Information Cascades on Arbitrary Topologies

Jun Wan, Yu Xia, Liang Li and Thomas Moscibroda

Presented by Oliver Richter


Influence
you can make a difference

Example Restaurant

Option A

Option B

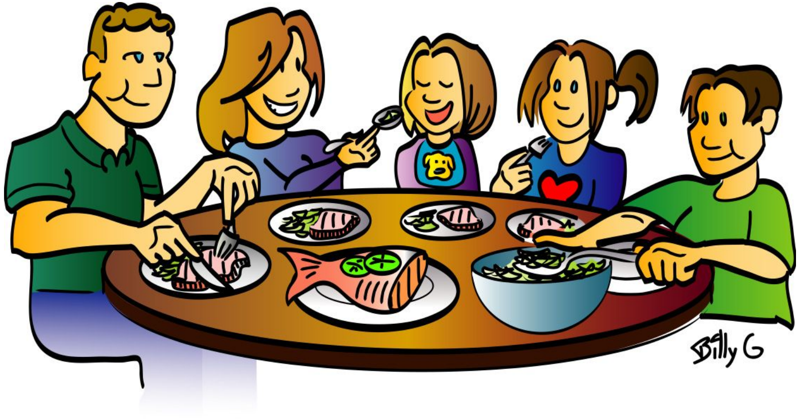

Other Examples

Mass movements

Fashion

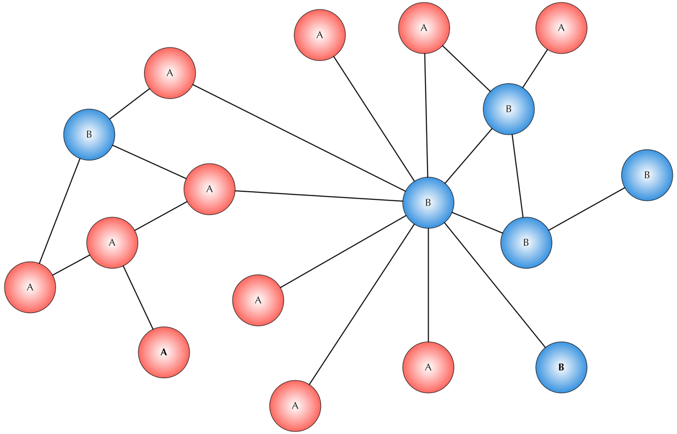

Voting

Books

The Model

50% chance

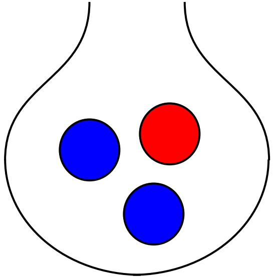

“Majority blue”

50% chance

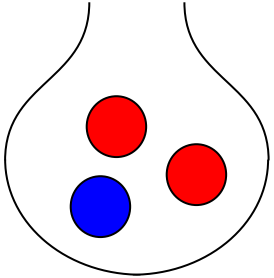

“Majority red”

The Model

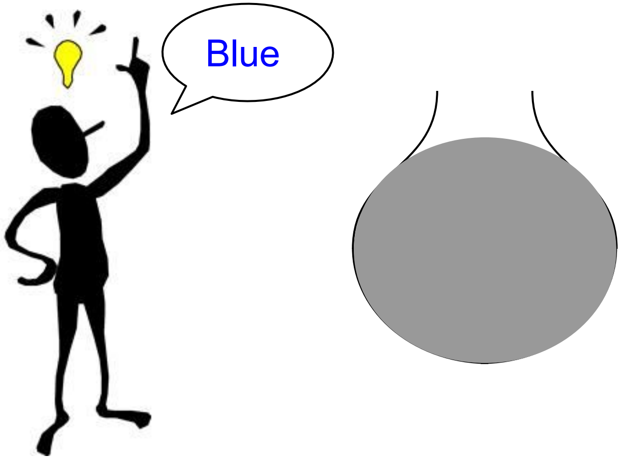

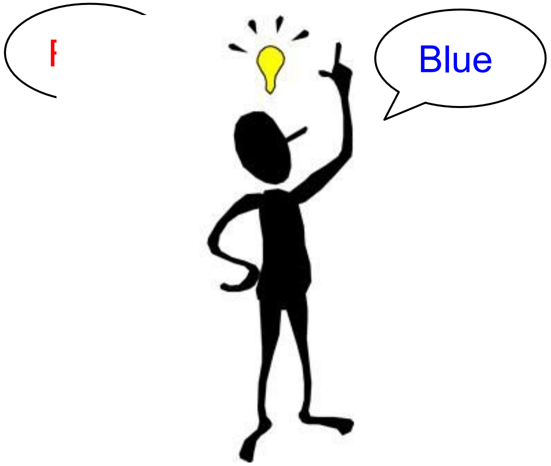

1. Student

- Announces what he/she sees
- Reveals private information

2. Student

- Announces what he/she sees
- Reveals private information

The Model

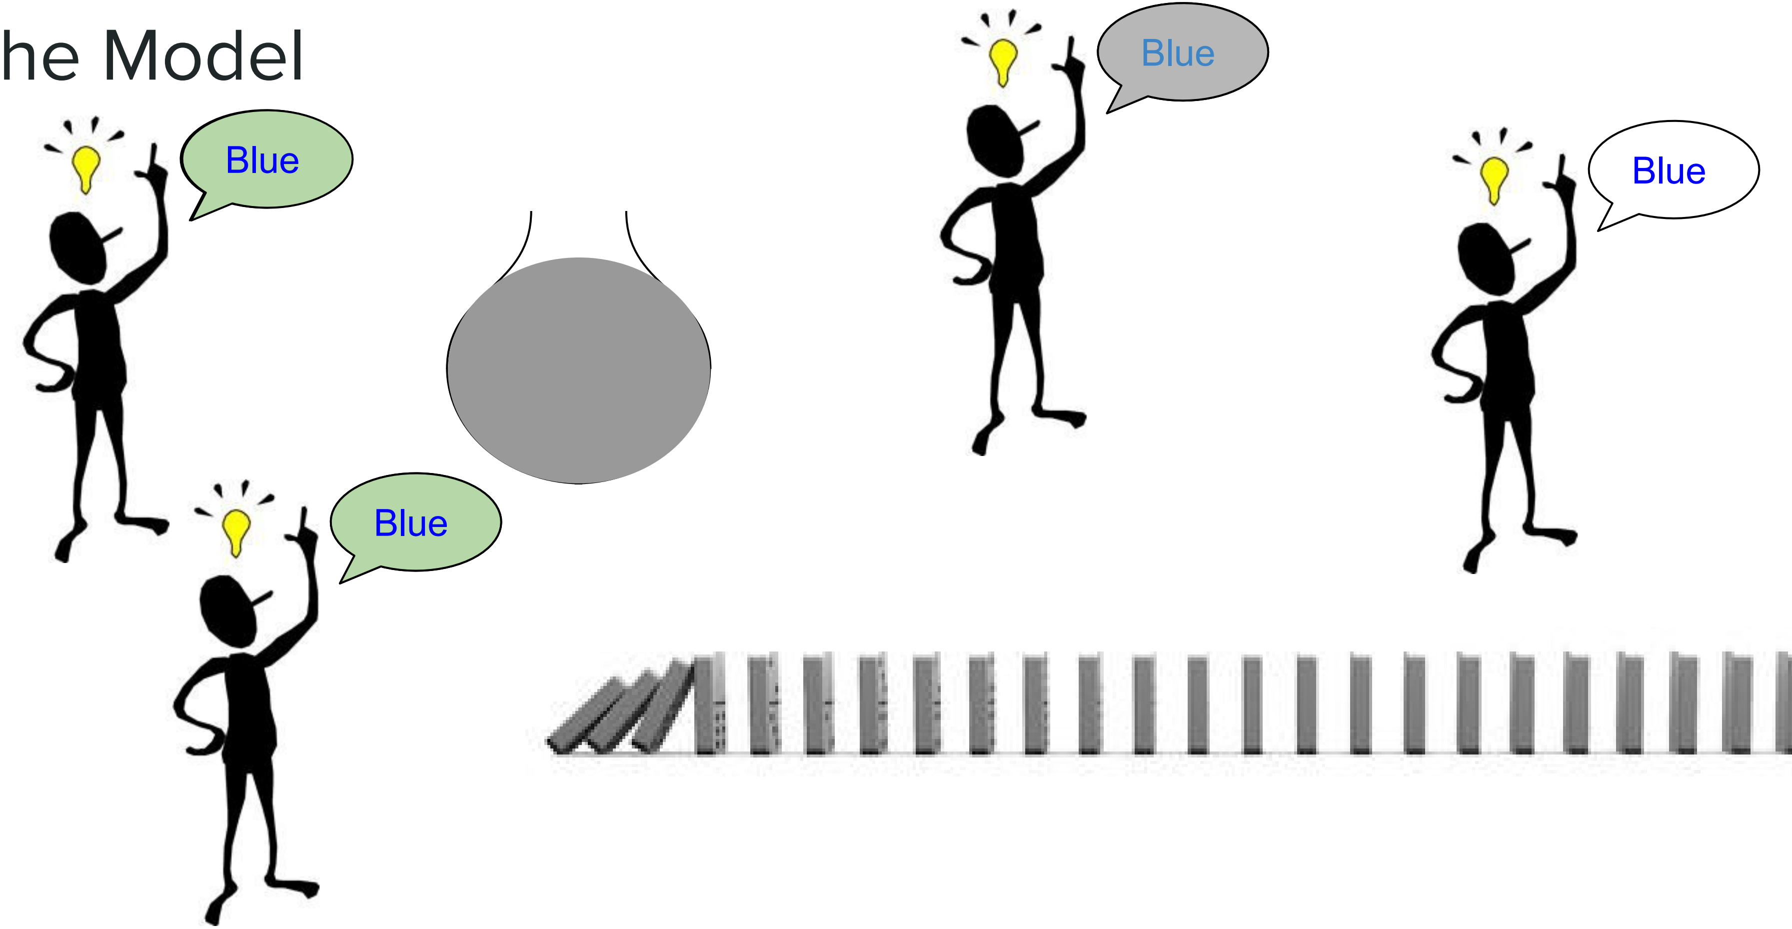

Information Cascades

- ...can be wrong
- ...are based on little information
- ...are fragile

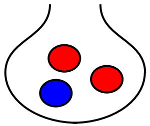

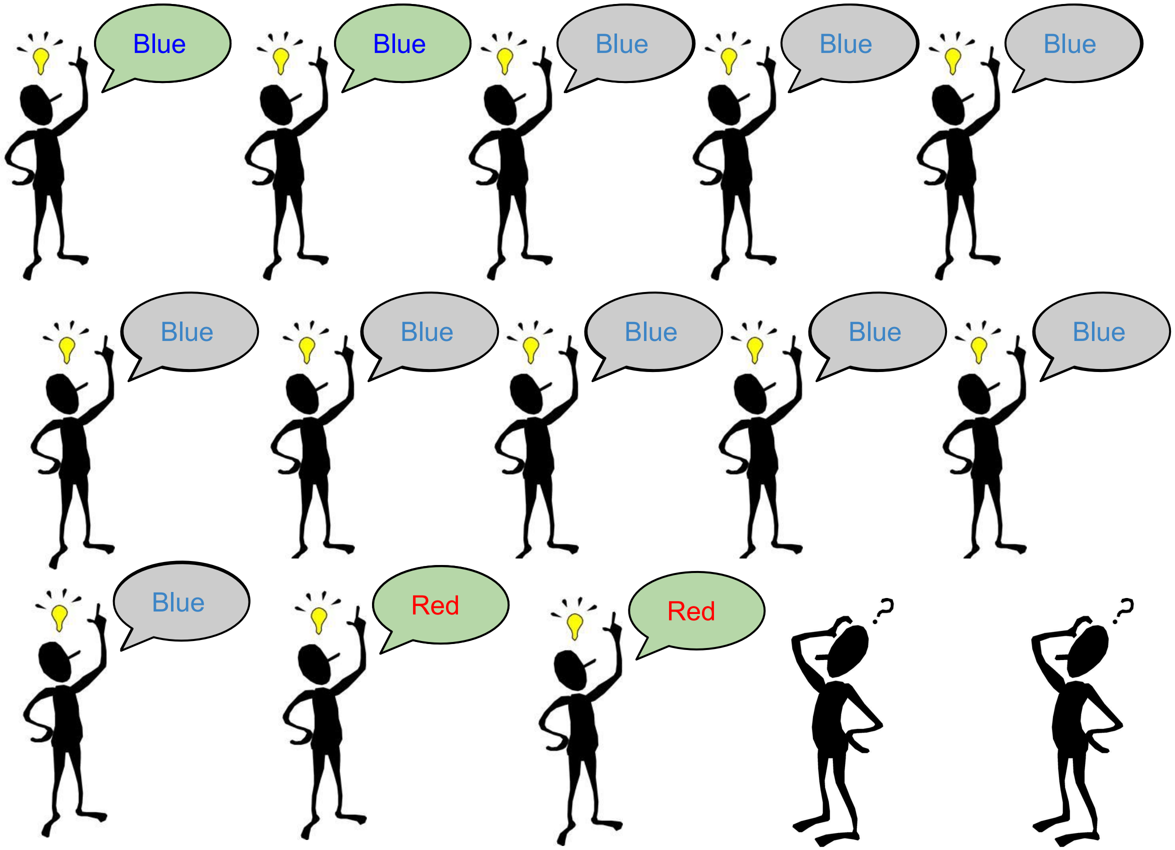

The General Model

- Each agent has to accept or reject a given option
  - Limited private information
  - Public announcements of others
  - Private information is correct with probability $q > \frac{1}{2}$
- The world is in one of two states
  - Accepting is good
  - Rejecting is good

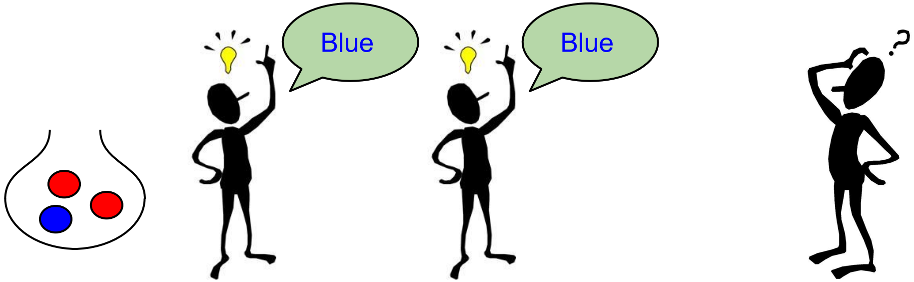

- Urne example
  - $q = \frac{2}{3}$
  - Accept “majority blue” or reject it (i.e. aim for “majority red”)

Related Work

- Sequential setting:
  - Kleinberg:[1] Majority algorithm on independent signals is optimal

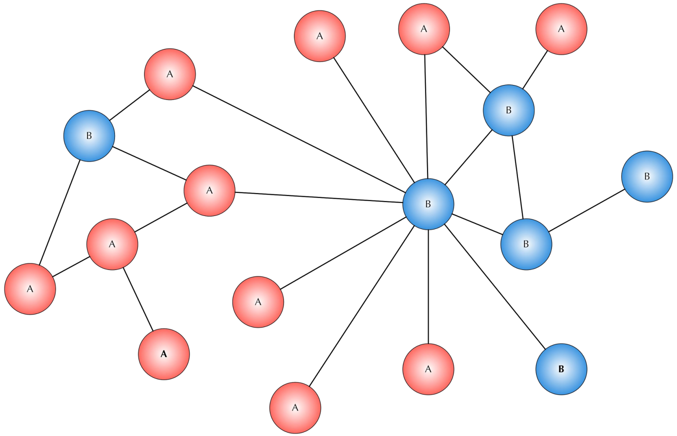

- Round based setting:
  - Golub and Jackson:[2] Convergence to truth if influence vanishes

1) *Networks, crowds, and markets: Reasoning about a highly connected world. Cambridge University Press, 2010.*
2) *Naïve Learning in Social Networks and the Wisdom of Crowds. American Economic Journal: Microeconomics 2010*

How good are we doing?

Let *n* be the number of students

Full information

- Chance of a wrong information cascade on 3. student: ⅓ * ⅓ = 1/9
- Chance of a wrong cascade in total:

$$\frac{1}{9} + \frac{1}{3}\sum_{k=1}^{n}\left(\frac{2}{3}\cdot\frac{1}{3}\right)^{n} \approx \frac{1}{5} \qquad \text{for } n \to \infty$$

- Expected number of wrong guesses $\approx \frac{1}{5}$ n

How good are we doing?

Let *n* be the number of students

No information

- Everybody follows private information
- Expected number of wrong guesses = ⅓ * n

In both cases:

➔ Ω(n) expected mistakes!

Can we do better?

Information Sharing Based on Graphs

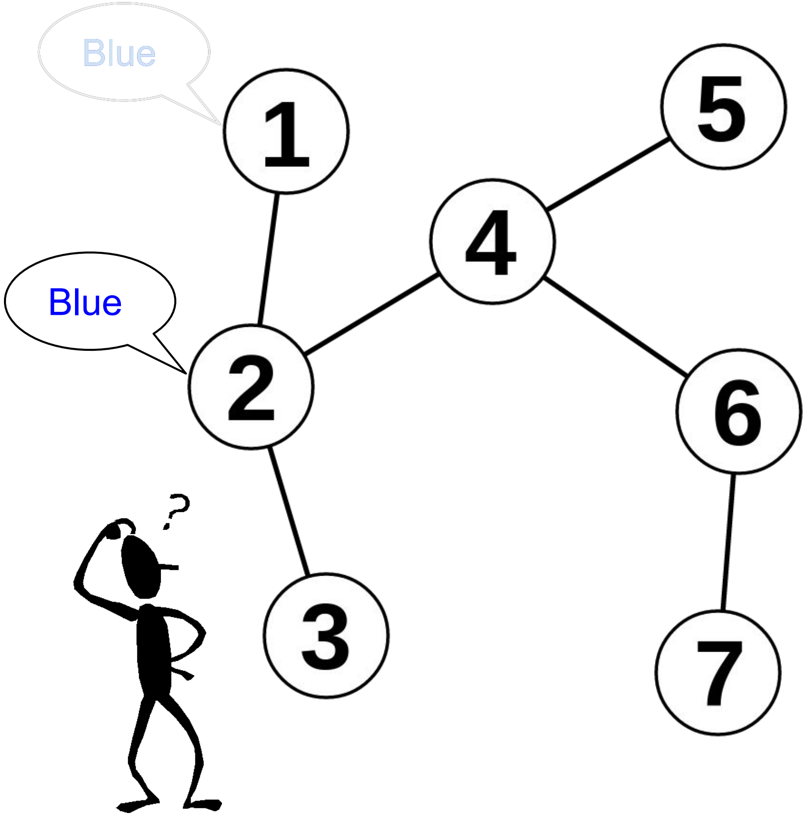

How much connectivity is needed?

Random Graph *G(n,p)*

- *n* is the number of nodes
- Each pair of nodes is connected with probability *p*

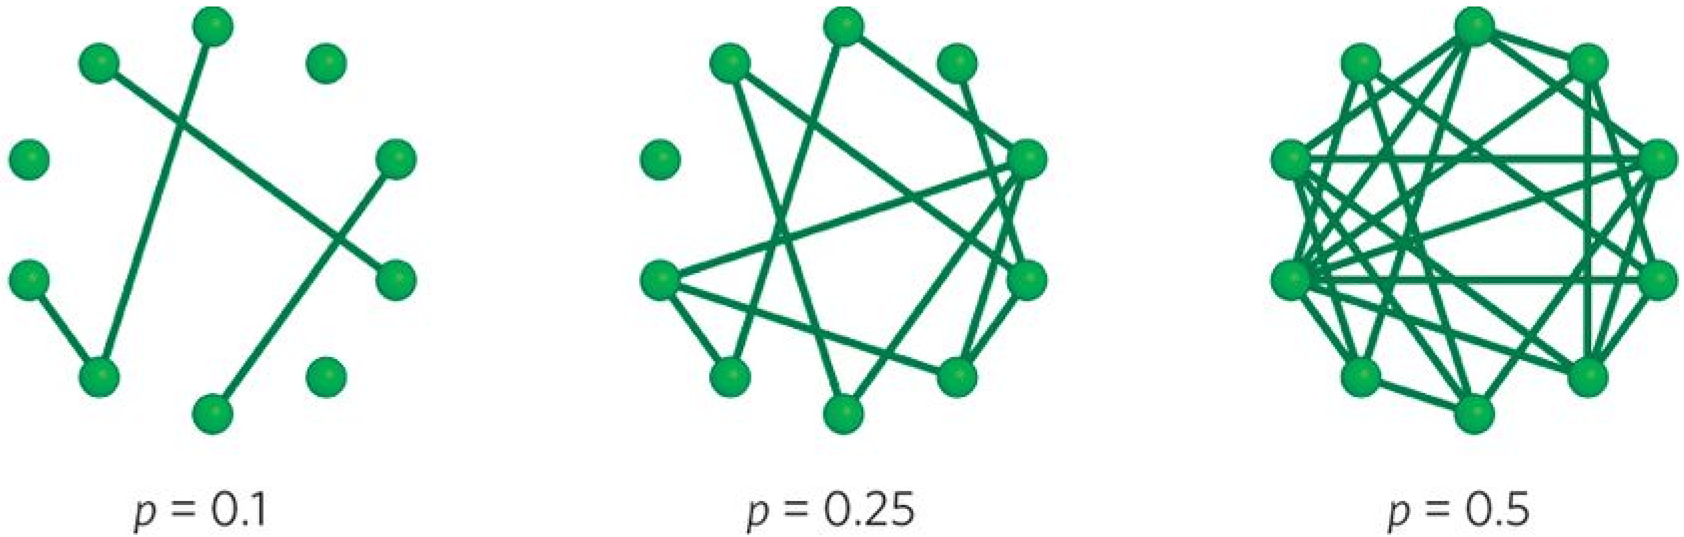

Random Graph *G(n,p)*

- $p = 1$ ➔ fully connected graph
- All nodes get information of all other nodes
- Chance of a wrong cascade $\approx \frac{1}{5}$
- Expected number of wrong guesses $\approx \frac{1}{5}$ n

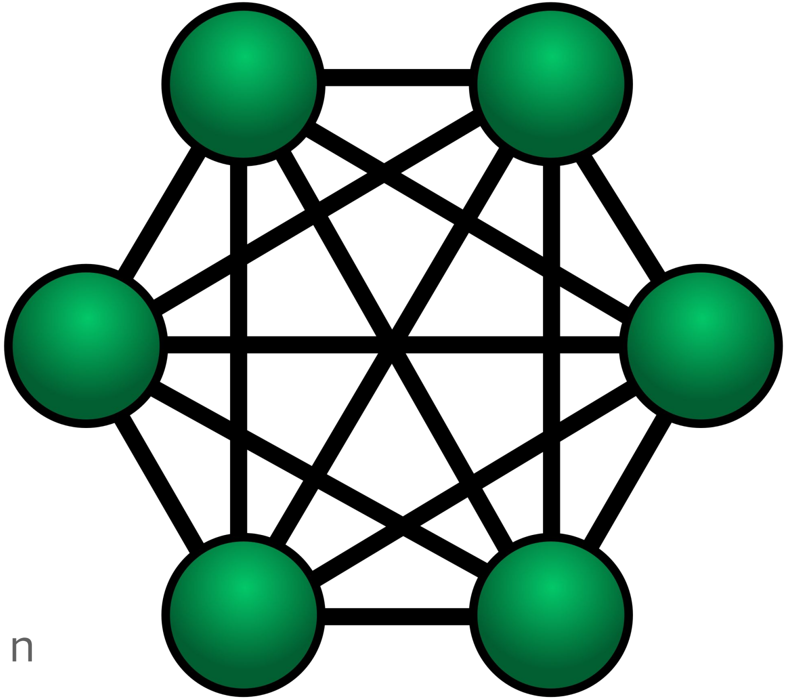

Random Graph *G(n,p)*

- $p = 0$ ➔ empty graph
- All nodes get only their private information
- Chance of being wrong = ⅓
- Expected number of wrong guesses = ⅓ n

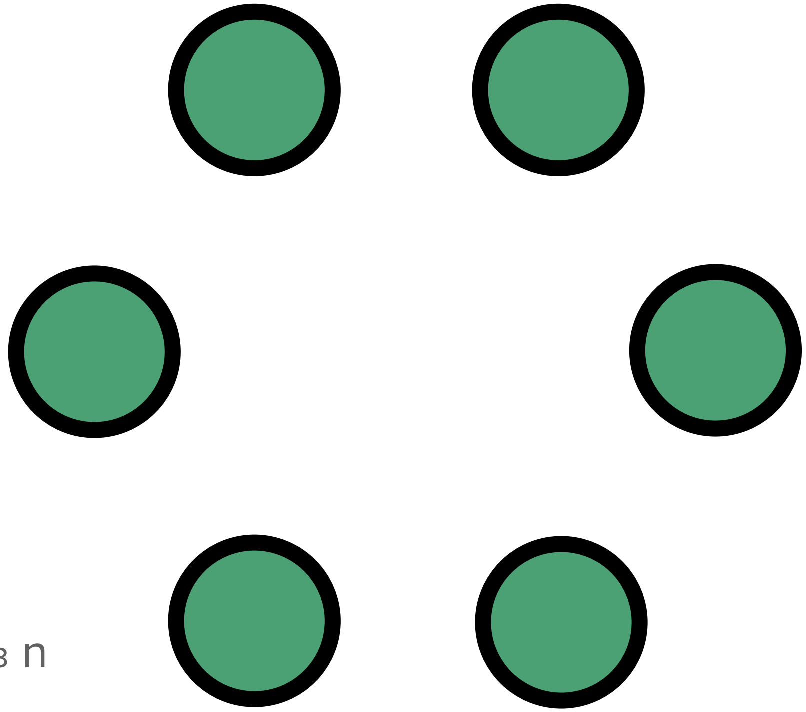

Random Graph *G(n,p)*

- *p* = 1 ➔ fully connected graph ➔ full information ➔ Ω(n) mistakes
- *p* = 0 ➔ empty graph ➔ no information sharing ➔ Ω(n) mistakes
- What happens in between?

Majority algorithm on random graphs with optimal connection probability *p* results in Θ(log n) mistakes in expectation

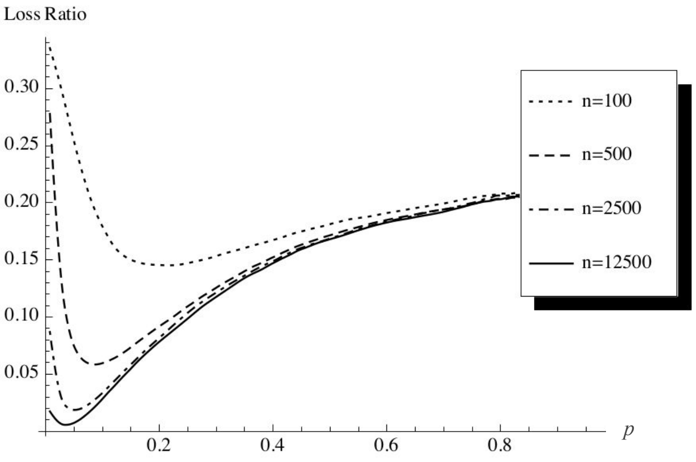

Figure 1 Performance of random graphs for different p and n.

How come?

- Basic idea: less neighbors ➔ more likely to reveal private information
  - ➔ Increased chance of correct information cascade
- If $p \in \Theta(1/\log n)$, w.h.p. each of the first $\Theta(\log n)$ nodes has at most one neighbor
- Nodes with at most one neighbor reveal private information

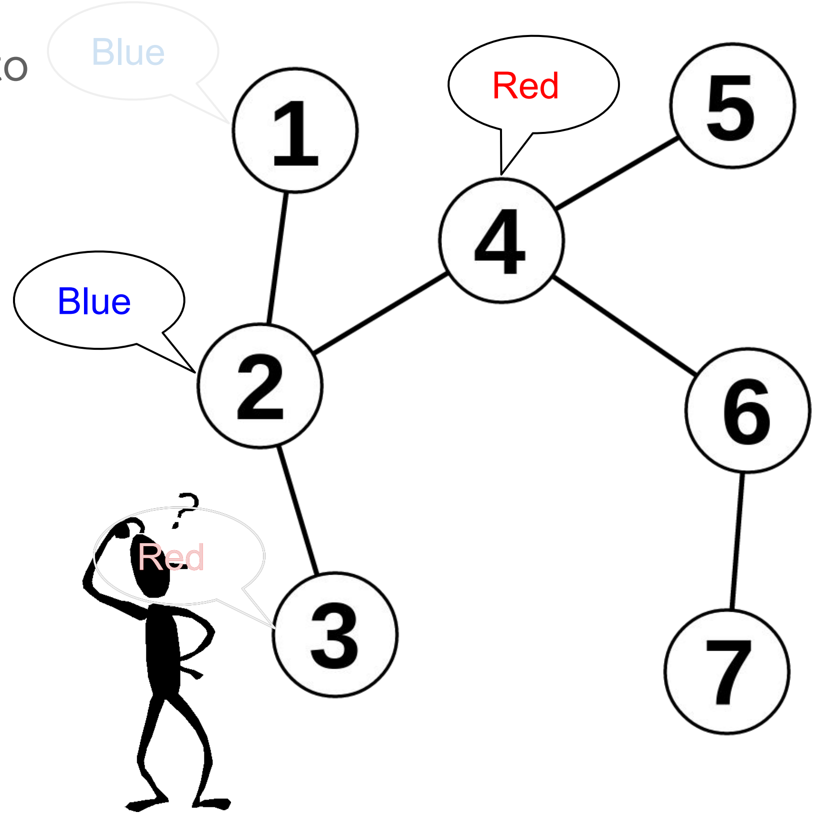

How come?

- If $p \in \Theta(1/\log n)$, w.h.p. each of the first $\Theta(\log n)$ nodes has at most one neighbor
- Nodes with at most one neighbor reveal private information

➔ The first $\Theta(\log n / p)$ nodes mostly correct

➔ The remaining n - $\Theta(\log n / p)$ nodes have expected number of mistakes in $O(1)$

$\Theta(\log n)$ expected mistakes for correct p

➔ Fraction of mistakes goes to 0 as $n \to \infty$

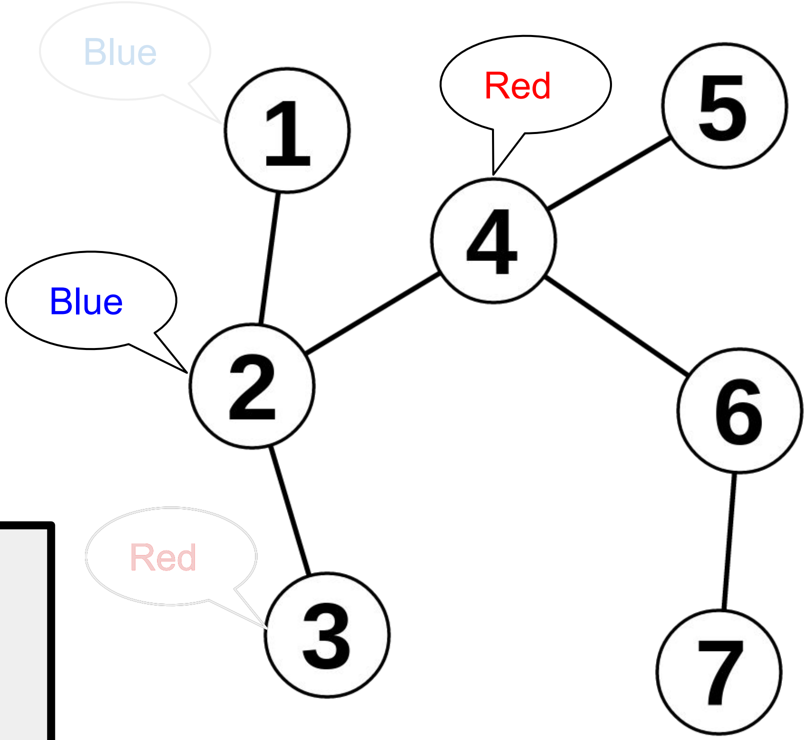

Can we do better?

- Each node *i* should try to minimize the expected number of wrong guesses
- Optimal algorithm:
  - Node n tries to minimize its own failure probability
  - Node n-1 tries to minimize the failure probability of itself and node n
  - …
  - Node 1 tries to minimize the overall failure probability
  - Given how node 1 decides, node 2 can determine how it will decide in each case
  - …
  - Given how nodes 1 to n-1 will decide in each case, node n can determine how it will decide

More concrete

Given full information

- Up to a given point, nodes reveal their private information
- From that point on, all nodes will chose the majority of previous outputs

➔ At a given switching point *m*, nodes switch from revealing private information to choosing the majority of previous outputs

➔ m > Ω(log n) or the expected number of mistakes is > sqrt(n) / 2

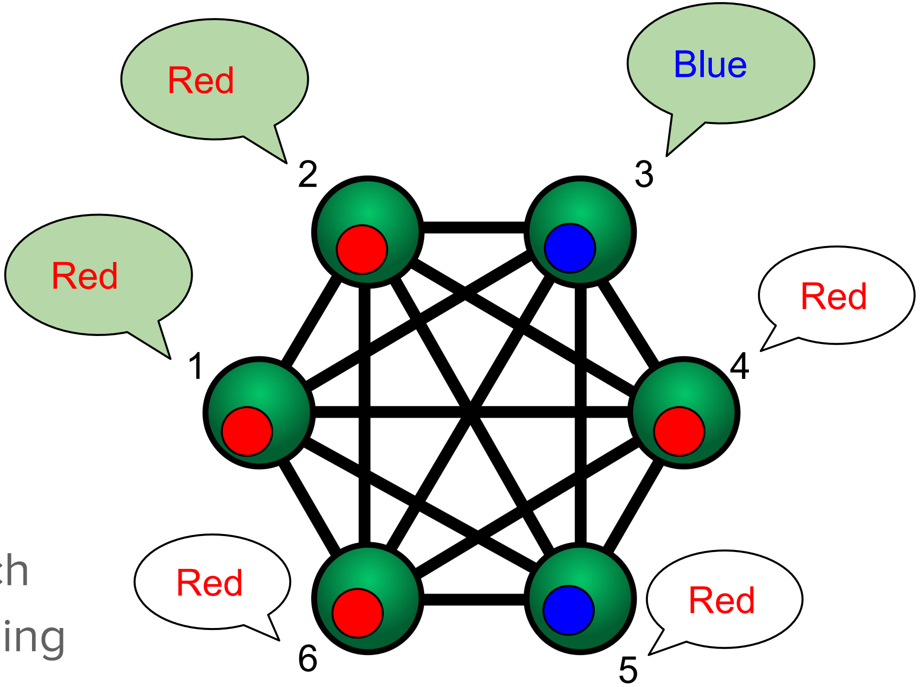

What we have seen so far...

- Random graphs achieve Θ(log n) expected mistakes
- Θ(log n) expected mistakes is asymptotically optimal
- Can we achieve it without randomness?

Layer Graphs

Idea:
Choose the number and sizes of the layers such that the expected number of mistakes is minimal, i.e. in $\Theta(\log n)$

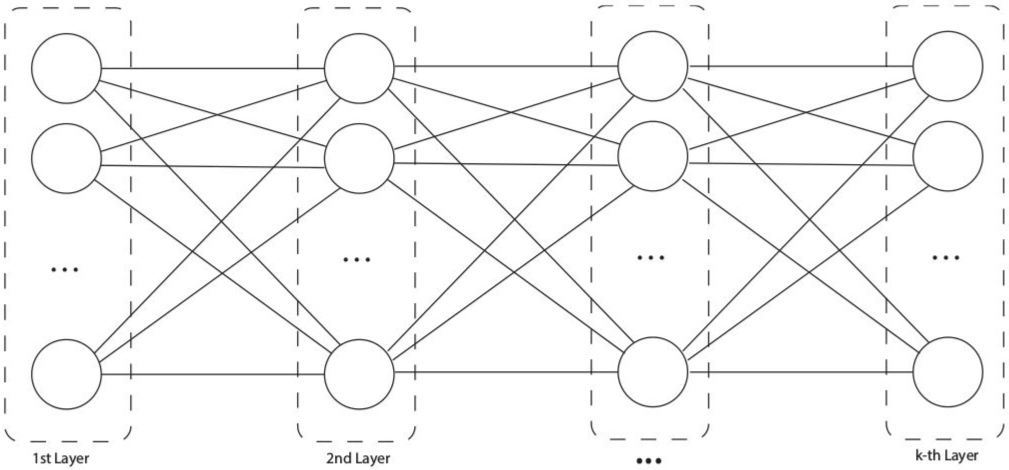

- Let $s = 1/(4q(1 - q))$, $q>½$ is the probability of a private information being correct
- The optimal layer topology has $k = n / \log_s(n) + o(n / \log_s(n))$ many layers
- The first layer has $\log_s(n)$ many nodes
- All following layers decrease in size as a staircase

Conclusion

- Too much information is (asymptotically) as bad as no shared information
- Just enough information can lead to less mistakes
- This can be achieved through:
  - Connectivity regulation
  - Algorithmic regulation
  - Structural regulation

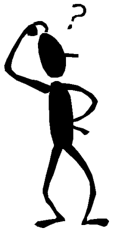

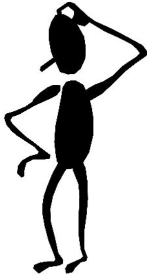

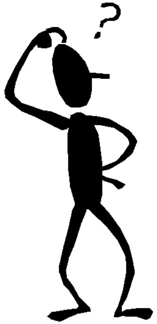

Questions?