Economics of Information Networks

Stephen Turnbull

Division of Policy and Planning Sciences

Lecture 1: November 16, 2023

Abstract

Introducing the instructor, the course, and the field of network economics.

Everything I need to know I learned in *Economics of Information Networks*

- This is one of the most important courses you will take in the Service Engineering program.
  - That's really just a joke.
  - I just like to say that because it *was* true of my *Basic Data Analysis* course.
- Of course, all the professors say that their course is really useful and interesting ... what's special about *this* course?
- In other words, "Is this course really going to be useful to me?"

I'm glad you asked!

- Information networks are a crucial aspect of economic progress.
- Of course, today the Internet generates a lot of value-added directly, through games and search engines, cloud services and improved efficiency of communication.
- But in fact, all social activities are mediated by networks: communication networks and transportation networks, networks of friends, networks of allies, networks of business contacts.
  - The old saying, "It's not *what* you know, it's *who* you know" (that leads to success in business and society) is a direct reflection of this fact.
- Network effects are an important cause of *agglomeration economies* (which is a fancy way of saying "people and businesses benefit by gathering in cities").

I'm *really* glad you asked!

- In economics, we started by taking a very abstract view. A network is simply a collection of economic agents (households, businesses, and governments) that can trade some good with each other.
- In many markets, simply increasing network size makes the network more valuable.
  - This isn't true for typical private goods: your costs don't decrease unless you actually sell more, and the pleasure of consumption is not dependent on how many other people consume the same good.
  - But for goods like telephones (the more of your friends have one, the happier you are), it is true.
- This simple fact has important implications for network growth and competition in that market.

I'm *still* glad you asked!

- By looking at network *structure*, we can learn a lot about costs of providing a distributed service such as communication and transportation.
- Network structure also has important implications for bargaining power (if everybody knows the same people you know, they don't need to ask for an introduction, and won't owe you any favors) and employment opportunities (if somebody you know knows somebody who has a job you can do, they can connect you).
- Network structure can help to predict alliances in politics and business.
- Network structure can help you avoid the high cost and low efficiency of broadcast media if you can identify a network that connects the people you want to reach.

Objects of analysis of networks

- Simple, abstract network externalities without concern for structure. We just count connected members. Analysis is done via the usual microeconomic foundations of utility functions, cost functions, and profit functions, plus markets.

- Network flows, usually in fixed networks such as communications and transportation networks. This is a specialized area of operations research which is highly developed.

  However, in modern electronic information networks, flow constraints are usually rather secondary, so we will treat this subject only briefly and leave a full discussion to specialized courses.

- Network structure, both fixed and endogenous. Network structure is modeled using *graph theory*.

Brief course description

Goal Understanding of the basic ideas of network analysis using economics, operations research, and graph theory.

Overview of the Lectures We consider basic ideas about modeling networks as graphs. Then we introduce simple economic models, followed by network flow analysis. In the second part (the majority) of the course, we apply these models to various kinds of information networks, including communications networks and the Internet. We will also discuss the nature of *security* in a network environment.

Prerequisites and Language

Prerequisites Although not absolutely necessary, for best results students should have taken college level calculus and linear algebra courses, and an introductory microeconomics course.

Language of Instruction I plan to lecture in English, and original course materials will generally be in English. I will accept and answer questions in Japanese to the extent possible (but my technical vocabulary is relatively weak; it's probably best to use English technical terms where possible).

Homework: Manual Calculation

- Calculation by hand will be a prominent feature of this class. N.B. “By hand” includes use of spreadsheets, but unfortunately I can’t permit that on examinations.
- Intended to improve your understanding and intuition about computations.
- Computers can do calculations more quickly, more accurately, and at far larger scale than any human is capable of, but they are a black box to any but expert software engineers. The “garbage in, garbage out” problem is perhaps the most dangerous fallacy in business research.

Homework: Computational Exercises

- Computational exercises *may* be assigned.
- Intended to help you discover the structure of some information networks.
- The intent is not to make you an expert at using computers.

Resources: URLS

- The primary resource for scheduling and access to resources like videos is the Teams Team for Economics of Information Networks.

- Just about anything you need to know about the class will be on the class home page, It may be easier to use than Teams for older information. `http://turnbull.sk.tsukuba.ac.jp/Teach/EconInfoNet/`. If it's posted on the class home page, "I didn't know (about the assignment, test, *etc.*)" will *not* be an acceptable excuse.

- Another useful URL is my personal calendar, `http://turnbull.sk.tsukuba.ac.jp/schedule.html`. (Note: I haven't been so good at keeping it current recently.)

Resources: Textbooks

- David Easley and Jon Kleinberg, *Networks, Crowds, and Markets* is *required* because it is available online (see home page for link). I recommend you buy it anyway if you have any interest in the technical analysis of information networks – this is a classic with comprehensive coverage of the state of the art when it was published. Exercises will be assigned from this book.
- Oz Shy, *The Economics of Network Industries*. Mostly about network externalities and strategic adaptation to them. Exercises will be assigned from this book. Probably not easily available. If interested, ask me.
- Hal Varian and Carl Shapiro, *Information Rules: A Strategic Guide to the Network Economy*. More an airport business book than a textbook, overlaps substantially with Shy's text, but oriented to the business practitioner than to the academic economist or staff mathematical analyst.

See the class home page for more references.

Introduction: Why Study Networks?

Please read Easley and Kleinberg, Networks, Crowds, and Markets*, Ch. 1. This section is complementary to that chapter – it is not the same!*

The word “network” has become widely used to describe “connectedness”:

- Data networks such as *The Internet*, organizational *intranets*, *local area networks* (LANs), *virtual networks* carried over the public Internet, *etc.*
- *Physical* networks for communication and transportation
- Social networks *mediated by websites* such as Facebook, Mixi, and Twitter
- Networks of companies: manufacturers’ supply chains (*keiretsu*), cooperative *corporate groups*
- *P2P* networks for file and network load sharing
- Networks of *personal contacts* (maintaining/extending them is “networking”)
- Even users of the same wordprocessing software are connected by their *ability to share documents*

Important: Networks are defined by the connections, *not* what is connected.

What is a network?

- A *network* is a *collection of entities connected to each other.* A group of related examples include:
  1. A group of computers connected to another by Ethernet.
  2. The users of the computers in (1) connected by email (a different network, especially if the computers are in a shared university lab).
  3. The users of the computers in (1) connected by Google chat (a different network from either (1) or (2)).
  4. The authority relationships in a company.
  5. A group of people connected by friendship.
  6. A group of Facebook pages connected by “friend” tags (a different network from (5)).
- And something very different is a set of railroad stations connected by track.

What is the economics of networks?

- Better expressed as "economics *on* networks."
- Classical and neoclassical economics focus on *market behavior*.
- The economist's model of a market is a degenerate hybrid network.
  - The *physical* aspect of the hybrid is the distribution of goods and services, and payments.
  - The *informational* aspect is the price.
  - It's *degenerate* because all traders are connected directly to all other traders. The network structure does not constrain behavior.
  - Example: geography may make it hard for two traders to interact (*e.g.*, a retail store and a consumer in different cities). (Neo)classical economics considers those to be *different markets*.
- The economics of networks brings network structure directly into the model of economic behavior. It allows us to think about how changes in network structure affect behavior.

Networks, links, and graphs

- It is often convenient to represent a network in terms of bilateral *links*, *i.e.*, a relation between specific pair of entities (usually called *objects* or *nodes*, and in math, *vertices*).
- Many of the connections in the examples above decompose easily into links.
  - *E.g.*, in the group of friends, each pair of people is friendly with the other. And if some pair of friends should fight, that doesn't dissolve the group: the other pairs are still friends.
  - Even though one email user can address a single message to several others, we can think of this as a set of pairwise links, each containing the author and one of the recipients.

Graphs

- A *graph* is a *collection of objects and links*.
- Other generic names for *object*: *node*, *vertex*. Domain-specific names are also used.
- Other generic names for *link*: *edge*, *arrow*. Domain-specific names are also used.
- Each link connects exactly two objects (which can be the same: this is called a *loop*).
- There are many kinds of graph, depending on the exact nature of links.
- Graph theory focuses on the *links*. The majority of terminology and analysis focus on link structure and properties.

Symmetry

- In some graphs, links are symmetric: if object A is linked to object B, then object B is linked (also, *connected*) to object A. These are *undirected graphs*. (The term *arrow* is never used for these links.)
- In other graphs, links are asymmetric. There can be a *link from A to B*, but no link from B to A. These graphs are called *directed graphs*. In directed graphs the term *arrow* is preferred to other names for *link*.
  - A directed graph can represent an undirected graph by *imposing* symmetry: for every link from A to B, there must be a link from B to A.
  - In discussing directed graphs, the meaning of *connected* may be surprising: A and B are *connected* if there is a link from A to B *or* there is a link from B to A (or both, but one or the other may be missing). This means that in $A \to C \leftarrow B$, A and B are connected even though there is no way to get from A to B

Multigraphs

- In many of the networks we discuss, it makes little sense to have multiple links between two nodes. What does it mean to have *two* (different) friendship relations with a person?
- On the other hand, the *data* may drive the definition. In the case of Facebook data we'll discuss later, there is *friendship* (as well as two kinds of *maintained relationships*.
  - In such cases, there may be multiple links (a long web page may link to another web page twice). A graph in which multiple links are allowed is often called a *multigraph*.
- Multigraphs may be directed or undirected; if it's important to distinguish them, we write "directed multigraph" and "undirected multigraph".

Categories

- In mathematics, an important kind of directed multigraph is a *category*. A category has two special properties: it is
  - *reflexive* (there is always a special link from each object to itself, called the *identity* link), and
  - *transitive*, such that if there is a link from A to B and a link from B to C, there must also be a link from A to C.
- Categories arise naturally when considering *paths* (sequences of links) through a graph.
- We won't be using category theory in this class, but the idea of paths is pervasive.

Abbreviations

- As usual, we abbreviate terms when the specific definition is clear from context. This is frustrating for students, but pervasive in technical discussion.
  - Specifically, *directed*, *undirected*, and *multi*- will be omitted frequently when discussing graphs.
  - The type of graph in examples will often be clear from the diagrams.

Link attributes

- In networks, links frequently have *attributes* or properties.
- Graphs with links possessing a quantitative attribute are called *weighted graphs*. In communications and transportation, *capacity* is very important. A link with low capacity becomes a *bottleneck*, impeding flow, if there are high-capacity links connected to it.
- Graphs with *types* of links are called *colored graphs*. In a multigraph, a group of Internet users may communicate by email, by Twitter, and by posting to their Facebook pages. A and B may be linked by email and Twitter, while B and C are linked by C's Facebook page where B makes comments and by email.
  - A *directed graph* can be represented by an *undirected colored graph* and an order on the objects, so that an "up" link goes from the lower to the higher, and a "down" link goes from the higher to the lower.
- Objects can have attributes, rarely as interesting as link attributes. (*Tokens*, which are "moving attributes" of objects, are a basic component of Petri net theory, used in the theory of *parallel computing*.)

The simplest graph

A single node with no links.

Set representation:

$\{A\}, \{\}$

Adjacency matrix:

	A
A	0

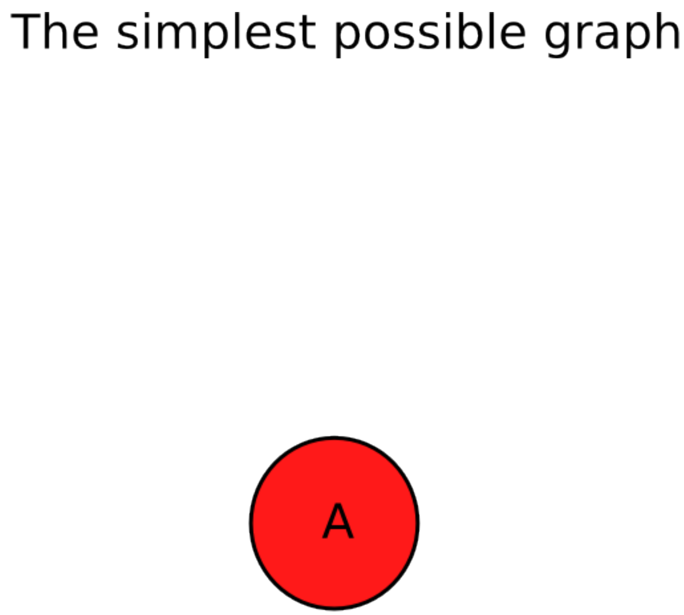

Not a graph

This is not a graph

Each link in a graph must connect two nodes (or possibly a single node with itself).

The simplest interesting graph

One: a simple directed graph
and the simplest category

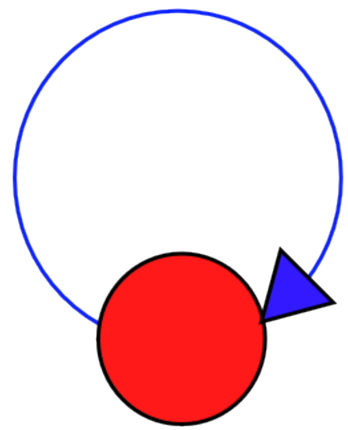

A directed graph with a single node and a single link. Set representation: $\{A\}, \{(A, A)\}$

Adjacency matrix:

	A
A	1

A graph with symmetric and asymmetric links

The notion of graph is abstract, and may be *represented* is several ways: a set of nodes plus a set of pairs of nodes from that set, an adjacency matrix, or a diagram.

Why is this graph called “Envy”?

(Previously it was called “Jealousy”.)

Set representation:

$\{A, B, C\}, \{(A, B), (B, C), (C, B)\}$

Adjacency matrix:

$$\begin{array}{cc|c|c|c|} & & & \text{head} & \\ & & A & B & C \\ \hline & A & 0 & 1 & 0 \\ \hline \text{tail} & B & 0 & 0 & 1 \\ \hline & C & 0 & 1 & 0 \\ \hline \end{array}$$

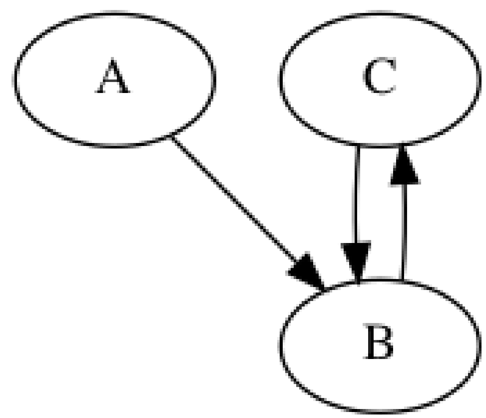

Another with symmetric and asymmetric links

Why is this graph called "infidelity"?

Set representation:

$\{A, B, C\}, \{(A, B), (B, C), (C, B)\}$

Adjacency matrix:

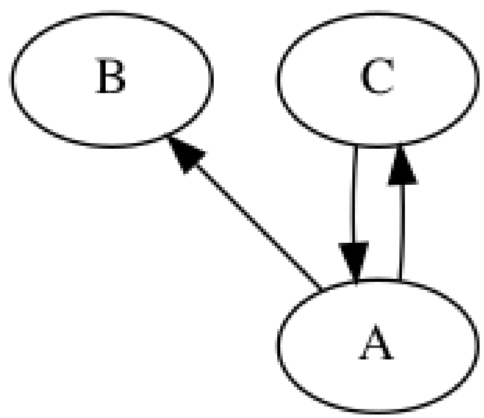

		head		
		A	B	C
	A	0	0	0
tail	B	1	0	1
	C	0	1	0

Homework Submission

1. Submit your homework *by email to*

   `"Economics of Information Networks"`
   `<netecon-hwturnbull.sk.tsukuba.ac.jp>`

   (Note: `netecon-hw@turnbull…`, not `turnbull@sk….`) The `Subject:` should be `0AL0200/01CN901 Homework #1`. (For assignments #2, #3, and so on, adjust the homework number.)

2. Without the class number and the homework assignment in hankaku romaji, your email may get lost due to spam filtering. Use the class number above, even if you are registered according to a different code.

3. Your email must contain your *name* and *student ID number*.

4. You should receive an automatic reply from the "network economics robot." Please save this reply as proof of submission, in case of a hardware or wetware problem in my systems.

5. For simple answers, I *strongly* prefer *plain text* or $\TeX$ notation for expressions and equations to *Word documents* and *HTML*. In plain text, you may write subscripts using functional or programming notation (*i.e.*, X_t becomes `X(t)` or `X[t]`), and superscripts using the caret (*i.e.*, X^t becomes `X^t`) or double-star (X^t becomes `X**t`).

Homework Hints

Many of the tasks assigned in homework are expressed using idioms specific to this class. A few of these words are mentioned below, along with the specific requirements they indicate.

solve Also **give a solution** or **derive**. You *must* show your work. Obvious calculations of common operations, such as the $6 \times 5 \times 4 \times 3 \times 2 \times 1$ in $6! = 6 \times 5 \times 4 \times 3 \times 2 \times 1 = 720$ may be omitted, but even slightly more complex operations such as ${}_6C_3 = \dfrac{6!}{3!3!} = 20$ should be written out. Fractions should be in lowest terms, but do not need to be reduced to decimals. Square roots of perfect squares should be reduced if you recognize them, but all roots *may* be left in the standard notations $\sqrt{x}$ or $x^{\frac{1}{2}}$ or similar as seems most appropriate.

discuss Most important, relate the computation to the real problem in economics (or physics or biology for some of the "toy" examples). Especially mention anything paradoxical, surprising, or extreme about the interpretation of the result in context of the real problem.

compare Like **discuss**, but more specific: you should use statements of the form "*this* is the same as *that*," "*this* is different from *that*," and (best) "*this* is similar to *that*, except ...."

When appropriate use quantitative or ordering comparisions: more/yes, sooner/later, *etc.*

show *expr* is *expr* Often you need to transform one of the expressions to the other. You must show your work, not just "expr $a =$ expr b (same!)"

notation You may define your own notation. For example, in Q#2 you may be asked to compare δ in Q#1 to δ in Q#2. This gets confusing and long winded (*i.e.*, because you write "δ of Problem 1" over and over again). It may be useful to rewrite one of the results by substituting γ for δ everywhere.

Homework 1

- Get Easley and Kleinberg [2010]. Download the current PDF version, or bookmark the HTML version.

  For proof, just send me an email saying you did so.

Network Externalities

We now look at the simplest interesting network model, using the simple "star" market with a network externality. We derive Metcalfe's law and the logistic growth model.

Network Structure of Markets

The simplest network model of a market simply connects traders to each other.

- No distinction between buyers and sellers, or between buying and selling.
- All traders are "in the same place"—as a network, all connected to each other.

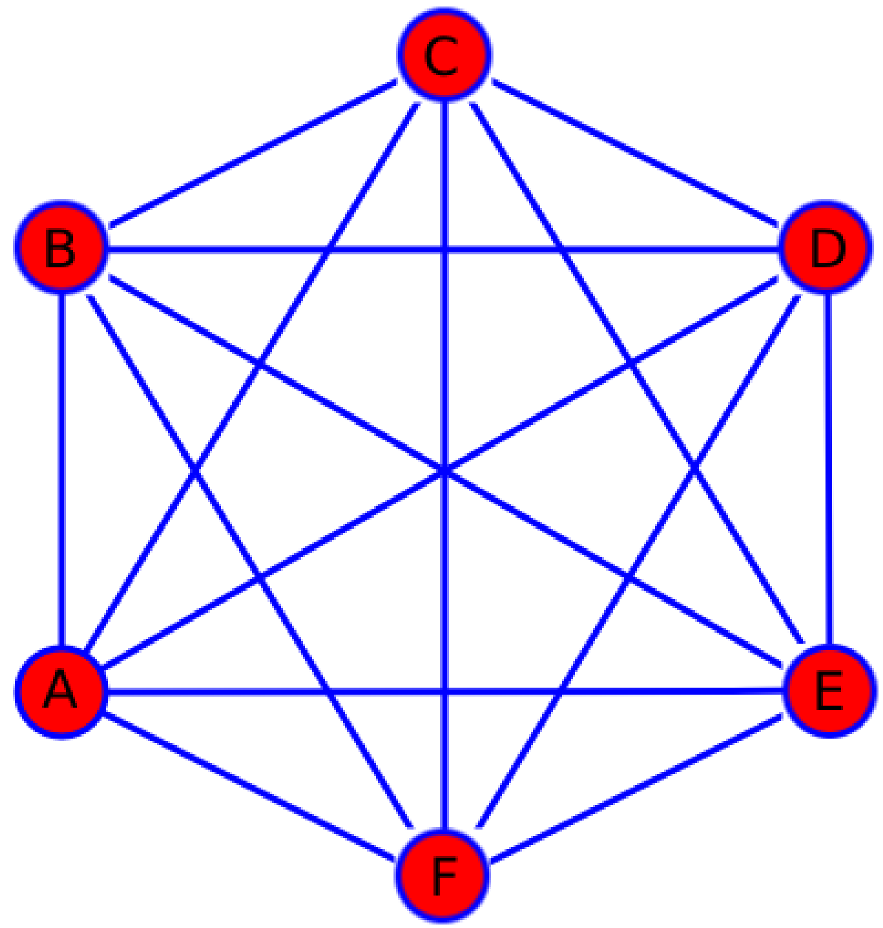

An Alternative Model of the Market

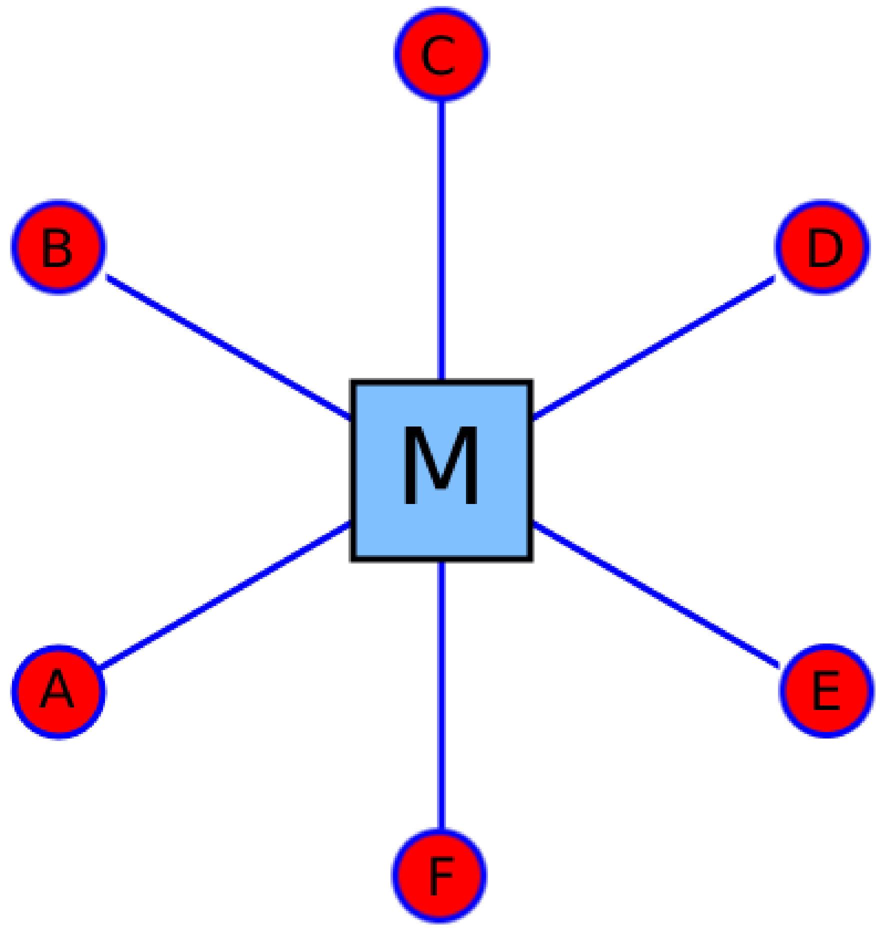

Another simple model connects the traders *through* the market.

The Logic of the Star Market

The links in the two market models are *symmetric*. It's easier to see what this means if we consider a much simpler structure: a set of *barter* relationships.

It's clear that something is missing from the symmetric link: it doesn't provide a way to talk about *equilibrium*, that is, the balance of value given for value received. Symmetry *assumes* balance.

The star arrangement allows the market to assure that balance, and achieve efficiency. Barter won't be efficient because *each* individual trade must balance.

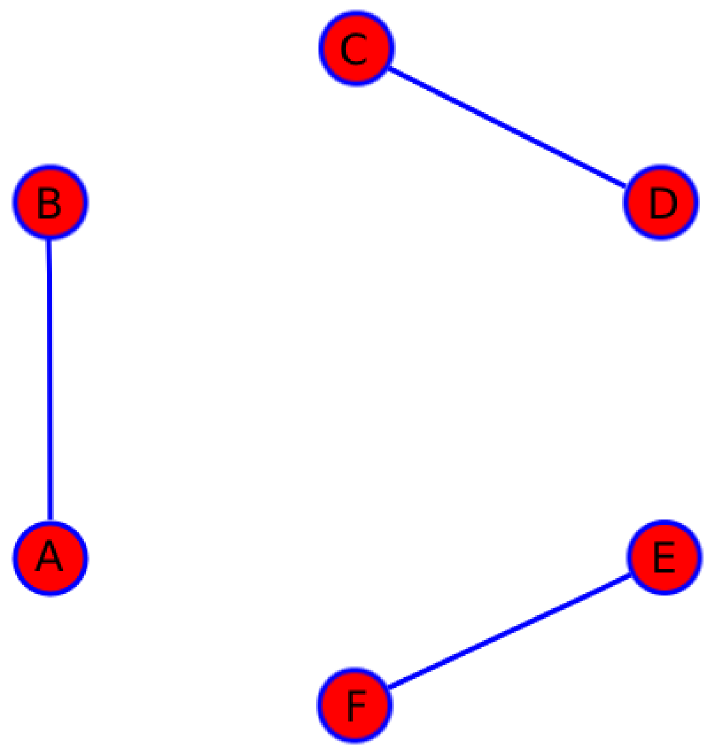

A Directed Graph Approach

Consider one pair in isolation, with directed links indicating transfer of a value from one to the other.

If one link were missing, then balance could not be achieved and the transaction fails.

The directed graph admits a representation of equilibrium as an equation.

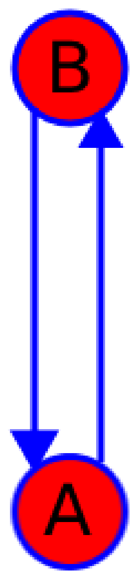

A Market with Money

Reintroducing the market and coloring the graph, blue links represent movement of goods and green the movement of money.

Of course the values must balance.

From the directions of the arrows, infer that A is a buyer and B a seller. There is no need to define *buyer* and *seller* trader types.

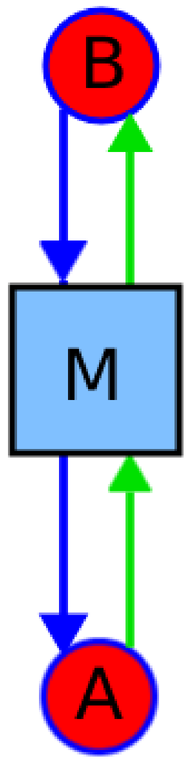

Matching Markets

In *matching markets*, the market develops *because* there are strong complementarities among types.

Agents, not links, are "colored," according to type. Often "color" is represented by position on the right or left of the graph.

The links are directed, representing preferences, not goods flows.

This graph is *two-sided*, meaning both men and women have preferences about partners. The housing market is an example of a *one-sided* matching market.

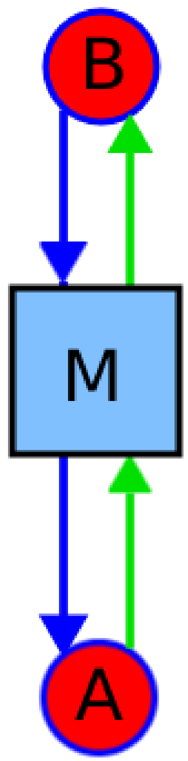

Using Graphs

- Graphs are helpful in understanding the relationships among actors (like buyers and sellers) and institutions (like markets).
- They make clear what equations need to be defined and solved.
- Institutions may contain more detailed graphs. For example, in the stock market not only investors and issuing firms participate, but also market makers, who keep inventories and smooth out variations in supply and demand.

Network Industries

This section loosely follows Shy, The Economics of Network Industries.

- A *network industry* is one which maintains connections among its clients.
  - A market can be thought of as such a service in pure form, allowing its members to compare prices and arrange trades.
  - Most networks are impure, providing connection plus other services.
- Transportation and communication services may be used or not, along with the conceptual connection.
- A software application's file format may be used by a lone user purely to store information, as well as permitting file sharing among users of the same software.
  - Any standard, whether "official" or simply popular, has the same effect of creating a network.
- Networks create markets.

Network Externalities vs. IRTS in Production

- IRTS in production implies that a single large producer is most efficient, by definition. However, with network externalities in consumption, it is both theoretically possible and seen in practice that several providers share a single network.
- A fixed cost with constant marginal cost implies unbounded increasing returns. The model that leads to Metcalfe's law is far less plausible.

Metcalfe's Law

- To the extent that a network merely *provides connections* between users, its value to each user i depends on the set of connections available. We simplify by assuming that it is not the particular set, but rather the size of the set that matters.

- The simplest estimate of the *value of the network* assumes
  - users are symmetric: $U_i(N) = U(N)$
  - users do not discriminate: $U(N) = u(n)$, where $n = |N|$
  - values are additive: $V = \sum_{i \in N} u(N) = nu(n)$
  - individual value is linear: $u(n) = vn$

  If u_i is nonconstant, we say *network externalities* are present. The linear form $u_i(n) = v_i n_i$ provides a very strong network externality.

- *Metcalfe's Law* is immediate:

$$V = vn^2.$$

A Simple Model with a Network Externality

- We assume a potential market of users M, with $|M| = m$.
- The network externality follows Metcalfe's Law:

$$V = n(nv - c),$$

  where V is the total surplus of the industry, n is the number of users connected to the network, v is the value per connection to each user, and there is a cost of c to stay connected to the network.

  – Unlike the usual theory of the firm, there is a dramatic difference between $c = 0$ and $c > 0$ cases.

- The externality is represented by the coefficient n on v (inside the parentheses).

The Initial Coordination Problem

- Consider the inequality

$$u(n) - c = vn - c < 0,$$

which is the condition where a potential user does not want to join the network.

- It's easy to solve for n:

$$n < \frac{c}{v}.$$

- When $c > 0$ and $v > 0$ is small enough, there may be sizeable populations $n > 0$ such that $u(n) - c < 0$, so the market may fail unless at least $\frac{c}{v}$ users can be convinced to join at the beginning.

- If the initial size of the network is at least $\frac{c}{v}$, the dynamics of the network are qualitatively similar for $c > 0$ and $c = 0$.

Industry Dynamics with Network Externalities

- More interesting than the *fact* that there are increasing returns to size of the market on the demand side is the *effect* of these returns on the dynamics of the industry.

- For example, many innovations start with a single inventor, and as others realize that the innovation is useful, it propagates (or diffuses) through the industry (or even the economy as a whole).

  But with a *pure network good* (one which only offers value by connecting to others) there may be a *minimum viable scale* below which the cost of production is not balanced by the value, even though a large network might have very high net value to each user.

- This means that starting the network requires *coordination* (enough users joining the network at introduction), and therefore the normal market mechanism can fail to support the innovation.

Diffusion Dynamics for a Network Good

- We model the dynamics as a differential equation. The *hazard rate* for joining the network is proportional to the net value to the new user: $\alpha(vn - c)$.

- With m the total population of potential users, multiplying by the *non*user population $m - n$ gives the rate of diffusion:

$$\frac{dn}{dt} = \alpha(vn - c)(m - n),$$

  which has the solution

$$n(t) = \frac{m - \frac{c}{v}}{1 + e^{-\alpha v(m - \frac{c}{v})(t - t_0)}}.$$

- In the special case of $c = 0$, we can rearrange to get

$$\frac{dn}{dt} = (\alpha v)n(m - n),$$

  which is the familar logistic model with solution

$$n(t) = \frac{m}{1 - e^{-\alpha m v t}}.$$

The Logistic Growth Path

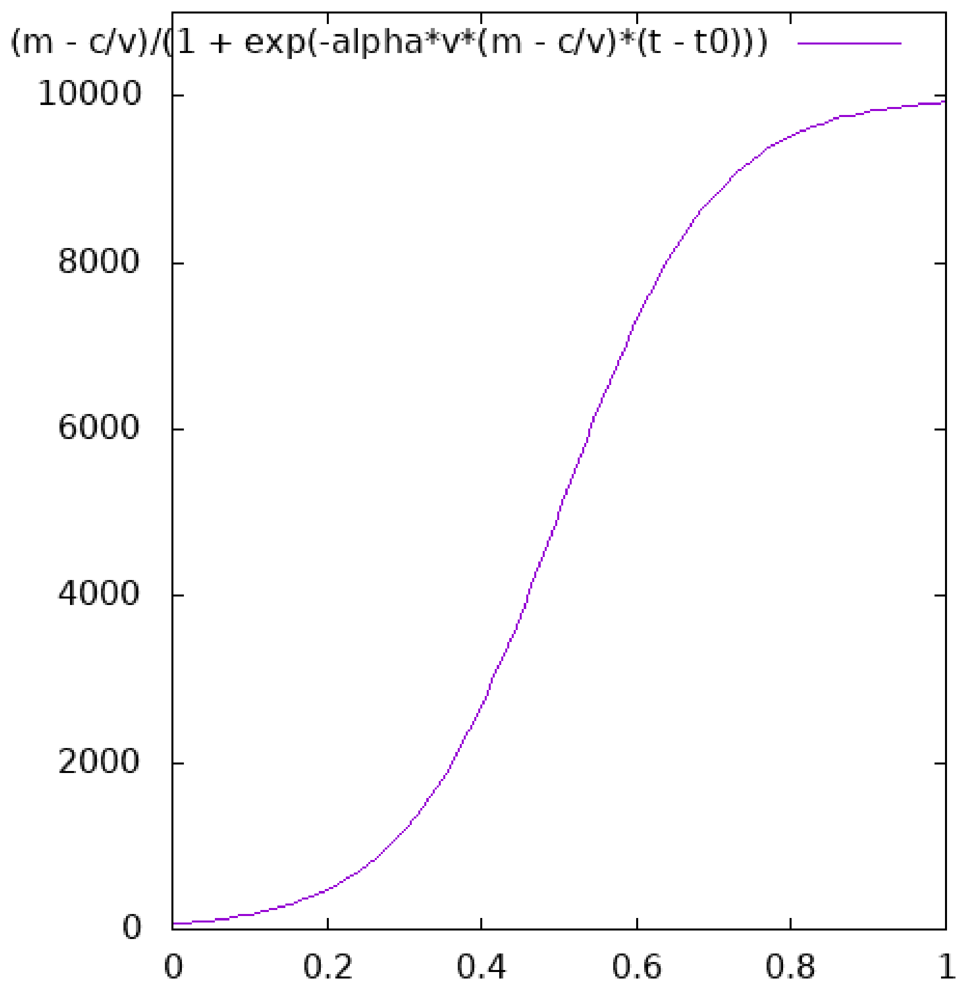

The S-shaped logistic growth path is bounded above and below.

Dynamic Competition between Incompatible Networks

- We consider the duopoly, but the principle applies to industries with more than two firms. We have a total population of potential users of m.

- Let the per user per connection values be $v_1 = v_2 = v$, the cost per connection be $p_1 = p_2 = c$, and the number of users (connections) for the two firms be n_1 and n_2.
  - The notations v_i and p_i ("p" for "price") indicate that in a more sophisticated model these might be differentiated or even strategic variables (especially p_i).

- *Incompatible* means that users on one network are *not* connected to the other. Thus for each user, the value of their network is $u_0 = 0$ if not connected to either, and $u_i(n_i) = v_i n_i - p_i$ if connected to network i.

Adoption Decisions of Users

- We suppose that the diffusion of non-users into each network is proportional to net value as in the monopoly case: $\alpha_i(v_i n_i - p_i)$. Once again we will assume symmetry: $\alpha_1 = \alpha_2 = \alpha$.
  - This assumption is more plausible than the assumptions for the "strategic" variables.
- We assume no switching cost, and that existing users switch from 2 to 1 according to the difference in net values: $\delta((v_1 n_1 - p_1) - (v_2 n_2 - p_2))$.
  - Note this hazard rate may be negative.
  - If you were wondering why the hazard rates for non-users have the same α, this switching can help justify that assumption.
  - In a course in economic dynamics, you'd be asked to show when the model with $\alpha_1 = \alpha_2$ and a high δ is equivalent to $\alpha_1 \neq \alpha_2$ and a lower δ.

The Diffusion Model

- Make all symmetry assumptions, and $n_1 > n_2$ at the start of time.
- Then we have

$$\begin{aligned}\dot{n}_1 &= \alpha(vn_1 - c)(m - n_1 - n_2) + \delta v(n_1 - n_2)n_2 \\ \dot{n}_2 &= \alpha(vn_2 - c)(m - n_1 - n_2) - \delta v(n_1 - n_2)n_2\end{aligned}$$

  Conceptually there are also terms $\pm\delta v \max\{n_2 - n_1, 0\}n_1$ in each differential equation, but on the assumption $n_1 > n_2$, they are zero. On that assumption, we can omit the max in the equations above.

- It is easy to see that $\dot{n}_1 > \dot{n}_2$, and for small enough c, $\dot{n}_1 > 0$. (The last is non-trivial to prove because in the limit non-users and n_2 go to 0.)
- Thus $\frac{d}{dt}(n_1 - n_2) > 0$. $\frac{d}{dt}(m - n_1 - n_2) < 0$ if $\dot{n}_2 \geq 0$, so eventually $\dot{n}_2 < 0$.
- Even with $\dot{n}_2 < 0$, $|\dot{n}_1| > |\dot{n}_2|$, so $\frac{d}{dt}(m - n_1 - n_2) < 0$. $n_2 \to 0$ and $n_1 \to m$.
- Symmetry implies that the opposite conclusions hold if $n_2 > n_1$, so this model is "tippy": whichever network starts out ahead soon crowds out the other.

Dynamic Games

- Mathematical analysis of even the simplest game is quite complex. It's easy to see that if the symmetric model is extended so that each firm can choose price p_i, the firm that starts with greater n_i has a big advantage.
  - As long as that firm is willing to match $p_i = p_j$, it will win the whole market.
  - If the monopoly is expected to continue for a long time, firms may even be willing to offer negative prices.
- If everything is symmetric, the game is very similar to the "War of Attrition", which is known to have only mixed strategy equilibria.

Compatible Networks

- As mentioned, for many networks an *interconnection standard* can be created. This means that (subject to quality of service considerations for the "foreign" users) the network externality is based on the sum of users of all networks in the "internet."
  - Large networks don't have a competitive advantage: several networks of different sizes can share the market.
  - Market structure (number of companies) is more stable.
  - The value to each user is greater (approximately double in the duopoly) so price increase may be more than enough to compensate the leader for allowing interconnection.
- Examples: "The" Internet, protocols such as the "World Wide Web," standards like the "DOM" for web browsers (allows Javascript to work on different browsers) and "ODF" for office automation
- Competition is based on price and service quality.

Standards and "Open Source"

- "Open" standards (no royalty to implement) lead to "open source" implementations.
  - "Poor" or hobbyist programmers write their own implementations and contribute them.
  - Business customers trying to avoid "lock-in" may write their own implementations and contribute them when they are not mission-critical or competitive advantage.
  - Open source businesses may implement to support a further value-added product or service.

Homework 2

1. Give two examples of networks you participate in or use. For each example, describe the network as a graph:
   (a) What are the objects?
   (b) What are the links?
   (c) Is the graph directed? Explain why or why not.
   (d) Is it colored? What do the "colors" represent.
   (e) Is it a multigraph? Explain why or why not.
   (f) Do the objects have important attributes (other than their links)? What are they?

   For full credit, use one example nobody else in the class uses. For each additional person who uses your less unique example, the point score will be multiplied by 0.9.

Homework 3

Mathematicians, like game theorists, like to give their examples cute names that are memorable. Explain the following names (or labels) for graphs shown in Lecture 1.

1. "Envy" (slide 23). (This was originally called "Jealousy," but "Envy" is more accurate in English.
2. "Infidelity" (slide 24).

Homework 4

Suppose that each user i has a random value of being connected to each other user j, $u_i(j)$, and the values are independently and identically distributed. Suppose that users are added to the network in a random sequence.

1. Does Metcalfe's law hold?
2. If so, prove it ...
3. ... and suggest generalizations where it still holds.
4. If not, give a counterexample ...
5. ... and suggest restrictions that might make it hold.