Economics of Information Networks

Stephen Turnbull

Division of Policy and Planning Sciences

Lecture 5: December 17, 2020

Abstract

In Lecture 5, Part 1 we look at information cascades*: dynamic processes where information spreads through a network. These processes are related to "flame wars" and "disinformation campaigns" on social networks, as well as to some kinds of behavior in markets.*

Behavior and information

- Networks connect people, and allow them to respond to others' activity.
- This is fundamental to *economics*, to "social engineering," and in fact to any activity that takes place in groups.
- In *social* behavior (including economics), the effect of another's behavior is not physical (except in war and competitive sports); it's informational. We respond to the *information* about others' behavior.
- Sometime the behavior itself is pure information: inspiring speeches, online bullying, or disgusting and dangerous disinformation about COVID vaccines.

Information cascades

- An example most of us have experienced is walking down a "food alley" in a shopping center, mall, arcade, or downtown, and noticing restaurants that have lines, sometimes implying a wait of hours.
- Of course, one response is "the wait is too long, let's go somewhere else." That's based on a physical (or regulatory) fact: the restaurant will allow only so many people in.
- Another, however, is "wow, the food must be delicious." Even if you pass it by today, maybe you put it on your "list" for when you have more time.
- But consider: you're not the only one who thinks that way. Probably, some of the people in line are there for that same reason! So you (who have not eaten there) think that it's delicious because other people (who have not eaten there!) think it's delicious because they joined a long line.
- *That* is an *information cascade.*

A simple numerical example

- Suppose all diners choose at random from the restaurants. Then no information cascade can occur. The distribution of line lengths will be approximately normal, according to the Central Limit Theorem.
- Suppose all the restaurants are different styles, and everyone chooses the food they want. No cascade is possible.
- Suppose all the restaurants are the same style, and everyone chooses the restaurant they think is most delicious. Then you can safely choose the restaurant with the longest line.
- Suppose half the diners know which they think is delicious, and half choose randomly among the restaurants with the longest line. This supports a *pure* cascade.
- Information cascades may be based on *statistical inference.*

Homework: due December 24, 2020

Consider a situation where there are *two* restaurants, A and B. You think that 50% of all people are **informed**: they know which restaurant is better. 50% are **uninformed**: they guess that the restaurant with the longer line is better. (If the lines are the same length, they choose with uniform probability.) The type of each person who arrives is independent of all the others.

How should you assess the probability that the long line waits at the good restaurant?

1. What are all the possible orders of types **informed** and **uninformed** arriving at the restaurants? (There are 8.)
2. What is the probability of each order of arrival?

3. Considering that the knowers choose A and the guessers may choose randomly, for each order, find the distribution of allocations of diners to restaurants. "Allocation" means which line each person joins. "Distribution" means that for each allocation you must compute the probability of that allocation.
4. Reduce the distribution of allocations (specific people) to the distribution of line lengths at Restaurant A (of course, Restaurant B has 3 minus that length waiting).
5. For each line length, what is the probability you choose the right restaurant?
6. What is the overall probability you choose the right restaurant? Is it better than 1/2?
7. Is choosing the long line a good strategy if you want to eat good food?
8. Are there *pure* cascades in this example?

“Herding” and information cascades

- *Herding* is a more general term used in psychology and sociology to denote imitative behavior.
- In the case of animals like sheep, herding is instinctive. These fields tend to interpret human herding as either instinctive or socially reinforced conformism.
- There may be some cases of instinctive herding in humans, even in economics, but as we see with the restaurant example an information cascade can induce *rational herding*.
- Kleinberg and Easley give many examples:

  Fashions and fads, voting for popular candidates, the self- reinforcing success of books placed highly on best-seller lists, the spread of a technological choice by consumers and by firms, and the localized nature of crime and political movements can all be seen as examples of herding[.] (p. 484)

Network externalities *vs.* information cascade

- Note that network growth also looks like herding behavior.
- Unlike conformism, it is based on rational choice.
- However, unlike information cascades, it is based on *direct benefits*. You can't go wrong by joining the bigger network under Metcalfe's assumptions about the benefits from the network. You *can* make a mistake in the information cascade case, and under some circumstances you will.

An experiment

Two urns are filled with red and blue colored balls. One is 2/3 red, the other 2/3 blue. We pick one with equal probability, and put the other away. We have different people independently draw a ball, look at it without showing anyone, guess the color of the urn so that all can hear, and replace the ball. Then go on to the next person. Each person who guessed correctly gets ¥1000, the others nothing.

- The first person draws a ball. Suppose they see blue. (The same argument follows with colors exchanged if they saw red.)
- To maximize expected profit, they should guess blue.
- Suppose the second person also sees blue. Obviously they also want to guess blue.
- If they saw red, then they have a split sample, and could guess either. Let's suppose they break the tie by guessing the color they saw.

- Now suppose the first two both saw blue. We know that profit maximizing subjects will both report blue. If the third person *also* sees blue, obviously they should guess blue.
- But suppose the third person saw red. In that case, two people say they saw blue, and they have a strong incentive to guess the color they saw. So it's as if the third person drew three balls, two of which are blue and one red.
- In this case, the evidence is that the urn is majority blue (two blue balls to one red ball), so the third person should guess blue—no matter which color they saw.

- How should the fourth person think? The first two guessed blue, and we know they're trustworthy. But the third always wants to guess blue, so the fourth wants to ignore their guess. But no matter what she sees, a majority of the balls she has seen or can deduce the color are blue, so she wants to guess blue regardless of the color she saw.

Homework: due December 24, 2020

1. What does the fifth person think if they see blue? What if red?
2. What happens if the first two split, one red, one blue, and the third sees blue?
3. What can you say about the sequence of guesses in general?

Cascades and math: Bayes' Law

- We know that in such situations the quantitive effect of receiving information (denoted **message**) on our belief about the unknown **state** is given for each possible state by *Bayes' Law*:

$$\mathcal{P}[\mathbf{state}|\mathbf{message}] = \frac{\mathcal{P}[\mathbf{message}|\mathbf{state}]}{\mathcal{P}[\mathbf{message}]}\mathcal{P}[\mathbf{state}]$$

- The logic is that if the probability of a **message** in a given situation **state** is higher than our current assessment of the overall probability that we receive **message**, then the fraction is bigger than 1, and tells us how much we should increase the probability we assess for that **state**.
- In this equation, $\mathcal{P}[\mathbf{state}]$ and $\mathcal{P}[\mathbf{message}]$ are not necessarily *base rates*. We can also think of them as *prior probabilities* based on repeated updating as we receive message after message, up to just before we receive the current **message**.

Cascades and math: Bayesian analysis

- Where do the probabilities come from?
  - *Base rates* are generally measurable by counting.
  - The conditional probability of message given state comes from some scientific model, and some measurement.
  - The conditional probability of state given message is computed from Bayes' Law.

Kleinberg and Easley, Ch. 16, gives an extended discussion of a realistic example of a witness with imperfect visibility testifying about an accident involving a taxicab.

Cascades and math: Bayesian games

- In a situation where people make *interacting decisions* (a *game*), we use optimization theory to characterize their *optimal decisions*, then place probabilities on their choice among optimal decisions.
- An *equilibrium* occurs when the probabilities we place on all optimal choices give all players no reason to change their behavior.
- In *Bayesian game theory*, we suppose people have different types (*e.g.* **informed** *vs.* **uninformed**), and we estimate probabilities of their types and informational state based on what we observe of their behavior and what we know of their incentives for different behavior in each situation.
- A very famous example is the so-called *Chain Store Paradox*, which analyzes an example in which a player strategically induces an information cascade in order to justify behavior that is irrational but profitable.

The chain store paradox stage game

	fight	share
in	$-1, -1$	$1, 1$
out	$0, 2$	$0, 2$

local store payoff listed first

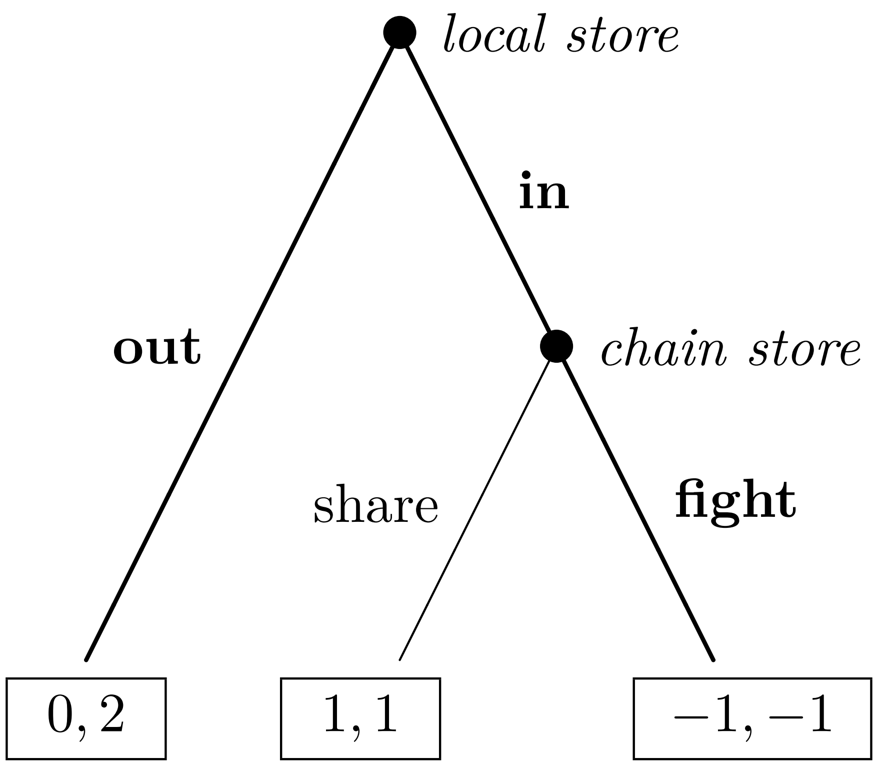

Summary of information cascades

This list includes that of Kleinberg and Easley section 16.7.

- Some situations make cascades very likely.
- Cascades can be wrong.
- Cascades can be based on very little information.
- Cascades can be fragile.