When Knowing Early Matters: Gossip, Percolation and Nash Equilibria

The abstract idea of "information spreading through a network" arises in many contexts:

- gossip algorithms (designed to keep nodes of a decentralized network updated about information used to maintain the network)
- first passage percolation (FPP)
- epidemics

as well as harder-to-model contexts such as the diffusion of technological innovations or of ideologies.

I will describe a new "economic game theory" aspect via non-rigorous arguments. First, we review FPP.

FPP as spread of information

- n agents
- rates $\theta_{i,j} \geq 0$ for agents i, j
- At $t = 0$ one agent i_0 has an item of information
- information spreads between agents: from i to j at (stochastic) rate θ_{ij}.

So this is FPP on a finite set with Exponential (rate θ_{ij}) random times; the set becomes a network with edges (i, j) such that $\theta_{ij} > 0$. Assuming that the network is connected, the information reaches each agent.

Such processes are well understood, from many viewpoints. Write $\theta = (\theta_{ij})$ and consider the random time $T = T^n_\theta$ that the information reaches a uniform random agent. For a sufficiently symmetric ("transitive") network the distribution of T^n_θ does not depend on i_0.

We will be interested in the order of magnitude of the **window width** w_θ^n over which $P(T \leq t)$ increases from near 0 to near 1.

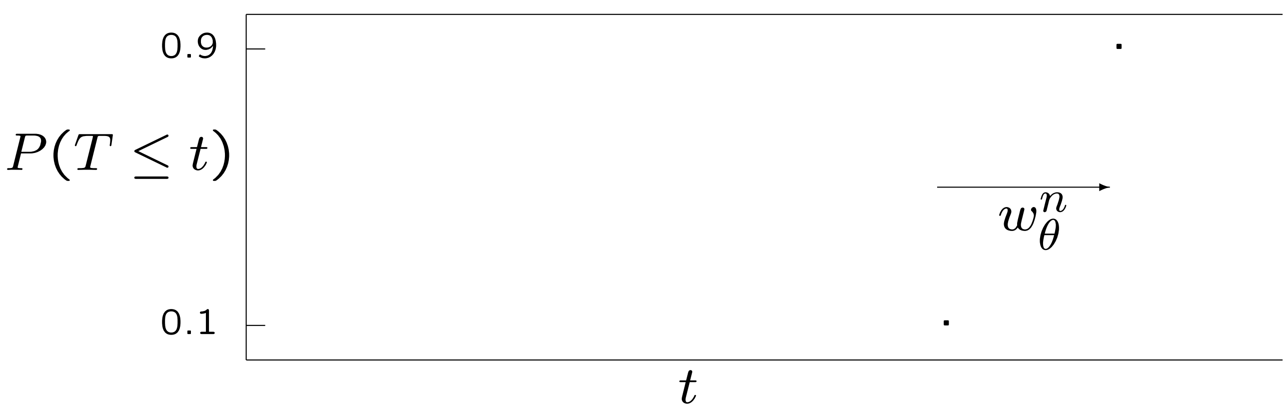

General framework: communication games with rank based rewards

• n agents (results in the $n \to \infty$ limit)

• New items of information arrive at times of a rate-1 Poisson process; each item comes to one random agent.

Information spreads between agents in ways to be described later [there are many variants], which involve communication costs paid by the *receiver* of information, but the common assumption is

• The j'th person to learn an item of information gets reward $R(\frac{j}{n})$.

Here $R(u), \ 0 < u \leq 1$ is a decreasing function with $R(1) = 0; \quad 0 < \bar{R} := \int_0^1 R(u)du < \infty$.

Intended as toy model for gossip or insider trading.

Assuming connectivity, the total reward from each item is $\sum_{j=1}^{n} R(\frac{j}{n}) \sim n\bar{R}$.

Write

$$\textbf{payoff} = \textbf{benefit} - \textbf{cost}$$

where these **boldface** quantities are "per agent per unit time". Note **benefit** $= \bar{R}$ is fixed. And **cost** can be made arbitrarily small by simply communicating less often (Nuance of model: i calls j at a given price and learns all items j knows – maybe 0 new items, maybe 2 new items).

Thus in the "social optimum" protocol, agents communicate slowly, giving payoff arbitrarily close to $\bar{R}$. But if agents behave selfishly then one agent may gain an advantage by paying to obtain information more quickly, and so we seek to study Nash equilibria for selfish agents.

Distinguish one agent as **ego**. Recall that the characterization of a **Nash equilibrium** is that **ego** cannot increase his payoff by deviating from the equilibrium strategy. Instead of allowing arbitrary strategies we allow only specified strategies depending on a parameter.

We will see that, depending on the structure of communication costs, any of the following possibilities may occur in the $n \to \infty$ limit.

- efficient: Nash payoff = social optimum payoff, that is Nash cost = 0.
- wasteful: 0 < Nash payoff < social optimum payoff
- totally wasteful: Nash payoff = 0.

Here are two basic examples.

Example: the complete graph case

Network communication model: Each agent i may, at any time, call any other agent j (at cost 1), and learn all items that j knows.

Poisson strategy. The allowed strategy for an agent i is to place calls, at the times of a Poisson (rate θ) process, to a random agent.

Result. In the $n \to \infty$ limit the Nash equilibrium value of θ is

$$\theta^{\text{Nash}} = \int_0^1 (1+\log(1-u))R(u)du = \int_0^1 r(u)g(u)du > 0 \quad (1)$$

where $g(u) = -(1-u)\log(1-u) > 0$
and $r(u) = -R'(u) \geq 0$.

In particular the Nash equilibrium payoff $\bar{R} - \theta^{\text{Nash}}$ is strictly less than the social optimum payoff $\bar{R}$ but strictly greater than 0. So this is a "wasteful" case.

Example: the nearest neighbor grid

Network communication model: Agents are at the vertices of the $N \times N$ torus (i.e. the grid with periodic boundary conditions). Each agent i may, at any time, call any of the 4 neighboring agents j (at cost 1), and learn all items that j knows.

Poisson strategy. The allowed strategy for an agent i is to place calls, at the times of a Poisson (rate θ) process, to a random neighboring agent.

Result. The Nash equilibrium value of θ is such that

$$\theta_N^{\text{Nash}} \sim N^{-1} \int_0^1 g(u) r(u) du \qquad (2)$$

where $g(u) > 0$ is a certain function and $r(u) = -R'(u) \geq 0$.

So here the Nash equilibrium payoff $\bar{R} - \theta_N^{\text{Nash}}$ tends to $\bar{R}$; this is an "efficient" case.

Discussion

These results suggest many questions . . . . . .

- Wouldn't it be better to place calls at regular time intervals?

- In the grid context, what about the case where an agent can call any other agent but the cost is a function of distance?

- What about the symmetric model where, when i calls j, they exchange information?

- Why in formulas (1,2), do we see decoupling between the reward function $r(u)$ and the function $g(u)$ involving the rest of the model?

- In the nearest-neighbor grid case, wouldn't it be better to cycle calls through the 4 neighbors?

- What about non-transitive models, such as social networks where different agents have different numbers of friends, so that different agents have different strategies in the Nash equilibrium?

- To model gossip, wouldn't it be better to make the reward to agent i depend on the number of other agents who learn the item from agent i?

- To model insider trading, wouldn't it be better to say that agent j is willing to pay some amount $s(t)$ to agent i for information that i has had for time t, the function $s(\cdot)$ not specified in advance but being a component of strategy and hence with a Nash equilibrium value?

Methodology for formulas

Allowing agents' behaviors to be completely general makes the problems rather complicated (e.g. a subset of agents could seek to coordinate their actions) so in each specific model we restrict agent behavior to be of a specified form with a parameter θ; the agent's "strategy" is just a choice of θ. If all agents use the same parameter value θ then the spread of each item of information through the network is as some model-dependent FPP process. So there is some function $F_{\theta,n}(t)$ giving the proportion of agents who learn the item within time t after the arrival of the information into the network. Now suppose one agent **ego** uses a different parameter value ϕ and gets some payoff-per-unit-time payoff(ϕ, θ). The Nash equilibrium value θ^{Nash} is the value of θ for which **ego** cannot do better by choosing a different value of ϕ, and hence is the solution of

$$\frac{d}{d\phi}\text{payoff}(\phi, \theta)\bigg|_{\phi=\theta} = 0. \tag{3}$$

Obtaining a formula for payoff(ϕ, θ) requires knowing $F_{\theta,n}(t)$ and knowing something about the geometry of the sets of informed agents at time t, but does not require any more.

I will outline heuristics for the order-of-magnitude behavior in the two basic examples and a third example.

The complete graph case. FPP on complete graph has elementary analysis (Kendall 1957). For us, key point is that given the realization of FPP ($\theta = 1$) the time T_n to reach random agent has $T_n - t_n \to T^*$ where the limit has logistic distribution

$$P(T^* \leq t) = \frac{e^t}{1 + e^t}, \ -\infty < t < \infty.$$

Write θ_n for Nash equilibrium rates; so window width w_n is order $1/\theta_n$. Suppose $w_n \to \infty$; then **ego** could call at fixed slow rate ϕ and have

$$\text{cost} = \phi; \quad \text{benefit} \to R(0).$$

But this is better than the payoff $\bar{R}$ to other agents.

Conclusion; w_n bounded, so θ_n bounded away from 0, so wasteful.

The $N \times N$ nearest neighbor grid

The *shape theorem* for FPP on the infinite lattice started at the origin says that the random set $\mathcal{B}_s$ of vertices reached before time s grows linearly with s, and the rescaled set $s^{-1}\mathcal{B}_s$ converges to a limit deterministic convex set $\mathcal{B}$.

So when agents call at rate $\theta = 1$ the window width w_N is order N. The time difference of learning item for two neighbors of **ego** is $\asymp$ 1, so the increased reward to **ego** by calling neighbors faster is at most $\asymp 1/N$. This is unaffected by making each agent's call rate $= \theta_N$. So **ego** is willing to call at rate $\asymp 1/N$, giving $\theta_N^{\text{Nash}} \asymp 1/N$.

Here's the first example that needs some thought.

The $N \times N$ torus with short and long range interactions

Network communication model. The agents are at the vertices of the $N \times N$ torus. Each agent i may, at any time, call any of the 4 neighboring agents j (at cost 1), or call any other agent j at cost $c_N \geq 1$, and learn all items that j knows.

Poisson strategy. An agent's strategy is described by a pair of numbers $(\theta_{\text{near}}, \theta_{\text{far}}) = \theta$:

at rate θ_{near} the agent calls a random neighbor

at rate θ_{far} the agent calls a random non-neighbor.

This model obviously interpolates between the complete graph model ($c_N = 1$) and the nearest-neighbor model ($c_N = \infty$).

First let us consider for which values of c_N the nearest-neighbor Nash equilibrium (θ_{near} is order N^{-1}, $\theta_{\text{far}} = 0$) persists in the current setting. When **ego** considers using a non-zero value of θ_{far}, the cost is order $c_N \theta_{\text{far}}$. The time for information to reach a typical vertex is order $N/\theta_{\text{near}} = N^2$, and so the benefit of using a non-zero value of θ_{far} is order $\theta_{\text{far}} N^2$. We deduce that

if $c_N \gg N^2$ then the Nash equilibrium is asymptotically the same as in the nearest-neighbor case; in particular, the Nash equilibrum is efficient.

Let us study the more interesting case

$$1 \ll c_N \ll N^2.$$

The argument has three steps.

1. Consider the *window width* w_N of the associated percolation process at the Nash equilibrium $(\theta_{\text{near}}^{\text{Nash}}, \theta_{\text{far}}^{\text{Nash}})$. Suppose **ego** deviates from the Nash equilibrium by setting his $\theta_{\text{far}} = \theta_{\text{far}}^{\text{Nash}} + \delta$. The increased benefit to **ego** is order δw_N and the increased cost is δc_N. At the Nash equilibrium these must balance, so

$$w_N \asymp c_N.$$

2. Now consider the difference ℓ_N between the times that different neighbors of **ego** are reached. Then ℓ_N is order $1/\theta_{\text{near}}^{\text{Nash}}$. Write $\delta = \theta_{\text{near}}^{\text{Nash}}$ and suppose **ego** deviates from the Nash equilibrium by setting his $\theta_{\text{near}} = 2\delta$. The increased benefit to **ego** is order ℓ_N / w_N and the increased cost is δ. At the Nash equilibrium these must balance, so $\delta \asymp \ell_N / w_N$ which becomes

$$\theta_{\text{near}}^{\text{Nash}} \asymp w_N^{-1/2} \asymp c_N^{-1/2}.$$

3. Finally we need to calculate how the window width w_N for FPP depends on $(\theta_{\mathsf{near}}, \theta_{\mathsf{far}})$, and it turns out

$$w_N \asymp \theta_{\mathsf{near}}^{-2/3} \theta_{\mathsf{far}}^{-1/3}.$$

We have 3 equations for 3 unknowns, and we solve to find

$$\theta_{\mathsf{near}}^{\mathsf{Nash}} \text{ is order } c_N^{-1/2} \text{ and } \theta_{\mathsf{far}}^{\mathsf{Nash}} \text{ is order } c_N^{-2}.$$

In particular the Nash cost $\asymp c_N^{-1/2}$ and the Nash equilibrium is efficient.