Social and Information Networks

University of Toronto CSC303
Winter/Spring 2019

Week 8: March 4,6 (2019)

Announcements and agenda

Announcements

- The lecture and tutorial schedule for this week and next week.
  - Lectures today and Wednesday as usual
  - No tutorial this Friday
  - Tutorial on Monday, March 11
  - Lectures on Wednesday and Friday, March 13 and 15.
- So far only about 20 proposals for critical review assignment. If you submitted a proposal and it was not acknowledged please resend.
- I repeat the regrading policy: Requests for regrading must be done within a week of the assignment being returned. For Assignment 1, we will accept requests up to this Wednesday, March 6 (11:59PM).

Todays agenda.

1. We will first review the direct benefits model up to where we left off last week. .
2. Discrete step dynamics to reach an equilibrium
3. Chapter 19: Cascades in a network

Market with large number of producers and quick review of the direct benefit effect model

- The model assumes that there is some (say industry wide) cost p^* at which a unit of the good can be produced.
  - Perhaps to make this more realistic, assume this cost includes an industry wide small profit/unit
  - In any case we are assuming that no producer is willing to supply the good at price below p^* per unit of good.
- Another (more substantial) assumption:

> There are enough producers capable of producing an unlimited supply of the good and no single producer can change the market. Implicitly the goods are identical, independent of the producer.

- Thus in aggregate these producers can supply as much of the good as desired at price p^* per unit but will not produce any of the good at price below p^* per unit.
- This also fixes the price at p^* since by assumption competition will not allow any producer to ask for more than p^* per unit.

Intrinsic value function and direct benefits function

- $r(x)$ is the intrinsic value that consumer x has for the good and we are assuming $r(x)$ is decreasing and continuous in $x \in [0, 1]$ and $r(1) = 0$.

- We are assuming that the reservation price of consumer x is $r(x)f(z)$ when there is a fraction z of existing users of the good where $f(z)$ is increasing and continuous in $z \in [0, 1]$. And (for now) assume $f(0) = 0$.

- So now a consumer x is willing to buy a unit of the good at price p^* if x believes a fraction z of users will also be using the good and $r(x)f(z)$ is at least p^*.

What if everyone makes a perfect (shared) prediction?

Self-fulfilling expectations equilibrium

If everyone makes the same prediction about the fraction z buying the good, and then every consumer x acts on this assumption and decides to buy based on whether or not $r(x)f(z)$ is at least p^*, then (eventually) the fraction of adopters will actually be this z.

- This z is called a self-fulfilling expectations equilibrium for the quantity z (at price $p^* > 0$).
- For a fixed z, as x increases, $r(x)f(z)$ decreases, so we have:

Fact

If $p^* > 0$ and z in $(0, 1)$ is a self fulfilling expectations equilibrium at p^*, then $p^* = r(z)f(z)$. Why? By the assumption that $f(0) = 0$, $z = 0$ is also a self-fulfilling expectations equilibrium.

- This is a more complex (and more interesting) situation than without direct benefits in which case high prices simply imply low demand.

The concrete case of $r(x) = 1 - x$ and $f(z) = z$

- As an example of the model, the text considers the decreasing reserve price (intrinsic value) function $r(x) = 1 - x$ and the increasing direct benefit function $f(z) = z$.
- Then in addition to $z = 0$, a self-fulfilling expectations equilibrium $z > 0$ must satisfy $p^* = (1 - z)z$.

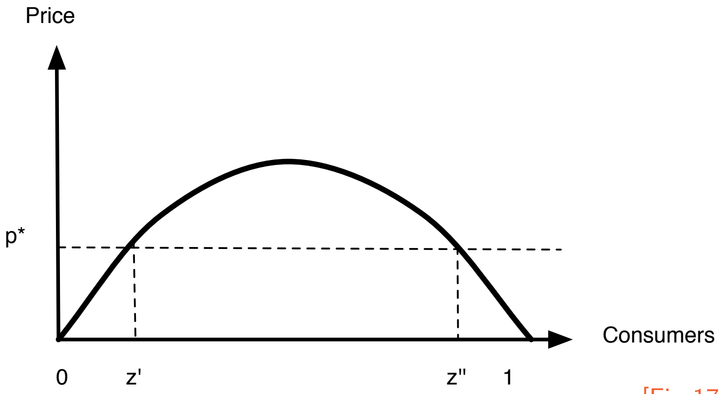

[Fig 17.3, E&K]

What are the equilibria for this example?

- By taking the derivative of $h(z) = r(z)f(z)$, we see that $h(z)$ has maximum value at $z = \frac{1}{2}$ (and hence $h(z) = \frac{1}{4}$) so that for $p^* > \frac{1}{4}$ there is no (real valued) solution to $p^* = r(z)f(z)$
- The case $p^* = 0$ is not interesting; we will soon consider the special case $p^* = \frac{1}{4}$.
- For any p^* in $(0, \frac{1}{4})$, there are exactly two distinct zeros z', z'' and at the points $z = 0$, z', z'', if everyone believes exactly a z fraction will be buying according to the reservation price, then precisely this fraction will do so.

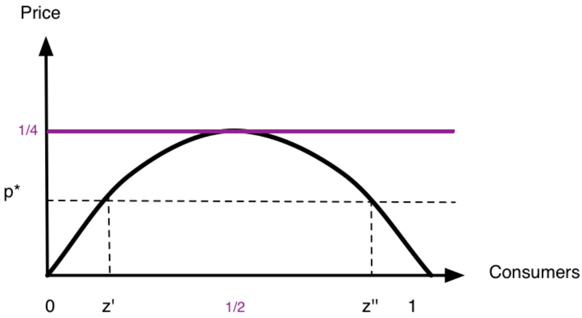

Why can't there be other equilibria?

- What happens when the demand z is not one of these equilibria points z', z'' (for a price $p^* < \frac{1}{4}$)?

- Three cases:
  1. If $0 < z < z'$, then $r(z)f(z) < p^*$ and there is downward pressure on the demand since the reservation price is less than p^*.
  2. If $z' < z < z''$, then there is upward pressure on demand since $r(z)f(z) > p^*$ and more purchasers are willing to buy.
  3. If $z'' < z$ then we again have $r(z)f(z) < p^*$ causing downward pressure on the demand.

- Note the qualitative difference between z' and z''.
  - Values of z near z'' will push the demand toward z''. That is, z'' is a very stable equilibrium.
  - In contrast, demand predictions around z' are very unstable in that the demand pressure can go either way.

More qualitative comments re equilibria

- The unstable equilibrium point z' is called a critical or tipping point. It is indeed critical for the producers to get past this tipping point in the demand.

- As the price p^* is lowered, the critical point z' (in this reasonably illustrative example) gets lowered and the eventual demand gets larger moving toward demand z''. This is why it is often in the interest of a company to lower initial prices to get past the tipping point.

- We now return to the special case of $p^* = \frac{1}{4}$. Now there is just one non zero equilibrium at $z = \frac{1}{2}$. Following the reasoning for the case of $0 < z < z'$, any deviation from $z = \frac{1}{2}$ will result in downward pressure so that this equilibrium is highly unstable.

What if everyone does not make a perfect (shared) prediction?

- We are now changing the story as to perfect shared predictions (but not the model). Assume that we are now tracking participation in a given activity, say a large online social network, or television series, or involvement in a political movement.
  - We view such participation as being more fluid than buying an object unless the cost for the object is minor.
  - In the case of participation there are maximum costs p^* (monetary or effort or reputation etc which can all be seen to be ultimately "costs") that a person will pay.

- We maintain the same model that x will participate if and only if $r(x)f(z)$ is at least p^*.

Discrete step dynamics

- We will assume that at some initial time $t = 0$, the observable demand level is z_0. This will cause the demand to change to some z_1 at time $t = 1$ and similarly the demand then changes to z_2 at time $t = 2$, etc.
- That is, if everyone observes demand z at some point of time, then the set of people participating at the next time step will be all those people x in $(0, \hat{z}]$ where $\hat{z}$ satisfies $r(\hat{z})f(z) = p^*$. (Recall that r(x) is continuous and decreasing in x.)
- That is, the next demand level will be the $\hat{z}$ satisfying:

$$\hat{z} = g(z) = r^{-1}\left(\frac{p^*}{f(z)}\right)$$

Since $r()$ is continuous and decreasing, such a solution will exist as long as $p^*/f(z)$ is at most $r(0)$. Otherwise $g(z) = 0$.

The same specific case of $r(x) = 1 - x, f(z) = z$

- For definiteness we will again take the specific case of $r(x) = 1 - x$ and $f(z) = z$. Hence $r(0) = 1$ and $p^*/f(z) = p^*/z$. So we want the shared demand observation z to satisfy $p^*/z = p^*/f(z) \le r(0) = 1$ or equivalently that z is at least p^*.
- It is easy to verify that $r^{-1}(y) = 1 - y$
- Hence in this case, $g(z) = r^{-1}(\frac{p^*}{z}) = 1 - \frac{p^*}{z}$ when $z \ge p^*$, and $g(z) = 0$ otherwise.

Change in demand for $r(x) = 1 - x$, $f(z) = z$

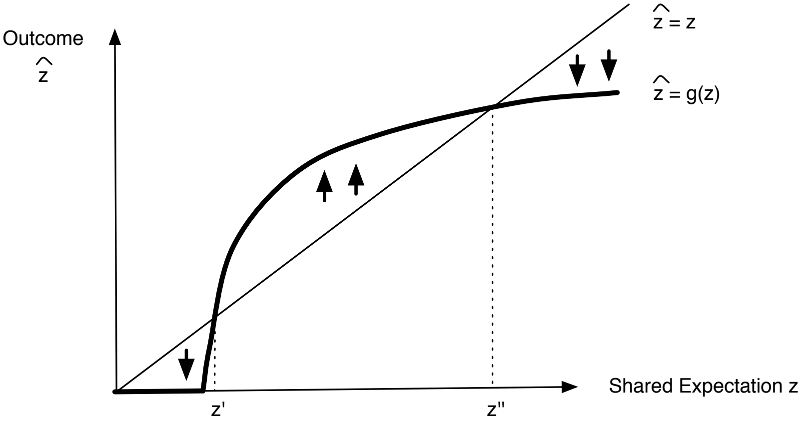

Figure: Fig 17.5 in E&K. For $r(x) = 1 - x, f(z) = z$, the curve for $g(z)$ is depicted. Here $g(z) = 1 - p^*/z$ if $z \geq p^*$ and 0 otherwise. When the curve crosses the line $z = z^*$, we have self-fulfilling expectations equilibria. Downward (resp. upward) arrows indicate downward (upward) pressure on demand. Here one can visualize why z' is unstable while z'' is a stable equilibtrium.

What we expect to see more generally

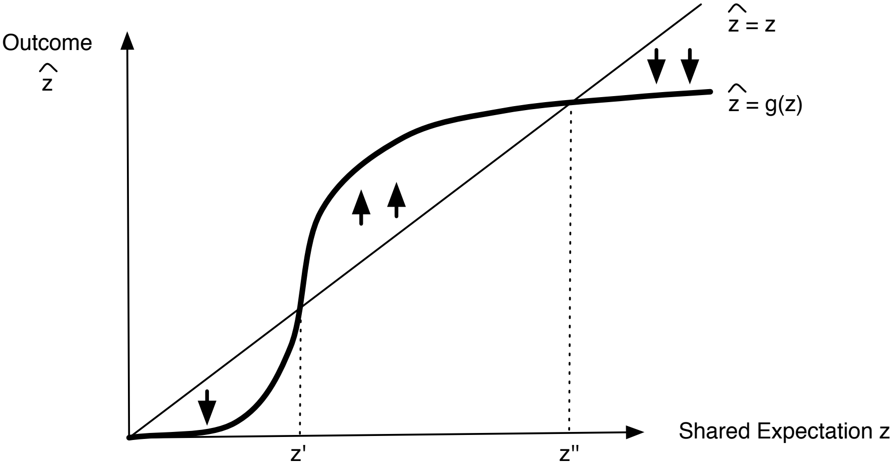

Figure: More generally, we would expect a smoother demand curve below the equilibrium z'.

[Fig 17.6, E&K]

Discrete step dynamic behaviour

Starting at some initial observable demand z_0, we generate future demands according to $z_{t+1} = g(z_t)$ for each time step $t = 0, 1, 2, \ldots$

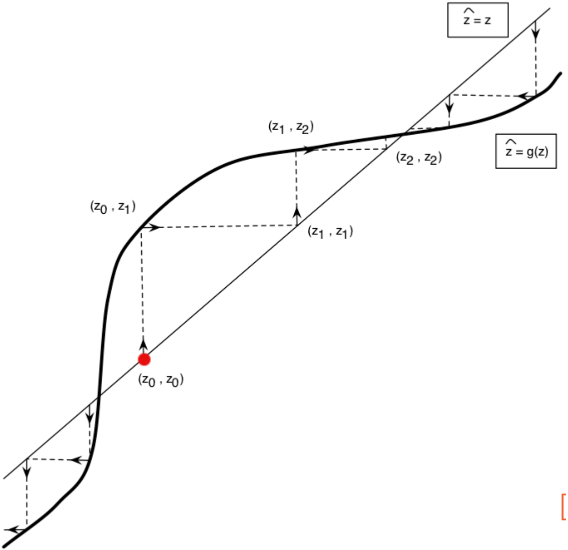

[Fig 17.9, E&K]

Discrete step dynamic behaviour

Starting at some initial observable demand z_0, we generate future demands according to $z_{t+1} = g(z_t)$ for each time step $t = 0, 1, 2, \ldots$

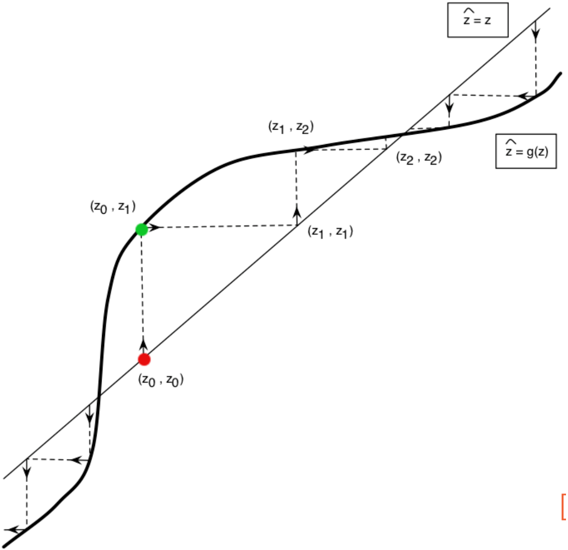

[Fig 17.9, E&K]

Discrete step dynamic behaviour

Starting at some initial observable demand z_0, we generate future demands according to $z_{t+1} = g(z_t)$ for each time step $t = 0, 1, 2, \ldots$

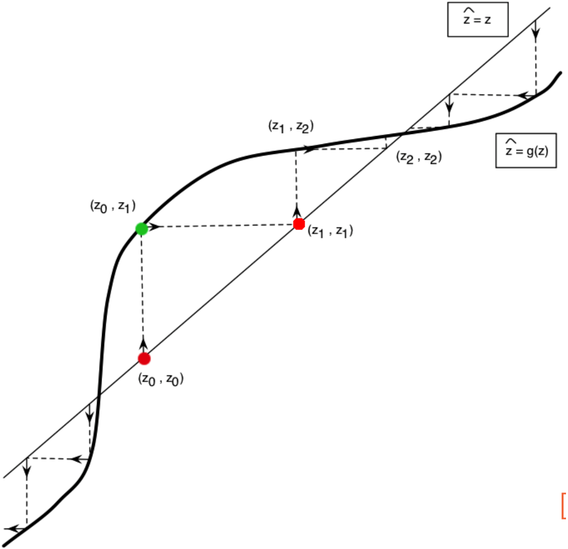

[Fig 17.9, E&K]

Discrete step dynamic behaviour

Starting at some initial observable demand z_0, we generate future demands according to $z_{t+1} = g(z_t)$ for each time step $t = 0, 1, 2, \ldots$

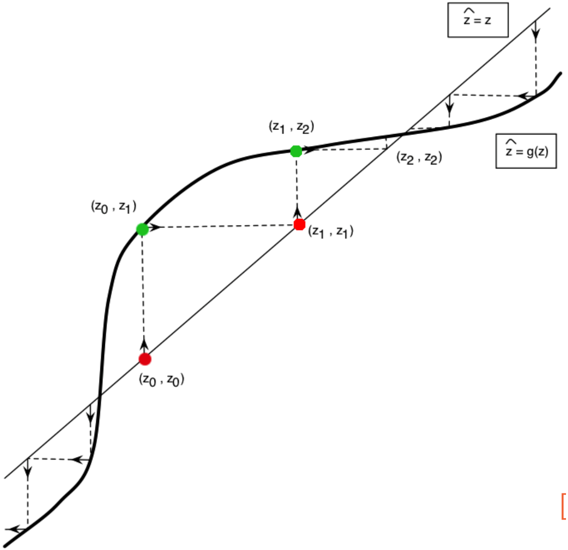

[Fig 17.9, E&K]

Discrete step dynamic behaviour

Starting at some initial observable demand z_0, we generate future demands according to $z_{t+1} = g(z_t)$ for each time step $t = 0, 1, 2, \ldots$

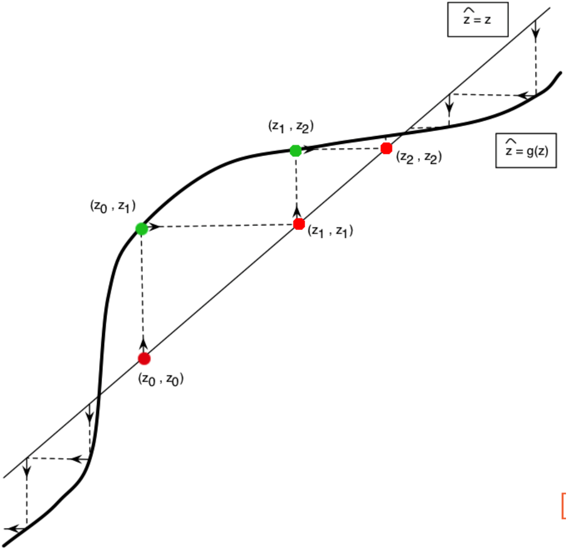

[Fig 17.9, E&K]

Concluding the discussion of Chapter17. What happens when $f(0) > 0$?

One of the model assumptions up to now is that $f(0) = 0$. We have been assuming that the reservation price for consumer x is $r(x)f(z)$ when $f(z)$ is the direct benefit factor for having a fraction $z \in [0, 1]$ of existing users of the good. This implies that $x = 0$ is an (uninteresting) equlibrium..

While this may be the case for say a FAX machine (assuming only one is ever sold), in general most goods have some (perhaps very small) value even without anyone else having it. So what happens when $f(0) > 0$?

The text again considers the example $r(x) = 1 - x$ but now assumes that $f(z) = a + z^2$ for some $a > 0$.

Now $z = 0$ is no longer an equilibrium point. In the specific example the new $g(z) = 1 - \frac{p^*}{a+z^2}$. As the text deomonstrates, the dynamics is essentially the same but now small changes in the price p^* can cause large changes in the equlibria that will be reached starting at $z = 0$.

Visualizing the dynamics for $f(0) > 0$

Letting $r(x)$= 1-x and $f(x) = a + z^2$.

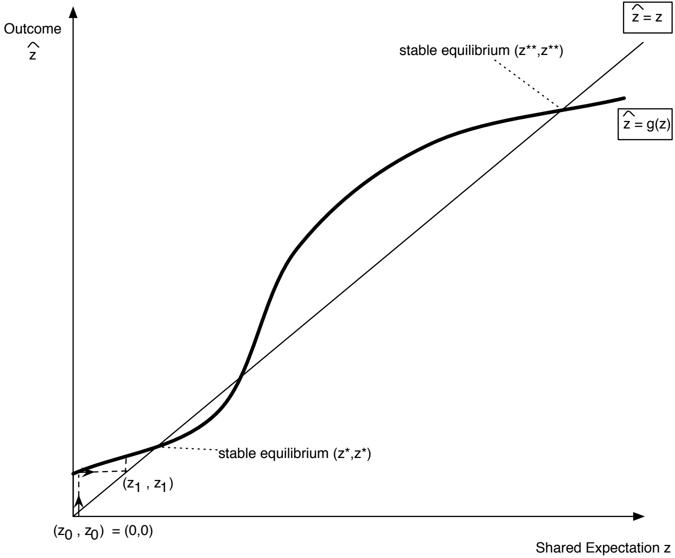

Visualizing the dynamics for $f(0) > 0$

Letting $r(x)$= 1-x and $f(x) = a + z^2$.

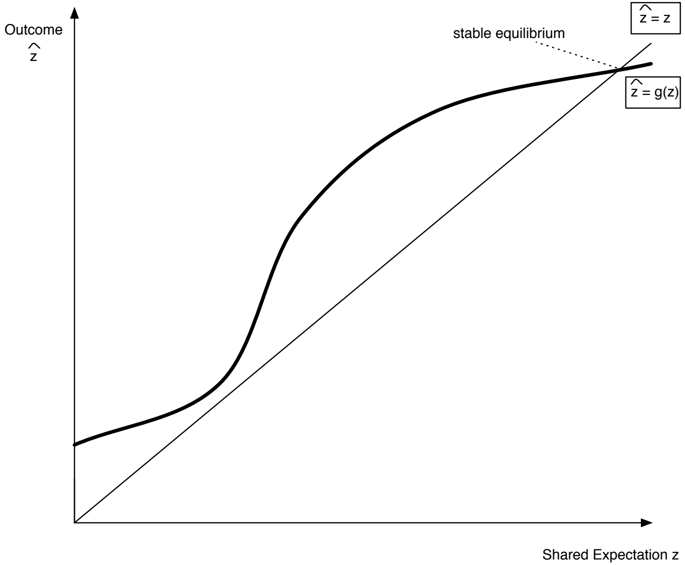

Visualizing the dynamics for $f(0) > 0$

Letting $r(x)$= 1-x and $f(x) = a + z^2$.

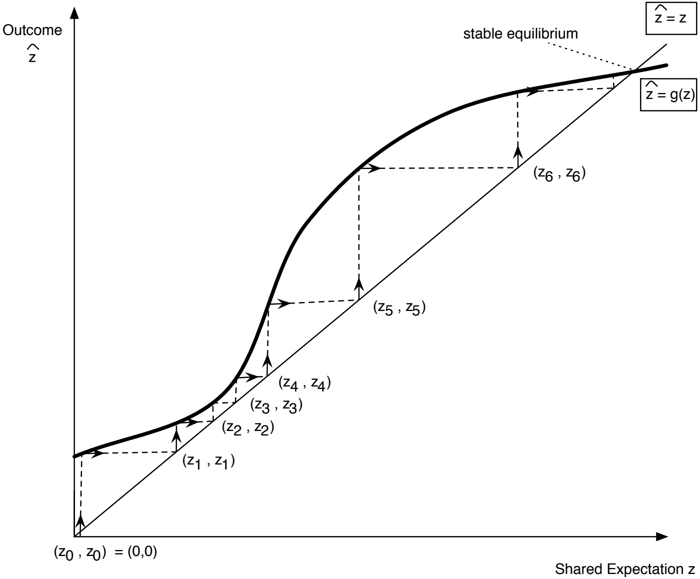

Chapter 19: Influence spread in a social network

- We begin a study of the spread/diffusion of products/influence in a social network (Chapter 19) in contrast to population wide spread phenomena as studied in Chapters 16, 17 and 18.

- The goal (as before) is to qualitatively understand the process in a highly stylized (but hopefully still interesting) setting.

- We will (as usual) be interested in what kind of general conclusions can be inferred from such an understanding?

Recap: population wide effects

- In Chapters 16 (herding or informational effects), 17 (direct benefit effects), and 18 (rich get richer models) we did not have a social network per se.
- These chapters dealt with population wide effects. Although :
  - One can construe Chapter 16 as taking place in a network where the ith individual is connected to all $i-1$ previous individuals.
  - Chapter 17 can be construed as taking place in a network where everyone is directly connected (the network is a complete graph).
  - Chapter 18 studies random processes by which networks can grow and and one can think of situations where the resulting network is a social network.
- But still . . . these are basically population wide effects absent from an existing social network.

Social network effects

- Now we wish to consider an existing social network where edges (ties) between individuals represent some sort of friendship/relationship.
- This takes us back to concepts introduced in Chapters 3 and 4.
- There we saw the contrast between
  - homophily (we tend to be friends with people of similar backgrounds, geography, interests)
  - social influence (we join clubs, are influenced) by our friends/relations.

Models of influence spread/diffusion

- One of the most important themes of the text (and CSC 303) is that we construct models to gain insight.
  - Our models are often (maybe always) very simplified given the complexity of real social and economic networks.
  - There is always a tradeoff between the adherence to reality and our ability to analyze and gain insight.
- How we model diffusion in a social network will clearly depend on what product, idea, membership, etc. we are considering.
- There are many assumptions as to how products, ideas, influence are spread in a social network and what are the set of individual alternatives.
- The main emphasis in Chapter 19 is on a very simple process of diffusion where each person has 2 alternative decisions:
  1. stay with a current "product" B
  2. or switch to a (new) product A.

A simple model of diffusion in a social network

- Let's assume that we are making decisions based on the direct benefit of being coordinated with our friends beyond any intinsic value associated with the decision (e.g. when the decision is the purchase of an item).

- A standard example is what laptop or cell phone we decide to buy to the extent that we are mostly influenced by our friends rather than by general population wide usage. What influences you most? Friends or general population benefits?

A simple model of diffusion in a social network

- Let's assume that we are making decisions based on the direct benefit of being coordinated with our friends beyond any intinsic value associated with the decision (e.g. when the decision is the purchase of an item).

- A standard example is what laptop or cell phone we decide to buy to the extent that we are mostly influenced by our friends rather than by general population wide usage. What influences you most? Friends or general population benefits?
  - Choosing between two weekly television shows that occur at the same time or who to vote for are other examples.

- In fact, the model given in this chapter dictates that certain decisions (i.e. to change from B to A) are irreversible.
  - The text calls this a "progressive process" in the sense that it progresses in only one direction. Any good examples of truly (or *essentially*) irreversible decisions?

A simple model of diffusion in a social network

- Let's assume that we are making decisions based on the direct benefit of being coordinated with our friends beyond any intinsic value associated with the decision (e.g. when the decision is the purchase of an item).

- A standard example is what laptop or cell phone we decide to buy to the extent that we are mostly influenced by our friends rather than by general population wide usage. What influences you most? Friends or general population benefits?
  - Choosing between two weekly television shows that occur at the same time or who to vote for are other examples.

- In fact, the model given in this chapter dictates that certain decisions (i.e. to change from B to A) are irreversible.
  - The text calls this a "progressive process" in the sense that it progresses in only one direction. Any good examples of truly (or *essentially*) irreversible decisions?
  - For example, the decision to get a tatoo.

A threshold model for spread

- We assume that some number of individuals are enticed (at some time $t = 0$) to adopt a new product A.
- Outside of these "initial adopters", we assume all other individuals in the network are initially using a different product B (or equivalently this is the first product in a given market).
- This is not really a competitive influence model as B is not really competing. (More comments later.)
- The first model we consider for diffusion is that every node v has a threshold q (in absolute or relative terms) for how many of its neighbors must have adopted product A before v adopts A.

Threshold model (continued)

- For simplicity the text initially assumes that every node v (i.e. individual) in the network has the same threshold but then later explains how to deal with individual thresholds.
- If at some time t, the threshold for a node v has been achieved, then by time time $t+1$, v will adopt product A.
- If the threshold has not been reached then v decides not to adopt A at this time.

Note

Although it is not explicitly stated, the initial adopters never reverse their adoption.

- Given these model assumptions, adopting A is irreversible for all nodes in the network.

Determining a (relative) threshold

- One way (some might say is usually the best way) to reason about a plausible threshold for a node is to view one's decision in economic terms.
- Specifically for every edge (v, w) in the network suppose
  - There is payoff a to v and w if both v and w have adopted product A.
  - There is payoff b to v and w if both v and w have adopted product B.
  - A zero payoff when v and w do not currently utilize the same product.
- This determines a simple coordination game.

		w: A	w: B
v	A	a, a	$0, 0$
	B	$0, 0$	b, b

Figure: $A - B$ coordination [Fig 19.1, E&K]

Coordination game induces threshold

- Suppose node v has not yet adopted A at time t, but a fraction p of the $d(v)$ neighbors of v have already adopted A, then:
  - By switching, the payoff to v is $p \times d(v) \times a$.
  - By staying with B, v has payoff $(1 - p) \times d(v) \times b$.
- Thus node v will switch to A if

$$p \times d(v) \times a \geq (1 - p) \times d(v) \times b$$

  (for simplicity say v switches when payoffs are equal).

- This is then equivalent to saying that v will switch whenever p is at least $\frac{b}{a+b} = q$ which is then the relative threshold.
- That is, whenever there is at least a (threshold) fraction q of the neighbours of node v that have adopted A, then v will also adopt A.

The process unfolds (example: $a = 3$ and $b = 2$)

[Fig 19.3, E&K]

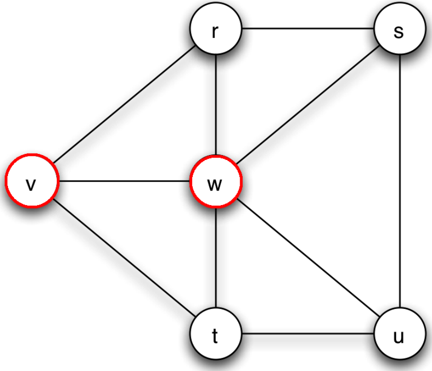

$t = 0$

- A node adopts A if and only if the threshold $q = \frac{b}{a+b} = 2/5$ is reached.
- Two nodes v and w are initial adopters.

The process unfolds (example: $a = 3$ and $b = 2$)

[Fig 19.3, E&K]

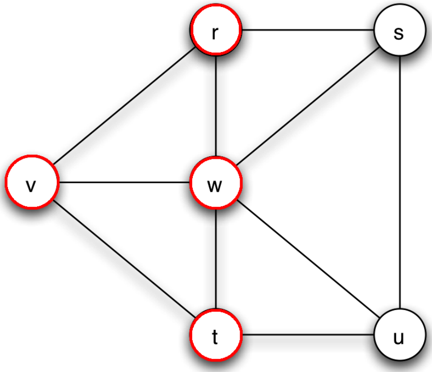

$t = 1$

- A node adopts A if and only if the threshold $q = \frac{b}{a+b} = 2/5$ is reached.
- Two nodes v and w are initial adopters.

The process unfolds (example: $a = 3$ and $b = 2$)

[Fig 19.3, E&K]

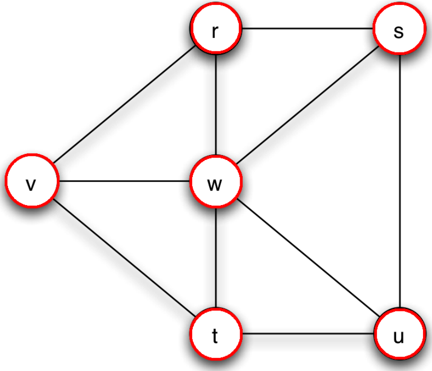

$t = 2$

- A node adopts A if and only if the threshold $q = \frac{b}{a+b} = 2/5$ is reached.
- Two nodes v and w are initial adopters.

Complete cascades vs tightly-knit communities (example: $a = 3$, $b = 2$, $q = 2/5$)

- The previous example showed a complete cascade where all nodes eventually adopt A.
- In the next example, "tightly-knit communities" block the spread.

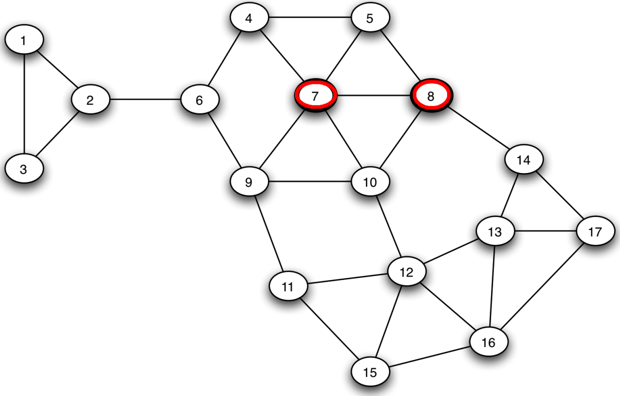

$t = 0$

[Fig 19.4, E&K]

Complete cascades vs tightly-knit communities (example: $a = 3$, $b = 2$, $q = 2/5$)

- The previous example showed a complete cascade where all nodes eventually adopt A.
- In the next example, "tightly-knit communities" block the spread.

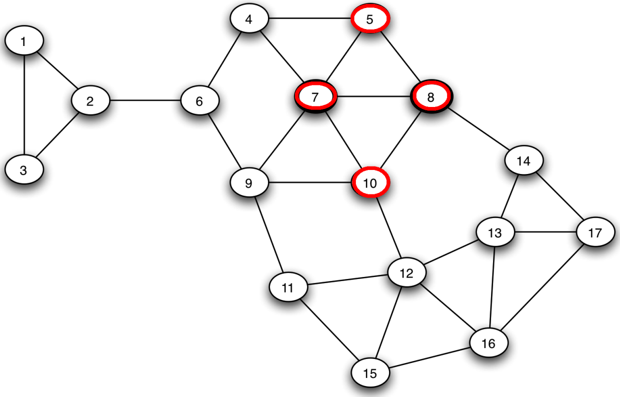

$t = 1$

[Fig 19.4, E&K]

Complete cascades vs tightly-knit communities (example: $a = 3$, $b = 2$, $q = 2/5$)

- The previous example showed a complete cascade where all nodes eventually adopt A.
- In the next example, "tightly-knit communities" block the spread.

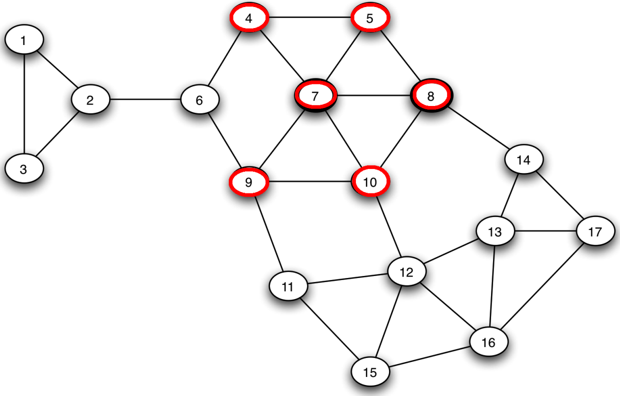

$t = 2$

[Fig 19.4, E&K]

Complete cascades vs tightly-knit communities (example: $a = 3$, $b = 2$, $q = 2/5$)

- The previous example showed a complete cascade where all nodes eventually adopt A.
- In the next example, "tightly-knit communities" block the spread.

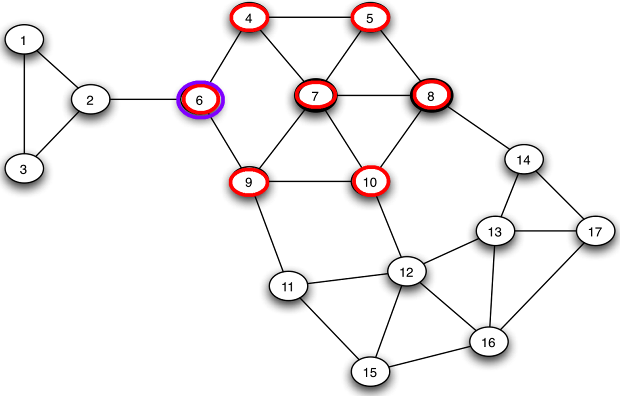

$t = 3$

[Fig 19.4, E&K]

End of Monday, March 4 lecture

We ended Monday lecture at the start of the discussion of Chapter 19 and the topic of cascades in networks on Wednesday. However, I omitted a discussion in Chapter 17 on direct network effects that I should have included. Namely, what happens when the benefits function $f(z) > 0$. I will add some slides (where they belong at the end of the Chapter 17 discussion) which now means that the Monday lecture in effect ended at slide 30.

Todays agenda

- Go over the slides concerning $f(0) > 0$
- Continue the discussion of cascasdes in a network
- Choosing influential initial adopters. This is material not in text but can be found in the article Kempe, Kleinberg and Tardos paper which has been posted on the web page.
- Finish discussion of chapter 19

Factors determining the rate and extent of diffusion in a social network

1. The structure of the network.

2. The relative payoffs vs costs for adopting a new product.
   - We haven't spoken of costs yet but we usually do have a cost for adopting a new product.
   - We can introduce such a cost into the model by saying that v will not adopt the new A unless

$$p \times d(v) \times a \geq (1 - p) \times d(v) \times b + cost$$

   - We could also add intrinsic values for A and B to both sides of the above inequality to determine the threshold for v adopting A.

3. The choice of initial adopters.
   - This raises an interesting computational question as to how to select the most influential nodes (within some budgetary constraint).

Defining a tightly-knit community

- We want to show that not only do tightly-knit communities cause a cascade to be blocked but moreover this is the only thing that can stop a cascade.

- To do so, we need a more precise definition.

Definition

A non-empty subset S of nodes is a blocking cluster of density p if every node $v \in S$ has at least a fraction p of its edges go to nodes in S.

Aside

- Clustering is a pervasive concept in many fields and contexts (beyond networks).
- It is an intuitive concept that can be defined in many ways.
- There does not appear to be any one definition that is always (or even usually) most preferred.

Clusters at different levels of granularity

- The given definition of a blocking cluster does not imply a unique way of clustering the nodes.

- Indeed if S and T are both clusters of density p, then the union of S and T is a cluster of density p.
  - Note: this is not generally true of the intersection of S and T.

- This clustering definition also implies that the set of all nodes is a cluster of density 1.

Clusters vs complete cascades

- Suppose we have a network threshold spread model with threshold q, an initial set of A adopters I and $V' = V - I$ is the set of nodes that are not initial adopters.

- Then we have the following (provable) intuitive result that characterizes when complete clusters will or will not form:

  - If V' contains a cluster C of density greater than $1 - q$, then the initial adopters will not cause a complete cascade. Furthermore, no node in C will adopt A.
  - If in a network with threshold q and an initial set I of adopters does not cause a complete cascade, then the non initial adopters nodes $V' = V - I$ must contain a cluster of density greater than $1 - q$.

When nodes have different thresholds

- As remarked before the assumption that all nodes have the same threshold is not essential.
- Consider a node v. Suppose now that for every adjacent edge (v, w), node v has payoff $a(v)$ (resp. $b(v)$) if both v and w have adopted product A (resp. B) and a zero payoff if v and w currently utilize different products.
- If node v has not yet adopted A at time t, but a fraction p of the $d(v)$ neighbours of v have already adopted A, then:
  - By switching, v has payoff $p \times d(v) \times a(v)$.
  - By staying with B, v has payoff $(1 - p) \times d(v) \times b(v)$.
- Thus node v will switch to A if

$$p \times d(v) \times a(v) \geq (1 - p) \times d(v) \times b(v).$$

- This is then equivalent to saying that v will switch whenever

$$p \geq \frac{b(v)}{a(v) + b(v)} = q(v)$$

which is then the threshold for node v.

Redefining blocking clusters

- A blocking cluster is now a set of nodes C such that every node $v \in C$ has more than a fraction $1 - q(v)$ of its adjacent nodes in C.
- It follows (as in the case of homogenous threshold nodes) that a given set of adopters I in a network will not cause a complete cascade iff $V - I$ contains a blocking cluster C.

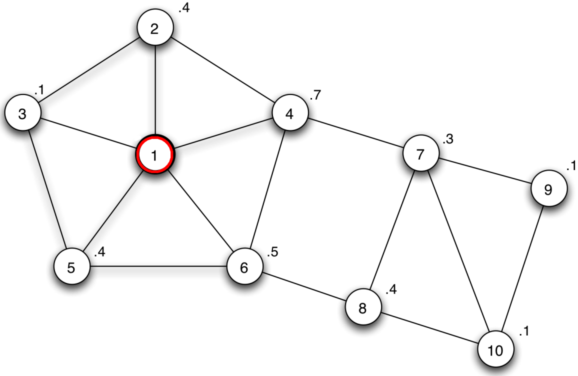

$t = 0$

[Fig 19.13, E&K]

Redefining blocking clusters

- A blocking cluster is now a set of nodes C such that every node $v \in C$ has more than a fraction $1 - q(v)$ of its adjacent nodes in C.
- It follows (as in the case of homogenous threshold nodes) that a given set of adopters I in a network will not cause a complete cascade iff $V - I$ contains a blocking cluster C.

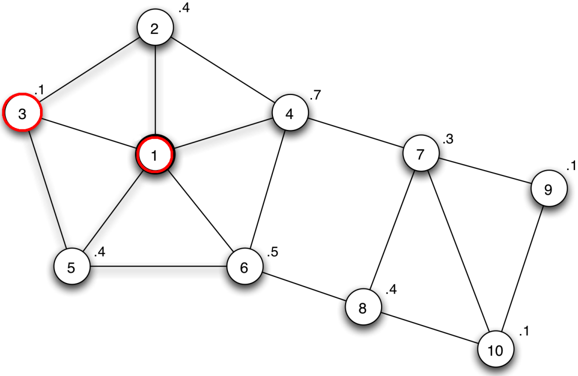

[Fig 19.13, E&K]

$t = 1$

Redefining blocking clusters

- A blocking cluster is now a set of nodes C such that every node $v \in C$ has more than a fraction $1 - q(v)$ of its adjacent nodes in C.
- It follows (as in the case of homogenous threshold nodes) that a given set of adopters I in a network will not cause a complete cascade iff $V - I$ contains a blocking cluster C.

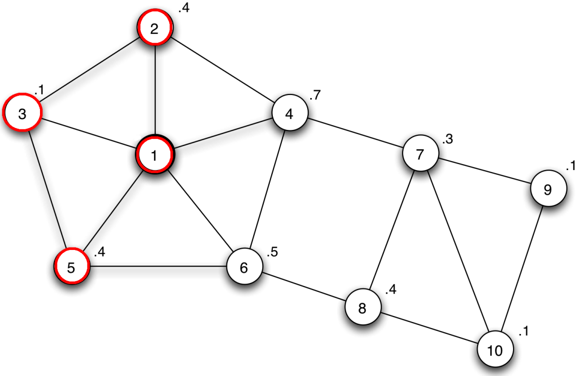

[Fig 19.13, E&K]

$t = 2$

Redefining blocking clusters

- A blocking cluster is now a set of nodes C such that every node $v \in C$ has more than a fraction $1 - q(v)$ of its adjacent nodes in C.
- It follows (as in the case of homogenous threshold nodes) that a given set of adopters I in a network will not cause a complete cascade iff $V - I$ contains a blocking cluster C.

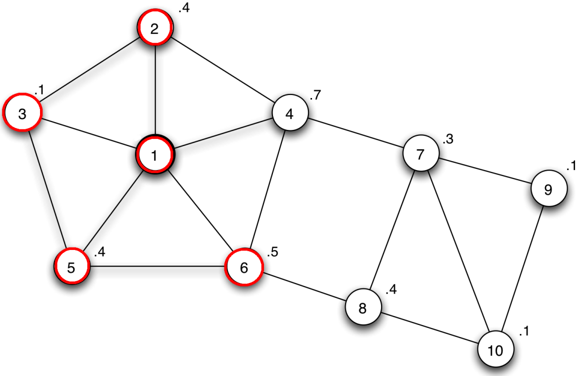

$t = 3$

[Fig 19.13, E&K]

Redefining blocking clusters

- A blocking cluster is now a set of nodes C such that every node $v \in C$ has more than a fraction $1 - q(v)$ of its adjacent nodes in C.
- It follows (as in the case of homogenous threshold nodes) that a given set of adopters I in a network will not cause a complete cascade iff $V - I$ contains a blocking cluster C.

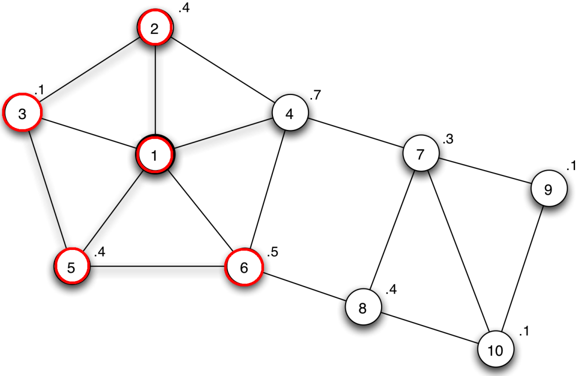

[Fig 19.13, E&K]

$t = 4$

Redefining blocking clusters

- A blocking cluster is now a set of nodes C such that every node $v \in C$ has more than a fraction $1 - q(v)$ of its adjacent nodes in C.
- It follows (as in the case of homogenous threshold nodes) that a given set of adopters I in a network will not cause a complete cascade iff $V - I$ contains a blocking cluster C.

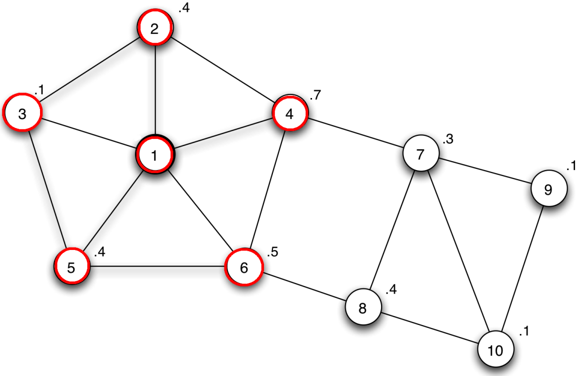

$t = 5$

[Fig 19.13, E&K]

Choosing influential adopters

- Suppose we wish to spread a new technology and to do so we have money to influence some "small" set of initial adopters (e.g. by giving away the product or even paying people to adopt it).

- Even in this simple model of (non-competitive) influence spread, and even if we have complete knowledge of the social network, it is not at all clear how to chose an initial set of adopters so as to achieve the largest spread.

- Furthermore the spread process could be much more sophisticated.
  - For example, adoption by a node might be a more random process (say adopting with some probability relative to the nodes threshold) and maybe the influence of neighbors first increases and then decreases over time.

Choosing influential adopters continued

- Suppose we have funds/ability to influence k nodes to become initial adopters.
  - We can try all possible subsets of the entire $n = |V|$ nodes and for each such subset simulate the spread process.
  - But clearly as k gets larger, this “brute force” becomes prohibitive especially for large networks.

- Earlier in the course, we mentioned that for many optimization problems (like the one being considered now), there is a widely held belief (with good supporting evidence from complexity theory):

“NP-hard problems” cannot be optimally solved in an efficient manner (and sometimes we cannot even get a good approximation to optimality).

Can we determine a “good” set of initial adopters?

- For even simple models of information spread as being discussed here, complexity theory (the P vs NP conjecture) argues that we cannot efficiently choose the best set of initial adopters. There is a class of networks for which (assuming the $P \neq NP$ conjecture) it is not possible to obtain an approximation within a factor n^c for any $c < 1$.

- Instead we will identify properties of a spread process that will allow a good approximation: a good set of initial adopters that will do “almost as well” as the best set.

Note: What follows is a discussion as to how to choose a set of initial adopters by a relatively efficient approximation algorithm when making some assumptions on the spread process. However, we would need much more efficient methods for very large networks.

Influence maximization models; monotone submodular set functions

- Some spread models have the following nice properties.

Let $f(S)$ be size (or more generally a real value benefit since some nodes may be more valuable) of the final set S of adopters satisfying:

1. Monotonicity: $f(S) \leq f(T)$ if S is a subset of T
2. Submodularity: $f(S+v) - f(S) \geq f(T+v) - f(T)$ if S is a subset of T

- We also usally assume that $f(\varnothing) = 0$. Such normalized, monotone, submodular functions arise in many applications.
- The simple threshold examples considered thus far are monotone processes but are not submodular in general. Are these contrived worst case network examples?
- But some variants of the threshold model and related models do satisfy these properties. We consider two such stochastic models.

Linear threshold model

- We have an edge weighted (undirected or directed) network where weight $w(u, v)$ represents the relative influence (e.g. quantitative version of weak and strong ties) of node u on node v.

- Now each nodes threshold $q(v)$ is chosen randomly in $[0, 1]$ to model lack of knowledge as to how easy it is to influence a given individual.

- A node v adopts A if the sum of all edge weights into v exceeds the randomly chosen $q(v)$.

- **Goal:** find an initial set of k adopters so as to maximize the expected number (or benefit) of eventual adopters. (This is a stochasitic process so that we are trying to optimize the expected value of the process.)

- **Aside:** We often use the language of disease spread and say "infected nodes" rather than "already influenced nodes".

The linear threshold model

- Each node v chooses a threshold t_v randomly from $[0, 1]$.
- Each edge (u, v) has assigned weight w_{uv} from $[0, 1]$ such that

$$\sum_{u \to v} w_{uv} \leq 1.$$

- In each step t, a node v is infected if the weighted sum of incident edges coming from infected neighbors exceeds threshold.

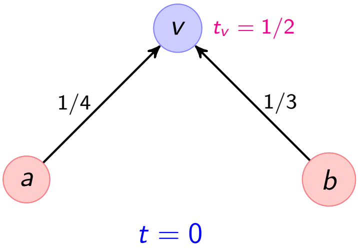

The linear threshold model

- Each node v chooses a threshold t_v randomly from $[0, 1]$.
- Each edge (u, v) has assigned weight w_{uv} from $[0, 1]$ such that

$$\sum_{u \to v} w_{uv} \leq 1.$$

- In each step t, a node v is infected if the weighted sum of incident edges coming from infected neighbors exceeds threshold.

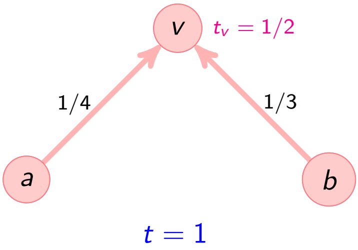

The linear threshold model

- Each node v chooses a threshold t_v randomly from $[0, 1]$.
- Each edge (u, v) has assigned weight w_{uv} from $[0, 1]$ such that

$$\sum_{u \to v} w_{uv} \leq 1.$$

- In each step t, a node v is infected if the weighted sum of incident edges coming from infected neighbors exceeds threshold.

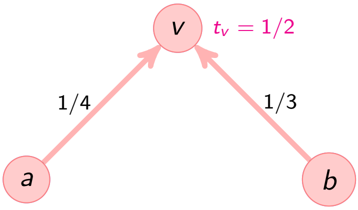

$t = 1$

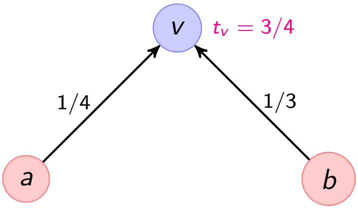

$t = 0$

The linear threshold model

- Each node v chooses a threshold t_v randomly from $[0, 1]$.
- Each edge (u, v) has assigned weight w_{uv} from $[0, 1]$ such that

$$\sum_{u \to v} w_{uv} \leq 1.$$

- In each step t, a node v is infected if the weighted sum of incident edges coming from infected neighbors exceeds threshold.

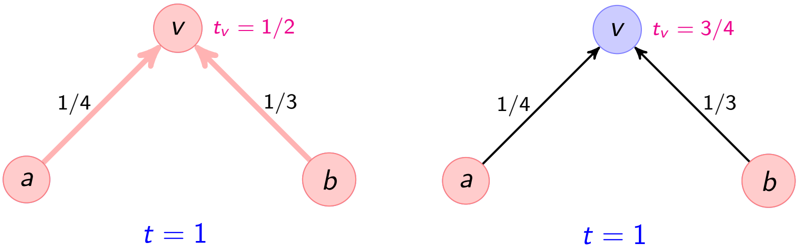

Independent cascade influence model

- We again have an edge weighted network (as in threshold model) but now the weights $p(u, v) \leq 1$ represent the probability that node u will influence node v given one and only one chance to do so.
- That is, if node u adopts A at time t, then with probability $p(u, v)$, node v will adopt v at time $t + 1$.
- After this, node u will *not* have another opportunity to influence v.
- **Goal for both threshold and cascade models:** to find initial set of adopters to maximize the expected number of eventual adopters.
- Threshold and (especially) cascade processes are motivated by models for the contagious spread of disease. Should disease spread and influence spread should be governed by similar processes?
  - See `http://www.economist.com/blogs/babbage/2012/04/social-contagion`

The Independent Cascade Process

- Each edge (u, v) has an associated probability p_{uv}.
- In each step t, nodes that adopted technology at step $t-1$ "infect" each of their uninfected neighbors with probability p_{uv}.

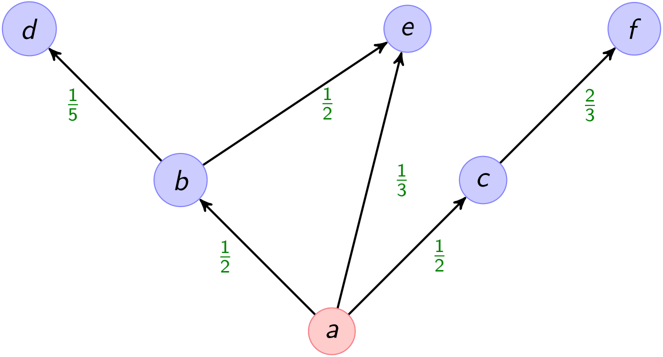

$t = 0$

The Independent Cascade Process

- Each edge (u, v) has an associated probability p_{uv}.
- In each step t, nodes that adopted technology at step $t-1$ "infect" each of their uninfected neighbors with probability p_{uv}.

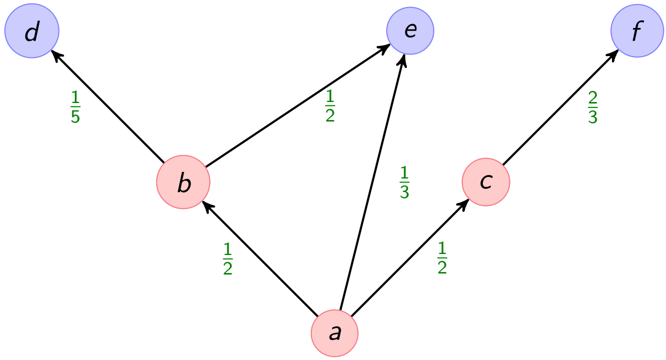

$t = 1$

The Independent Cascade Process

- Each edge (u, v) has an associated probability p_{uv}.
- In each step t, nodes that adopted technology at step $t - 1$ "infect" each of their uninfected neighbors with probability p_{uv}.

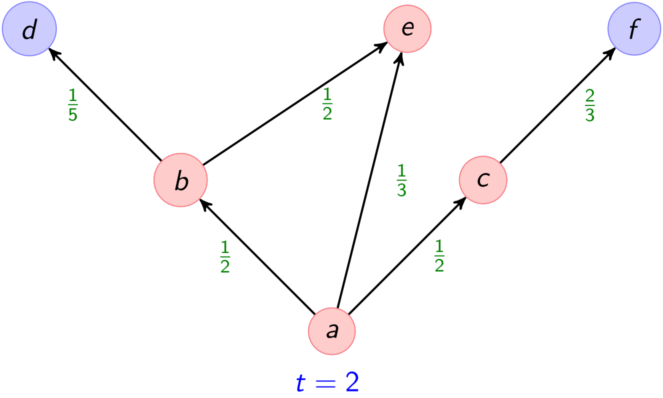

The Independent Cascade Process

- Each edge (u, v) has an associated probability p_{uv}.
- In each step t, nodes that adopted technology at step $t-1$ "infect" each of their uninfected neighbors with probability p_{uv}.

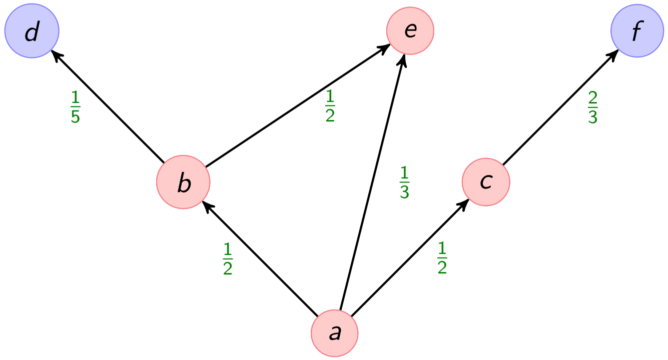

$t = 3$