Economics of Information Networks

Stephen Turnbull

Division of Policy and Planning Sciences

Lecture 2: November 15, 2018

Abstract

We review Metcalfe's Law and consider how it affects industry adoption dynamics. Then we look at network flow problems.

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

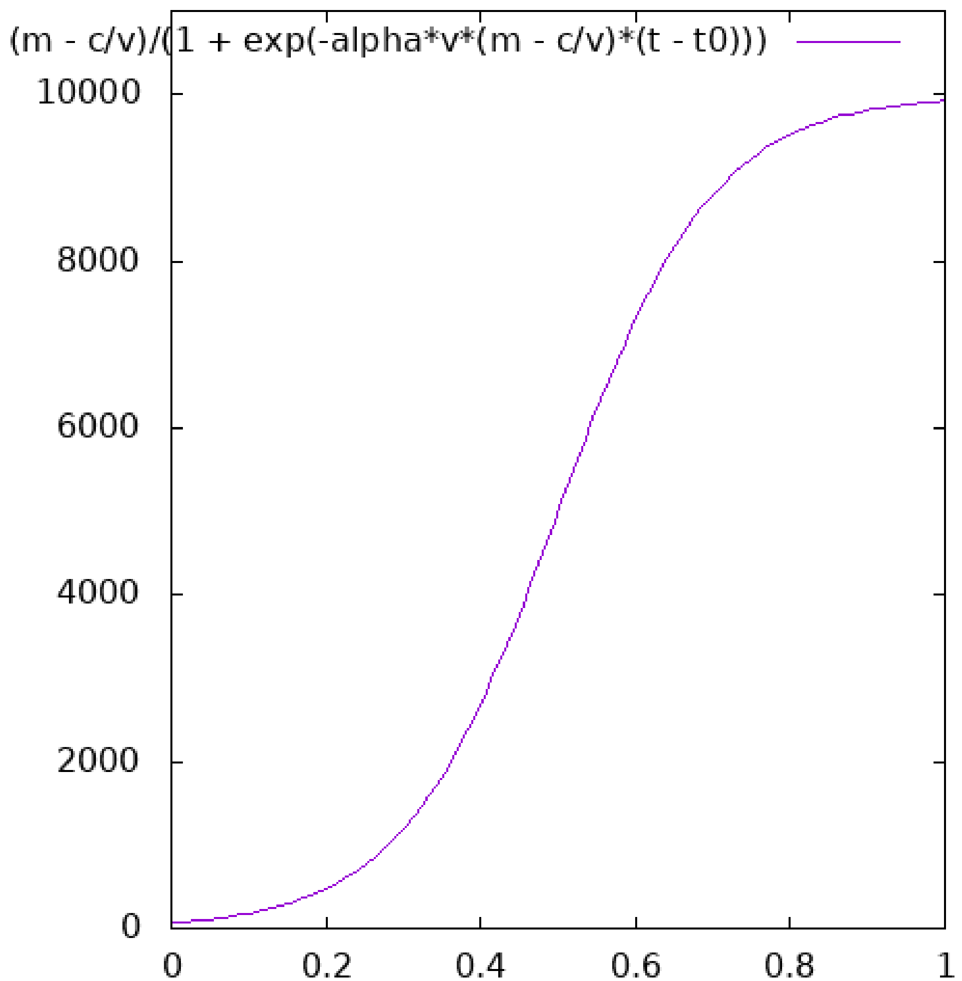

The S-shaped logistic growth path is bounded above and below.

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 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

Network Flow Problems

- From the pure network externality of cellphones and the Internet, we turn to a congestion externality on a network.
- Congestion is a familiar concept: too many users of a facility impose time costs (at a bridge or tollgate, for example) or degraded quality (pollution).
- First we look at the no congestion case.

No Congestion Case

- We use a network model of a *weighted graph*, where each link has a number associated with it. The generic term is "weight," but we will interpret it as *capacity*.
  - Both *directed* and *undirected* graphs are useful. Directed graphs are used in controlled situations of production, such as an oil refinery, and for automobile traffic. Undirected graphs are used in situations such as the effect of tsunamis on river systems.
- Each link has a variable *flow* associated with it. If the graph is undirected, it may be negative. If directed, it must be positive. The capacity is the upper limit on the absolute value of the flow through the link.
- Two kinds of analysis are done:
  - Maximum flow (logistics, transportation and communication networks), where the *demand* is given.
  - Actual flow (transportation planning, flood prevention), where the *supply* is given.

Network Flow Example

"0" and "1" are the labels for the *source* and the *sink* (In general there could be several of each.) The link weights are capacities. The problem is to maximize flow from 0 to 1. In this simple example it's easy to see that the upper path can handle $\min\{40, 30\}$ and the lower $\min\{20, 45\}$, for a total maximum flow of 50.

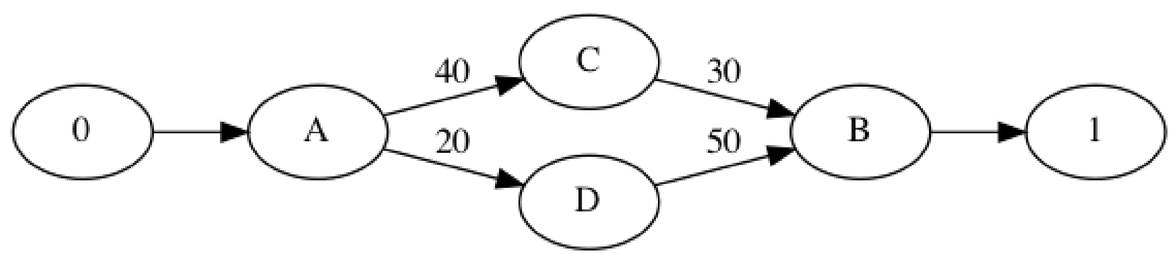

Linear Programming Solution

This optimization can be analyzed as a *linear programming problem* (LP).

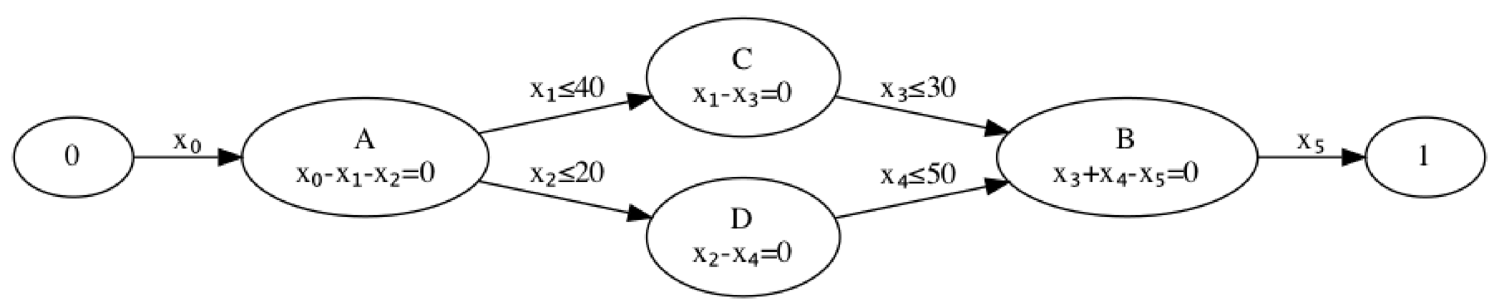

Transform the Network Flow Problem to LP

1. For each node, there is a constraint that the difference between the sum of flows *into* that node and the sum of flows *out of* that node equals zero.
2. For each *internal* link i, add two constraints on link flow x_i: $x_i \leq w_i$ and $x_i \geq 0$ (directed graph) or $x_i \geq -w_i$ (undirected graph).
3. The object is to Maximize the sum of flows into the sink, 1.

Fundamental Theorems

Assuming the graph of the network is connected,

1. Because the number of equality constraints is not greater than the number of links (flow variables), and the variable limits are non-empty, the constraint set is always non-empty.
2. A variant of the simplex method solves the problem efficiently.

The Braess Paradox

- We now consider a *game* on a network which leads to the paradoxical result that adding capacity to a network can cause it to have *lower* flow in equilibrium.
- This requires congestion in the links.

Congestion

- Congestion is a familiar concept: too many users of a facility impose time costs (at a bridge or tollgate, for example) or degraded quality (pollution).
  - We typically model this as a cost to each user based on the total number of users:

    $$U_i(n) = v_i - c_i n$$

    which is very similar to the model leading to Metcalfe's law, except that the coefficient on n is negative.
  - In some cases (*e.g.*, pollution) it's useful to consider the users' quantity choices as well:

    $$U_i(q) = v_i(q_i) - c_i \sum_{j=0}^{n} q_j.$$

Latency

- In the network flow problems we solved by linear programming, the capacities are given, and the "liquid" automatically flows from the source(s) into the network and out to the sink(s) at rates up to the maximum.
- We did not consider *how long* it takes to get from a source to a sink. This parameter is called *latency*.
- In a production context, we often trade latency for capacity. By waiting a while for output, we may be able to increase *throughput*.
- This will increase the rate of output once the process gets started, and this is the foundation of the economies of scale due to "continuous process production" technologies.

The Braess Example

- In *consumption* (*e.g.*, streaming videos or commuting to school or work) we care about latency. This drives Braess's example.
- There is road network used by 4000 drivers to get from A to B.
- There are two routes, one via C and one via D. $A \to C$ and $D \to B$ are high-capacity highways that take 45 minutes regardless of traffic.
- $C \to B$ and $A \to D$ are bridges of small capacity such that the time to cross increases with the number of cars trying to cross the bridge. We denote the number of cars crossing A to D and continuing to B (path ADB) as x, and the number crossing C to B and continuing to B (path ACB) as y. The bridges have the same capacity, and the time to cross is the same: $x/100$ for ADB, and $y/100$ for ACB.
  - We'll see later why we name the paths rather than the links. It's a notational convenience.

Braess Example Network

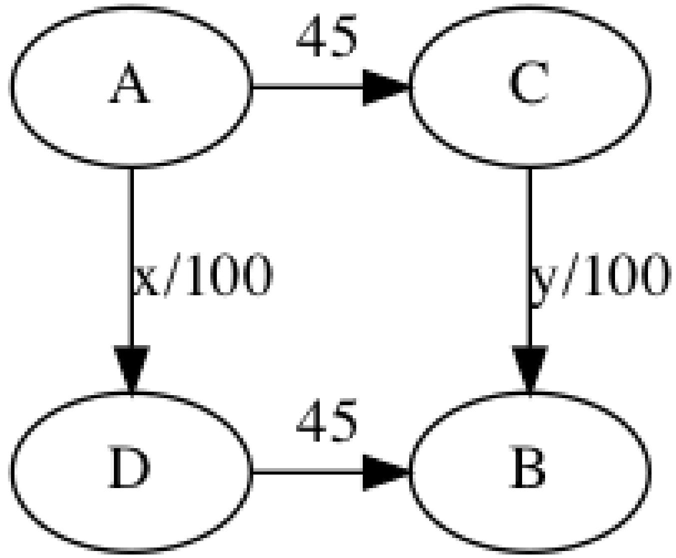

Here's a diagram of the network.

Nash Equilibrium

- We want each driver to choose a route (path) that minimizes her travel time.
- Then we have one constraint: $x + y = 4000$, and two travel time formulæ: time(ADB) $= 45 + x/100$ and time(ACB) $= y/100 + 45$.
  - Suppose $x > y$. Then time(ADB) $>$ time(ACB), and some of the drivers will want to change route from ADB to ACB. The opposite is true if $y > x$.
    * This is true for $x = 4000$ or $y = 4000$, no "corner" equilibria.

    So the equilibrium share is $x = y$.
  - Then $4000 = x + y = 2x$, so $x = y = 2000$.
  - This makes sense for society, as it has the smallest maximum travel time over all commuters. In fact everybody takes the same amount of time: $45 + 2000/100 = 65$ minutes.

The Paradox

- Now suppose that a new bridge of very high capacity is built between D and C. It takes 0 time to cross the bridge. (Relaxing this unrealistic assumption will be a homework problem.)
  - There are three routes: A to D to B (taken by x drivers), A to C to B (taken by y drivers), and A to D to C to B crossing all three bridges (taken by z drivers).

The Paradox Network

Here's a diagram of the network inducing the paradox.

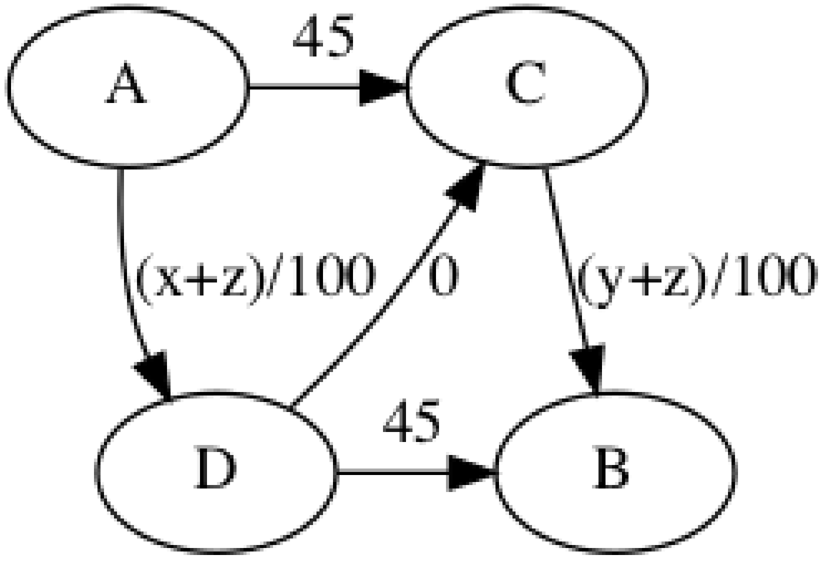

Equilibrium in the Paradox

- The "everybody goes to work" constraint is $x + y + z = 4000$.
  - time(ADB) $= (x + z)/100 + 45$, time(ACB) $= (y + z)/100 + 45$, and time(ADCB) $= (x + z)/100 + (y + z)/100$.
  - Suppose $x > 0$. Then time(ADCB) $<$ time(ADB), and some of those drivers will want to switch from ADB to ADCB. This is true for *any* $x > 0$, so the only equilibrium x is $x = 0$. Similarly, equilibrium y is $y = 0$.
  - time(ADB) = time(ACB) = time(ADCB) = 80 minutes!

Homework #6

Due: November 22, 2018 at 11:00. Submit to `turnbull@sk.tsukuba.ac.jp` with `Subject: Homework #7 01CN901`.

Consider the following variants on the Braess model.

1. Consider the original network flow problem (without the new bridge), but with asymmetric capacities for the bridges: the time to cross $C \to B$ is longer: $y/80$. Is there an equilibrium? If so, what is it? If not, why not?

2. Generalize the formula for equilibrium in part 1 to the case where the time to cross $A \to D$ is αx and to cross $C \to B$ is βy. Ignore boundary conditions (*alpha* and *beta* are "too different" for the formula to work).

3. Discuss the boundary conditions for the formula in part 2.

4. In the three-bridge network, formulate the relevant equations when, instead of labelling paths with traffic levels, we label links with traffic levels:

$A \to D$: x

$C \to B$: y

$D \to C$: z Compare this notation with the notation used in class. Which

$A \to C$: u

$D \to B$: v

is easier to solve? Which provides more information, or are they the same?

You can save yourself a lot of annoyance by choosing appropriate notation! You're also likely to make your readers (including advisors and examiners!) happier.

5. Realistically, bridges take time to cross. Suppose in the three-bridge problem it takes one minute to cross the bridge $D \to C$. How does the equilibrium change?

6. Is there a cost (in minutes) to crossing $D \to C$ large enough so that no driver will use that bridge in equilibrium? If so, what is the smallest such time? What are the equilibria for slightly smaller times?

A Brief History of Compatibility in Computers

- The earliest computers were one-of-a-kind, custom built for a specific customer.
  - An interesting example of early mass production were the code breaking machines (partly mechanical and partly electric) used by the U.S. Navy to break the "Enigma" code used by German submarines in WWII.
- In the late 1950s, IBM strategists realized that there was a large market for electronic computers, and determined to capture it with the design of System/360 and the companion software OS/360.
  - The head of Univac Co. which made the first digital computer thought the world would never need more than 100 computers.
  - System/360 was a family of CPUs, memory banks, mass storage, and other peripherals that could be matched in different ways: a fast CPU with small memory and storage for mathematical computations, or a slower CPU with fast I/O and larger memory and storage for business database applications, and so on.
  - System/360 hit 100 orders in 4 months.

History of Compatibility in Computers, cont.

- This not only created a large market for IBM, but also opened the door to *plug-compatible manufacturers*, who could not produce a whole competitive computer, but did produce better or cheaper components.
  - Amdahl in the U.S. and Hitachi and Fujitsu in Japan even went so far as to produce complete compatible hardware that ran OS/360 and application software.
- Apple popularized the personal computer with the Apple II (largely aided by Dan Bricklin's invention of the electronic spreadsheet), but again IBM created several large markets with the open architecture of the IBM PC, spawning both a huge software industry as well as the "clone" manufacturers, and many innovative peripherals.
- Today such open architectures are the rule, although typically they are protected by "patent pools" of the IP rights of several companies.
  - Consider standards like WiFi, Bluetooth, and USB (to name only those familiar to consumers).

Adoption with network externalities

- Describe the situation with a simple game:

	New	Old
New	α, α	γ, δ
Old	δ, γ	β, β

Table 1: Static technology adoption game

where the network externalities are indicated by the parameters $\alpha > \gamma$, $\alpha > \delta$, $\beta > \gamma$, and $\beta > \delta$ (that is, matched technology is always better than mismatched technology for both players).

- Two inefficient equilibria can arise:

Excess inertia $\alpha > \beta$ but (Old, Old) is observed.

Excess momentum $\alpha < \beta$ but (New, New) is observed.

Concrete instances

- With the constraints on the parameters, it is approximately a game of pure coordination, with the equilibria described later.
- Without the constraints, it is a generalized game of coordination, where there is a social benefit if both players play the "right" strategy. (Normally in game theory we associate "right" with "equilibrium," or sometimes "efficient." Here the meaning is stronger: "efficient and fair.")
- This schema fits both the *game of pure coordination* (left) and the *prisoners' dilemma* (right):

	New	Old
New	4, 4	0, 0
Old	0, 0	1, 1

	New	Old
New	4, 4	0, 5
Old	5, 0	1, 1

Table 2: Two generalized coordination games

Solutions

- It is easy to solve these games.

Pure coordination
- By process of elimination, both mismatched cases are not equilibrium: both players want to deviate.
- It doesn't matter that if they both deviate at the same time, they end up mismatched in disequilibrium again. This is the "best response (or Cournot) dynamic," but it is a poor algorithm for finding solutions. Alternating Cournot works in this case, but not necessarily in others.
- The two matched cases are both strict equilibria.
- The equations for the mixed strategy equilibrium have a unique solution, also an equilibrium.
- The (New, New) equilibrium is "good," the others are "bad."

Prisoners' dilemma
- Old is a strictly dominant strategy for both players, so the (Old, Old) equilibrium is unique.
- It's "bad," in fact, the worst possible.

Stories

- More important than solving games and assessing their equilibria is "telling stories" about how those payoffs might come to be.
- Both example games are symmetric, so we tell the story from the point of view of one. The other is derived by exchanging the strategies and payoffs.

Stories

Pure coordination Here, we are probably looking at *consumers* of the technology. The basic benefits of the two variants are similar, so $\gamma = \delta = c$, and for convenience of computation we set $c = 0$. Of course we could just add c to all parameters, get the same basic schema, and it would have the same equilibria.

Then we suppose the New variant of the technology has better communication capabilities, or potentially attracts more users, giving higher network externality benefits.

Prisoners' dilemma How can we tell this story?

We need to look at the *other* side of the market. This is a case of the *network providers* who can adopt new technology.

But how is it that the Old variant can have higher payoffs in competition with the New? Note that New has lower payoffs in competition with Old than sticking with Old. So perhaps in competition Old is attracting all the customers away from New (New probably starts smaller), and ends up as a (near-) monopolist. Metcalfe's Law justifies the huge payoff.

Homework #7

Due: November 22, 2018 at 11:00. Submit to `turnbull@sk.tsukuba.ac.jp` with `Subject: Homework #7 01CN901`.

In class, we discussed how the *game of pure coordination* and the *prisoners' dilemma* could be interpreted as models of technology adoption in markets with network externalities. Now let's consider the *asymmetric* game *battle of the sexes*, with payoffs

	New	Old
New	$4, 1$	$0, 0$
Old	$0, 0$	$1, 4$

Table 3: Battle of the sexes

1. What are the "good" and "bad" outcomes of this game?

2. How is this game similar to the *game of pure coordination*?

3. How is this game different from the *game of pure coordination*?

4. Is it useful to compare this game to the *prisoners' dilemma*? If so, what lessons do you draw from the comparison? If not, why isn't it useful for comparison?

5. What are the equilibria of this game?

6. "Tell a story" about how these payoffs might arise if two *end users* of the technology are playing the game.

7. "Tell a story" about how these payoffs might arise if two *providers* of the technology are playing the game.

Note: your answer to the last two questions may be "I can't tell a sensible story." If so, describe your difficulty.