How Robust is the Wisdom of the Crowd?

Noga Alon, Michal Feldman, Omer Lev & Moshe Tennenholtz

IJCAI 2015

Buenos Aires

Is this new movie any good?

the world

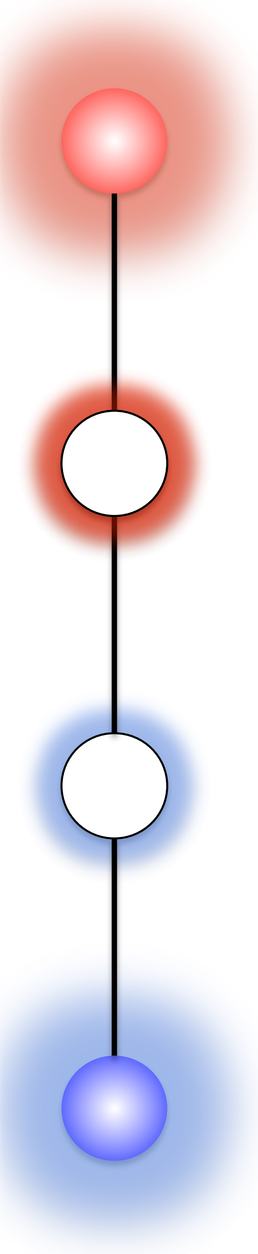

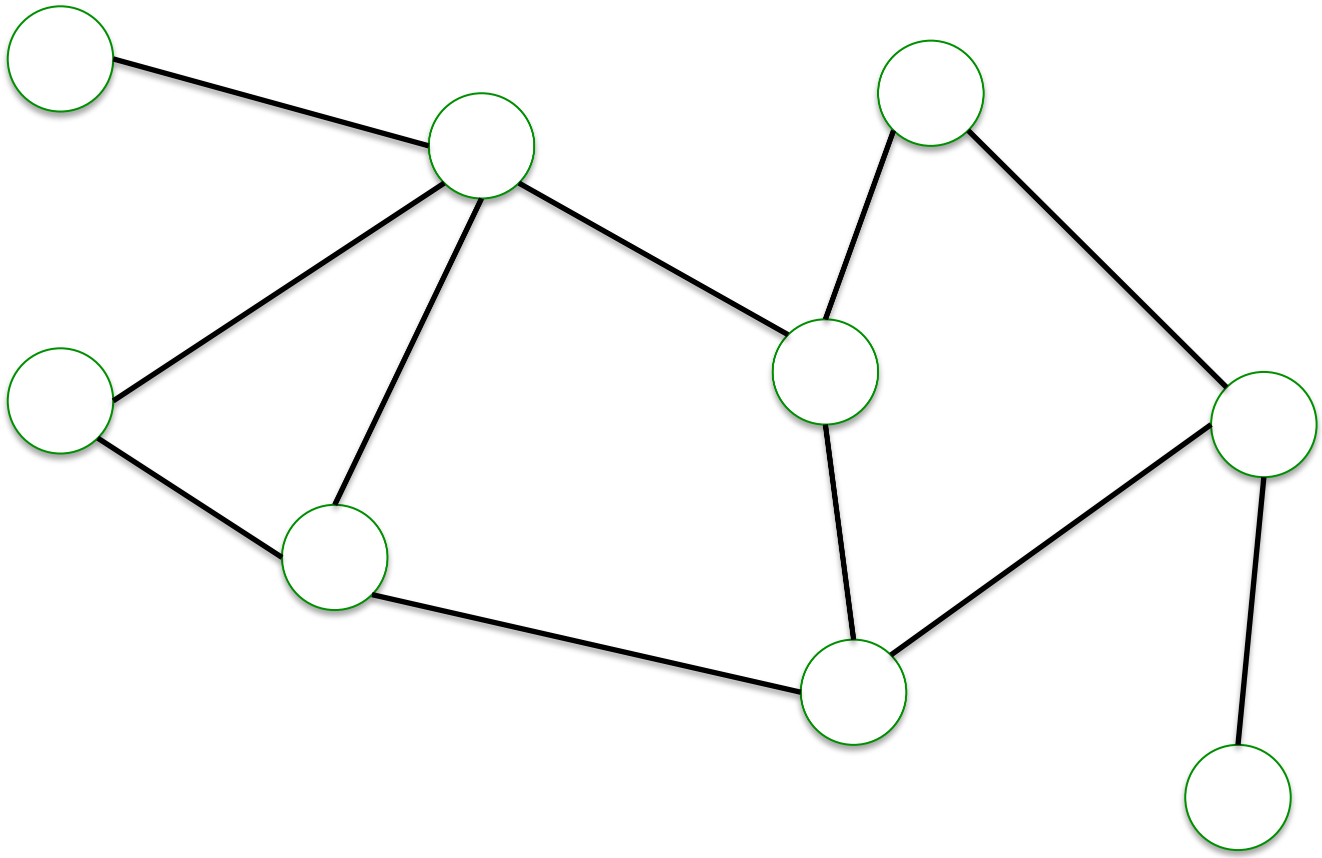

Is this new movie any good?

film critics

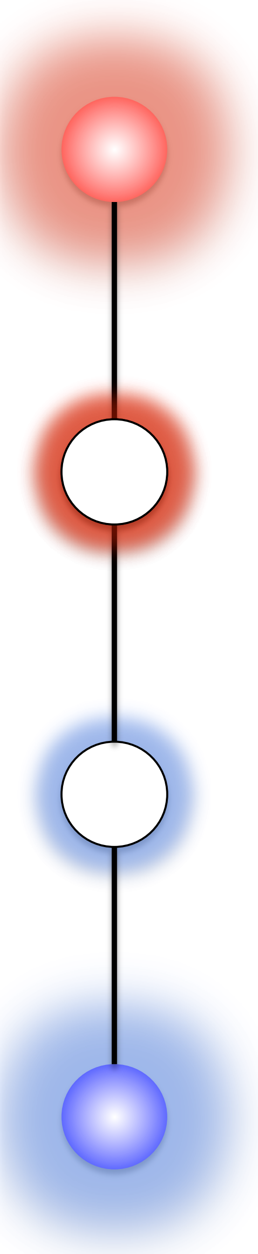

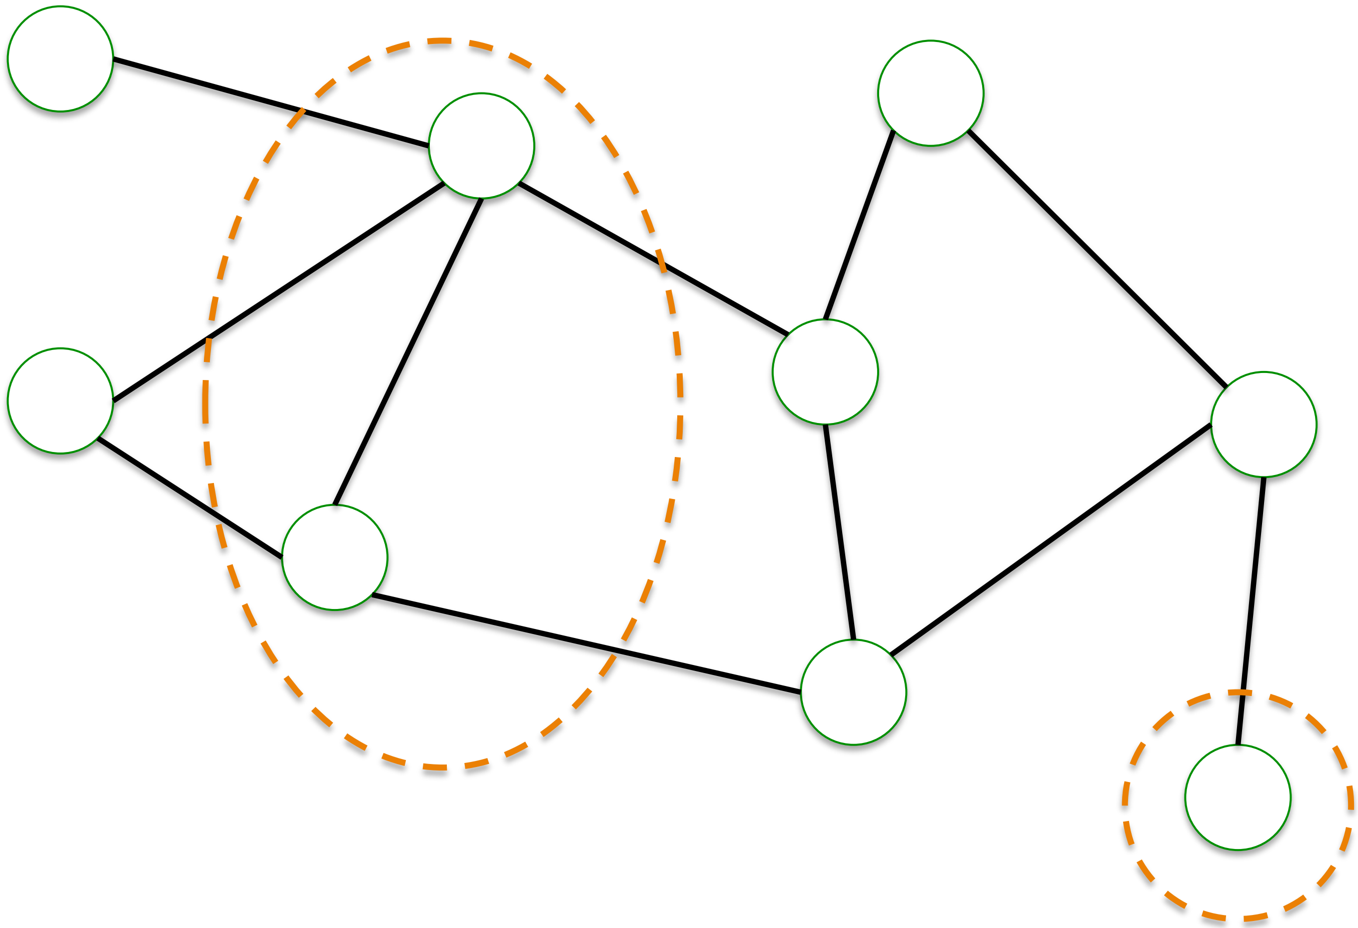

Is this new movie any good?

film critics

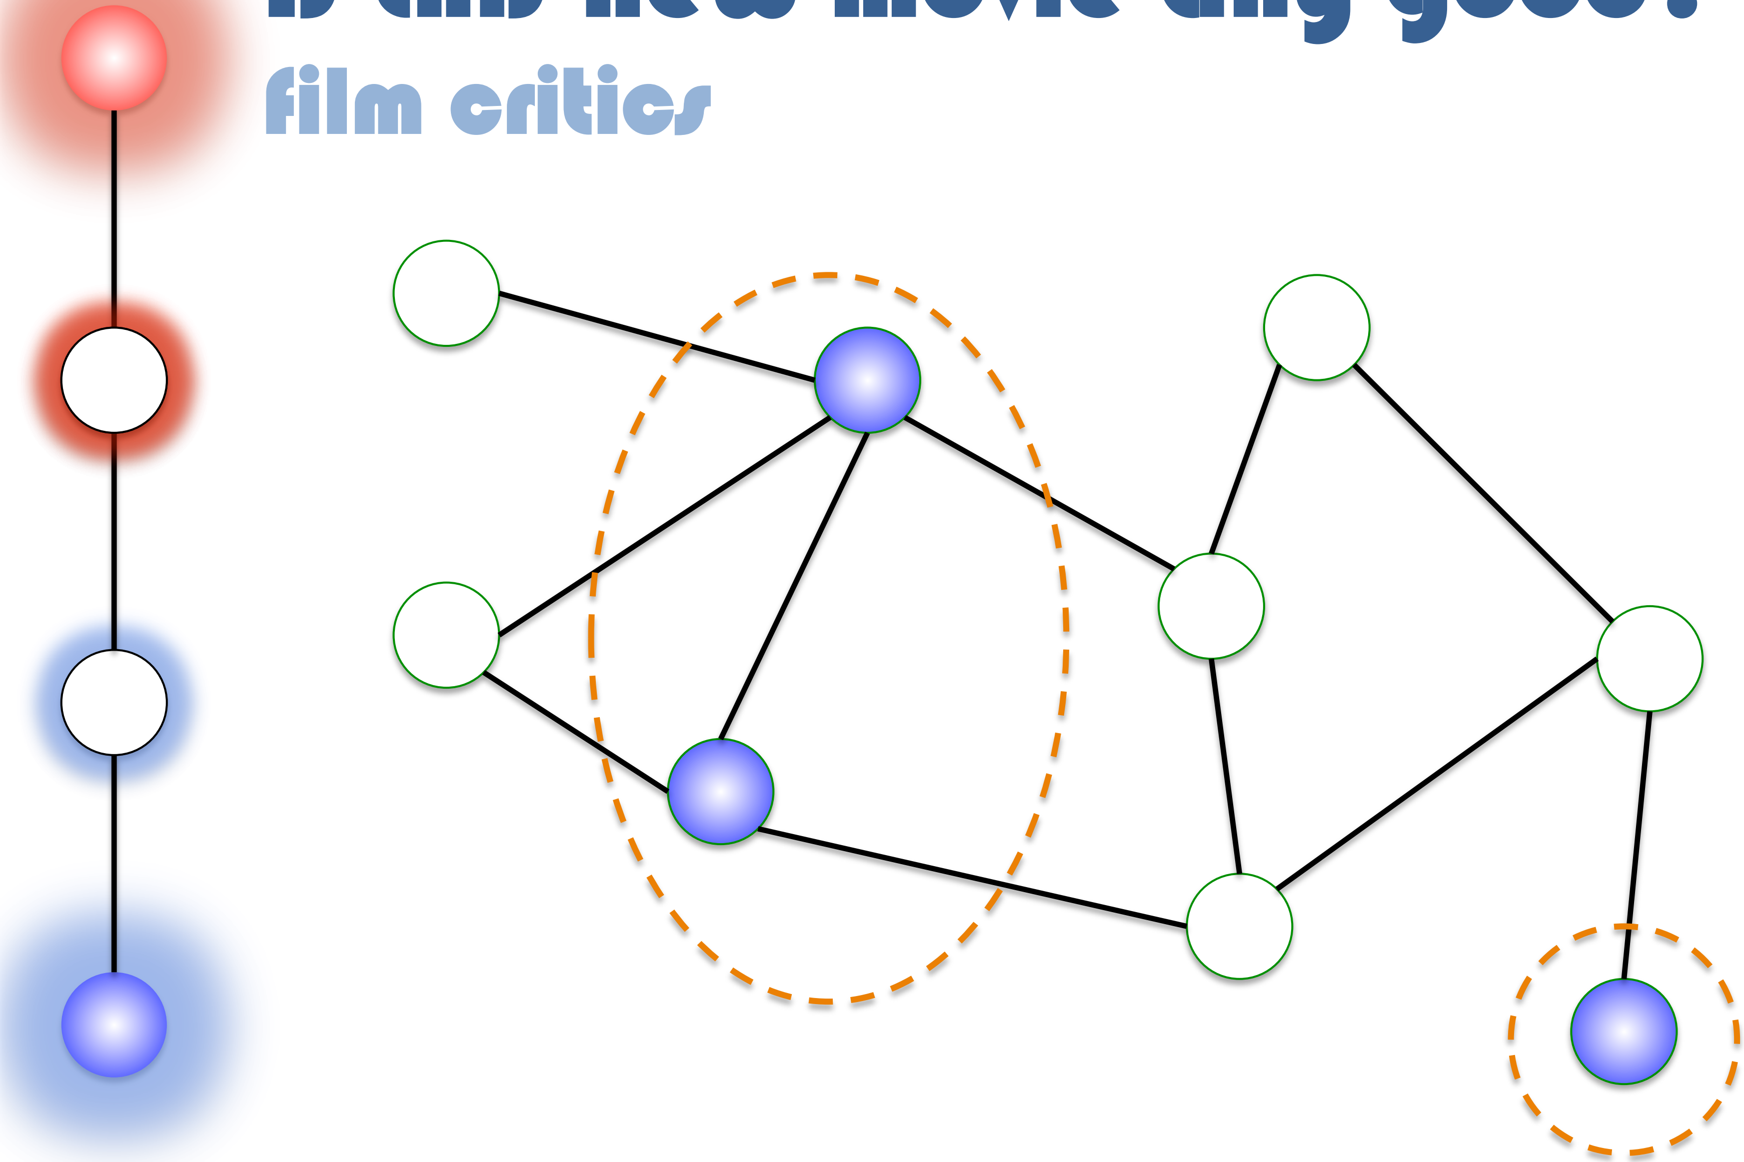

Is this new movie any good?

film critics

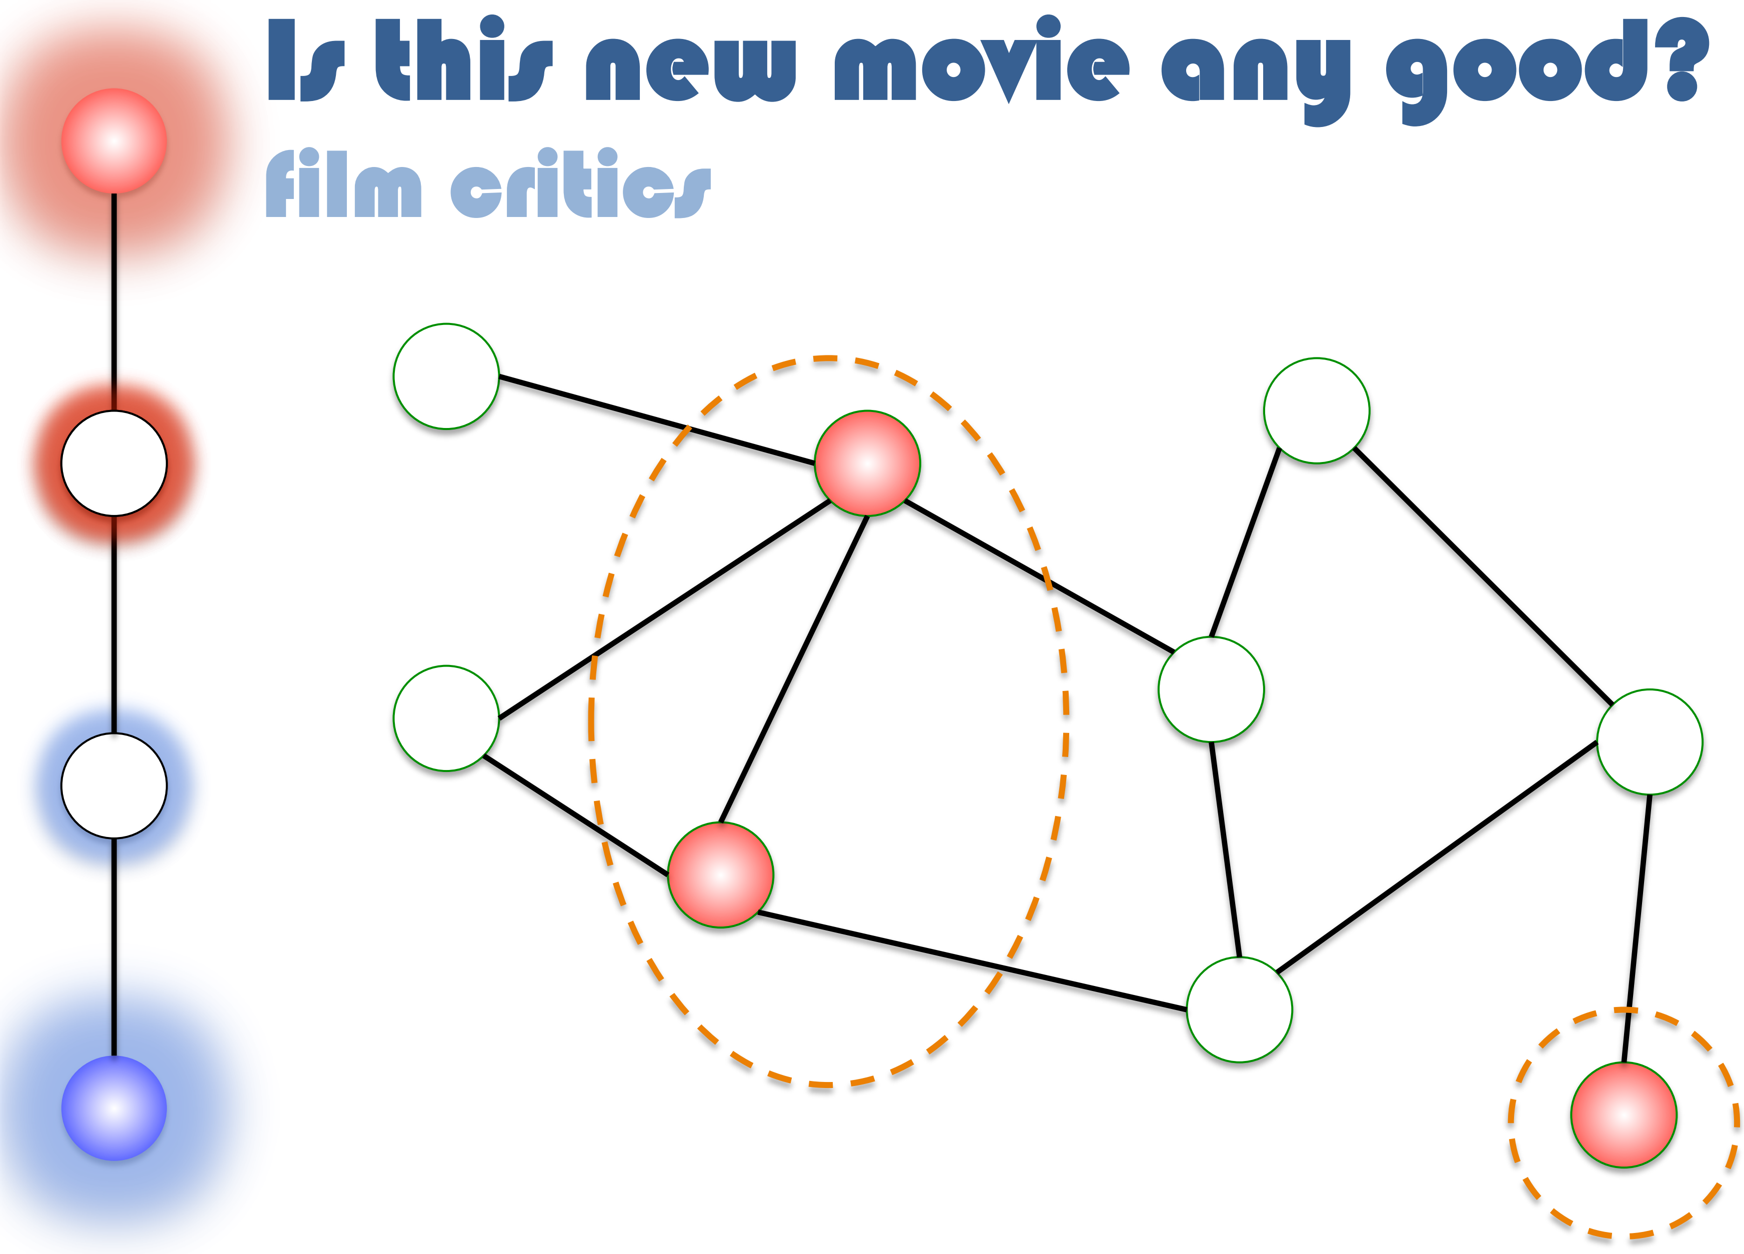

Is this new movie any good?

film critics

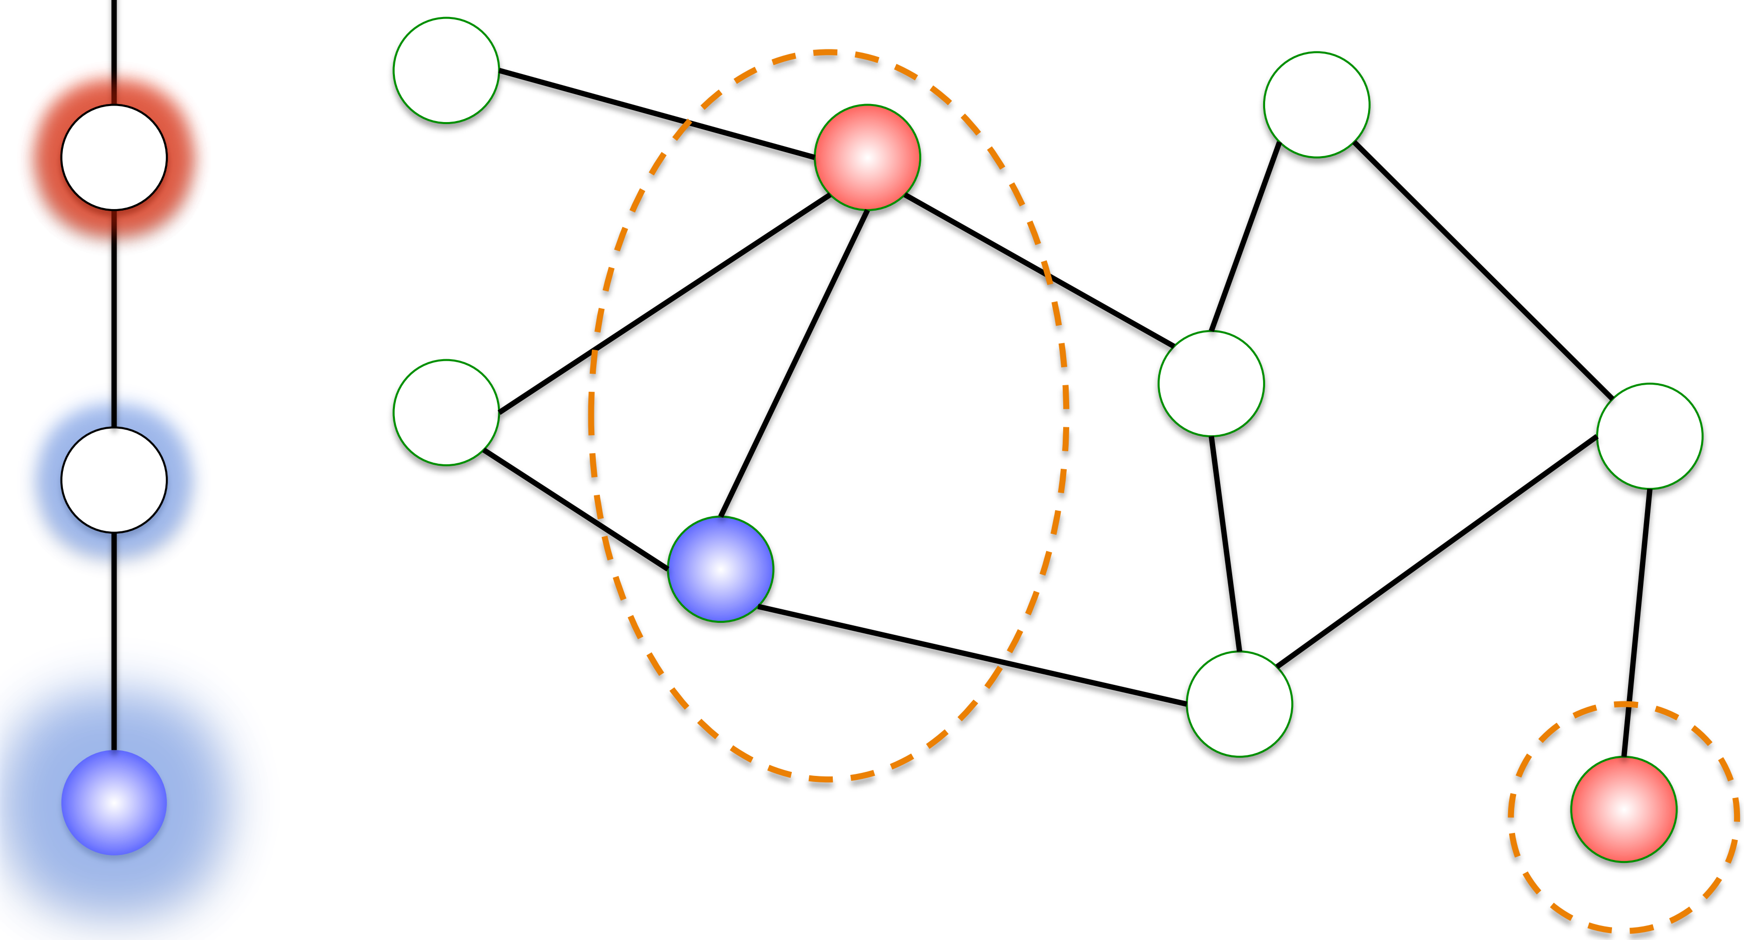

Is this new movie any good?

film critics' influence

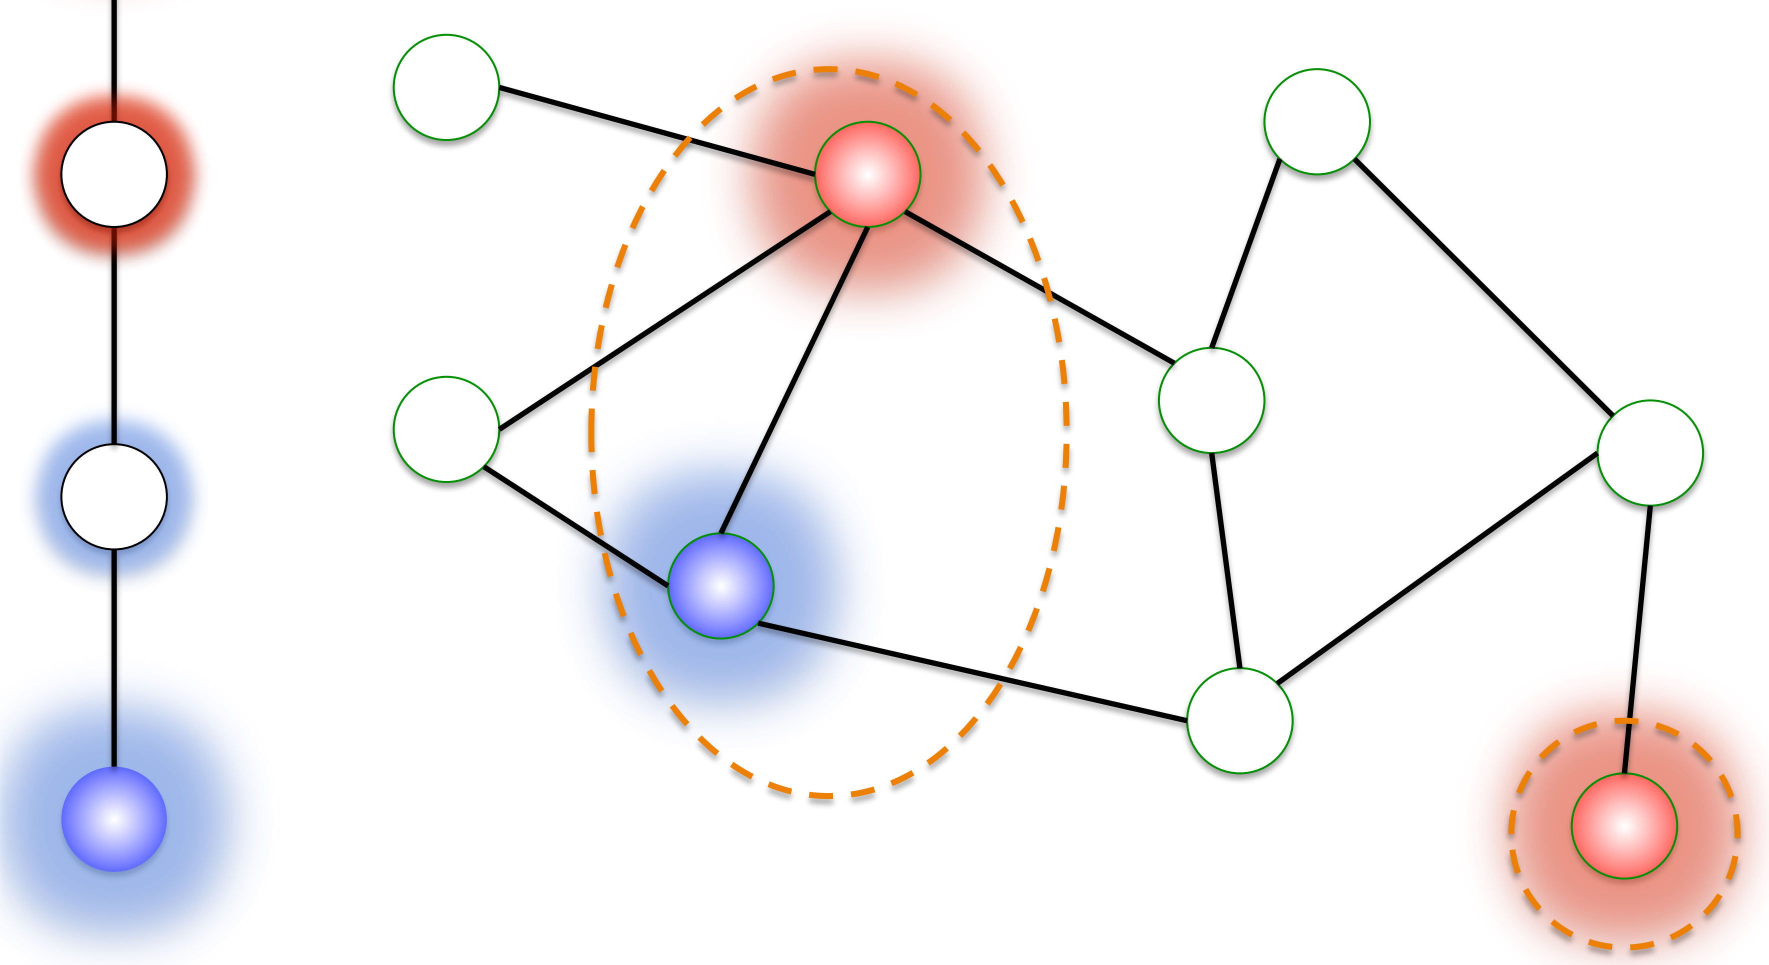

Is this new movie any good?

experts' influence

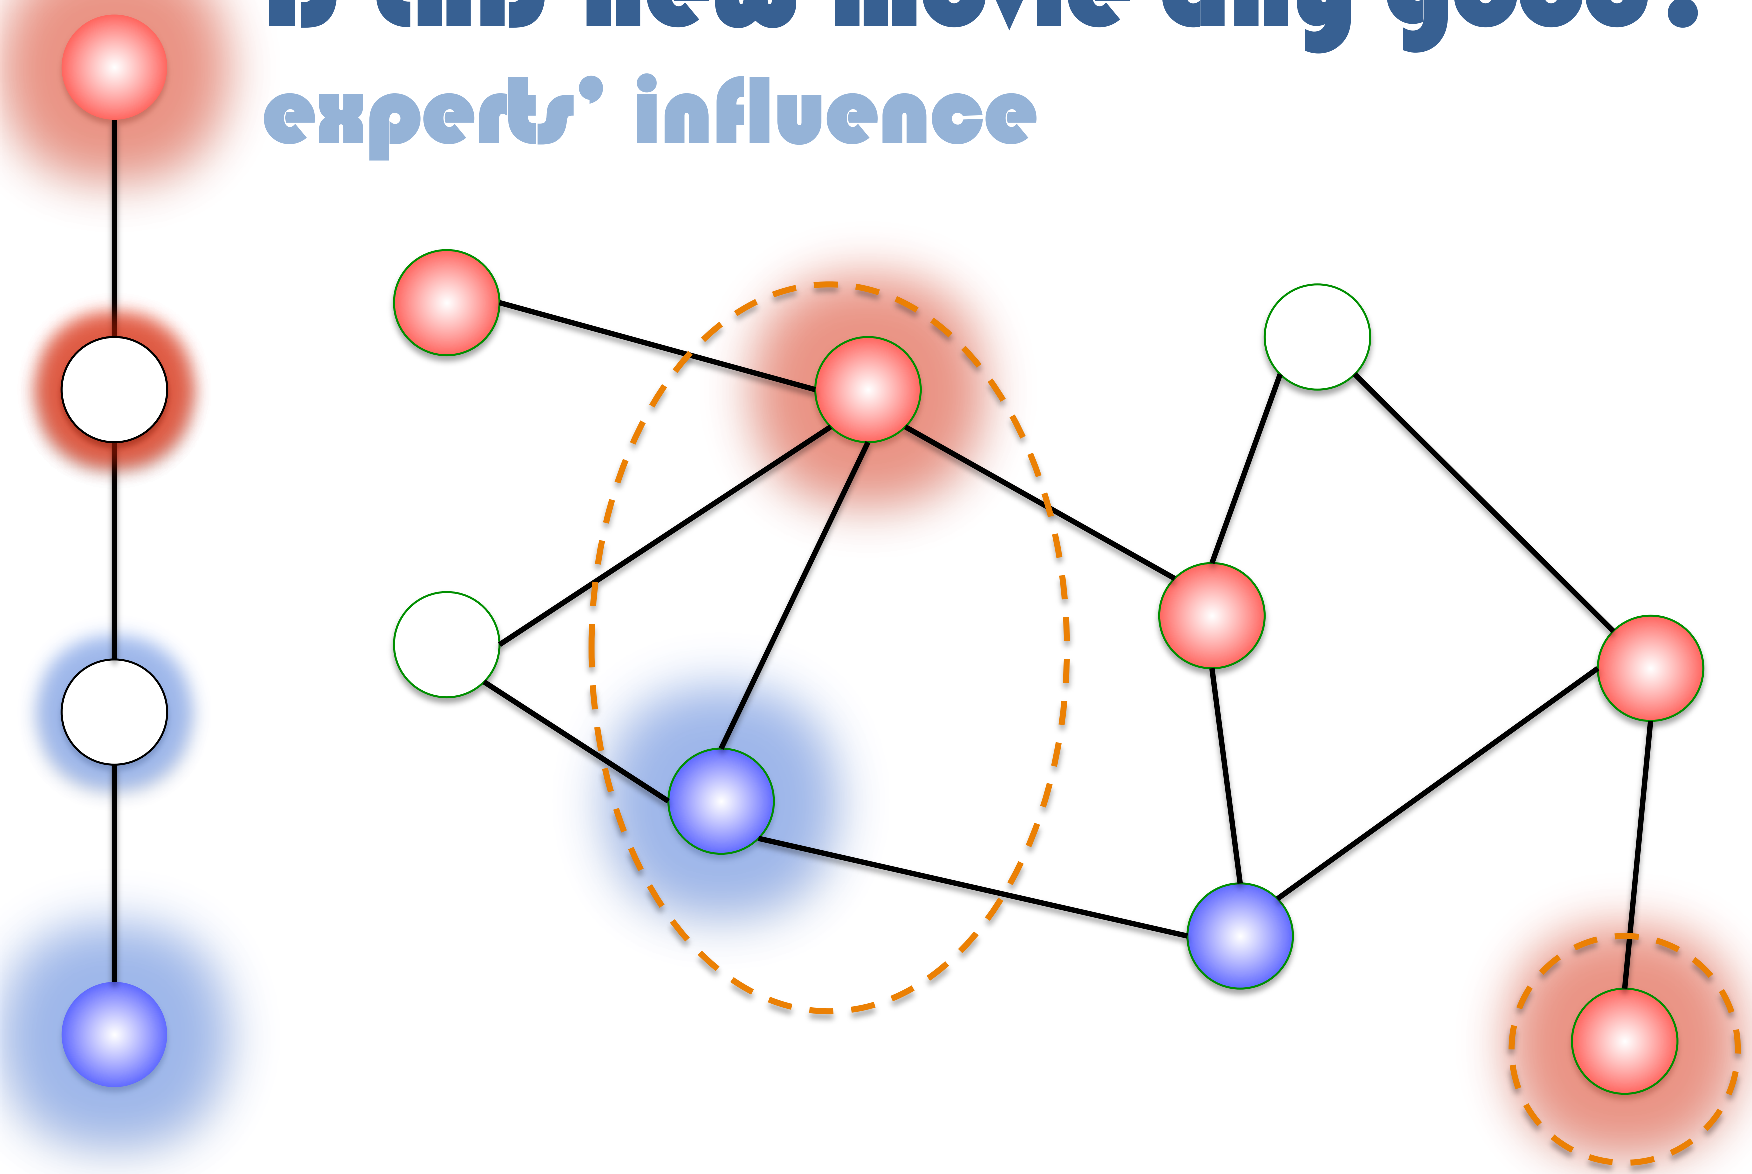

Is this new movie any good?

aggregate

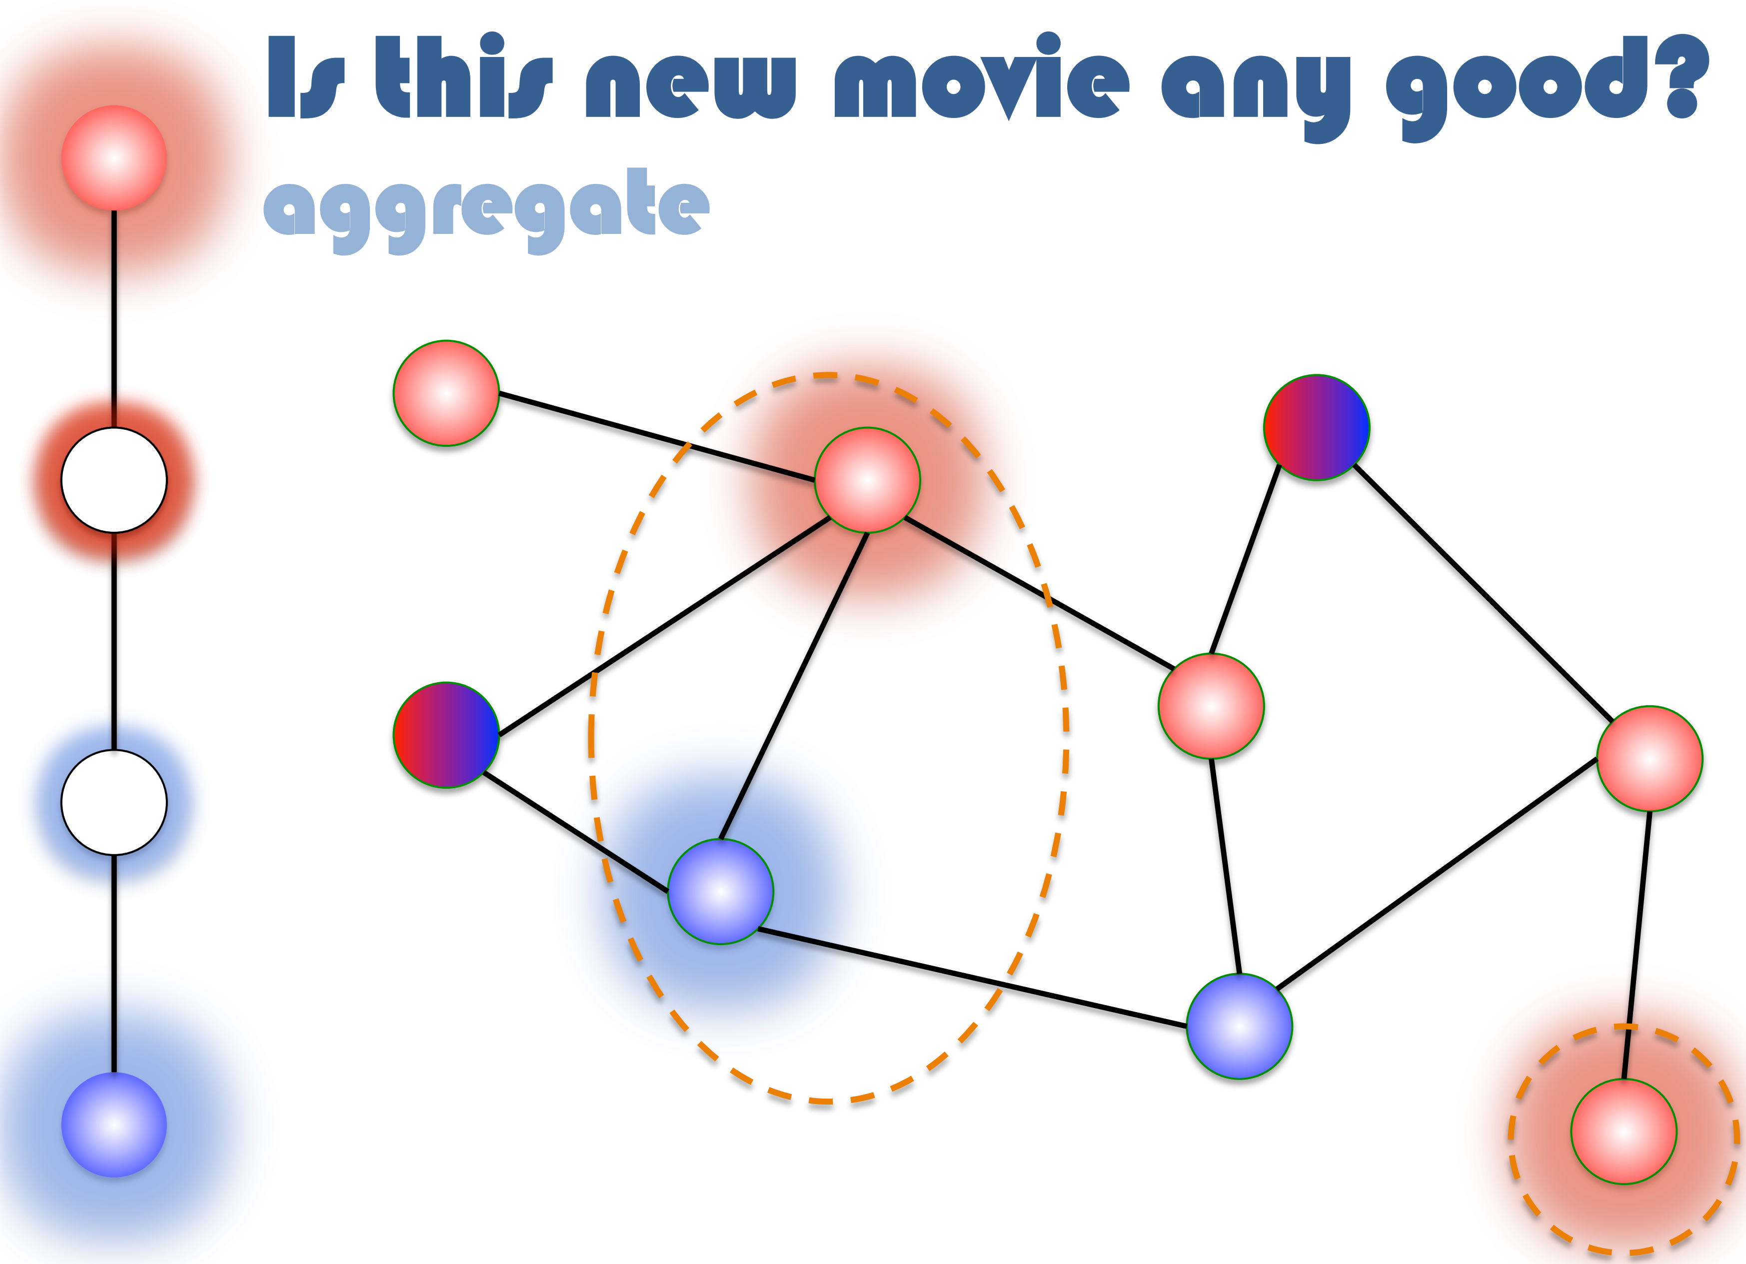

Is this new movie any good?

aggregate

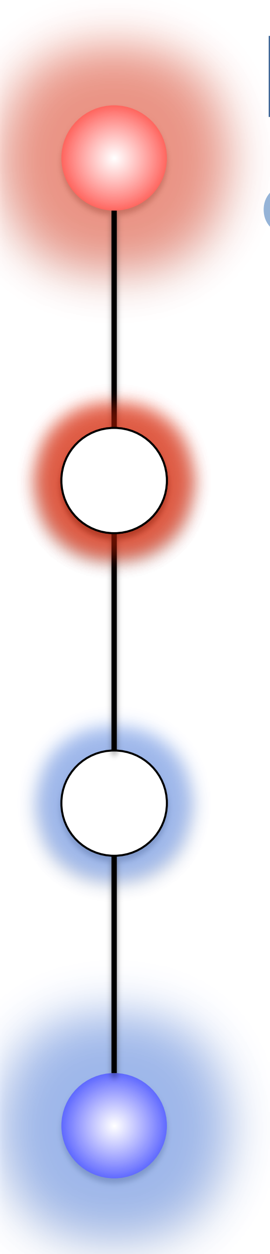

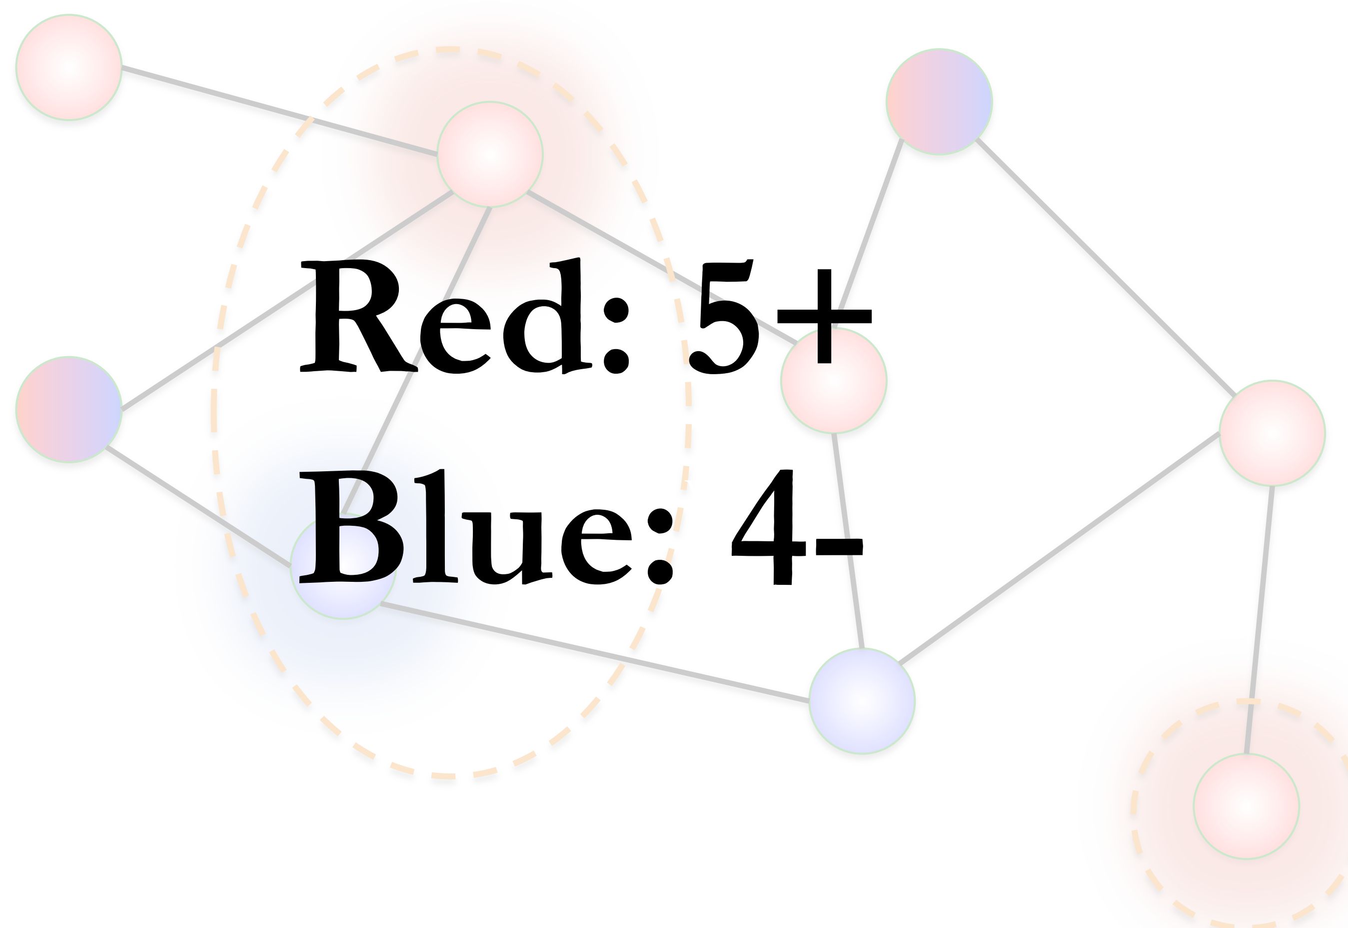

Our model social graph

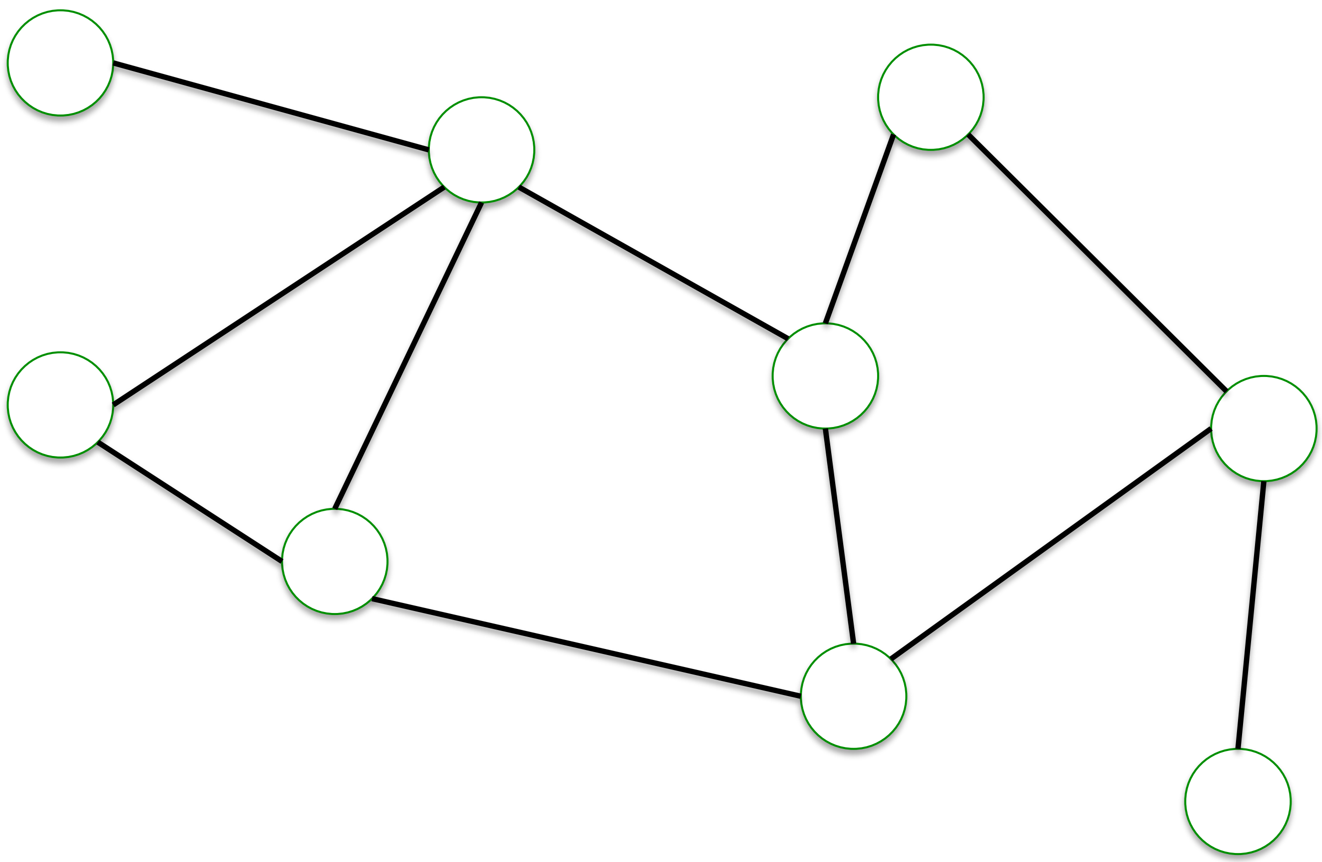

Our model experts

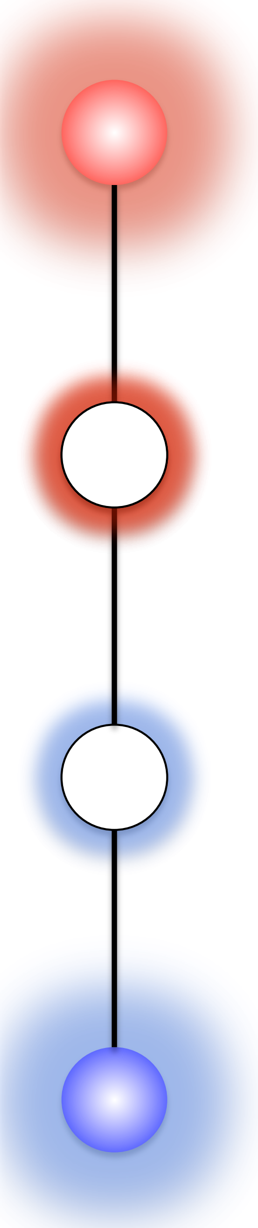

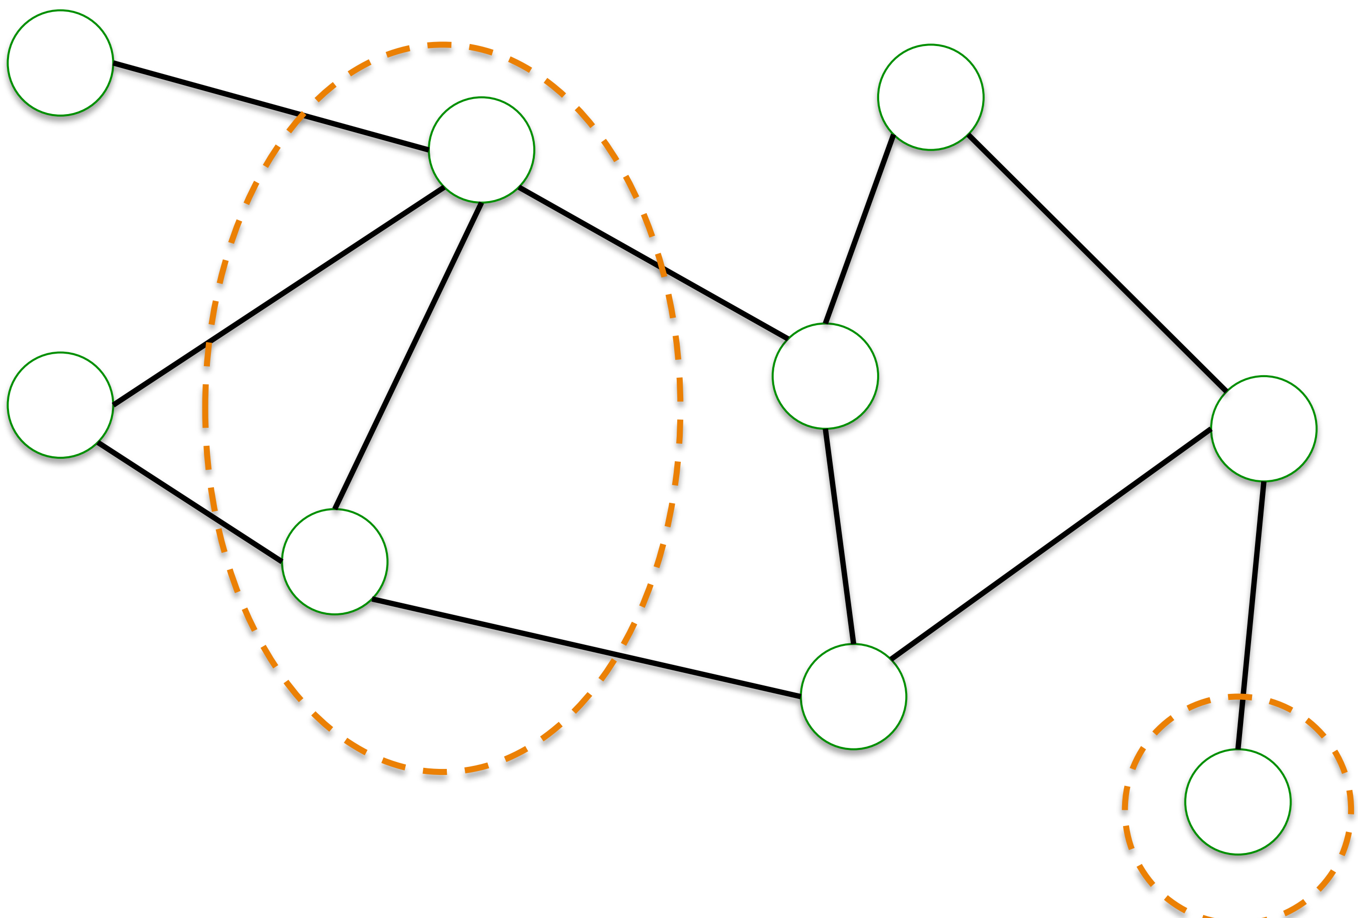

Our model experts' opinions

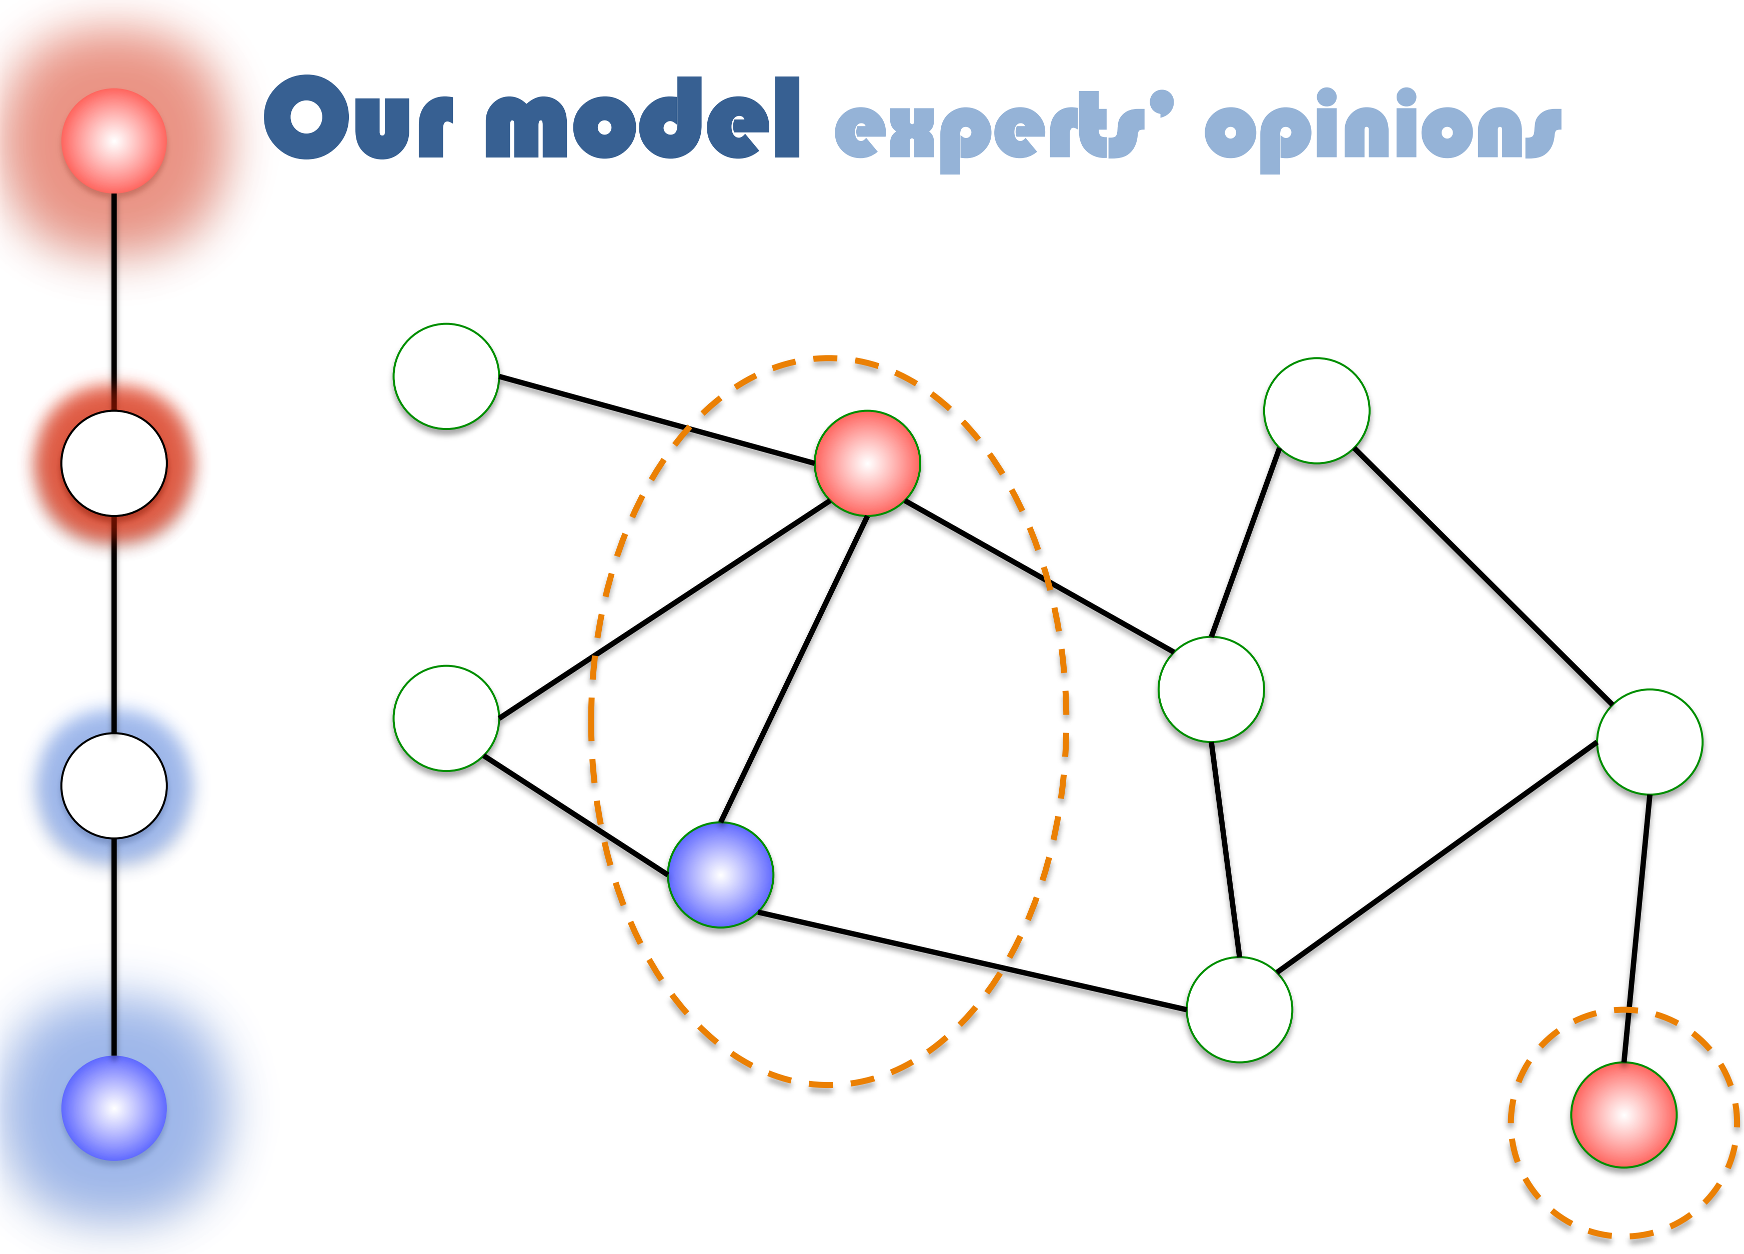

Our model experts' opinions

We assume underlying truth is **Red**

Regular people can mistake truth for **Blue** with probability ½

But experts will mistake truth for **Blue** with probability ½-$\boldsymbol{\delta}$

Our model outcome

We don’t know who the experts are: we only see the aggregate number of **Blues** and **Reds**

Adversary types

Weak Adversary

A *weak adversary* is one that can choose the set of experts, but has no other power on the experts' ultimate choice.

Weak Adversary

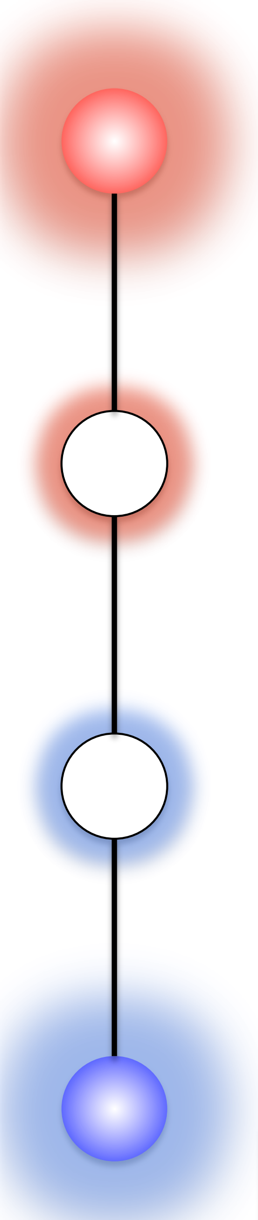

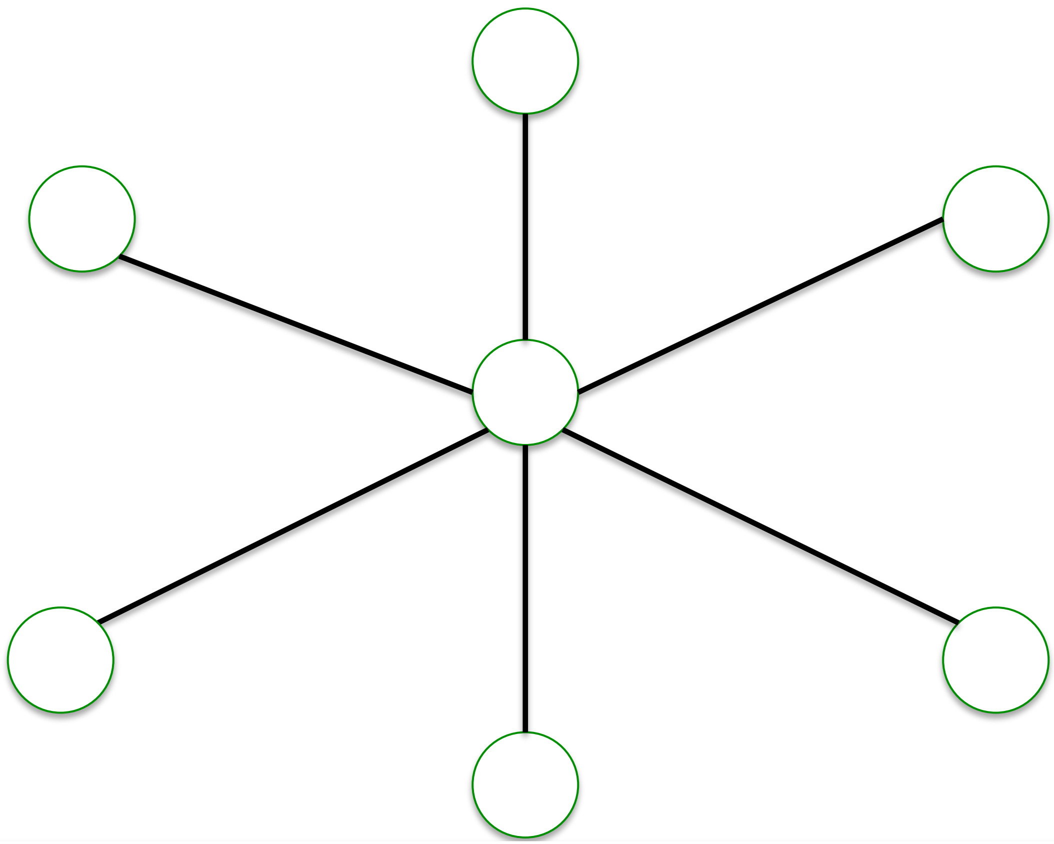

Choosing the middle vertex as expert means **Blue** wins with probability $\frac{1}{2}$-$\boldsymbol{\delta}$

Weak Adversary

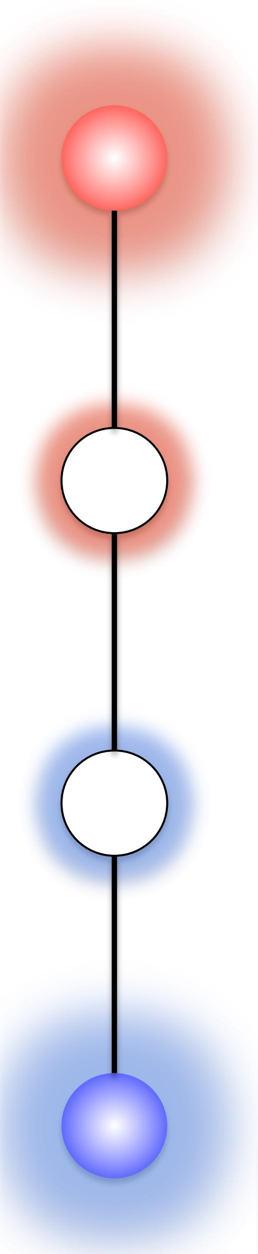

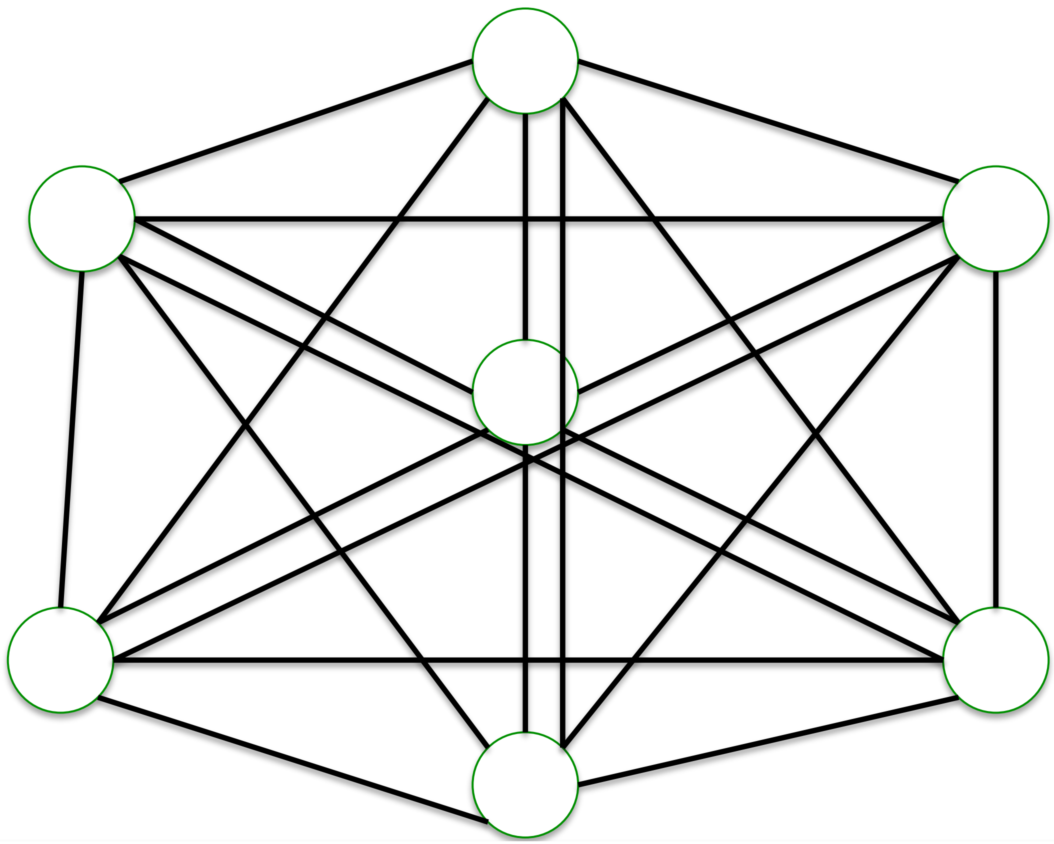

Probability of **Blue** is probability of majority **Blue** within experts – $\binom{n}{\frac{n}{2}}(\frac{1}{2}-\delta)^{\frac{n}{2}}$

Theorem 1

If experts' size is μn, for $\varepsilon<\mu$, for large enough n, there is an absolute constant c such that if highest degree Δ satisfies:

$$\Delta < c\frac{\epsilon\delta^4\mu n}{\log(\frac{1}{\epsilon})}$$

Then majority over vertices gives truth with probability at least 1-ε

Strong Adversary

A *strong adversary* is one that can choose the set of experts as well as what each experts says (but at the appropriate ratio).

Expander

An *expander* (n,d,λ) is a d-regular graph on n vertices, in which the absolute value of every eigenvalue besides the first is at most λ.

Theorem 2

Let G be a (n, d, λ)-graph, and suppose

$$\frac{d^2}{\lambda^2} > \frac{1}{\delta^2 \mu (1 - \mu + 2\delta\mu)}$$

Then for strong adversaries the majority answers truthfully.

Theorem 5 proof

A known theorem states that in a (n, d, λ)-graph:

$$\sum_{v \in V} (|N(v) \cap A| - \frac{d|A|}{n})^2 \leq \lambda^2 |A| (1 - \frac{|A|}{n})$$

(where $N(v)$ is the set of neighbors of vertex v)

Theorem 5 proof

Using this when A is the set of **Red** experts, and for B, the set of **Blue** ones, we add the equations, getting:

$$\sum_{v\in V}(|N(v)\cap A| - \frac{d|A|}{n})^2 + (|N(v)\cap B| - \frac{d|B|}{n})^2 \leq \lambda^2[|A|(1-\frac{|A|}{n}) + |B|(1-\frac{|B|}{n})].$$

Theorem 5 proof

We are interested in vertices which turn **Blue,** so have more **Blue** neighbors than **Red.**
These are set X.

$$\sum_{v \in X} (|N(v) \cap A| - \frac{d|A|}{n})^2 + (|N(v) \cap B| - \frac{d|B|}{n})^2$$

Theorem 5 proof

However, for $a>b$, $x \geq y$: $(x-b)^2+(y-a)^2 \geq (a-b)^2/2$, so:

$$\sum_{v \in X} (|N(v) \cap A| - \frac{d|A|}{n})^2 + (|N(v) \cap B| - \frac{d|B|}{n})^2 \geq |X| \frac{d^2(|A|-|B|)^2}{2}$$

Theorem 5 proof

Hence

$$|X|\frac{d^2(|A|-|B|)^2}{2} \leq \lambda^2[|A|(1-\frac{|A|}{n})+|B|(1-\frac{|B|}{n})]n^2$$

And we need X to be less than $(\frac{1-\mu}{2}+\delta\mu)n$

Random Graphs

A *random graph G(n,p)* is one which contains *n* vertices and each edge has a probability *p* of existing.

Theorem 3

There exist a constant *c* such that if $\mu < ½$, in a random graph G(n,p), if

$$d = np \geq c \cdot \max\{\frac{\log(\frac{1}{\mu})}{\delta^2}, \frac{1}{\mu\delta}\}$$

The majority will show the truth with high probability even with a strong adversary

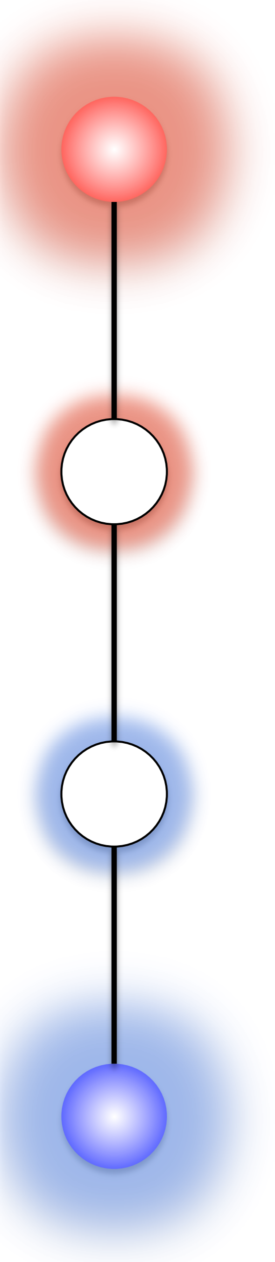

Iterative Propagation

Iterative Propagation

Allowing propagation to be a multi-step process rather than a "one-off" step can be both harmful and beneficial for some adversaries

Weak Adversary

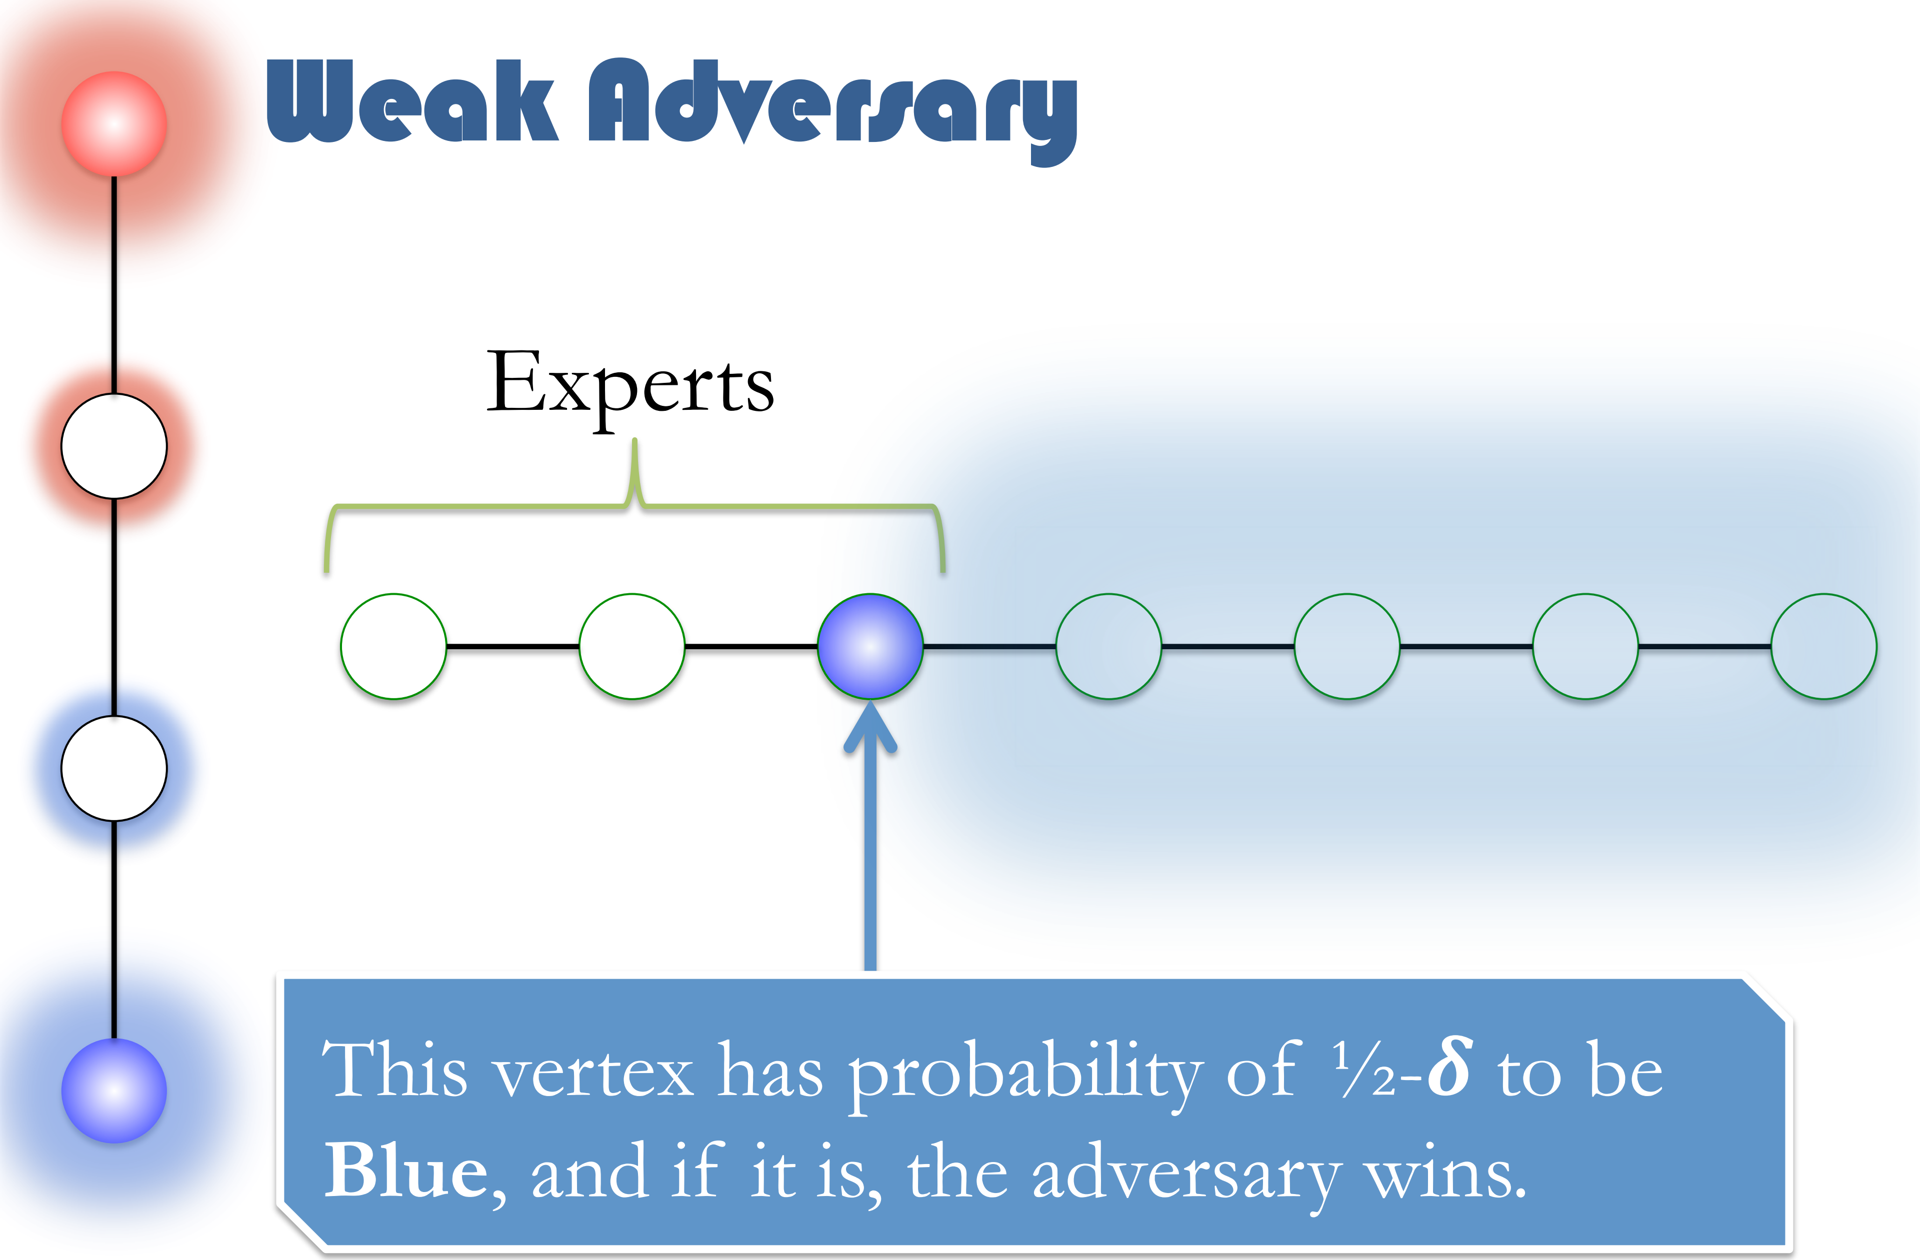

Random Process

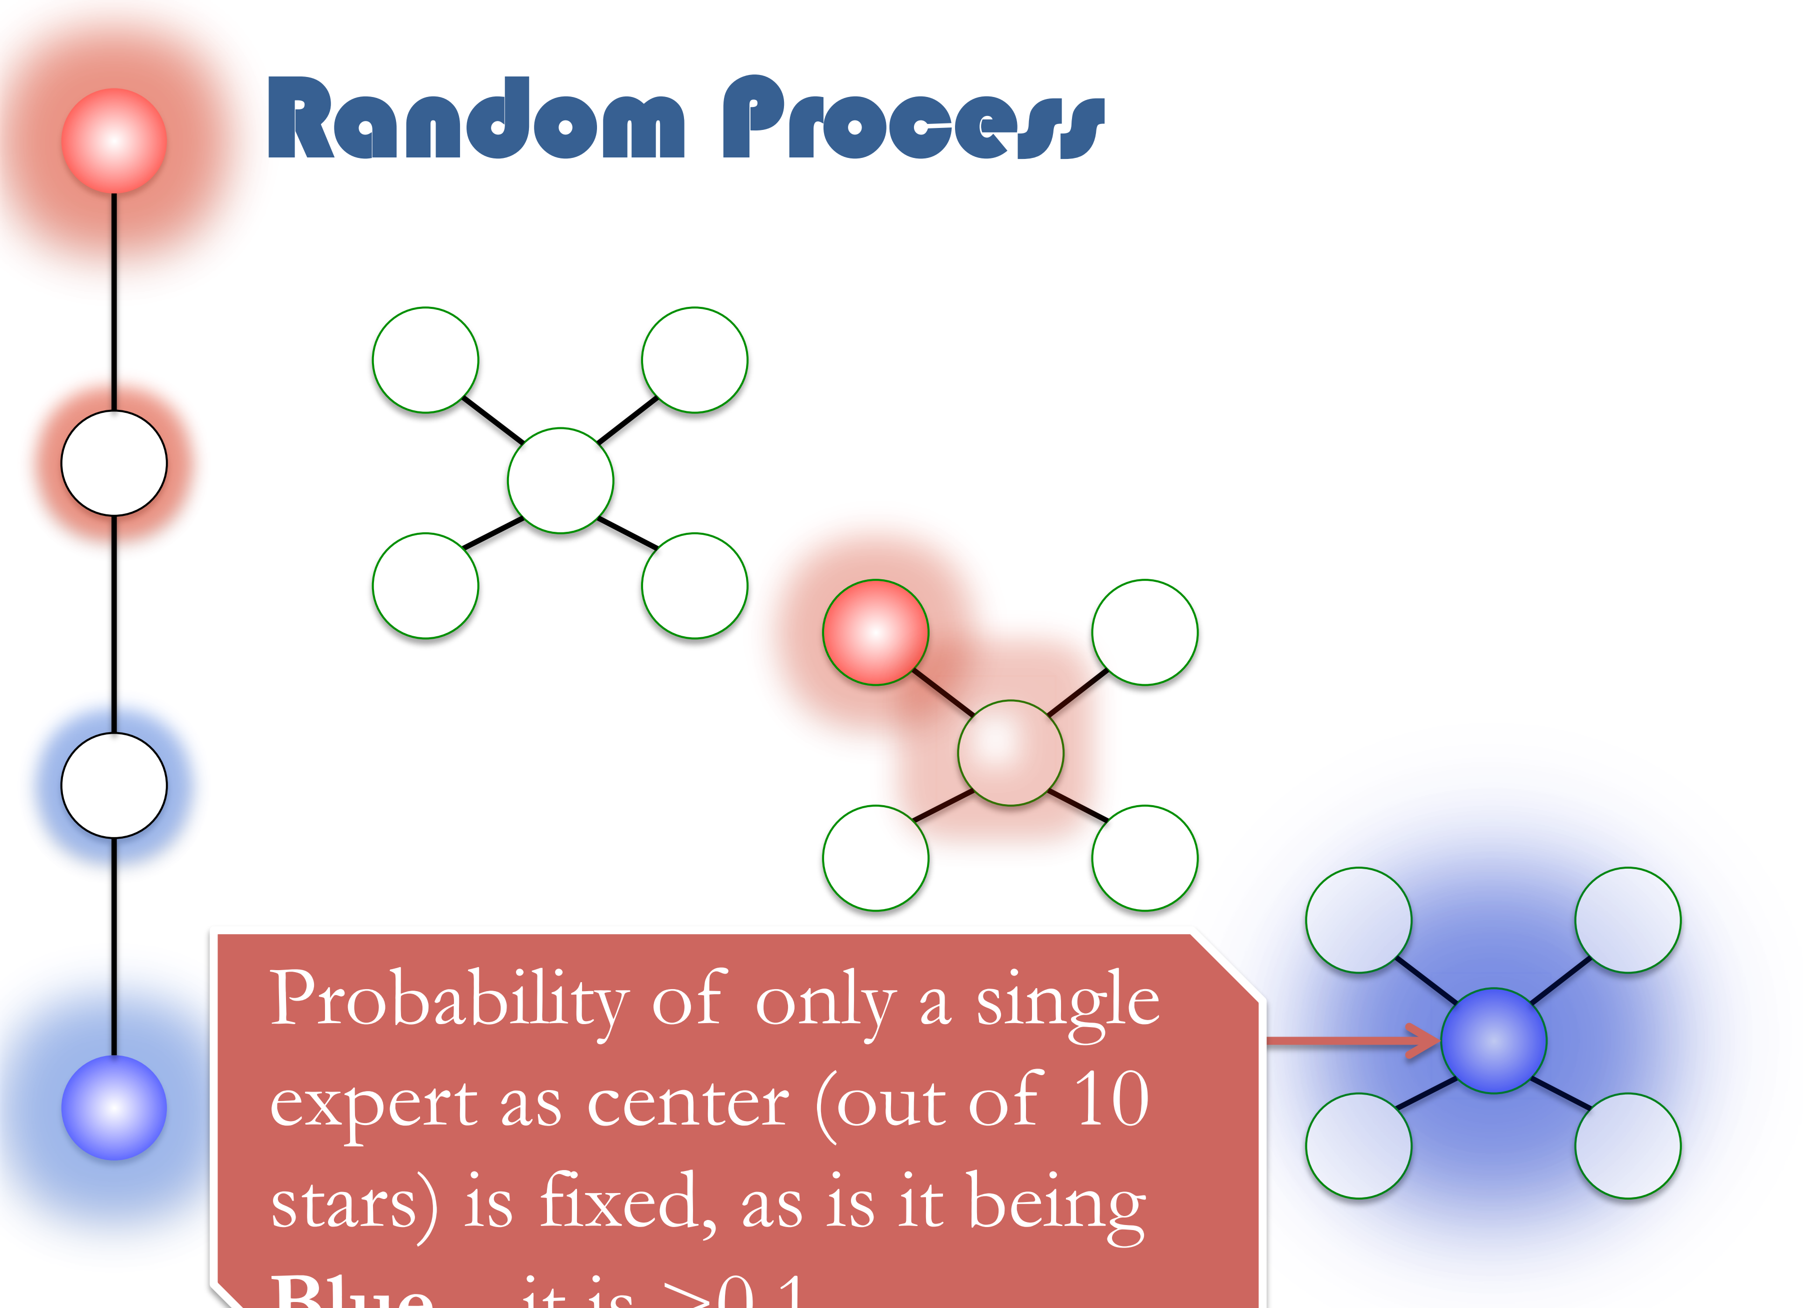

Probability of only a single expert as center (out of 10 stars) is fixed, as is it being **Blue** – it is >0.1

Random Process

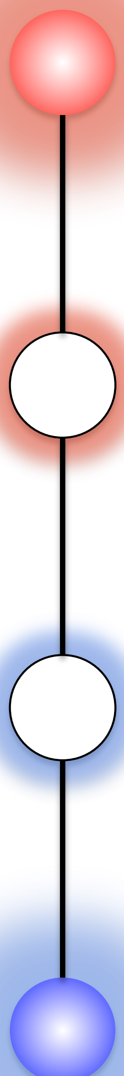

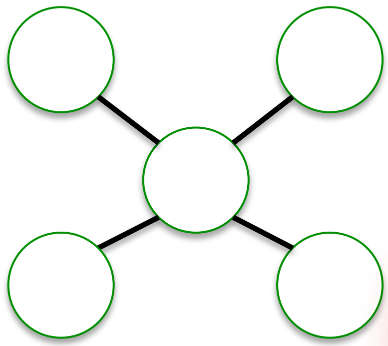

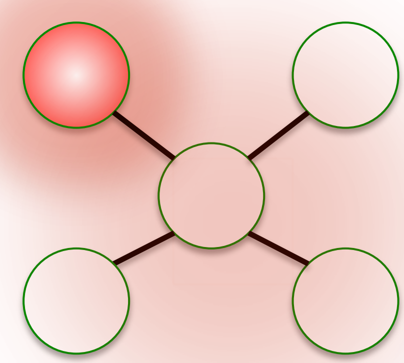

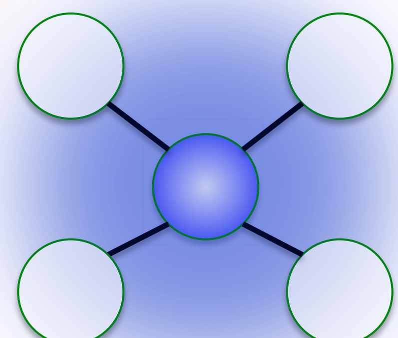

Now, regardless of location, a **Red** in a star colors the star **Red**

Future Research

Other ways to aggregate social graph information may result in different bounds

Hybrid capabilities of adversaries

More specific types of graphs

Multiple adversaries

Thanks for listening!