The Logic of Social Influence in Networks

-An Introduction-

Zoé Christoff

university of groningen

NWO

TOPOS Seminar

-17 March 2022 -

Logical Foundations of social influence

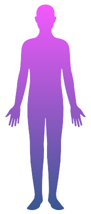

- understand how human agents reason and behave (empirical)

- understand how ideally rational agents would behave (theoretical)

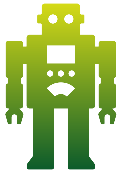

Big picture

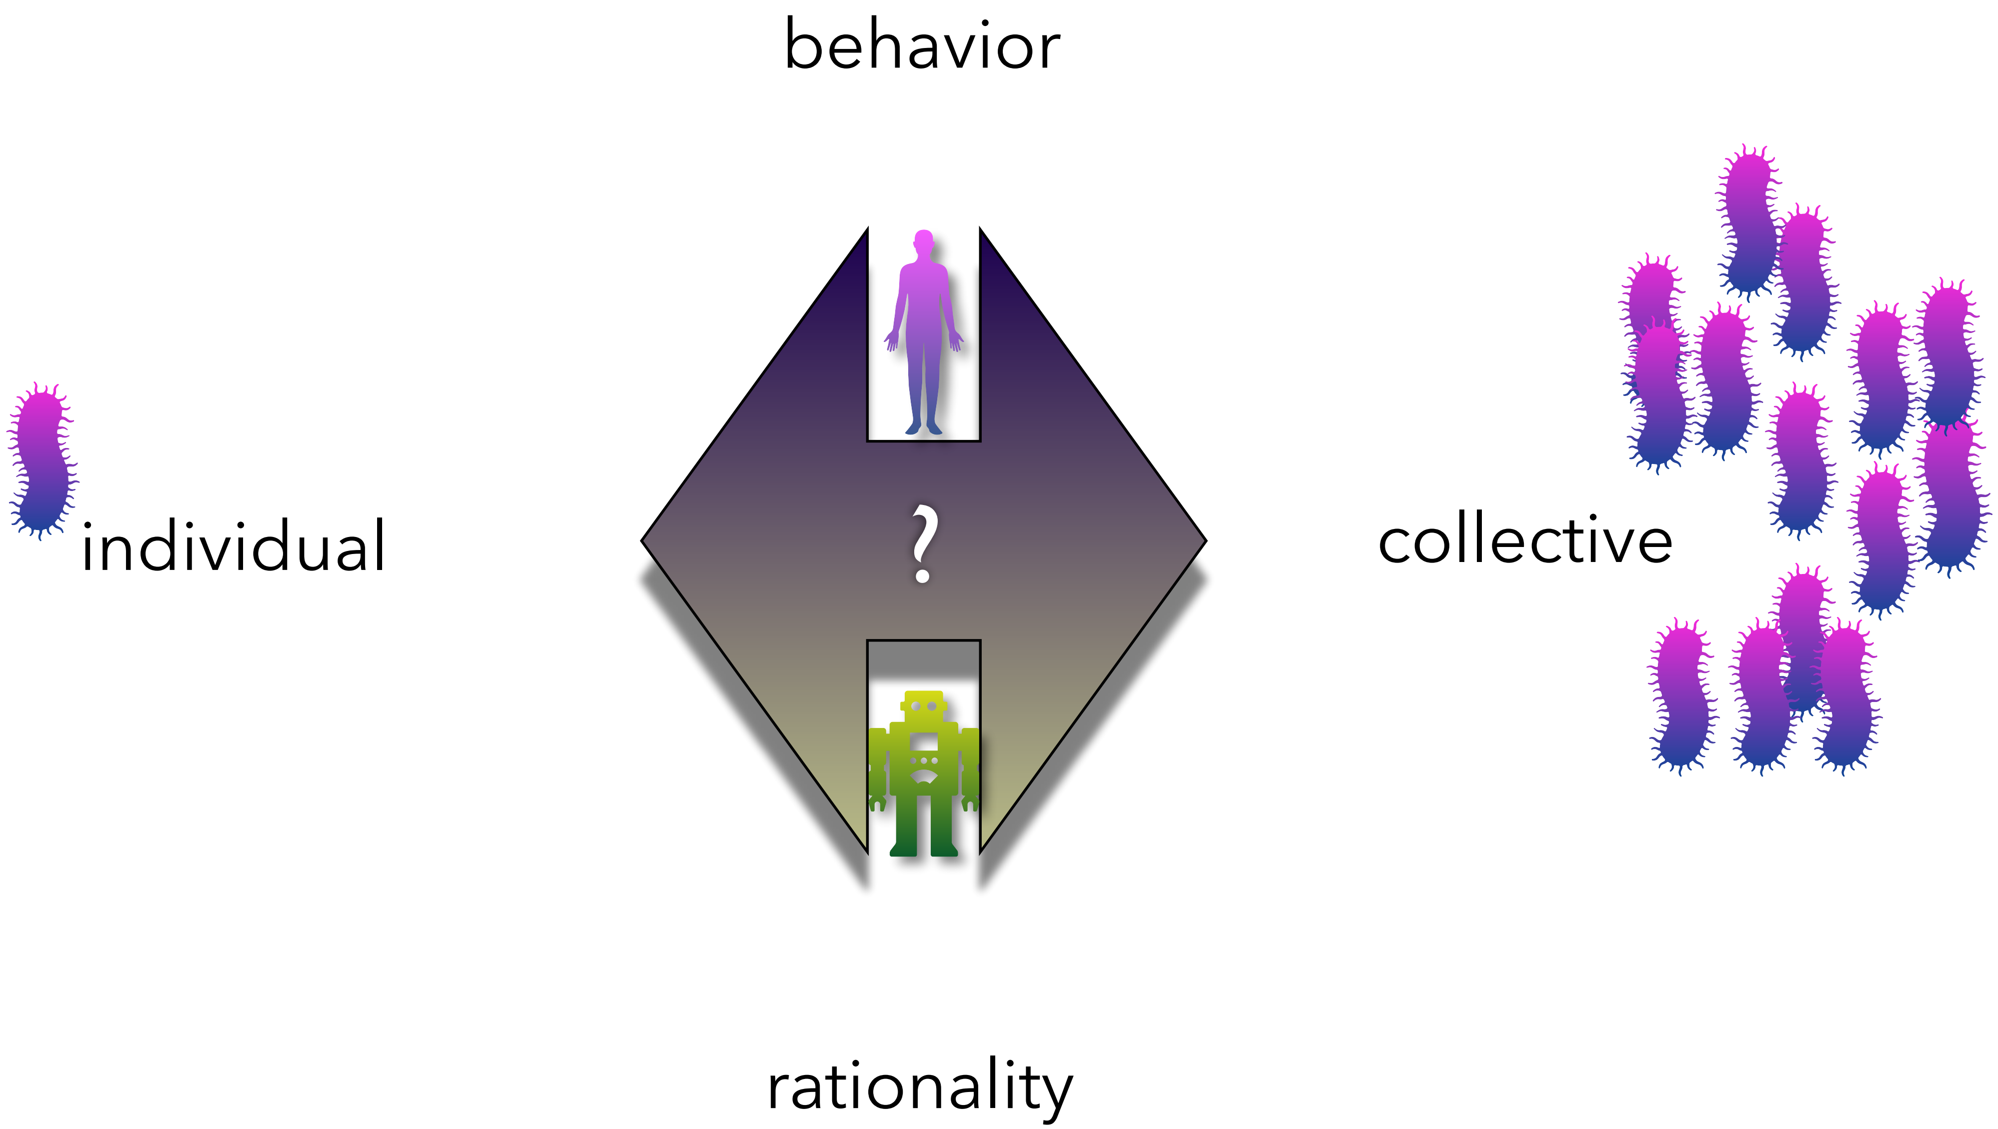

Example: Schelling's Segregation Model

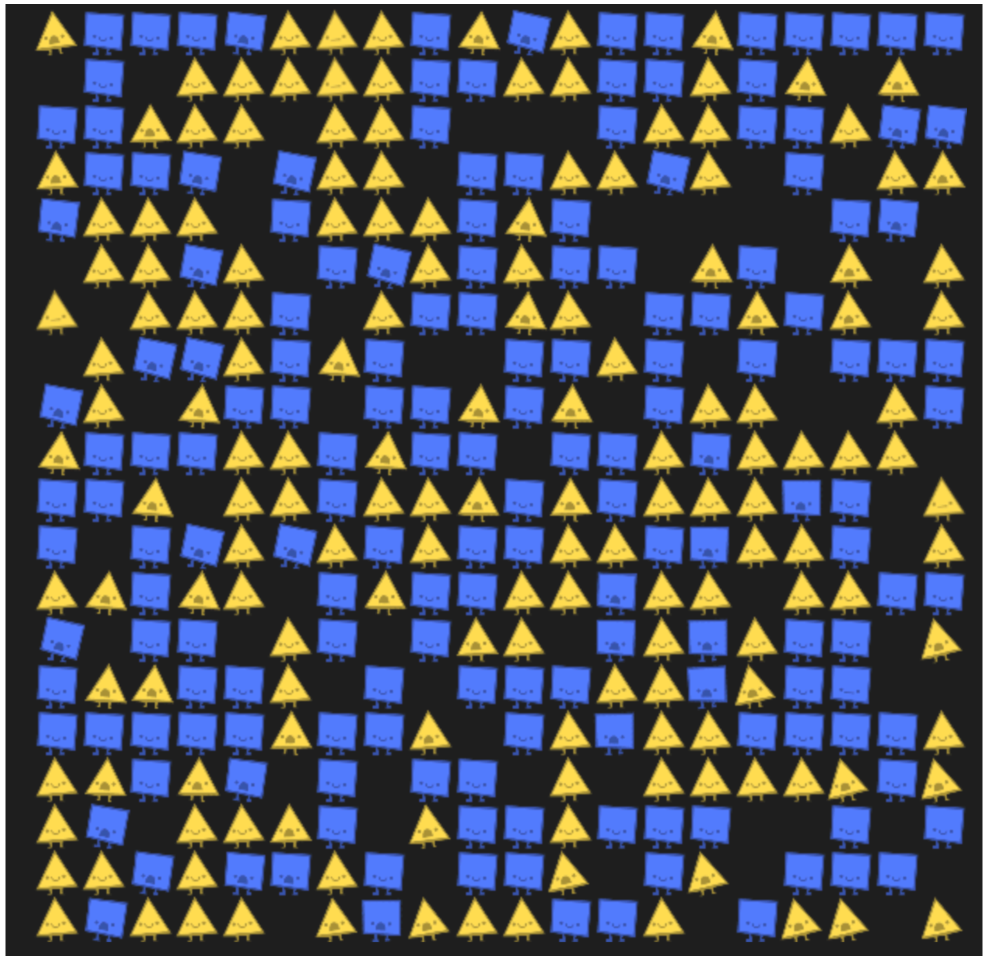

....

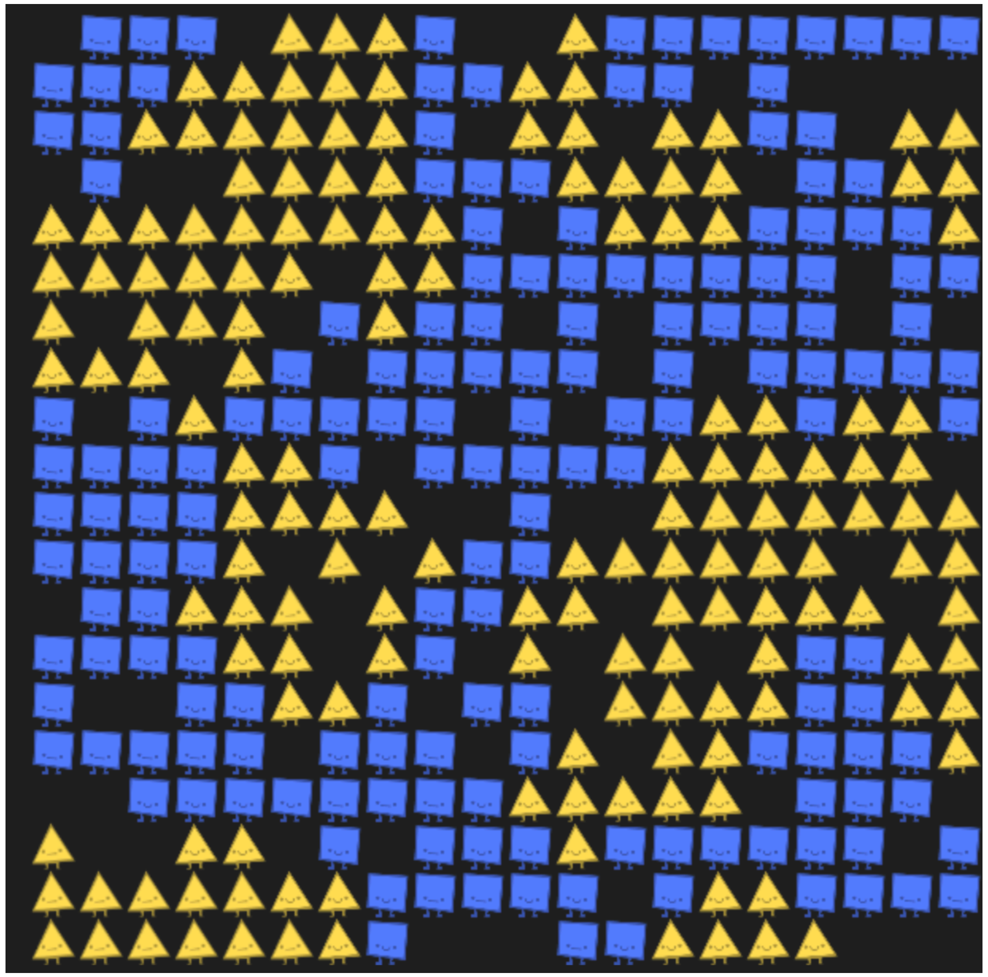

- uniform rule: prefer to move if less than 1/3 of your neighbors are of your type

A small individual bias has a huge collective impact.

Pictures: Nicky Case's simulation, try it <u>here</u>

What does the model show?

I. How do humans behave collectively?

Take a bunch of humans

Independent guesses

Galton F, Vox Populi, Nature, 1907..

Independent guesses

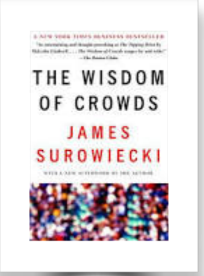

Independent guesses

The Condorcet's Jury Theorem
(Marquis of Condorcet,1784)

- The most basic jury theorem in social choice
- N = the number of jurors
- p = the probability of an individual juror being right
- μ = the probability that a jury gives the correct answer

$$\mu = \sum_{i=m}^{N} \left(\frac{N!}{(N-i)!\,i!} \right) (p)^i (1-p)^{N-i}$$

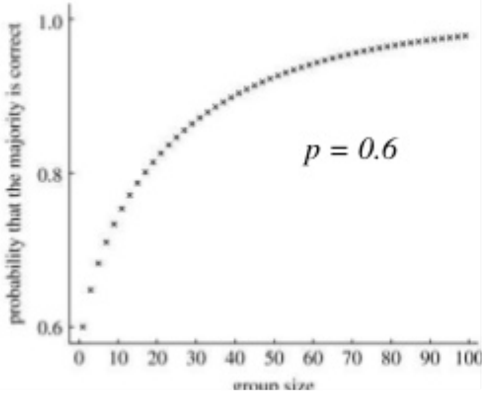

- $p > 0.5$ implies $\mu > p$.
- and $\mu \rightarrow 1$ when $N \rightarrow \infty$.

Take a bunch of … agents

The Condorcet's Jury Theorem
(Marquis of Condorcet,1784)

- The most basic jury theorem in social choice
- N = the number of jurors
- p = the probability of an individual juror being right
- μ= the probability that a jury gives the correct answer

$$\mu = \sum_{i=m}^{N} \left(\frac{N!}{(N-i)!\,i!} \right) (p)^i (1-p)^{N-i}$$

- $p > 0.5$ implies $\mu > p$.
- and $\mu \rightarrow 1$ when $N \rightarrow \infty$.

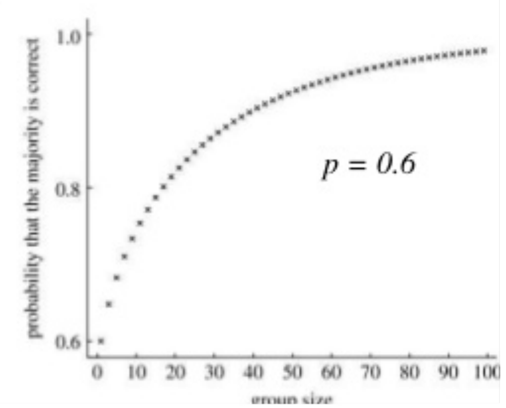

Without social influence

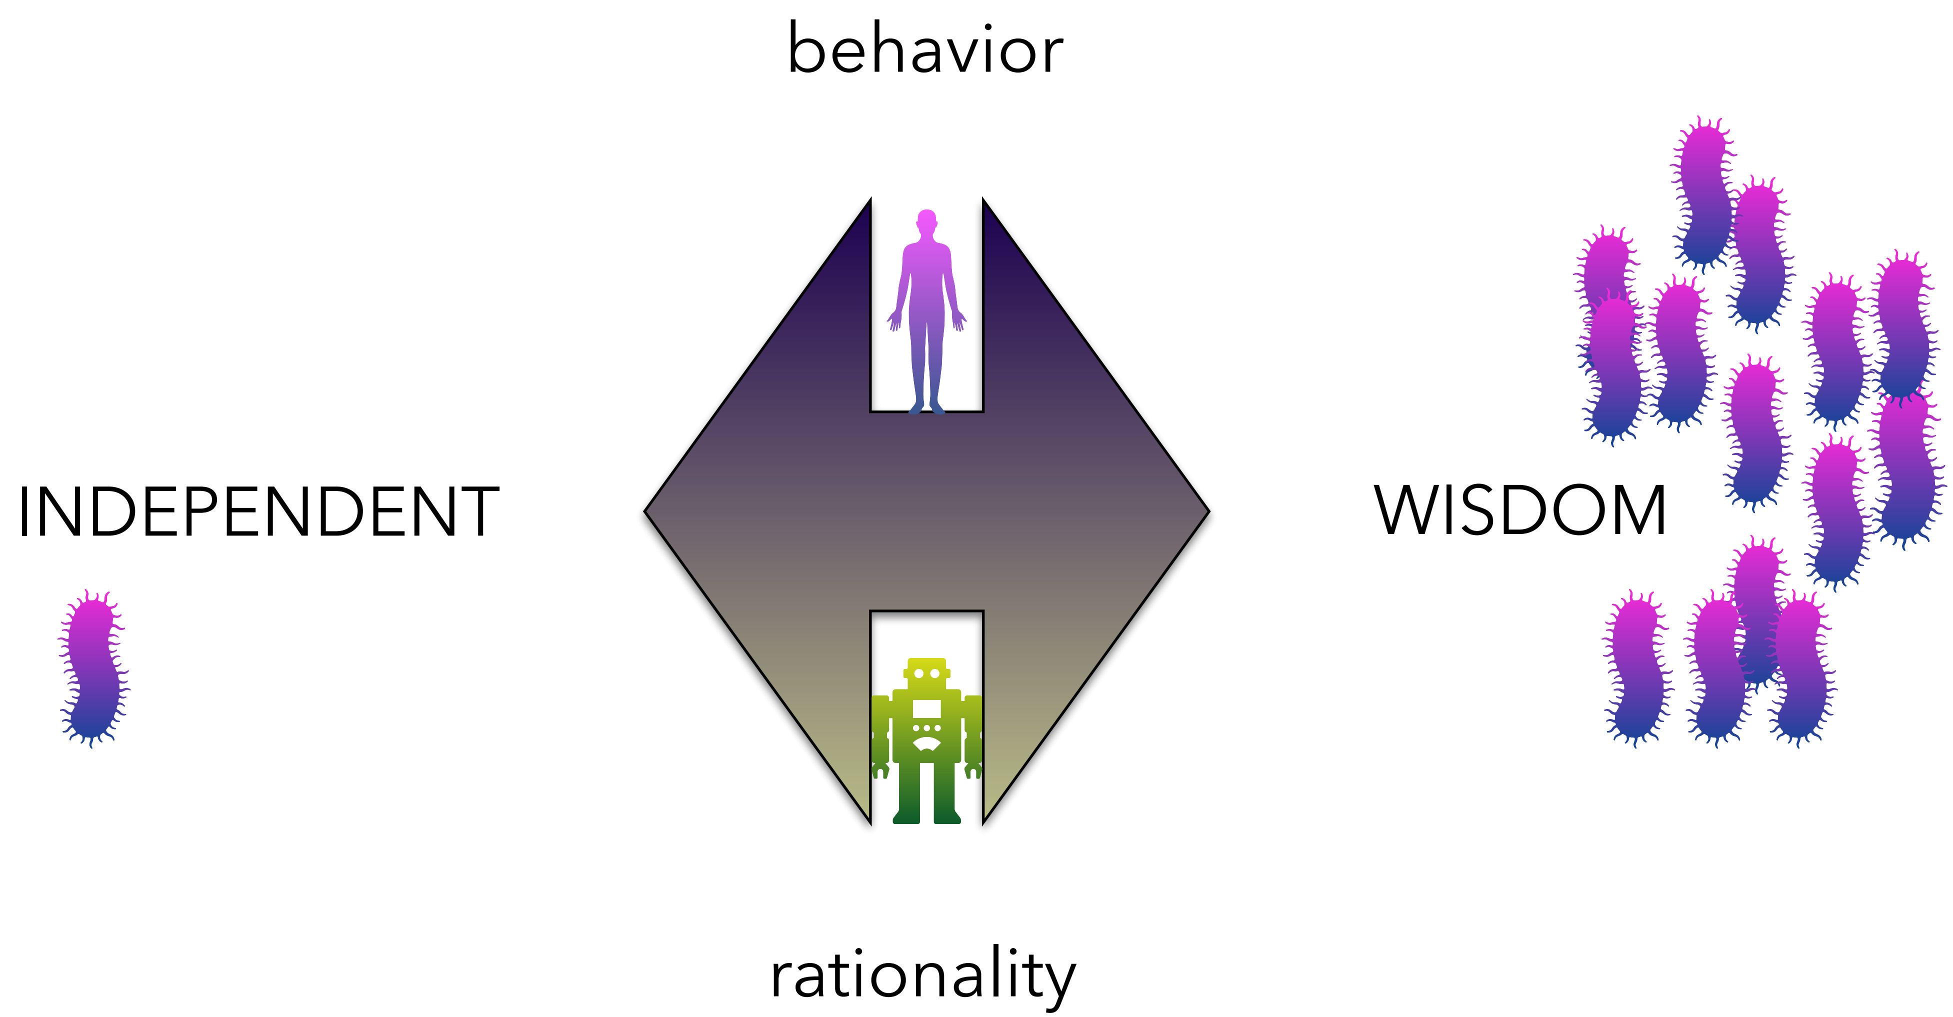

With social influence?

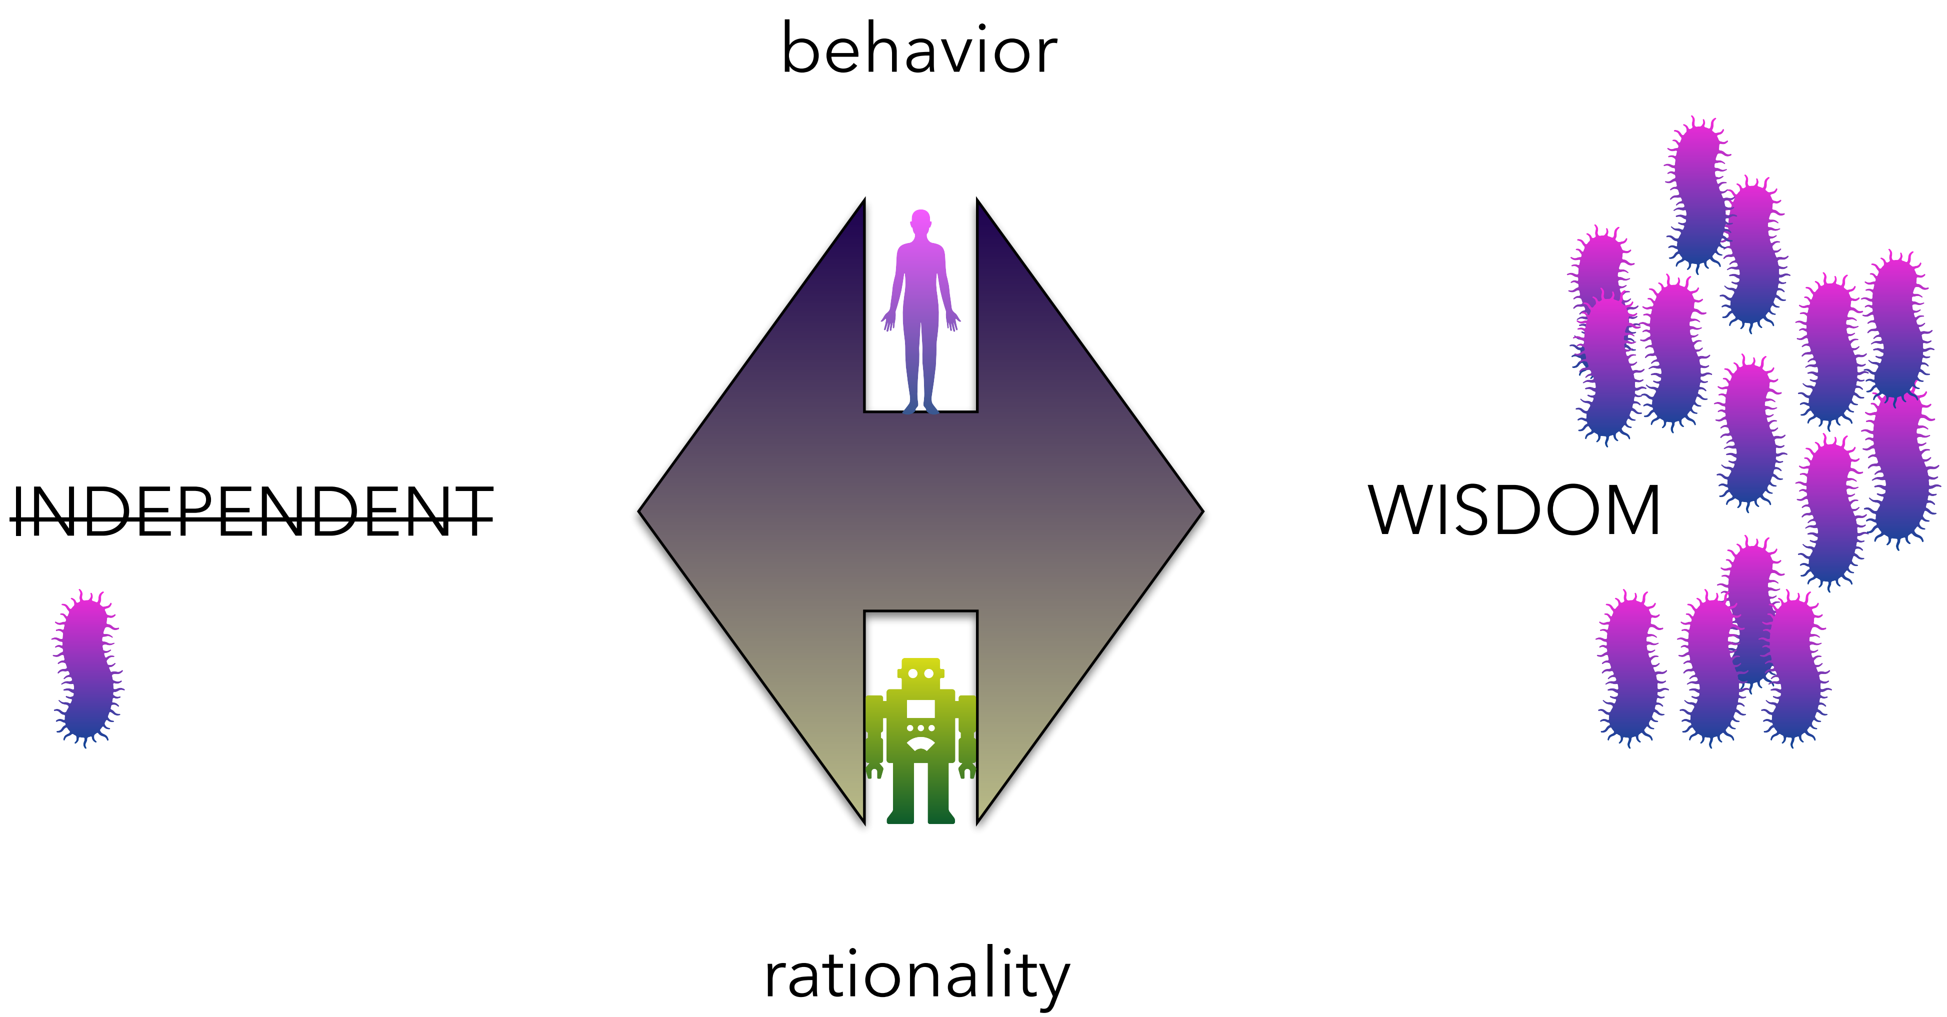

Sometimes…

…individuals lead each other in the wrong direction

Cartoon by Kevin Kallaugher (KAL) for *The Economist*

Correlated behavior and collective failures

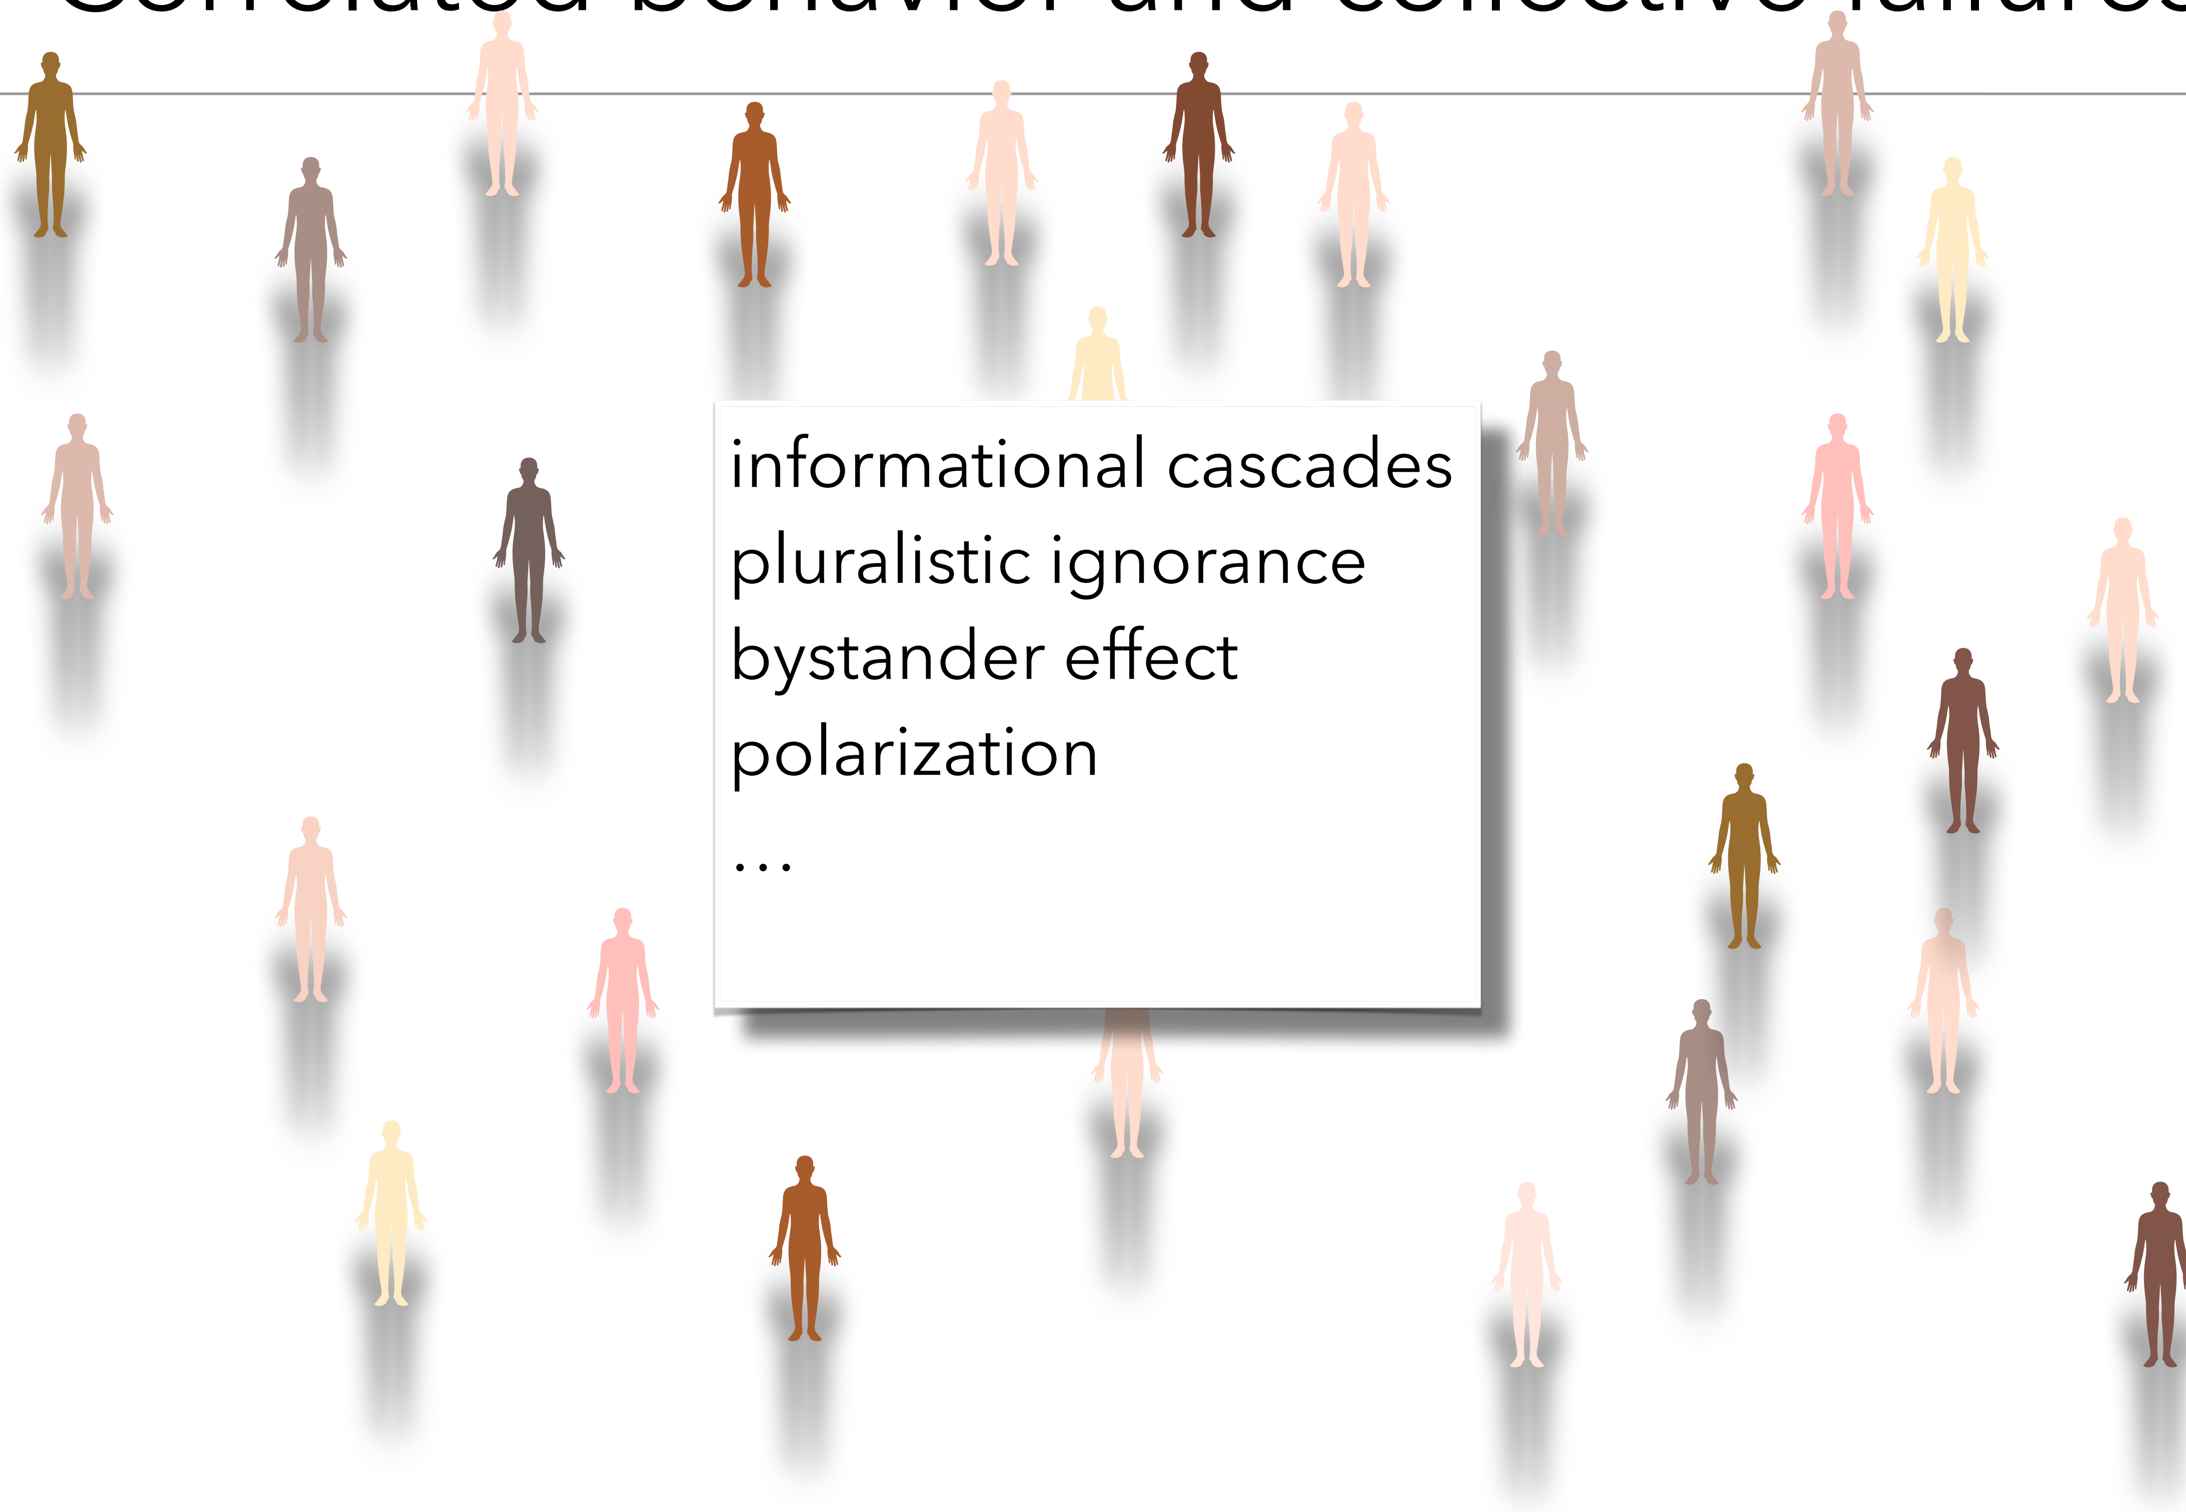

informational cascades
pluralistic ignorance
bystander effect
polarization
…

1) Pluralistic ignorance

"a situation where the majority of group members privately rejects the norm, but assumes (incorrectly) that most others accept it" *(Katz & Allport 1931:152)*

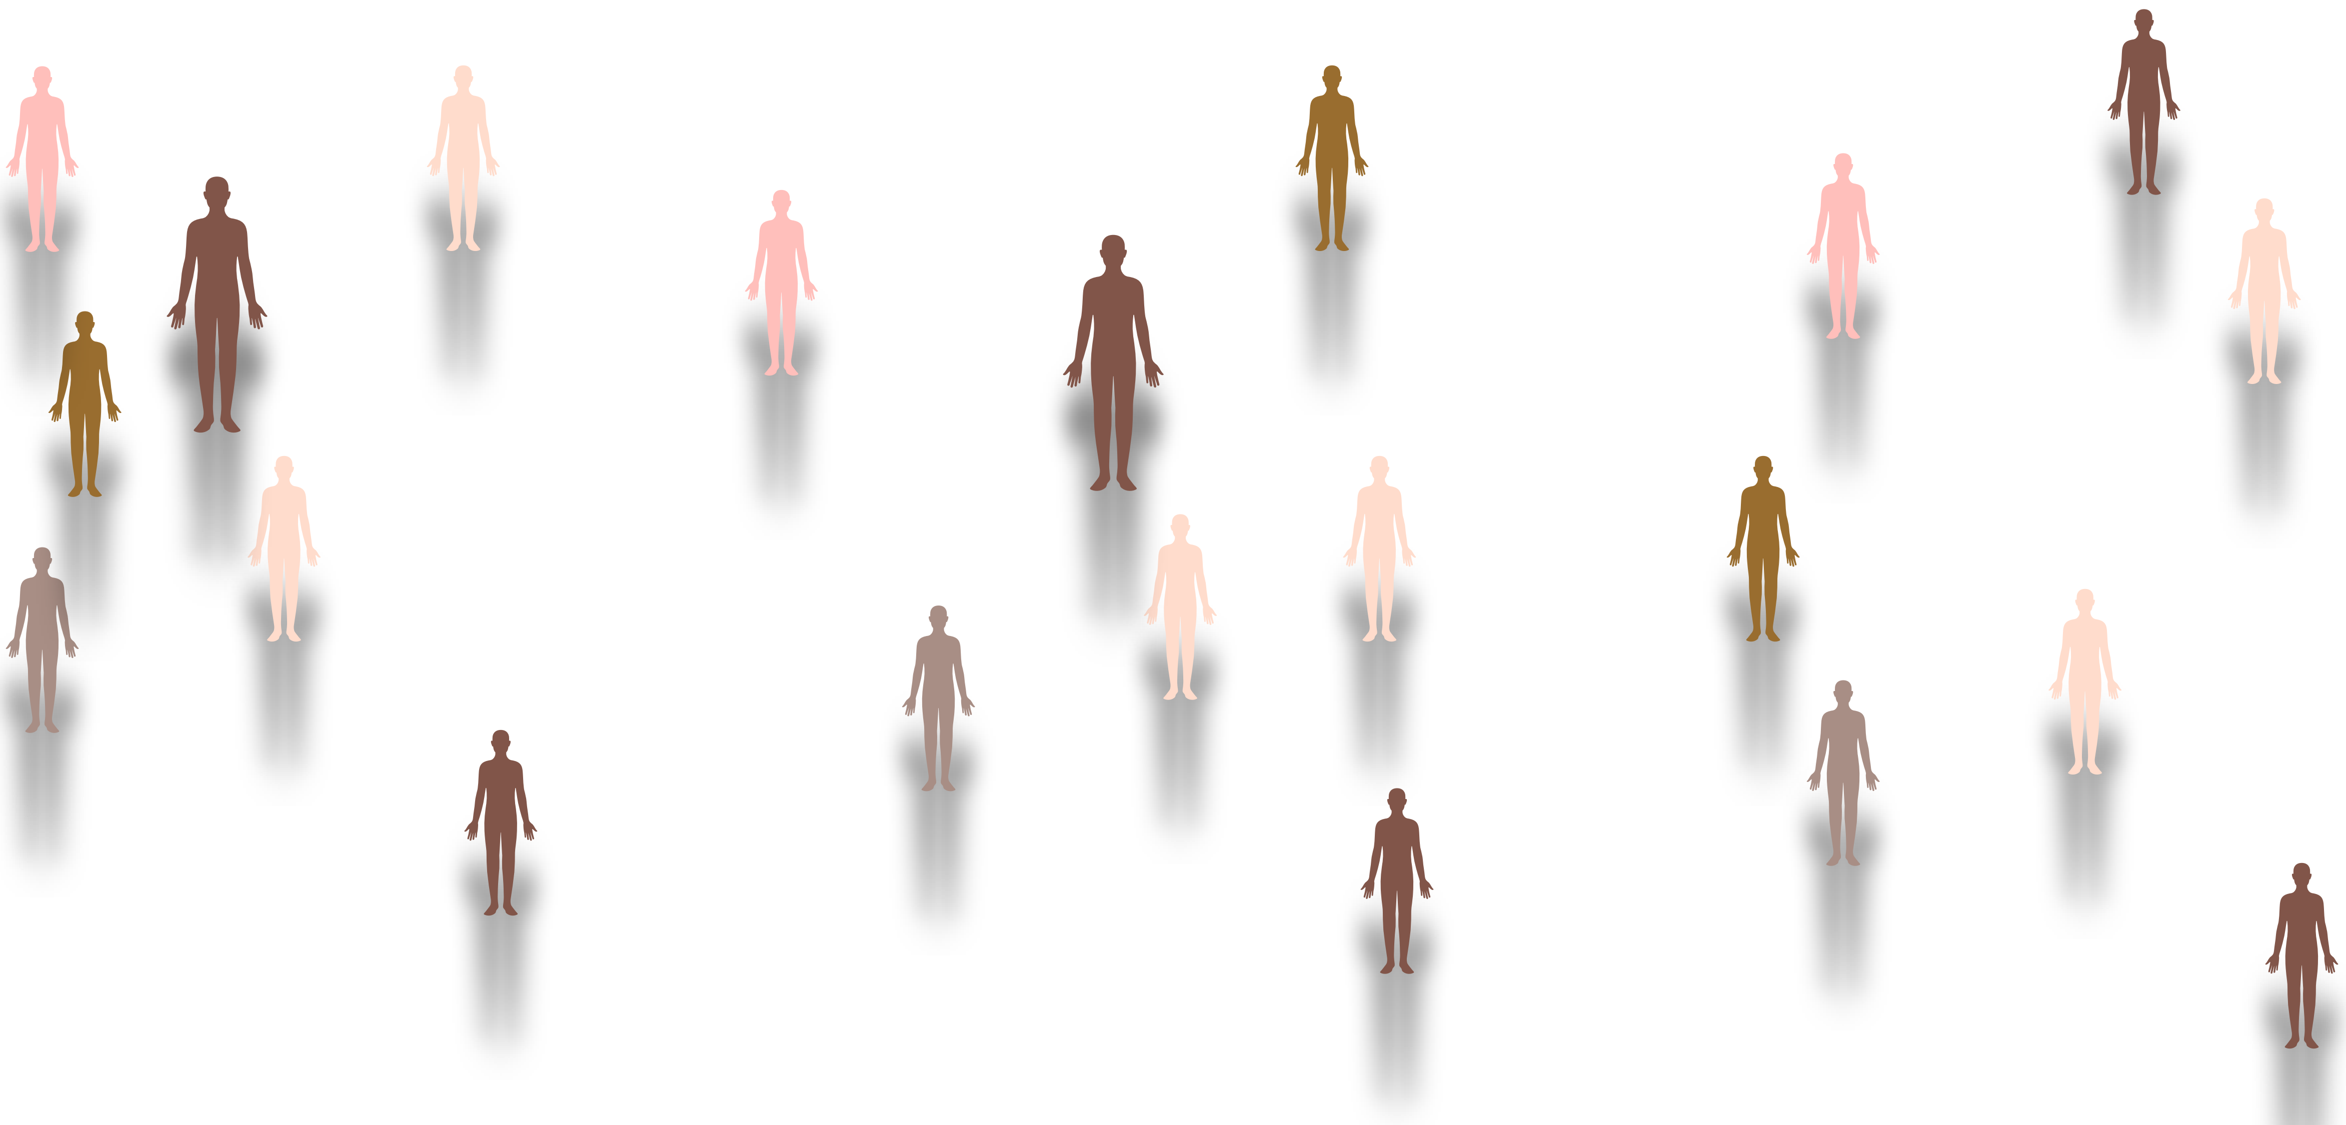

The Emperor's new clothes

(H.C. Anderson, Fairy Tales Told for Children, 1837)

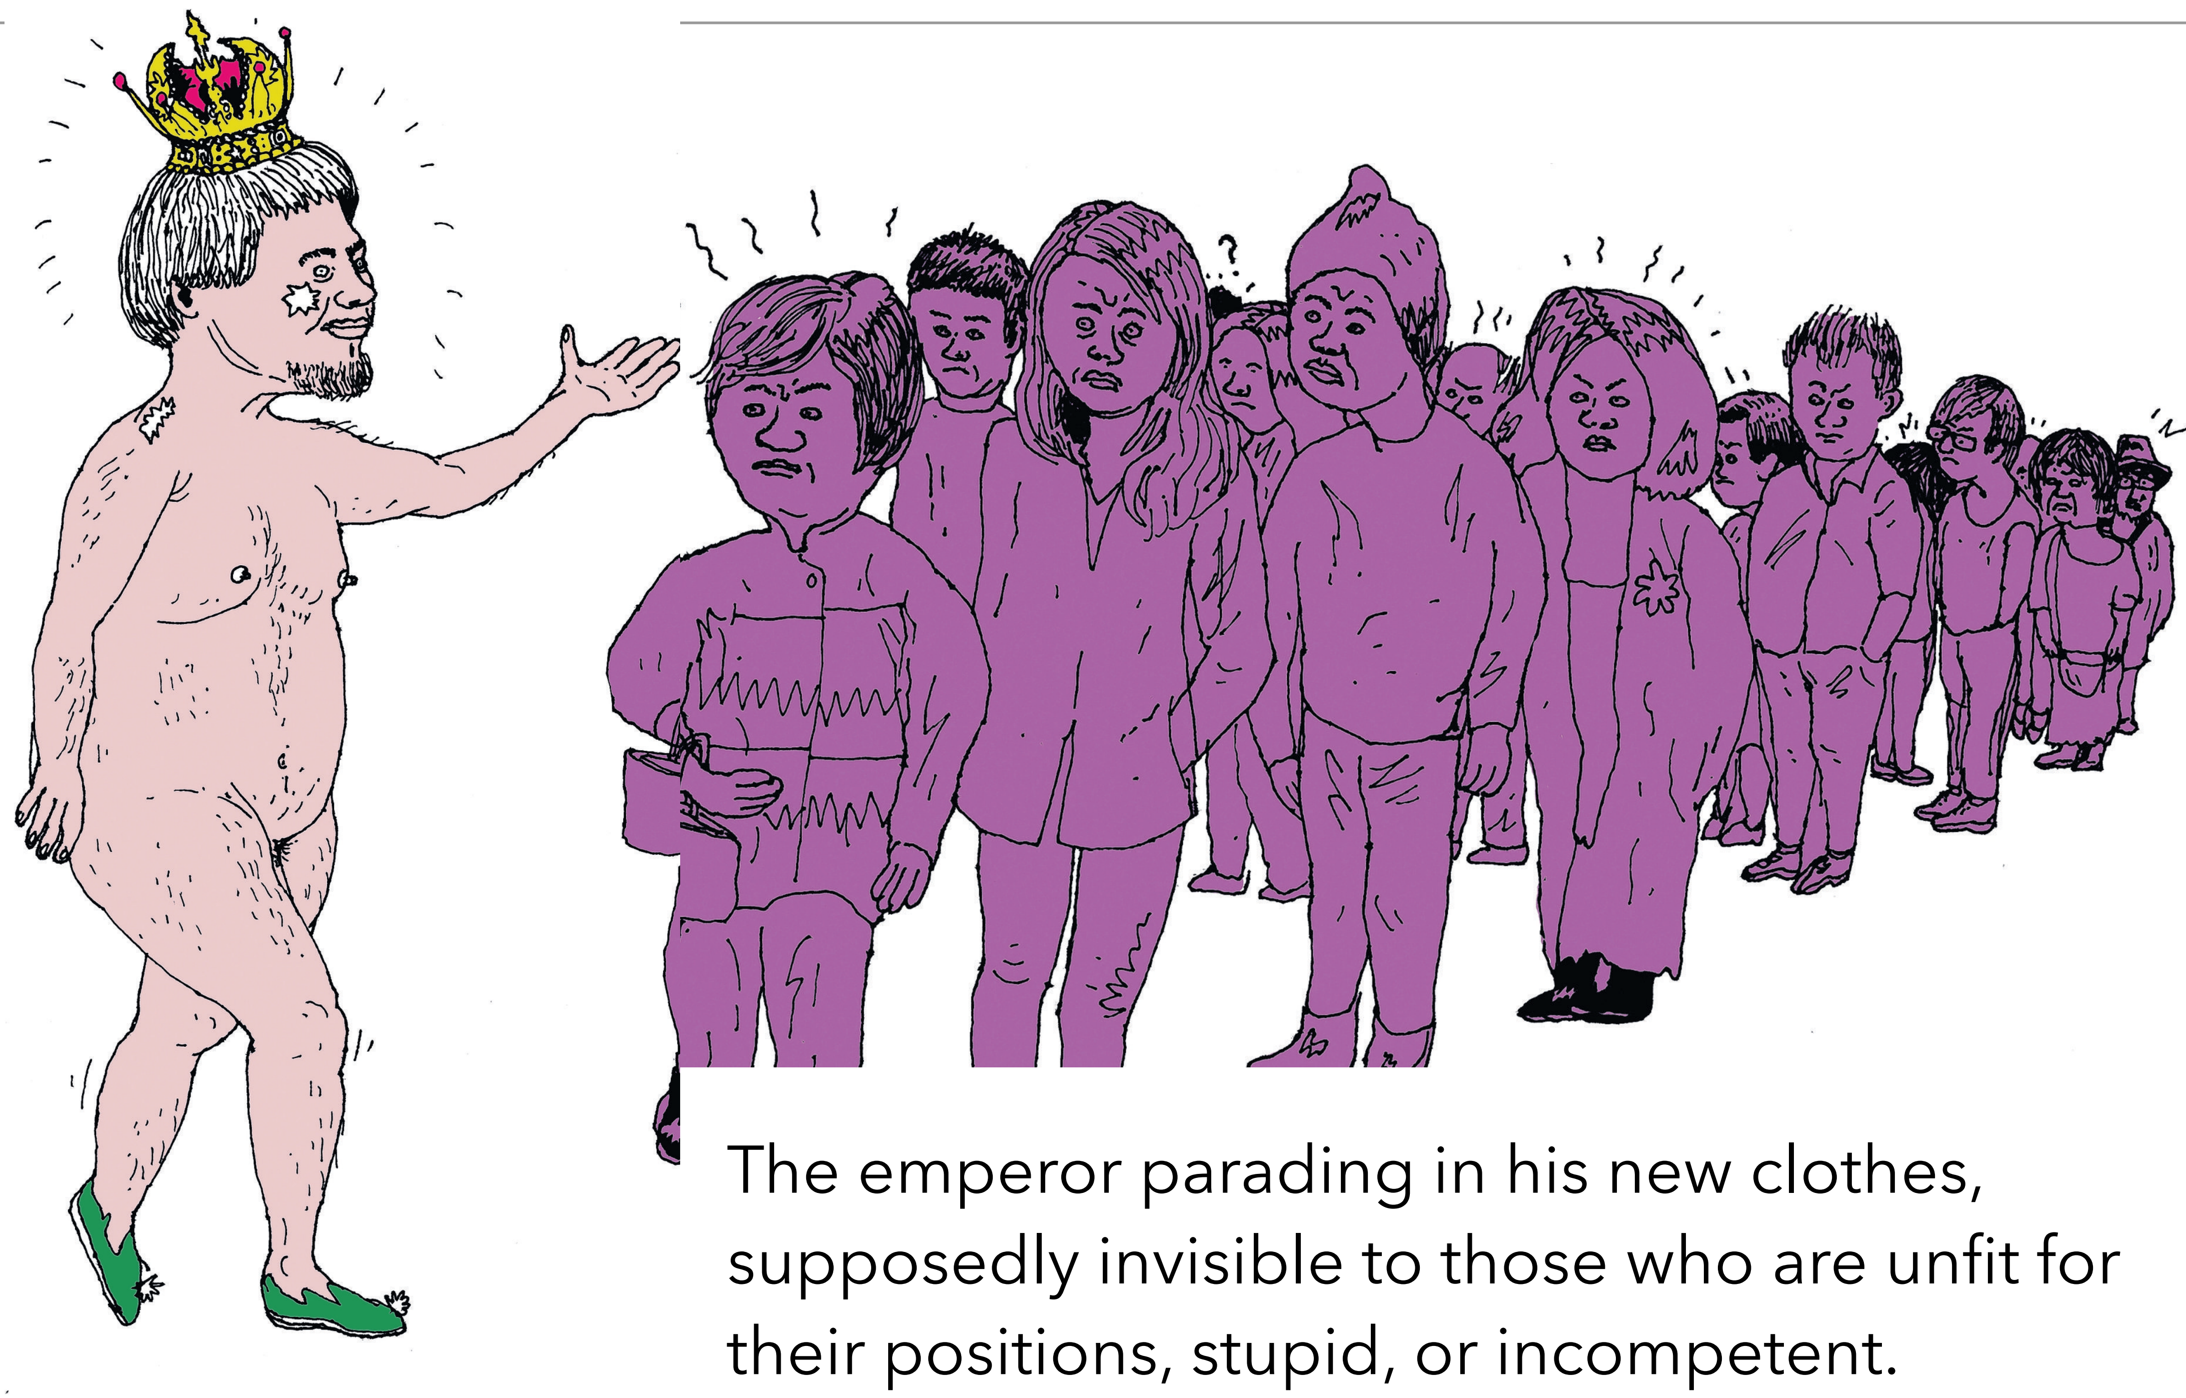

The emperor parading in his new clothes, supposedly invisible to those who are unfit for their positions, stupid, or incompetent.

Illustration by Johan Kleinjan ® Dutch newspaper Trouw, published April 2nd 2016

Pluralistic ignorance

“a situation where the majority of group members privately rejects the norm, but assumes (incorrectly) that most others accept it” *(Katz & Allport 1931:152)*

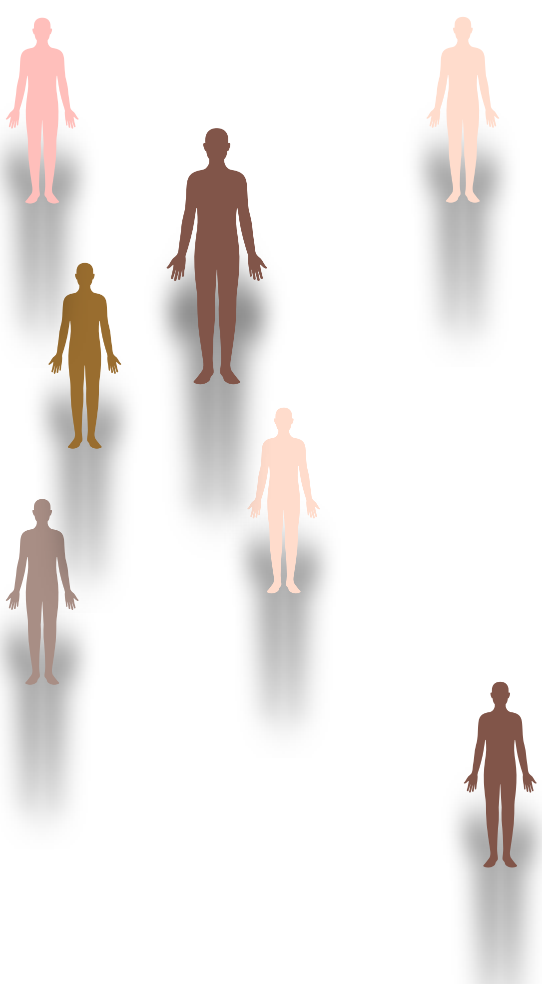

Documented examples include:

- ☐ classroom situations
- ☐ college drinking norms
- ☐ segregation norms

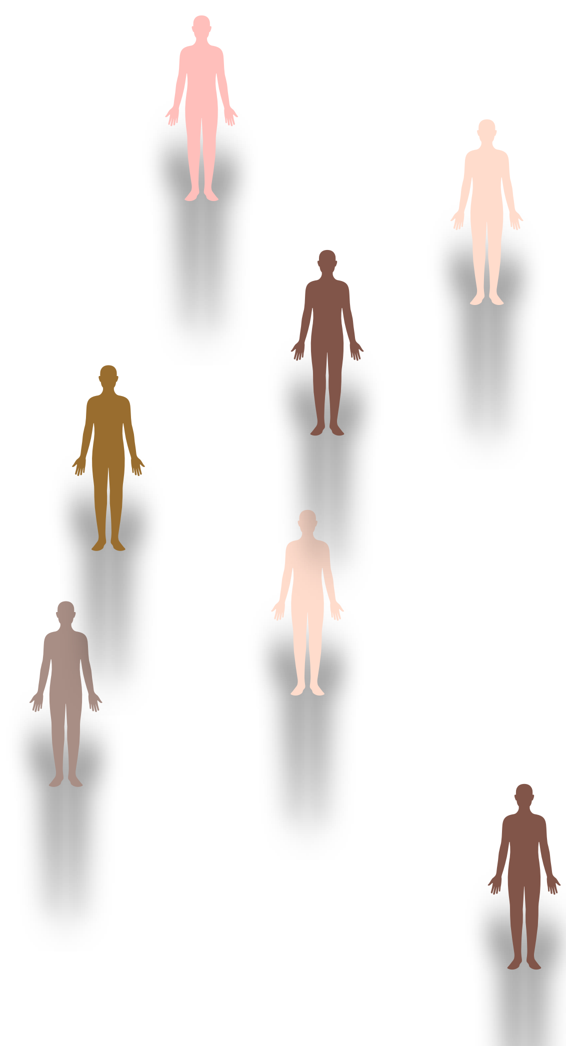

Private *vs* Expressed Opinions

"All those in favor say 'Aye.'"
"Aye." *"Aye."* *"Aye."* *"Aye."* *"Aye."*

2) Informational Cascades

- Agent disregard their own private information to follow the choice of some preceding agents
- All remaining agents pick the same option, even if they have diverging private evidence
- This imitation effect might lead the whole community to make the worst possible choice

Example: Choose a restaurant

A

B

Better reviews

worse reviews

Private information

Public information

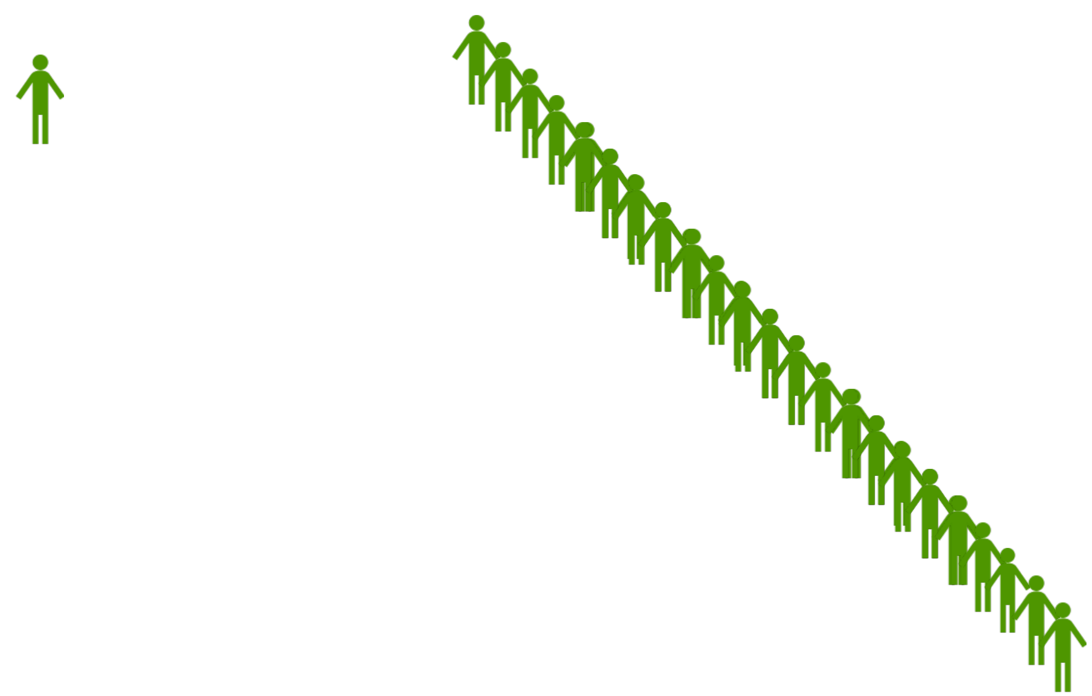

Mindless imitation effects? (again)

Cartoon by Kevin Kallaugher (KAL) for *The Economist*

What type of results?

Social psychologists are likely interested in:

- what are the conditions for such a phenomenon to arise?
- where is the individual "error" leading to such collective catastrophic results?
- Is the phenomenon preventable/correctable?

Can logic help with any of these?

II. Models of informational cascades

- ***Logical Models of Informational Cascades*** (pdf), Alexandru Baltag, Zoé Christoff, Jens Ulrik Hansen and Sonja Smets, in J. van Benthem and F. Liu (Eds.): Logic across the University: Foundations and Applications, — Proceedings of the Tsinghua Logic Conference, Beijing, 14-16 October 2013, Studies in Logic, Volume 47, pp.405-432, College Publications, London, (2013).

Analysis of a cascade

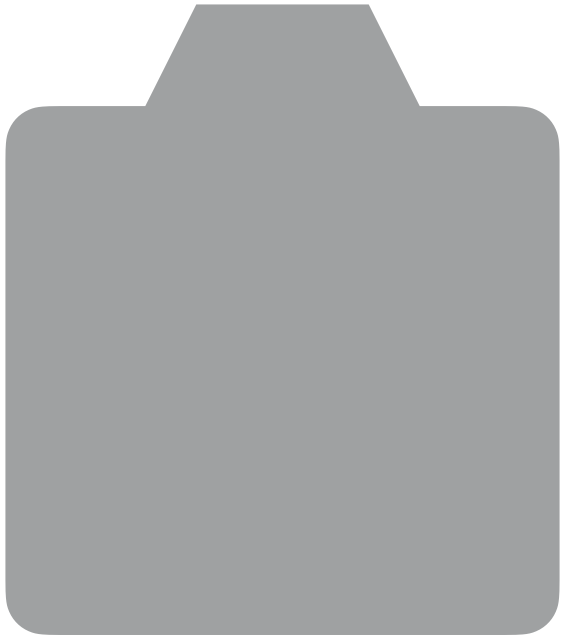

1 (opaque) urn, containing three marbles

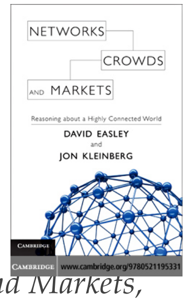

Chapter 16 of: Easley D. & Kleinberg J. Networks Crowds and Markets, CUP, 2010.

What is the content of the urn ?

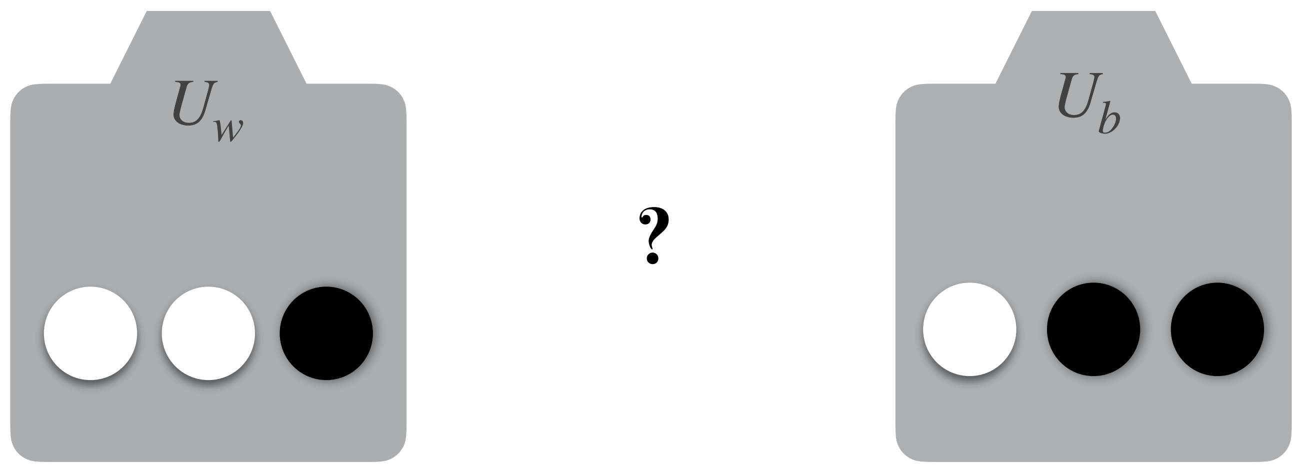

Goal: guess correctly the content of the urn, given:

your own secret observation 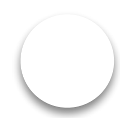 or 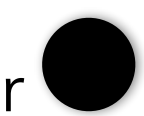

AND

the visible guesses of previous players

First observation and guess

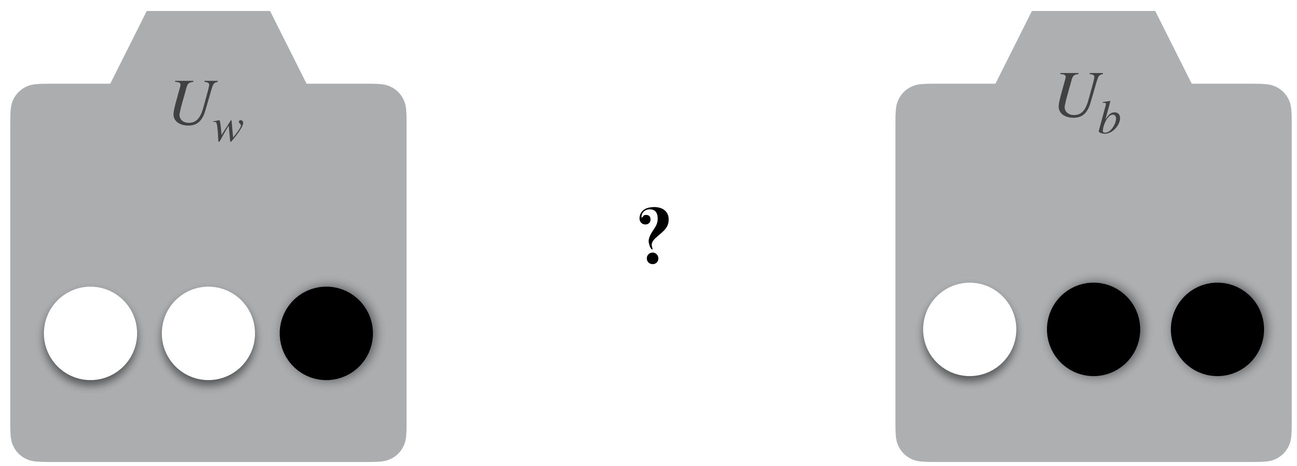

- if observe 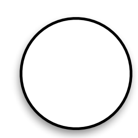 guess U_W
- if observe 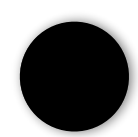 guess U_B

First observation and guess

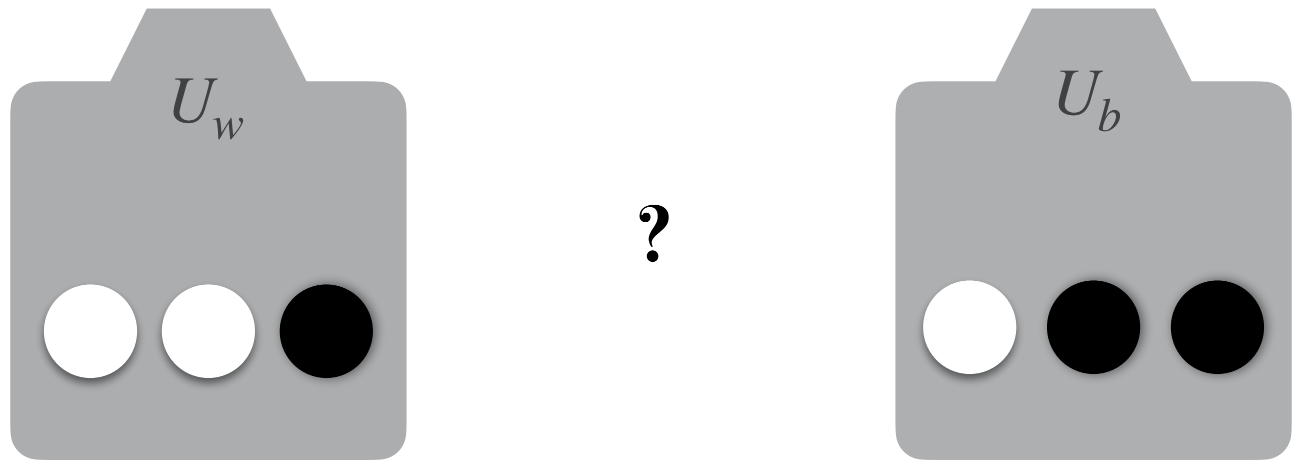

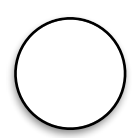

guess U_W

Second observation and guess

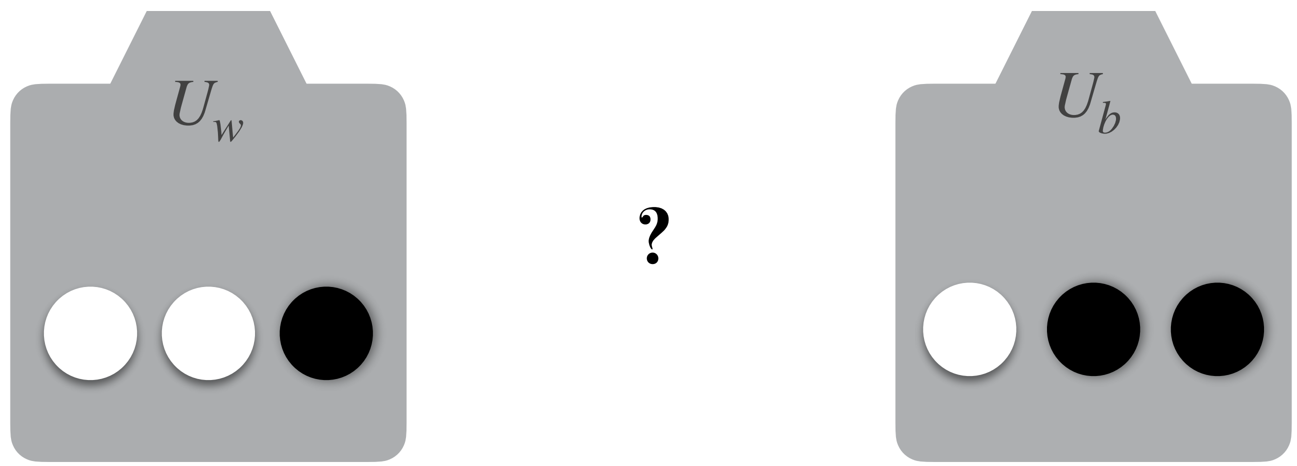

- observe U_W + 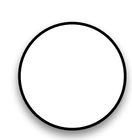 : guess U_W

Third observation and guess

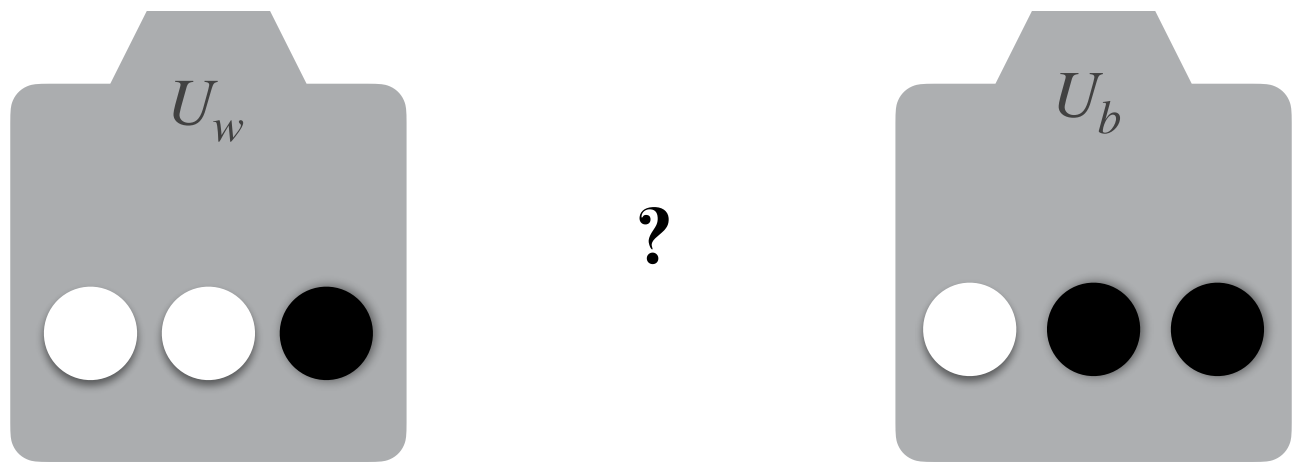

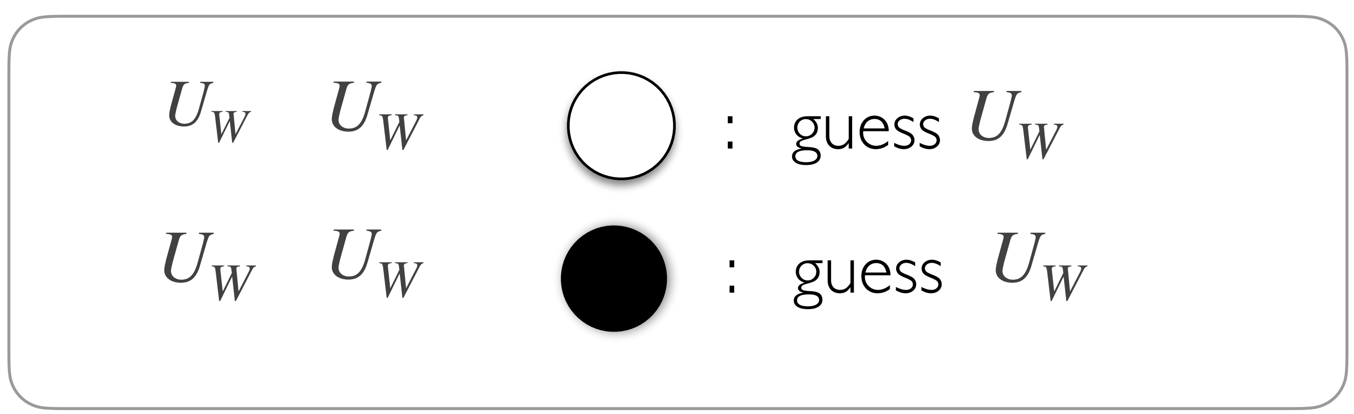

Conclusion: this is rational!?

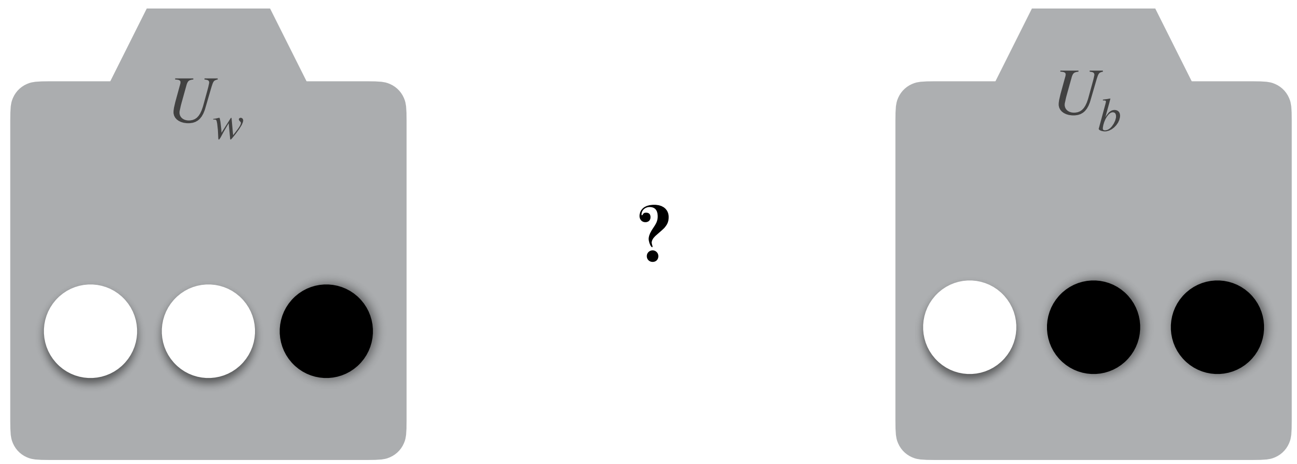

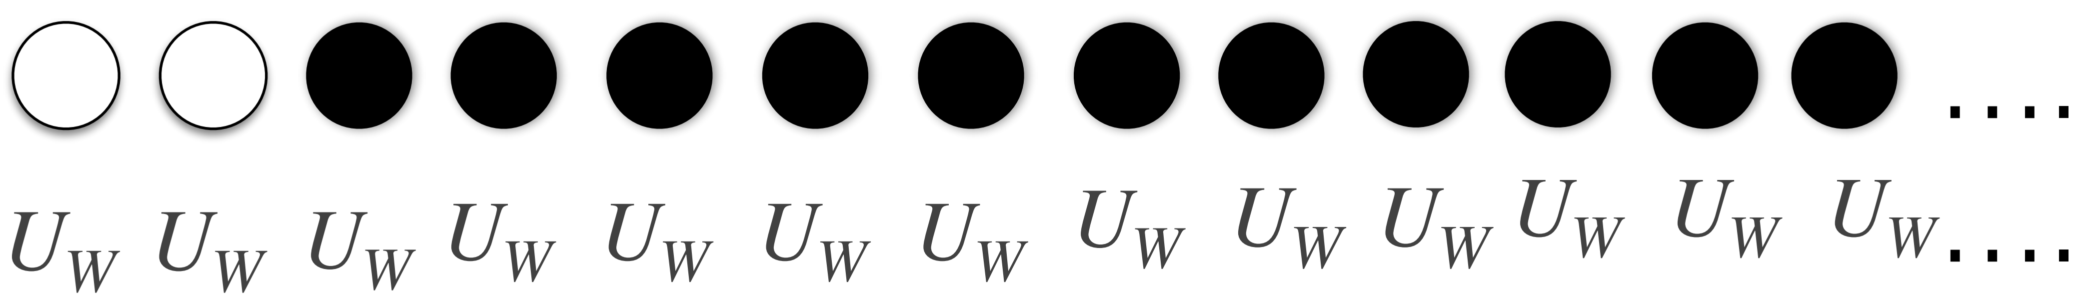

Objection?

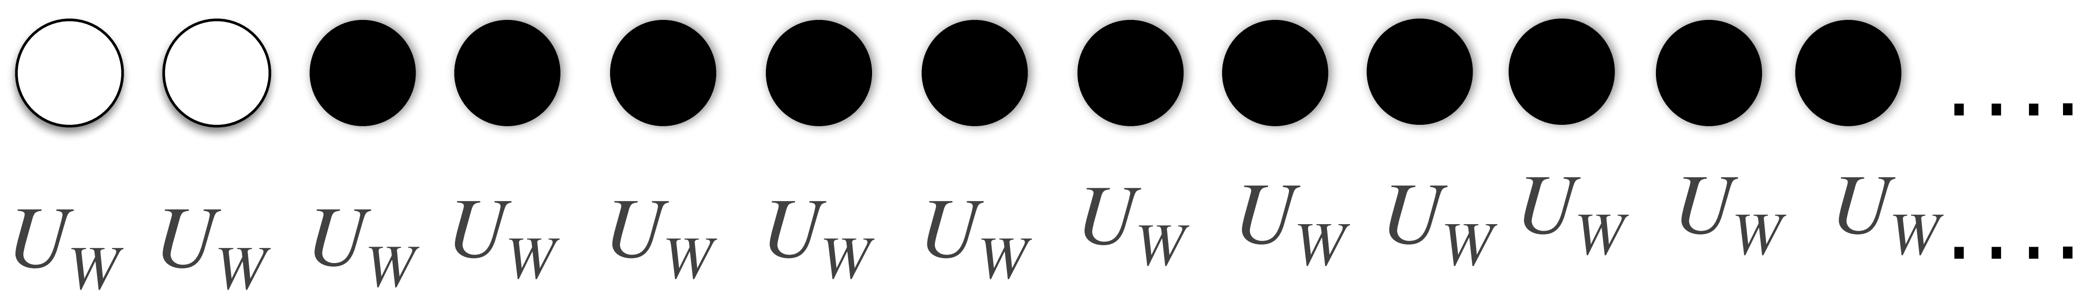

But…what if agents were really **really** smart ?

Answer: Make agents maximally smart

Probabilistic DEL
Model Bayesian reasoning AND
(unbounded) higher-order reasoning.

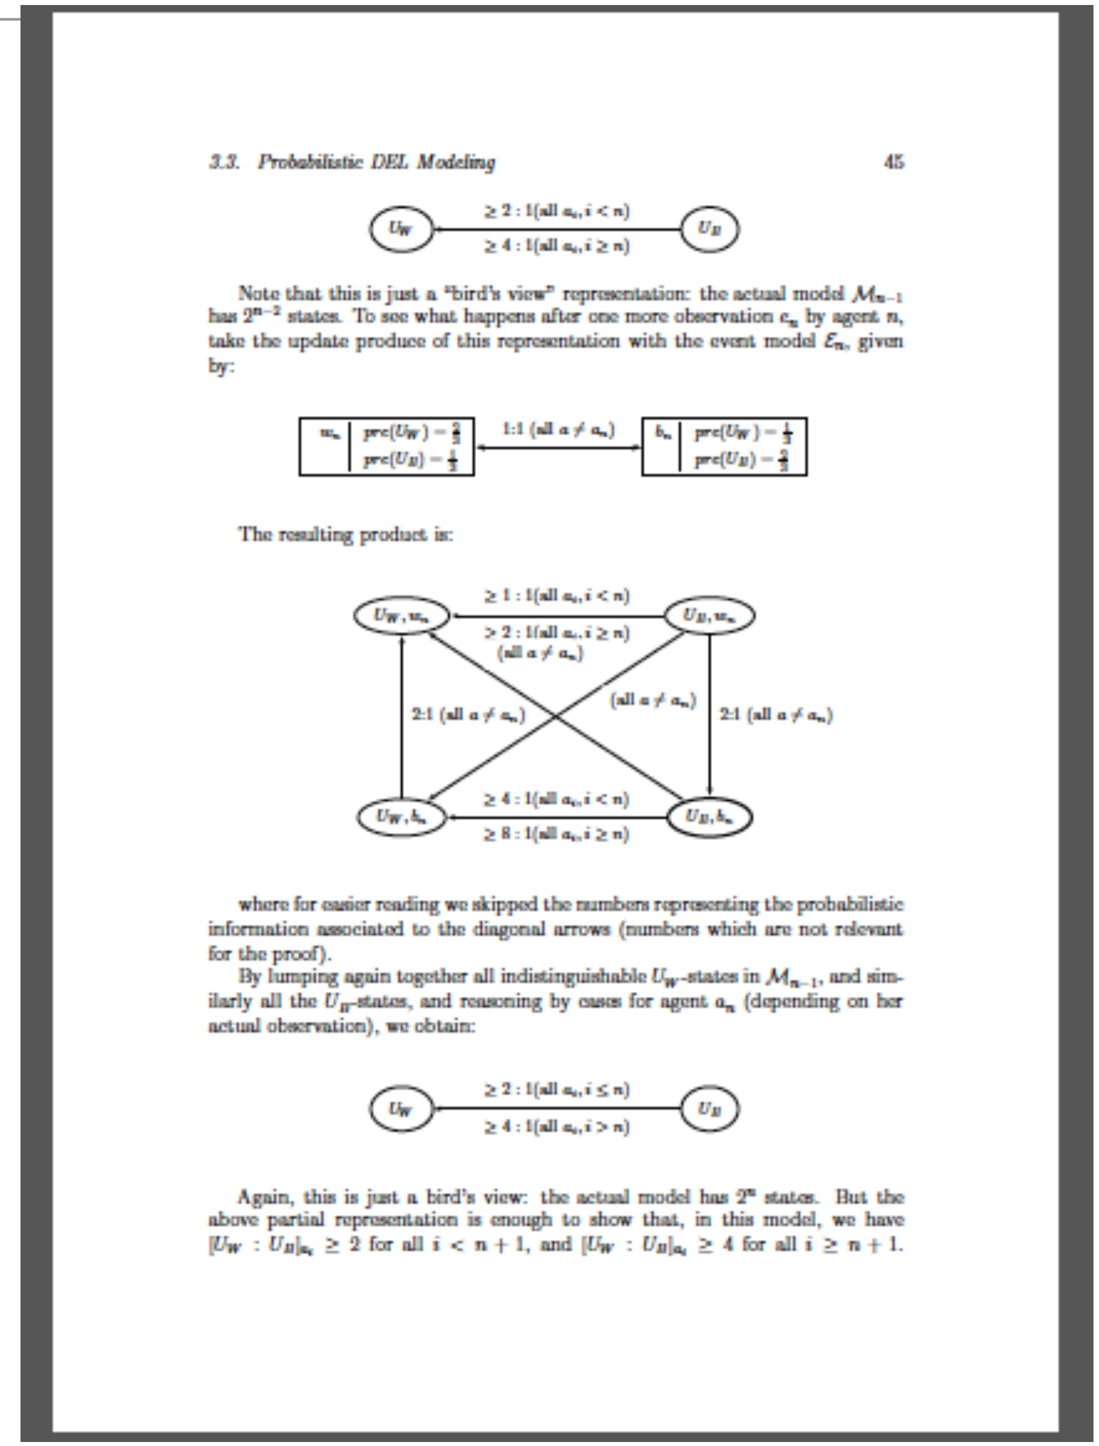

3.3. *Probabilistic DEL Modeling* 45

U_W — $\geq 2 : 1(\text{all } a_i, i < n)$ / $\geq 4 : 1(\text{all } a_i, i \geq n)$ — U_B

Note that this is just a "bird's view" representation: the actual model $\mathcal{M}_{n-1}$ has 2^{n-2} states. To see what happens after one more observation e_n by agent n, take the update produce of this representation with the event model $\mathcal{E}_n$, given by:

$w_n \mid pre(U_W) = \frac{2}{3}, pre(U_B) = \frac{1}{3}$ — $1{:}1\ (\text{all } a \neq a_n)$ — $b_n \mid pre(U_W) = \frac{1}{3}, pre(U_B) = \frac{2}{3}$

The resulting product is:

U_W, w_n — $\geq 1 : 1(\text{all } a_i, i < n)$ / $\geq 2 : 1(\text{all } a_i, i \geq n)$ — U_B, w_n

$(\text{all } a \neq a_n)$; $2{:}1\ (\text{all } a \neq a_n)$; $(\text{all } a \neq a_n)$; $2{:}1\ (\text{all } a \neq a_n)$

U_W, b_n — $\geq 4 : 1(\text{all } a_i, i < n)$ / $\geq 8 : 1(\text{all } a_i, i \geq n)$ — U_B, b_n

where for easier reading we skipped the numbers representing the probabilistic information associated to the diagonal arrows (numbers which are not relevant for the proof).

By lumping again together all indistinguishable U_W-states in $\mathcal{M}_{n-1}$, and similarly all the U_B-states, and reasoning by cases for agent a_n (depending on her actual observation), we obtain:

U_W — $\geq 2 : 1(\text{all } a_i, i \leq n)$ / $\geq 4 : 1(\text{all } a_i, i > n)$ — U_B

Again, this is just a bird's view: the actual model has 2^n states. But the above partial representation is enough to show that, in this model, we have $[U_W : U_B]_{a_i} \geq 2$ for all $i < n+1$, and $[U_W : U_B]_{a_i} \geq 4$ for all $i \geq n+1$.

Logical Models of Informational Cascades (pdf), Alexandru Baltag, Zoé Christoff, Jens Ulrik Hansen and Sonja Smets, in J. van Benthem and F. Liu (Eds.): Logic across the University: Foundations and Applications, — Proceedings of the Tsinghua Logic Conference, Beijing, 14-16 October 2013, Studies in Logic, Volume 47, pp.405-432, College Publications, London, (2013).

Initial situation model

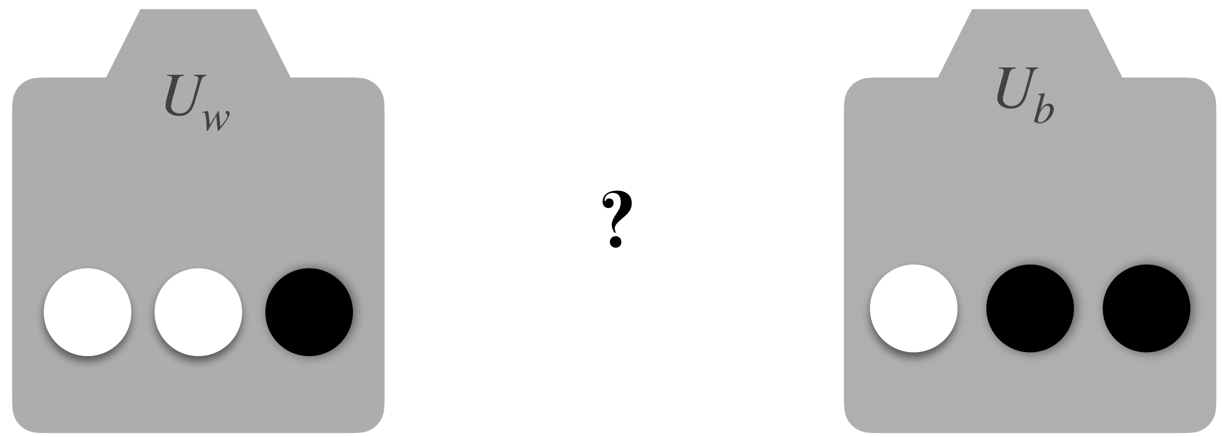

$$P(U_W) = P(U_B) = \frac{1}{2}$$

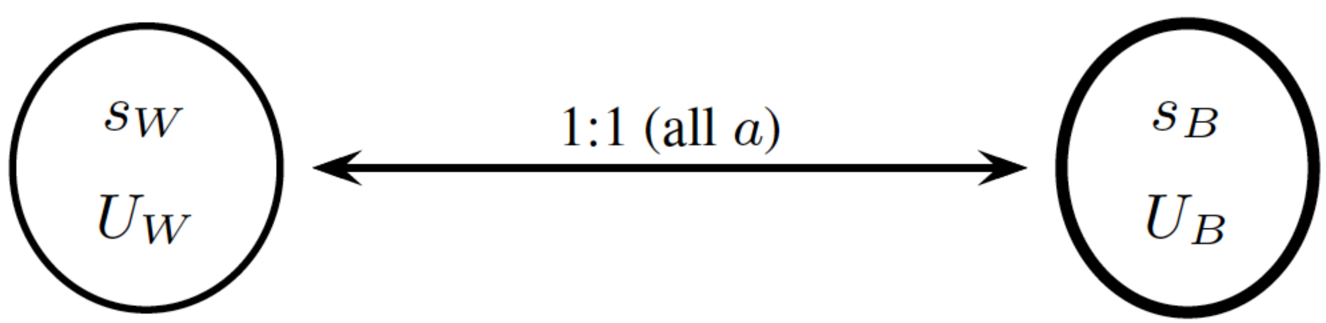

(first) marble observation: action model

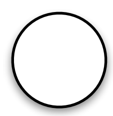

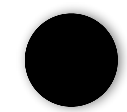

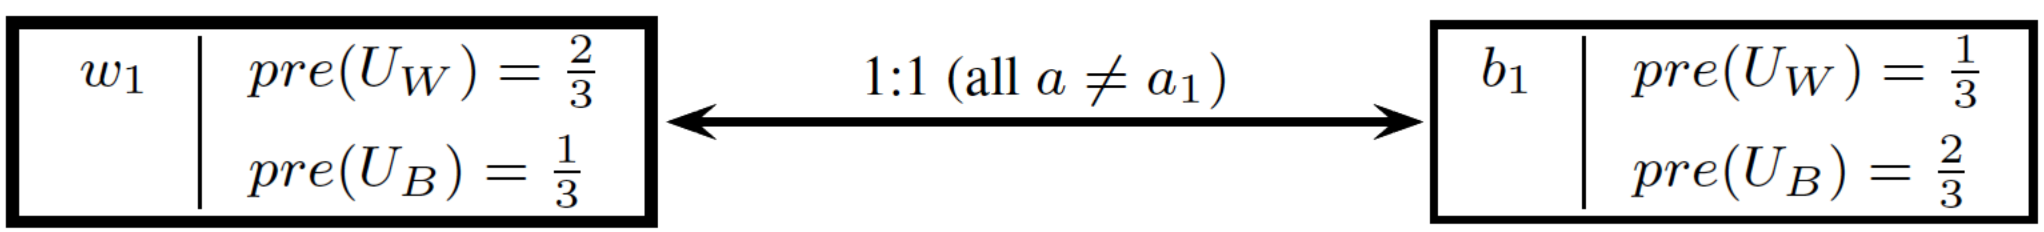

After the first observation

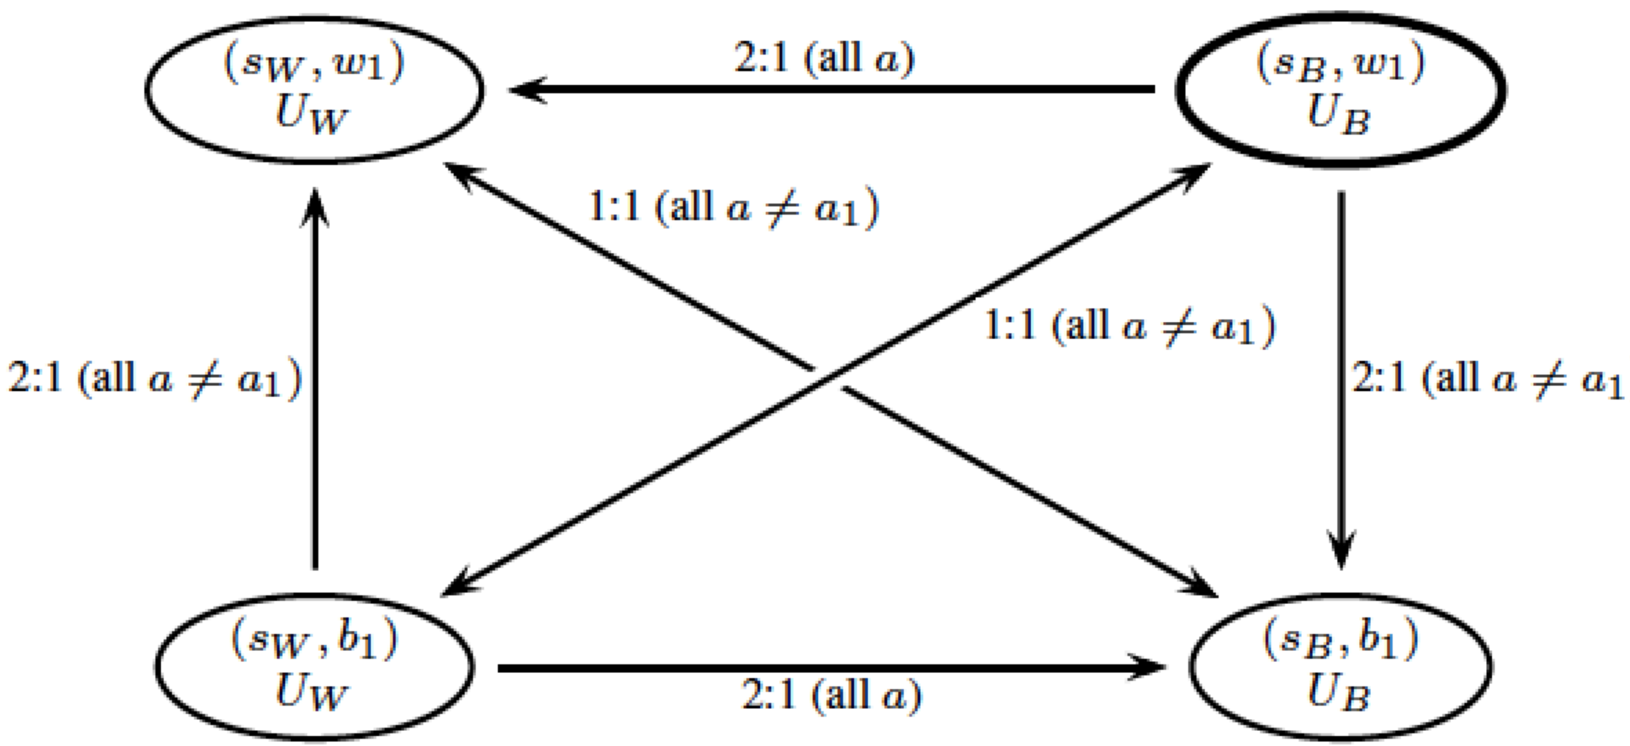

After the first agent announces her guess

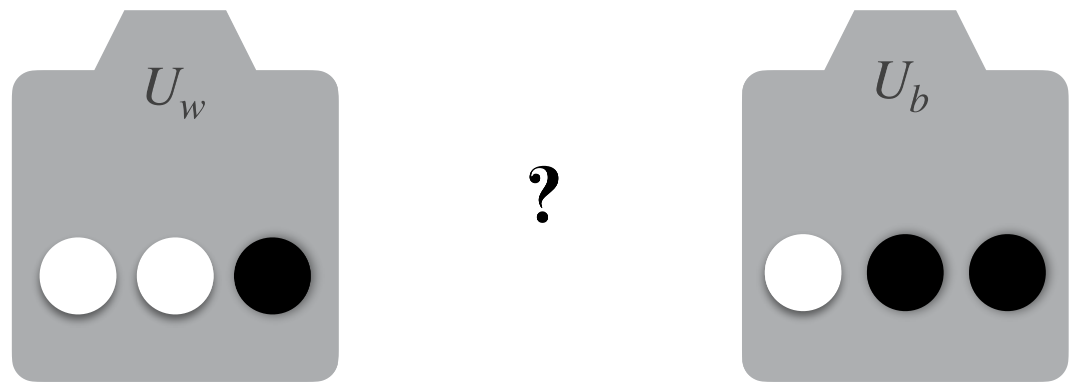

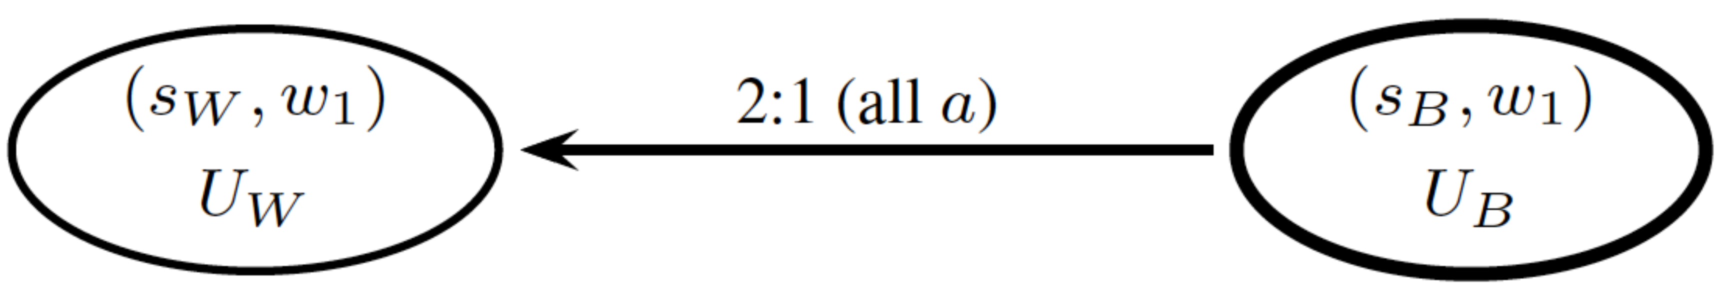

After the second agent announces her guess

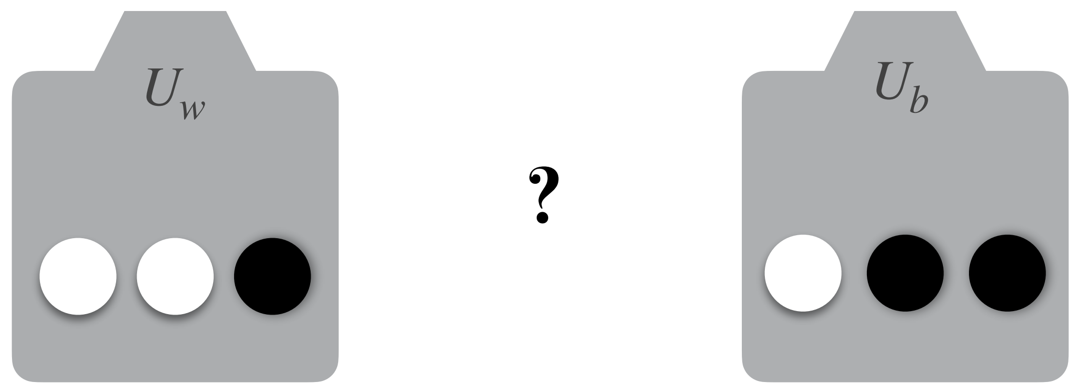

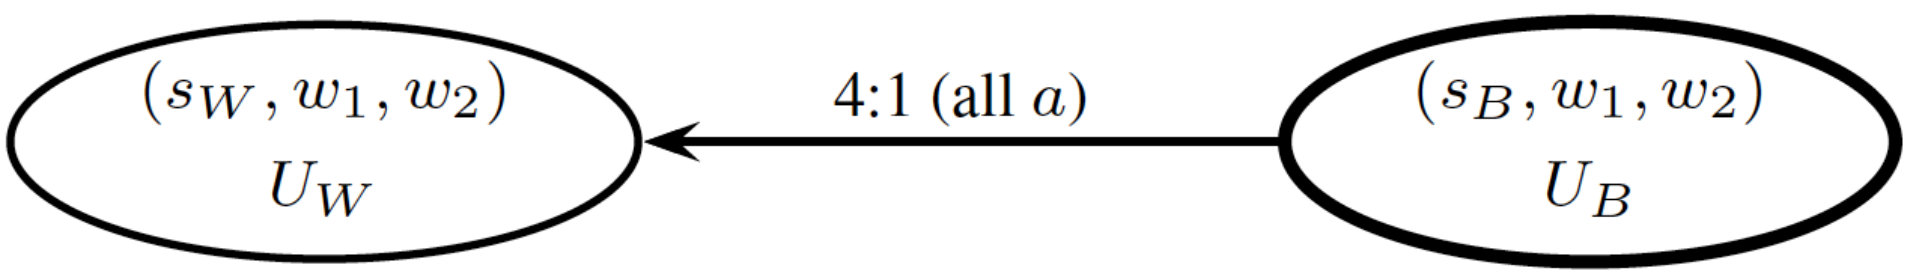

After the third observation

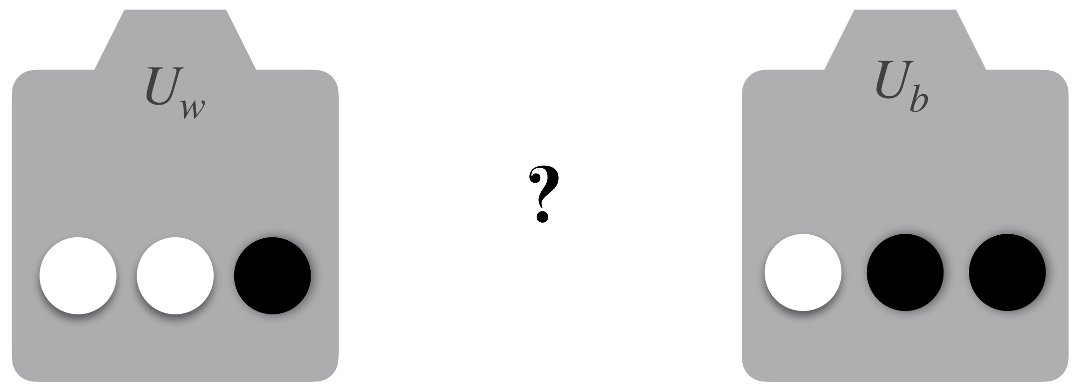

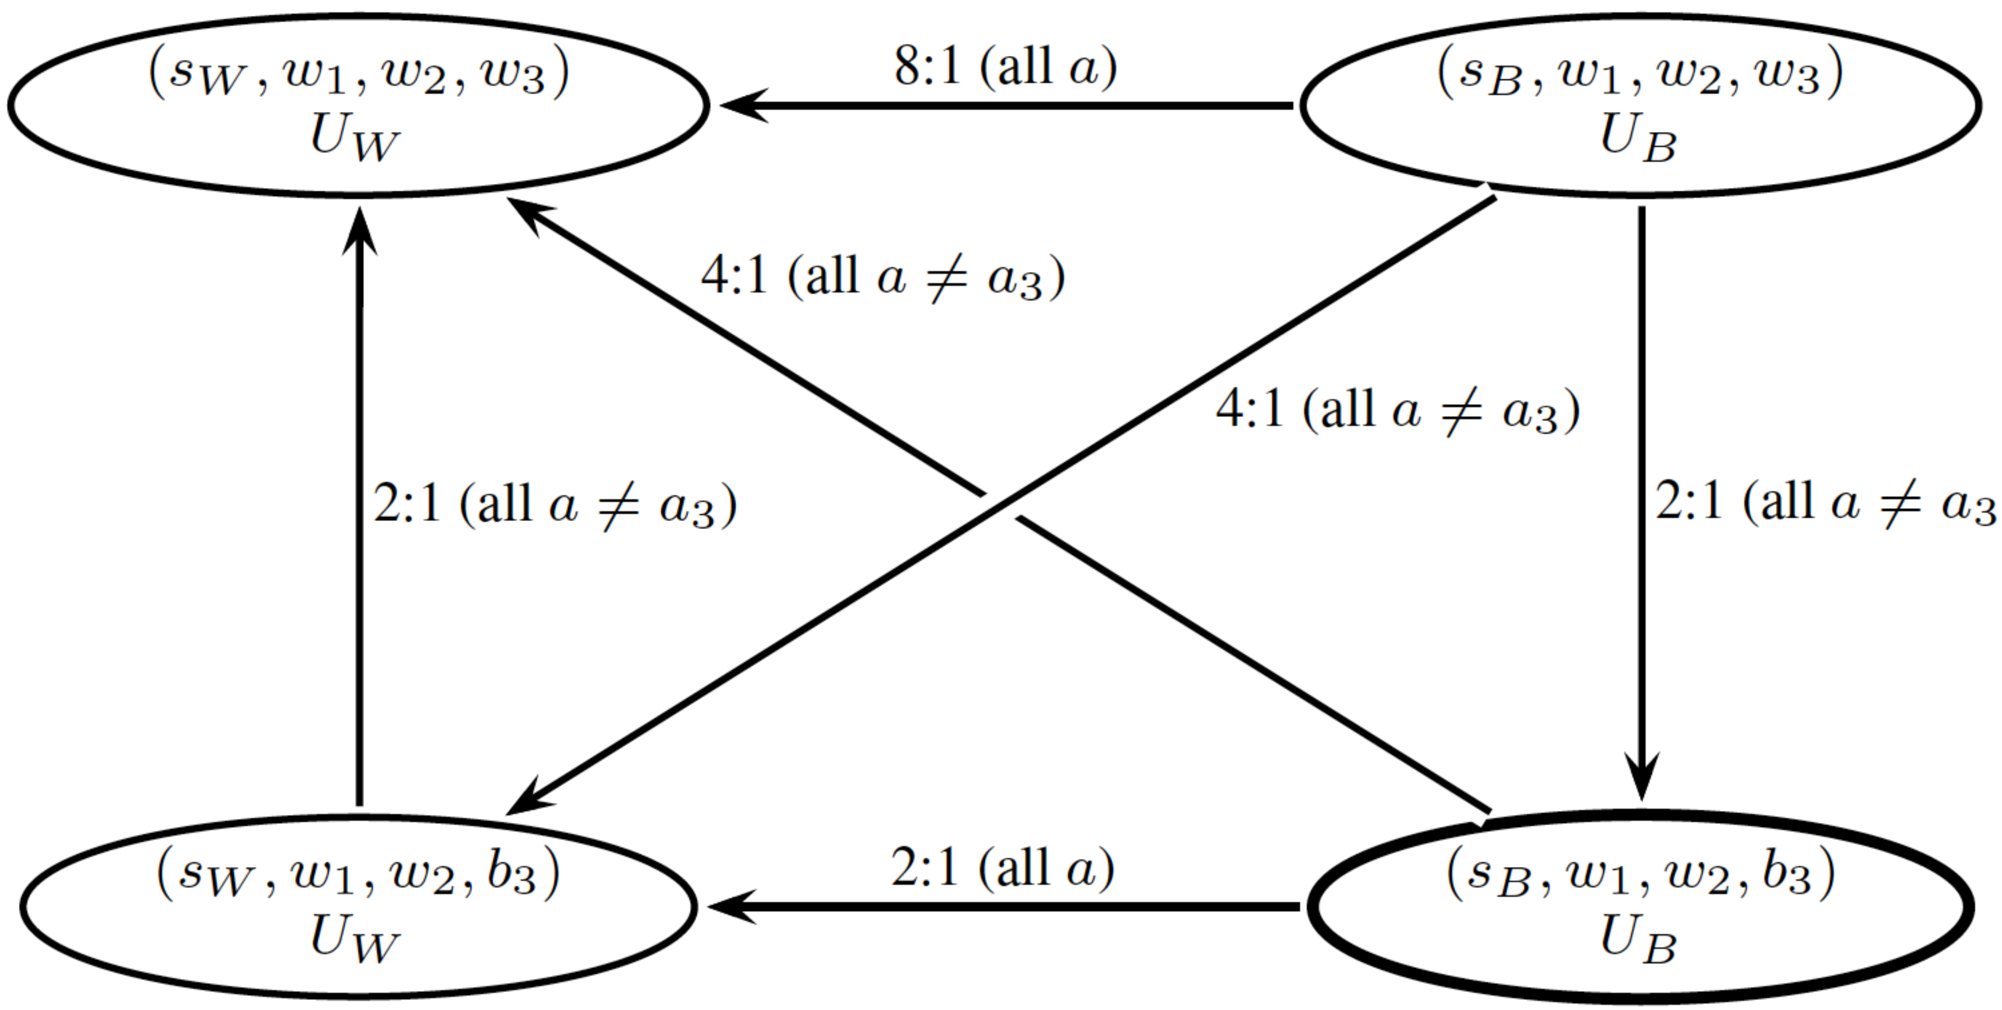

$[U_W : U_B]_{a_3} > 1$ is now *common knowledge*

Conclusion: Still perfectly rational!

Captures the inescapability of rational cascades, even for **very** smart agents.

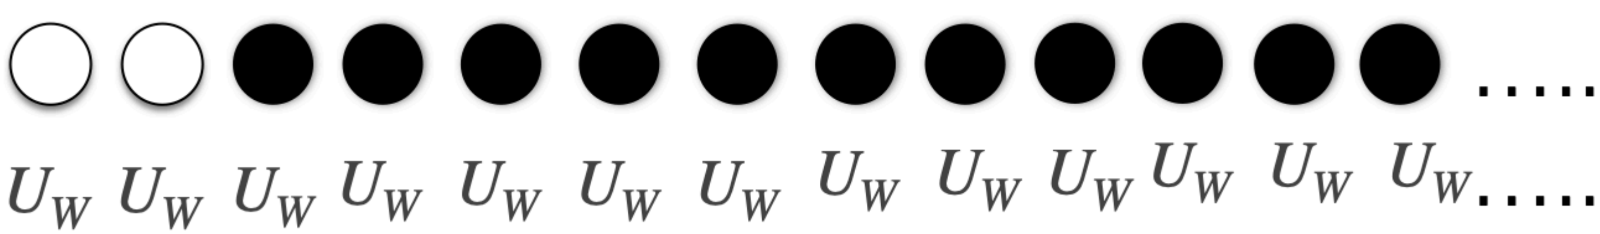

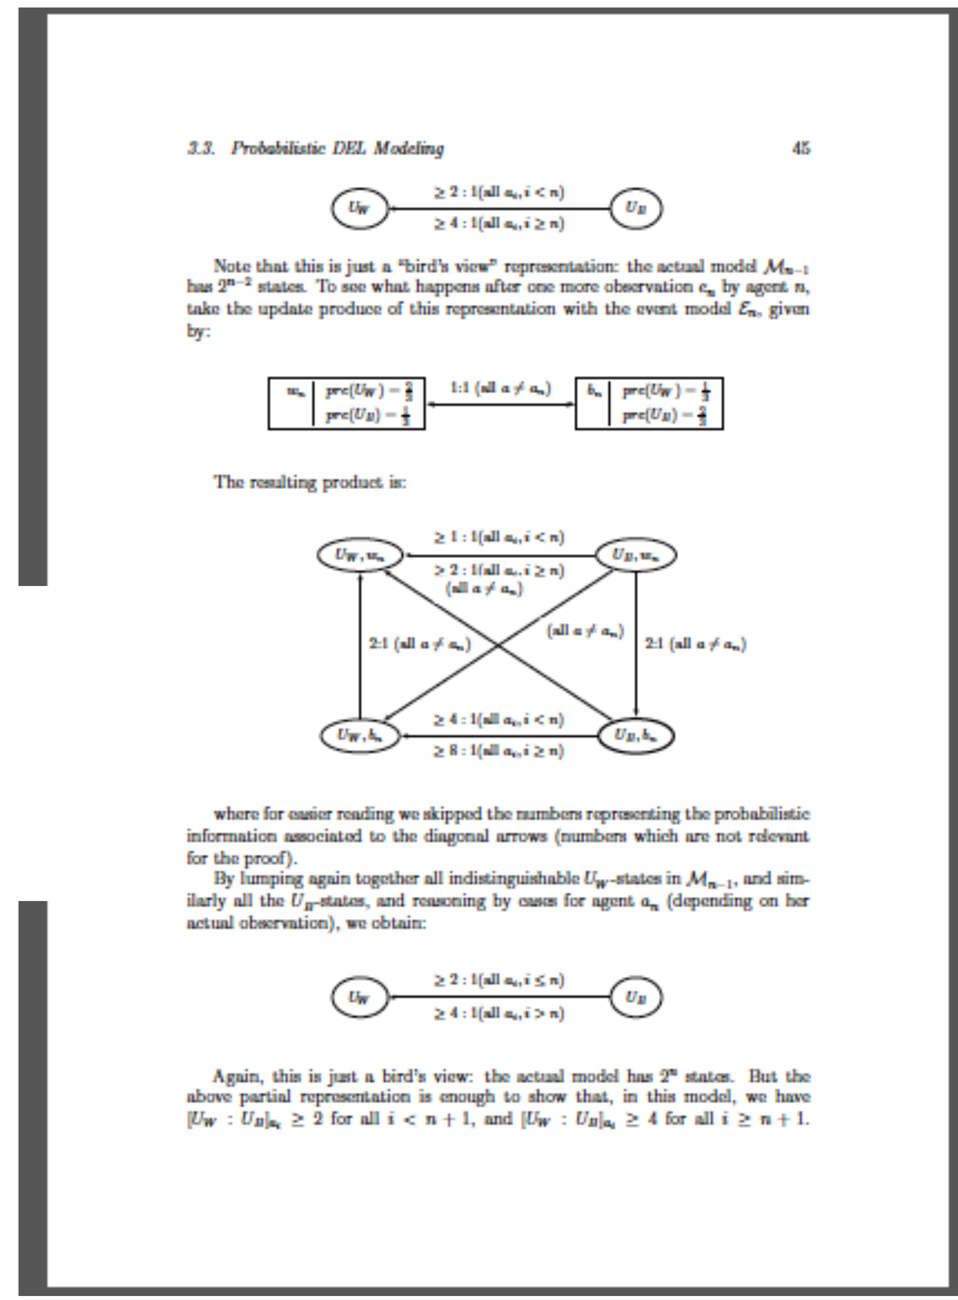

3.3. Probabilistic DEL Modeling 45

U_W — $\geq 2:1(\text{all } a_i, i < n)$ / $\geq 4:1(\text{all } a_i, i \geq n)$ — U_B

Note that this is just a "bird's view" representation: the actual model $\mathcal{M}_{n-1}$ has 2^{n-2} states. To see what happens after one more observation e_n by agent n, take the update produce of this representation with the event model $\mathcal{E}_n$, given by:

w_n | $pre(U_W) = \frac{2}{3}$, $pre(U_B) = \frac{1}{3}$ — 1:1 (all $a \neq a_n$) — b_n | $pre(U_W) = \frac{1}{3}$, $pre(U_B) = \frac{2}{3}$

The resulting product is:

U_W, w_n — $\geq 1:1(\text{all } a_i, i < n)$ / $\geq 2:1(\text{all } a_i, i \geq n)$ — U_B, w_n; (all $a \neq a_n$); 2:1 (all $a \neq a_n$); (all $a \neq a_n$); 2:1 (all $a \neq a_n$); U_W, b_n — $\geq 4:1(\text{all } a_i, i < n)$ / $\geq 8:1(\text{all } a_i, i \geq n)$ — U_B, b_n

where for easier reading we skipped the numbers representing the probabilistic information associated to the diagonal arrows (numbers which are not relevant for the proof).

By lumping again together all indistinguishable U_W-states in $\mathcal{M}_{n-1}$, and similarly all the U_B-states, and reasoning by cases for agent a_n (depending on her actual observation), we obtain:

U_W — $\geq 2:1(\text{all } a_i, i \leq n)$ / $\geq 4:1(\text{all } a_i, i > n)$ — U_B

Again, this is just a bird's view: the actual model has 2^n states. But the above partial representation is enough to show that, in this model, we have $[U_W : U_B]_{a_i} \geq 2$ for all $i < n+1$, and $[U_W : U_B]_{a_i} \geq 4$ for all $i \geq n+1$.

What does our DEL model show ?

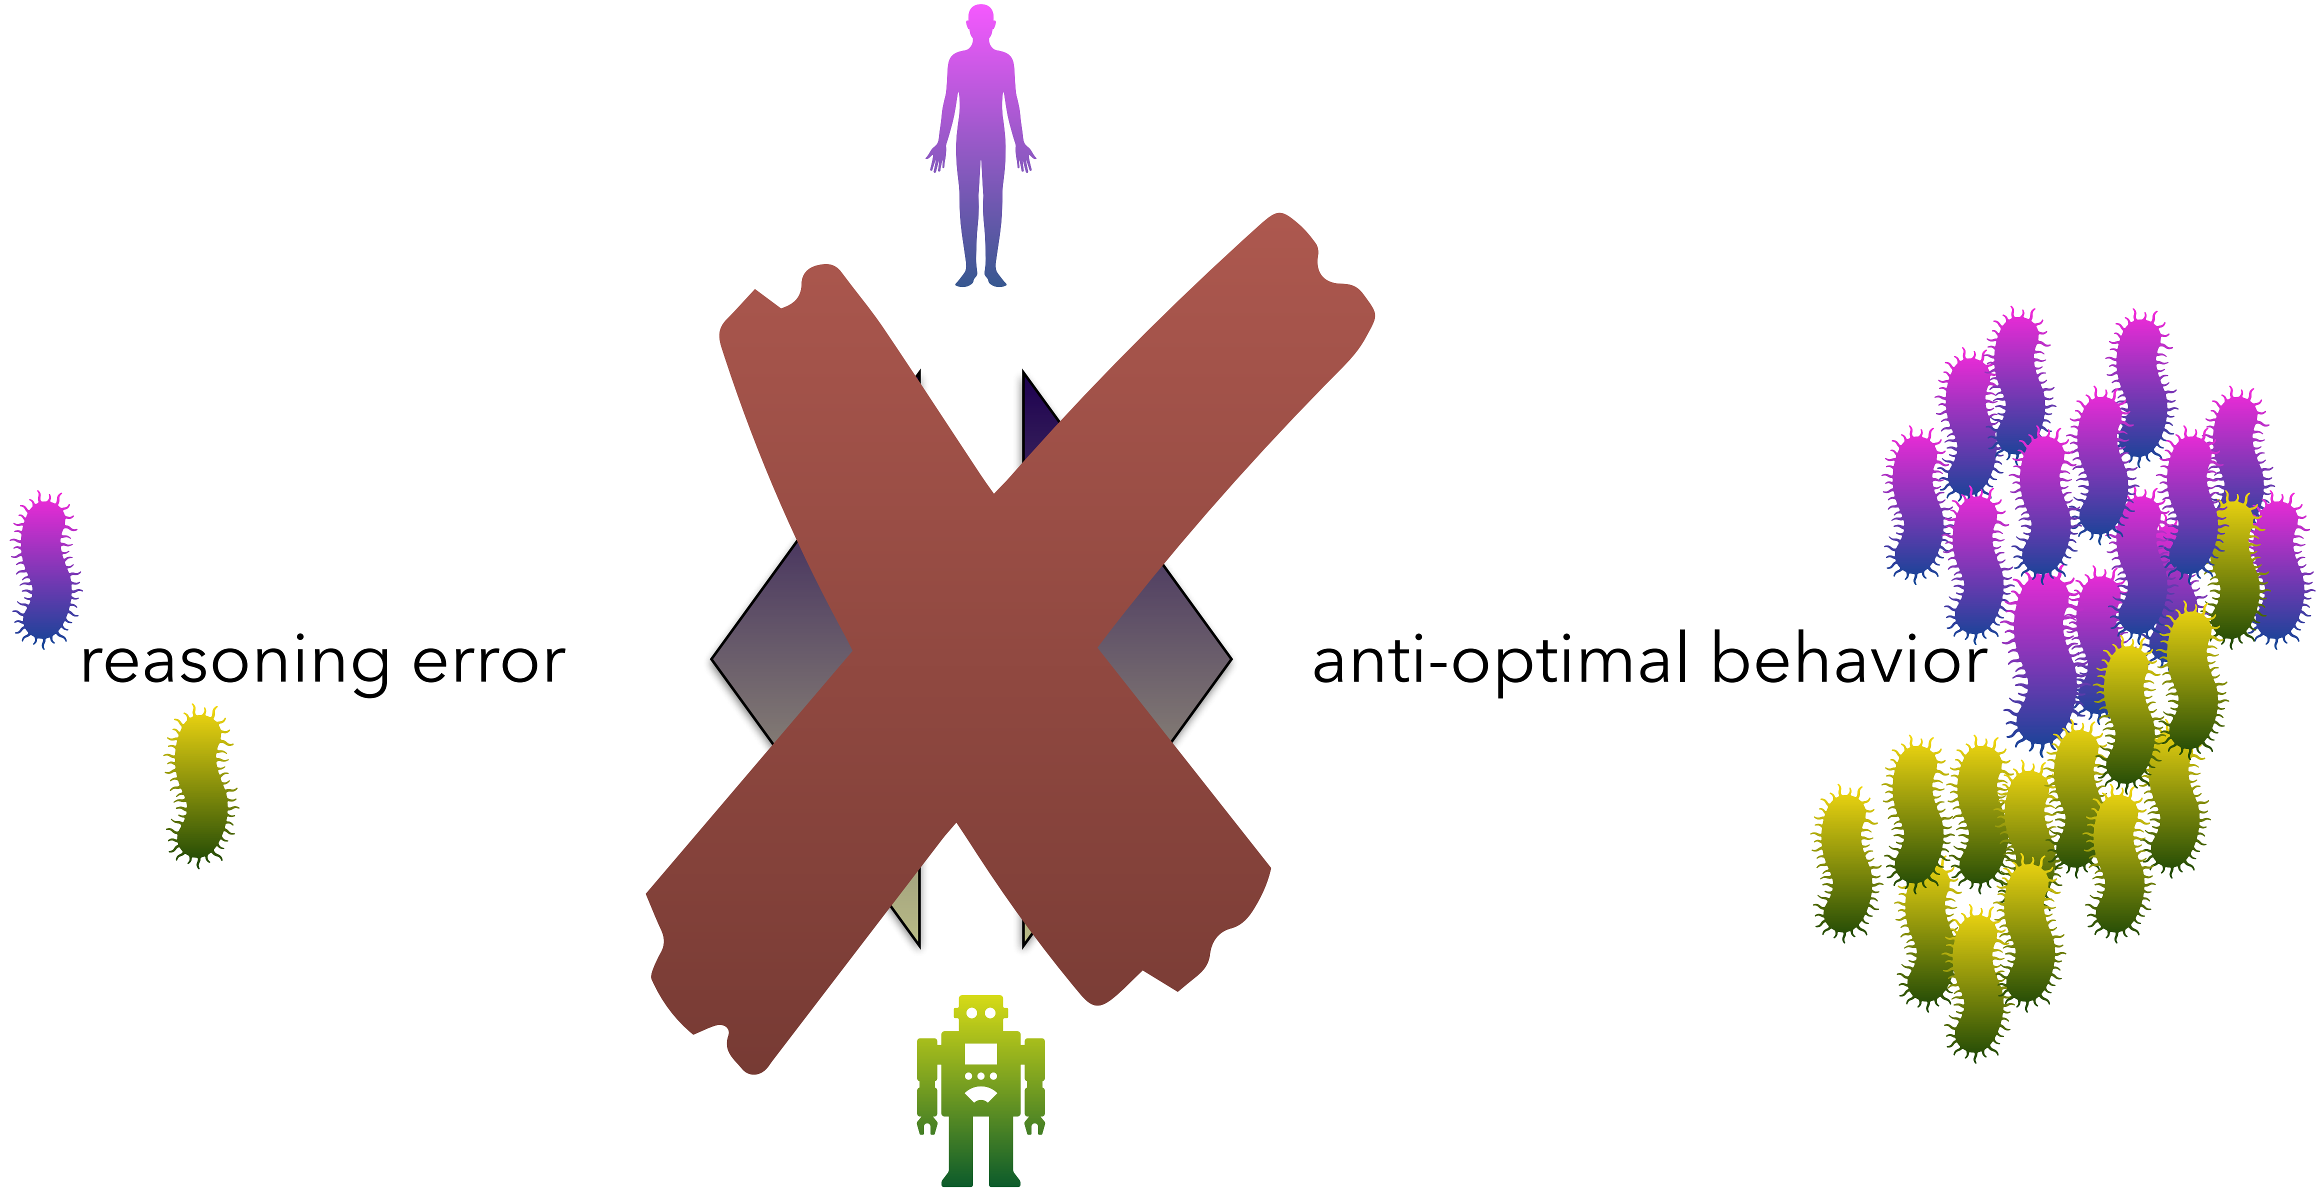

III. Social Influence in networks: simple diffusion models

Section slides adapted from *Social Networks for Logicians*, course at ESSLLI 2018, 30th European Summer School in Logic, Language and Information, Sofia (with Pavel Naumov)

1) Threshold Models

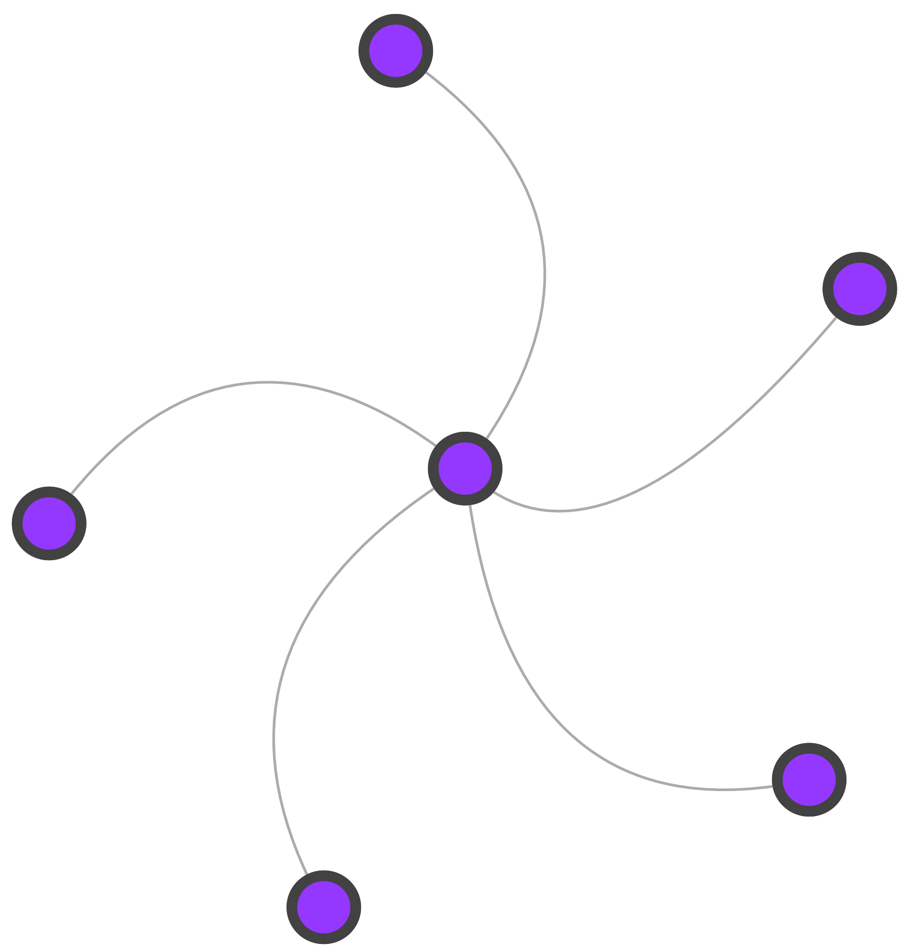

$\theta = 0.5$

θ is a (uniform) threshold value

Rule: agent adopts (a new color) if the ratio of her neighbors who already display it is at least θ

“Complete cascade”

$\theta = 0.5$

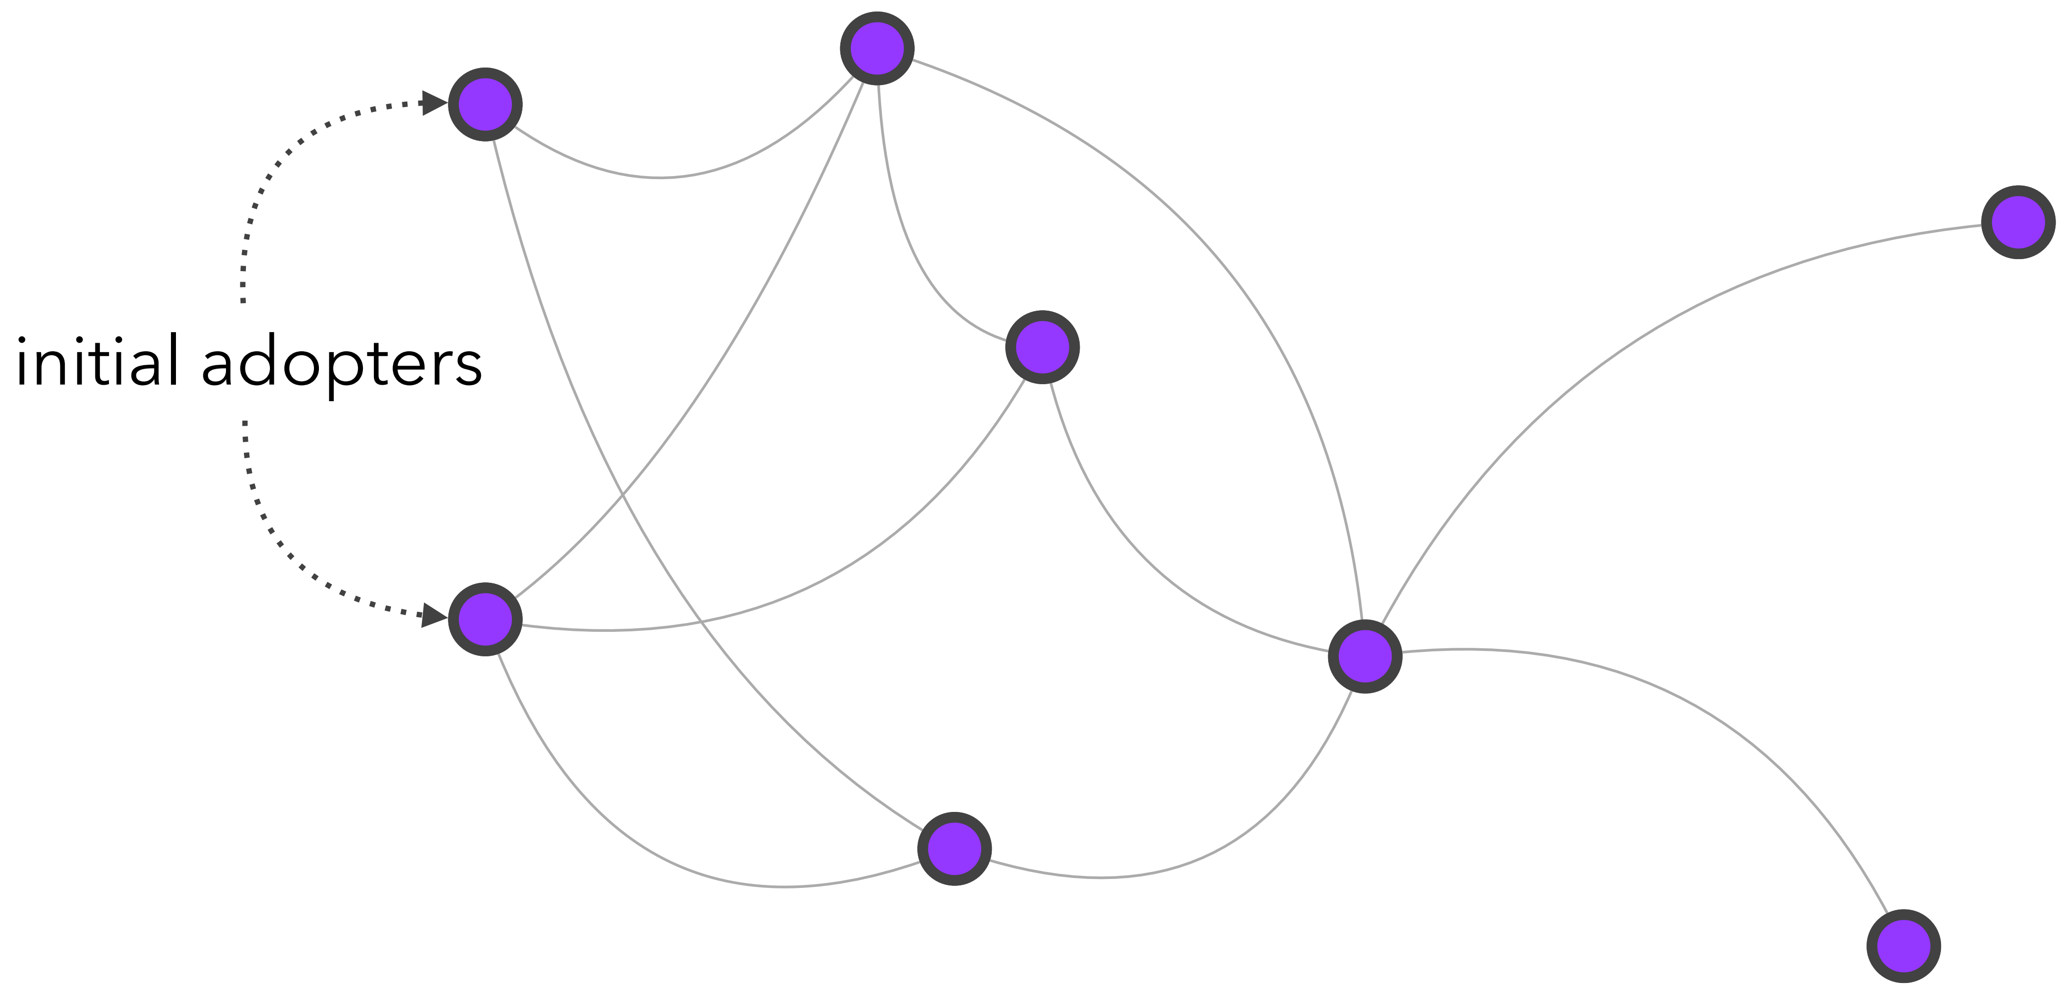

“Cluster-Cascade” Theorem

Theorem (folklore)
All nodes will eventually be infected if and only if among nodes who are not infected there is no cluster of density higher than 1-θ.

2) DeGroot Model

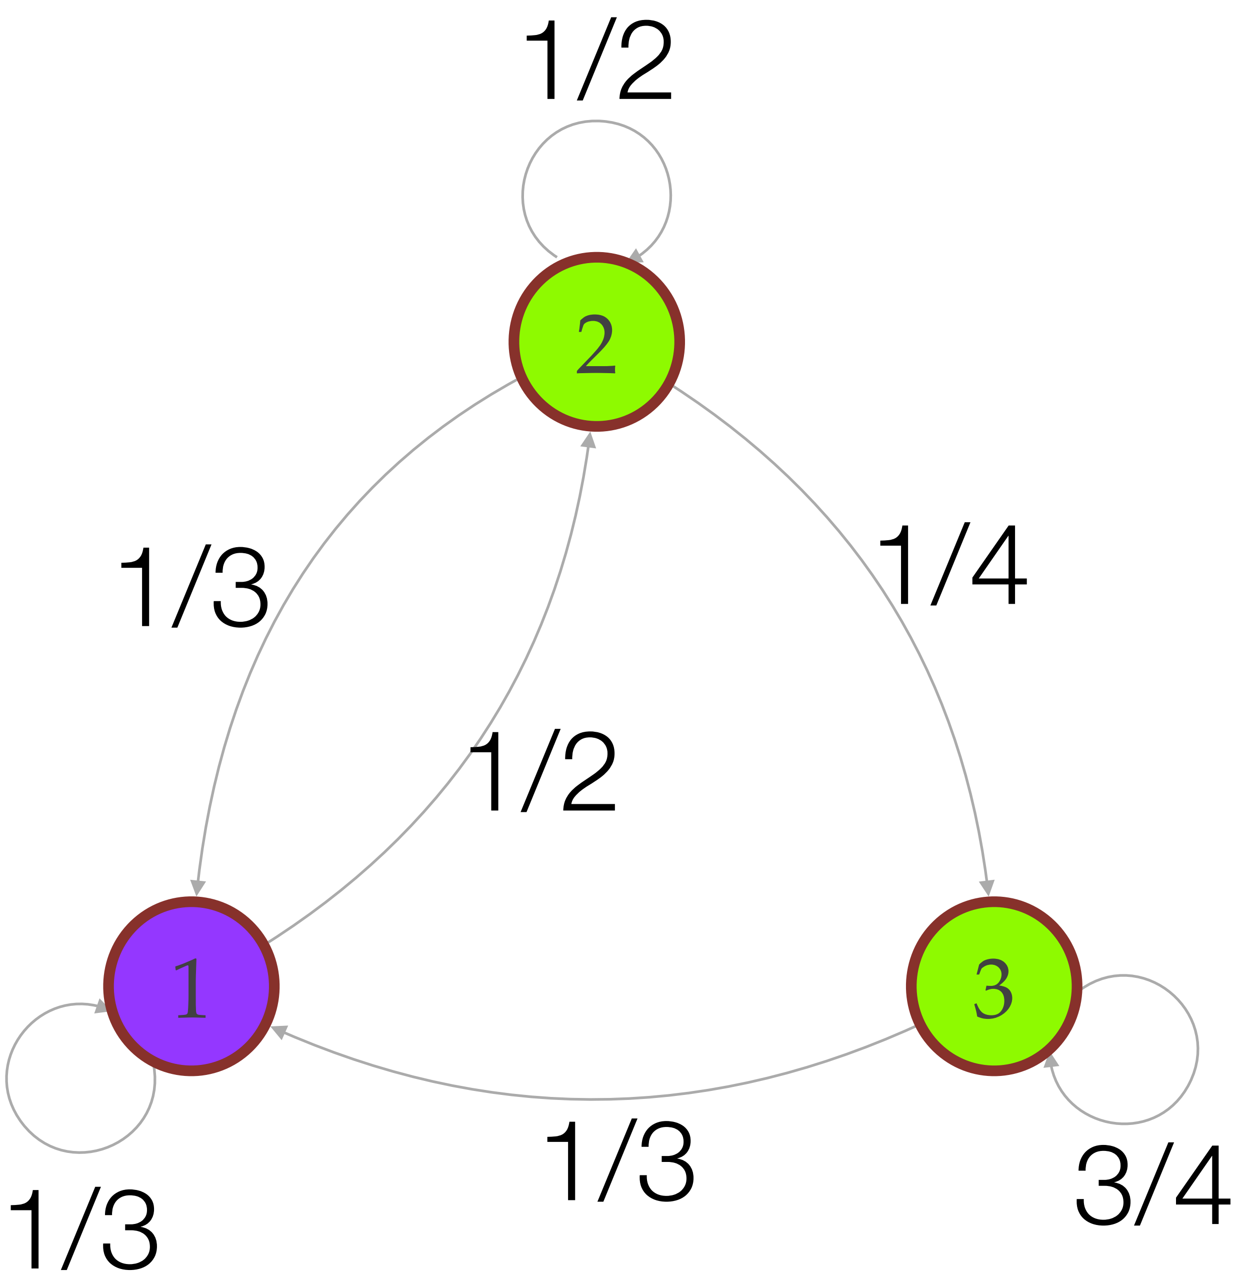

$$T = \begin{pmatrix} 1/3 & 1/3 & 1/3 \\ 1/2 & 1/2 & 0 \\ 0 & 1/4 & 3/4 \end{pmatrix}$$

$$\bar{p}_1 = T\bar{p}_0 \begin{pmatrix} 1/3 & 1/3 & 1/3 \\ 1/2 & 1/2 & 0 \\ 0 & 1/4 & 3/4 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 1/3 \\ 1/2 \\ 0 \end{pmatrix}$$

$$\bar{p}_2 = T\bar{p}_1 \begin{pmatrix} 1/3 & 1/3 & 1/3 \\ 1/2 & 1/2 & 0 \\ 0 & 1/4 & 3/4 \end{pmatrix} \begin{pmatrix} 1/3 \\ 1/2 \\ 0 \end{pmatrix} = \begin{pmatrix} 5/18 \\ 5/12 \\ 1/8 \end{pmatrix}$$

$$\ldots$$

$$\bar{p}_n = T\bar{p}_{n-1} = T^n\bar{p}_0$$

Convergence in the DeGroot Model

A DeGroot model converges if and only if in each strongly connected closed set, the greatest common divisor of all cycles is equal to 1.

See Theorem 8.1 in Matthew O. Jackson, *Social and Economic Networks*, Princeton University Press, 2008.

What type of results? (again)

Type of results network analysis is typically interested in (given one class of networks, and one class of rules):

- which diffusion states are reachable from which?
- which diffusion processes stabilize?
- what graphs guarantee stabilization?

Can logic help with any of these?

IV. Logical models of influence in networks

The idea

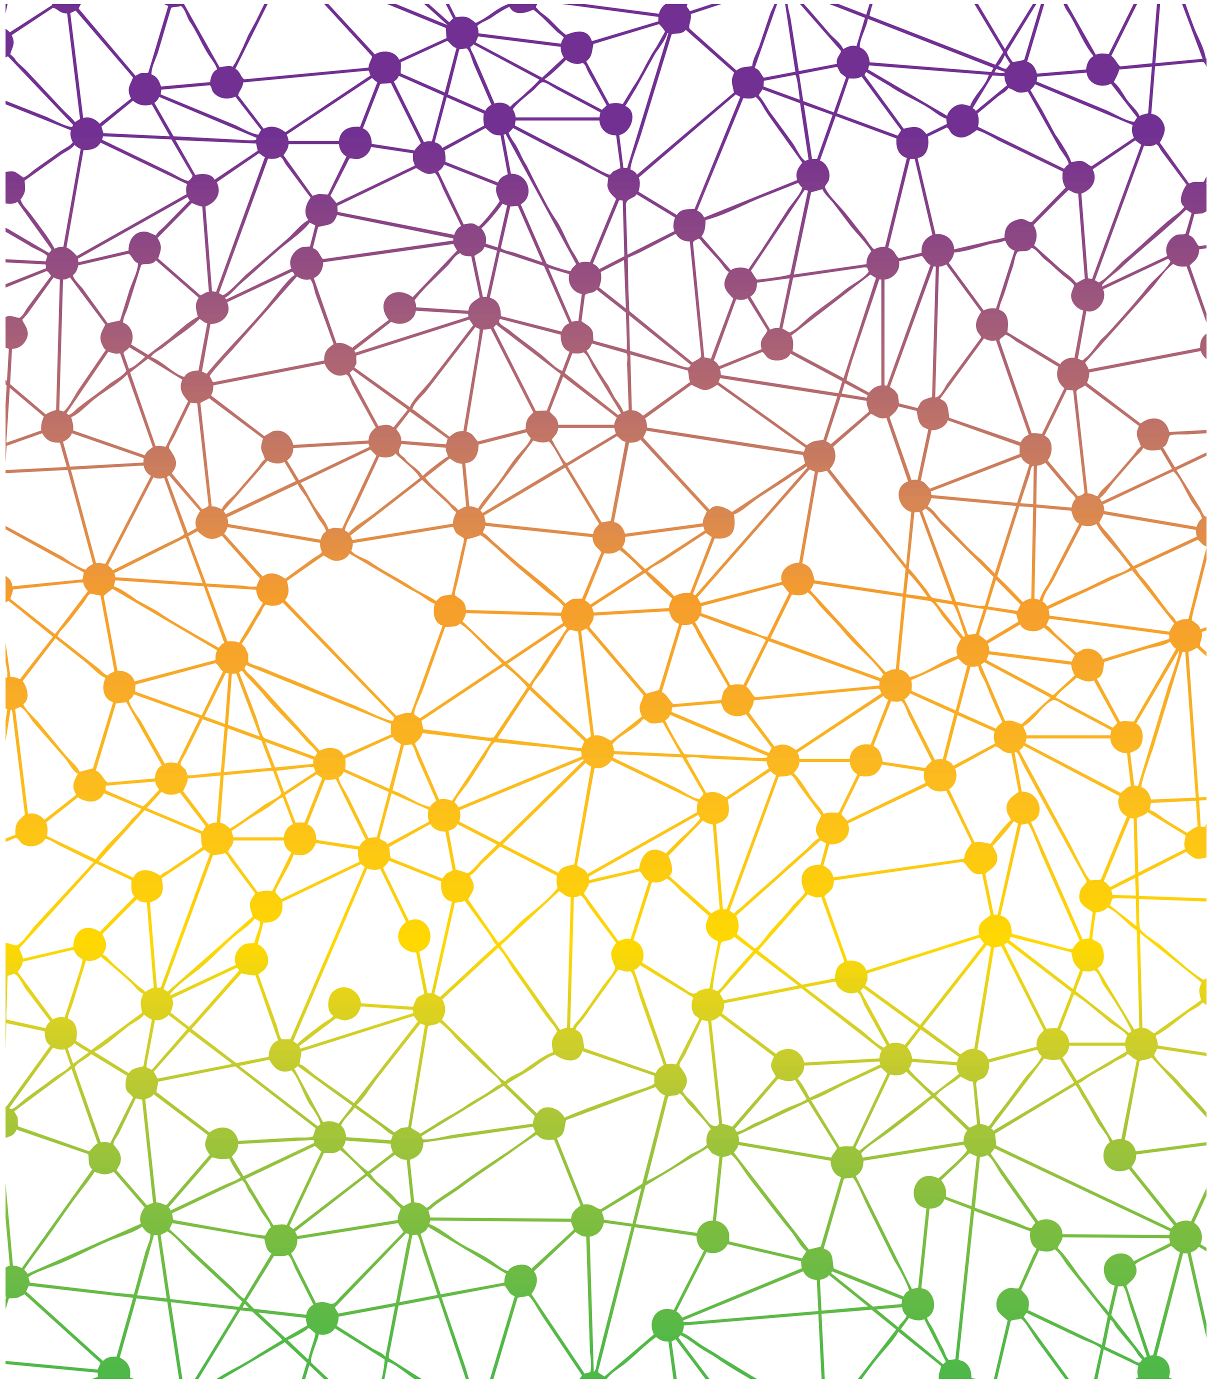

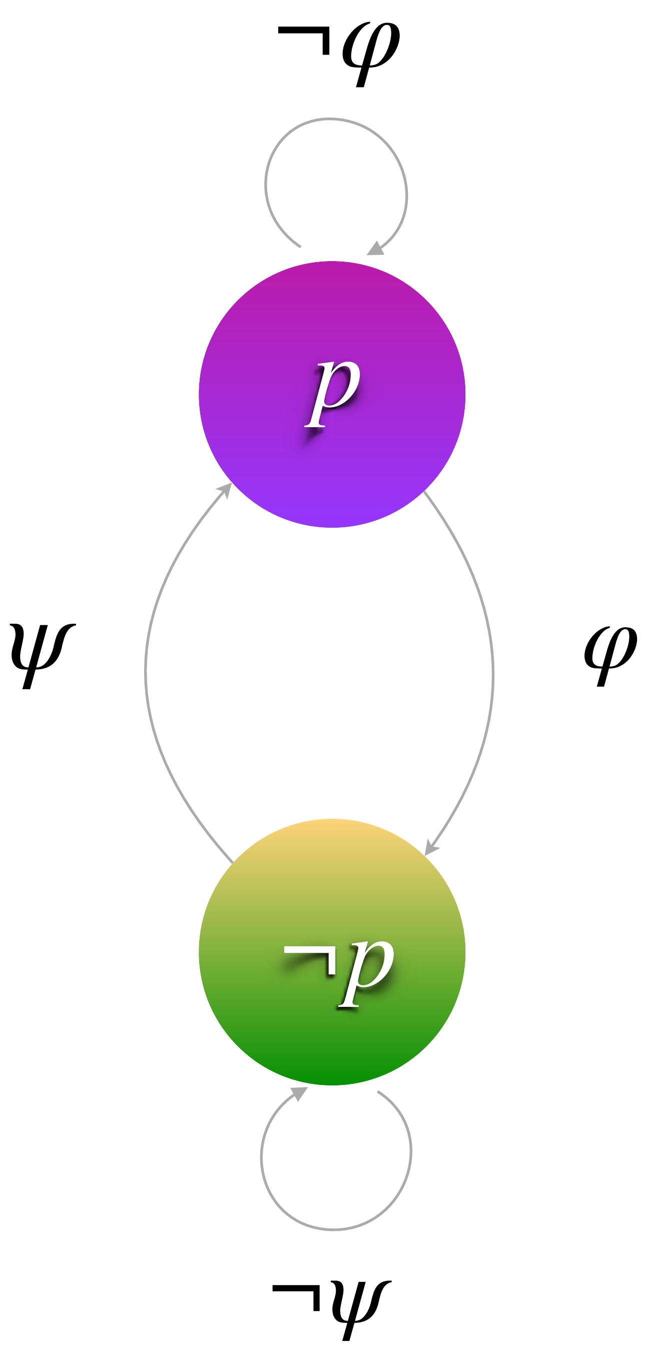

Minimal example: 2 states/colors

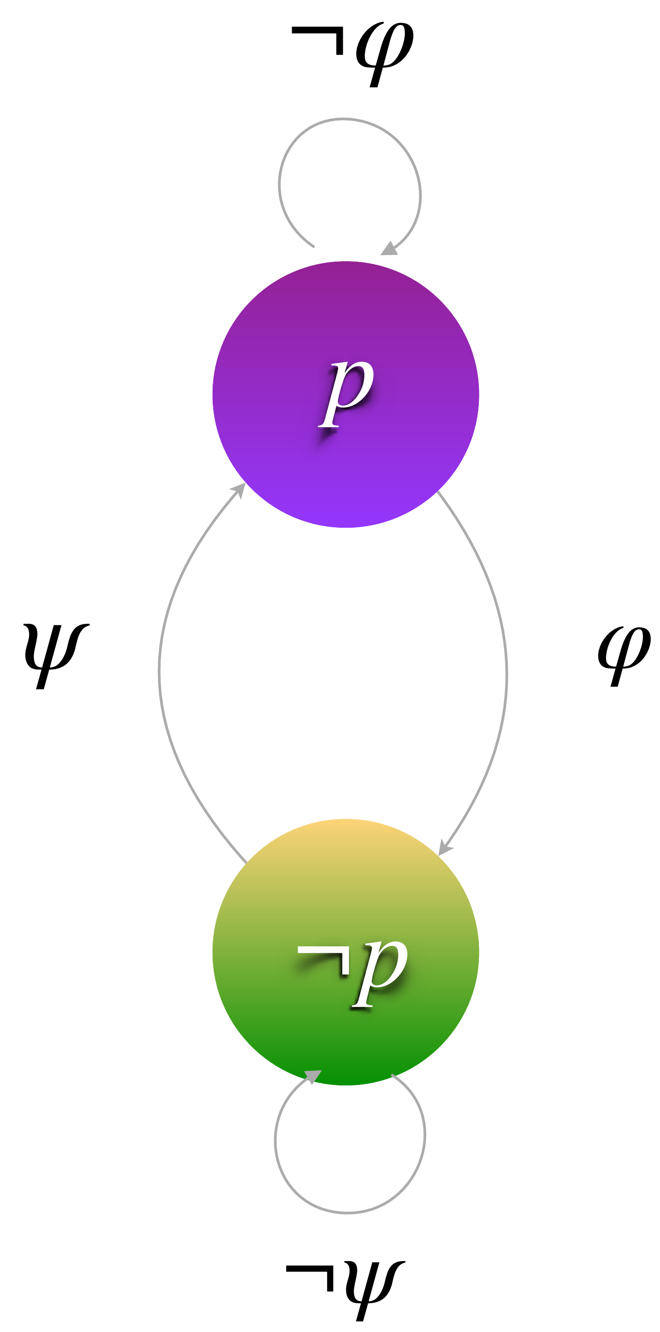

Minimal example: unanimity

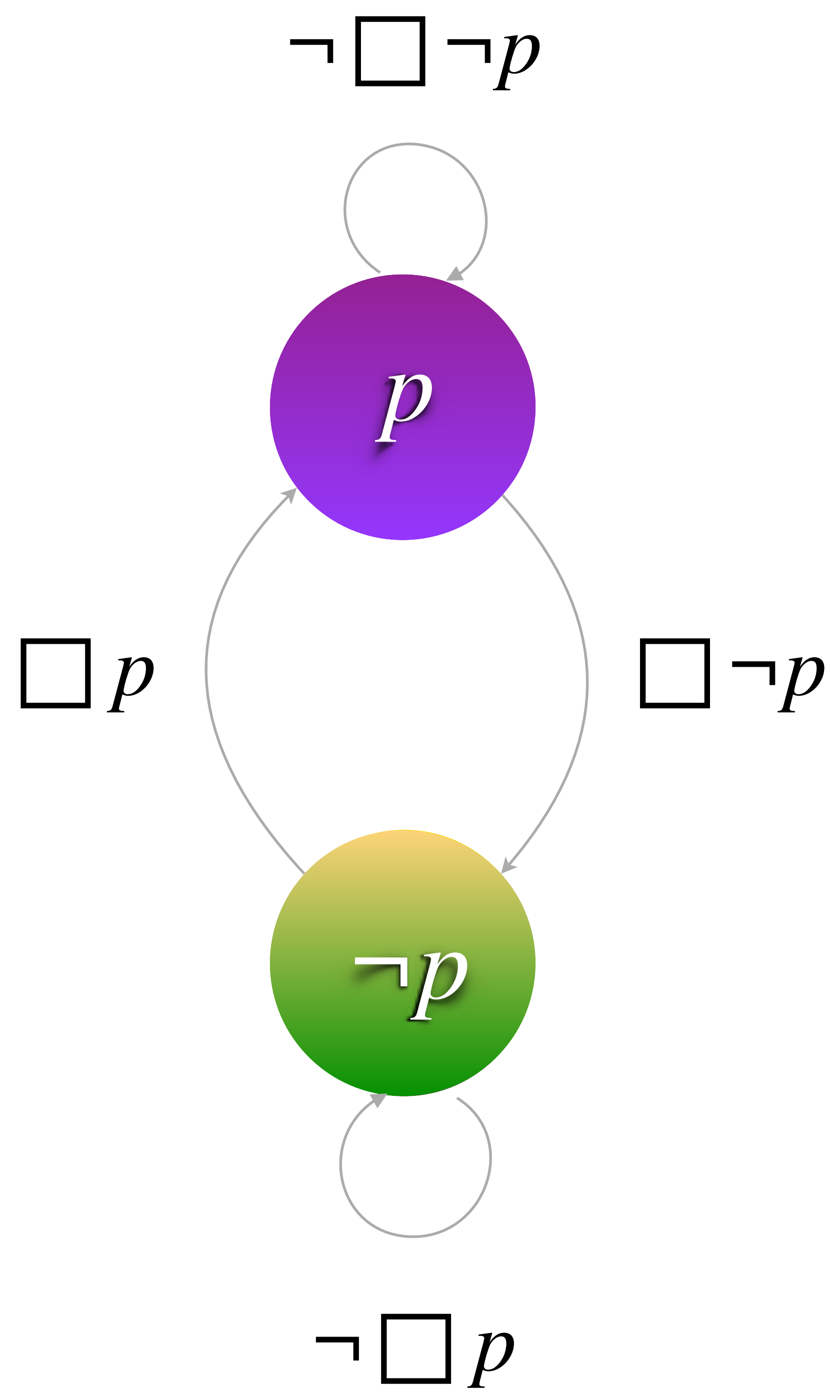

Diffusion dynamics under unanimity

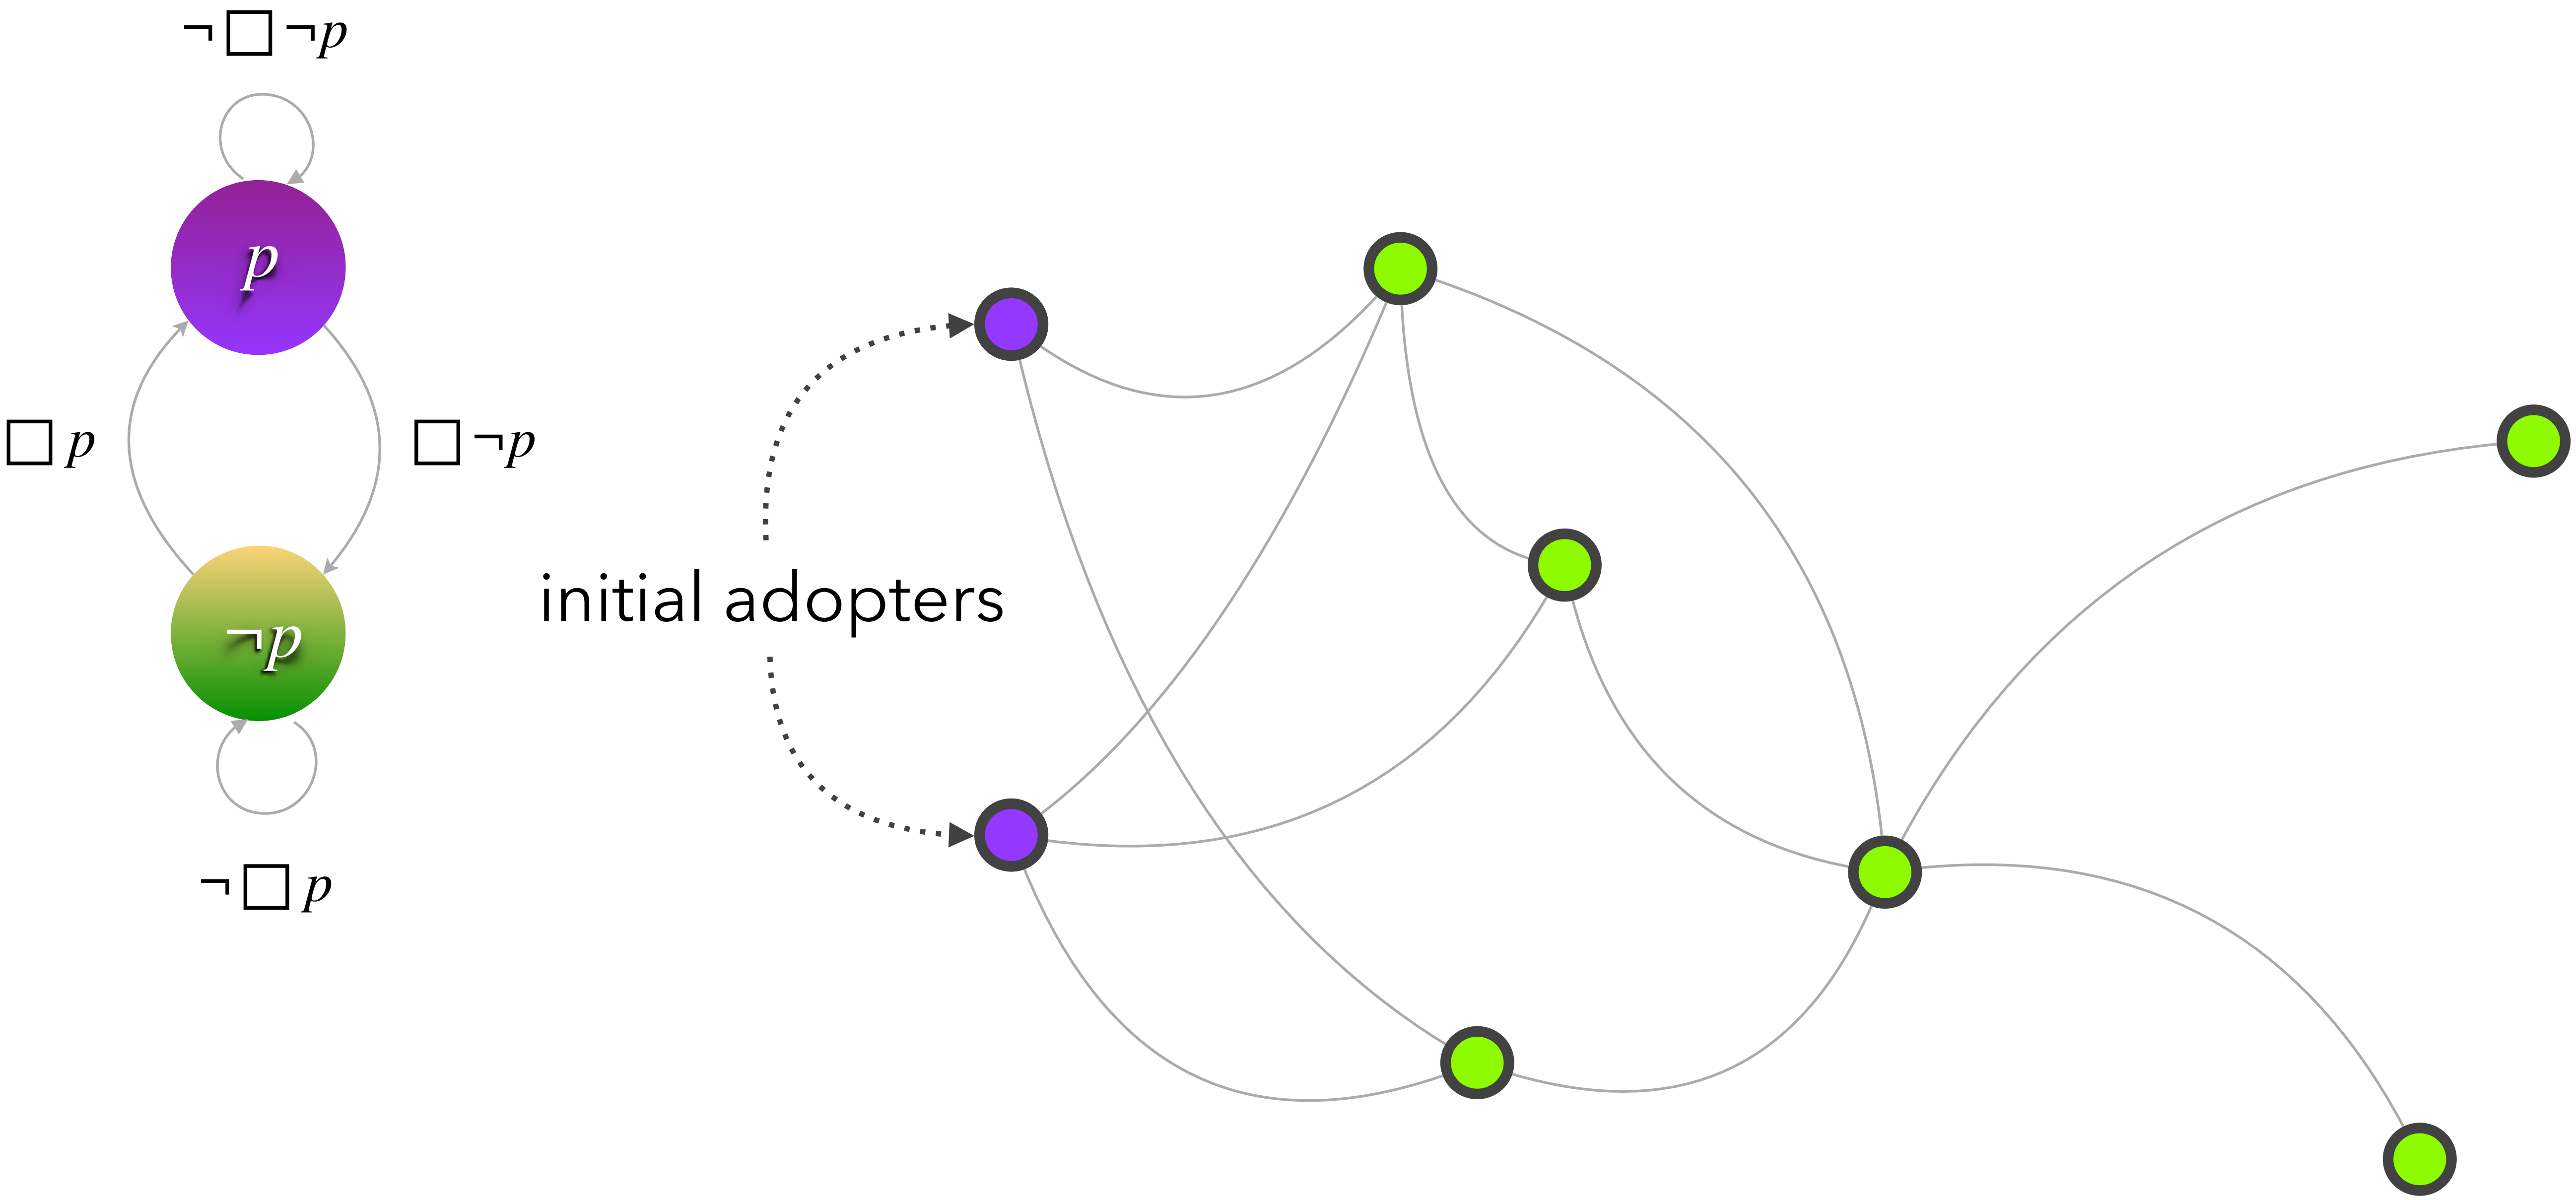

Diffusion dynamics under unanimity

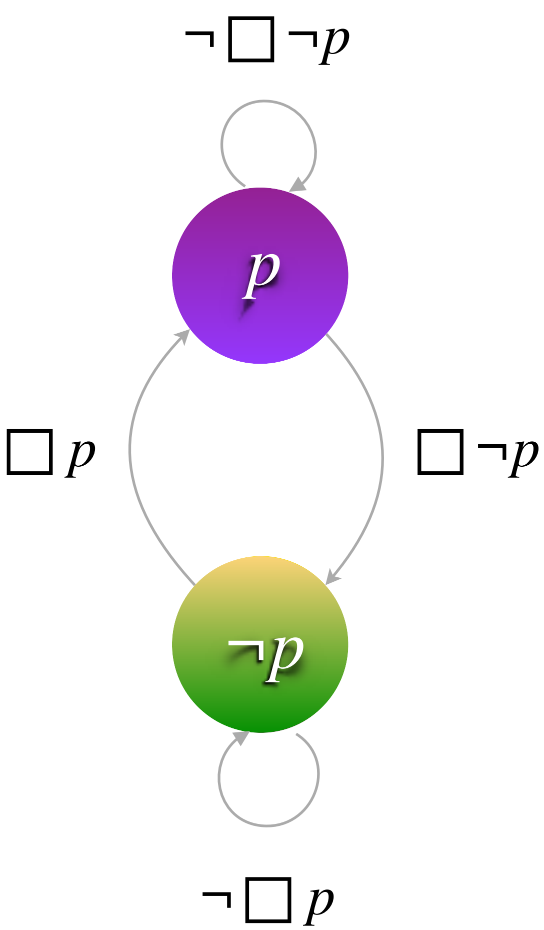

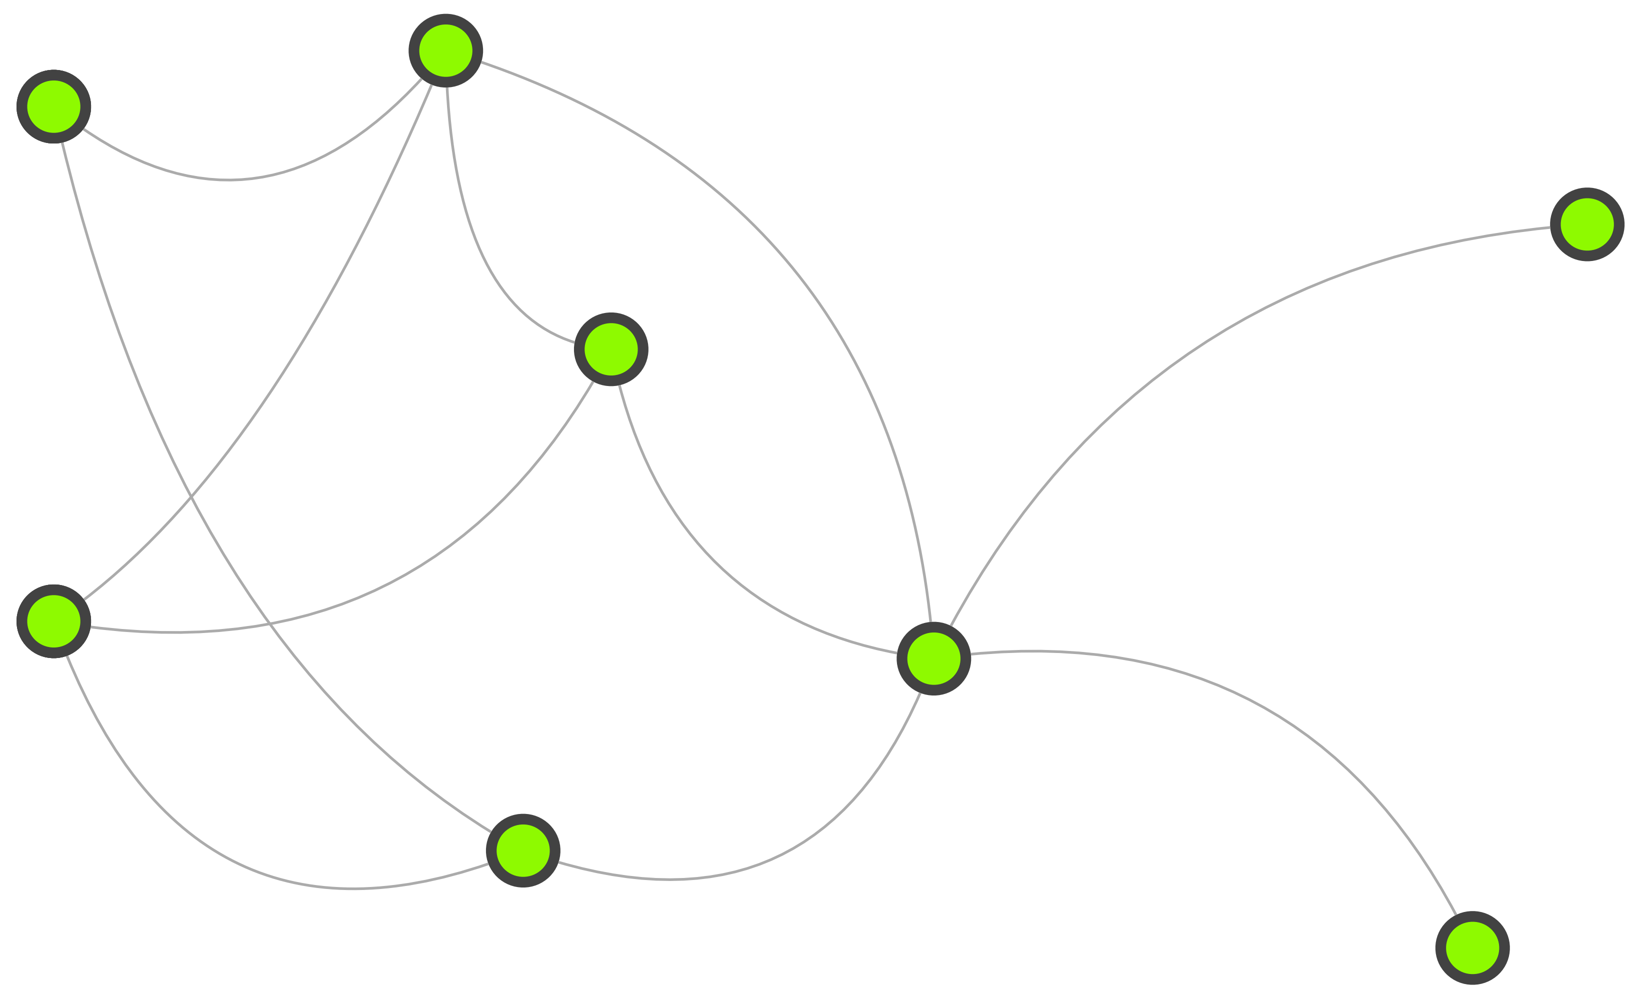

Stabilization

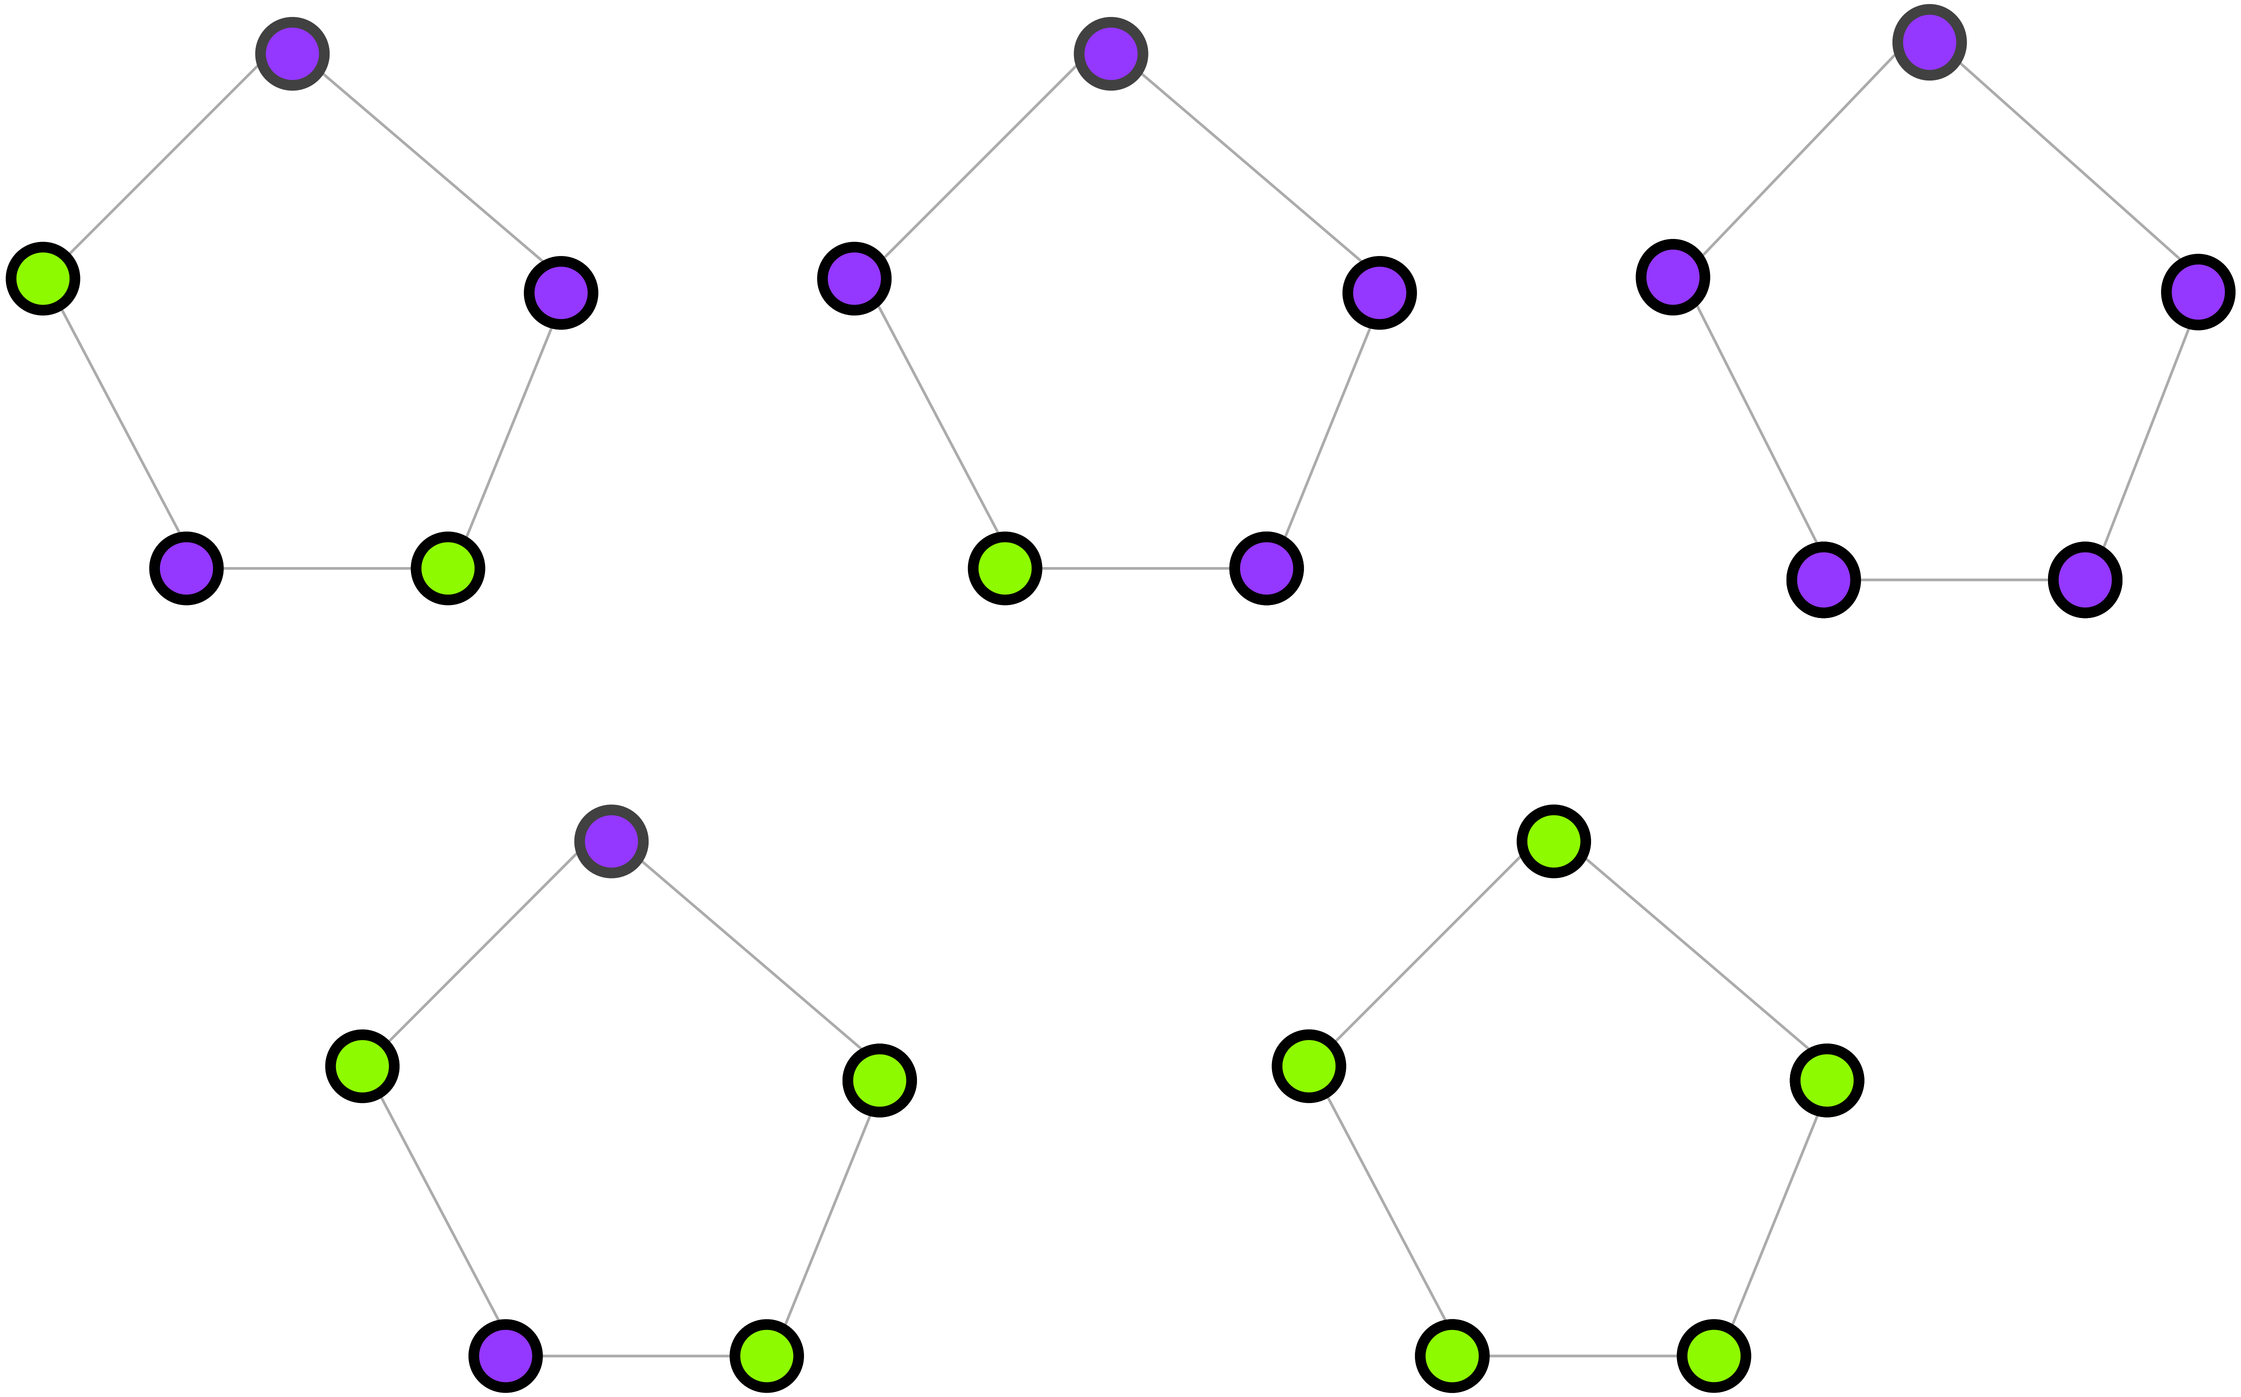

Oscillation

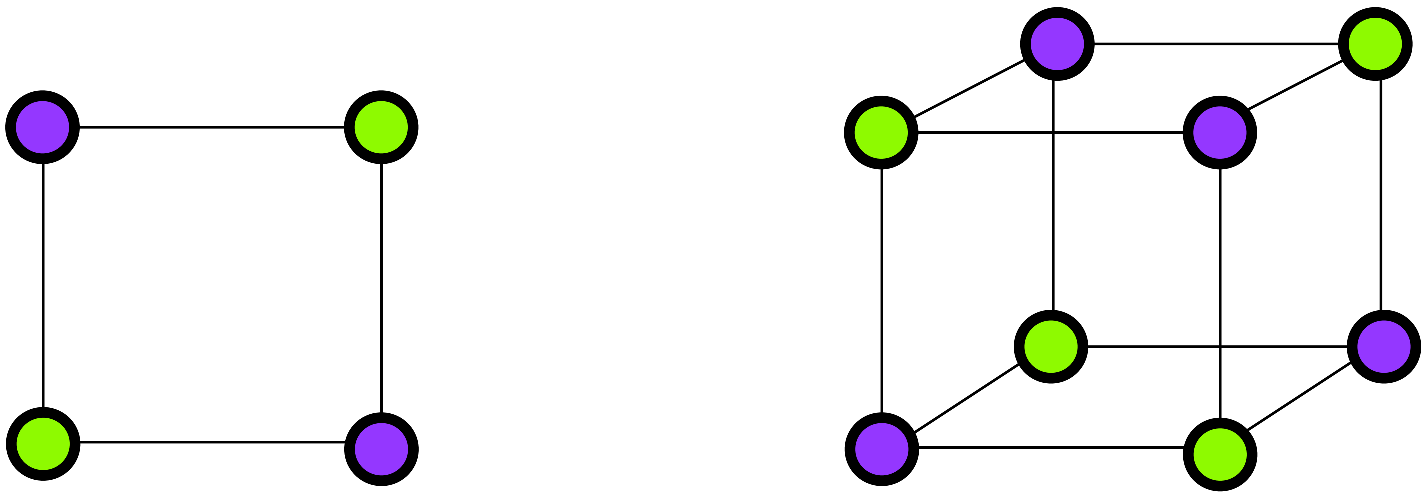

Stabilization conditions

Graph has 2-coloring if and only if it has no odd-length cycles.

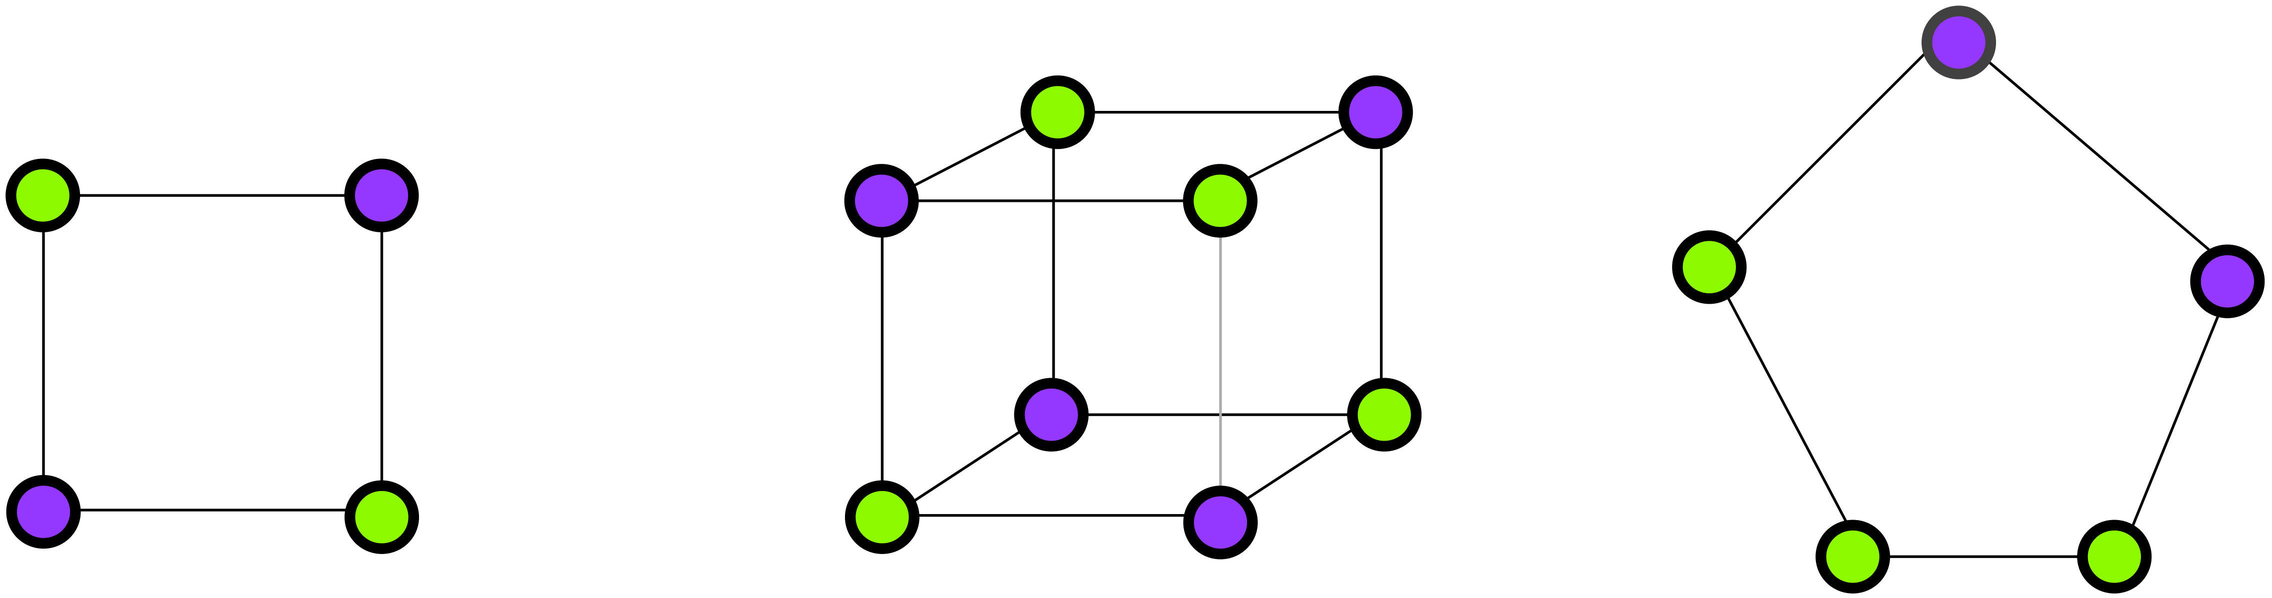

More fine-grained agents: 3 states?

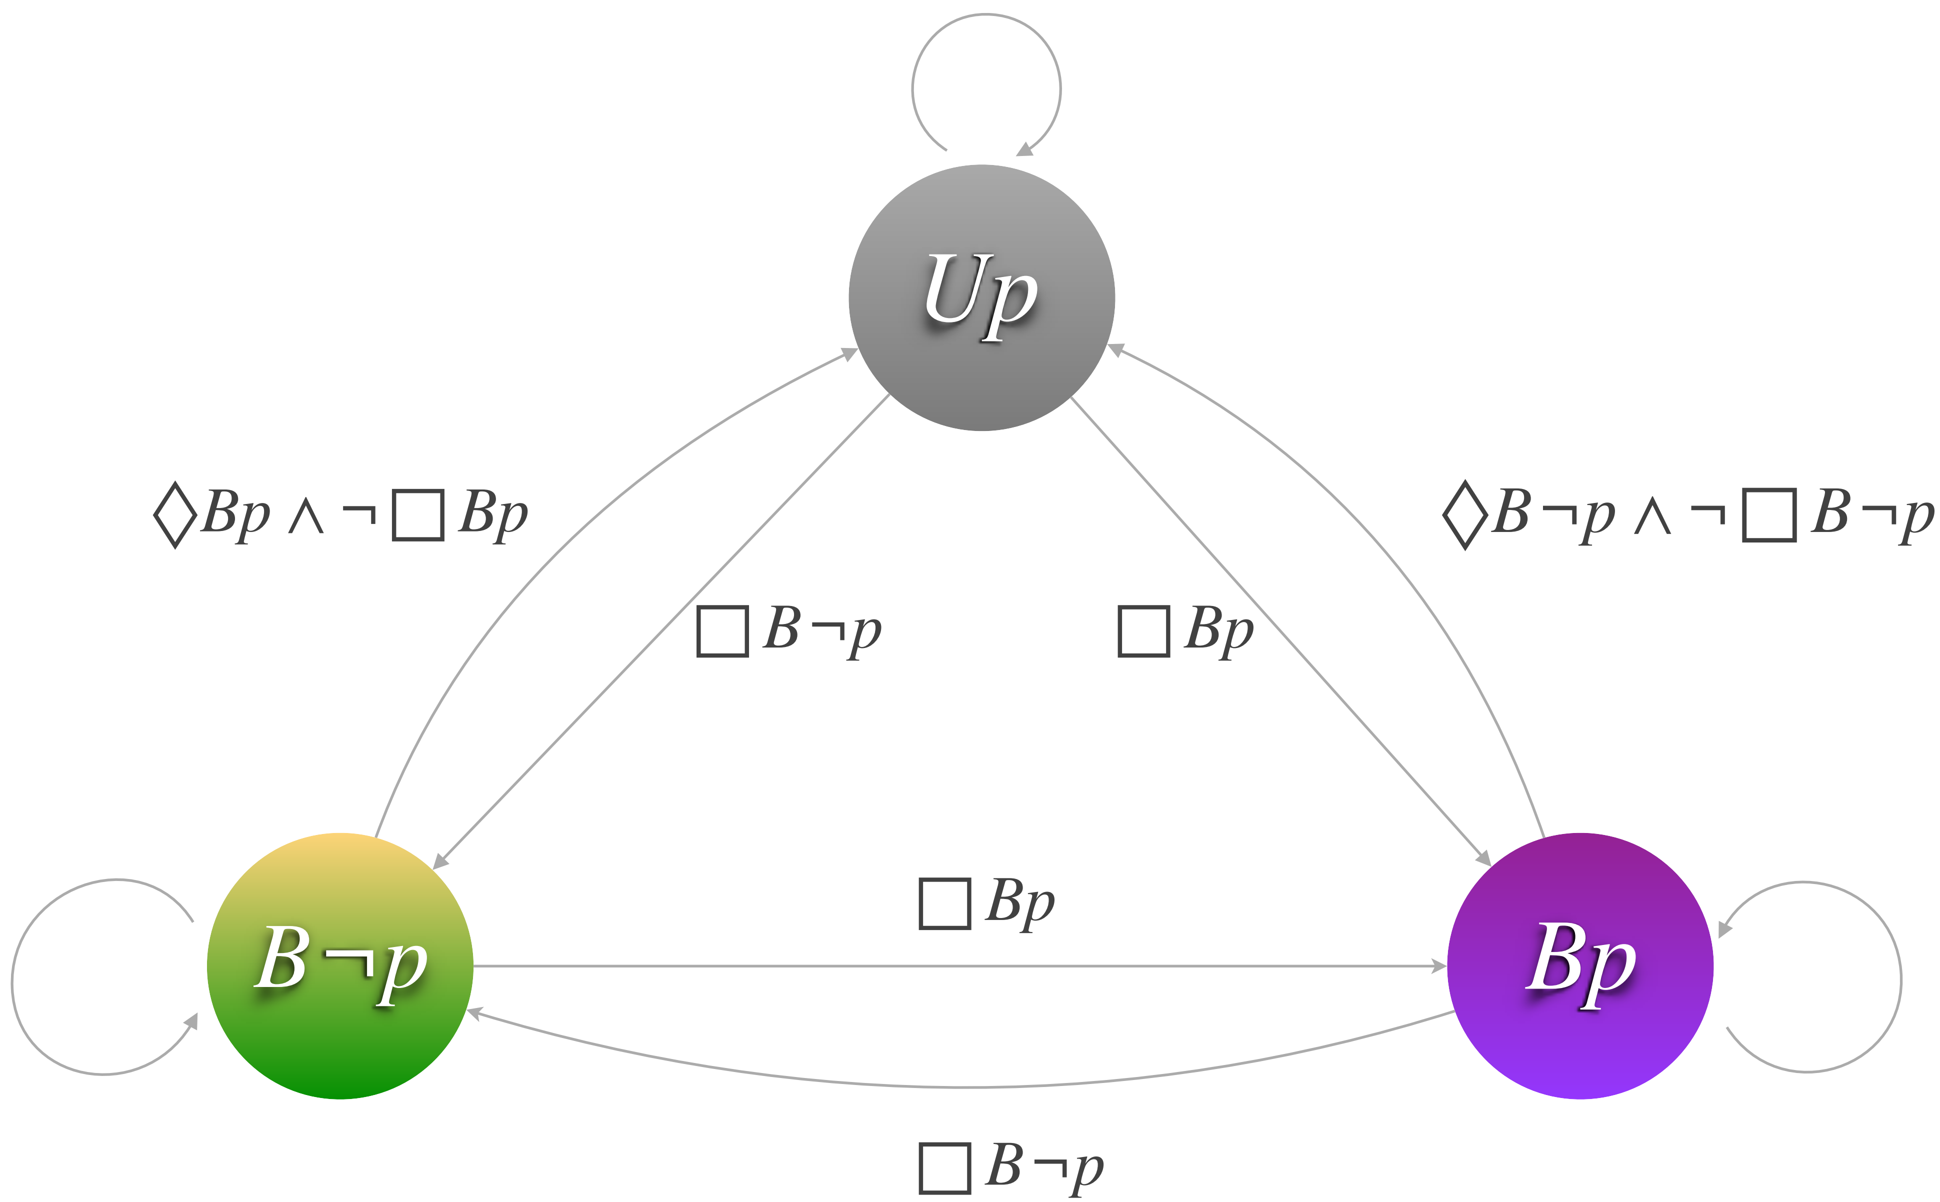

Logical Dynamics of Belief Change in the Community, Fenrong Liu, Jeremy Seligman & Patrick Girard, Synthese (2014)

Oscillation?

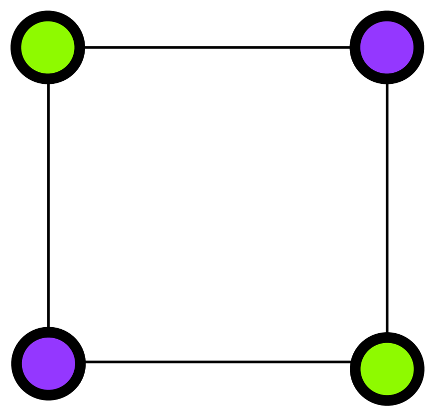

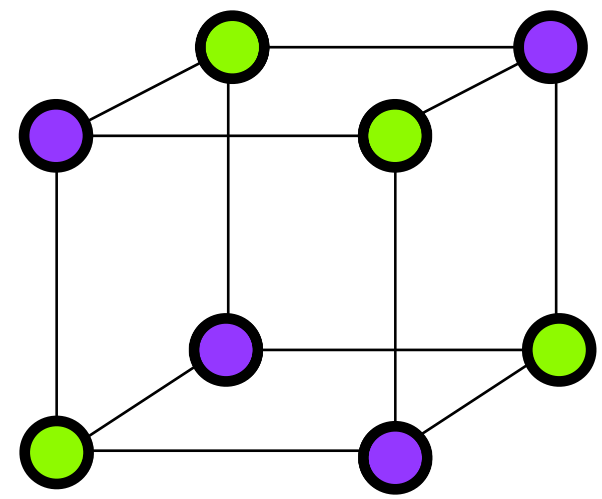

More theorems about diffusion stabilization (from the perspective of judgment aggregation) in :

Grandi, Lorini & Perrussel, Propositional Opinion Diffusion, AAMAS15

Christoff & Grossi, Stability in binary opinion diffusion, LORI 2017

Even more fine-grained models?

"All those in favor say 'Aye.'"
"Aye." *"Aye."* *"Aye."* *"Aye."* *"Aye."*

Even more fine-grained models?

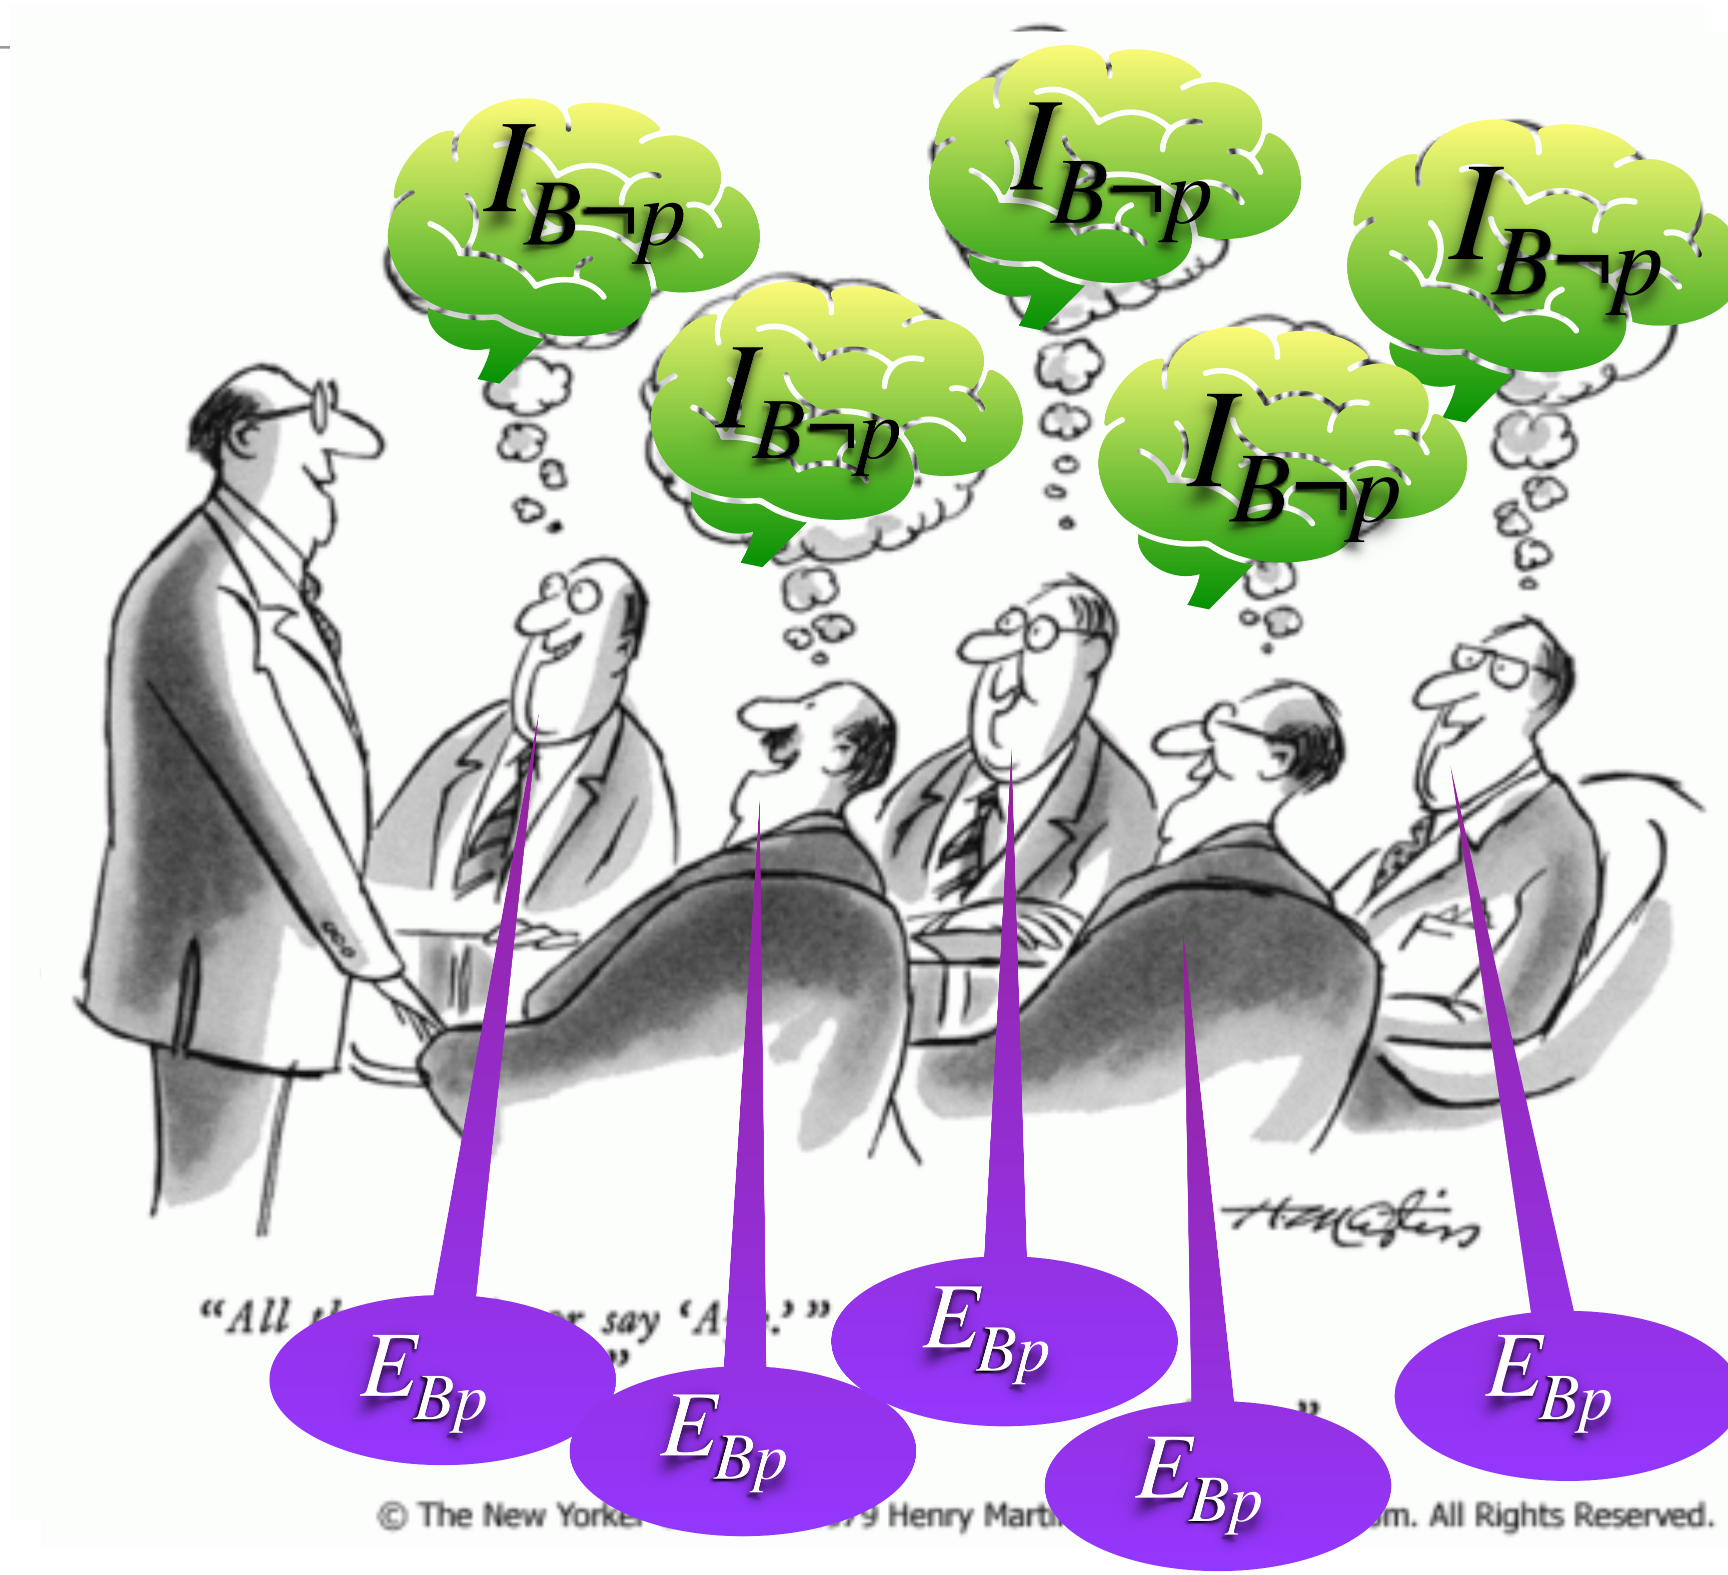

Even more fine-grained agents: 9 states?

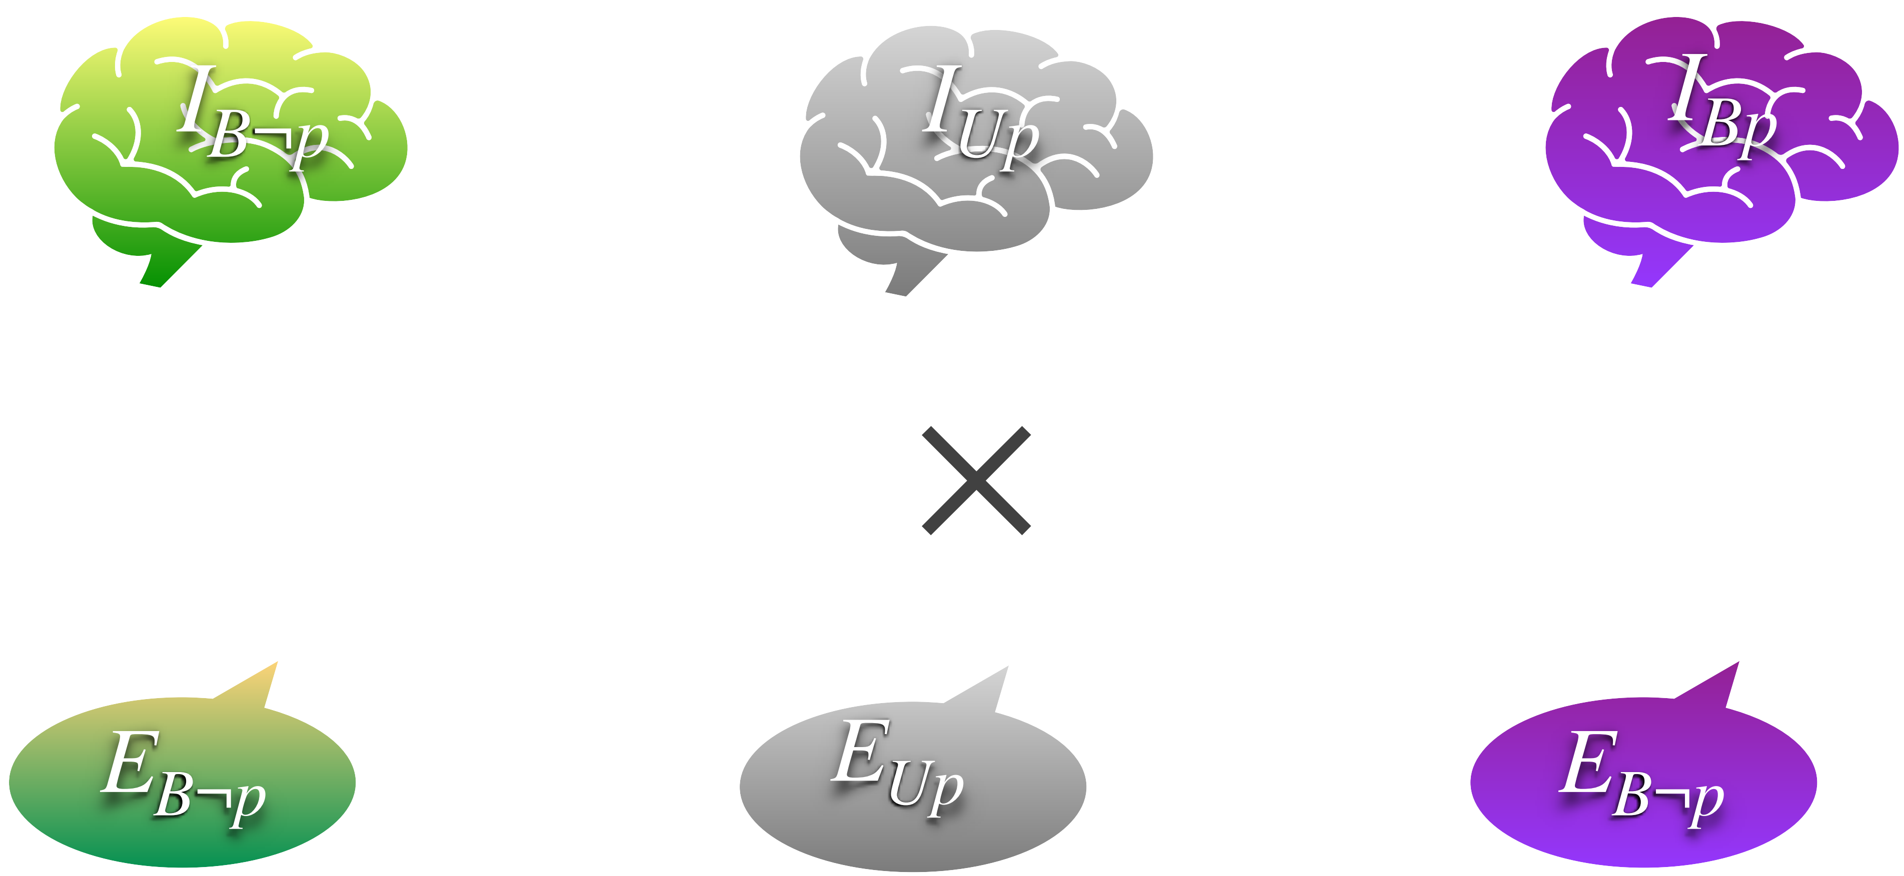

Christoff & Hansen "A two-tiered formalization of social influence" (LORI 2013)

Christoff, Hansen & Proietti "Reflecting on social influence in Networks" (2016) *Journal of Logic, Language & Information*

	Inner state	$\langle F\rangle E_B\varphi$	$\langle F\rangle E_B\neg\varphi$	$\langle F\rangle E_U\varphi$	Type 1
1	$I_B\varphi$				$\rightsquigarrow E_B\varphi$
2	$I_B\neg\varphi$	1	1	1	$\rightsquigarrow E_B\neg\varphi$
3	$I_U\varphi$				$\rightsquigarrow E_U\varphi$
4	$I_B\varphi$				$\rightsquigarrow E_B\varphi$
5	$I_B\neg\varphi$	1	1	0	$\rightsquigarrow E_B\neg\varphi$
6	$I_U\varphi$				$\rightsquigarrow E_U\varphi$
7	$I_B\varphi$				$\rightsquigarrow E_B\varphi$
8	$I_B\neg\varphi$	1	0	1	$\rightsquigarrow E_U\varphi$
9	$I_U\varphi$				$\rightsquigarrow E_U\varphi$
10	$I_B\varphi$				
11	$I_B\neg\varphi$	1	0	0	$\rightsquigarrow E_B\varphi$
12	$I_U\varphi$				
13	$I_B\varphi$				$\rightsquigarrow E_U\varphi$
14	$I_B\neg\varphi$	0	1	1	$\rightsquigarrow E_B\neg\varphi$
15	$I_U\varphi$				$\rightsquigarrow E_B\neg\varphi$
16	$I_B\varphi$				
17	$I_B\neg\varphi$	0	1	0	$\rightsquigarrow E_B\neg\varphi$
18	$I_U\varphi$				
19	$I_B\varphi$				$\rightsquigarrow E_B\varphi$
20	$I_B\neg\varphi$	0	0	1	$\rightsquigarrow E_B\neg\varphi$
21	$I_U\varphi$				$\rightsquigarrow E_U\varphi$
22	$I_B\varphi$				$\rightsquigarrow E_B\varphi$
23	$I_B\neg\varphi$	0	0	0	$\rightsquigarrow E_B\neg\varphi$
24	$I_U\varphi$				$\rightsquigarrow E_U\varphi$

Example of result: Dissolving pluralistic ignorance in networks

Proposition 2. *Let $\mathcal{M} = (A, \sim, g, \nu)$ be a finite, connected, symmetric network model in a state of unstable pluralistic ignorance. Then the following are equivalent:*

(i) After a finite number of updates by the influence event $\mathcal{I}$, $\mathcal{M}$ will end up in a stable state where pluralistic ignorance is dissolved, i.e. there is a $k \in \mathbb{N}$ such that $\mathcal{M} \otimes^k \mathcal{I} \models G(I_B\varphi \wedge E_B\varphi)$ and $\mathcal{M} \otimes^k \mathcal{I} = \mathcal{M} \otimes^{k+1} \mathcal{I}$.

(ii) There is an agent that expresses her true belief in φ for two rounds in a row, i.e. there is an $a \in A$ and a $k \in \mathbb{N}$ such that $\mathcal{M} \otimes^k \mathcal{I}, a \models E_B\varphi$ and $\mathcal{M} \otimes^{k+1} \mathcal{I}, a \models E_B\varphi$.

(iii) There are two agents that are friends and both express their true beliefs in φ in the same round, i.e. there are $a, b \in A$ and a $k \in \mathbb{N}$ such that $a \sim b$, $\mathcal{M} \otimes^k \mathcal{I}, a \models E_B\varphi$, and $\mathcal{M} \otimes^k \mathcal{I}, b \models E_B\varphi$.

(iv) There are two agents that are friends and have paths of the same length to the agent named by i, i.e. there are agents $a, b \in A$ and a $k \in \mathbb{N}$ such that $a \sim b$, $\mathcal{M}, a \models \langle F \rangle^k i$, and $\mathcal{M}, b \models \langle F \rangle^k i$.

(v) There is a cycle in $\mathcal{M}$ of odd length starting at the agent named by i, i.e. there is a $k \in \mathbb{N}$ such that $\mathcal{M} \models @_i \langle F \rangle^{2k-1} i$.

(vi) There is a cycle in $\mathcal{M}$ of odd length, i.e. there is a $k \in \mathbb{N}$ and $a_1, a_2, ..., a_{2k-1} \in A$ such that $a_1 \sim a_2, a_2 \sim a_3, ..., a_{2k-2} \sim a_{2k-1}, a_{2k-1} \sim a_1$.

Even more fine-grained? *n* layers, *m* values

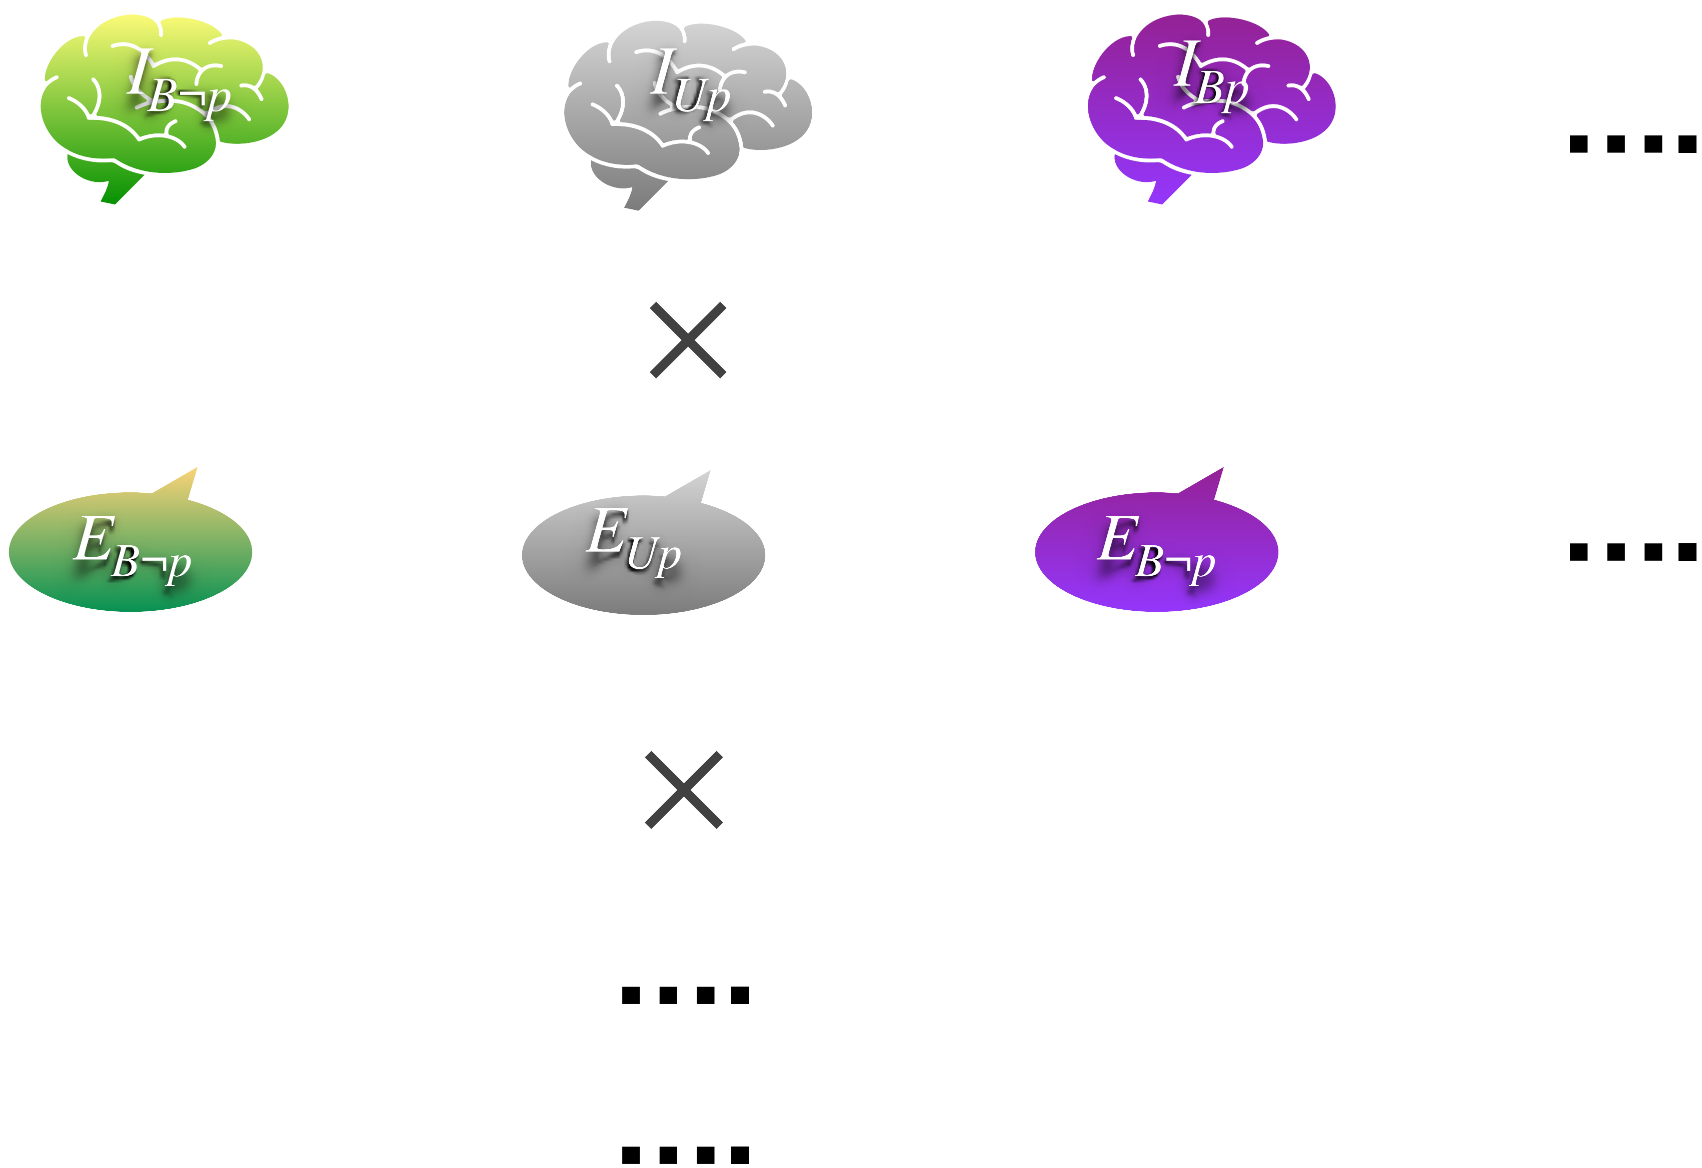

Christoff & Hansen, 2016, A logic for diffusion in social networks, *Journal of Applied Logic*

V. Combining epistemic and network directions

DEL of threshold diffusion

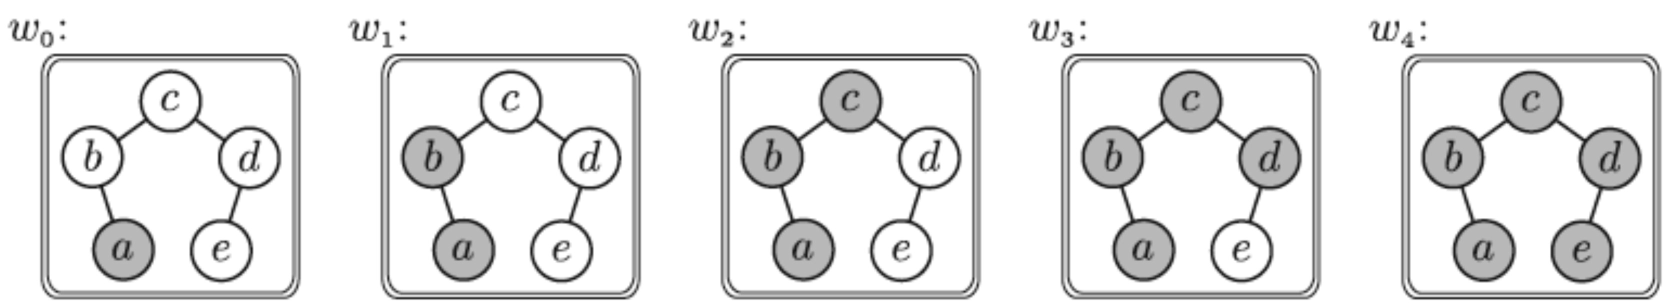

Model 3 diffusion policies:

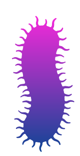

- adopt whenever enough of your neighbors have adopted

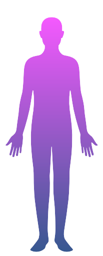

- adopt whenever YOU KNOW THAT enough of your neighbors have adopted

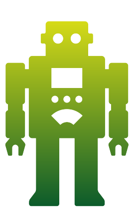

- adopt whenever YOU KNOW THAT enough of your neighbors WILL have adopted (at some point)

Compare the 3 diffusion policies: not that big a difference!

ALEXANDRU BALTAG
ZOÉ CHRISTOFF
RASMUS K. RENDSVIG
SONJA SMETS
Dynamic Epistemic Logics of Diffusion and Prediction in Social Networks
(2019) *Studia Logica*

Logic as modeling tool

I have presented 3 examples of insight from logical perspective:

- Inescapability of cascades for agents with unbounded higher-order rationality (probabilistic DEL)
- Diffusion dynamics and network structures relation (modal/hybrid/fixed point logics for social networks)
- Insight on how the behavior of “very smart” agents might not differ so much from the ones of “bacteria-like” agents (diffusion epistemic logic)

Beyond what I mentioned so far…

Rich logic toolbox by now, to capture for instance:

- what happens in diffusion in the long run directly ("ability-logics"):

  Christoff & Naumov (2019), Social Networks Diffusion with Recalcitrant Agents, *Journal of Applied Logic*

  Ågotnes & Christoff (2020), Reasoning about cascading abilities in Networks, Netreason@ECAI

- how the network structure evolves:
  - how networks with friends and enemy to tend towards balance:

    Xiong & Ågotnes (2020), On the Logic of Balance in Social Networks, JOLLI

    Hoek, Kuijer, & Wáng (2020), Logics of Allies and Enemies: A Formal Approach to the Dynamics of Social Balance Theory, IJCAI

  - how links tend to be created/deleted based on agents similarity:

    Smets & Velázquez-Quesada (2020), A Closeness and Priority-Based Logical Study of Social Network Creation, JOLLI

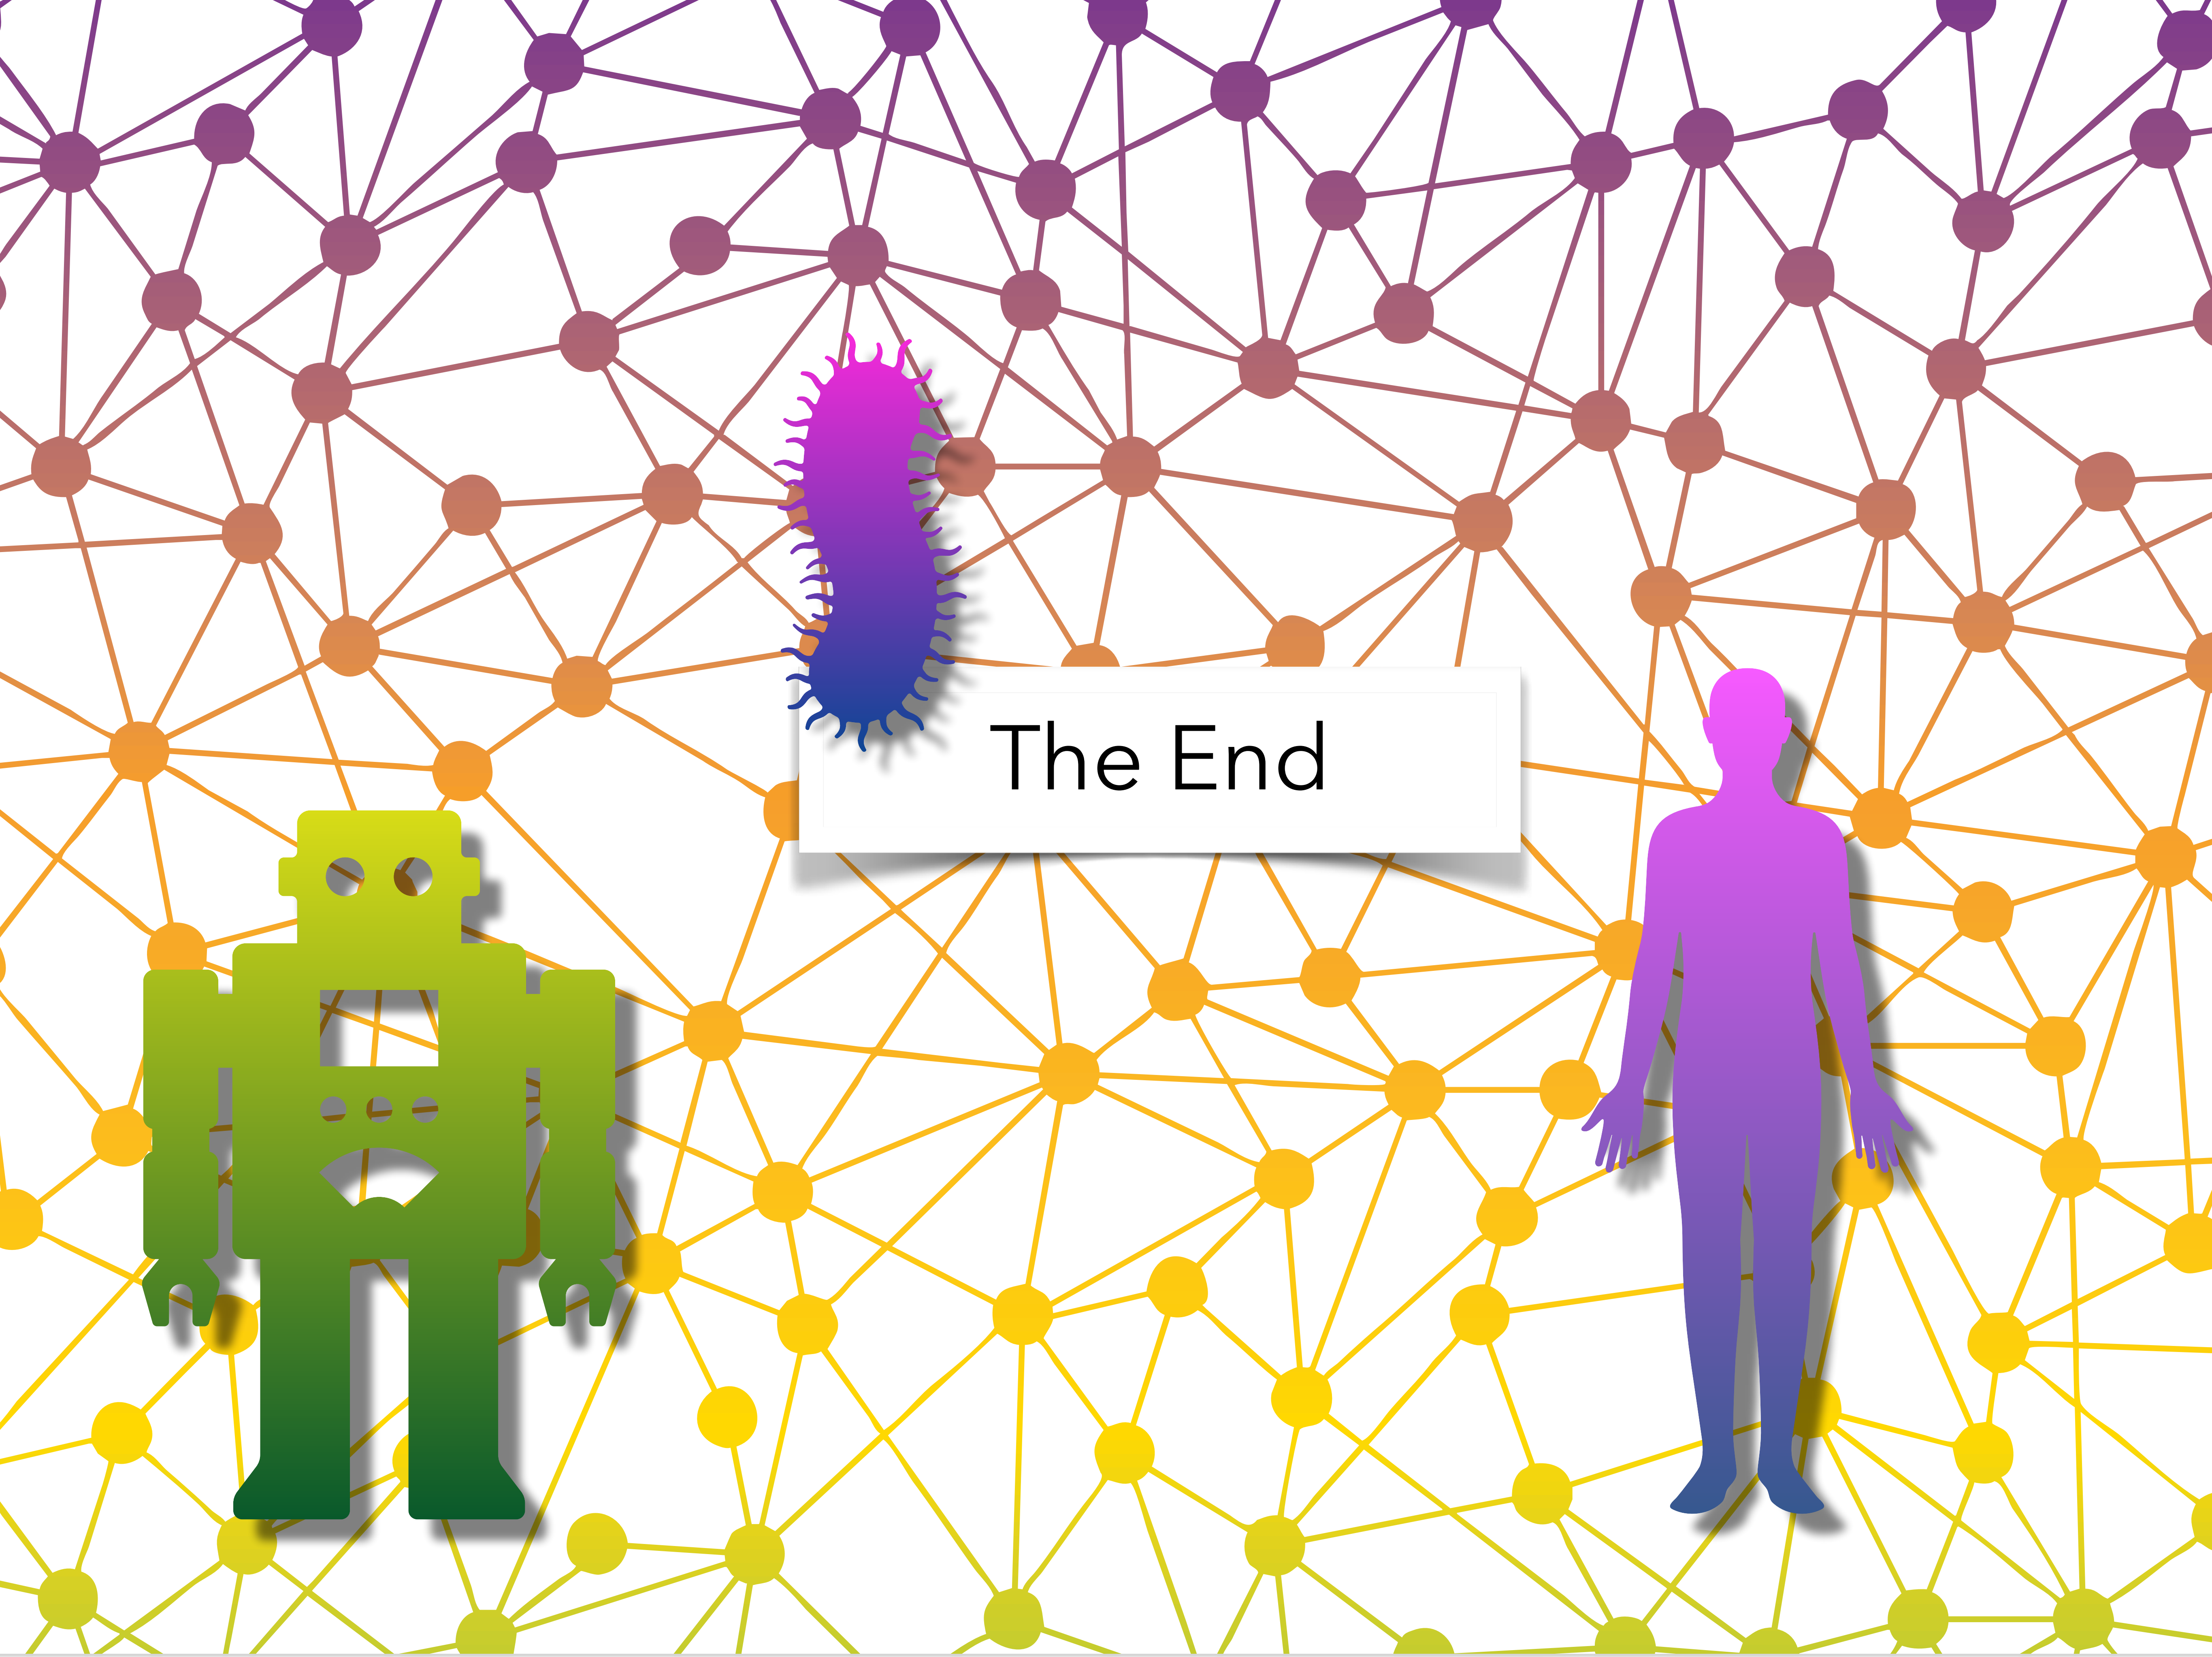
The End