Analytical Sociology: Norms, Actions and Networks

Chapter 7: Online Networks and the Diffusion of Protest

Sandra González-Bailón (University of Oxford)
Javier Borge-Holthoefer (University of Zaragoza)
Yamir Moreno (University of Zaragoza)

Keywords: collective action, political protests, diffusion, networks, thresholds, social influence, complex contagion, Twitter

Subjects: online networks, diffusion models, threshold models, social influence

Further readings:

EASLEY, D. & KLEINBERG, J. 2010. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*, New York, NY, Cambridge University Press.
KATZ, E. & LAZARSFELD, P. 1955. *Personal Influence. The Part Played by People in the Flow of Mass Communications*, New York, NY, Free Press.
ROGERS, E. M. 2003. *Diffusion of Innovations*, New York, NY, Free Press.
SCHELLING, T. C. 1978. *Micromotives and Macrobehavior*, London, Norton.
WATTS, D. J. 2011. *Everything is Obvious.*Once you know the answer.*, New York, NY, Crown Business.

One of the programmatic aims of analytical sociology is to uncover the individual-level mechanisms that generate aggregated patterns of behaviour (Hedström, 2005, Hedström and Bearman, 2009). The connection between these two levels of analysis, often referred to as the micro-macro link, is characterised by the complexity and non-linearity that arises from interdependence, that is, from the influence that actors exert on each other when taking a course of action (Schelling, 1978). Networks capture the structure of that interdependence, opening the channels through which threshold effects, tipping points, cumulative causality and path dependency take place (Hedström, 2005: 99); in other words, networks help trigger the chain reactions that transform individual decisions and actions into the collective dynamics we observe on the aggregate. For this reason, networks have been central in the analysis of a wide range of phenomena (Watts, 1999, Watts, 2004), including social movements and the dynamics of protest diffusion.

Collective action and the emergence of social movements offer an ideal empirical setting to analyse how networks mediate the link between individuals and groups, and test theories that are central to the analytical sociology tradition (Baldassarri, 2010).
The increasing availability of digital data, based on real-time transactions between people, provides new empirical evidence with which to assess theories that so far had to deal with weak proxies to the structure of human interaction. This chapter makes use of that sort of data to test some of the mechanisms long theorised in the study of collective action. We use one instance of digitally coordinated protest – the Spanish Occupy movement as it emerged in May 2011 (see Borge-Holthoefer et al., 2011, González-Bailón et al., 2011) – to illustrate how interdependence and social influence drive the emergence and growth of collective action. We use this case to analyse the way in which individual communication patterns concatenate in complex networks of interaction that can ultimately lead to an explosion of activity on the aggregate level, hence boosting the global visibility of social movements and political protests.

The chapter is organised as follows. We start our argument at the collective level: we review previous research on the dynamics of protest diffusion, and we show the aggregated patterns of activity and recruitment that characterise our case. We then move to the individual level to discuss threshold mechanisms and the way in which they operate, both in theory and in practice, according to our data. We finally shift focus to the networks of interdependence that act as a bridge between the individual-level mechanisms and the aggregated patterns, and we provide some evidence of how the local contexts that networks create change the dynamics of social influence. We assess future lines of research in the concluding section and outline ways in which digital data can help us push forward the analytical sociology programme.

1. Diffusion Dynamics

Models of diffusion see society as a collective learning system where new practices and behaviour spread through interpersonal networks (Rogers, 2003). Most research has focused on the diffusion of innovations and on how communication networks allow assessing the risk involved in adopting a new product: early adopters help overcome uncertainty by setting up "community laboratories" from which the late majority gain information and experience prior to their adoption (Ryan and Gross, 1943: 18). In sociological research, Coleman, Katz and Menzel's study on the diffusion of a new drug among physicians is one of the most influential and replicated: they showed that communication ties to colleagues who had used the drug had a positive impact on the decision to prescribe it, especially during the first months of the drug's availability, when there was greater uncertainty (Coleman et al., 1957). The interpersonal networks amongst physicians triggered a snowball process in which those who had introduced the drug passed on the innovation to their colleagues, allowing them to learn from each other instead of learning from journals or other exogenous sources.

Diffusion models, however, go beyond the spread of innovations to apply to a wider range of examples that include fashion, voting, music sales, or the popularity of restaurants (Easley and Kleinberg, 2010: chapter 19). These instances of collective behaviour also exhibit the cascading effects of interdependent decision making, expressed in the form of network externalities and herding behaviour: the first refer to situations where there is a direct benefit from aligning with the actions of others, the second to situations where there is a tendency to make decisions based on inferences from what earlier people have done. Often, these two effects coexist, as it is usual in the adoption of one technological standard over another: people benefit indirectly from previous adoptions because they can draw relevant information about the technology from those actions; but they can also benefit directly if by adopting the same standard they can enjoy compatibility with a larger number of people (Easley and Kleinberg, 2010: 426-427). Collective action falls in this wider range of examples: when the decision to participate in a collective effort spreads through a population, the dynamics often reproduce the features typical of cascades and network externalities.

The emergence of collective action has often been theorised as a social dilemma that requires rational actors to solve the free rider problem, usually by means of selective incentives or the enforcement of sanctions in small groups (Olson, 1965). However, when interdependence is taken into account, and the decision to join a collective effort is seen through the lens of political efficacy, the costs and benefits associated to a contribution to the public good shift with the number of others that are already contributing (Gould, 1993, Oliver et al., 1985). Sequential decision making and the way in which social influence spreads through networks hold the key to these dynamics: by looking at what other people decided to do before, actors can better determine their decision to join; and this triggers a chain reaction the size of which will depend on how people monitor each other's behaviour through their personal networks (Marwell and Oliver, 1993). The success or failure of a collective mobilisation – or of emergent fads, new books and restaurants – depends on whether those activation cascades reach a critical mass of people. Reaching a critical number matters because only then the diffusion process becomes self-sustaining (Schelling, 1978: chapter 3); when a critical mass is not reached we have an instance of failed diffusion or, in the collective action context, failed mobilisation.

Studies of diffusion have generally tried to identify why some practices spread while others languish (Strang and Soule, 1998). They have usually focused on the attributes of early and late adopters, the characteristics of the thing being diffused, and the relative influence on adoption of network peers compared to exogenous sources of information like mass media (for a review, see Rogers, 2003: chapter 2). According to these studies, early adopters tend to be better connected and exhibit greater centrality in their networks. The attributes of the thing or practice being diffused also matter: compared to viral products, for instance, “political mobilisation moves more sluggishly, needing to gain momentum within neighbourhoods and small communities” (Easley and Kleinberg, 2010: 511). This illustrates why networks are so important for diffusion processes: they help identify community boundaries and the ties that bridge across them, facilitating or hampering the global diffusion of, in this case, political mobilisation.

Disentangling the effects of social influence in networks from the effects of exogenous factors is one of the biggest challenges in diffusion research (Aral, 2011). One of the studies reanalysing Coleman et al.’s data found that adoptions that were taken as the outcome of contagion were in fact the effect of an aggressive marketing campaign, external to the network tracked in the original study (van den Bulte and Lilien, 2001). What the study concluded, in other words, is that common exposure to the same global source of information, rather than local influence or contagion through peer networks, explained the adoption rates of the new drug in this particular community. Even when the dynamics of diffusion are assumed endogenous to a social system determining the individual-level mechanism that drives the process is not straightforward. Contagion, influence or learning point to different behavioural drivers and leave a distinctive imprint on diffusion curves and their acceleration rates, which can be used to extrapolate hypotheses about the underlying mechanisms (Young, 2009); however, once social networks are taken into account – as opposed to just assuming random encounters between individuals – the effects of those mechanisms on the aggregated dynamics are increasingly more complex (for a review of the literature, see DiMaggio and Garip, 2012).

Event history methods have long been proposed to shift the focus of analysis to the micro-level, especially to the effects that the possibility of interpersonal communication between adopters and non-adopters have on diffusion. Researchers of social movements have often applied this framework to the study of protest diffusion (Andrews and Biggs, 2006, Hedström, 1994, Hedström et al., 2000, Myers, 2000), analysing the relative effects of social influence in diffusion patterns controlling for demographic and population density factors. This research assumes that communication networks open the pathways that encourage and help sustain protest waves: the social influence hypothesis requires an internal feedback loop connecting earlier and later participants, the type of communication that networks facilitate. However, they differ in the proxies used to capture communication network and in their findings. There is positive evidence that spatial proximity, and the networks that such proximity breeds, helped trade unions grow (Hedström, 1994), and that the network formed by the travel routes of political agitators boosted the diffusion of political parties (Hedström et al., 2000). The networks related to mass media distribution have also been associated to the diffusion of collective violence, in this case, the racial rioting that spread in the US in the sixties (Myers, 2000); however, the wave of sit-ins that also spread around the same time did not seem to be driven by communication networks, as captured by affiliation to the same college associations: mass media was a more significant factor in spreading the news of these protests (Andrews and Biggs, 2006). The question these studies leave unanswered is how good the networks they analyse are as an approximation to actual communication (as opposed to just the potential to communicate), and how many of the reported effects are an artefact of the way in which networks are measured.

Digitally-enabled protests are facilitating the type of network data that was difficult to obtain before. The case study considered here – the protests that erupted in Spain in May

2011 as a reaction to the financial crisis and cuts to public spending – tracks the exchange of protest messages in an online network, Twitter, which allows us to reconstruct broadcasting channels of communication (the following-follower network) and how these were activated in the particular context of the protests. The data was collected as follows. The stream of information related to the mobilisations was captured using a list of relevant hashtags, which are user-generated labels that summarise the content of the messages sent and flag them as part of a stream of specific information exchange. Most of these labels echoed the motto of the protests or the main mobilisation date (like '15M' or 'real democracy now', the name of the online-born platform coordinating the campaign). The exchange of these messages, and the number of people sending them, were tracked for the period of one month, with the first big demonstration day at the centre of the observation window (for a full description of the data see Borge-Holthoefer et al., 2011 and González-Bailón et al., 2011).

-- Figure 1 about here –

Figure 1 presents aggregated data around the protests for the observed period. Panels A and B display levels of activity in terms of number of users that were active in the exchange (N ~ 88,000) as frequencies (A) and cumulative numbers (B); panels C and D show activity levels measured as the number of protest messages sent (N ~ 500,000), again in frequency and cumulative form. The cumulative curves show the S-shape typical of diffusion processes, entering the exponential growth phase when around 10% of the population was active, and approximating saturation towards the end of the period. Users underlying this curve can be classified using the same categories applied in other diffusion studies: innovators (or leaders), early adopters, early majority, late majority and laggards. This classification is useful because it tells us who are the actors with higher intrinsic motivation and who need more information from their social circles to be motivated to join. Individuals that are part of the same group tend to have many attributes in common (Rogers, 2003: chapter 7), which adds a richer sociological dimension to the classification. These categories of adopters are based on chronological time of activation; the question, under the social influence hypothesis, is how early participants managed to lead late participants into the protest.

Traditional models of diffusion that focus on the aggregated level have usually assumed that all members of the population are equally susceptible, that contacts between all pairs of spreaders-potential adopters are equally likely and equally contagious, and that the rate and contagiousness of contacts do not vary with time (Strang and Tuma, 1993: 617). These assumptions are very unlikely for most examples of empirical diffusion, but particularly so in the context of political mobilisations like the example considered here: these mobilisations are characterised by sudden peaks of activity that draw from the cumulative effects of feedback dynamics (see Biggs, 2005). The online trails left by the Spanish mobilisation allow us to zoom into individual level behaviour and open the black box of those feedback effects. The data allow us to reconstruct the network connecting protesters and track patterns of communication over time, which in turn allows us to measure the relative exposure of users to protest information as time went by. Using this information, we can capture the propensity to join of each individual by looking at how many others had joined before they decided to follow; in other words, we can offer the continuous equivalent of the five user categories drawn from diffusion studies. The following section considers the dynamics of these individual activations before the effects of local networks are taken into account.

2. Thresholds and Critical Mass

As the previous section showed, diffusion studies define critical mass as the number of people that need to join the adoption curve before the diffusion becomes self-sustaining. The notion of ‘threshold’ appeals to the same idea, but at the individual level: it is the critical number of people that need to have adopted a behaviour before an actor decides to adopt as well (Granovetter, 1978b, Schelling, 1978). Thresholds can be defined as proportions or absolute numbers. For some behaviour, like adopting a language or following a new fashion trend, it is proportions that influence people, for instance in the form of a majority rule; whereas for other types of behaviour, like attendance to a discussion group or tournament, what matters are absolute numbers: these activities require a minimum number of participants to take place, and they usually benefit from network externalities (DiMaggio and Garip, 2012; see also Schelling, 1978: chapter 3). Regardless of how the group of reference is operationalised, in both cases individuals monitor the actions of others before deciding which course of action to take.

Thresholds describe the variation among individuals in their propensity to do something: they focus on differences in individual preferences rather than on where those preferences come from (Granovetter, 1978a). In the context of collective action, some people decide to contribute only when a significant number of others are already participating; in this sense, thresholds can also be seen as the degree of responsiveness to social influence (Young, 2009: 1905): low threshold actors are first movers, and high threshold actors are followers who will not act without their peers. The way in which these thresholds are measured, however, makes implicit assumptions about the mechanisms driving individual activation. Measuring thresholds with absolute numbers assumes that actors are driven by notions of political efficacy: the larger the number of protesters, the higher the chances of success and the likelier that actors will overcome their resistance to participate. Measuring thresholds with relative numbers, on the other hand, implies that actors are susceptible to normative behaviour, that is, to the preponderance of a given course of action relative to a population. In this latter case, both adopters and non-adopters are influential, whereas in the former, only adopters have an impact on individual decisions. This difference is subtle, but simulation studies show that it affects the dynamics of diffusion (Centola and Macy, 2007: 711). Which of these mechanisms offers a more accurate description of individual actions is ultimately an empirical question.

The way in which thresholds distribute in the population is also crucial to understand how behaviour cascades: the domino effect underlying diffusion processes depends on the relative weight of first movers and followers. However, thresholds have been difficult to establish empirically because of lack of appropriate data. Most threshold studies are based on simulation models that make assumptions about threshold distributions (Watts and Dodds, 2010). The standard has been to assume that they are normally distributed, in line with the also conventional assumption in diffusion literature on how adopters distribute between the extremes of ‘innovators’ and ‘laggards’: most of them fall in the intermediate area of early and late majority (Rogers, 2003: chapter 7). It is also an empirical matter to determine if this assumption is realistic or if the way in which thresholds distribute changes from population to population or with the type of behaviour being spread. To the extent that chain reactions depend on these distributions, establishing their empirical shape is a crucial step in the explanation of diffusion dynamics.

The Spanish protest data allow us to infer the threshold distribution empirically using the time of activation or users, that is, the moment when they start sending protest messages. Individual thresholds are defined as the number of other users that had already sent a protest message when a user gets activated. Users that sent messages when none or a small number of others had done so are the leaders and early participants in the movement: whatever their reasons to activate the chain reaction, their early behaviour reveals intrinsic preferences to pioneer the protest. Users requiring a higher number of other users participating before joining the protest, on the other hand, are the late majority and laggards, and display higher susceptibility to social influence and to the actions of other people: they only become active when they register that many other users are already protesting. The empirical distribution of these thresholds is summarised in figure 2, where thresholds are measured both as numbers (panel A) and proportions (panel B).

-- Figure 2 about here –

The distribution of thresholds as absolute numbers is bounded by the size of the population we analyse here: there are approximately 90 thousand users messaging about the protests and this gives the upper bound for the threshold range. Users approximating this upper bound joined only when the vast majority of users in the network had already joined the protest; and vice versa: users approximating the lower bound got activated when only a small number of other users were sending messages. The histogram with proportions gives us the same information, but this time in relative terms: for each value in the horizontal axis (binned in 0.05 intervals in the upper subpanel, and as a logarithmic transformation in the lower subpanel), the calculation of thresholds takes into account the influence of both adopters and non-adopters.

These thresholds capture a global measure of social influence. While it is unrealistic to assume that a direct communication channel exists between all these users, they can still monitor general levels of activity by other means. One of these are trending topics, which is the Twitter way of telling users which issues are hot or most salient in the online exchange of information. Trending topics depend on the number of people using certain hashtags and the number of messages broadcasted with those labels. As such, they give a rough but very visible measure of how many users are actively participating in the exchange of protest information. Another means of global social influence – particularly after the first big demonstration day identified as 15-M in figure 1 – is common exposure to mainstream media: most news reporting of the events highlighted the use of social network sites like Twitter to coordinate the protests. These are the global sources of social influence that the distributions of figure 2 capture: although these thresholds do not discriminate which channel is more relevant (global information obtained from Twitter or from mainstream media), they measure how the overall volume of protesters helped trigger the decision of still inactive users to start sending protest messages as well.

Figure 3 shows the same information, this time highlighting the nonlinear effects and cumulative causality triggered by the activation of thresholds. The figure plots the accumulated number of activations in time t as a function of the accumulated activations in $t - 1$, where t is defined as days. Given the way in which thresholds are distributed in this population of users, the chain reaction starts when users with threshold up to 0.1 get activated: a day later, twice as many are active, and this starts the bandwagon effect that ultimately reaches the entire population. The curve illustrates how the minority of users whose behaviour does not depend on numbers start the diffusion process: their contributions encourage subsequent contributions, which encourage other contributions, and so on. These dynamics reflect the two fundamental properties of collective action, namely that decision making is interdependent and that it takes place in a sequential fashion (Gould, 1993, Oliver et al., 1985, Macy, 1991). This is why the distribution of thresholds is so important to understand aggregated dynamics: an absence of a large enough number of users with low to intermediate thresholds (the early adopters and early majority of diffusion studies) would have acted as a firewall and the chain reaction would have stopped before reaching the full population.

-- Figure 3 about here --

The way in which the distribution of thresholds shapes chain reactions underscores the importance of individual heterogeneity in modelling collective action: not everybody is equally susceptible to the behaviour of others, and this variability drives the domino effect that underlies the diffusion of protest: at each stage of the process, individuals with higher resistance are increasingly mobilised because they register, sequentially, that their thresholds are being satisfied. The effects of cumulative causality, and the exponential growth that results from it, means that in certain periods behaviour will spread faster, creating bursts of activity. This process, however, is also characterised by a second factor that adds another source of heterogeneity to the dynamics: social structure, or the way in which people are connected or related to each other (Granovetter, 1978a: 1429). Structure matters because it mediates individual activations: much in the same way as not everybody is equally susceptible to the behaviour of others, not everybody inhabits the same local contexts, or is exposed to the same information. The assumption made so far is that social influence works globally; but heterogeneity also arises from individual interactions, and from differences in the group of reference that each user monitors. This structure affects not only the information that individuals access but also the paths that chain reactions will follow. This source of heterogeneity and its effects on diffusion dynamics are considered in the next section.

3. Networks and Social influence

Networks play a key role in diffusion processes because they facilitate threshold activation at the local level. Individual actors are not always able to monitor accurately the behaviour of everyone else (as system-level thresholds assume) or they might be more responsive or easier to be influenced by those they feel closer to, represented in their personal networks. Networks establish direct communication links and determine the exposure to the behaviour of local neighbours, creating for each individual a different group of reference (Valente, 1996). The way in which networks mediate diffusion according to this local approach is different from what the diffusion studies explored above assume: they incorporate networks to the event history framework as approximations to global channels of communication through which information spreads, so they cannot capture the nuanced mechanisms of individual exposure and how this affects activation times. Two individuals with the same threshold but connected to local networks of different size and composition might be activated, and contribute to the diffusion, at different stages: it is not only that their networks change; the threshold distribution of their neighbours is also likely to be different.

The heterogeneity contained in local networks makes the activation of chain reactions a complex process. Even when actors have the same thresholds, those with larger personal networks will register a critical number later than those with smaller networks (Watts and Dodds, 2010: 486). This matters not only because it defines the tempo of local activations, but also because the domino effect will escalate to global proportions only if it manages to reach the connections that bridge local communities or clusters. These connections, called weak ties because they tend to link socially distant people (Granovetter, 1973) and span structural holes (Burt, 1992), facilitate the diffusion of activations beyond the local personal networks where they start. A gap in the threshold distribution can stop a chain reaction, but structural holes in a network (or the lack of connections between local clusters) can also stop global diffusion. Simulation results show that chain reactions require this type of structural bridges linking socially distant actors (Macy, 1991), and also that, under certain conditions, those bridges need to be structurally wide, that is, contain multiple ties of repeated interactions to facilitate diffusion (Centola and Macy, 2007). Simulation models also show that having a small group of highly connected actors in the network helps tip the system into a critical mass of activations (Marwell and Prahl, 1988); highly connected actors are, by virtue of their centrality, more likely to span structural holes.

Networks matter, then, not only because they shape influence on the local level but also because they open the diffusion trails that help chain reactions go global. The empirical analysis of these dynamics requires having access to the activation time of each actor as a time series, and to the configuration of their personal networks when activation takes place. Figure 4 illustrates the way in which individual thresholds can be inferred from the Spanish protests data. In this example, the threshold of a focal user (the black node in the centre) is defined as the proportion of users activated (k_a) over the total number of neighbours (k_{in}), which is the group of reference that influences the decision to join the protests. The first two example networks on the top have the same size, but the focal users differ in the critical number of activations they need to register to become activated: the user on the left needs less local pressure than the user on the right. When thresholds are defined as proportions, sometimes two users will have the same threshold even when they are surrounded by a different number of activations, as the last two example networks show: in both cases, the focal user has a threshold of 0.5, that is, they will not send their first protest message until at least 50% of their neighbours have already been activated; however, this percentage requires two previous activations for the user on the right, but four for the user on the left; all else equal, the activation of this second user will take longer to happen.

-- Figure 4 about here –

The main difference of this definition compared to the global definition presented in the previous section is that now every user has a different group of reference they monitor to determine their own critical number. The empirical distribution of these local thresholds is displayed in figure 5, again defined as numbers (panel A) and proportions (panel B). Compared with the previous distributions of global thresholds, this time there is greater heterogeneity: the upper bound of the distribution for thresholds as numbers is determined by the size of local networks; since the degree distribution of these networks has a very heavy tail (Borge-Holthoefer et al., 2011), it is not surprising to find that a few number of users needed the activation of a disproportionate number of neighbours before sending their first message. The relative definition of thresholds smoothes these differences, this time resembling more closely the normal distribution often assumed in simulation models; compared to the thresholds defined globally (figure 2, panel B) this local definition classifies more actors as mid and late adopters.

-- Figure 5 about here –

The effects of global and local information on adoption rates have already been explored in other online contexts (Onnela and Reed-Tsochas, 2010). The relative weight of global and local influence on the decision to adopt reflects different personal preferences and population heterogeneity. Figure 6 illustrates how local networks act as a source of individual heterogeneity. It plots the increase in the percentage of local activations at time t – the moment when a focal user sends the first protest message – compared to the previous time period t – 1 (where is t again measured as days). Panel A shows that most users got activated after an increase of 10 to 20% in the number of neighbours already active; only a very small number of users reacted after a sudden change in the configuration of their personal networks (i.e. a difference greater than 0.5 percentage points in the period of one day). Panel B shows that users classified in the intermediate values of the threshold distribution required higher increases in the number of activations in their personal networks during the day leading to their decision to join. The dispersion captured by the boxplots reflects the heterogeneity that different local networks add to the activation of same–threshold individuals.

-- Figure 6 about here –

Users with a different propensity to be active respond to a different critical number but, as Figure 6 shows, users with the same propensity or threshold might also react at different times because they inhabit different networks. This refers back to the importance of heterogeneity in network connections identified by simulations of critical mass in collective action (Marwell and Oliver, 1993). The data from the Spanish protests suggest that heterogeneity does not depend just on personal preferences but on how they concatenate with the preferences of the other actors present in personal networks. An important aspect of this dynamic is what has been called ‘complex contagion’, that is, the need to receive stimuli from multiple sources prior to the adoption of a given behaviour (Centola and Macy, 2007). Global influence offers a unique signal that people can interpret to determine if their critical number has been reached: the aggregate number of people doing something, or an approximation to it (like Twitter trends), helps determine if the volume of contributions is large enough to make one's contribution worthwhile. However, when the main diffusion mechanism is the network of personal contacts or direct communication, the question arises of how often and from how many contacts an actor needs to receive signals from before activation. Our inference of local thresholds factors in the number of active sources to which users are exposed when they send their first protest message; and the empirical distributions that result from this inference suggest that complex contagion (or multiple reinforcement) might be more effective for some users than for others.

Figure 7 offers additional evidence of the heterogeneity that local networks introduce in the activation process. Panel A plots the chronological day of activation against global thresholds, and panel B against local thresholds. The assumption made in the case of global information is that all users react to the same number of previous activations, which means that users with the same threshold will be activated around the same time. In the local case, however, each user is exposed to different information – this explains the variance for same-threshold users. Although the trend suggests that users with higher thresholds take, on average, longer to get activated, any given day will see the activation of people with different thresholds. This is the main difference with the scenario that assumes global information: there is no straight relationship between thresholds and the chronological stage at which individuals will join the diffusion curve. This is why it is so important to analyse how local networks mediate political behaviour and the emergence of collective action: it makes it impossible to predict when a given actor will be activated if we only have information about their intrinsic inclination. The trigger will ultimately depend on their neighbours (and on the neighbours of their neighbours, and so on).

-- Figure 7 about here –

4. Conclusion: Digital Data and Analytical Sociology

This chapter has illustrated some of the ways in which digital data can help test models that so far relied mostly on simulations. The theoretical framework of our exercise is rooted in diffusion studies and in efforts to identify the individual level mechanisms that operate underneath aggregated adoption curves. Diffusion models offer a more appropriate lens to understand the emergence of collective action than rational action models that do not take into account the sequential nature of decision making and the cascading effects that arise from interdependence. Threshold models of social influence have been explored using data collected around a case of spontaneous political protest. The digital traces left by participants in the protest have been used to reconstruct the channels through which information about the protest flowed, and analyse how that communication network helped trigger a chain reaction of activations. The richness of online data, and in particular, the ability to capture activation rates in local networks as a continuous time series, allowed us to infer the threshold distribution, using both global and local definitions, and absolute and relative numbers. This allowed us to assess the heterogeneity of same-threshold individuals in terms of chronological activation times, shedding empirical light into the effects of interdependence in decision-making. This heterogeneity is partly explained by the effects of network size (which mediates exposure and the registration of a critical number) and by the effects of complex contagion (which requires diverse sources of exposure, but is likely to differ amongst actors).

In their classic diffusion study, Coleman and co-authors concluded with a methodological note in which they emphasised the importance of explicitly taking into account "the structuring of single persons into larger units" (Coleman et al., 1957: 269). As they anticipated, the analysis of interpersonal structures of communication has become increasingly relevant in sociological research, a process facilitated by the increasing availability of digital data. The potential of such data to revolutionise our understanding of social behaviour has been acknowledged for a few years now (Lazer et al., 2009, Watts, 2007). This potential derives from the richness and higher quality of the data, but also from the collaborative nature of the research that such data promotes, which transcends old disciplinary boundaries and results in more creative research questions and designs (good examples can be found in Easley and Kleinberg, 2010 and Watts, 2011). The richness of online data also allows sociologists to be better positioned to explore one of the big questions of social science research: how to go from individual actions to collective outcomes, or why the whole cannot be reduced to any of its parts (Coleman, 1990, Schelling, 1978). This question has been difficult to tackle so far because of lack of appropriate empirical data and the complexity of the connection between the individual and the collective; this is precisely what the analysis of digital trails allows us to understand better.

Using online data to advance sociological research will, however, still require solving important challenges; for instance, how to best model continuous large networks and the way in which they co-evolve with individual behaviour (Snijders, 2011); how to compare the long studied properties of networks measured as snapshots with their evolving counterparts, easier to get from online interactions (Kossinets and Watts, 2009); or how to control for self-selection in the study of social influence (Aral, 2011; Aral, Muchnik and Sundararajan 2009) when demographic information is not available or reliable, as it often happens with online data. One of the greatest challenges of future research on online diffusion, in the context of collective action or in other empirical settings, is how to discriminate between the competing forces of exogenous and endogenous influence, or between the effects of mass media and interpersonal communication (Katz and Lazarsfeld, 1955). Added to that problem is the multiplex nature of networks; that is, the fact that people are simultaneously embedded and interacting in several networks at a time, creating parallel layers of interaction that are connected through feedback effects. None of these issues is specific to online data, but access to that data can help build better models and push forward the frontiers of social science research.

In spite of the difficulties associated to the analysis of online data, their analysis promises to shed new light into the complexities of human interactions. This is particularly important in the study of collective action and social mobilisations, not only because classic studies have often made virtue out of necessity in dealing with imperfect network data, but also because recent events have given rise to many unfounded claims about how online networks facilitate the explosion of political protest (i.e. Andersen, 2011). The models and theories that have been developed for decades within the analytical sociology tradition, in conjunction with the better data that online interactions make available, can help build stronger theories and make better sense of social change and the forces that make it happen.

References

ANDERSEN, K. 2011. The Protester. *Time*, 14 December.

ANDREWS, K. T. & BIGGS, M. 2006. The Dynamics of Protest Diffusion: Movement Organisations, Social Networks, and News Media in the 1960 Sit-Ins. *American*

ARAL, S. 2011. Identifying Social Influence: A Comment on Opinion Leadership and Social Contagion in New Product Diffusion. *Marketing Science,* 30, 217-223.

ARAL, S., MUCHNIK, L. & SUNDARARAJAN, A. 2009. Distinguishing Influence-Based Contagion from Homophily-Driven Diffusion in Dynamic Networks. *PNAS,* 106**,** 21544-49.

BALDASSARRI, D. 2010. Collective Action. *In:* BEARMAN, P. & HEDSTRÖM, P. (eds.) *Handbook of Analytical Sociology.* Oxford: Oxford University Press.

BIGGS, M. 2005. Strikes as Forest Fires: Chicago and Paris in the late nineteenth century. *American Journal of Sociology,* 110, 1684-1714.

BORGE-HOLTHOEFER, J., RIVERO, A., GARCÍA, I., CAUHÉ, E., FERRER, A., FERRER, D., FRANCOS, D., IÑIGUEZ, D., PÉREZ, M. P., RUIZ, G., SANZ, F., SERRANO, F., VIÑAS, C., TARANCÓN, A. & MORENO, Y. 2011. Structural and Dynamical Patterns on Online Social Networks: the Spanish May 15th Movement as a Case Study doi:10.1371/journal.pone.0023883. *PLoS ONE* 6, 6e23883.

BURT, R. S. 1992. *Structural Holes. The Social Structure of Competition,* Cambridge, MA, Harvard University Press.

CENTOLA, D. 2010. The Spread of Behavior in an Online Social Network Experiment. *Science,* 329, 1194-97.

CENTOLA, D. & MACY, M. W. 2007. Complex Contagions and the Weakness of Long Ties. *American Journal of Sociology,* 113, 702-34.

COLEMAN, J. S. 1990. *Foundations of Social Theory,* Cambridge, Mass., Belknap Press of Harvard University Press.

COLEMAN, J. S., KATZ, E. & MENZEL, H. 1957. The diffusion of an innovation among physicians. *Sociometry,* 20, 253-270.

DIMAGGIO, P. & GARIP, F. 2012. Network Effects and Social Inequality. *Annual Review of Sociology,* DOI: 10.1146/annurev.soc.012809.102545.

EASLEY, D. & KLEINBERG, J. 2010. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World,* New York, NY, Cambridge University Press.

GONZÁLEZ BAILÓN, S., BORGE-HOLTHOEFER, J., RIVERO, A. & MORENO, Y. 2011. The Dynamics of Protest Recruitment through an Online Network. *Scientific Reports,* 1.

GOULD, R. 1993. Collective Action and Network Structure. *American Sociological Review,* 58, 182-196.

GRANOVETTER, M. 1973. The strength of weak ties. *American Journal of Sociology,* 78, 1360-1380.

GRANOVETTER, M. 1978a. Threshold Models of Collective Behavior. *American Journal of Sociology,* 83, 1420-43.

GRANOVETTER, M. 1978b. Threshold Models of Diffusion and Collective Behaviour. *Journal of Mathematical Sociology,* 9, 165-179.

HEDSTRÖM, P. 1994. Contagious collectivities: on the spatial diffusion of Swedish trade unions, 1890-1940. *American Journal of Sociology,* 99, 1157-79.

HEDSTRÖM, P. 2005. *Dissecting the Social: on the Principles of Analytical Sociology,* Cambridge, Cambridge Univesity Press.

HEDSTRÖM, P. & BEARMAN, P. 2009. What is Analytical Sociology All About? An Introductory Essay. *Handbook of Analytical Sociology* Oxford: Oxford University Press.

HEDSTRÖM, P., SANDELL, R. & STERN, C. 2000. Mesolevel Networks and the Diffusion of Social Movements: the Case of the Swedish Social Democratic Party. *American Journal of Sociology,* 106, 145-72.

KATZ, E. & LAZARSFELD, P. 1955. *Personal Influence. The Part Played by People in the Flow of Mass Communications,* New York, NY, Free Press.

KOSSINETS, G. & WATTS, D. J. 2009. Origins of Homophily in an Evolving Social Network. *American Journal of Sociology,* 115**,** 405-50.

LAZER, D., PENTLAND, A., ADAMIC, L., ARAL, S., BARABÁSI, A.-L., BREWER, D., CHRISTAKIS, N. A., CONTRACTOR, N. S., FOWLER, J. H., GUTMANN, M., JEBARA, T., KING, G., MACY, M. W., ROY, D. & VAN ALSTYNE, M. 2009. Computational Social Science. *Science*, 323, 721-723.

MACY, M. W. 1991. Chains of Cooperation: Threshold Effects in Collective Action. *American Sociological Review,* 56, 730-747.

MARWELL, G. & OLIVER, P. 1993. *The Critical Mass in Collective Action,* Cambridge, Cambridge University Press.

MARWELL, G. & PRAHL, R. 1988. Social networks and collective action. A theory of critical mass III. *American Journal of Sociology*, 94, 502-34.

MYERS, D. J. 2000. The Diffusion of Collective Violence: infectiousness, susceptibility, and mass media networks. *American Journal of Sociology,* 106, 173-208.

OLIVER, P., MARWELL, G. & TEIXEIRA, R. 1985. A theory of the critical mass. I. Interdependence, group heterogeneity, and the production of collective action. *American Journal of Sociology,* 91, 522-56.

OLSON, M. 1965. *The Logic of Collective Action: public goods and the theory of groups,* Cambridge, MA, Harvard University Press.

ONNELA, J.-P. & REED-TSOCHAS, F. 2010. Spontaneous Emergence of Social Influence in Online Systems. *Proceedings of the National Academy of Sciences,* 107, 18375-18380.

ROGERS, E. M. 2003. *Diffusion of Innovations,* New York, NY, Free Press.

ROMERO, D. M., MEEDER, B. & KLEINBERG, J. Year. Differences in the Mechanics of Information Diffusion Across Topics: Idioms, Political Hashtags, and Complex Contagion on Twitter. *In:* International World Wide Web Conference, 2011 Hyderabad, India.

RYAN, B. & GROSS, N. C. 1943. The Diffusion of Hybrid Corn in Two Iowa Communities. *Rural Sociology,* 8, 15-24.

SCHELLING, T. C. 1978. *Micromotives and Macrobehavior,* London, Norton.

SNIJDERS, T. A. B. 2011. Statistical Models for Social Networks. *Annual Review of Sociology*, 37**,** 129-151.

STRANG, D. & SOULE, S. A. 1998. Diffusion in Organizations and Social Movements: From Hybrid Corn to Poison Pills. *Annual Review of Sociology,* 24, 265-90.

STRANG, D. & TUMA, N. B. 1993. Spatial and Temporal Heterogeneity in Diffusion. *American Journal of Sociology 99,* 99, 614-39.

VALENTE, T. W. 1996. Social network thresholds in the diffusion of innovations. *Social Networks*, 18, 69-89.

VAN DEN BULTE, C. & LILIEN, G. L. 2001. Medical Innovation Revisited: Social Contagion versus Marketing Effort. *American Journal of Sociology,* 106, 1409-35.

WATTS, D. J. 1999. Networks, Dynamics and the Small World Phenomenon. *American Journal of Sociology*, 105, 493-527.

WATTS, D. J. 2004. The 'new' science of networks. *Annual Review of Sociology,* 30, 243-270.

WATTS, D. J. 2007. A twenty-first century science. *Nature,* 445, 489.

WATTS, D. J. 2011. *Everything is Obvious.*Once you know the answer.,* New York, NY, Crown Business.

WATTS, D. J. & DODDS, P. S. 2010. Threshold Models of Social Influence. *In:* BEARMAN, P. & HEDSTRÖM, P. (eds.) *Handbook of Analytical Sociology*. Oxford: Oxford University Press.

YOUNG, H. P. 2009. Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning. *American Economic Review,* 99, 1899-1924.

List of figures: