more exposed to the vast majority of non-adopters than they are to the small number of early adopters. In contrast, it is much more likely that the adoption threshold of someone on the network periphery (with a more modest number of connections) will be met. It is therefore possible for an innovation to take hold, gain momentum, and spread through the periphery until it becomes impossible to ignore, even for people at the network's centre. This is what happened with Twitter; once a 'critical mass' of ordinary people had been reached, social stars who had previously been reluctant became convinced of its legitimacy and turned into dedicated users.

In 1944, sociologist Paul Lazarsfeld had coined the term 'opinion leaders' to refer to members of the population who were far more attuned than average to media advertising. This challenged traditional broadcast media, but his 'law of the few' suggested that advertisers need only concentrate resources on well-connected social stars, who would then pass the messages on to the rest of the population. This has become broadly accepted because it can work well in some situations, and today it is referred to as 'influencer marketing'.

The myth of the influencer suggests that it is necessary to find and recruit these special people in order to spread an idea, trend, or movement. However, as seen above, highly connected people are often roadblocks, and are in fact the last step in the change process. Instead of special people, what matters are special places in the network.

Studies of the fall of the Berlin Wall, the US civil rights movement, and the Arab Spring have shown that people do not join revolutions just because they are unhappy. Rather, they join them because of their social networks; if people see that other, similar citizens are taking a stand, they come to believe that they can also make a difference, and they want to take part. This can actually be a problem for dedicated, high-profile activists, who are typically younger, have fewer responsibilities, and are more willing to put themselves at risk than the average citizen. All these traits serve to highlight how different they are from others.

To recruit ordinary people into activism, what matters is the network periphery. Despite social media, the patterns of people's social networks have actually not changed very much; most people's ties are primarily with friends, family members, colleagues, and neighbours. Data from the Arab Spring showed that the greatest predictor of real-world activism (such as protest attendance) was coordination in the large, apparently unremarkable network periphery. As such, it cannot be claimed that social media was responsible for the Arab Spring.

The myth of the influencer appeals to a cultural love of heroes, but they are often just fairly ordinary people located in the right part of the network at the right time, which allows their actions to have a decisive impact. The 'right time and place' is in fact a measurable feature of social networks, patterns which can be identified and targeted, online and offline.

Viral spreading applies to both the Black Death and Covid-19, but whilst the former spread through Europe over many years, the latter took only weeks. Although the difference in their pathways is due to modern transport networks, the core concept of viral spreading applies equally to both. This suggests that, when it comes to influencing contagion, the reach characteristic of modern-day networks dominates the redundancy of mediaeval networks. Sociologist Mark Granovetter concluded that the people with whom you have weak ties are the best ones to recruit for job-hunts, advertising, and change initiatives. These ties allow us to connect with many more, and more diverse, people.

However, none of the data on how social media technologies, movements, or norms spread support the importance of weak ties. Although Twitter spread incredibly rapidly, it did so through pathways completely different to those associated with viral spreading; through redundant ties. It initially spread from neighbourhood to neighbourhood in San Francisco, and then across the California countryside. However, before it reached critical mass in Los Angeles, it jumped across to Cambridge, Massachusetts, following the strong ties between young, tech-minded graduates from each coast's top schools; a defining feature of the modern era is that there are networks of people who are socially close but geographically distant. It then resumed its geographic spread, eventually reaching New York and Chicago.

This pattern overturned half a century of conventional wisdom on the importance of weak ties, revealing redundancy to be crucial (rather than inefficient). Social change (such as the spread of writing and religions) has always happened this way, but it is only recently - with a globally networked world - that it has become evident that this process is distinct from viral spreading.

Chapter 3: The Myth of Stickiness: Why Great Innovations Fail

Inferior innovations have often beaten their competitors in the marketplace because their producers have found better ways to use networks, particularly to gain incumbency. The natural inclination is to improve the product itself, making it easier to use, more striking and exciting, and less expensive. However, catchy advertising, aggressive marketing and compelling science are not enough to overcome the enduring opposition to change that can arise from social and cultural norms embedded in our networks, even amongst the people who are most in need of the change.

Ideas of product 'stickiness' - features such as practicality, novelty, tangibility, and emotional charge - as a means to success can be misguided, leading a product to not only fail, but backfire. For example, the Google Glass launch campaign was based on recruiting high-status early-adopters to spread the remarkable, sophisticated, memorable, and discussion-worthy technology to everyone else. Individually, the two ingredients of *awareness* and *differentiation* can be very effective, but they do not mix well. When Glass came out, everyone heard about it, but the only people wearing it were techies, who were culturally, economically, and socially distinct from people who were not wearing it (but knew they were supposed to want to). The exclusivity did not rouse aspiration, but rather resentment. It solidified latent social differentiation between the kinds of people who would

and wouldn't wear it, creating a normative backlash. Google's image was transformed from cool search company to a tech giant making surveillance technology for the rich.

The demographic transition that many developing countries entered in the 1960s resulted in the threat of devastating overpopulation for countries that could not rapidly change their norms around family planning. Whilst some countries (like Pakistan) relied on broadcast media strategies and failed, others (like South Korea) achieved all of their goals ahead of schedule by using complex contagion strategies. The redundant connections of strong ties within village social clusters (such as tightly-knit 'women's groups') were used to achieve social acceptance of contraception amongst early-adopters. The strong ties between clusters then enabled it to spread throughout the rest of a village. Although contraception also challenged social norms, its success stands in sharp contrast to Google's Glass.

Google also attempted to break into the social media market. People with accounts were signed up automatically for Google+. Although this ubiquity generated massive awareness, it backfired, as with Google Glass. However, this time the two conflicting ingredients were awareness and *lack of use*; the more that awareness outpaces uptake, the stronger the implicit signal that there must be something wrong with the innovation. Everyone knew about Google+, and knew that everyone else knew, but also knew that no-one was using it. Google accidentally created enormous social proof against its own platform, a particular problem for a social technology that relies on social coordination (i.e. people moving together). By attempting to mobilise the entire population in one go, Google effectively guaranteed that the countervailing influences would be insurmountable, making it impossible for Google to challenge the powerful incumbent (Facebook). Following the results in Chapter 1 on mobilising an insurgency, a much better strategy would have been to start at the network periphery.

In 2005, researchers unveiled PrEP, a miracle HIV drug. Its marketing campaign in Zimbabwe followed the viral-marketing strategy, with a message that PrEP was free, easy to use, and lifesaving. However, it was actively resisted, as people feared that taking it would indicate to their community that they were infected. Given the stigma of HIV, there were clear social reasons not to take it, regardless of what medical authorities or marketing campaigns said.

The myth of stickiness suggests that key product features and marketing are what ensures an innovation's success. However, social and cultural norms mean that the less familiar and the more disruptive an innovation is, the more resistance it will encounter. This is the main reason why social change is so hard. The solution is not influencers, viral marketing, or stickiness, but rather the infrastructure of contagion that can trigger a chain of diffusion and an explosion of social change.

Part II - The Changemaker's Playbook: How to Build Contagion Infrastructure

Chapter 4: How Change Happens: the Discovery of Complex Contagions

Although simple contagions have become established as the default model for how social change occurs, the reality is that any change that involves risk (financial, psychological, reputational) will induce resistance that prevents people from being 'infected' simply by contact with a single 'carrier'. This resistance creates complexity in how the contagion spreads, with the barriers taking four main forms:

- **Coordination:** many innovations only make sense to adopt if others are also doing so.
- **Credibility:** if there is scepticism about the risk or effectiveness of an innovation, repeated confirmation by others is necessary to create the social proof that it is not as risky as feared.
- **Legitimacy:** adoption can pose a risk of embarrassment or tarnished reputation if things do not go as planned. The more people who adopt, the stronger the expectation that others will approve.
- **Excitement:** some innovations are only attractive in an emotionally energising setting.

However, this resistance is also an opportunity to achieve enduring change. Although people require reinforcing contact with multiple adopters to be convinced, the barriers identified above can - once overcome - make an innovation truly sticky. In other words, innovations that face the greatest resistance are also the ones that people often become most committed to once they have chosen to adopt them. Sociologists refer to this as 'entrenchment'. For example, recruitment for the British Army in the first world war targeted the strong ties in residential neighbourhoods and professional communities, transforming social norms against enlistment by professionals into an expectation (via social proof) that they do so.

The network geometries associated with weak and strong ties look very different. The former looks like a fireworks display, with each person at the centre of an 'explosion' of radiating connections. Each tie is a leap to a random, sometimes distant part of the network, and there is very limited redundancy (people are not connected to each others' friends). The latter looks like a fishing net, composed of interlocking triangles and rectangles. This 'clustering' corresponds to a high degree of redundancy, with people having many friends in common. Not surprisingly, real-world social networks are a combination of these two types, but they have very different effects. Whereas fishing net patterns foster trust, intimacy, and solidarity, fireworks patterns do not provide the accountability to underpin solid cooperation.

A 'health-buddy' experiment among strangers in social networks with these two geometries resulted in widely divergent use of a health recommendations forum. 'Early adopters' (who created an account after receiving a single message) in a fireworks network typically used the technology once, but then never logged in again. In a fishing net, those who required multiple invitations ('laggards') were three hundred times more likely to keep using the forum than the early adopters in the fireworks network, and continued to use it long after the

experiment finished. In this instance, network redundancy was serving two roles. Firstly, reinforcing messages from multiple health buddies provided social confirmation that the forum had the necessary coordination value and credibility for adoption. Secondly, the reasons for adoption (particularly coordination) also kept people coming back, as more neighbours meant more health recommendations.

Chapter 5: Complex Contagion in Action: Memes, Bots, and Political Change

For 'idiom' hashtags on Twitter, viral diffusion works. However, political hashtags (due to the perceived social risk for users in aligning themselves with a particular stance) are complex contagions, with people typically waiting until they have seen several people in their network using them before adopting them themselves. In fact, any potentially contentious idea requires the kind of network geometry that provides redundant social confirmation.

Studies by two British mathematicians of the top 26 viral memes of 2014 (particularly the 'ice bucket challenge') found that no common theme, feature, or trigger was shared by all of them. The only mathematical distinction was that the successes had almost universally benefited from socially reinforcing networks; although they spread rapidly, they were complex contagions. The mathematicians then created a model that successfully predicted the growth, peak, and decay of the ice bucket challenge, along with other memes.

A Danish computer scientist then took the idea further, and attempted to use the principles of complex contagion to spread their own memes using 'bots'. The memes were designed to promote social cooperation and good feelings, and they wanted to know whether they would spread as simple or complex contagions, and what the best strategies for spreading them would be. For six weeks, thirty-nine bots linked together into a 'botnet' posted messages, accumulating more than 25 000 human followers. Because the bots were connected to each other, two kinds of social reinforcement were created; firstly, the bots' followers received similar messages from multiple accounts. Secondly, people observed the bots' liking and forwarding of each others' messages (boosting perceived legitimacy), which created a third-party social approval effect. The botnet proved successful in spreading feel-good memes amongst its network, with each one propagating as a complex contagion through redundant social ties. Sending repeated messages from a single source was ineffective; the signals had to come from multiple sources, with the more reinforcing social confirmation, the better. The experiment showed that not only is the social process predictable, it can also be automated, with a relatively small number of bots required to set it off. In short, contrary to years of received wisdom, reinforcement within redundant social clusters is more important than emotional triggers or sticky messages.

Chapter 6: Contagion Infrastructure: the Importance of Wide Bridges

Ever since groundbreaking work by Granovetter in the 1970s, connections across social clusters have been referred to as 'bridges', a term that became interchangeable with the idea of weak ties. Their value was considered to be proportional to their length, or 'reach', and this assumption is still widely held in social sciences, industry, and advocacy. However,

bridges can also be characterised by their 'width' - that is, how many ties they contain. A weak tie is a narrow bridge, whereas a wide bridge implies many redundant strong ties, reflecting a true collaboration. Because these wide bridges are often found in residential neighbourhoods, the complex contagions they support usually spread geographically. But as was discussed in the case of Twitter, it is actually redundancy rather than geography itself that matters. In short, 'Narrow bridges speed information over weak ties. Wide bridges bear social change over strong ones.'

In an organisational setting, this means that whilst someone in a 'broker' position connecting two social clusters (such as engineering and sales teams) via weak ties to both might be in a strong position to aggregate and share lots of *information* from across the organisation, they will probably not be capable of actually *transferring knowledge* or *change*. Their weak ties cannot overcome the barriers of trust and credibility (particularly since the two clusters will be aware that the broker stands to make a reputational gain from their mere participation), risk (associated with disruptive adoption of the innovation, and uncertainty about its relevance), and, most importantly, the need for the entire cluster to coordinate on adoption. To actually pull the innovation transfer off, wider bridges are needed. This could be achieved by creating settings in which a fishing net pattern of relationships between the two clusters can form. Although the broker loses some of their structural advantage as the sole information point between the two clusters, the wide bridges can provide the necessary trust, risk-reduction, and coordination infrastructure to actually achieve the change. Wide bridges also increase organisational stability by ensuring that channels of communication and exchange survive the natural turnover of individuals.

Wide bridges are particularly important for cross-organisational collaboration. Up until the late 1970s, firms had maintained rigid hierarchical structures with well-defined boundaries. However, in response to growing complexity and interdependence in competitive markets, big players in the Japanese electronics industry started to integrate specialised subcontractors into their R&D, manufacturing, and production timelines. By the early 1980s, the industry was an engine of innovation powered by wide bridges. Industrial networks in Silicon Valley copied the strategy, with the resulting trust, reciprocity, and reduced risk leading to an explosion in new technology marking the start of the era of 'open innovation'.

As an unprecedented collaboration across over a dozen competing scientific institutions and nations (without the leadership from a single government agency that the Manhattan or Apollo projects had relied upon), the Human Genome Project had to replicate this network pattern as well. The consortium arrangement - that included regular conferences, reciprocal laboratory tours, shared databases, and electronic peer-to-peer exchange networks as its innovation infrastructure - was a model of collaborative science, and advanced towards its aims (far more ambitious than anything happening in Silicon Valley) at incredible speed.

Contagion infrastructure can also emerge spontaneously among the diverse communities of modern society. In July 2014, a month before the police murder of Michael Brown in Ferguson, Missouri, conversations on Twitter around civil rights, Black activism, and police violence were taking place in independent communities connected only by narrow bridges. The vast majority of communications were internal to each group, just like firms before the era of open innovation. Soon after the shooting, however, Missouri citizens were using the hashtags Ferguson and Black Lives Matter as tools of strategic coordination, as well as

emotional solidarity. This brought previously well-defined groups into conversation with each other, with the shared language enabling cooperation and solidarity across the country. The widening bridges were used to successfully influence the language used by the mainstream media, reframing the narrative of the protests from ‘mob’ to ‘citizens’ and ‘community’.

Following further police shootings, geographically and culturally diverse groups, national news media, and government officials were all able to coordinate on the Black Lives Matter message. Within less than a year, the movement had sparked a national and international conversation that even the White House had been drawn into. By the time of the murder of George Floyd in May 2020, the vast network of wide bridges created by the Black Lives Matter movement was able to spread and transform public outrage into the most widespread solidarity campaign in history, in complete contrast to the isolation of communities previously suffering from police brutality.

Chapter 7: The Principle of Relevance: the Power of People Like Us and Unlike Us

The other crucial element for the success of a complex contagion is relevance. People are far more likely to overcome the inertia necessary to adopt a new behaviour if they see other people they perceive as relevant doing so. ‘Homophily’ is the term for the observation that similar people tend to associate, but only certain kinds of similarity create relevance. No single trait (such as gender, race, fitness, status, age, income, or political ideology) is always influential; context is key. Even so, there are three principles for understanding how relevance is established in a given context:

- If people need social proof that something will be useful for them, similarity with previous adopters helps. For example, in a medical setting a ‘role model’ such as an authoritative professional may actually make people feel devalued if they are perceived as too dissimilar; advice from patients’ peers with similar conditions and who are familiar with the associated challenges may be seen as more credible. This also applies to business governance and strategy; in this case there can also be safety in numbers.
- Emotional excitement, loyalty, or solidarity also require similarity among reinforcing contacts to inspire change. Activating a shared sense of identity can be effective, but this is surprisingly malleable; solidarity can be induced among strangers simply by assigning them entirely fictitious identities. In addition, wide bridges composed of people who share some characteristics of two antagonistic groups (such as Arab citizens of Israel) can act as a ‘bridging group’, the mere existence of which can change each group’s feelings towards the out-group.
- If change is based on legitimacy (i.e. a sense that something is widely accepted), then diversity among adopters is key. The more similar adopters are to each other, the more obvious it becomes that they are different to everyone else (e.g. in the case of Google Glass). The resulting countervailing influences can prevent diffusion. In the case of gay marriage, it has been shown that just ten contacts from diverse segments of someone’s community are enough to persuade them that the social

risks of showing support are minimal. Likewise, because political campaign donations are also complex contagions, diversity amongst early supporters is essential for a new political candidate. Although ‘mobilising the base’ is important later on, what matters initially is demonstrating that the candidate does not just represent a homogenous special interest group. If they can draw broad support, then they send an implicit message of viability, drawing further contributions in a self-fulfilling prophecy. In short, the quality of early support is more important than the quantity. The same applies to social technologies; adoption and continued use of Facebook was driven by how diverse a given person’s contacts were who were on it, not by multiple signals from the same peer group.

It is therefore relevance, rather than similarity per se, that matters. These principles are a guide to the strategies necessary to activate the most relevant sources of social influence in different contexts.

The previous chapters identified wide bridges and relevance as essential components of a contagion infrastructure. The next part of the book looks at how best to deploy resources to kickstart a change campaign, what critical mass is necessary to give it a life of its own, and how to overturn a competing norm that has already taken hold.

Part III: The 25% Tipping Point

Chapter 8: In Search of a New Normal

In September 1967, driving in Sweden switched from left-hand to right-hand. On ‘H-Day’, even though people knew what to do, there was chaos; people did not know what *other* people were going to do. If you saw someone approaching on the wrong side of the road, should you switch to accommodate them, or stick to your guns? At that moment, laws and broadcast information are irrelevant. Solving these ‘coordination dilemmas’ requires mind-reading, but if you have no idea what to expect from other people, it is impossible to coordinate with them. Even though everyone knew the rules on H-Day, they couldn’t read each others’ minds. Such coordination problems are solved by intuitively-known social norms, but these can change (or be changed).

The social power of a witch hunt comes from the requirement for people to protect themselves by concealing their preferences. Even if people strongly disagree with the new norm, it is now impossible for them to read each others’ minds, and their best guess about what they should expect from others becomes based on a shared illusion of consent. The more people conceal their true thoughts, the stronger the incentives for everyone to enforce the norm, lest they be seen as deviant.

Even when the norm is harmless (such as shaking hands), or people conform to it without realising, coordinating on an alternative can be extremely hard. Norms are useful exactly because they automatically solve coordination problems, and breaking something that works

(such as learning a new language) is difficult. Thomas Kuhn showed that paradigm shifts in science are often accompanied by social bewilderment; once-respected experts become irrelevant or seen as incompetent, with wider society affected as well. Although social norms can advance faster than 'one funeral at a time', there are similar obstacles to those faced by the proponents of a new scientific paradigm. A movement must help people through uncertainty to a new set of expectations and a new sense of competence. The key to doing this is to understand how language works and what it reveals about how norms take root.

Wittgenstein's philosophy was based on the idea that language is key to understanding the world, and that language itself (and the resulting social life as a whole) could be understood as a series of social coordination 'games'. This has become central to research in psychology, sociology, philosophy, and computer science, and made it possible to design an experiment to observe a social norm taking hold. The experiment was based on the idea that every coordination game has a 'tipping point' - a threshold at which everyone's ideas about the acceptability of a novel idea could suddenly change. If a tipping point could be reached, an entire population could be efficiently ferried from one social norm to another just by activating a critical number of early adopters. This in turn would make it possible to make reliable predictions about social change.

Chapter 9: Wittgenstein, MeToo, and the Secret of Cultural Change

Although the term 'tipping point' was popularised by Malcolm Gladwell, it is used here to refer to a scientific theory that there is a measurable critical mass that can trigger wholesale changes in a population's behaviour once it has been reached. Business professor Rosabeth Kanter identified the hallmarks of an organisation or institution (such as national politics) in which a critical mass of say, female politicians, had not yet been reached. In these instances, women played a token role, and were seen as emblematic of women generally. They were expected to conform to strong expectations about their behaviour, and following these norms enabled women to avoid failures of coordination. In short, they were not seen as legitimate political actors. Kanter believed that reaching a critical mass of female legislators could cause a major cultural shift, with a hypothesised 20-35% of the population needed. Mathematical modelling of complex contagions suggested that the tipping point would be at 25%.

However, to truly understand this dynamic, it is necessary to show that tipping points really do exist, that they are the cause (among many candidates) of social change, and that the threshold can be quantified. Based on Wittgenstein's idea that the way people make sense of the world is through a series of coordination games, the tipping point can be defined as the moment at which people could no longer coordinate without changing a norm; in other words, the point at which people's need for socially-coordinated behaviour overcomes their desire for continuity and tradition.

An experiment using a 'naming game' was conducted to answer these questions. People in ten independent groups of twenty to thirty people were connected into a network, within which they had to come up with an appropriate name for a random person. In each round, they were paired with another participant, and received monetary rewards or punishments if

they chose the same or different names respectively. People could not use population-level data to inform their choices (they could only see what their current partner had guessed in the last round), and of course there was no right answer. Although there was chaos initially, within fifteen rounds people always converged on an answer. This happened because as soon as two people guessed the same name by chance, they would feel pleasantly surprised and stick with it for at least a few more rounds. When someone encountered both of these players in succession, they would likely start using the same name, as would anyone else with the same contacts. This spark of coordination could then spread throughout the network, being further reinforced with every round. People immediately knew what to expect from a new contact, and since people were rewarded for coordination, there was every incentive to stick with it. Although in each group the name was different, convergence was observed in all cases.

Each group was then seeded with a group of 'committed activists', who were tasked with overturning the norm by relentlessly sticking with an alternative name. When these activists made up 17-21% of the population, their efforts over dozens of rounds made no difference. But once they reached 25% or above, the committed minority succeeded every time in abruptly shifting the norm.

Chapter 10: The Blind Spot in the Mind's 'I': Unexpected Triggers for Tipping Points

People generally attribute the choices of others to a desire to conform with expected norms, but rarely explain their own behaviour in the same way. Instead, they are sure that their own decisions are based on their own intelligent reasoning and preferences. This is known as the 'introspection illusion', and illustrates that most of the time, the social factors that influence people's behaviours are in a blind spot, with attention given to internal states. This strongly implies that people's beliefs about what will motivate them are unreliable.

This was demonstrated in a study to determine which factors householders believed would be most effective in encouraging them to save energy versus which actually were. Out of a choice of 1) environmental and social values, 2) money-saving, and 3) social norms amongst peers, people in the study group thought that only the first two would be effective. This was even the case for the participants who had been given information about their peers' behaviour, and who were in fact the only group that showed (due to this peer influence) any significant reductions in energy usage. Aside from the (unsurprising) conclusion that people's opinions about what motivates their behaviour are inaccurate, the main finding was that the factor that people believed would be least likely to influence them was in fact the most important.

The difficulty in triggering a tipping point is that not infrequently, everyone is waiting for someone else to go first. The most common strategy to overcome this is to deploy financial incentives. However, technology adoption, from televisions, the internet, and smartphones to solar panels is primarily influenced by what people's neighbours are doing, rather than financial incentives, availability, or informational awareness.

For someone tasked with spreading a new norm by inducing a tipping point on a limited budget (say \$10) that can be allocated in any manner between 10 individuals in the network, there are three core strategies:

- The 'shotgun' strategy, based on viral marketing principles of simple contagions. This allocates \$1 to 10 individuals, who are selected to be as far apart from each other as possible in order to create maximum exposure. However, since redundancy is rejected in favour of reach, each 'early adopter' is surrounded by a sea of non-adopters/countervailing influences. There will be no creation of social reinforcement, legitimacy, credibility, or social currency for adoption, and furthermore even well-incentivised change-agents are not immune to peer pressure to themselves abandon the innovation. Finally, once people have been exposed but declined to adopt (and have seen that nobody else who was aware of it has adopted either), the effort will backfire (as happened with Google+); people notice that the innovation failed, and feel a need to find a justification for why they did not adopt it and would not in the future. For all these reasons, the complex contagion will fail.
- The 'silver bullet' is a popular alternative strategy, and concentrates all resources on a single 'influencer' - someone who is so well-connected and influential that they can single-handedly initiate a chain reaction that transforms their community's behaviour. As was discussed earlier, a highly-connected individual faces overwhelming countervailing influences and is therefore unlikely to come out against the status quo. However, if the \$10 incentive works, then the influencer will promote the new norm to all of their contacts. Since (by definition) this person has many widely-distributed contacts, this appears to be an even more efficient way to spread it than the shotgun strategy. However, the influencer's contacts themselves would find themselves in the same position as the shotgun strategy's change agents, but without even the \$1 incentive to protect them against countervailing influences. If the influencer's contacts are tightly clustered enough to support each other, then this defeats the purpose of using an expensive influencer in the first place. Finally, backfire effects can be even worse than for the shotgun strategy; with Google Glass, high-status 'influencer' early adopters violated non-adopters' social norms so badly that implicit expectations about face-to-face interaction turned into an explicit culture war, stigmatising the technology. In general, any challenging social innovation should prioritise local support over widespread awareness.
- The final strategy is the 'snowball'. Instead of targeting special *people*, special *places* where the innovation can take hold are utilised. Rather than trying to convince everyone to adopt at once, the innovation is carefully incubated as support grows towards a critical mass. Similar to the shotgun strategy, the \$10 would be distributed between ten people, but they should be selected from the same social cluster, rather than spread as widely as possible. Their mutual reinforcement allows them to not just stick with the innovation to give it a foothold, but to spread it to neighbouring clusters across wide bridges through coordinating to increase its legitimacy, credibility, and social currency until the tipping point is reached.

As well as successfully spreading contraception in Korea, the snowball strategy was demonstrated by a study of adoption of a sustainable farming practice (pit planting) in

Malawi. Despite years of government and NGO efforts at outreach and promotion, by 2009 fewer than 1% of Malawi farmers were using pit planting. In 2010, a tipping points strategy was tested by a group of scientists led by economist Lori Beaman. After exhaustively mapping the social connections in 200 villages (each with about 200 people), they worked with the government to train 'change agents' for each village. The villages were divided into four groups, and a different strategy was tested on each group of 50: shotgun, snowball, snowball-neighbourhood (in which residential neighbourhoods were used to select change agents, with the hope that this would provide enough reinforcement by chance, i.e. without the expense of producing a full social network map), and silver bullet (which, since this was the outreach strategy already being used by the government, served as the control group).

It quickly became clear that pit planting was a complex contagion, with willingness to even learn about it strongly dependent on social reinforcement (farmers connected to more than one change agent were more than 200% more likely to adopt it than those connected to only one). The study's conclusions were clear: the government's silver bullet/influencer strategy had come in last, with almost no impact on knowledge or uptake. The shotgun strategy was only slightly better. Snowball-neighbourhood had a significant impact, with a 50% increase in adoption compared to the shotgun strategy but no meaningful effect on social norms. Snowball was by far the most effective; even with only two change agents (the absolute minimum required), it produced an almost 300% increase in adoption compared to the shotgun strategy, as well as spreading knowledge. The snowball-neighbourhood strategy proved most sensitive to the number of change agents, since selecting more within the same geographic area would increase the odds that they would share social clusters. The snowball strategy was also boosted by extra change agents, whilst shotgun and silver bullet were not.

The same dynamics also apply in modern industrialised settings, as demonstrated by the spread of hybrid corn in the US. Despite the vast amounts spent in the 1930s developing and marketing it, and the Dustbowl conditions that it was better suited for than conventional varieties, by 1933 fewer than 1% of farmers had adopted it, despite a traditional media and viral marketing campaign that meant almost 70% had heard of it. There were multiple obstacles: expense, fear of the unknown, its odd appearance, and the Dustbowl that made people even more risk-averse than usual. Particularly in tough times, people rely on a reputation for solid decision-making. If they take a bad risk on something that others have already rejected, it makes them look foolish and incompetent. Because these influences fell in farmers' blind spot, rumours that something must be wrong with the seed began to circulate, justifying non-adoption. Ironically, these rumours spread through the same word-of-mouth channels that the earlier advertisements had targeted, with new evidence circulated to counter the rumours only serving to deepen suspicions.

By 1934, the companies promoting hybrid corn had almost given up. By chance, however, a small group of early-adopters in Iowa formed a 'community laboratory' to experiment with the new corn. Once it took hold there, the same factors that caused farmers to initially resist it now caused them to adopt it; they could see the success their peers were having with it, making it seem more credible - social confirmation which reduced the perceived risk of using it, leading to legitimisation. Hybrid corn became an increasingly widespread innovation to help farmers through the drought, and within a decade it had achieved 98% market penetration in Iowa. It then spread across the US, ultimately reaching 100% saturation.

The spirit of the snowball can be seen in Germany's '1000 roofs' strategy, which aimed to redress the gap that existed in 1990 between its ambitious renewable energy goals and the reality of limited adoption. Within a short space of time, the government oversaw the installation of roof-mounted solar on two thousand grid-connected homes. Even though this is a tiny fraction of homes, the selected houses were tightly clustered. A 2016 study reported that in neighbourhoods where a critical mass had been achieved, a dense concentration of adoption had proliferated throughout the entire region and even across state lines and national borders. By this point, Germany had the world's highest per-capita solar energy production, with the analysis finding that neighbourhood-level social influences were essential for the speed and scale of the transformation. Similarly in Japan, the strongest predictor of whether a household installed rooftop solar was the number of people in their neighbourhood who already had it, as opposed to the financial incentives on offer. These case studies show that the snowball-neighbourhood strategy can be effective, but this depends on two key factors: in addition to the requirement for a sufficient density of early-adopters on a specific street or block to pressure others into coordinating on the new norm, the behaviour also has to be visible.

Part IV: Discord, Disruption, and Discovery

Chapter 11: Optimising Innovation: Social Networks for Discovery

The concepts already introduced (such as bridge width, relevance, reinforcement, and social clustering) can be applied in an organisational context to boost creative capacity.

Generally, innovation is aimed at producing something new yet commercially successful. Broadway musicals can be used to illustrate this. At the start of a run, there is no way of predicting which productions will be hits, and which flops, as both will (almost always) have the key ingredients for success. Despite this, there is a huge gap between the failures and successes.

Although marketers have tried to understand this for decades, it is only recently that network science-based studies have begun to shed light on this. The findings show that neither the individual artists involved, or elements of the production such as songs, colour palettes, or even the themes developed are predictive. Instead, hits come about when collaborative teams balance common learning from past experiences with novel ideas generated in the course of working together. In short, they require social networks that balance coordination with creativity.

As in engineering, innovative successes arise when a team of people with complementary skills come together in a 'social cluster'. For such teams, what matters most of all is how well connected they are to their larger creative network. The most intense periods of innovation on Broadway have occurred when the network has consisted of social clusters connected by wide bridges. Each cluster could incubate a distinct style, whilst the bridges enabled experienced artists and newcomers to coordinate to explore new ideas. In contrast, firework

display networks (where everyone works with everyone, and a few people and themes dominate) have historically produced very high rates of flops. As with the initial failure to introduce hybrid corn, the status quo had a reasonable track record and was firmly entrenched, with innovators facing overwhelming countervailing influences, forcing them to conform.

Applying this to small-scale teams challenges the conventional management wisdom around optimising for informational flows through weak ties, weekly meetings, regular check-ins, and high-contact spaces. The intention is that everyone is kept in the loop on the latest breakthroughs at all times. These strategies may work for solving simple problems, but not for complex ones.

This is illustrated by research on the 2006 Netflix Prize, a competition to find the best film predictor algorithm, and subsequent public competitions for solving difficult problems. By designing different team network structures, different approaches to innovation management were tested. In the eight fireworks network teams, everyone was connected to everyone to optimise for immediate sharing of discoveries. In the eight fishing net teams, however, members could only see solutions produced by the few people they were immediately connected to, with solutions originating further away required to travel across wide bridges to get to them.

In a data set with more than 15 000 possible solutions, the fireworks teams started off well, with 'good' solutions reaching everyone almost as soon as they were found. The teams then coalesced on these solutions and continued to improve them, but at the expense of wider exploration of the solution landscape. By all looking at the problem in the same way, they missed out on the chance to find the best solutions. In contrast, the clustered teams preserved a diversity of approaches by limiting the spread of information, protecting researchers long enough to do truly innovative things. The result was that each and every fishing net team found a solution that was better than the best one produced by any of the fireworks teams.

This can be explained by observing that familiar solutions are easy to understand and spread, and therefore behave as simple contagions that - in a fireworks network - easily outcompete unfamiliar ideas that must spread in a complex manner.

When AI algorithms were applied to the same problem, they outperformed the fireworks human teams, but typically not the fishing net teams. They had the same problem as the fireworks teams of failing to find globally optimum solutions by prematurely latching onto single, strongly-predictive variables at the expense of exploring more unusual combinations.

This demonstrates that well-designed teams are characterised by coordination across diversity of approach. This enables them to both explore unusual ideas and reinforce promising innovations, something that - for complex problems - is not best supported by maximising information flow.

Another illustration of this is the apparent paradox that although China was far ahead of Europe in developing and spreading innovations for most of history, it was relatively backward European states which ended up dominating much of the world. However, China

had a highly centralised pattern of informational and political networks, which - although they enabled innovations to spread extremely rapidly at first, and gave China a head start - also resulted in a long-term tendency to stamp on radical changes. By contrast, Europe was a patchwork of comparatively small states, much more like a fishing net. Innovations were frequently impeded by national borders, which - although slowing down development - also enabled unique cultures and approaches to develop. In the long run, European states were able to invent - and adapt from outside sources, such as China - things that the others could not have imagined, and that the Chinese state would not have allowed.

Chapter 12: Bias, Belief, and the Willingness to Change

As well as the discovery of new ideas, their acceptance is also influenced by whether a network structure supports contagion. This is particularly true when groups are predisposed towards some ideas rather than others.

'Motivated reasoning' is the term that refers to the tendency to interpret neutral information in the light of existing preferences and biases. This can skew processing of facts around issues such as climate change, an effect that can be amplified by social network structure, which effectively acts as a filter on how new information is interpreted. 'Echo chambers' or 'silos' are homogenous and polarised spaces which reinforce partisan bias through exposing people to repeated interactions with like-minded peers.

However, simply increasing bridge width across echo chambers to enable discussion is not necessarily enough to enable people to process new information on issues such as climate change more objectively. When bias is present, framing also matters - if people are put in a setting where they are reminded of their prior political commitments (even just by party logos or memes), then they typically fail to learn anything or revise their views, as in- and out-group identities strongly determine perceived relevance and therefore influence. In social change, it is therefore important not only to craft a compelling message and to have it supported by the right contagion infrastructure, but also to pay attention to the framing effects that determine what people actually hear; diverse participants must be made to seem relevant to each other.

Finally, network centralisation also determines the effects of bias. In highly centralised networks, the 'influencers' introduced at the start of the book are highly connected to everyone else, and have outsized influence on the flow of information when compared to people on the periphery. Fishing net networks are egalitarian, with everyone effectively part of a large but interconnected periphery and able to have their voice heard. In these networks, ideas flow between small clusters across wide bridges, with new ideas able to emerge and spread from anywhere without being silenced by a powerful, central social star.

Not surprisingly, centralised networks tend to allow the people at the centre to spread their biases across entire populations. Combined with cultural, political, or social homogeneity, this is a particular problem - ideas that reinforce existing beliefs spread as simple contagions, which a centralised network pattern is perfectly suited to support. In echo chambers, influencers can very easily spread misleading information that plays to

established prejudices. Challenging ideas, as complex contagions, will not get a look in, as they typically emerge at the egalitarian, moderately-connected network periphery (away from the overwhelming countervailing influences faced by those at the centre). The influence of influencers therefore has a natural asymmetry - they can spread simple contagions, but not complex ones, and this is particularly important when it comes to the spread of misinformation. Bias is not something that exists only in individuals' heads, but also in the social network structure they are a part of.

These findings apply even in settings where objective evidence is supposed to be the basis for decision-making, such as medicine. In fact, doctors may be even more susceptible to bias, as high-stakes decisions in uncertain situations make unpopular decisions even riskier, making conformity more likely. One study showed that in the case of a new cancer treatment, neither the nature of patients' illness, the doctors' background and experience, or money explained differences in whether the treatment was used or not. It was in fact determined by reinforcement within the doctors' social network - whether it had gained the critical mass that made it socially legitimate.

Given the right network structures, relying on peer assessments can be effective when faced with uncertainty. However, in the wrong network structures, bias can be self-reinforcing - as correlations between wealth, diet, race, gender etc. become known in the medical community, they can become accepted as inevitable, with the resulting social norms used as the basis for discrimination in choice of treatments.

These biases can be hard to remove because they are often unconscious, or *implicit* in the professional norms and reinforcing social networks of individual doctors. Encouraging them to reflect on their decisions before making final treatment recommendations failed to decrease bias. However, egalitarian network structures for exchanging opinions can have incredibly powerful effects - one study showed that it was capable of eliminating the factor of two difference in referrals for emergency treatment between white men and black women, and made doctors eleven times more likely to switch from under-treatment to providing the correct care than those in the control group. This result was replicated over one hundred times in different clinical cases.

Furthermore, clinicians who scored highly in some quizzes did not do well on others. In a centralised network, an authoritative person is likely to be influential even outside their particular area of expertise. Countering this requires an active effort to bring in voices from the network periphery.

Professional sports scouting shows the same tendency as medicine to be a hierarchical old boys' club, with entrenched norms favouring certain kinds of players at the expense of others, even if the outcomes don't justify it. Although modern data science approaches are widespread in professional sports, there is a lot of tacit knowledge that is omitted as unimportant. Empowering coaches at the network periphery to have their voices heard without being dominated by higher-ranking coaches significantly improved predictions of players' performances, even in a very short period.

Biases cause us to favour familiar answers over correct ones, even when this is expensive. This tendency is reinforced by centralised networks, which further entrench the biases and

allow ideas that conform with them to spread as simple contagions. This makes it very hard to interpret new information accurately, much less solve complex, challenging problems. Across a variety of contexts, the network periphery is needed to spark and support meaningful social change.