Reliable Inference at Scale using Graph Structure

15/04/2026 @ IISc/CNI Seminar

Mansi Sood

Schmidt Science Fellow

MIT LIDS

PhD, ECE @CMU

B. Tech, M. Tech, IITB

A.G. Jordan Award for *Outstanding PhD Thesis* '25

Info Theory & Applications Graduation Day Award '24

Networks in context: How representation informs what we can answer?

What are the nodes and edges?

Networks in context: How representation informs what we can answer?

Distributed Systems

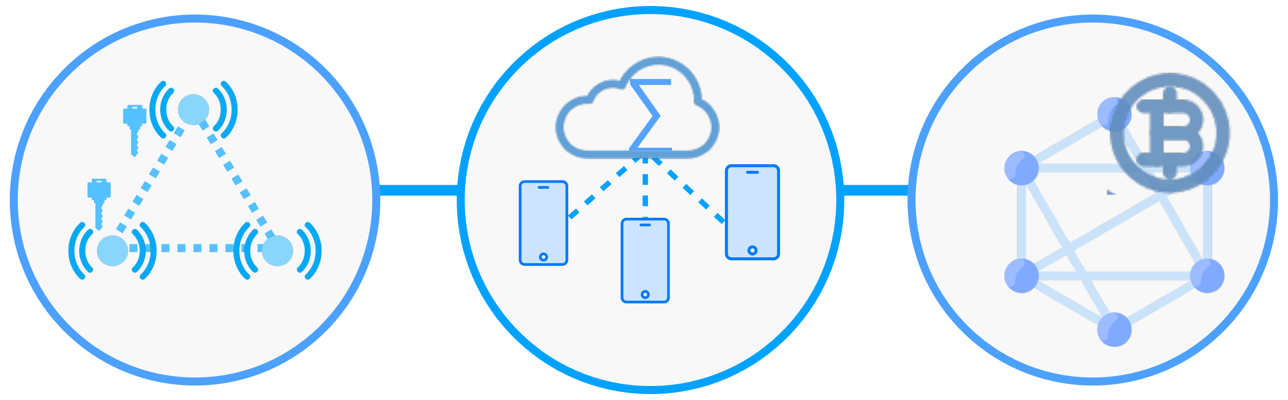

Probabilistic Graphical Models

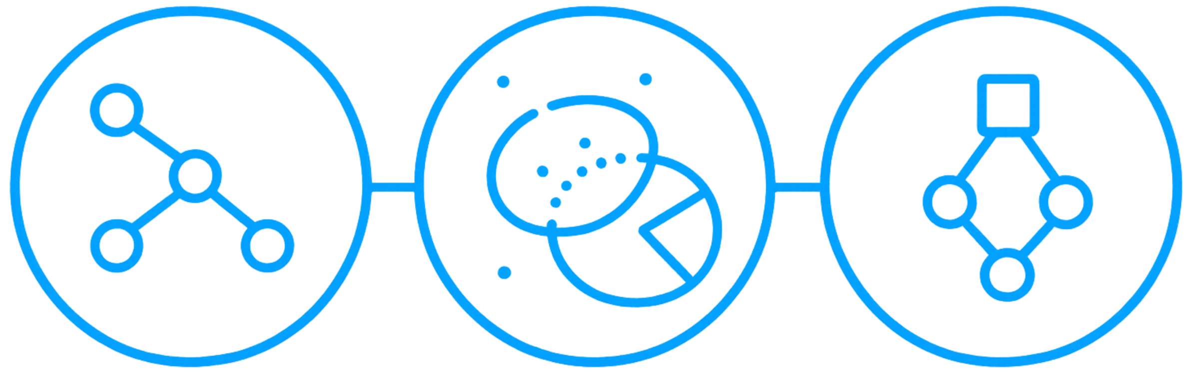

Social Networks

Networks in context: How representation informs what we can answer?

Distributed Systems

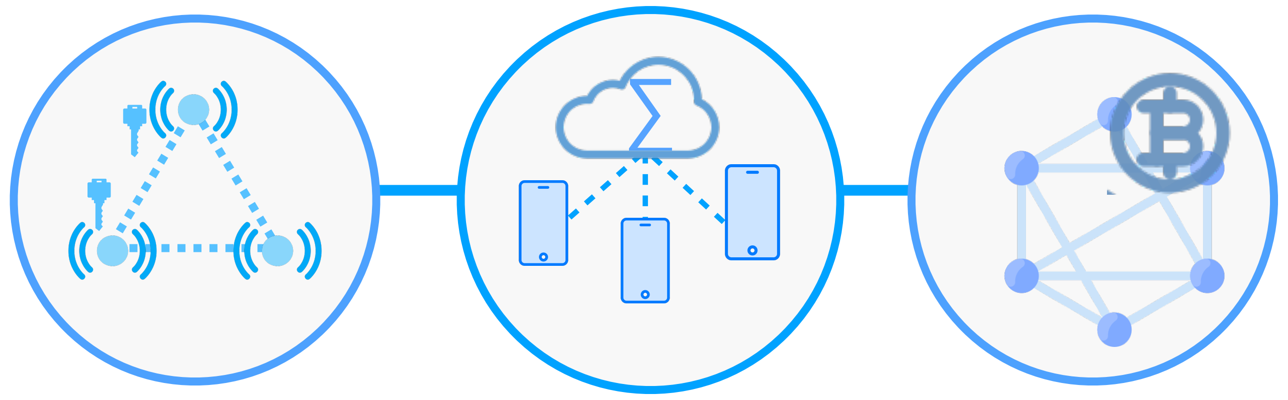

'secure' communication channel between devices

Probabilistic Graphical Models

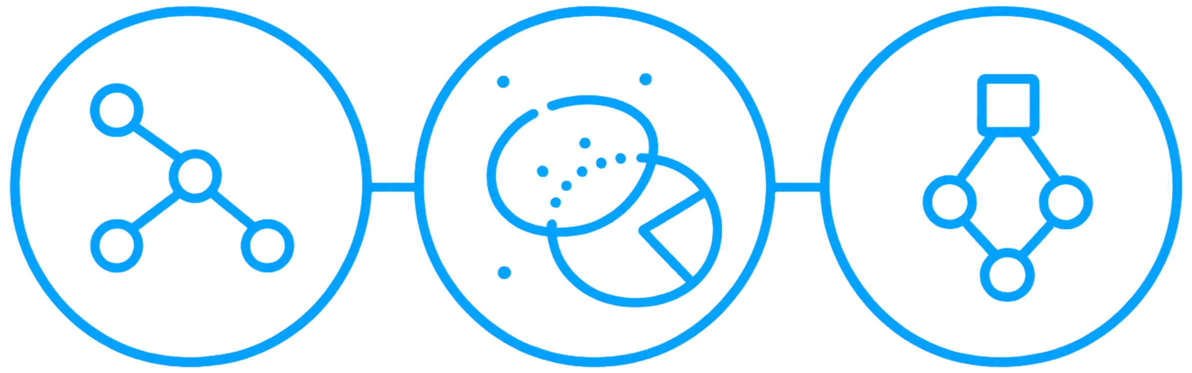

conditional dependencies in observational data

Social Networks

pathway for contagion to travel between people

Networks in context: How representation informs what we can answer?

Setting:

Distributed Systems

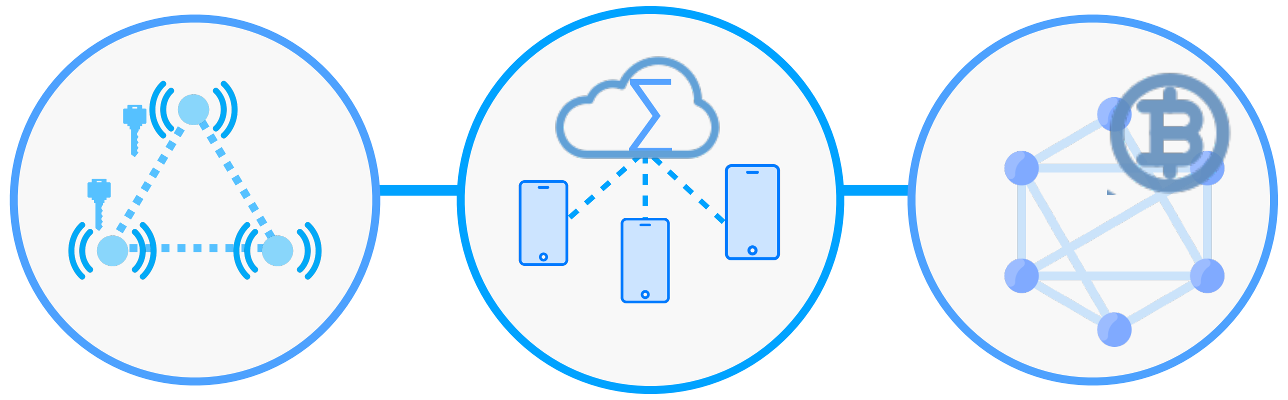

'secure' communication channel
between devices

Probabilistic Graphical Models

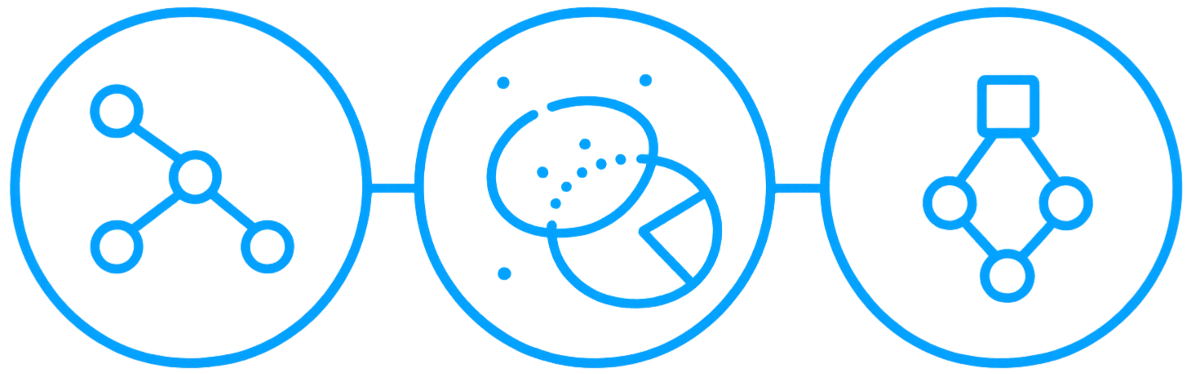

conditional dependencies
in observational data

Social Networks

pathway for contagion to travel
between people

Question:

[network design]

[distribution learning]

[dynamical processes]

Networks in context: How representation informs what we can answer?

Distributed Systems

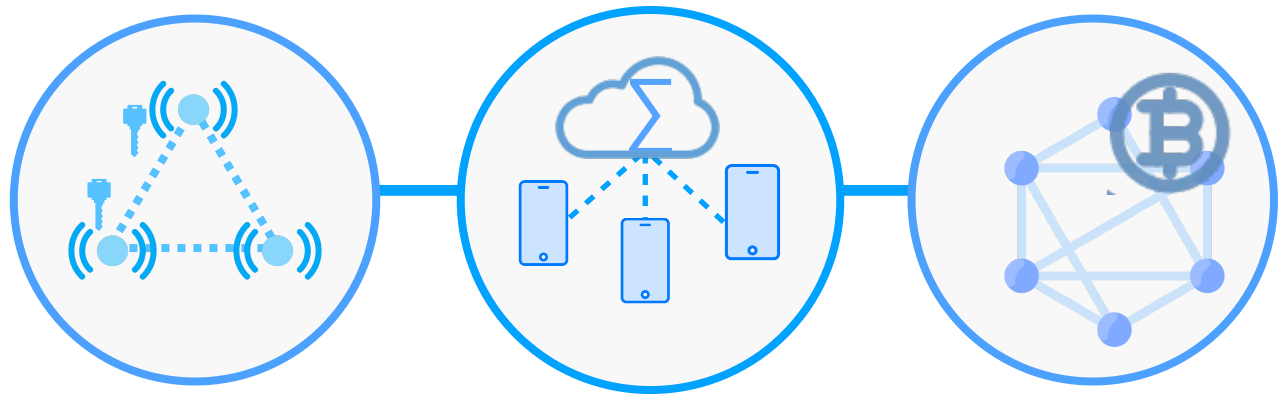

‘secure’ communication channel
between devices

[network design]

M.S., O.Y., **IEEE Trans. on Information Theory** ’23
M.S., O.Y., **IEEE Trans. on Information Theory** ’21
M.S., O.Y., **IEEE ICC** ’21 **(Best Paper Award)**
+IEEE ISIT ’21, ISIT ’20, CDC ’20, Globecom ’19

+ arXiv:2508.11863

Probabilistic Graphical Models

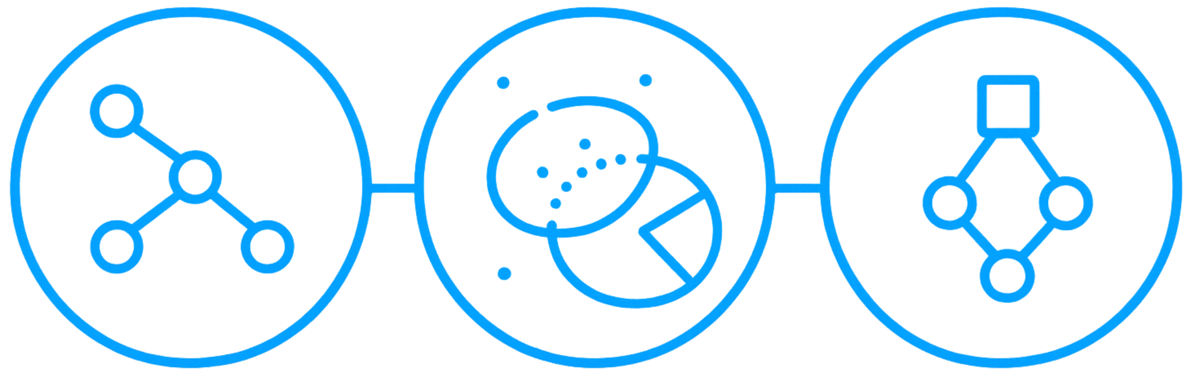

conditional dependencies
in observational data

[distribution learning]

M.S. D.S., to appear at **ISIT** ’26
M.S., D.S., Allerton, ’25

Social Networks

pathway for contagion to travel
between people

[dynamical processes]

M.S., et al., **Proceedings of the National Academy of Sciences**, ’23
+IEEE ICC ’23, Netsci/Networks ’21

+ arXiv:2409.17352

Reliable Inference at Scale Using Graph Structure

Distributed Systems

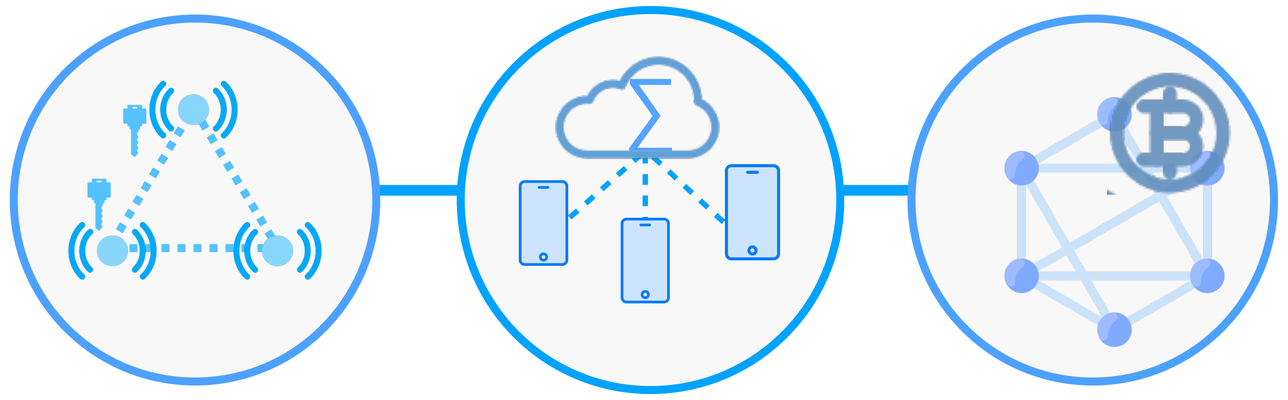

'secure' communication channel
between devices

Probabilistic Graphical Models

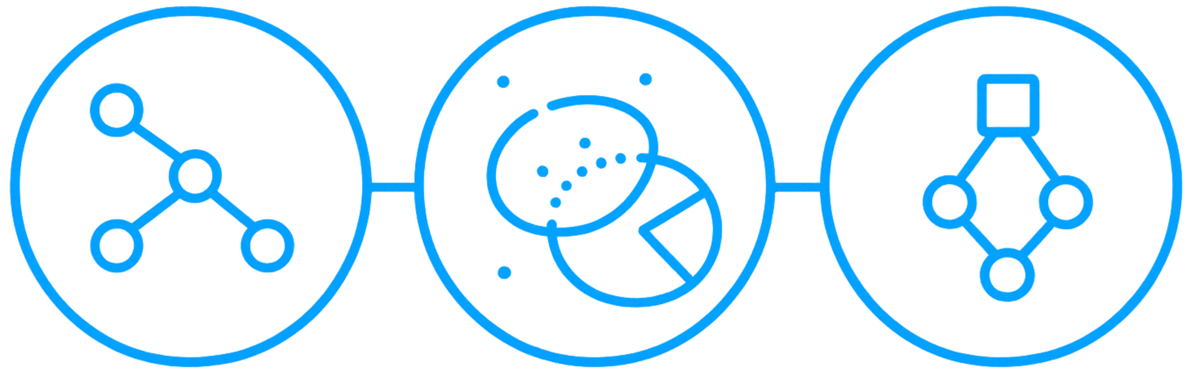

conditional dependencies
in observational data

Reliable Inference at Scale Using Graph Structure

notions of connectivity

Distributed Systems

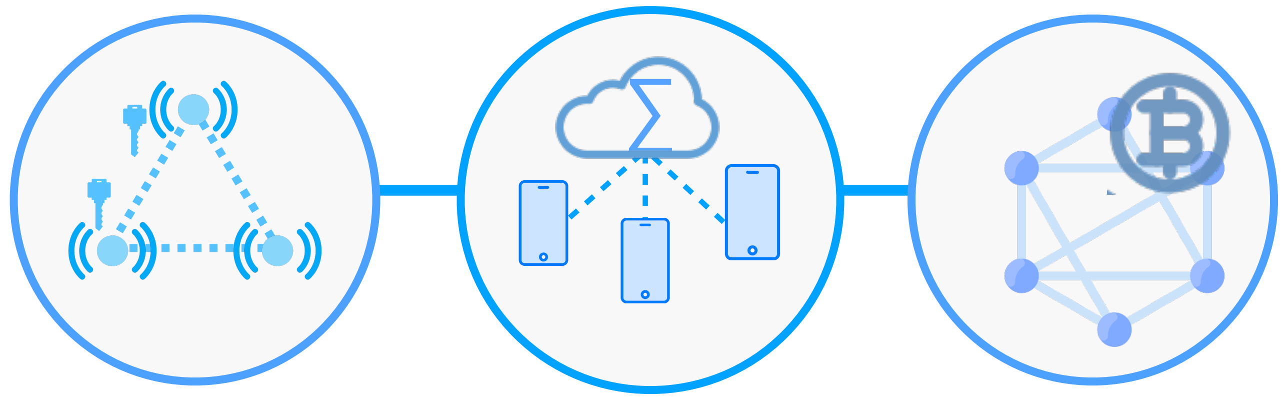

'secure' communication channel
between devices

Probabilistic Graphical Models

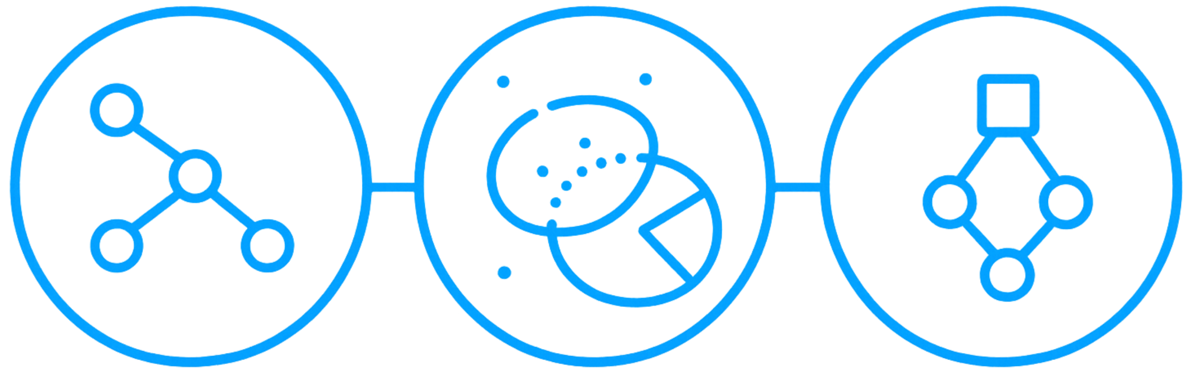

conditional dependencies
in observational data

efficiency

Reliable Inference at Scale Using Graph Structure

notions of connectivity

Distributed Systems

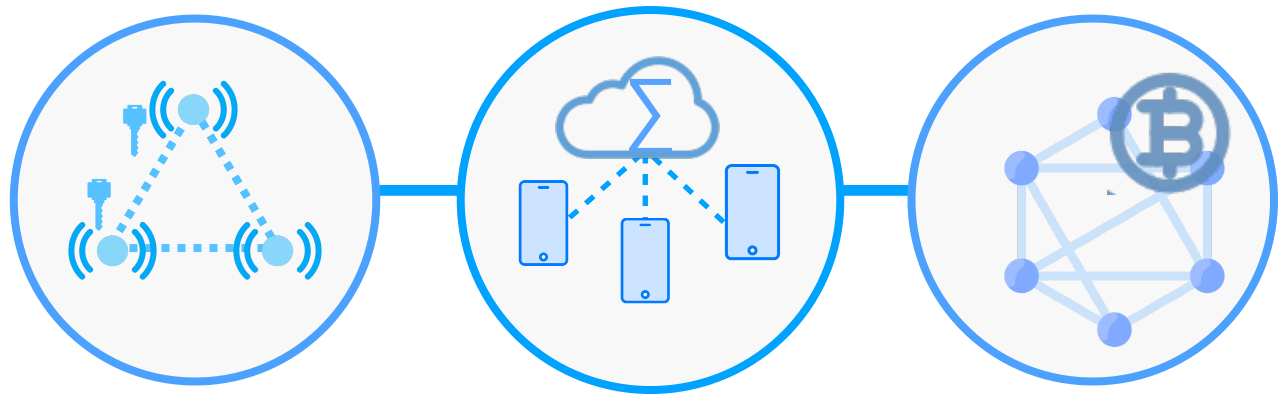

'secure' communication channel

between devices

communication

Probabilistic Graphical Models

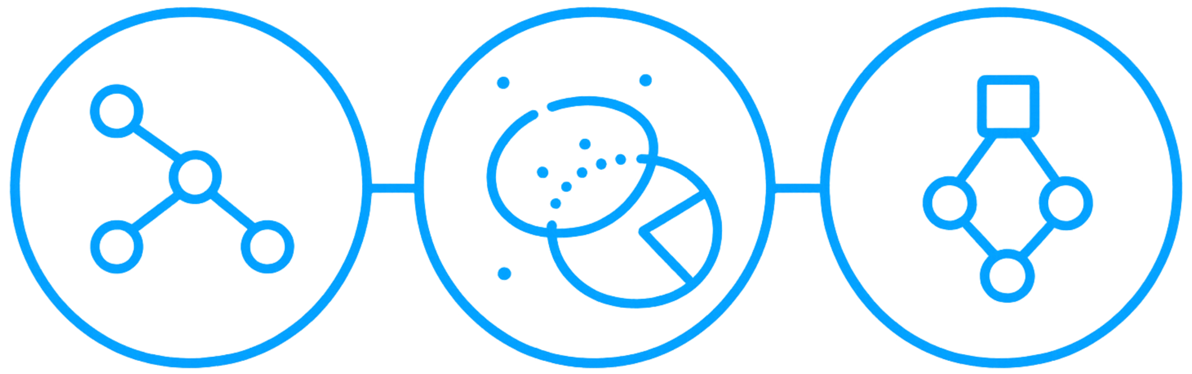

conditional dependencies

in observational data

computation

efficiency

Reliable Inference at Scale Using Graph Structure

notions of connectivity

Distributed Systems

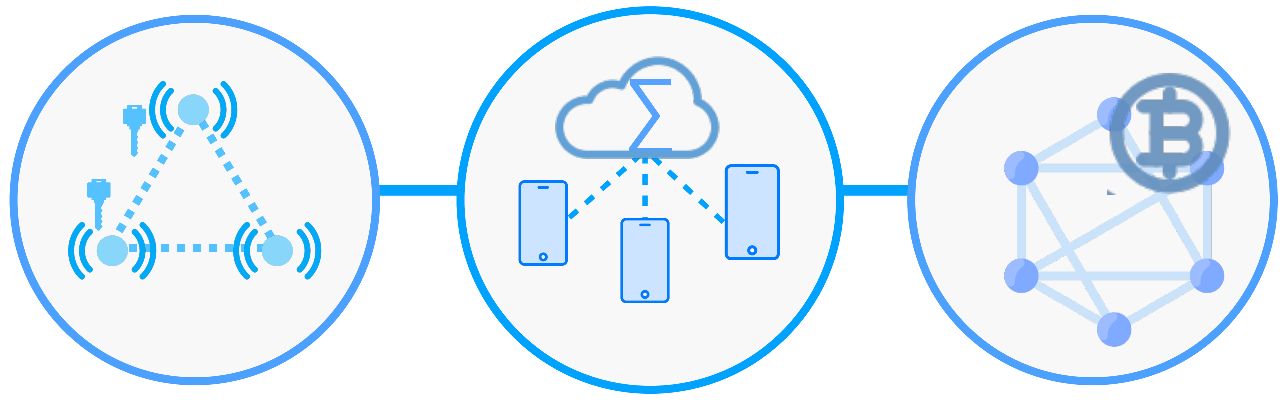

'secure' communication channel
between devices

communication

Probabilistic Graphical Models

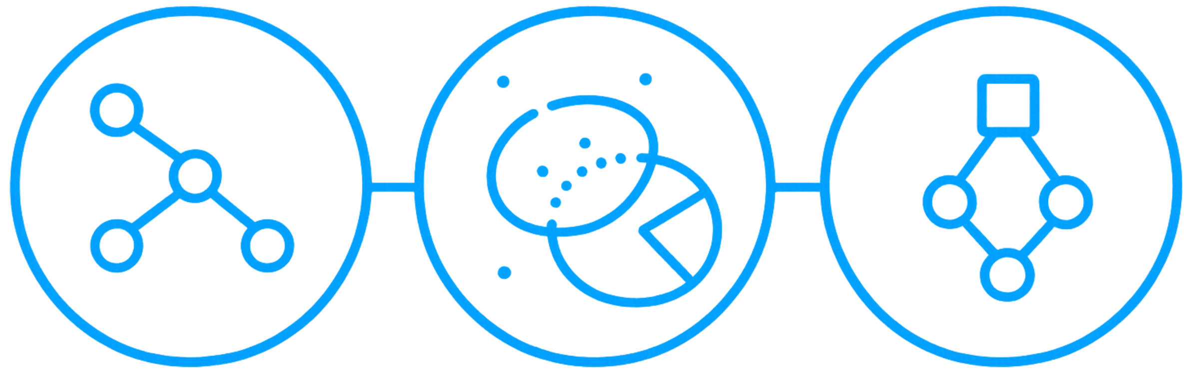

conditional dependencies
in observational data

computation

efficiency

heterogeneity

Reliable Inference at Scale Using Graph Structure

notions of connectivity

Distributed Systems

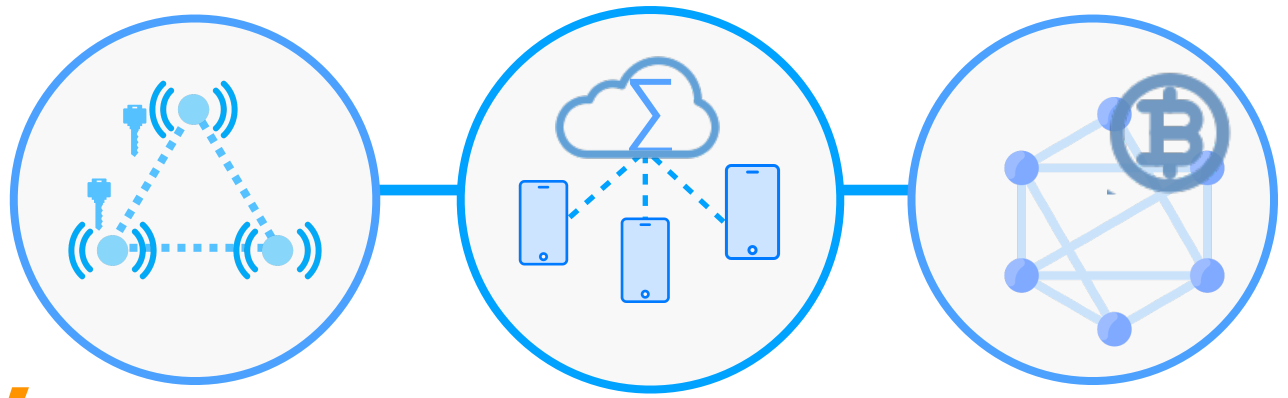

Today

'secure' communication channel
between devices

communication

Probabilistic Graphical Models

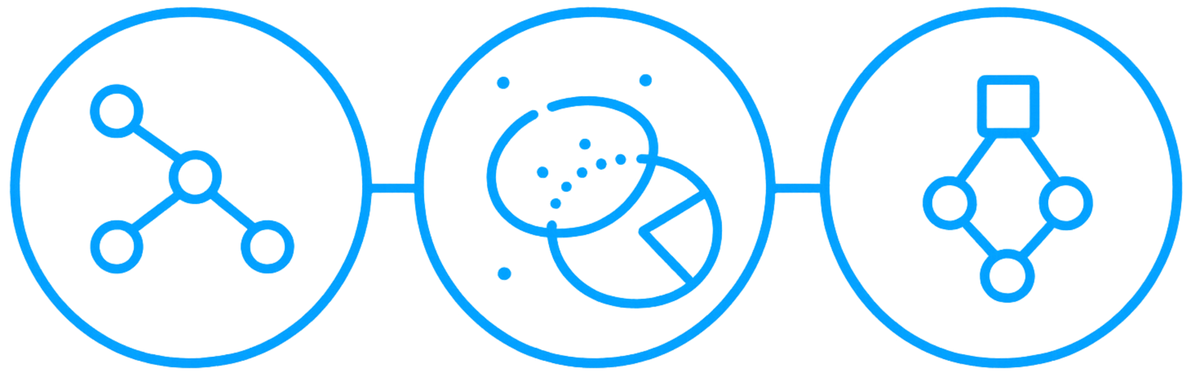

conditional dependencies
in observational data

computation

efficiency

heterogeneity

Acknowledgements

Devavrat Shah
MIT

Osman Yagan
CMU

Chai Wah Wu
IBM Research

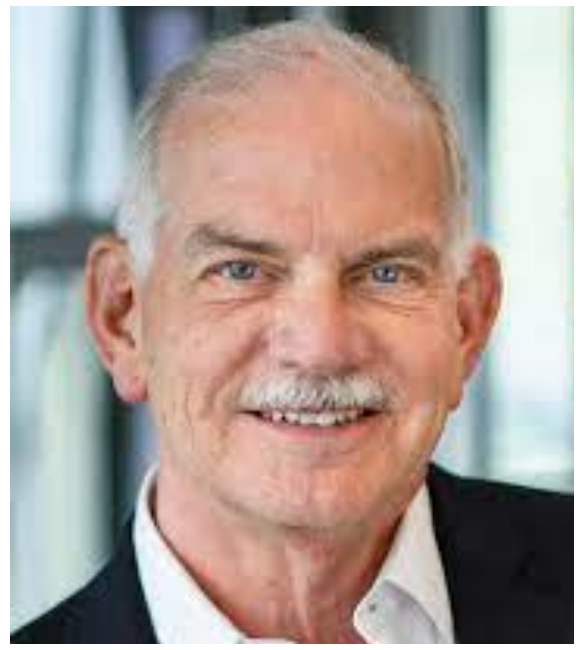

H. Vincent Poor
Princeton

Simon A. Levin
Princeton

SCHMIDT SCIENCE FELLOWS

Eray Can Elumar
Intel

Ani Sridhar
NJIT

Rashad Eletreby
Walmart Labs

Hejin Gu
CMU

Part 1: Road-map

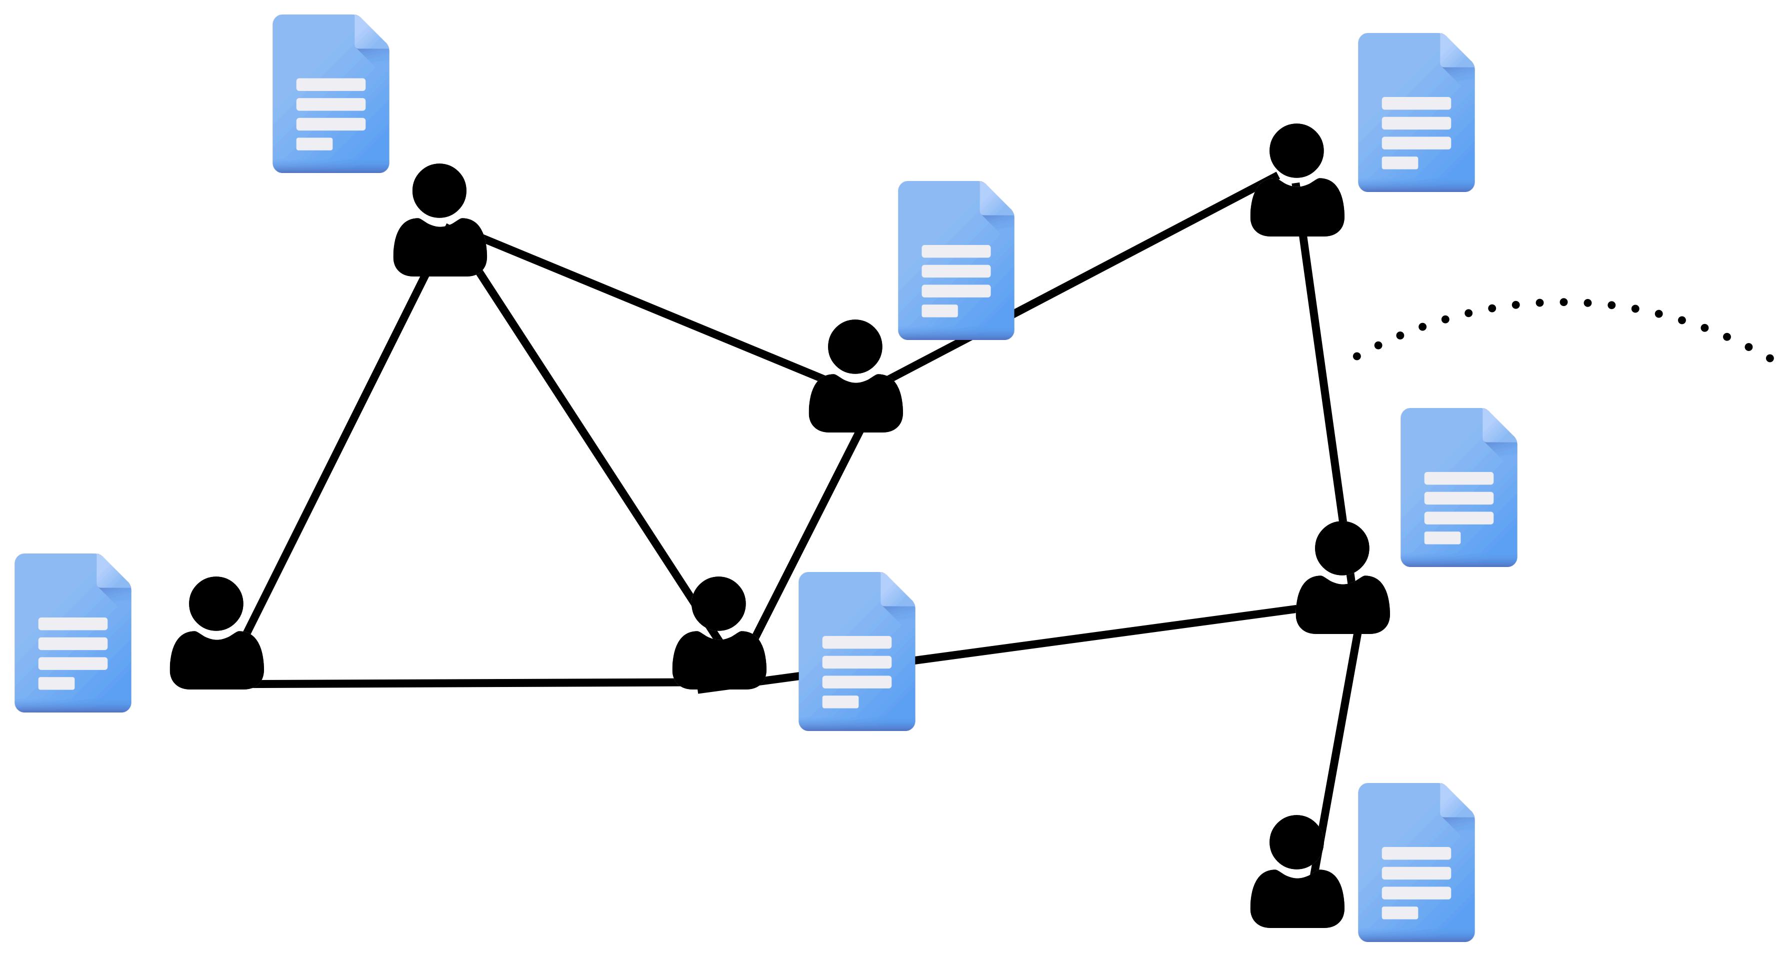

How to securely aggregate data?

Random
[network design]

- [ ] From distributed systems to random graphs
  The sparsity connectivity trade-off
- [ ] What are random K-out graphs?
  What makes them a useful design tool?
- [ ] Random K-out graphs in action
  What graph properties matter?
- [ ] Review: Notions of Connectivity
- [ ] Preview of contributions
  Key proof techniques

From distributed systems to random graphs…

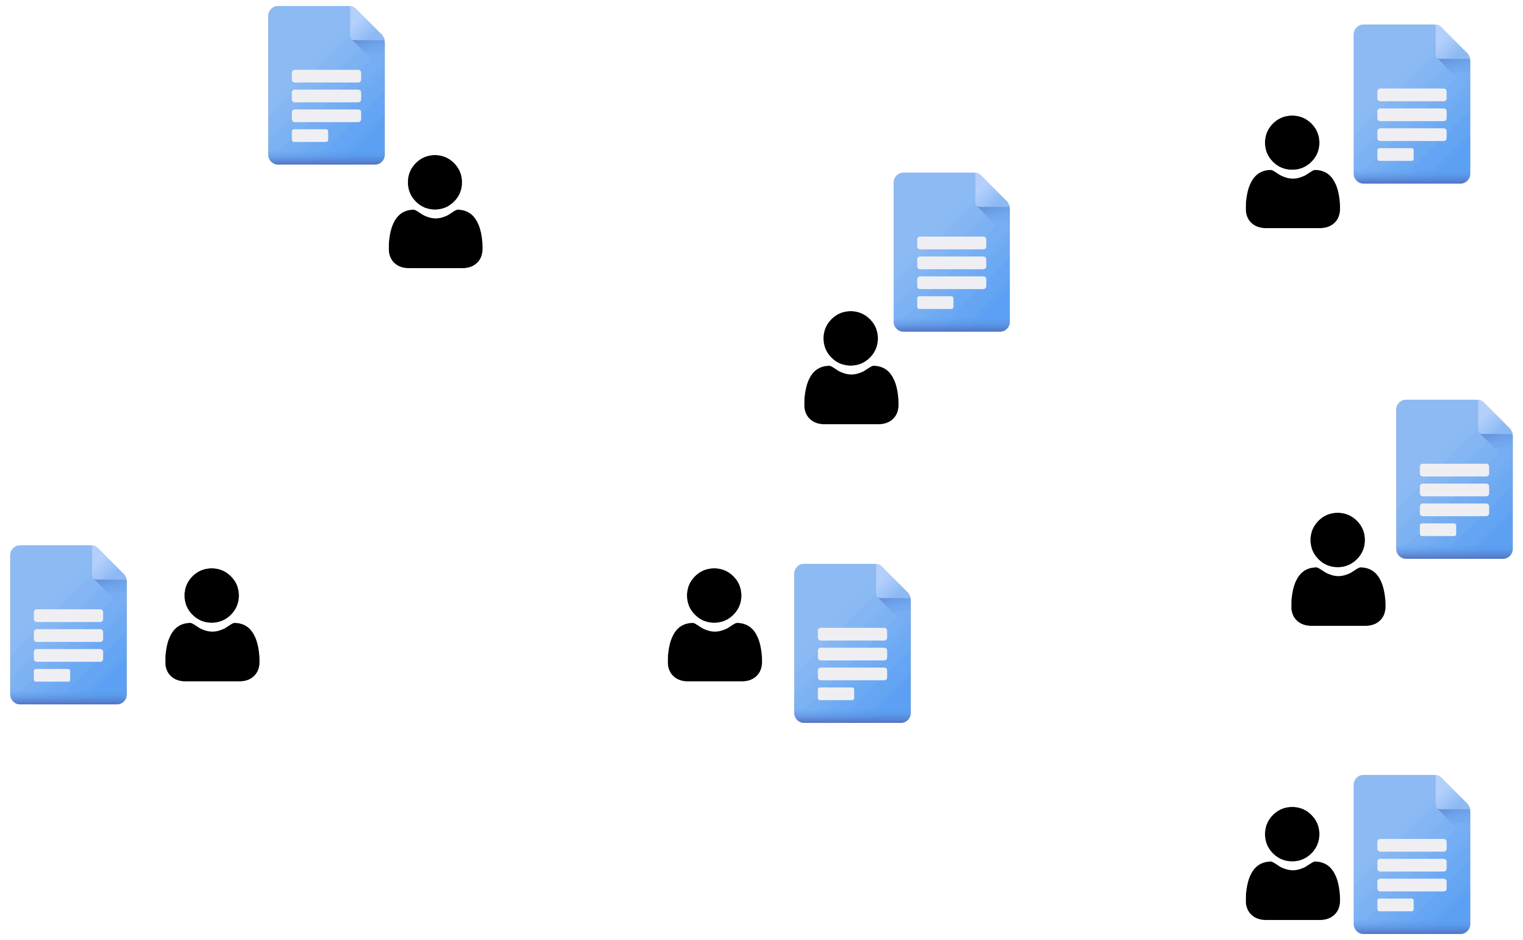

How to securely aggregate data?

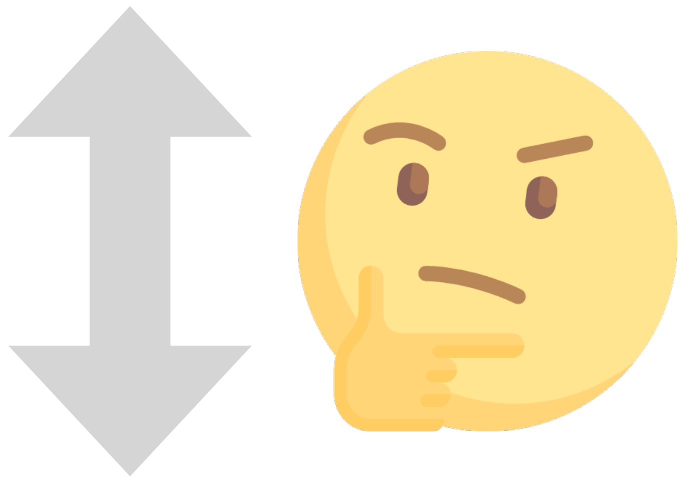

Random
[network design]

From distributed systems to random graphs…

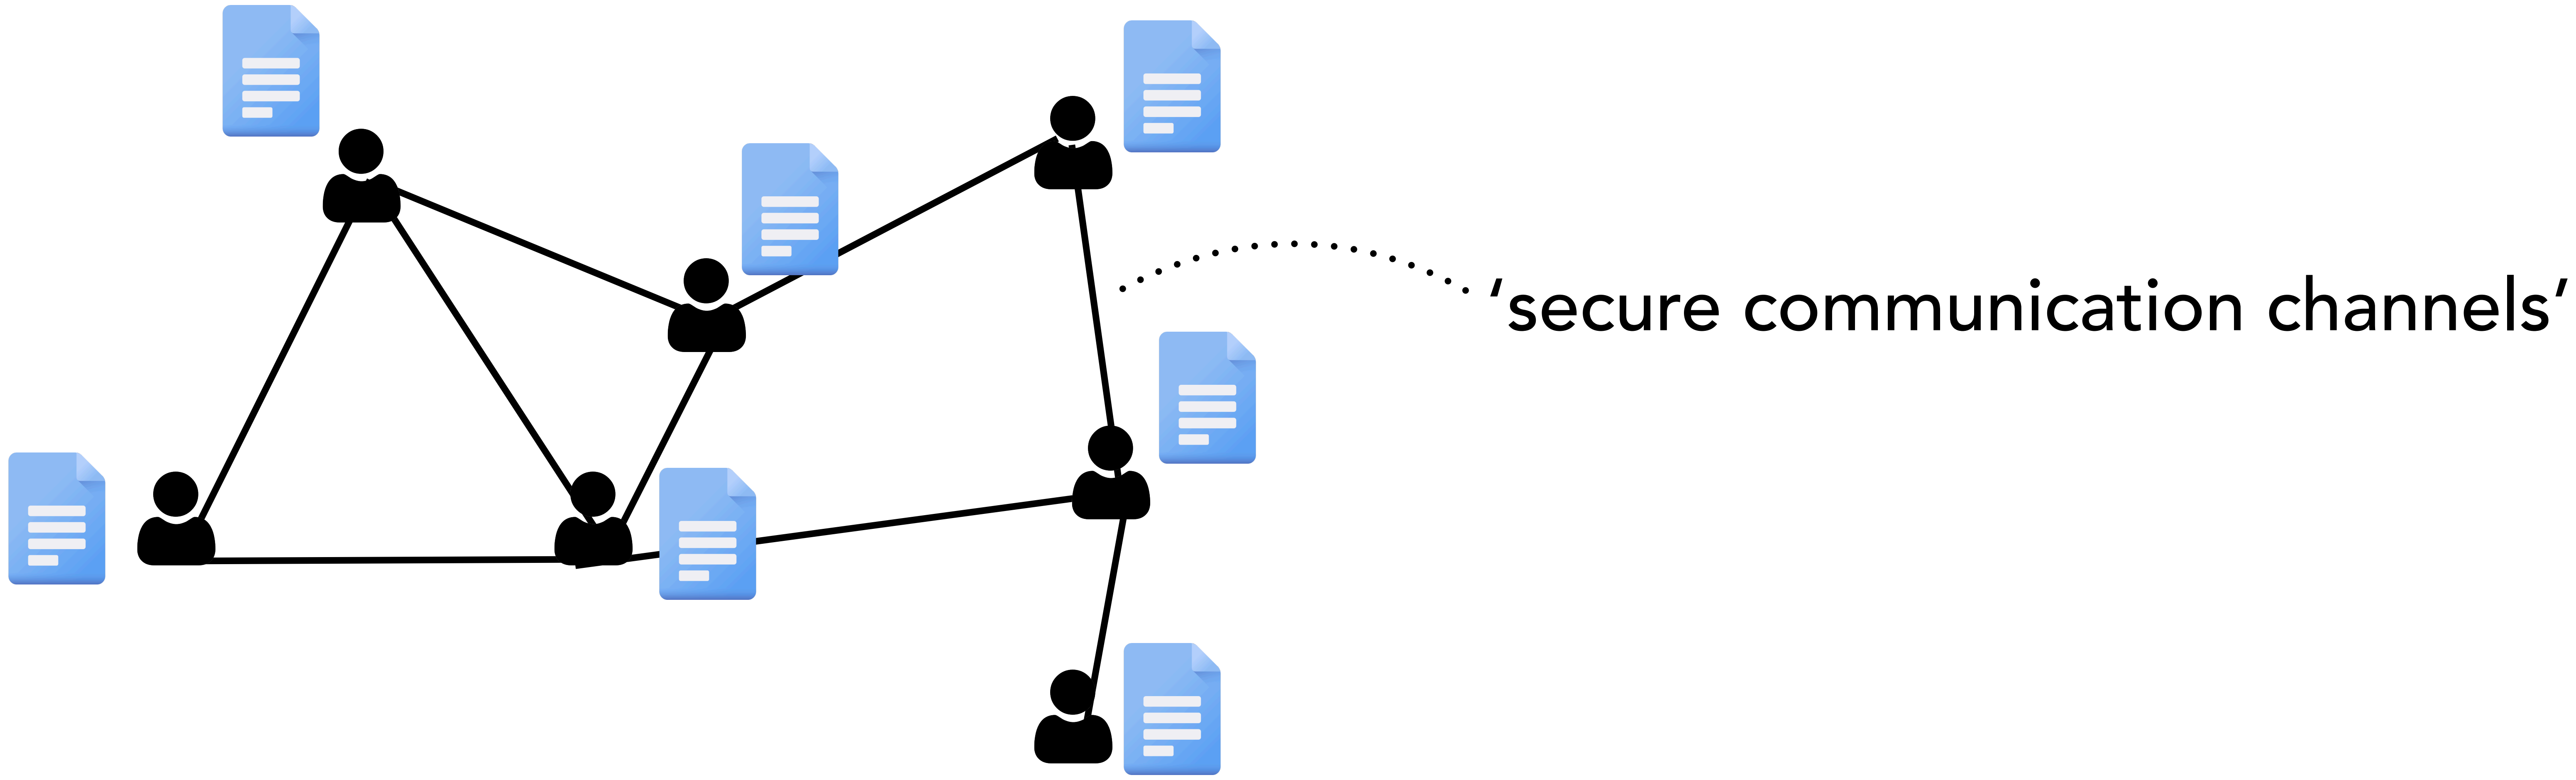

How to securely aggregate data?

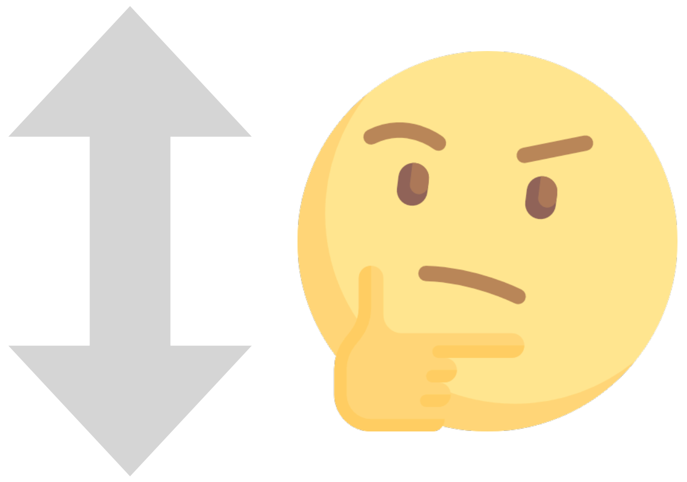

Random
[network design]

Securely aggregating information over a network

Setting:

Edge: communication channels

Node: private high-dim feature

No trusted orchestrator

Goal:

Learn $\Sigma_i w_i$

without revealing w_i

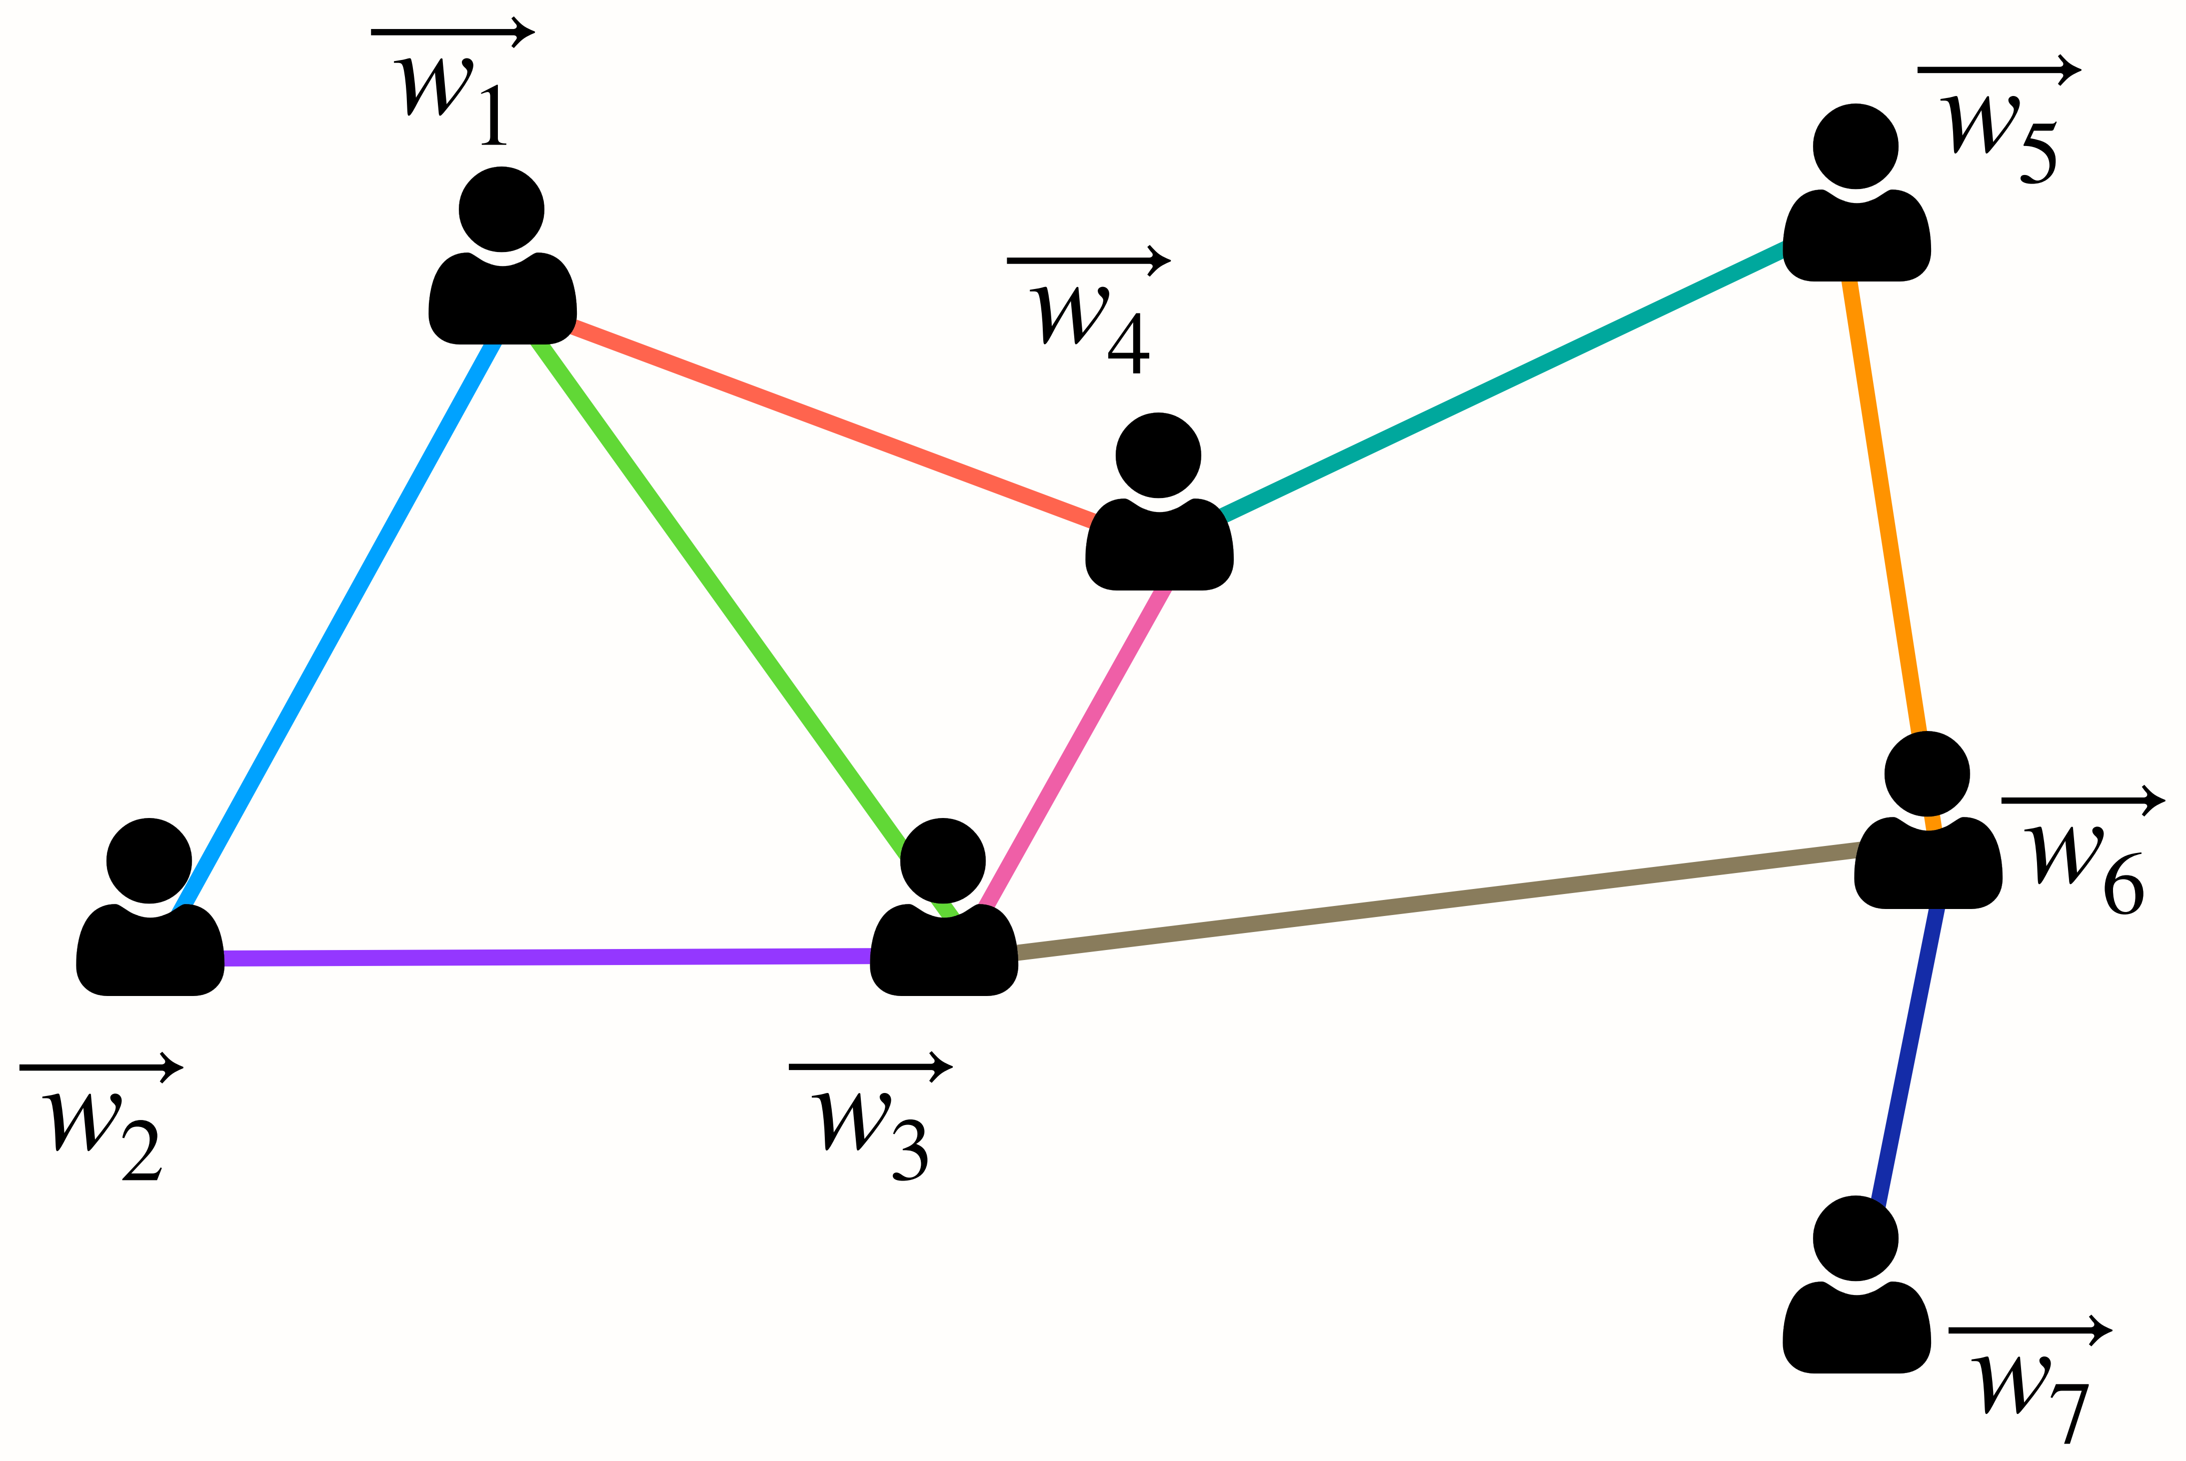

Securely aggregating information over a network

Setting:

Edge: communication channels

Node: private high-dim feature

No trusted orchestrator

Goal:

Learn $\Sigma_i w_i$

without revealing w_i

Approach:

add pairwise masks (that cancel in aggregate)

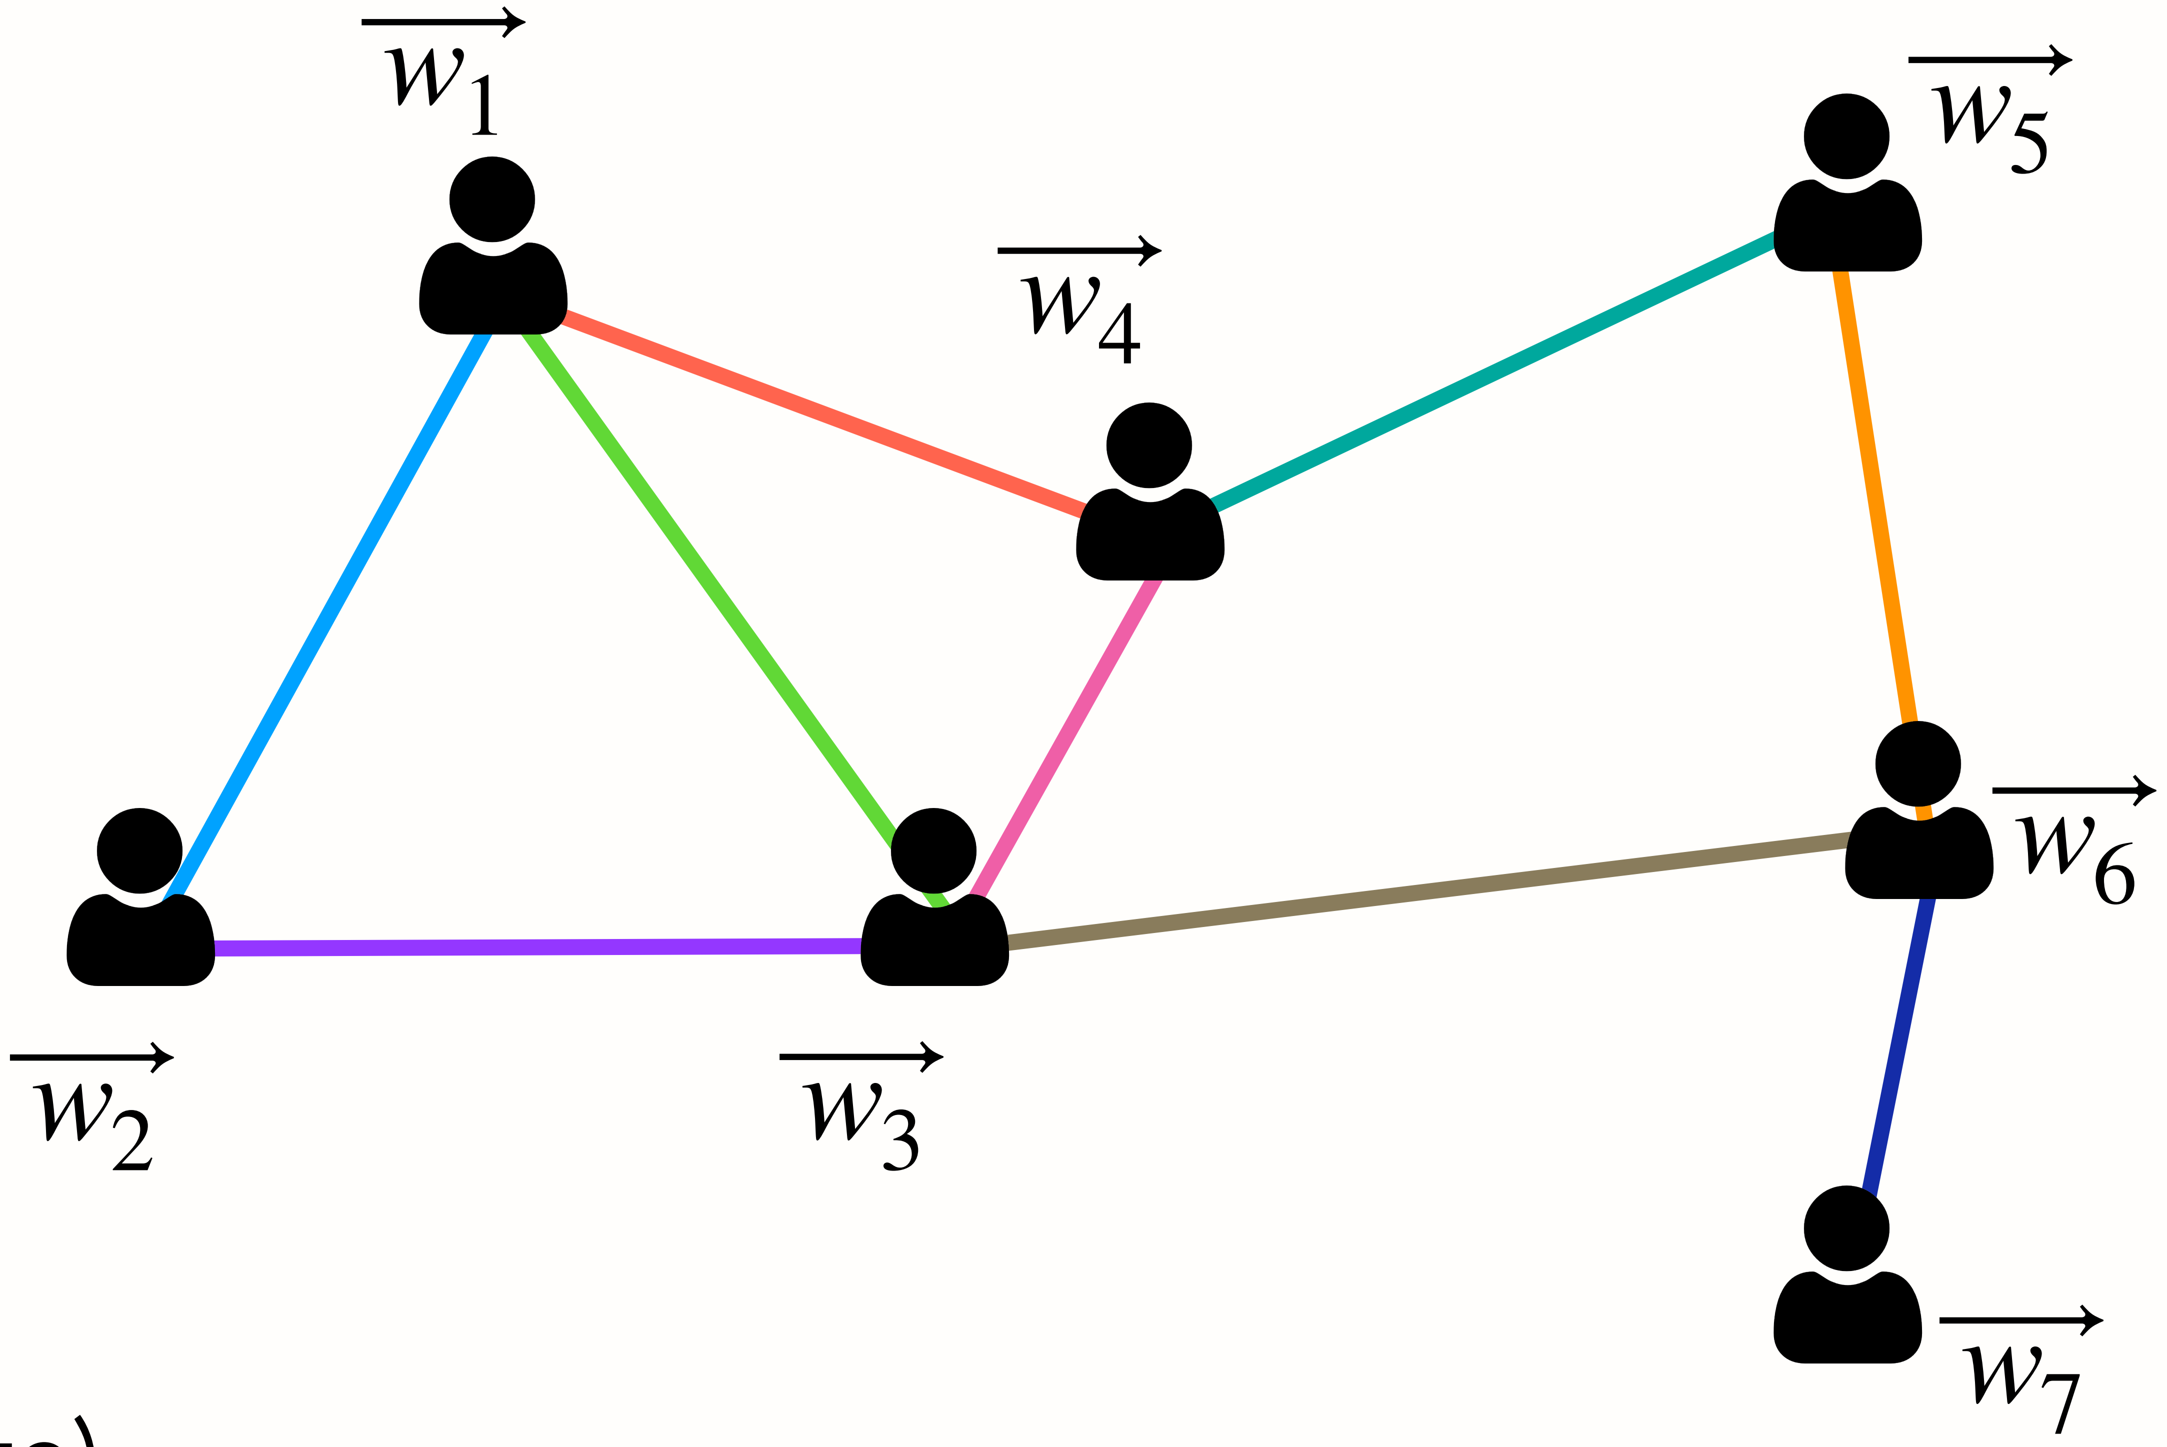

[Sabater et al. '22],
[Bell et al. '20],
[Bonawitz et al. '17]
⋮
[Chaum '88]

Securely aggregating information over a network

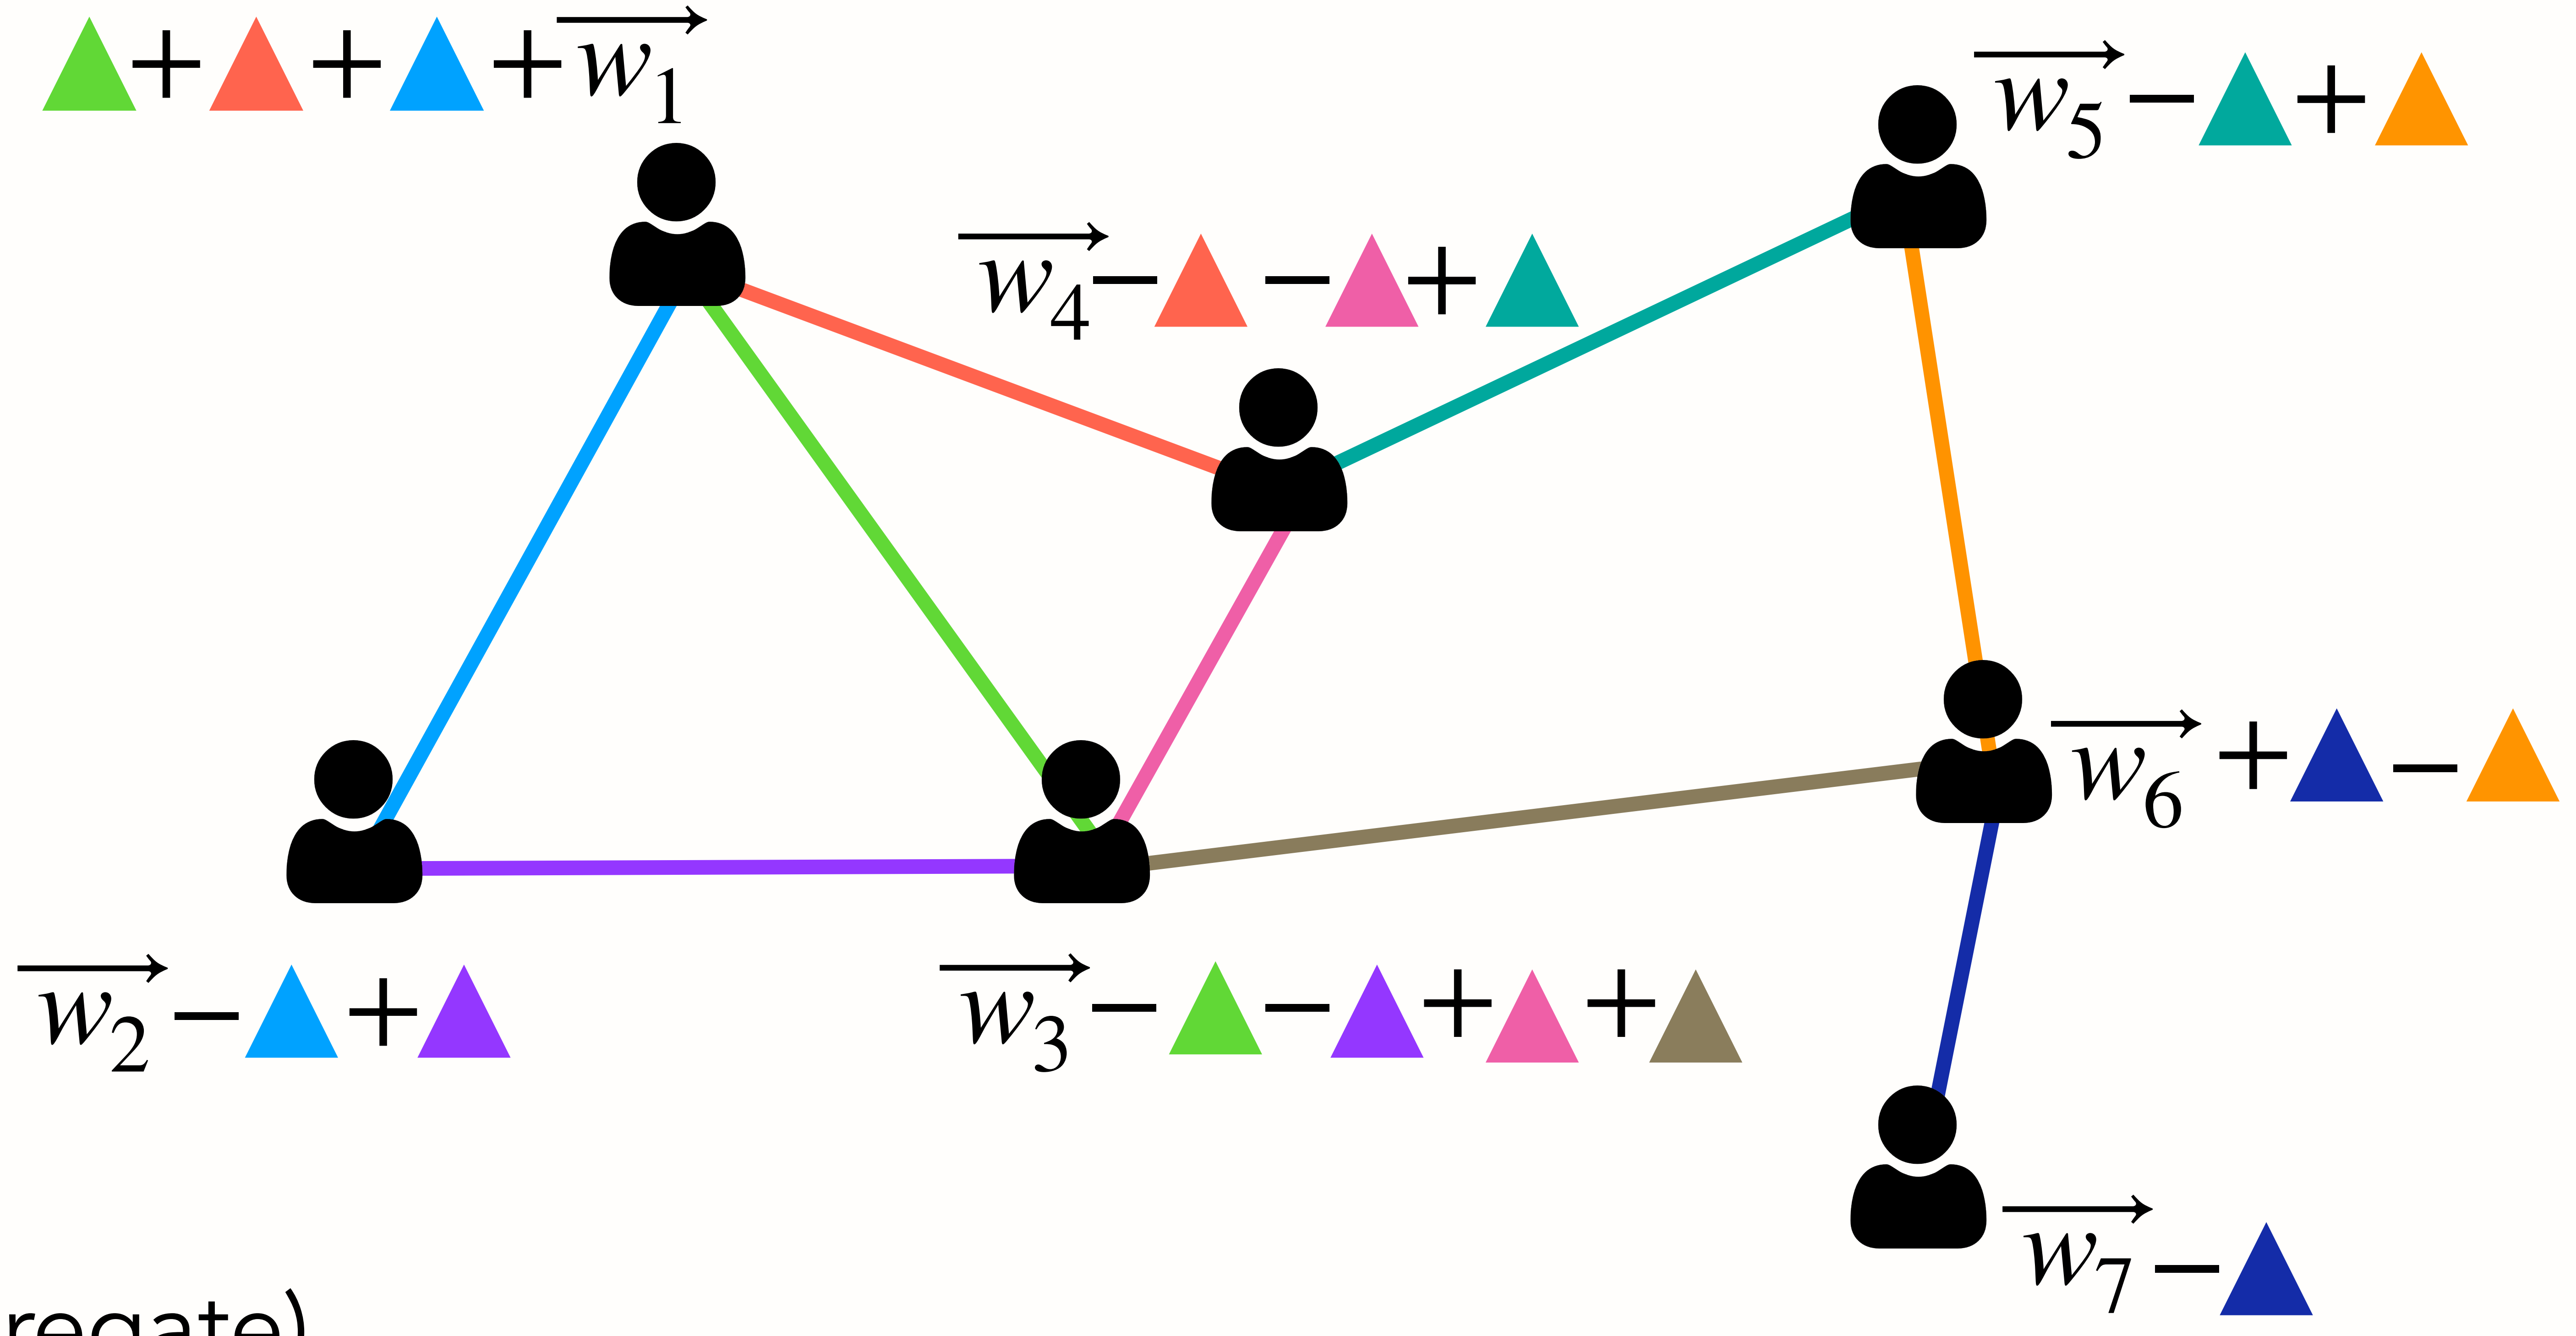

Goal:

Learn $\Sigma_i w_i$

without revealing w_i

Approach:

add pairwise masks (that cancel in aggregate)

broadcast masked features to network

[Sabater et al. '22], …

Securely aggregating information over a network

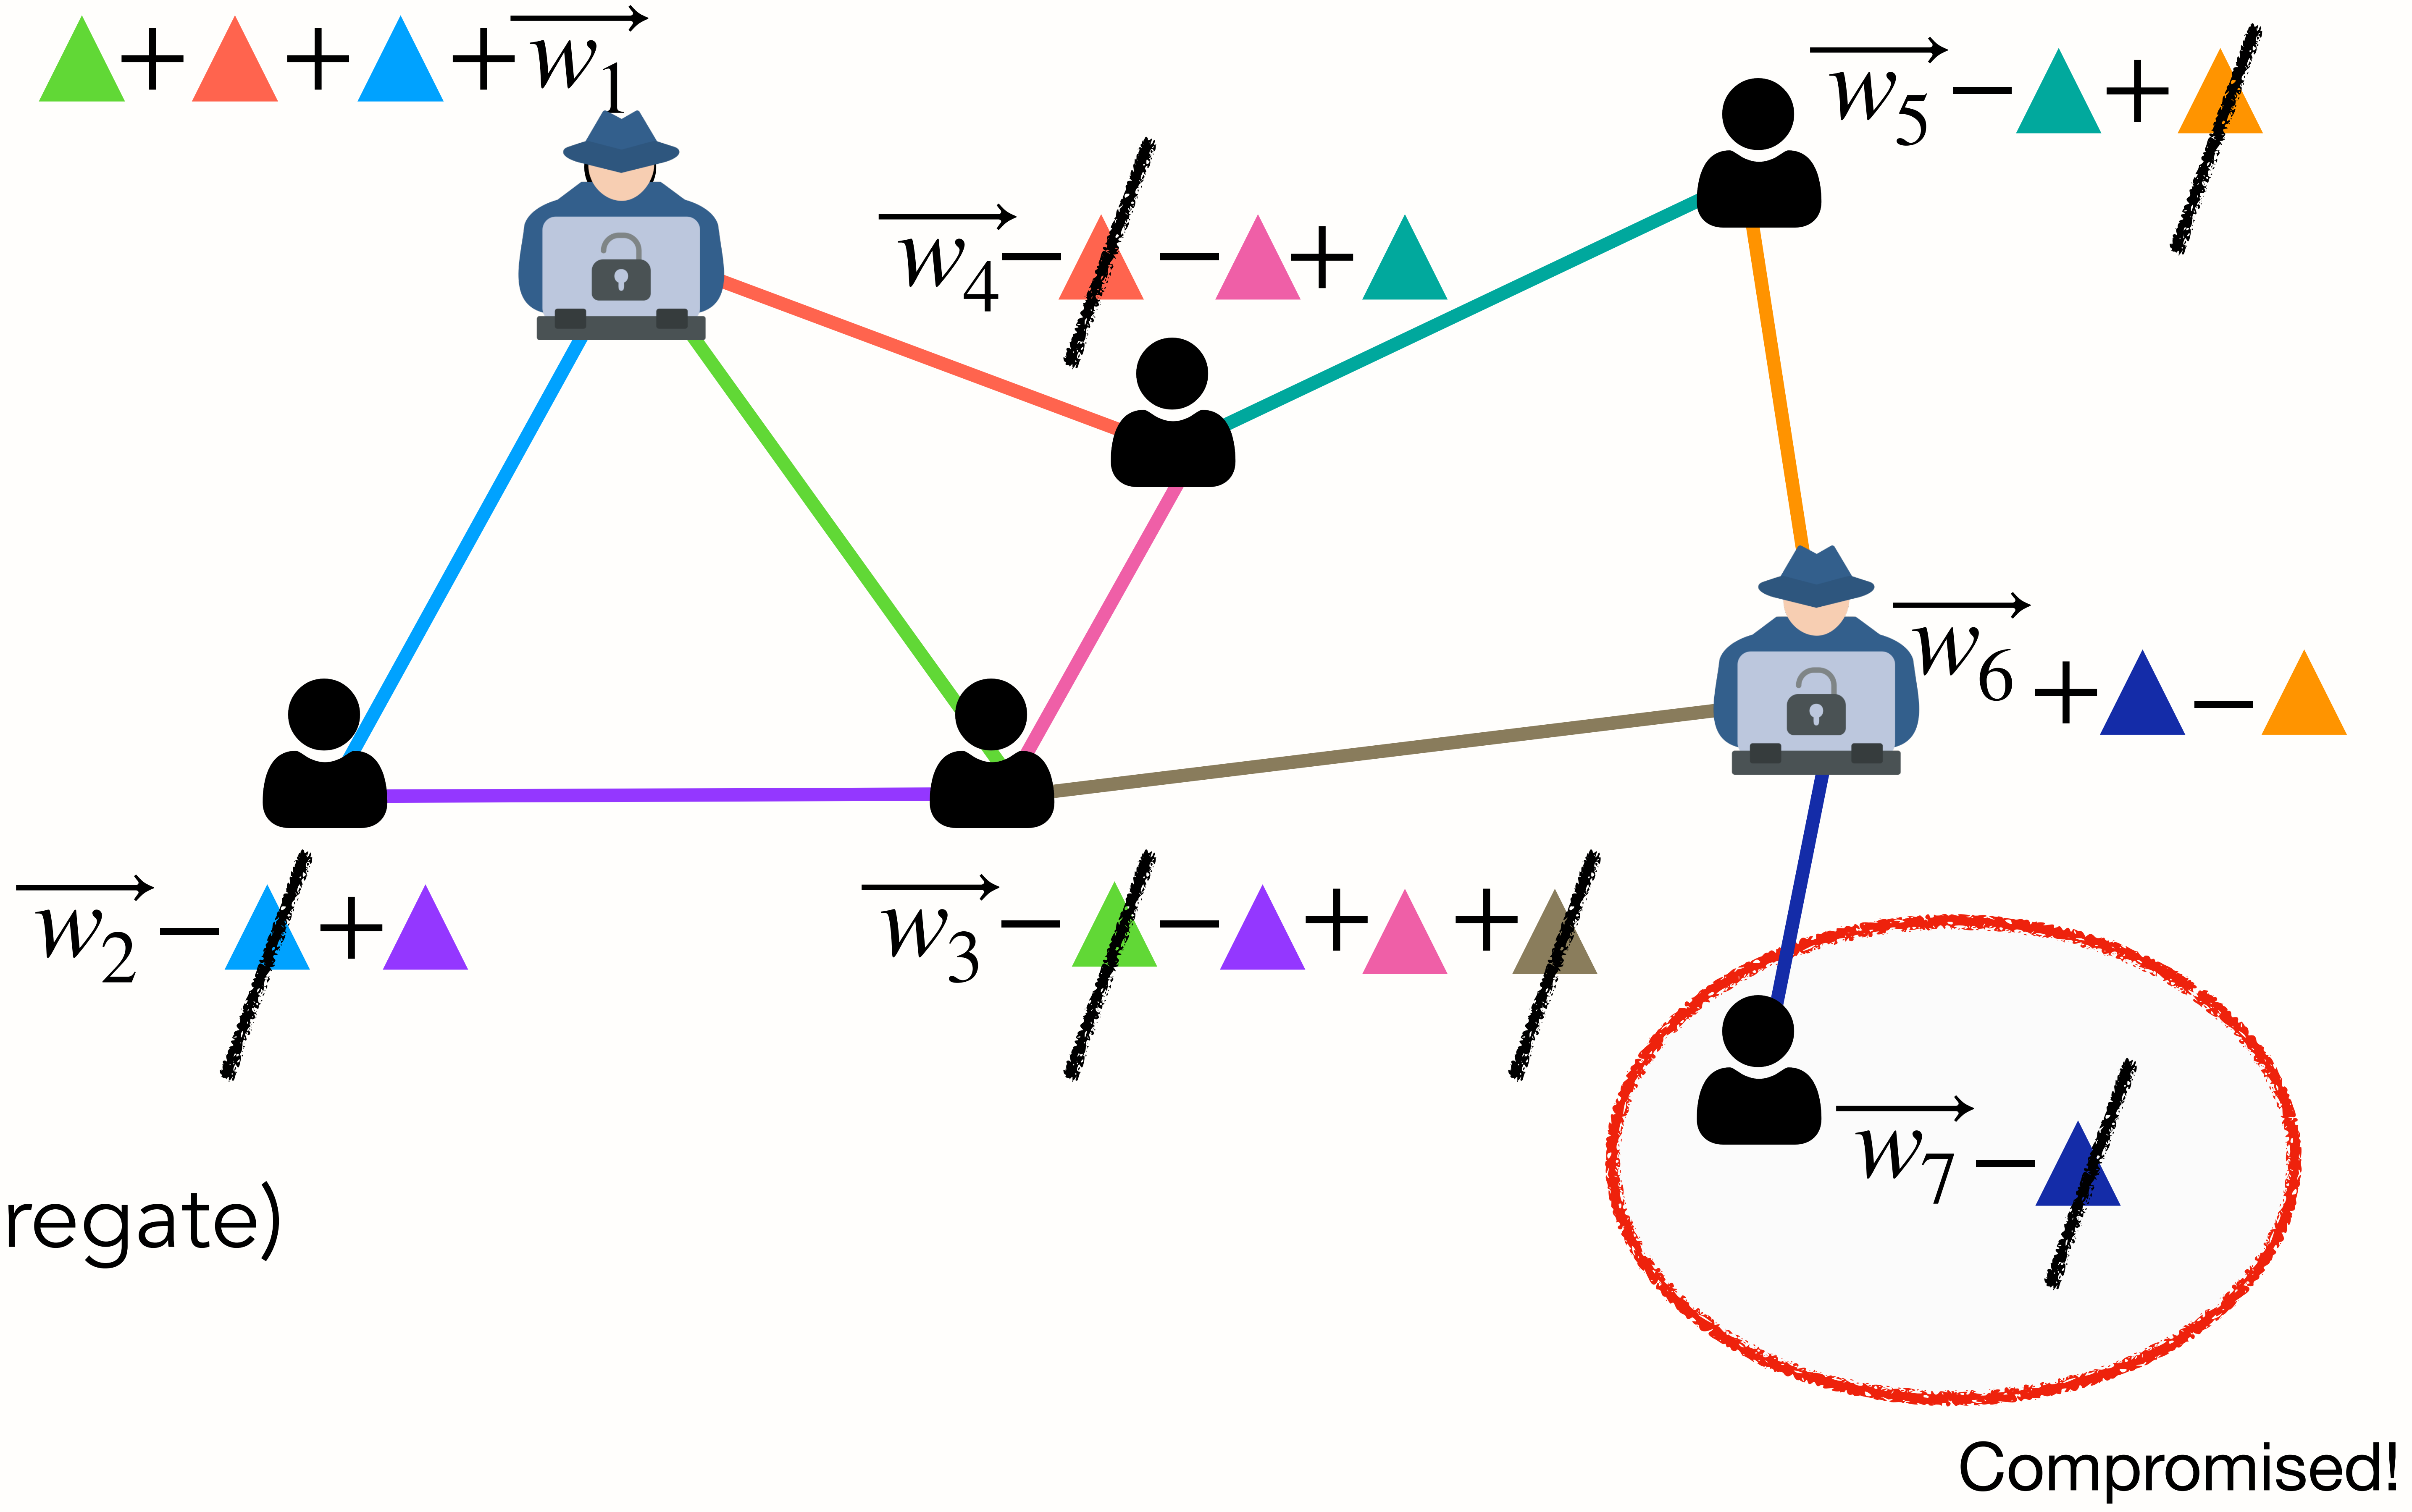

Goal:

Learn $\Sigma_i w_i$

without revealing w_i

Approach:

add pairwise masks (that cancel in aggregate)

Challenge:

What if some nodes collude, i.e.,

Share masks for their neighbors?

Securely aggregating information over a network

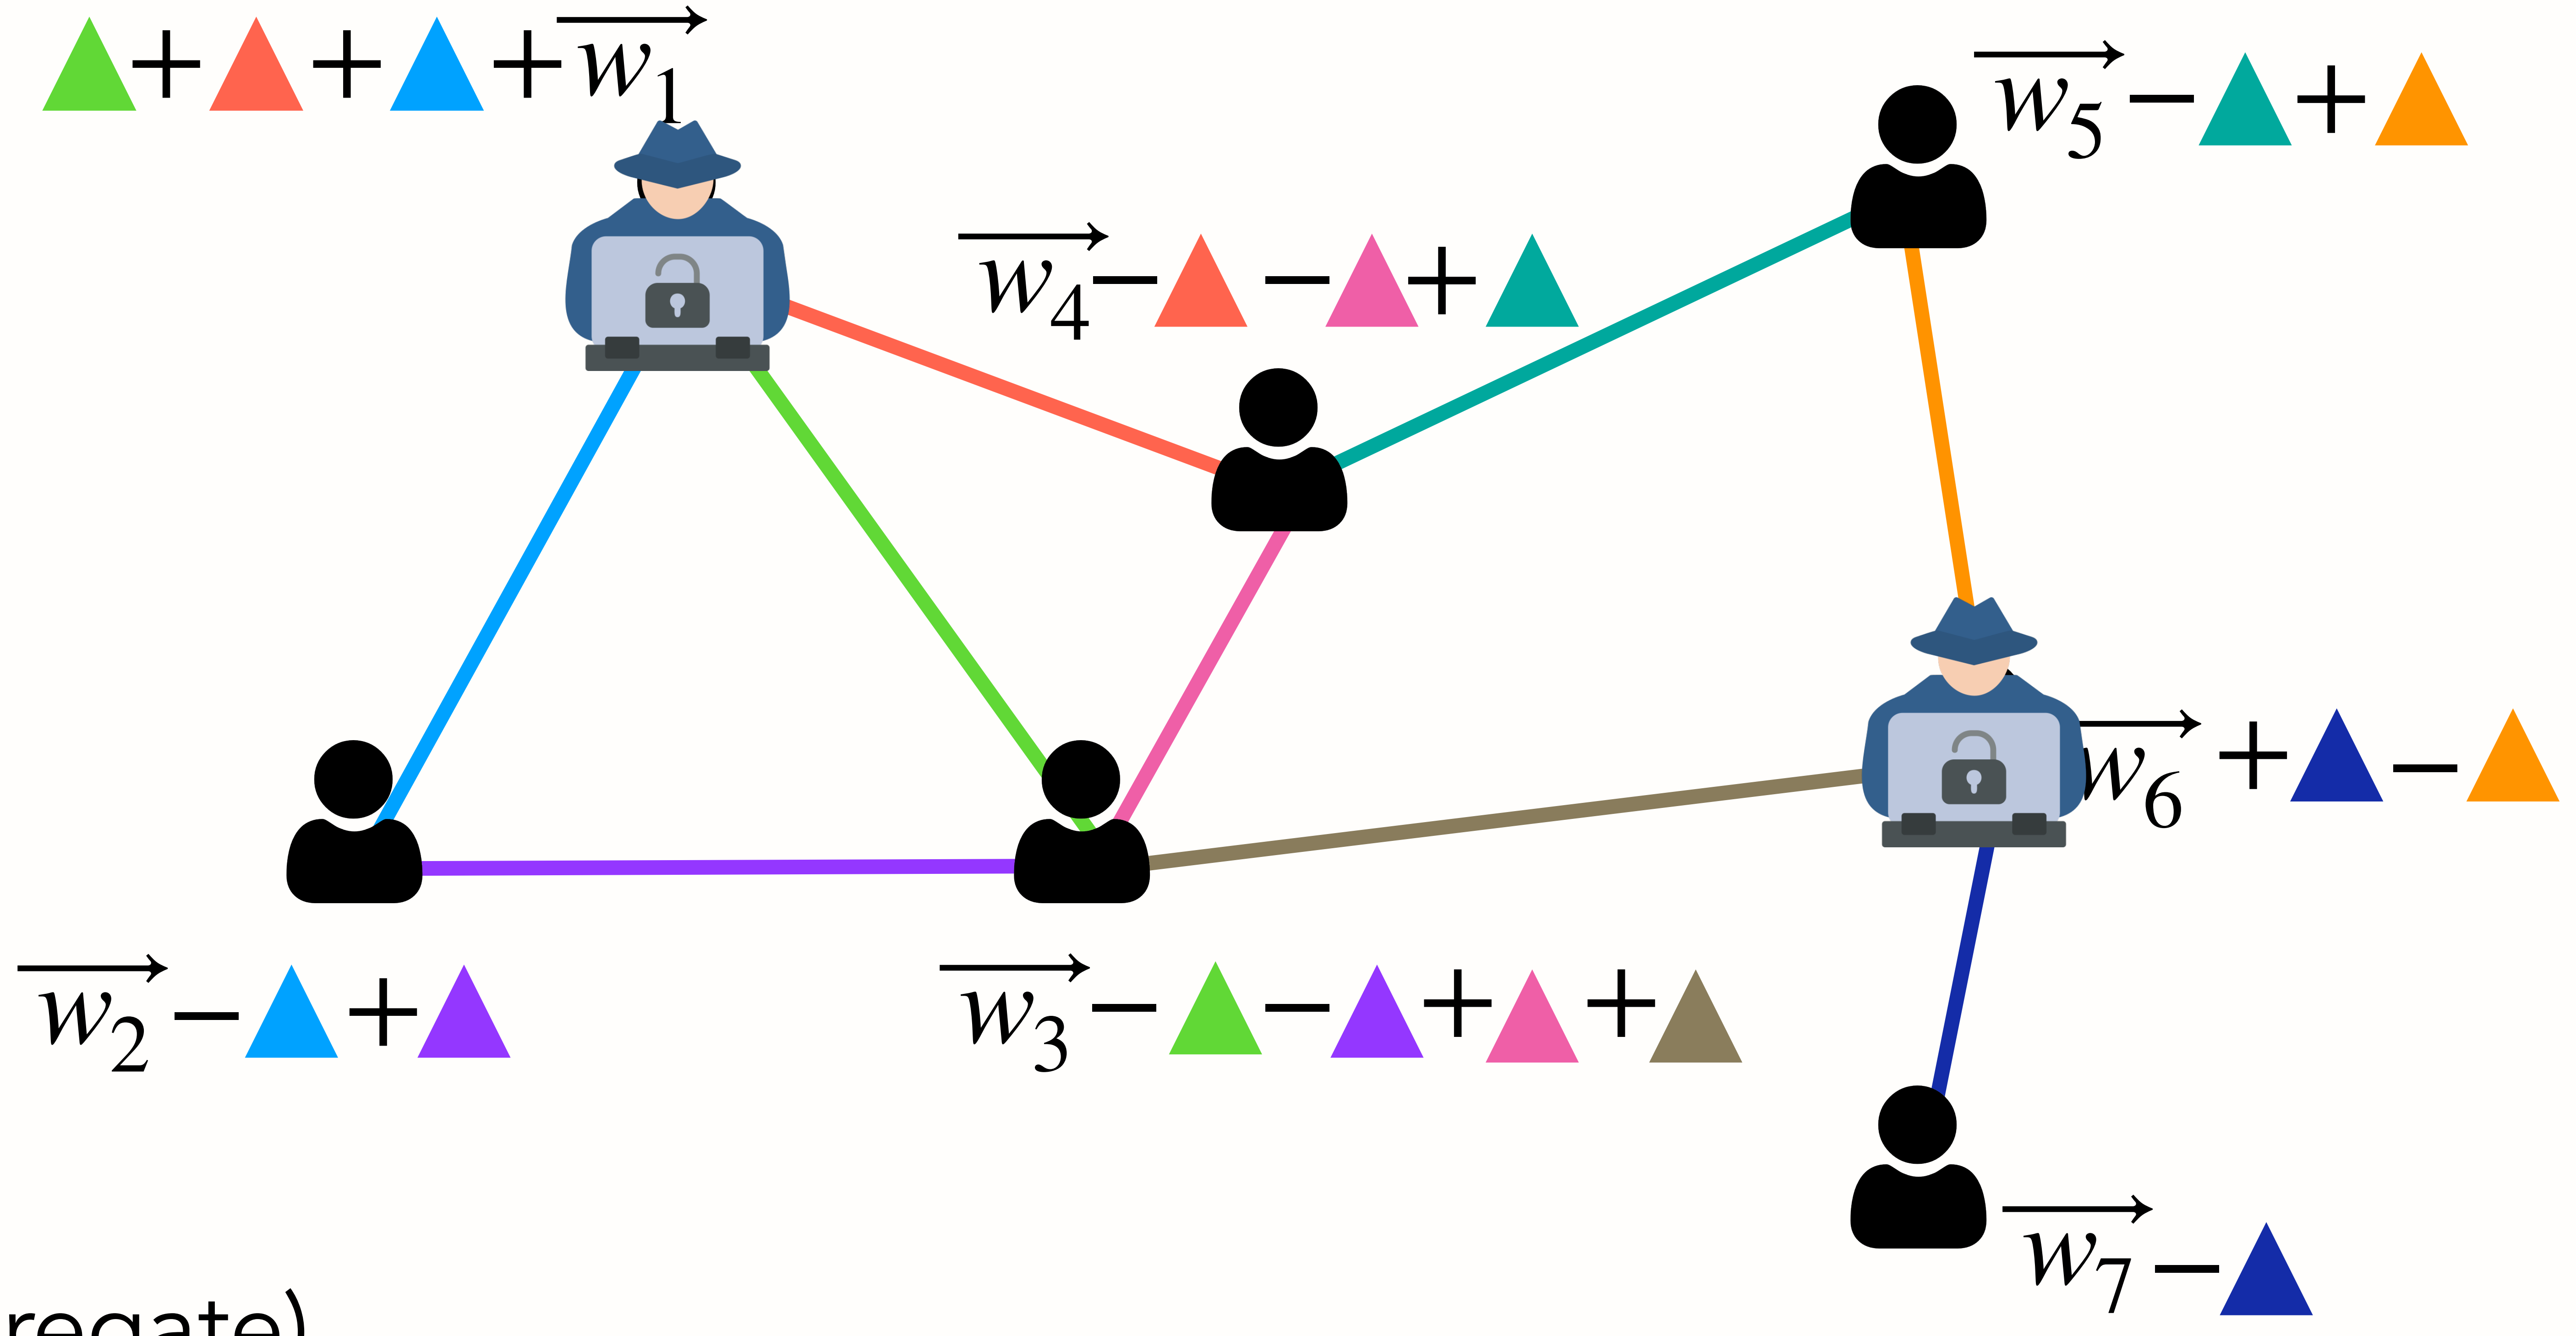

Goal:

Learn $\Sigma_i w_i$

without revealing w_i

Approach:

add pairwise masks (that cancel in aggregate)

Challenge:

What if some nodes collude, i.e.,
Share masks for their neighbors?

Securely aggregating information over a network

Performance trade-off:

better connectivity $\Longrightarrow$ better masking but raises communication costs!

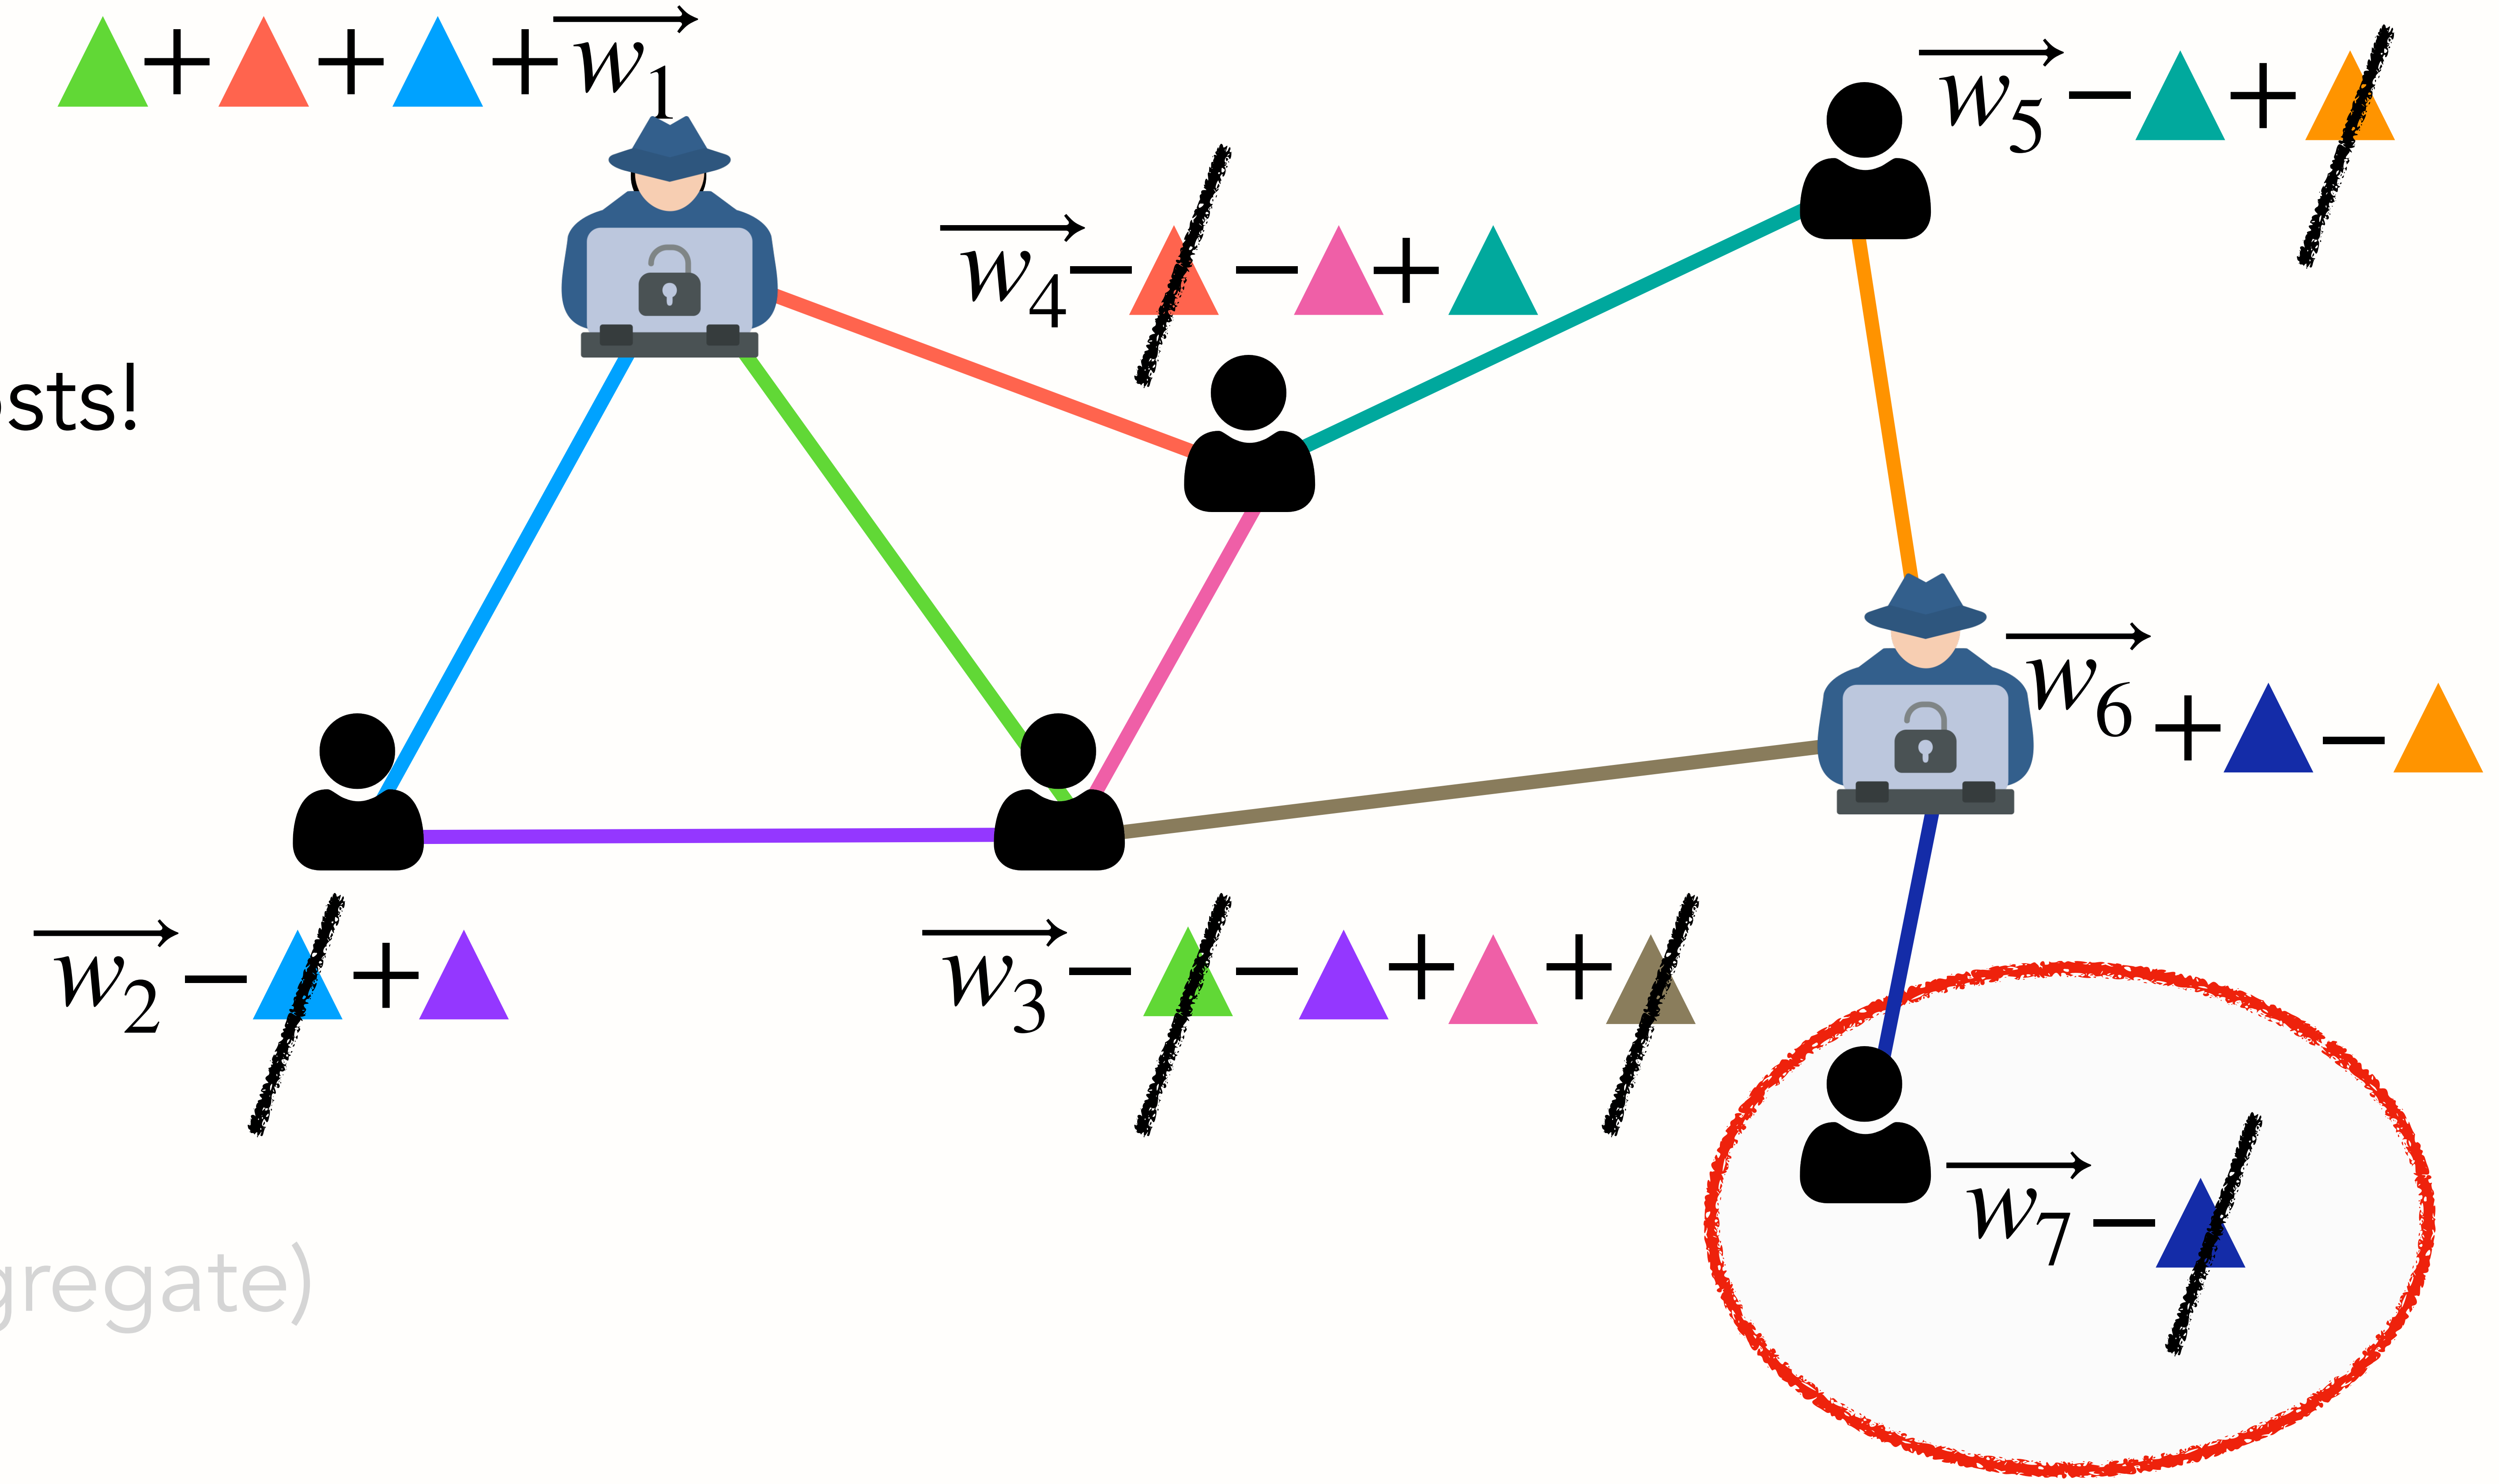

Goal:

Learn $\Sigma_i w_i$ without revealing w_i

Approach:

add pairwise masks (that cancel in aggregate)

Challenge:

What if some nodes collude, i.e.,
Share masks for their neighbors?

From distributed systems to random graphs…

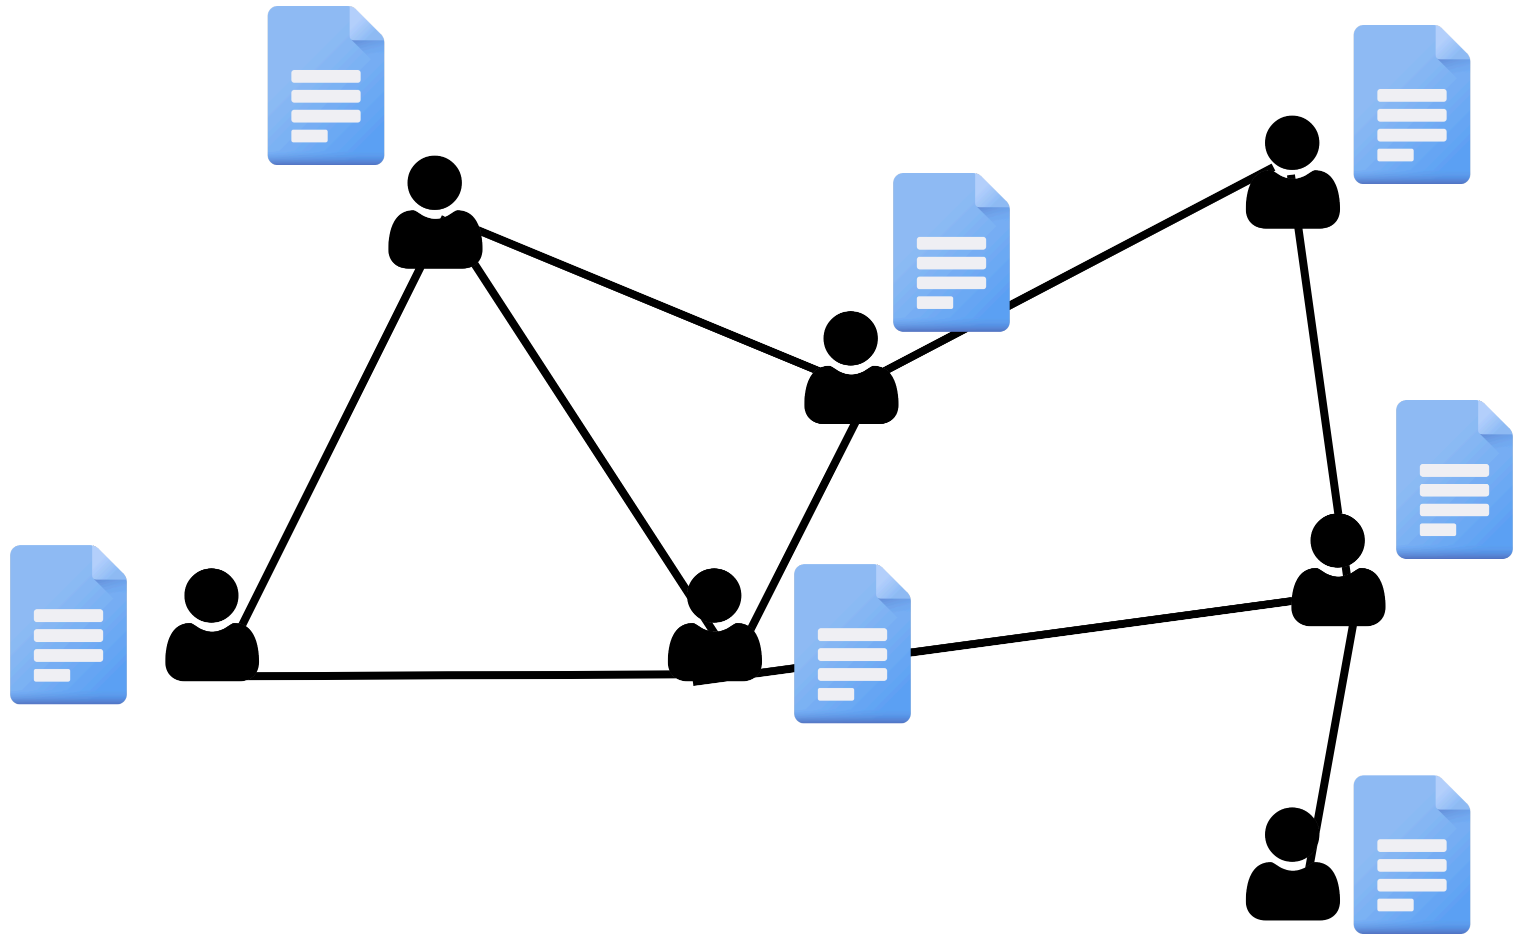

How to securely aggregate data?

Random
[network design]

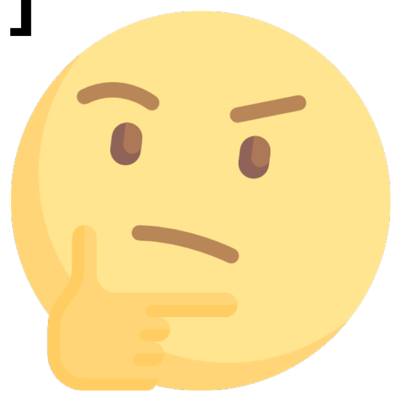

Performance trade-off:

'better' connectivity $\Longrightarrow$ better masking but raises communication costs!

From distributed systems to random graphs…

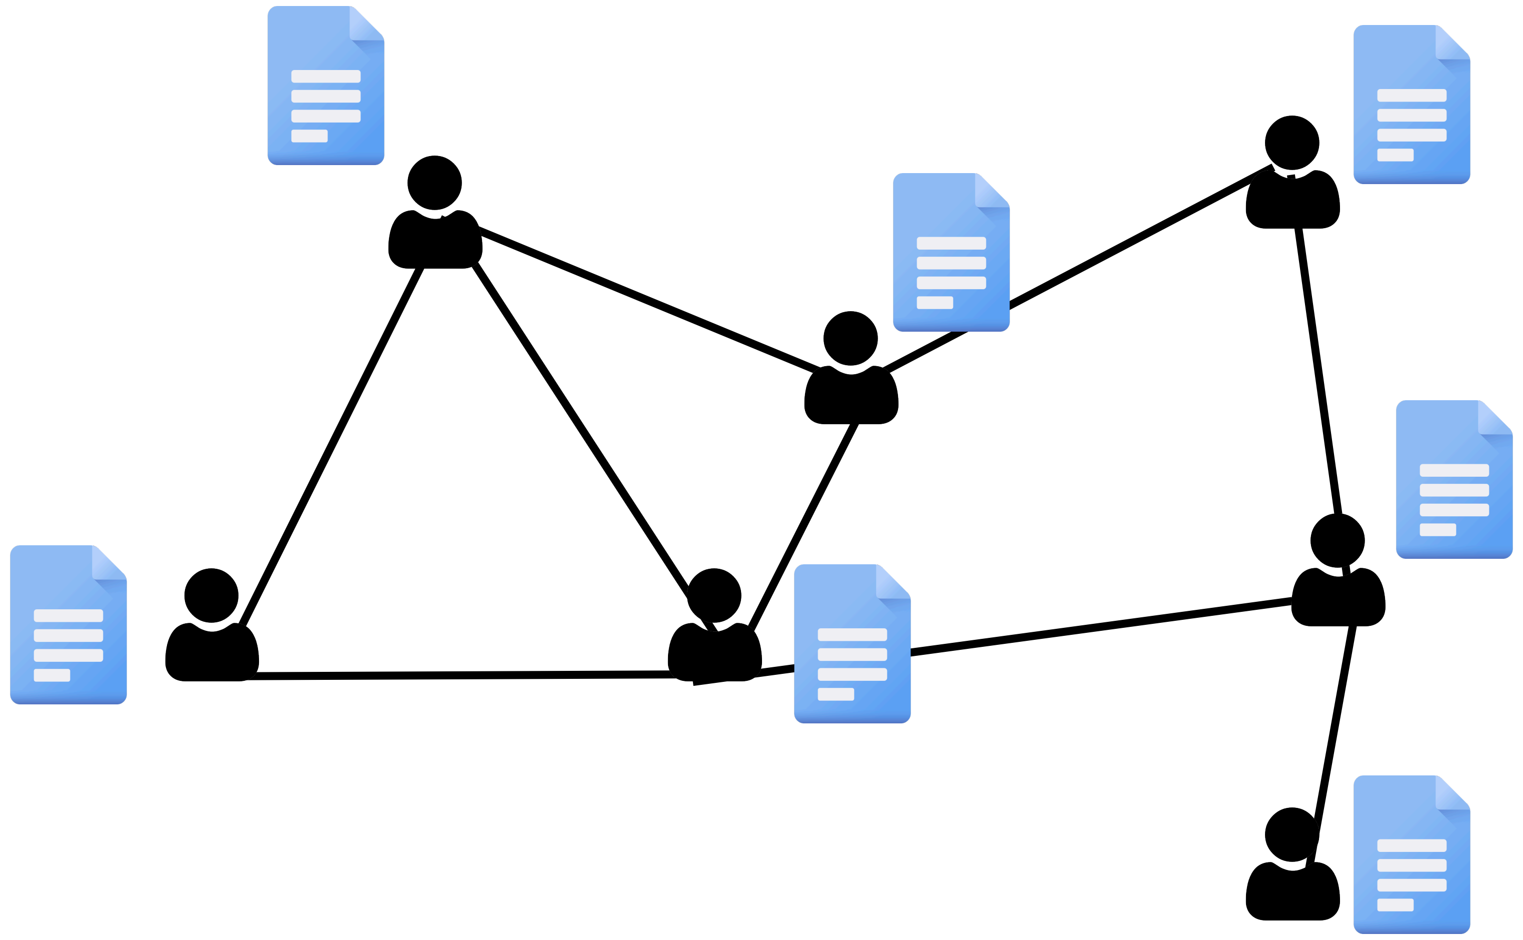

How to securely aggregate data?

Random
[network design]

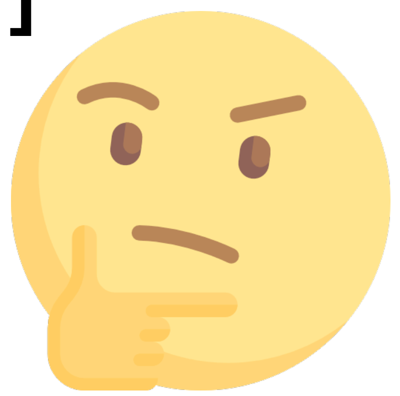

Performance trade-off:

'better' connectivity $\Longrightarrow$ better masking but raises communication costs!

A network that is just right!
well-connected yet sparse,
allows distributed, lightweight
construction…

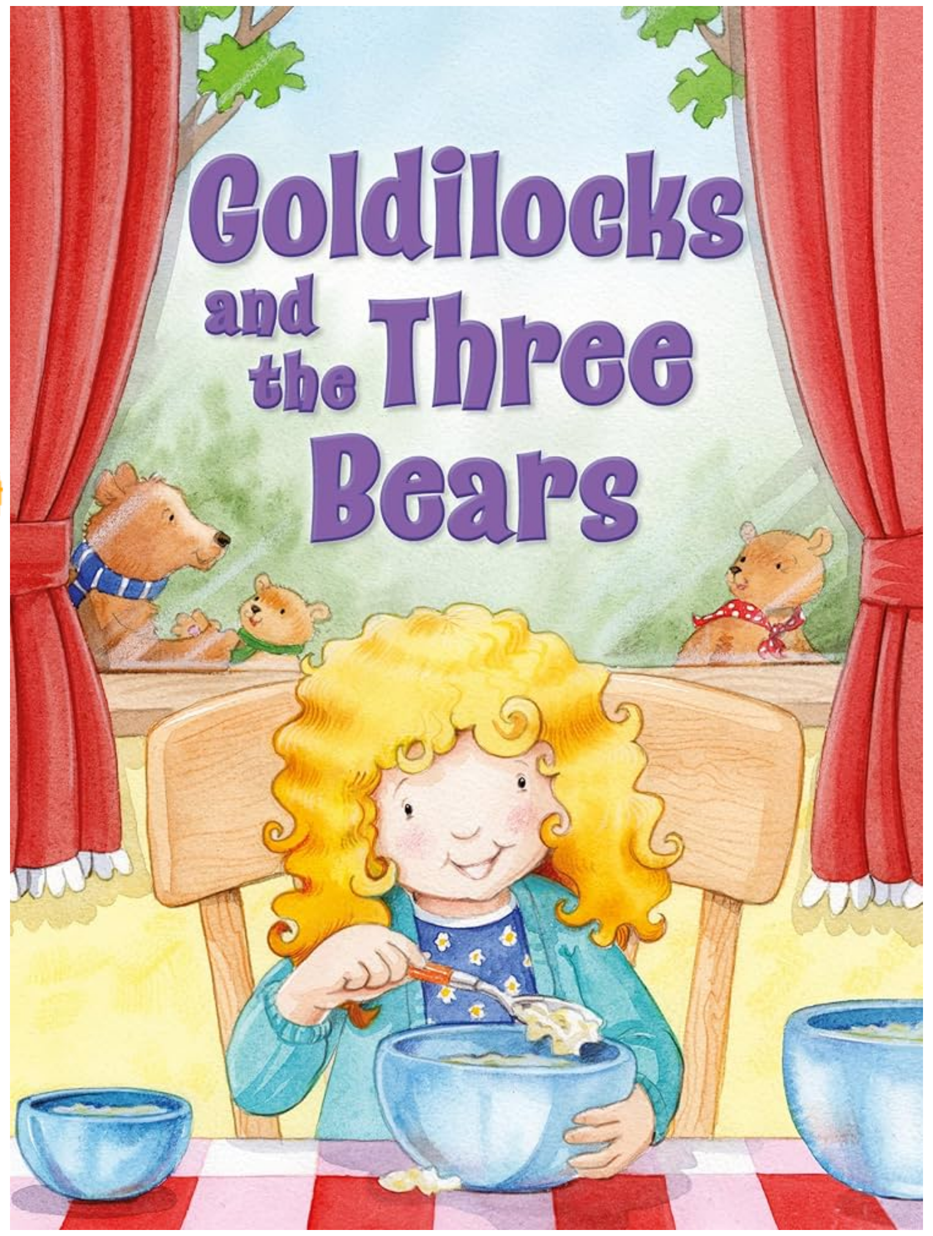

From distributed systems to random graphs…

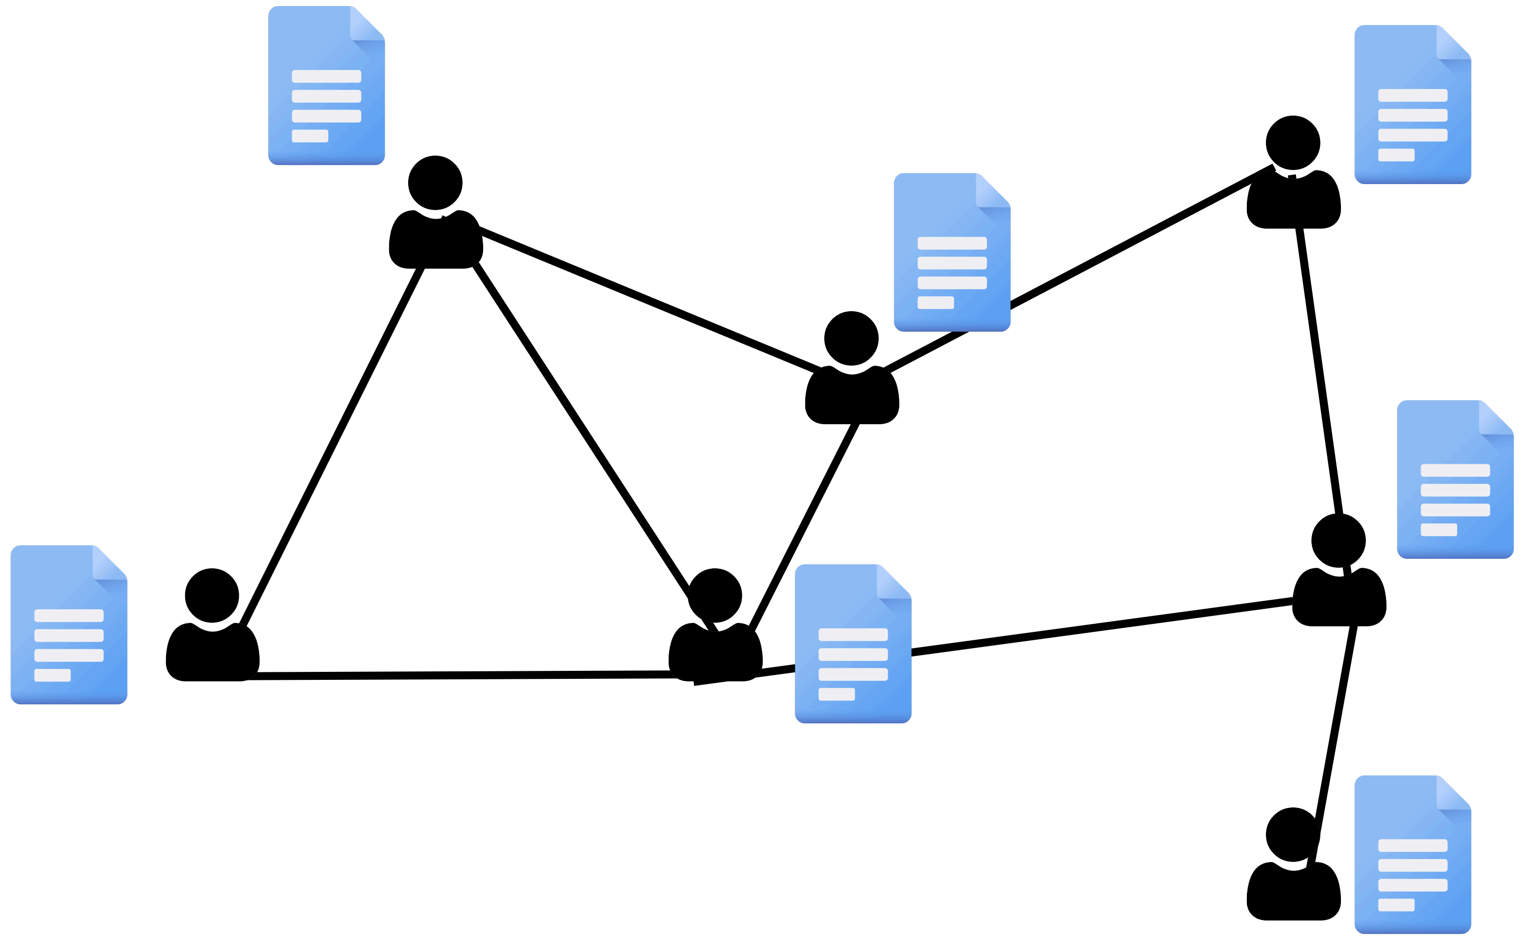

How to securely aggregate data?

Random
[network design]

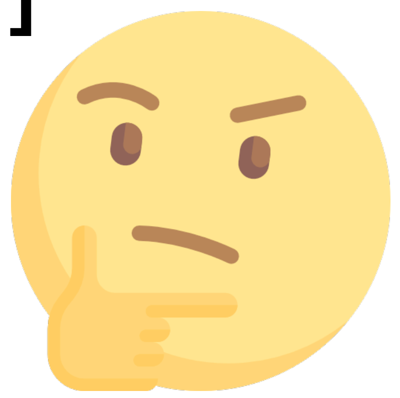

Performance trade-off:

'better' $\Longrightarrow$ better masking
connectivity but
raises communication costs!

A network that is just right!
well-connected yet sparse,
allows distributed, lightweight
construction.

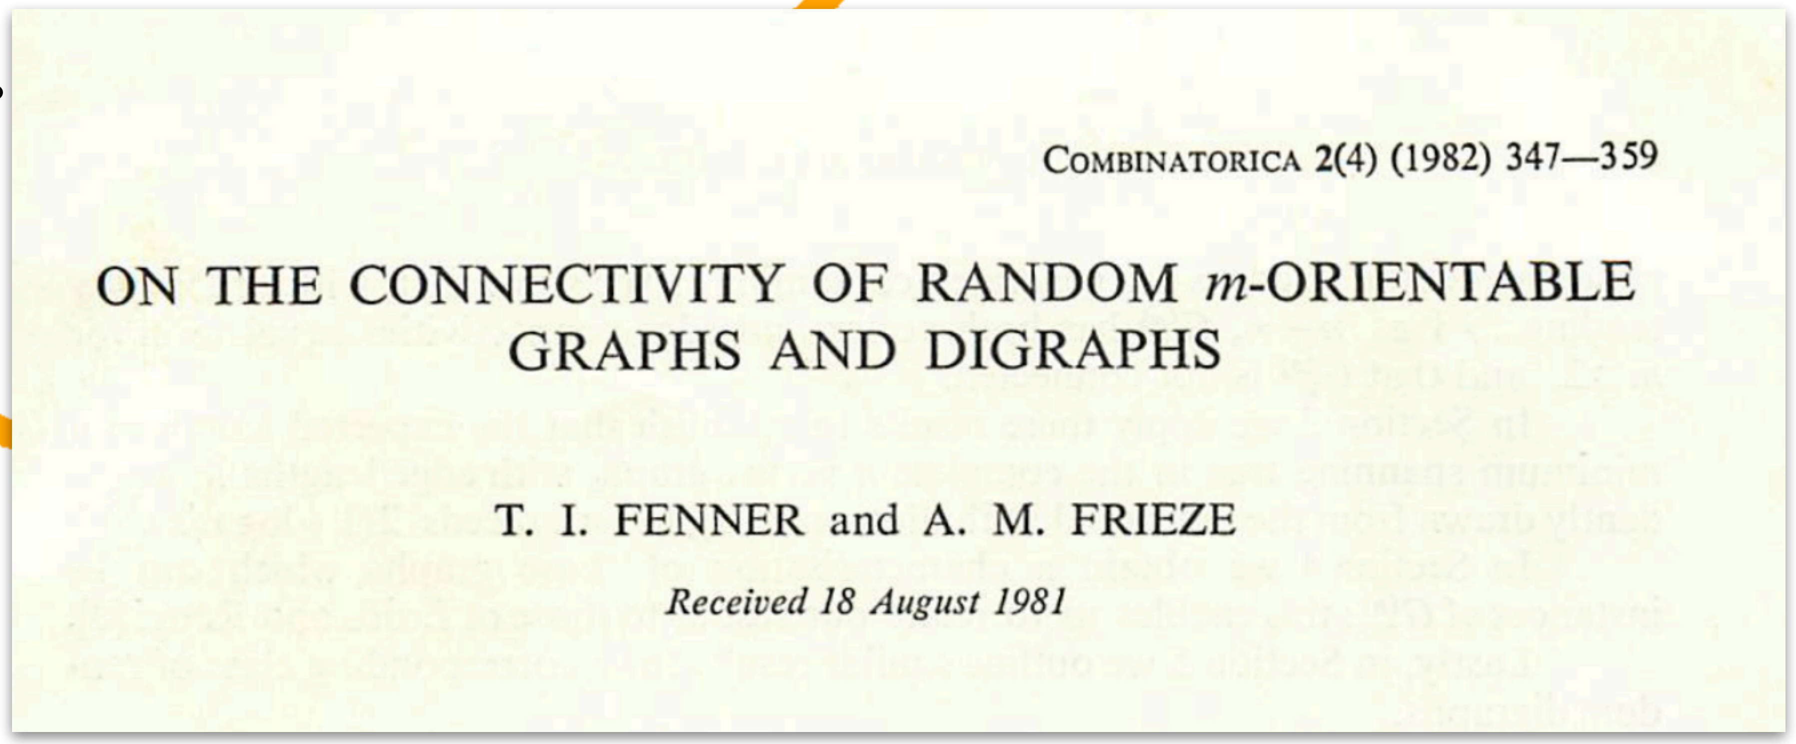

COMBINATORICA 2(4) (1982) 347—359

ON THE CONNECTIVITY OF RANDOM *m*-ORIENTABLE GRAPHS AND DIGRAPHS

T. I. FENNER and A. M. FRIEZE

Received 18 August 1981

Part 1: Road-map

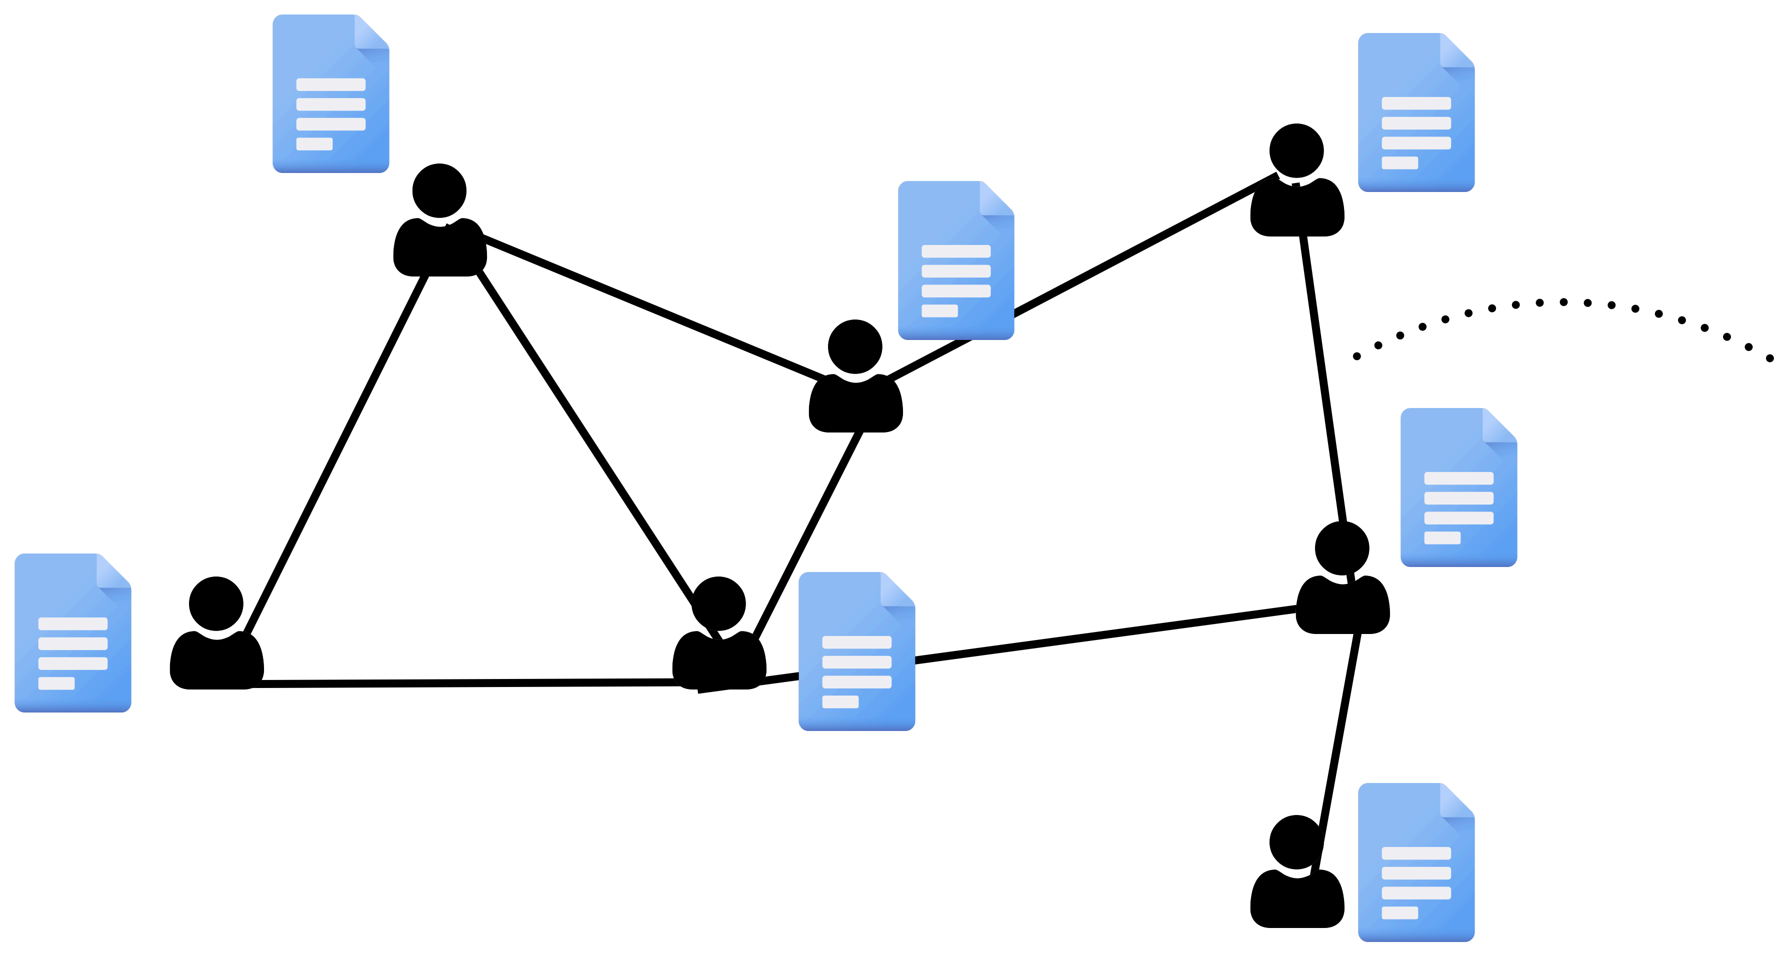

How to securely aggregate data?

[network design]

- ☑ From distributed systems to random graphs
  The sparsity connectivity trade-off
- ☐ What are random K-out graphs?
  What makes them a useful design tool?
- ☐ Random K-out graphs in action
  What graph properties matter?
- ☐ Review: Notions of Connectivity
- ☐ Preview of contributions
  Key proof techniques

Random K-out Graphs

$\mathbb{H}(n, K)$

Each node selects K neighboring nodes

chosen uniformly at random from all n-1 nodes

Edge (i, j) exists if

node i selects node j

or

node j selects node i

Random K-out Graphs

$\mathbb{H}(n, K)$

Each node selects K neighboring nodes

chosen uniformly at random from all n-1 nodes

Edge (i, j) exists if

node i selects node j

or

node j selects node i

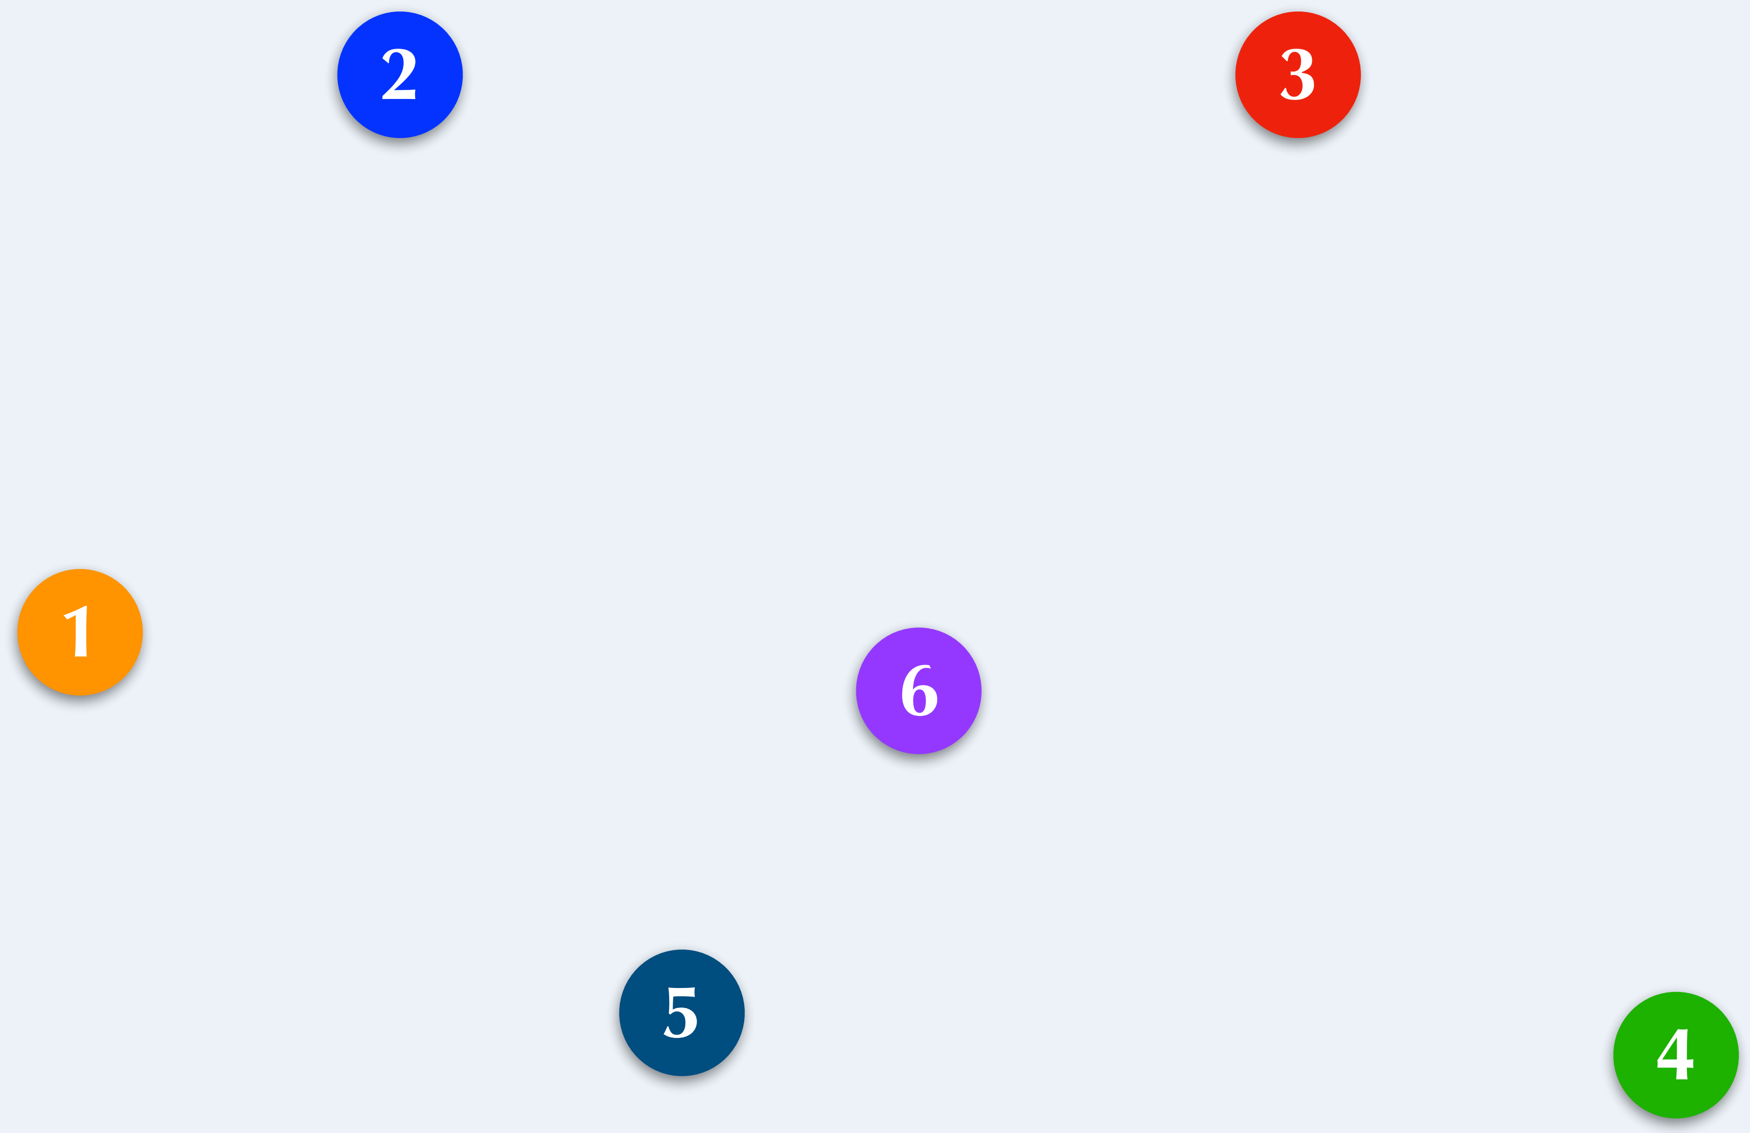

Random 2-out graph on 6 nodes

$(K = 2,\ n = 6)$

Random K-out Graphs

$\mathbb{H}(n, K)$

Each node selects K neighboring nodes

chosen uniformly at random from all n-1 nodes

Edge (i, j) exists if

node i selects node j

or

node j selects node i

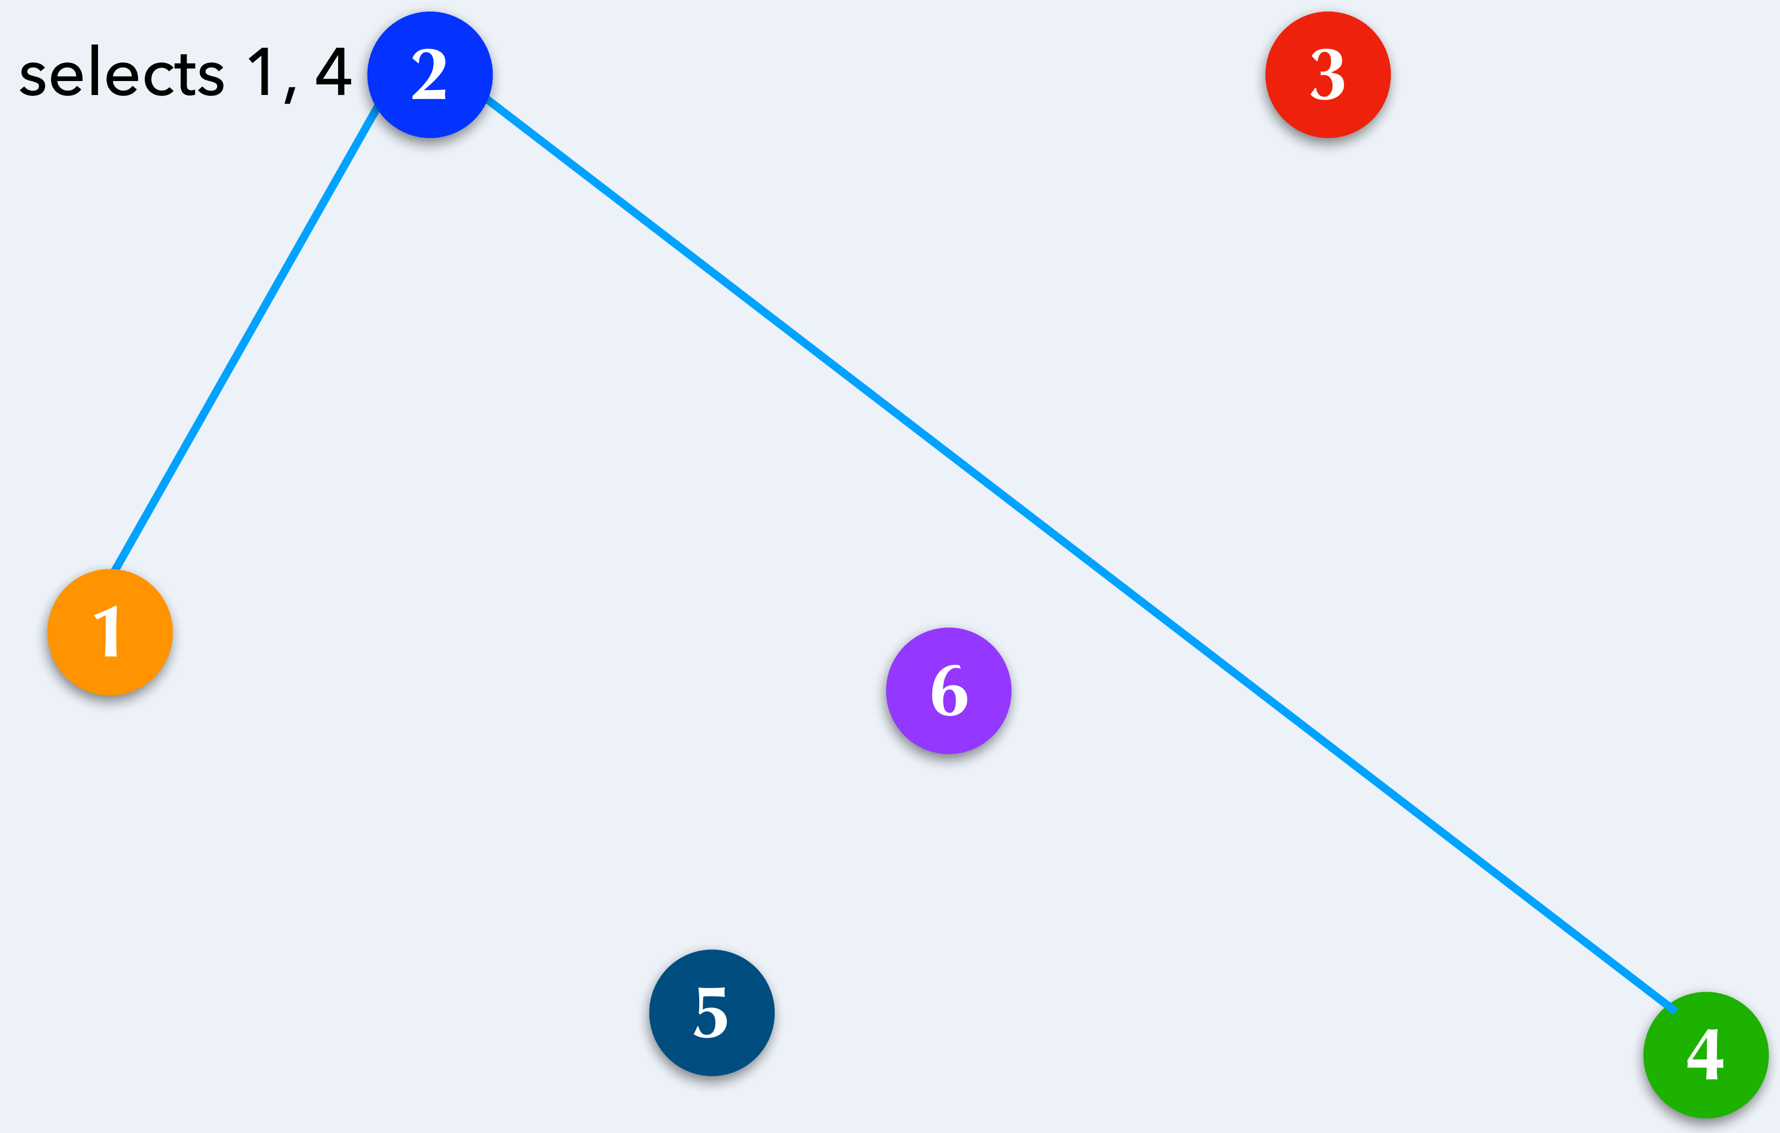

Random 2-out graph on 6 nodes

$(K = 2,\ n = 6)$

Random K-out Graphs

$\mathbb{H}(n, K)$

Each node selects K neighboring nodes

chosen uniformly at random from all n-1 nodes

Edge (i, j) exists if

node i selects node j

or

node j selects node i

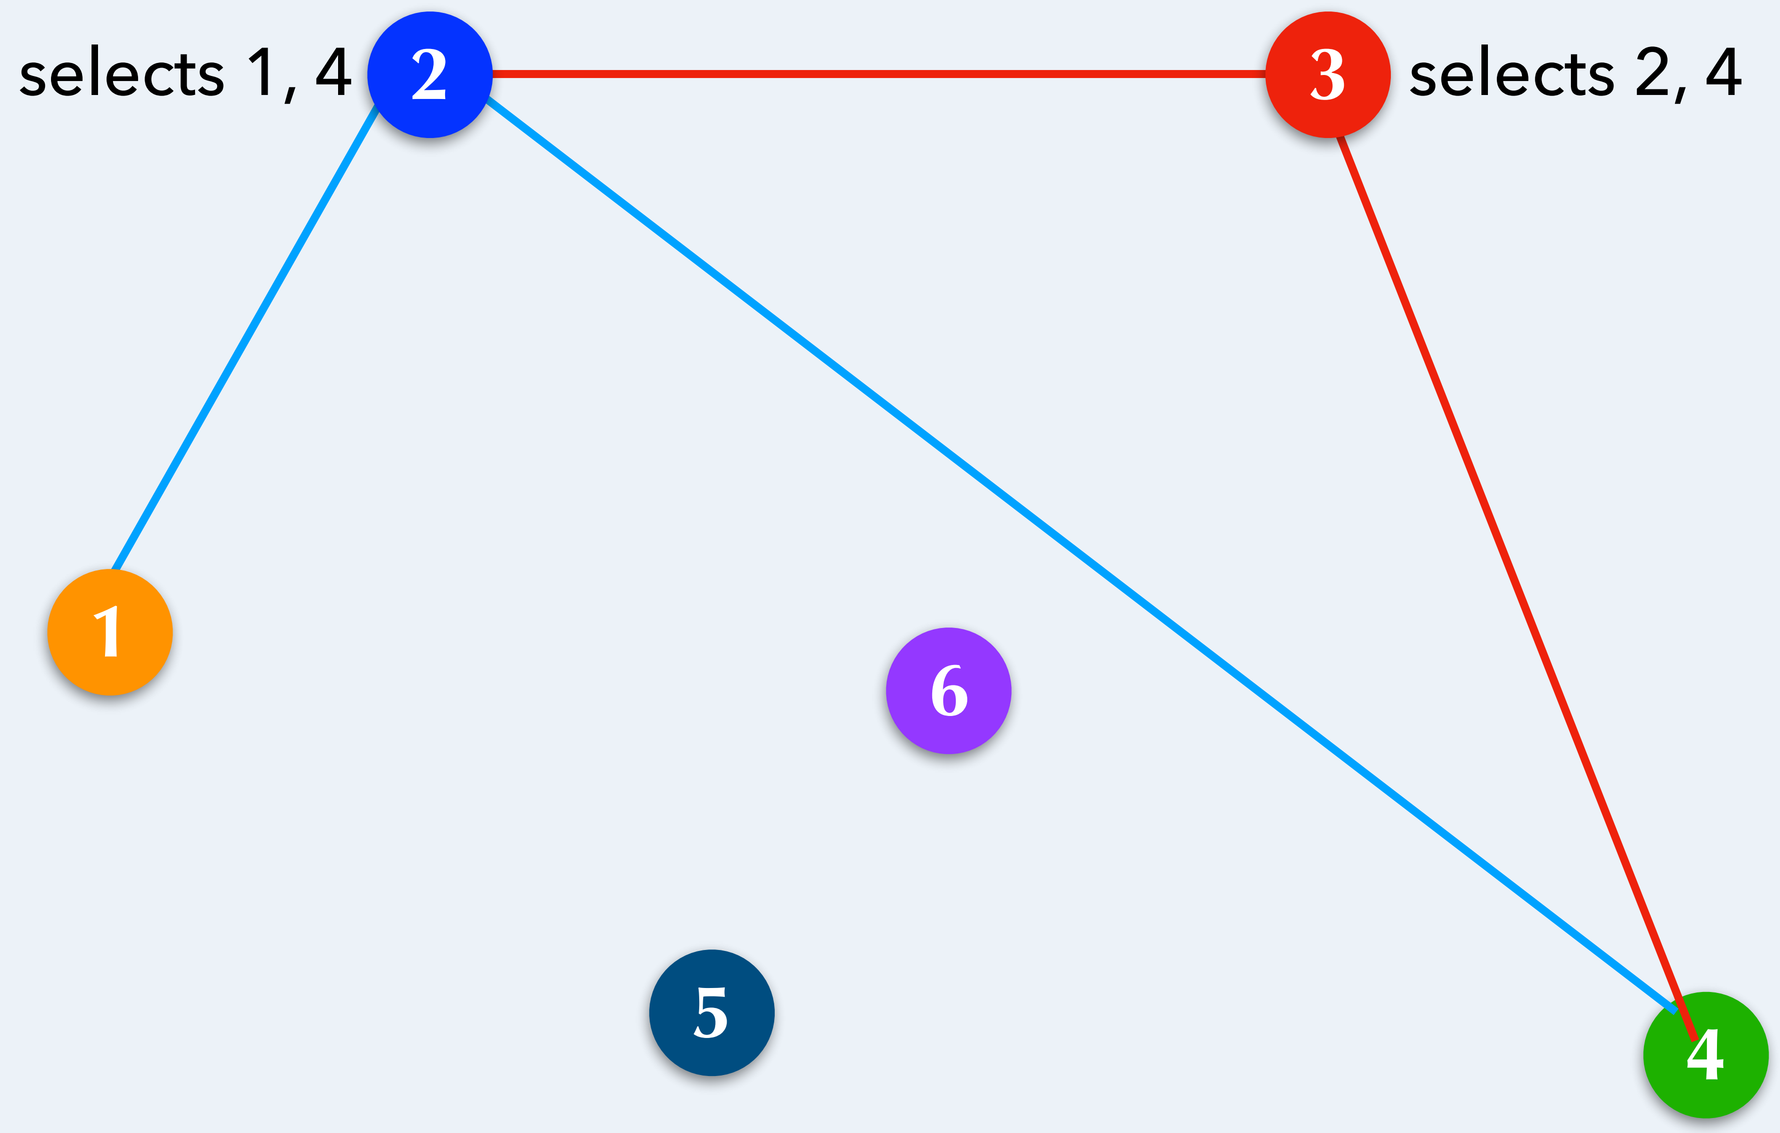

Random 2-out graph on 6 nodes

$(K = 2,\ n = 6)$

Random K-out Graphs

$\mathbb{H}(n, K)$

Each node selects K neighboring nodes

chosen uniformly at random from all n-1 nodes

Edge (i, j) exists if

node i selects node j

or

node j selects node i

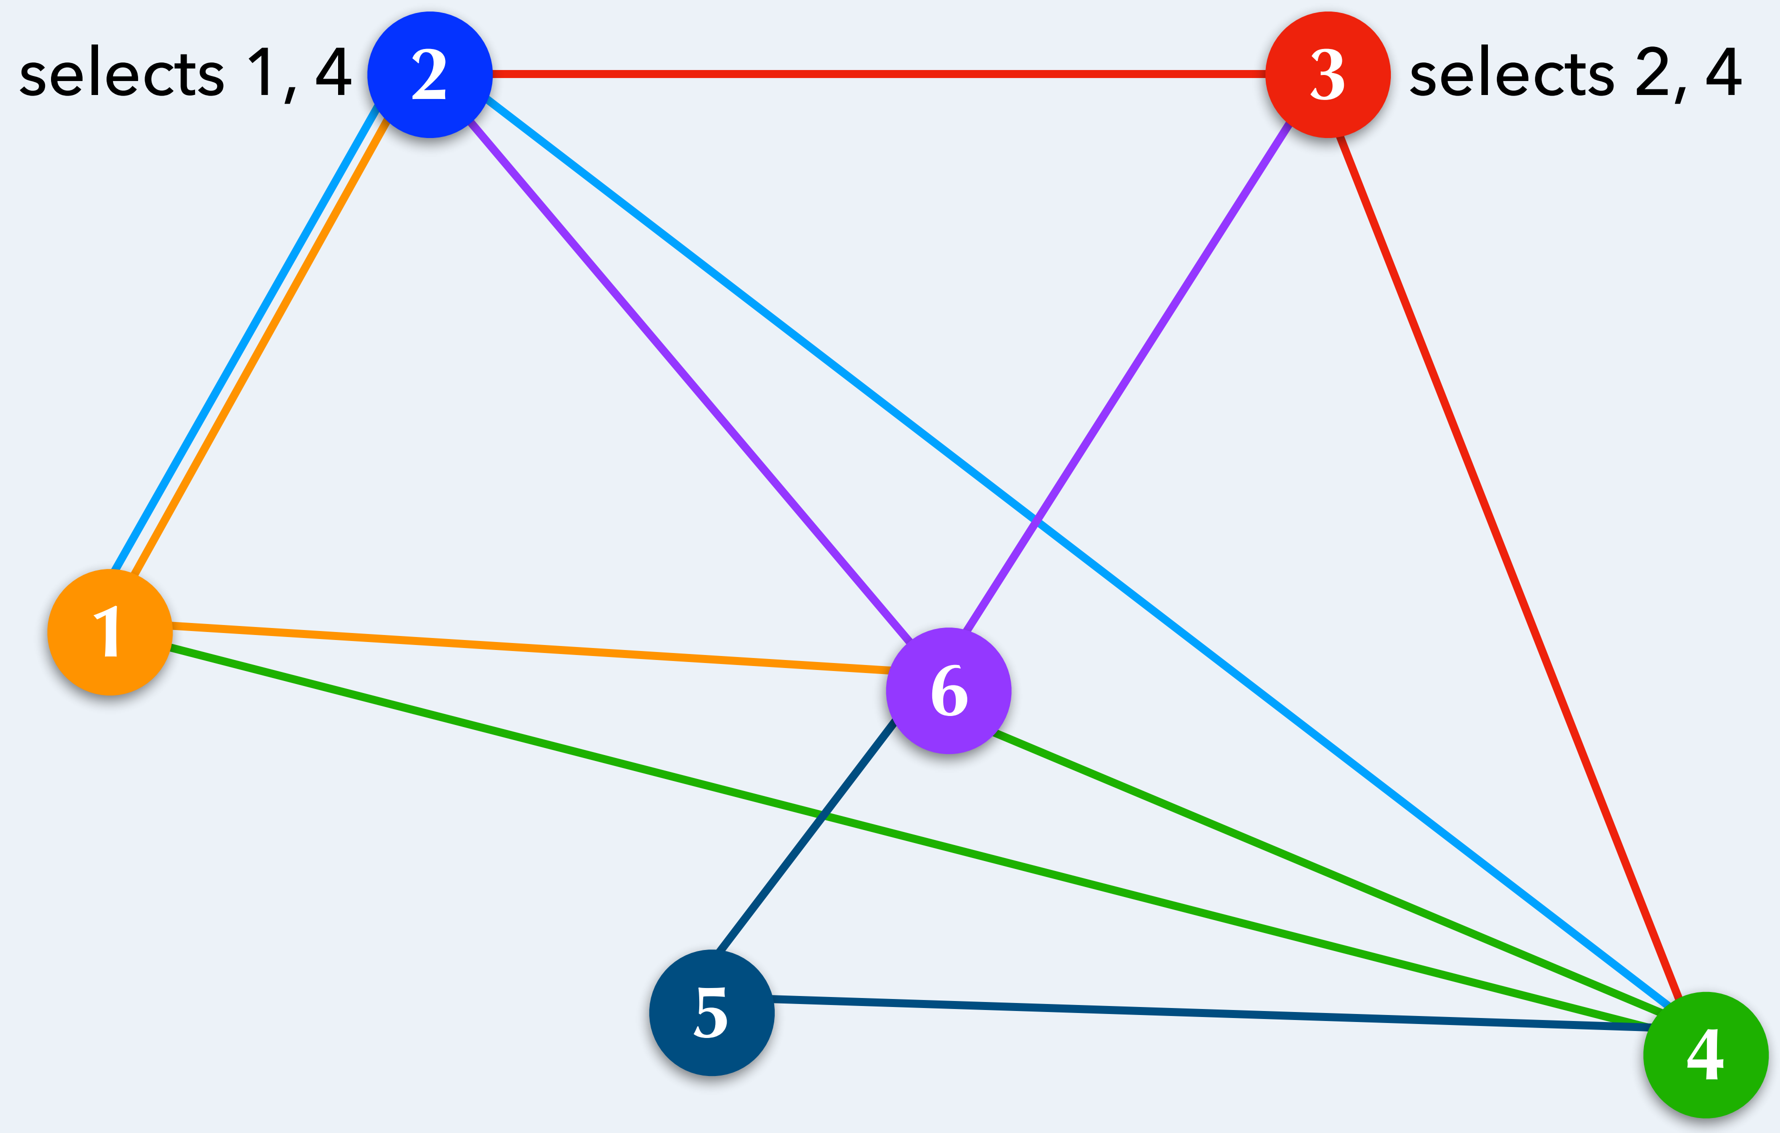

Random 2-out graph on 6 nodes

$(K = 2,\ n = 6)$

Random K-out Graphs

$\mathbb{H}(n, K)$

Each node selects K neighboring nodes

chosen uniformly at random from all n-1 nodes

Edge (i, j) exists if

node i selects node j

or

node j selects node i

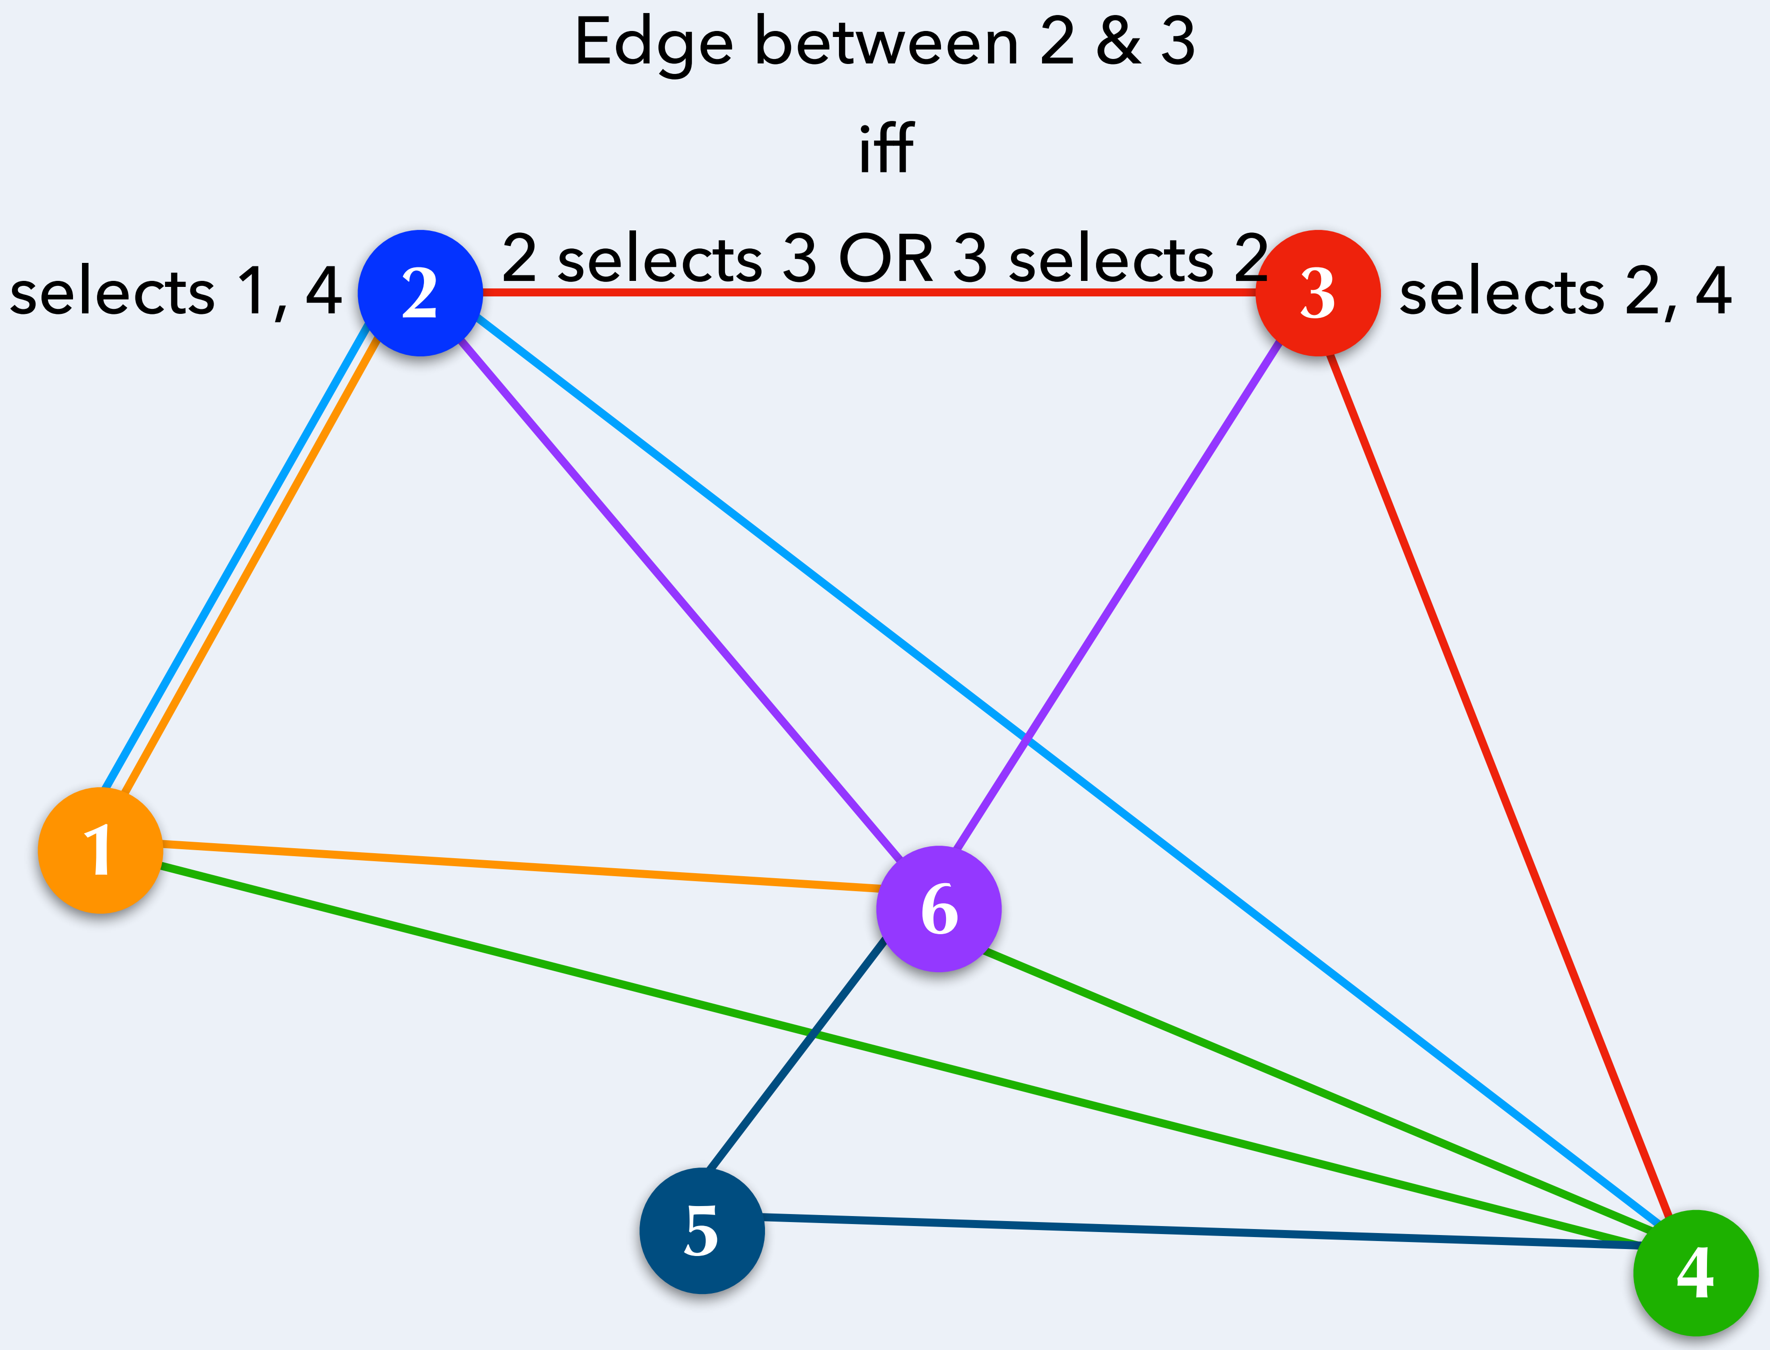

Random 2-out graph on 6 nodes

$(K = 2,\ n = 6)$

Random K-out Graphs

$\mathbb{H}(n, K)$

Each node selects K neighboring nodes

chosen uniformly at random from all n-1 nodes

Edge (i, j) exists if

node i selects node j

or

node j selects node i

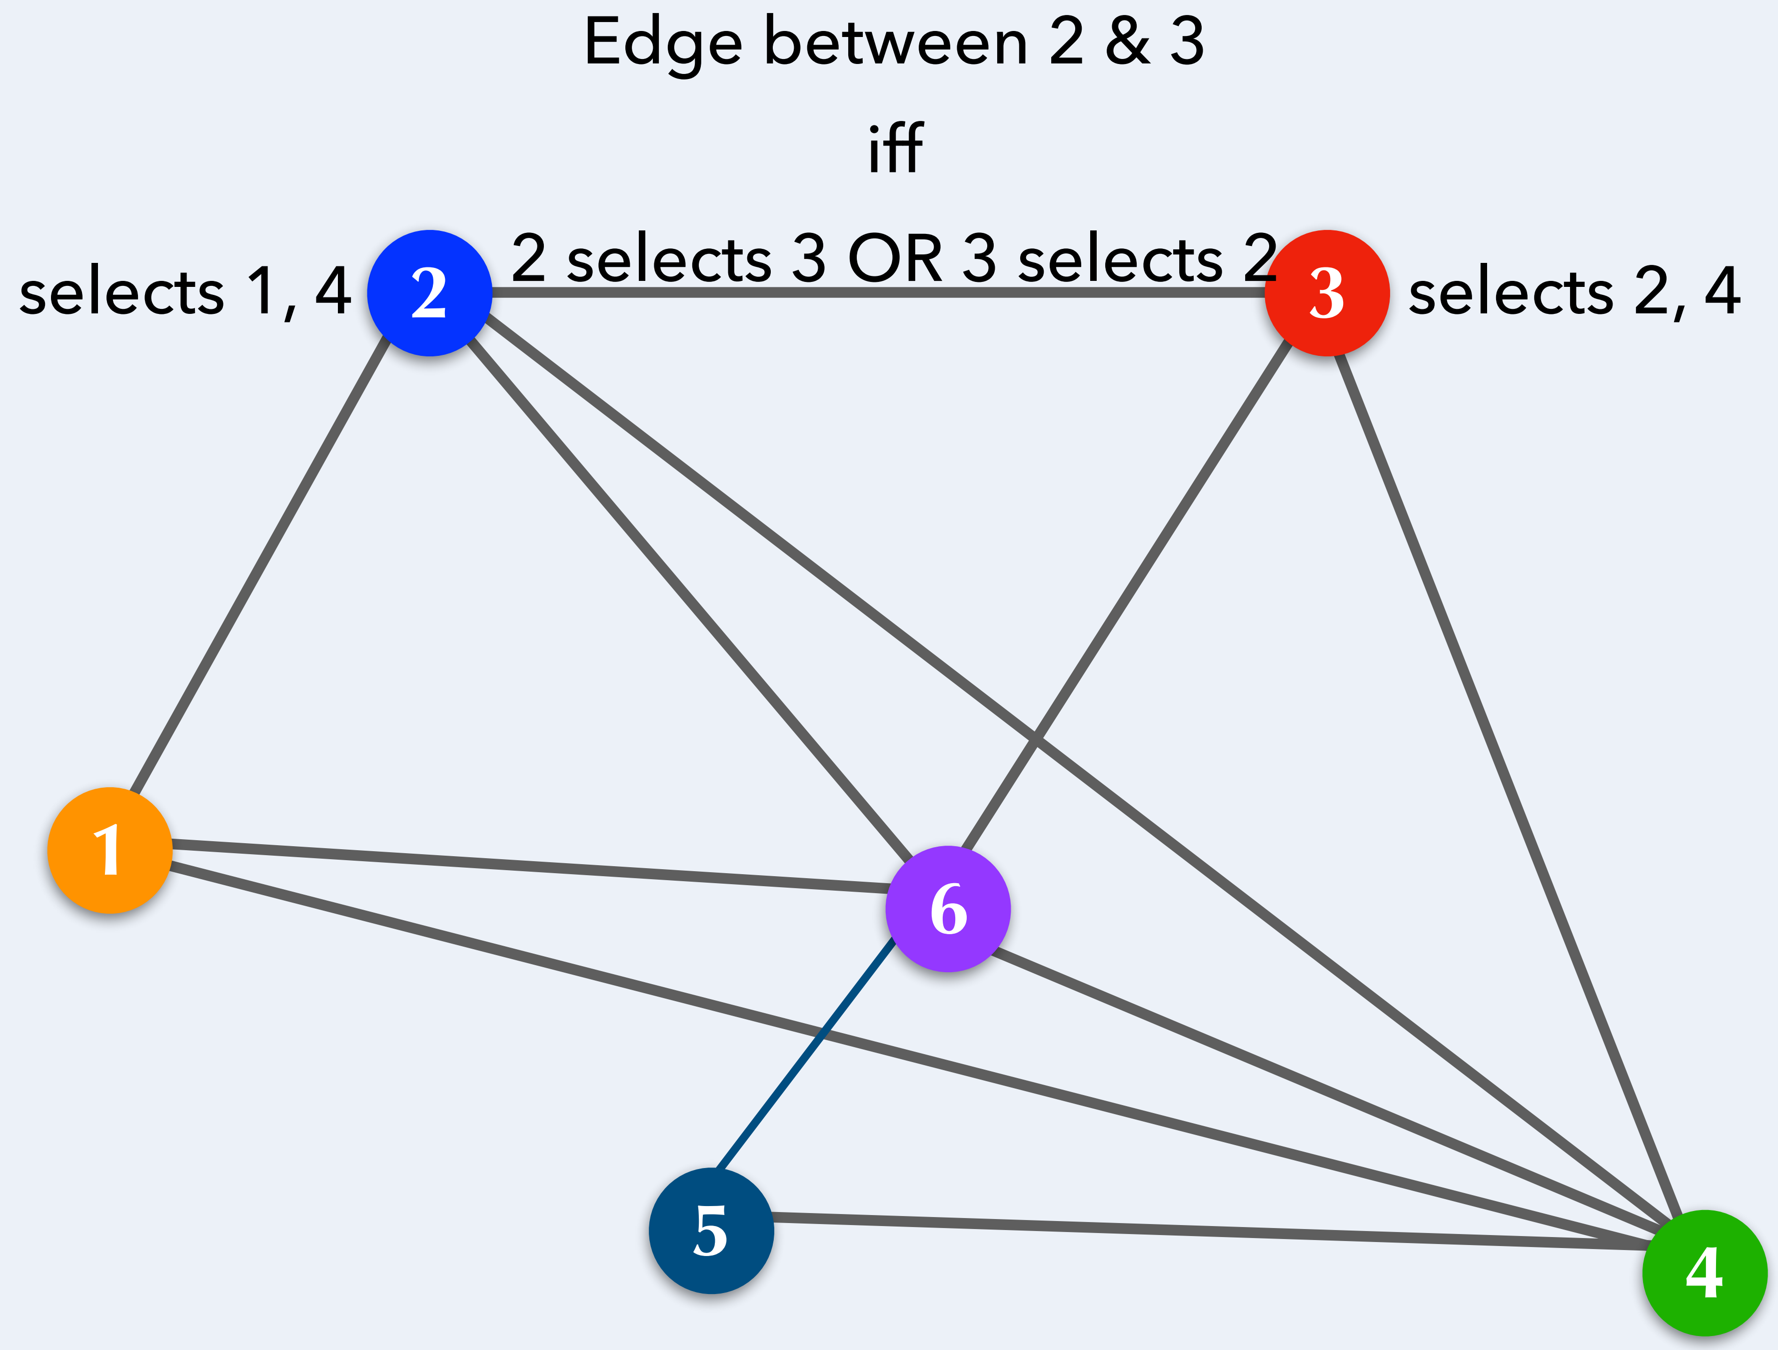

Random 2-out graph on 6 nodes

$(K = 2,\ n = 6)$

Connectivity

[Fenner and Frieze '82]

Random K-out Graphs $\mathbb{H}(n, K)$

For $K \geq 2$, *connected* with high probability (with probability $\rightarrow 1$ as # nodes $\rightarrow \infty$).

(for $K = 1$, *disconnected* with high probability)

Connectivity

Random K-out Graphs $\mathbb{H}(n, K)$

[Fenner and Frieze '82]

For $K \geq 2$, *connected* with high probability (with probability $\rightarrow 1$ as # nodes $\rightarrow \infty$).

(for $K = 1$, *disconnected* with high probability)

With average degree ~ 4, we get connectivity whp

P [edge between i and j] = $1 - \left(1 - \frac{K}{n-1}\right)^2$

Average node degree = $2K - \left(\frac{K^2}{n-1}\right)$

Connectivity

[Fenner and Frieze '82]

Random K-out Graphs $\mathbb{H}(n, K)$

For $K \geq 2$, *connected* with high probability (with probability $\rightarrow 1$ as # nodes $\rightarrow \infty$).

(for $K = 1$, *disconnected* with high probability)

With average degree ~ 4, we get connectivity whp

Erdos Renyi Random Graphs $\mathbb{G}(n, p)$

In contrast Erdos Renyi random graphs require average degree ~ $\log n$ for connectivity whp

scales with n

Connectivity

[Fenner and Frieze '82]

Random K-out Graphs $\mathbb{H}(n, K)$

For $K \geq 2$, *connected* with high probability (with probability $\to 1$ as # nodes $\to \infty$).

(for $K = 1$, *disconnected* with high probability)

Theorem [Sood and Yagan, ICC'21*]

$$\mathbb{P}[\ \mathbb{H}(n, K) \text{ is } \textit{connected}\ ] = 1 - \Theta(1/n^{K^2-1}),\ K \geq 2$$

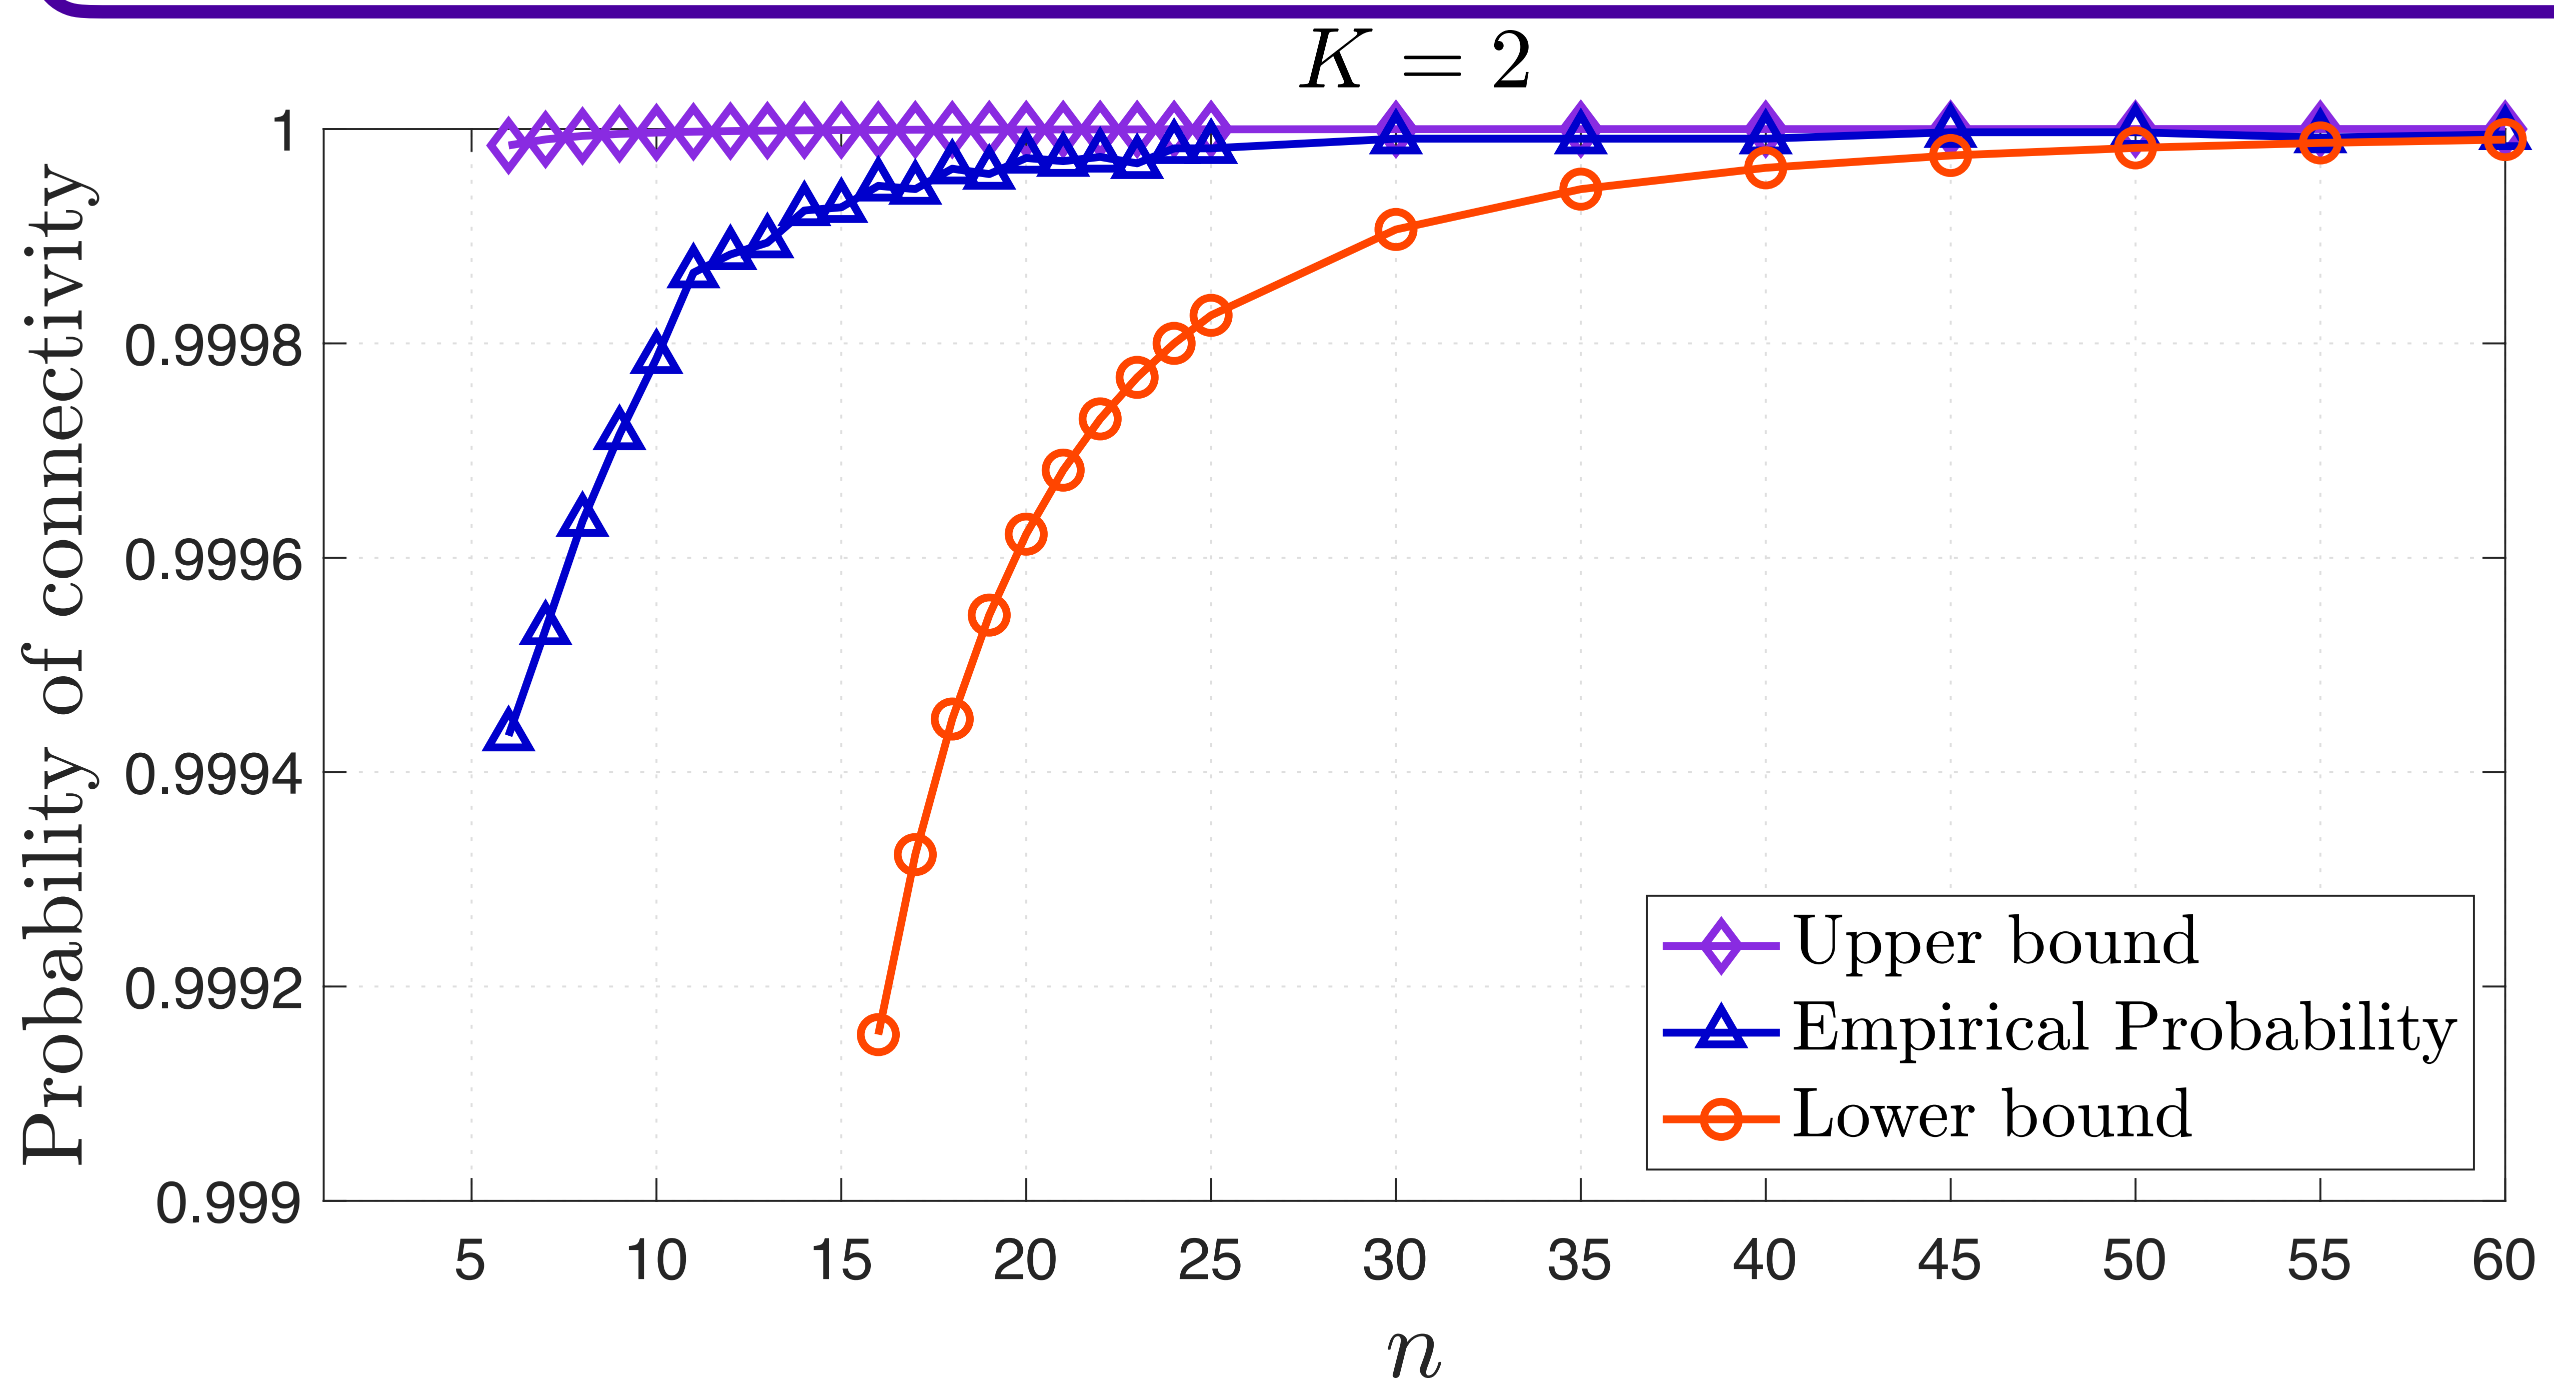

*ICC' 21 Best Paper Award, Communication Theory

What if K is not same for all nodes?

So far…

For (homogeneous) random K-out graphs, $p_{\text{connectivity}} = 1 - \Theta(1/n^{K^2-1}),\ K \geq 2$

What if some nodes make fewer than 2 selections?

Inhomogeneous Random K-out Graphs

- Each node is assigned a type which determines the number of selections
- Nodes can make fewer than 2 selections

Inhomogeneous Random K-out Graphs

$\mathbb{H}(n, \mu, K_n)$

$\boldsymbol{n}$: number of nodes

Label nodes independently as
Type-I wp $\boldsymbol{\mu}$ (>0), Type-II wp $\mathbf{1}$-$\boldsymbol{\mu}$

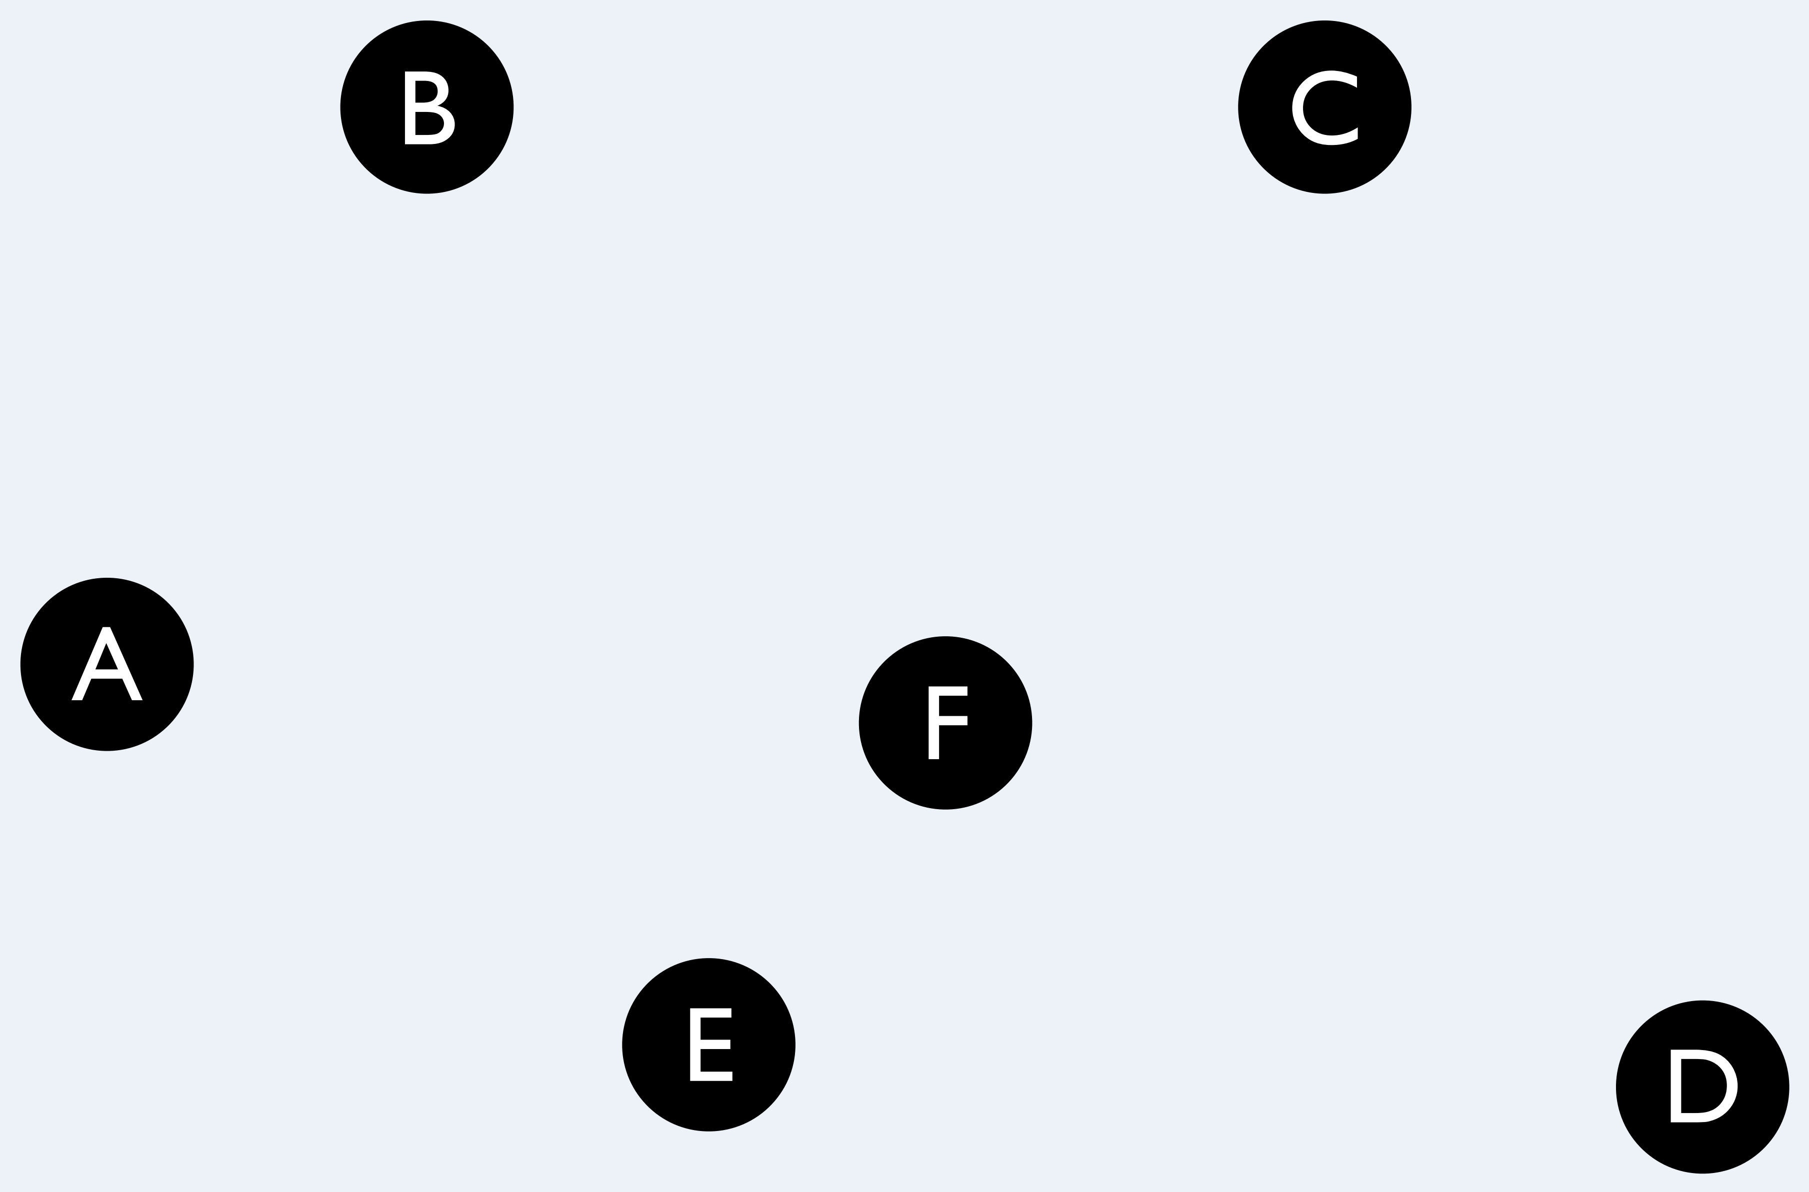

Inhomogeneous K-out Random graph (n = 6, $\boldsymbol{K_n}$ = 3)

Inhomogeneous Random K-out Graphs

$\mathbb{H}(n, \mu, K_n)$

$\boldsymbol{n}$: number of nodes

Label nodes independently as
Type-I wp $\boldsymbol{\mu}$ (>0), Type-II wp $\mathbf{1}$-$\boldsymbol{\mu}$

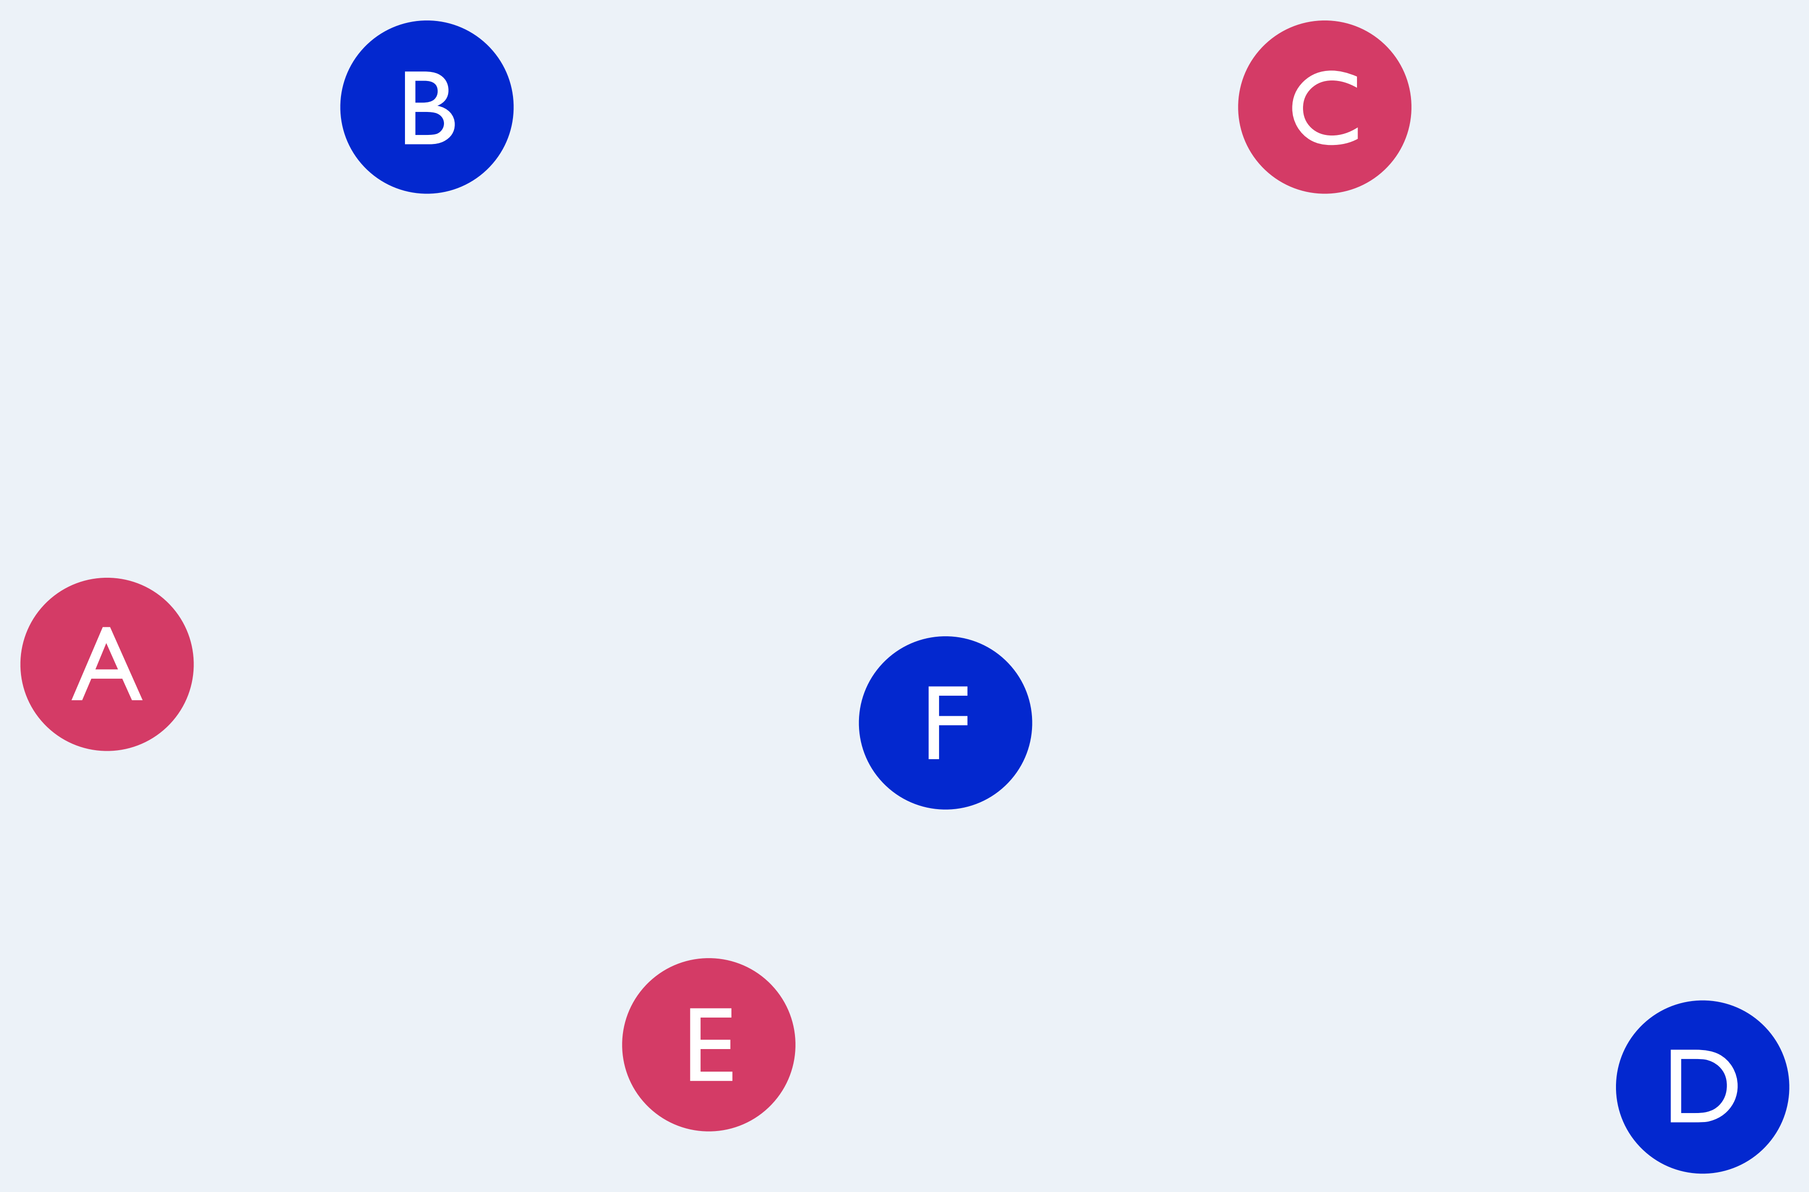

Inhomogeneous K-out Random graph (n = 6, $\boldsymbol{K_n}$ = 3)

Inhomogeneous Random K-out Graphs

$\mathbb{H}(n, \mu, K_n)$

$\boldsymbol{n}$: number of nodes

Label nodes independently as
Type-I wp $\boldsymbol{\mu}$ (>0), Type-II wp $\mathbf{1}$-$\boldsymbol{\mu}$

Type-I nodes select $\mathbf{1}$ node,
Type-II nodes select K_n (≥2) nodes
(uniformly at random from all n-1 nodes)

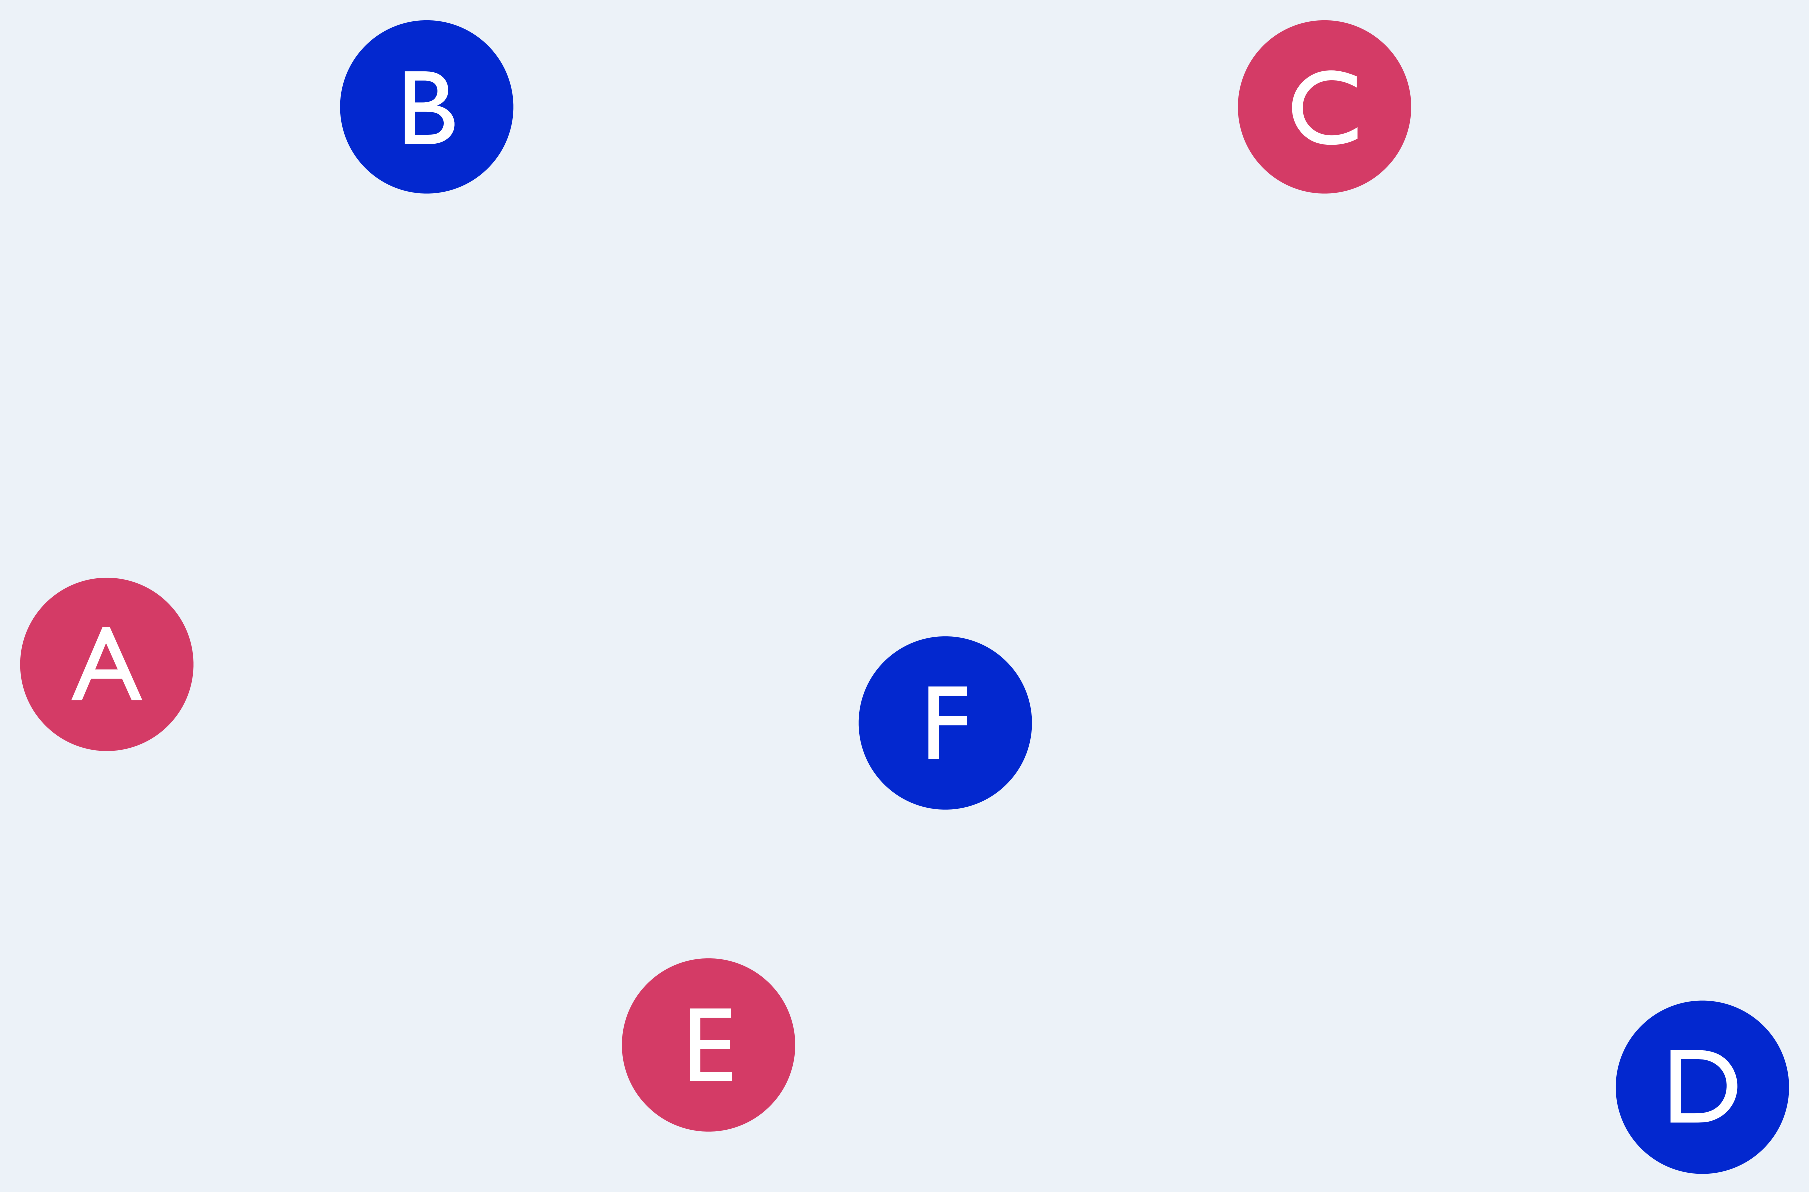

Inhomogeneous K-out Random graph (n = 6, $\boldsymbol{K_n}$ = 3)

Inhomogeneous Random K-out Graphs

$\mathbb{H}(n, \mu, K_n)$

$\boldsymbol{n}$: number of nodes

Label nodes independently as
Type-I wp $\boldsymbol{\mu}$ (>0), Type-II wp $\mathbf{1}$-$\boldsymbol{\mu}$

Type-I nodes select $\mathbf{1}$ node,
Type-II nodes select K_n (≥2) nodes
(uniformly at random from all n-1 nodes)

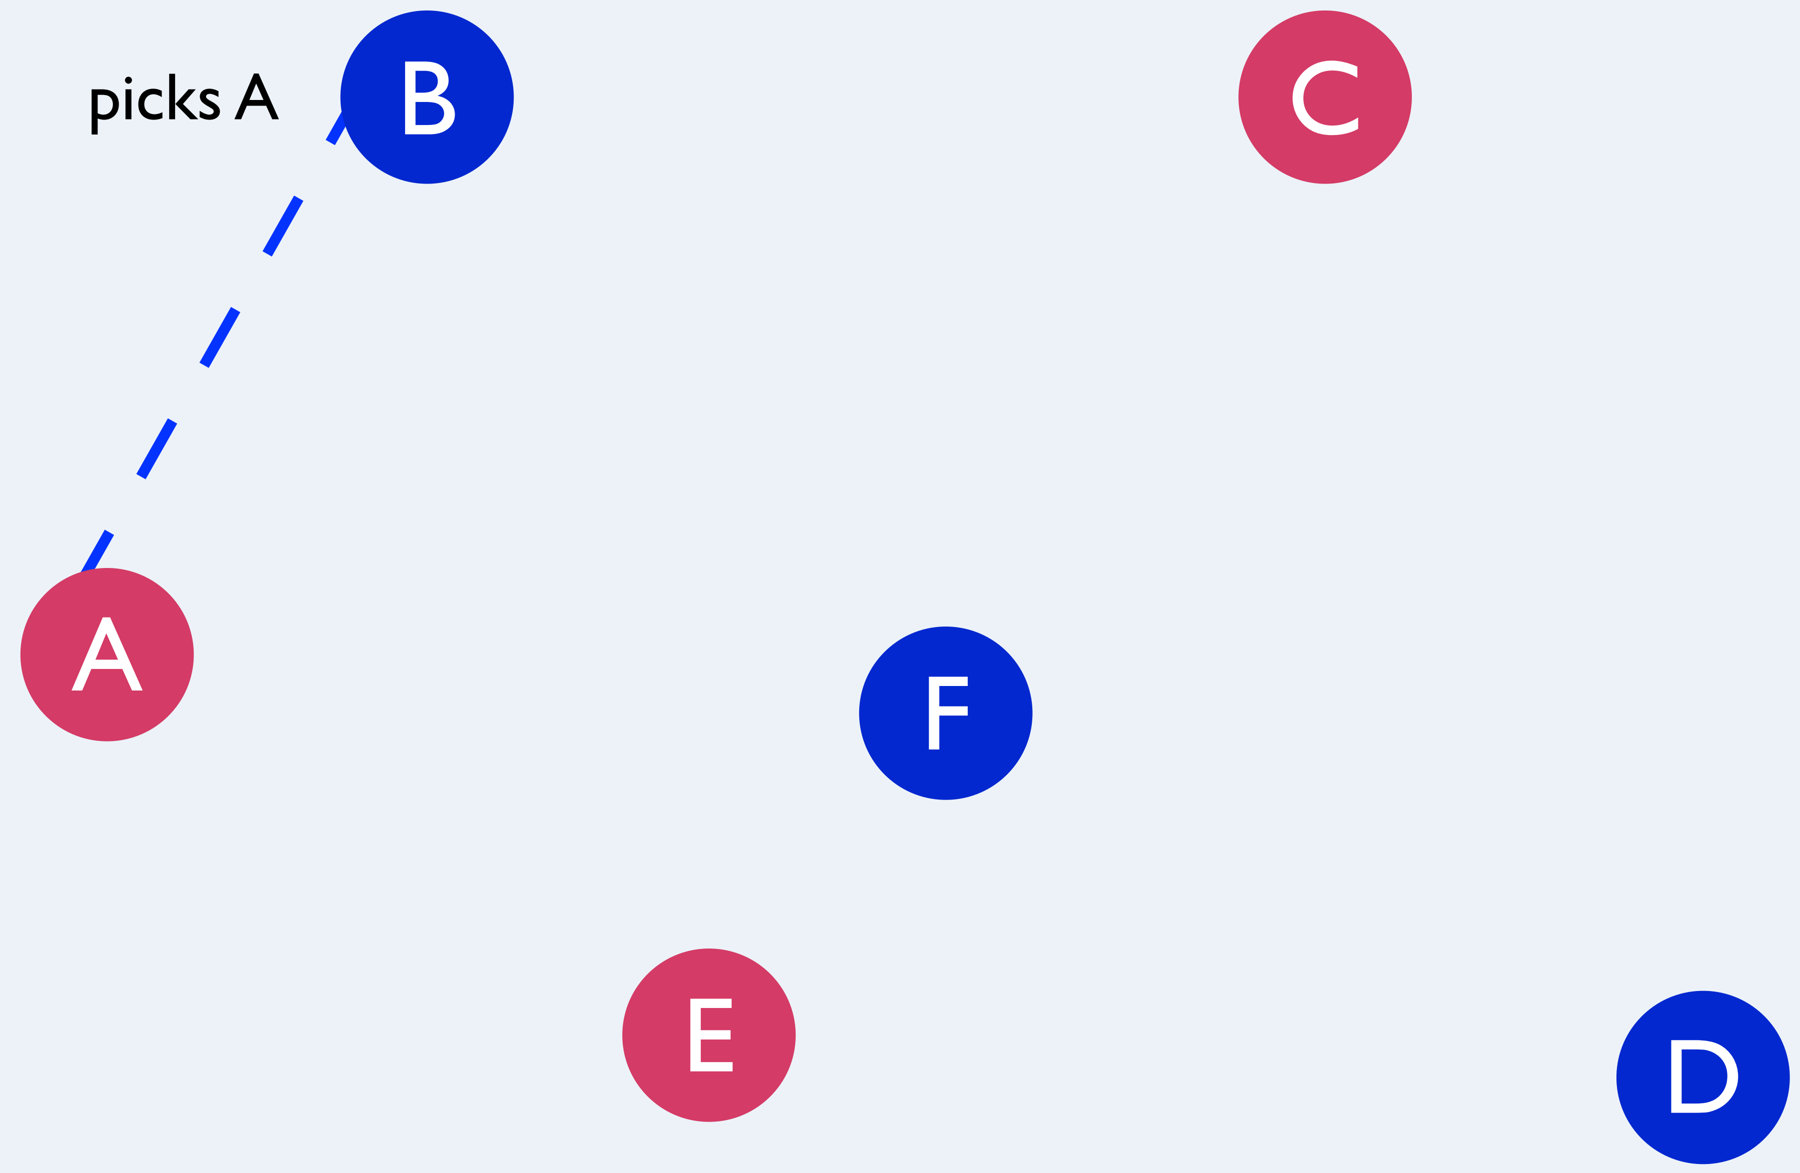

Inhomogeneous K-out Random graph (n = 6, $\boldsymbol{K_n}$ = 3)

Inhomogeneous Random K-out Graphs

$\mathbb{H}(n, \mu, K_n)$

$\boldsymbol{n}$: number of nodes

Label nodes independently as
Type-I wp $\boldsymbol{\mu}$ (>0), Type-II wp $\mathbf{1}$-$\boldsymbol{\mu}$

Type-I nodes select $\mathbf{1}$ node,
Type-II nodes select K_n (≥2) nodes
(uniformly at random from all n-1 nodes)

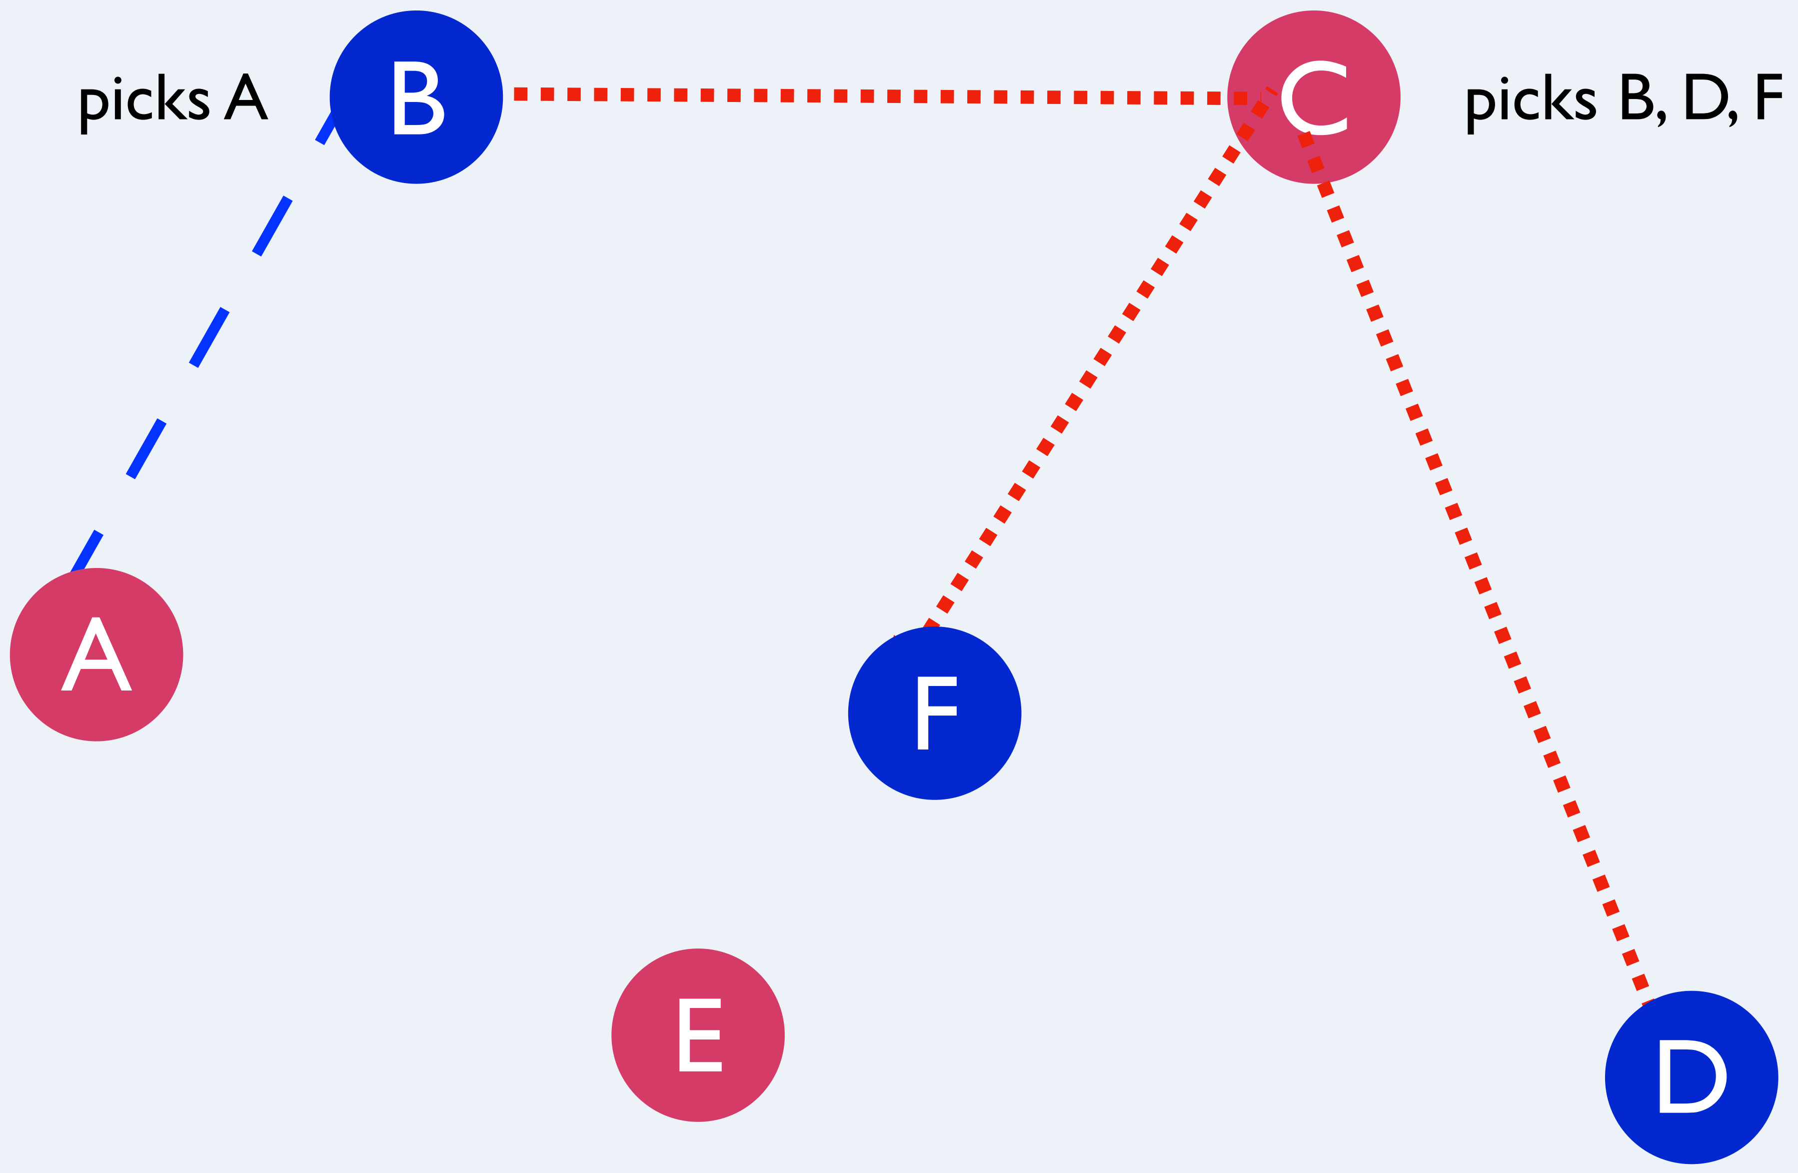

Inhomogeneous K-out Random graph (n = 6, $\boldsymbol{K_n}$ = 3)

Inhomogeneous Random K-out Graphs

$\mathbb{H}(n, \mu, K_n)$

$\boldsymbol{n}$: number of nodes

Label nodes independently as
Type-I wp $\boldsymbol{\mu}$ (>0), Type-II wp $\mathbf{1}\text{-}\boldsymbol{\mu}$

Type-I nodes select $\mathbf{1}$ node,
Type-II nodes select K_n (≥2) nodes
(uniformly at random from all n-1 nodes)

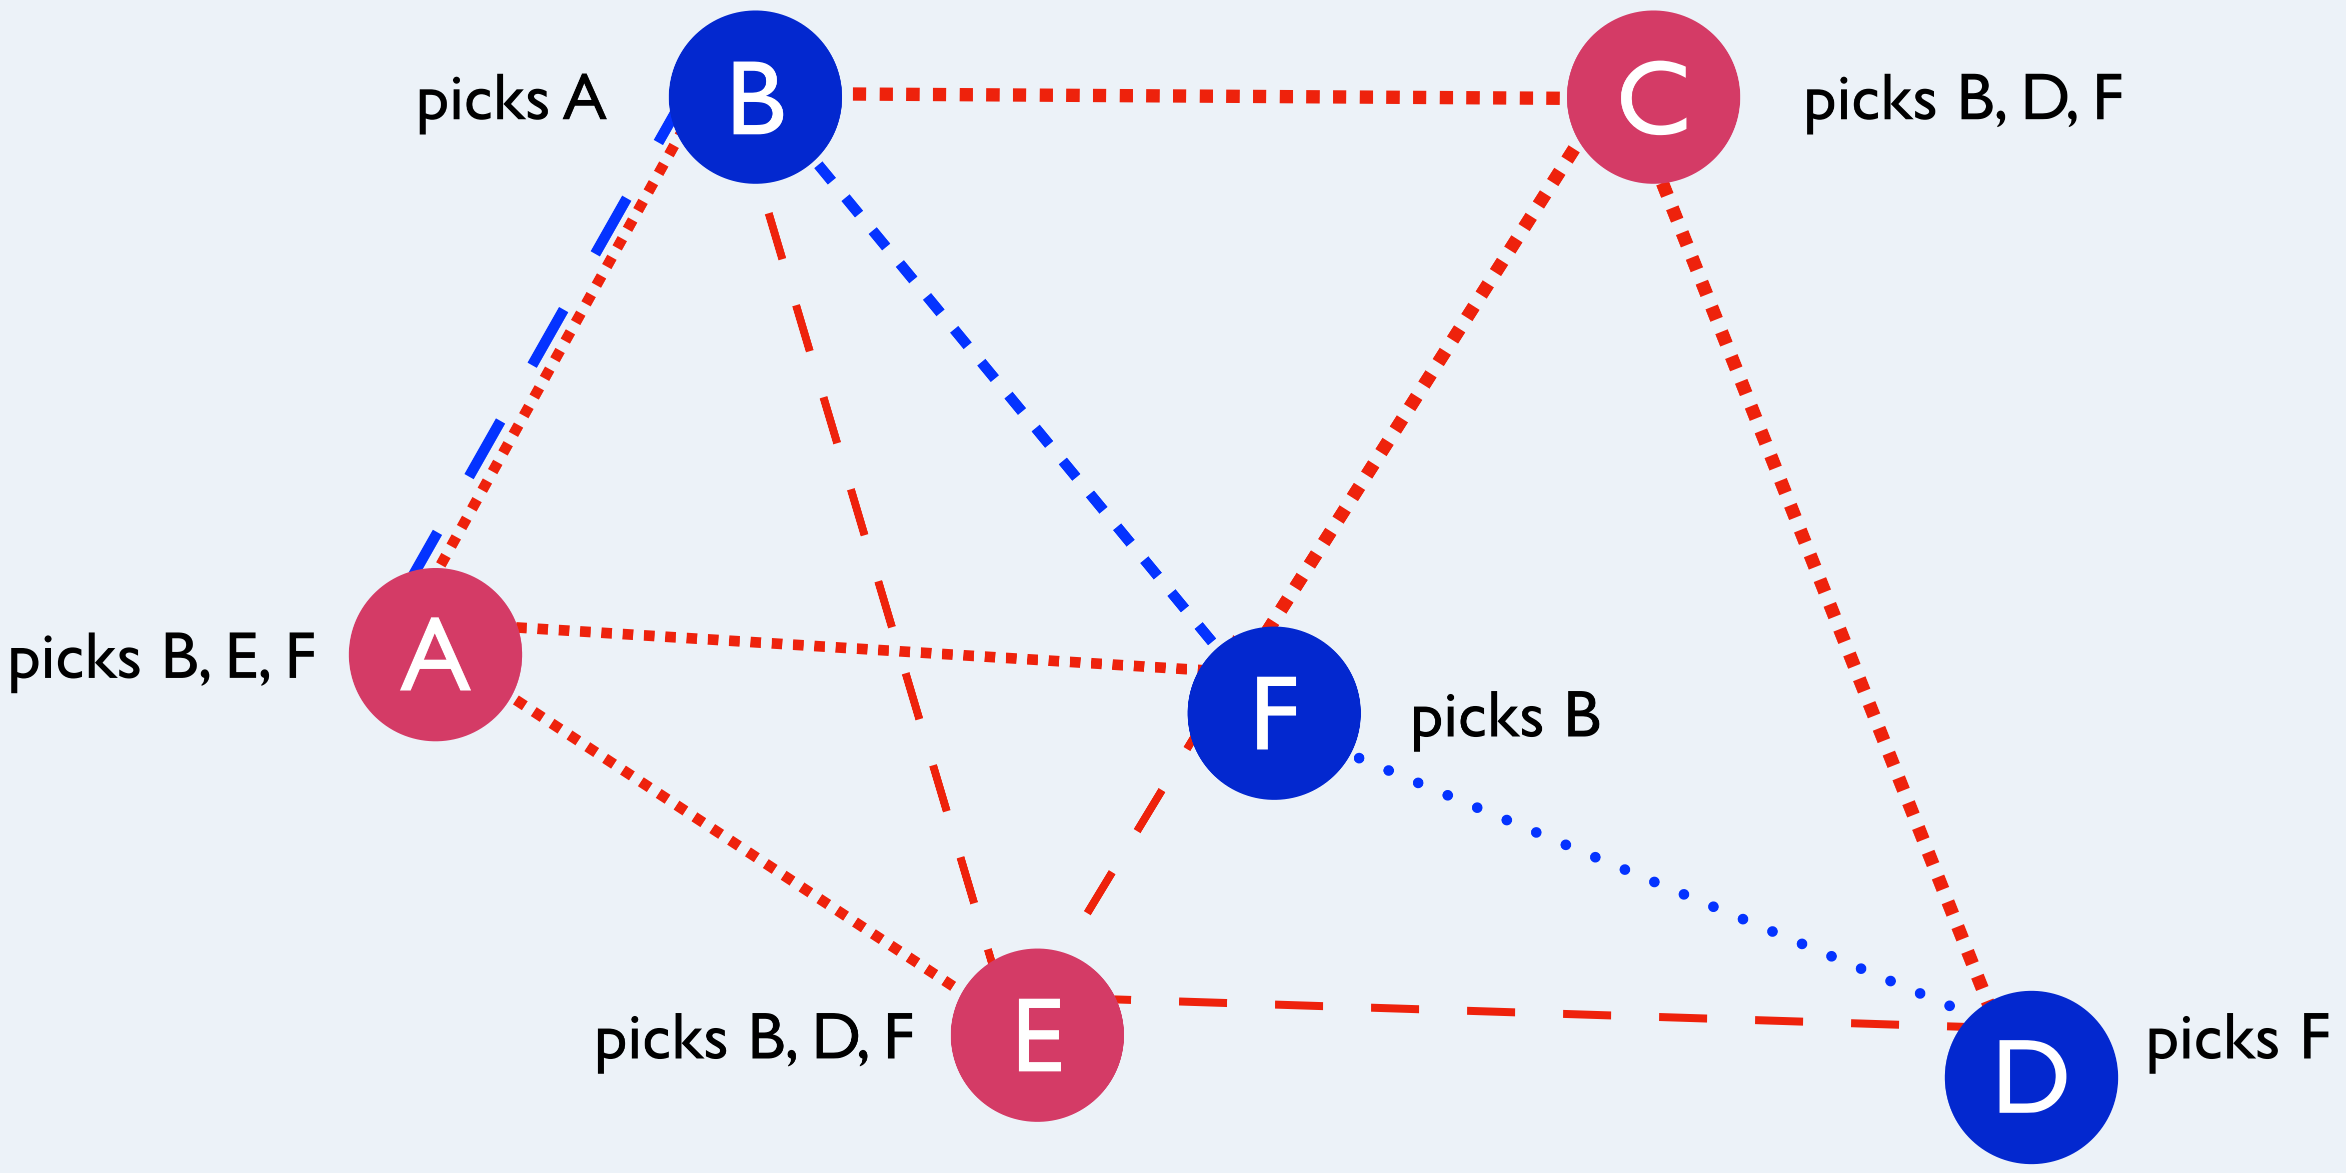

Inhomogeneous K-out Random graph (n = 6, $\boldsymbol{K_n}$ = 3)

Inhomogeneous Random K-out Graphs

$\mathbb{H}(n, \mu, K_n)$

$\boldsymbol{n}$: number of nodes

Label nodes independently as
Type-I wp $\boldsymbol{\mu}$ (>0), Type-II wp $\mathbf{1}$-$\boldsymbol{\mu}$

Type-I nodes select $\mathbf{1}$ node,
Type-II nodes select K_n (≥2) nodes
(uniformly at random from all n-1 nodes)

Edge (i, j) exists if
node i selects node j
or
node j selects node i

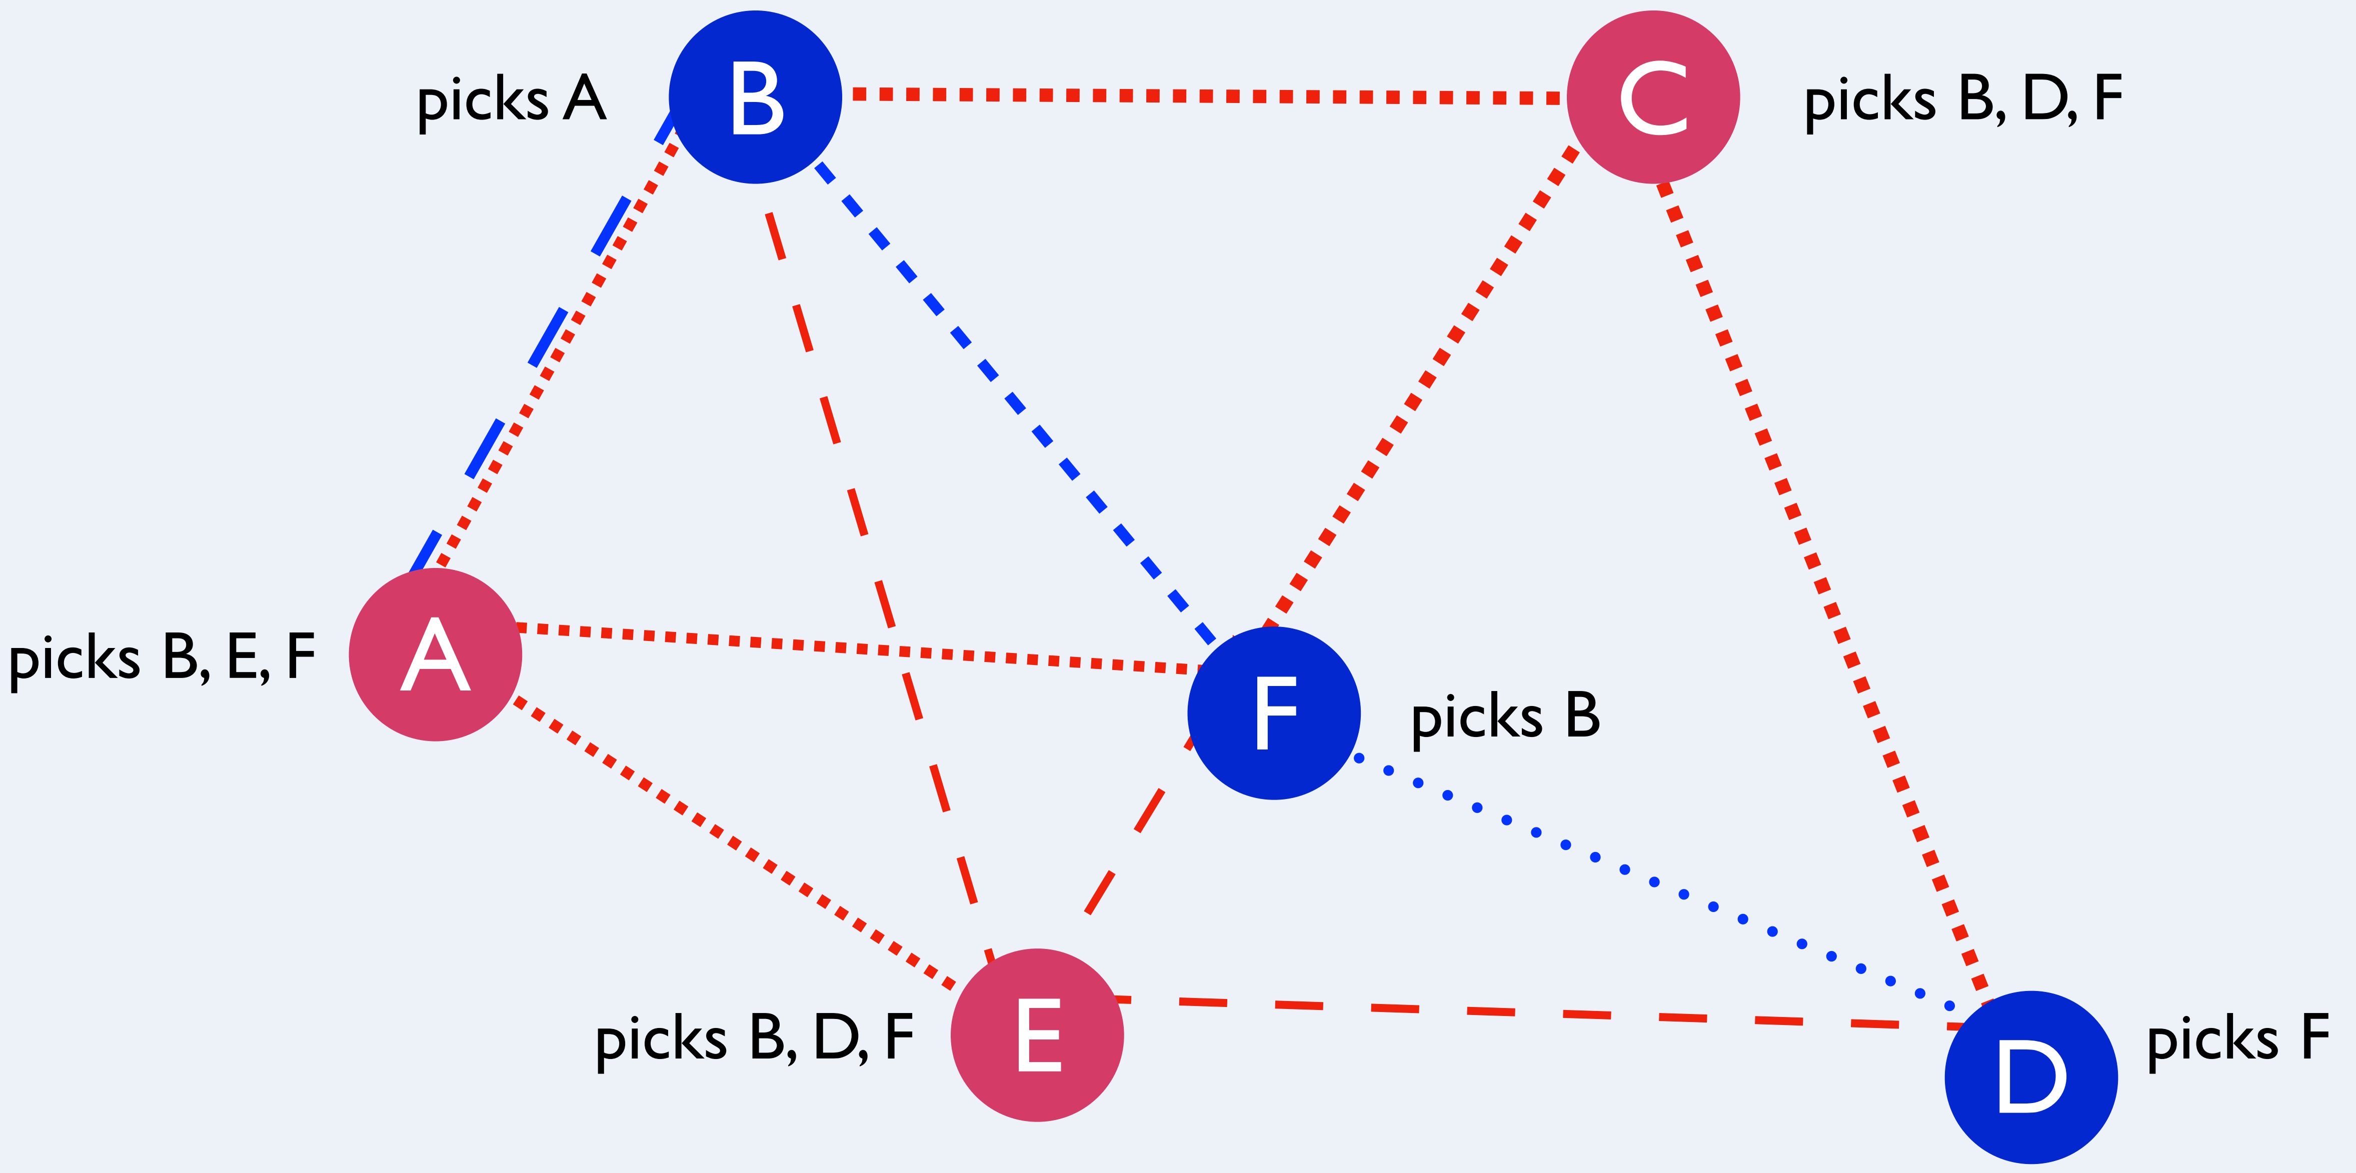

Inhomogeneous K-out Random graph (n = 6, $\boldsymbol{K_n}$ = 3)

Inhomogeneous Random K-out Graphs

$\mathbb{H}(n, \mu, K_n)$

$\boldsymbol{n}$: number of nodes

Label nodes independently as
Type-I wp $\boldsymbol{\mu}$ (>0), Type-II wp $\mathbf{1}$-$\boldsymbol{\mu}$

Type-I nodes select $\mathbf{1}$ node,
Type-II nodes select K_n (≥2) nodes
(uniformly at random from all n-1 nodes)

Edge (i, j) exists if
node i selects node j
or
node j selects node i

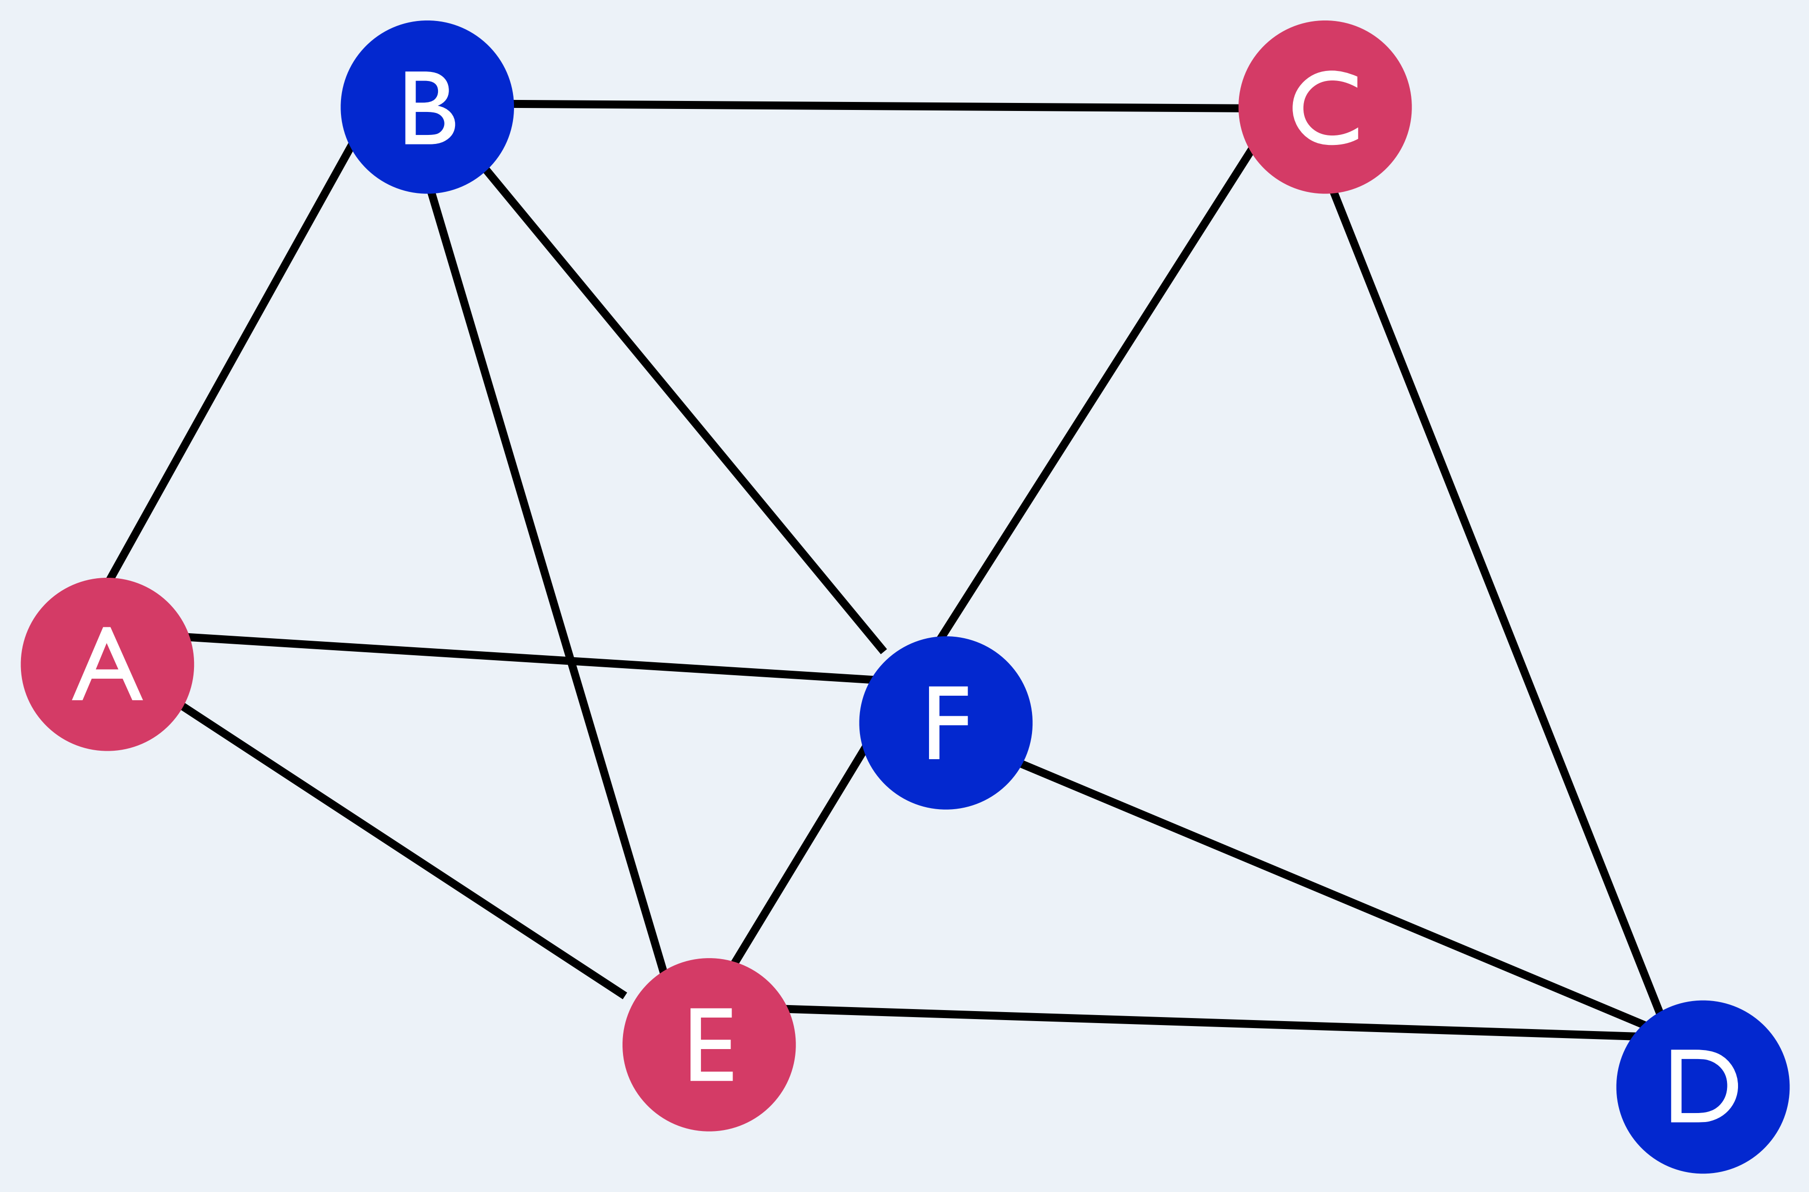

Inhomogeneous K-out Random graph (n = 6, $\boldsymbol{K_n}$ = 3)

Part 1: Road-map

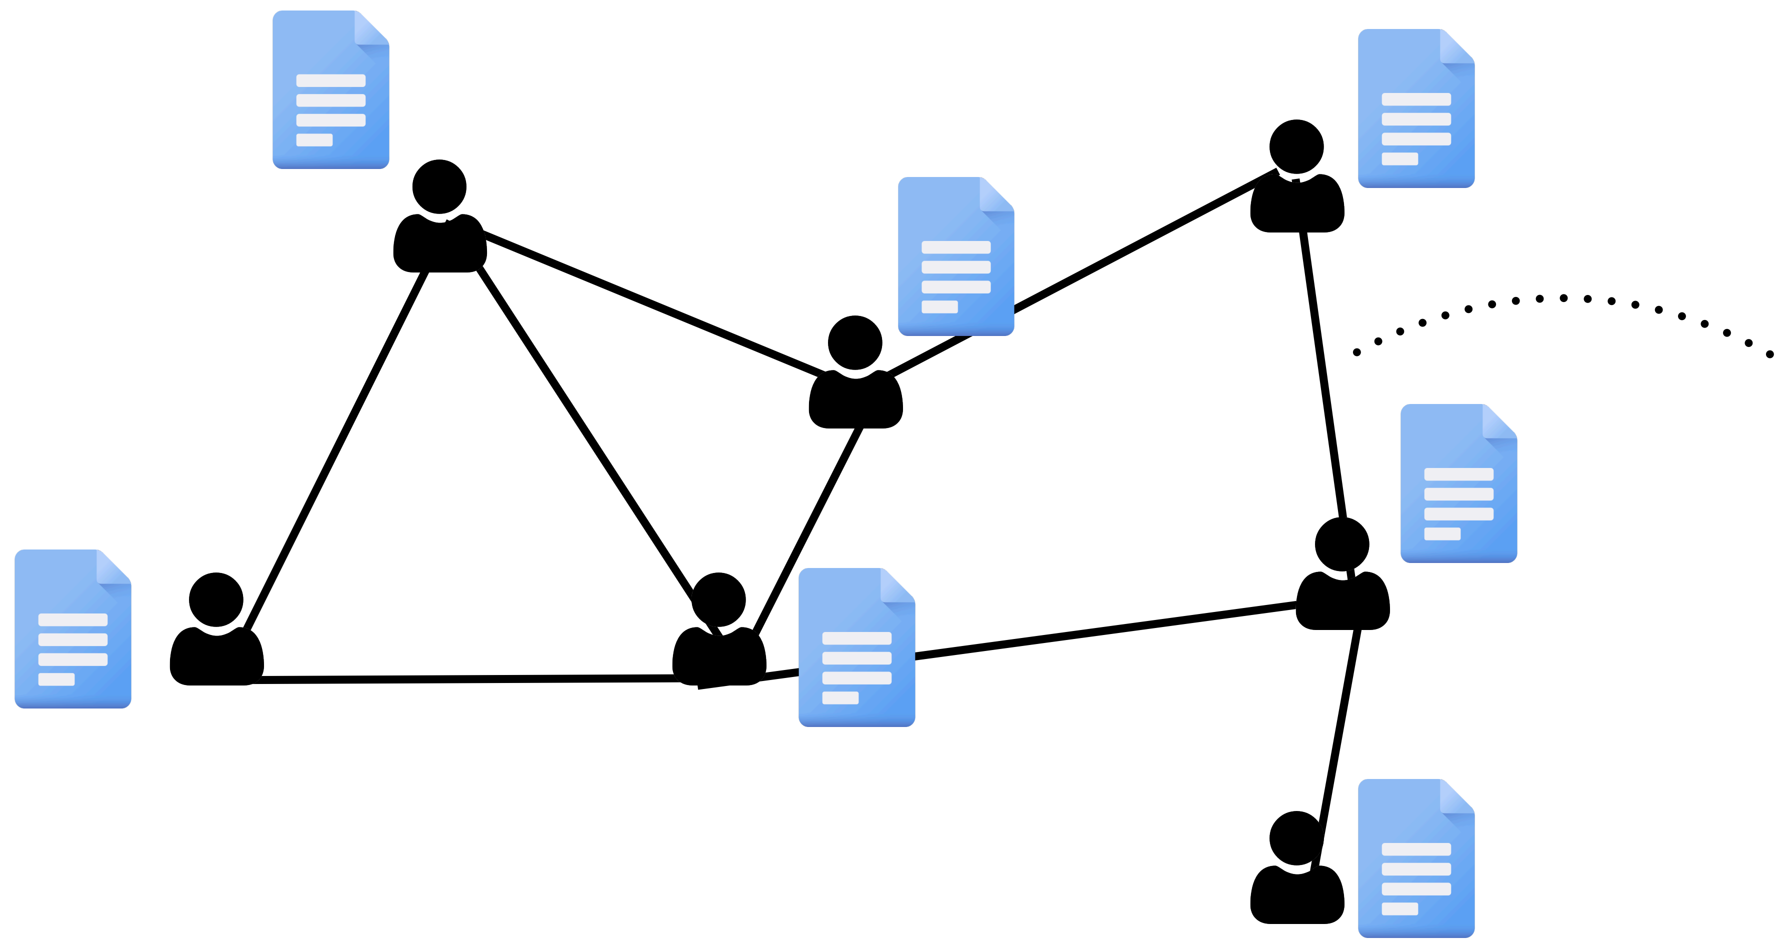

How to securely aggregate data?

Random

[network design]

- ☑ From distributed systems to random graphs
  The sparsity connectivity trade-off
- ☑ What are random K-out graphs?
  What makes them a useful design tool?
- ☐ Random K-out graphs in action
  What graph properties matter?
- ☐ Review: Notions of Connectivity
- ☐ Preview of contributions
  Key proof techniques

Balancing sparsity with connectivity in distributed systems

communication ↔
computation ↔

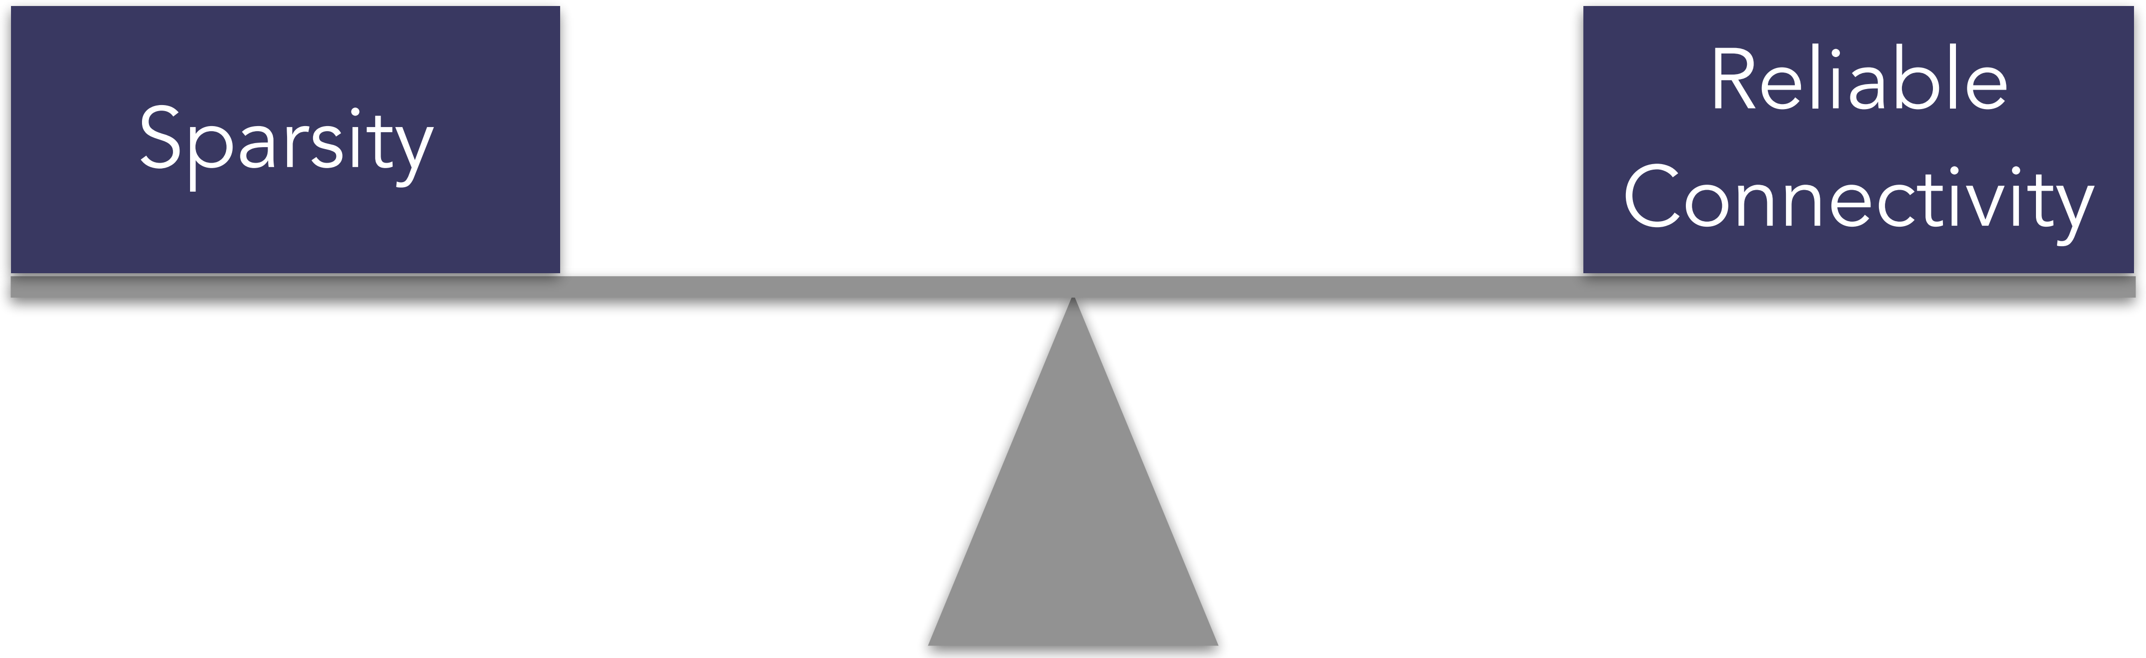

↔ privacy guarantees
↔ convergence

'Random K-out graphs'

well-connected yet sparse, allows distributed construction…

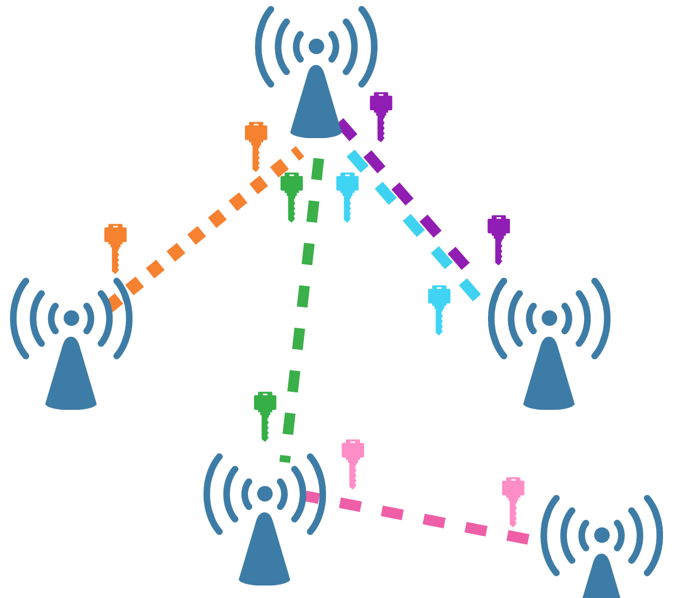

Key Predistribution in Wireless Sensor Networks

Chan et al., 2003

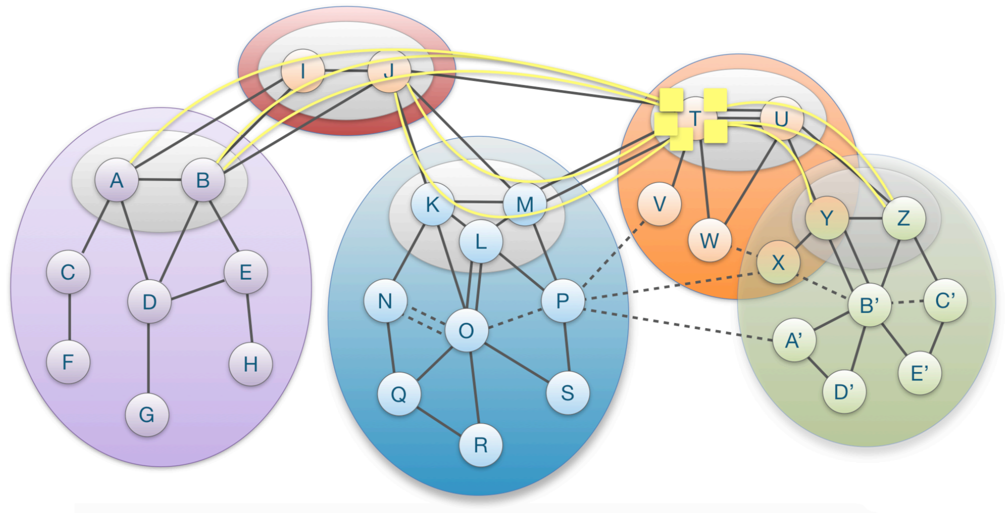

Next Generation Internet Architectures

Perrig et al., 2017

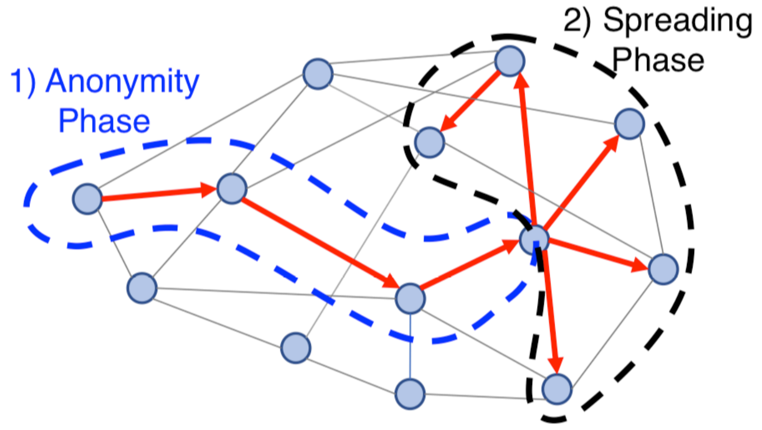

Cryptocurrency Networks

Bojja et al., 2017

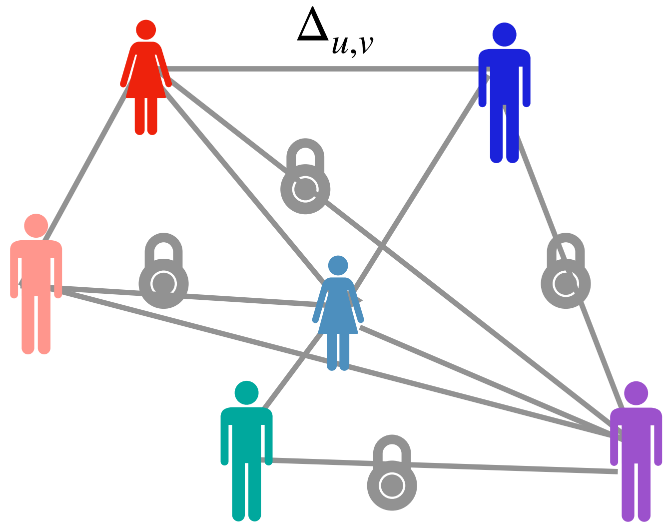

Distributed Differential Private Averaging

Sabater et al., 2022

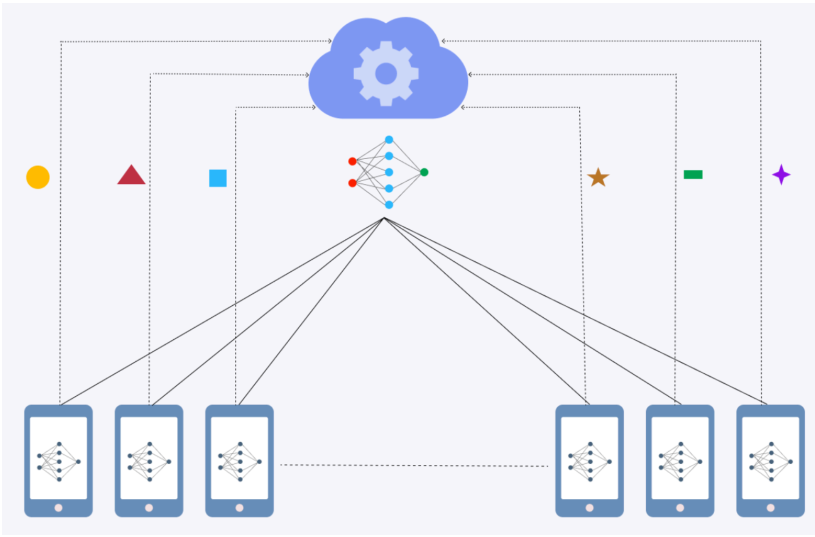

Secure Aggregation

Bonawitz et al.,
Bell et al., 2020

Balancing sparsity with connectivity in distributed systems

communication ↔
computation ↔

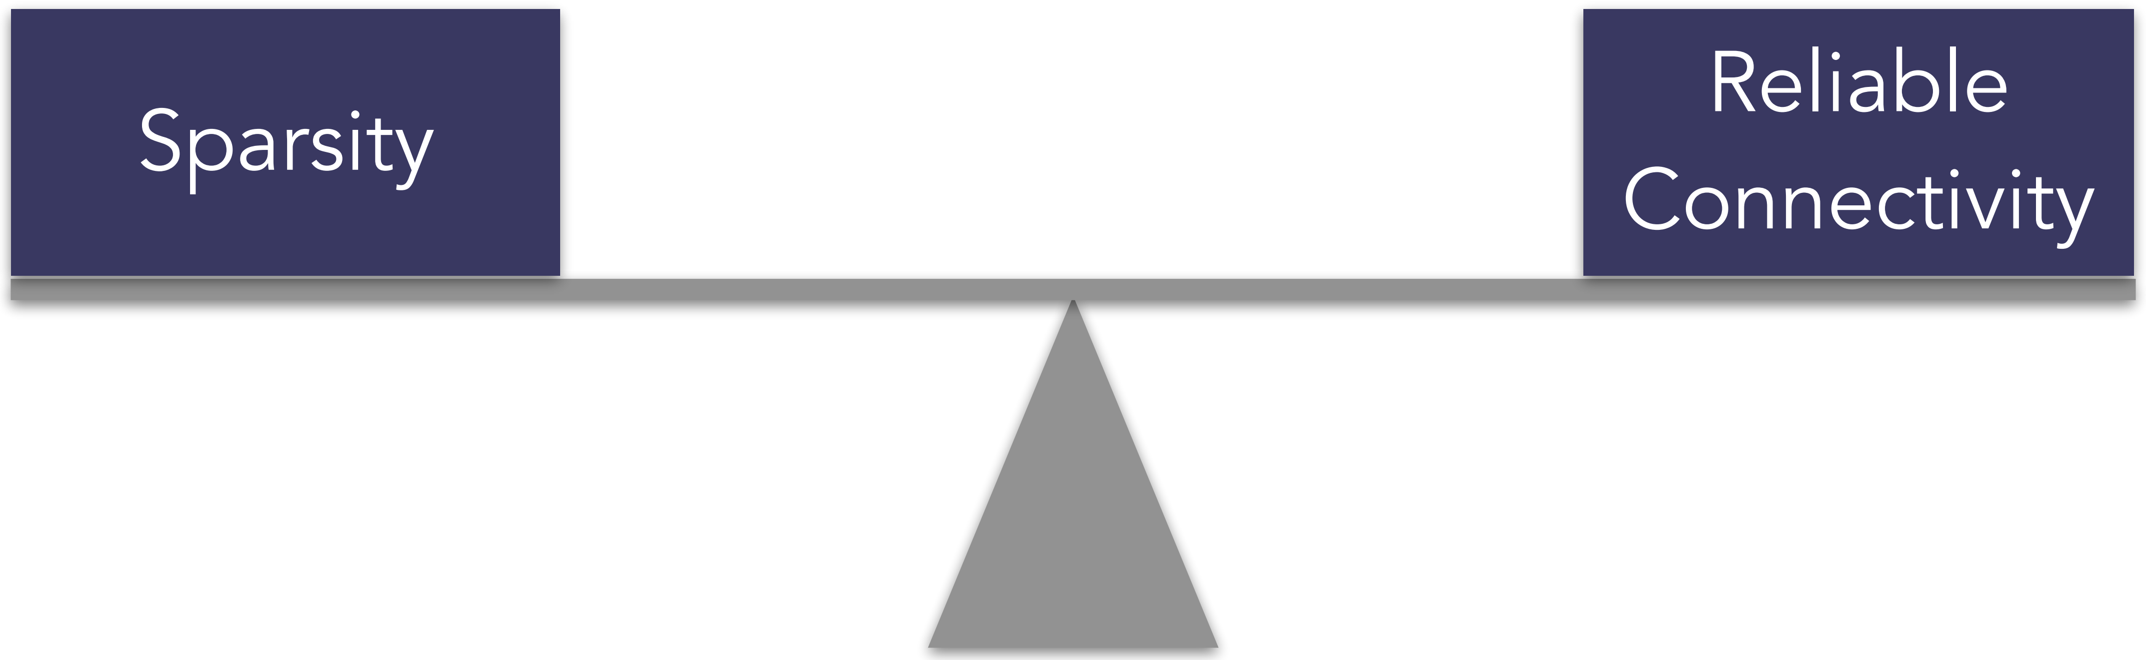

↔ privacy guarantees
↔ convergence

'Random K-out graphs'

well-connected yet sparse, allows distributed construction…

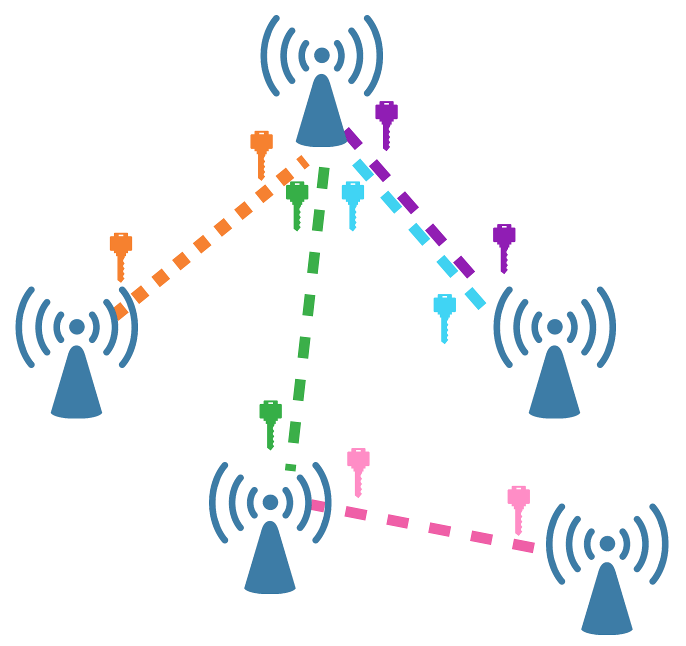

Key Predistribution in Wireless Sensor Networks

Chan et al., 2003

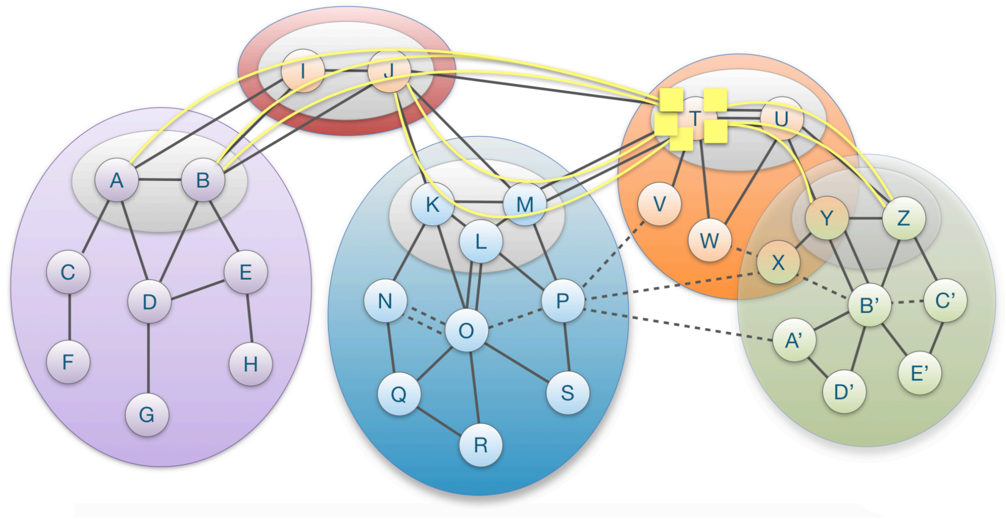

Next Generation Internet Architectures

Perrig et al., 2017

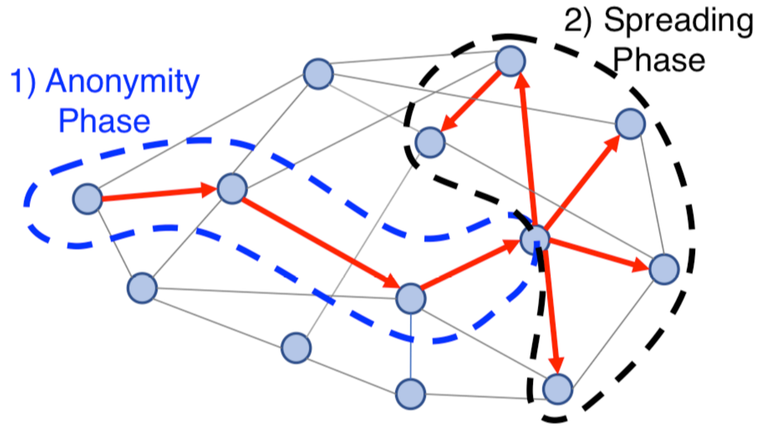

Cryptocurrency Networks

Bojja et al., 2017

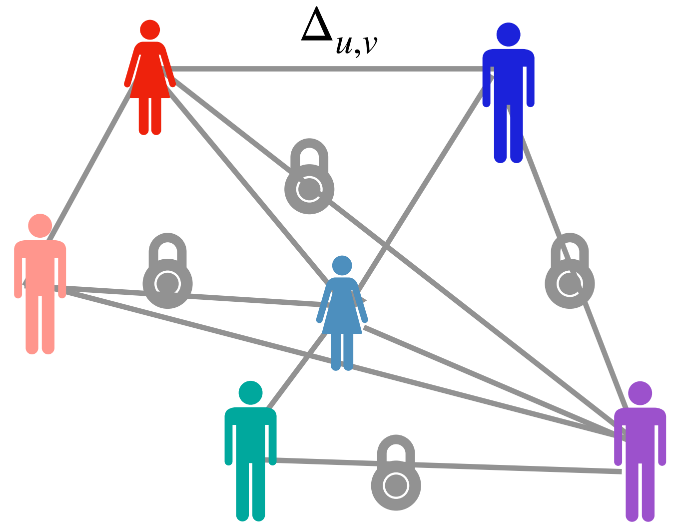

Distributed Differential Private Averaging

Sabater et al., 2022

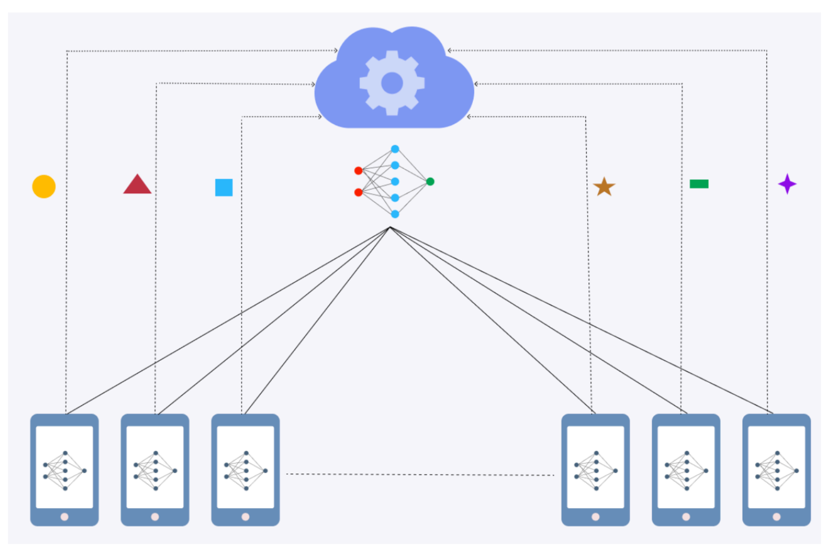

Secure Aggregation

Bonawitz et al.,

Balancing sparsity with connectivity in distributed systems

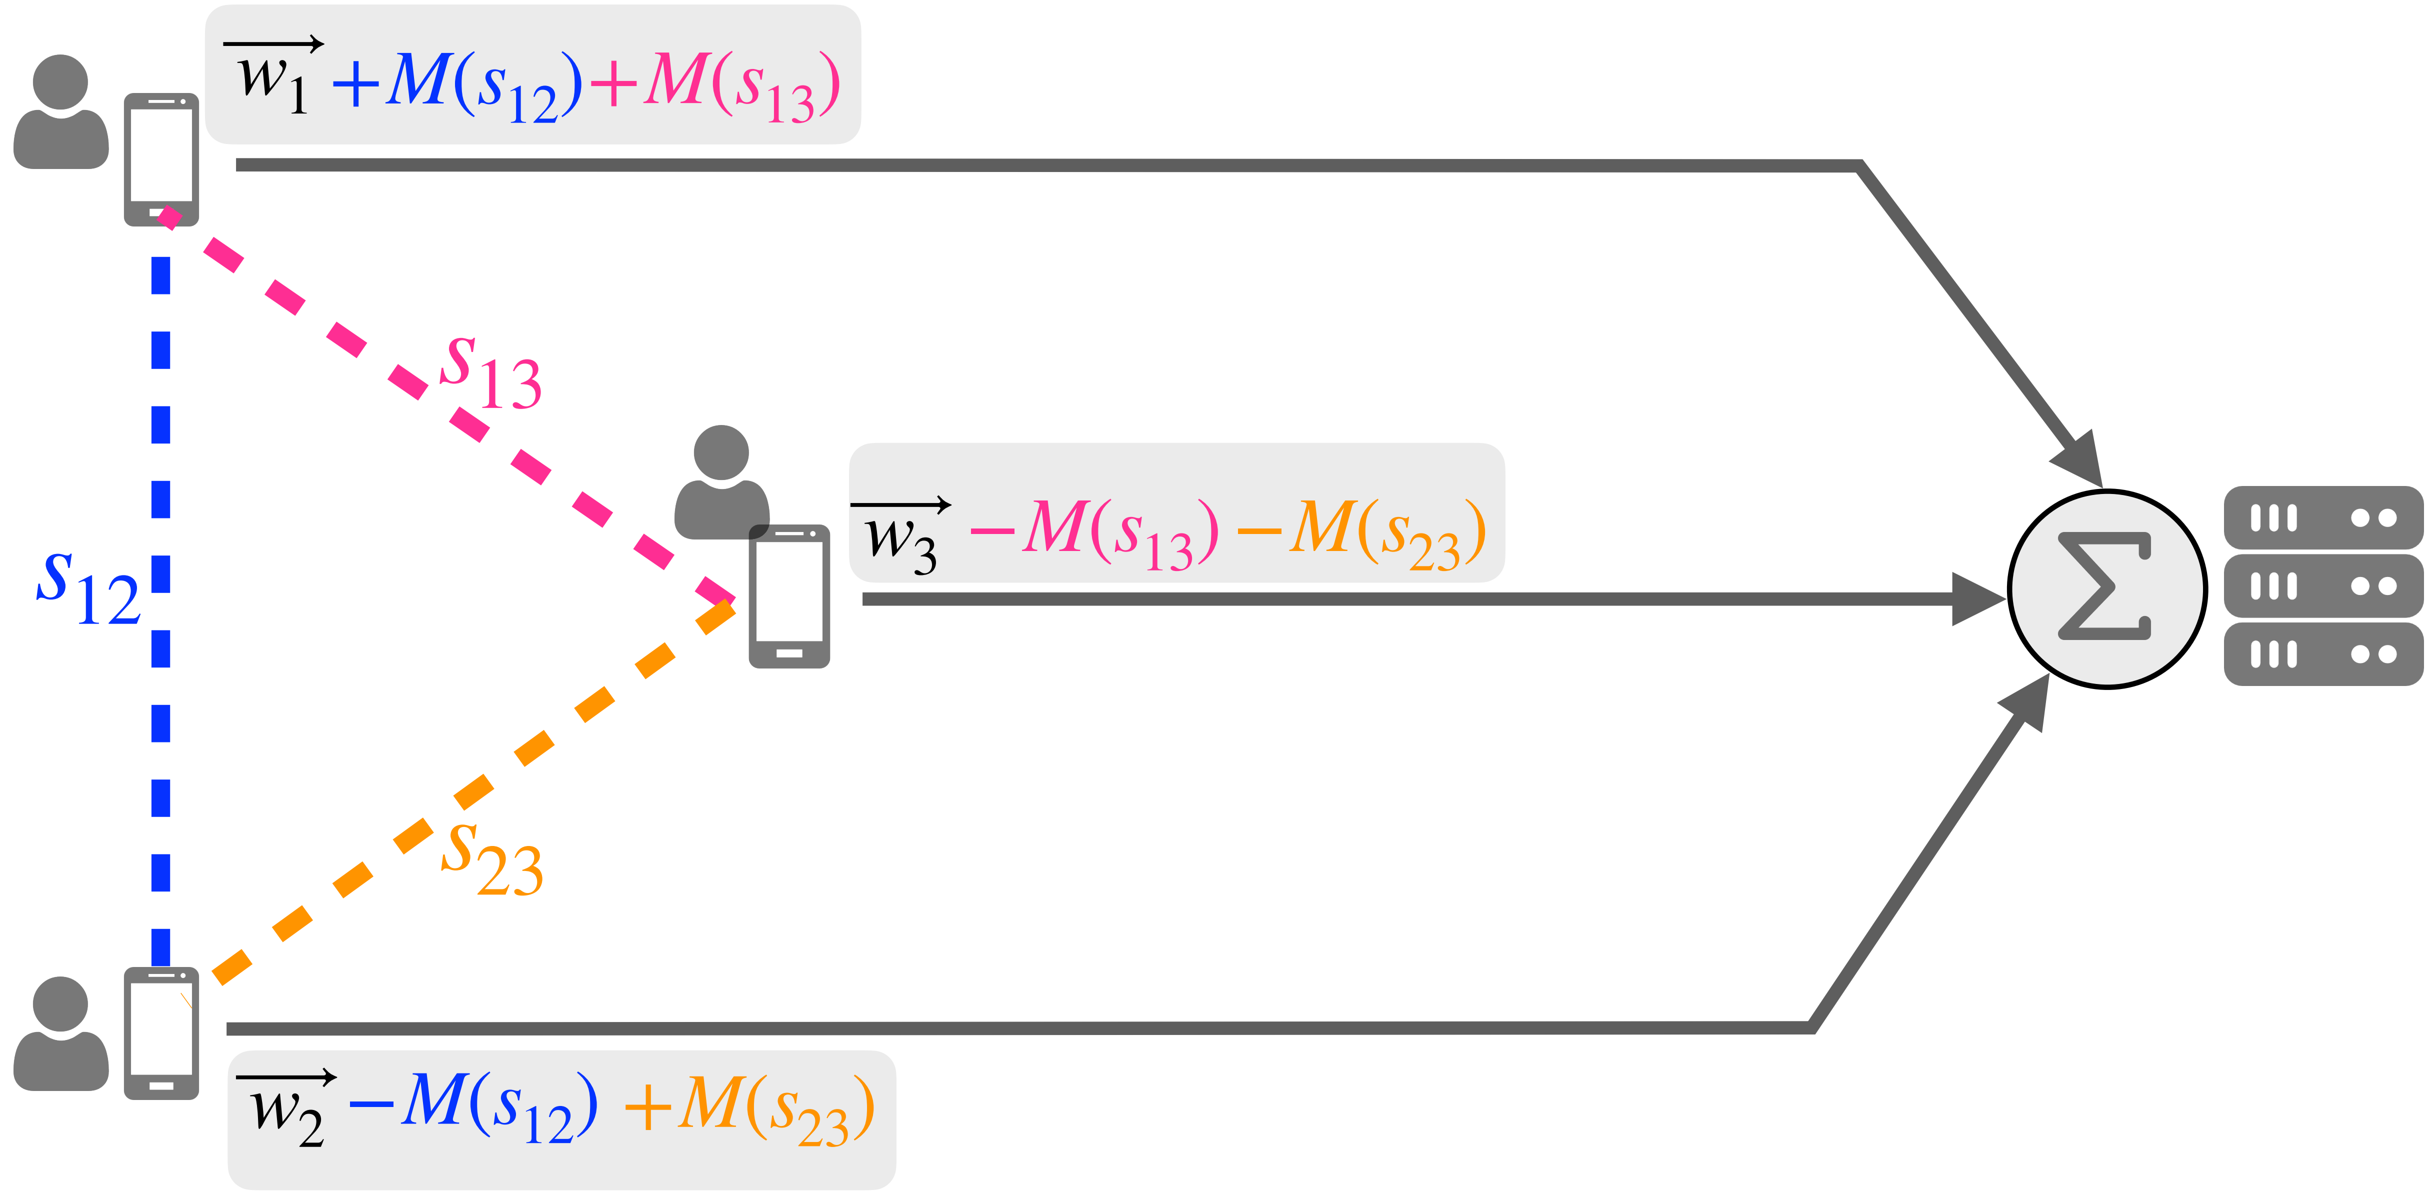

Flamingo: Multi-Round Single-Server Secure Aggregation with Applications to Private Federated Learning

Yiping Ma* Jess Woods* Sebastian Angel*† Antigoni Polychroniadou‡ Tal Rabin*
*University of Pennsylvania †Microsoft Research ‡J.P. Morgan AI Research & AlgoCRYPT CoE

Secure Single-Server Aggregation with (Poly)Logarithmic Overhead

James Bell, The Alan Turing Institute, London, UK, jbell@turing.ac.uk
K. A. Bonawitz, Google, New York, US, bonawitz@google.com
Adrià Gascón, Google, London, UK, adriag@google.com
Tancrède Lepoint, Google, New York, US, tancrede@google.com
Mariana Raykova, Google, New York, US, marianar@google.com

In the second step, the graph is established as follows. Each client $i \in [n]$ randomly chooses γ other clients in $[n]$ as its neighbors, and tells the server about their choices. After the server collects all the clients' choices, it notifies each client of their neighbors indexes in $[n]$ and public keys. The neighbors of client i, denoted as $A(i)$, are those corresponding to vertices that have an edge with i (i.e., i chose them or they chose i).

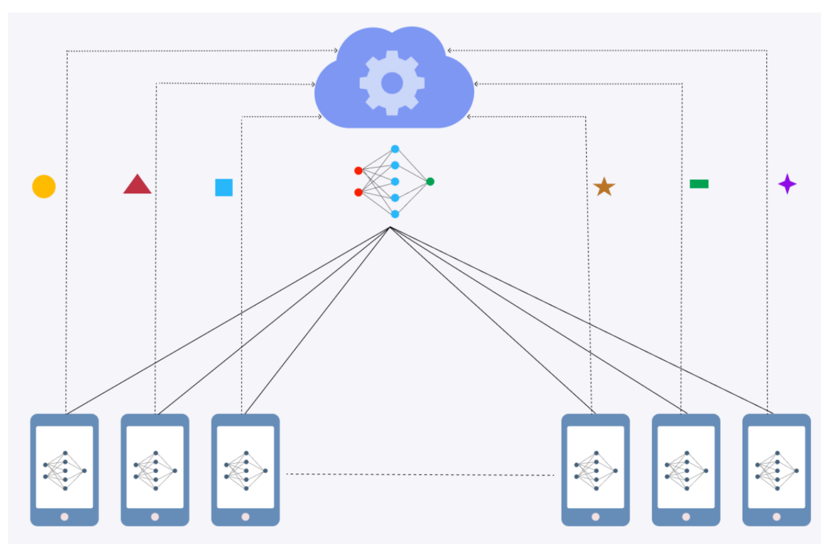

Secure Aggregation

Bell et al., 2020

Balancing sparsity with connectivity in distributed systems

communication ↔
computation ↔

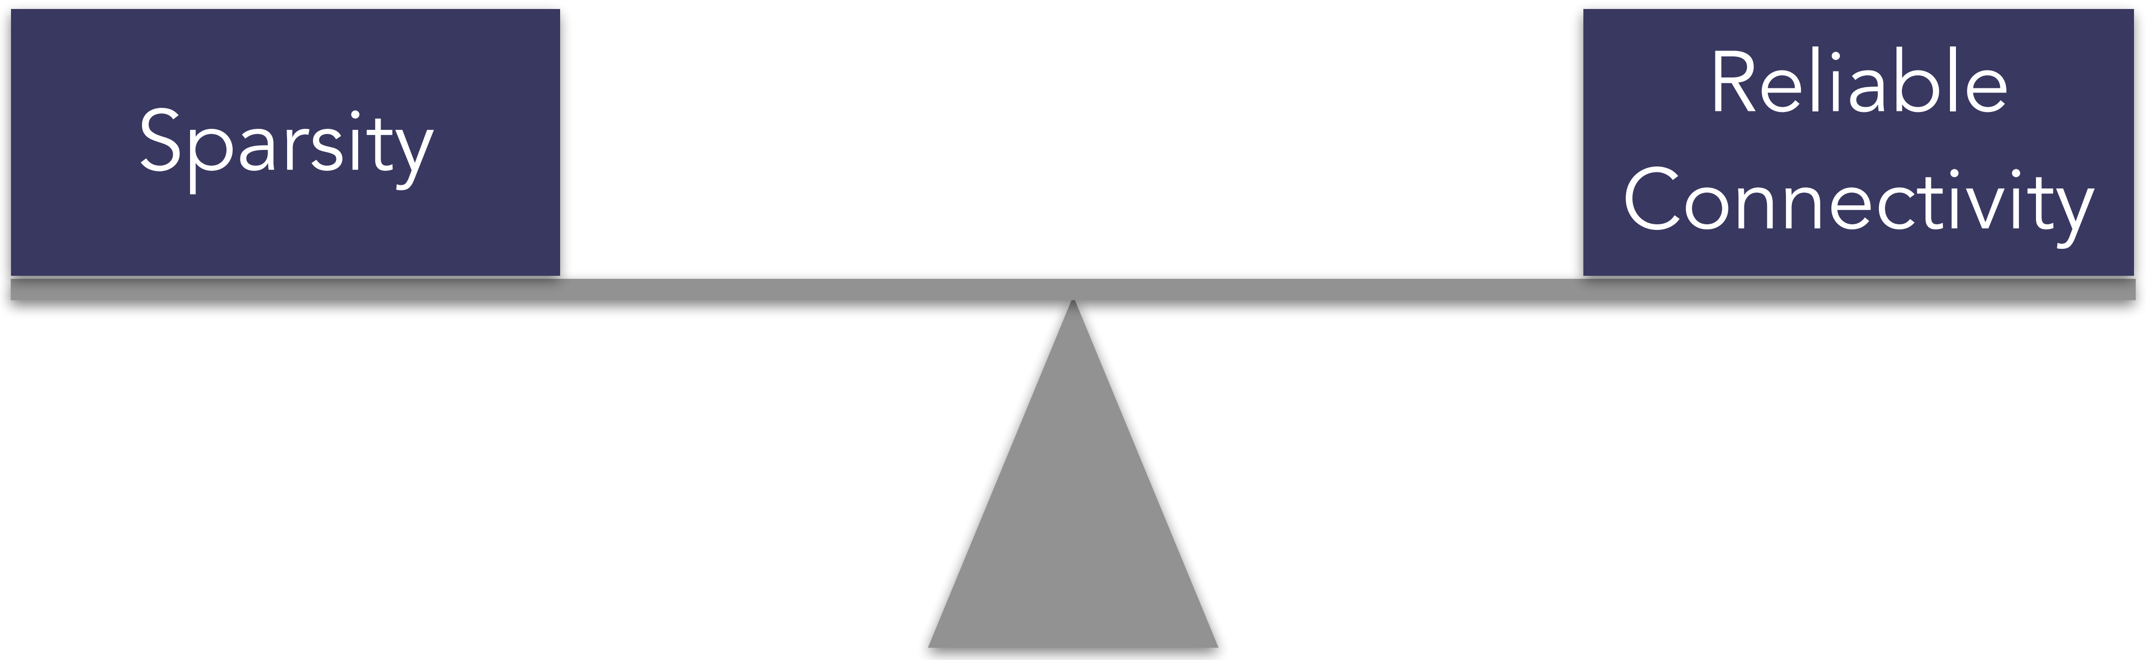

↔ privacy guarantees
↔ convergence

'Random K-out graphs'

well-connected yet sparse, allows distributed construction…

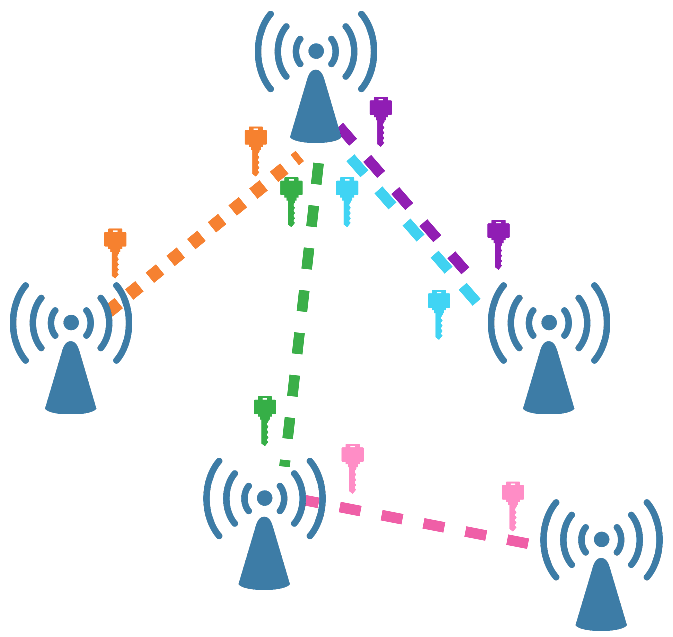

Key Predistribution in Wireless Sensor Networks

Chan et al., 2003

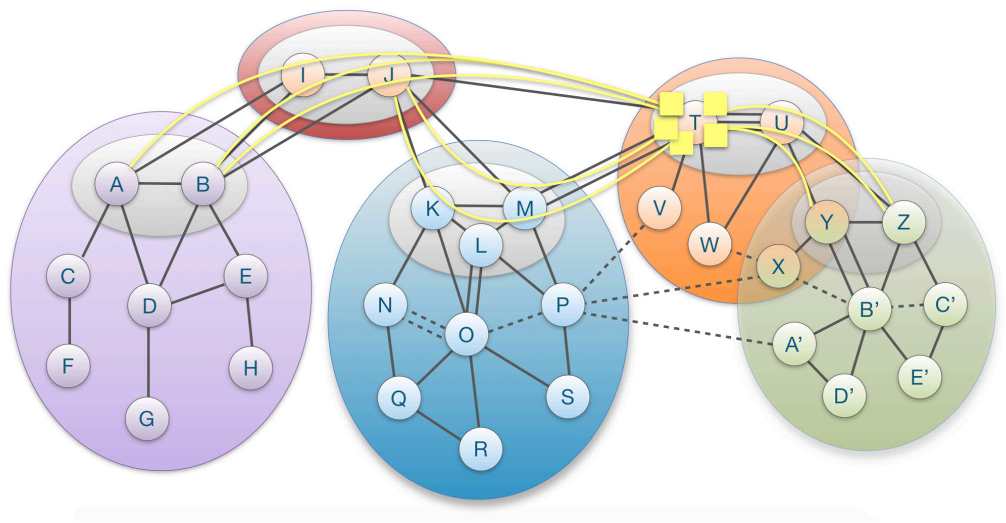

Next Generation Internet Architectures

Perrig et al., 2017

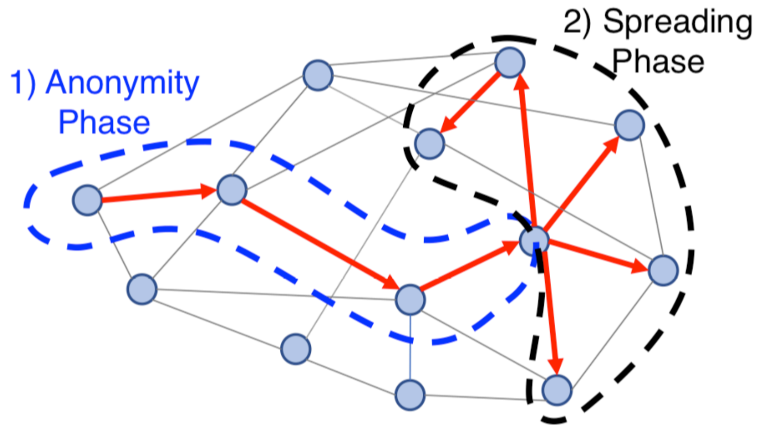

Cryptocurrency Networks

Bojja et al., 2017

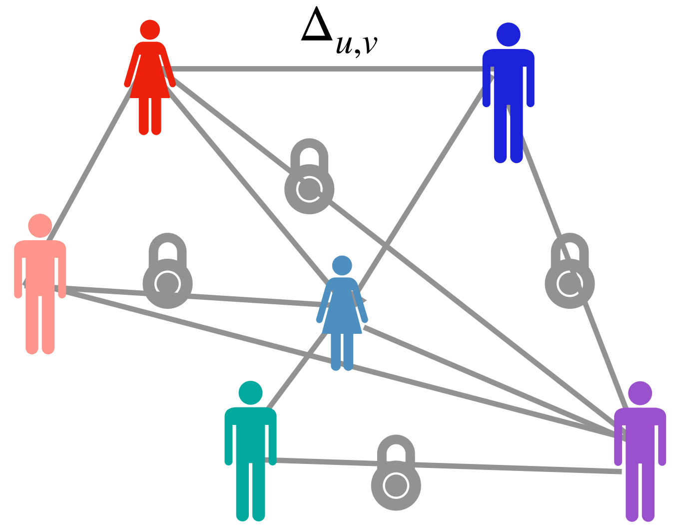

Distributed Differential Private Averaging

Sabater et al., 2022

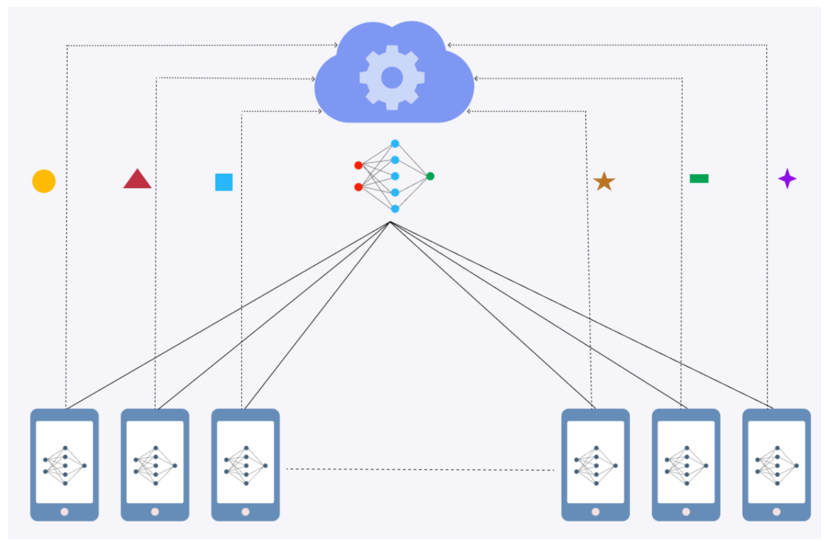

Secure Aggregation

Bonawitz et al., Bell et al., 2020

Securely aggregating information over a network

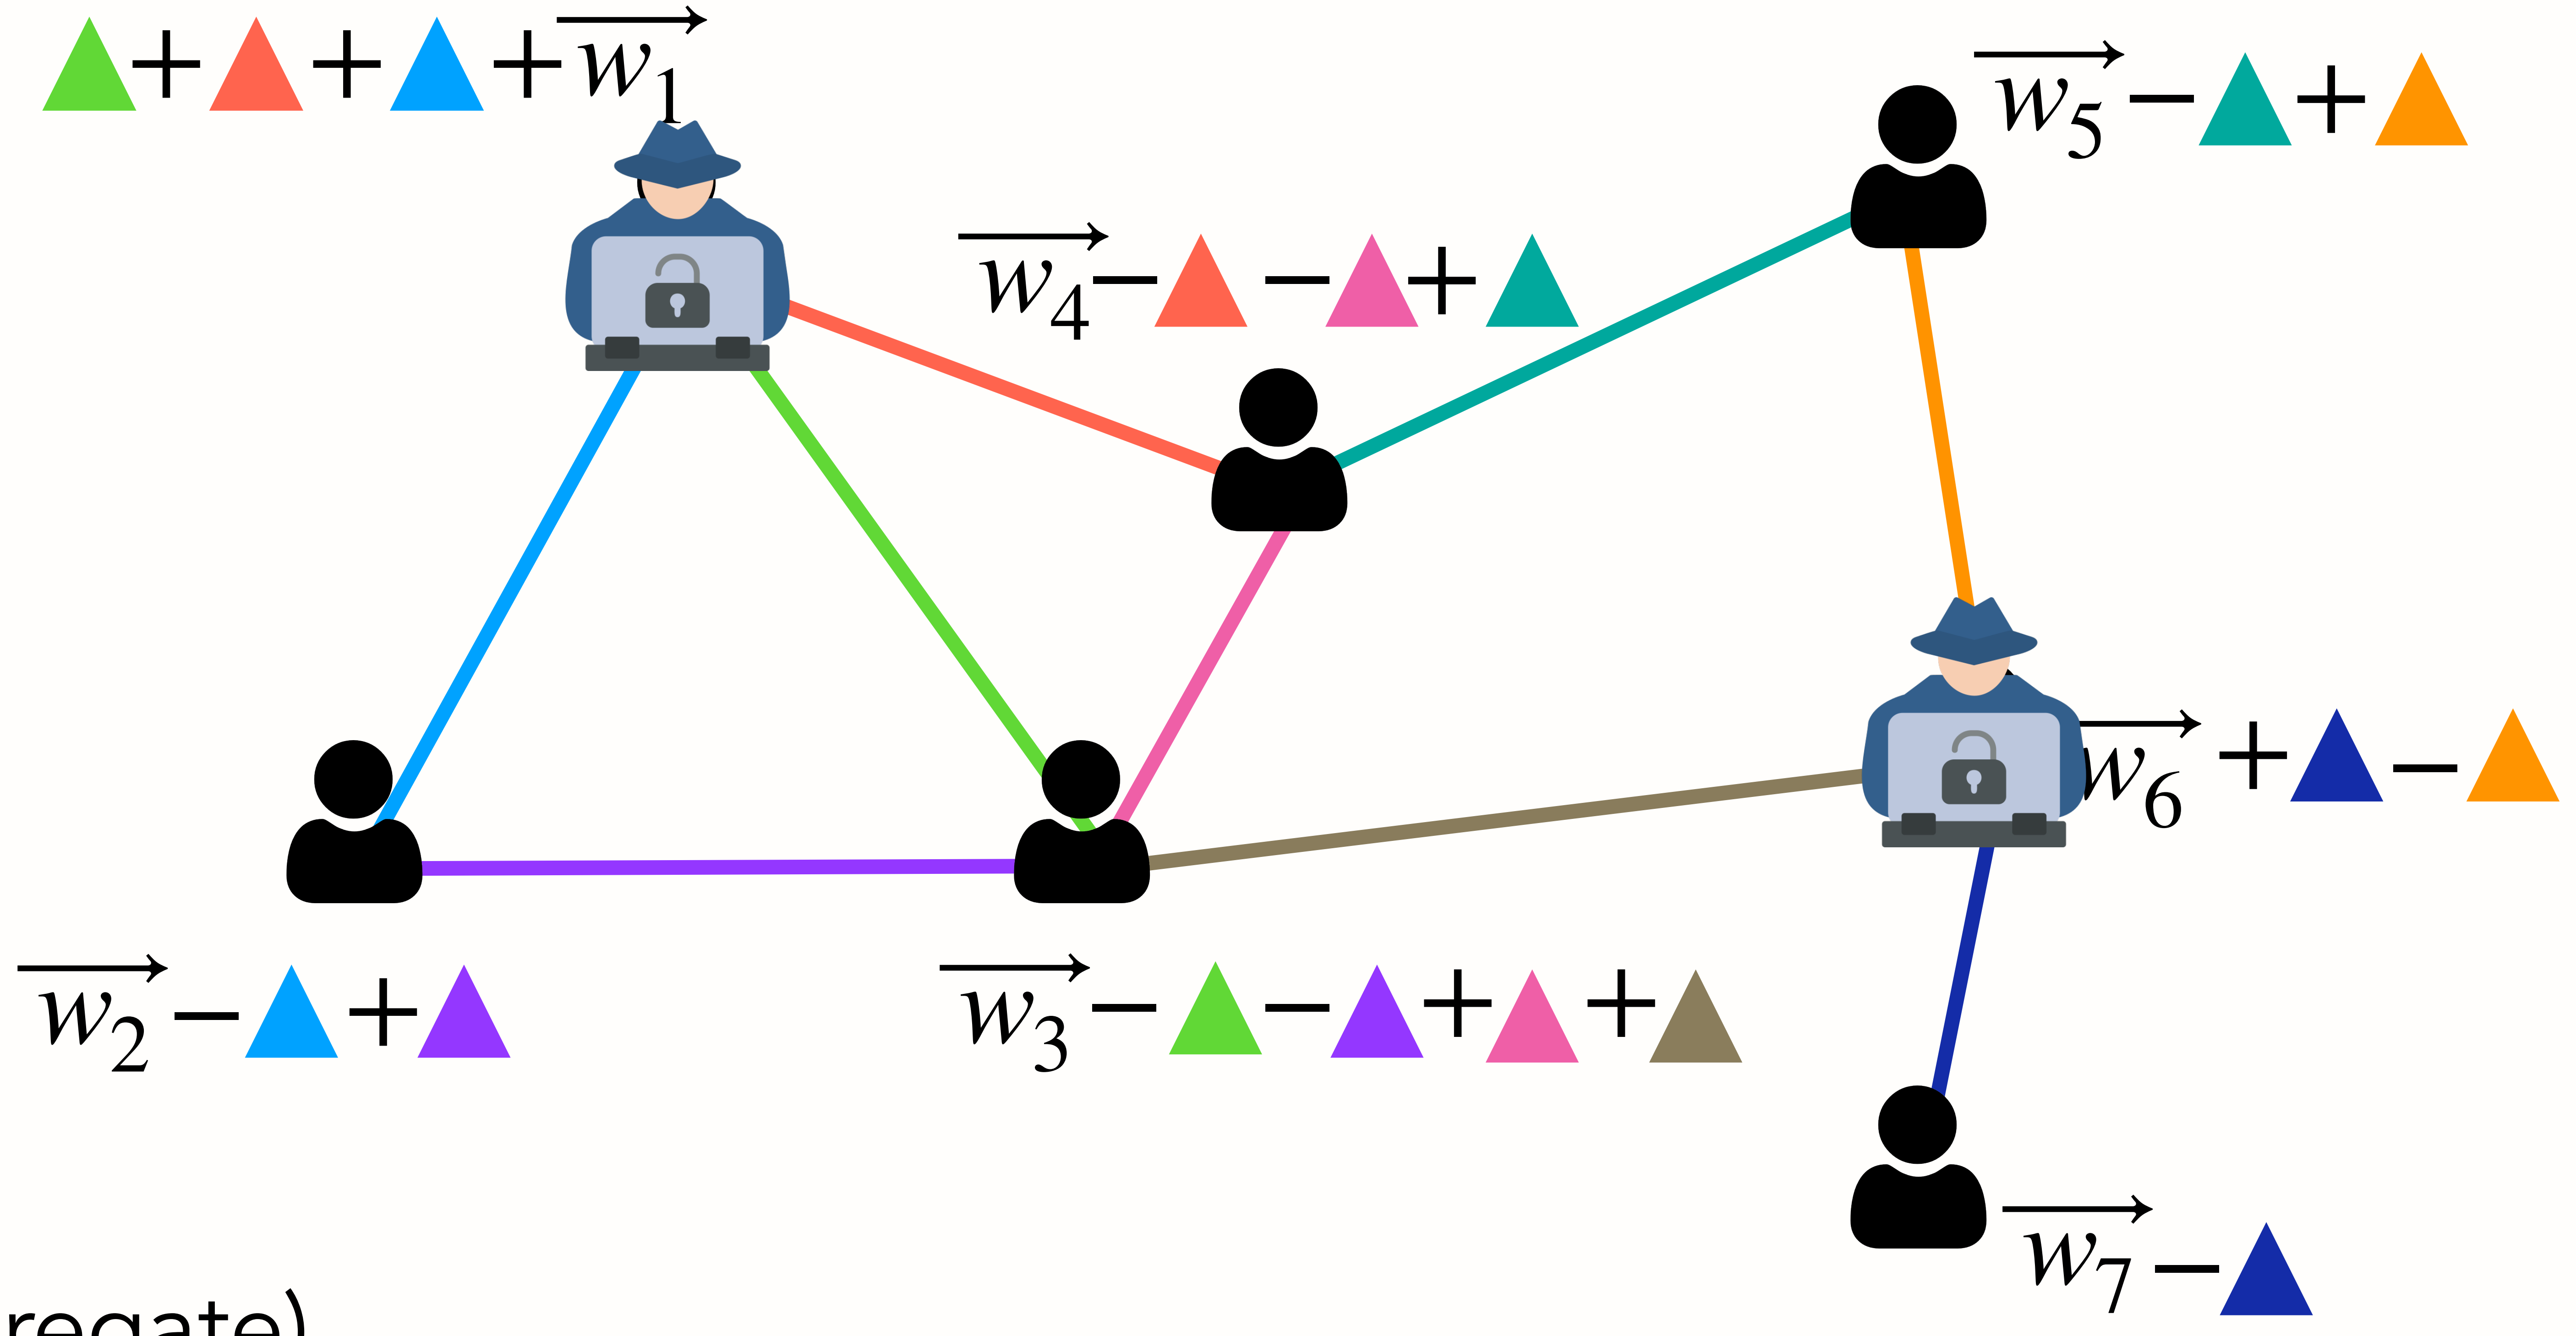

Goal:

Learn $\Sigma_i w_i$

without revealing w_i

Approach:

add pairwise masks (that cancel in aggregate)

Challenge:

What if some nodes collude, i.e.,

Share masks for their neighbors?

Securely aggregating information over a network

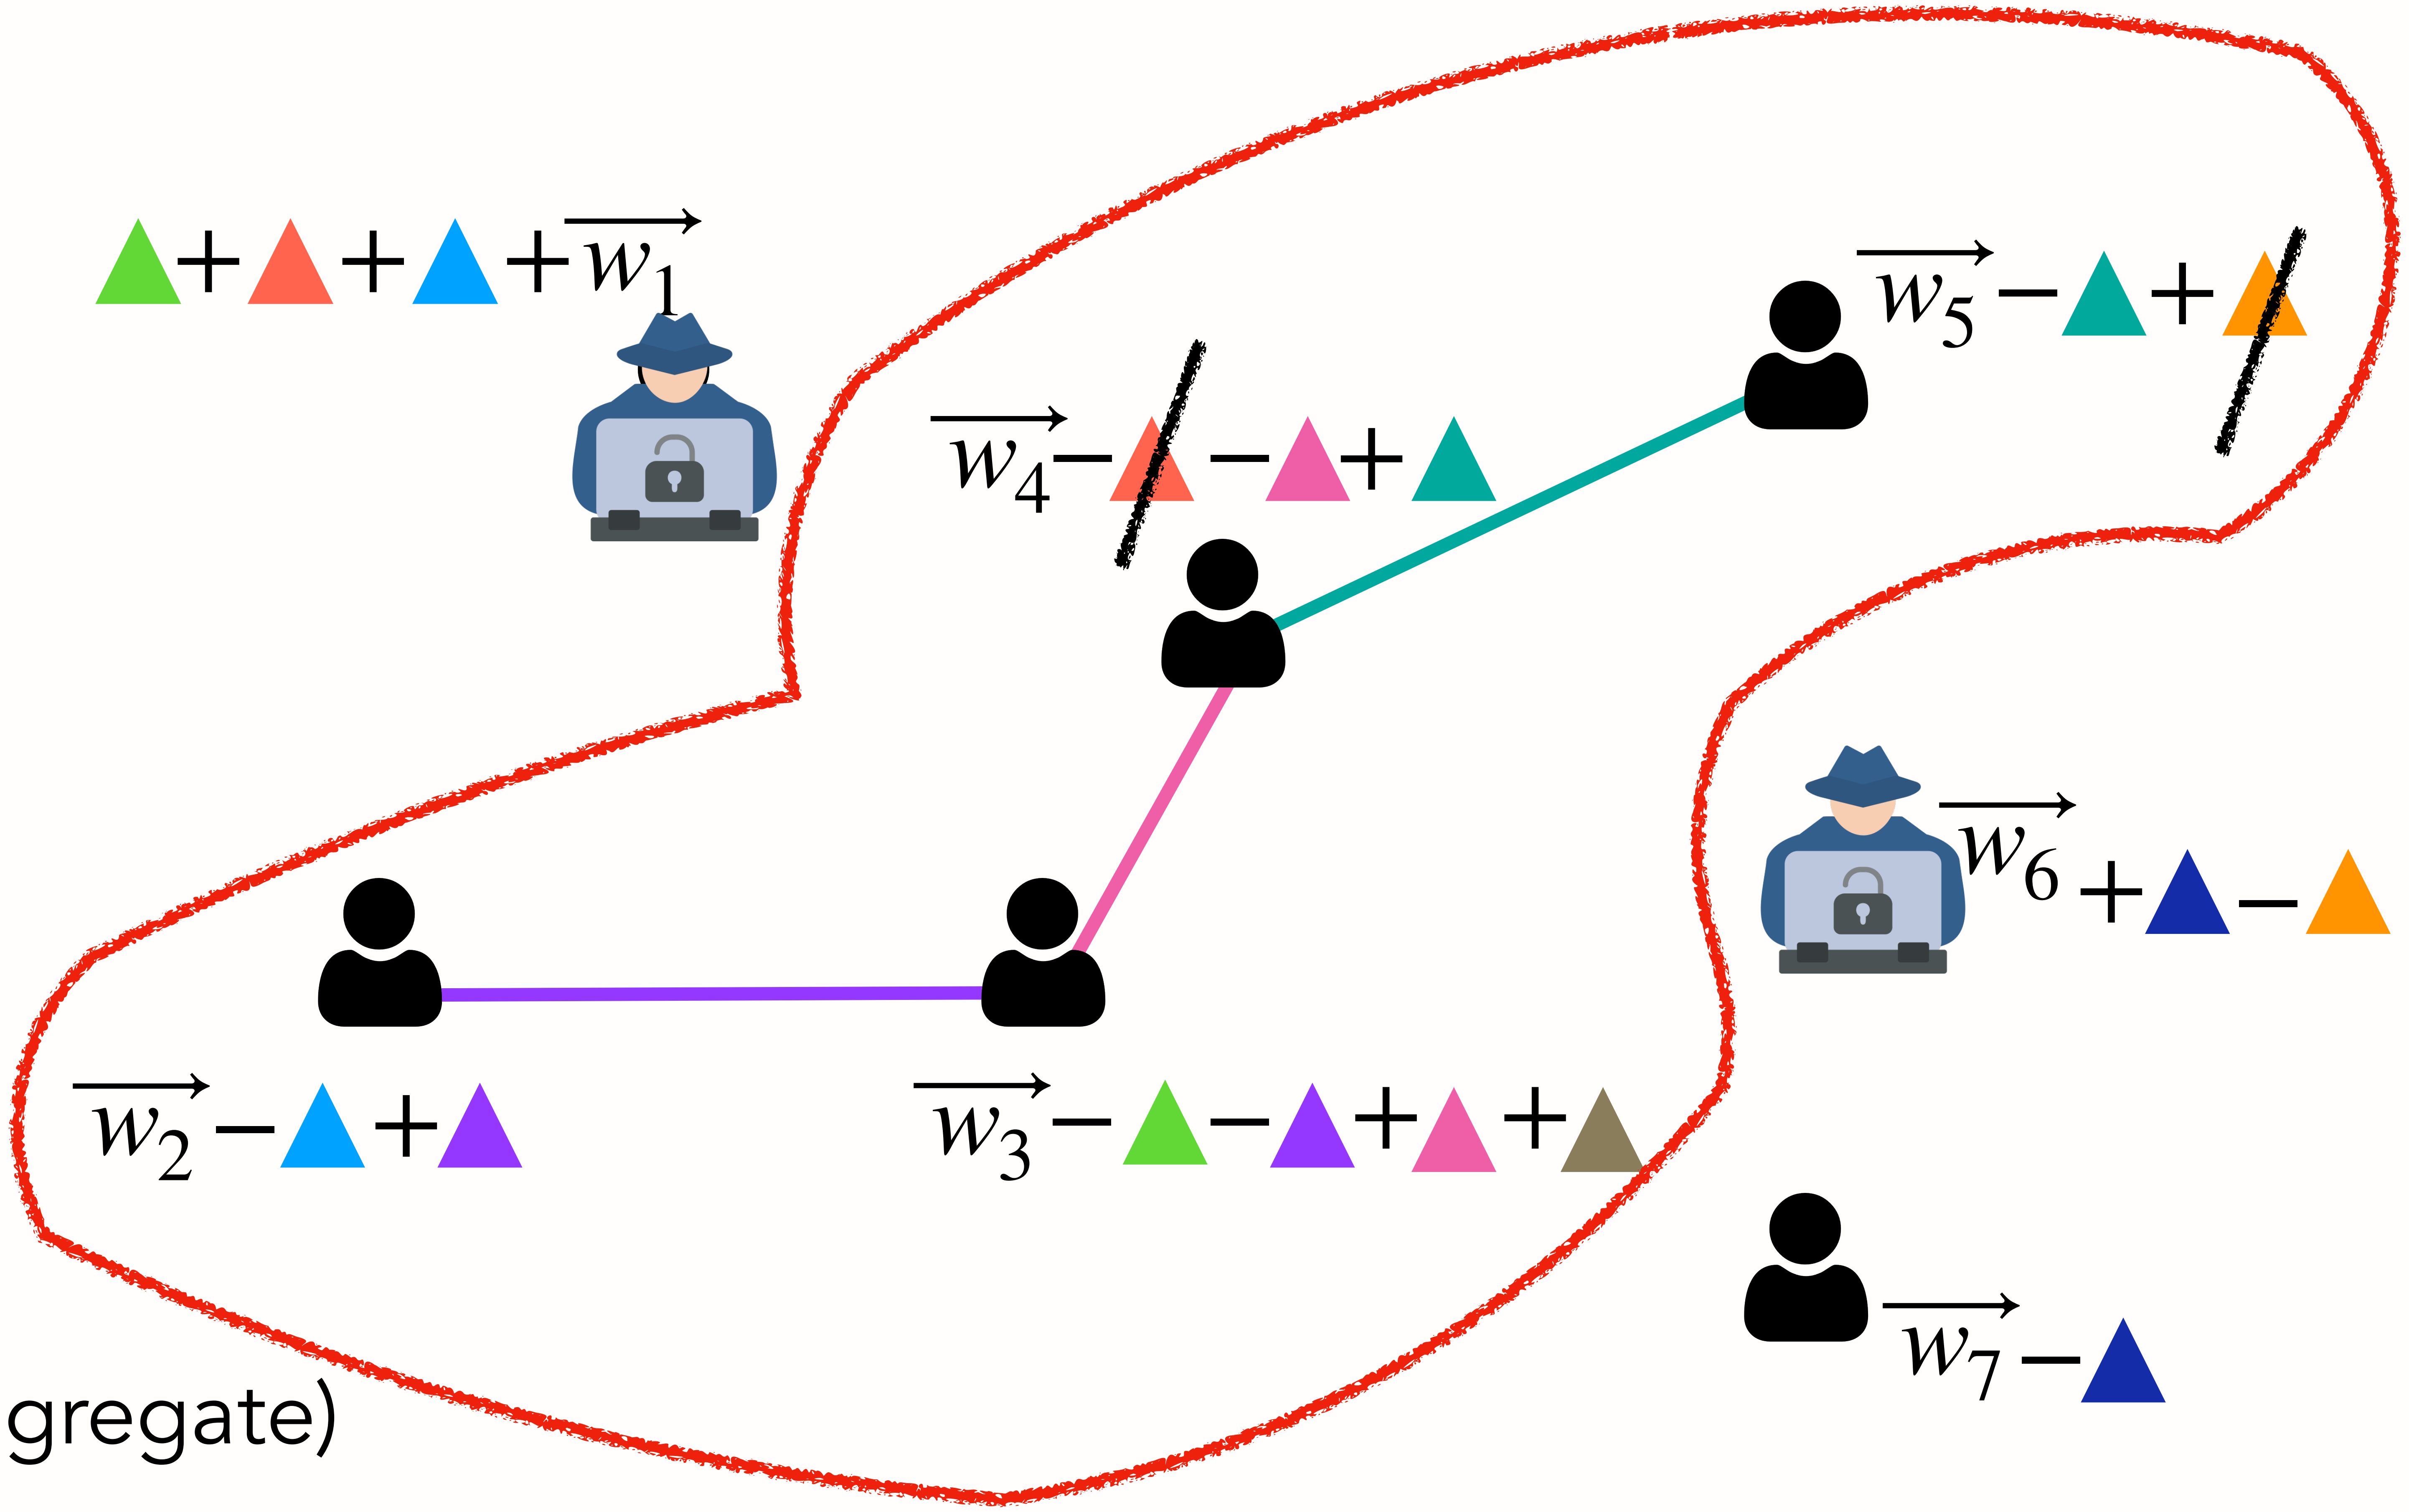

Goal:

Learn $\Sigma_i w_i$

without revealing w_i

Approach:

add pairwise masks (that cancel in aggregate)

Challenge:

What if there are multiple corrupt nodes? → Model corruptions as deletions

How to set K for for reliable connectivity after node deletions?

$\mathbb{H}(n, K)$

Secure information aggregation → Graph properties?

Connectivity after node failures/corruptions

What if there are multiple corrupt nodes? → Model corruptions as deletions

How to set K for reliable connectivity after node deletions?

Select K so that connectivity persists after node deletions

(if connectivity is not feasible, quantify size of connected subgraphs after deletions)

targeted subset of nodes

random subset of nodes

Part 1: Road-map

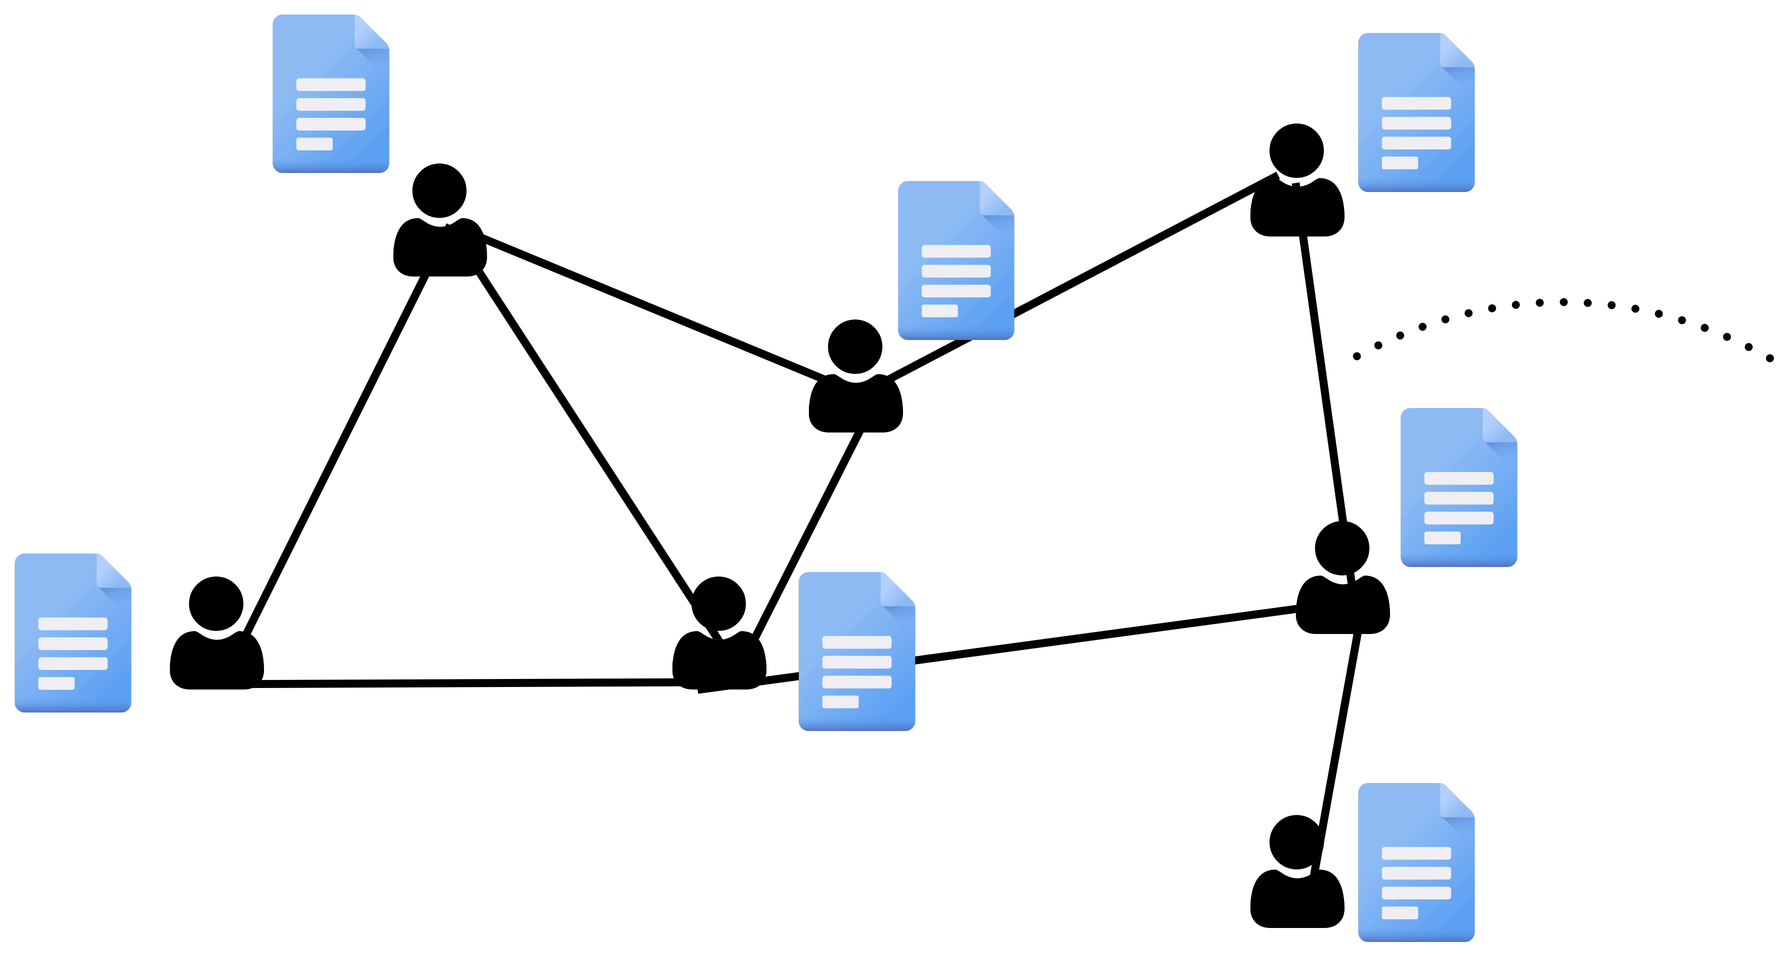

How to securely aggregate data?

Random
[network design]

- [x] From distributed systems to random graphs
  The sparsity connectivity trade-off
- [x] What are random K-out graphs?
  What makes them a useful design tool?
- [x] Random K-out graphs in action
  What graph properties matter?
- [ ] Review: Notions of Connectivity
- [ ] Preview of contributions
  Key proof techniques

How to quantify strength of connectivity?

Increasing *strength* of connectivity

a 'giant' component

$|C_{max}| = \Omega(n)$

1-connectivity

$|C_{max}| = n$

***k*-connectivity**

ensures connectivity despite <u>any</u> k-1 node failures*

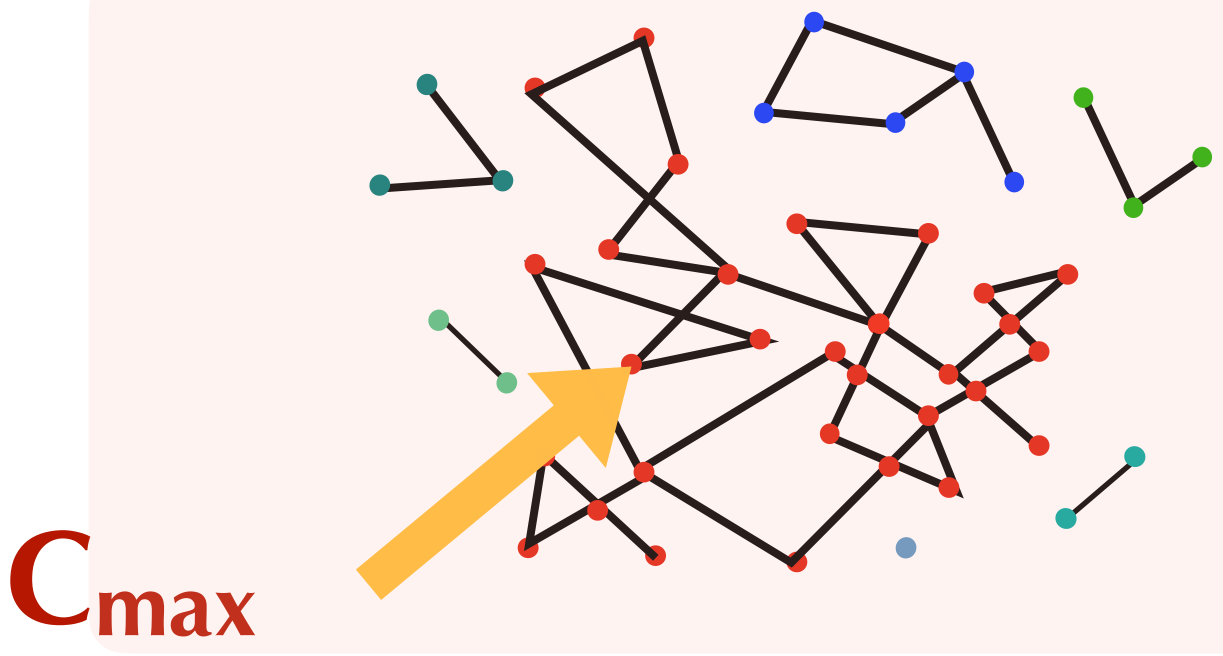

+ whether these properties are retained upon component failures

Part 1: Road-map

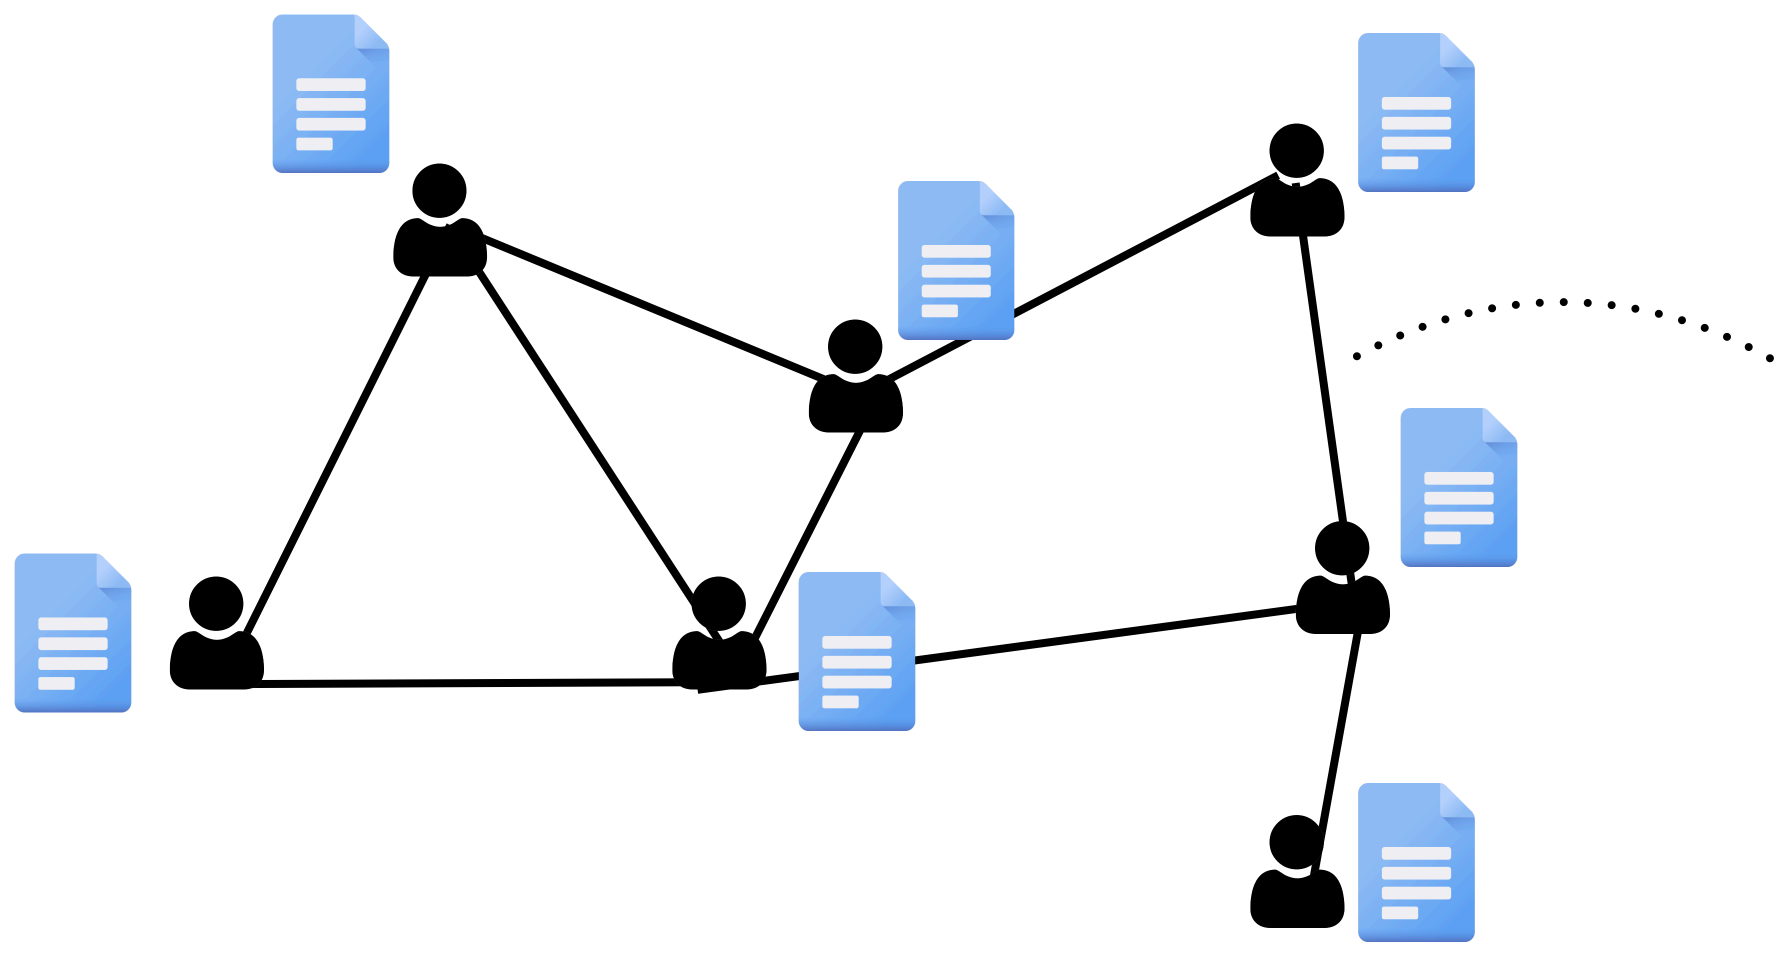

How to securely aggregate data?

Random

[network design]

- ✓ From distributed systems to random graphs
  - The sparsity connectivity trade-off
- ✓ What are random K-out graphs?
  - What makes them a useful design tool?
- ✓ Random K-out graphs in action
  - What graph properties matter?
- ✓ Review: Notions of Connectivity
- → Preview of contributions
  - Key proof techniques

Performance metrics

Performance metrics

Increasing *strength* of connectivity (results whp)

a 'giant' component

1-connectivity

k-connectivity
($k \geq 2$)

Performance metrics

Increasing *strength* of connectivity (results whp)

Weaker

Stronger

a 'giant' component

1-connectivity

***k*-connectivity**
($k \geq 2$)

Inhomogeneous Random K-out Graphs

(A node makes
1 selection wp μ,
else $K_n \geq 2$ selections)

Homogeneous Random K-out Graphs

(All nodes make
K_n selections)

Performance metrics

Increasing *strength* of connectivity (results whp)

Weaker → Stronger

	a 'giant' component	**1-connectivity**	***k*-connectivity ($k \geq 2$)**
Inhomogeneous Random K-out Graphs (A node makes 1 selection wp μ, else $K_n \geq 2$ selections)			
Homogeneous Random K-out Graphs (All nodes make K_n selections)		[Fenner & Frieze '82] requires $K_n \geq 2$	[Fenner & Frieze '82] requires $K_n \geq 2k$

Performance metrics

Increasing *strength* of connectivity (results whp)

Weaker → Stronger

	a 'giant' component	1-connectivity	k-connectivity ($k \geq 2$)
Inhomogeneous Random K-out Graphs (A node makes 1 selection wp μ, else $K_n \geq 2$ selections)		[Eletreby & Yagan, '19] requires K_n unbounded (no matter how slowly growing)	
Homogeneous Random K-out Graphs (All nodes make K_n selections)		[Fenner & Frieze '82] requires $K_n \geq 2$	[Fenner & Frieze '82] requires $K_n \geq 2k$

Performance metrics

Increasing *strength* of connectivity (results whp)

Weaker → Stronger

	a 'giant' component	1-connectivity	k-connectivity ($k \geq 2$)
Inhomogeneous Random K-out Graphs (A node makes 1 selection wp μ, else $K_n \geq 2$ selections)	?	[Eletreby & Yagan, '19] requires K_n unbounded (no matter how slowly growing)	?
Homogeneous Random K-out Graphs (All nodes make K_n selections)	? what if a random subset of nodes fail?	[Fenner & Frieze '82] requires $K_n \geq 2$ Tighter bounds? ?	[Fenner & Frieze '82] requires $K_n \geq 2k$

Performance metrics

Increasing *strength* of connectivity (results whp)

Weaker → Stronger

	a 'giant' component	**1-connectivity**	***k*-connectivity ($k \geq 2$)**
Inhomogeneous Random K-out Graphs (A node makes 1 selection wp μ, else $K_n \geq 2$ selections)	?	[Eletreby & Yagan, '19] requires K_n unbounded (no matter how slowly growing)	?
Homogeneous Random K-out Graphs	?	[Fenner & Frieze '82] requires $K_n \geq 2$ Tighter bounds? ?	[Fenner & Frieze '82] requires $K_n \geq 2k$

Note: Rich literature on variants of K-out graphs/distributed constructions and related properties

[Frieze, Karonski '24], [Pirani et al., 23], [Penrose '18], [Frieze, Johannson '15], [Cooper, Frieze '88] and more

Preview of contributions (informal)

Increasing *strength* of connectivity (results whp)

Weaker → Stronger

	a 'giant' component	1-connectivity	k-connectivity ($k \geq 2$)
Inhomogeneous Random K-out Graphs (A node makes 1 selection wp μ, else $K_n \geq 2$ selections)	$\lvert C_{\max} \rvert = n - O(1)$ whp $\lvert C_{\max} \rvert = n(1 - o(1))$ after $o(n)$ nodes fail IEEE Trans. on Info. Theory '23	[Eletreby & Yagan, '19] requires K_n unbounded (no matter how slowly growing)	requires K_n of order $\log(n)$ IEEE Trans. on Info. Theory '21
Homogeneous Random K-out Graphs (All nodes make K_n selections)	Provide K_n required for a given $\lvert C_{\max} \rvert$ whp given the size of (random) node failures IEEE ISIT '21	[Fenner & Frieze '82] requires $K_n \geq 2$ $p_{\text{con}} = 1 - \Theta(1/n^{K^2-1})$ even when $o(n)$ random node failures, $p_{\text{con}} \to 1$ IEEE ICC'21 (Best Paper)	[Fenner & Frieze '82] requires $K_n \geq 2k$

Preview of contributions (informal)

Increasing *strength* of connectivity (results whp) →

Weaker → Stronger

	a 'giant' component	**1-connectivity**	**k-connectivity ($k \geq 2$)**
Inhomogeneous Random K-out Graphs (A node makes 1 selection wp μ, else $K_n \geq 2$ selections)	$\lvert C_{\max} \rvert = n - O(1)$ whp $\lvert C_{\max} \rvert = n(1 - o(1))$ after $o(n)$ nodes fail IEEE Trans. on Info. Theory '23	[Eletreby & Yagan, '19] requires K_n unbounded (no matter how slowly growing)	requires K_n of order $\log(n)$ IEEE Trans. on Info. Theory '21
Homogeneous Random K-out Graphs (All nodes make K_n selections)	Provide K_n required for a given $\lvert C_{\max} \rvert$ whp given the size of (random) node failures IEEE ISIT '21	[Fenner & Frieze '82] requires $K_n \geq 2$ $p_{\text{con}} = 1 - \Theta(1/n^{K^2-1})$ even when $o(n)$ random node failures, $p_{\text{con}} \to 1$ IEEE ICC'21 (Best Paper)	[Fenner & Frieze '82] requires $K_n \geq 2k$

A closer look

Increasing *strength* of connectivity (results whp)

Weaker — Stronger

a 'giant' component

1-connectivity

k-connectivity ($k \geq 2$)

Inhomogeneous Random K-out Graphs

(A node makes 1 selection wp $0 < \mu < 1$, else $K_n \geq 2$ selections)

requires K_n of order $\log(n)$

IEEE Trans. on Info. Theory '21

A closer look

Increasing *strength* of connectivity (results whp)

Weaker

Stronger

a 'giant' component

1-connectivity

***k*-connectivity ($k \geq 2$)**

Inhomogeneous Random K-out Graphs

(A node makes 1 selection wp $0 < \mu < 1$, else $K_n \geq 2$ selections)

requires K_n of order $\log(n)$

IEEE Trans. on Info. Theory '21

Threshold for k-connectivity

Recall: Each node is Type-I w.p. $\boldsymbol{\mu}$, Type-II w.p. $\mathbf{1}$-$\boldsymbol{\mu}$

(pick $\mathbf{1}$ node) (pick $\boldsymbol{K_n}$ nodes)

Theorem: k-connectivity

$$\lim_{n\to\infty} P\,[H(n;\,\mu,\,K_n) \text{ is } k\text{-connected}] = \begin{cases} 1 & \text{if } K_n = \dfrac{\log n + (k\text{-}2)\log\log n}{(1\text{-}\mu)} + \omega(1) \\ 0 & \text{if } K_n = \dfrac{\log n + (k\text{-}2)\log\log n}{(1\text{-}\mu)} - \omega(1) \end{cases}$$

where $0 < \mu < 1, \quad k = 2, 3, \ldots.$

Sood and Yagan, IEEE Transactions on Information Theory '21

Threshold for k-connectivity

1000 nodes, averaged over 1000 independently generated networks

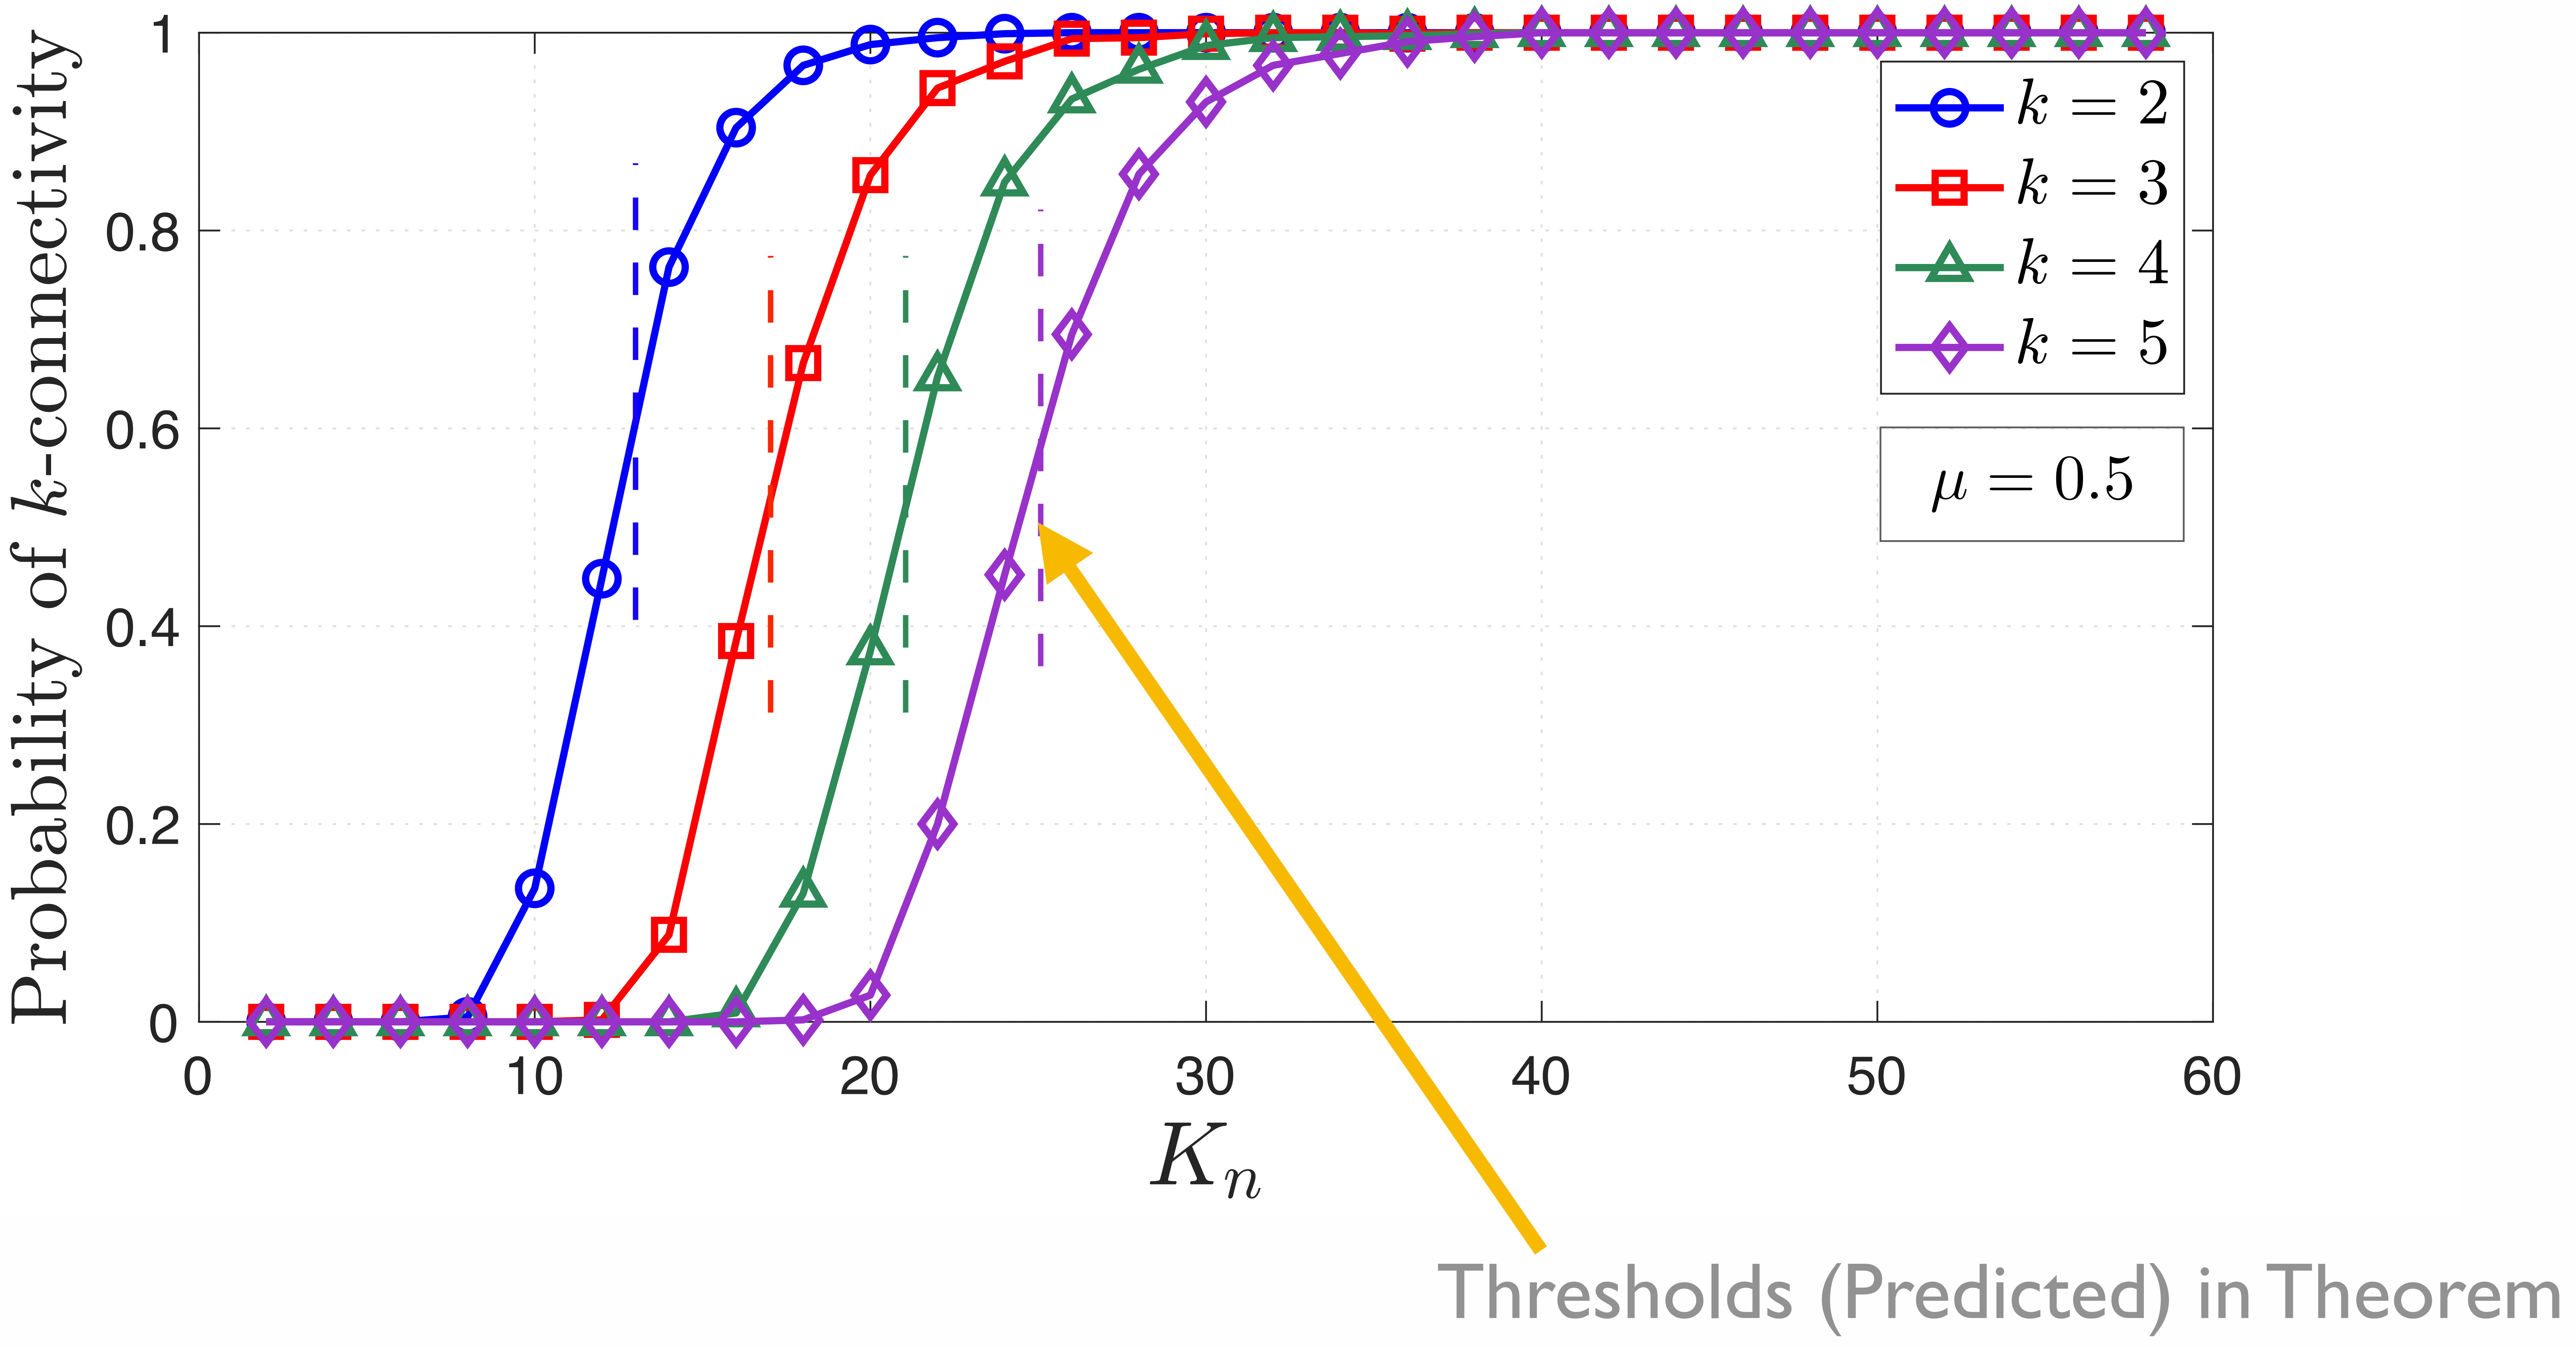

Method of moments: a recipe for reasoning about thresholds

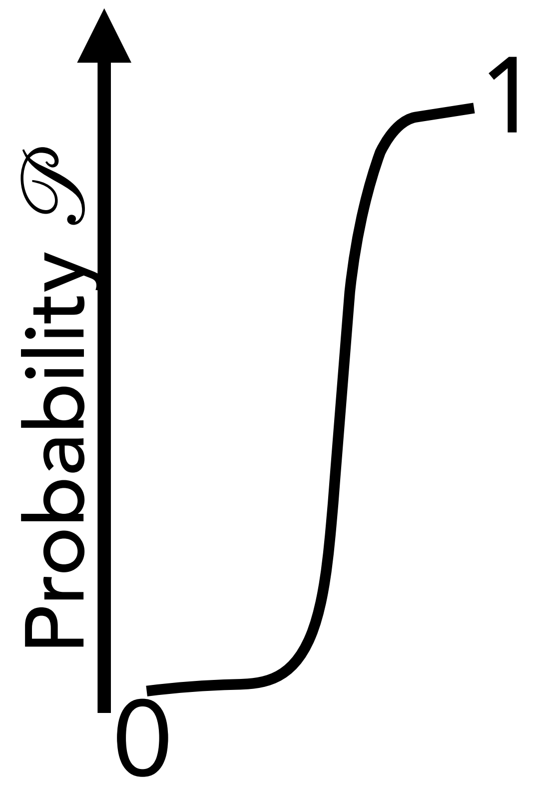

Method of moments: a recipe for reasoning about thresholds

- **Given:** A random graph $\mathbb{G}(n, \theta_n)$ and a property $\mathscr{P}$ that is *monotone* in edge addition

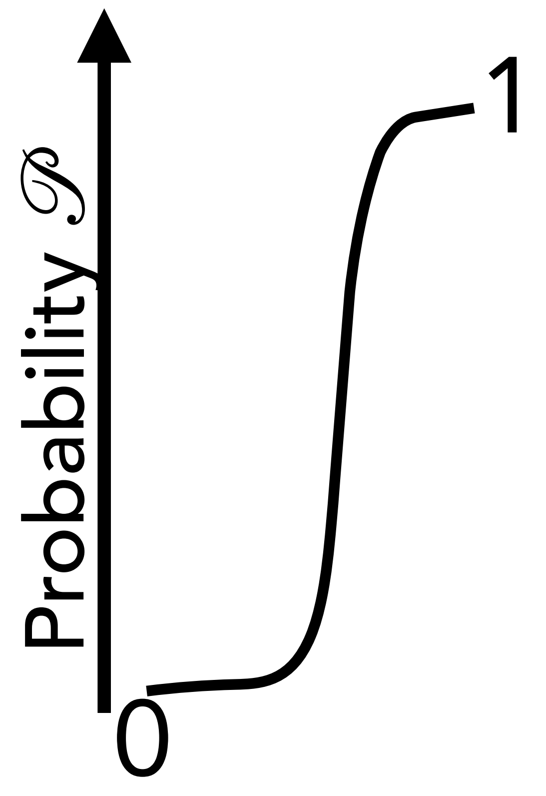

Method of moments: a recipe for reasoning about thresholds

- **Given:** A random graph $\mathbb{G}(n, \theta_n)$ and a property $\mathscr{P}$ that is *monotone* in edge addition
- **Goal:** Under what conditions (on θ_n) does $\mathbb{G}(n, \theta_n)$ admit a property $\mathscr{P}$ wp $\to 1$ or 0 as $n \to \infty$

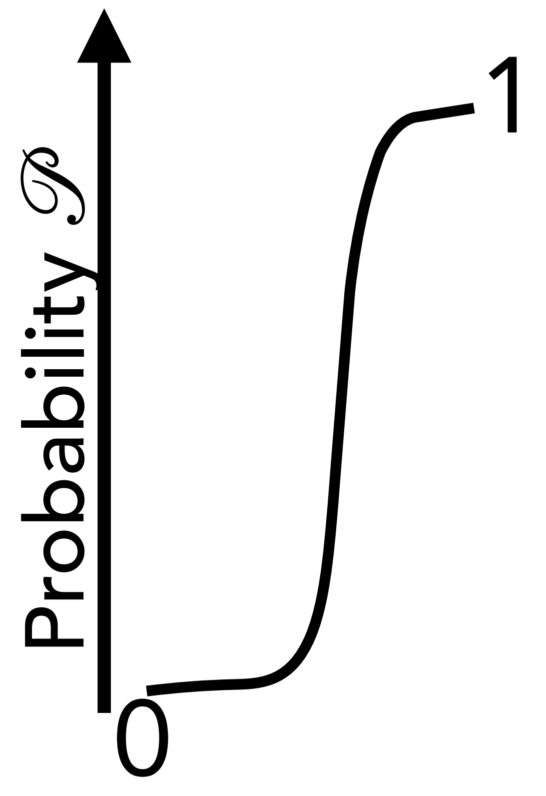

Method of moments: a recipe for reasoning about thresholds

- **Given:** A random graph $\mathbb{G}(n, \theta_n)$ and a property $\mathscr{P}$ that is *monotone* in edge addition
- **Goal:** Under what conditions (on θ_n) does $\mathbb{G}(n, \theta_n)$ admit a property $\mathscr{P}$ wp $\to 1$ or 0 as $n \to \infty$

  (Often defined in terms of a θ_n^* 'around which $\mathscr{P}$ abruptly changes from 0 to 1' "Threshold")

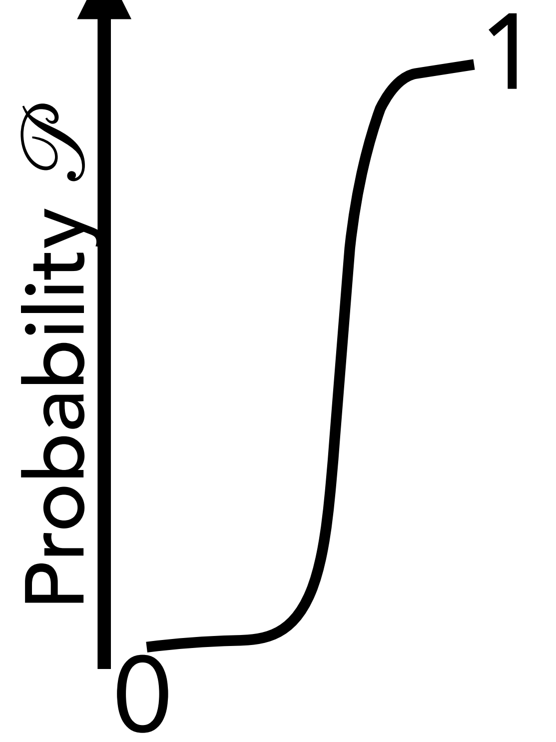

Method of moments: a recipe for reasoning about thresholds

- **Given:** A random graph $\mathbb{G}(n, \theta_n)$ and a property $\mathscr{P}$ that is *monotone* in edge addition
- **Goal:** Under what conditions (on θ_n) does $\mathbb{G}(n, \theta_n)$ admit a property $\mathscr{P}$ wp $\to 1$ or 0 as $n \to \infty$

  (Often defined in terms of a θ_n^* 'around which $\mathscr{P}$ abruptly changes from 0 to 1' "Threshold")

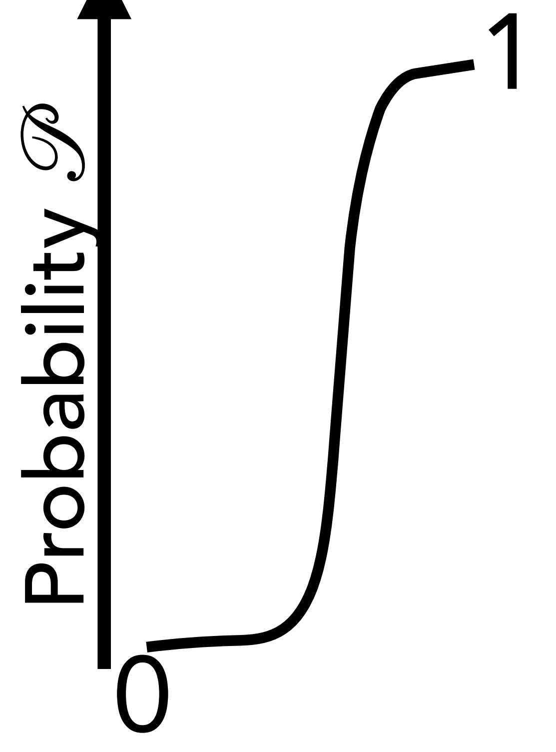

Method of moments: a recipe for reasoning about thresholds

- **Given:** A random graph $\mathbb{G}(n, \theta_n)$ and a property $\mathscr{P}$ that is *monotone* in edge addition
- **Goal:** Under what conditions (on θ_n) does $\mathbb{G}(n, \theta_n)$ admit a property $\mathscr{P}$ wp $\to 1$ or 0 as $n \to \infty$
  (Often defined in terms of a θ_n^* 'around which $\mathscr{P}$ abruptly changes from 0 to 1' "Threshold")

- Identify 'substructures' whose occurrence are 'bad events' that preclude property $\mathscr{P}$.

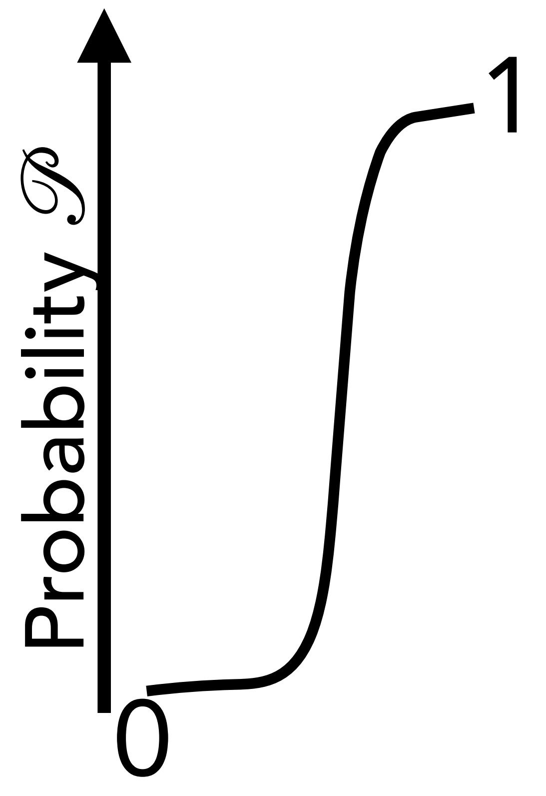

Method of moments: a recipe for reasoning about thresholds

- **Given:** A random graph $\mathbb{G}(n, \theta_n)$ and a property $\mathscr{P}$ that is *monotone* in edge addition
- **Goal:** Under what conditions (on θ_n) does $\mathbb{G}(n, \theta_n)$ admit a property $\mathscr{P}$ wp $\to 1$ or 0 as $n \to \infty$
  (Often defined in terms of a θ_n^* 'around which $\mathscr{P}$ abruptly changes from 0 to 1' "Threshold")

- Identify 'substructures' whose occurrence are 'bad events' that preclude property $\mathscr{P}$.

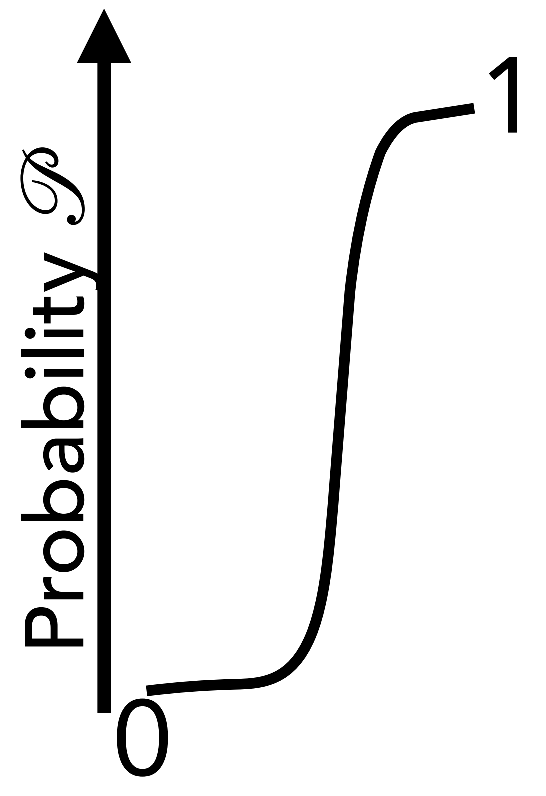

Method of moments: a recipe for reasoning about thresholds

- **Given:** A random graph $\mathbb{G}(n, \theta_n)$ and a property $\mathscr{P}$ that is *monotone* in edge addition
- **Goal:** Under what conditions (on θ_n) does $\mathbb{G}(n, \theta_n)$ admit a property $\mathscr{P}$ wp $\to 1$ or 0 as $n \to \infty$
  (Often defined in terms of a θ_n^* 'around which $\mathscr{P}$ abruptly changes from 0 to 1' "Threshold")

- Identify 'substructures' whose occurrence are 'bad events' that preclude property $\mathscr{P}$.

- Suppose $Z_i,\ i = 1,2,...$ are indicator random variables for 'bad events' and further

$$\mathbb{P}\left[\Sigma_i Z_i \geq 1\right] \to 0 \implies \mathscr{P} \text{ wp } \to 1$$
$$\mathbb{P}\left[\Sigma_i Z_i \geq 1\right] \to 1 \implies \mathscr{P} \text{ wp } \to 0$$

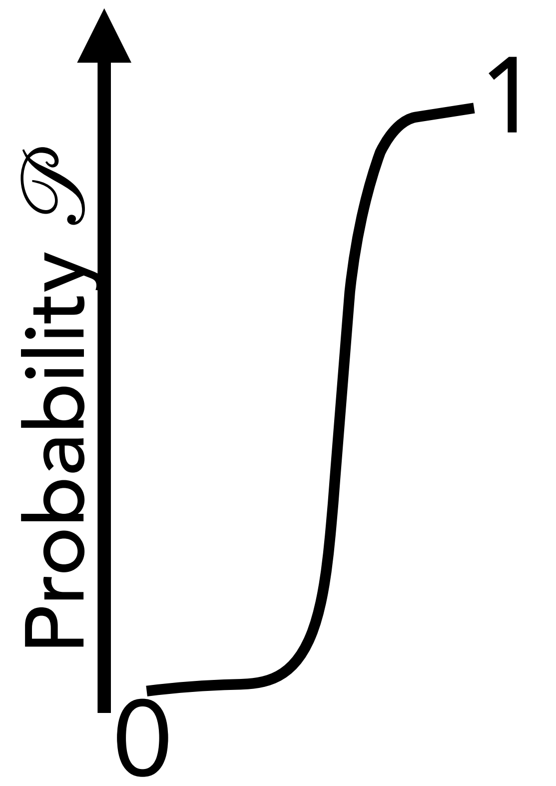

Method of moments: a recipe for reasoning about thresholds

- **Given:** A random graph $\mathbb{G}(n, \theta_n)$ and a property $\mathscr{P}$ that is *monotone* in edge addition
- **Goal:** Under what conditions (on θ_n) does $\mathbb{G}(n, \theta_n)$ admit a property $\mathscr{P}$ wp $\to 1$ or 0 as $n \to \infty$

  (Often defined in terms of a θ_n^* 'around which $\mathscr{P}$ abruptly changes from 0 to 1' "Threshold")

- Identify 'substructures' whose occurrence are 'bad events' that preclude property $\mathscr{P}$.

- Suppose $Z_i,\ i = 1,2,...$ are indicator random variables for 'bad events' and further

$$\mathbb{P}\left[\Sigma_i Z_i \geq 1\right] \to 0 \implies \mathscr{P} \text{ wp } \to 1$$

$$\mathbb{P}\left[\Sigma_i Z_i \geq 1\right] \to 1 \implies \mathscr{P} \text{ wp } \to 0$$

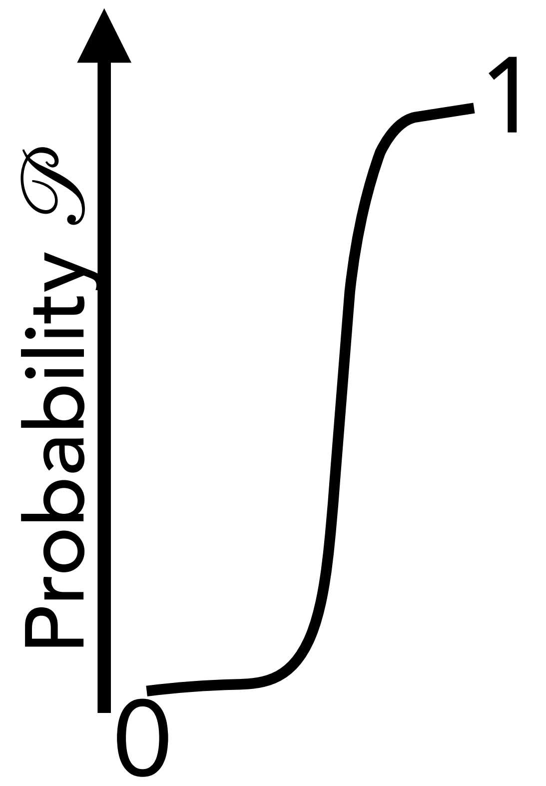

- In the method of moments, we bound the above two events using moments of Z_is.

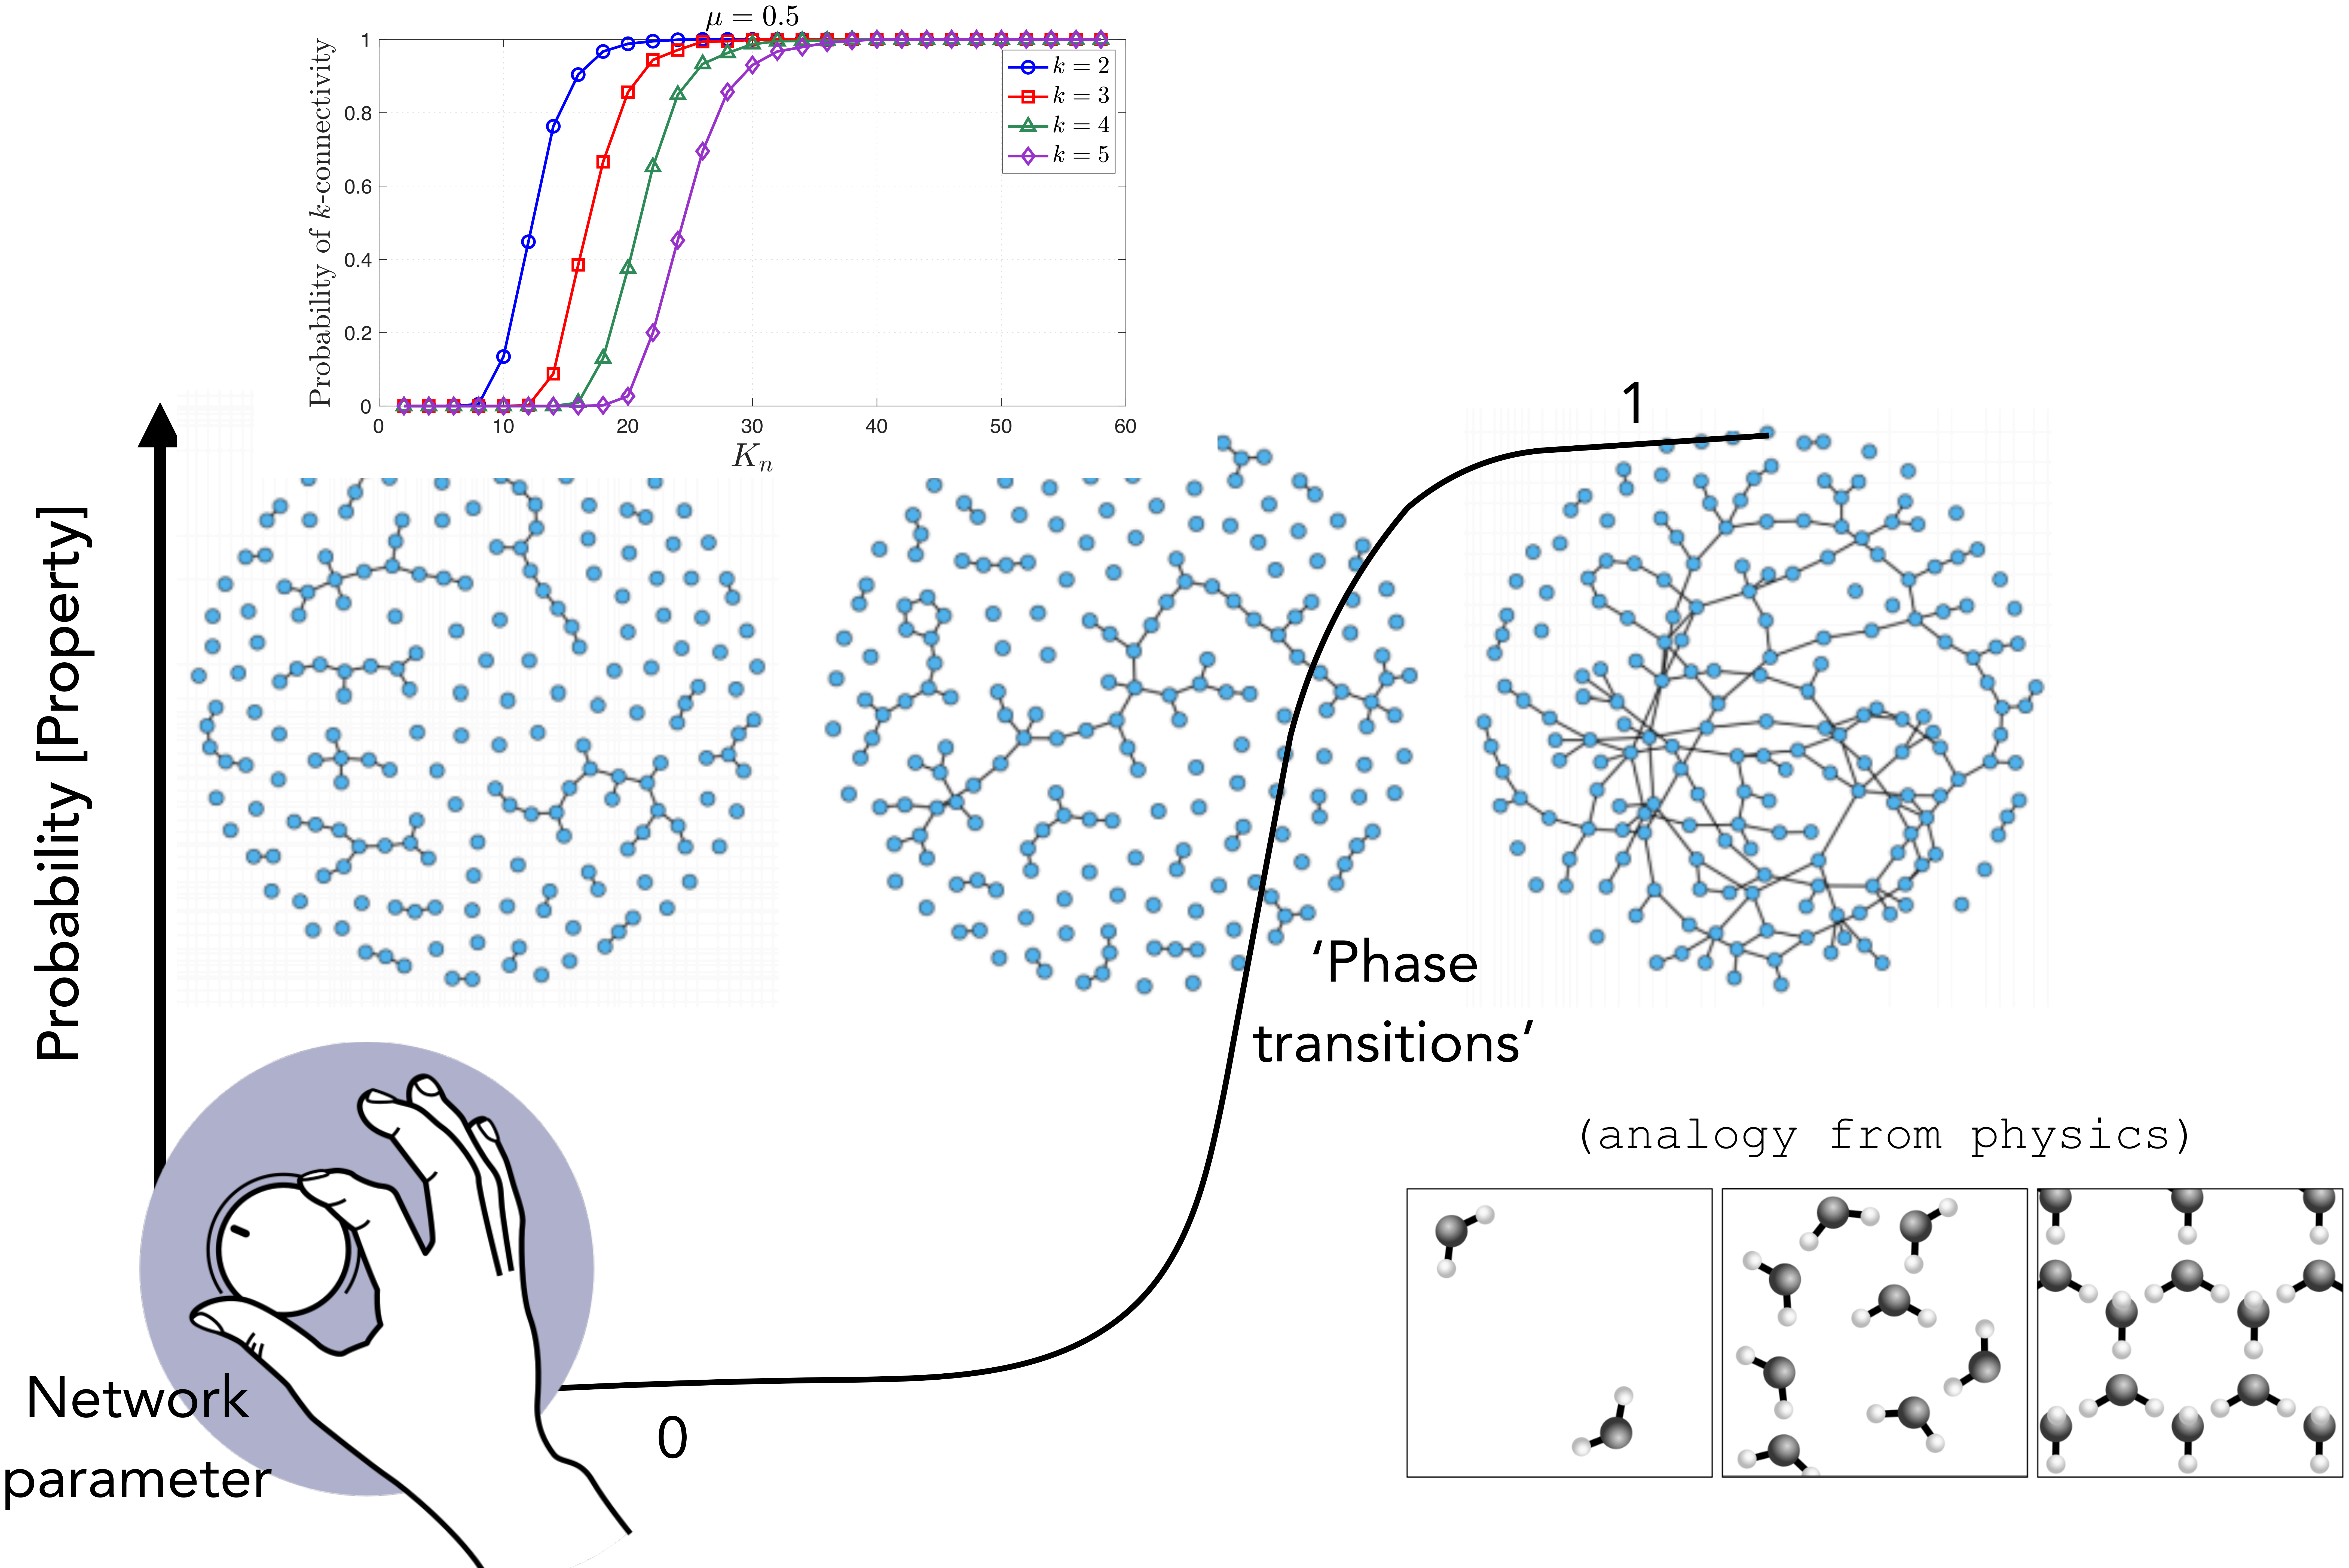

$\mu = 0.5$
Probability of k-connectivity
K_n
$k = 2$
$k = 3$
$k = 4$
$k = 5$
Probability [Property]
1
0
'Phase transitions'
(analogy from physics)
Network parameter

Method of moments: a recipe for reasoning about thresholds

- Suppose Z_i, $i = 1,2,\ldots$ are indicator random varables for 'bad' events / substructures

'First-moment' technique:

$$\mathbb{P}\left[\Sigma_i Z_i \geq 1\right] \to 0 \implies \mathscr{P} \text{ wp } \to 1$$

$$\leq \mathbb{E}[\Sigma_i Z_i]$$

suffices to argue that $\mathbb{E}[\Sigma_i Z_i] \to 0$

requires only 'first-moment'

'Second-moment' technique:

$$\mathbb{P}\left[\Sigma_i Z_i \geq 1\right] \to 1 \implies \mathscr{P} \text{ wp } \to 0$$

$$\mathbb{P}[\Sigma_i Z_i \geq 1] \geq \frac{\mathbb{E}^2[\Sigma_i Z_i]}{\mathbb{E}[(\Sigma_i Z_i)^2]}$$

e.g., argue with CS

$$\begin{aligned}\mathbb{E}^2[\Sigma_i Z_i] &= \mathbb{E}^2[(\Sigma_i Z_i) \times \mathbb{I}\{\Sigma_i Z_i \geq 1\}] \\ &\leq \mathbb{E}[(\Sigma_i Z_i)^2] \times \mathbb{P}[\Sigma_i Z_i \geq 1]\end{aligned}$$

Method of moments: a recipe for reasoning about thresholds

- Suppose Z_i, $i = 1,2,\ldots$ are indicator random varables for 'bad' events / substructures

'First-moment' technique:

$$\mathbb{P}\left[\Sigma_i Z_i \geq 1\right] \to 0 \implies \mathscr{P} \text{ wp } \to 1$$

$$\leq \mathbb{E}[\Sigma_i Z_i]$$

suffices to argue that $\mathbb{E}[\Sigma_i Z_i] \to 0$

requires only 'first-moment'

'Second-moment' technique:

$$\mathbb{P}\left[\Sigma_i Z_i \geq 1\right] \to 1 \implies \mathscr{P} \text{ wp } \to 0$$

$$\mathbb{P}[\Sigma_i Z_i \geq 1] \geq \frac{\mathbb{E}^2[\Sigma_i Z_i]}{\mathbb{E}[(\Sigma_i Z_i)^2]}$$

e.g., argue with CS

$$\mathbb{E}^2[\Sigma_i Z_i] = \mathbb{E}^2[(\Sigma_i Z_i) \times \mathbb{I}\{\Sigma_i Z_i \geq 1\}]$$
$$\leq \mathbb{E}[(\Sigma_i Z_i)^2] \times \mathbb{P}[\Sigma_i Z_i \geq 1]$$

Method of moments: a recipe for reasoning about thresholds

- Suppose Z_i, $i = 1,2,\ldots$ are indicator random varables for 'bad' events / substructures

'First-moment' technique:

$\mathbb{P}\left[\Sigma_i Z_i \geq 1\right] \to 0 \implies \mathscr{P}$ wp $\to 1$

$\leq \mathbb{E}[\Sigma_i Z_i]$

suffices to argue that $\mathbb{E}[\Sigma_i Z_i] \to 0$

requires only 'first-moment'

'Second-moment' technique:

$\mathbb{P}\left[\Sigma_i Z_i \geq 1\right] \to 1 \implies \mathscr{P}$ wp $\to 0$

$$\mathbb{P}[\Sigma_i Z_i \geq 1] \geq \frac{\mathbb{E}^2[\Sigma_i Z_i]}{\mathbb{E}[(\Sigma_i Z_i)^2]}$$

This tends to be challenging to prove (dependence)

e.g., argue with CS

$$\begin{aligned}\mathbb{E}^2[\Sigma_i Z_i] &= \mathbb{E}^2[(\Sigma_i Z_i) \times \mathbb{I}\{\Sigma_i Z_i \geq 1\}] \\ &\leq \mathbb{E}[(\Sigma_i Z_i)^2] \times \mathbb{P}[\Sigma_i Z_i \geq 1]\end{aligned}$$

Method of moments: a recipe for reasoning about thresholds

- Suppose Z_i, $i = 1,2,\ldots$ are indicator random varables for 'bad' events / substructures

'First-moment' technique:

$\mathbb{P}\left[\Sigma_i Z_i \geq 1\right] \to 0 \implies \mathscr{P}$ wp $\to 1$

$\leq \mathbb{E}[\Sigma_i Z_i]$

suffices to argue that $\mathbb{E}[\Sigma_i Z_i] \to 0$
requires only 'first-moment'

'Second-moment' technique:

$\mathbb{P}\left[\Sigma_i Z_i \geq 1\right] \to 1 \implies \mathscr{P}$ wp $\to 0$

$$\mathbb{P}[\Sigma_i Z_i \geq 1] \geq \frac{\mathbb{E}^2[\Sigma_i Z_i]}{\mathbb{E}[(\Sigma_i Z_i)^2]}$$

This tends to be challenging to prove (dependence)

e.g., argue with CS

$\mathbb{E}^2[\Sigma_i Z_i] =$

K-out graph mechanism induces dependencies in edge generation

Method of moments: a recipe for reasoning about thresholds

- Suppose Z_i, $i = 1,2,...$ are indicator random varables for 'bad' events / substructures

'First-moment' technique:

$\mathbb{P}\left[\Sigma_i Z_i \geq 1\right] \to 0 \implies \mathscr{P}$ wp $\to 1$

$\leq \mathbb{E}[\Sigma_i Z_i]$

suffices to argue that $\mathbb{E}[\Sigma_i Z_i] \to 0$

requires only 'first-moment'

'Second-moment' technique:

$\mathbb{P}\left[\Sigma_i Z_i \geq 1\right] \to 1 \implies \mathscr{P}$ wp $\to 0$

$$\mathbb{P}[\Sigma_i Z_i \geq 1] \geq \frac{\mathbb{E}^2[\Sigma_i Z_i]}{\mathbb{E}[(\Sigma_i Z_i)^2]}$$

This tends to be challenging to prove (dependence)

e.g., argue with CS

$\mathbb{E}^2[\Sigma_i Z_i] =$

K-out graph mechanism induces dependencies in edge generation

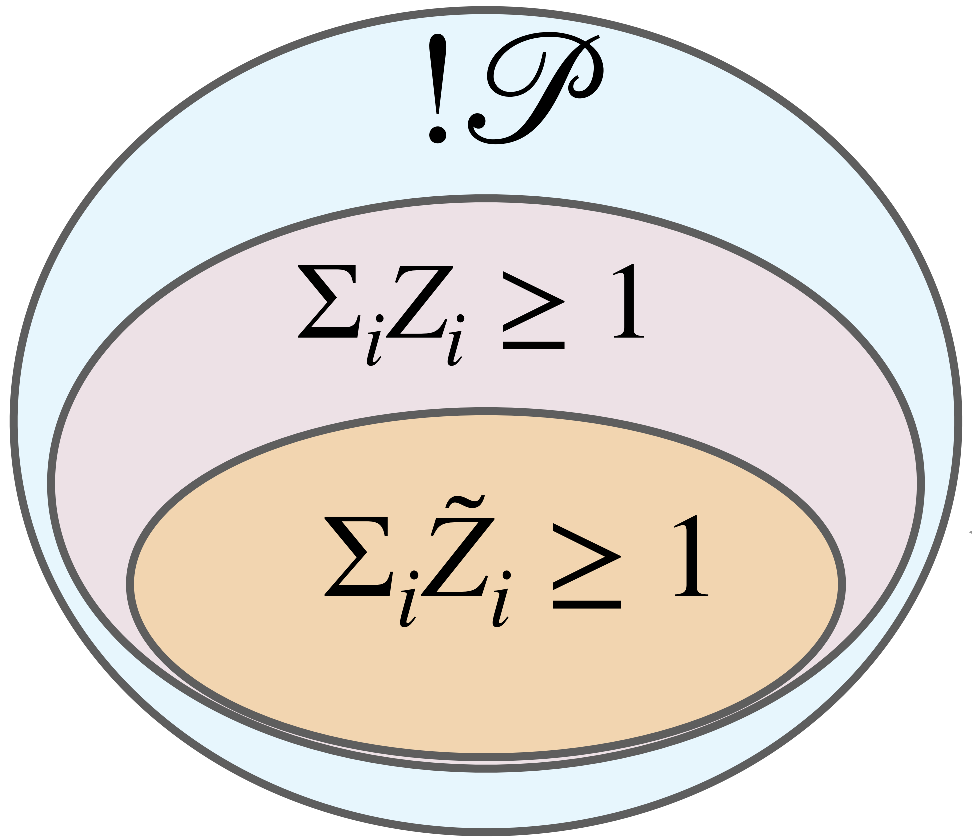

Often useful to define another bad event which is a subset, easier to compute, & yet drives the limits suitably

k-vertex connectivity

Minimum node degree (δ)

Minimum no. of edges to be deleted to disconnect the network (k_E) (edge connectivity)

Minimum no. of vertices to be deleted to disconnect the network (k_V) (vertex connectivity)

$$k_V \leq k_E \leq \delta$$

$$k_V \geq k \Rightarrow k_E \geq k \Rightarrow \delta \geq k$$

(Strongest)

minimum node degree ≥ k does not imply k-connectivity

An illustration $k_V < k_E < \boldsymbol{\delta}$

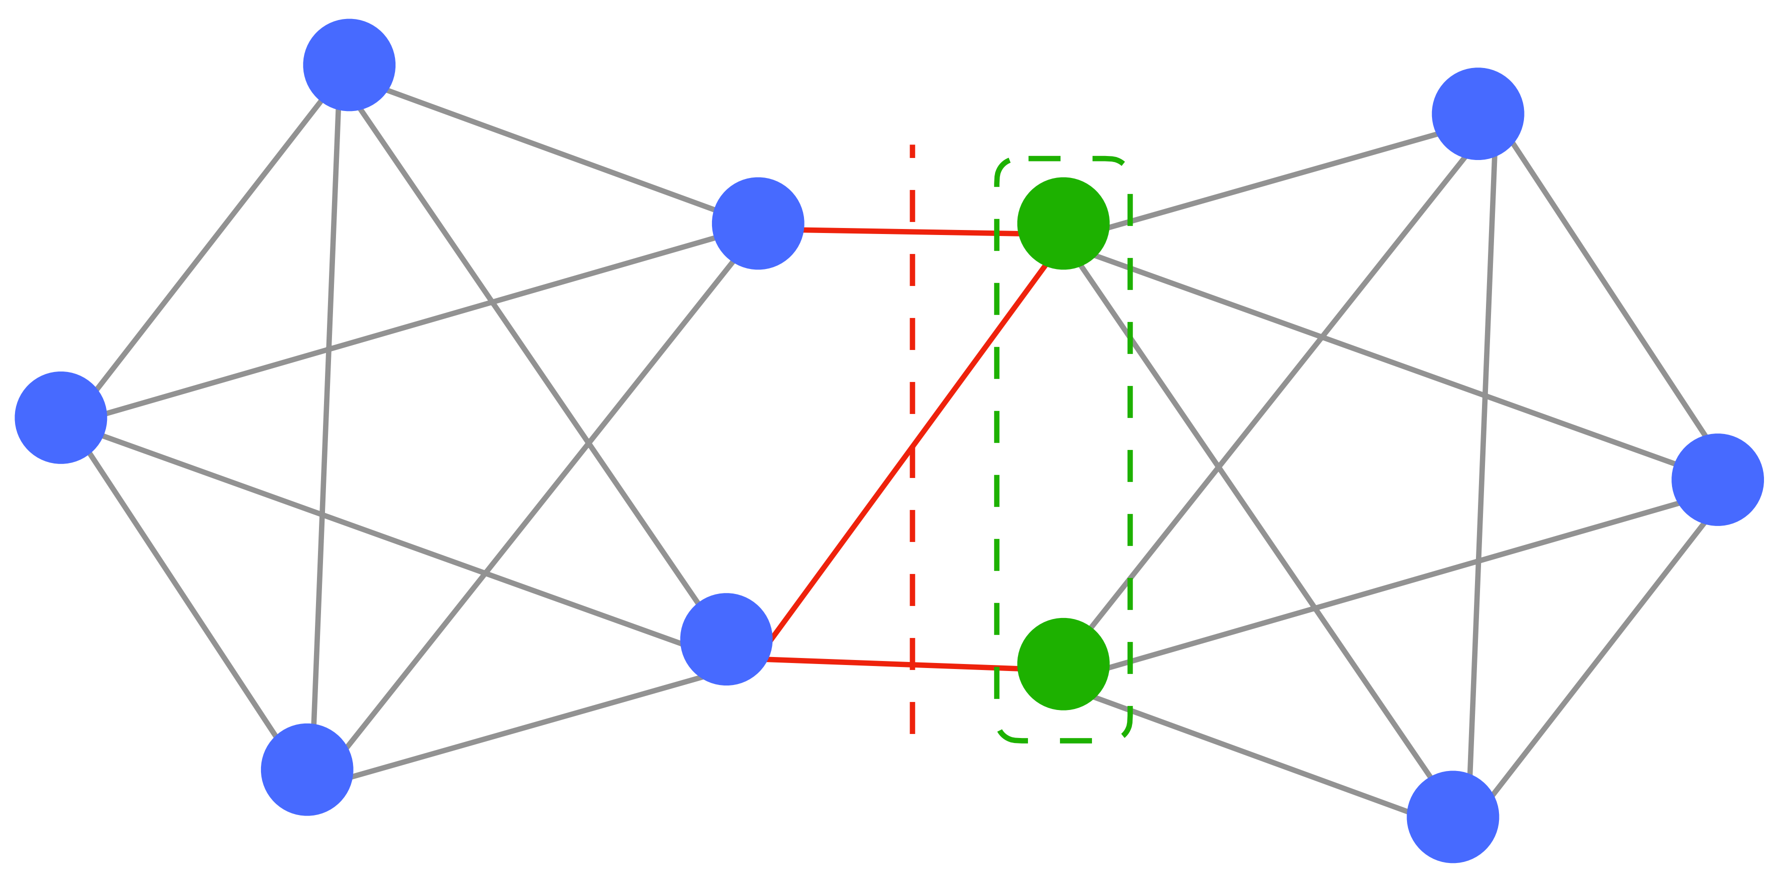

Minimum node degree ($\boldsymbol{\delta}$) = 4

Minimum no. of edges to be deleted to disconnect (k_E) = 3

Minimum no. of nodes to be deleted to disconnect (k_V) = 2

Threshold for Minimum node degree $\delta \geq k$

Recall: Each node is Type-I w.p. $\boldsymbol{\mu}$, Type-II w.p. $\mathbf{1}\text{-}\boldsymbol{\mu}$

(pick $\mathbf{1}$ node) (pick $\boldsymbol{K_n}$ nodes)

Theorem: minimum node degree

$$\lim_{n\to\infty} P\,[H(n;\, \mu,\, K_n) \text{ has } \delta \geq k] = \begin{cases} 1 & \text{if } K_n = \dfrac{\log n + (k\text{-}2)\log\log n}{(1\text{-}\mu)} + \omega(1) \\ 0 & \text{if } K_n = \dfrac{\log n + (k\text{-}2)\log\log n}{(1\text{-}\mu)} - \omega(1) \end{cases}$$

where $0 < \mu < 1, \quad k = 2, 3, \ldots.$

*Sood and Yagan, IEEE Transactions on Information Theory '21

k-connectivity (zero law)

Minimum node degree $\delta \geq k$ is necessary for k-connectivity

Theorem: minimum node degree*

$$\lim_{n\to\infty} P\,[H(n;\mu, K_n) \text{ has } \delta \geq k] = \begin{cases} 1 & \text{if } K_n = \dfrac{\log n + (k\text{-}2)\log\log n}{(1\text{-}\mu)} + \omega(1) \\ 0 & \text{if } K_n = \dfrac{\log n + (k\text{-}2)\log\log n}{(1\text{-}\mu)} - \omega(1) \end{cases}$$

where $0 < \mu < 1, \;\; k = 2, 3, \ldots.$

*Sood and Yagan, IEEE Transactions on Information Theory '21

k-connectivity (one law)

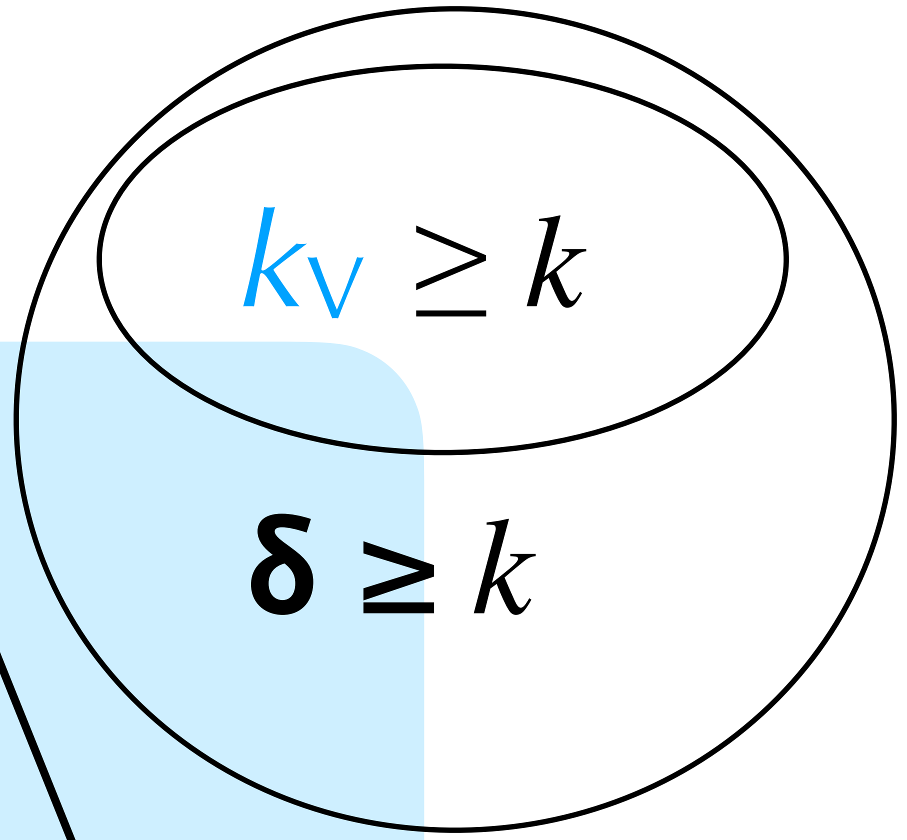

$0 < \mu < 1, \quad k = 2, 3, \ldots .$ if $K_n = \dfrac{\log n + (k\text{-}2) \log\log n}{(1\text{-}\mu)} + \gamma(n)$

$\lim_{n\to\infty} P\,[\boldsymbol{\delta} \geq k] = 1$ $\quad k_V \leq \boldsymbol{\delta}$

To show: $\lim_{n\to\infty} P\,[k_V \geq k] = 1$

$$P\,[k_V \geq k] = P\,[\boldsymbol{\delta} \geq k] - P\,[k_V < k, \boldsymbol{\delta} \geq k]$$

Prove that → 0

$$P[\{\kappa_v = \ell, \delta > \ell\}] \to 0, \ell = 0, 1, \ldots, k-1$$

Prove that → 0

Focus on sequences $\gamma_n = \omega(1), \gamma_n = O(\log n)$

Use coupling argument to show that the monotone property also holds for the general case

$0 < \mu < 1, \quad k = 2, 3, \ldots .$ if $K_n = \dfrac{\log n + (k\text{-}2) \log\log n + \gamma_n}{(1-\mu)}$

Show that

$$P[\{\kappa_v = \ell, \delta > \ell\}] \to 0, \ell = 0, 1, \ldots, k-1$$

Step 2: Focus on sequences $\gamma_n = \omega(1), \gamma_n = O(\log n)$

Use coupling argument to show that if the monotone property (k-con) holds for this case, then it also holds for the general case

Necessary condition for $P[\{\kappa_v = \ell, \delta > \ell\}] = B_{UT} \cap C_T \cap D_{UT}$

$$P[\{\kappa_v = \ell, \delta > \ell\}] \to 0, \ell = 0,1,\ldots,k-1$$

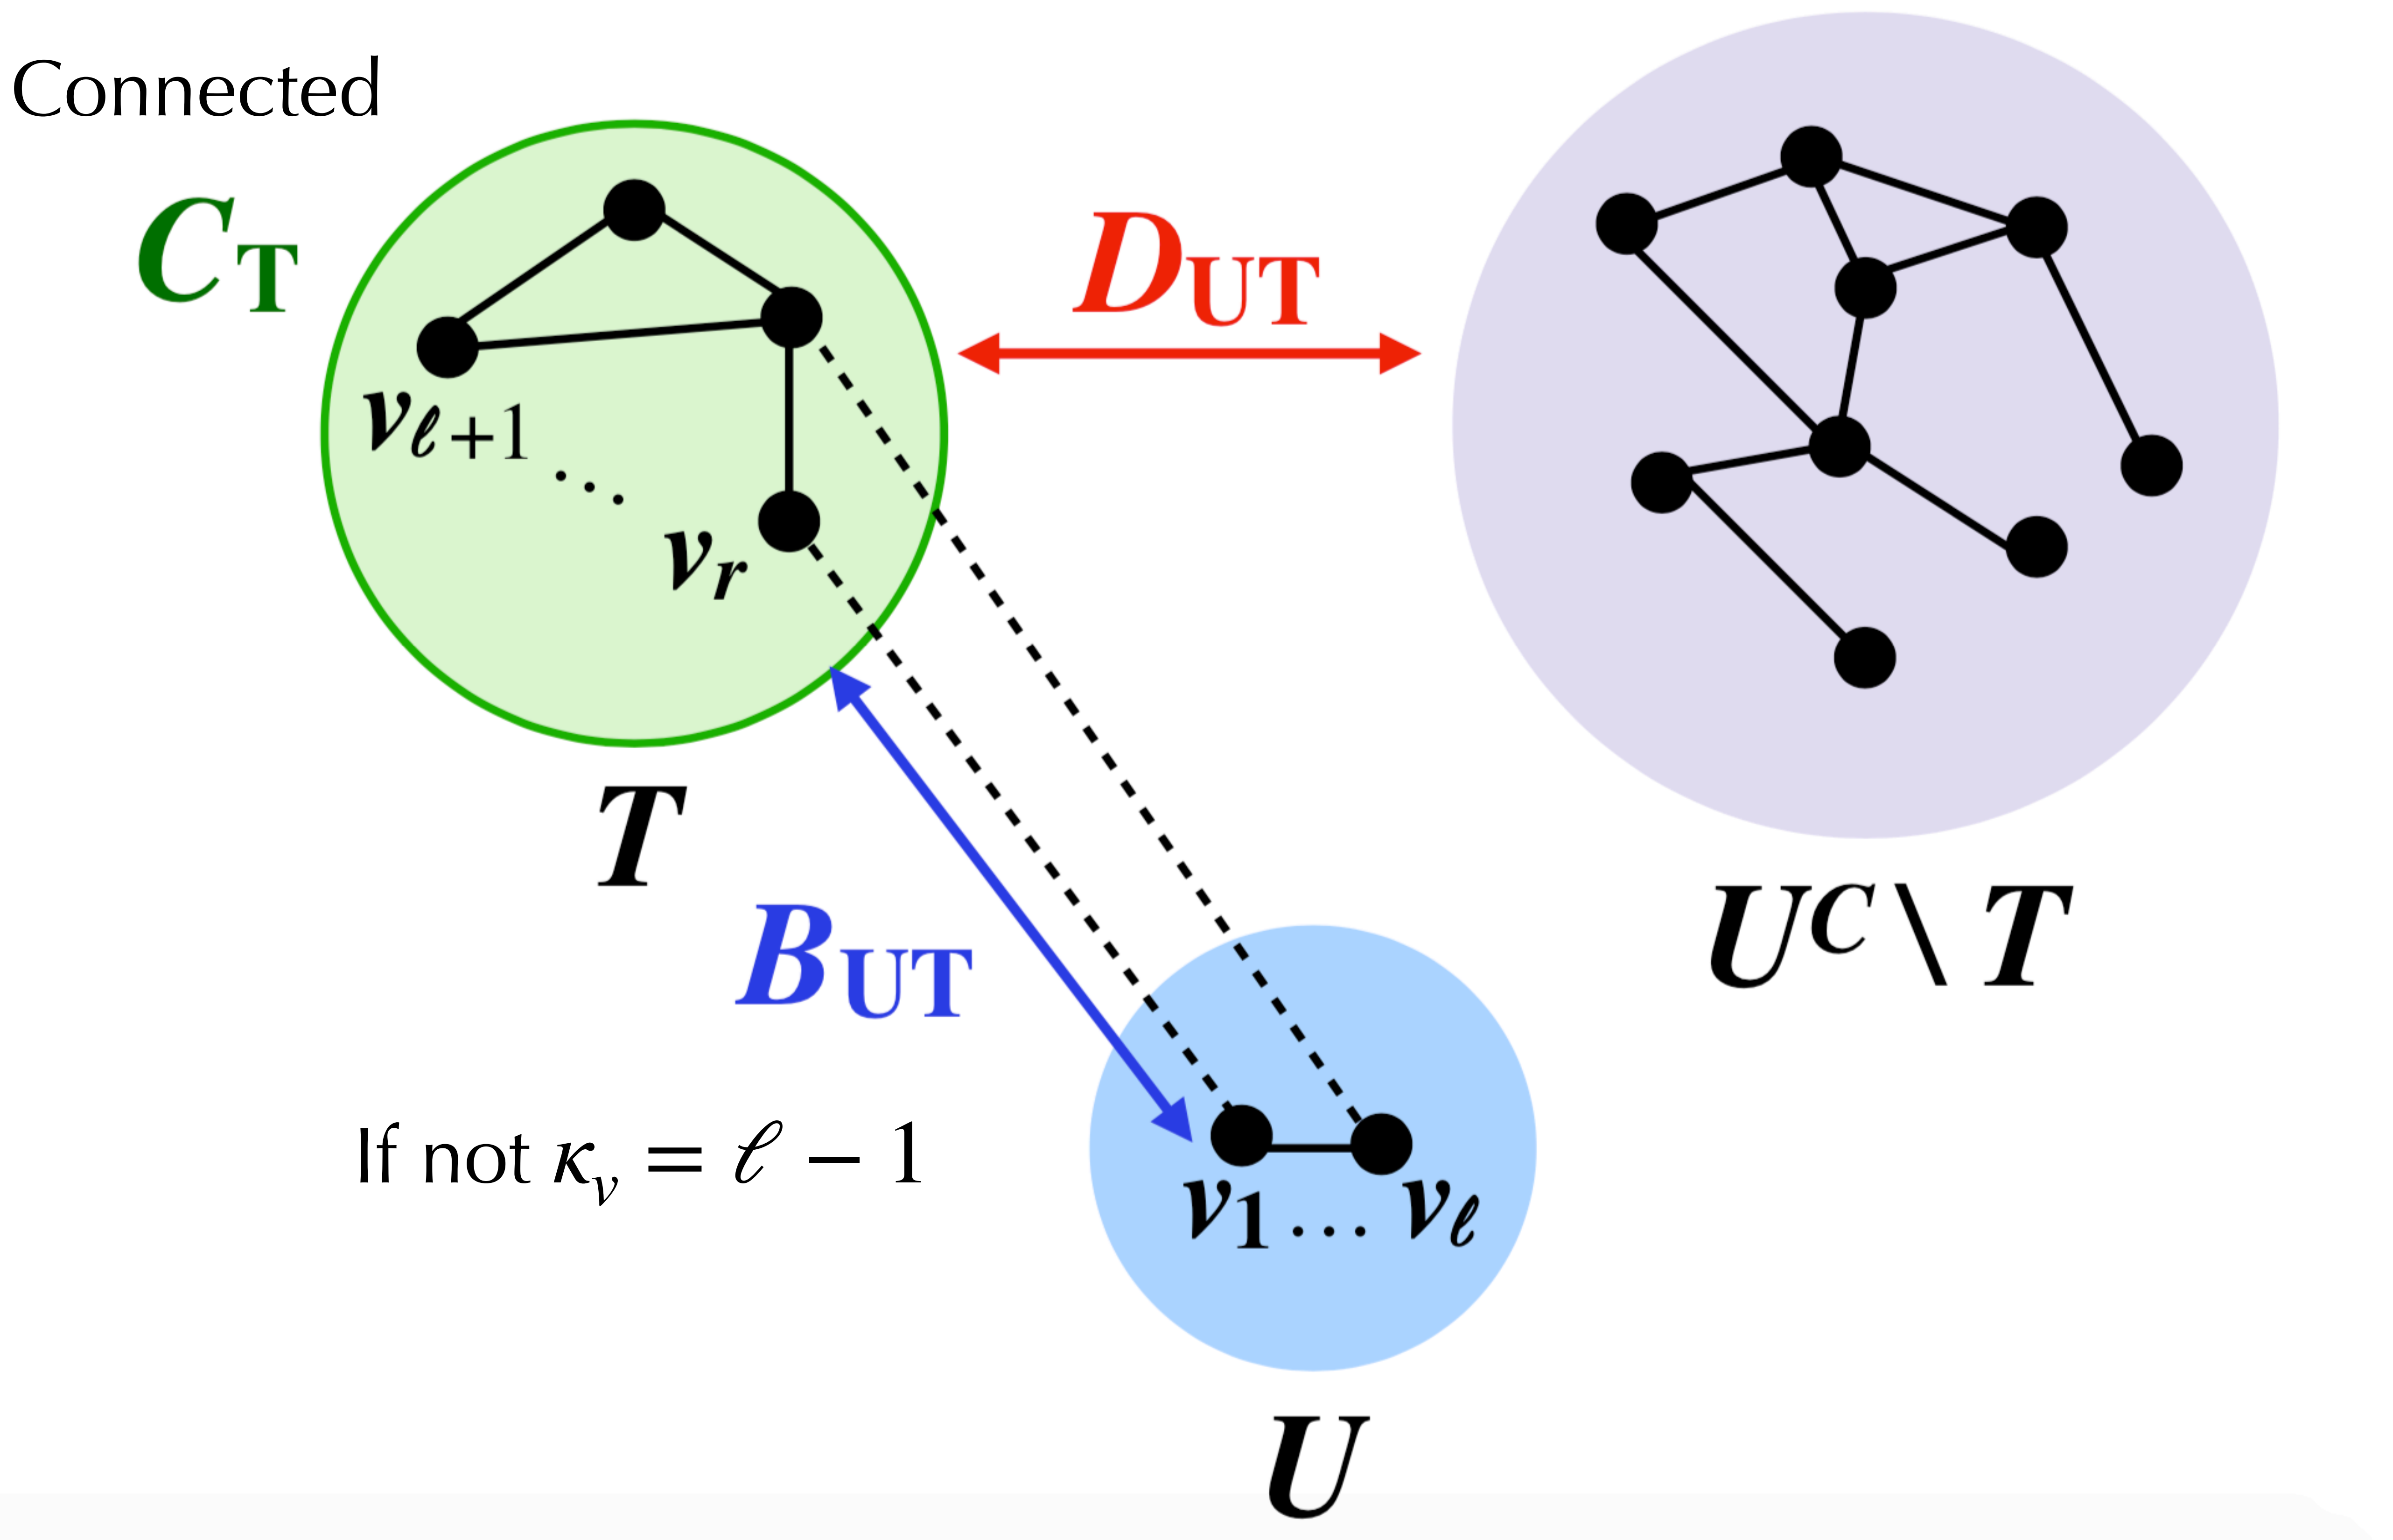

$|T| \geq 2$
$(\delta \geq \ell + 1)$

For $\mathbb{H}(n;\mu,K_n)$, if $\delta > \ell$ and $\kappa_v = \ell$ then $\exists U, T \subset \mathcal{N}$ such that $|U| = \ell$ and $2 \leq |T| \leq \lfloor \frac{n-\ell}{2} \rfloor$ and the following events occur.

1. $\mathcal{C}_T$: $\mathbb{I}(n;\mu,K_n)(T)$ is connected,
2. $\mathcal{B}_{UT}$: All nodes in U have a neighbor in the set T,
3. $\mathcal{D}_{UT}$: T is *isolated* in $\mathbb{I}(n;\mu,K_n)(U^c)$; i.e., there are no edges in $\mathbb{H}(n;\mu,K_n)$ between nodes in T and nodes in $U^c \setminus T$.

$$P[\{\kappa_v = \ell, \delta > \ell\}] \leq P[\bigcup_{U,T} B_{UT} \cap C_T \cap D_{UT}]$$

Goes to 0

Part 1: Road-map

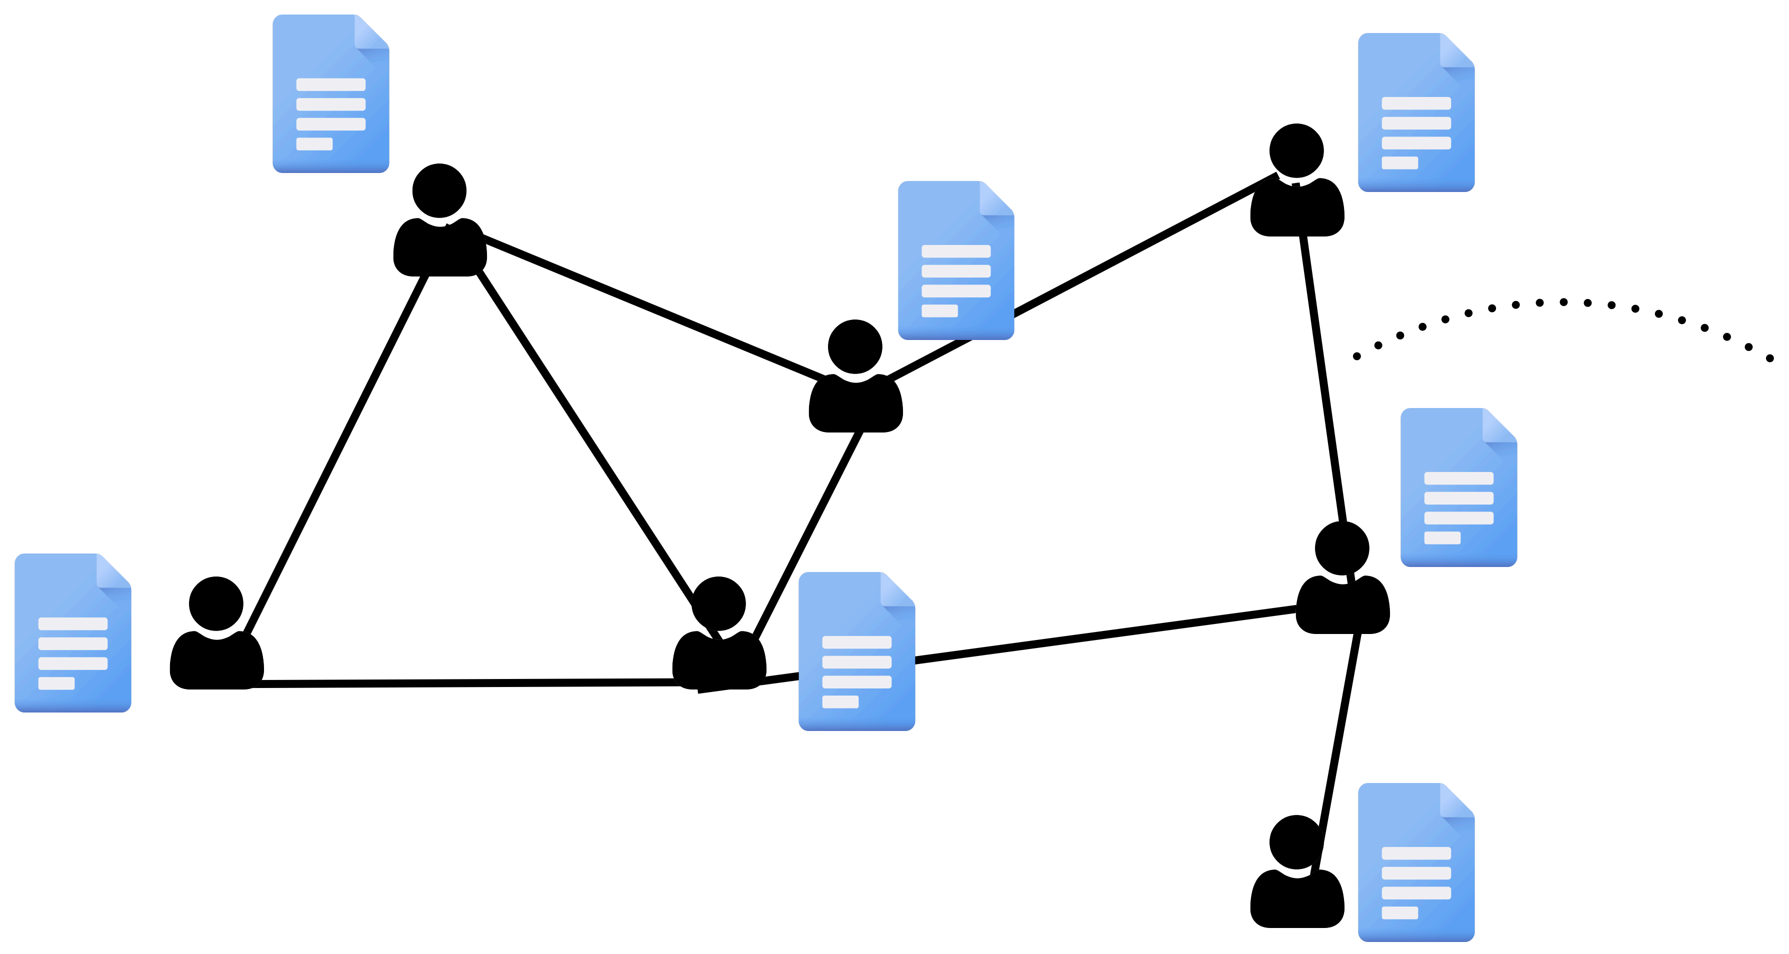

How to securely aggregate data?

[network design]

- From distributed systems to random graphs
  - The sparsity connectivity trade-off
- What are random K-out graphs?
  - What makes them a useful design tool?
- Random K-out graphs in action
  - What graph properties matter?
- Review: Notions of Connectivity
- Preview of contributions
  - Key proof techniques

Sparsity-connectivity in distributed network design

Sparsity-connectivity in distributed network design

- Balancing sparsity with connectivity is a fundamental challenge in distributed network design
  - Surprising manifestations of connectivity (e.g., privacy implications)

Sparsity-connectivity in distributed network design

- Balancing sparsity with connectivity is a fundamental challenge in distributed network design
    - Surprising manifestations of connectivity (e.g., privacy implications)

Sparsity-connectivity in distributed network design

- Balancing sparsity with connectivity is a fundamental challenge in distributed network design
  - Surprising manifestations of connectivity (e.g., privacy implications)

- Random K-out graphs
  - sparse yet well-connected topologies without a central coordinator
  - order wise comparable efficiency as fundamental limits, more efficient than ER graphs
  - widely appearing heuristic in application such as privacy mechanisms, with privacy $\leftrightarrow$ connectivity
  - analytical tractability to provide finer-grained guarantees against complex adversarial models

Sparsity-connectivity in distributed network design

- Balancing sparsity with connectivity is a fundamental challenge in distributed network design
  - Surprising manifestations of connectivity (e.g., privacy implications)

- Random K-out graphs
  - sparse yet well-connected topologies without a central coordinator
  - order wise comparable efficiency as fundamental limits, more efficient than ER graphs
  - widely appearing heuristic in application such as privacy mechanisms, with privacy $\leftrightarrow$ connectivity
  - analytical tractability to provide finer-grained guarantees against complex adversarial models

Sparsity-connectivity in distributed network design

- Balancing sparsity with connectivity is a fundamental challenge in distributed network design
    - Surprising manifestations of connectivity (e.g., privacy implications)

- Random K-out graphs
    - sparse yet well-connected topologies without a central coordinator
    - order wise comparable efficiency as fundamental limits, more efficient than ER graphs
    - widely appearing heuristic in application such as privacy mechanisms, with privacy $\leftrightarrow$ connectivity
    - analytical tractability to provide finer-grained guarantees against complex adversarial models

- We provide a characterization of properties that arise in different operational contexts (finite nodes, different failure modes, heterogeneity); challenging to analyze due to correlations in edge generation

Sparsity-connectivity in distributed network design

- Balancing sparsity with connectivity is a fundamental challenge in distributed network design
    - Surprising manifestations of connectivity (e.g., privacy implications)

- Random K-out graphs
    - sparse yet well-connected topologies without a central coordinator
    - order wise comparable efficiency as fundamental limits, more efficient than ER graphs
    - widely appearing heuristic in application such as privacy mechanisms, with privacy ↔ connectivity
    - analytical tractability to provide finer-grained guarantees against complex adversarial models

- We provide a characterization of properties that arise in different operational contexts (finite nodes, different failure modes, heterogeneity); challenging to analyze due to correlations in edge generation

Sparsity-connectivity in distributed network design

- Balancing sparsity with connectivity is a fundamental challenge in distributed network design
  - Surprising manifestations of connectivity (e.g., privacy implications)

- Random K-out graphs
  - sparse yet well-connected topologies without a central coordinator
  - order wise comparable efficiency as fundamental limits, more efficient than ER graphs
  - widely appearing heuristic in application such as privacy mechanisms, with privacy ↔ connectivity
  - analytical tractability to provide finer-grained guarantees against complex adversarial models

- We provide a characterization of properties that arise in different operational contexts (finite nodes, different failure modes, heterogeneity); challenging to analyze due to correlations in edge generation

- Limitations: Other modes of heterogeneity, sharper results for component sizes, dynamic graphs

Reliable Inference at Scale Using Graph Structure

notions of connectivity

Distributed Systems

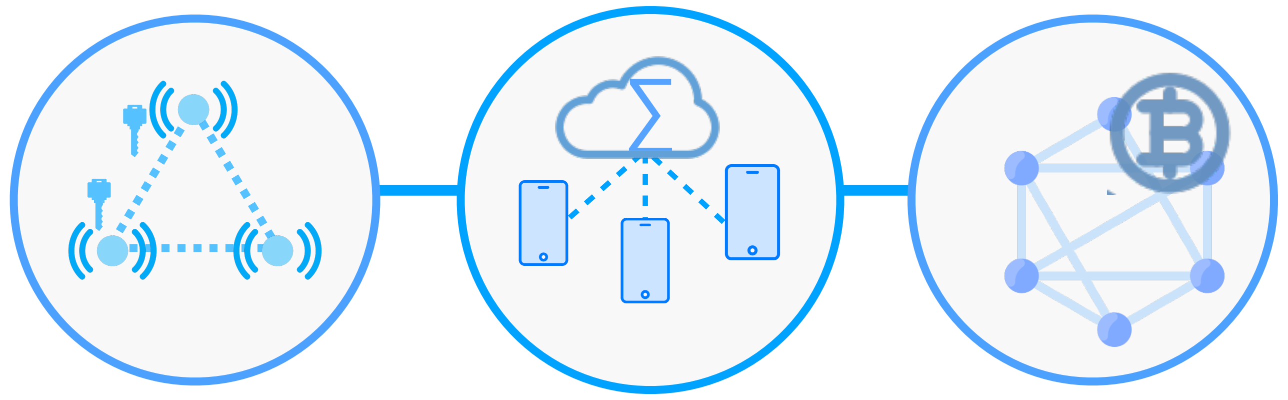

'secure' communication channel
between devices

Probabilistic Graphical Models

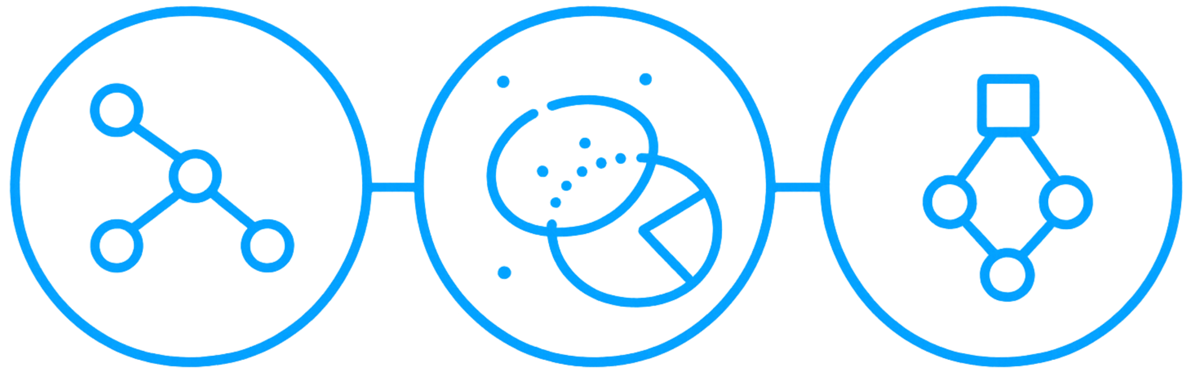

conditional dependencies
in observational data

efficiency

Combinatorics x distribution learning via graphical models

Graph encodes statistical dependencies between system variables

graph over random variables ↔ family of distributions consistent with combinatorial structure

~ absence of edges as independence relationships

$X_{I_1} \perp\!\!\!\perp X_{I_2} \mid \perp X_{I'}$ whenever I' separates I_1, I_2

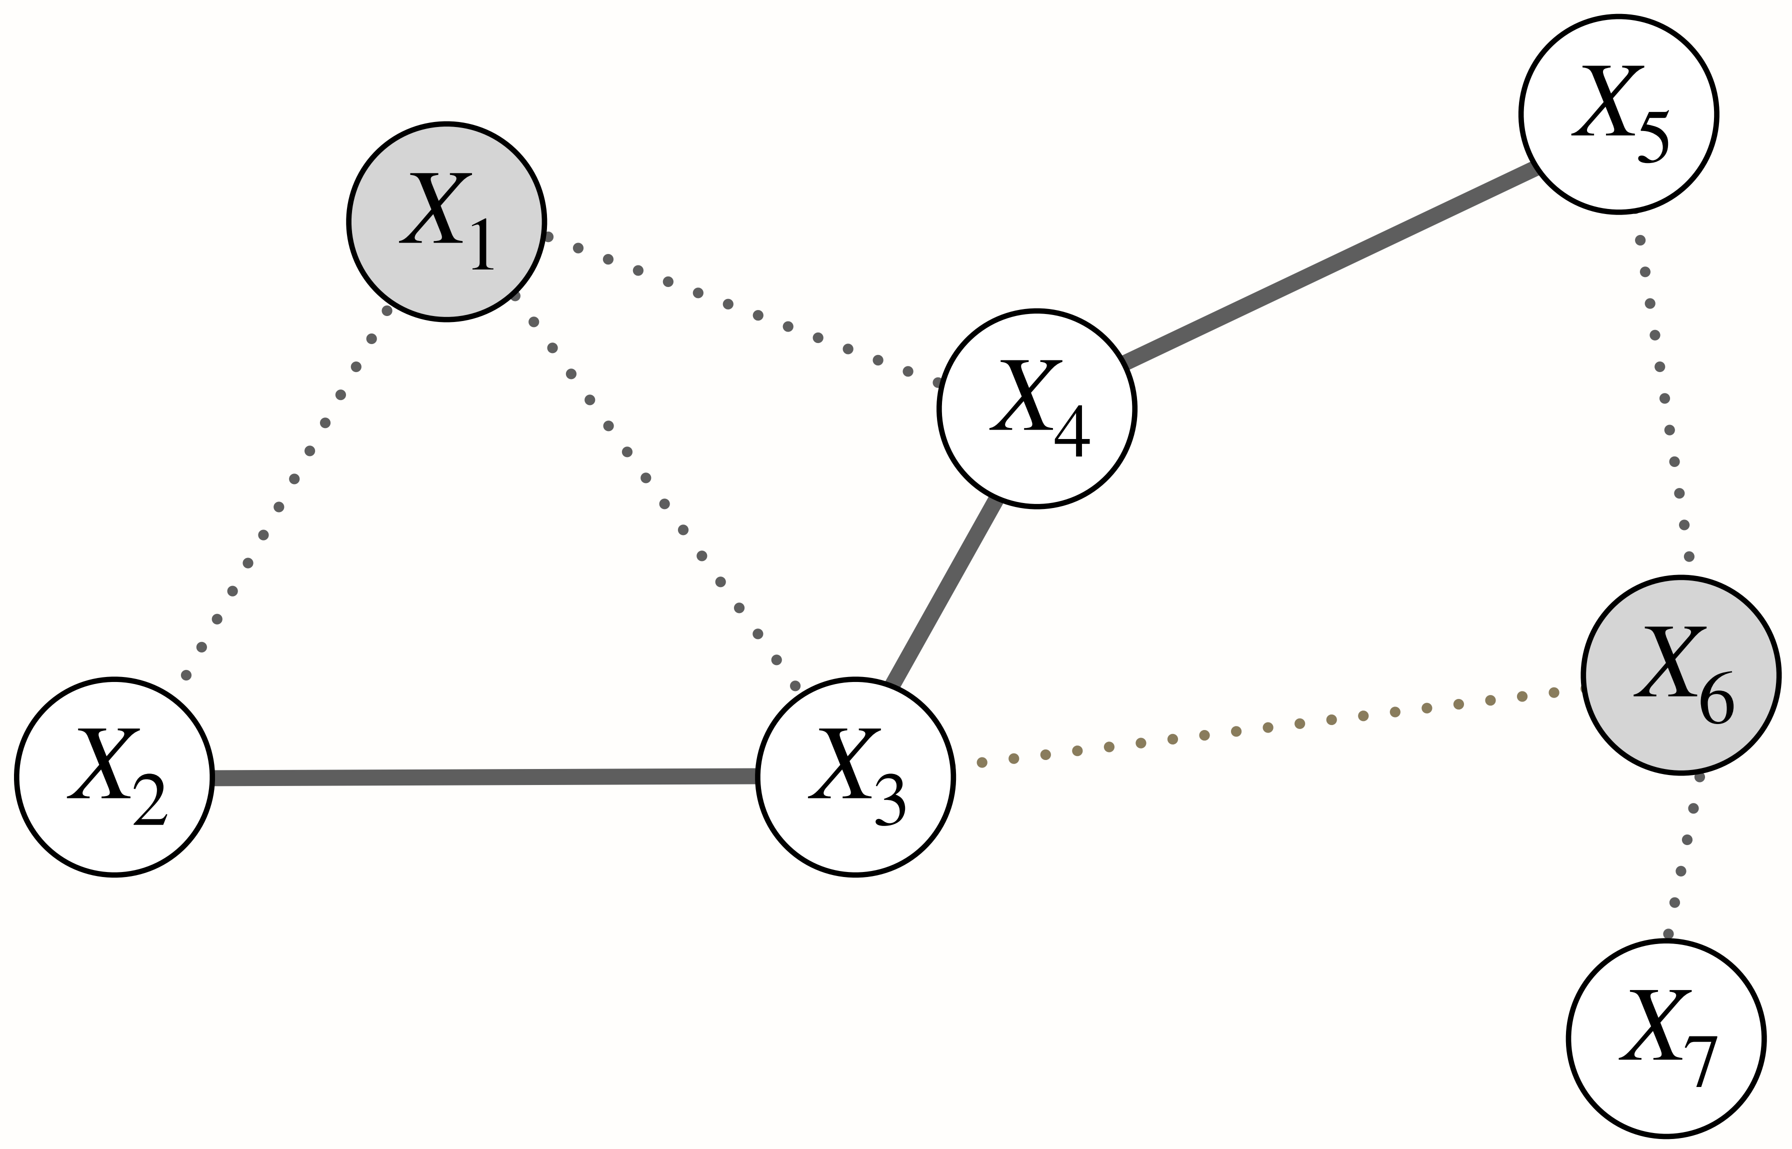

↑ edges: ↓ dependency assumptions, but ↑ parameters to learn

Combinatorics x distribution learning via graphical models

Graph encodes statistical dependencies between system variables

graph over random variables ↔ family of distributions consistent with combinatorial structure

~ absence of edges as independence relationships

$X_{I_1} \perp\!\!\!\perp X_{I_2} \mid \perp X_{I'}$ whenever I' separates I_1, I_2

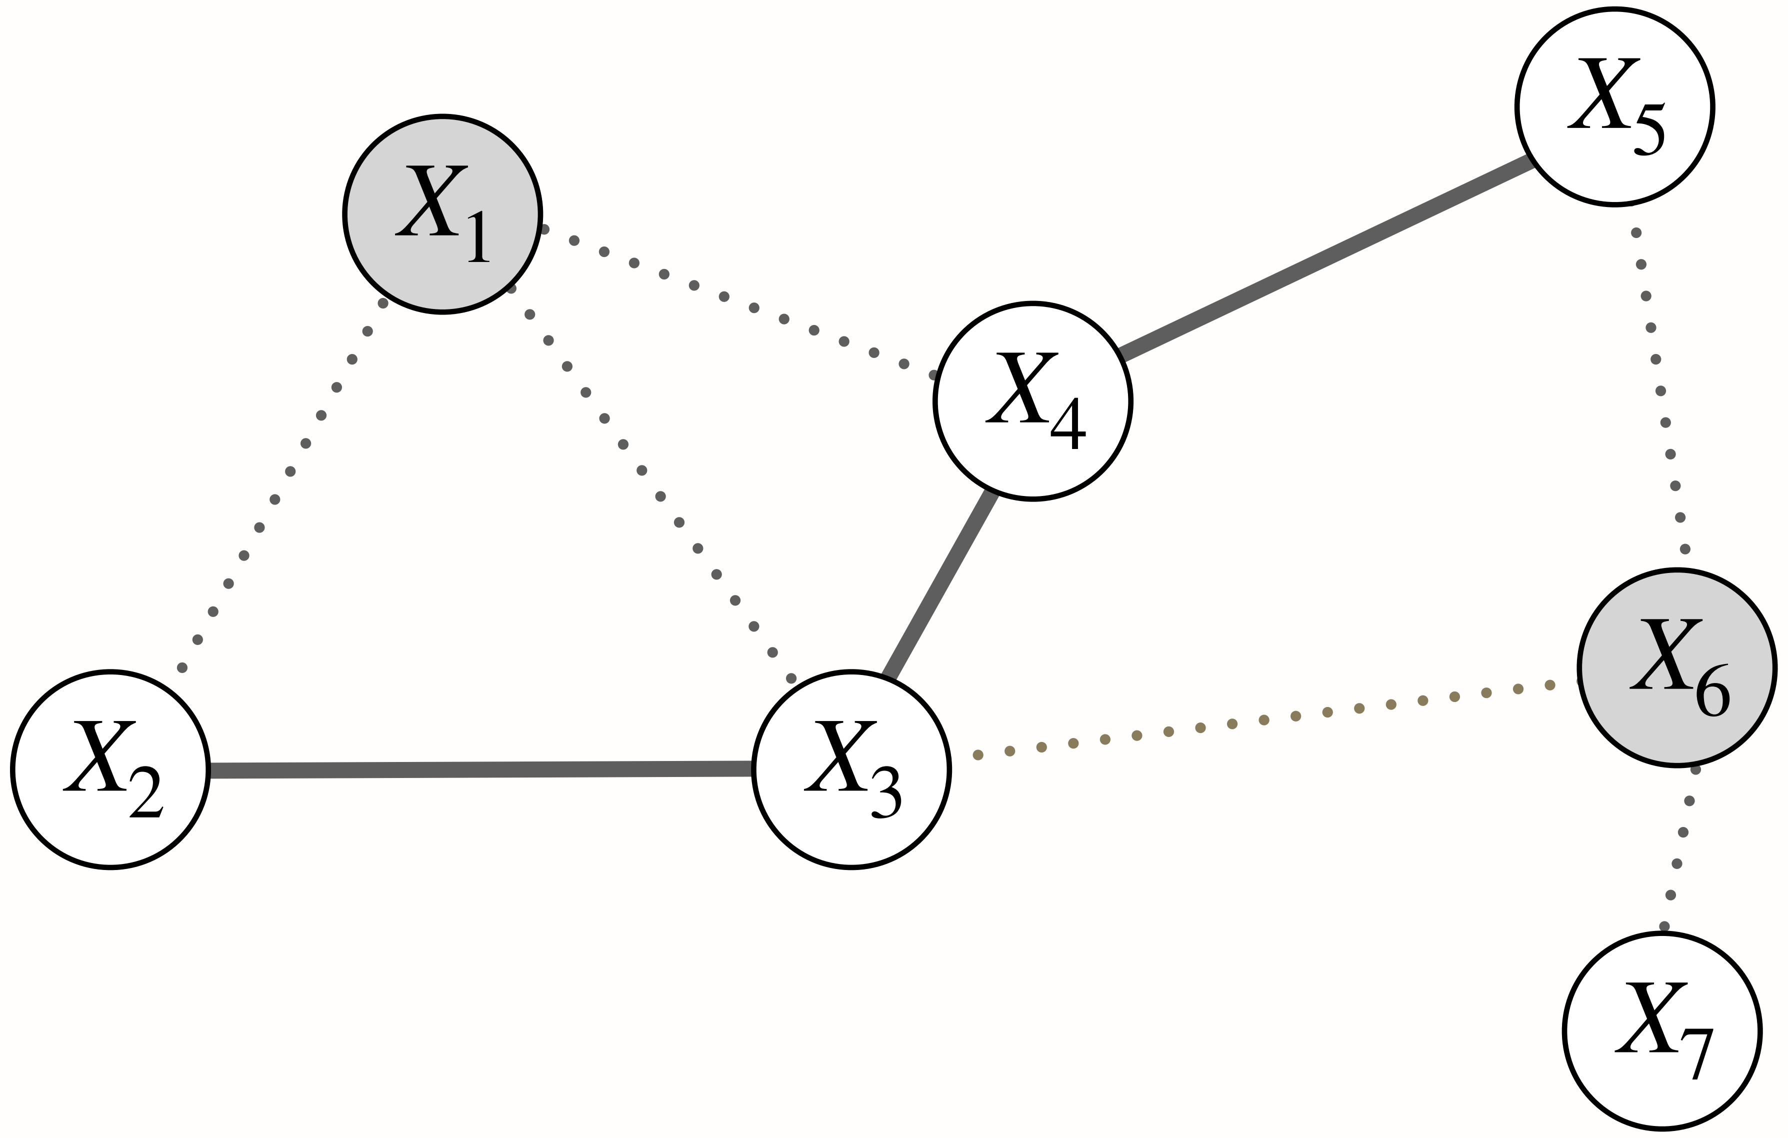

↑ edges: ↓ dependency assumptions, but ↑ parameters to learn

Reliable Inference at Scale Using Graph Structure

Distributed Systems

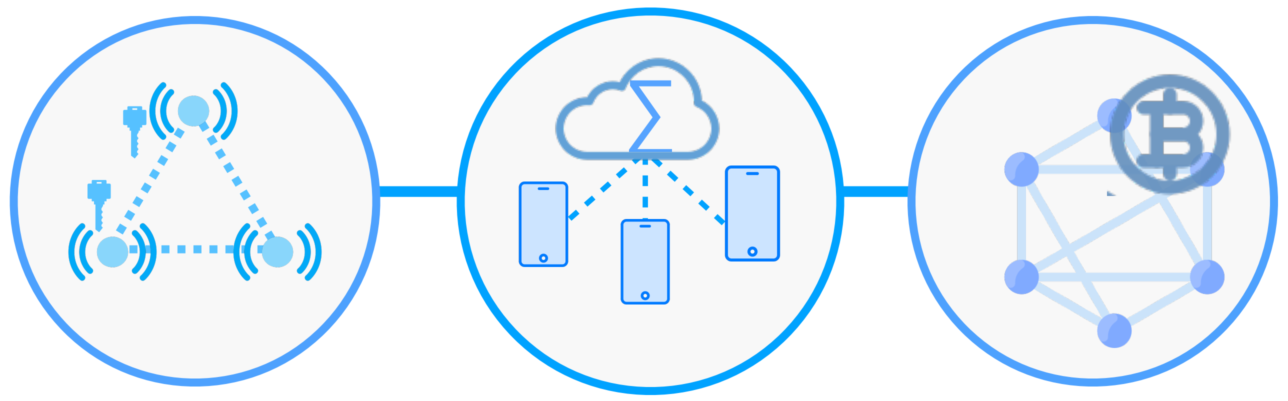

'secure' communication channel
between devices

Probabilistic Graphical Models

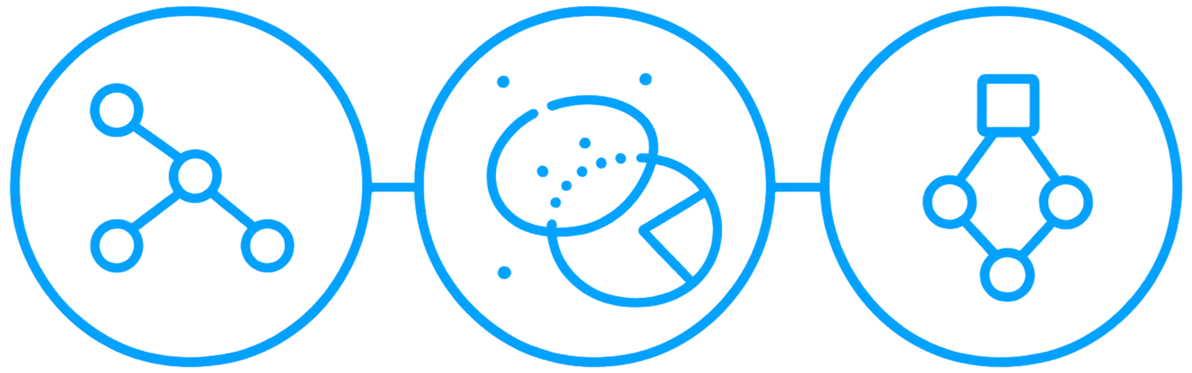

conditional dependencies
in observational data

computation

heterogeneity

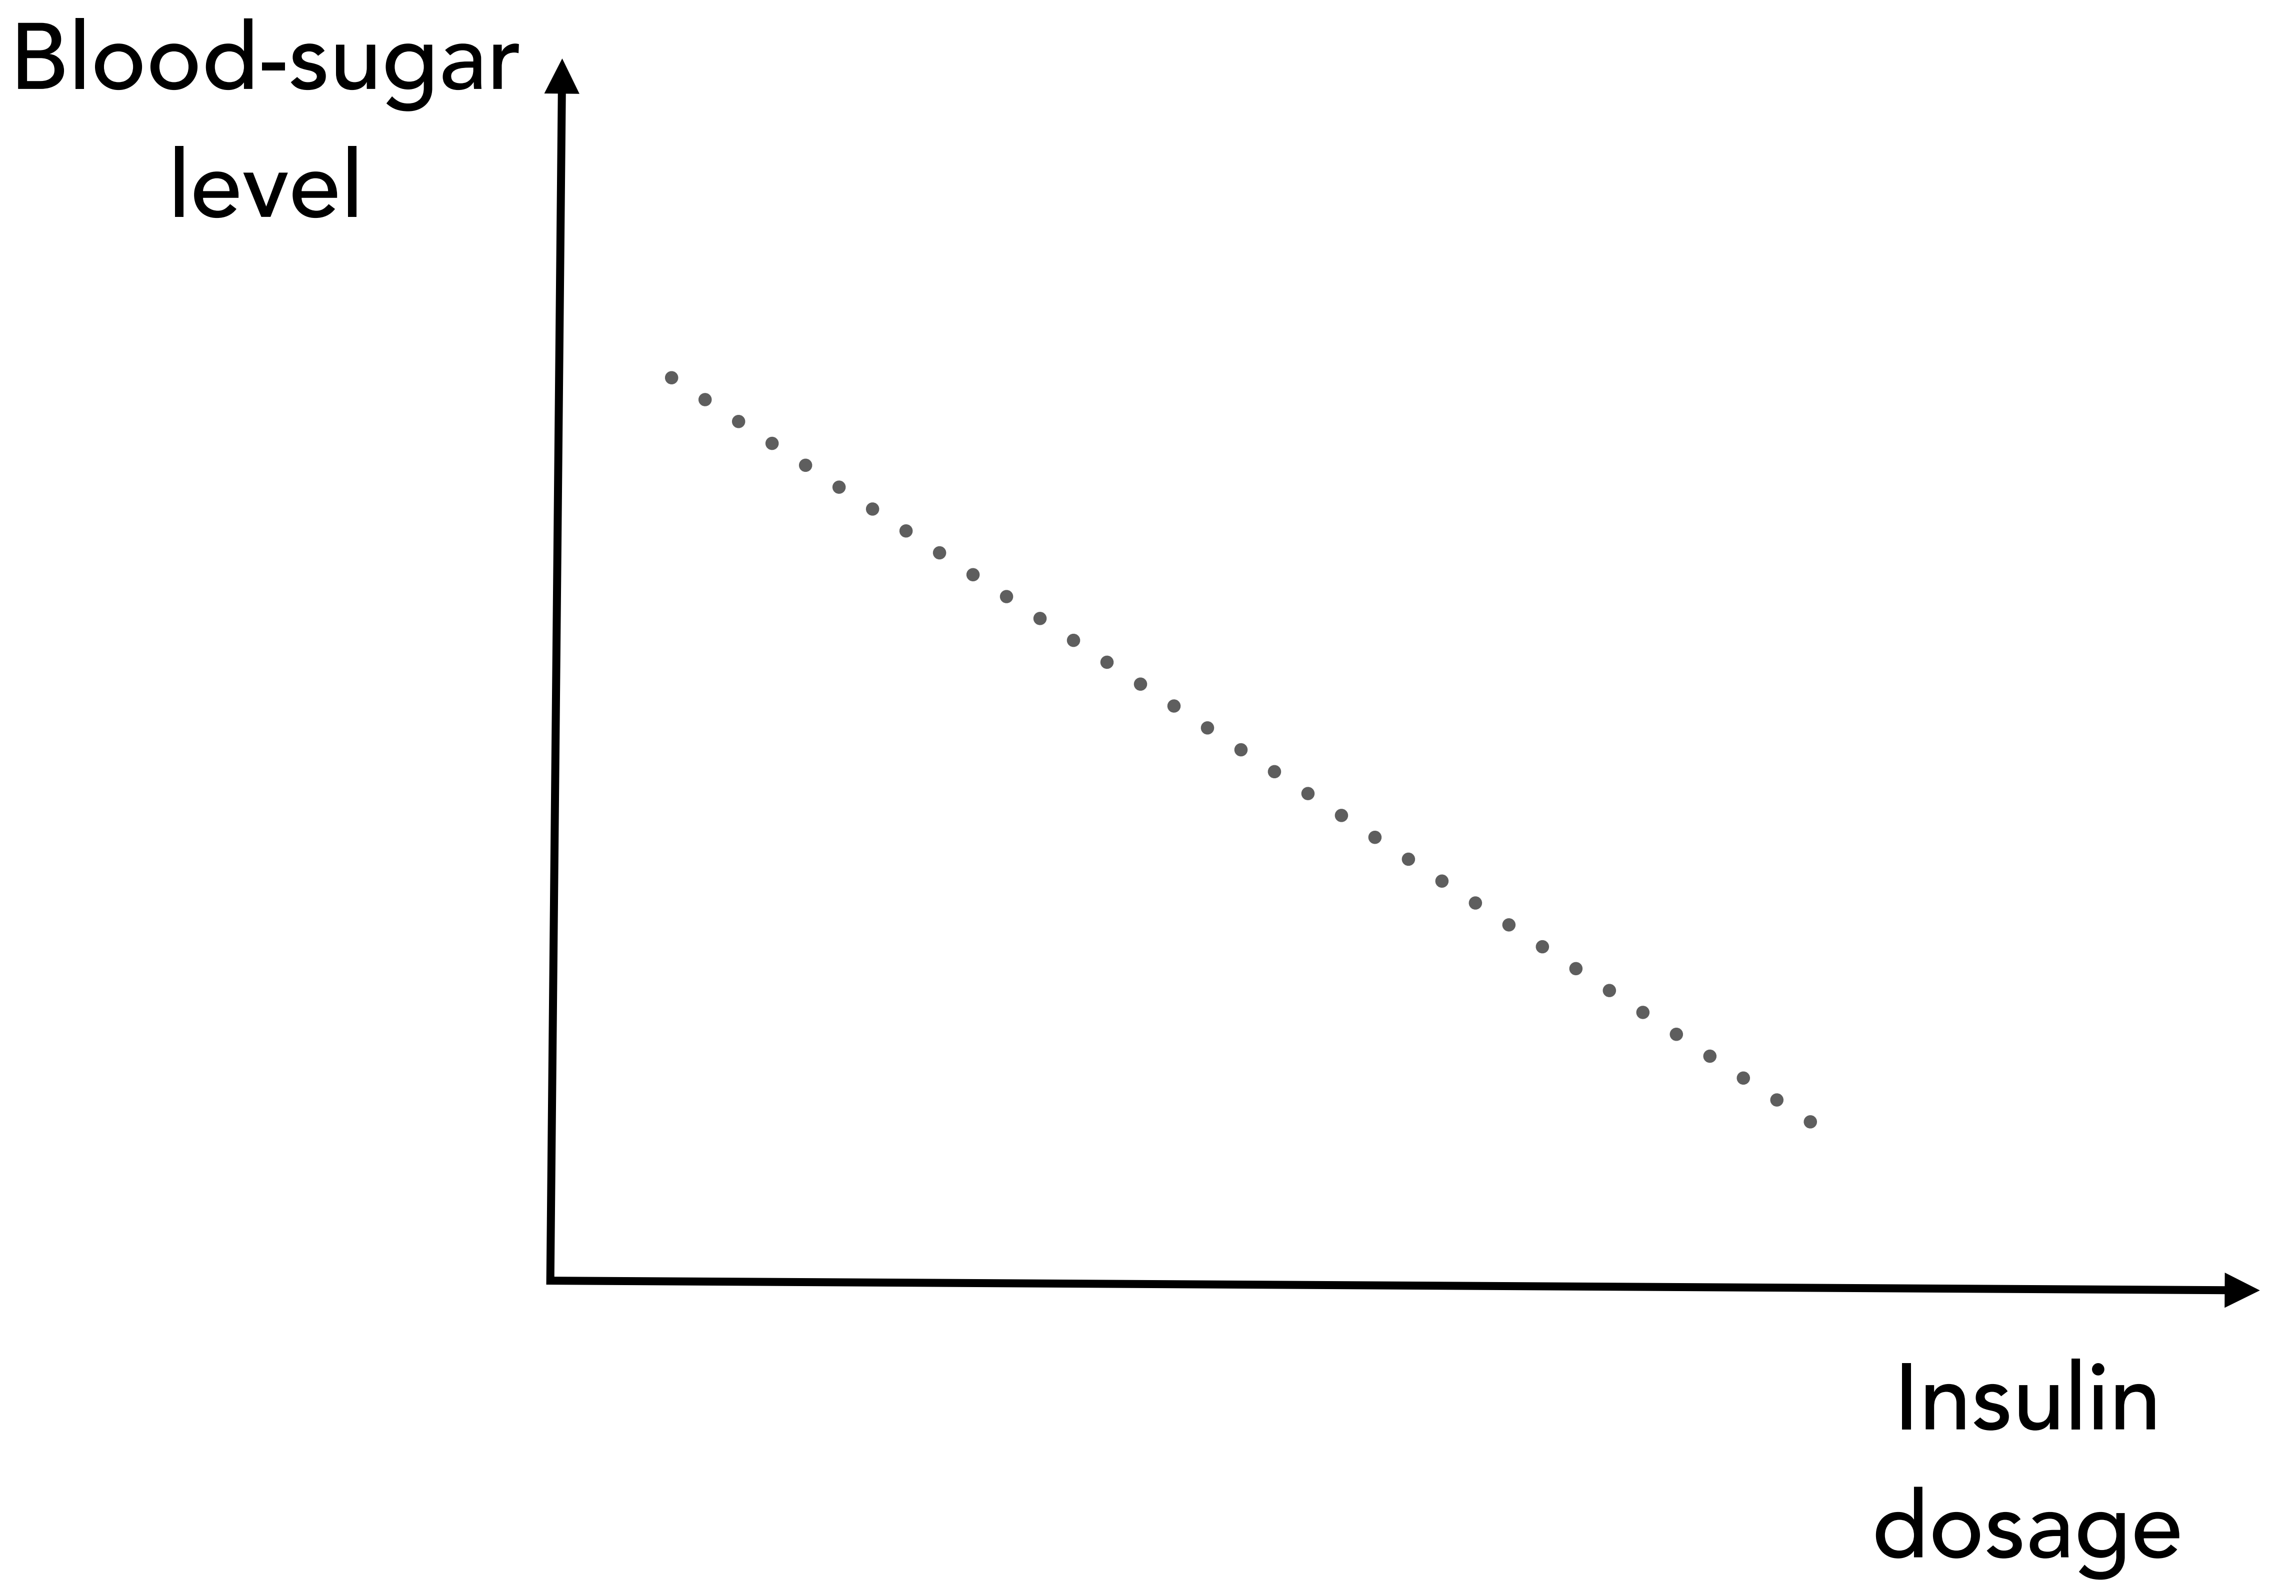
Blood-sugar
level
Insulin
dosage


Blood-sugar
level
Insulin
dosage

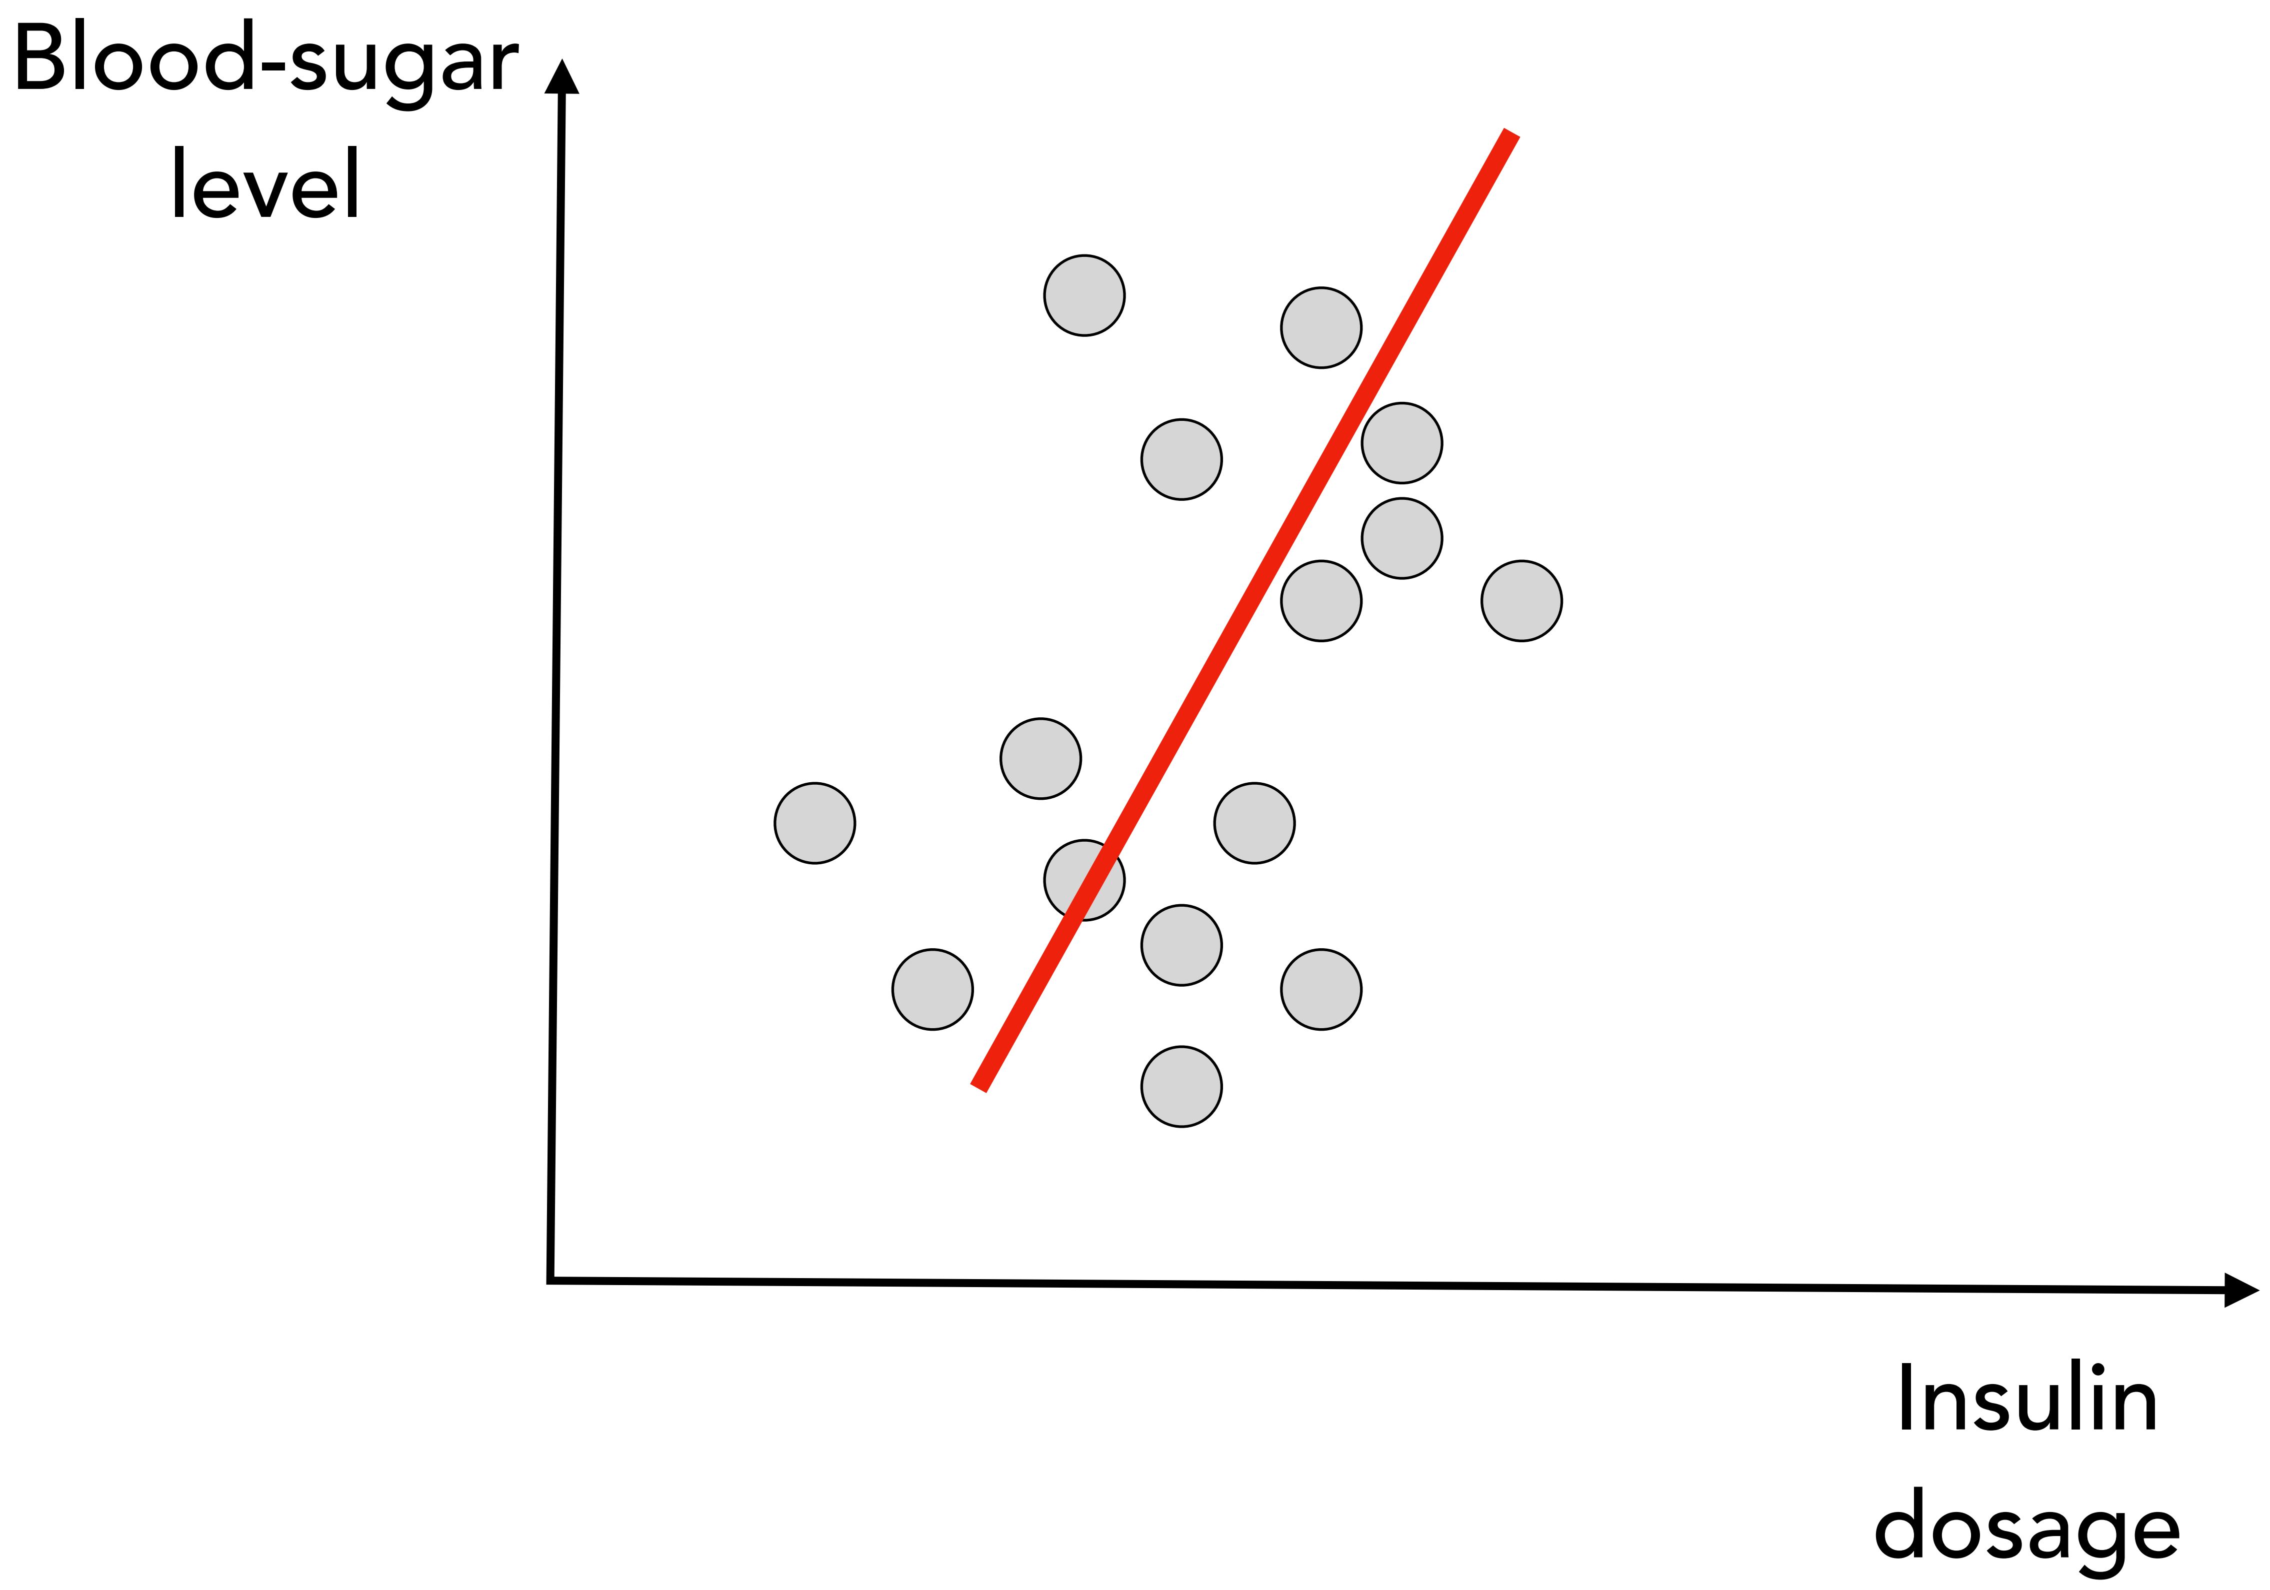
Blood-sugar level
Insulin dosage

Simpson's Paradox

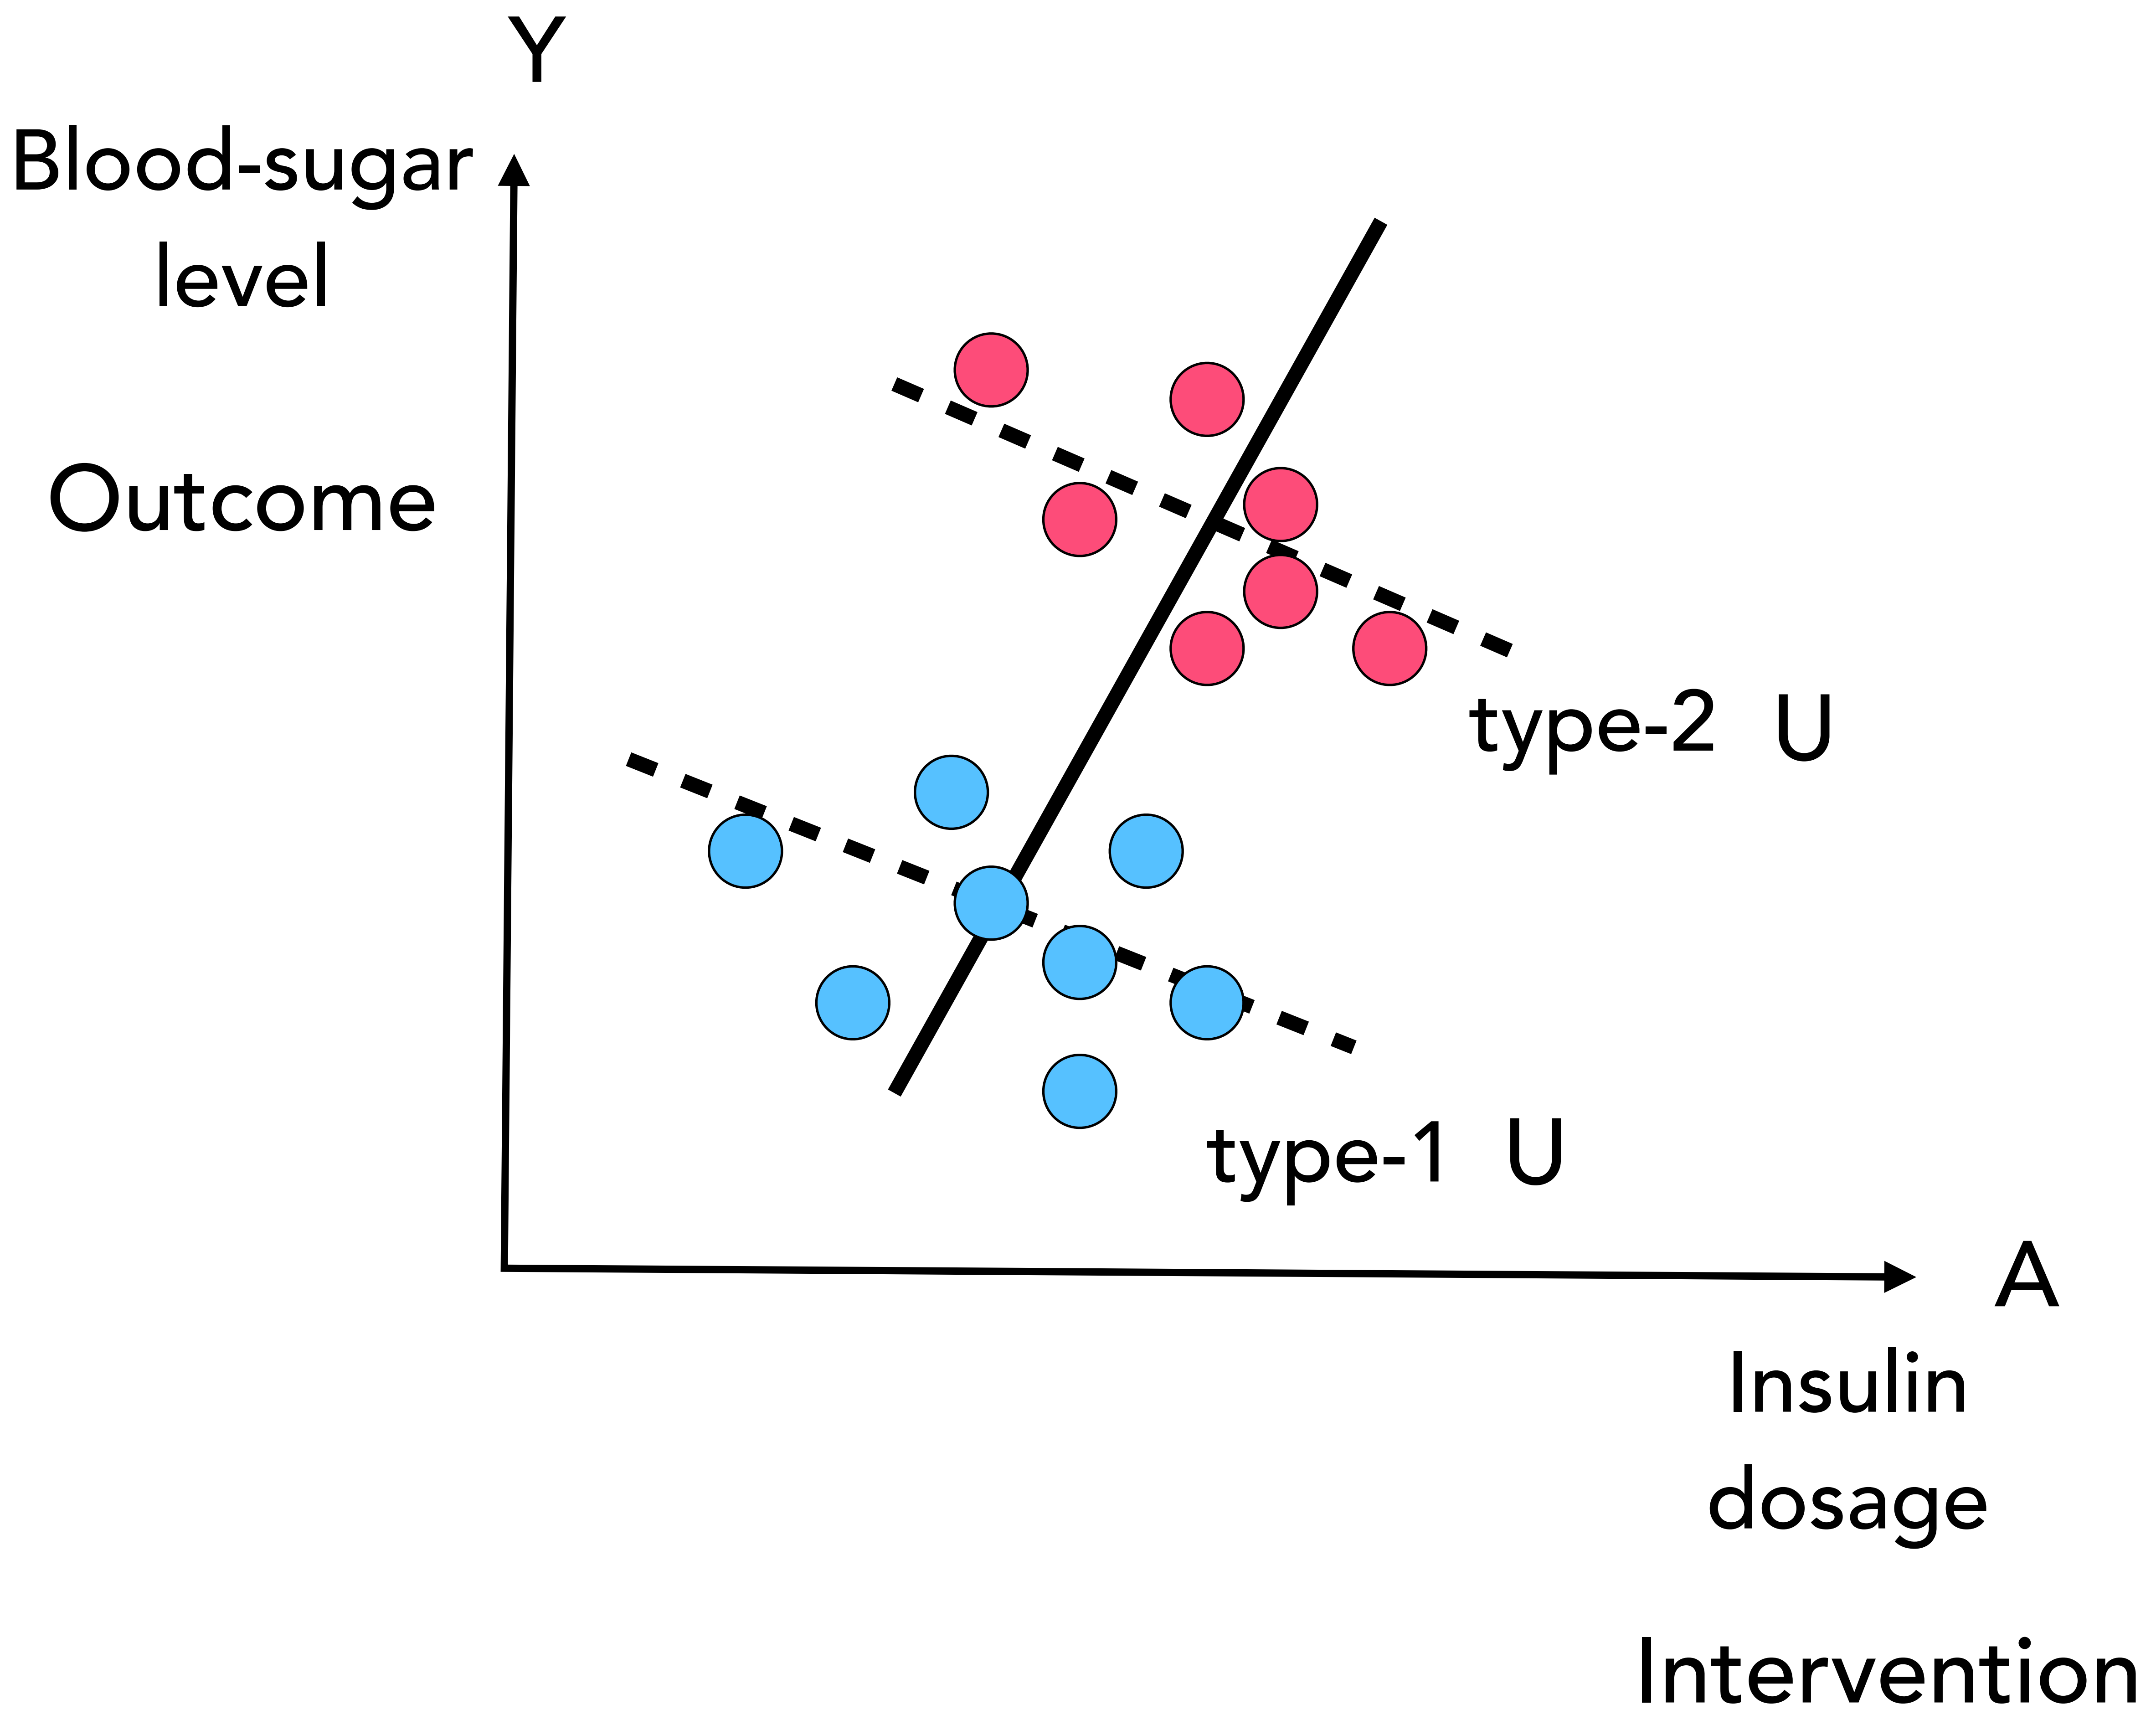

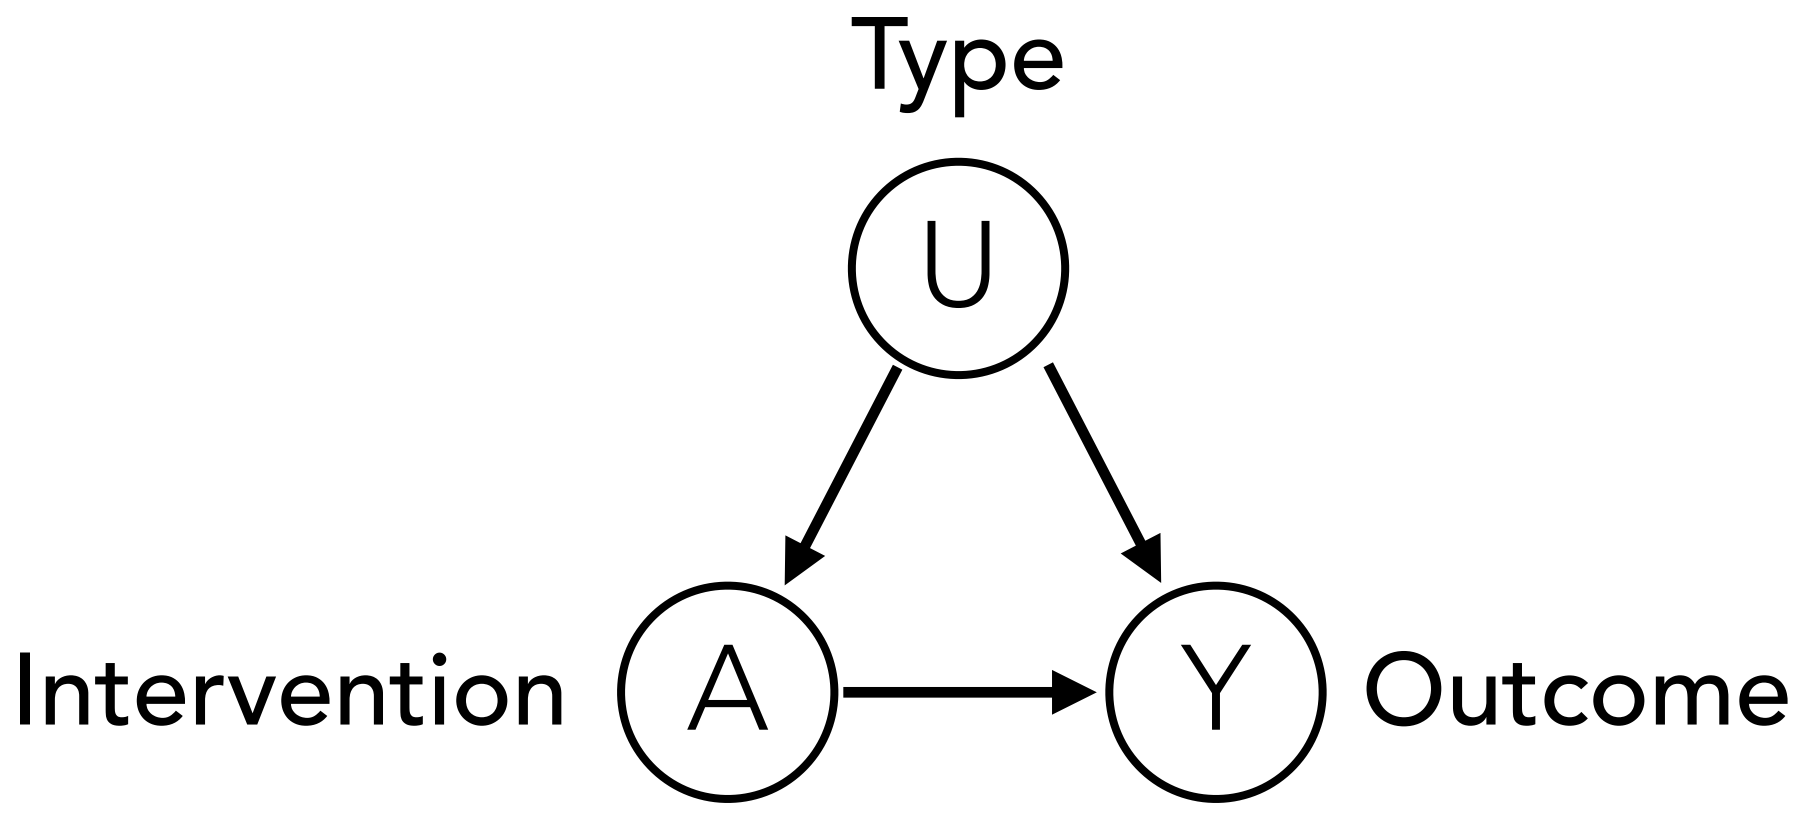

Simpson's Paradox

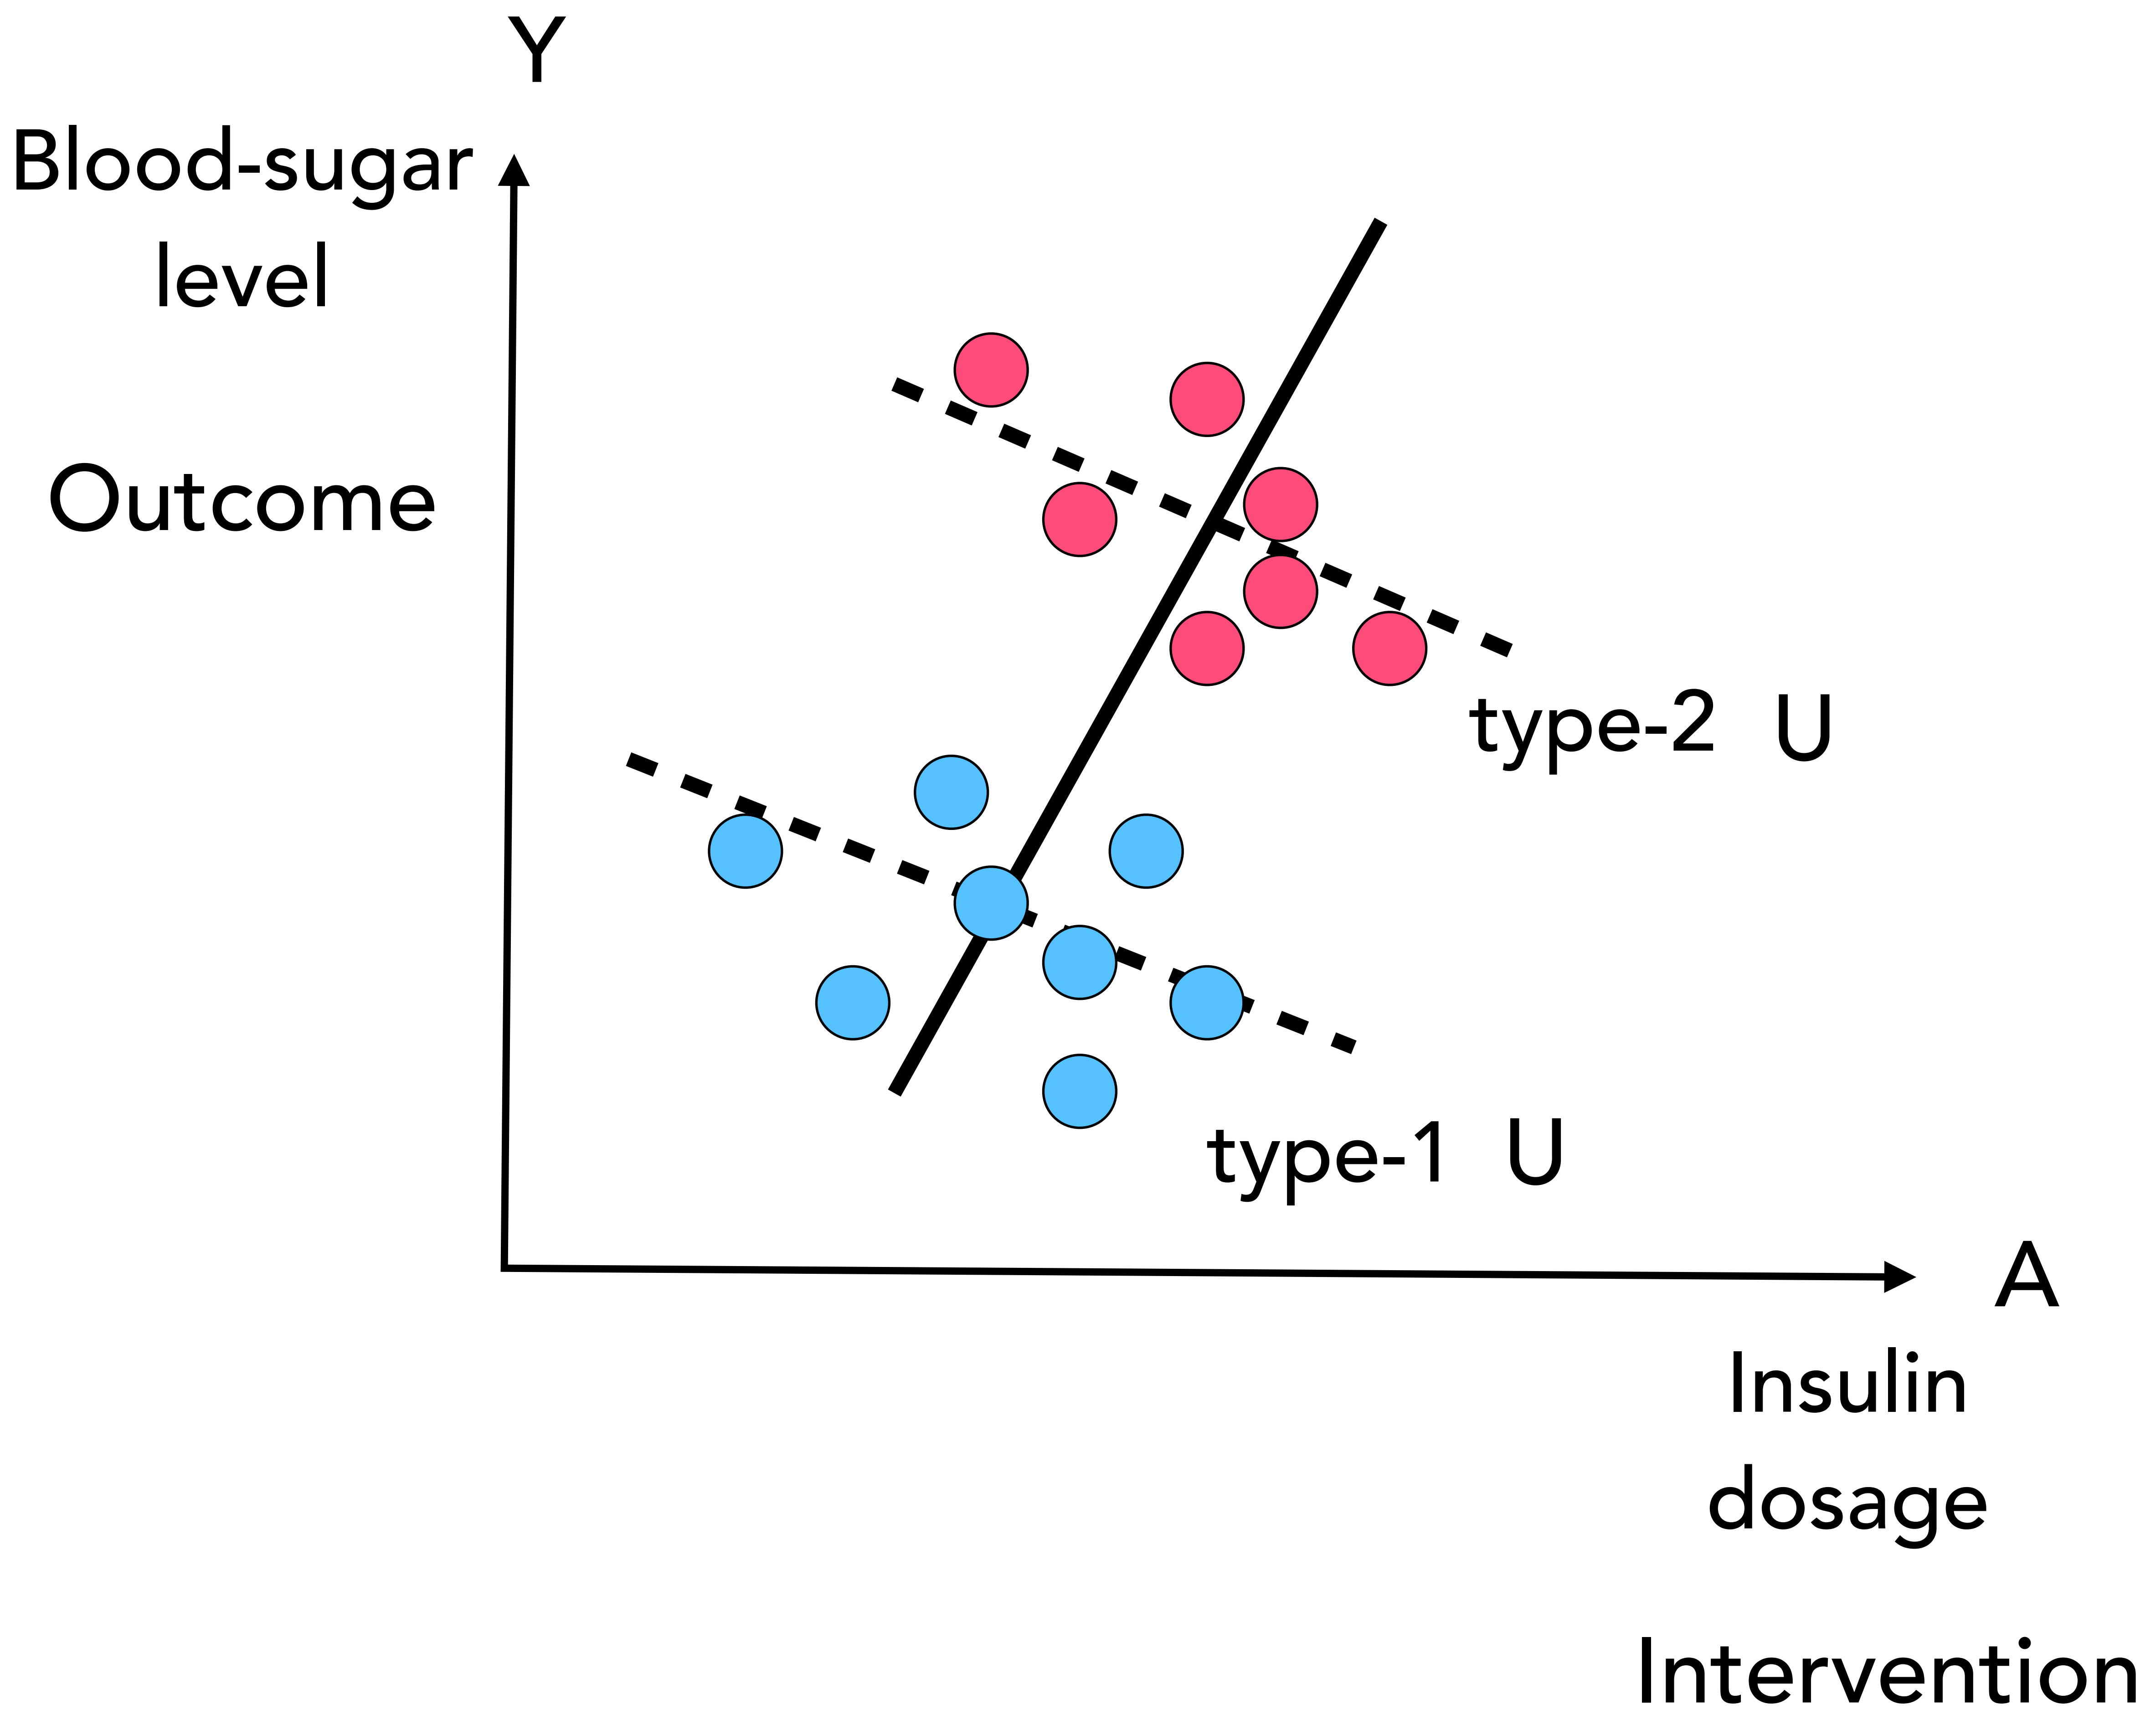

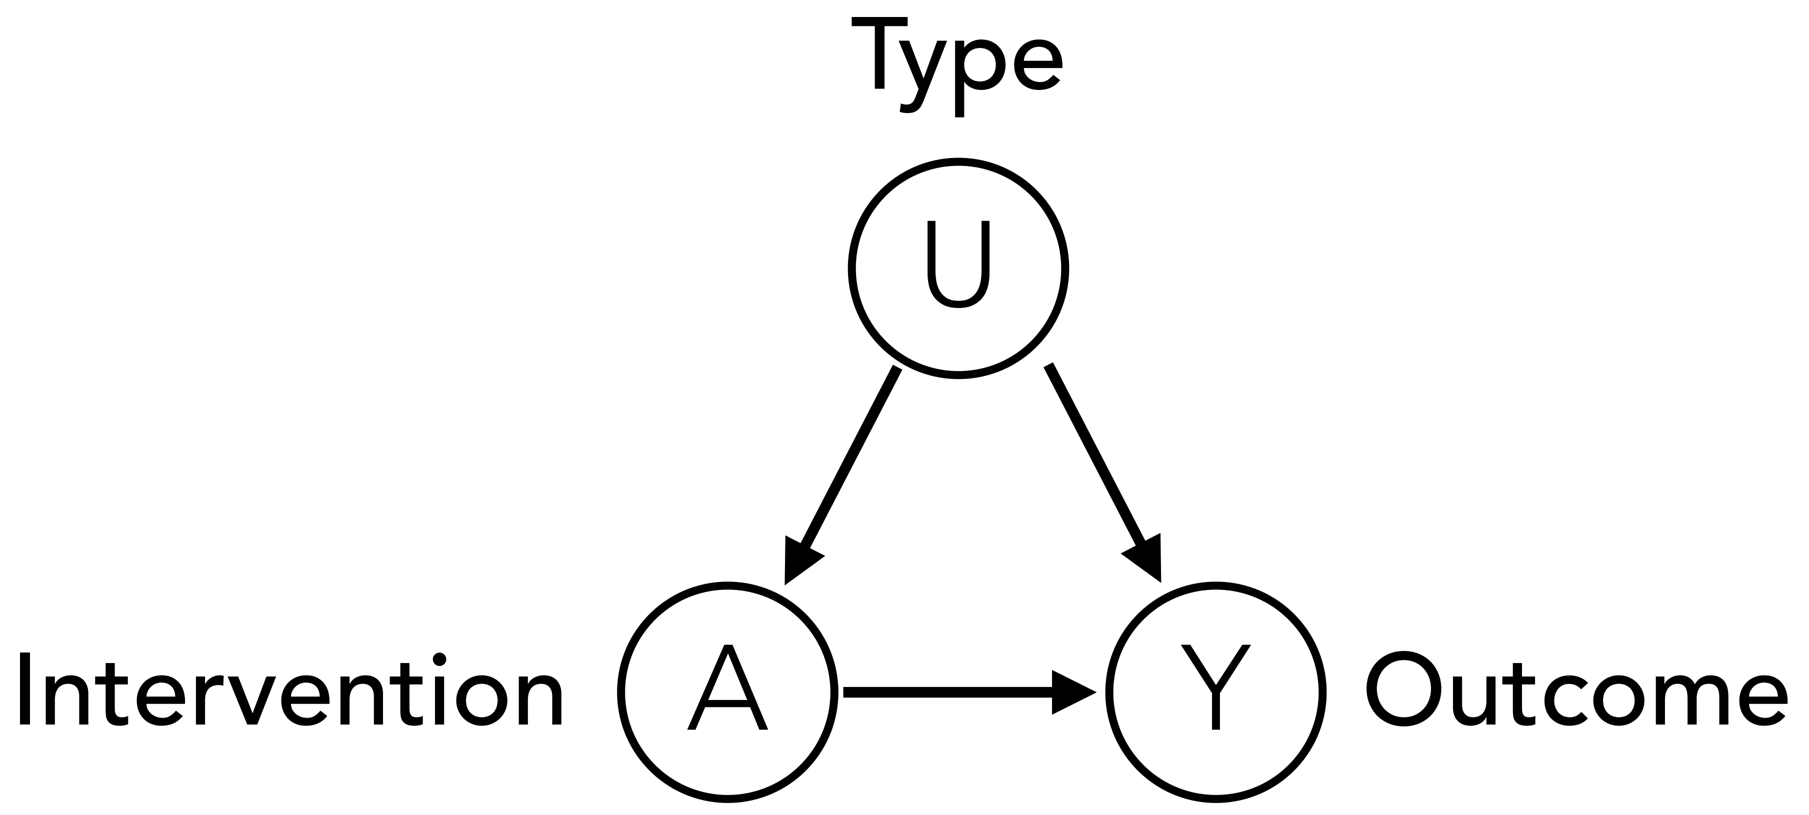

Given samples $(a^{(k)}, y^{(k)})_{k=1}^{n}$ from observational distribution $p(y, a)$,

Learn interventional distribution $p(y \mid \mathrm{do(a)})$ which factorizes as

$$p(y \mid a, u = 1) \cdot p(u = 1) + p(y \mid a, u = 2) \cdot p(u = 2)$$

component densities | mixing density

Our observational world is full of mixtures…

- Cells as vectors of gene expressions (d~20,000)
- **Cell-types** described via mixing weights for base distributions

How to go from one cell-type (mixing distribution) to another?

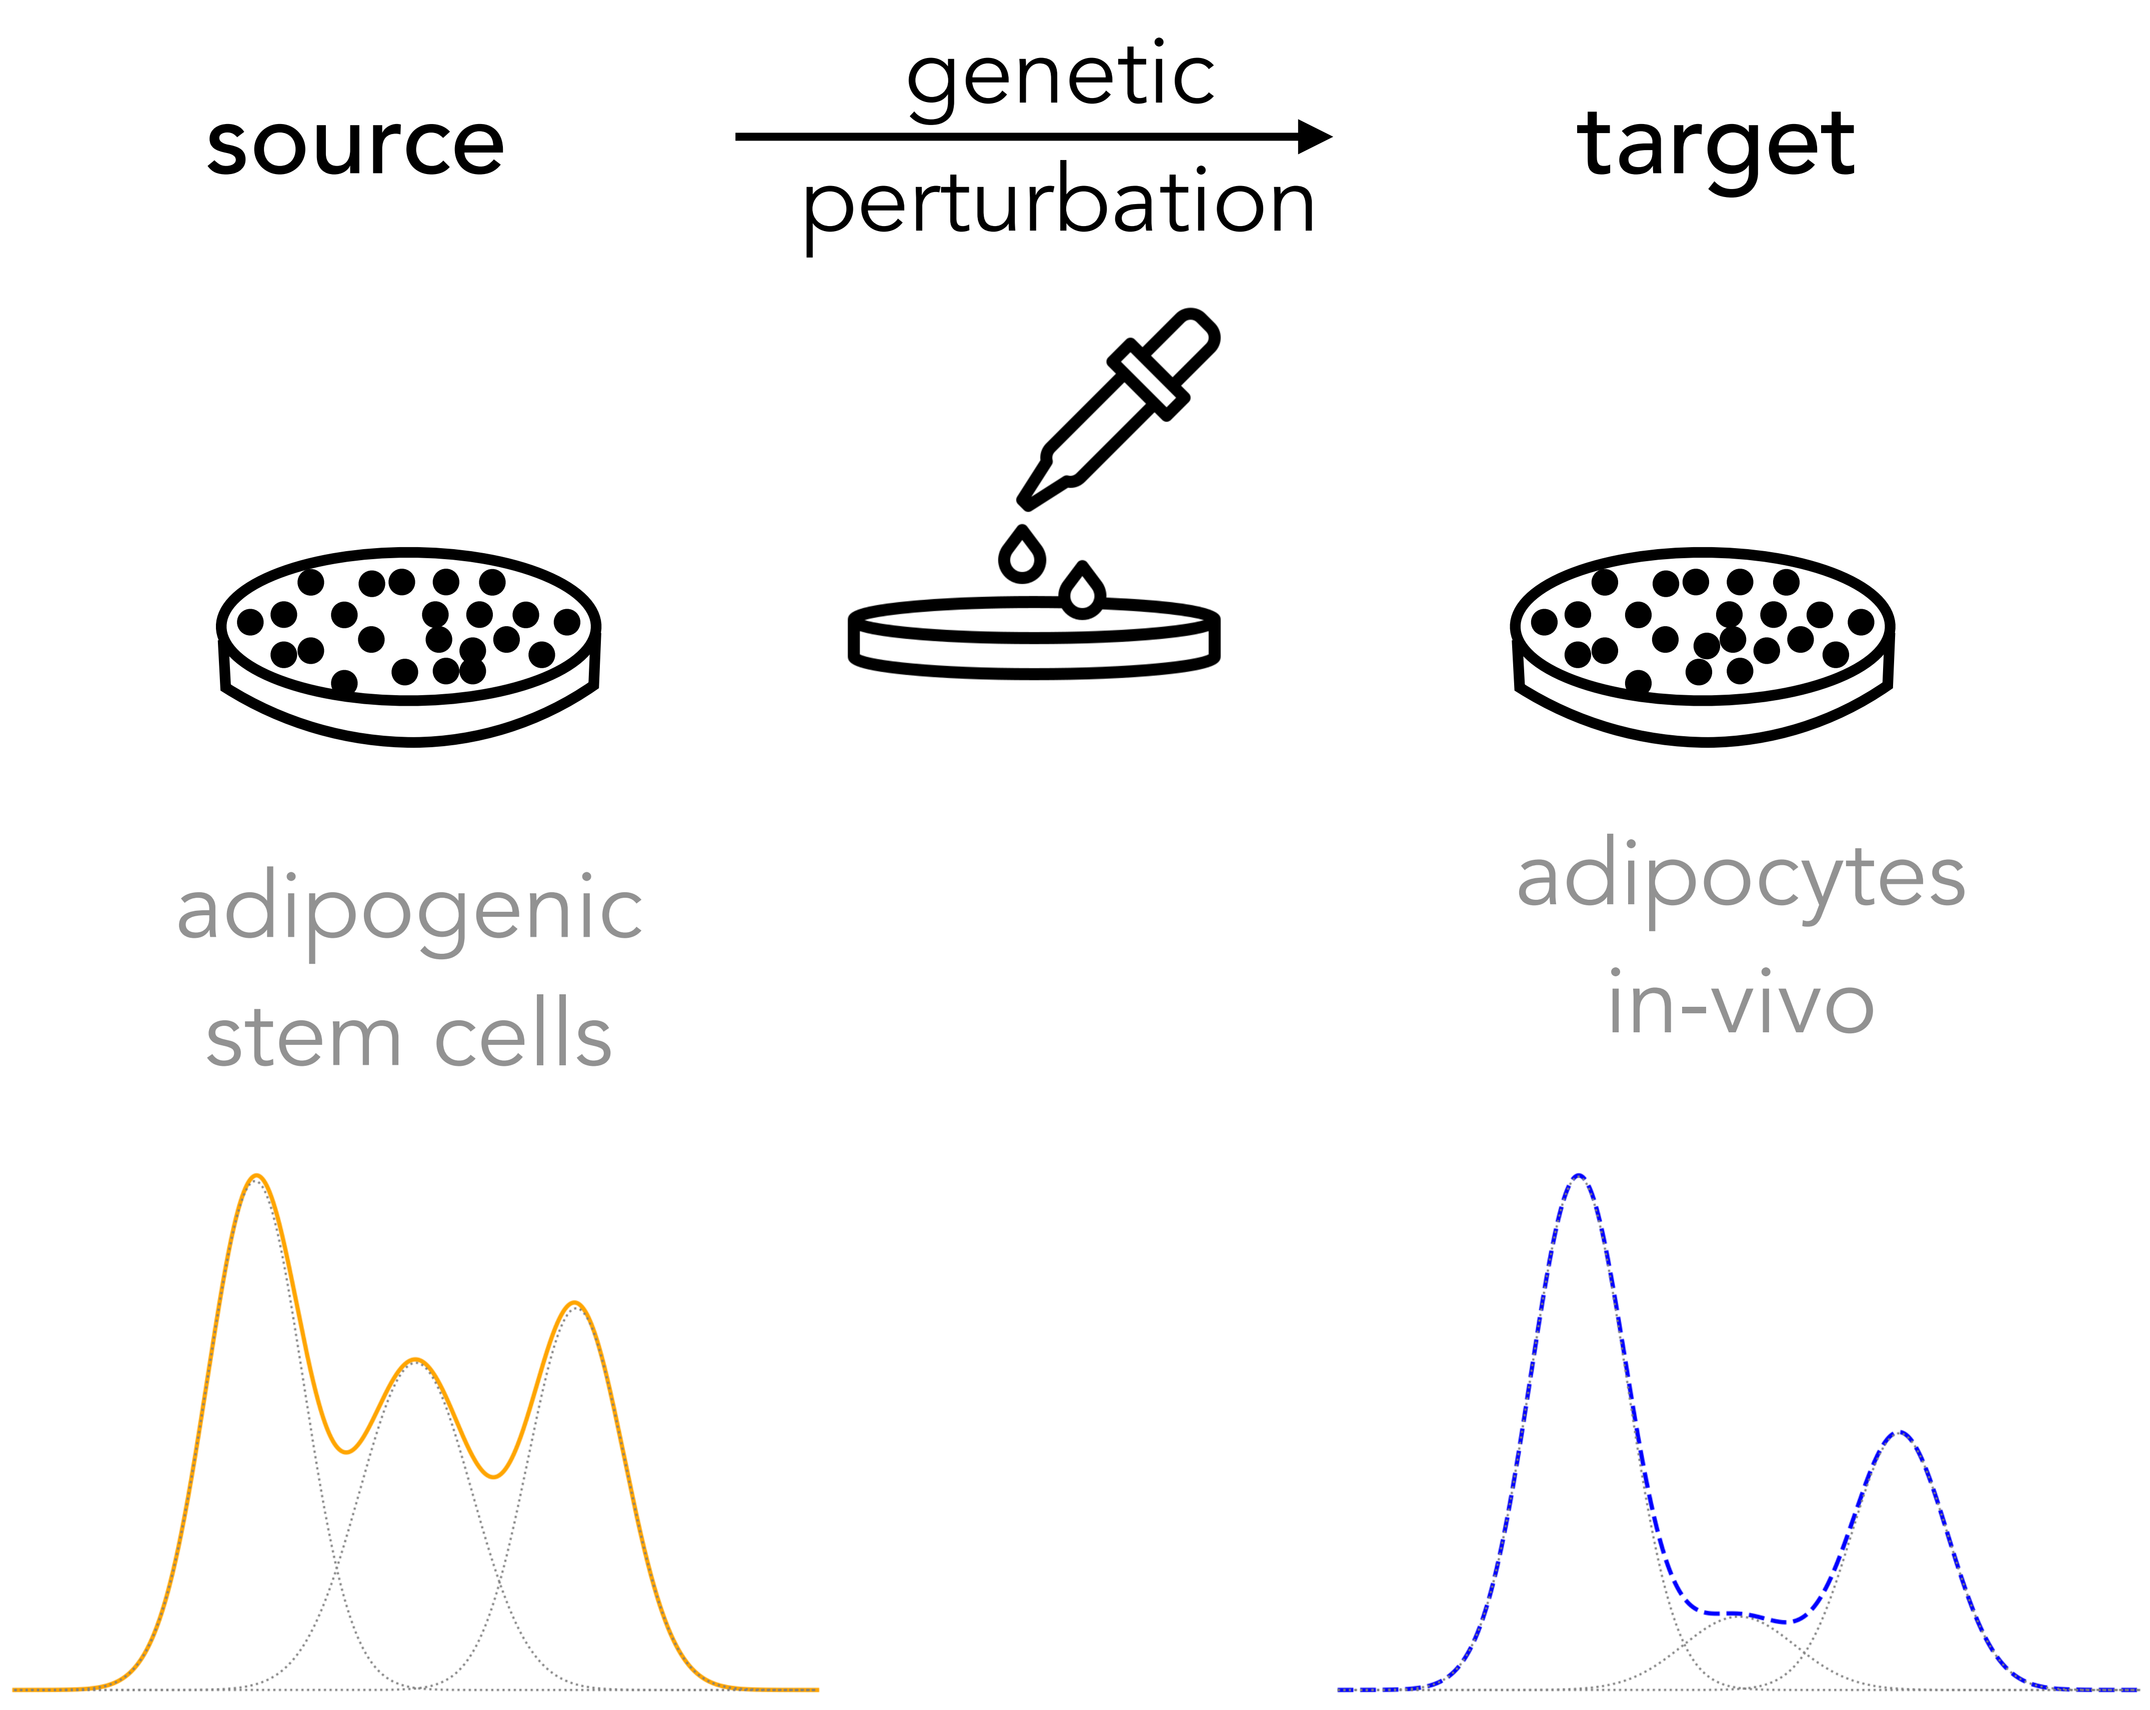

Source and target cell-types as mixtures of few base distributions

Heterogeneity in observational data

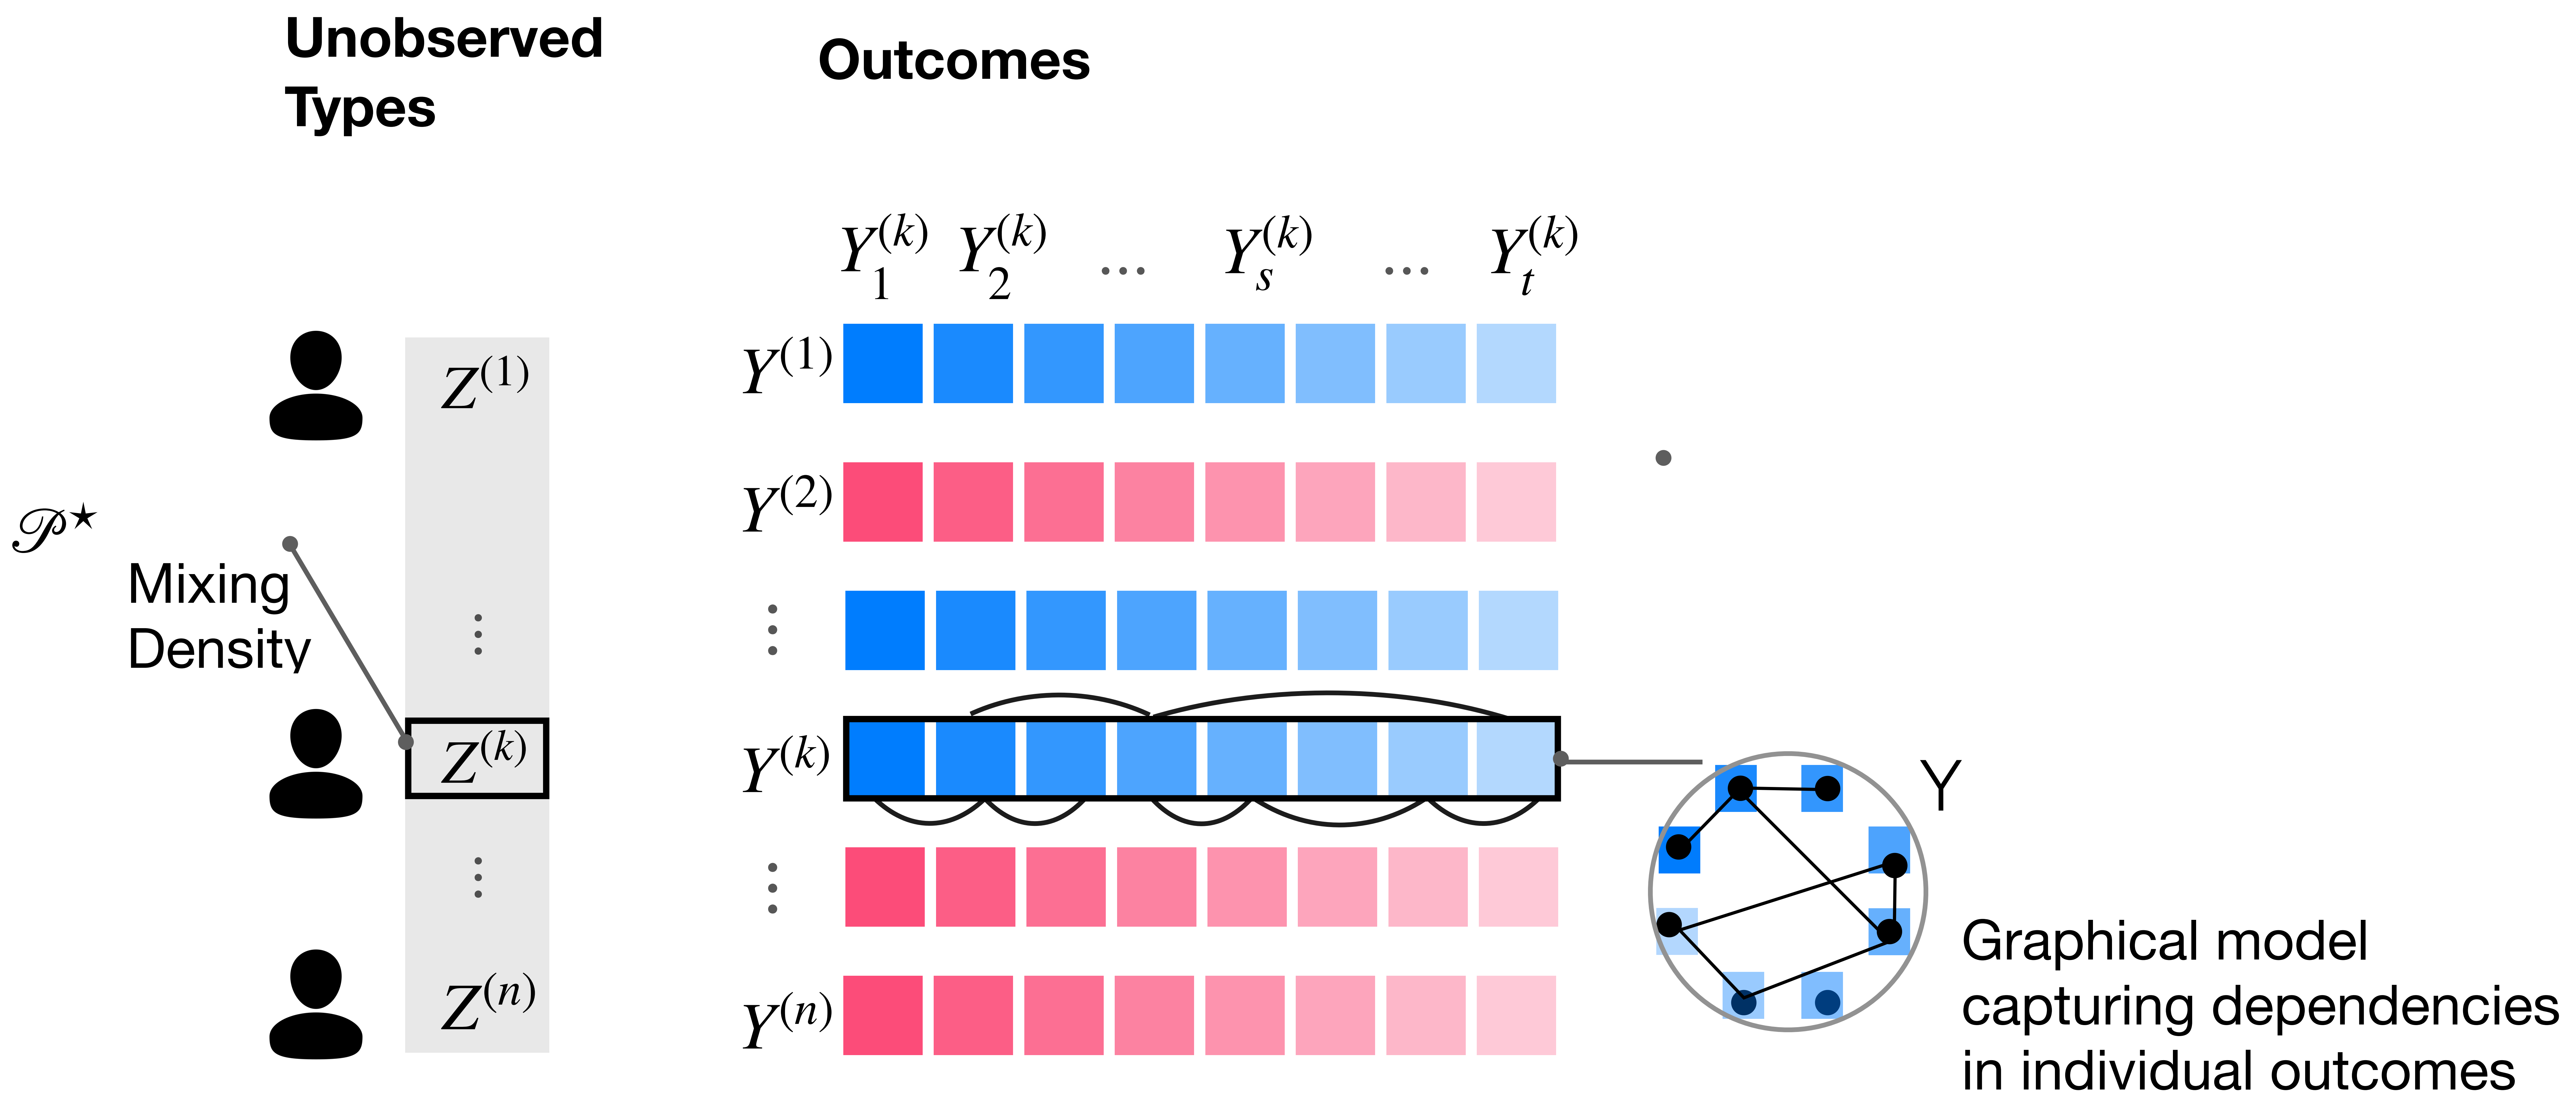

Formalizing mixtures

(Log-likelihood)

$$p(y) = \sum_{\theta} p(y; \theta) \; p(\theta)$$

mixture components (typically parametric)

mixing distribution

$$\sum_{k \in [n]} \log \left(\sum_{\theta} p(y^{(k)}; \theta) \;\; p(\theta) \right)$$

- log sum computation
- non convex objective

mixtures x causal

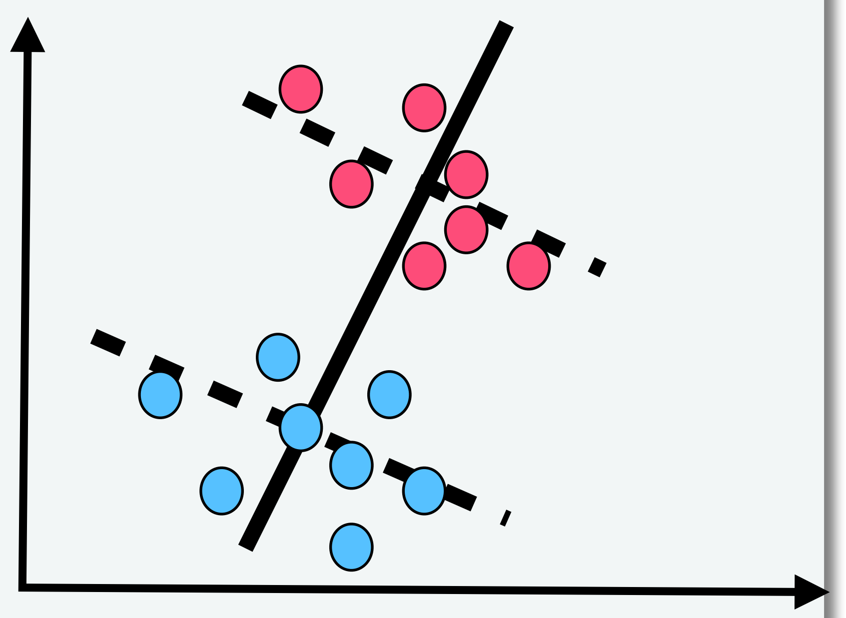

high-dimensions

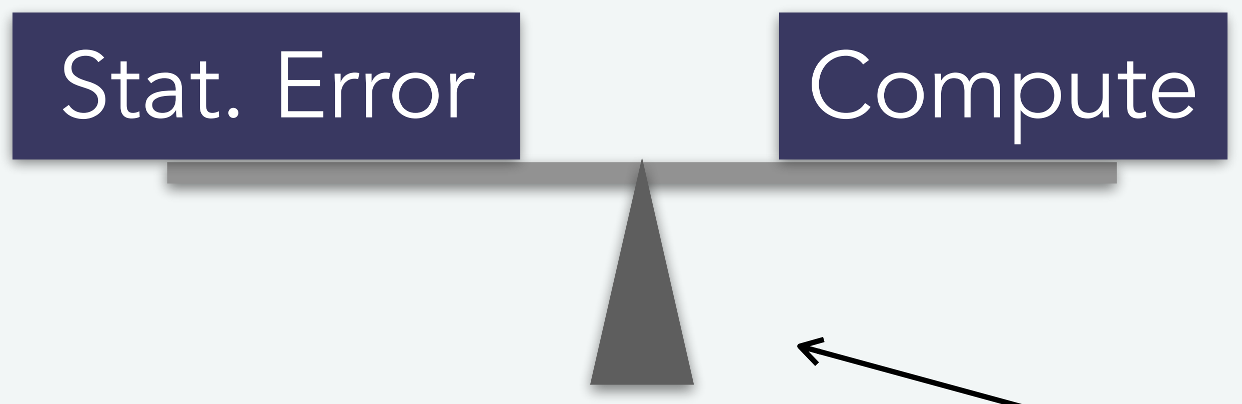

convex loss
exp family geometry

dependencies

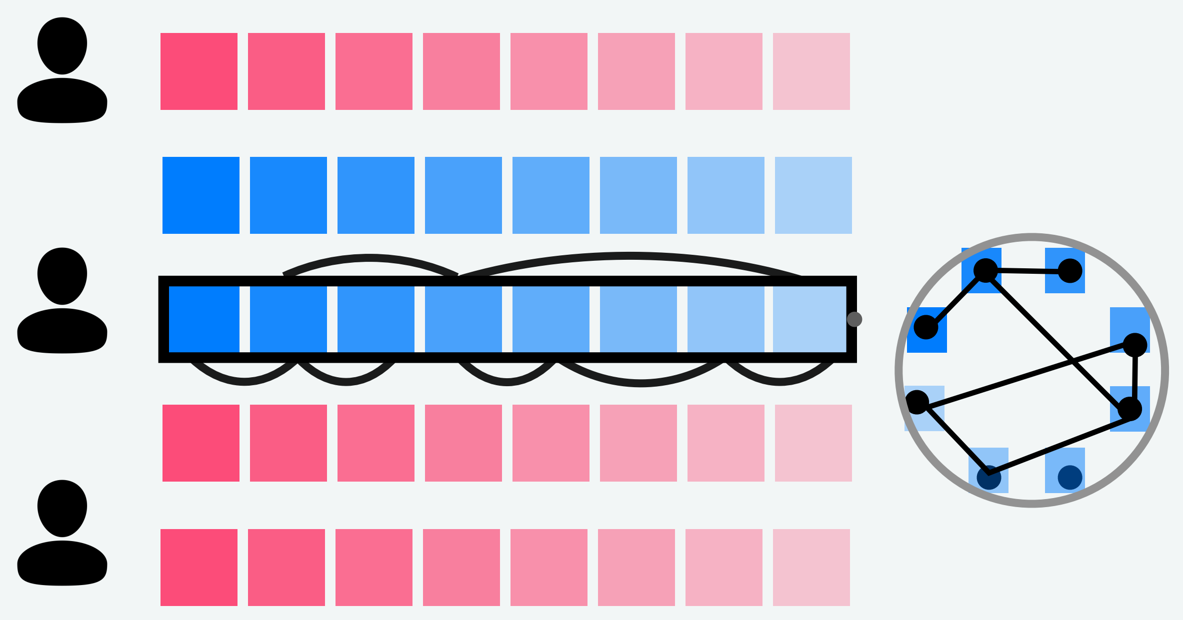

Mixtures of
Exponential family

M.S. D.S., to appear at **ISIT** ‘26

A new lens for learning exponential family mixtures

convex loss minimization → clustering

(error rates ↔ metric entropy of parameter space)

Summing: Reliable Inference at Scale Using Graph Structure

notions of connectivity

Theme: leverage structure (encoded as graphs) to improve efficiency

implications of **heterogeneity** for large-scale inference,

Focus: algorithms appearing as canonical subroutines in inference pipelines

securely aggregating information over a network of devices

'secure' communication channel between devices

communication

Learning parameters in graphical models with latent structure

conditional dependencies in observational data

computation

Reasoning about our networked world…

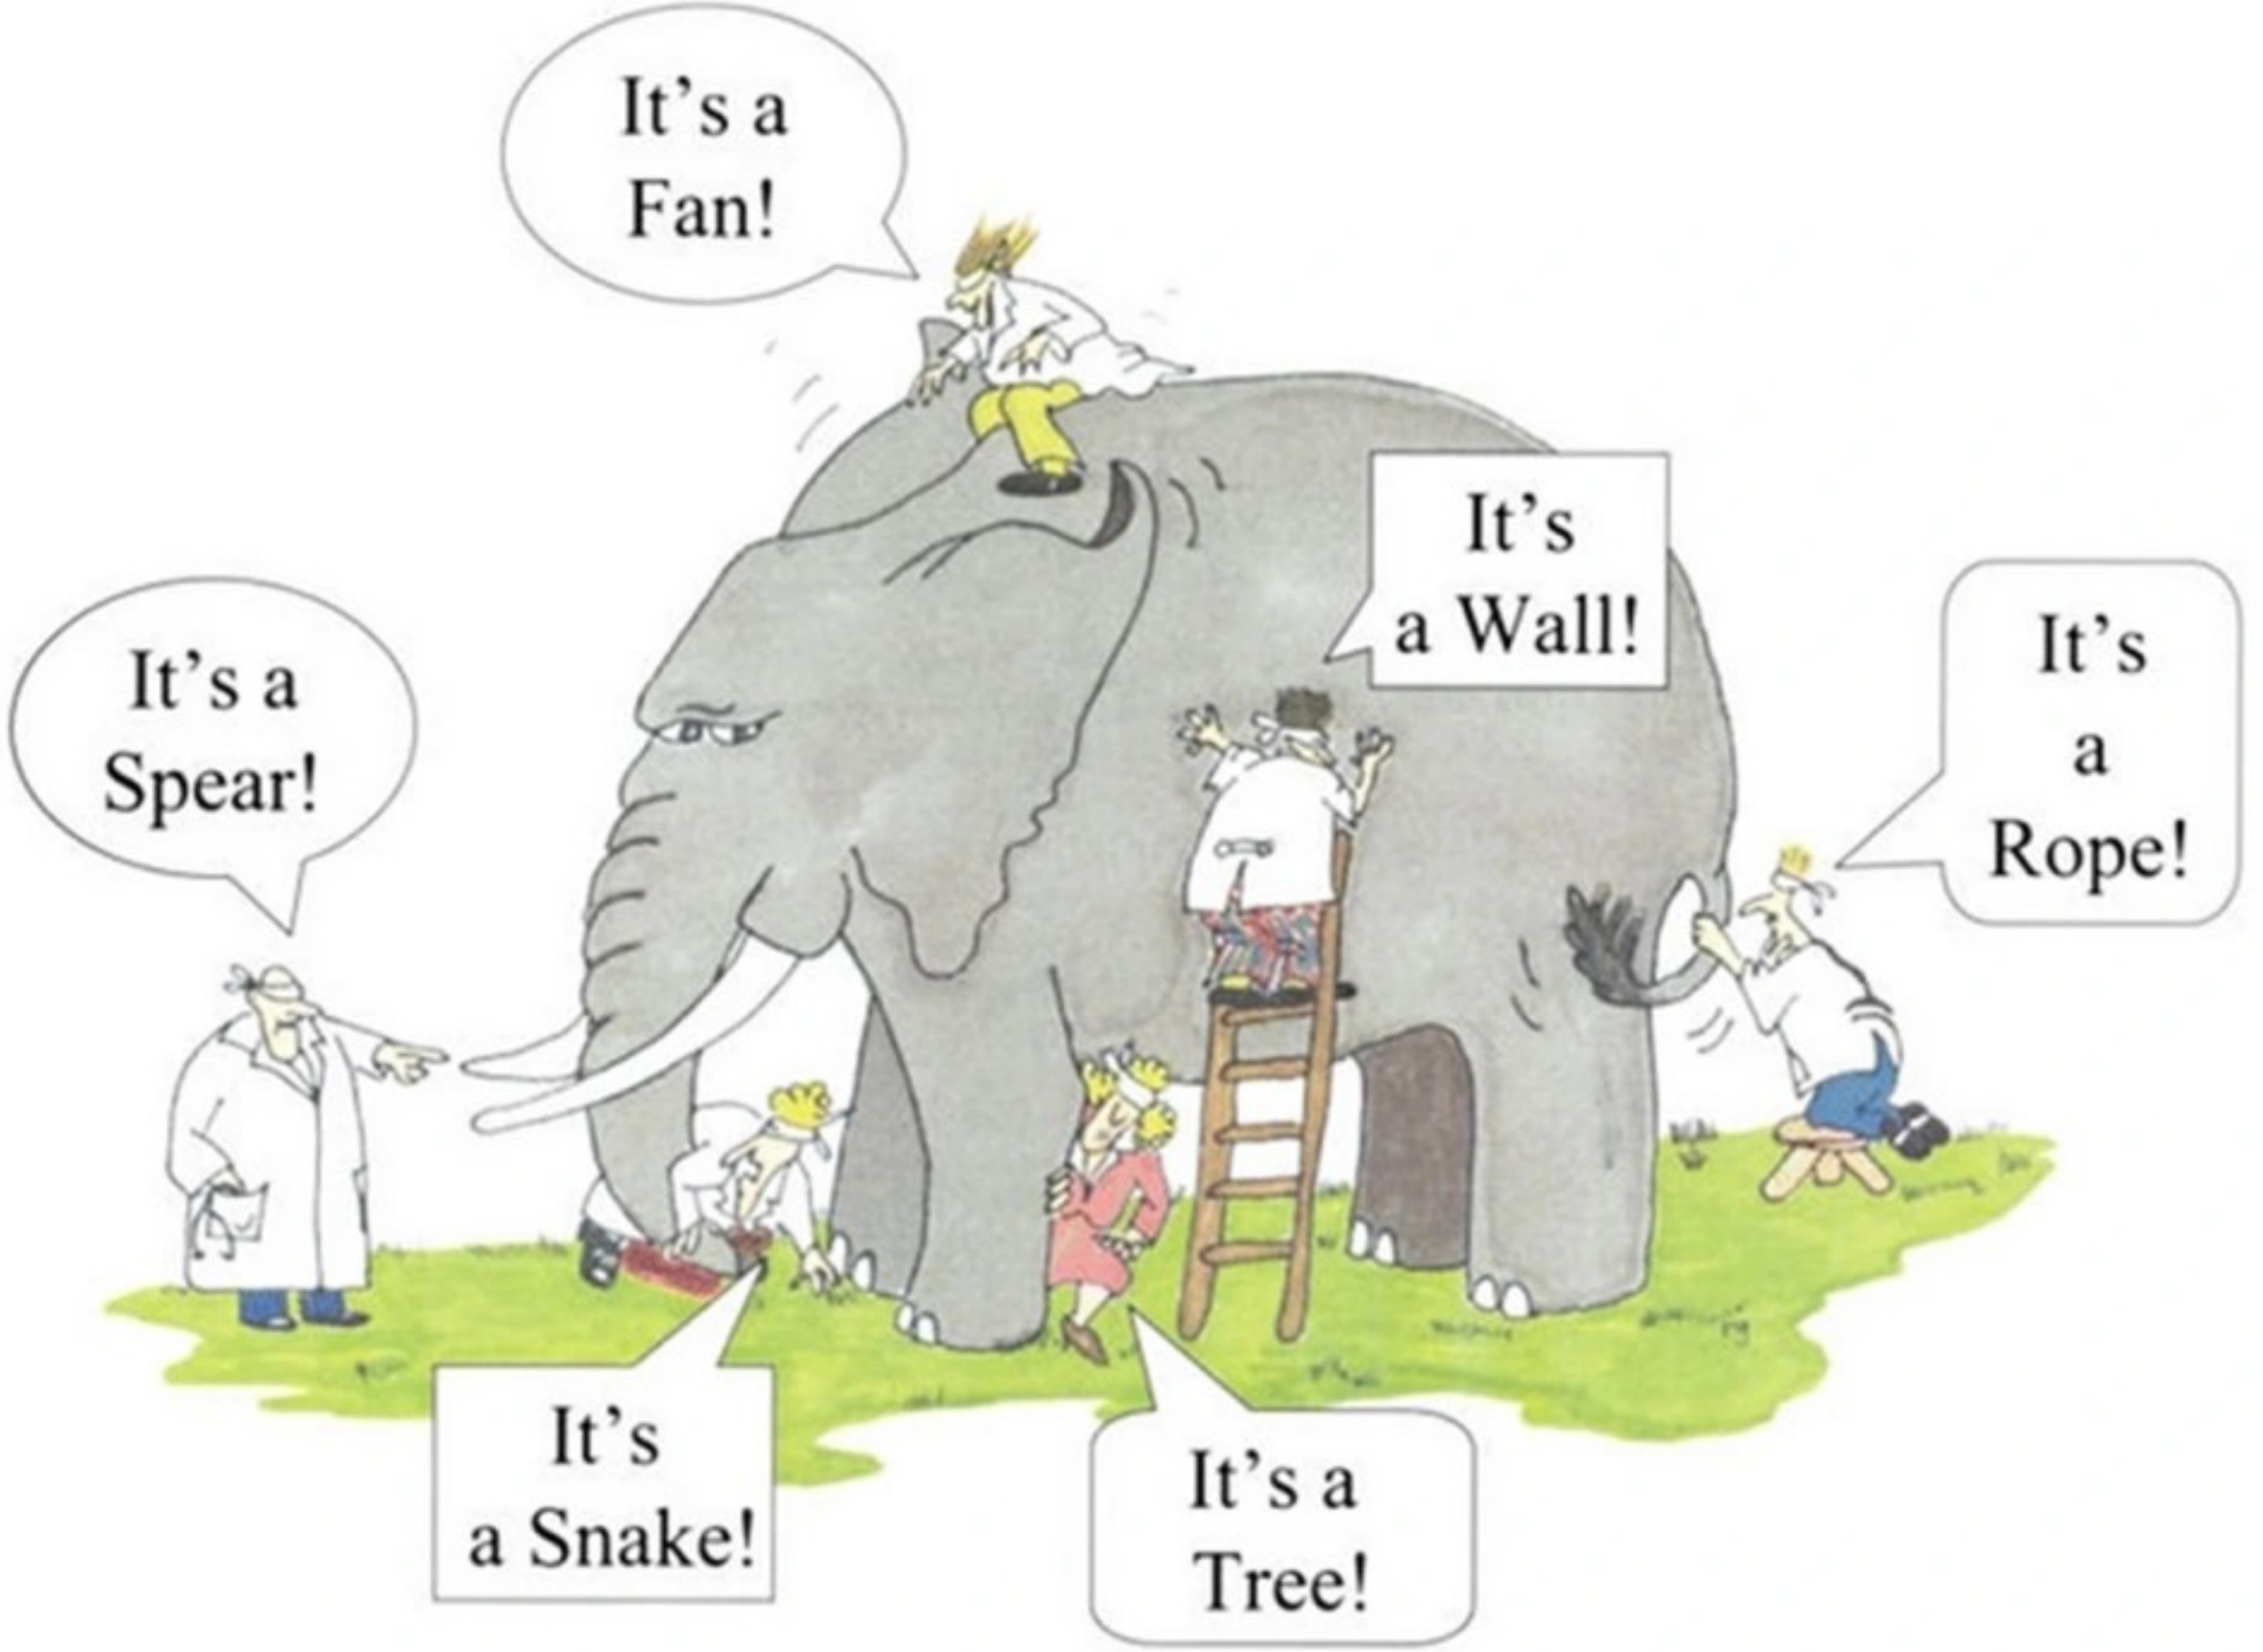

Source: https://speakingofresearch.com/2017/03/02/understanding-the-animal-not-just-its-parts/

Mansi Sood MIT

msood@mit.edu

https://arxiv.org/pdf/2508.11863

Acknowledgements: Schmidt Science Fellowship; Philip and Marsha Dowd Fellowship, Cylab Presidential Fellowship, David H. Barakat and LaVerne Owen-Barakat Fellowship, Lee-Stanziale Ohana Endowed Fellowship, Knight Fellowship, NSF CCF #1617934, Office of Naval Research (ONR) #N00014-23-1-2275, the Air Force Office of Scientific Research (AFOSR) # FA9550-22-1-0233, and IoT@CyLab