LinkedIn

A/B Testing in Dense Large-Scale Networks: Design and Inference

Oct 16, 2019




Preetam Nandy
Kinjal Basu
Shaunak Chatterjee
Ye Tu

Network Effect in A/B Testing

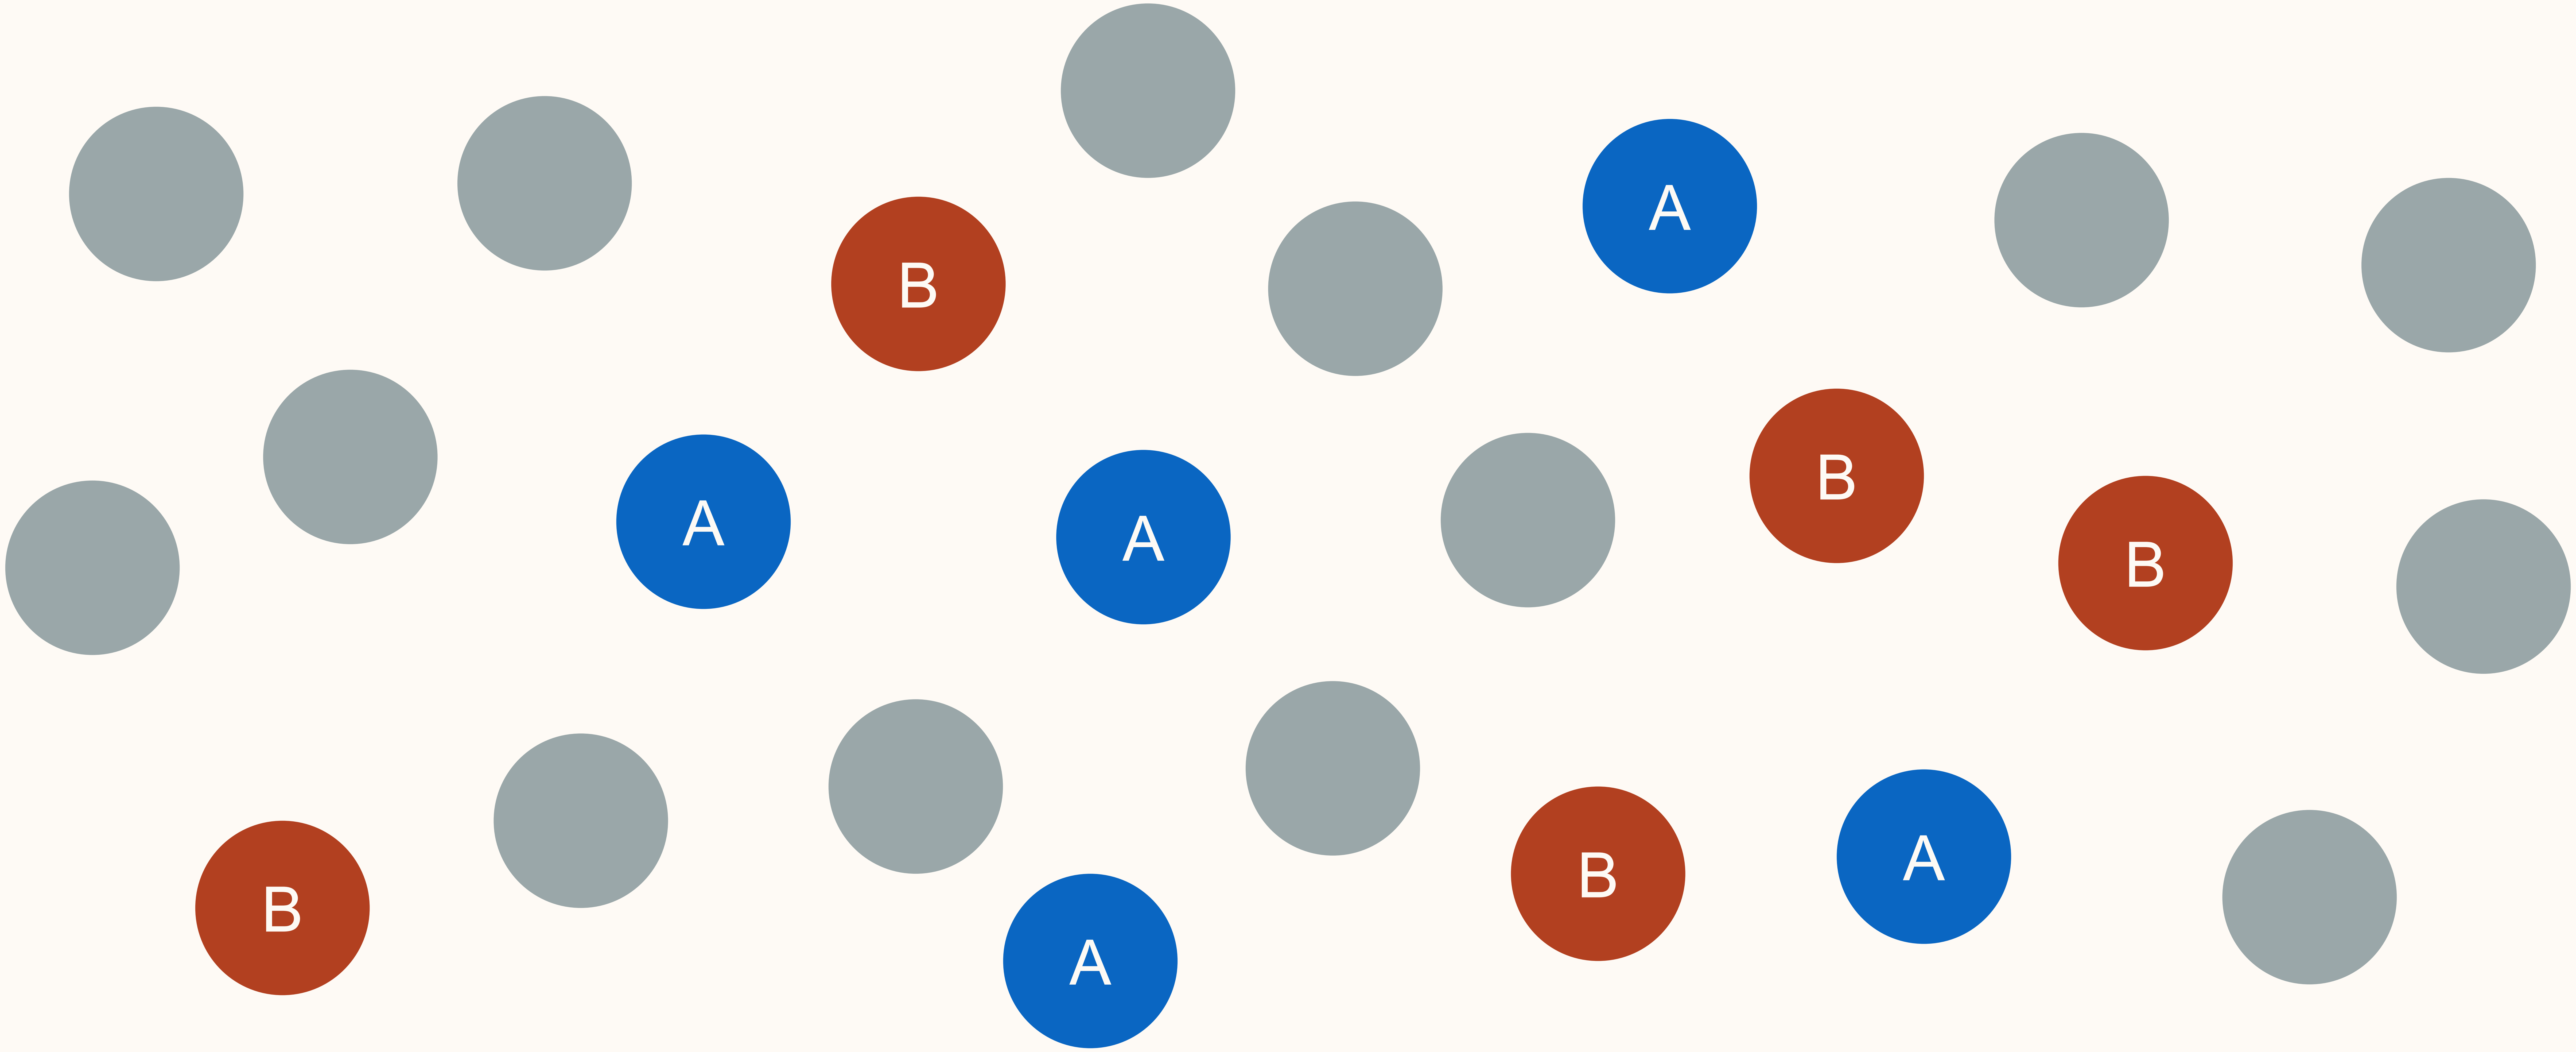

Network Effect in A/B Testing

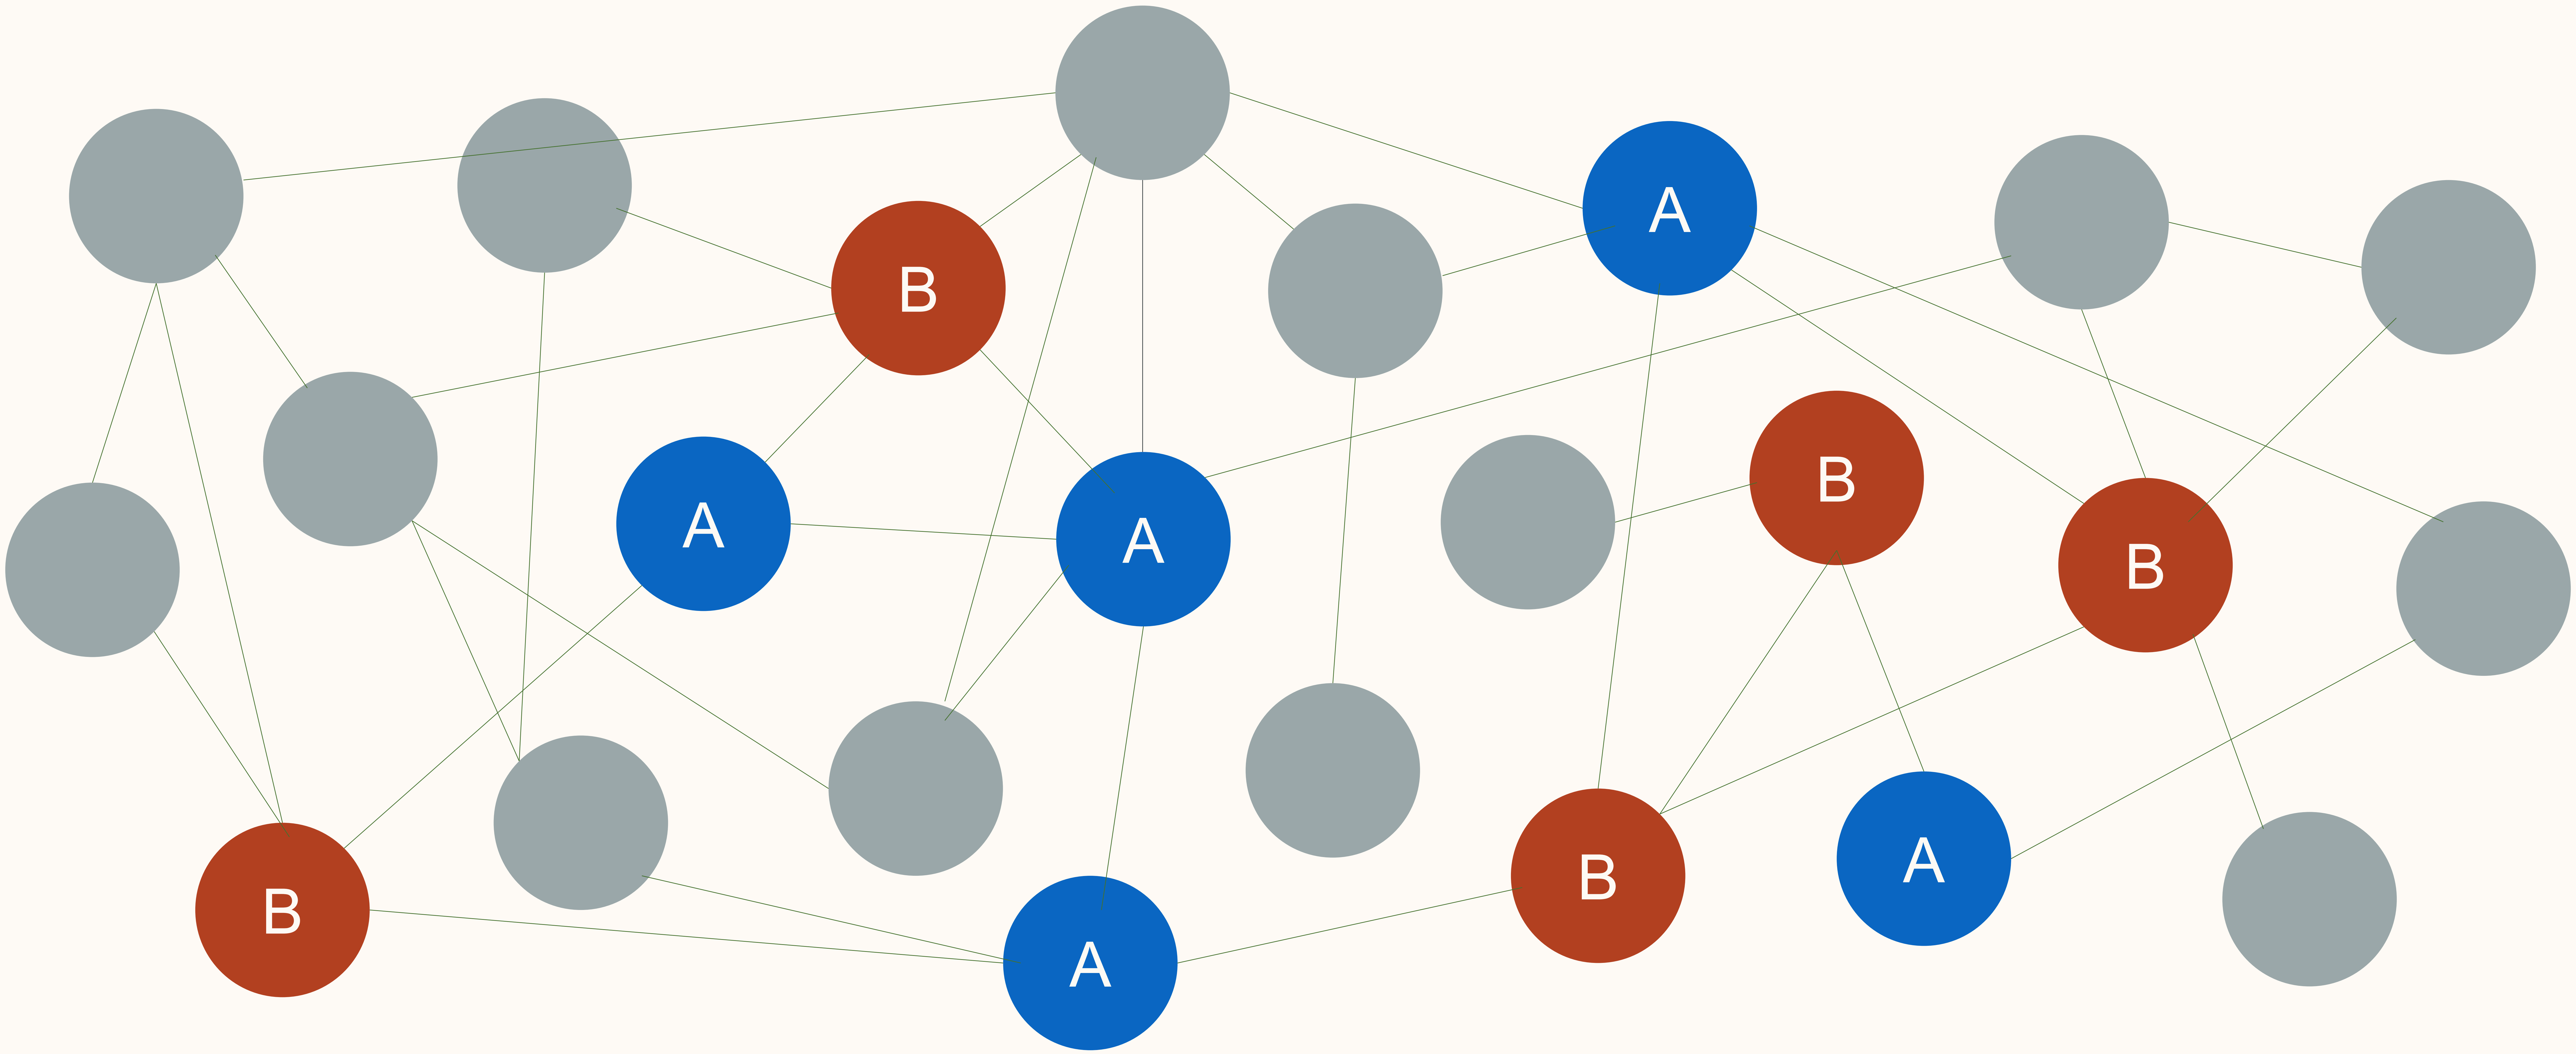

All neighbors of a treated node are *not* receiving the same treatment as the treated node

Network Effect in A/B Testing

Expectation

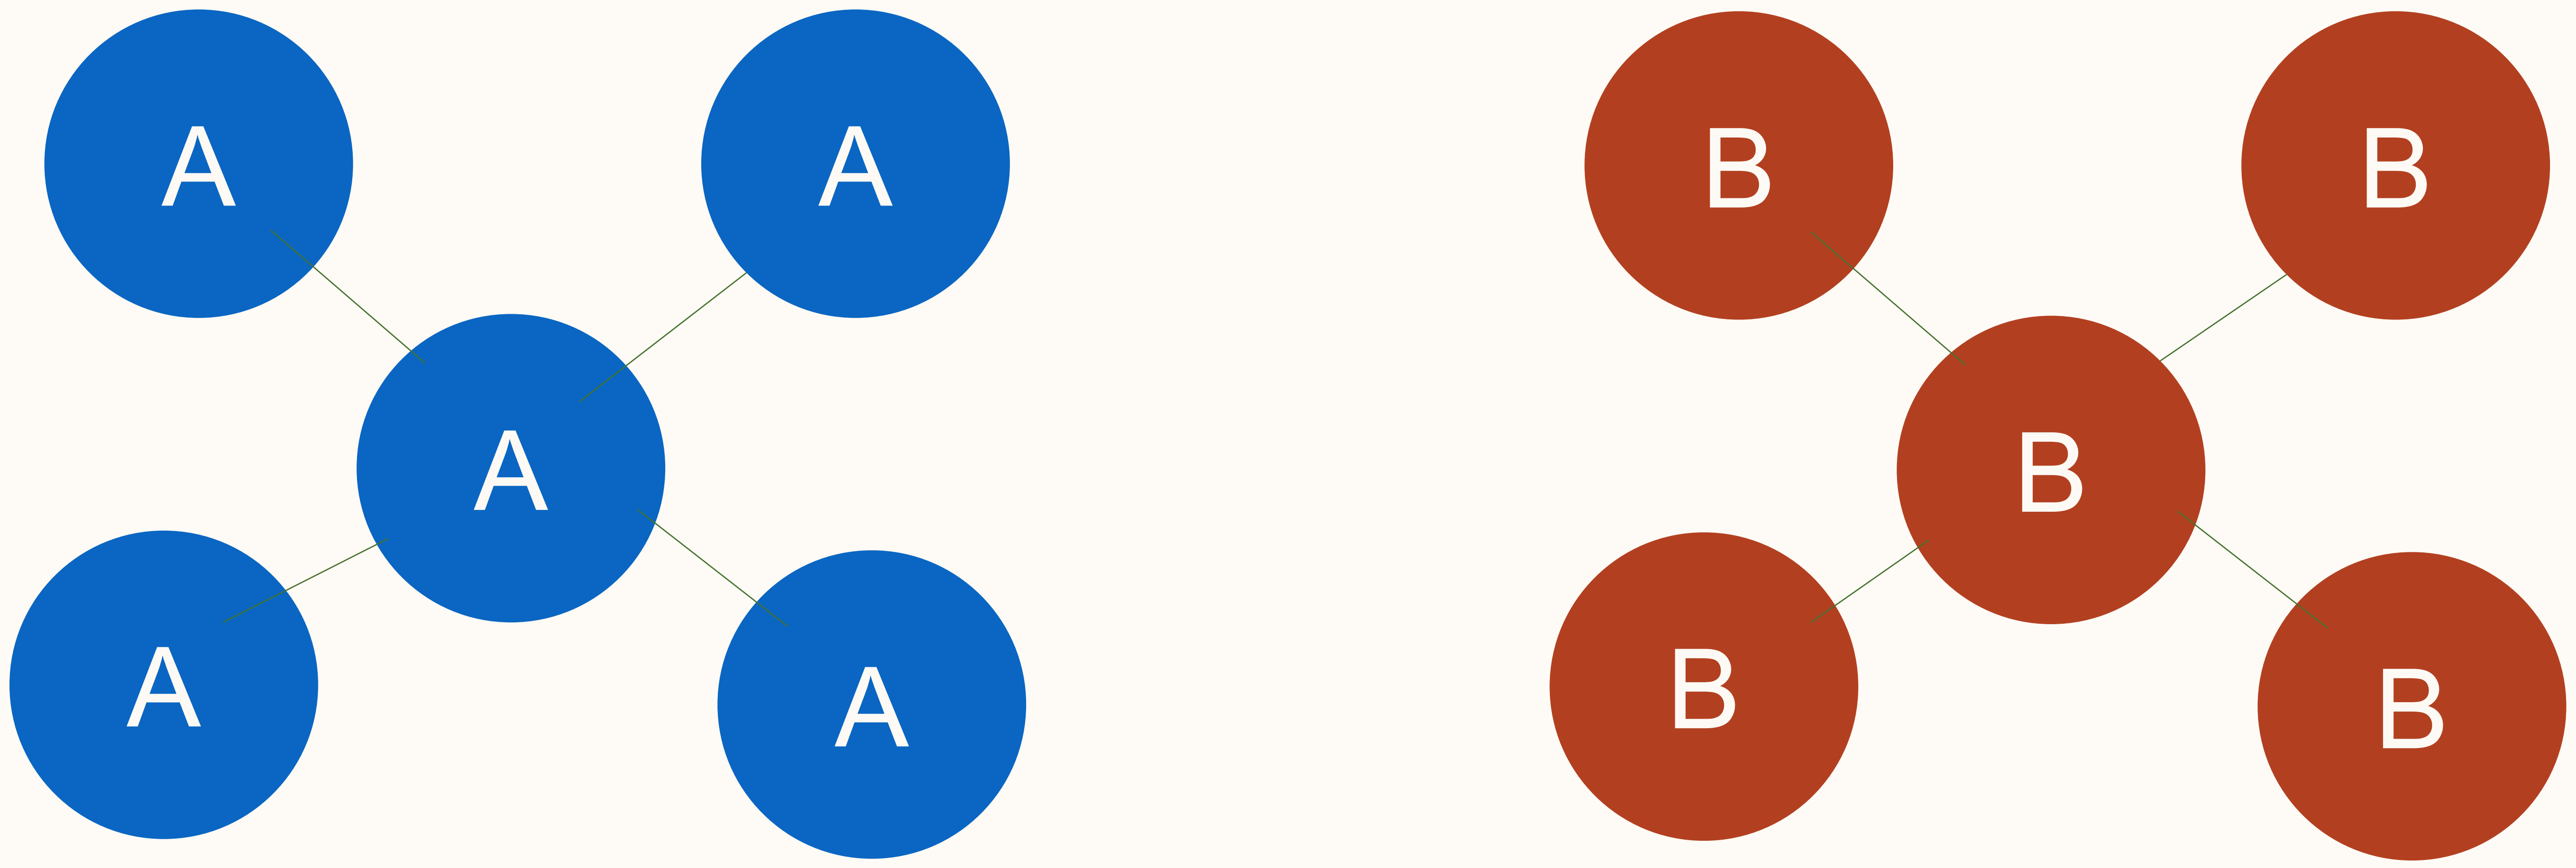

Reality

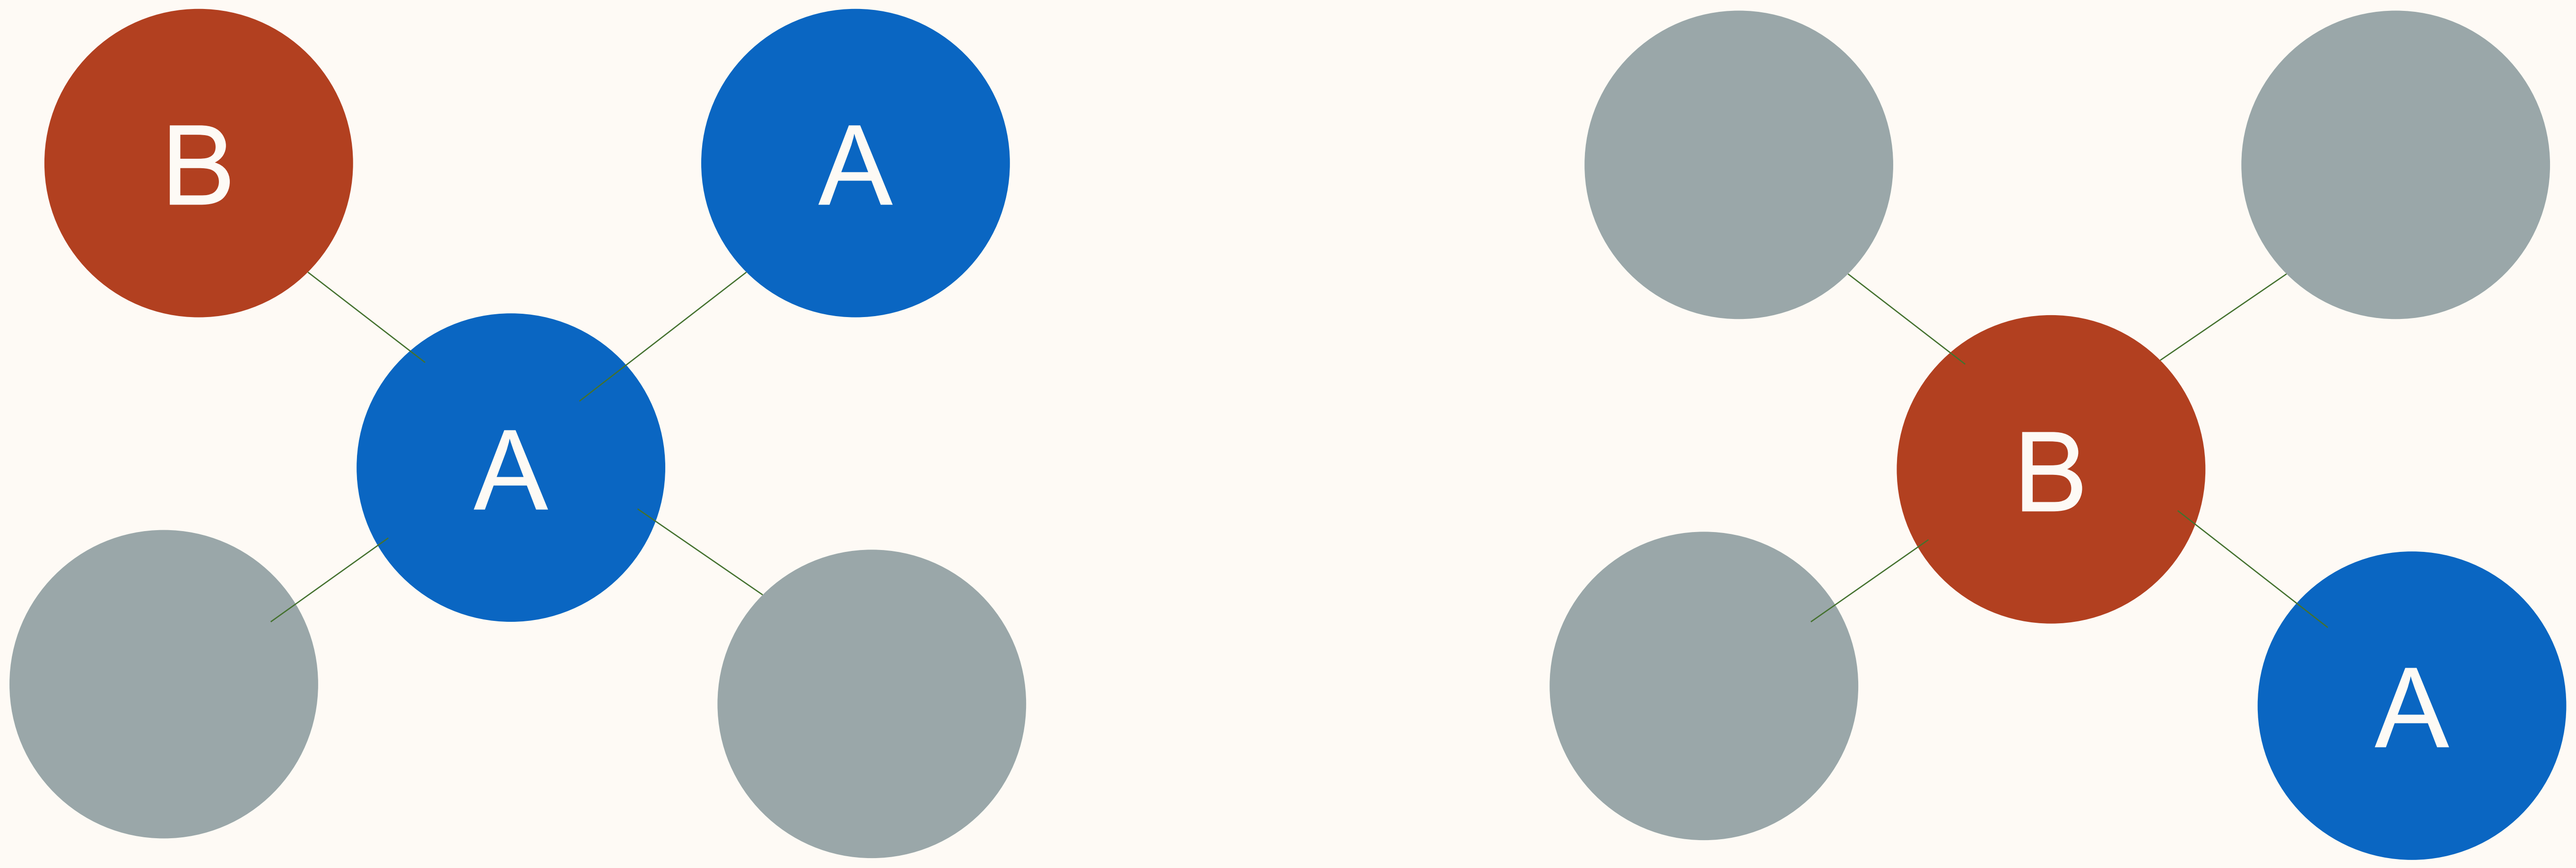

Network Effect in A/B Testing

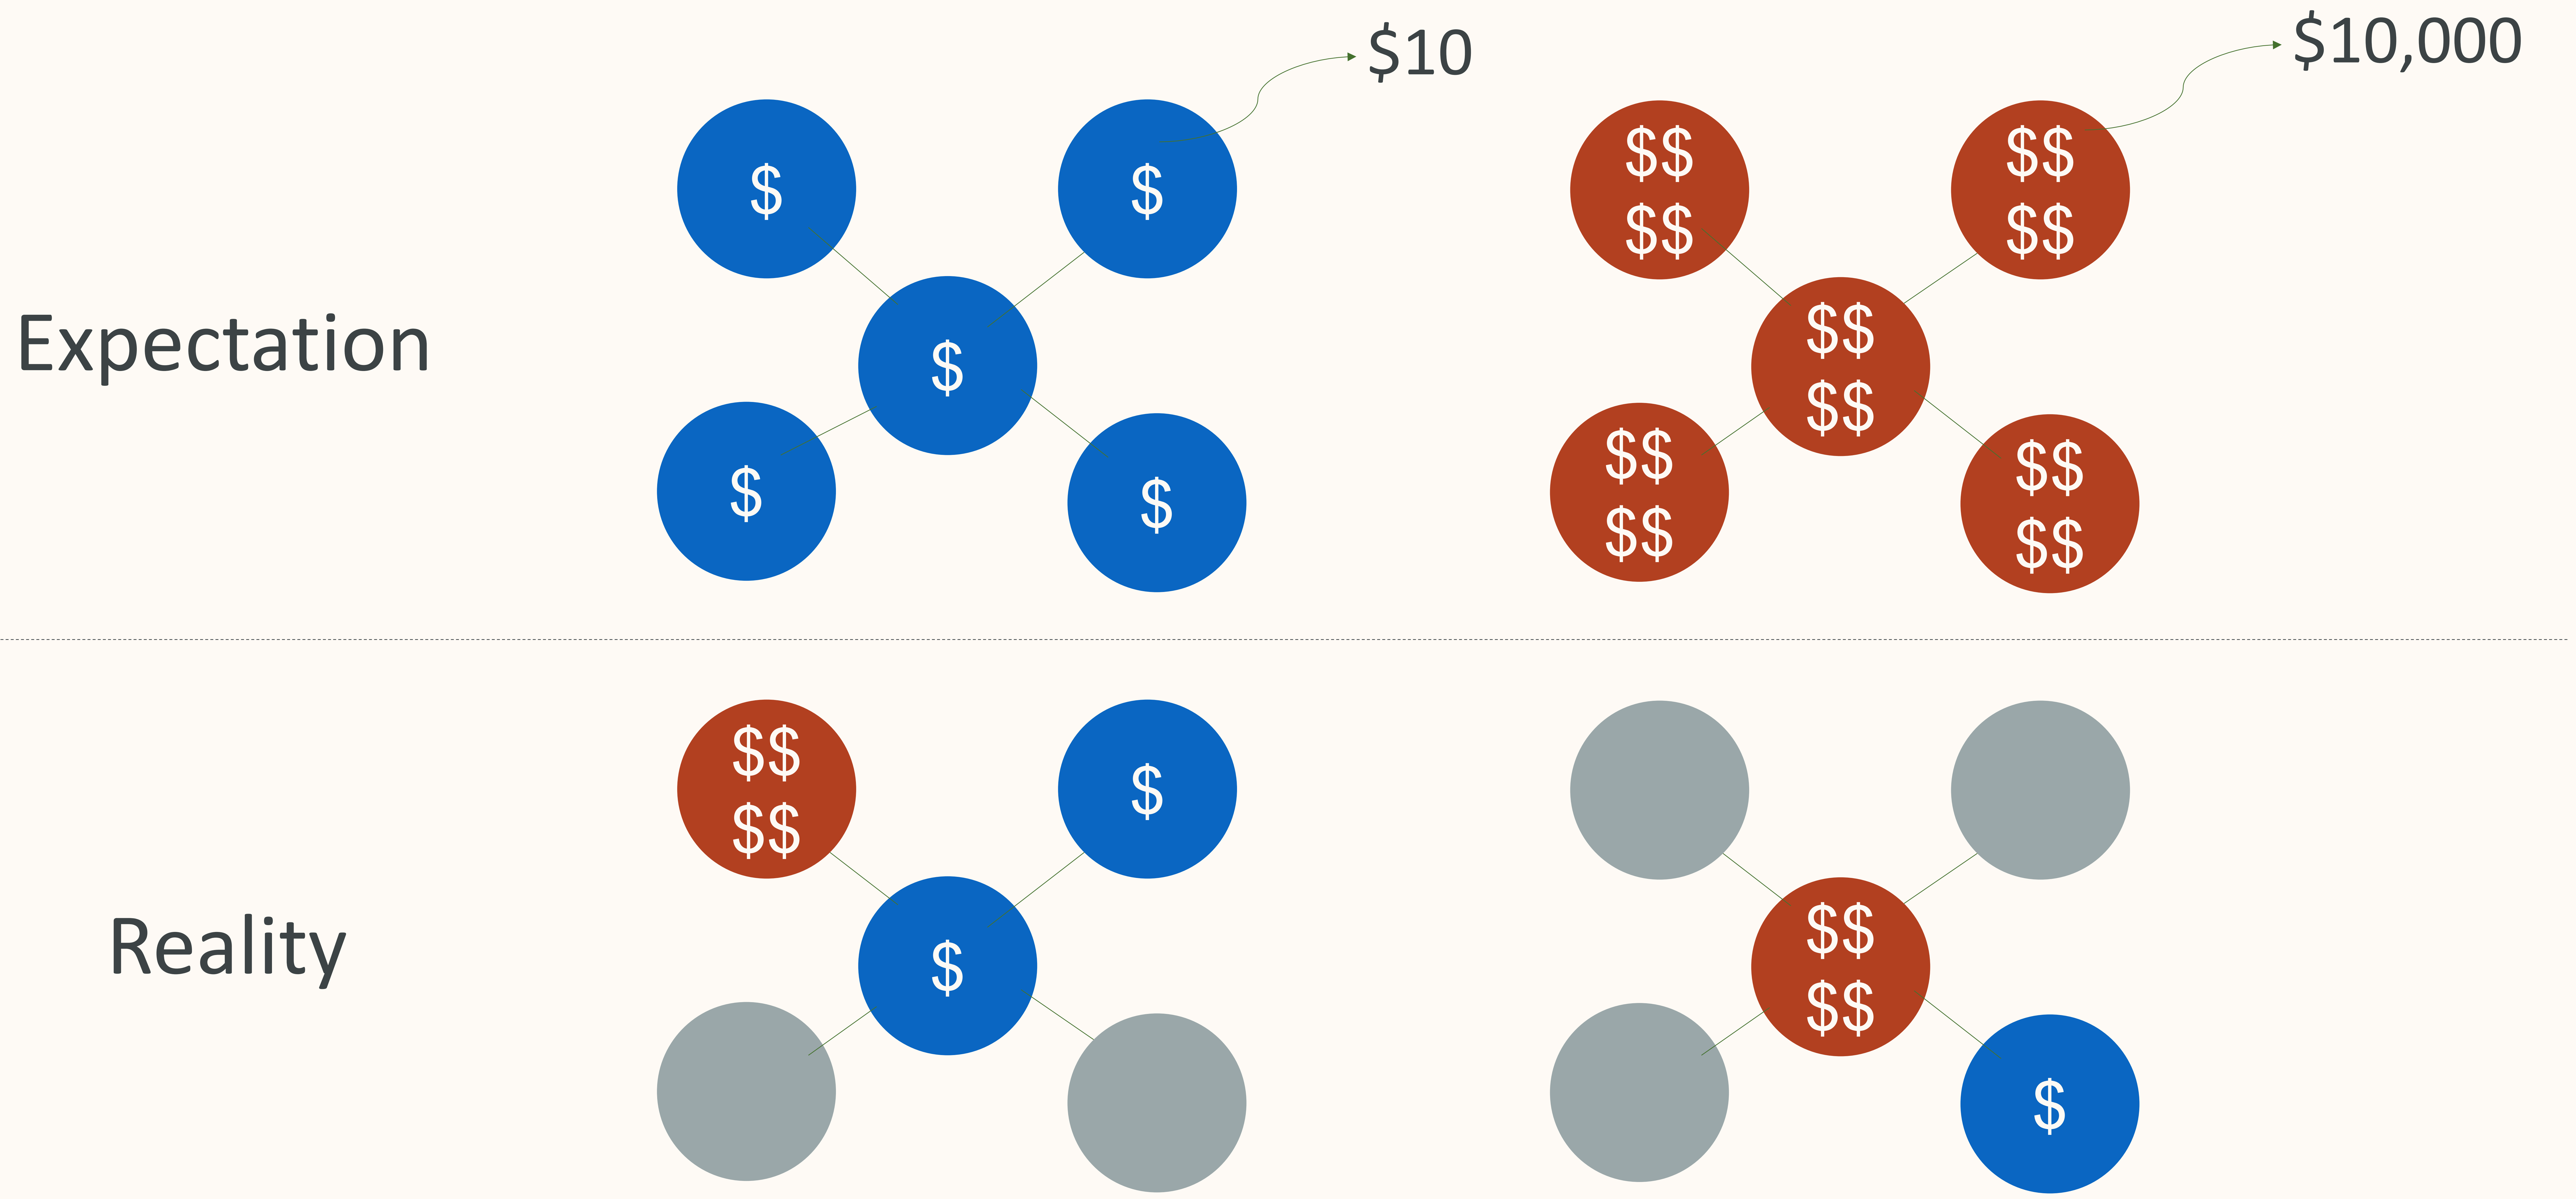

Ego-cluster Experiment: When Reality Meets Expectation

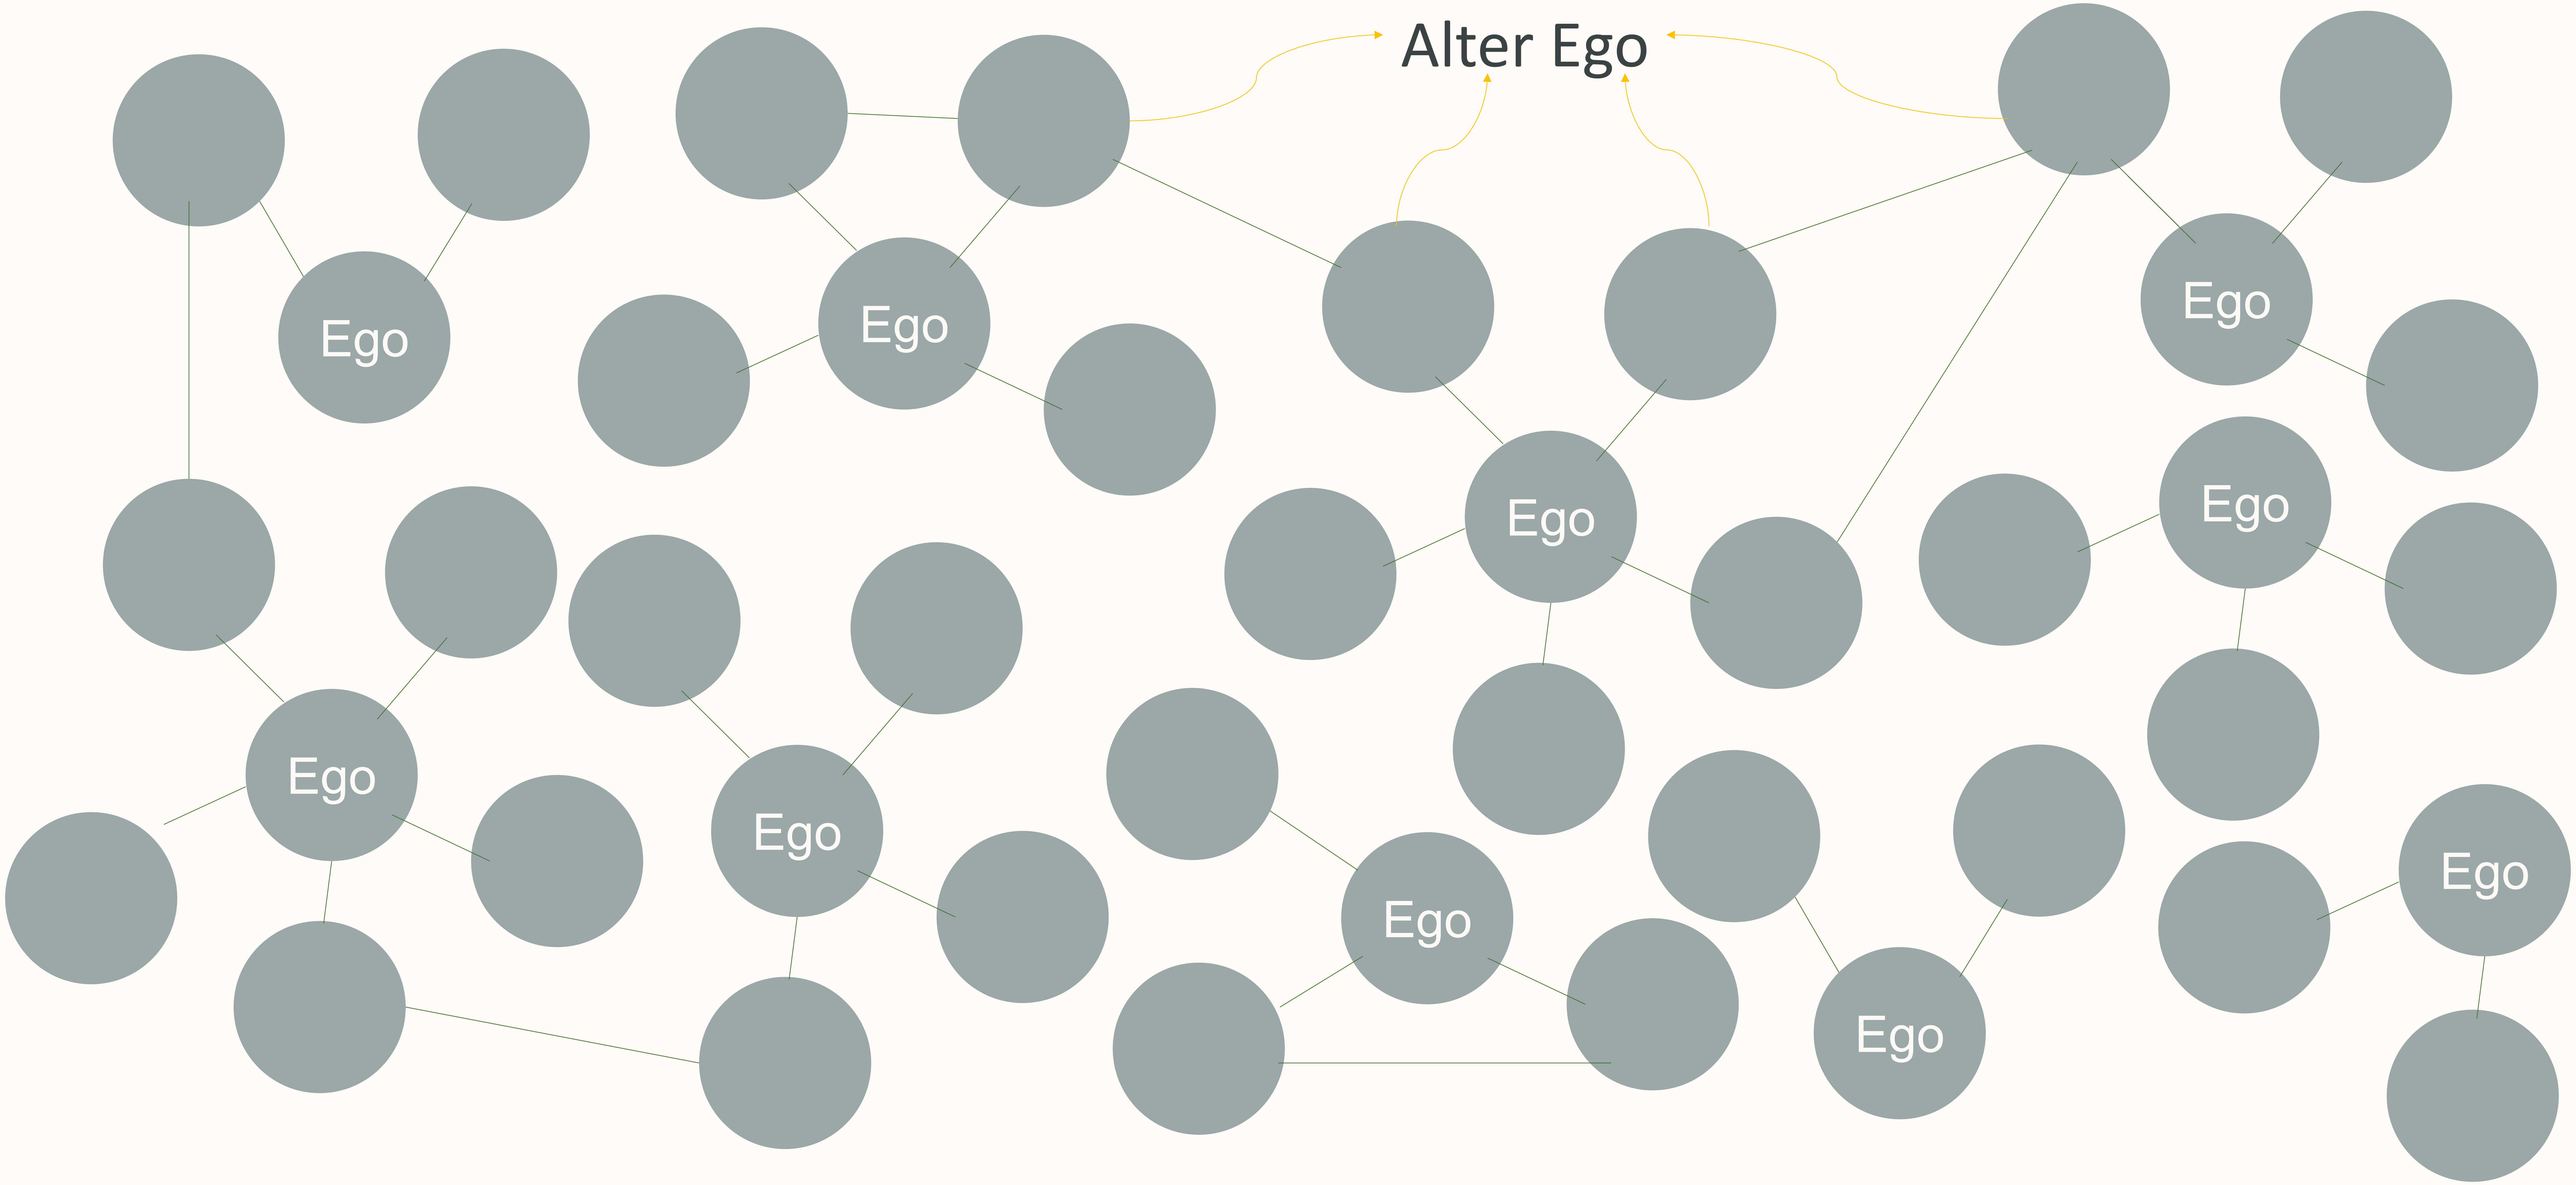

- G. Saint-Jacques, M. Varshney, J. Simpson, and Y. Xu. Using ego-clusters to measure network effects at LinkedIn. arXiv preprint arXiv:1903.08755, 2019.

Ego-cluster Experiment: When Reality Meets Expectation

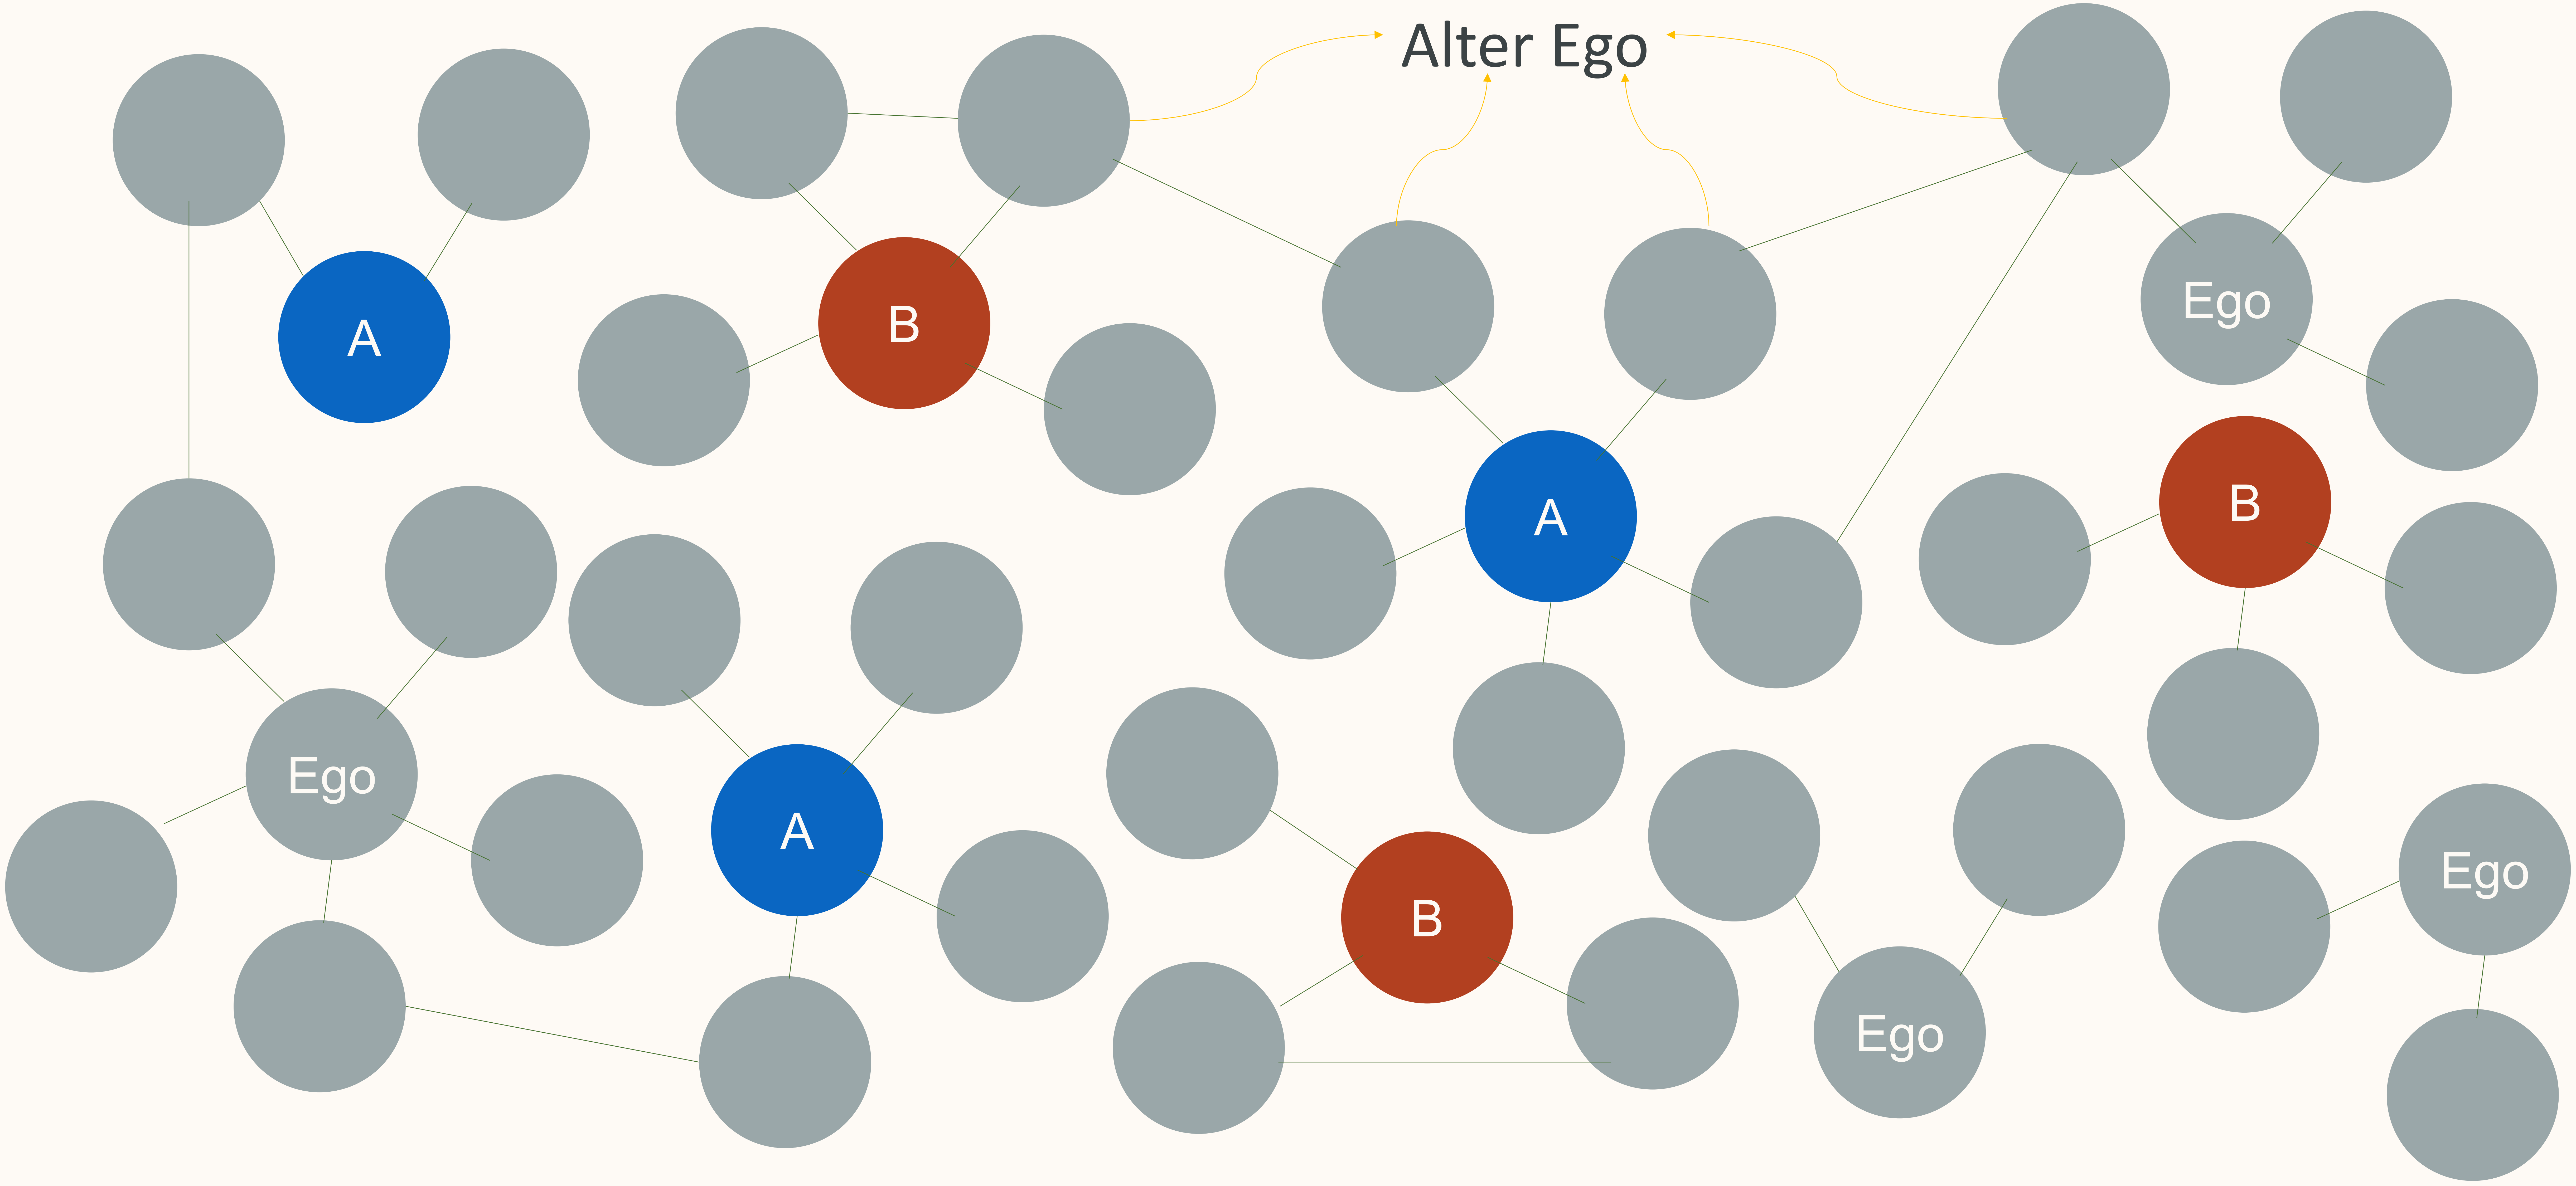

- G. Saint-Jacques, M. Varshney, J. Simpson, and Y. Xu. Using ego-clusters to measure network effects at LinkedIn. arXiv preprint arXiv:1903.08755, 2019.

Ego-cluster Experiment: When Reality Meets Expectation

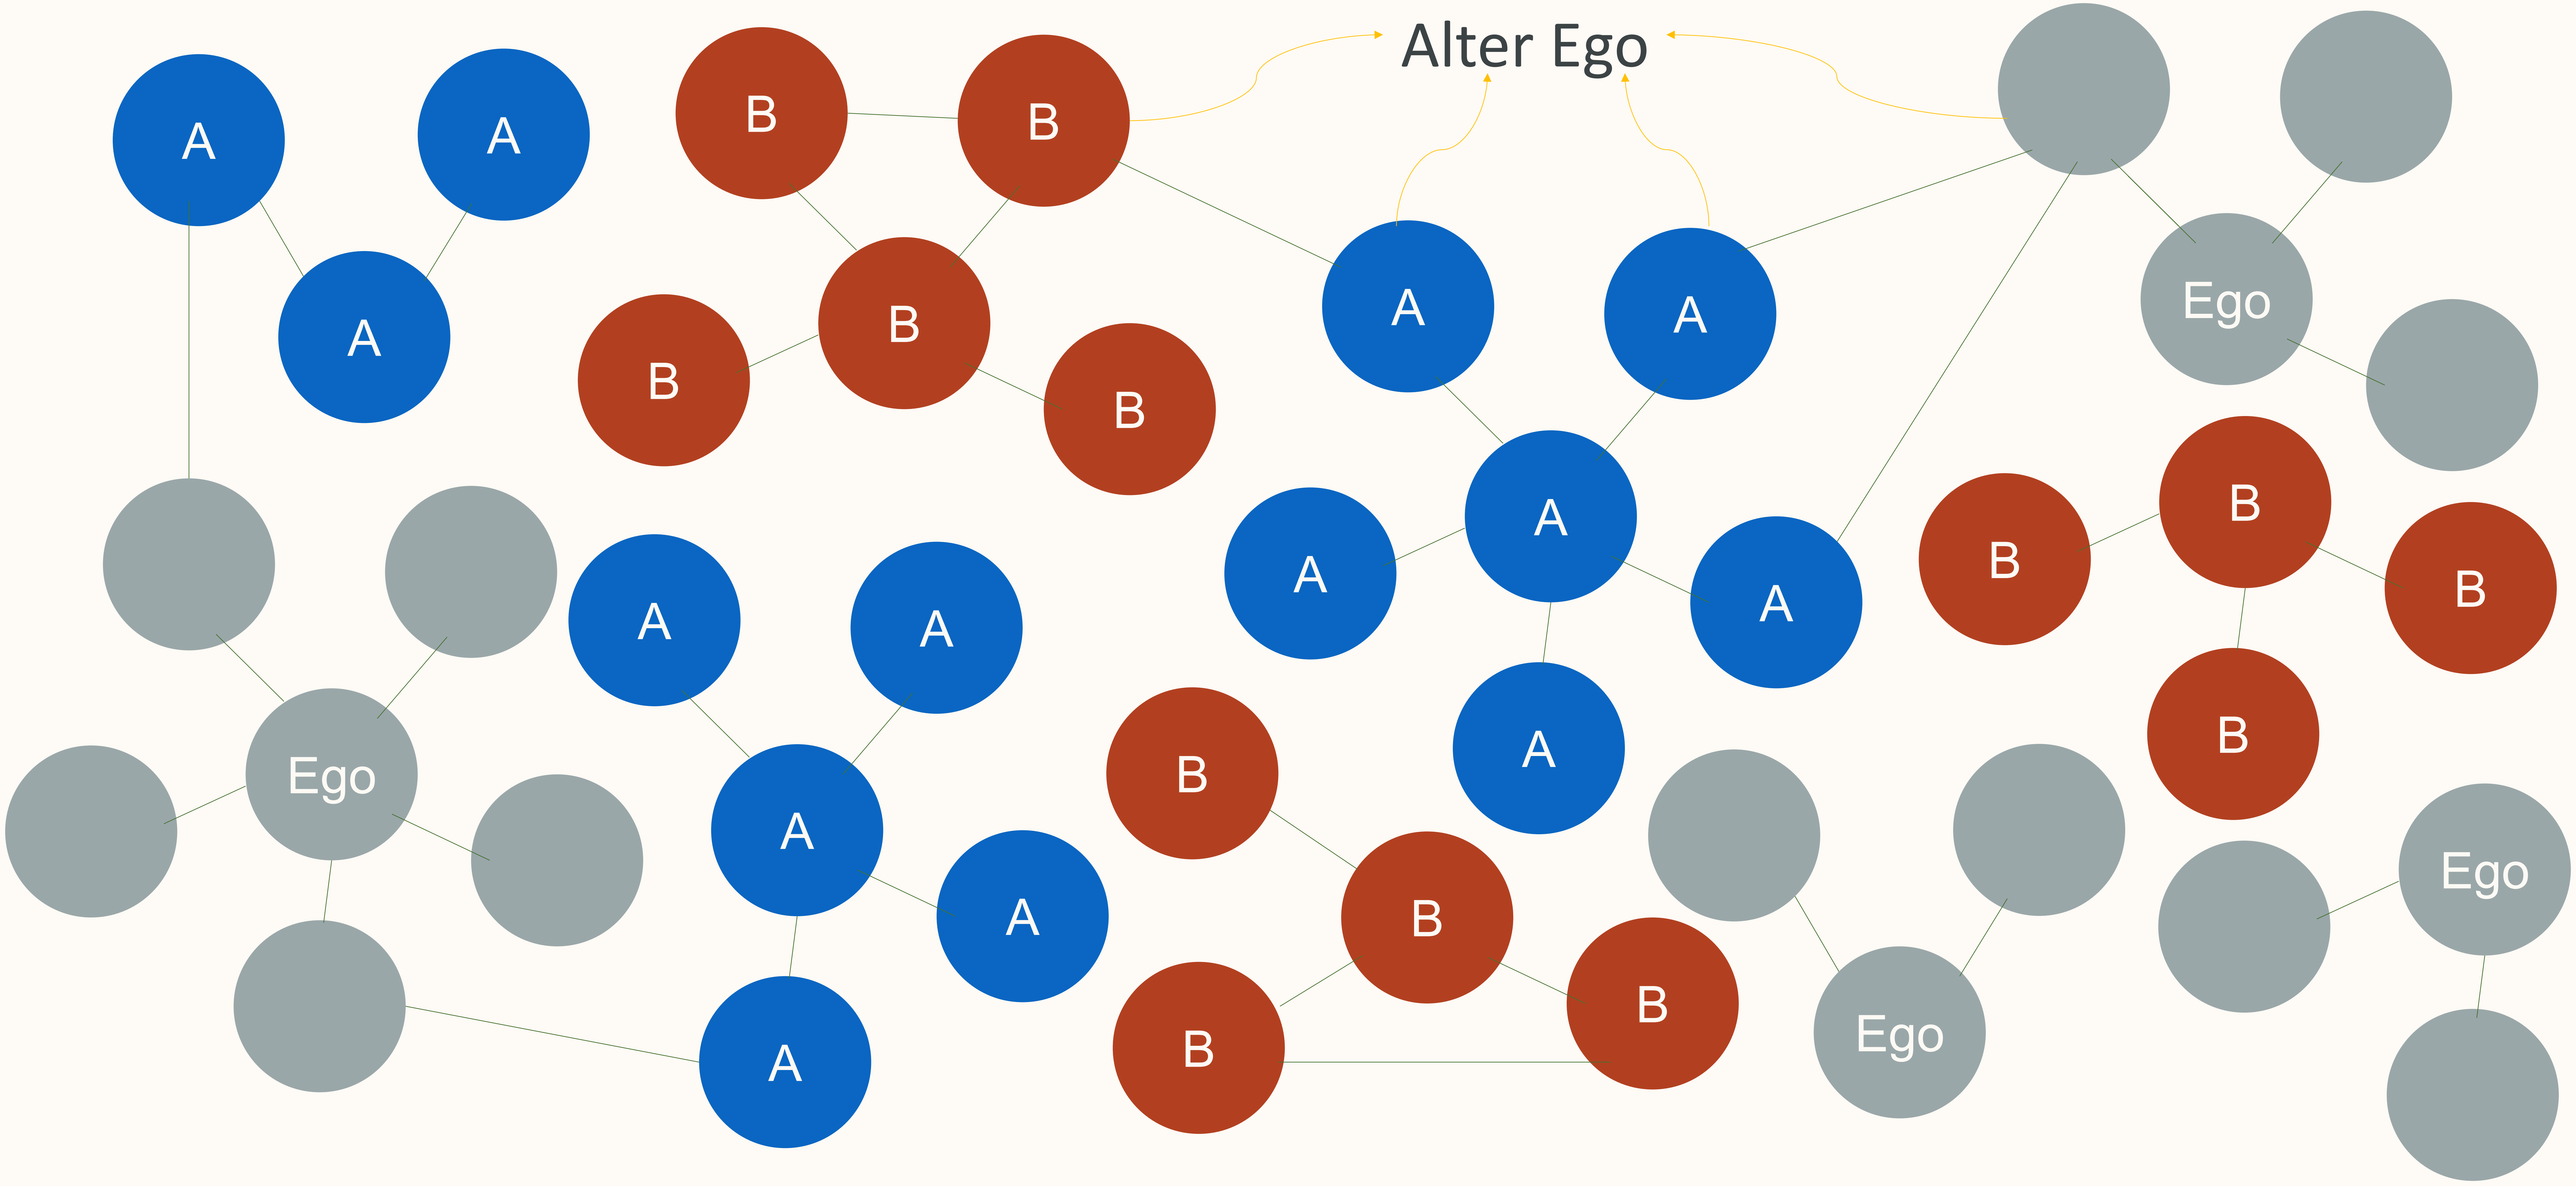

• G. Saint-Jacques, M. Varshney, J. Simpson, and Y. Xu. Using ego-clusters to measure network effects at LinkedIn. arXiv preprint arXiv:1903.08755, 2019.

Ego-cluster Experiment: When Reality Meets Expectation

Pros:

- A general solution for taking care of network effects (up to the first degree)
- Works under minimal assumptions

Cons:

- It may not be possible to find an ideal set of ego clusters even for moderately dense networks: Need to allow for a small amount of contamination
- Does not work well for dense networks

- G. Saint-Jacques, M. Varshney, J. Simpson, and Y. Xu. Using ego-clusters to measure network effects at LinkedIn. arXiv preprint arXiv:1903.08755, 2019.

Network Effect in Exposure Redistribution Experiments

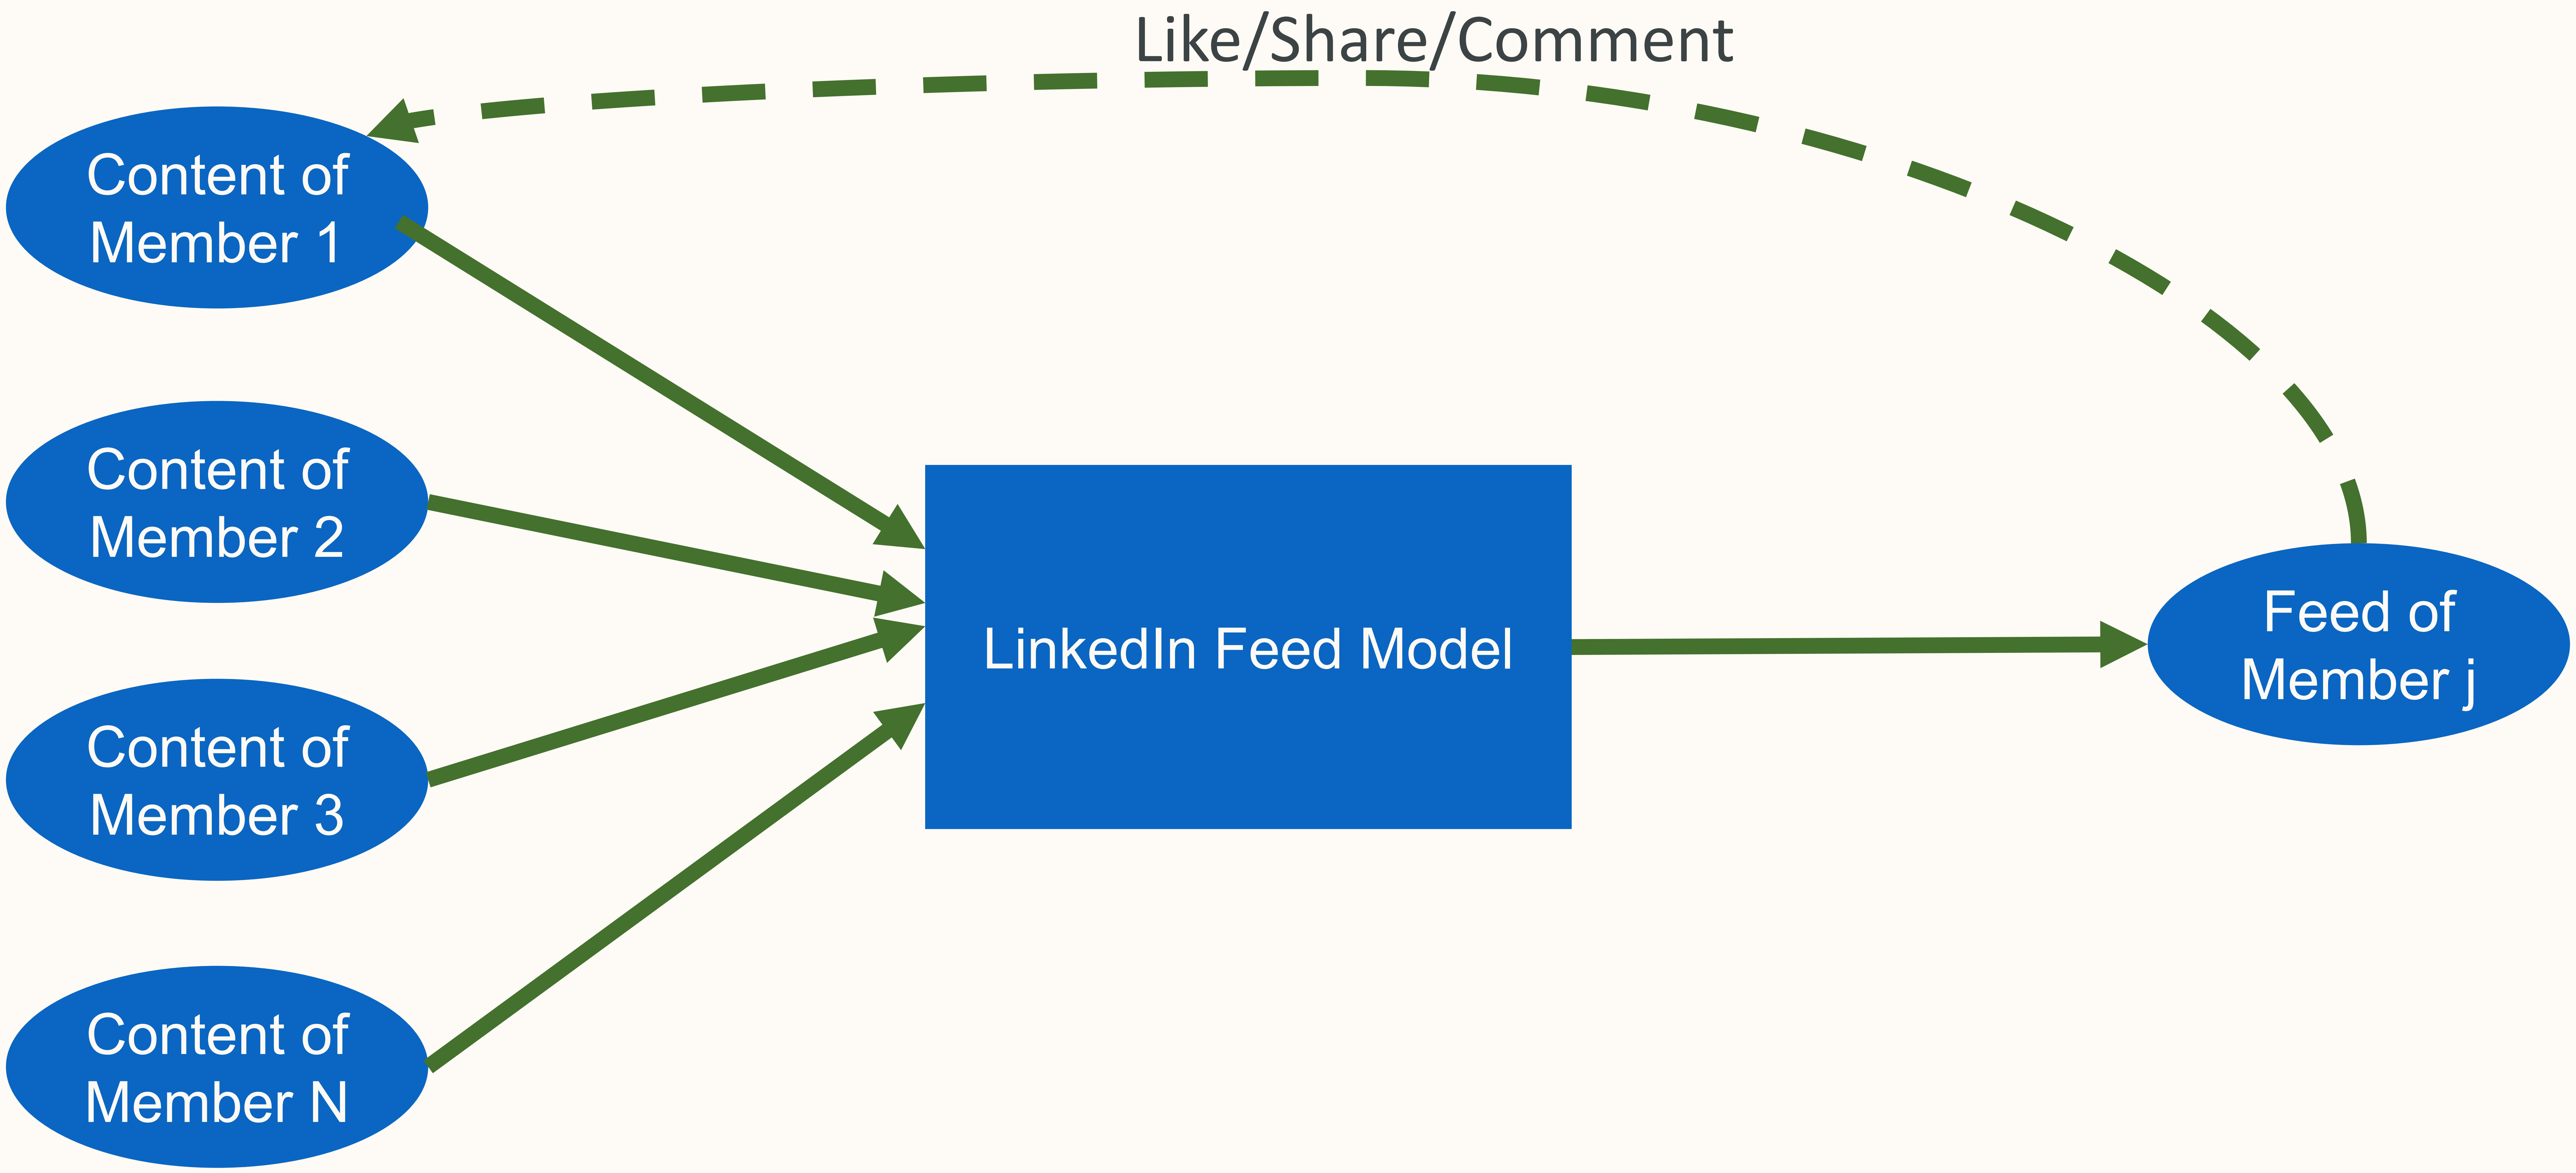

Network Effect in Exposure Redistribution Experiments

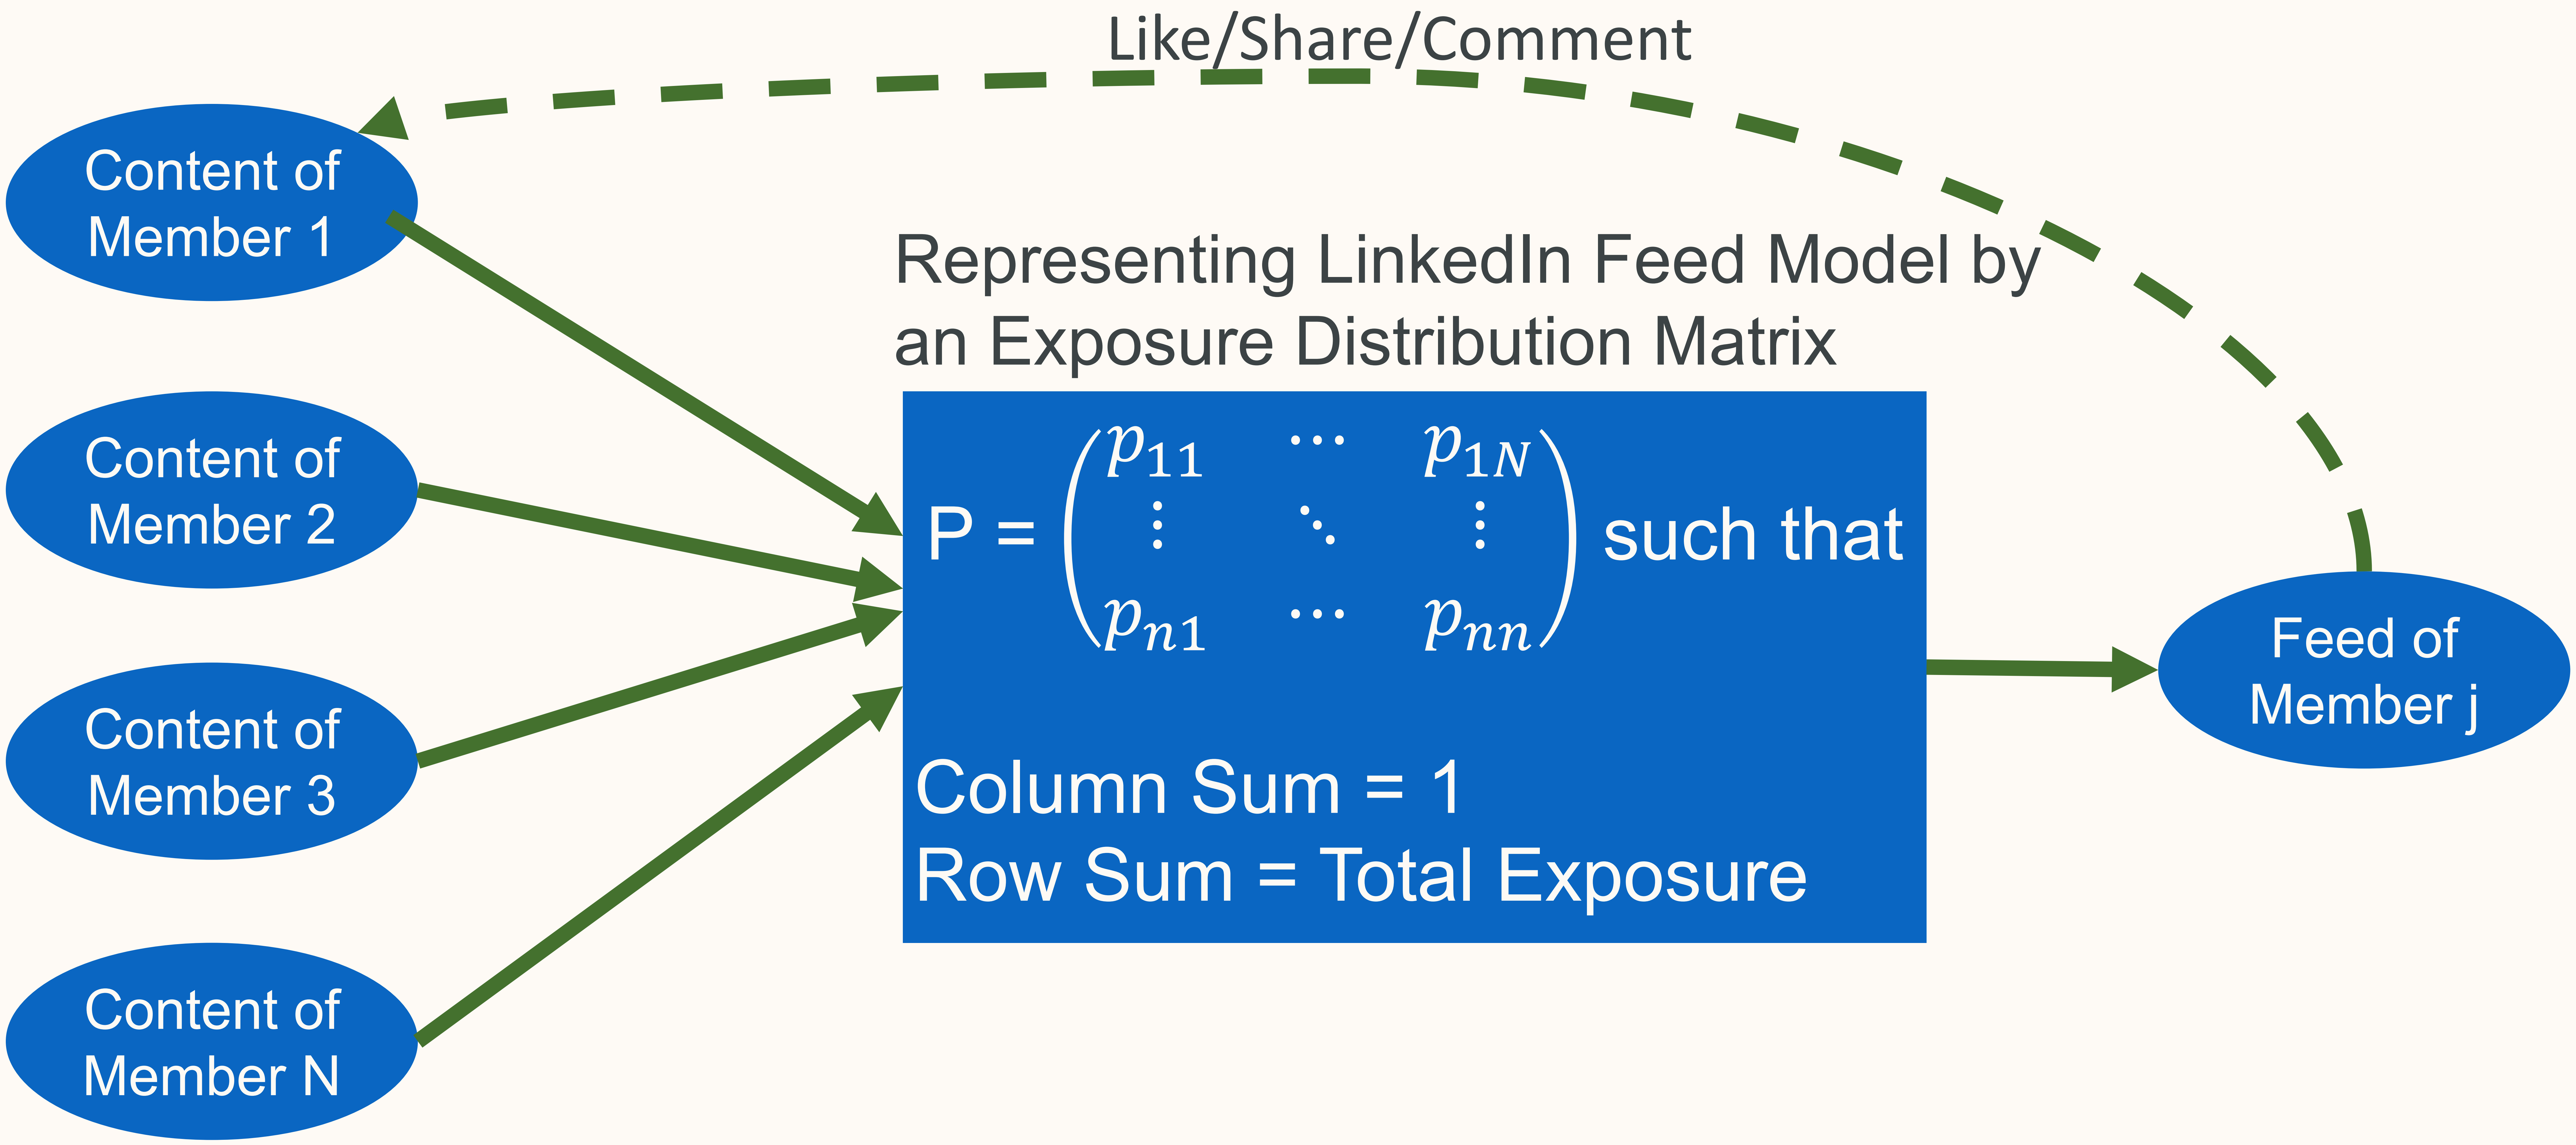

Network Effect in Exposure Redistribution Experiments

Exposure Distribution A

0	0.30	0.20	0	0.60	0
0	*	*	*	*	0
0.50	*	*	. *	*	0.20
0.25	*	*	*	*	0.60
0	*	*	*	*	0.20
0.25	0.10	0	0.10	0.10	0

Exposure Distribution B

0	0.40	0.10	0	0.25	0
0	*	*	*	*	0.30
0.20	*	*	*	*	0.40
0.45	*	*	*	*	0
0	*	*	*	*	0.30
0.35	0	0	0.30	0.20	0

A/B Testing Design

0	0.20	0.10	0	0.60	0
0	*	*	*	*	0.30
0.50	*	*	*	*	0.40
0.25	*	*	*	*	0
0	*	*	*	*	0.30
0.25	0.20	0	0.40	0.10	0

Network Effect in Exposure Redistribution Experiments

Exposure Distribution A

0	0.30	0.20	0	0.60	0	= 1.10
0	*	*	*	*	0	
0.50	*	*	. *	*	0.20	
0.25	*	*	*	*	0.60	
0	*	*	*	*	0.20	
0.25	0.10	0	0.10	0.10	0	

Exposure Distribution B

0	0.40	0.10	0	0.25	0	
0	*	*	*	*	0.30	
0.20	*	*	*	*	0.40	
0.45	*	*	*	*	0	
0	*	*	*	*	0.30	
0.35	0	0	0.25	0.25	0	= 0.85

A/B Testing Design

0	0.20	0.10	0	0.60	0	$= 0.90 \neq 1.10$
0	*	*	*	*	0.30	
0.50	*	*	*	*	0.40	
0.25	*	*	*	*	0	
0	*	*	*	*	0.30	
0.25	0.20	0	0.40	0.10	0	$= 0.95 \neq 0.85$

Curing Network Effect with Additional Assignments

Exposure Distribution A

0	0.30	0.20	0	0.60	0	= 1.10
0	*	*	*	*	0	
0.50	*	*	. *	*	0.20	
0.25	*	*	*	*	0.60	
0	*	*	*	*	0.20	
0.25	0.10	0	0.10	0.10	0	

Exposure Distribution B

0	0.40	0.10	0	0.25	0	
0	*	*	*	*	0.30	
0.20	*	*	*	*	0.40	
0.45	*	*	*	*	0	
0	*	*	*	*	0.30	
0.35	0	0	0.25	0.25	0	= 0.85

A/B Testing Design with Additional Assignments

0	0.20	0.10	0.20	0.60	0	= 1.10
0	*	*	*	*	0.30	
0.50	*	*	*	*	0.40	
0.25	*	*	*	*	0	
0	*	*	*	*	0.30	
0.25	0.20	0	0.30	0.10	0	= 0.85

The Main Assumption

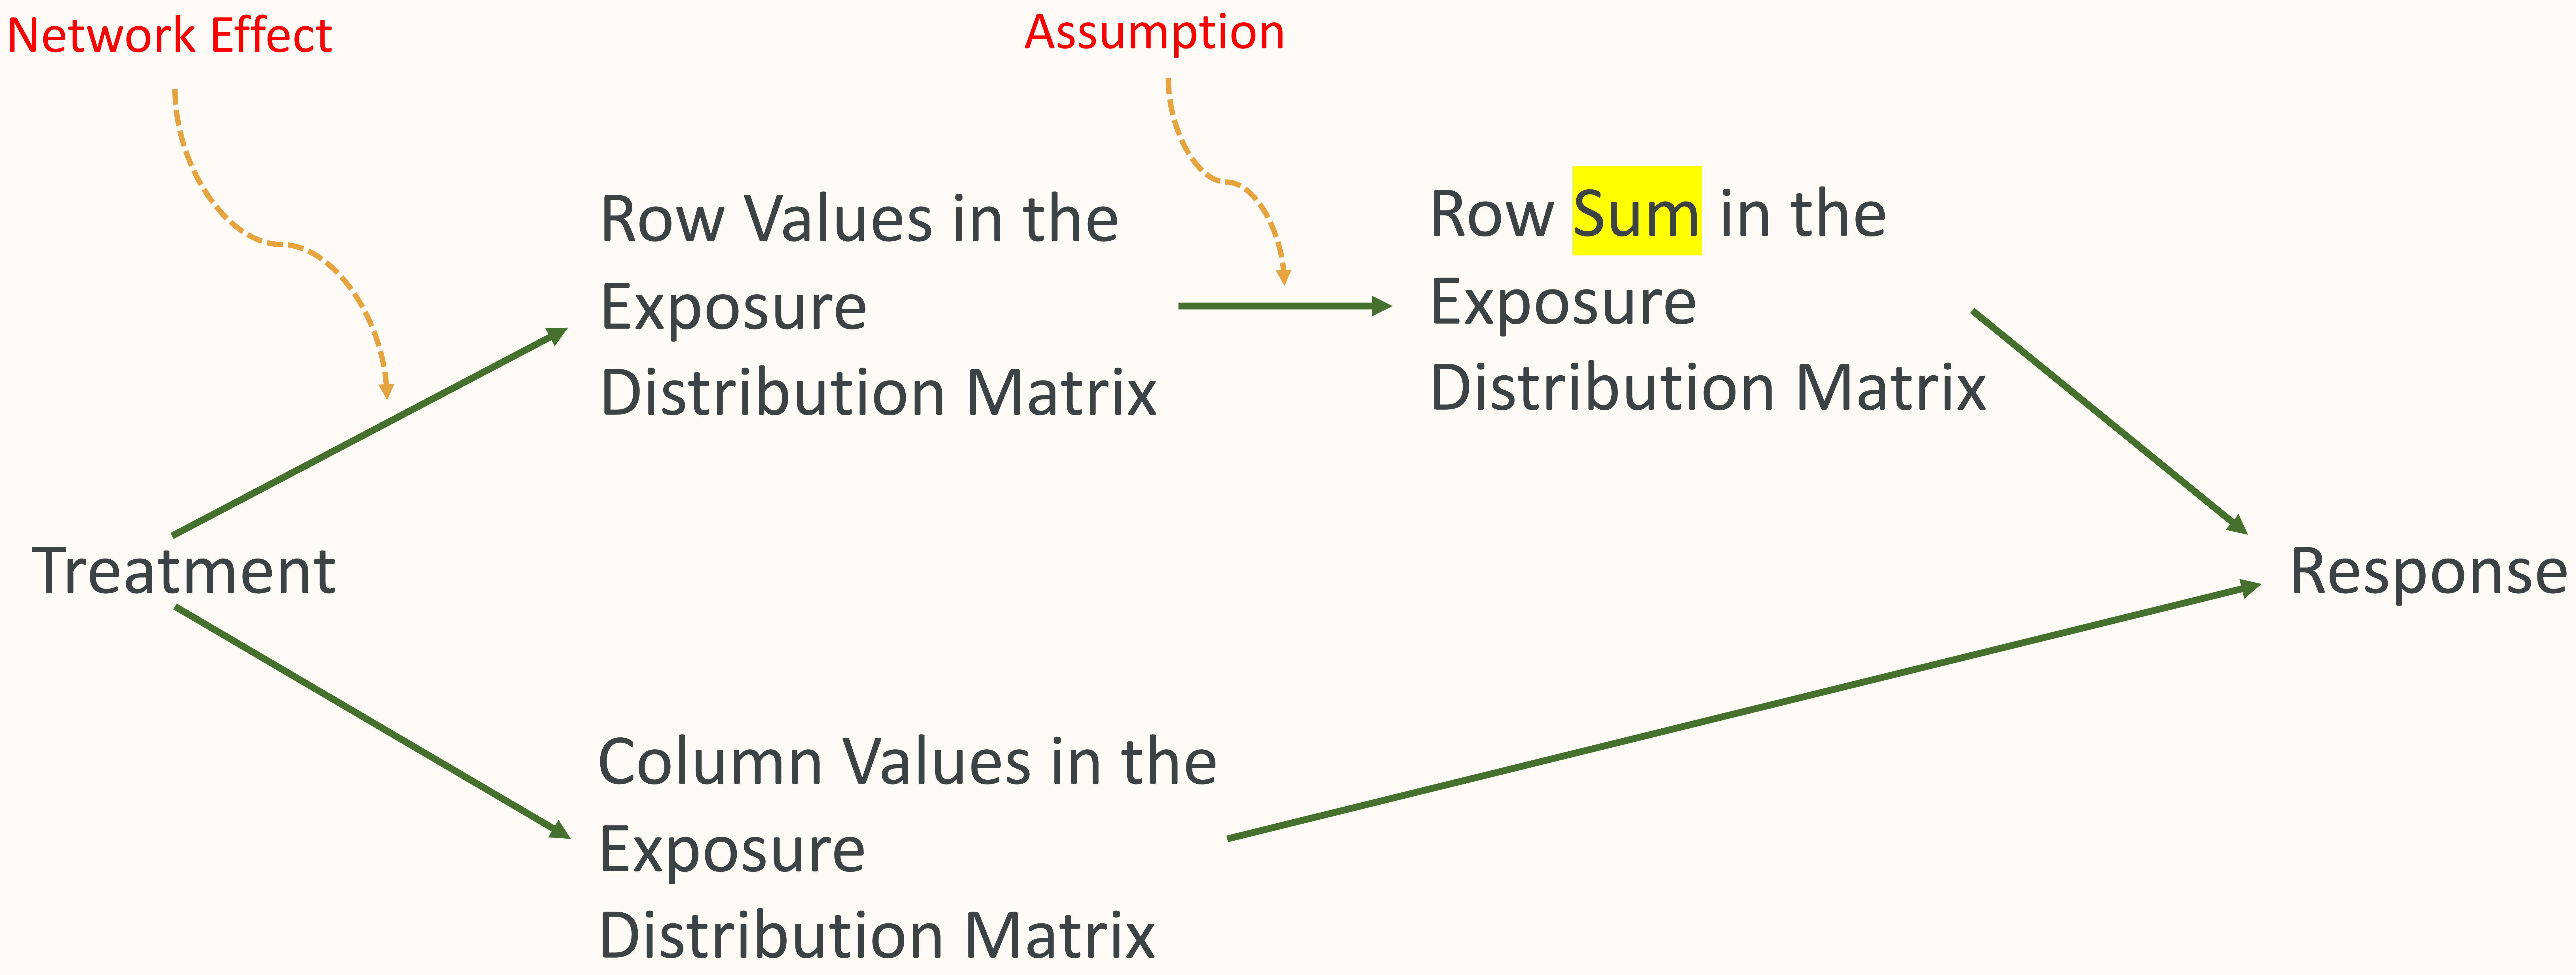

Step 1: Randomly Assign A and B

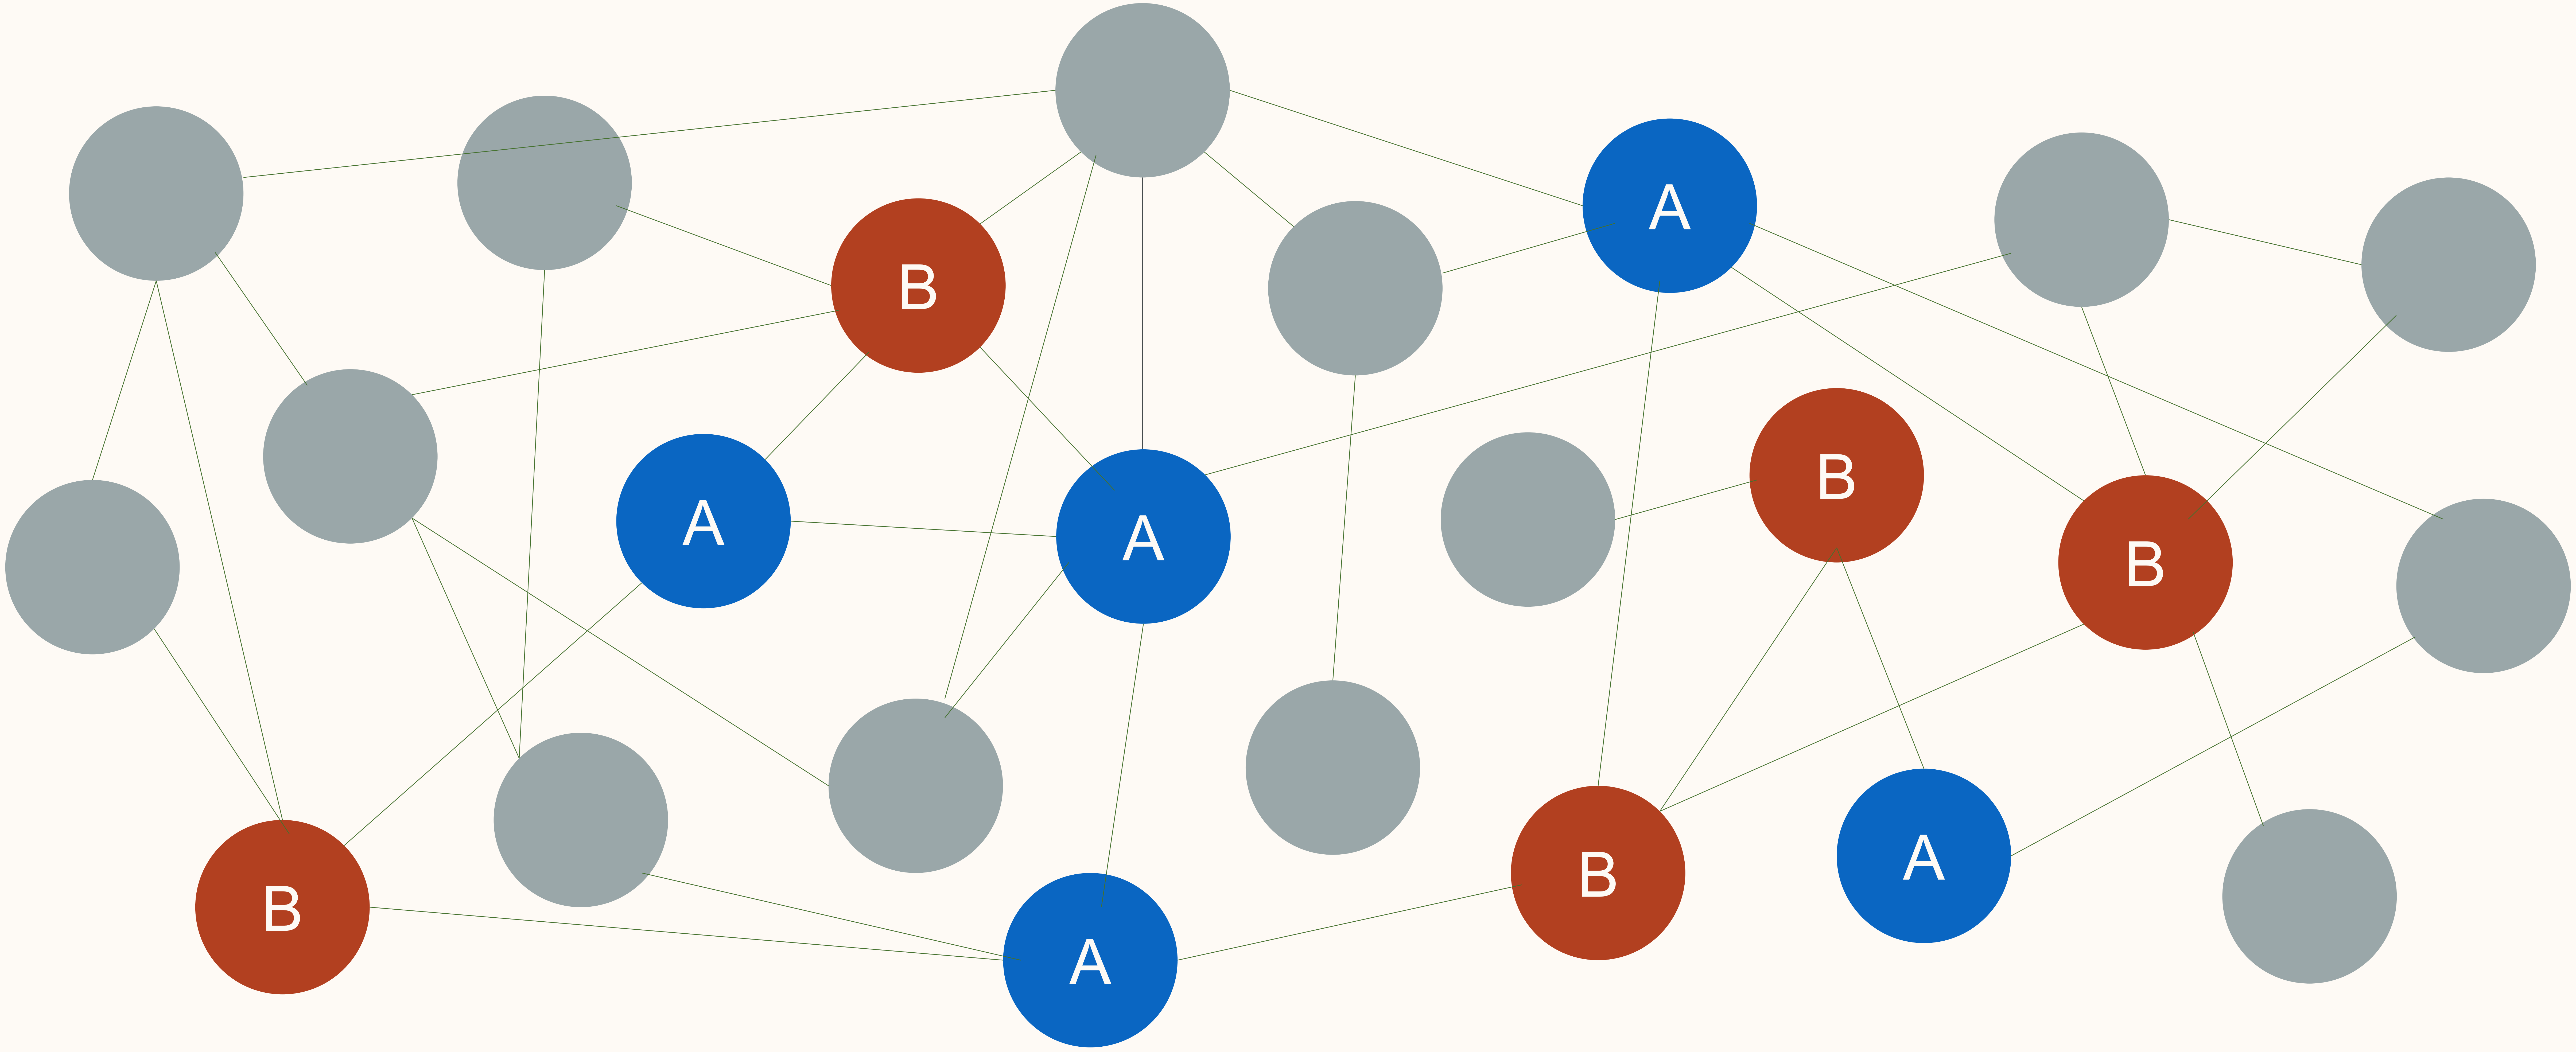

Step 2: Randomly Choose Additional Nodes to Assign C

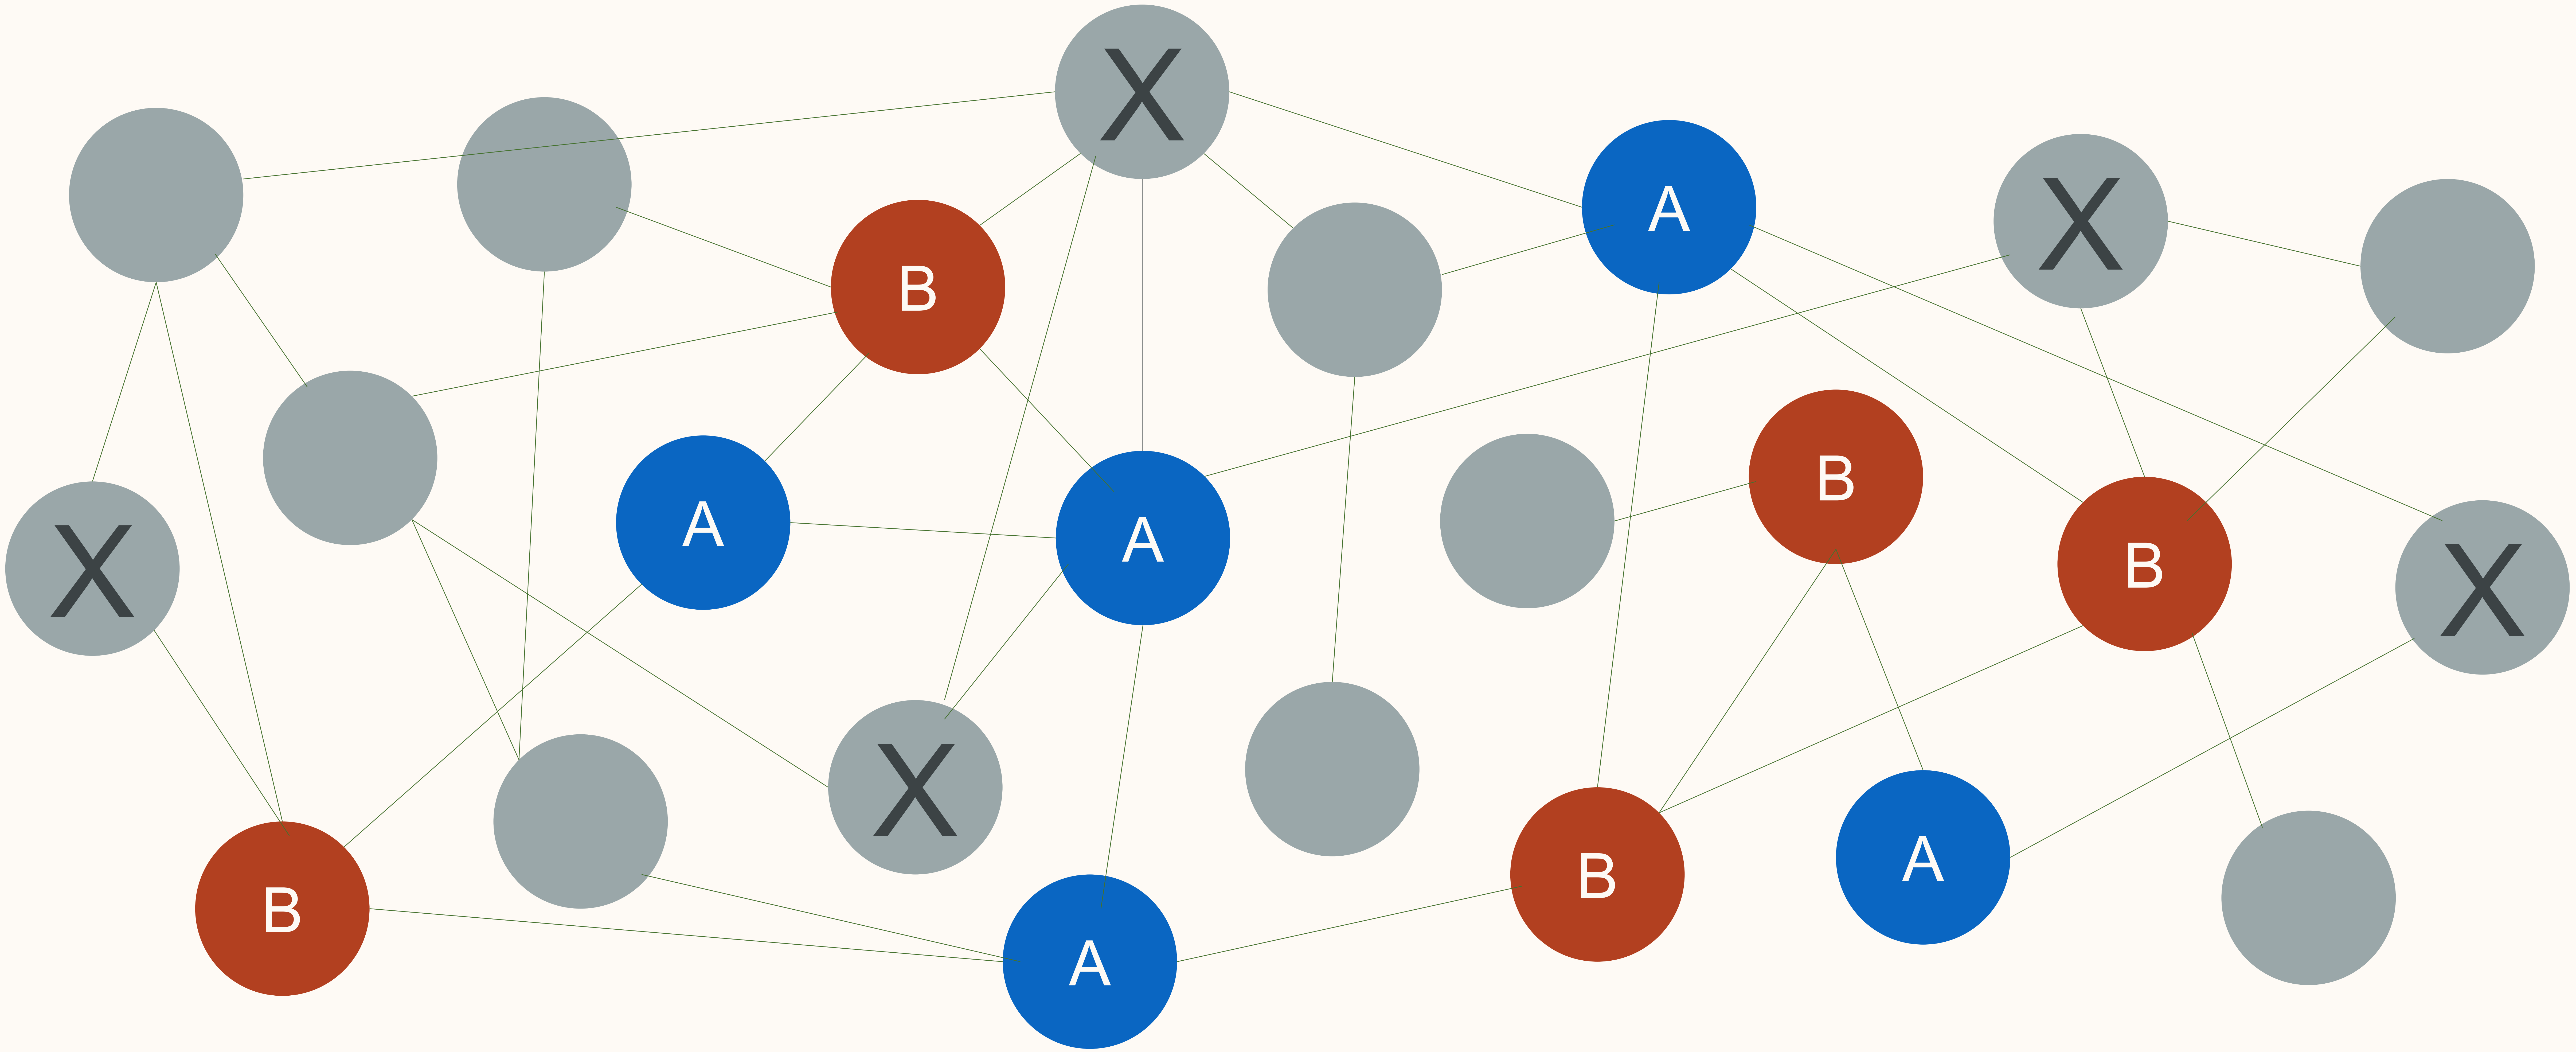

Step 3: Solve a Constrained Optimization Problem

Obtain C by minimizing

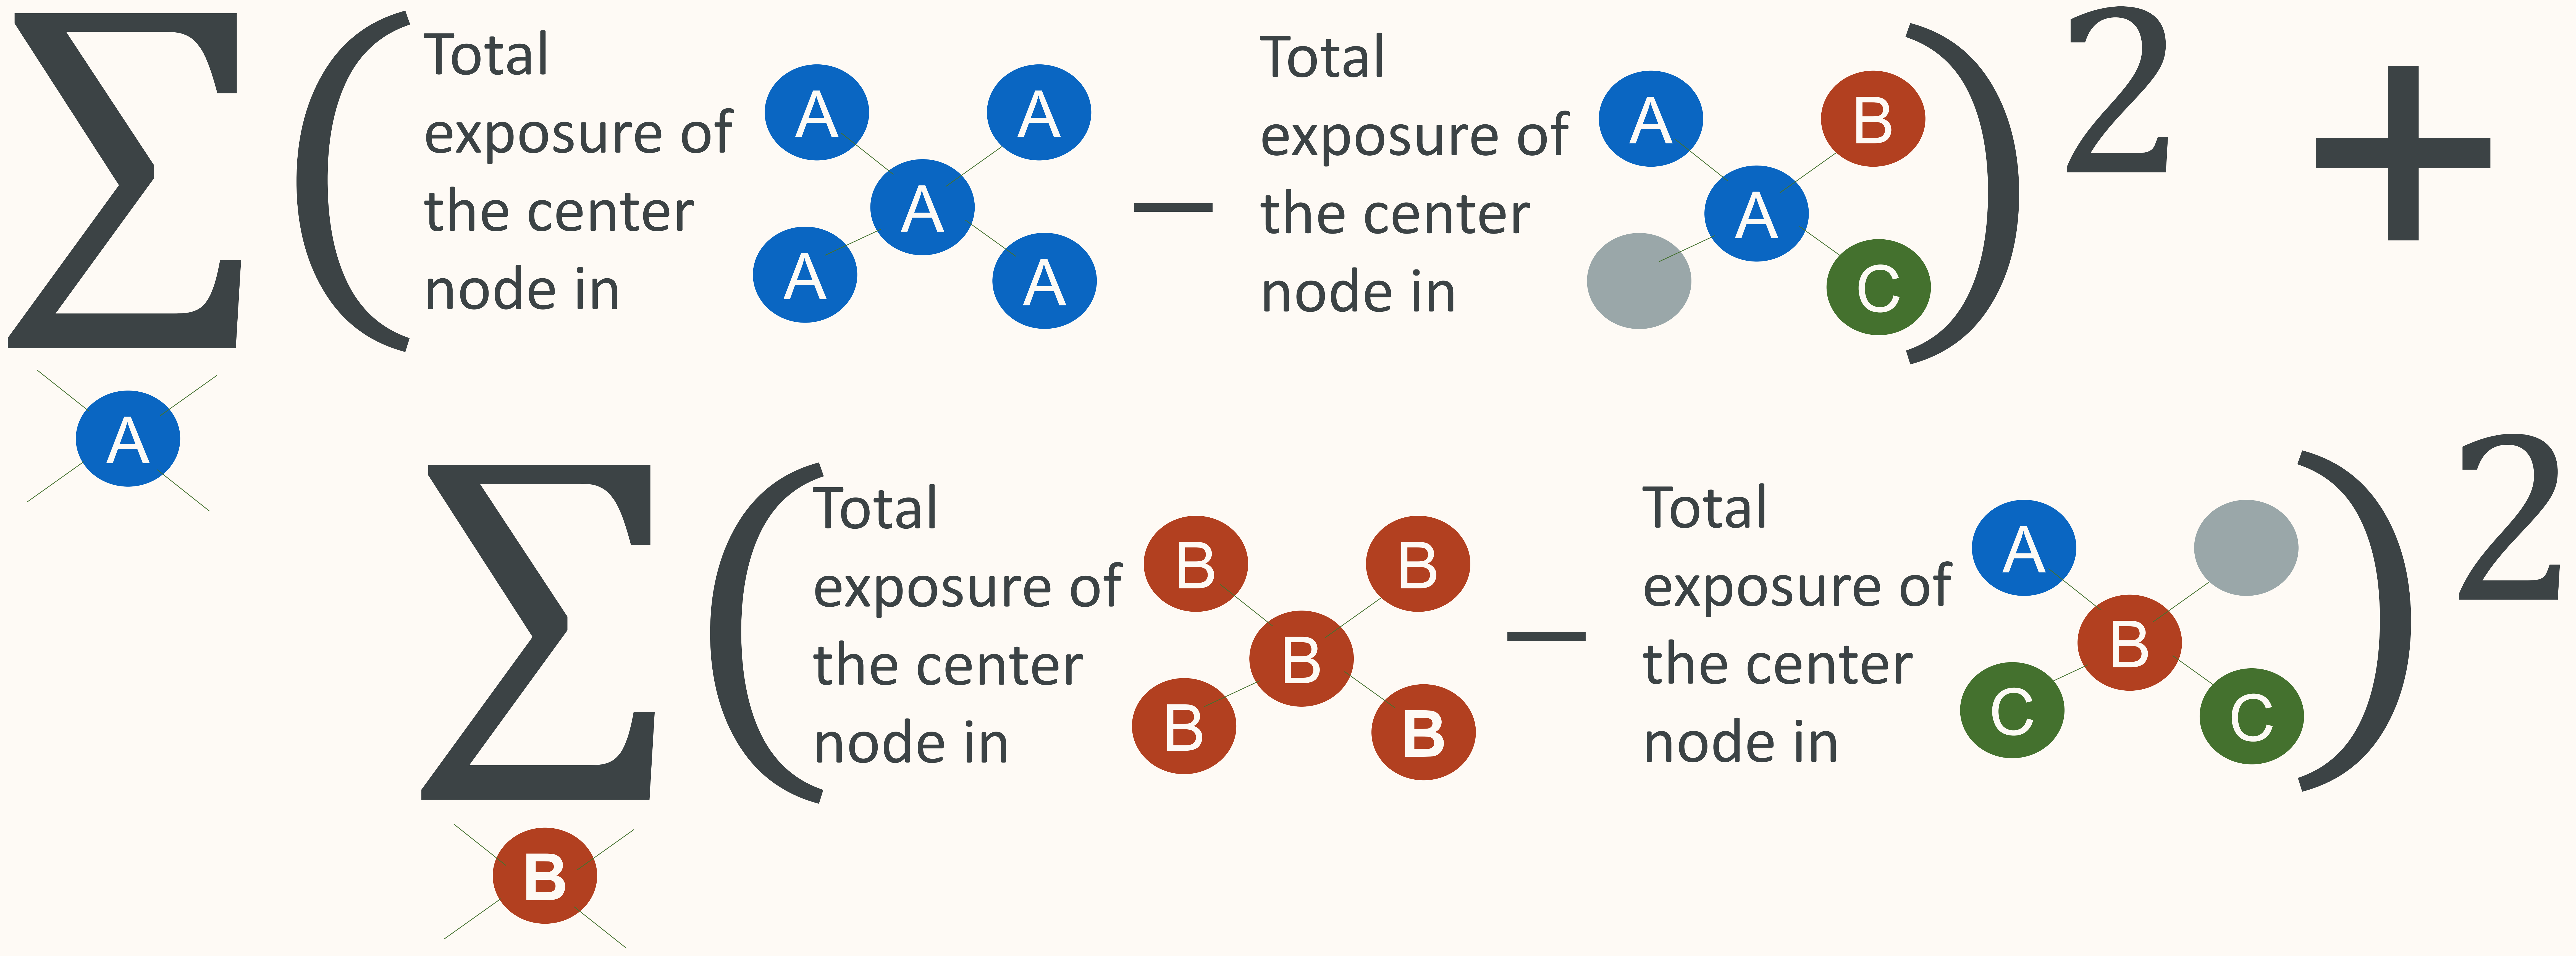

with respect to certain constraints controlling the risk of the experiment

Step 4: Run Experiment and Collect Data

Response:

$$\{Y(A)\} \text{ and } \{Y(B)\}$$

Observed Total Exposure:

$$\{X(A)\} \text{ and } \{X(B)\}$$

Expected Total Exposure:

$$\{X(A)\} \text{ and } \{X(B)\}$$

Step 5: Importance Sampling Correction

Average Treatment Effect:

$$\frac{1}{|B|}\sum_{B}\left\{Y(B)\,\frac{f_{expected}(X(B))}{f_{observed}(X(B))}\right\} - \frac{1}{|A|}\sum_{A}\left\{Y(A)\,\frac{f_{expected}(X(A))}{f_{observed}(X(A))}\right\}$$

OASIS: Optimal Allocation Strategy and Importance Sampling

1. Randomly assign A and B
2. Randomly choose additional nodes X
3. Solve a constrained optimization to assign C to X
4. Run experiment and collect data
5. Importance sampling correction

- P. Nandy, K. Basu, S. Chatterjee, Y. Tu. A/B Testing in Dense Large-Scale Networks: Design and Inference. arXiv preprint arXiv:1901.10505, 2019.

OASIS: Optimal Allocation Strategy and Importance Sampling

Implementation:

- Implemented for LinkedIn Feed experiments and using it for experiments targeted toward creator experience enhancement
- C = A $*$ boost factors (normalized to have each column sum equals 1)
- Solve optimization to get boost factors, where we control risk by setting a lower and an upper bound for boost factors
- Update boost factors regularly to handle dynamic network/treatment
- Correct bias with importance sampling

- P. Nandy, K. Basu, S. Chatterjee, Y. Tu. A/B Testing in Dense Large-Scale Networks: Design and Inference. arXiv preprint arXiv:1901.10505, 2019.

OASIS: Optimal Allocation Strategy and Importance Sampling

Validation:

- Theoretical results for robustness and verified in simulation
- No statistically significant result in A/A test
- Uniform p-values in A/A test
- Comparing with ego-cluster results (where an ego-cluster experiment is possible)
- Other special experimental designs for online validation

- P. Nandy, K. Basu, S. Chatterjee, Y. Tu. A/B Testing in Dense Large-Scale Networks: Design and Inference. arXiv preprint arXiv:1901.10505, 2019.

OASIS: Optimal Allocation Strategy and Importance Sampling

Pros:

- Theoretically sound under certain assumptions
- Works well for dense networks
- Can handle multiple treatments simultaneously
- Can handle dynamic networks and dynamic treatments
- Can control the risk of the experiment explicitly by adding constraints in the optimization

Cons:

- Relies on a number of assumptions
- Works only for a certain type of experiments

- P. Nandy, K. Basu, S. Chatterjee, Y. Tu. A/B Testing in Dense Large-Scale Networks: Design and Inference. arXiv preprint arXiv:1901.10505, 2019.

References

- P. Nandy, K. Basu, S. Chatterjee, Y. Tu. A/B Testing in Dense Large-Scale Networks: Design and Inference. arXiv preprint arXiv:1901.10505, 2019.

- G. Saint-Jacques, M. Varshney, J. Simpson, and Y. Xu. Using ego-clusters to measure network effects at LinkedIn. arXiv preprint arXiv:1903.08755, 2019.

Acknowledgement

- Kinjal Basu (co-author)
- Shaunak Chatterjee (co-author)
- Ye Tu (co-author)
- Guillaume Saint-Jacques
- Yav Bojinov
- Ya Xu
- Souvik Ghosh
- Ajith Muralidharan and many others

Thank you