Kickstarting Memes and Movements: A Distribution-Based Threshold Model

Emily Robertson

August 15, 2014

Abstract

In social movements, the decision of individuals to participate depends both on their own preferences as well as the behavior of others. Three examples of this type of interaction are the Arab Spring protests, Kickstarter campaigns, and internet memes. One way of modeling this is Granovetter's threshold model of collective action, in which each agent needs to see a "threshold" level of participation before joining in. Crucial to this model is the distribution of thresholds. Small changes to the population's mean threshold, as well as the standard deviation of thresholds, lead to very different outcomes. This inquiry seeks to establish that the shape of the threshold distribution determines the ultimate success or failure of a social movement. We analyze both symmetric and skewed distributions and identify these critical parameters of success. The results provide key implications for individuals seeking to either promote or deter social movements.

1 Introduction

This inquiry seeks to establish that the shape of the threshold distribution determines the ultimate success or failure of a social movement. In social movements, individual agents must decide whether or not to participate. These decisions have collective effects: if enough agents join in, the movement is successful. However, the costs and benefits of participation change with the number of participants. Popular movements are more likely to succeed, and so may seem less risky.

One way of modeling this process is Granovetter's [1978] threshold model. In this model, each agent must see a "threshold" level of participation before joining a movement. Crucial to this process is the distribution of thresholds within a population. Small changes to either the mean threshold level or the variance in thresholds can have strong effects. This explains why some countries seemingly on the brink of social change maintain the status quo, while others break into full on revolution.

Similar themes have been explored in many areas of social science. Kuran [1989] theorizes that the apparently instant success of some movements is due to preference falsification, in which an individuals's personal preferences and public expressions often differ. Lohmann [1994] explained the East German Leipzig demonstrations as an information cascade, in which protest activity revealed previously unknown information about the regime. Ginkel and Smith [1999] applied a game theoretic approach in modeling interactions between governments, dissidents and the mob. Their approach suggests that the dissidents within a population play a key role in determining the success of the movement, both in engaging the mob and in providing information of the stability of a regime. Myatt and Wallace [2008] applied the threshold model as a coordination game, with the addition of "bad apples" who can destabilize a movement. Bhavnani and Ross [2003] explored the role of "announcements" in informing the public about the credibility of a movement. Recently, work on threshold models has focused on their application to social networks [Macy, 1991, Centola and Macy, 2007, Hedström, 1994, Borge-Holthoefer et. al., 2013].

Both Granovetter and Yin [1998] stress the importance of threshold distributions in determining the success of a movement. Yin's analysis of equilibria in various distributions suggests that strategies for those seeking to grow or shrink a movement can be very different depending on the current makeup of the population. However, this analysis focuses on a nonzero level of initial participation in determining the ultimate success of a movement. While this reveals patterns across distributions, this initial population of dissidents necessarily runs into the coordination problem described above. All movements must start somewhere, and each population may require a different proportion of instigators in order to inspire widespread participation.

This approach to modeling social movements can be applied to many seemingly unrelated situations, from historically significant revolutions to the everyday use of social media. Three scenarios which exhibit this interaction are the Arab Spring protests of 2010-2011, Kickstarter campaigns, and the spread of internet memes. In the Arab Spring, protesters faced high participation costs, including the risk of violence. However, the success of the Tunisian protests and their high visibility on social media influenced action in other countries. In Kickstarter campaigns, projects are only funded when they reach a specific goal. Potential donors can view the progress toward that goal before deciding to donate. Internet memes tend to spread primarily through social media, and the "utility" of continuing a meme includes its perceived popularity among a particular social group.

The recent Arab Spring protests are a classic example of a social movement. Of particular interest is research exploring the use of social media during the protests. Social media made protesters more visible, perhaps inspiring others to join in. Using survey data collected during the Egyptian Tahrir Square protests, Tufecki and Wilson [2012] found that social media was a major source of information regarding protest activity, and that social media use was a good indicator of participation in the movement. Lotan et. al. [2011] analyzed information flows on Twitter, and Kavanaugh et. al. [2011] compared Twitter use during the Arab Spring to that of other emergency situations.

Empirical data on Kickstarter campaigns and internet memes is more sparse, but some early work [Mollick, 2014, Kuppuswamy and Bayus, 2013, Etter, Grossglauser and Thiran, 2013, Bauckhage and Kersting, 2013] suggests that Kickstarter campaigns exhibit both free rider problems as well as bandwagon effects as they progress. Initial success is typically due to the participation of eager friends and family, followed by a lull in the middle of the campaign. As a campaign nears its goal, the rate of new donations spikes. Analysis by Mitra and Gilbert [2014] of the language used in Kickstarter campaigns found that phrases which indicated confidence and strength, such as "has pledged" and "project will be" were positive predictors of success, while weak phrases such as "hope to get" and "not been able" were negative ones.

The goal of this paper is to outline how population characteristics drive participation within a threshold model. Section 2 of this paper introduces the threshold model, and Section 3 applies this model to a population with normally distributed thresholds. My analysis focuses on the case where initial participation is zero, so that only "instigators" who wish to participate no matter what anyone else does will start the movement. We identify critical values of sigma for the population, such that movements are either always or never successful. We then compare these results to small samples of populations to examine their robustness. Section 4 applies this same analysis to right and left-skewed distributions, Section 5 is a discussion of policy implications from these results, and finally Section 6 concludes with several areas of future research.

2 The Threshold Model

In its simplest form, the threshold model involves a population of n agents, each of which faces a decision: whether or not to participate in a certain movement. In making this decision, agents consider only one factor: how many of their peers have already participated. Each agent has a "threshold" level of participation they need to see before joining in. These thresholds are randomly assigned to each agent according to some set probability distribution.

For example, an agent with a threshold of 10 would need to see 10% of their peers participate in the movement before joining. Agents with thresholds at or below 0 can be considered to always participate, no matter what their peers are doing. These are the "instigators," and no social movement can exist without at least one of these people. As Granovetter [1978] points out, the difference between a threshold of 0 and -10 may be conceptually important, but in action there is no difference. Individuals with negative thresholds may vary in the strength of their views, but they all face the same simple decision: to participate or not. Likewise, agents with thresholds of 100 or above are the societal "sticks-in-the-mud." These individuals never join the movement, no matter how many others they see joining in.

We start with some exogenously determined initial level of participation, p_0. From there, participation is determined by an iterative process, so that $p_{t+1} = P(X \leq p_t) * 100$. In other words, in the first round any agent with a threshold $\leq p_0$ joins in. This percentage of the population then becomes p_1. In the second round, agents with a threshold $\leq p_1$ join in, and this percentage of the population becomes p_2. This allows us to define an equilibrium state, in which $p_t = p_{t+1}$. At this point, no agent has incentive to wither join or leave the movement.

3 Normally distributed thresholds

The first distribution we will explore is the normal distribution, with mean μ and standard deviation σ. The normal distribution is a reasonable place to start, since this represents a population which has a general tendency toward one opinion, even though individual agents have different thresholds. Agents are randomly assigned thresholds from the normal probability density function (PDF), which is given by:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(x-\mu)^2}{2\sigma^2}} \tag{1}$$

We can see how participation evolves over time by examining the cumulative distribution function (CDF) for a given distribution, which gives the probability of an agent having a threshold at or below a certain level. The CDF for the normal distribution is:

$$F(x) = \frac{1}{2}\left[1 + erf\left(\frac{x-\mu}{\sigma\sqrt{2}}\right)\right] \tag{2}$$

3.1 Equilibrium for the normal distribution

Figure 1 shows the progression of participation for one distribution and two different choices of p_0. For the same distribution (in this case $\mu = 30$ and $\sigma = 10$), the choice of p_0 dramatically changes the outcome. When $p_0 = 30$, participation quickly climbs to 100%. However, when p_0 is set to 20, participation actually falls over time.

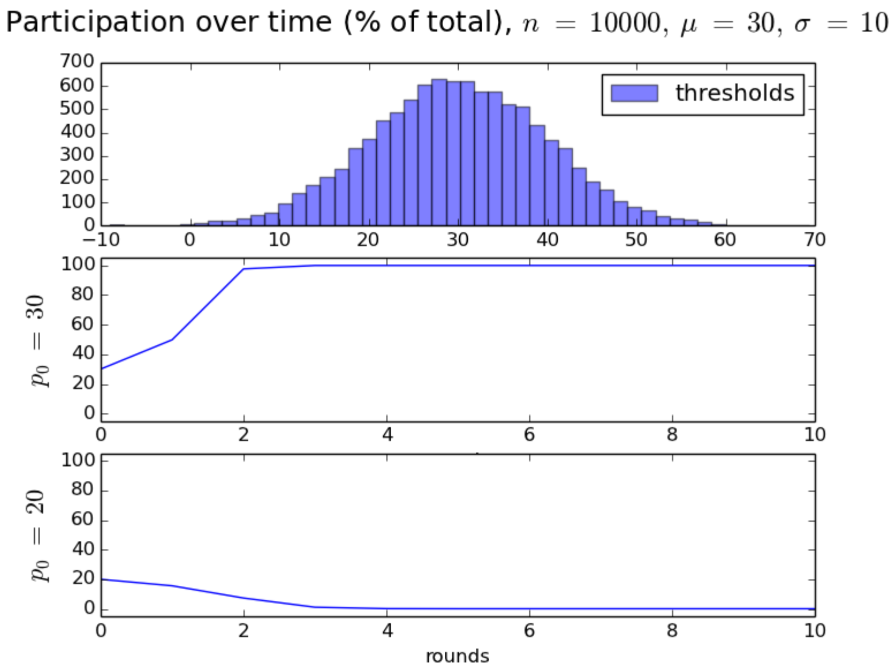

Figure 1: participation over time with $p_0 = 30$ and $p_0 = 20$

We see why this happens when we examine the CDF of this distribution. If we start at $p_0 = 30$ we are strictly on the part of the curve above the 45° line. Thus, the movement is always growing. When we start at $p_0 = 20$, we find that only about 17% of the population has a threshold of 20 or less. We are now in the section of the curve below the 45° line, and so the movement shrinks over time. This process is shown graphically in Figure 2a.

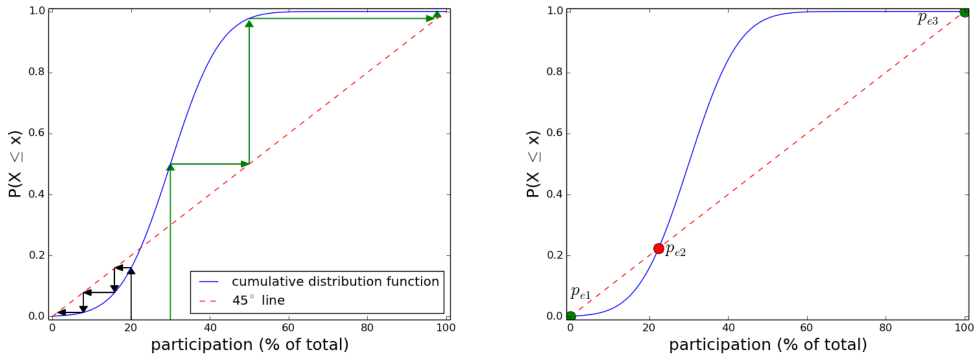

(a) Participation, $p_0 = 20, p_0 = 30$ (b) Equilibria for $\mu = 30$, $\sigma = 10$

This allows us to understand the 45° line as the line of equilibrium values, since by definition along this line a particular threshold value exactly matches the percentage of the population with that threshold. Thus, we will have an equilibrium level of participation whenever the CDF intersects this 45° line.

Figure 2b shows the equilibria for this distribution. The first and third equilibria are stable and the second is unstable, since any perturbation around this specific value will result in the movement growing or shrinking to another equilibrium.

3.2 Growing a movement from scratch

This equilibrium analysis yields some interesting conclusions, but the assumption that an initial group of people will spontaneously decide to participate in a protest is not necessarily a sound one. In a population with little to no instigators, each agent is only interested in participating if they can be assured that others will do so as well. With no prior level of participation to guide their decision, this group of first-round participants must simultaneously agree to join together.

In this situation, initial participation can be thought of as an n-player collective action game [Anthony, 2005, Ginkel and Smith, 1999, Oliver, Marwell and Teixeira, 1985]. While this is an interesting area of study, for the remainder of this paper we will focus on a simple threshold model, with the provision that initial participation is set at 0.

What if we applied a p_0 of 0 to the earlier distribution? Since this population has an equilibrium at 0, participation would remain flat, but this is not so for every distribution. $p_0 = 0$ represents a special case of the threshold model, in which only the "instigators" are starting the movement. Consider the case in which $\mu = 30$, $p_0 = 0$, but now we have raised σ to 20. This more heterogeneous population now has many instigators, whereas earlier there were little to none. Thus, in this instance the entire population eventually chooses to participate in the movement.

Again, we turn to the CDF for the explanation of this. Figure 3 shows CDFs for several distributions of $\mu = 30$. Where $\sigma = 10$ there are three equilibria, as established earlier. But, when $\sigma = 20$ there is only one equilibrium. This is a key feature of the normal distribution: it intersects the 45° line either one, two, or three times. This allows us to define a critical value of σ for a given μ: the value of σ for which there are exactly two equilibria.

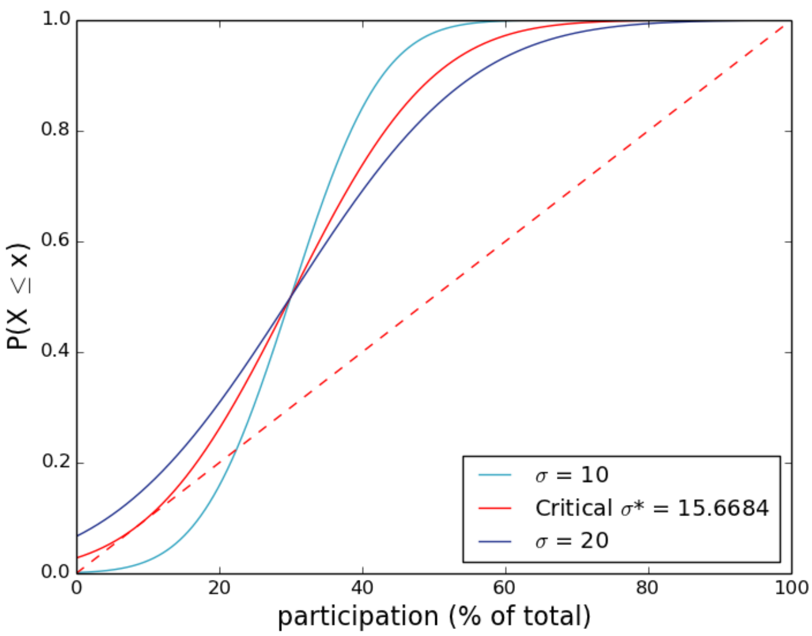

Figure 3: CDFs for $\mu = 30$

The middle line of Fig 3 shows the critical $\sigma*$ for $\mu = 30$. Above this value, there is only one equilibrium, which in this case is near 100% participation. Below this value, there are three equilibria, two stable and one unstable. Thus, $\sigma*$ represents an interesting property: if σ is above $\sigma*$, any initial level of participation will yield the same equilibrium. It is only under this circumstance that a movement can grow completely from scratch.

3.3 Deriving $\sigma*$ for a given μ

We can determine $\sigma*$ for a given μ (or a critical μ for a given level of σ, if we so desired) by satisfying two conditions. First, the CDF must intersect the 45° line, which has the equation $y = .01x$. Second, the 45° line must be tangent to the CDF. Since the PDF gives the slope of the CDF, we can satisfy this condition by setting the PDF equal to .01, the slope of the 45° line. The required system of equations is shown below.

$$.01 = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(x-\mu)^2}{2\sigma^2}} \tag{3}$$

$$.01x = \frac{1}{2}\left[1 + erf\left(\frac{x-\mu}{\sigma\sqrt{2}}\right)\right] \tag{4}$$

Fig 4 shows the critical $\sigma*$ for mu. As μ increases, $\sigma*$ increases at an increasing rate. This is for two reasons. First, as the population becomes more moderate, a larger variation in thresholds is needed in order to have "instigators" in the population. Second, as μ rises the CDF shifts to the right. Increasing σ stretches the CDF out, so 3-equilibrium solutions are less likely when σ is large.

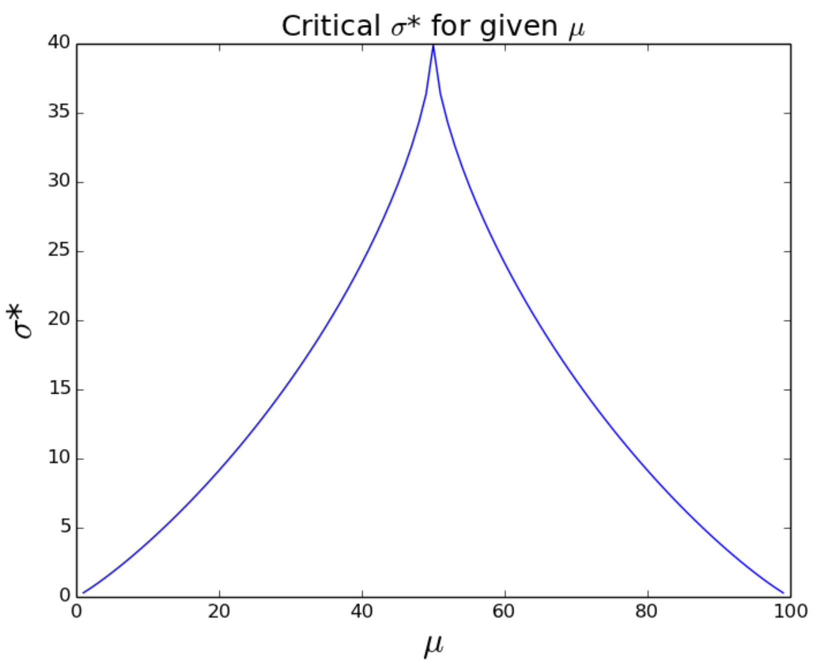

Figure 4: Critical $\sigma*$ for given μ

3.4 Differences when μ is low vs. when μ is high

We have seen that for low-μ populations, a critical $\sigma*$ exists which identifies where the population can grow a social movement from the ground up. In identifying $\sigma*$, we see that its value is symmetric around $\mu = 50$ (see Fig. 4. This is due to the symmetric nature of the normal distribution. However, $\sigma*$ takes on a very different meaning when $\mu > 50$. In this case, when μ is large and $\sigma > \sigma*$, the one equilibrium point occurs at a very low level of participation, instead of a high one as is the case when μ is low. Again, this is because raising μ has shifted the CDF to the right, so now the CDF for $\sigma*$ is tangent to the 45° line from below. Above this critical value, the majority of the CDF is now beneath the 45° line, not above it. This can be seen in Fig. 5.

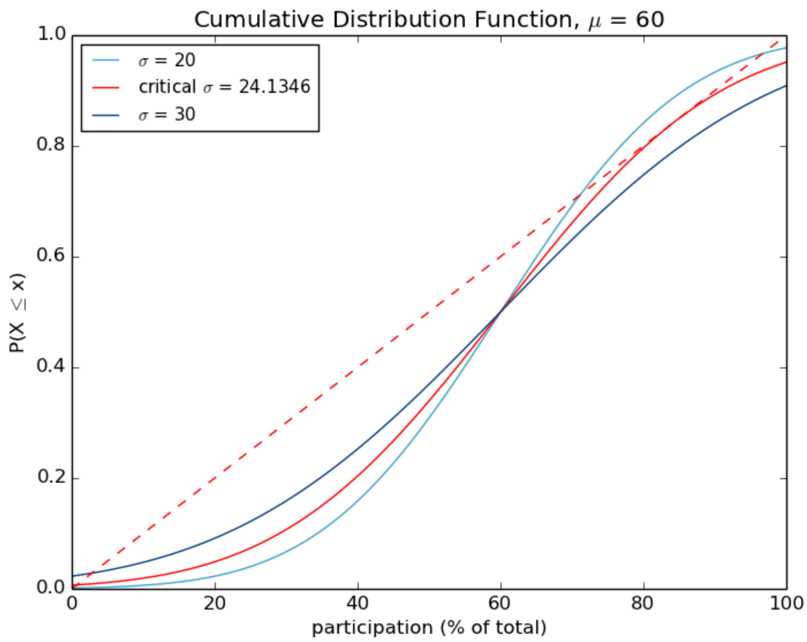

Figure 5: CDFs for $\mu = 60$

As with the case of low μ, for high μ when $\sigma > \sigma*$ there are three equilibria. However, the distribution is now more weighted toward the lower equilibrium, and the region in which the movement can grow is much smaller. Ultimately, for a high μ population there is little to no chance to grow a social movement unless it is begun with an extraordinarily high level of initial participation.

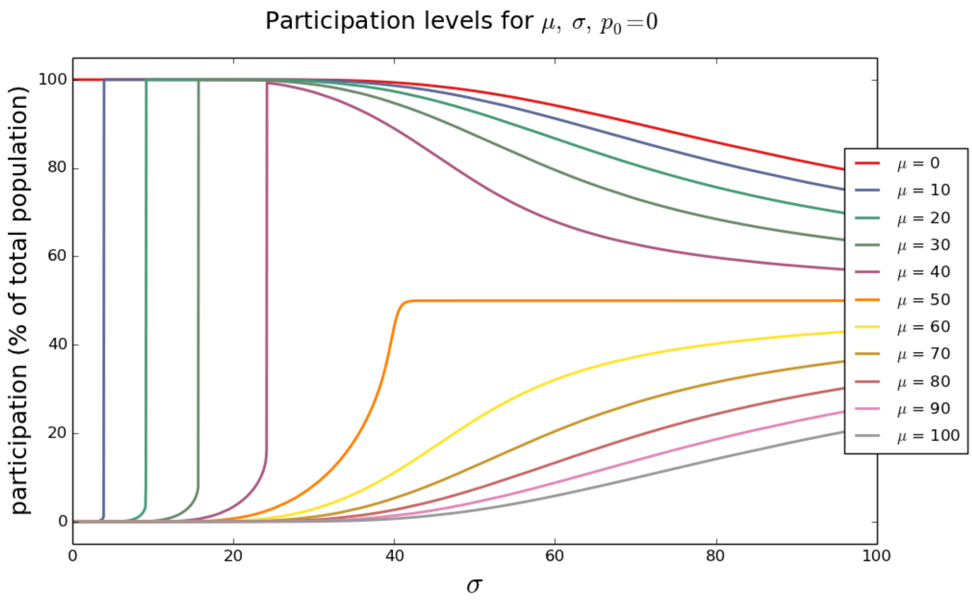

Figure 6: Participation after 100 rounds

Interestingly, for low μ populations participation actually starts to decrease as σ gets very high. This is because as σ increases, the proportion of "sticks-in-the-mud" goes up. For high μ populations, as σ rises the lower equilibrium point begins to rise, again because the CDF stretches horizontally as σ increases. The limit for all μ is 50% as $\sigma \to \infty$. This pattern can be seen in Fig. 6.

3.5 Equilibrium when n is small

So far we have looked only at large ($n = 10000$) populations of agents. How does this analysis differ when only a small sample of agents is taken?

When n is small, say 100 agents, sampling variation becomes a large factor. These effects are especially pronounced in cases like those in Fig. 7a, where the CDF very closely follows the 45° line. In this situation, the pattern of participation has a great deal of variance around the critical level of sigma. While a large sample would have captured the population's true tendency toward a low equilibrium, random small samples often produce situations in which the empirical CDF looks very different than the distribution it was drawn from.

This variation can work both for and against a movement. In cases where $\sigma < \sigma^*$ for the population as a whole, sampling with small n occasionally produces situations with much higher participation than in a large n sample.

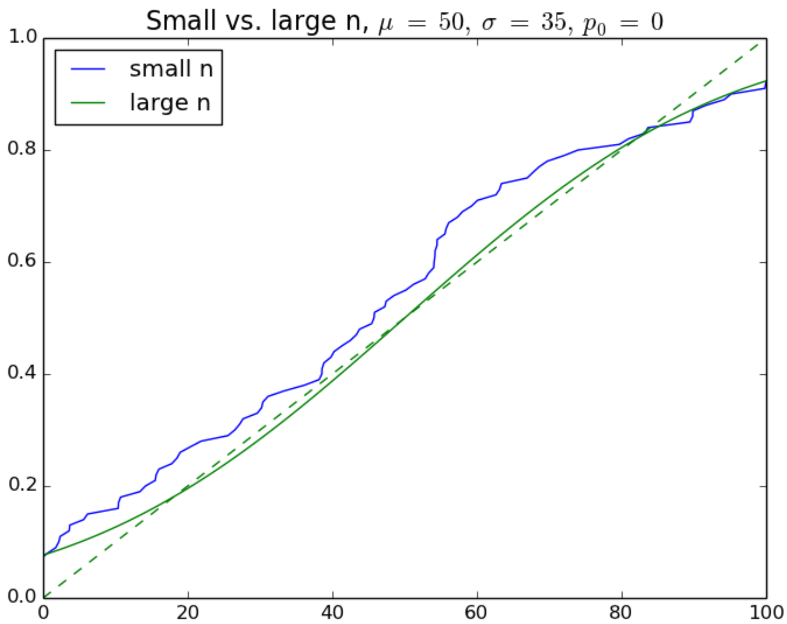

(a) CDF of $\mu = 50$, $\sigma = 35$, $n = 100$

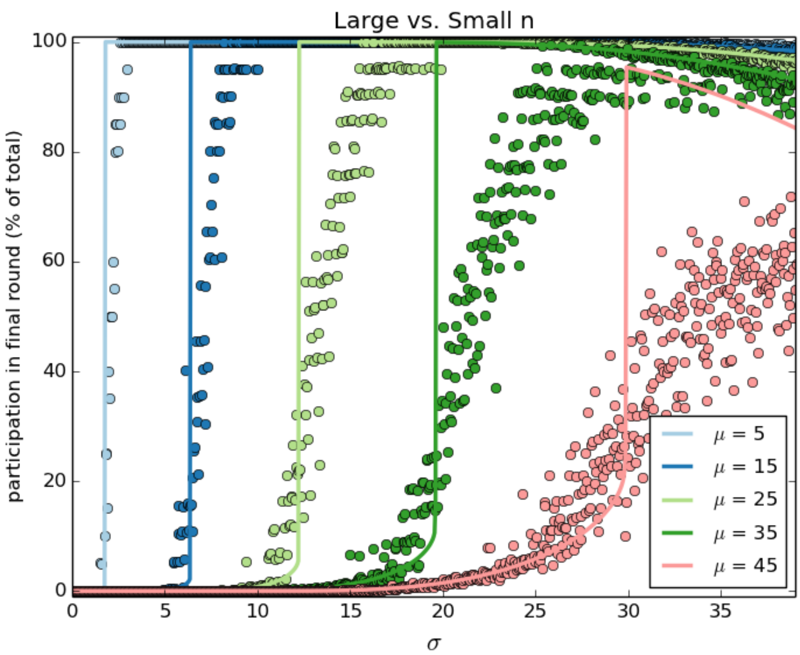

(b) Mean participation of 20 simulations, with each simulation running for 100 rounds

Figure 7: Effects of small n

This phenomenon is captured in Fig. 7b, which compares average levels of participation for $n = 100$ to the mathematical model of participation. For low levels of μ, as σ increases past $\sigma*$, participation is similar for both the model and the random sample. However, for larger μ populations, the behavior around $\sigma*$ becomes much more varied, so that when σ is just below $\sigma*$, small n samples actually do better on average than the model. Likewise, when σ is just above $\sigma*$, small n samples do worse on average than the model.

The effects of this sampling variation are more pronounced as μ increases because, as we saw earlier, $\sigma*$ increases along with μ. As $\sigma*$ increases, the CDF becomes more stretched out until it lies quite close to the 45° line. This makes the distribution more vulnerable to sampling variation, as thresholds can easily end up higher or lower than they would be otherwise.

In a practical sense, we can think of this as a case in which the population as a whole has a certain tendency, but the adoption of a movement is confined to some particular area. A good example is presidential primary elections. It is common for candidates in primaries to appear to be more extreme in their views in order to drum up support from within their parties. However, voters in primary elections tend to feel more strongly about "core" party issues than voters in general elections, and so candidates who seemed like perfectly logical choices in the primary often end up having to moderate their positions in order to win over the country as a whole. Views which seem too extreme for the general election thus appear to be perfectly logical when faced with a small sample.

4 Skewed distributions

So far we have examined how the relationship between μ, σ and p_0 determines the adoption of a social movement over time within a normal distribution. Now, we will examine how this analysis can be applied to a population in which the threshold distribution is not symmetric. Here We will present an analysis of the Weibull distribution, but the techniques outlined in the previous section can be applied to any probability distribution.

The Weibull distribution has three parameters, α, β and μ. The degree of skewness is determined by α, with maximum skew occurring as $\alpha \to 0$. The spread of the data is determined by β, with the distribution covering a wider range as β rises. Finally, μ is a location parameter, which defines a lower bound for the data. The basic Weibull distribution is right-skewed, but it can easily be adapted to left-skewness by transforming the PDF across the y-axis, in which case μ then becomes an upper bound. The PDF and CDF for the (right-skewed) Weibull distribution are:

$$f(x) = \frac{\alpha}{\beta}\left(\frac{x-\mu}{\beta}\right)^{\alpha-1} e^{-\left(\frac{x-\mu}{\beta}\right)^{\alpha}} \tag{5}$$

$$F(x) = 1 - e^{-\left(\frac{x-\mu}{\beta}\right)^{\alpha}} \tag{6}$$

Results from the skewed distribution are shown in Fig. 8. In order to isolate the effects of skewness, here μ is kept constant (at either -1 or 101, depending on the direction of skewness) while α varies. Critical values of β were identified for $\alpha \in [1.2, 4]$. As with the normal distribution, populations with a lower mean threshold were more likely to have a successful movement. Because of the translation of the distribution over the y-axis, $\beta > \beta*$ for a right-skewed distribution meant that a movement was never successful, while $\beta > \beta*$ for a left-skewed distribution meant that a movement was always successful. The shaded area of Fig 8 shows the "winning" combinations of α and β for both distributions.

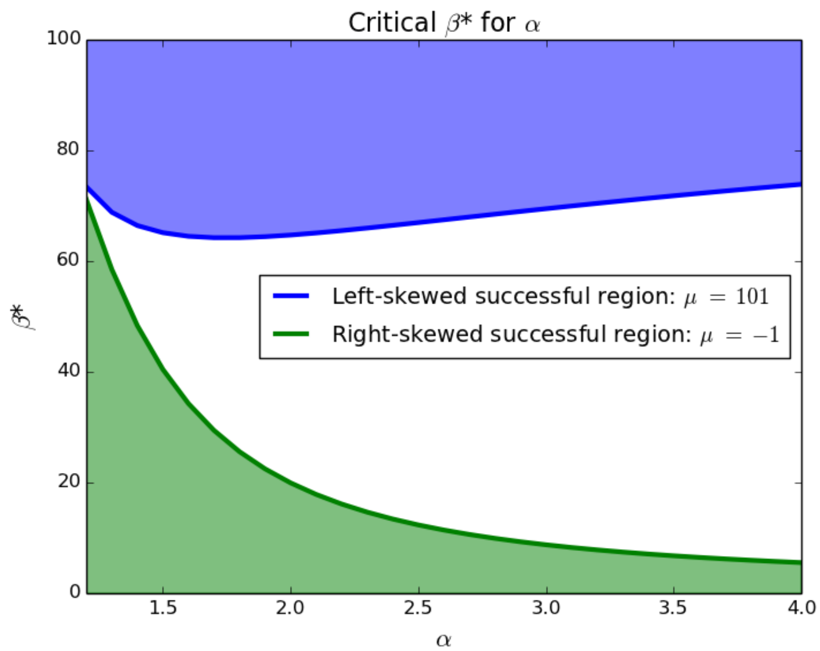

Figure 8: Results from the Weibull Distribution

Both left- and right-skewed distributions had combinations of α and β which produced successful movements, but $\beta*$ was significantly higher for left-skewed distributions. Importantly, $\beta*$ needed to be so large for these populations that the "tail" of the distribution had to extend well beyond zero. Thus, while β^* does exist for left-skewed distributions within the definition of the model, in a practical sense (where thresholds at or below zero are equivalent to zero) these "successful" distributions cannot really be considered skewed at all.

This leads us to the right-skewed distributions, for which a large β was actually detrimental to success. This was because, when μ is held constant at -1, the proportion of instigators falls as skewness falls. Keeping β low maintains this crucial segment of the population, without which no movement can pick up momentum. It should not be surprising, then, that more negatively skewed distributions could tolerate a much greater variation in β.

5 Policy implications

What do these results tell us about real-life social movements? Perhaps the most important lesson from the threshold model is that widespread social movements are actually quite rare. Growing a movement from scratch requires a perfect storm of conditions: a distribution of thresholds with low μ, as well as sufficient σ such that there are a large number of instigators, but not so high that there are many "sticks-in-the-mud" who will prevent full participation.

Policy implications then depend on the ultimate goal: to encourage the movement or to prevent it from spreading. Individuals who want to grow a movement first must act as instigators, so that they can lessen the risk for individuals who have low but non-zero thresholds. If μ is low, large heterogeneous (high σ) populations are more likely to succeed. However, if μ is high there is little chance of success. Though changing the average threshold of a population may seem daunting, attempts to do just this are prevalent within mass media, both from those seeking to inspire social movements and those trying to prevent them. In the Arab Spring protests, participants used social media to spread pictures and video of every transgression by the regime, hoping to turn the world sympathetic to their cause. Within the local region, these updates served as a means of communication and a show of solidarity, assuring would-be protesters that they would not be alone.

On the other side, those wanting to prevent a movement from happening would want to decrease σ, uniting the population around a moderate view and removing the probability of instigators within the population. A good example of this is government propaganda within authoritarian regimes. These messages reinforce the idea that the status quo is desirable, that citizens who agree with the government are in the majority, and that any rebellion would be unusual and quickly quashed. However, as Bhavnani and

Ross [2003] showed, governments employing these tactics walk a fine line between reinforcing their strength and undermining their credibility by making announcements that are so optimistic they are not believable. When the population's μ is low but close to 50, opponents of a movement benefit from restricting information so that citizens can only observe a small portion of the participation. This was the goal of the Egyptian government's shutdown of the internet during the Tahrir square protests, though this tactic ended up being ineffective.

6 Conclusion

This inquiry has sought to establish that the shape of a population's threshold distribution determines the ultimate success or failure of a social movement. Since the true thresholds of individuals are difficult to measure, tactics to increase or shrink participation often amount to a guessing game for the actors involved. In real life, not only is this distribution of thresholds difficult to estimate, but so is the actual participation within the population at any given time. The threshold model becomes more complex with imperfect information, and a good direction for future work would be exploring how these dynamics change if agents only know the participation of those in their social circle, or those in their neighborhood, or those they see on mass media.

Another assumption made in this model is that each agent makes an equal contribution to the movement, namely through their own participation. This is not the case in real life. Memes spread faster with the help of individuals with large platforms, such as popular bloggers. Likewise, Kickstarter campaigns gain greater visibility when well-connected individuals promote them, and those with greater resources can contribute more toward the campaign[1]. Another possible avenue for research would be to explore how inequality of resources affects participation in movements.

Though the threshold model is a simplification, it nonetheless provides key insights into the nature of social movements. Small changes to the distribution of thresholds can lead to vastly different outcomes. Further, the success or failure of a movement often is driven by the presence of instigators. It is their interaction with the population that determines mass behavior. In social movements, it may indeed be the case that the tail of the distribution "wags the dog."

[1]Kickstarter actually has a \$10,000 limit on contributions, which is easily large enough to influence any small to medium-sized campaign. Individuals can also offer to match any contributions raised through successful campaigns, as happened with the Reading Rainbow campaign when celebrity Seth McFarlane offered to match up \$1 million in contributions.

References

Anthony, Denise. "Cooperation in microcredit borrowing groups: identity, sanctions, and reciprocity in the production of collective goods." *American Sociological Review* 70, no. 3 (2005): 496-515.

Bauckhage, Christian, Kristian Kersting, and Fabian Hadiji. "Mathematical Models of Fads Explain the Temporal Dynamics of Internet Memes." In ICWSM. 2013.

Bhavnani, Ravi, and Michael Ross. "Announcement, credibility, and turnout in popular rebellions." *Journal of Conflict Resolution* 47, no. 3 (2003): 340-366.

Borge-Holthoefer, Javier, Raquel A. Baños, Sandra González-Bailón, and Yamir Moreno. "Cascading behaviour in complex socio-technical networks." *Journal of Complex Networks* 1, no. 1 (2013): 3-24.

Centola, Damon, and Michael Macy. "Complex contagions and the weakness of long ties." *American Journal of Sociology* 113, no. 3 (2007): 702-734.

Etter, Vincent, Matthias Grossglauser, and Patrick Thiran. "Launch hard or go home!: predicting the success of kickstarter campaigns." In Proceedings of the first ACM conference on Online social networks, pp. 177-182. ACM, 2013.

Ginkel, John, and Alastair Smith. "So You Say You Want a Revolution: A Game Theoretic Explanation of Revolution in Repressive Regimes." *The Journal of Conflict Resolution* 43, No. 3 (Jun 1999): 291-316.

Granovetter, Mark. "Threshold Models of Collective Behavior." *The American Journal of Sociology* 83, No. 6 (May, 1978): 1420-1443.

Hedström, Peter. "Contagious Collectivities: On the Spatial Diffusion of Swedish Trade Unions, 1890-1940." *American Journal of Sociology* 99, No. 5 (Mar., 1994): 1157-1179.

Kavanaugh, Andrea, Seungwon Yang, Steve Sheetz, L. T. Li, and E. Fox. "Microblogging in crisis situations: Mass protests in Iran, Tunisia, Egypt." In ACM. CHI. 2011.

Kuppuswamy, Venkat and Barry L. Bayus. "Crowdfunding Creative Ideas: the Dynamics of Projects Backers in Kickstarter: SSRN Working Paper." (2013).

Kuran, Timur. “Sparks and prairie fires: A theory of unanticipated political revolution.” *Public Choice* 61, no. 1 (1989): 41-74.

Lohmann, Susanne. “The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, EastGermany, 1989-91.” *World Politics*, vol. 47, no. 1 (Oct 1994): 42-101

Lotan, Gilad, Erhardt Graeff, Mike Ananny, Devin Gaffney, and Ian Pearce. “The Arab Spring— the revolutions were tweeted: Information flows during the 2011 Tunisian and Egyptian revolutions.” *International Journal of Communication* 5 (2011): 31.

Macy, Michael W. “Chains of cooperation: Threshold effects in collective action.” *American Sociological Review* (1991): 730-747.

Mitra, Tanushree and Eric Gilbert. “The Language That Gets People to Give: Phrases That Predict Success on Kickstarter.” Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing. (Feb 2014): 49-61.

Mollick, Ethan. “The Dynamics of Crowdfunding: An Exploratory Study.” *Journal of Business Venturing* 29, No. 1 (Jan 2014): 1-16.

Myatt, David P. and Chris Wallace. “When Does One Bad Apple Spoil the Barrel? An Evolutionary Analysis of Collective Action” *The Review of Economic Studies* 75, No. 2 (Apr., 2008): 499-527.

Oliver, Pamela, Gerald Marwell, and Ruy Teixeira. “A theory of the critical mass. I. Interdependence, group heterogeneity, and the production of collective action.” *American journal of Sociology* (1985): 522-556.

Tufecki, Zeynep, and Christopher Wilson. “Social Media and the Decision to Participate in Political Protest: Observations From Tahrir Square.” *Journal of Communication* 62 (2012): 363-379.

Yin, Chien-Chung. “Equilibria of collective action in different distributions of protest thresholds.” *Public Choice* 97, no. 4 (1998): 535-567.