Day Four: Dynamic Networks

Dr Jamsheed Shorish

The Australian National University
jamsheed.shorish@anu.edu.au
http://be.linkedin.com/in/jshorish

18 July 2013 - Beihang University

Overview

The Bass Toy Model

Bass, F. M. (1969) "A New Product Growth Model for Consumer Durables", *Management Science* v. 15 no. 5, pp. 215-227.

Agent Based Simulation

Compartment Models of Diffusion

- The SIR Model
- The SIS Model

Examples

- The Banking Crisis
- Meme Diffusion

The Bass Toy Model

- The Bass model is an abstract model of social interaction, based upon two characteristics: *innovation* and *imitation*.
- It can be used by firms to estimate the consumption behaviour of their product demand.
- Although very abstract, it introduces in a clean manner the following characteristics of dynamic networks:
  - **Time**, both discrete and continuous;
  - Dependence upon **other individuals** (as a neighbourhood, or generally);
  - Convergence to a **stationary solution**, or 'steady state'.
- The Bass model can also be used to investigate how firms adopt their products according to their competitors, resulting in a supply-side model of innovation vs. imitation.

The Bass Toy Model: Discrete

- The Bass model assumes that there are two different forms of social behaviour:
  - **Innovation**, which occurs at a rate $\tilde{\gamma}$;
  - **Imitation**, which occurs at a rate $\tilde{\delta}$.
- In *discrete* time ($t = 1, 2, \ldots$), the **law of motion** for the population of consumers is:

$$p(t+1) = p(t) + \tilde{\gamma}(1 - p(t)) + \tilde{\delta}(1 - p(t))p(t),$$

  where $p(t)$ is the fraction of the population at time t who are currently consumers of the firm's product.
- Thus, it is the *proportion* of the consumer population that is changing over time.

The Bass Toy Model: Continuous

- Often, it is easier mathematically to investigate models where the time interval of interaction is very much smaller than the time interval of analysis.
- In the limit of an extremely small (actually: infinitesimal) time interval Δt, we assume that time is continuous.
- In continuous time, $\tilde{\gamma} = \gamma \Delta t$, $\tilde{\delta} = \delta \Delta t$ and the Bass model becomes:

$$p(t + \Delta t) = p(t) + \tilde{\gamma}(1 - p(t)) + \tilde{\delta}(1 - p(t))p(t) \Rightarrow$$
$$p(t + \Delta t) - p(t) = (\gamma \Delta t + \delta p(t) \Delta t)(1 - p(t)) \Rightarrow$$
$$\frac{p(t + \Delta t) - p(t)}{\Delta t} = (\gamma + \delta p(t))(1 - p(t)).$$

- Taking the limit as $\Delta t \to 0$ yields an **ordinary differential equation** or ODE:

$$p'(t) = (\gamma + \delta p(t))(1 - p(t)).$$

The Bass Toy Model: Solving the ODE

- Let's solve the resulting ODE for the Bass model: this means finding a function $p(t)$ that depends upon γ, δ, and t.
- To do this we need an *initial condition* for the distribution of product consumers–let's assume that at time $t = 0$, there are no consumers, i.e. $p(t) = 0$.
- Solving this ODE isn't particularly easy, because it is non-homogeneous and non-linear in p.
- But we can get some idea of the solutions we obtain by looking at the *implicit solution* of the system–the following graphic displays the *qualitative* shape of the $p(t)$ solution using Romberg interpolation of the ODE, for a few γ and δ values:

The Bass Toy Model: Implicit Solution

<1><2>(1cm, 4cm)(0.5cm, 2cm)

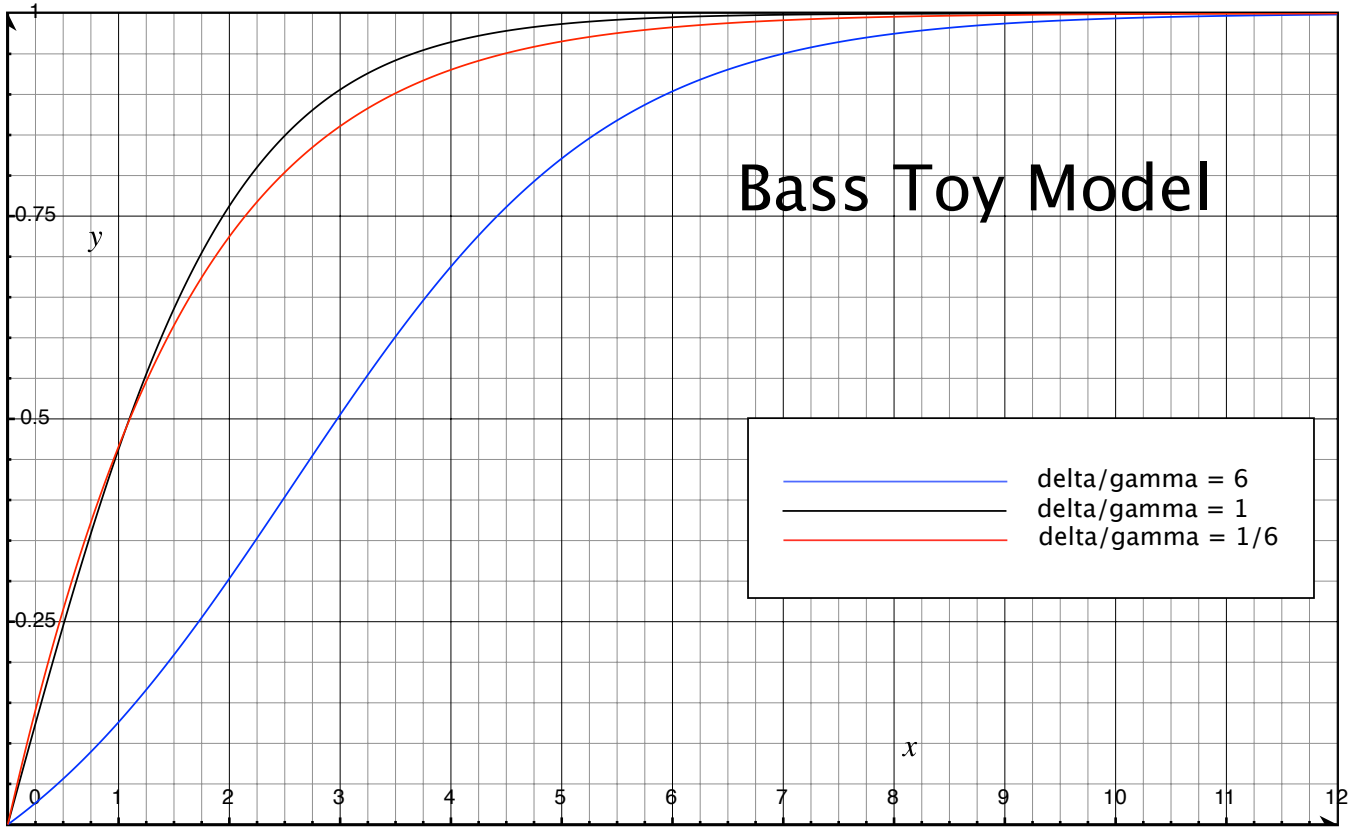

- Note here that the shape of the implicit solution remains essentially the same, but shifts according to the *fraction* δ/γ.

The Bass Toy Model: Qualitative Features

- We learn a great deal about the system even before explicitly finding a solution.
- For example, we can see that more consumers will purchase the product as time goes by.
- In addition, in the limit ('infinite time'), everyone will consume the product.
- Finally, we can see that when there is a higher rate of imitation (δ) over innovation (γ), convergence is *slower* than the opposite–it takes longer for everyone to purchase the product. A producer, then, would like to influence consumers to imitate more.

The Bass Toy Model: Explicit Solution

- The explicit solution of the ODE is:

$$p(t) = \frac{1 - e^{-(\delta+\gamma)t}}{1 + \frac{\delta}{\gamma} e^{-(\delta+\gamma)t}} \Leftrightarrow p(\tilde{t}) = \frac{1 - e^{-(1+q)\tilde{t}}}{1 + q e^{-(1+q)\tilde{t}}},$$

  where $q := \frac{\delta}{\gamma}$ is the ratio of imitation to innovation, and $\tilde{t} := \gamma t$ (we'll just continue to call this t in what follows).

- Note that as expected from the qualitative analysis, the ratio q is what matters–time is rescaled according to the value of γ, but this does not affect the solution.

Convergence to a Stationary State

- Now that we have the explicit solution we can confirm what we expected from the qualitative results–that starting from a consumer base of zero, eventually all consumers will purchase the product.
- We can see this by taking the *limit* of the explicit solution for $p(t)$, as time goes to infinity:

$$p^\star := \lim_{t\to\infty} p(t) = \frac{1-e^{-\infty}}{1+\frac{\delta}{\gamma}e^{-\infty}} \Rightarrow$$

$$p^\star = 1.$$

- We can also verify $p^\star = 1$ is a stationary solution in the ODE.

Stability of the Stationary State

- Let's modify the model *slightly* by allowing for a very small chance that consumers return the product immediately after purchasing it. We call this chance η and assume that $\delta, \gamma \gg \eta$. How does this change our possible convergence to $p^\star = 1$?
- This exercise is known as a *perturbation* of the dynamical system. Perturbations are important for stability analysis in addition to understanding the impact of using a model as an approximation to a more complicated environment. Perturbation analysis is used extensively in the natural and social sciences.
- We modify the Bass model appropriately:

$$p'(t) = (\gamma + \delta p(t))(1 - p(t)) - \eta p(t) \Rightarrow$$
$$p'(t) = \gamma + (\delta - \gamma - \eta)p(t) - \delta(p(t))^2. \qquad (1)$$

Stability of the Stationary State

- Note that the effect of a small perturbation is to (in effect) slightly decrease the impact of both innovation and imitation. Part of the existing consumer base decides to return the product and hence 'undo' the positive effect of consumer adoption, but there is a continuing second-order impact of imitation that is unaffected.
- The explicit solution of the resulting ODE is actually significantly more complicated. We also see from the ODE that that the stationary state $p^\star = 1$ only exists when $\eta \to 0$, as $p'(t) = -\eta$.

Stability of the Stationary State

- The perturbation could equally be called another trajectory $\eta(t) := p(t) - p^\star = p(t) - 1$. Key: If $p^\star$ is stable, then $\eta(t)$ goes to zero as $t \to \infty$.
- Taking the derivative $d\eta(t)/dt$ leads to a condition on the derivative of the original ODE system, namely:

The stationary state $p^\star = 1$ is stable if $f'(p^\star) < 0$, where

$$f(p) := \gamma + (\delta - \gamma)p - \delta(p(t))^2$$

- But we know that

$$f'(p^\star) = (\delta - \gamma) - 2\delta p^\star = -\gamma - \delta < 0.$$

Social Interaction and Rules

- The Bass Model is a simple, yet powerful abstraction of real social interactions. Its main weakness is that it does not allow for individuals to interact–rather, the population acts 'on average' to describe the behaviours of innovation and imitation.
- By contrast, **agent-based models** (ABMs) are designed from the beginning to facilitate *strategic* interactions between individuals, known as *agents*.
- ABMs are useful across a wide spectrum of social phenomena, but can also be used as expert systems, classifiers, and even symbolic computers.
- This is because the agents in ABMs exhibit *rules-based* behaviour–if a condition (or conditions) are met, a rule (or rules) are executed. Rules can be both deterministic or random.

Rules-based Behaviour: The Prisoner's Dilemma

- Consider a set of agents with only two actions: *cooperate* and *defect*. Agents select one of these actions every time they meet another agent (agents only meet pairwise).
- An agent receives a *payoff* based upon their action and the action of the opposing agent.
- The highest payoff occurs when the agent selects *defect* when the other agent selects *cooperate*.
- But the agent receives the **worst** payoff when selecting *defect* while the other agent selects *defect*.
- Somewhere in the middle is the payoff when both agents select *cooperate*.

The Prisoner's Dilemma

- This action-to-payoff mapping can be summarised in a *payoff matrix*–the following is a representative example:

		Agent 2	
		C	D
Agent 1	C	$8, 8$	$-10, 16$
	D	$16, -10$	$-5, -5$

- So if e.g. Agent 1 selects "[C]ooperate" and Agent 2 selects "[D]efect", then Agent 1 receives a payoff of -10 (the first entry in the top right box) while Agent 2 receives a payoff of 16 (the second entry in that box).

The Prisoner's Dilemma–Static

- If the agents met only once, and selected their actions only once, then there is an equilibrium concept known as a *Nash equilibrium* which allows one to predict what the agents will do.

> The driving force behind the Nash equilibrium is that agents will do as well as they can, conditional upon what other agents do.

- Thus, for example, Agent 1 will select *defect* conditional upon Agent 2 selecting *cooperate*, because this is better than choosing *cooperate* when Agent 2 selects *cooperate*.
- The Nash equilibrium criterion selects only one set of actions as the equilibrium of the agent's interaction: *defect* for both agents! From our payoff matrix, each agent will receive a payoff of -5.

The Prisoner's Dilemma–Cooperate?

- The outcome of the static case seems artificial somehow–couldn't the agents figure out a way to cooperate instead?
- Indeed, there is a wealth of evidence that agents do just that: in the real world, somehow agents decide to cooperate.
- A large research agenda on e.g. trust, fairness, etc. focuses on these additional attributes of agents as a way to foster (or explain) cooperation.
- Other research has focused on allowing agent interaction to be *dynamic*.

Repeated Interaction and Cooperation

- Consider a network of agents who meet other agents randomly. We can think of this as an agent interacting with someone in their neighbourhood randomly, and engaging in the Prisoner's Dilemma game.
- Then the payoff of the interaction can be seen as affecting an **attribute** (such as wealth). The "sugarscape" model of Axtell and Epstein (Brookings Institution) is one such wealth-attribute network environment.
- In a similar vein as the sugarscape model, suppose that agents can "die" if their wealth falls below a threshold value.
- Will agents still select defection as the Nash equilibrium?

Repeated Interaction and Cooperation

- To answer this question we need a *rule* for our ABM, specifying the condition(s) under which an action is taken.
- Let the rule be:

> "An agent in a period t will cooperate *if* in period $t-1$ the opposing agent cooperated. But if in period $t-1$ the other agent defected, then in period t the agent will defect."

- This is known as the "tit-for-tat" rule and was first implemented by Anatol Rapaport.

Tit-For-Tat

- Let's see how this rule affects a simulated network created by 50 agents, where agents are coloured by strategy (orange for defect, blue for cooperate) and the first action is random.
- Each agent has an initial wealth of 50–if their wealth falls below zero at any time, they die and stop participating. We use our payoff matrix given earlier for each simulation.
- The following film shows 100 periods of interaction for a *complete network*, i.e. a network where every agent is connected to every other agent. Agents who have died are coloured black.
- Agents meet randomly through their neighbourhoods, and their initial wealth level is augmented (or diminished) according to their fortunes in the game. Run film

Tit-For-Tat

- As another example, consider a preferential attachment network (still 50 agents for 100 periods, initial wealth 50).
  Run film
- The results are about the same, although the preferential attachment characteristic of the network makes it more likely for agents to cooperate than to defect (why?).
- Let's keep the preferential attachment model, but only have one cooperator in a sea of defectors. Will we get a "cooperation infection"? A little thought will lead you to conclude that a single cooperator is very unlikely to survive.
  Run film
- How about a small group of defectors? A single defector will spawn a single defector indefinitely (why?), but will a group of, say, 10% of the total population increase, decrease, or stay the same? Run film

Mathematical Epidemiology

- **Networks describe interactions**, and it was realised in the mid-20th-century that valuable insight can be gained from studying the *transmission* of an attribute from one node to another, when an interaction is assumed to take place.
- Perhaps the earliest systematic study was in epidemiology, in which networks were (and are) used to describe the pathways for the spread of an infection, or a disease.
- Current research on e.g. SARS, MERS-CoV, HIV/AIDS, tuberculosis and malaria all rely heavily upon the specification of how diseases spread from person to person (or from person to animal to person).
- In turn, network models of mathematical epidemiology have proven valuable in advertising and marketing (for product information dissemination, product adoption etc.), online "meme" diffusion (from e.g. YouTube etc.), and financial contagion (e.g. the banking crisis of 2008-2012).

The SIR Model

- Consider a population of agents that are *susceptible* to a particular infection. If an agent is *infected*, they are contagious for some time and may spread the infection to others. Eventually, the agent either 1) *recovers* and is no longer infected, is no longer contagious, and cannot become infected again, or 2) dies and is *removed*.

  Such an environment is described by the Susceptible - Infected - Removed/Recovered (SIR) Model.

- The driving question that the SIR model can help answer is: given an initial infection, will there be a 'pandemic', i.e. an infection which is widespread?

The SIR Model: Definitions

- Consider a network $\mathcal{N}$ with N nodes $\{1, \ldots, N\}$ and an edge set $\mathcal{E}$. Time t is discrete.
- Each node i at time t has a **state** s_i^t, with $s_i^t \in \{S, I, R\}$, denoting [S]usceptible, [I]nfected and [R]emoved, respectively (here we treat recovered and removed nodes identically).
- An edge $e_{ij} \in \mathcal{E}$ indicates that nodes i and j are connected in a fashion that makes infection possible from one to the other (the edges are undirected).
- If a possible edge $e_{ij} \notin \mathcal{E}$ then i cannot directly infect j, and vice-versa.

The SIR Model: Definitions

- For simplicity we assume that the probabilities of making a *transition* from one state to another are independent of the topology of the network.
- This is an abstraction and is not usually true.
- Given an edge $e_{ij} \in \mathcal{E}$ and $s_i^t = I$, i.e. node i is infected, there is a probability that node j will become infected which depends upon the current state s_j^t:

$$\mathbb{P}(s_j^{t+1} = I \mid s_j^t, s_i^t = I) = \begin{cases} 0 & \text{if } s_j^t = R, \\ \chi & \text{if } s_j^t = I, \\ \phi & \text{if } s_j^t = S. \end{cases}$$

- χ and ϕ are both positive and constant–the relationship between the two drives the dynamics, much as the relationship between γ and δ in the Bass Model defined *its* dynamics.

The SIR Model: Definitions

- To close the model we must specify other transitions from the states S, I and R, formalising the intuition that $S \to I \to R$.
  1. First, we denote the probability of removal once infected by $\mathbb{P}(s_j^{t+1} = R | s_j^t = I) = 1 - \chi$.
  2. This implies that $\mathbb{P}(s_j^{t+1} = S | s_j^t = I) = 0$, i.e. an infected node does not spontaneously become susceptible [this is the main distinguishing feature separating the SIR model from the Susceptible-Infected-Susceptible (SIS) model, discussed shortly].
  3. We assume that spontaneous infection is impossible, i.e. $\mathbb{P}(s_j^{t+1} = I | s_j^t = S, s_i^t \in \{S, R\}) = 0, \forall i$.
  4. Next, $\mathbb{P}(s_j^{t+1} = R | s_j^t = R) = 1$, i.e. a removed node stays removed.
  5. Finally, $\mathbb{P}(s_j^{t+1} = R | s_j^t = S) = 0$, i.e. a susceptible node is not removed unless it has been infected first.

The SIR Model: Dynamics

- The network $\mathcal{N}$ is initialised with a random assignment of one infected node ("patient zero").
- Each time step:
  1. All edges e_{ij} with one infected node are checked to see if the other node is susceptible–if it is, then it is infected with probability ϕ.
  2. All infected nodes are checked to see if they remain infected, with probability χ, or if they are removed, with probability $1 - \chi$.
- The model can be analysed analytically and numerically via simulation.

Simulating the SIR Model

- The SIR Model is simulated by creating a network of relationships and then infecting one individual.
- The propagation of the illness depends upon the probabilities χ and ϕ.
- Let's examine an SIR model using a built-in social network from `networkx`, the *Karate Club Graph*, with $\chi = 0.9$, $\phi = 0.1$. Run film
- By contrast, let's look at a more highly connected random graph, the *Newman-Watts-Strogatz* small world graph, with the same χ and ϕ values. Run film

The SIS Model

- A richer characterisation of many infectious diseases, such as the common cold or influenza, can be obtained by relaxing the assumption of a removed state.
- Rather, there are only two states, infectious (I) and susceptible (S). Nodes which have been infected in the past but are now susceptible may be re-infected.

  This is known as the Susceptible - Infected - Susceptible (SIS) Model.

- To investigate this richer environment we will be more formal than our treatment of the SIR model and examine the model analytically.

The SIS Model: Definitions

- As a first approximation, assume that the probability of infection is dependent upon the *degree* $d_i = d$ of a node i:

$$\mathbb{P}(s_i^{t+1} = I | s_i^t = S) = \phi dp,$$

  where ϕ is, as before, the probability of becoming infected conditional upon meeting an infected node, and p is the probability of meeting an infected node (so that the product dp is the fraction of nodes connected to i which are infected).
- For dp to make sense it must be the case that (on average) nodes are highly connected, so that $d \gg 0$.
- Note also that $\phi \max_i\{d_i\} \Leftrightarrow \phi d_{\max} < 1$ must hold, otherwise $\mathbb{P}(s_i^{t+1} = I | s_i^t = S)$ may not be a probability.

The SIS Model: Definitions

- To derive p it is necessary to examine the characteristics of the network $\mathcal{N}$. We need the *degree distribution* $P(d)$, and the fraction of nodes of degree d who are infected, $i(d)$.
- Then the probability of meeting an infected individual is simply the probability of meeting an infected node of degree d, summed over all degrees in the population:

$$p := \frac{\sum_d i(d)P(d)d}{\overline{d}},$$

  where the **average degree** $\overline{d} := \sum_d P(d)d$.
- To close the model, we assume that individuals recover from being infected with a probability

$$\mathbb{P}(s_i^t = S | s_i^t = I) := \mu.$$

The SIS Model: Dynamics

- The model has many interesting analytical features which can be addressed using network analysis.
- Question: is there a steady-state level of infection in the population? Assume the number of nodes is finite.

  Answer: The only steady-state for infection is 0, i.e. everyone is healthy.

- Question: Now suppose we have an infinite number of agents. Is there a steady-state level of infection now?

  Answer: This depends upon the relative impact of ϕ vs. μ.

The SIS Model: Steady State Dynamics

- Assuming an infinite number of agents, we can use an equilibrium condition to estimate the fraction of nodes that are infected, $i(d)$, for a steady state:

$$i(d)\mu = (1 - i(d))\phi dp \Rightarrow$$
$$i^{\star}(d) = \frac{\beta pd}{\beta pd + 1},$$

where $\beta := \phi/\mu$.

The SIS Model: Steady State Dynamics

- Using what we already know about p, we find the following:

$$p^\star = \sum_d \frac{p^\star P(d)\beta d^2}{\overline{d}(p^\star \beta d + 1)}.$$

- This is a nonlinear equation with solutions which depend upon β. For a given degree distribution, this equation can have zero, one or more steady-state levels of infection.
- For example, consider a regular network $\mathcal{N}_r$, where the degree of every node is the average degree of the network, $\overline{d}$.
- Then the steady state infection rate $p^\star$ has two possible values:

$$p^\star \in \{1 - \frac{1}{\overline{d}\beta}, 0\}.$$

The SIS Model: Steady State

- One steady state solution is the 'trivial' solution, where the infection rate is zero and hence never deviates from that point.
- The other steady state connects the average degree $\overline{d}$ to the ratio of the transmission-to-recovery ratio β. Note that $p^\star > 0$ only when

$$\overline{d} > \frac{1}{\beta} = \frac{\mu}{\phi},$$

  i.e. when the number of connections for each individual is sufficiently high that it can overwhelm the relative (to transmission) recovery rate.
- Otherwise, the relative recovery rate is too high to sustain an infection, and it dies out. This is demonstrated in the following numerical plot of $p^\star$ vs. β:

The SIS Model: Steady State

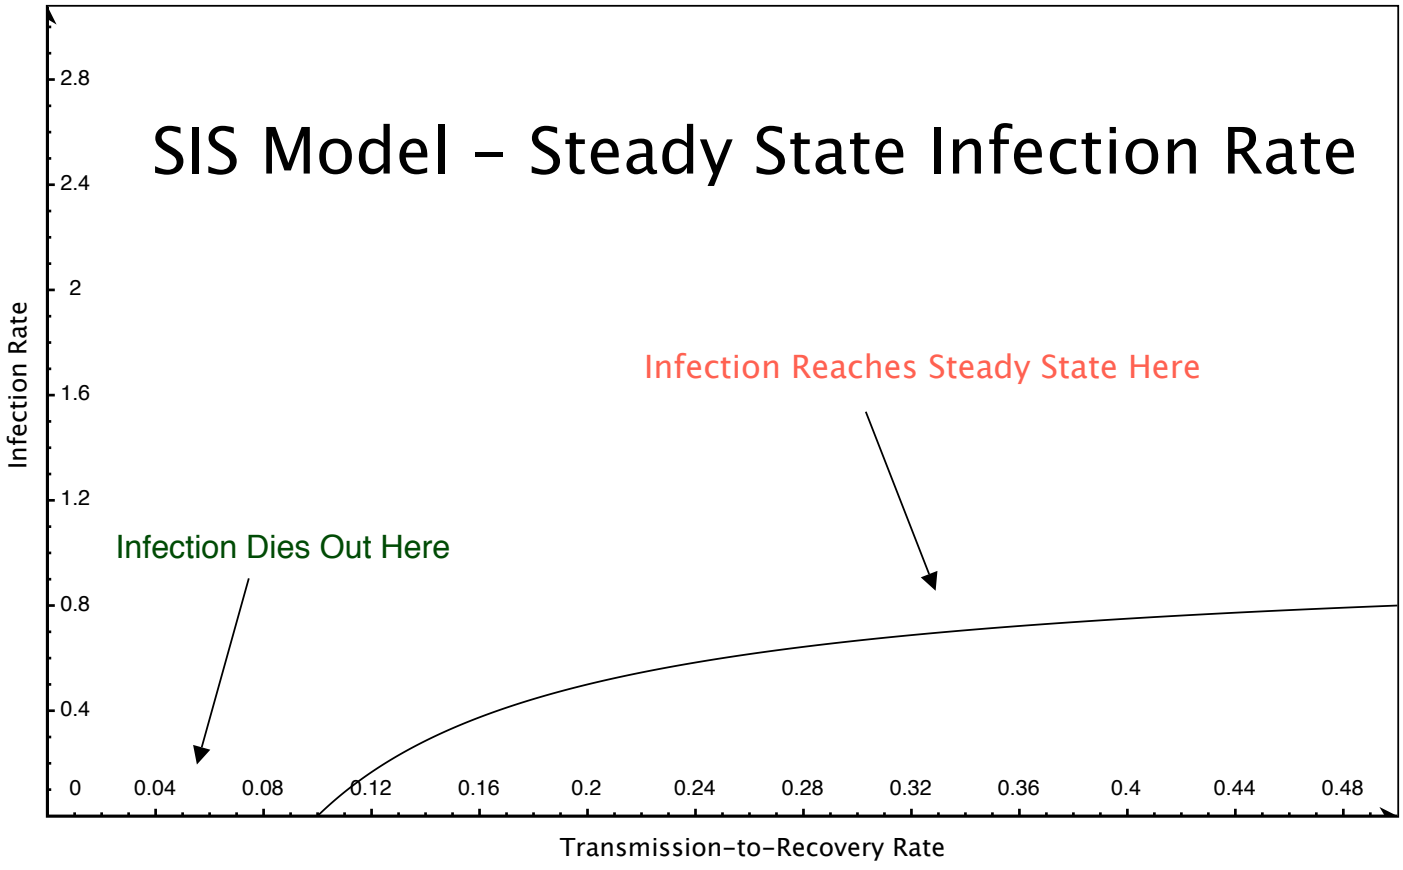

The SIS Model: Simulation

- To simulate the SIS model we use a large power law graph of 500 nodes, with an average degree $\overline{d} = 50$. We start with $phi = 0.6$ and $\mu = 0.6$, implying that $\beta = 1$ and hence that $p^{\star} = 1 - 1/ < d > \simeq 0.98$, which in turn implies $i^{\star}(d) = 0.98$. Hence around 490 nodes should be infected in the steady-state. Run film
- As seen from the simulation, the steady-state prediction is largely confirmed– there are around 460 infected nodes.
- Note that this is *not* proof! Simulations are suggestive but not conclusive to study steady state properties of dynamical systems.

The SIS Model: Degree Distribution

- In order to get a steady state solution, we assumed that the network was regular.
- In addition, the transmission rate ϕ was assumed to be much smaller than the largest degree value $d_{\max}$.
- Neither assumption is particularly realistic, and as a final example they will be relaxed.
- First, let us suppose that there is a simple degree distribution $P(d)$, with two degree categories: d_1 and d_2.
- Half of the nodes have degree d_1 and half d_2, i.e. $P(d_1) = P(d_2) = 0.5$.
- Finally, let $d_1 \gg d_2$.

The SIS Model: Degree Distribution

- Next, we relax the assumption that $\phi d_{\max} \leq 1$ and replace the probability of infection defined earlier by:

$$\mathbb{P}(s_i^{t+1} = I | s_1^t) = \phi d p + \frac{1}{2}(\phi p)^2 d(d-1).$$

- Although this is rather mysterious, it is in fact just the next term in the expansion of the general Binomial distribution of meeting an infected individual, conditional upon a degree d.
- This now allows for two degrees of freedom, i.e. pairs of numbers (ϕ, d) can be selected to satisfy this probability.

The SIS Model: Non-Steady-State

- Armed with these relaxed assumptions, we can rewrite the law of motion for the infected proportions of the population.
- In actuality the system is now *coupled* for $i(d_1)$ and $i(d_2)$, but we will abstract away from $i(d_2)$ for the sake of exposition and focus only upon $i(d_1)$:

$$\Delta i^t(d_1) := i^{t+1}(d_1) - i^t(d_1) \simeq (1 - i^t(d_1))\phi d_1 i^t(d_1) + \frac{1}{2}(\phi)^2 d_1(d_1 - 1)(i^t(d_1))^2 - \mu i^t(d_1).$$

- This is the non-steady-state law of motion for $i(d_1)$. Like the steady state equilibrium equation, it specifies inflows to and outflows from the infected population.

The SIS Model: Chaotic Dynamics

- We can reduce the complexity of the system by looking at one specific set of parameter values (recall we have two degrees of freedom to choose from).
- The reduced form of the law of motion becomes:

$$i^{t+1}(d_1) = 3.5(1 - i^t(d_1))i^t(d_1) + 0.1(i^t(d_1))^2 + 0.19i^t(d_1).$$

- This has a *very* interesting form. The first term is a so-called "logistic" equation and describes, among other things, the interaction between predator and prey species. The second and third terms are the nonlinear and linear, respectively, "forcing" terms–changes in $i(d_2)$ are actually propagated here (although they are not written).

The SIS Model: Chaotic Dynamics

- How does this system behave? We can plot the sequence of infected proportions $i^0(d_1), i^1(d_1), \ldots, i^t(d_1), i^{t+1}(d_1), \ldots$ and see if there appears to be any convergence:

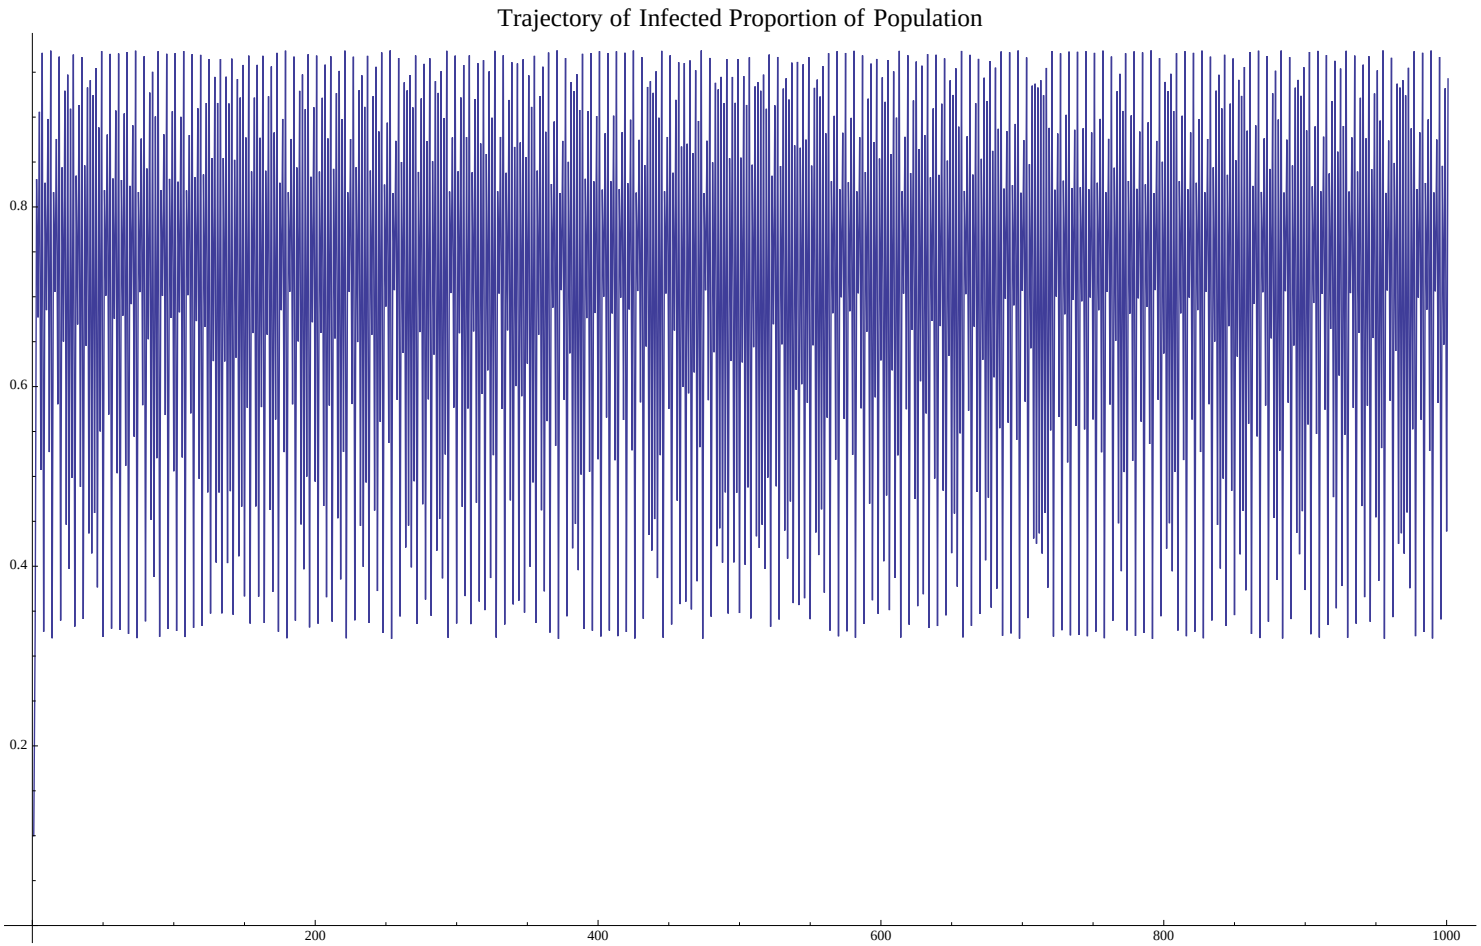

The SIS Model: Chaotic Dynamics

- There is no convergence. Rather, the trajectory of the infected population appears to wander within a subset of the $(0, 1)$ interval, never quite wiping out the population (where $i(d) \to 1$) but never dying out, either–in fact, around one third of the population is *always* infected.
- The trajectory is *chaotic*, and the SIS system described here exhibits **chaotic dynamics**.

> Chaotic systems exhibit a lack of stationary state, random trajectories, and sensitive dependence on initial conditions.

- For continuous time systems at least three degrees of freedom are required, but for discrete time systems one or two is often enough.
- Note that we have said nothing about the "reasonableness" of the parameters chosen to exhibit chaos.

The Banking Crisis

- The banking crisis of 2008-2012 highlighted how interconnected global financial markets really are.
- In this case, the interconnectedness resulted from *mutual indebtedness* of financial institutions to each other.
- Such indebtedness was due to banks borrowing from each other, using (in many cases) loans to home owners as collateral.
- When the value of these loans fell–the so-called "subprime mortgage crisis"–the value of bank debt collateral also collapsed.
- This collapse sent 'shock waves' through the network of financial institutions, creating **contagion** as many banks became worth less than the debt they owed, and in turn requested that other banks pay back what *they* were owed.

The Banking Network

- We can create an abstract network representation of the banking sector by assuming, first, that banks are nodes.
- Each bank has an attribute, its **balance sheet**, showing the excess of assets over liabilities. If a bank's debt is called by another bank, it is distressed. Otherwise it is either healthy, or it is bankrupt, i.e. in default–the sum of the debt calls exceeds the assets of the bank.
- Each weighted edge between nodes is a measure of indebtedness: $e_{ij} \in \mathcal{E}$ if bank i is in debt to bank j. We assume for simplicity that each bank owes the same amount on their debt.
- Finally, each node has a set of probabilities: the probability of becoming distressed, and the probability of going into default, conditional upon being distressed. Bankrupt banks do not participate further.

The Banking Network

- In this way we have modelled the banking sector as an SIR model–a healthy bank is Susceptible to becoming distressed, i.e. Infected. And when a bank defaults, it is Removed from the network.
- The *economics* is contained in how banks transmit distress–when a bank becomes distressed it 'calls in' their loans, i.e. it requires that those banks who are indebted to it must pay their debts.
- Given an initial distribution of banks and their indebtedness, the dynamic network can be simulated.

Banking Network Simulation

- First consider a scale-free directed network with 100 nodes. When a single node is distressed, the contagion spreads to other nodes far removed from the original location. Ironically, the first distressed node is likely to survive as it calls in its loans before the panic begins. Run film
- Another scale-free example demonstrates how a core nucleus of nodes cannot sustain the multiple requests to call in their debt, and they go bankrupt. Survivors include banks on the periphery. Run film
- As a final example, consider a directed 'growing graph' with the same number of agents. A growing graph is similar to preferential attachment in that nodes are added sequentially with an edge formed randomly based upon degree. There are few pathways between clusters of financial institutions: one solvent bank is enough to prevent the contagion from spreading to an entire sub-cluster of the graph. Run film

Meme Diffusion

- A *meme* is a small amount of media content which propagates via social networks, usually connecting to that network's social self-perception or embedded social norms, traditions, etc.
- Although memes have always existed, the rise of social media in the mid-2000s has greatly increased their scope and diffusion.
- As two well-known examples we have rick-rolling, the surreptitious hyperlinking to a 1980s pop tune rather than the link's stated destination, and the subtitling of a clip from the film Der Untergang, during a particularly intense scene depicting Adolf Hitler.
- In all cases, a meme is characterised by its **speed of transmission** and its **network penetration**–it rapidly diffuses through a network and is known by a large fraction of the network.

The Meme Diffusion Network

- We can create a network to examine memes by assuming, first, that individuals are nodes.
- Each node has an attribute, its knowledge of one (and only one) meme. If a meme A is known, the node's attribute is A. Otherwise it is considered to be uninterested in the memes it is exposed to.
- Each edge between nodes denotes a relationship between individuals–e.g. linked as friends on Facebook, joint followers on Twitter, etc.
- We assume for simplicity that the graph is undirected, but a richer model would assume that friends/followers are directed and hence individuals with 'meme influence' could be identified.

The Meme Diffusion Network

- Finally, we associate a *rule* to each node–given a neighbourhood $\mathcal{N}_i$ of node i, node i will on occasion transmit the most popular meme in $\mathcal{N}$. The 'on occasion' part means we allow for a $1 - \mu$ chance that agent i is uninterested in passing on any of the memes listed in $\mathcal{N}$, for a parameter μ.
- This is a network similar to the models of strategic interaction discussed earlier–the rule is in place to allow the individual (i.e. the agent) to respond to conditions as they arise in their social setting (here, their neighbourhood).
- Given an initial distribution of individuals and their social relationships, the dynamic network can be simulated and the ABM analysed.

Meme Diffusion: Simulations

- For example, here is a simulated Barábasi - Albert network of 100 agents, when an agent passes on the most popular meme in their neighbourhood with probability $\mu = 0.5$. Run film
- And here is a Newman-Watts-Strogatz network of 200 agents, where $\mu = 0.1$. The lower value of μ can 'stem the tide' of a popular meme, giving lesser memes more time to propagate, but in the end 'there can be only one'. Run film
- Finally, here is an Erdős - Rényi random graph with 200 agents, with an edge creation probability of 0.1. With relatively few edges, the first meme to meet a well-connected node has an early-mover advantage. Run film

The Bass Toy Model: Implicit Solution

<1><2>(1cm, 4cm)(0.5cm, 2cm)

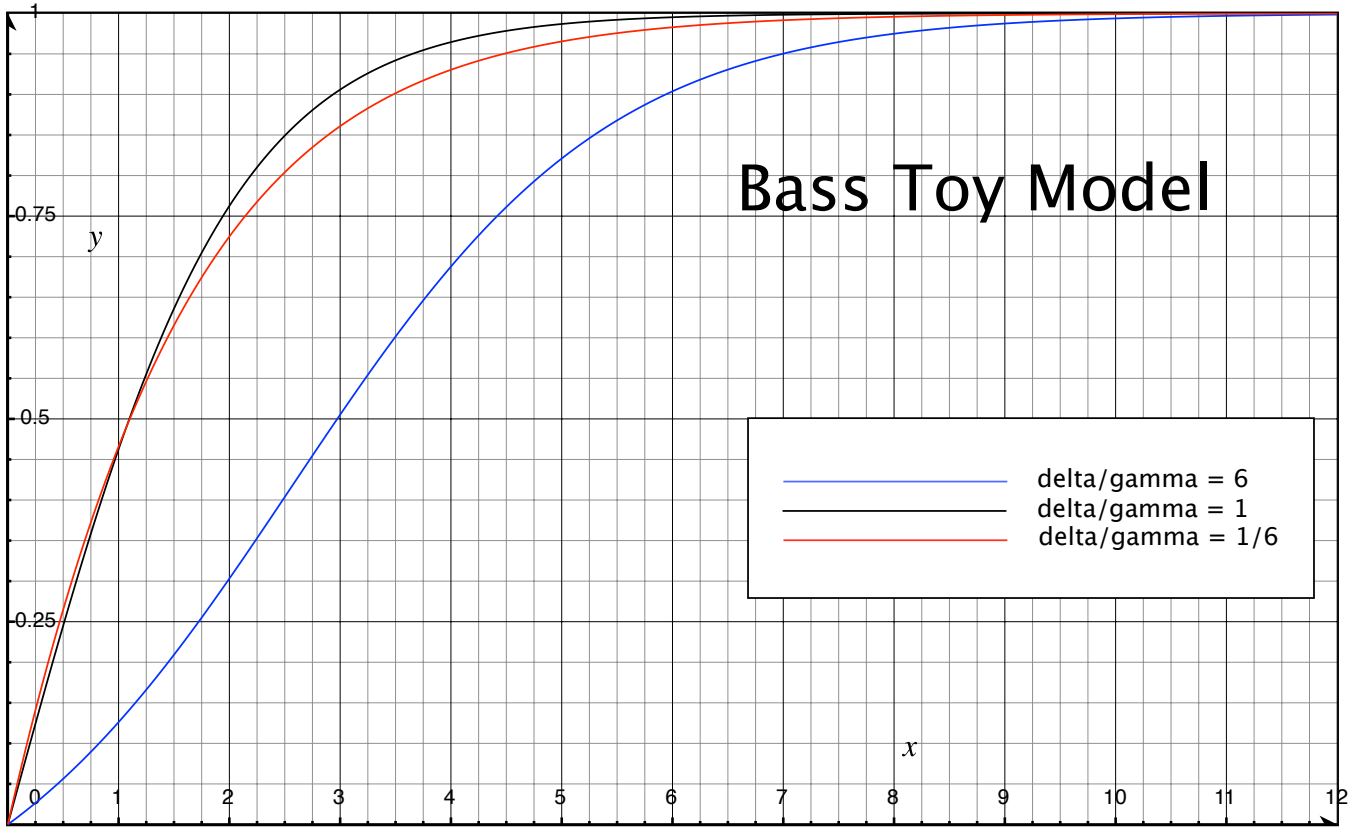

- Note here that the shape of the implicit solution remains essentially the same, but shifts according to the *fraction* δ/γ.