DISPLEXITY 2nd Workshop on Distributed Computing: Computability and Complexity

Imperfect Best-Response Mechanisms

Paolo Penna
LIAFA Univ. Paris Diderot

joint work with
Diodato Ferratioli
Univ. Rome “La Sapienza”

Game Theory

Ideally

In practice…

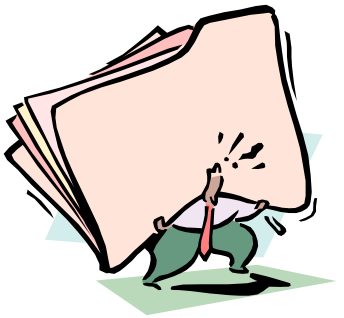

The Internet

No central authority, open, self organized, anarchic

Different “entities” which

- have their own goal
- may not follow the “protocol”

Rational (selfish)

This talk

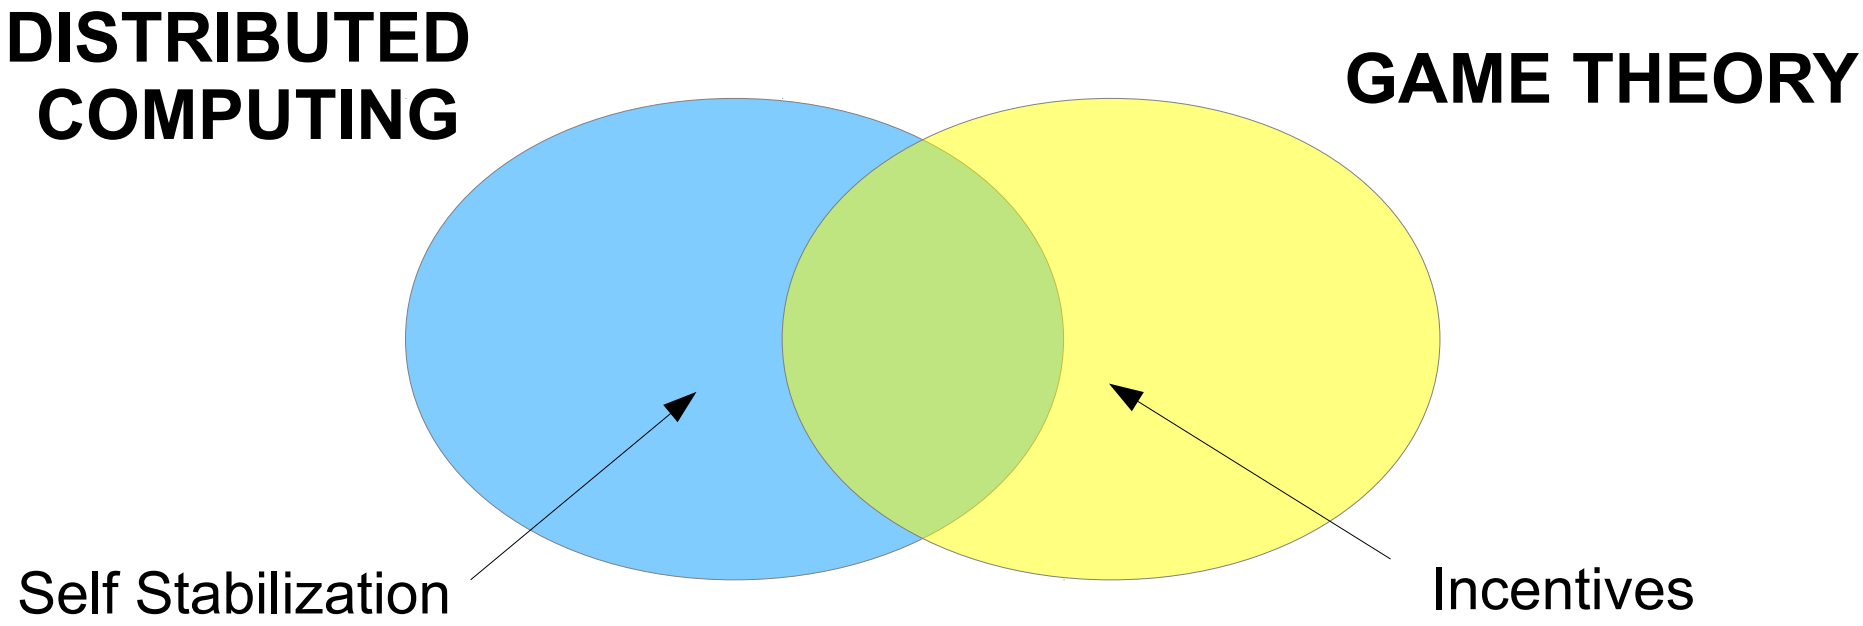

Best-Response Mechanisms
(Nisan, Schapira, Valiant, Zohar, ICS11)

Imperfect Best-Response Mechanisms
(Ferraioli&P., SAGT13)

TCP

Probe-increase educated-decrease

(increase rate if no packet lost, decrease othw)

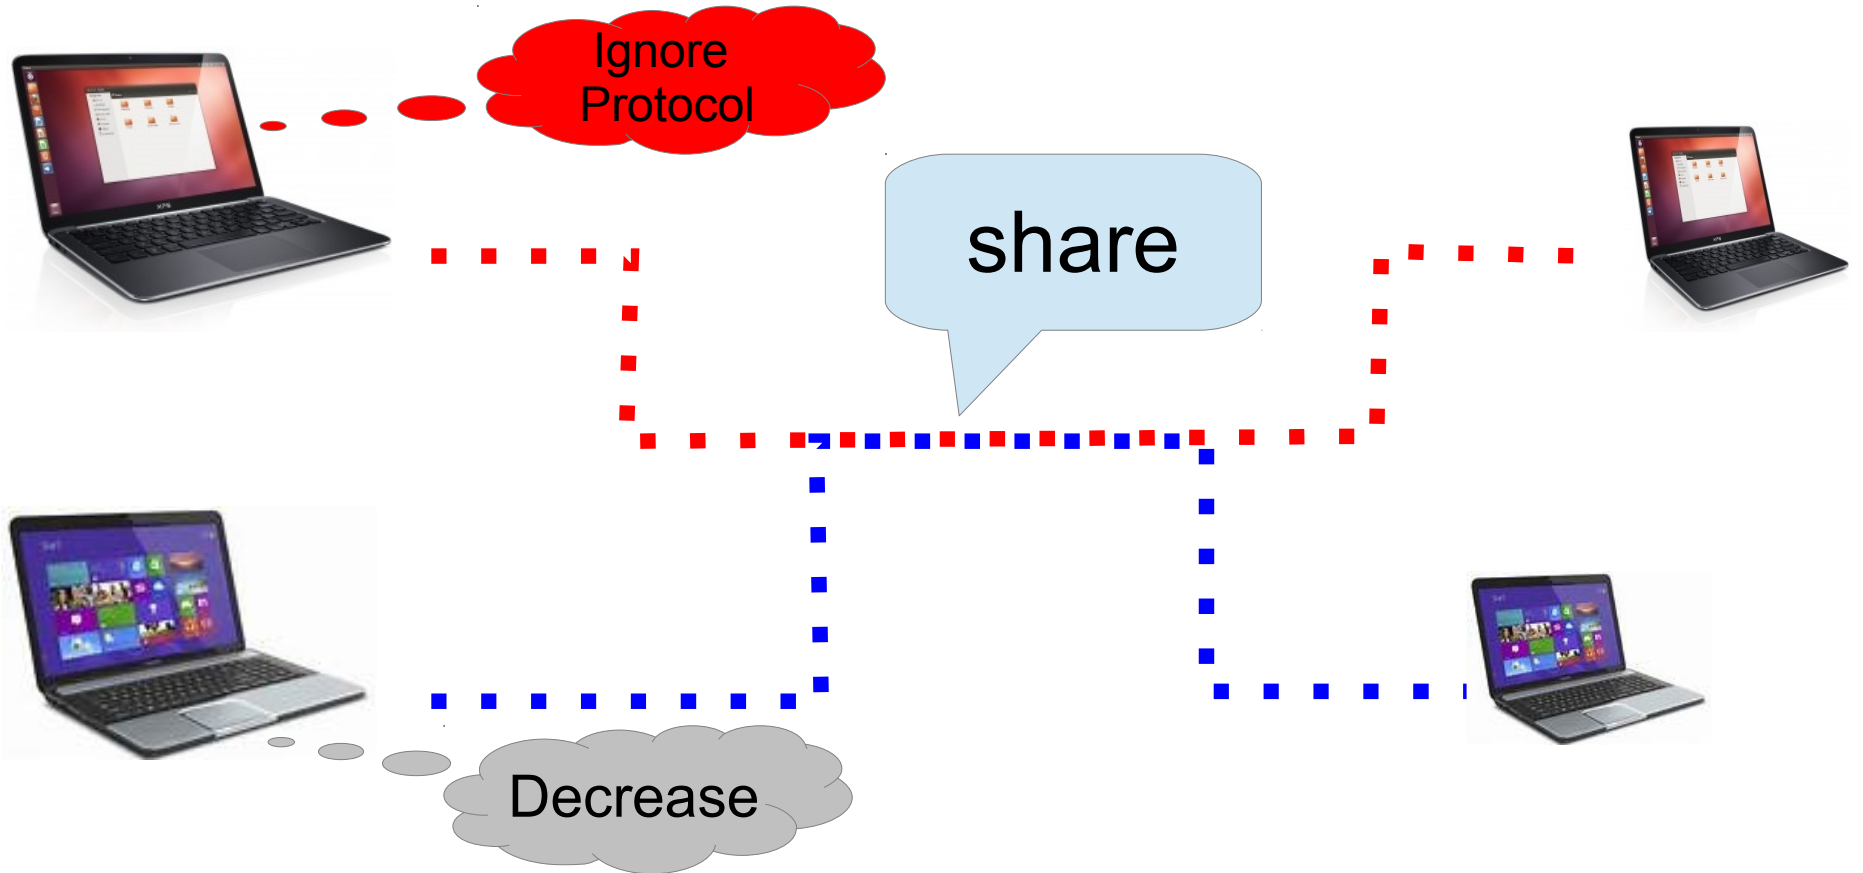

Border Gateway Protocol

Routes to destination

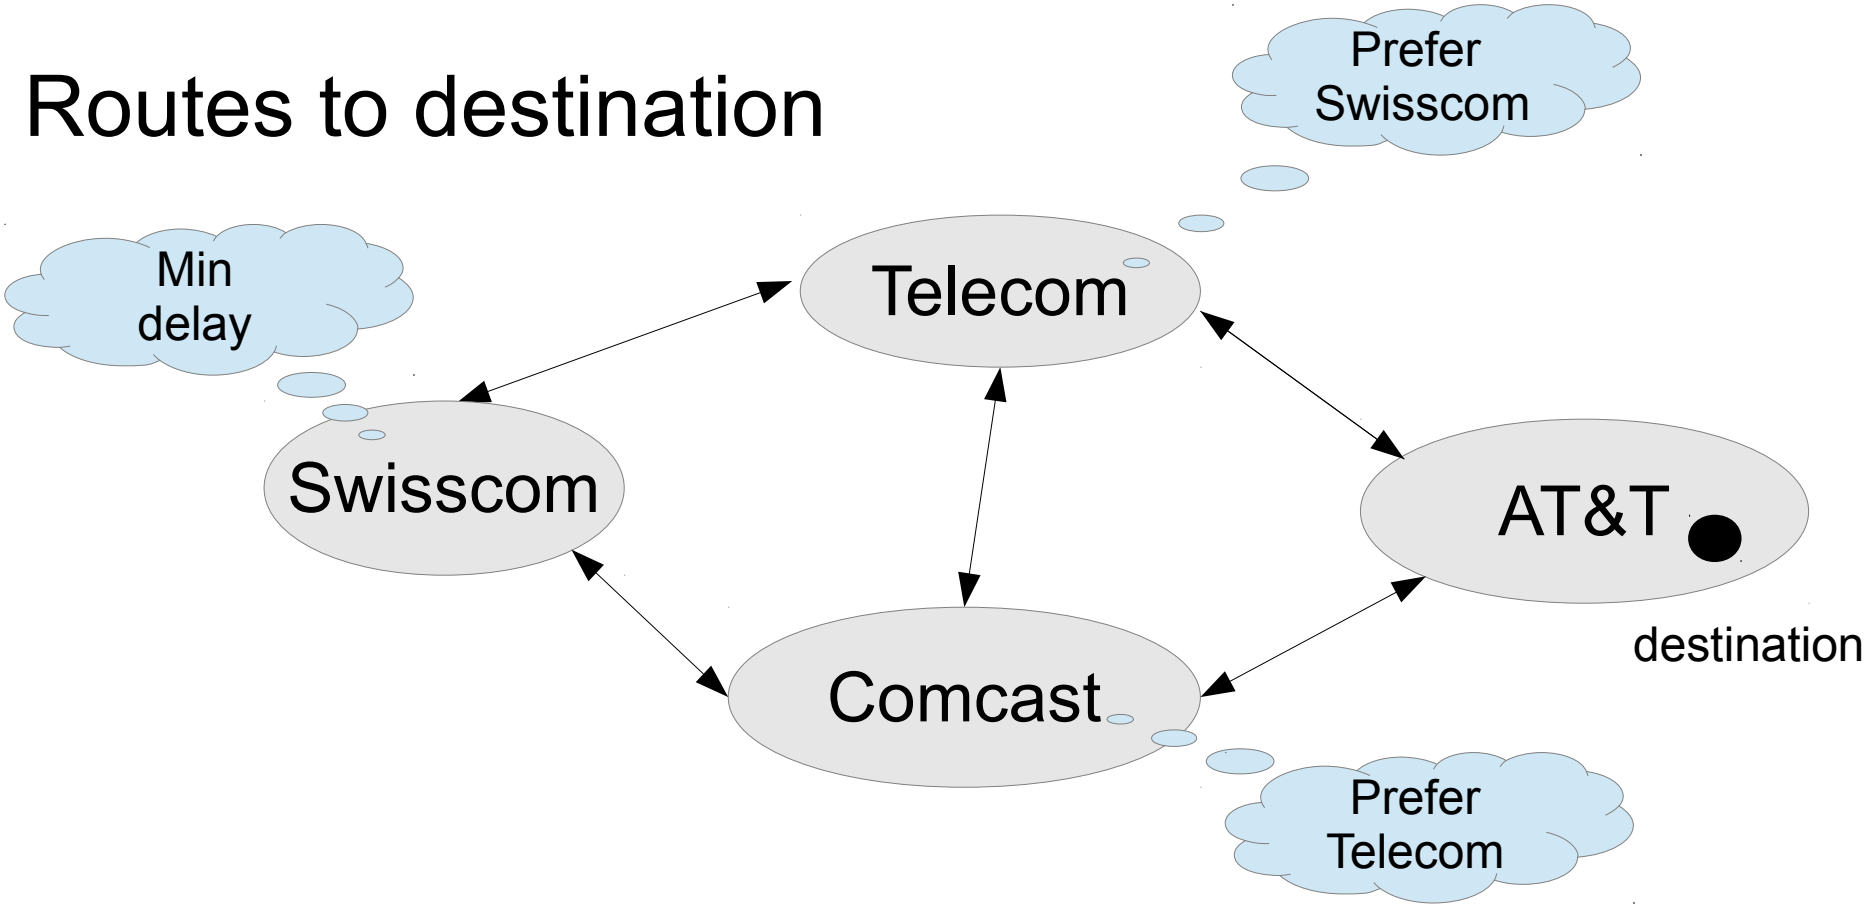

Local choice (“next hop”)
Autonomous (“best for me”)

Best-Response Mechanisms
(Nisan et al 2011)

Protocol: “repeatedly best respond” (greedy)

Asynchronous setting (adversarial schedule)

Example

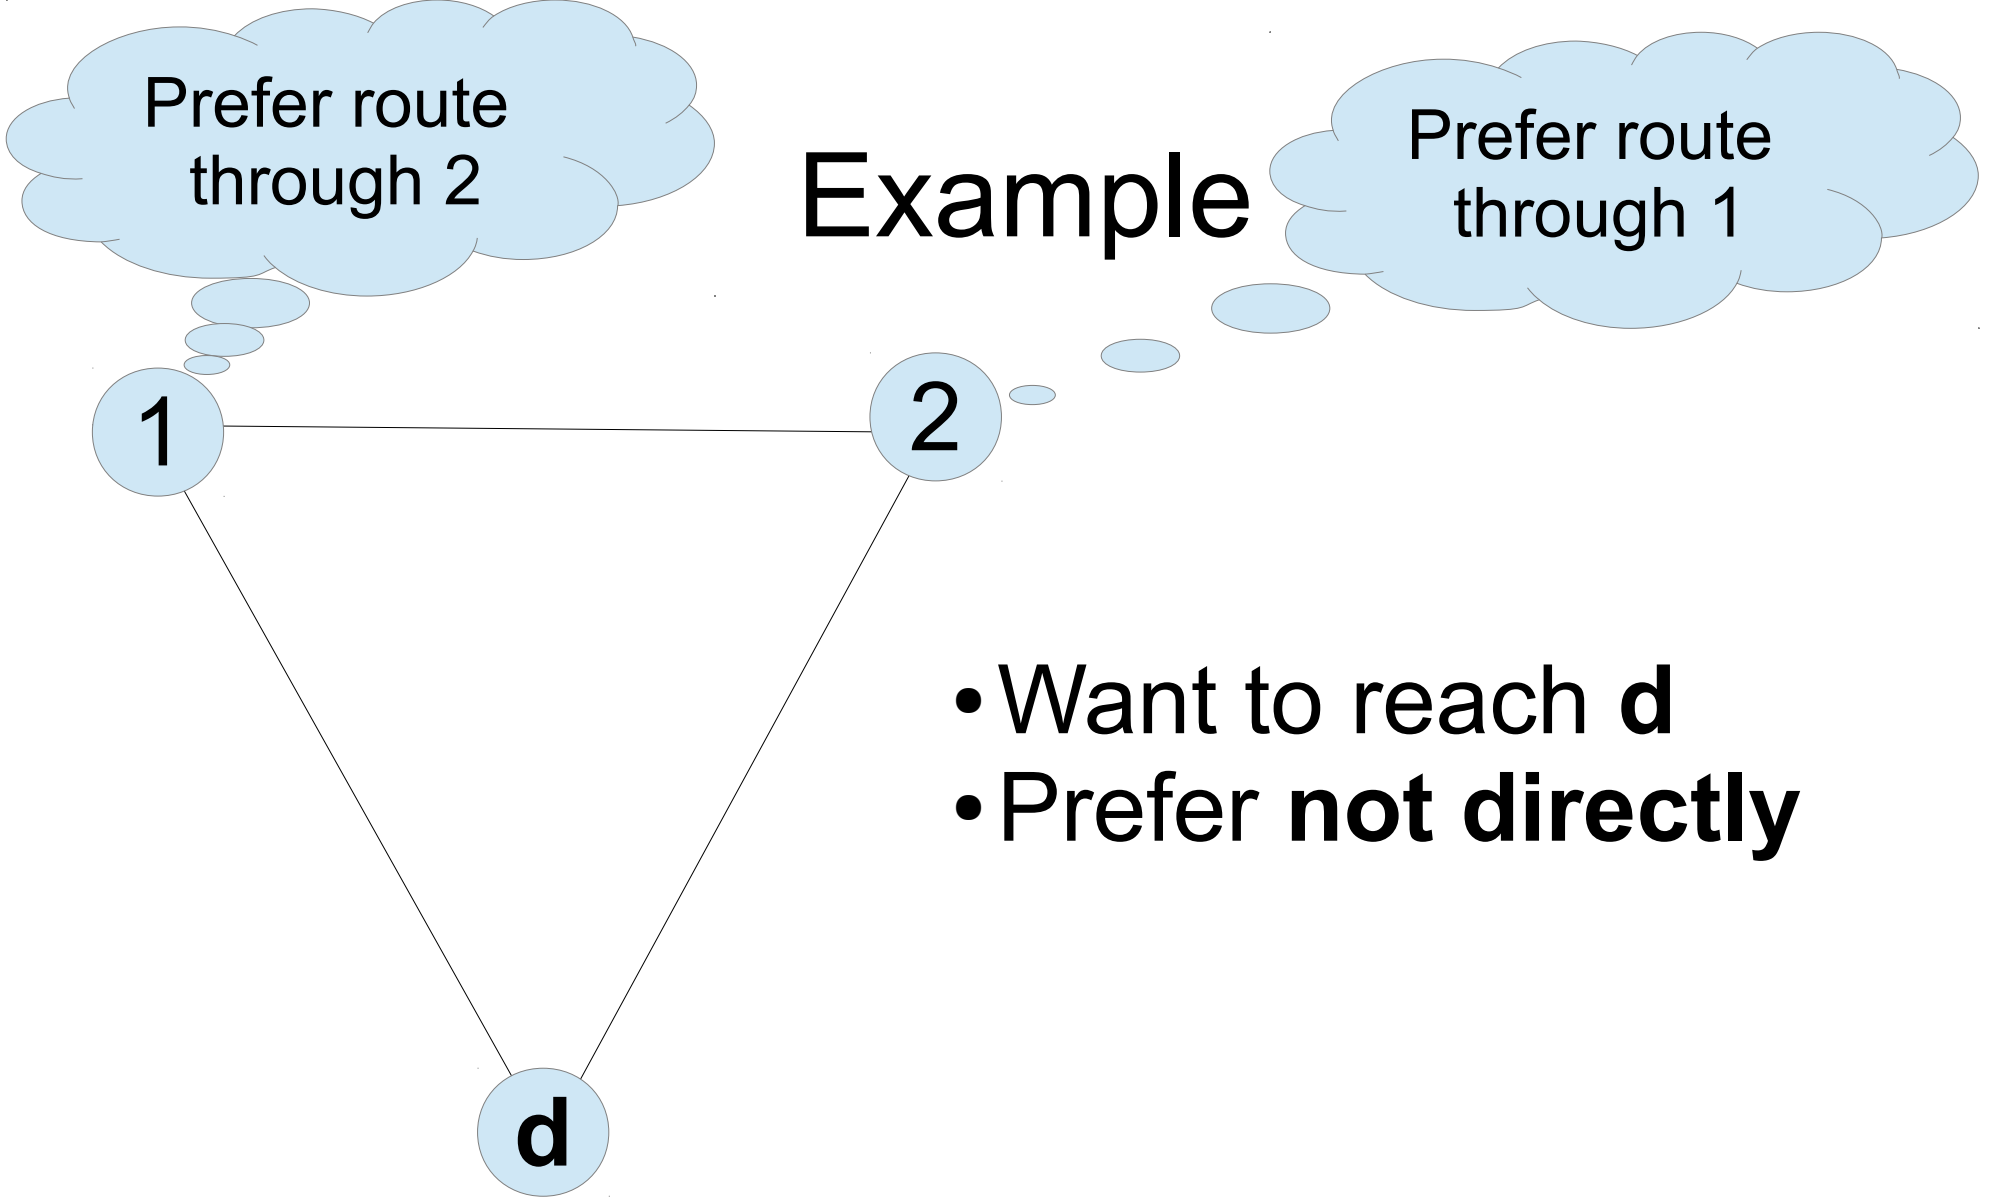

- Want to reach **d**
- Prefer **not directly**

Example

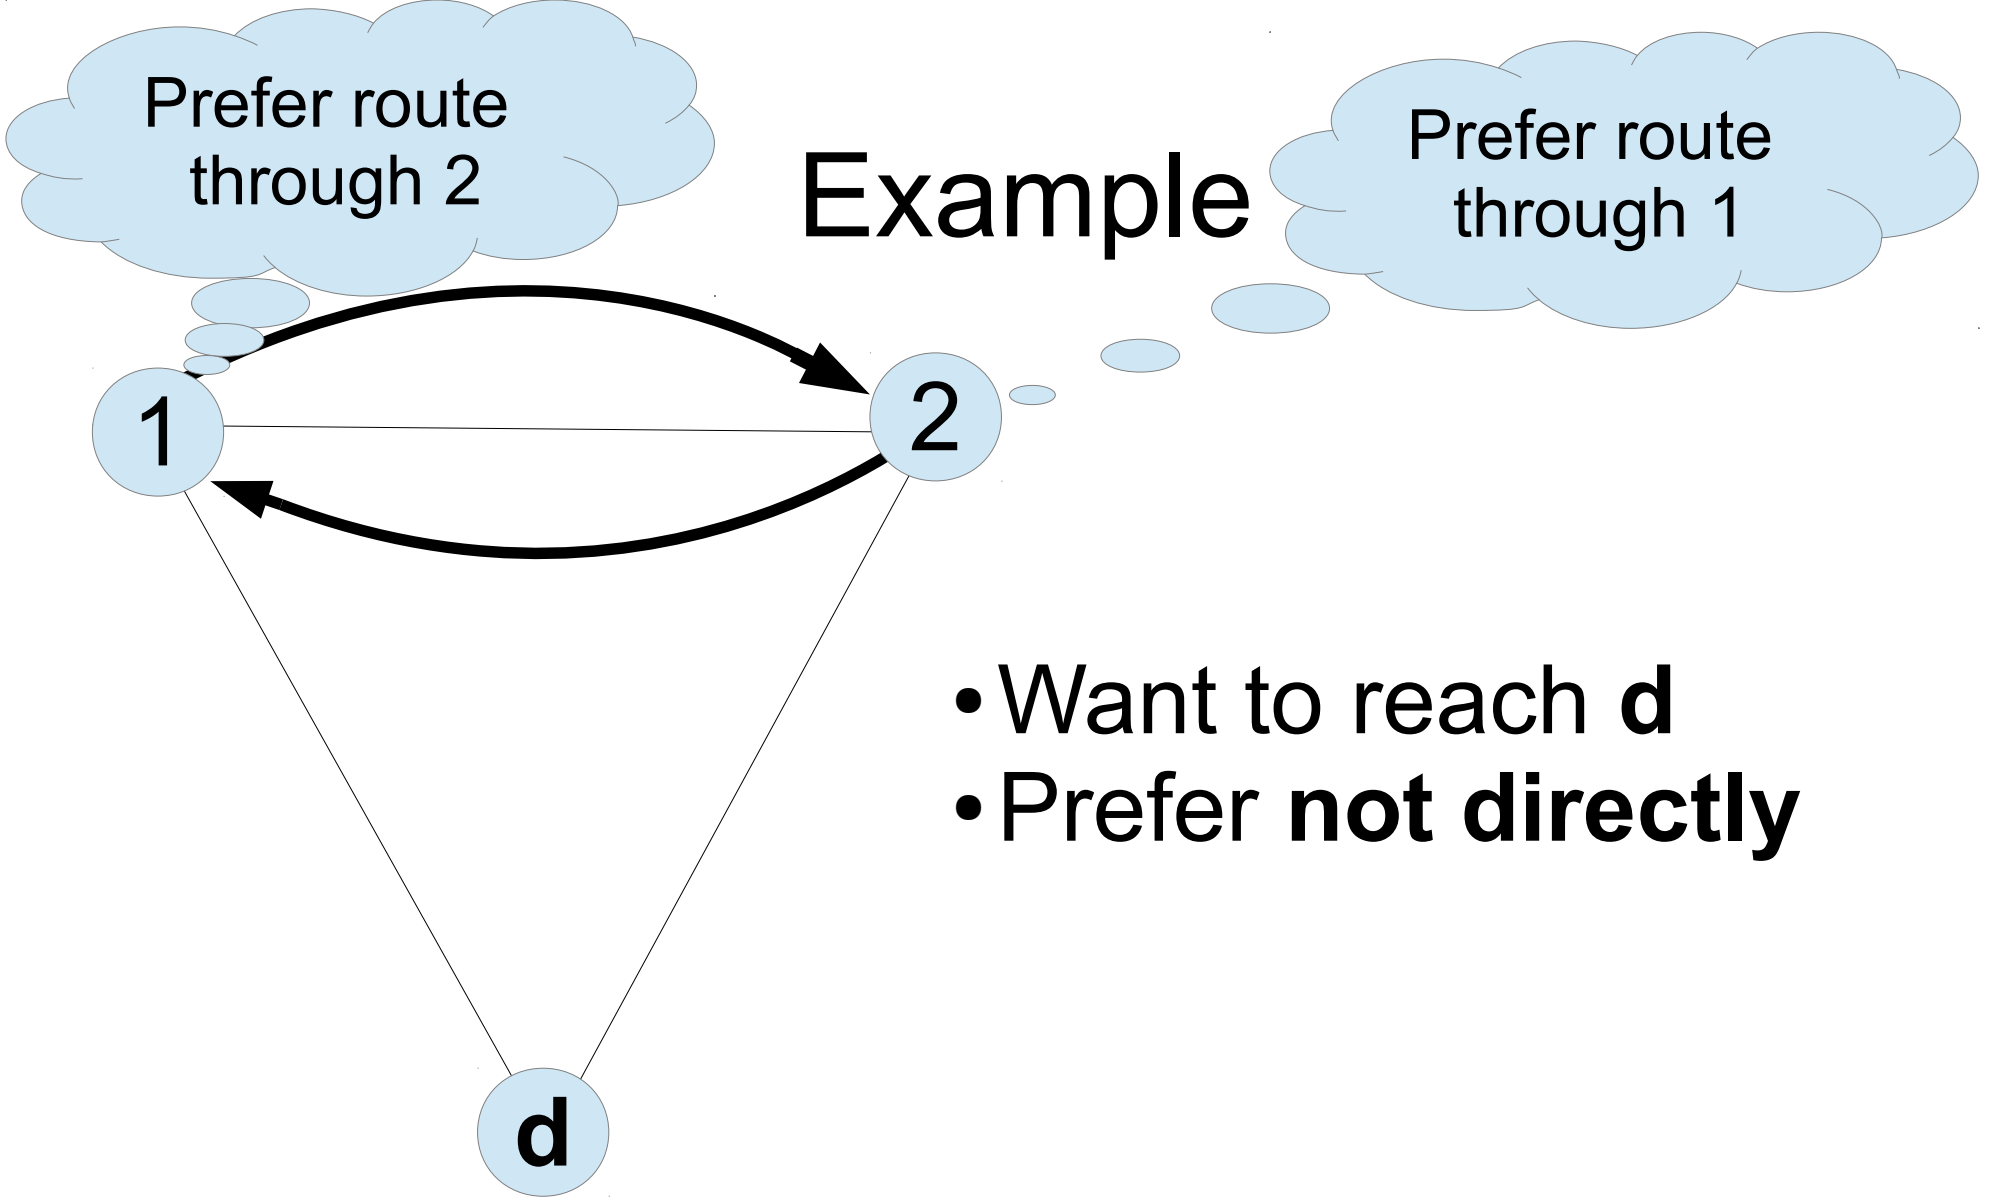

- Want to reach **d**
- Prefer **not directly**

Example

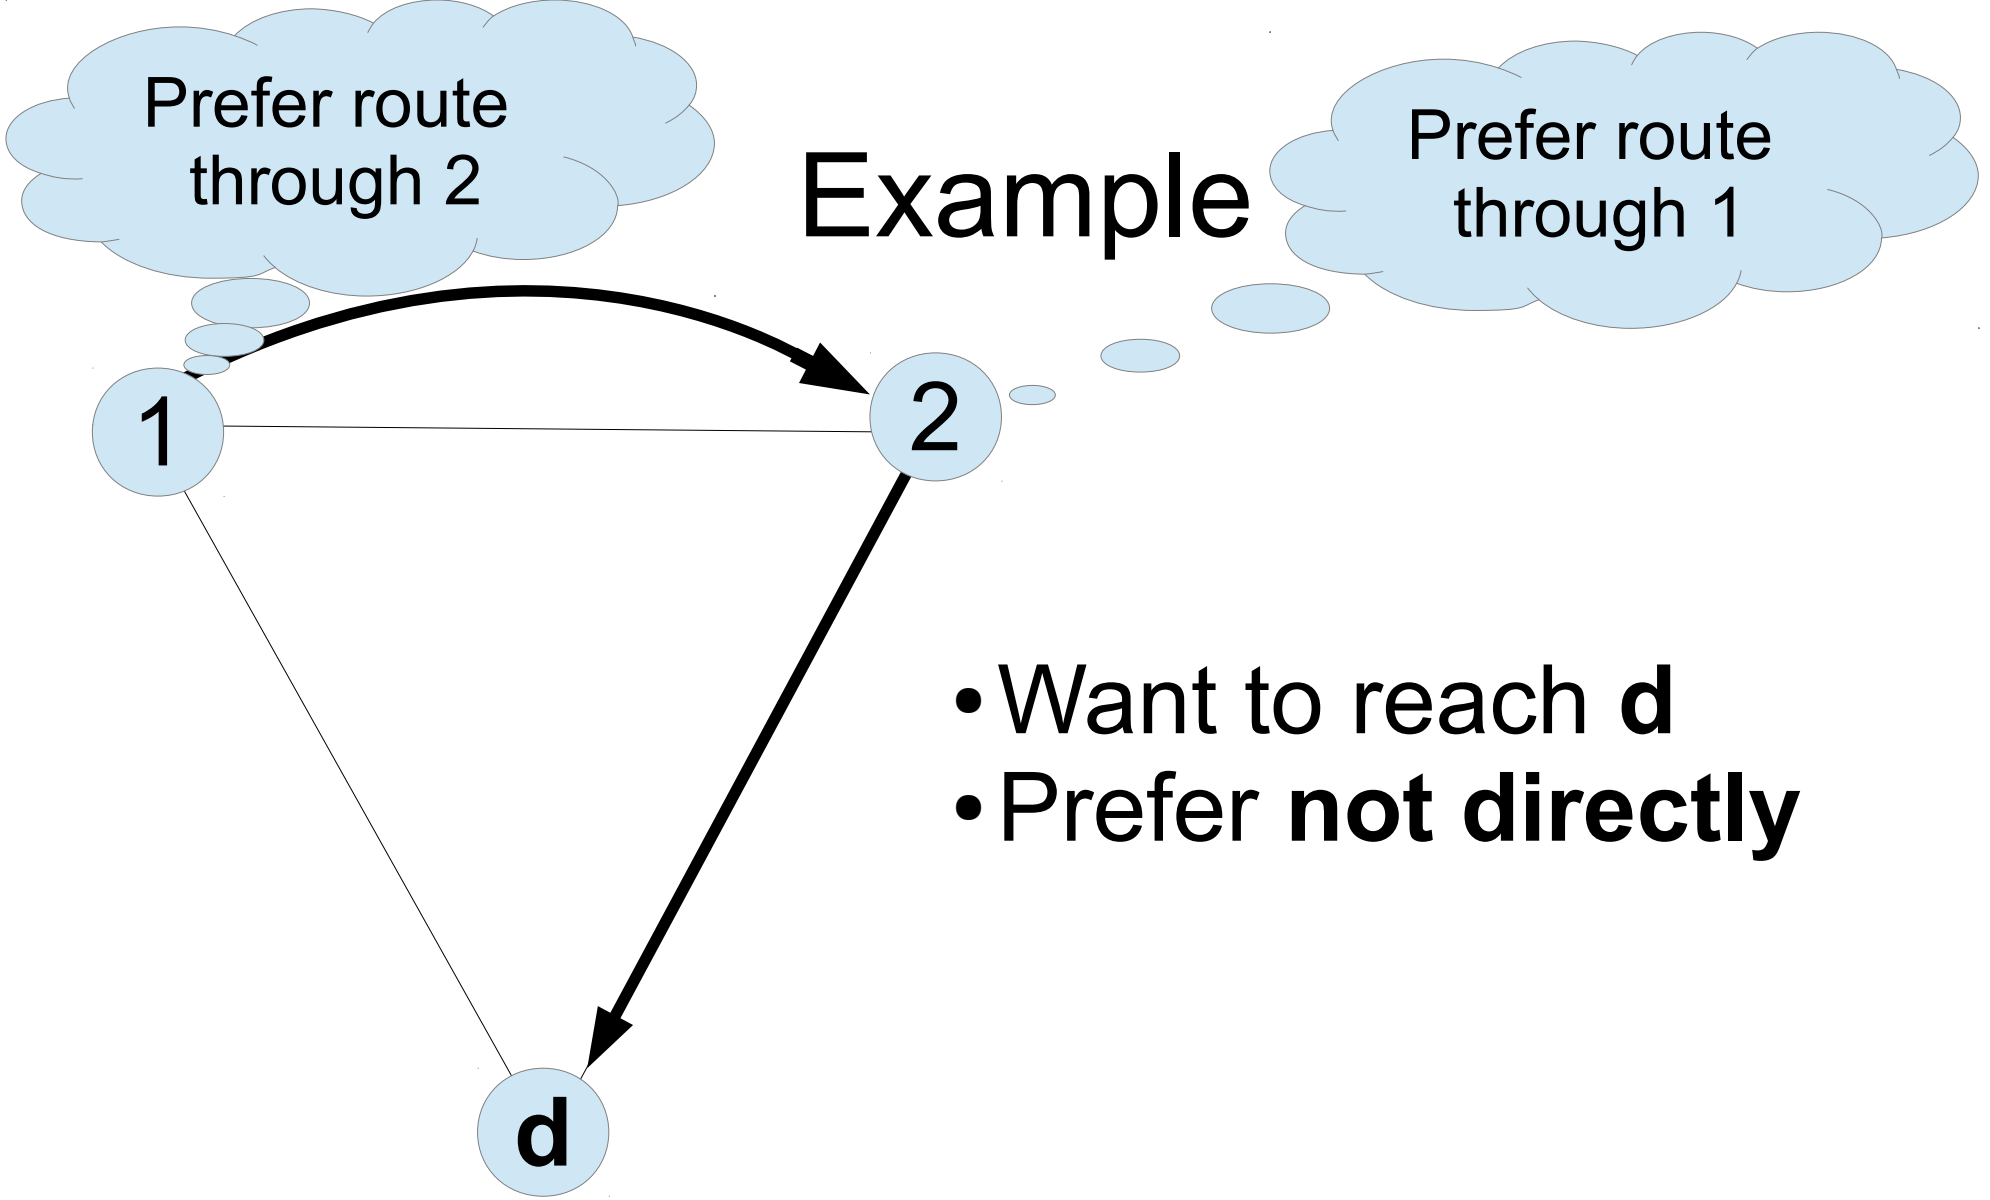

- Want to reach **d**
- Prefer **not directly**

(Nash) Equilibrium: no reason to change

Example

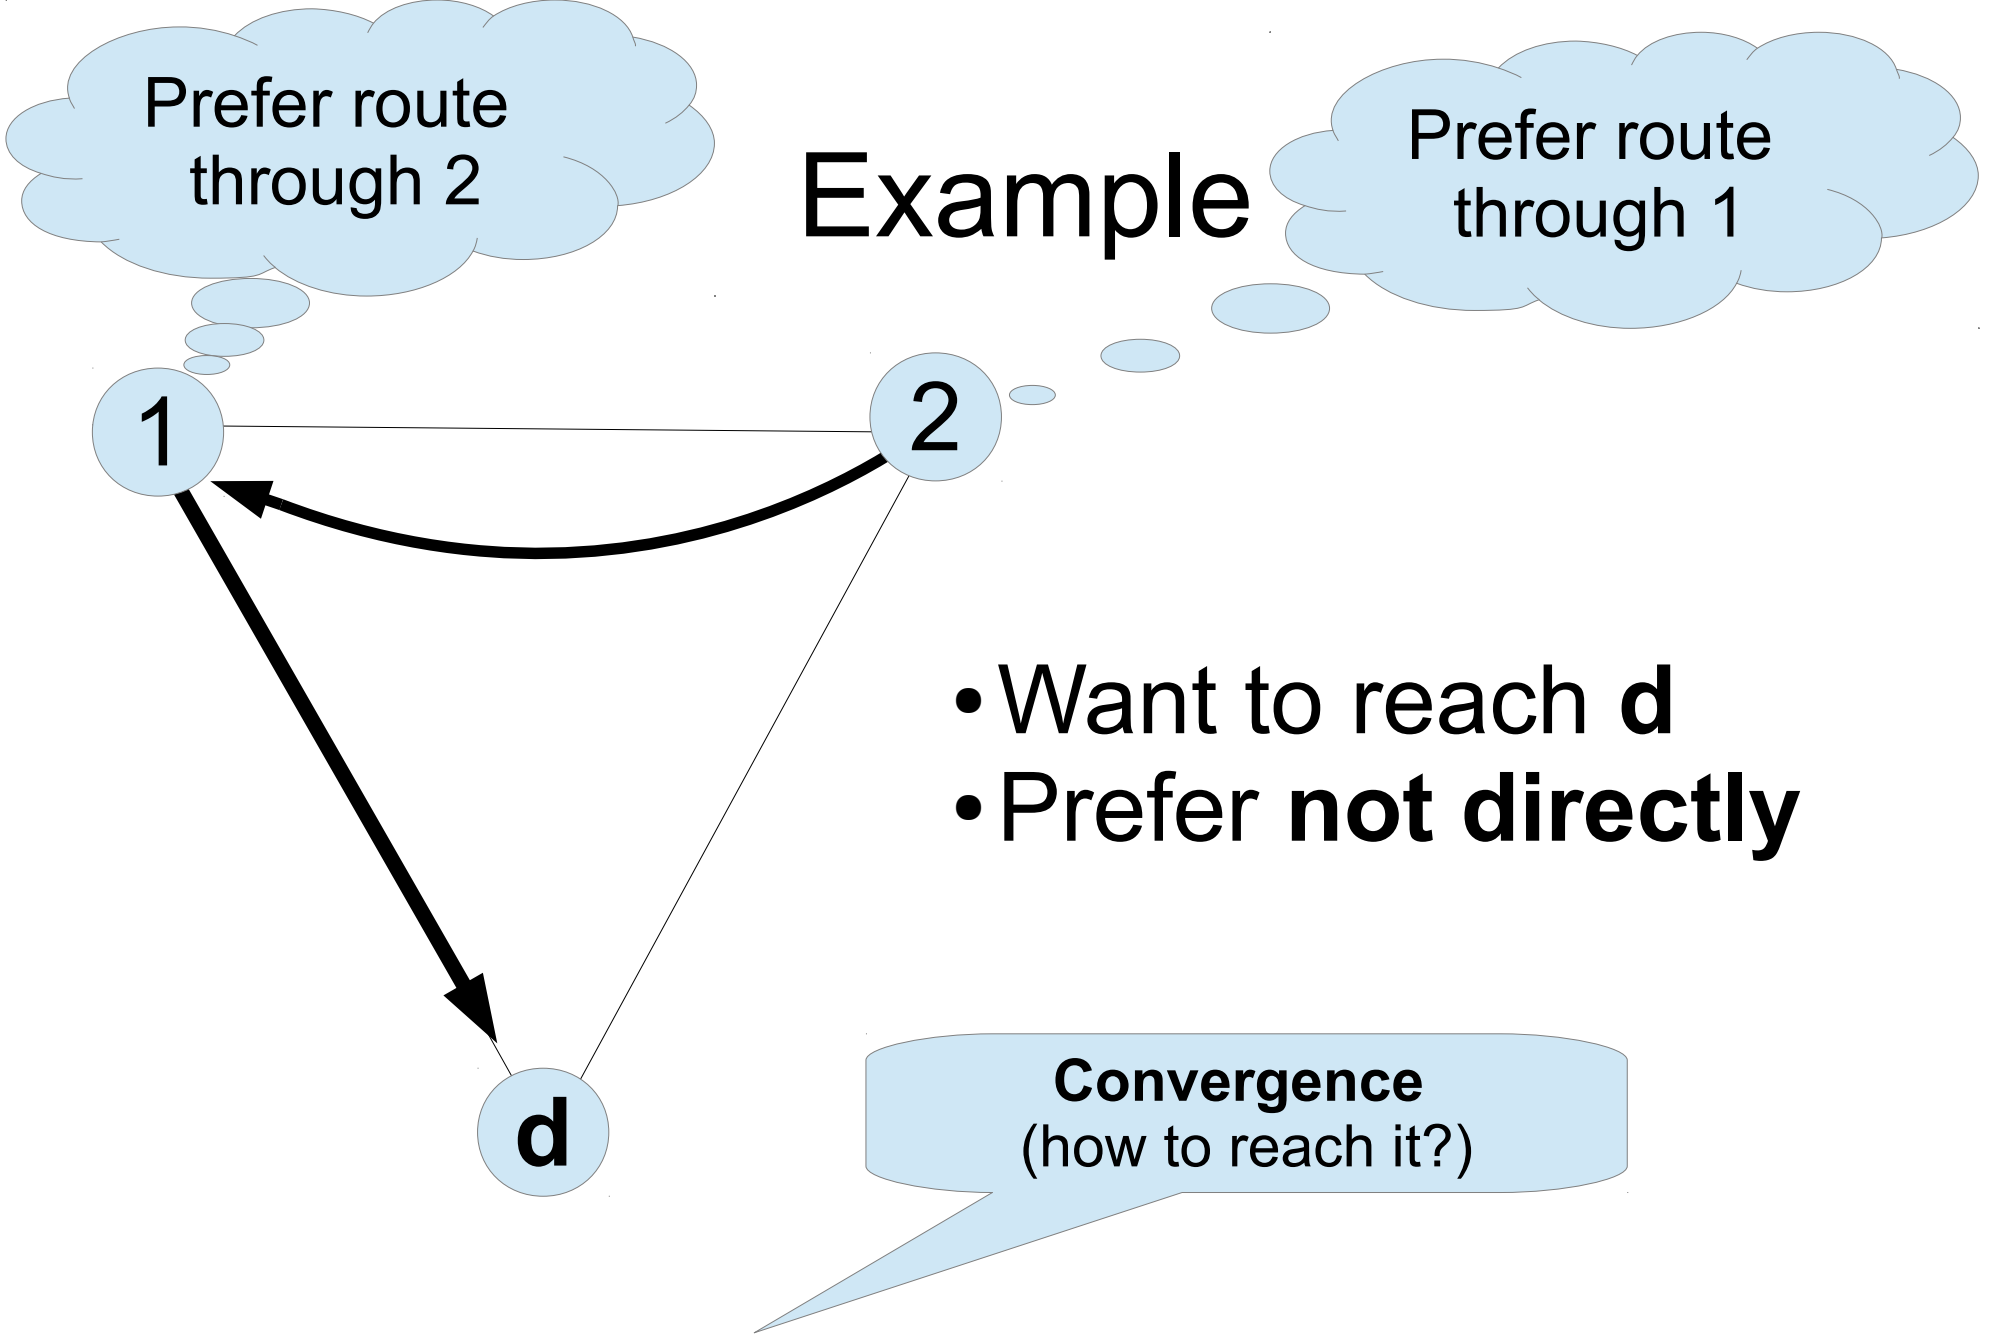

- Want to reach **d**
- Prefer **not directly**

Convergence
(how to reach it?)

(Nash) Equilibrium: no reason to change

Example

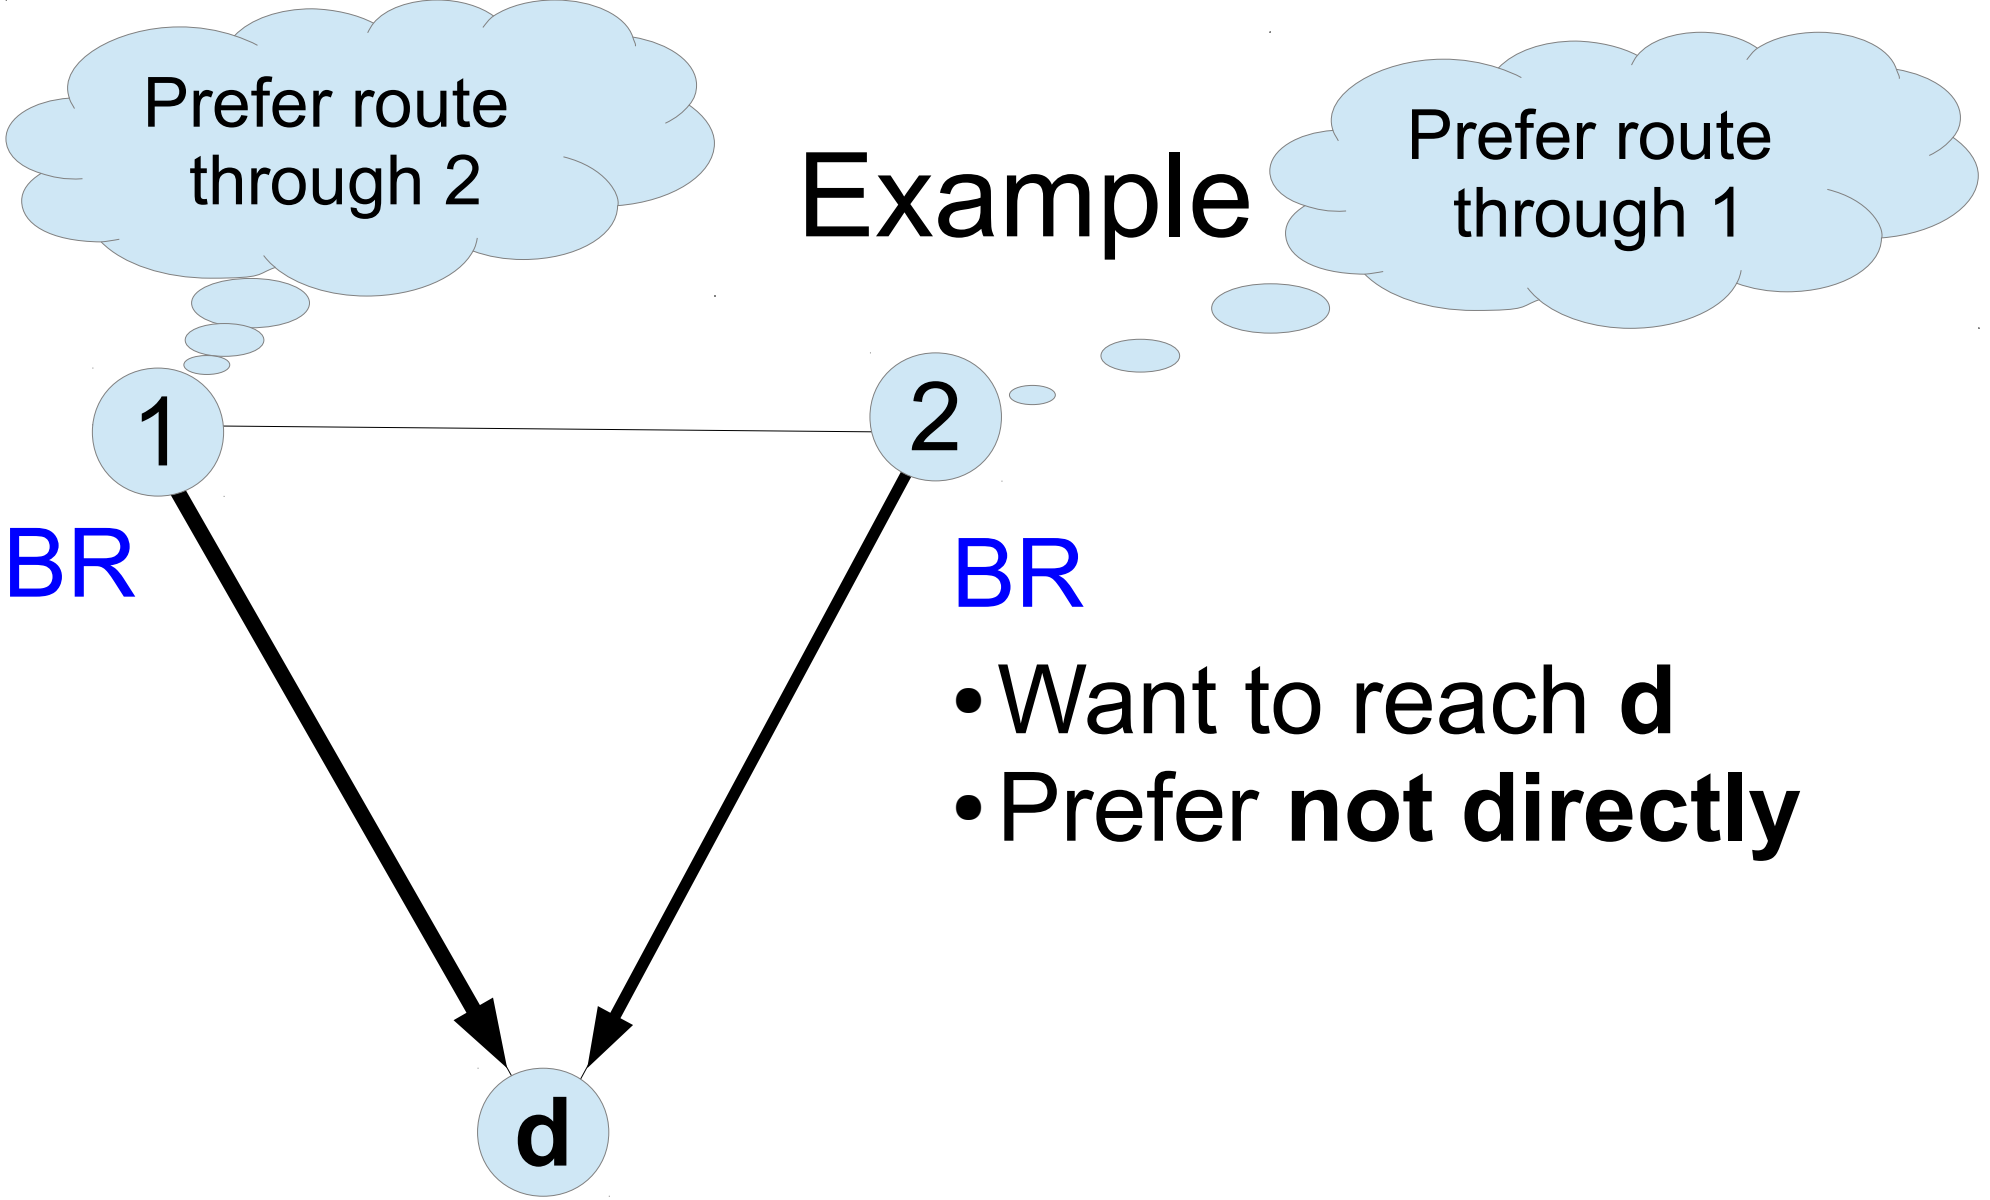

- Want to reach **d**
- Prefer **not directly**

Unstable

Example

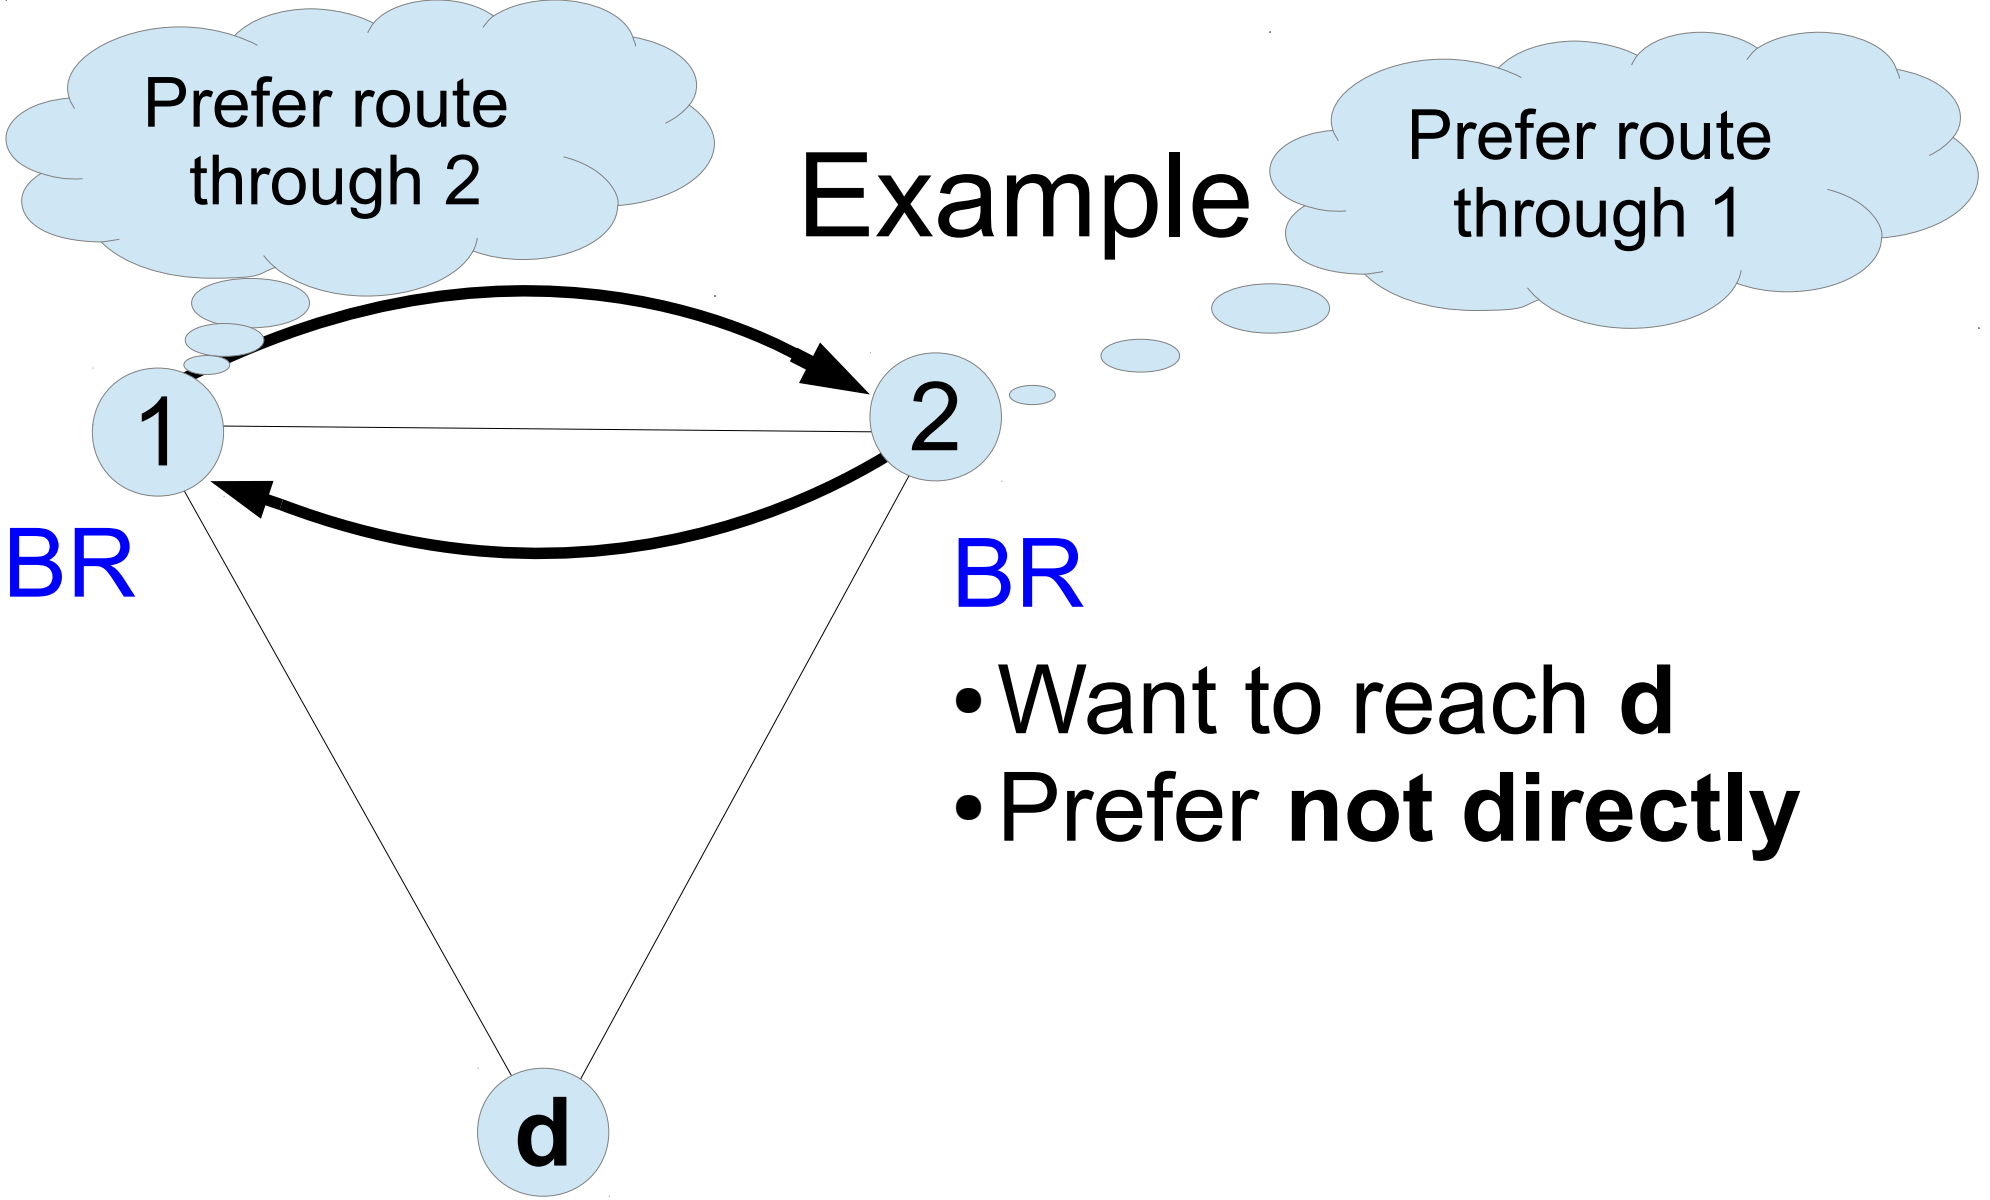

- Want to reach **d**
- Prefer **not directly**

Unstable

Example

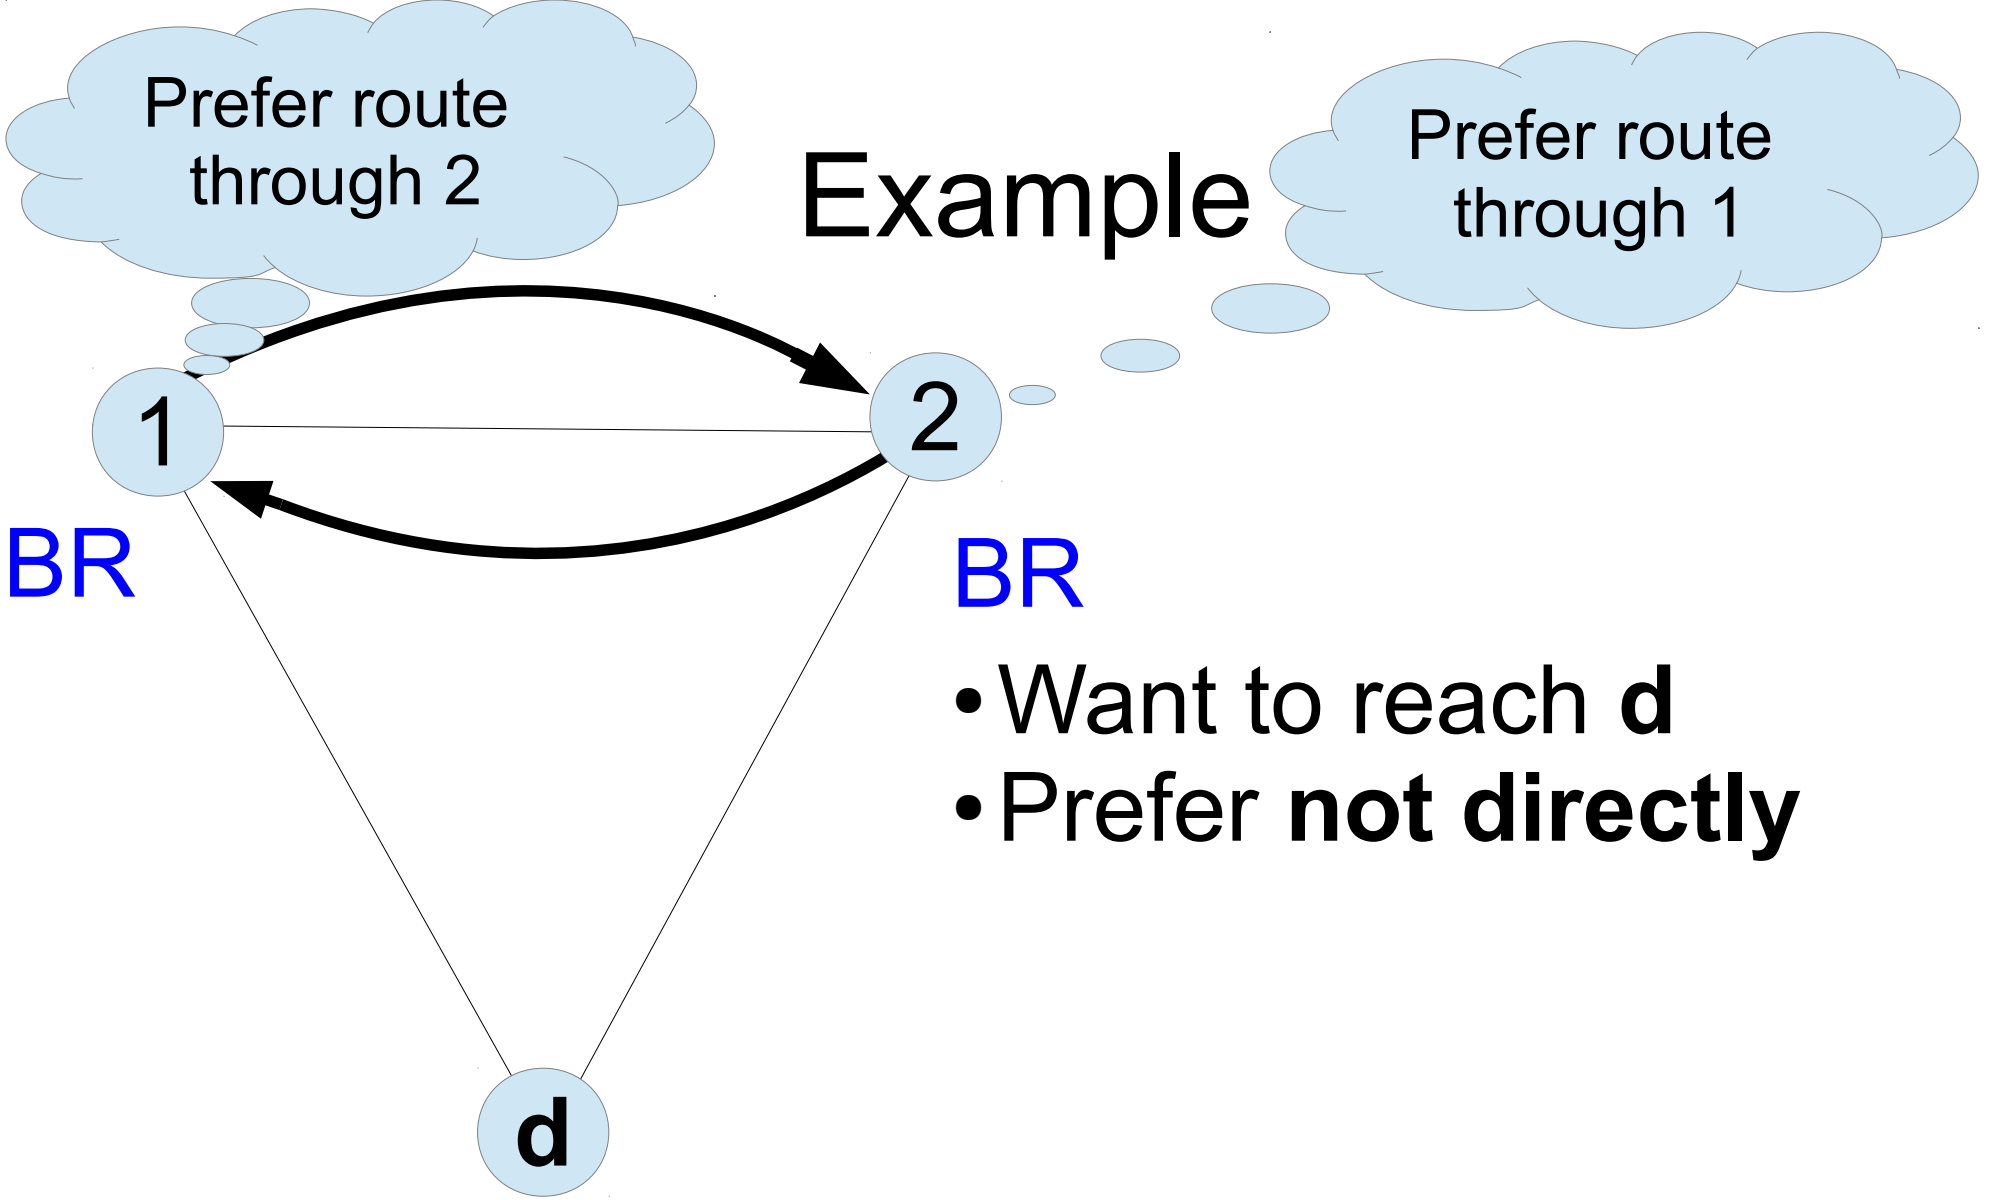

- Want to reach **d**
- Prefer **not directly**

No convergence!!

BGP does not work...

..yes it does!! (Gao-Rexford'01)

When do best-respond converge? (self-stabilization)

Converge?

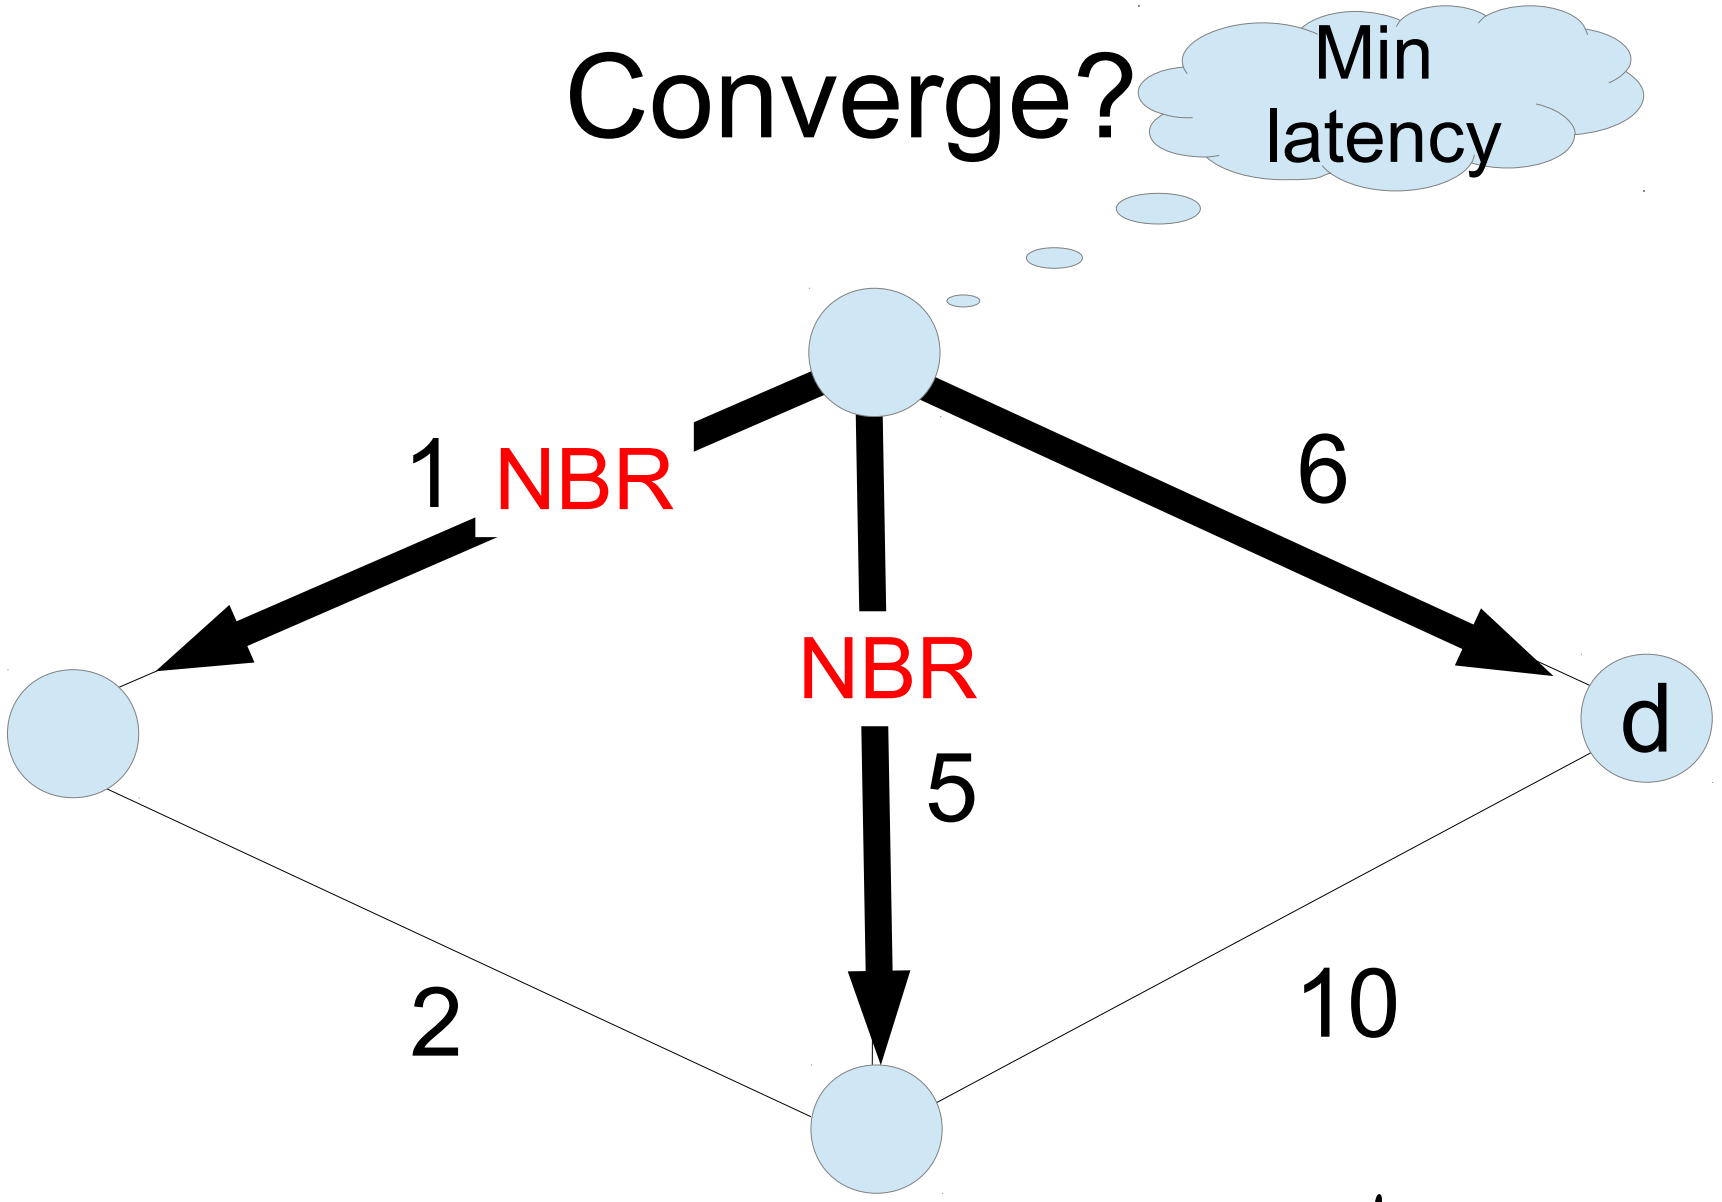

NBR = Never Best Response

Converge?

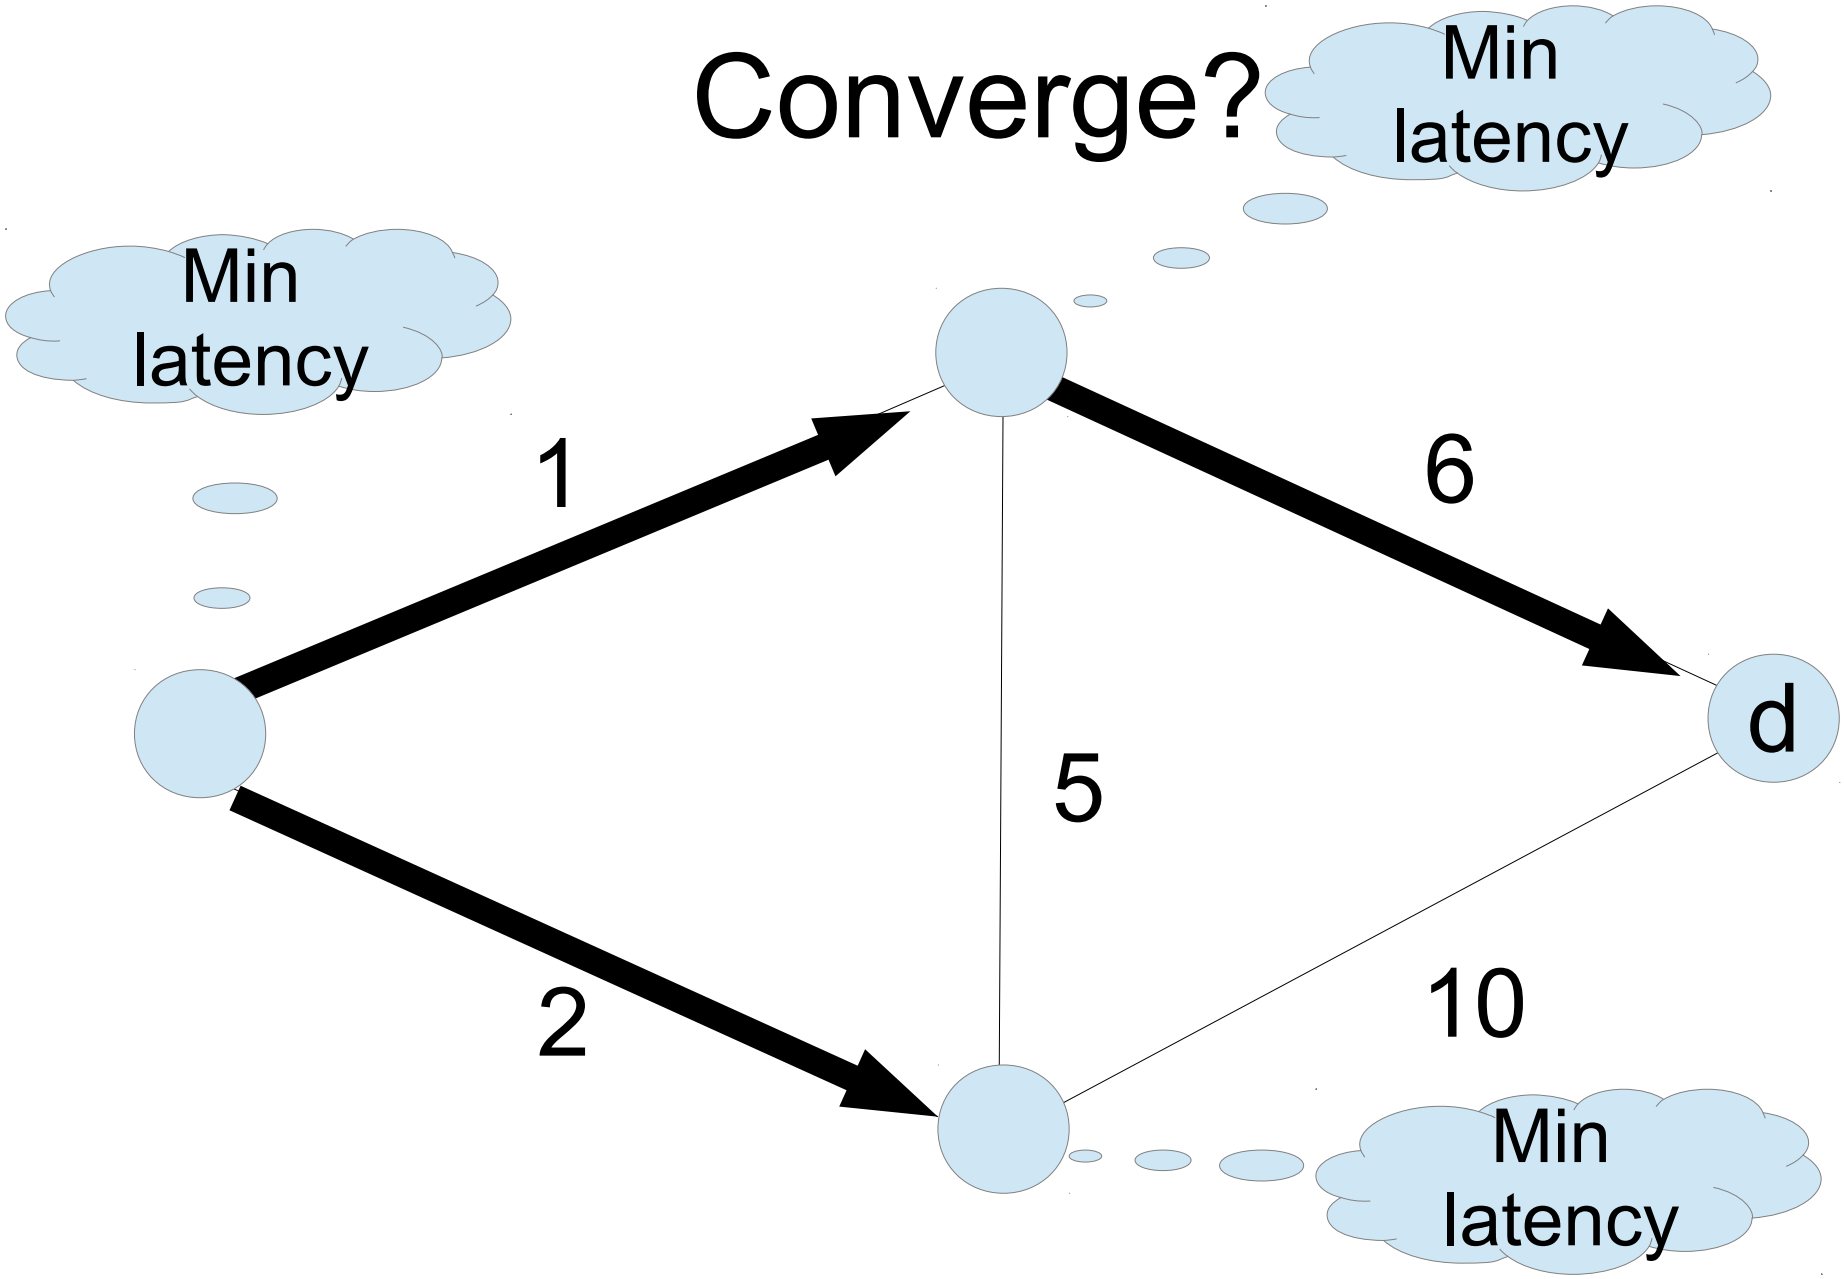

NBR = Never Best Response

Converge?

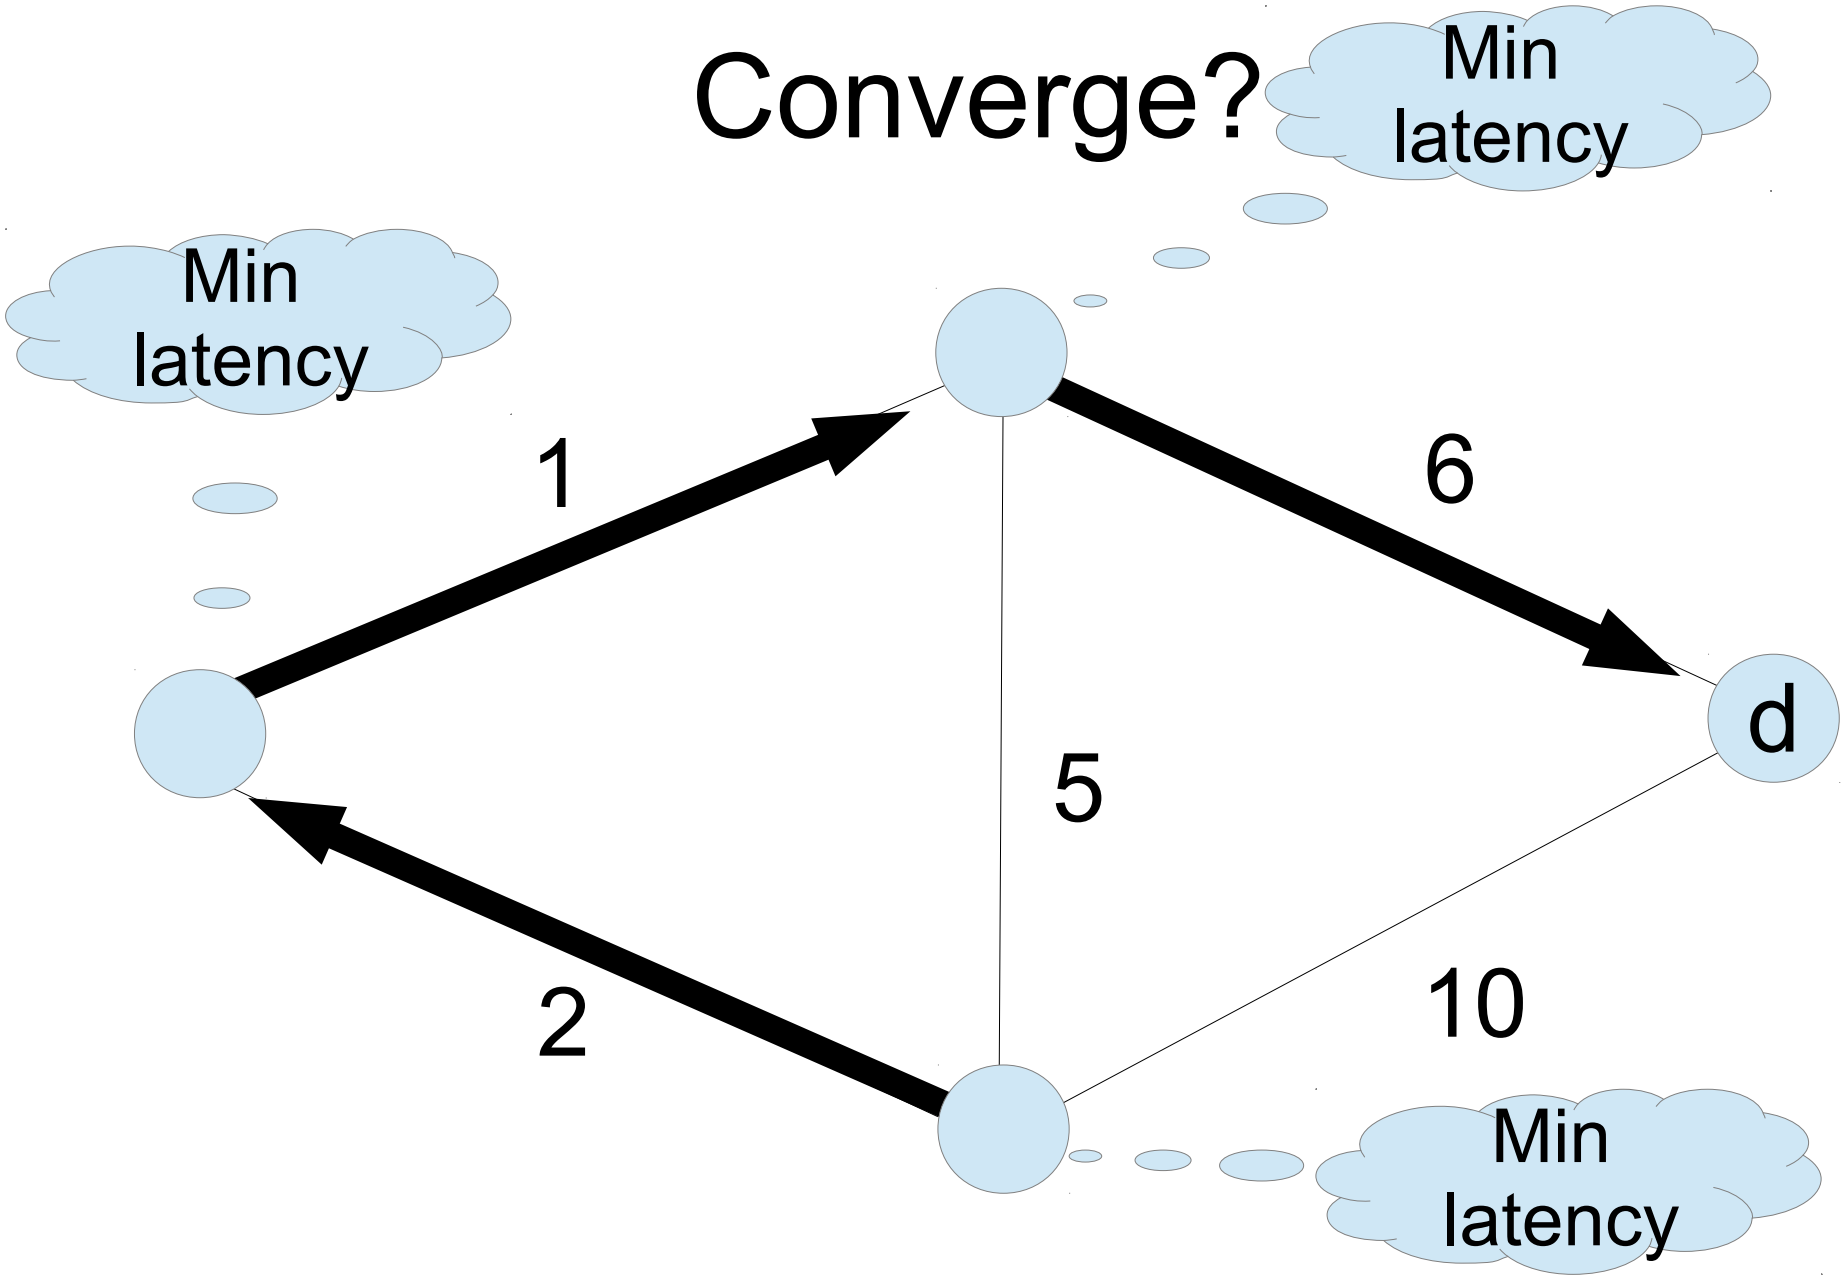

NBR = Never Best Response

Converge?

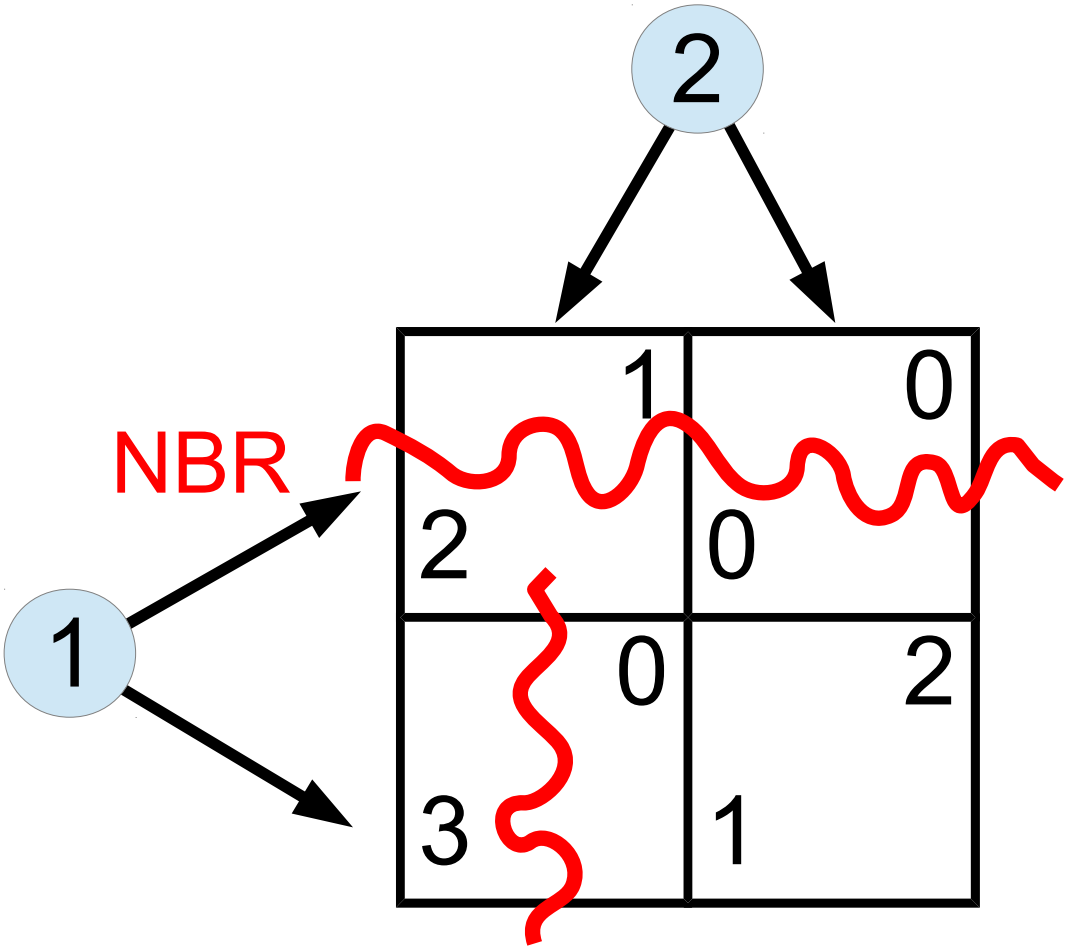

In 2 rounds:
equilibrium

Converge?

Adversary: Initial state and activation sequence

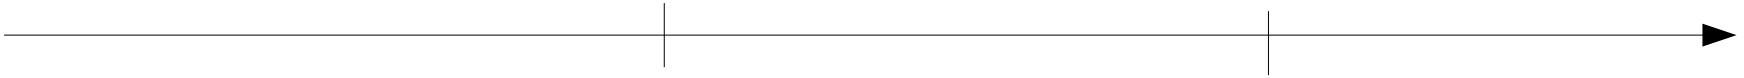

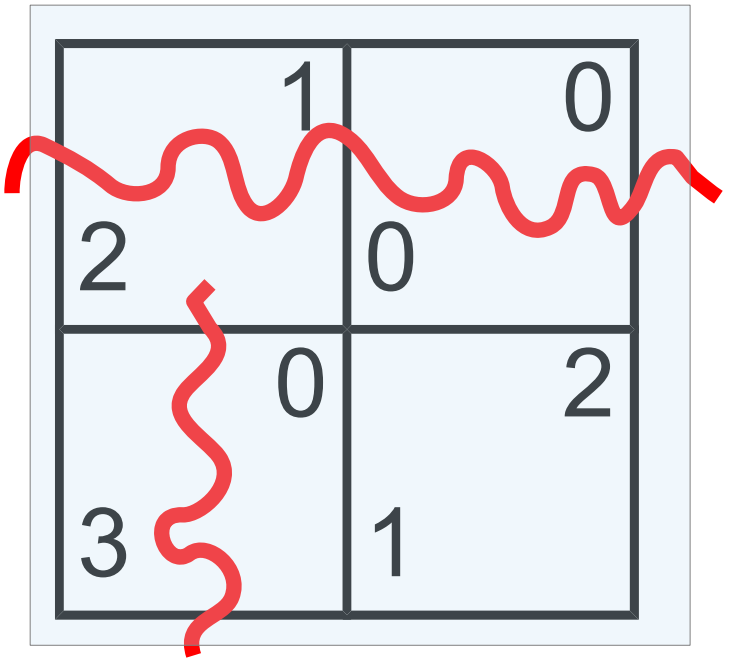

In 2 rounds: equilibrium

Round: All players activated at least once

Converge?

Adversary: Initial state and activation sequence

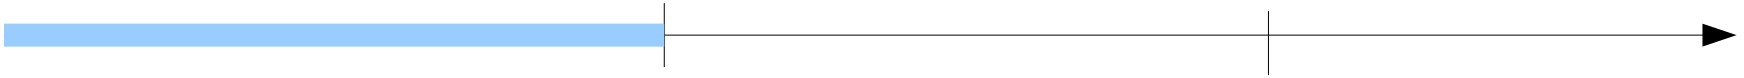

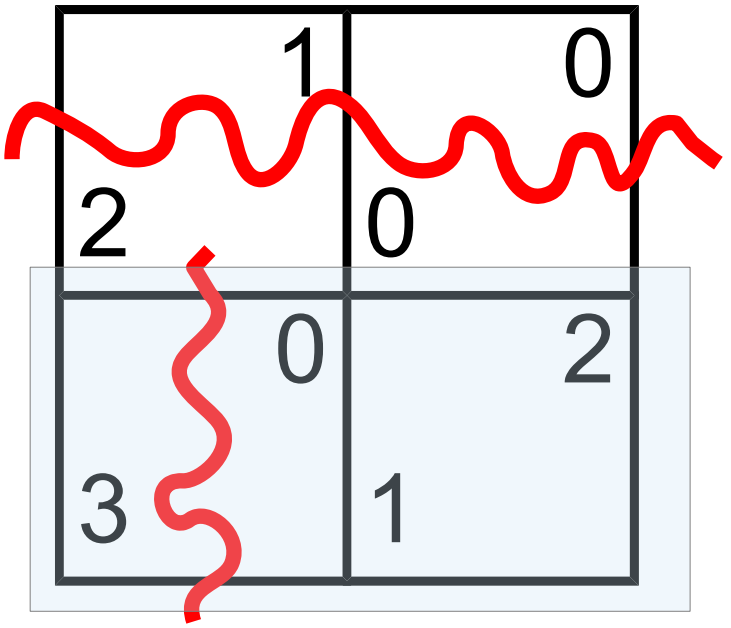

In 2 rounds: equilibrium

Round: All players activated at least once

Converge?

Adversary: Initial state and activation sequence

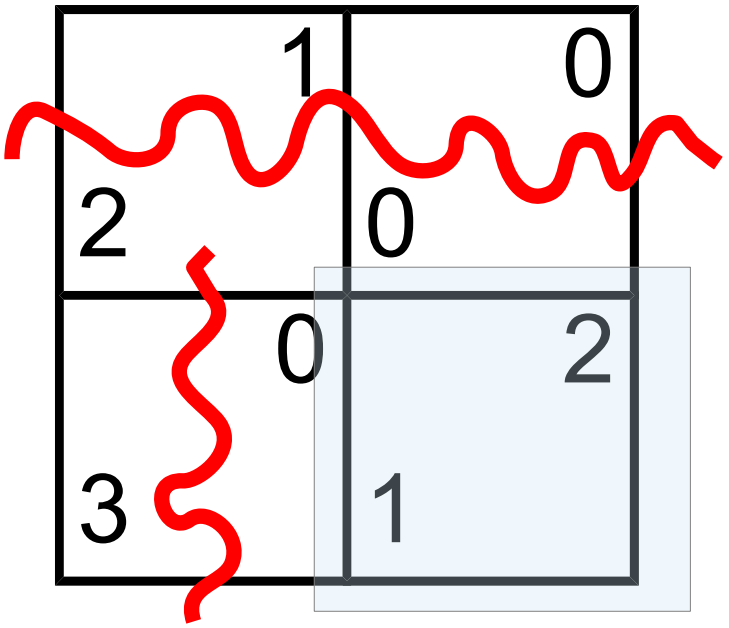

In 2 rounds: equilibrium

Round: All players activated at least once

Convergence

For NBR-solvable games repeated best-response converge to (Nash) equilibrium

Number of rounds = number of NBR eliminations

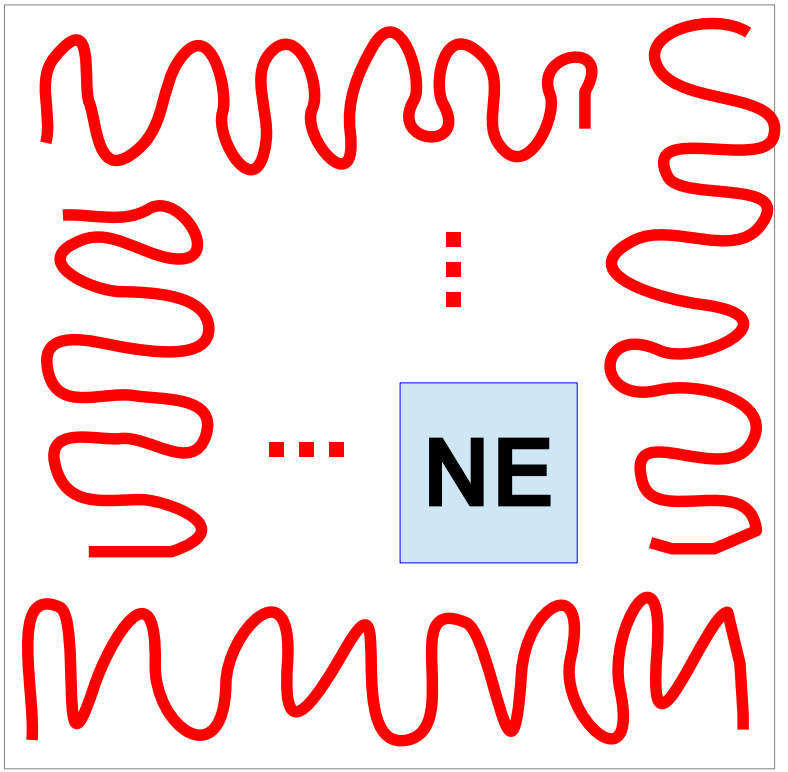

Best-Response Mechanisms
(Nisan et al 2011)

Protocol: “repeatedly best respond” (greedy)

Asynchronous setting (adversarial schedule)

Convergence for NBR-solvable games.

- Many applications (BGP, TCP-games, Intern-Hospital Matching, Auctions)

Incentive Compatibility

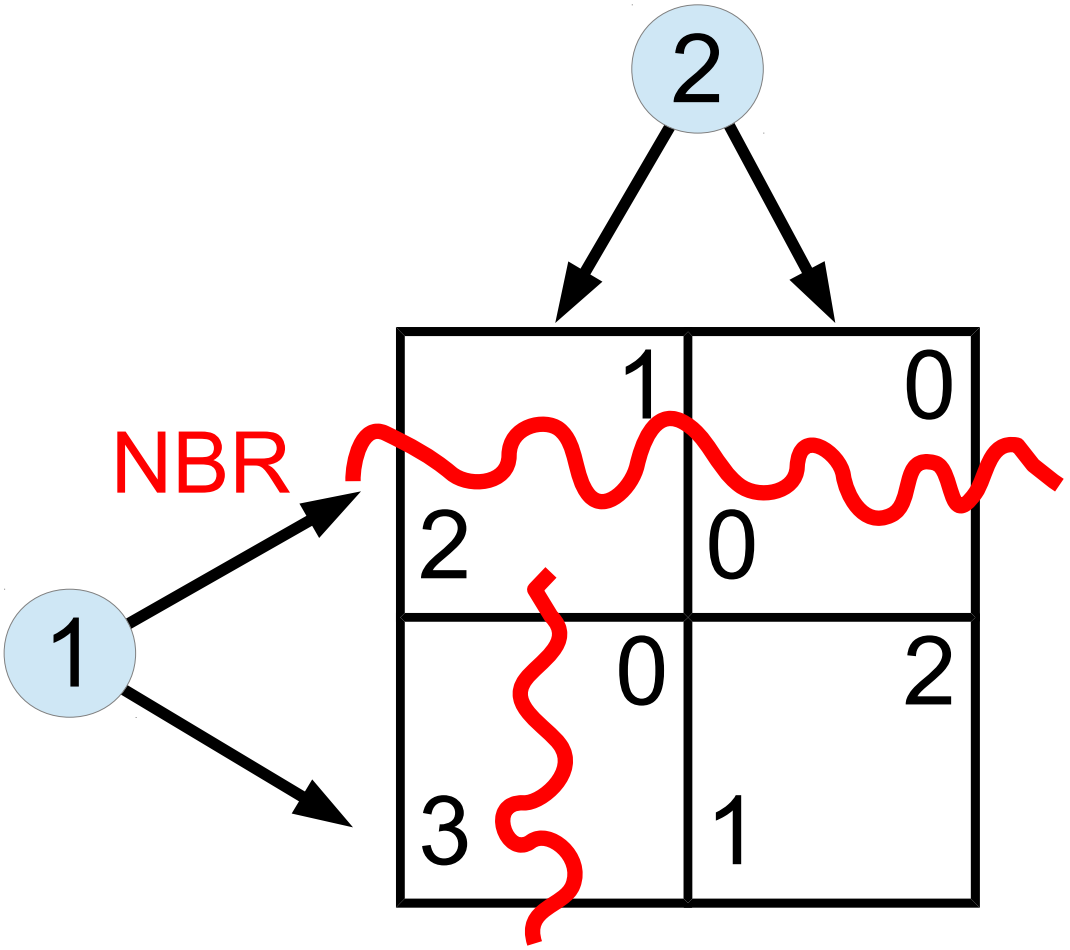

Incentive Compatibility

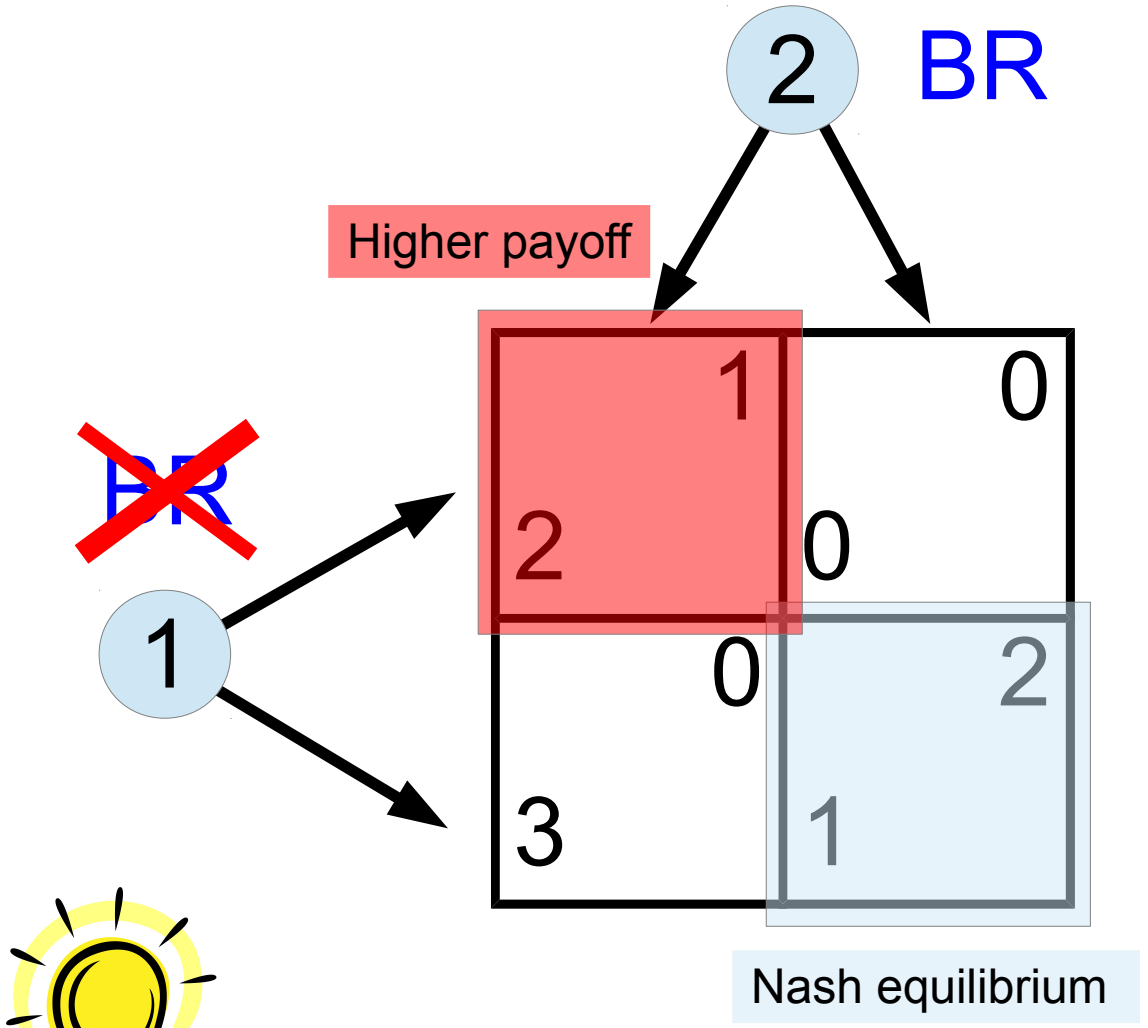

Clear outcome: payoff Nash at least payoff “discarded”

Best-Response Mechanisms
(Nisan et al 2011)

Protocol: “repeatedly best respond” (greedy)

Asynchronous setting (adversarial schedule)

Convergence and Incentive Compatibility together (NBR-solvable games with clear outcome).

- Many applications (BGP, TCP-games, Intern-Hospital Matching, Auctions)

Part 2
Mistakes and faults...what happens?

Imperfect Best-Response Mechanisms

(Ferraioli & Penna, 2013)

Protocol: “repeatedly best respond” (greedy)

Asynchronous setting (adversarial schedule)

Each time they respond, with small probability *p* do something else

Length of round is R with “good” probability

What if players sometimes take wrong decision?

Are these protocol robust to faults?

Convergence with mistakes

Convergence: Reach the (Nash) equilibrium with "good" probability

Obs: Probability *p* must be small enough...

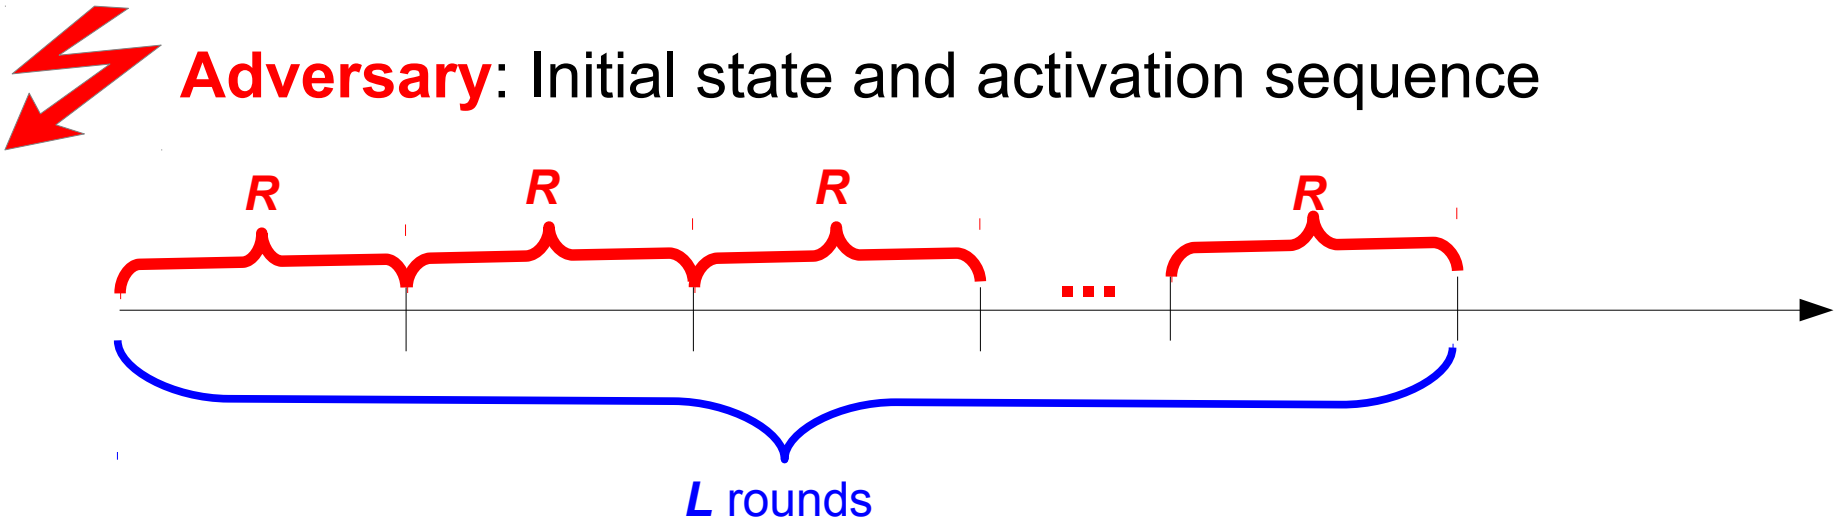

Convergence with mistakes

Convergence: Reach the (Nash) equilibrium with "good" probability

Obs: Probability *p* must be small enough...

Thm (Lower Bound). Even for deterministic non-adaptive adversary, convergence may require *p* exponentially small in the number of players

Proof of Lower Bound

Game with a **fragile** equilibrium

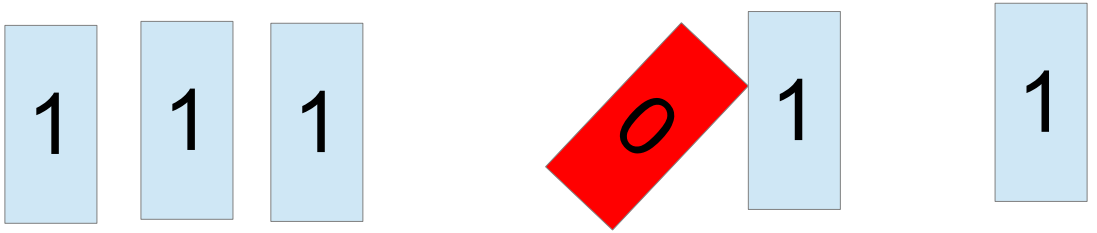

Adversary:

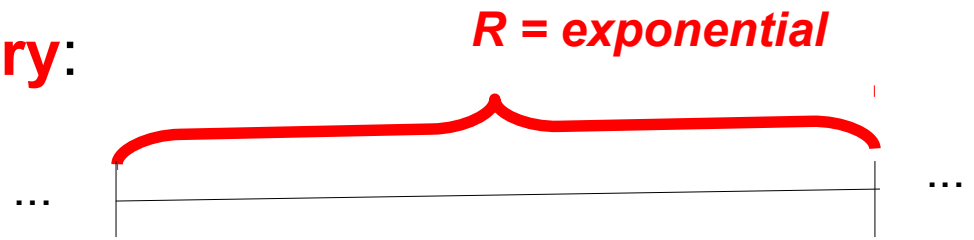

...first the game

$$\text{payoff}_i(S_1,\ldots,S_i,\ldots,S_n) = \begin{cases} 1 & \text{if } S_i = \text{AND}(1,S_1,\ldots,S_{i-1}) \\ 0 & \text{otherwise} \end{cases}$$

Play 1 if all before you play 1
0 otherwise

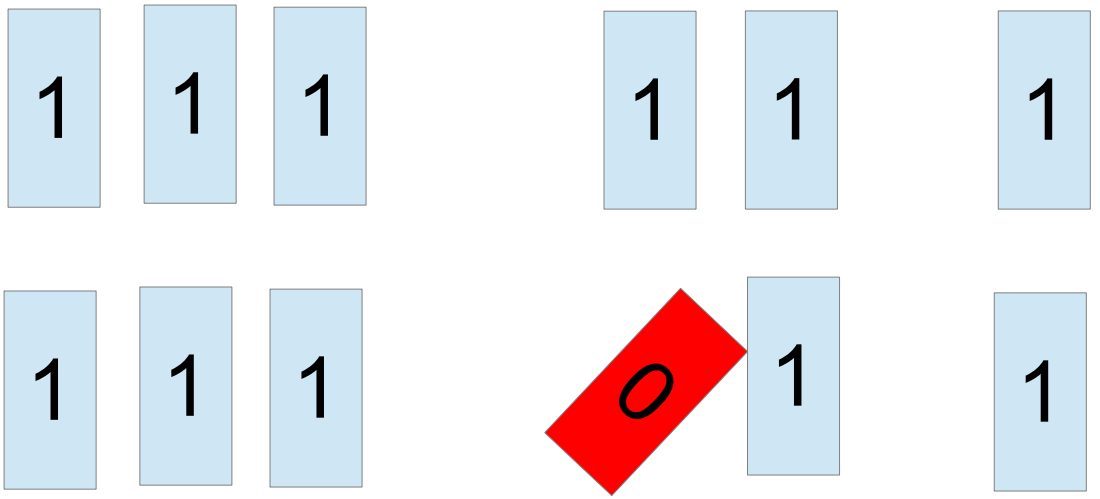

Proof of Lower Bound

Adversary:

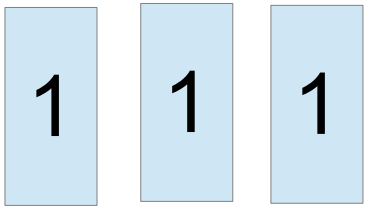

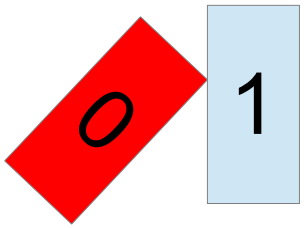

1

12

1213

12131214

1213121412131215

121312141213121512131214121312 16

$p \leq 1/R$

$R = 2^{n-1}$

Lower Bound

For some games, there is an adversary, such that convergence requires *p* exponentially small

$$p \leq 1/R = 1/2^{n-1}$$

Upper Bound

For convergence always enough *p* small in the inverse of "total time"

$$p \leq 1/(mRL)$$

An application to BGP

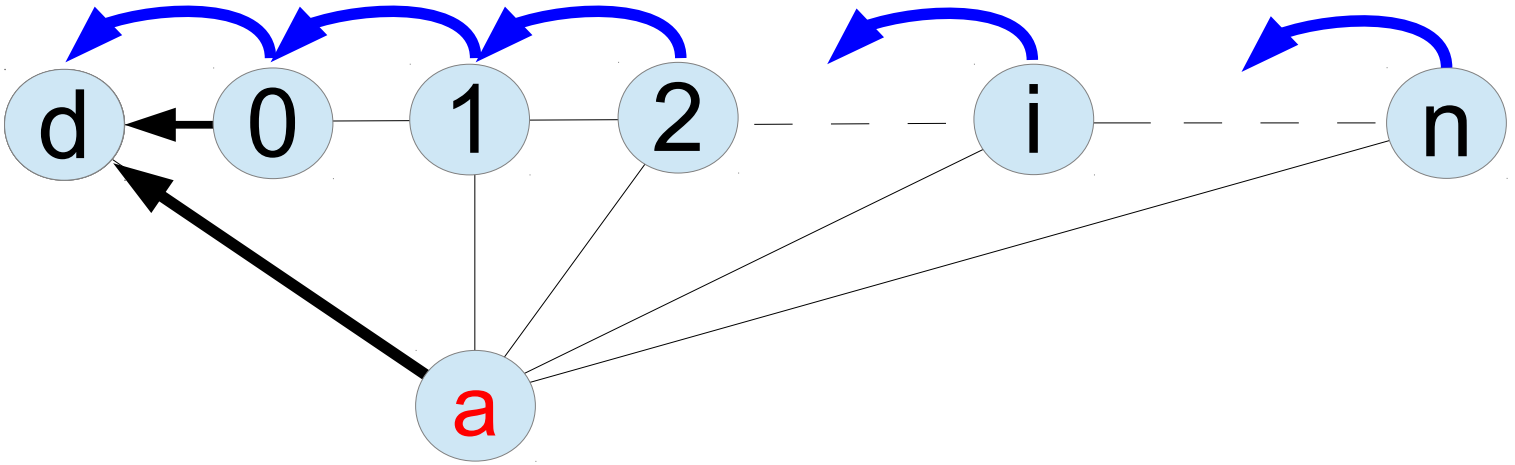

Lower bound applies to “real” BGP instances (Gao-Rexford model)

- No faults (p=0): BGP converges and Incentive Compatible (Levin,Schapira, Zohar'11)(Nisan et al'11)
- Faults (p>0): BGP does not converge unless p exponentially small

Incentive Compatibility and mistakes

- Need stronger condition
- Some games are not robust (TCP games)

Open Questions

Tighter bounds (specific games, adversaries)

More general games (restricted dynamics)

"Equilibria selection"

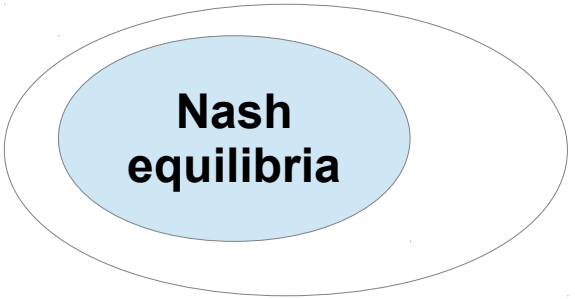

Two-layer games

routing → BGP

Commercial relationships → Gao-Rexford

Thank You!