Social and Information Networks

University of Toronto CSC303
Winter/Spring 2019

Week 7: February 25,27 (2019)

Announcements and agenda

Announcements

- The term test is this Friday, March 1 in this room MP 134. The scope of the term test will include:
  1. A question regarding the different closures (similar to question on assignment).
  2. A question on structural balance.
  3. A question on power laws.
- Next assignment is the written critical review of a current article due on March 15. I will elaborate on this now. I will also post something on the course web page.

Agenda

- Piazza postings: An interesting question regarding triadic closure and a comment re power law statement.
- We will review the immediate road map for what we will be doing in the next couple of weeks.
- We continue with the discussion on information cascades that we just began before reading week.

Question posed on piazza

Triadic closure and homophily
I just wanted to quickly double check my understanding. In a social-afflication network triadic closure (let us say A-B and A-C induce the closing of B-C) is not necessarily social influence nor selection, as we could argue that A has influenced B to become friends with C (thus social influence), or we could argue that B has selected C to become friends with due to their similarity in being friends with A (thus selection). Is this understanding correct? Thanks!

It is indeed a question that I also raised (and did not answer) on slide 30 of the Week 3 slides. You may have seen my reply to the question.

But what do you think? Is triadic closure due more to influence or selection?

Roadmap: Different forms of influence

- This will be the beginning of a several week discussion of
  - influence/technology/disease spread
    i.e. "contagion" in a very general sense in social networks;
- We will first be discussing Chapter 16 (information cascades) where (as we have seen before) sequential decisions are influenced by previous decisions. The chapter argues that being infuenced by previous decisions is *rational* and not necessarily mindless. Here the benefit is indirect in the sense that the probability of making a better decision can be improved by following othetrs.
- Then in Chapter 17 , influence comes in the form that there is a *direct benefit effect* (i.e. a change in the reward) for following others.
- That is, for the next several weeks we will be studying

various social processes that channel individual behaviour into collective behaviour.

Information cascades (herding) and Direct Benefits

- Chapter 16 concerns the phenomena of information cascades
  - whereby individuals observe and then make decisions sequentially based on the behaviour of people having made decisions earlier;
  - e.g. deciding on a restaurant by observing how many people are currently eating there, what clothes you buy, other fashions/fads.

- Chapter 17 discusses decisions based on direct benefit (e.g., using a popular operating system/ laptop because wide use implies more software support).

- Clearly both phenomena can be interacting when people make decisions (e.g. busy restaurants are more able to use fresh ingrediants); the text organization is to try to first isolate and model these phenomena so as to gain insight.

A simple information cascade model

Assumptions for an information cascade:

- Individuals make decisions sequentially and can observe the decision of those who have acted earlier.
- Each individual has some private information that can be used in making their decision.
- Individuals only observe the behavior of earlier people but do not know their private information beyond any inferences that can be made from the previous decisions/actions.
- Note how the musice evaluation experiment fits into this model.

The majority balls and bins example

- In the following experiment we will observe the formation of very uniform (albeit perhaps fragile, perhaps wrong) cascades.

- The majority blue vs majority red *balls in a bin* experiment.
  - Assume we have a bin which (with equal probablility) either has 2 red balls and one blue ball, or has 2 blue balls and one red ball.
  - Each individual in turn blindly (i.e. randomly) picks out a ball from the bin, observes its colour, and then places it back into the bin.
  - Now each individual votes (i.e. announces their decision as to which colour is in the majority) based on what they have observed and what others have previously reported.

- In a different experiment, one can ask what would happen if one votes their private info unless there is an opposing "clear majority"? What is a clear majority? Following the Supreme Court in decision regarding a Quebec independence vote, clear majority is not defined.

Possible behaviour for the majority balls experiment

- In the following experiment we will observe the formation of definite (albeit perhaps fragile, perhaps wrong) cascades.

- The majority blue vs majority red *balls in a bin* experiment.
  - Assume we have a bin which (with equal probablility) either has 2 red balls and one blue ball, or has 2 blue balls and one red ball.
  - Each individual in turn blindly (i.e. randomly) picks out a ball from the bin, observes its colour, and then places it back into the bin.
  - Now the experiment assumes that one votes (i.e. announces their decision as to which colour is in the majority) their own private information (the colour of the ball drawn) *unless* the difference in previous colours drawn is at least 2 in which case they vote with the majority.

- In a different experiment, one can ask what would happen if one votes their private info unless there is a "clear majority"? What is a clear majority? Following the Supreme Court, not defined.

Conditional probability and what we tend to believe

- Why in the simple balls and bin text example (with 3 balls), would we vote for blue if the first two decisions were blue even if we saw a red ball?

- Simply stated, conditioned on the information provided by these two decisions, we can infer that it is more likely that the urn contains 2 blue and 1 red ball. We now want to make precise such a statement in the language of probability.

- We have the following concepts and notation:
  - $Pr[A]$ = the "prior prob." of event A; e.g. urn has 2 blue,1 red
  - $Pr[A|B]$ = the "posterior or conditional prob." of event A given that event B occurred; e.g. B is the event that I drew a red ball and the previous decisions were $D_1, D_2, \ldots$

Bayes rule

- By definition of conditional prob., $Pr[A|B] = Pr[A \text{ and } B]/Pr[B]$
- By definition of conditional prob., $Pr[B|A] = Pr[A \text{ and } B]/Pr[A]$
- Therefore $Pr[A|B] \times Pr[B] = Pr[B|A] \times Pr[A]$
- So we have **Bayes Rule**

$$Pr[A|B] = Pr[A] \times \frac{Pr[B|A]}{Pr[B]}$$

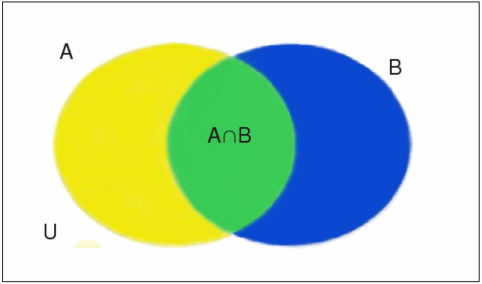

Figure: Two events A and B in a sample space, and the joint event $A \cap B$.

The black and yellow taxi example

- Eyewitness of a hit and run reports taxi was yellow
- Let A be the event that the true colour is yellow.
- Let A' be the event that the true colour is black.
- Let B is the event that the witness reports taxi was yellow. Given what is known, How likely do you think it is that the taxi was yellow?

The black and yellow taxi example

- Eyewitness of a hit and run reports taxi was yellow
- Let A be the event that the true colour is yellow.
- Let A' be the event that the true colour is black.
- Let B is the event that the witness reports taxi was yellow. Given what is known, How likely do you think it is that the taxi was yellow?
- The desired probability is the posterior prob $Pr[A|B]$
- We are given that the prior prob Pr[A] = .2;
- Also given that this witness mixes up colours 20% of the time:

$$Pr[\text{reports yellow} \mid \text{taxi yellow}] = Pr[\text{reports black} \mid \text{taxi black}] = .8$$

- Bayes rule gives

$$Pr[A|B] = Pr[A] * Pr[B|A]/Pr[B]$$

- We have all the information we need here except the denominator: the prior probability $Pr[B] = Pr[\text{report yellow}]$; this can happen if

The black and yellow taxi example

- Eyewitness of a hit and run reports taxi was yellow
- Let A be the event that the true colour is yellow.
- Let A' be the event that the true colour is black.
- Let B is the event that the witness reports taxi was yellow. Given what is known, How likely do you think it is that the taxi was yellow?
- The desired probability is the posterior prob $Pr[A|B]$
- We are given that the prior prob $\mathrm{Pr[A]} = .2$;
- Also given that this witness mixes up colours 20% of the time:

$$Pr[\text{reports yellow} \mid \text{taxi yellow}] = Pr[\text{reports black} \mid \text{taxi black}] = .8$$

- Bayes rule gives

$$Pr[A|B] = Pr[A] * Pr[B|A]/Pr[B]$$

- We have all the information we need here except the denominator: the prior probability $Pr[B] = Pr[\text{report yellow}]$; this can happen if
  1. the taxi is really yellow and witness reports yellow or

The black and yellow taxi example

- Eyewitness of a hit and run reports taxi was yellow
- Let A be the event that the true colour is yellow.
- Let A' be the event that the true colour is black.
- Let B is the event that the witness reports taxi was yellow. Given what is known, How likely do you think it is that the taxi was yellow?
- The desired probability is the posterior prob $Pr[A|B]$
- We are given that the prior prob $\text{Pr[A]} = .2$;
- Also given that this witness mixes up colours 20% of the time:

$$Pr[\text{reports yellow} \mid \text{taxi yellow}] = Pr[\text{reports black} \mid \text{taxi black}] = .8$$

- Bayes rule gives

$$Pr[A|B] = Pr[A] * Pr[B|A]/Pr[B]$$

- We have all the information we need here except the denominator: the prior probability $Pr[B] = Pr[\text{report yellow}]$; this can happen if
  1. the taxi is really yellow and witness reports yellow or
  2. the taxi is really black and witness reports yellow

Eye witness and the black & yellow taxi example

- We are trying to compute the $\Pr[B] = \Pr[\text{reports yellow}]$ which is

$$\Pr[\text{reports yellow}|A] \times Pr[A] + Pr[\text{reports yellow}|A'] \times Pr[A']$$
$$= .8 * .2 + .2 * .8 = .32$$

- Returning to Bayes rule:

$$Pr[A|B] = Pr[A] * Pr[B|A]/Pr[B] = .2 * .8/.32 = .5$$

- This is the same thing as if someone just guesses; that is,
  1. knew nothing about the distribution of yellow and black taxis (and therefore assumes $\Pr[A] = .5$) and
  2. didn't see anything (and therefore has no information upon which to condition the probability) and reports yellow by flipping a fair coin.
- Another way to say it: the taxi is equally likely to be yellow or black.
- The fact that the witness reports yellow has, however, significantly raised the chances of the taxi being yellow (from .2 to .5)

Another example: spam filtering

- We receive an email with subject heading that more or less says "check this out". Should we treat this as spam?
- Let A be the event that the message is spam.
- Let A' be the event that the message is not spam
- Let B is the event that the message heading is "check this out".
- We are given prior probability $Pr[A] = .4$, $Pr[B|A] = .01$ and $Pr[B|A'] = .004$. We want to compute $Pr[A|B]$.
- By Bayes Rule,

$$Pr[A|B] = Pr[A] \times Pr[B|A]/Pr[B]$$

- We need to calculate $Pr[B]$ and as before; that is, $P[B]$ is

$$Pr[B|A] \times Pr[A] + Pr[B|A'] \times Pr[A'] = (.01) \times (.4) + (.004) \times (.6) = .0064$$

- Hence $Pr[A|B] = Pr[A] \times Pr[B|A]/Pr[B] = .625$

Additional signals for spam, etc.

- Of course, spam filters (or most uses of Bayes rule) do not rely on only one simple signal (conditional event) such as the occurrence of a particular phrase say in the subject line.
- Rather multiple signals are used to detect spam (e.g. bold face in the subject, occurrence of large sums of money in the email content, etc.).
- These could be called positive signals/events (general or specific to an individual) that raise the probability that a message is spam.
- There may also be negative signals; that is events (usually specific to a given individual) that raise the probability that the message is not spam (e.g. "request to submit assignment late" in subject header).

Bayes with positive and negative signals

- If $B_1, B_2, \ldots, B_r$ are positive signals and $C_1, C_2, \ldots, C_k$ are negative signals then we can just let B (the conditional event) be a (complex) joint event, namely B_1 and B_2 and ... and B_r and not C_1 and not C_2 and ... and not C_k.
- However, we cannot simply expect that the probability of this joint event is the product of the individual probabilities as it is very likely that these positive and negative signals are not independent.
- With supervised learning, we can attempt to train a spam filter to experimentally determine all the prior and conditional probabilities.

Returning to balls and bins cascade.

- Bayes rule now gives us a quantitative explanation for why a cascade can (and will) form (if the individuals are making decisions implicitly or explicitly on such reasoning).
- The example given made an important assumption (so as to more easily illustrate the cascade or herding phenomena) implying:

The second person will rely on their own personal information (the ball they see) to break a tie. It will follow that that first two votes honestly report the first two balls drawn.

- So again using the same notation let A be the event that the urn contains two blue balls and one red ball.
- (Since the situation is completely symmetric the same analysis will hold for $\bar{A}$, the event of two red balls and one blue ball in the bin.)
- For the ith individual drawing a ball, we have a different event B depending on the first $i-1$ decisions (different from first $i-1$ balls drawn) and the ith ball drawn.

The first and second individuals

- For the prior probabilities we have $Pr[A] = Pr[\bar{A}] = 1/2$.
- Suppose the first person draws a blue ball. Now the desired conditional probability is $Pr[A|B]$ where B is the event that a blue ball is drawn. Note that the prior probability $Pr[B] = Pr[B|A] \times Pr[A] + Pr[B|\bar{A}] \times Pr[\bar{A}] = 1/2$.
- Using Bayes rule, the conditional probability $Pr[A|B]$ is $(1/2) * (2/3)/(1/2) = 2/3$ and the mathematically correct thing to do (i.e. decide or vote what is most probable) is to decide **majority blue**.
- Now suppose the second person draws a red ball. This person can infer that the first person saw a blue ball and now the event B we are conditioning on is the sequence of draws ($blue$, red).
- Note we are assuming that each draw is independent (sampling with replacement) but these are not independent events as they both depend on the bin majority and therefore $Pr[B]$ is not simply a product of event probabilities.

Second and third persons

- Using Bayes rule again

$$Pr[A|B] = Pr[A] \times Pr[B|A]/Pr[B] = (1/2) \times (2/9)/(2/9) = 1/2$$

where B is the sequence (*blue*, *red*) so that

$$Pr[B] = Pr[B|A] \times Pr[A] + Pr[B|\bar{A}] \times Pr[\bar{A}] =$$

$$(1/2)[(2/3) \times (1/3)] + (1/2)[(1/3) \times (2/3)] = 2/9$$

- The second player is now indeed (quantitatively) indifferent and by the tie-breaking assumption of this majority ball example, the second player would decide to say that the bin is **majority red**.
- We are (intuitively) in a situation where the third player can ignore the previous decisions.
- We will instead assume that the second player also saw blue and then will clearly decide blue.
- So now what does the third player decide if he/she draws a red ball?
- Now the event B we are conditioning on is the sequence of draws (*blue*, *blue*, *red*) as the third player can be sure that the first two draws were blue.

The cascade has begun

- So far the first two players decided according to the ball that they have drawn.
- In contrast when we calculate Bayes rule for the conditional probability

$$Pr[A|B] = Pr[A] \times Pr[blue, blue, red|A]/Pr[B]$$

- Once again, the denominator cannot simply be calculated by a product of the individual events and a simple calculation shows that $Pr[B] = 1/9$ (and not $1/8$); hence $Pr[A|B] = \frac{1}{2}\frac{4}{27}/\frac{1}{9} = 2/3$ so that the third player will ignore their signal and decide blue.
- Further players will infer that only the first two decisions were based on what occurred and will then also act just like the third player and will declare **majority blue**.
- **Note:** A more mindless 4th individul may think that the first three decisions were made indepndently and be even more persuaded to think that the majoriy color is blue.

Eventually a cascade must begin

- Section 16.6 shows why a cascade will almost surely eventually begin (if there are enough people in this experiment). Let's do this analysis a little differently than the text.

- What we have seen is that the third person will follow their signal if and only if the first two signals are different and otherwise a cascade will form.

- Note that even if the urn is **majority blue**, if the first two balls drawn are red, then the third player will declare the urn is **majority red** and we get a cascade (albeit an incorrect) sequence of decisions.

- Recall, the symmetry in this experiment.

A cascade almost surely forms

- What is the probability that a cascade will not form after say the first $2k$ players have drawn balls?
- This can only happen if each pair $(i, i+1)$ of drawings results in an opposite pair. For definiteness lets say the urn is **majority blue**.
- Then the probability that the first pair are opposites is $2 \times (2/3 \times 1/3) = 4/9$.
- Since we choose to have **majority blue** and **majority red** with equal probability, no matter what is the majority color we have probability $4/9$ to get opposite draws.
- Similarly, the next two colors drawn are opposites with probability $4/9$ and thus the probability that the first k pairs will have opposite colors is $(4/9)^k$ which limits to 0 as k increases.

Lessons from the urn cascade

- Cascades can be wrong and based on little information.
- Cascades can be fragile; e.g. if somehow you see the immediately preceding two (breaking ties as before) or more signals, then the cascade will become that of this latest information.
- Cascades are in contrast to "the wisdom of the crowd" (crowd sourcing) where we hopefully rely on independent observations (or explanations for decisions) before making our own decision.
- The text gives an example of a hiring committee perhaps being subject to a cascade but in reality such decisions are made (let's hope) in a more informative way;
  - Namely, we just don't listen to the votes of previous members but also hear their reasoning. This is analogous to seeing all the balls that have been drawn rather than just hearing all the votes.
  - Or we can vote by closed ballot. More generally, there will be rounds of discussion.

Ened of Monday, February 25 lecture

We ended at slide 22 which coincides with the end of the discussion of cascades.

Announcements:

- Tyrone is hoping to have Assignment 1 graded by tonight or tomorrow AM.
- Grading requests must come within one week of the graded assignments being made availalble. Each such request must come with a full explanation of why you think you may not have received proper credit. Caluclation errors can be immedietly dealt with by me.
- Test This Friday.
- Any question re the critical review assignment? I have posted some notes about this on the web page (in the Assignment section).

Direct benefit effects

- We now turn to Chapter 17 and a study of direct benefit effects
- For example, I might choose a social network (e.g. Linkedin) or an operating system just because its wide spread use gives me a direct benefit (e.g. more people know about me and that might improve my professional standing; more software is available).
- The text refers to this phenomena both as a direct benefit effect and as a "network effect". I think the latter is not the best choice of terminology as neither Chapter 16 nor 17 is in the context of say a social network. Indeed the text says

  > *...here, payoffs depend on the number of others who use a good and not on the details of how they are connected.*

- We will discuss social network influence issues later in Chapter 19.

Externalities

- A direct benefit effect is an example of a positive externality. The concept of an **externality** is a little vague but it is an important concept.
- The text explains

  > *An externality is any situation in which the welfare of an individual is affected by the actions of other individuals, without a mutually agreed-upon compensation.*

- Wikipedia gives a "definition" which I find less informative:

  > *In economics, an externality is the cost or benefit that affects a party who did not choose to incur that cost or benefit.*

  The cost of an externality is a negative externality,. . . while the benefit of an externality is a positive externality. . .
- Let's elaborate on what is and what is not an externality by some other examples.

What is and what isn't an externality

- In a combinatorial auction:
  - Individual items are allocated to different agents.
  - Obviously, another agent valuing an item more than me will affect my utility.
  - But this is not considered an externality. Why?

What is and what isn't an externality

- In a combinatorial auction:
  - Individual items are allocated to different agents.
  - Obviously, another agent valuing an item more than me will affect my utility.
  - But this is not considered an externality. Why?
  - Once I know my allocation (and the price I paid for it), then my utility is determined no matter who gets the other items.

What is and what isn't an externality

- In a combinatorial auction:
  - Individual items are allocated to different agents.
  - Obviously, another agent valuing an item more than me will affect my utility.
  - But this is not considered an externality. Why?
  - Once I know my allocation (and the price I paid for it), then my utility is determined no matter who gets the other items.

- On the other hand, let's consider the allocation of ad word slots.
  - This problem can be just considered as a special case of a combinatorial auction and hence assumed there are no externalities.
  - But it could be that my utility for getting a given slot can substantially depend on who gets the other slots.
  - If say competitors get those slots that might very well be a negative externality, whereas if complimentary advertisers get those slots it could be a positive externality.

Markets with a huge number of consumers

The assumption throughout Chapter 17

Any single individual won't affect the aggregate behaviour of the market.

- That is, whether or not I buy a few shares of a particular stock will not impact prices or overall demand.
- But if many people want to buy or sell a given stock then prices will be impacted which in turn will impact further demand.

- As we saw in the analysis of power laws, a common way to deal with a large but finite system of individuals is to consuder a continuous abstraction of the system. In particular, here we we will abstract the system as if individuals are just points on say the real line segment $[0, 1]$. (Apologies for reducing you to a point on the line.)

- Then each individual has no "mass" but subintervals do have proportional mass.

Consumers as points on the line $[0, 1]$

- We assume

> Each consumer is a point on the line segment $[0, 1]$
> wanting to buy one unit of a good.

- We also assume the consumer's willingness to buy the item depends both on
    1. their intrinsic interest in (i.e. value for) the item and
    2. the number of other people using the good (i.e. the direct benefit effect); the more users the more the item is worth.

- That is, we will be considering a positive externality. What is an example where more "users" will reduce the worth of an item?

- To start, the text first assumes no direct benefit effect and then studies how direct benefits change things.

Intrinsic interest, the reservation price

- We let the intrinsic interest be specified by a single reservation price $r(x)$ for individual x in $[0, 1]$. An individual x will buy the item at a price p if and only if $r(x)$ is at least p.

Intrinsic interest, the reservation price

- We let the intrinsic interest be specified by a single reservation price $r(x)$ for individual x in $[0, 1]$. An individual x will buy the item at a price p if and only if $r(x)$ is at least p.
- Without loss of generality, we arrange the individuals so that

$$x < y \text{ implies } r(x) \geq r(y)$$

Intrinsic interest, the reservation price

- We let the intrinsic interest be specified by a single reservation price $r(x)$ for individual x in $[0, 1]$. An individual x will buy the item at a price p if and only if $r(x)$ is at least p.
- Without loss of generality, we arrange the individuals so that

$$x < y \text{ implies } r(x) \geq r(y)$$

- For analysis we further assume $x < y$ implies $r(x) > r(y)$ and also that r is a continuous function on $[0, 1]$. (Once we have made the abstraction to the real line these are not critical assumptions.)

Intrinsic interest, the reservation price

- We let the intrinsic interest be specified by a single reservation price $r(x)$ for individual x in $[0, 1]$. An individual x will buy the item at a price p if and only if $r(x)$ is at least p.
- Without loss of generality, we arrange the individuals so that

$$x < y \text{ implies } r(x) \geq r(y)$$

- For analysis we further assume $x < y$ implies $r(x) > r(y)$ and also that r is a continuous function on $[0, 1]$. (Once we have made the abstraction to the real line these are not critical assumptions.)
- It follows that (except for the single point $x = 0$)

1. no one will buy the good at price $r(0)$ or more
2. at a price $r(1)$ or less everyone will buy the good.

Market demand for the good

- By continuity, for every price p with $r(0) > p > r(1)$, there is a unique $x \in (0, 1)$ (called the market demand at p) such that $r(x) = p$.
- That is, an x fraction will want to buy a unit of the good at price p.

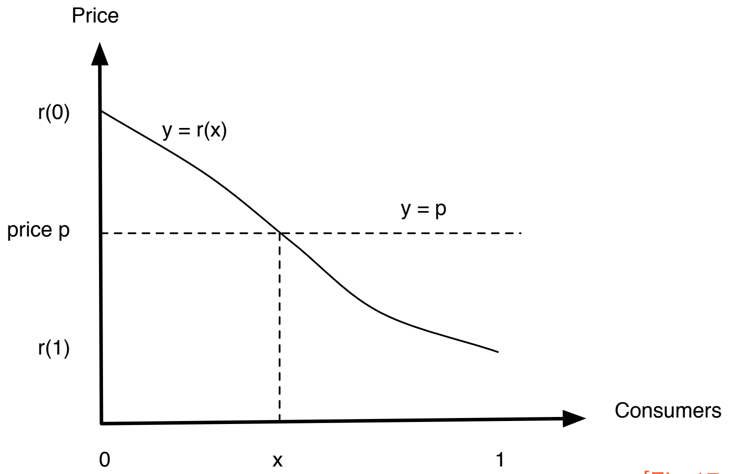

[Fig 17.1, E&K]

Market with large number of producers

- Now the discussion proceeds to assume 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.

Equilibrium quantity of good (at p^*)

- This fixed (non negotiable) cost of p^* per unit, which we can assume to be between $r(0)$ and $r(1)$, determines a unique x^* such that $r(x^*) = p^*$.

- This x^* is an equilibrium in the overall consumption of the good in the following sense.
  - If less than a fraction x^* purchased the good, there would be excess consumers with reservation prices above p^* and hence they would want to buy the good and thereby drive up consumption.
  - By assumption, consumers will not pay more than their reservation price meaning that it is not sustainable to have more than a fraction x^* of purchasers.

What are we assuming about the "good" when we are discussing how people react and converge to equilibrium?

And now what happens with the addition of direct benefit effects?

- We are assuming that having a large fraction of existing users of the good makes the good that much more desirable.

- This is modeled by now saying that the reservation price for consumer x is $r(x)f(z)$ when there is a fraction z of current users for some function $f(z)$ that is increasing in z.

- A few more assumptions, mainly to simplify the discussion; namely, assume f is continuous and $r(1) = 0$. And for now also 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 $p^* = (1 - z)z$.

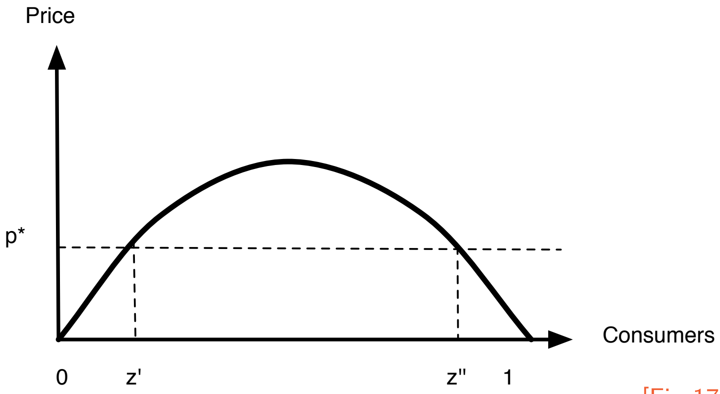

[Fig 17.3, E&K]

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.

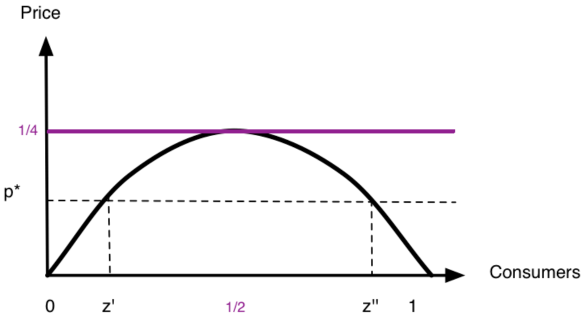

Why can't there be other equilibria?

- 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.

End of week 7

We ended week 7 with an discussion of a direct benefit model and, in particular, the concept of a self-fulfilling expectations equilibirium. Next week we will review the model and consider what happens if everyone does *not* make a perfect shared prediction. That is, if we are not in one of the equilibrium states.

We will then go on to Chapter 19 where we reintroduce networks into the discussion of cascades.