Unit 2: Contagion

READINGS AND RESOURCES

- Barnett GA (2011) Mathematical models of the diffusion process. In *The diffusion of innovations: A communication science perspective,* pp 103-122. New York: Peter Lang Publishing.
- Bass FM (1969) A new product growth for model consumer durables. *Management Science* 15: 215-227.

What is contagion?

Things spread from one person to another. Diseases for sure, but also information, behaviors, and technologies. Your friends buy a new kitchen appliance, and you buy the same one the next week. Your cousin starts doing hot yoga, and you find yourself at a Bikram studio a few days later. What does this look like from a dynamical perspective? What is the time course of adoption?

In 1962, the first edition of Everett Rogers' now-classic book, *The Diffusion of Innovations* was published. Rogers studied how various innovations were adopted, everything from hybrid seed corn among Midwest farmers to the adoption of ham radio among tech enthusiasts to the adoption of new ideas among French intellectuals. and found that most of them exhibited a sigmoidal or S-shaped relationship of cumulative adoption over time. That is, adoption increase slowly at first, then increased rapidly, then slowed down again. What explains this apparently universal pattern of adoption?

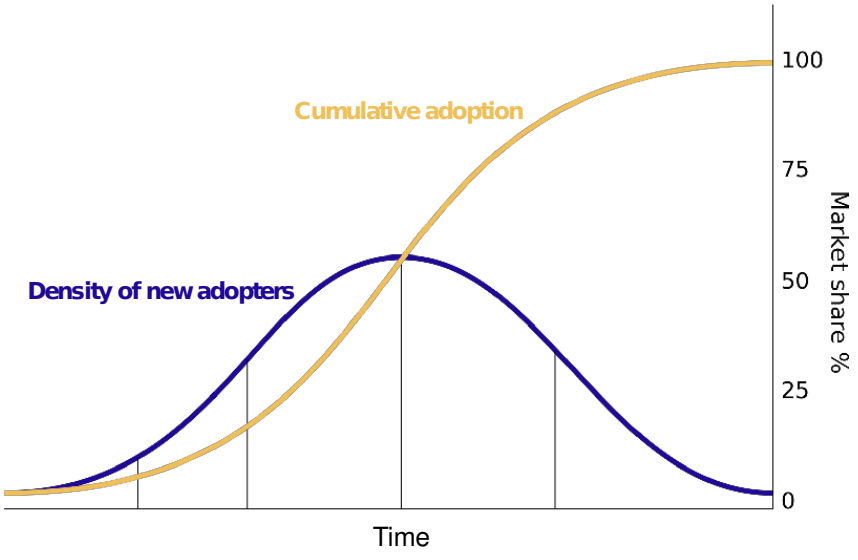

A simple model: spontaneous adoption

CODE: contagion_spontaneous.nlogo

We're going to assign each agent one of two possible states: infected or not. This terminology is taken from the idea of disease contagion, but it can be applied to the adoption of products or behaviors if you think of "infected" as simply indicating that the agent has adopted whatever it is. Agents will move around in space. Sometimes, an uninfected agent becomes infected. This can be thought of in several ways. We could be talking about a literal disease infection. But we could also be talking about the adoption of an idea, a behavior, or a product. Now, if there's no way for an agent to become *uninfected* (and we'll consider that later one), then obviously this model is always going to end up with the whole population infected. However, exactly how that happens isn't necessarily obvious. We'll focus here on what the model predicts about the temporal dynamics of infection.

Let's assume that at the beginning, a few agents introduce the behavior into the population. After that, each individual has a fixed probability of adopting it at any given time. We can do this by introducing a new variable called *spontaneous-infect*, the probability of adopting the behavior at each time. What we're really assuming here is that agents all have the same information and the same psychology – the same proclivity for adoption. What we'll looking for, then, is the dynamic pattern of how they adopt over time.

Let's take the simple model of agent movement we built at the end of the last lecture, in which we have a bunch of circular agents moving randomly in space. We'll get rid of the pens, but keep the rest. In that model, the agents didn't really do much. We're going to change that here.

For each component of the model description below, follow along by looking at the provided code to see how it each part is implemented.

SETTING UP THE MODEL

- Create a new slider for the probability of spontaneous adoption, *spontaneous-infect*, which can vary between 0 and 1.
- In the code window, give each turtle their own parameter (i.e., a turtles-own variable) called infected?, which simply indicates whether the agent is infected. The question mark indicates that it's a Boolean variable, this is a NetLogo convention.

INITIALIZATION

- Create N turtles. Place them randomly in 2-D space.
- Infect some of them.

DYNAMICS

- Stop simulation is everyone or no one is infected.
- ASK TURTLES: Become infected with probability *spontaneous-infect.*
- ASK TURTLES: Move.

PLOTTING:

- We will plot the number of infected agents over time. NetLogo makes creating this sort of plot very easy.

RESULTS

- If everyone has the same information and the same psychology – the same proclivity for adoption – what do we get? An **r-shaped curve** of asymptotic growth, in which the rate of adoption is maximal at first and then slows as there are increasingly few individuals who haven't yet adopted.
- This is definitely not the S-shaped adoption curve we see empirically. So we can confidently say that innovations don't diffuse this way, with everyone having the same proclivity for adoption and the same information. What would give us those S-shaped curves?

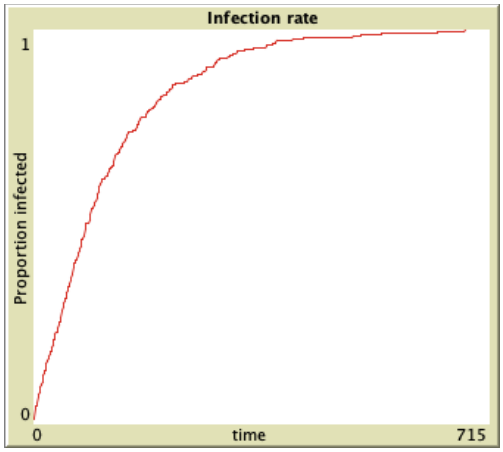

Rogers knew back in 1962 that an explanatory theory required something more than the assumption that there was one type of person, all responding to the same information. His proposal solution was based primarily on individual differences: that individuals differ in their proclivity for adoption. In addition to a small set of initial innovators, he partitions the population into the early adopters, early majority, late majority, and laggards. To make this taxonomy work, he also had to posit very specific proportions of the population that fall into each category. Which works in terms of when they adopt, but it assumes an awful lot about why people adopt and why those differences might occur, much of which is not really justified by what we know about humans. Of course, there might be some truth to it – there are certainly differences between individuals. But there's also a simpler explanation that can explain the S-shaped adoption curve.

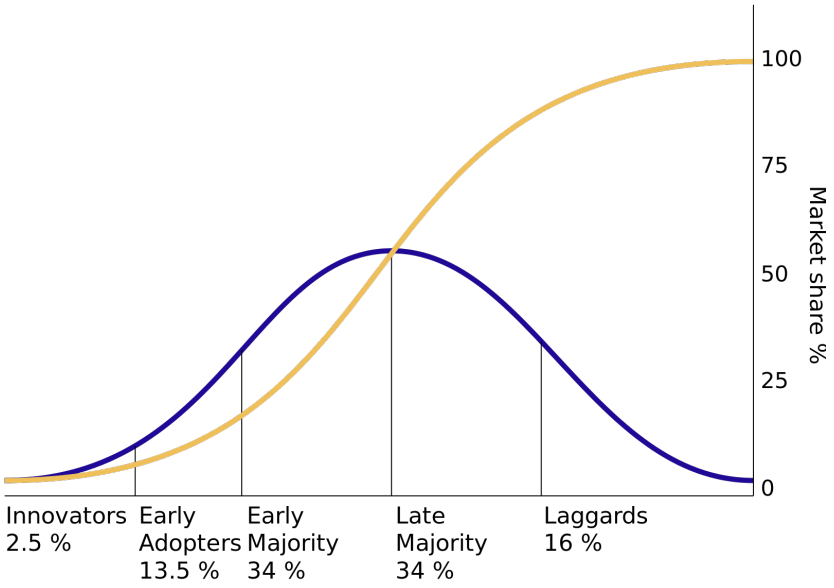

A model of social influence

CODE: Contagion_SI.nlogo

What if, instead of positing a very specific distribution of proclivities for adoption, one which has questionable support based on what is known about human psychology, we instead return to our original model in which everyone has pretty much the same proclivity to adopt. Wait a minute, you say, we already showed that wouldn't work. Ah, but the other assumption our first model made was that all individuals had access to the same information. What if *that* assumption no longer held?

This model was first introduced by Frank Bass in 1969, just a few years after the first edition of Rogers' book was published. What if, Bass proposed, innovations spread sort of like diseases? What if every exposure carried some risk that this time, an individual would be convinced and adopt.

So, let's return to our model. Now the fact that the agents are moving around in space *will* matter. At every tick of the simulation clock, each agent will consider all the other agents around itself nearby. NetLogo allows us to easily have an agent check all of the other agents within some radius. For each neighbor who has adopted the behavior, there is a fixed probability of adopting. This probability is the *transmissibility* of the behavior, similar to how contagious a disease is.

Transmissibility. Figuring out exactly how to model this will require some mathematics of probability. Let τ be the transmissibility, the probability of

being influenced to adopt by someone who has already adopted. If there's only one such neighbor, then we're set: the probability of adopting is τ. However, what if there are multiple neighbors who have adopted? Probabilities aren't additive, so we can't just sum up the τ's. Instead, their multiplicative. But τ^2 is less than τ, so that doesn't make sense. Instead, we can take a different tack: calculate the probability that the contagion *doesn't* spread. With multiple neighbors adopting, we need the joint probability that all of them fail to transmit the contagion. Say there are three neighbors who have adopted. The probability of the contagion not spreading from any of neighbors is $(1-\tau)(1-\tau)(1-\tau)=(1-\tau)^3$. More generally, with *n* infected neighbors, the probability of the contagion not spreading is $(1-\tau)^n$. So, the probability of spreading is simply the inverse: $1-(1-\tau)^n$.

SETTING UP THE MODEL

- Add *transmissibility* slider

INITIALIZATION

- Create N turtles. Place them randomly in 2-D space.
- Infect some of them.

DYNAMICS

- ASK TURTLES: Count the number of infected neighbors. **Become infected with a probability that depends on this number and on the *transmissibility* of the contagion*.***
- ASK TURTLES: Move.

RESULTS

- So what happens when adoption spreads via direct social influence rather than by a fixed response to global information? In the beginning, few agents have adopted, so most interactions will not lead to new adoptions. The rate of growth will be slow. As most agents adopt, the adoption rate will increase until a majority of agents have adopted. After this, adoption will slow down again since most interactions will be between agents who have already adopted. In other words, we get our **S-shaped curve**.

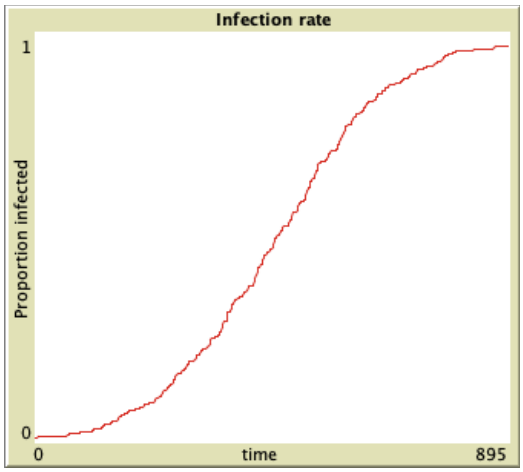

- Now, we've got a much simpler model that shows how the nature of social influence can generate sigmoid curves.
- Play around with the sliders on this model. It's often valuable to simply play with a model to get a feel for the sort of dynamics that occur when we change the parameters. Even experiments where the results are obvious are valuable as a check to make sure your model is running as it should. For example, vary the *transmissibility* of the contagion. We can see that the infection spreads much faster when transmissibility is 0.1 than when it is 0.01! Obviously, innovations that are more intrinsically attractive will diffuse more rapidly through a population. Other factors are perhaps less obvious.
- The fewer agents there are – the less dense the population – the fewer interactions will occur per unit time, and the longer it will take for the innovation to diffuse. This is perhaps counterintuitive. It might seem like the more people who need to adopt, the longer it will take. But if adoptions events are independent and rely on the density of social networks, then denser populations will adopt more rapidly.
- Relatedly, the speed of diffusion depends on how much mixing goes on – the rate at which adopters will interact with new people who haven't yet adopted. We can use the speed and turning-angle parameters to control the contact rate and ultimately the speed of diffusion.

The most important result is that the social influence model fits the data better than the spontaneous adoption model, and makes far fewer arbitrary assumptions than Rogers' model. The model shows us how social influence and local interactions can give rise to a sigmoidal adoption curve.

Recovery: The SIS model

Our model so far has assumed that once people are infected, or once they have adopted whatever product or behavior we're talking about, that's it. They always have it. That's a pretty strong assumption. In terms of infection,

sometimes people recover, and are no longer infected. In terms of products or behaviors, sometimes their use runs out or people decide they are no longer interested, and stop for a while. They might give it up forever, or they might simply take a hiatus. Let's focus on the latter case. What happens if we let people recover (or disadopt), at which point they return to a state of susceptibility? We can do this by assuming that individuals who are infected recover with a fixed probability γ : the recovery rate. What sort of dynamics does this yield? Will everyone eventually recover?

SETTING UP THE MODEL

- Add *recovery-rate* slider

INITIALIZATION

- Create *N* turtles. Place them randomly in 2-D space.
- Infect some of them.

DYNAMICS

- ASK TURTLES: Count the number of infected neighbors. Become infected with a probability that depends on this number and on the *transmissibility* of the contagion.
- **ASK INFECTED TURTLES: Recover with probability *recovery-rate.***
- ASK TURTLES: Move.

RESULTS

- So what happens? Does the contagion eventually fizzle out? Do we get wild oscillations? In most cases, no. Rather, the population settles to an equilibrium where the rate of new infections equals the rate of recovery. Note that we can get the any particular equilibrium rate of infection for almost any level of transmissibility as long as the recovery rate adjusts accordingly.
- Movement -- the effective contact rate -- decrease the equilibrium level of infection. Show this if we change speed from 0.5 to 0.1, and change the turning angle from 60 to 360, the equilibrium infection rate goes way down.

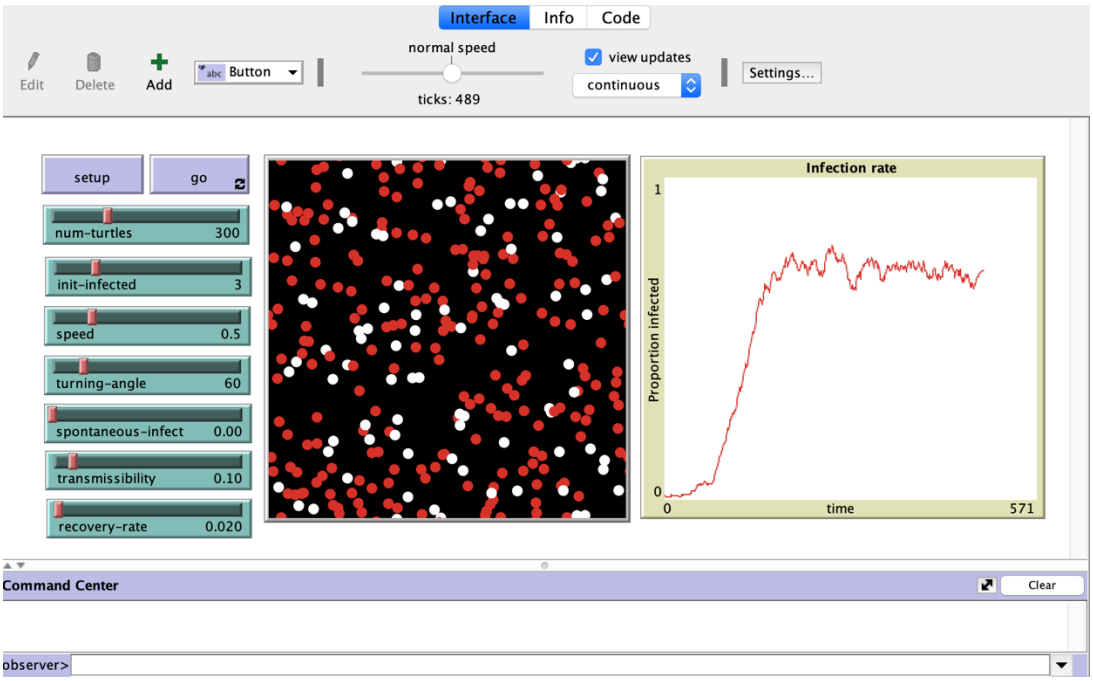

What's the use of models like this? SIS models like this can be particularly useful when considering disease dynamics. If we have estimates about the contact rate of a population and the transmissibility and recovery rate of a disease, we can estimate the proportion of the population that needs to be vaccinated to prevent an infection from spreading. However, they can also be used for understand any behaviors that spread through social influence but aren't necessarily permanent. For example, SIS models have been used to study health behaviors related to diet and exercise, as well as the spread of emotions or sentiments in social networks.

Further directions

- Transmission biases. We've assumed that individuals simply adopt behaviors from anyone they contact. But there may be good reasons for not doing so. In an uncertain environment, conformity to the majority behavior can be beneficial. If the value of behaviors is opaque, copying successful or prestigious individuals can be a good strategy. An extensive literature has considered the dynamics and evolutionary implications of various transmission biases, summarized in these papers:
    - Laland KN (2004) Social learning strategies. *Learning & Behavior* 32: 4-14.
    - Kendal RL et al. (2018) Social learning strategies: Bridge-building between fields. *Trends in Cognitive Sciences* 22: 651-665.
    - Smaldino et al. (2018) Sigmoidal acquisition curves are good indicators of conformist transmission. *Scientific Reports* 8: 14015.

- Complex contagion. The SI(S) models assume the probability of adoption increases linearly with the number of exposures. But, as special sort of transmission bias, social behaviors might require reinforcement from multiple sources before they are adopted. Centola and colleagues have studied how behaviors spread in social networks and applied models creatively to generate hypotheses.
    - Centola D, Macy M (2007) Complex contagions and the weakness of long ties. *American Journal of Sociology* 113: 702-734.
- Emotion contagion. Some researchers have used a modified SIS model to explore contagion of emotions or sentiments. That is, they have modeled how social interactions may cause particular feelings to spread on social networks.
    - Hill AL, Rand DG, Nowak MA, Christakis, NA (2010) Emotions as infectious diseases in a large social network: The SISa model. *Proceedings of the Royal Society B* 277: 3827–3835.
- Ingroup-bias and adoption. What happens if there is more than one social group, who view adoption by in-group and out-group members differently? There may be a bias to adopt a population product or behavior, but not if it's popular primarily among members of the out-group. We explore how such a bias can affect adoption dynamics and interacts with parameters influence group differences and population structure.
    - Smaldino PE et al. (2017) Adoption as a social marker: Innovation diffusion with outgroup aversion. *Journal of Mathematical Sociology* 41: 26-45.

Exercises

- *Getting used to it.* Plot equilibrium adoption in the SIS model as a function of *recovery-rate* for different values of *transmissibility*, *turning-angle*, and *speed*. Describe the relationship(s).
- *Vaccinate!* Modify the SIS model so that a fixed proportion *V* of the population is vaccinated so they cannot be infected. Start with a single infection, and consider whether the infection spreads (and reaches an equilibrium) or dies out (so that no agents are infected any more). Run a small batch of runs to consider the proportion of runs for each parameter condition in which the contagion failed to spread, varying the *transmissibility* and *recovery-rate* for at least two values each for arbitrary movement parameters. What is the relationship between *V, transmissibility,* and *recover-rate* in whether an infection spreads?
- *Don't copy those people.* Consider how to extend the model to include two groups who both adopt the product via social influence (and maybe even spontaneously), but are inclined to dis-adopt if they perceive it to be overly represented in the outgroup. How might you modify the contagion models we covered to do this?