Edge Dynamics and Opinion Polarization in Social Networks

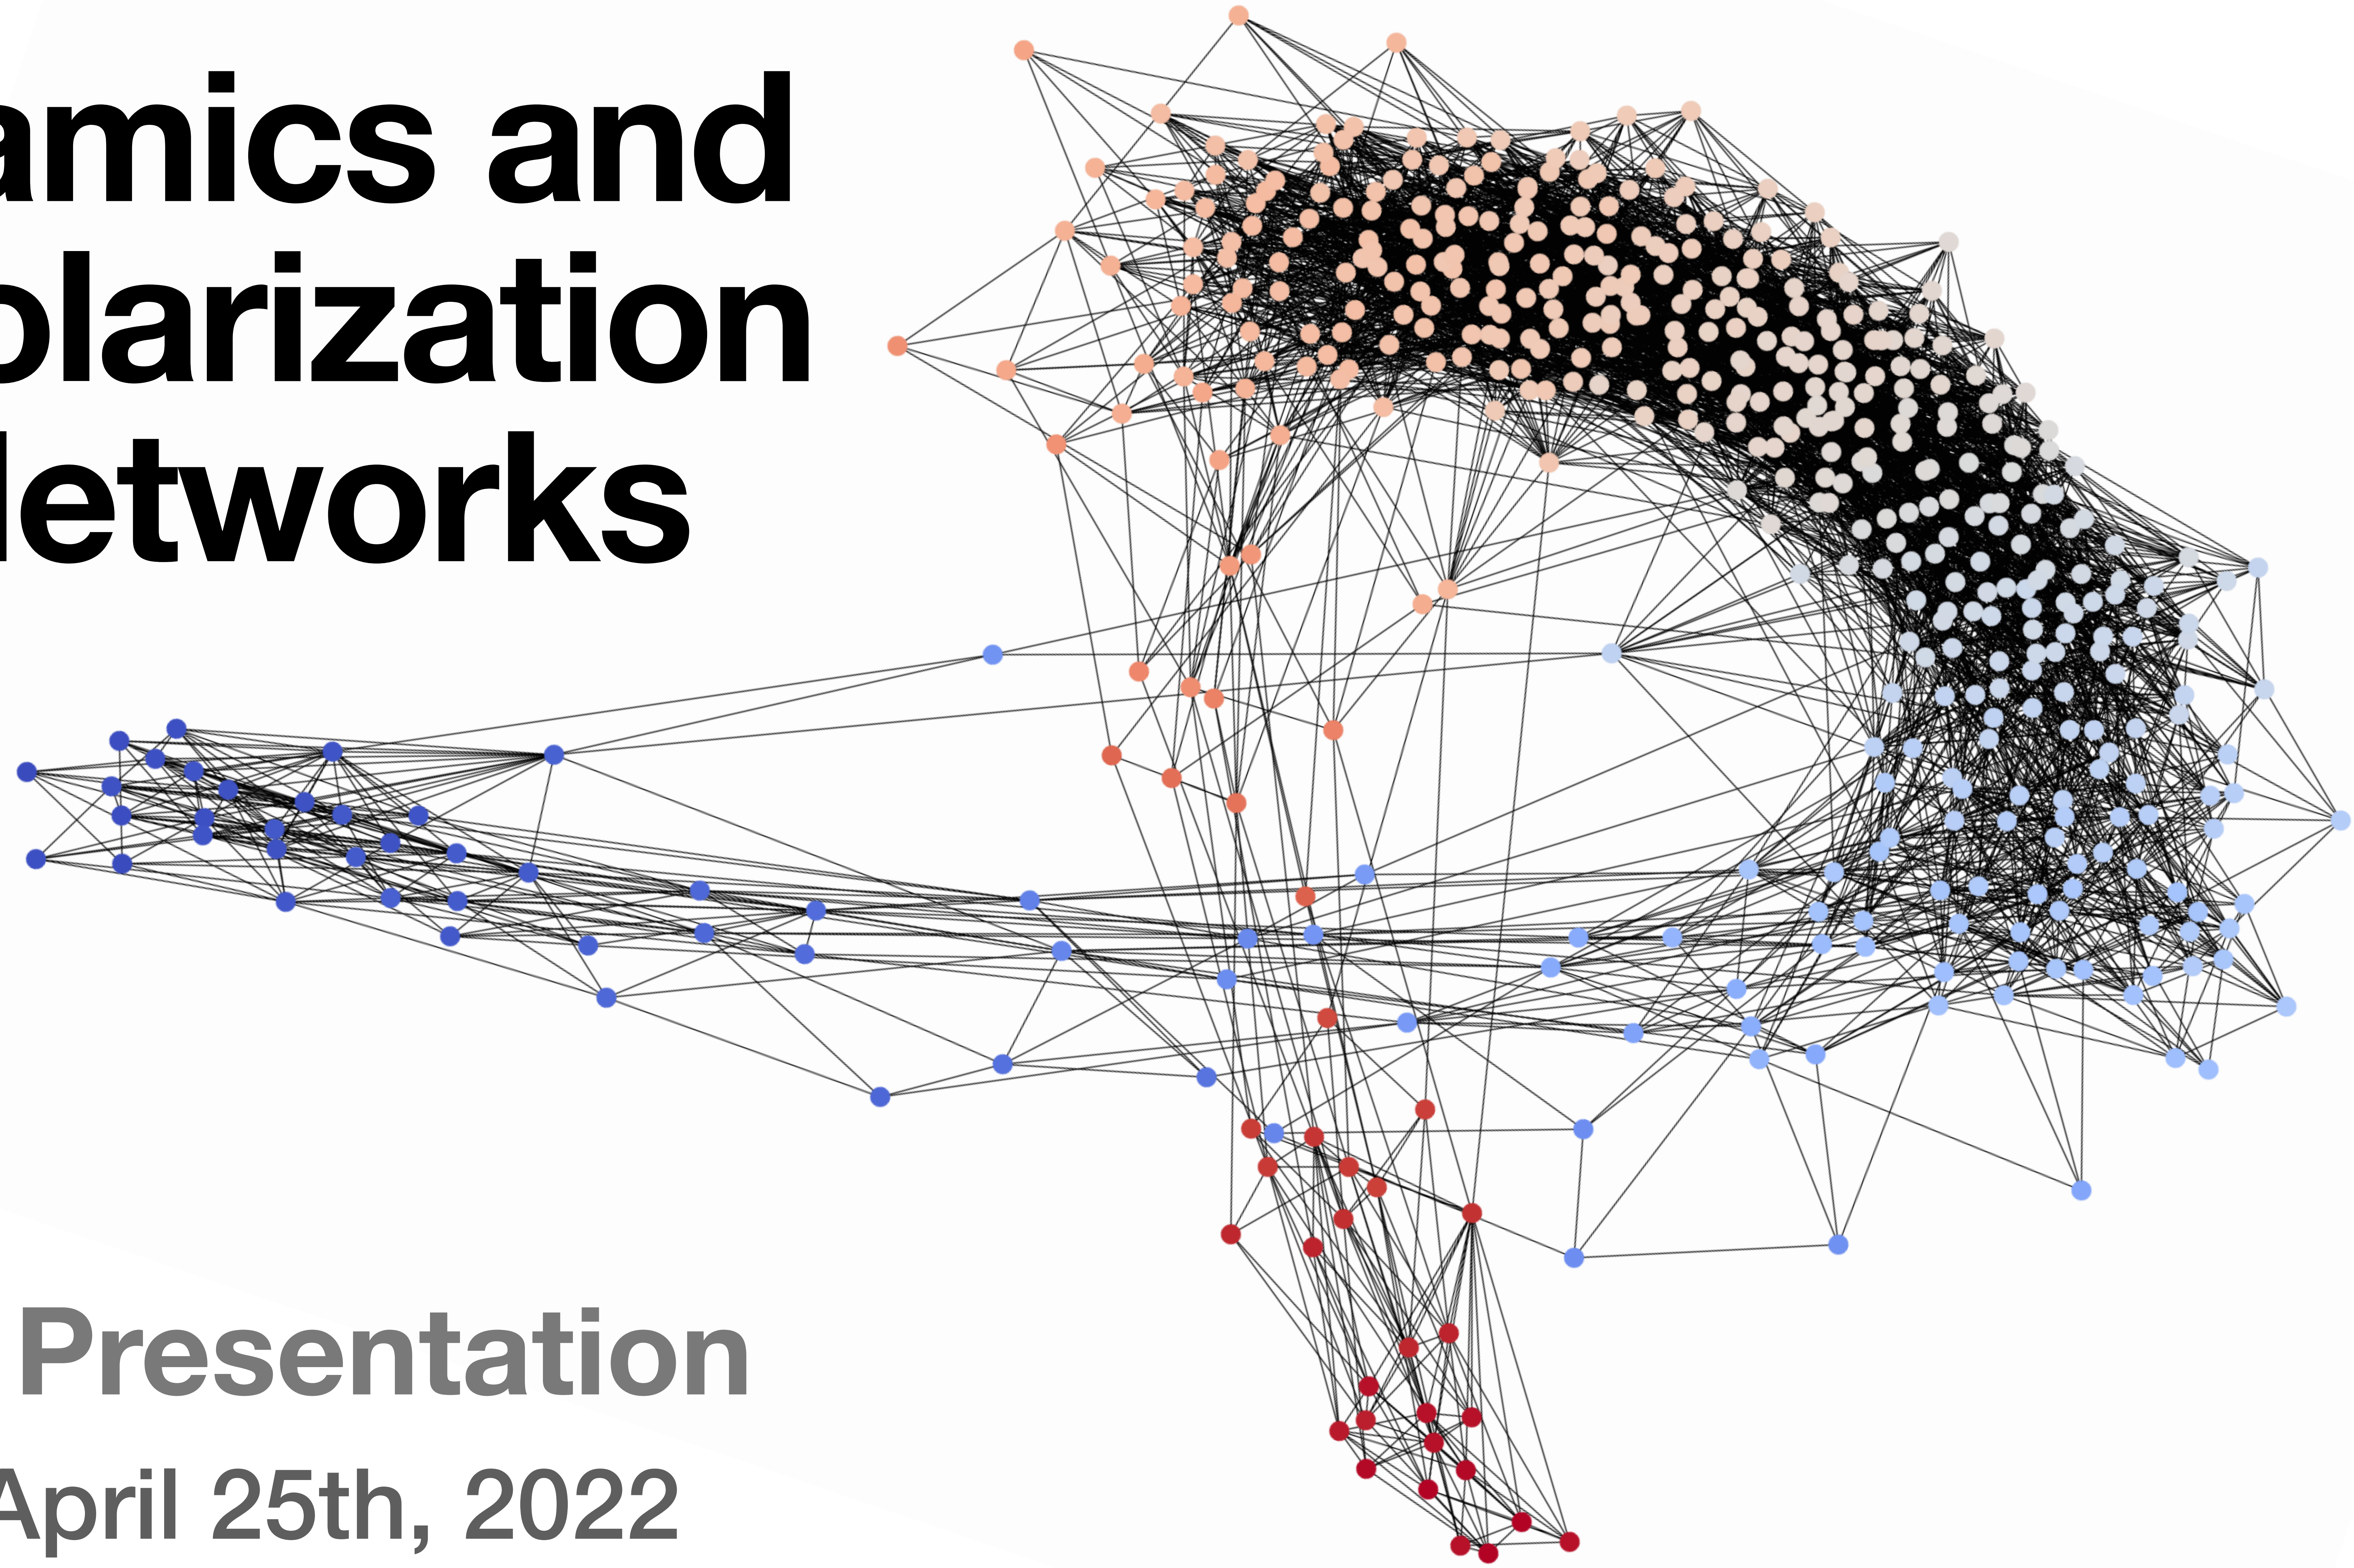

Honors Thesis Presentation

Adam Lechowicz, April 25th, 2022

University of Massachusetts

Overview

1. Opinion & Edge Dynamics Model
2. Characteristic Behavior
3. Theoretical Understanding
4. Verifying the Model's Realism
5. Conclusions / Future Steps

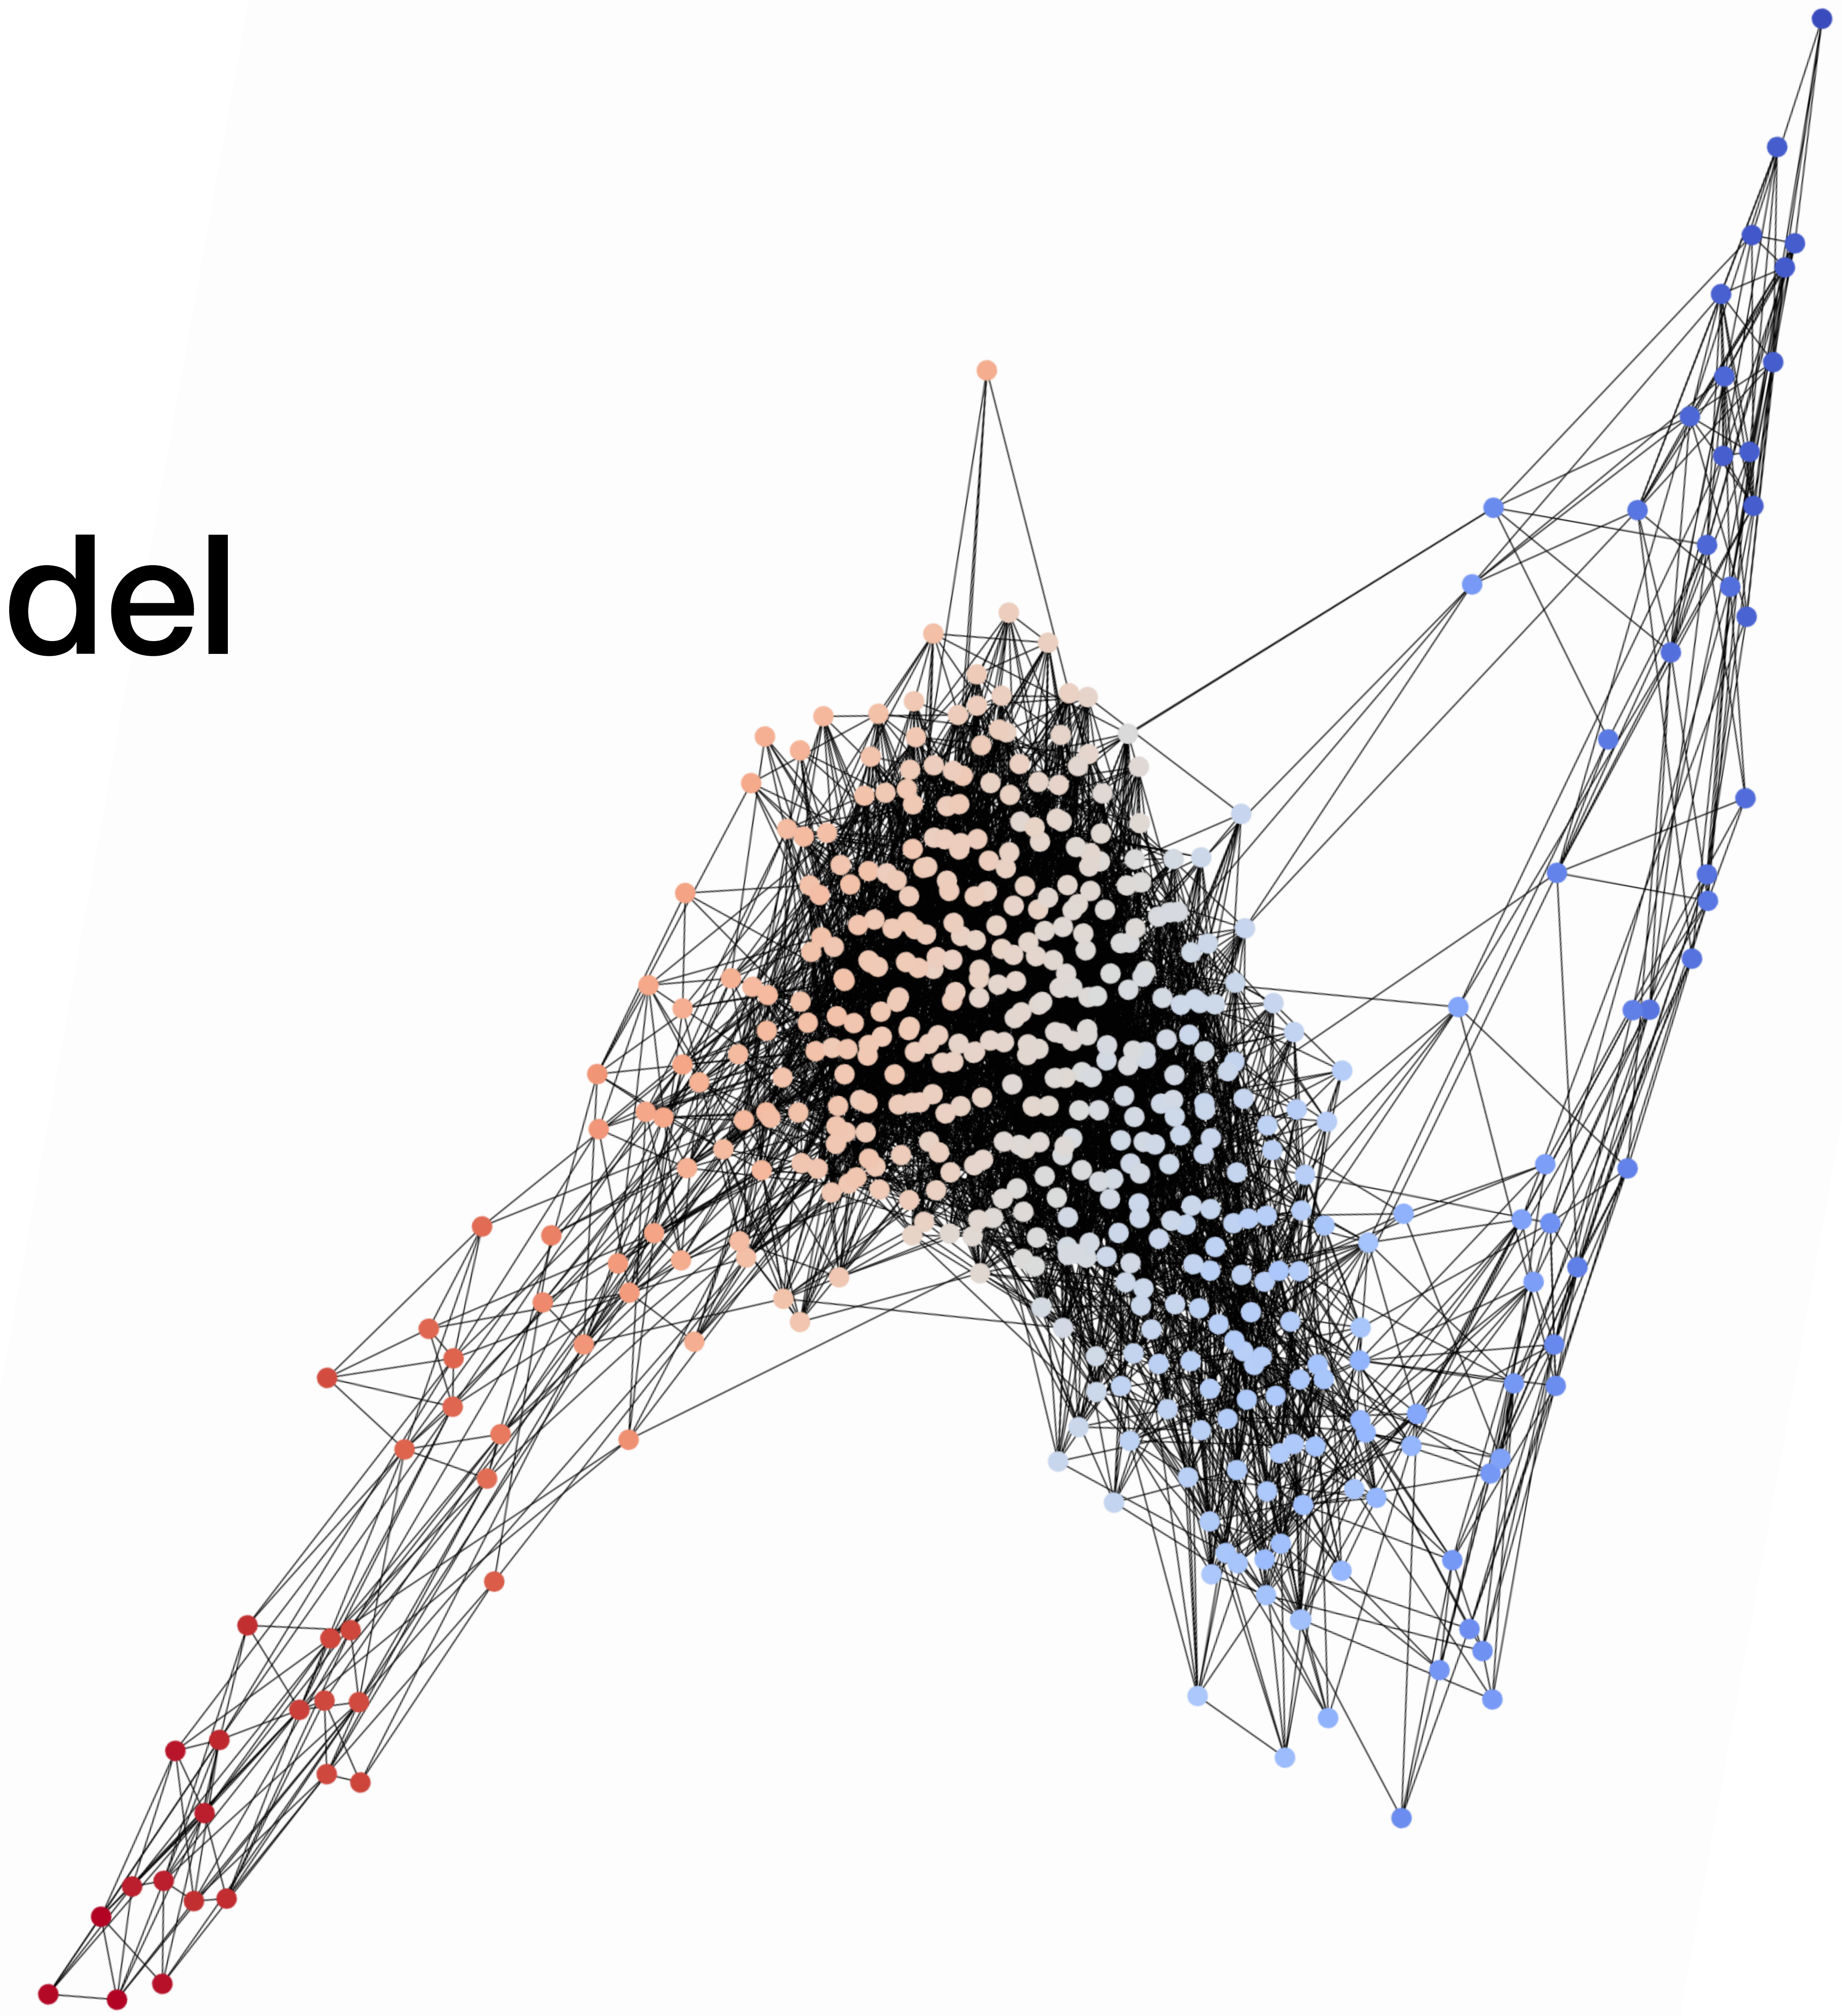

Representing a Social Network as a Graph

- We have some undirected graph $G = (V, E)$.
- In the synthetic setting (i.e., not a real social network), we use generative graph models:
  - $G(n, p)$ random graph model [Erdős and Rényi ’59]
  - “Preferential attachment” model [Barabási and Albert ’99]
  - “Small world” model [Newman, Watts, and Strogatz ’99]
- In the theoretical setting we often focus on standard E-R graphs, but others are useful to verify our model.

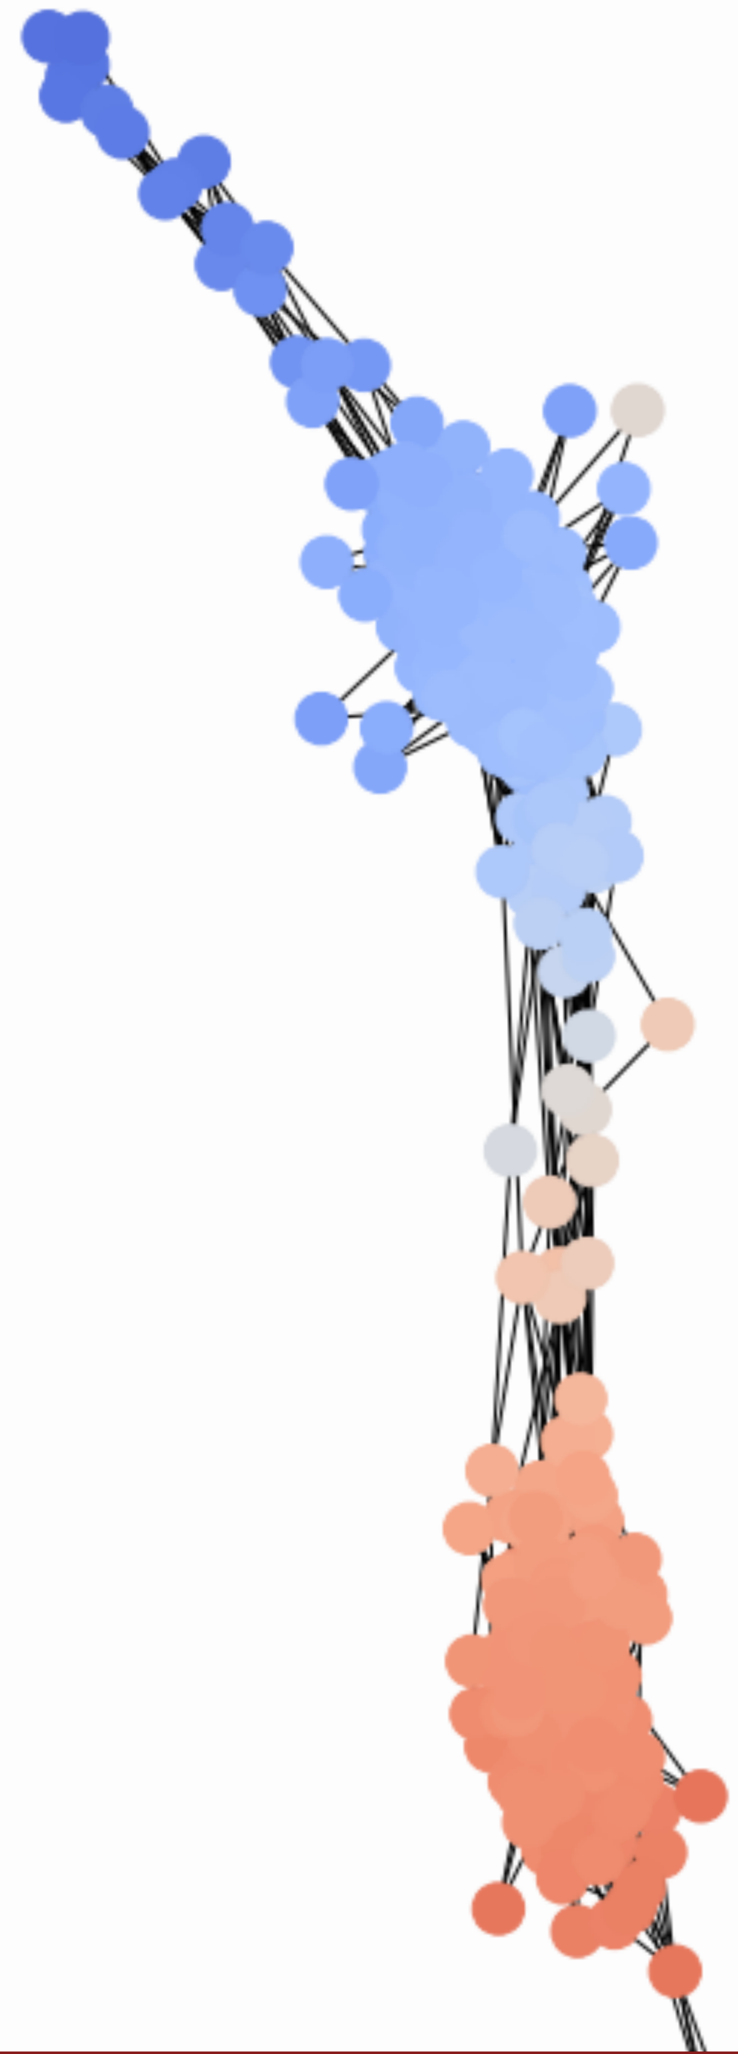

Opinions in the Graph

- Each node has an *opinion,* which is assigned using the F-J opinion dynamics model [Friedkin-Johnsen '91]
  - Each node has an **innate opinion (fixed)**, and an **expressed opinion (updated over time).**
- We set innate opinions to be either *discrete* or *continuous*:

$$\text{discrete} \in \{-1,1\}\ ,\ \ \text{continuous} \in [-1,1]$$

- This graph evolves over time. At each time step, we manipulate roughly 10% of the edges in the graph using *our edge dynamics model*.

Graph Evolution Over Time (Edge Dynamics)

- We approximate *confirmation bias* by stochastically removing disagreeable edges in the graph.

  - i.e. — if edge is *more* disagreeable, it is *more* likely to be removed

- For each removed edge, we add a *friend-of-friend edge* — approximates the design of some recommender systems.

- **Intuition:** from a sociological perspective, despite differences in opinion, many connections in the real-world remain strong. We set some percentage of *fixed edges,* which *cannot be deleted.*

Notation / Definitions

- Let $\vec{s}$ denote the *innate opinion vector*, with length n (num. of nodes)
- $\mathbf{I}$ is the $n \times n$ identity matrix
- L is the graph Laplacian, defined as $D - A$
- F-J opinion update, where $\vec{z}$ denotes *expressed opinion vector*:

$$\vec{z}(i) = \frac{\vec{s}(i) + \sum_j w_{ij}\vec{z}(j)}{1 + \sum_j w_{ij}} \quad \longrightarrow \quad \vec{z} = (\mathbf{I} + L)^{-1}\vec{s}$$

Notation / Definitions (cont.)

- **Polarization**, where $\vec{z}$ denotes *expressed opinion vector*:

$$P(L, \vec{s}) \;=\; \vec{z}^T \vec{z} \;=\; \vec{s}^T (\mathbf{I} + L)^{-2} \vec{s}$$

  - Note: this works because innate opinions are mean-0 vectors

- **Disagreement**, where $\vec{z}$ denotes *expressed opinion vector*:

$$D(L, \vec{s}) \;=\; \sum_{i,j} w_{i,j} \cdot (\vec{z}(i) - \vec{z}(j))^2 \;=\; \vec{z}^T L \vec{z} \;=\; \vec{s}^T (\mathbf{I} + L)^{-1} \, L \, (\mathbf{I} + L)^{-1} \vec{s}$$

“Standard” Experimental Setup

- Start by presenting the setup and results for a “benchmark” edge dynamics experiment:
  - Erdős–Rényi random graph with $n = 1000$
  - Fixed Edge percentage $\mathbf{p}_f \in [0\,\%, 5\,\%, 10\,\%, 15\,\%, 25\%]$
  - 500 “time steps” (iterations) of edge dynamics
  - Connection probability $p = 25/n$
  - friend-of-friend recommendations
  - “*Confirmation bias*” disagreeable edge deletion
  - Continuous innate opinions assigned to each node on [-1,1]
  - 10% of edges removed/added each time step.

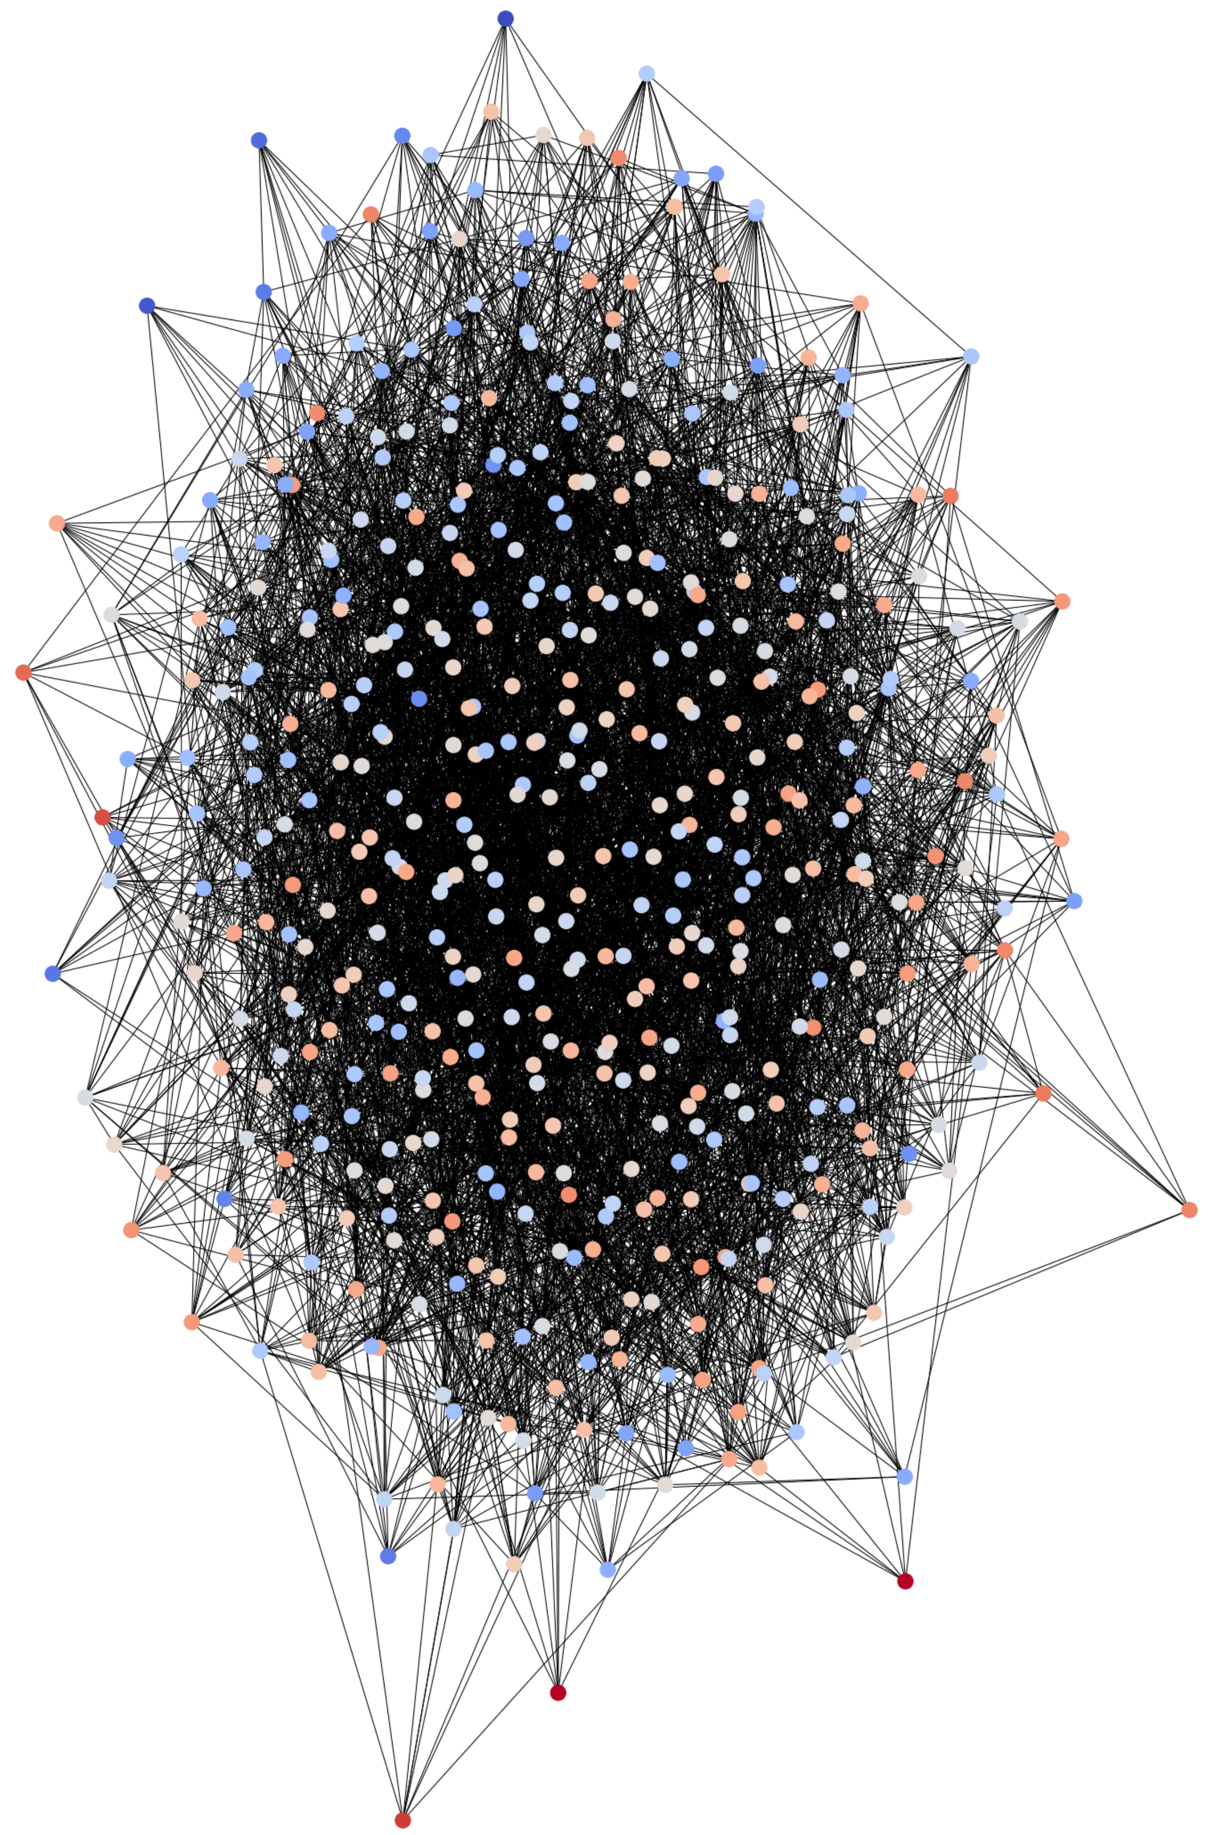

Characteristic Results

$$P(L,\vec{s}) = \vec{z}^T\vec{z} = \vec{s}^T(\mathbf{I}+L)^{-2}\vec{s} \qquad D(L,\vec{s}) = \vec{z}^T L\vec{z}$$

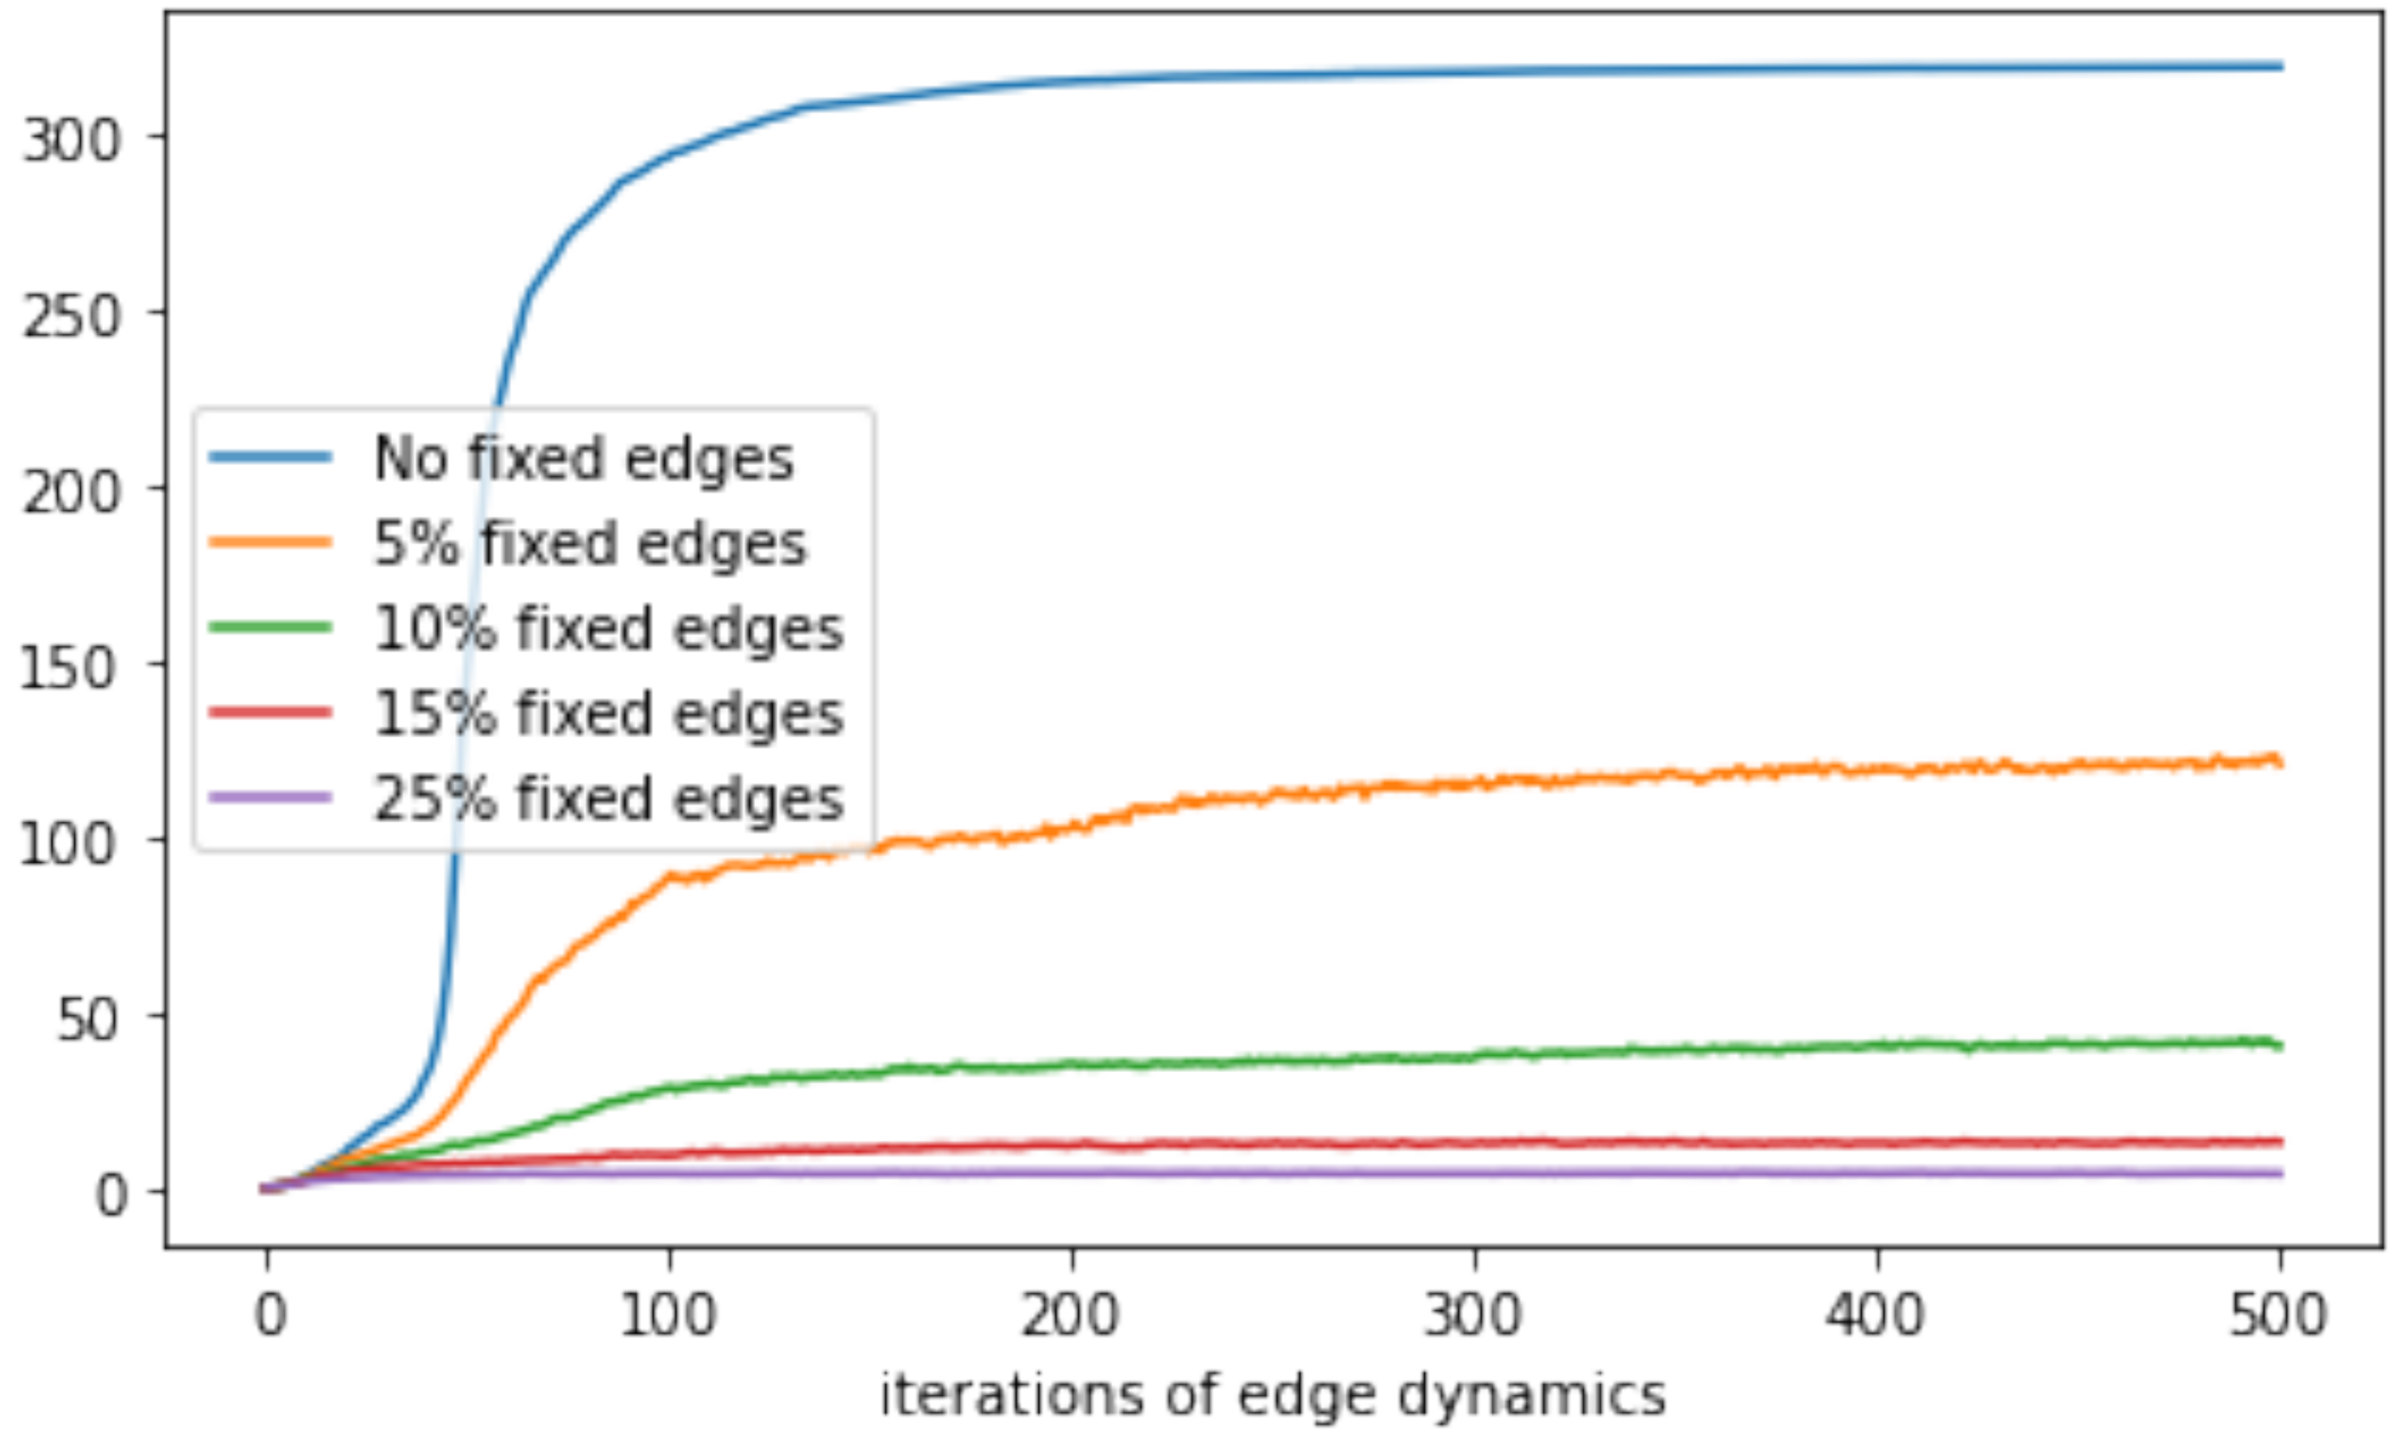

Polarization over time for standard experiment

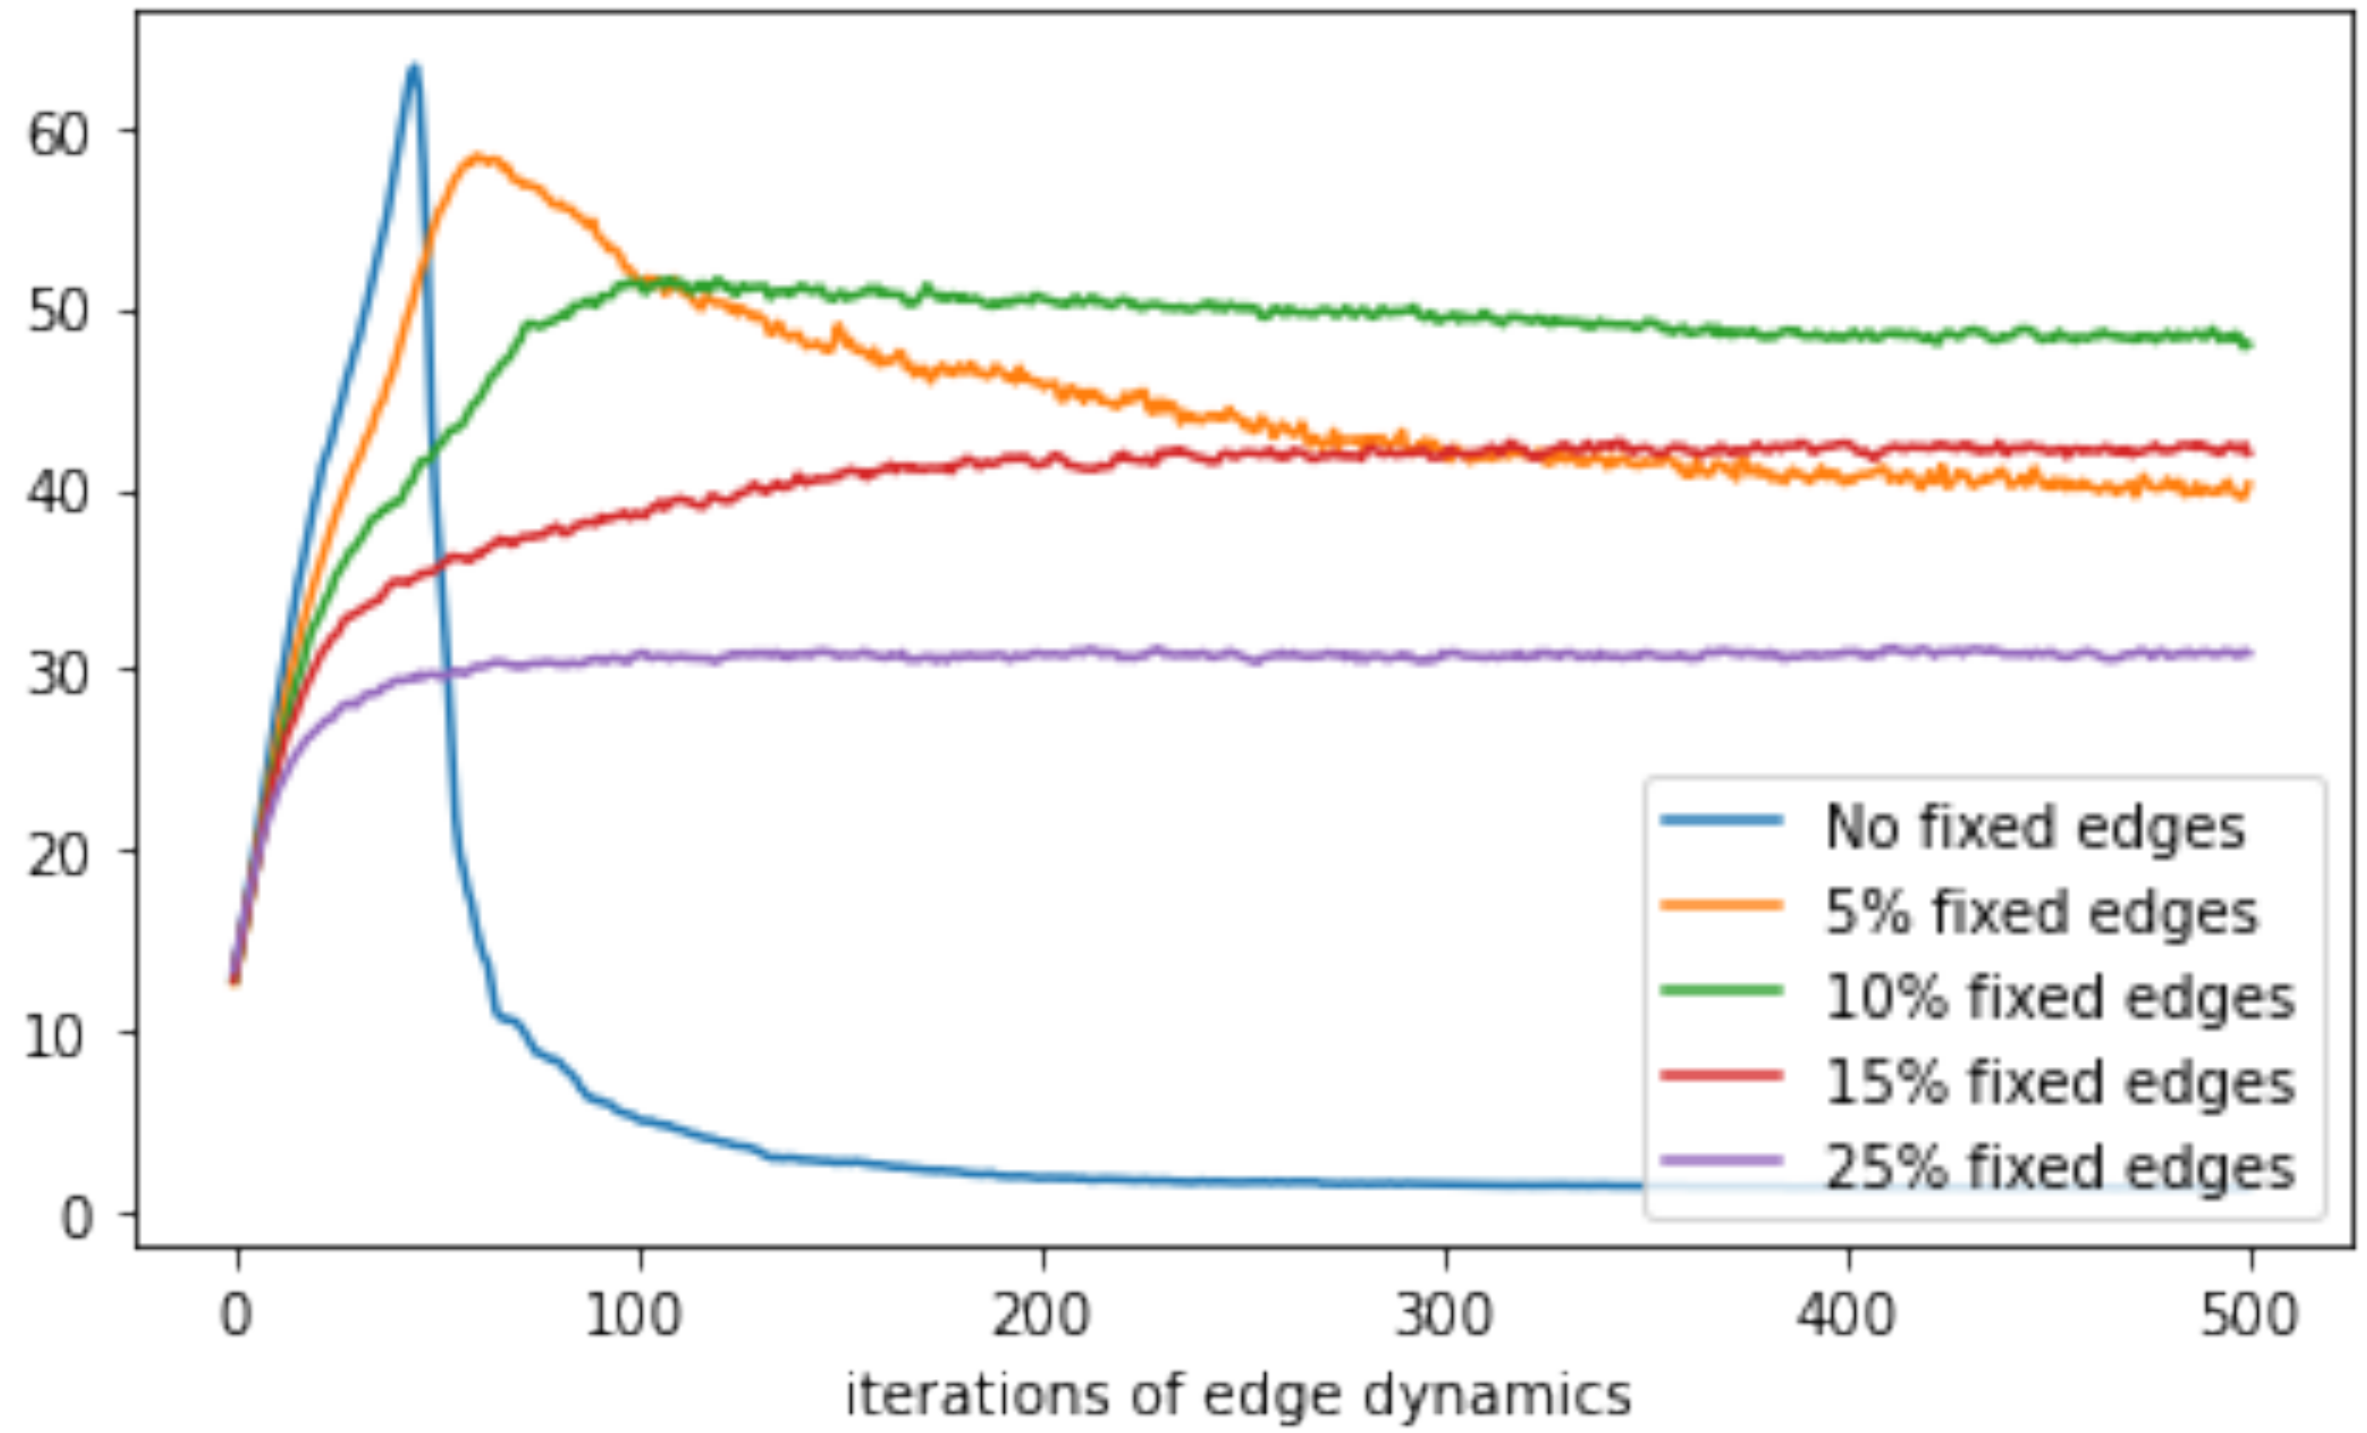

Disagreement over time for standard experiment

Visual / Intuitive Understanding

- We show the evolution of the standard experiment using NetworkX visualization tools.

- The color of each node is set according to its *expressed* opinion from the F-J equilibrium

  - [-1, … , 0 , … , 1] maps to [blue … white … red]

- Edges are thin black lines that connect nodes. Nodes are clustered near neighboring nodes.

- Show for both 0% fixed edges and 10% fixed edges.

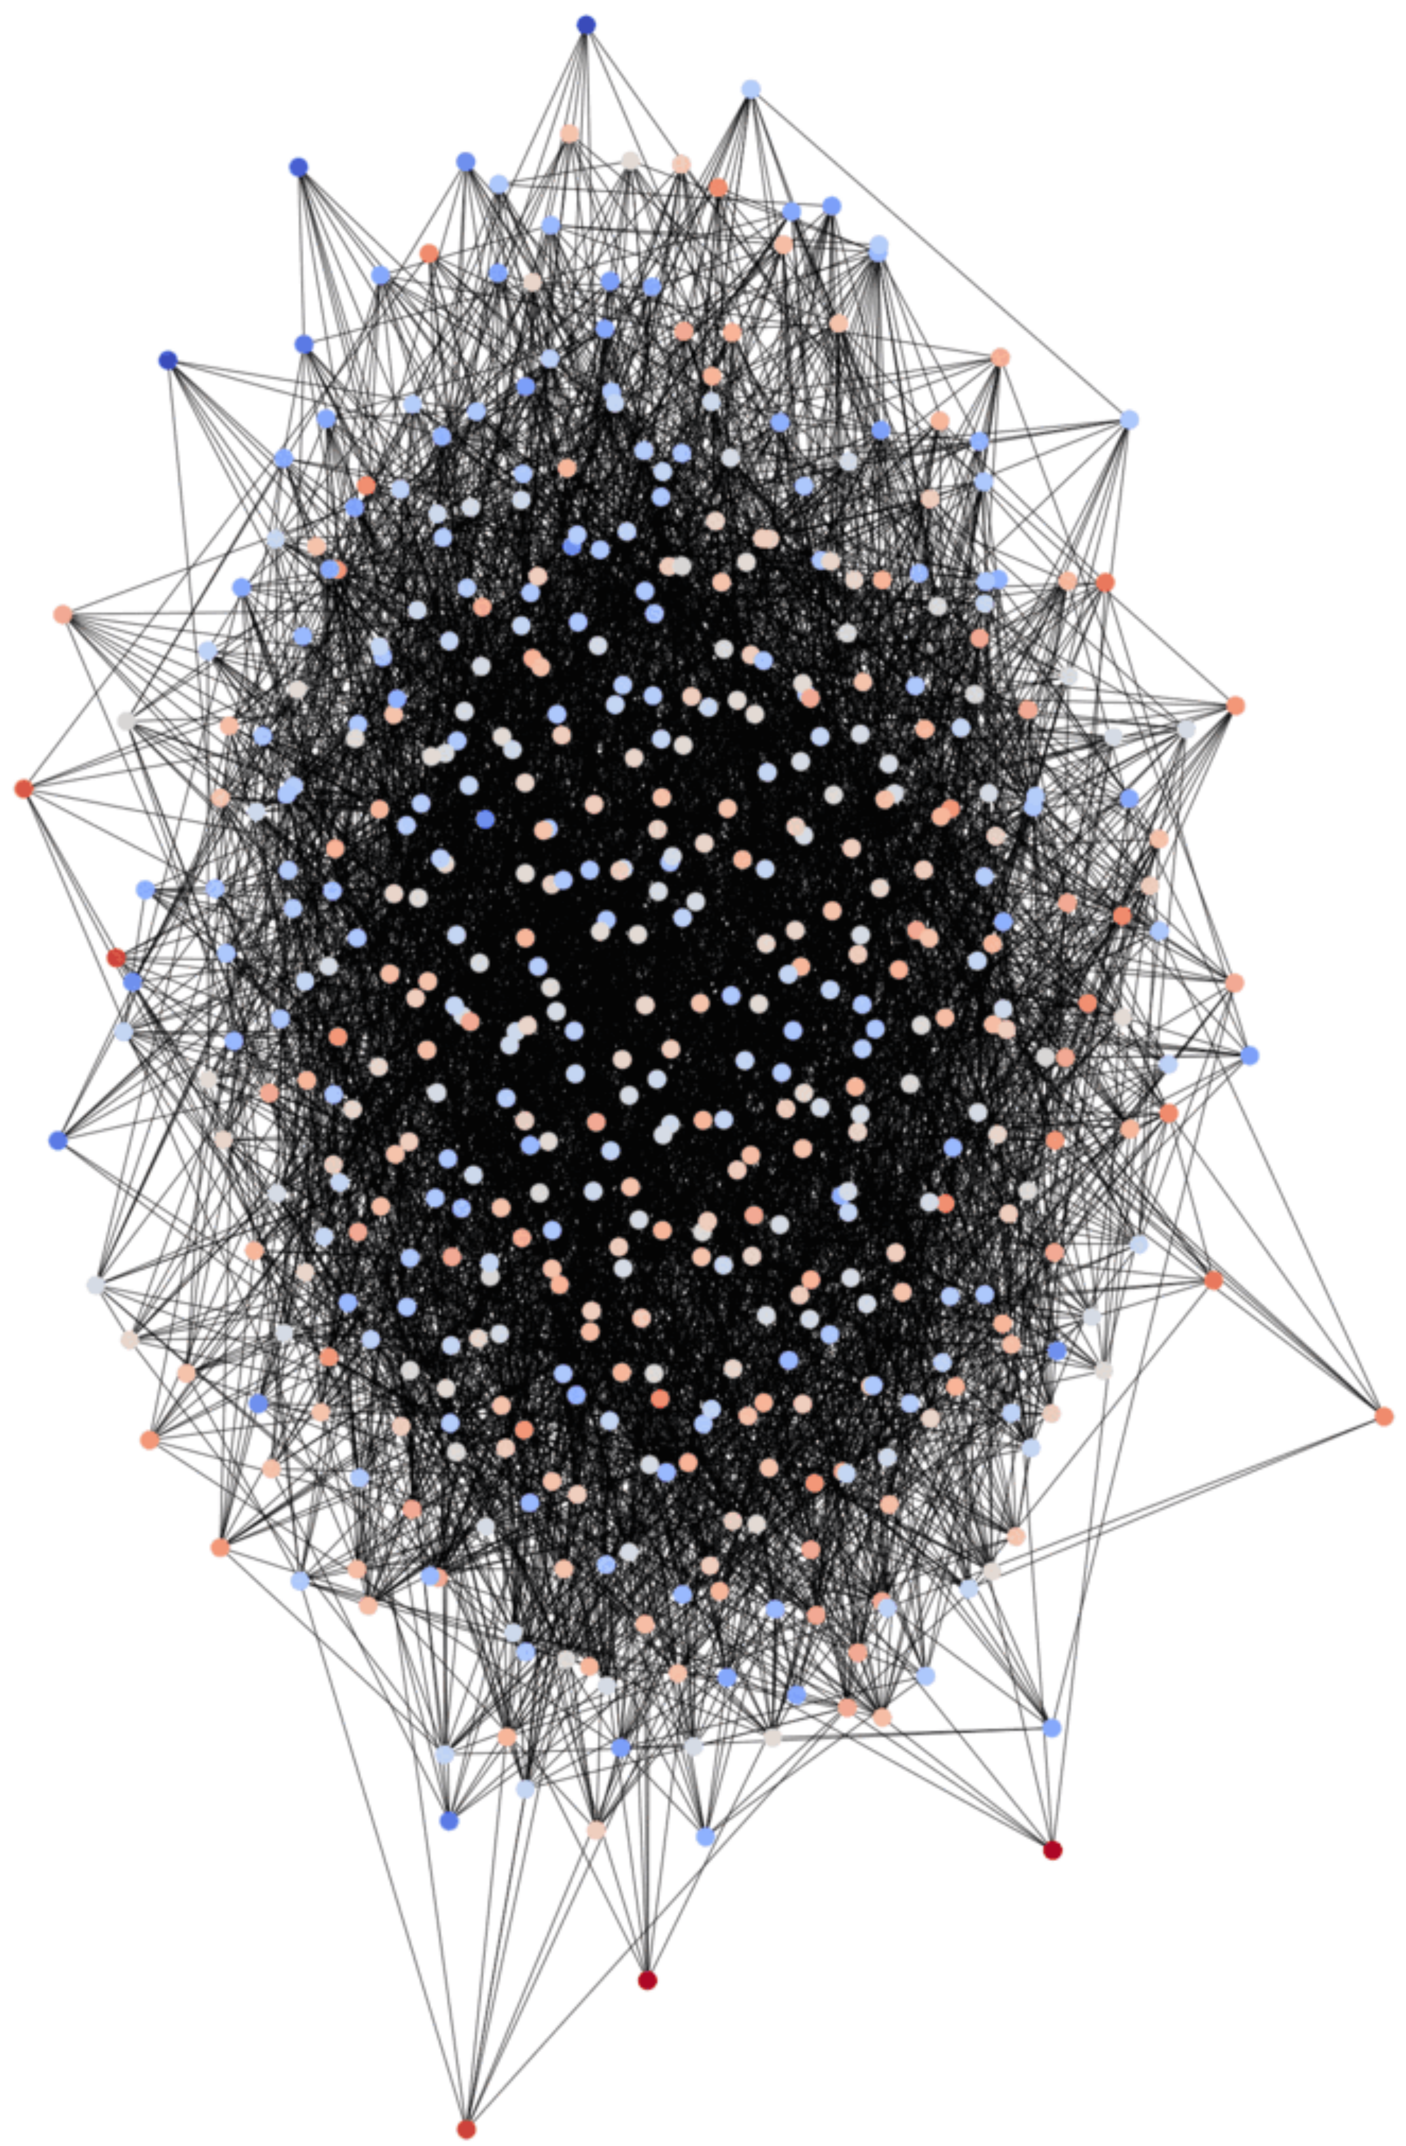

No Fixed Edges, E-R graph, $n = 1000$, standard experiment

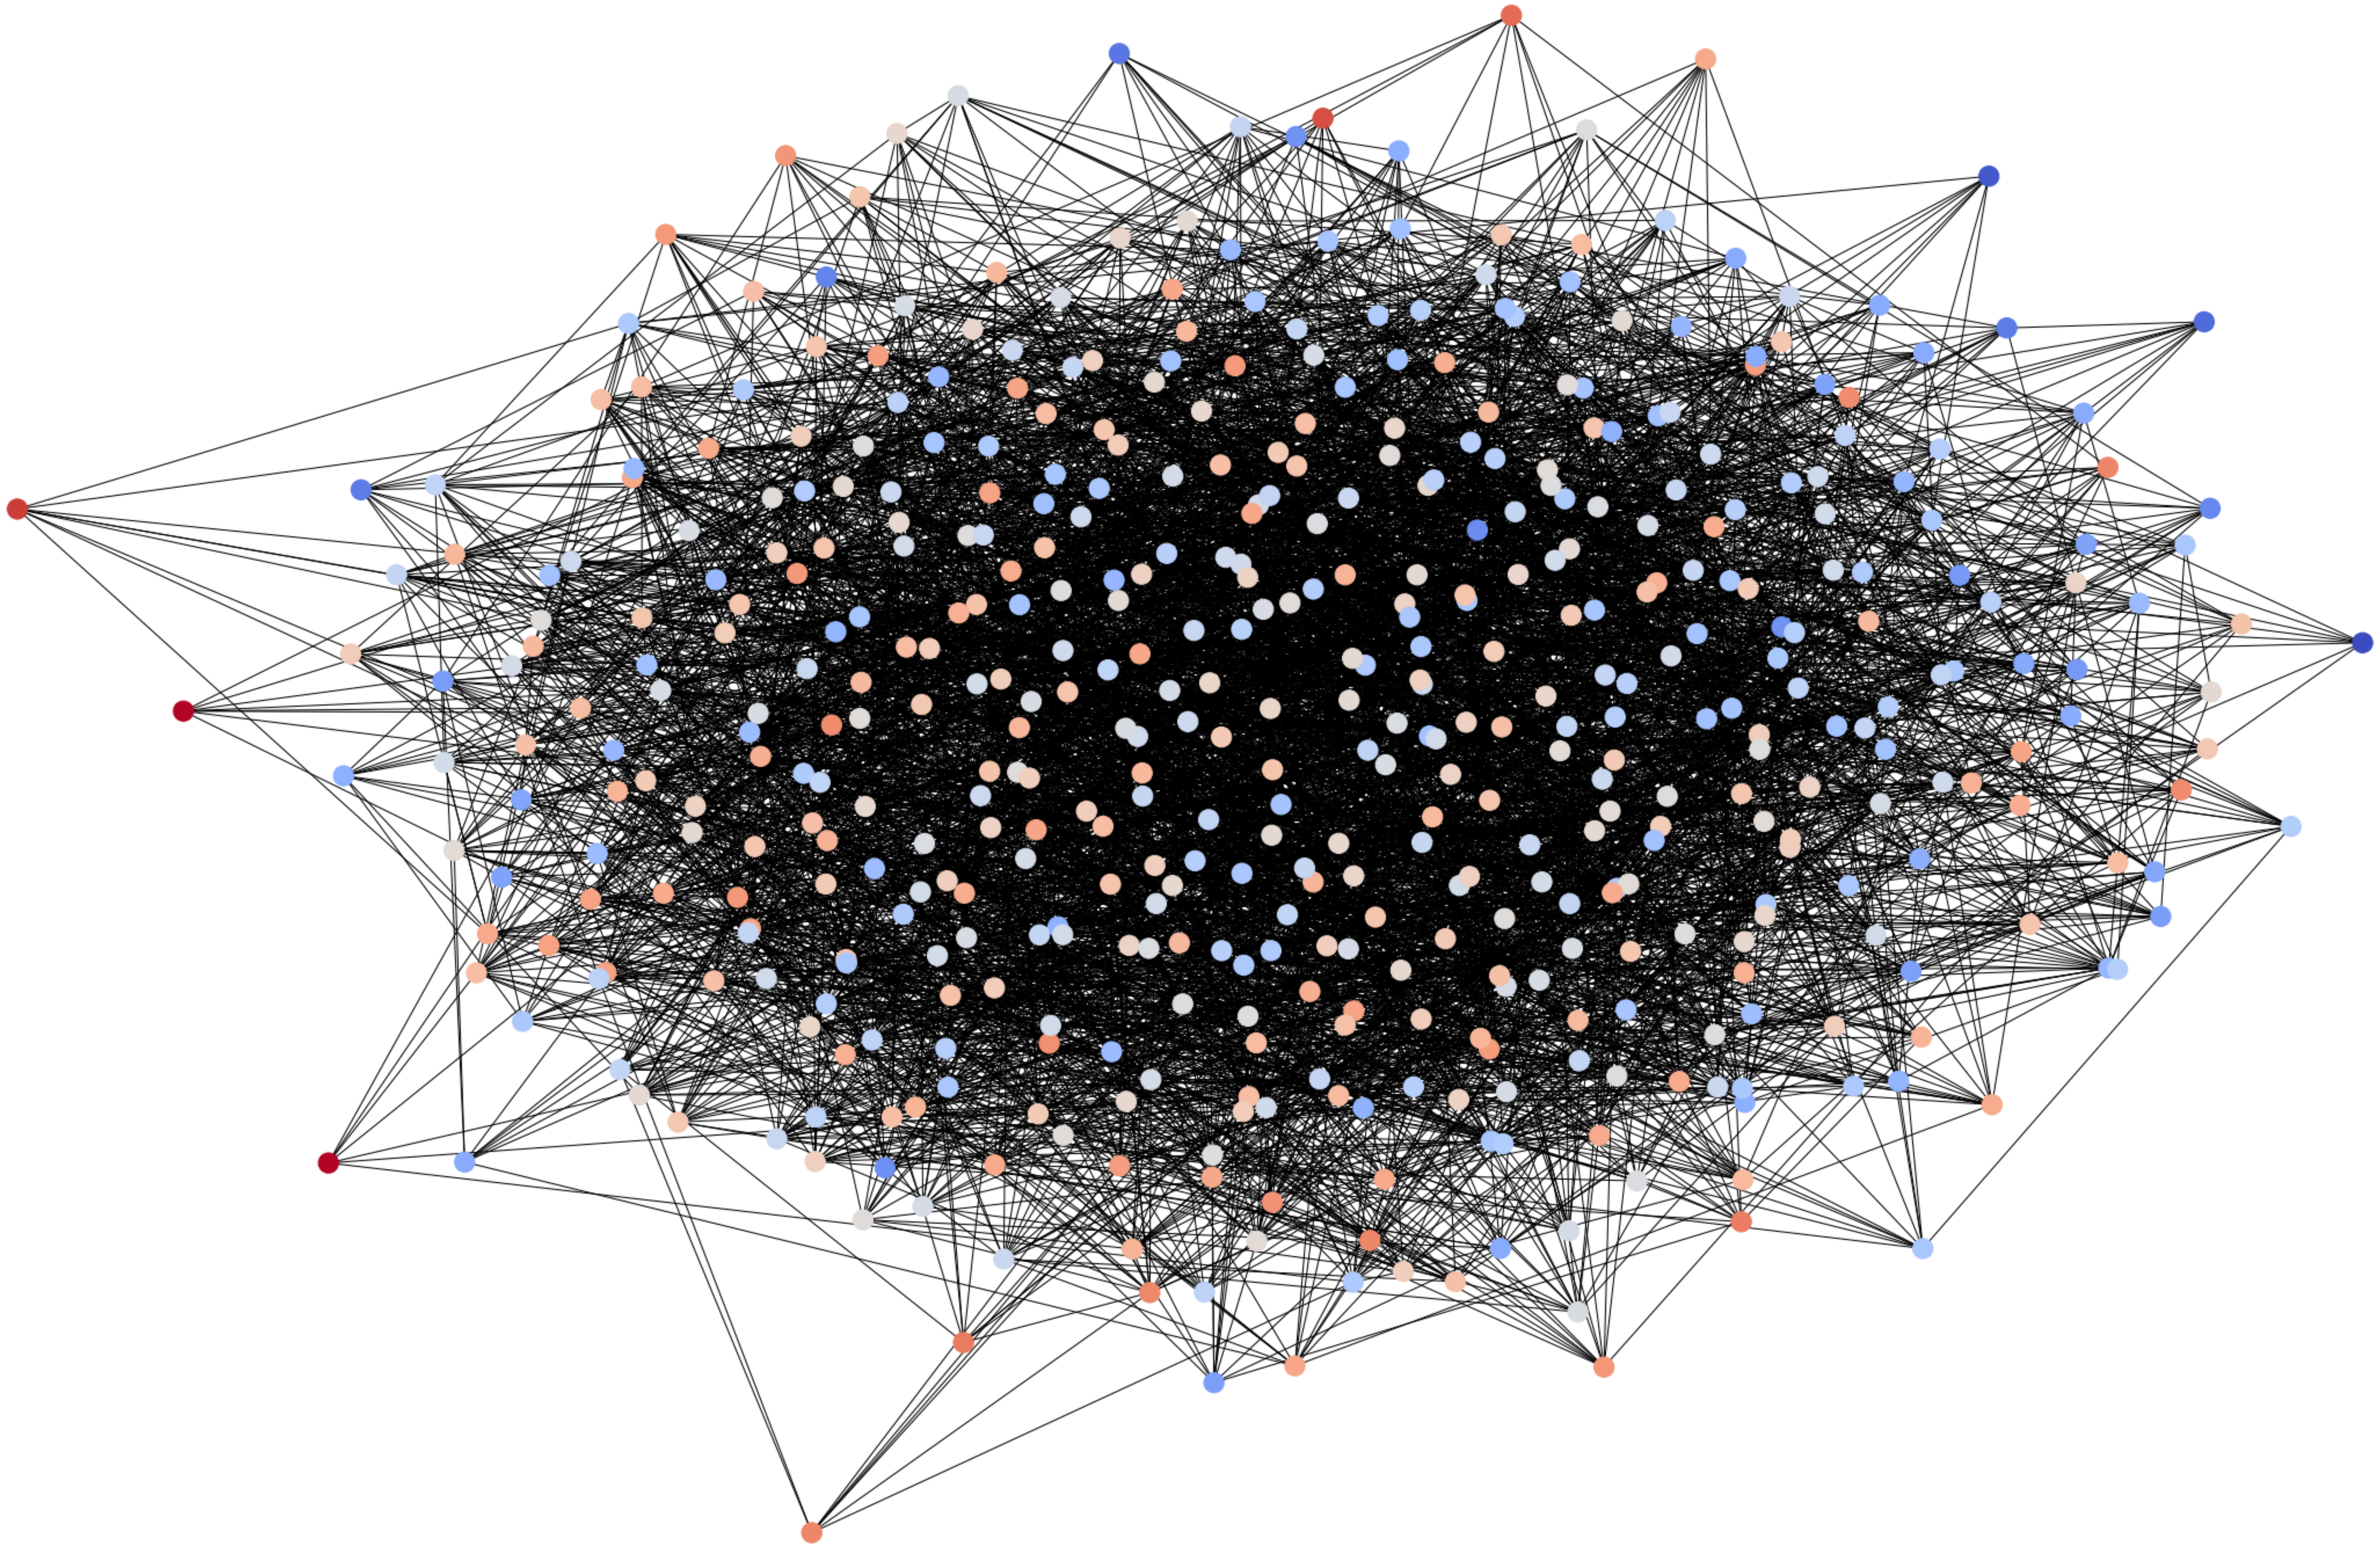

$t = 0$

No Fixed Edges, E-R graph, $n = 1000$, standard experiment

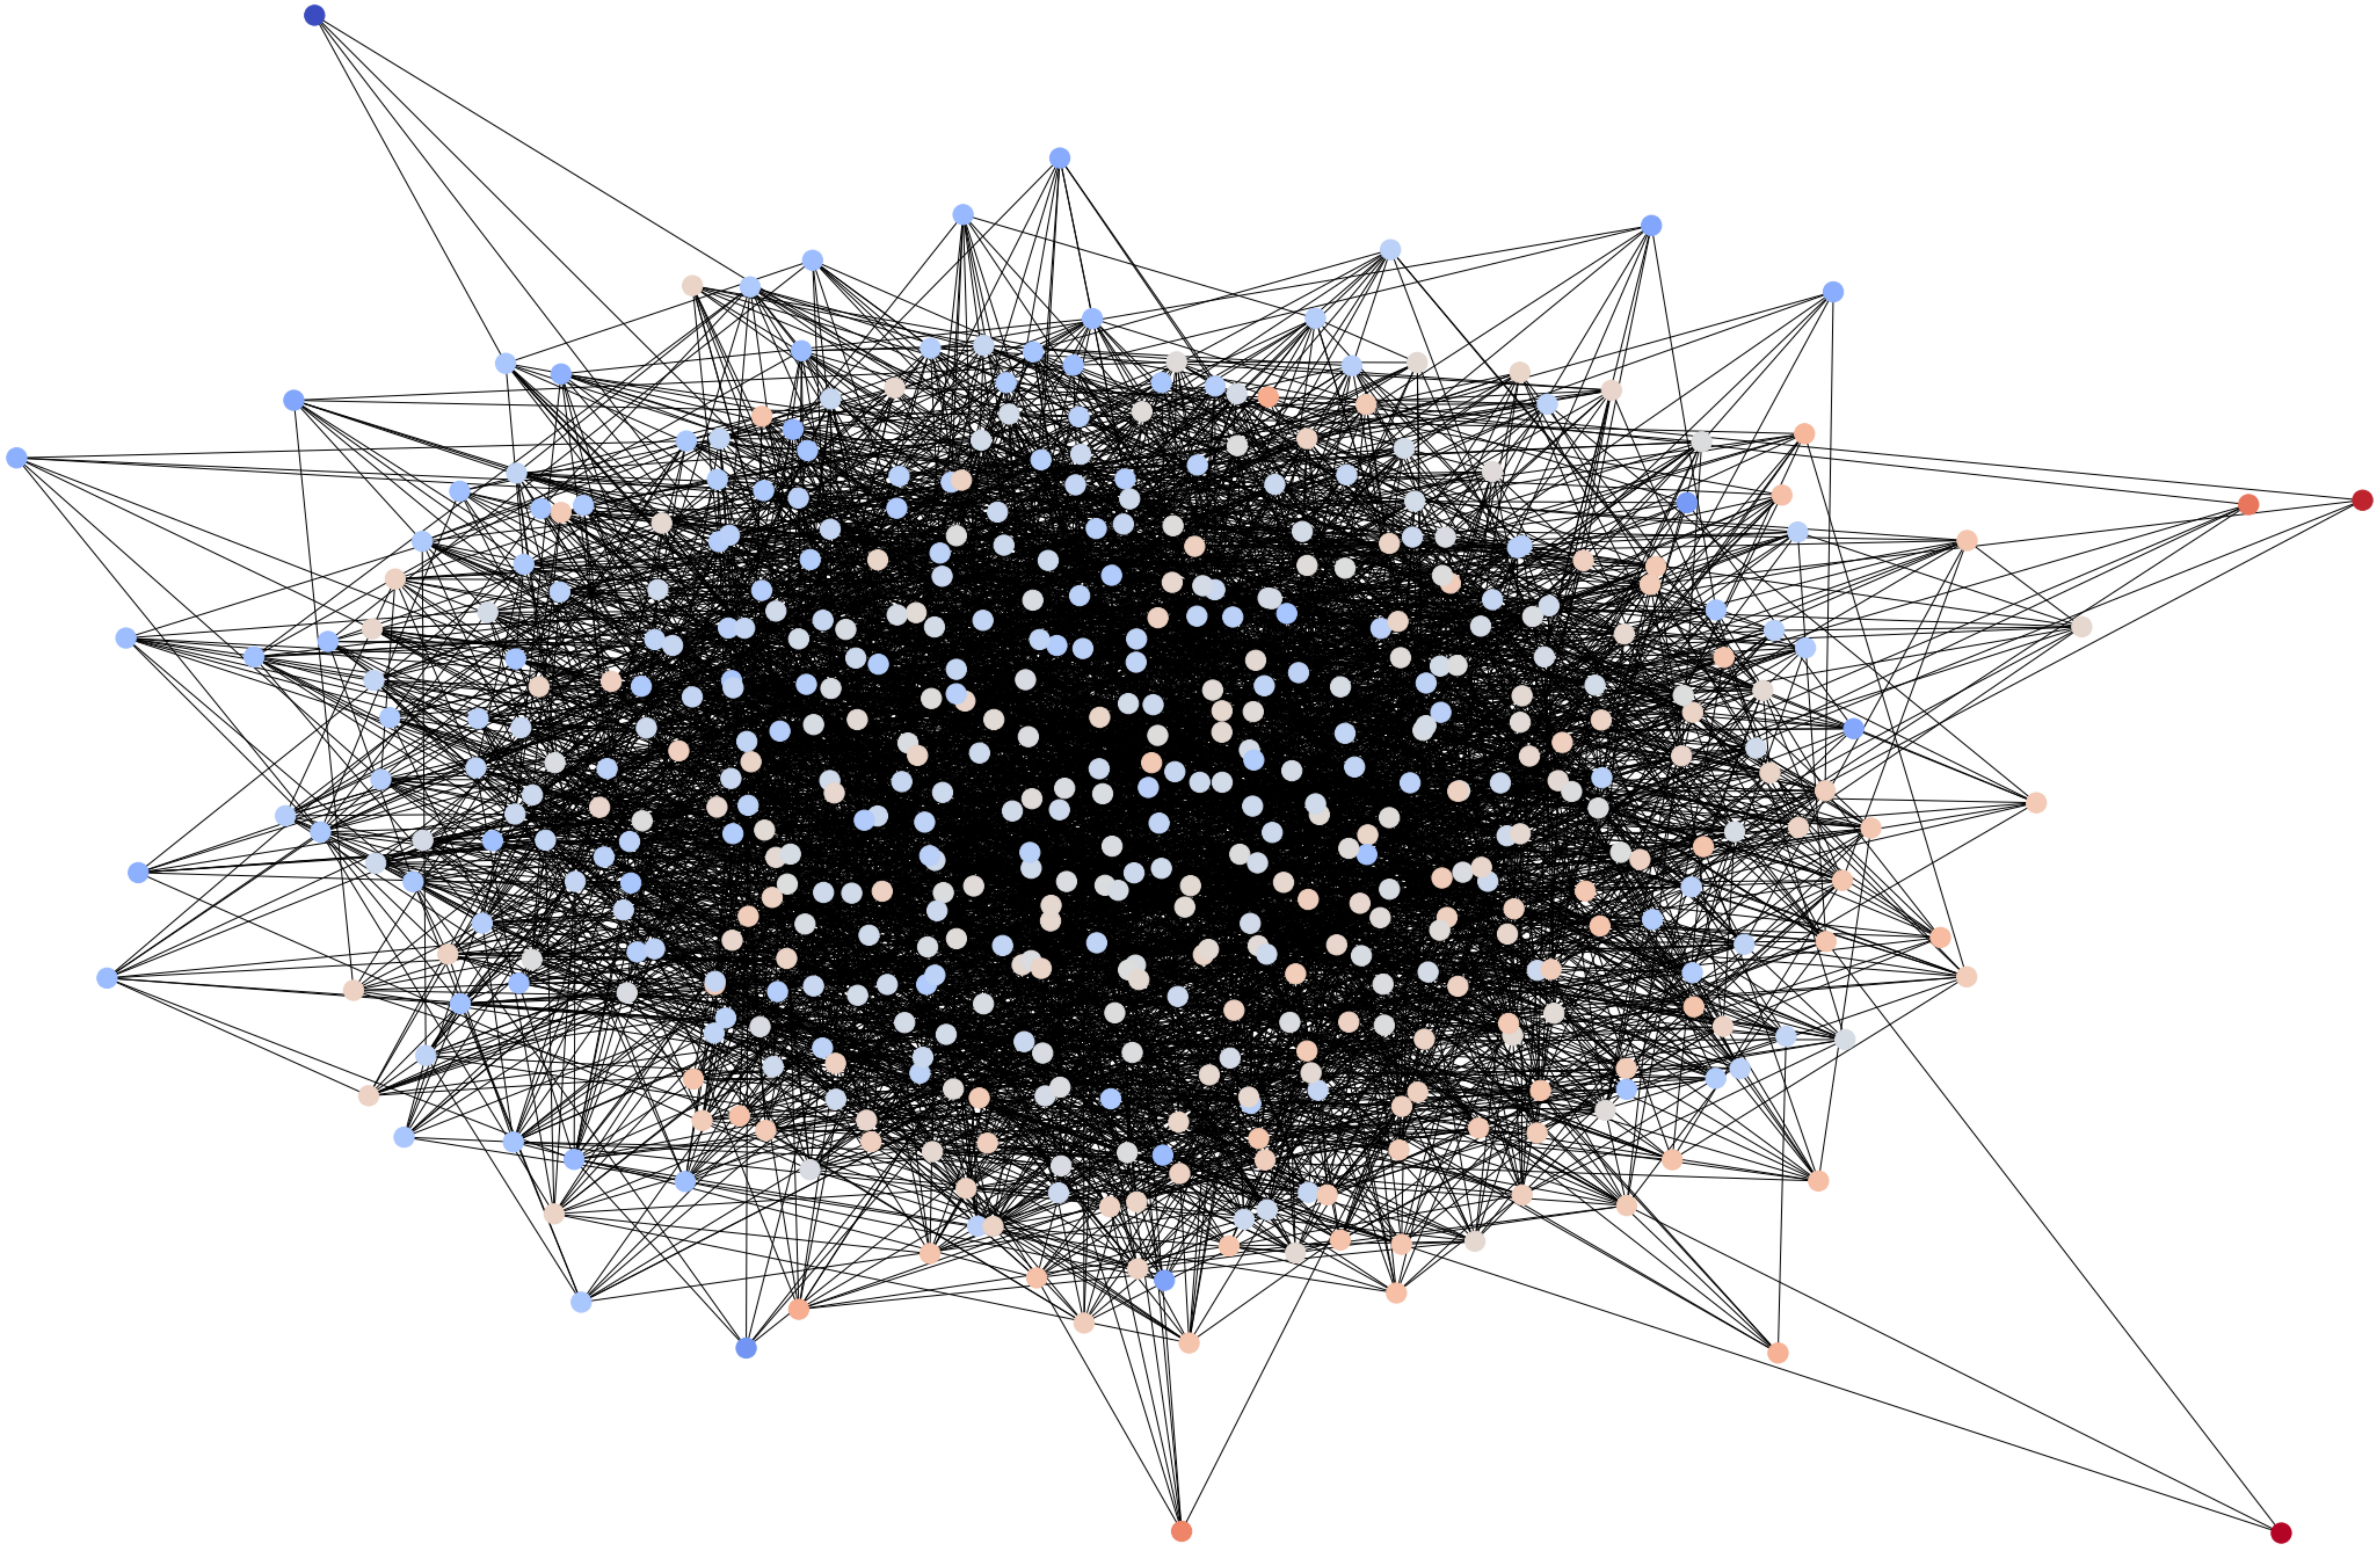

$t = 1$

No Fixed Edges, E-R graph, $n = 1000$, standard experiment

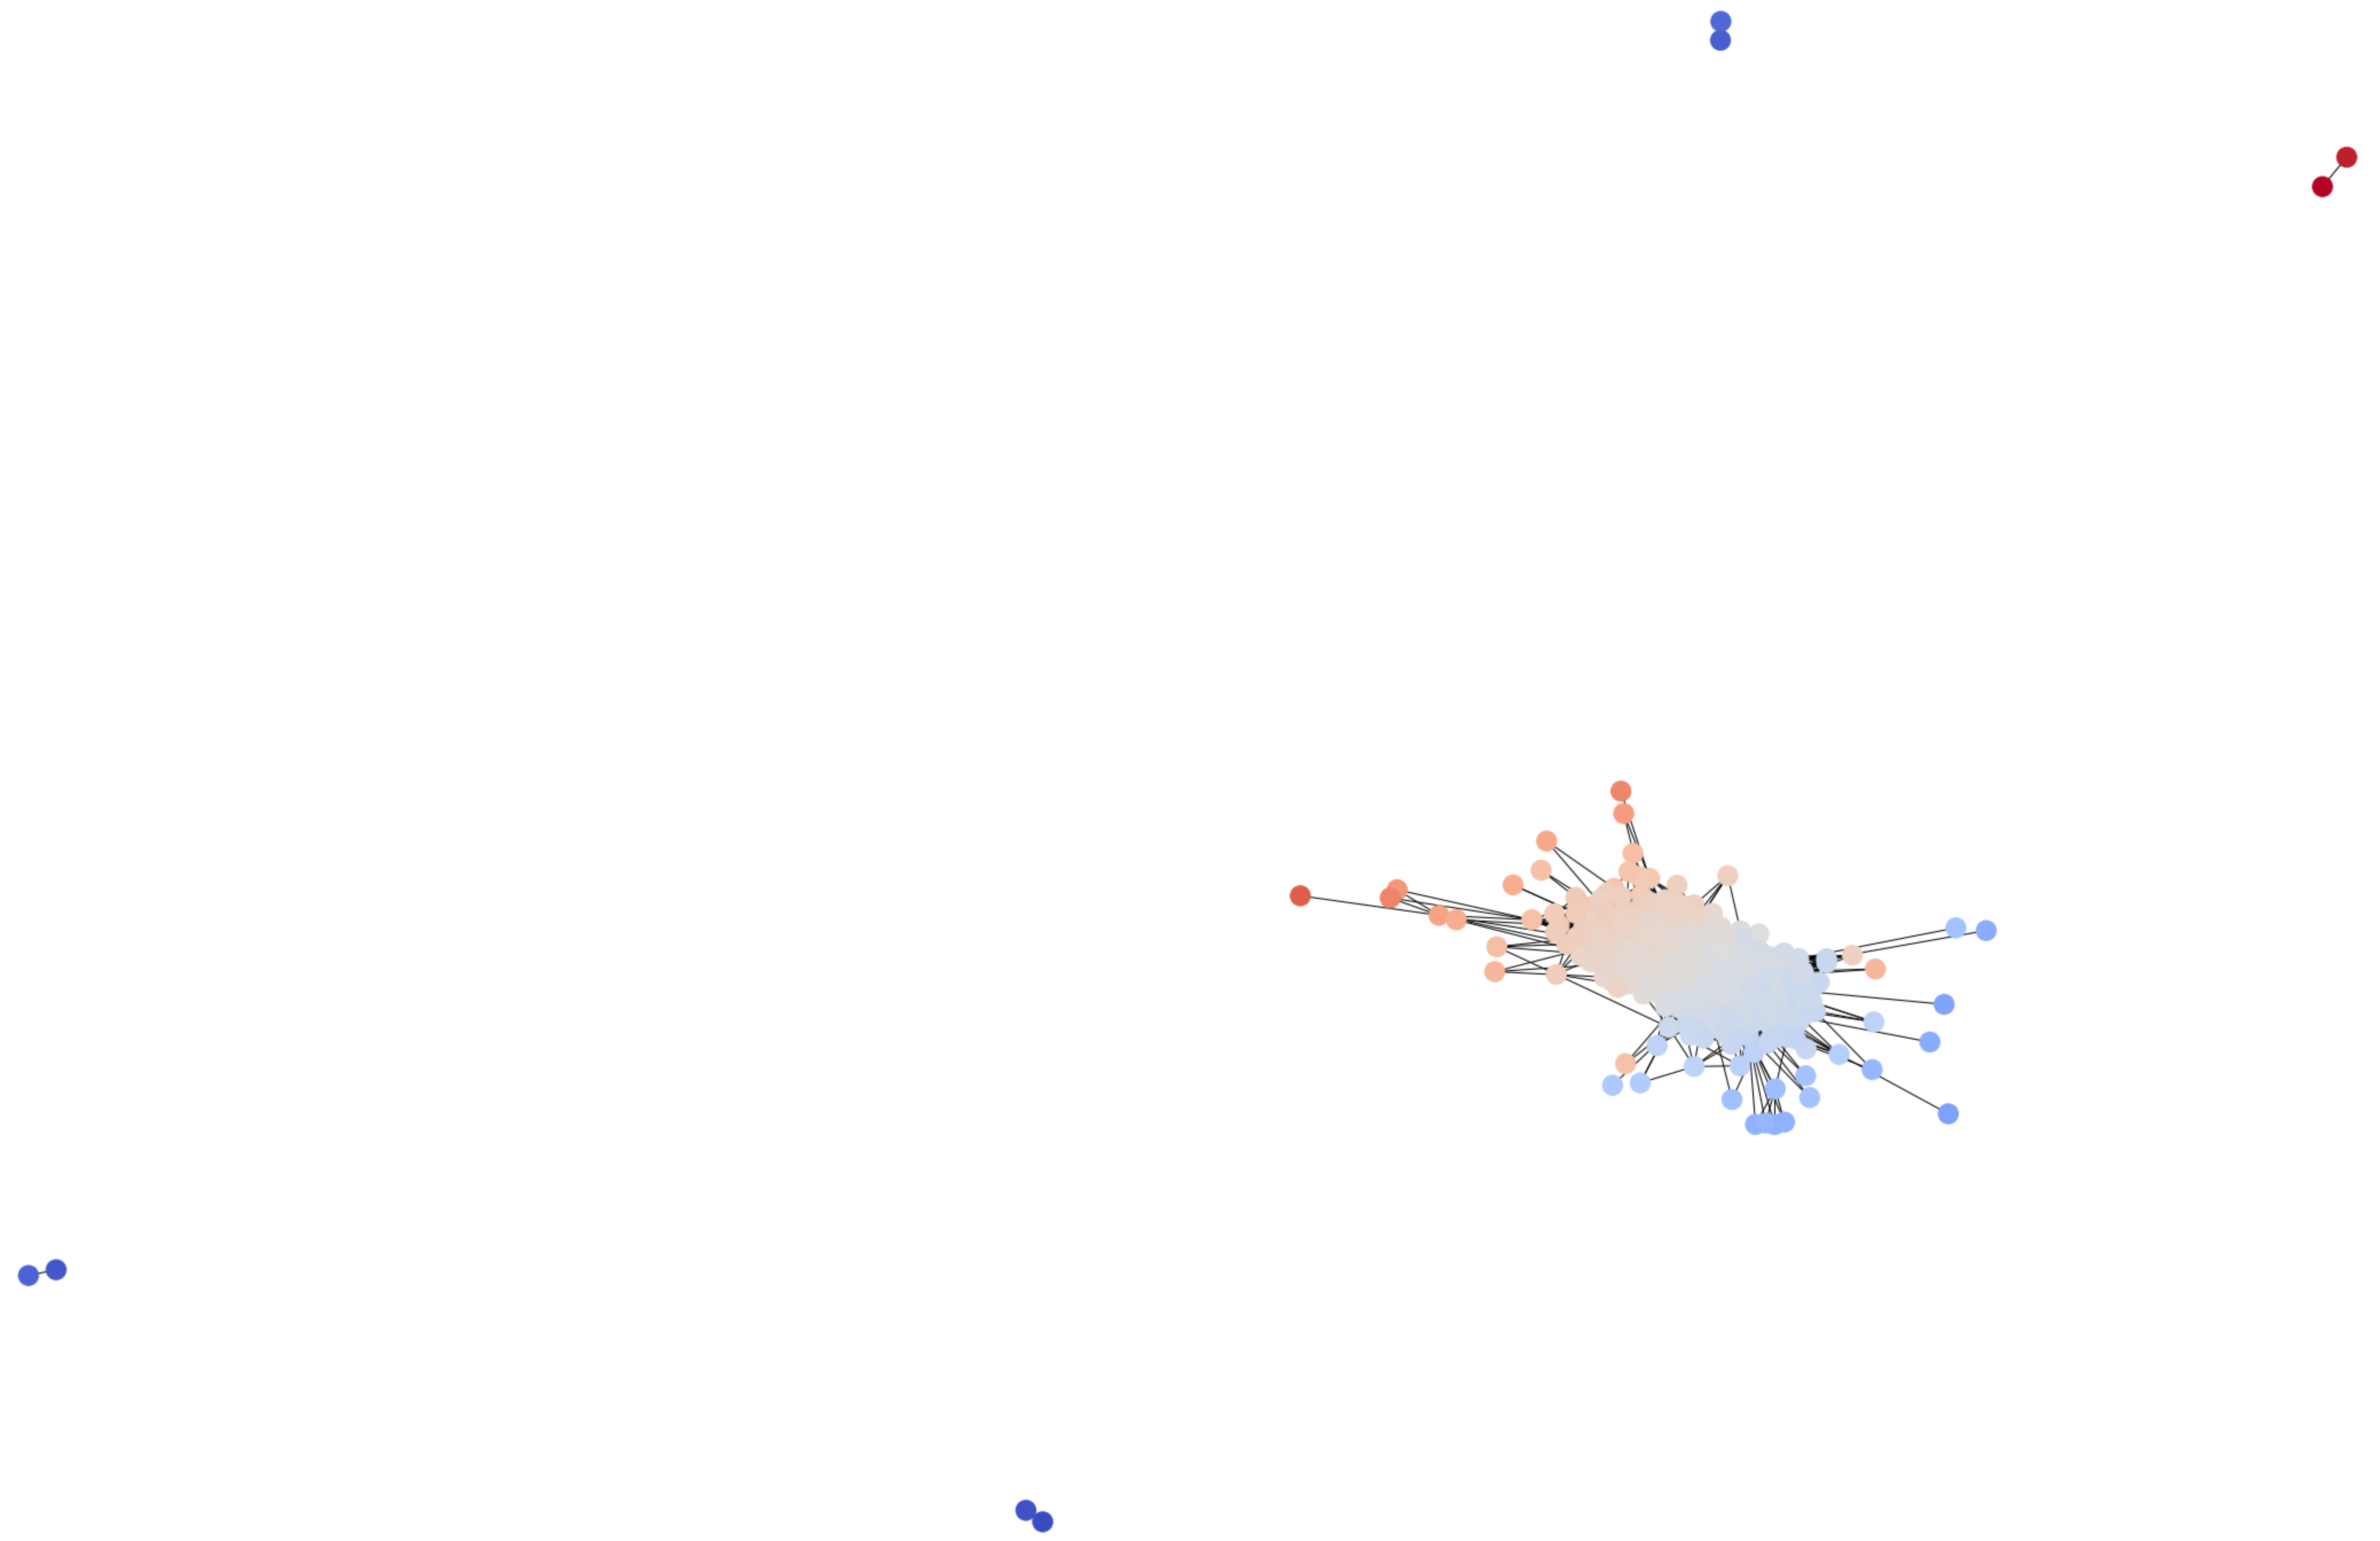

$t = 9$

No Fixed Edges, E-R graph, $n = 1000$, standard experiment

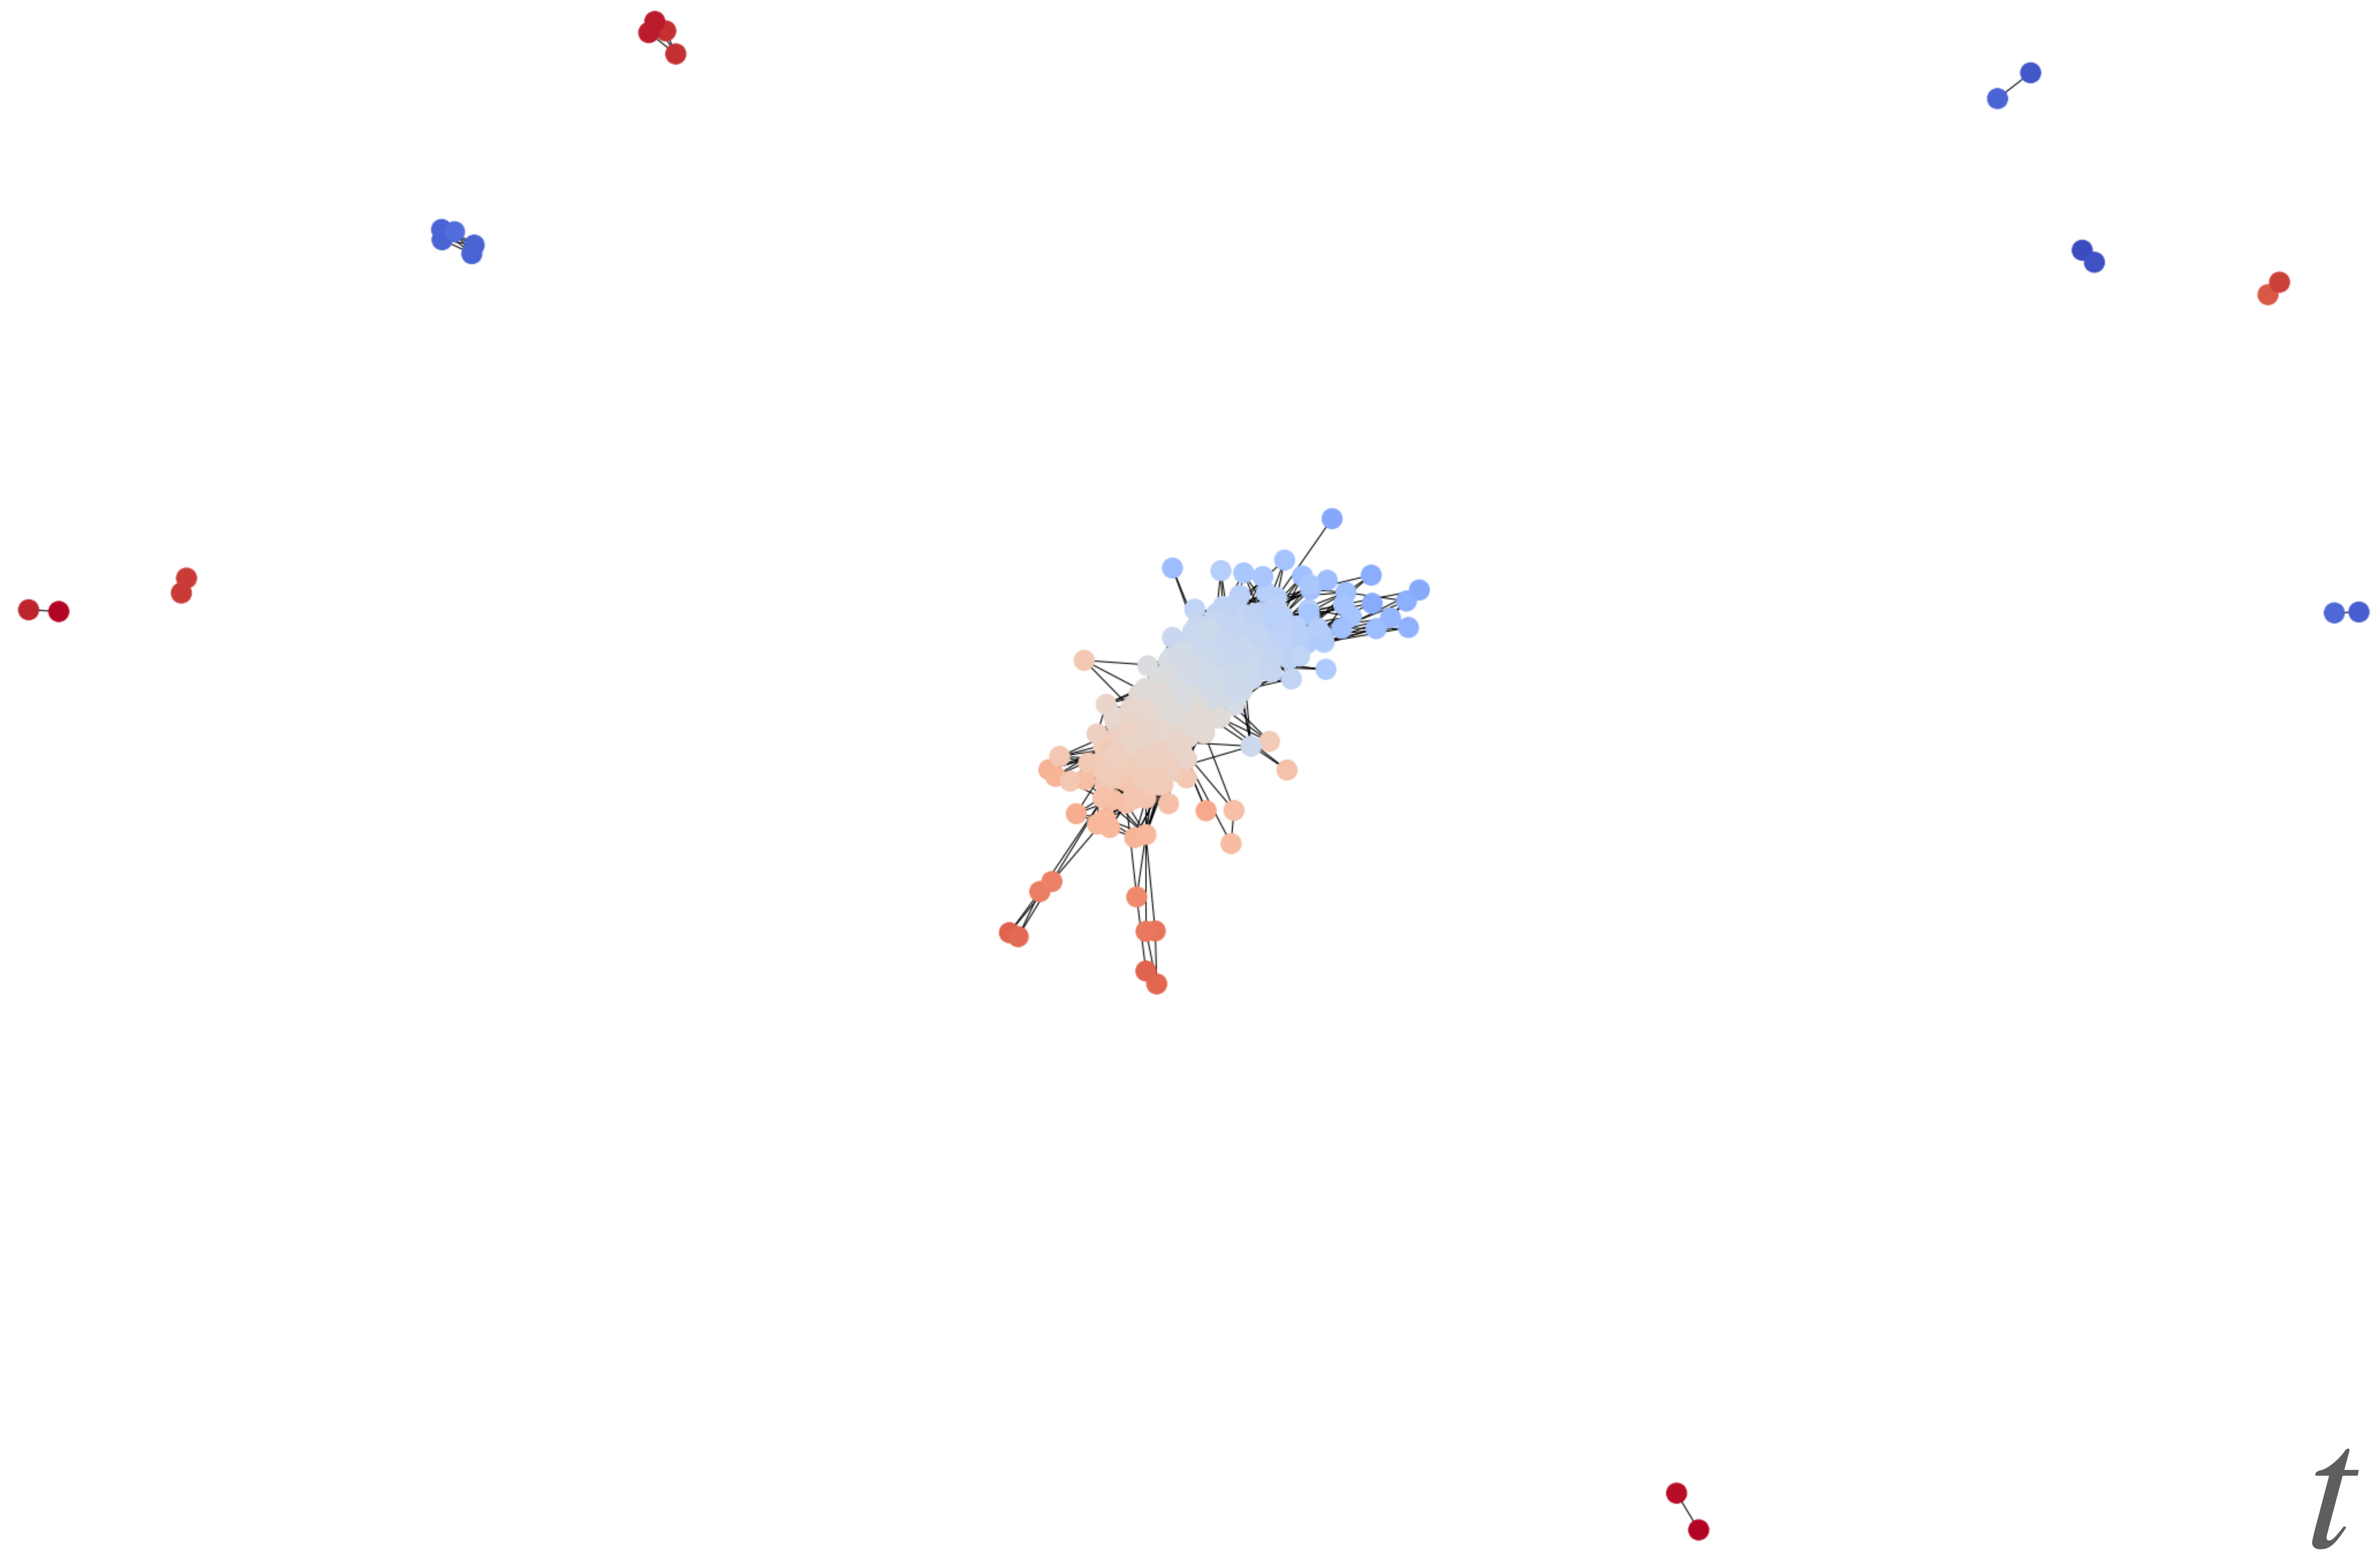

$t = 16$

No Fixed Edges, E-R graph, $n = 1000$, standard experiment

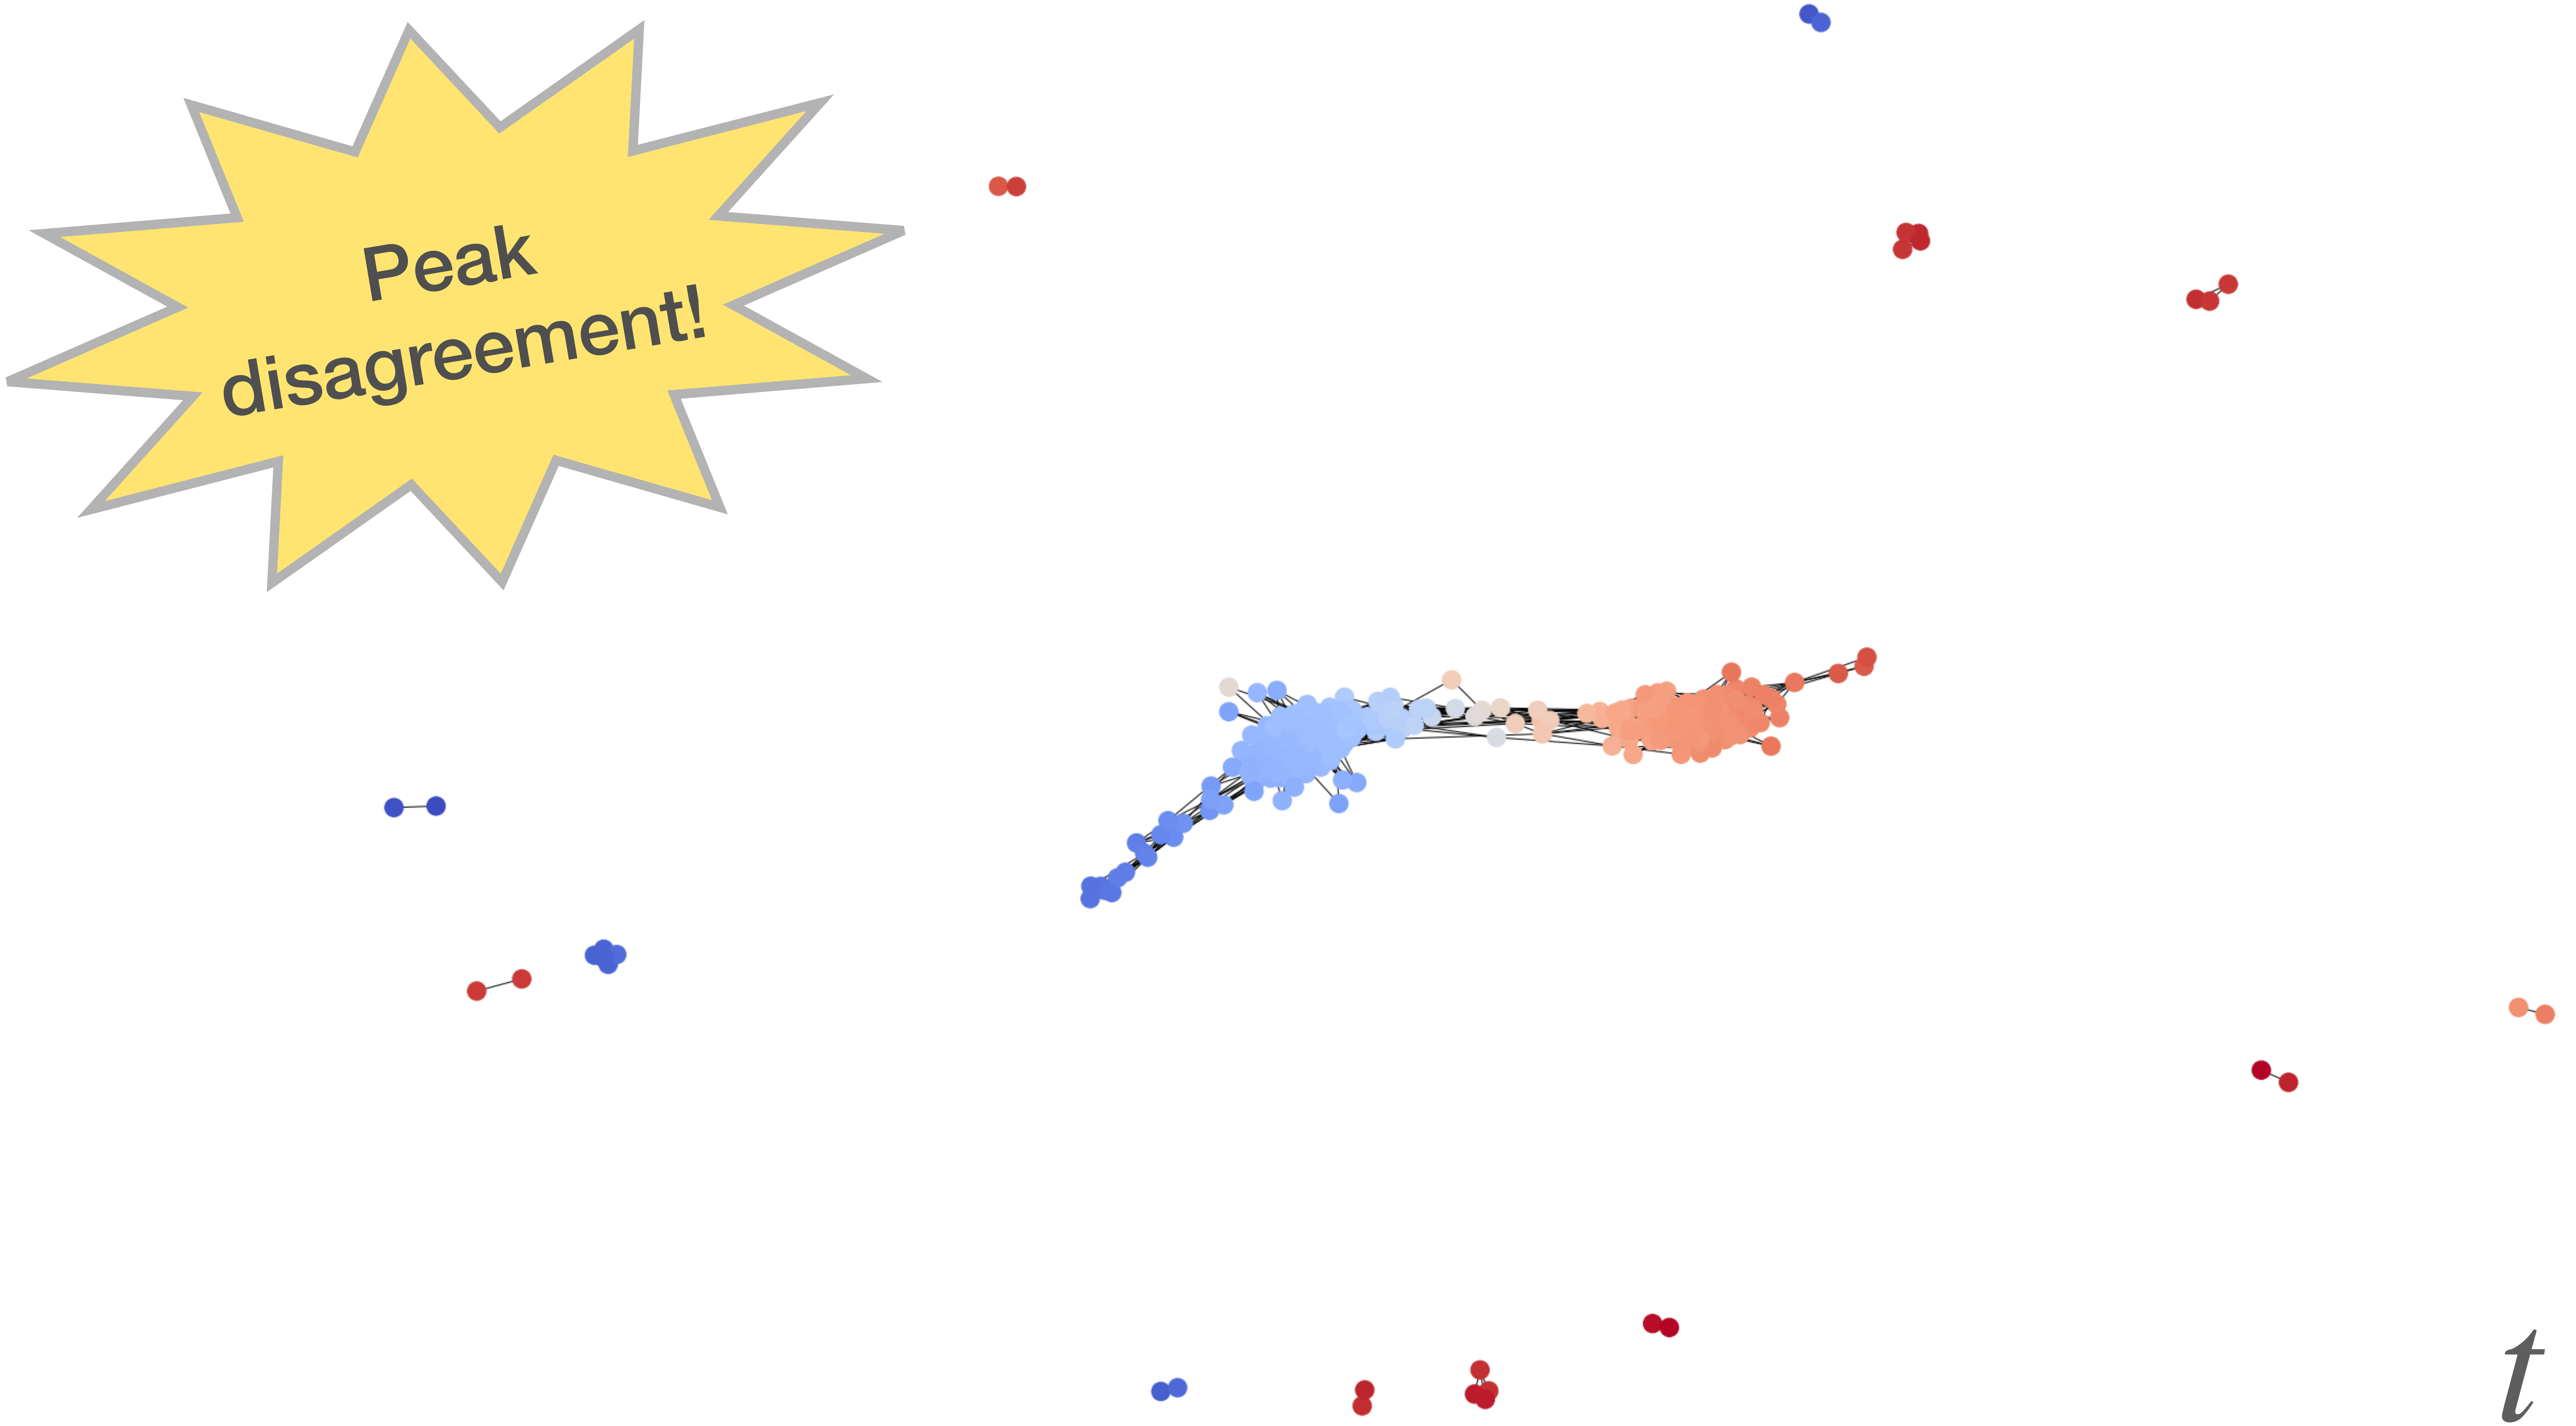

$t = 23$

No Fixed Edges, E-R graph, $n = 1000$, standard experiment

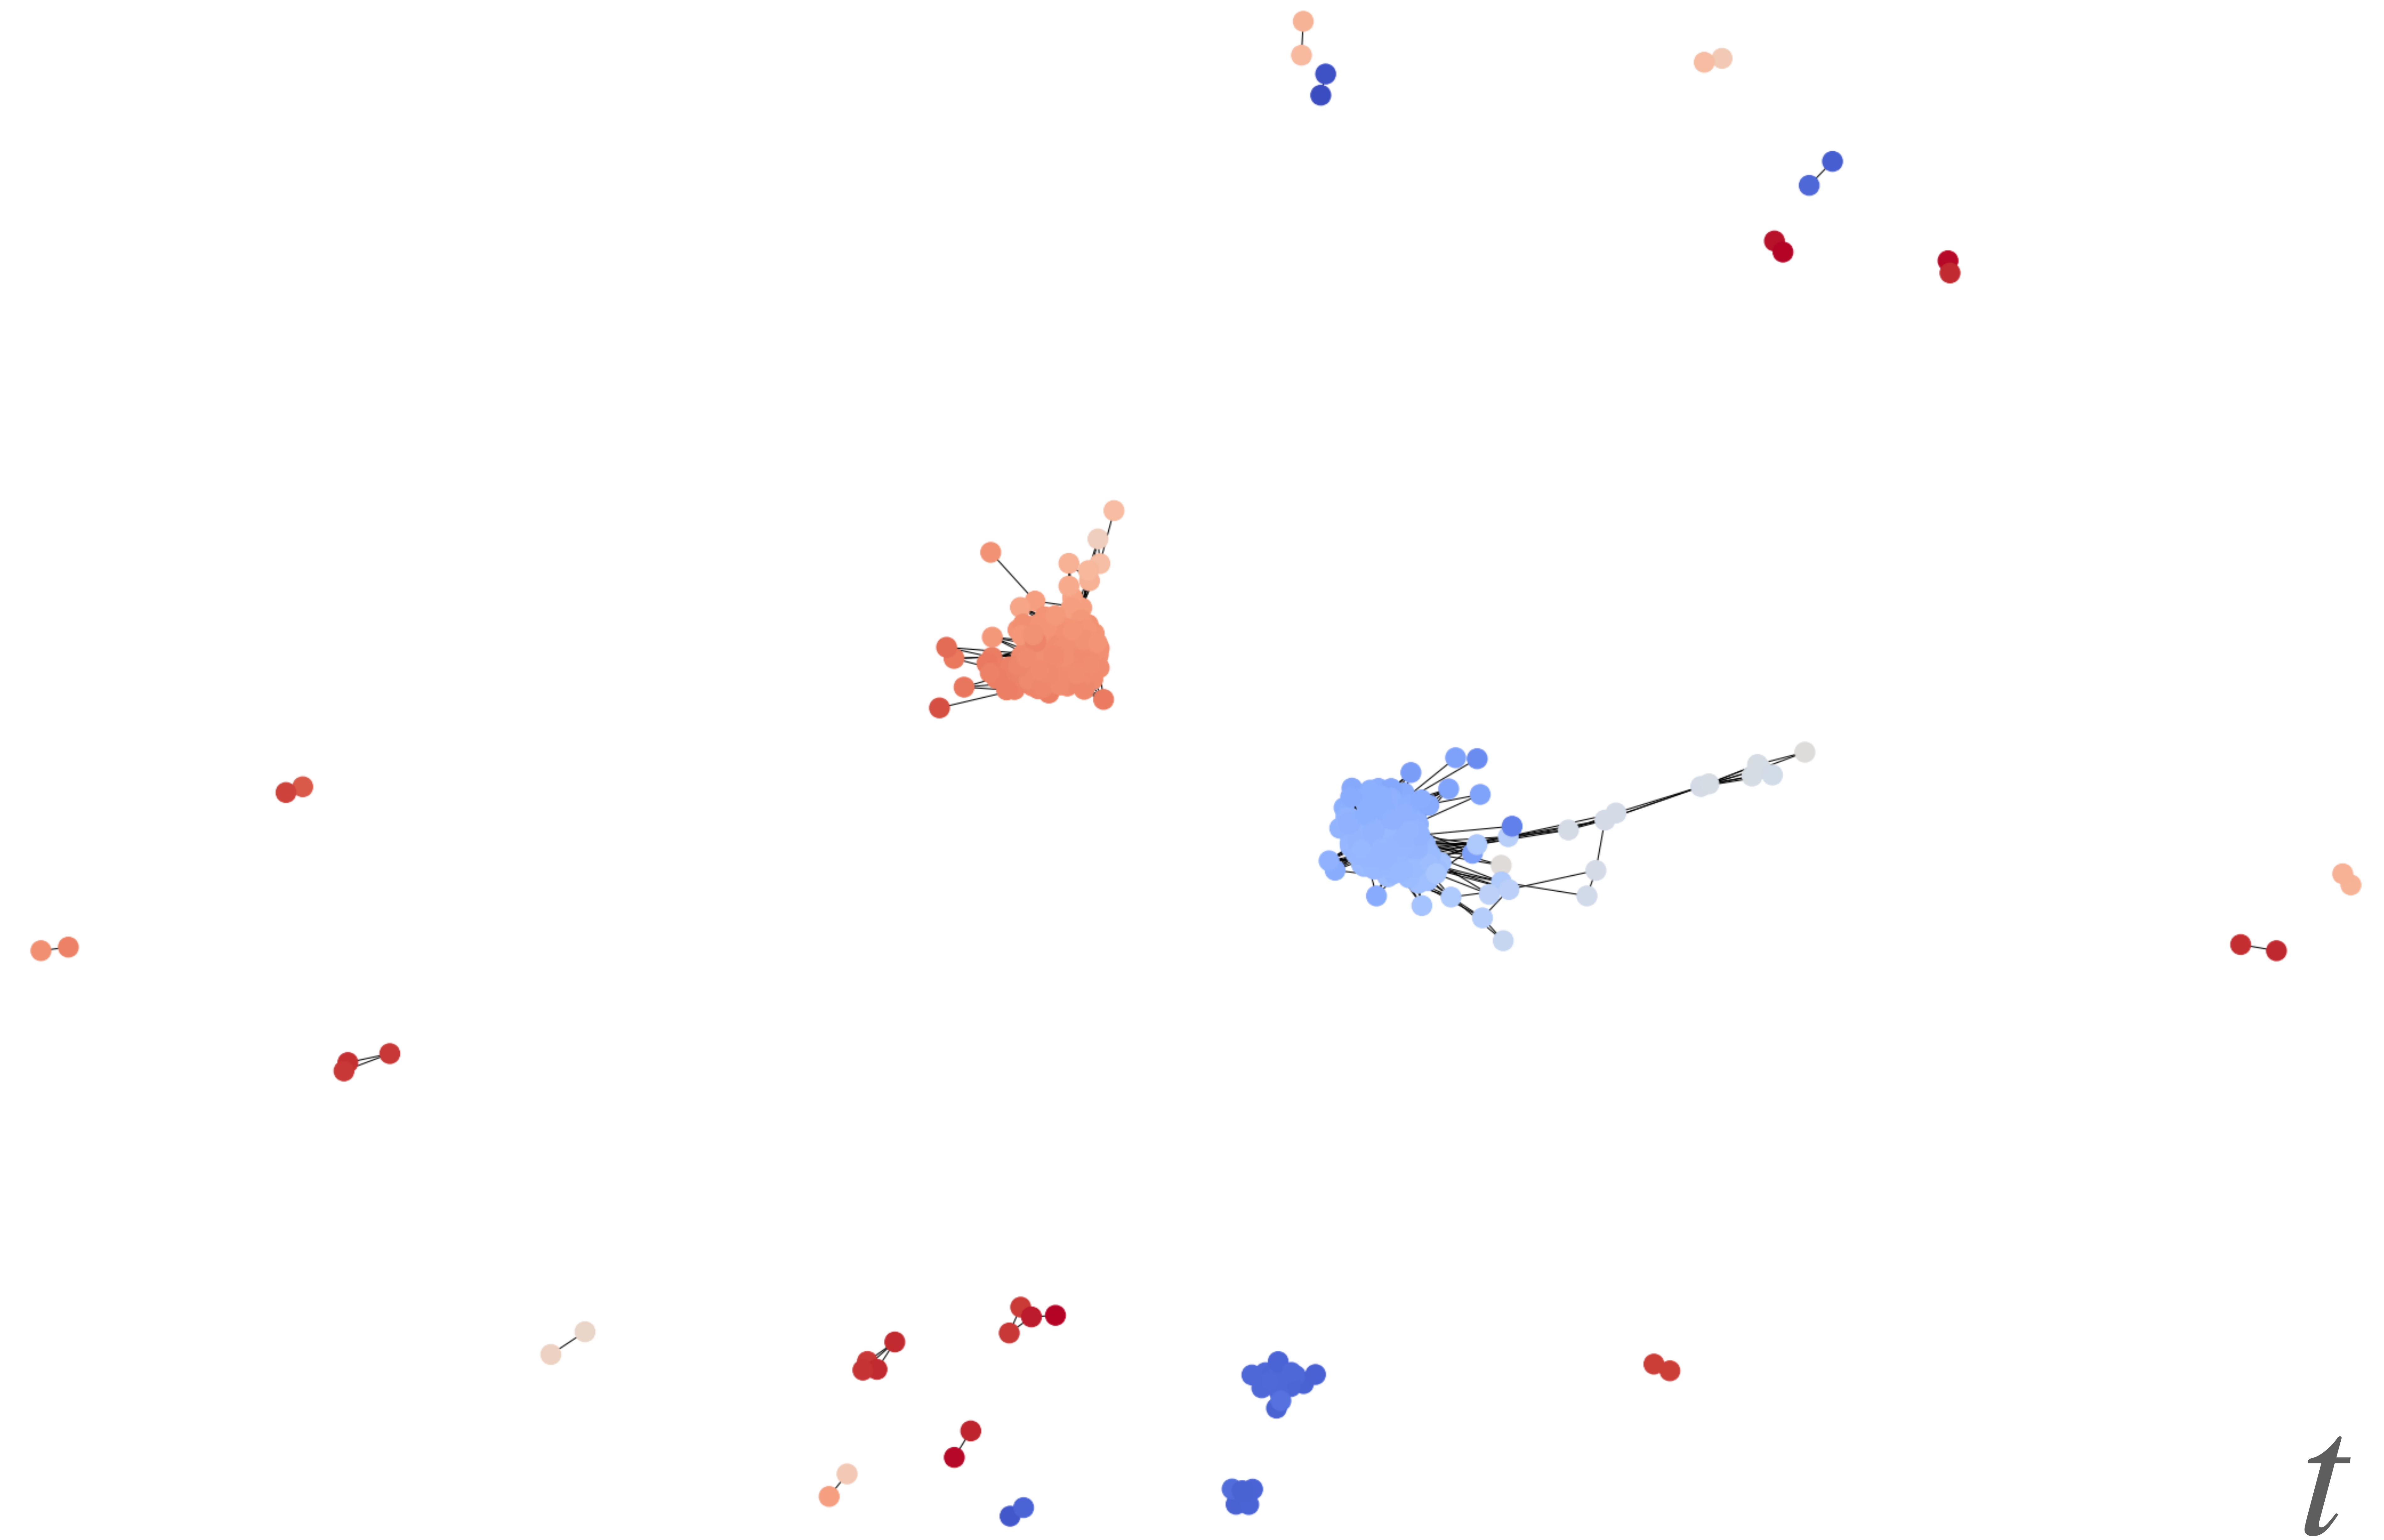

$t = 25$

No Fixed Edges, E-R graph, $n = 1000$, standard experiment

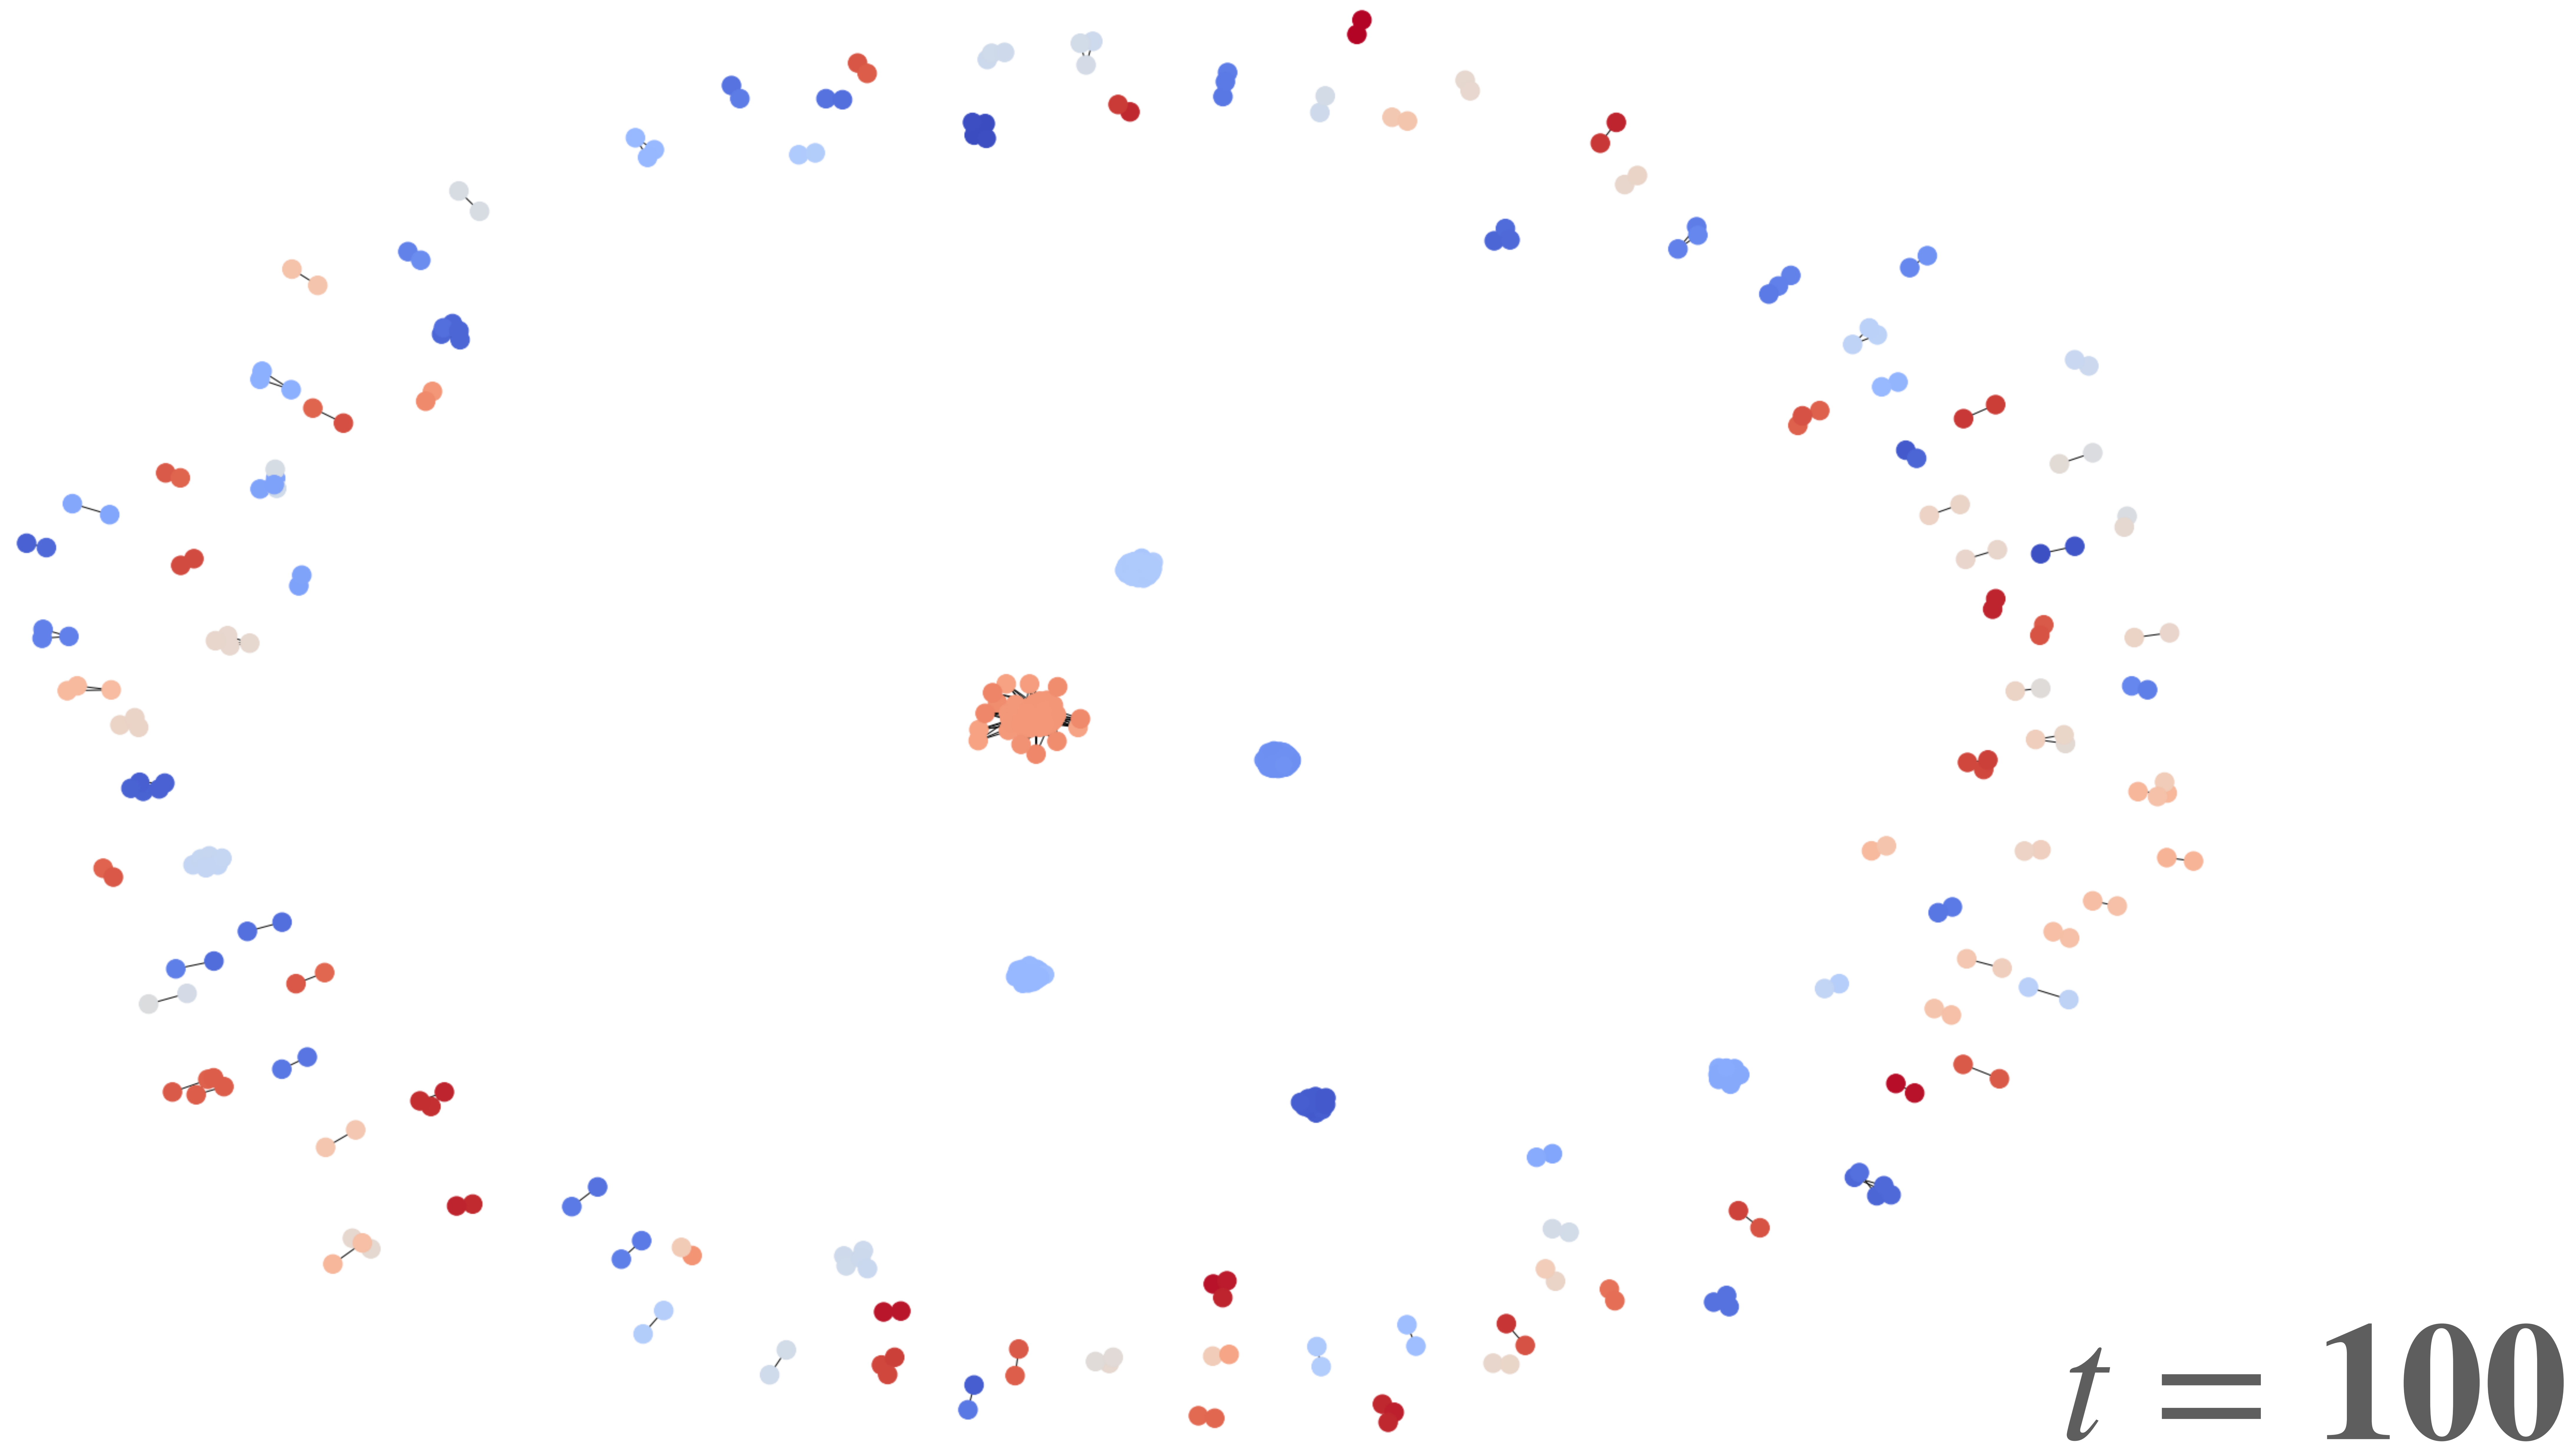

10% Fixed Edges, E-R graph, $n = 1000$, standard edge dynamics

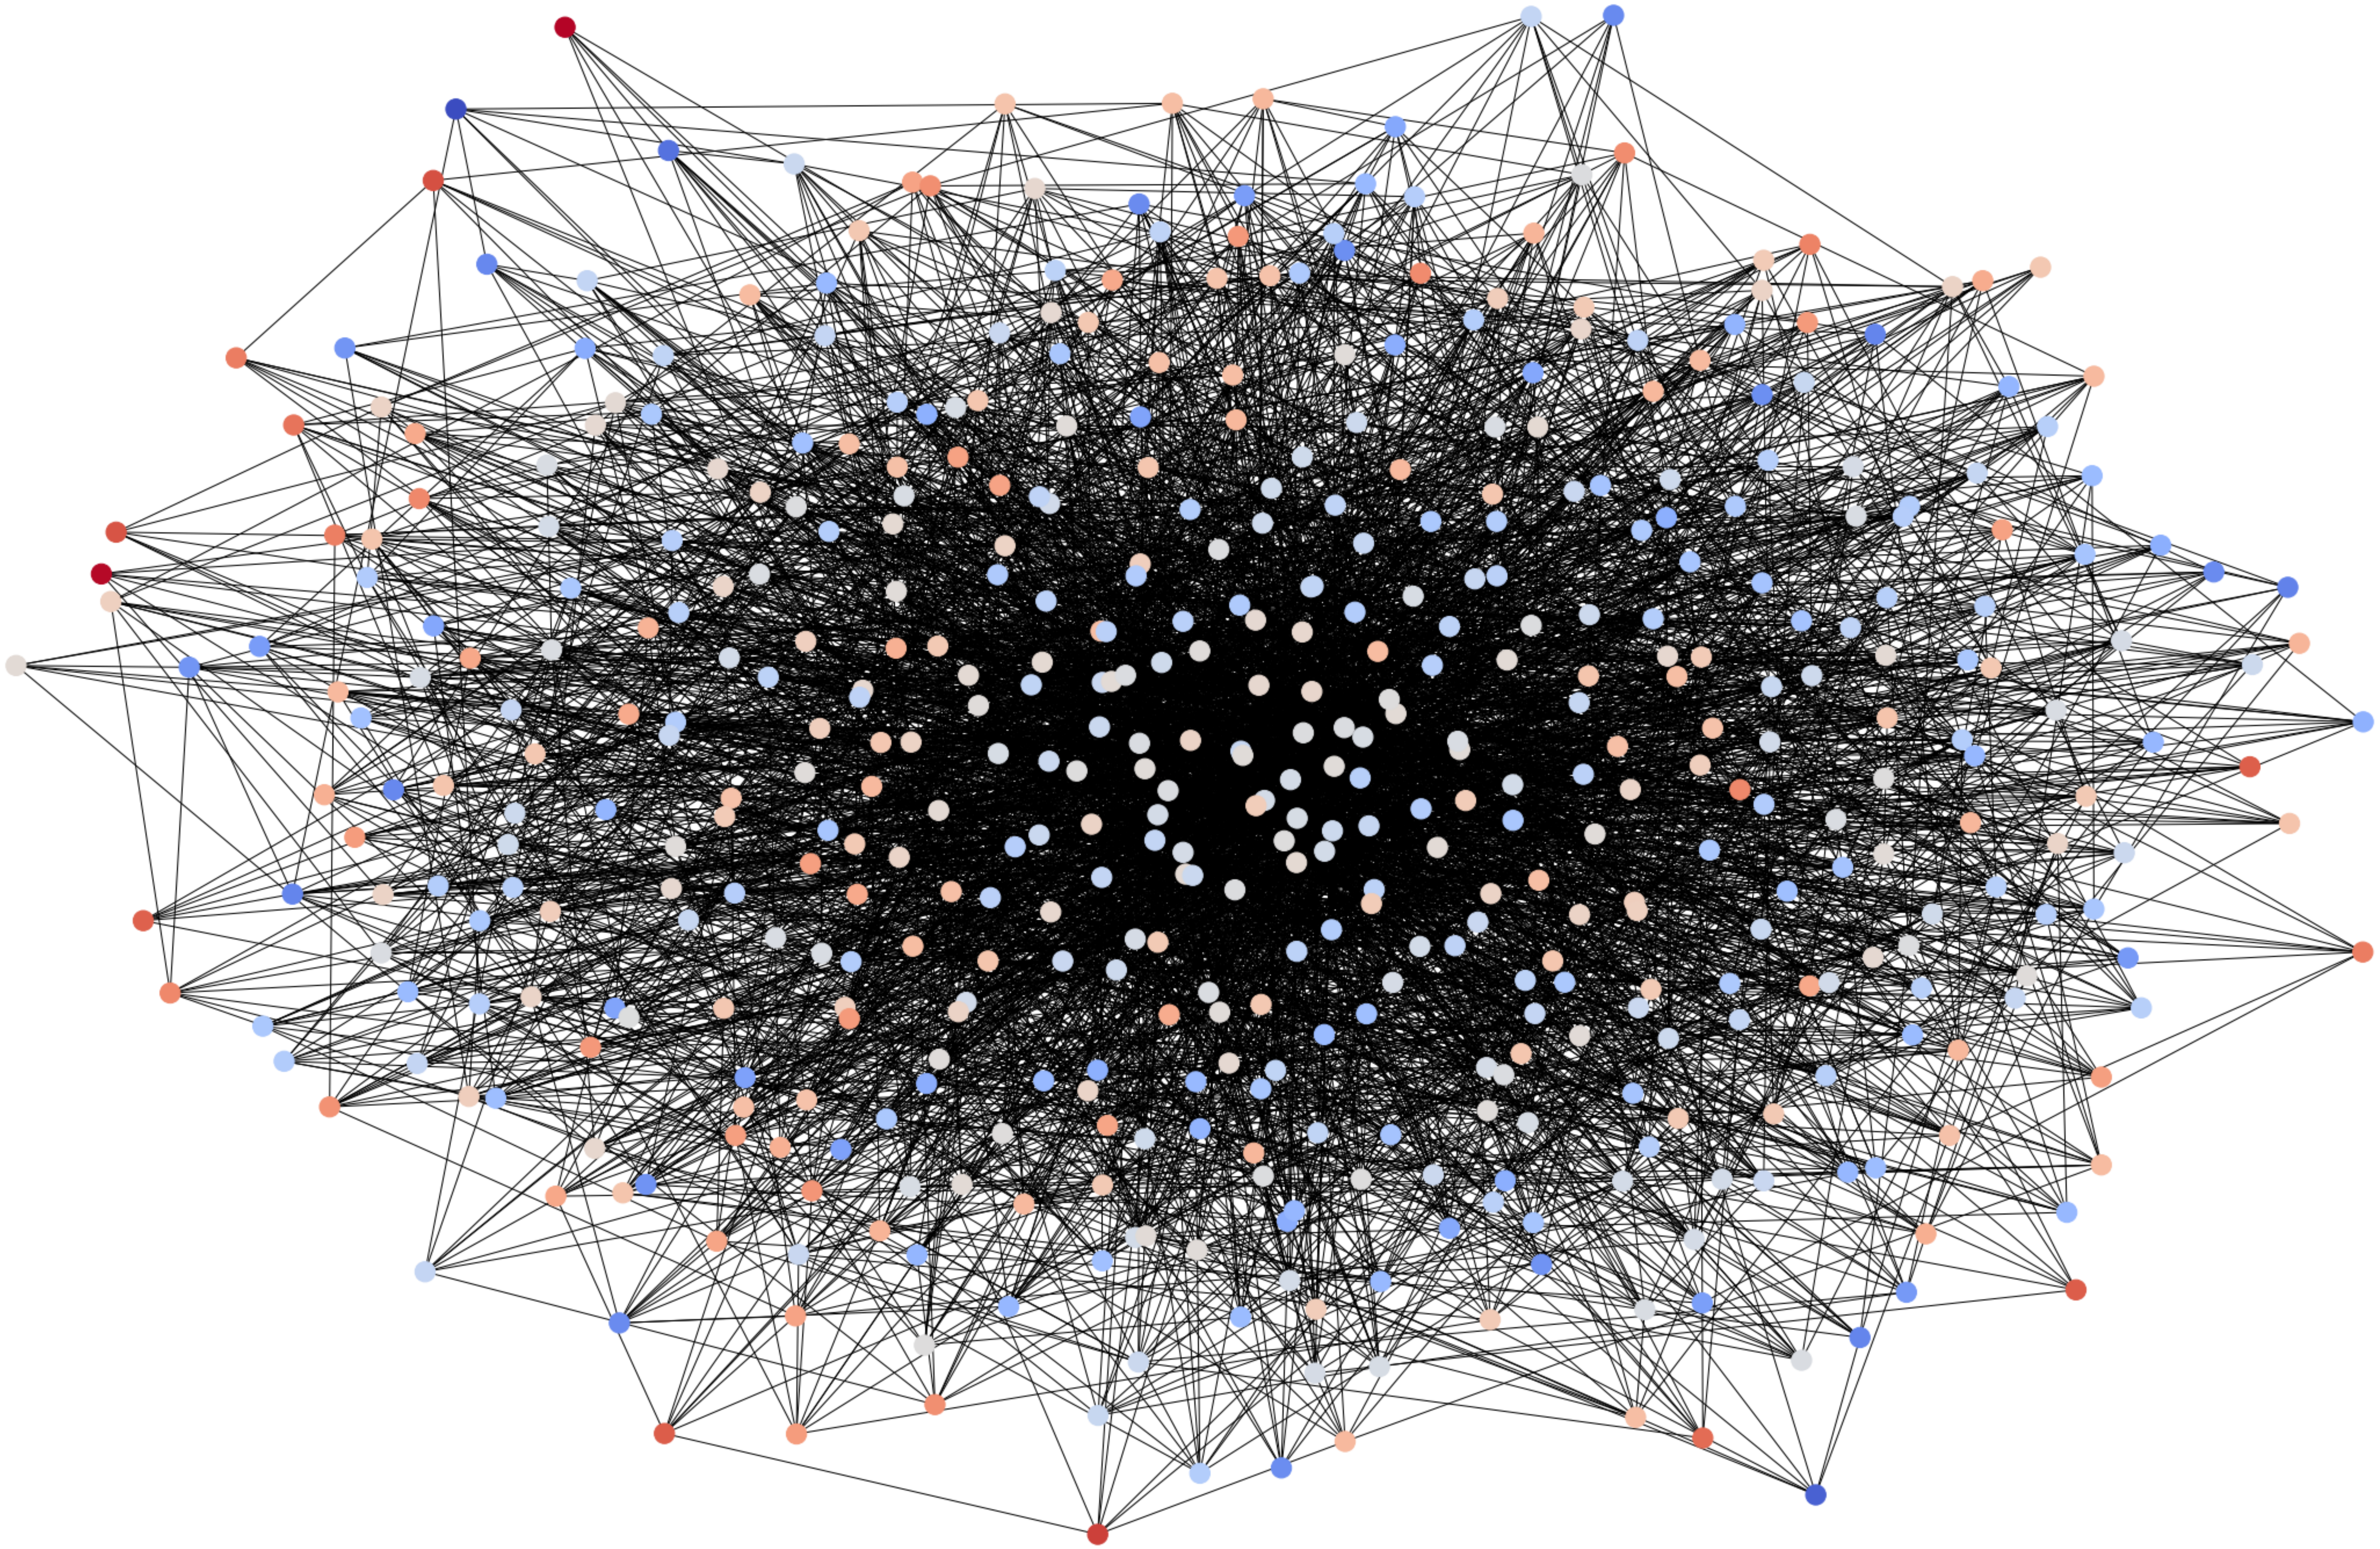

$t = 0$

10% Fixed Edges, E-R graph, $n = 1000$, standard edge dynamics

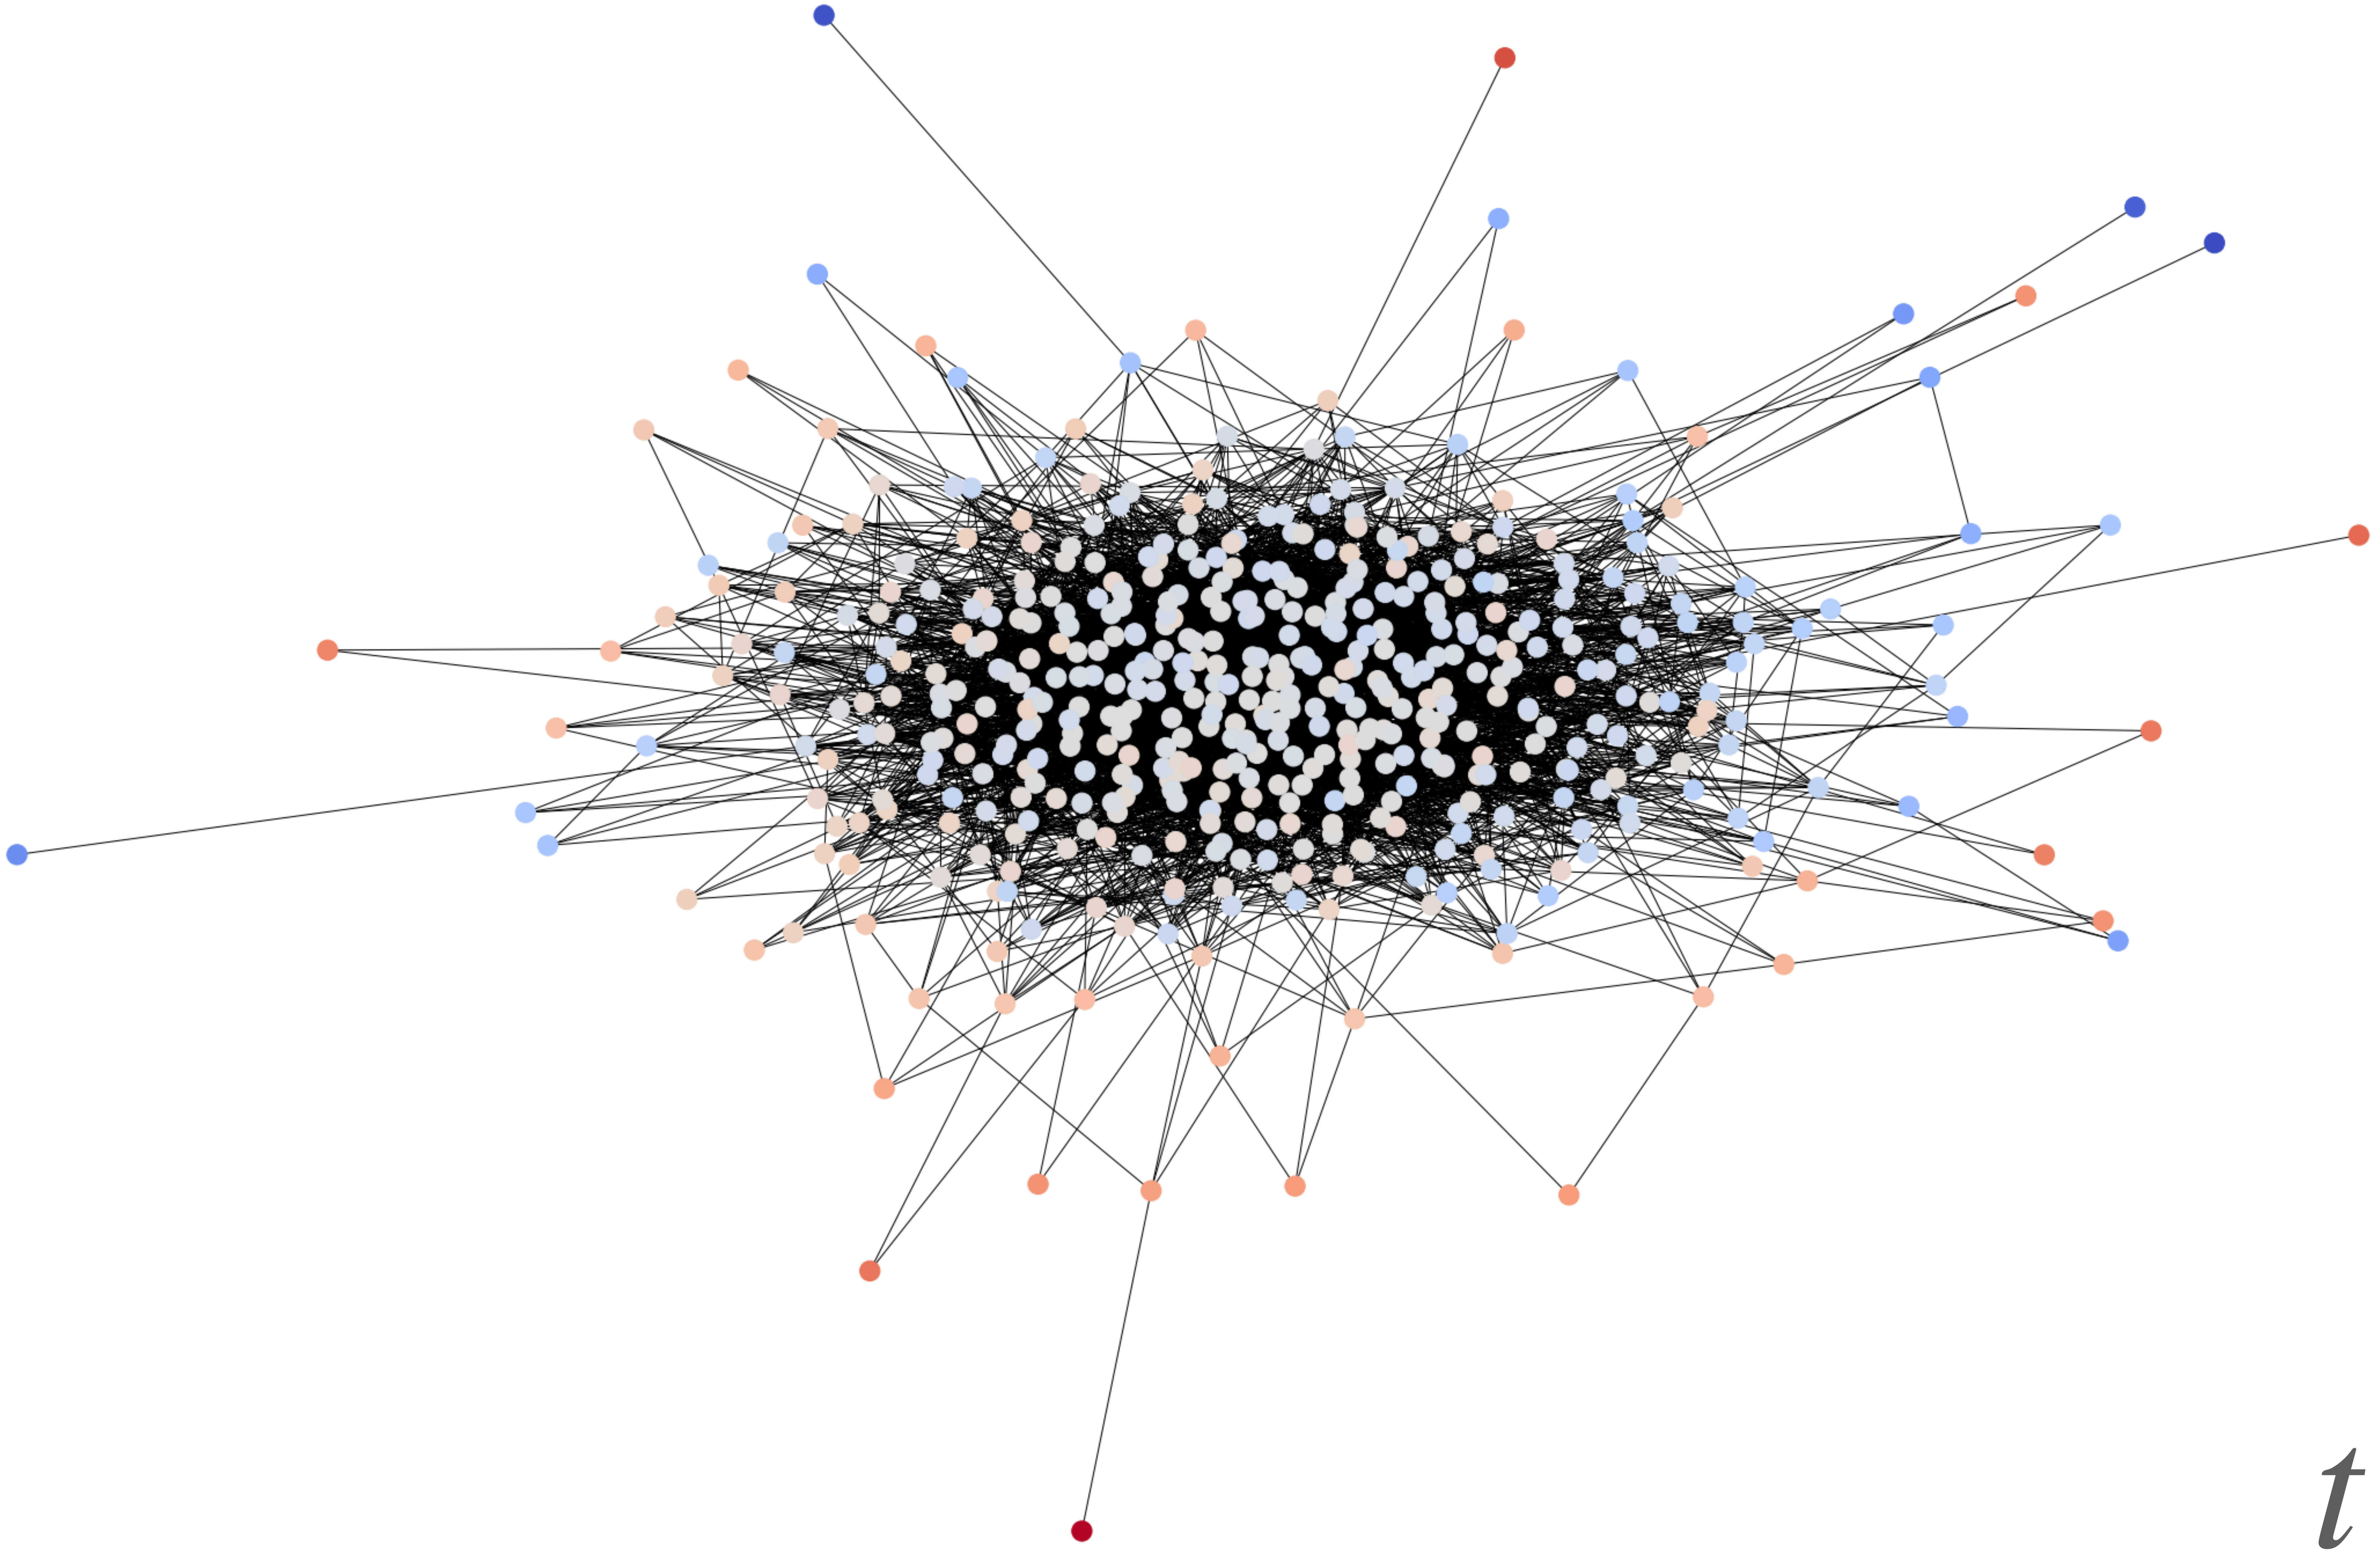

$t = 11$

10% Fixed Edges, E-R graph, $n = 1000$, standard edge dynamics

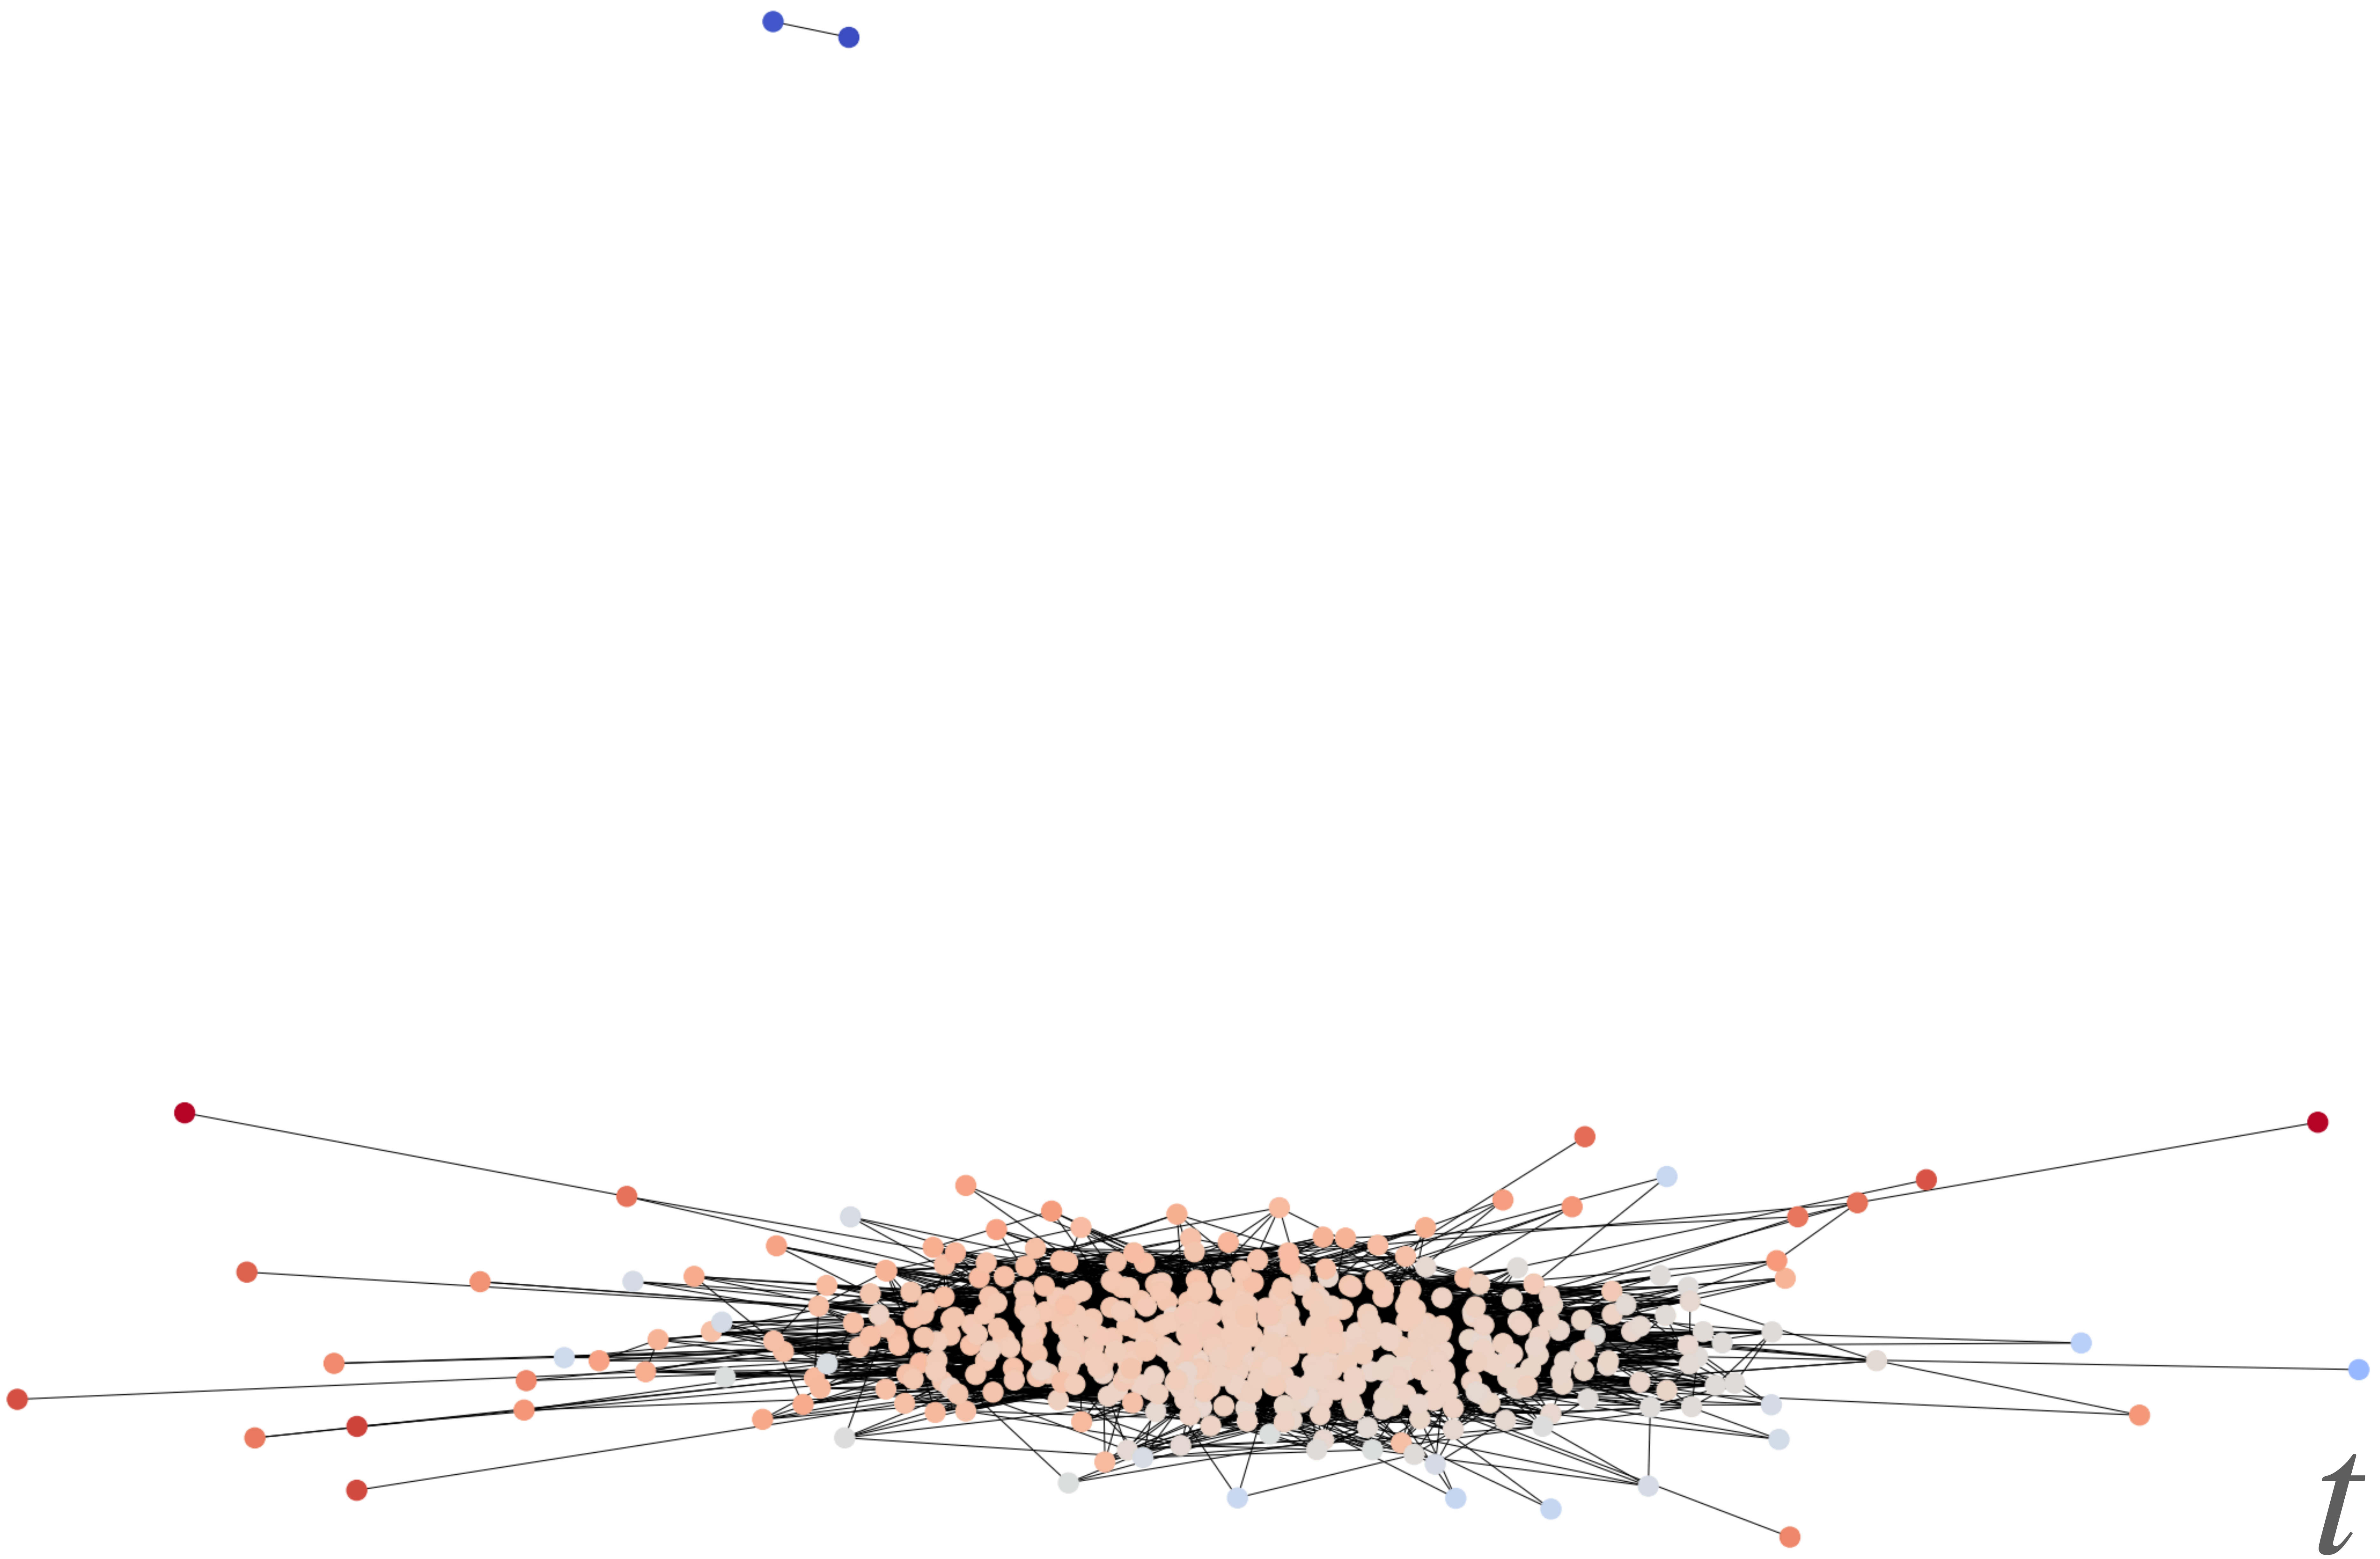

$t = 14$

10% Fixed Edges, E-R graph, $n = 1000$, standard edge dynamics

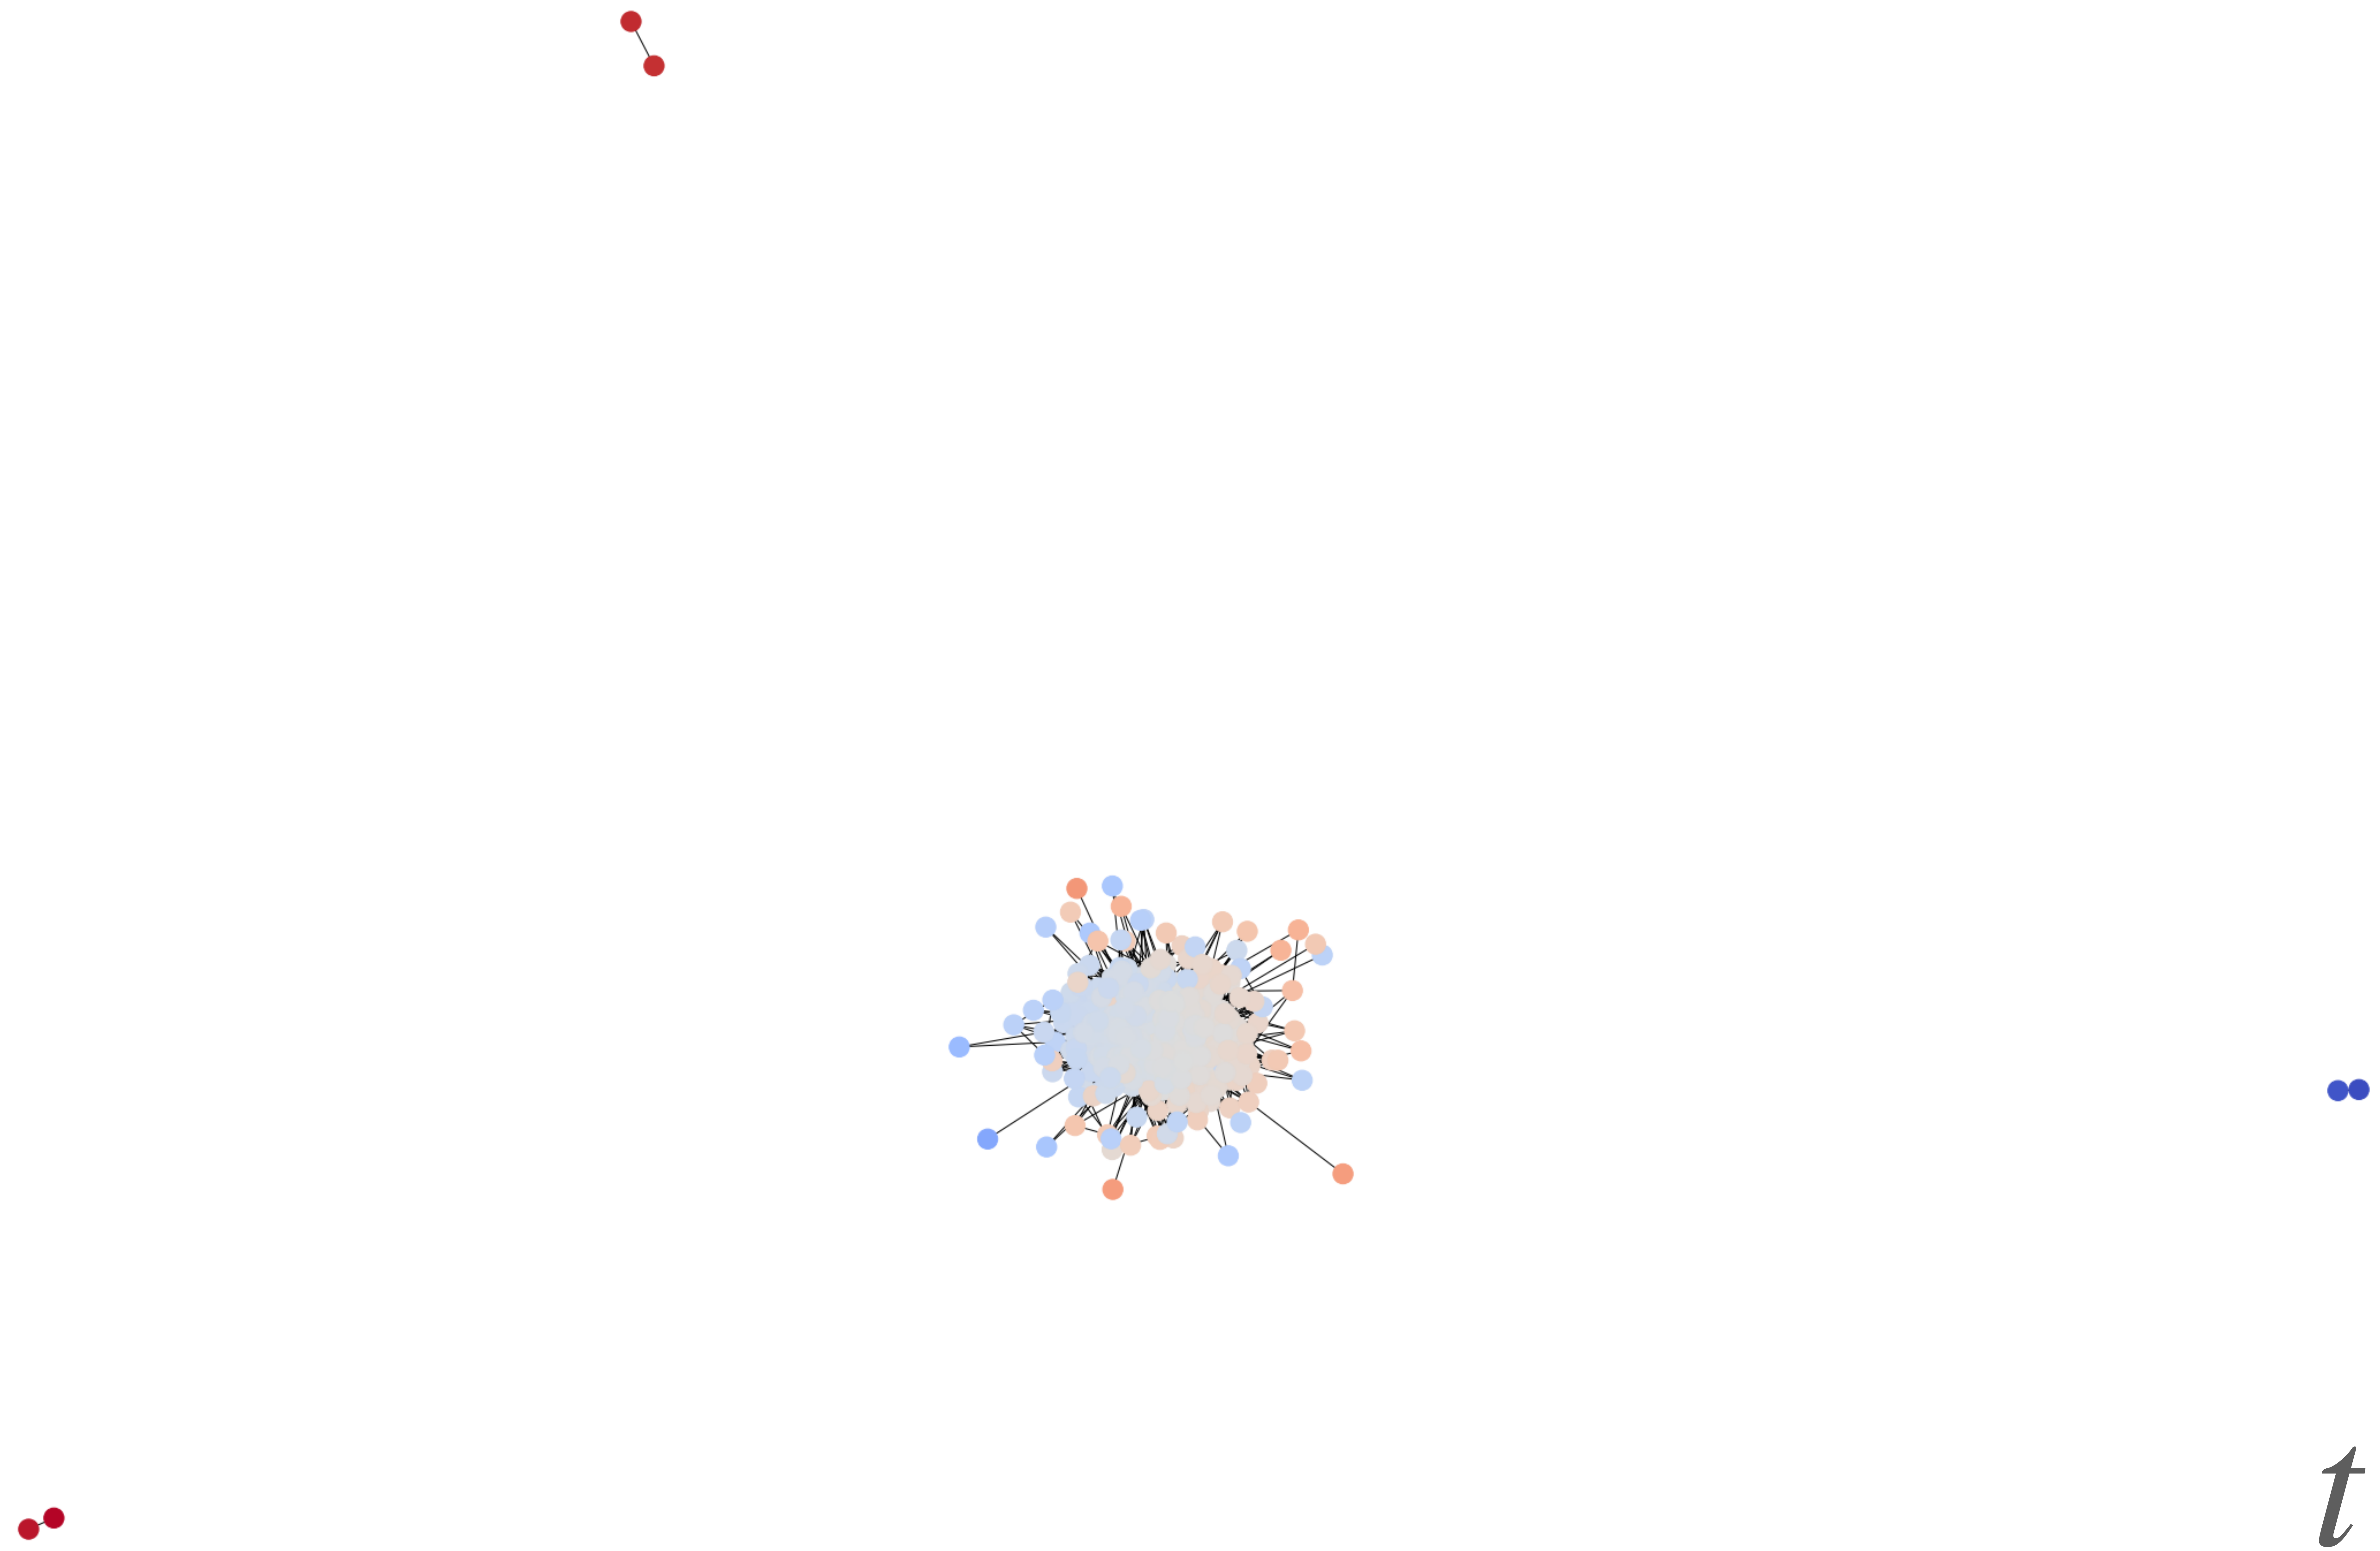

$t = 19$

10% Fixed Edges, E-R graph, $n = 1000$, standard edge dynamics

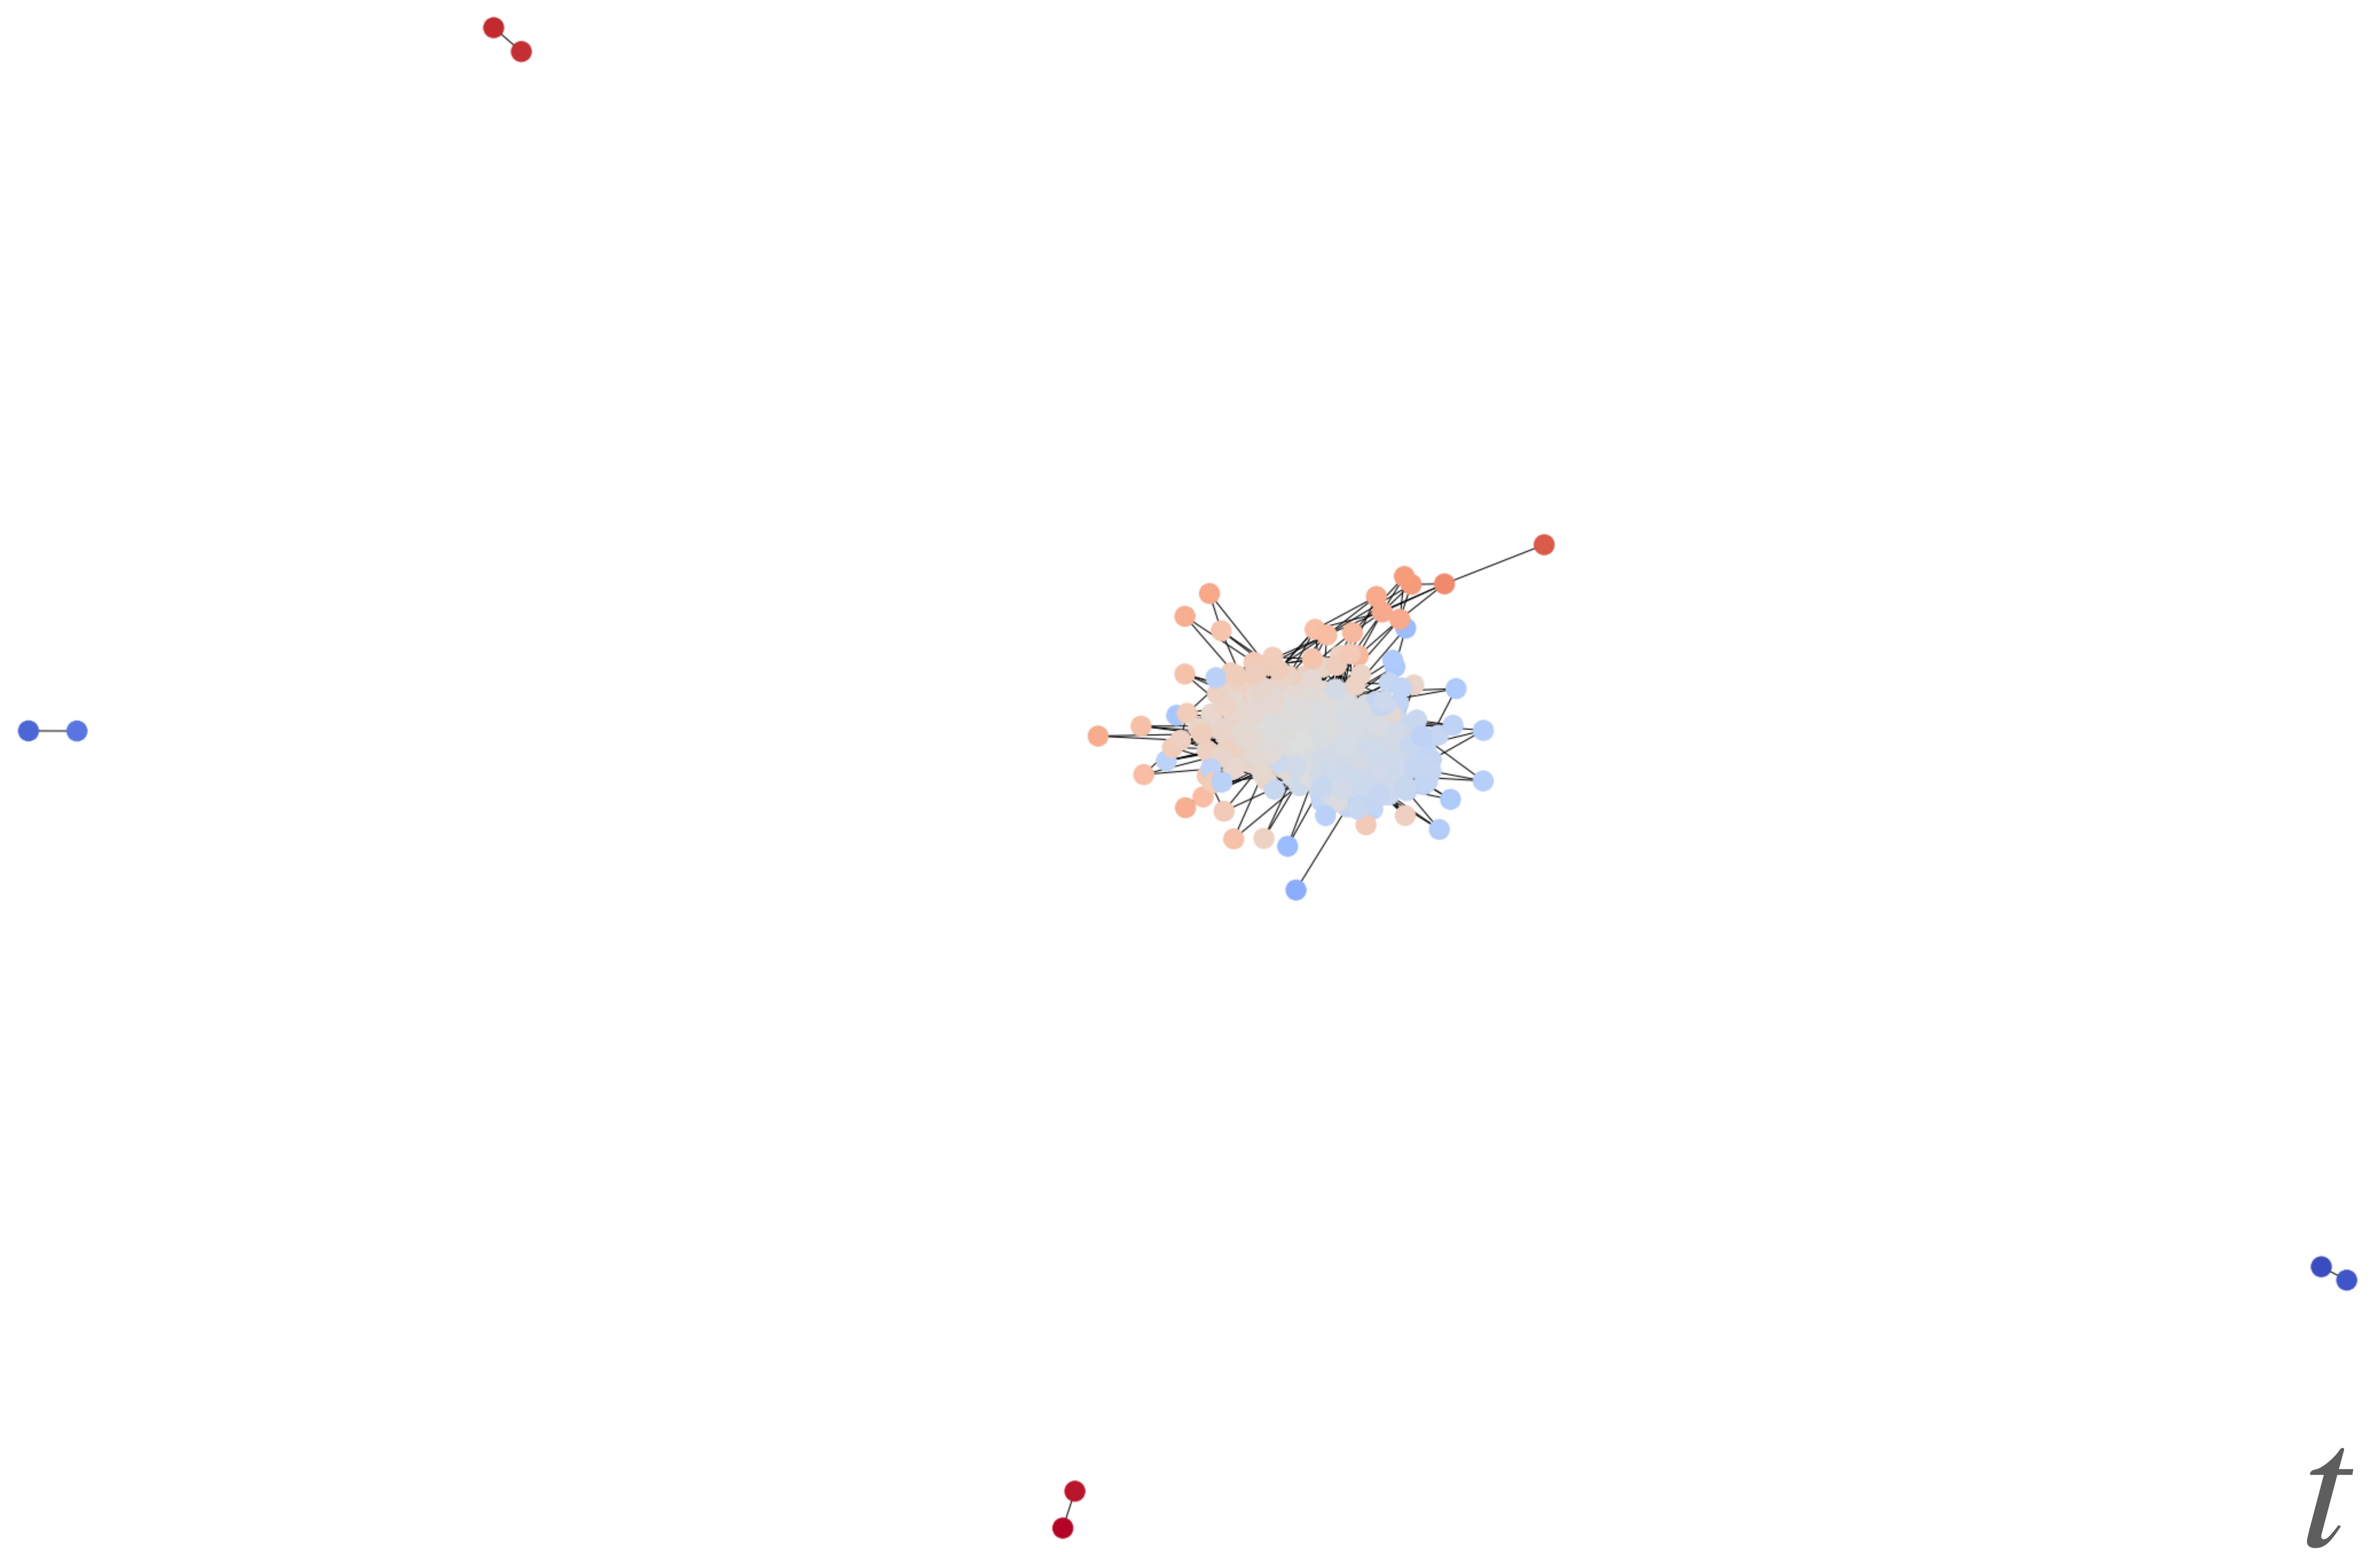

$t = \mathbf{50}$

10% Fixed Edges, E-R graph, $n = 1000$, standard edge dynamics

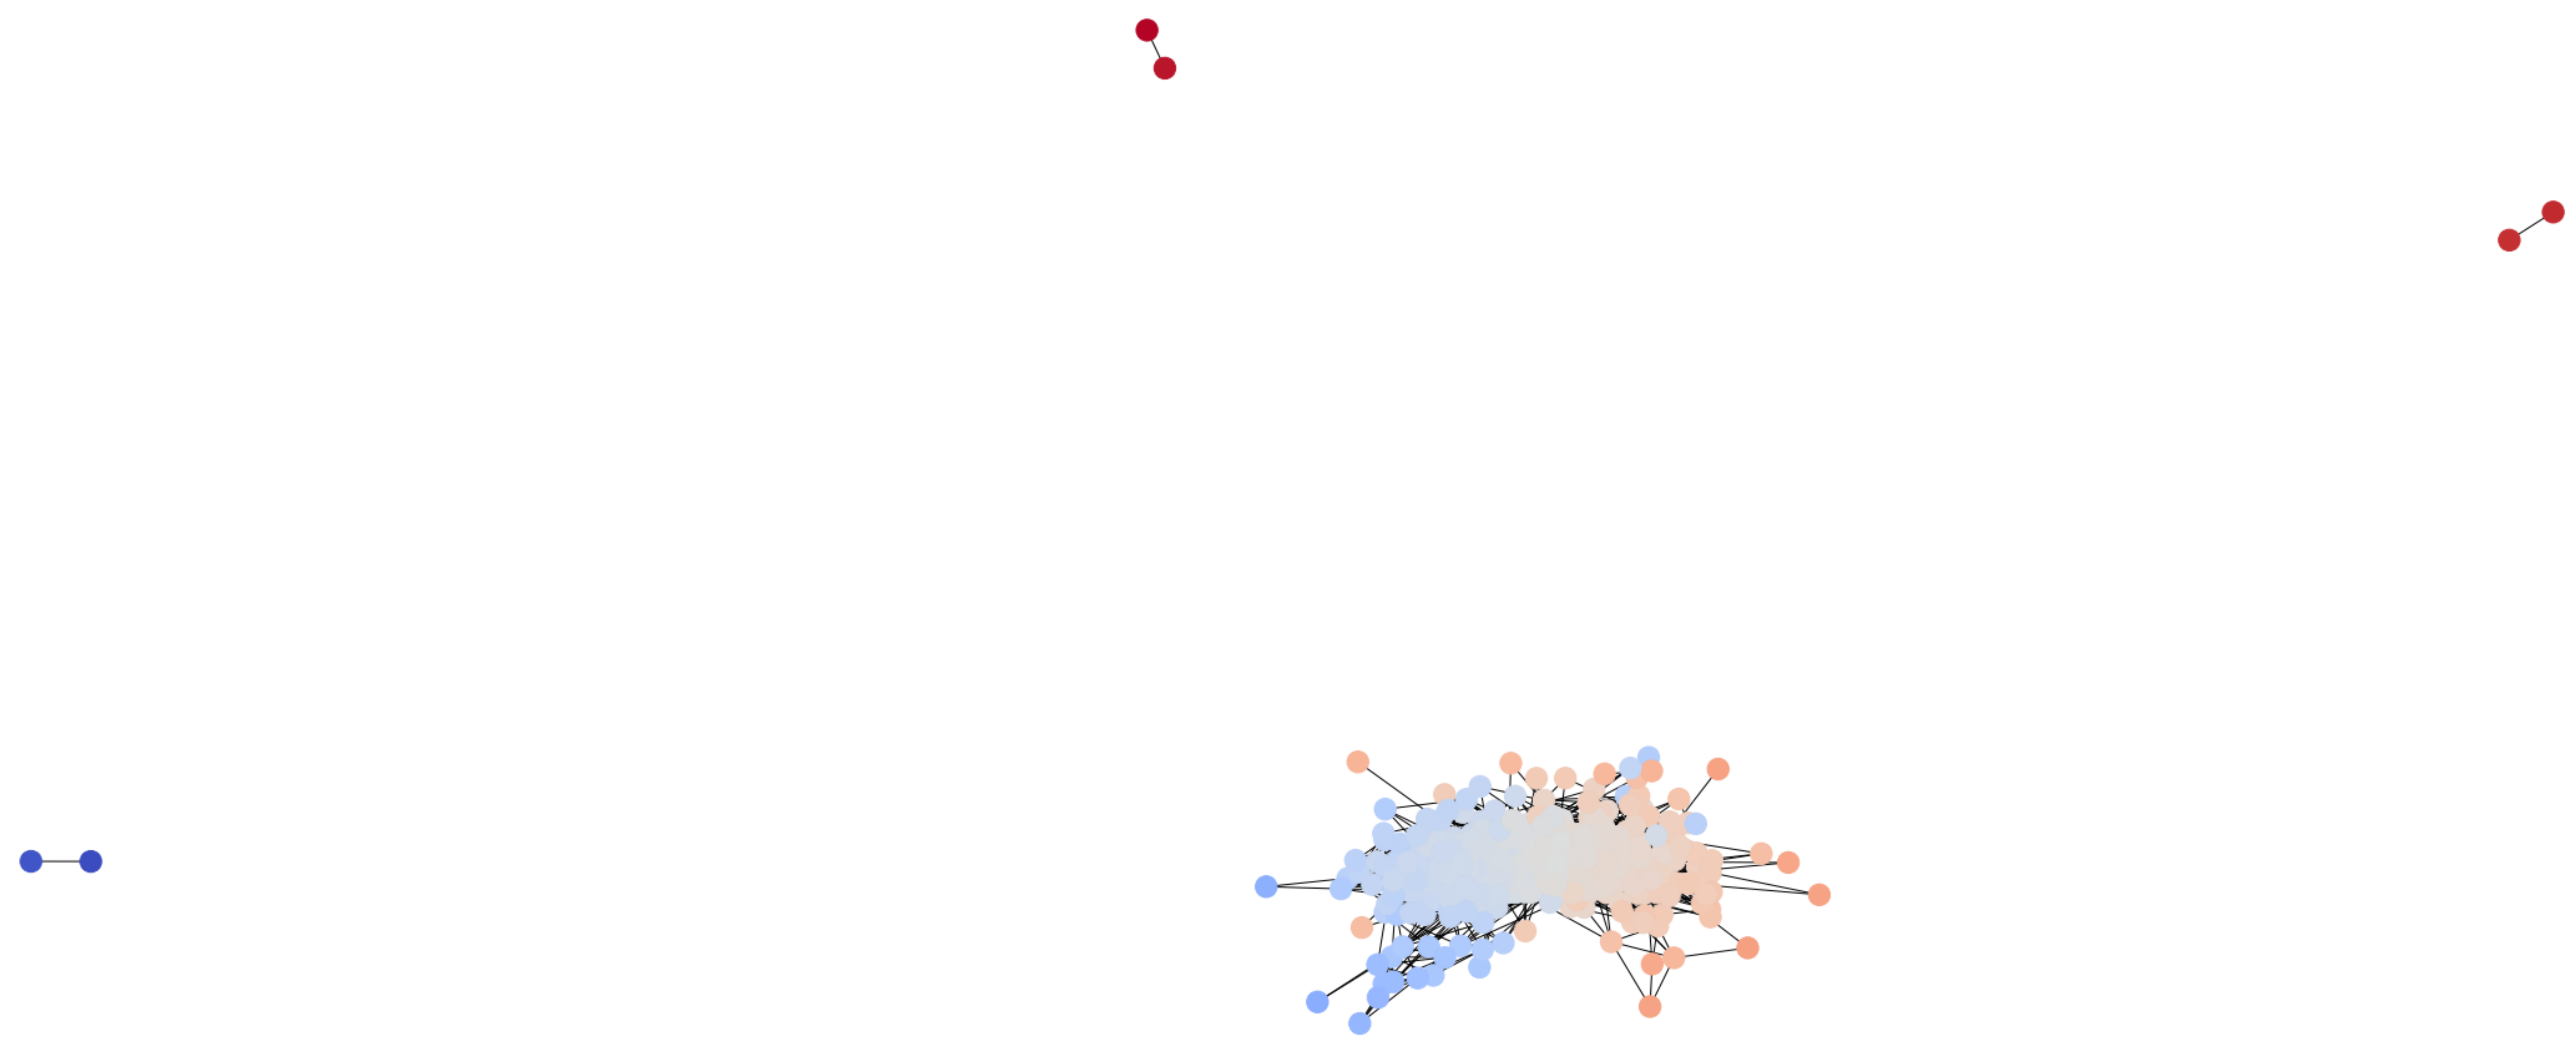

$t = 100$

10% Fixed Edges, E-R graph, $n = 1000$, standard edge dynamics

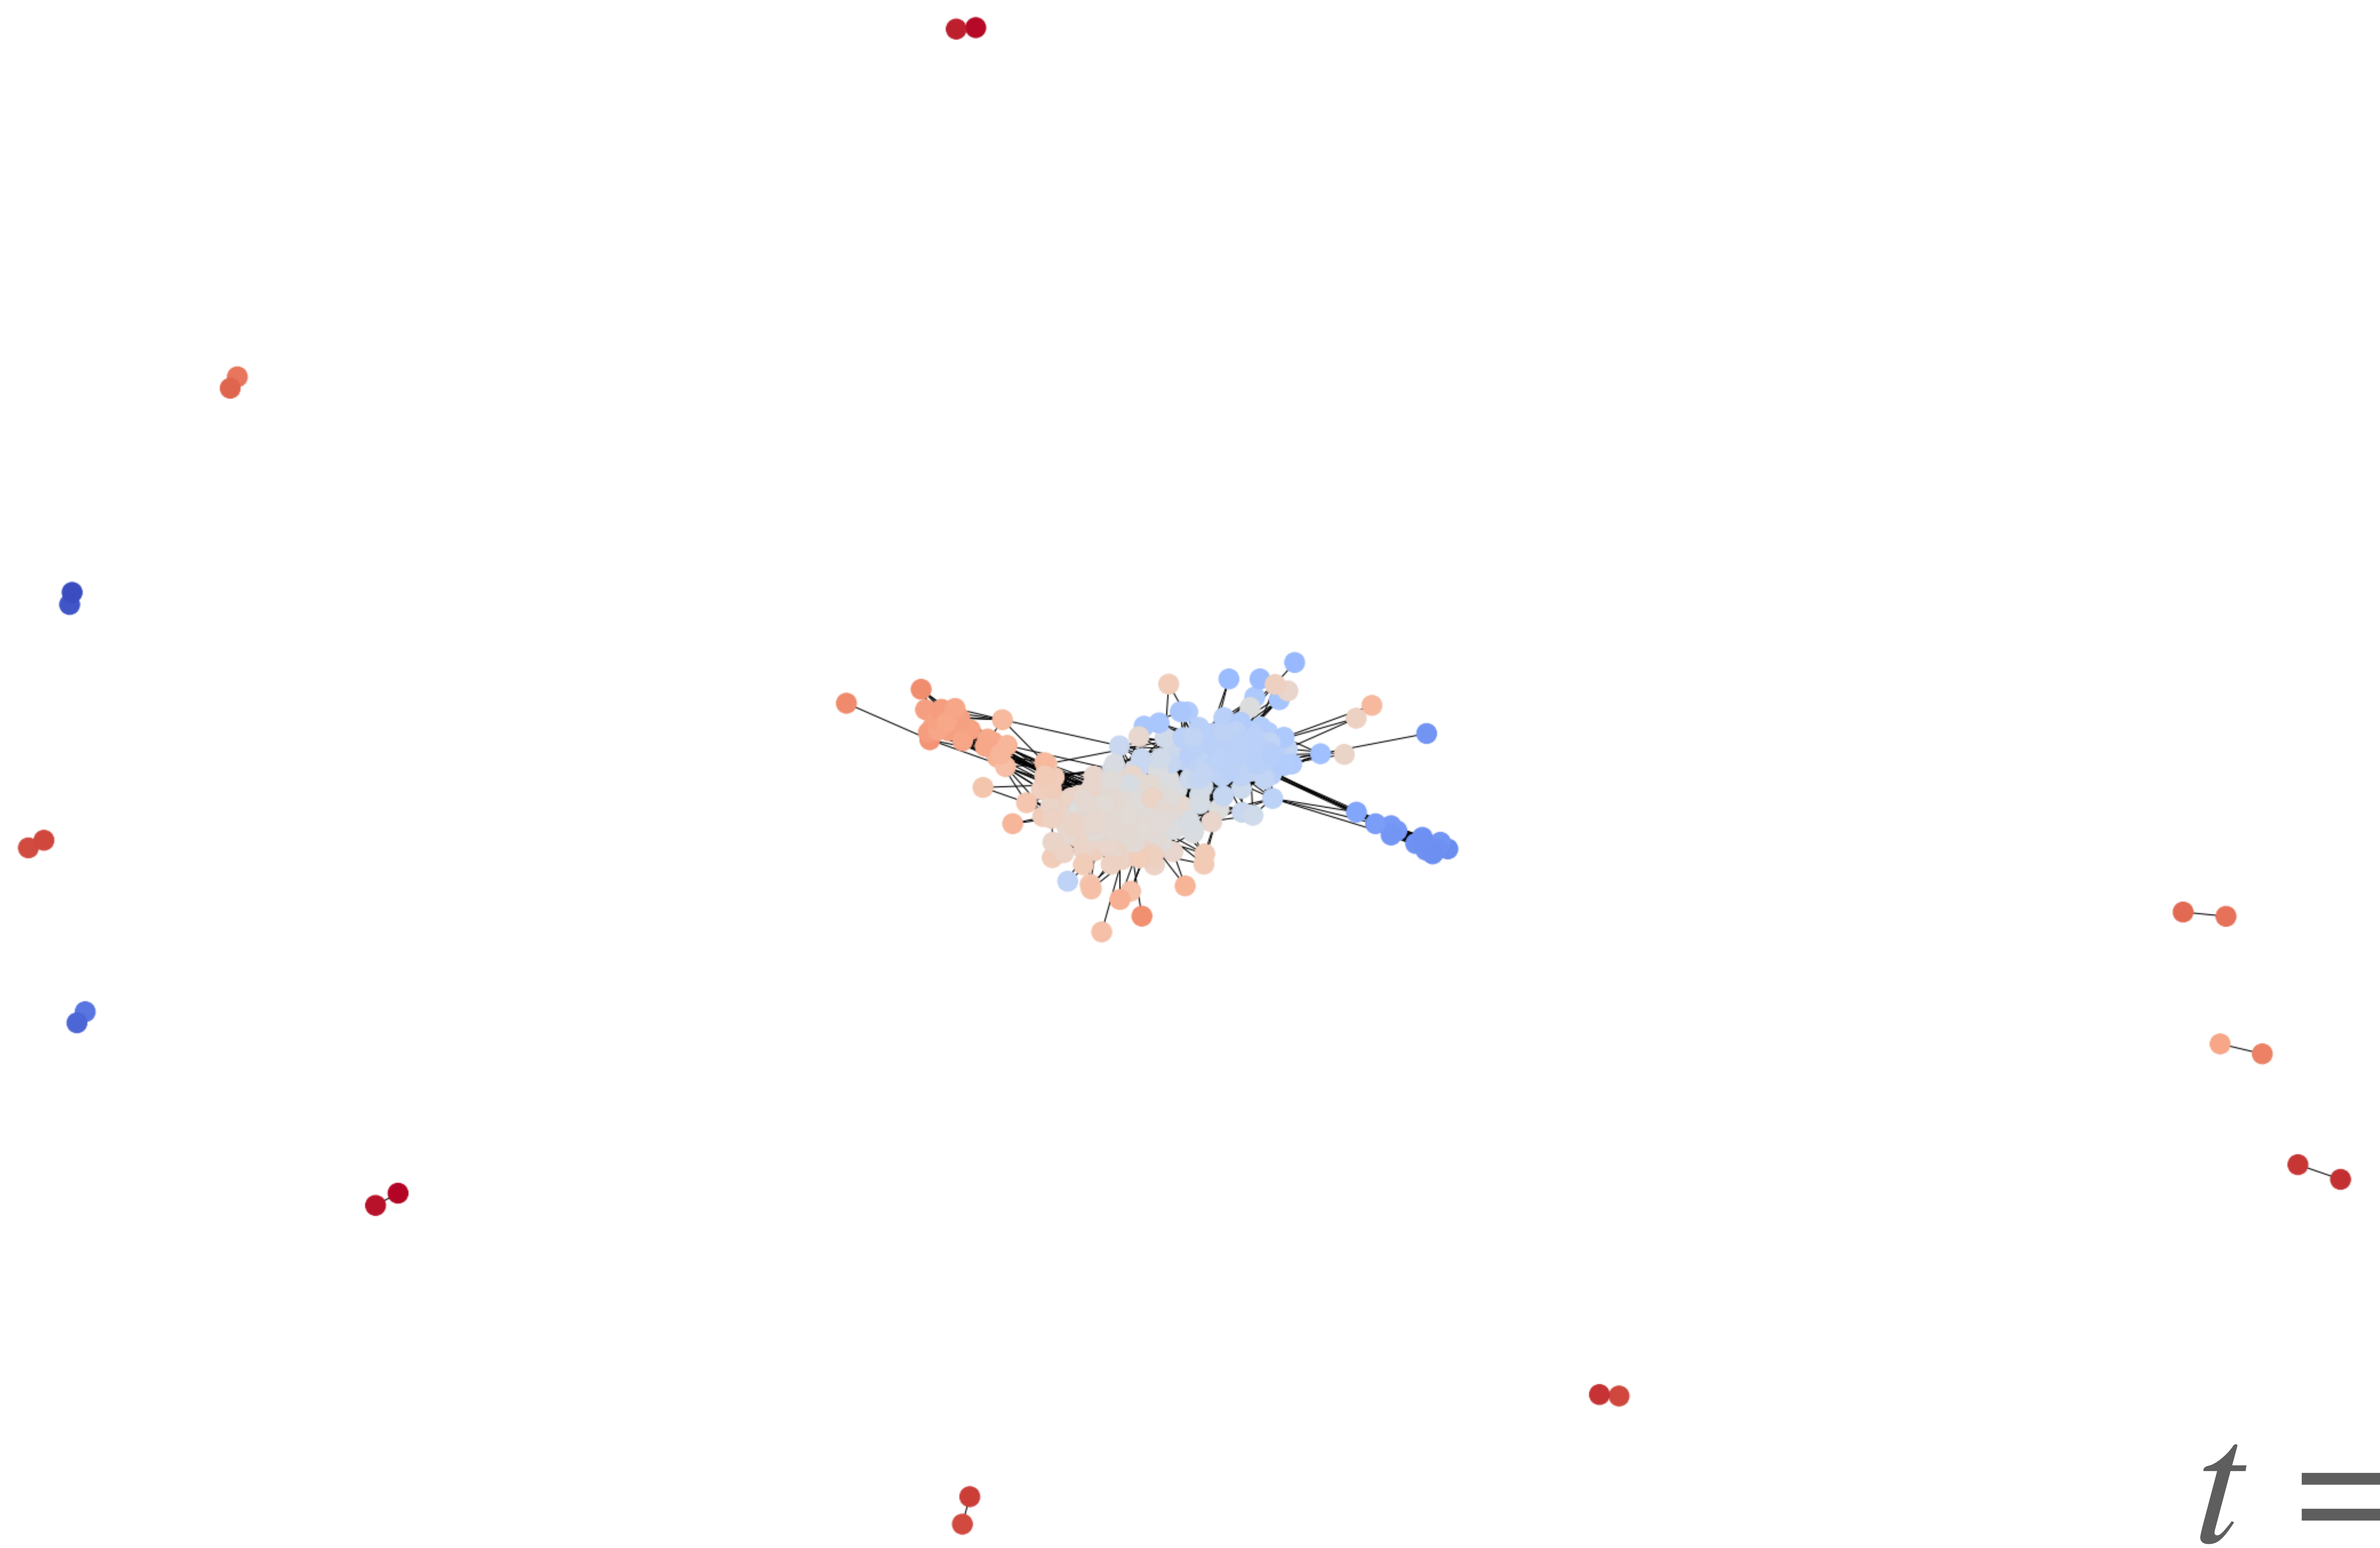

$t = 500$

Theoretical Analysis through Stochastic Block Model

- We view the underlying network as a Stochastic Block Model graph with two blocks, where each block holds a different innate opinion.
    - SBM allows us to specify two separate connection probabilities, p and q.
- Ties in with the classic problems of *community detection* and *spectral clustering.*

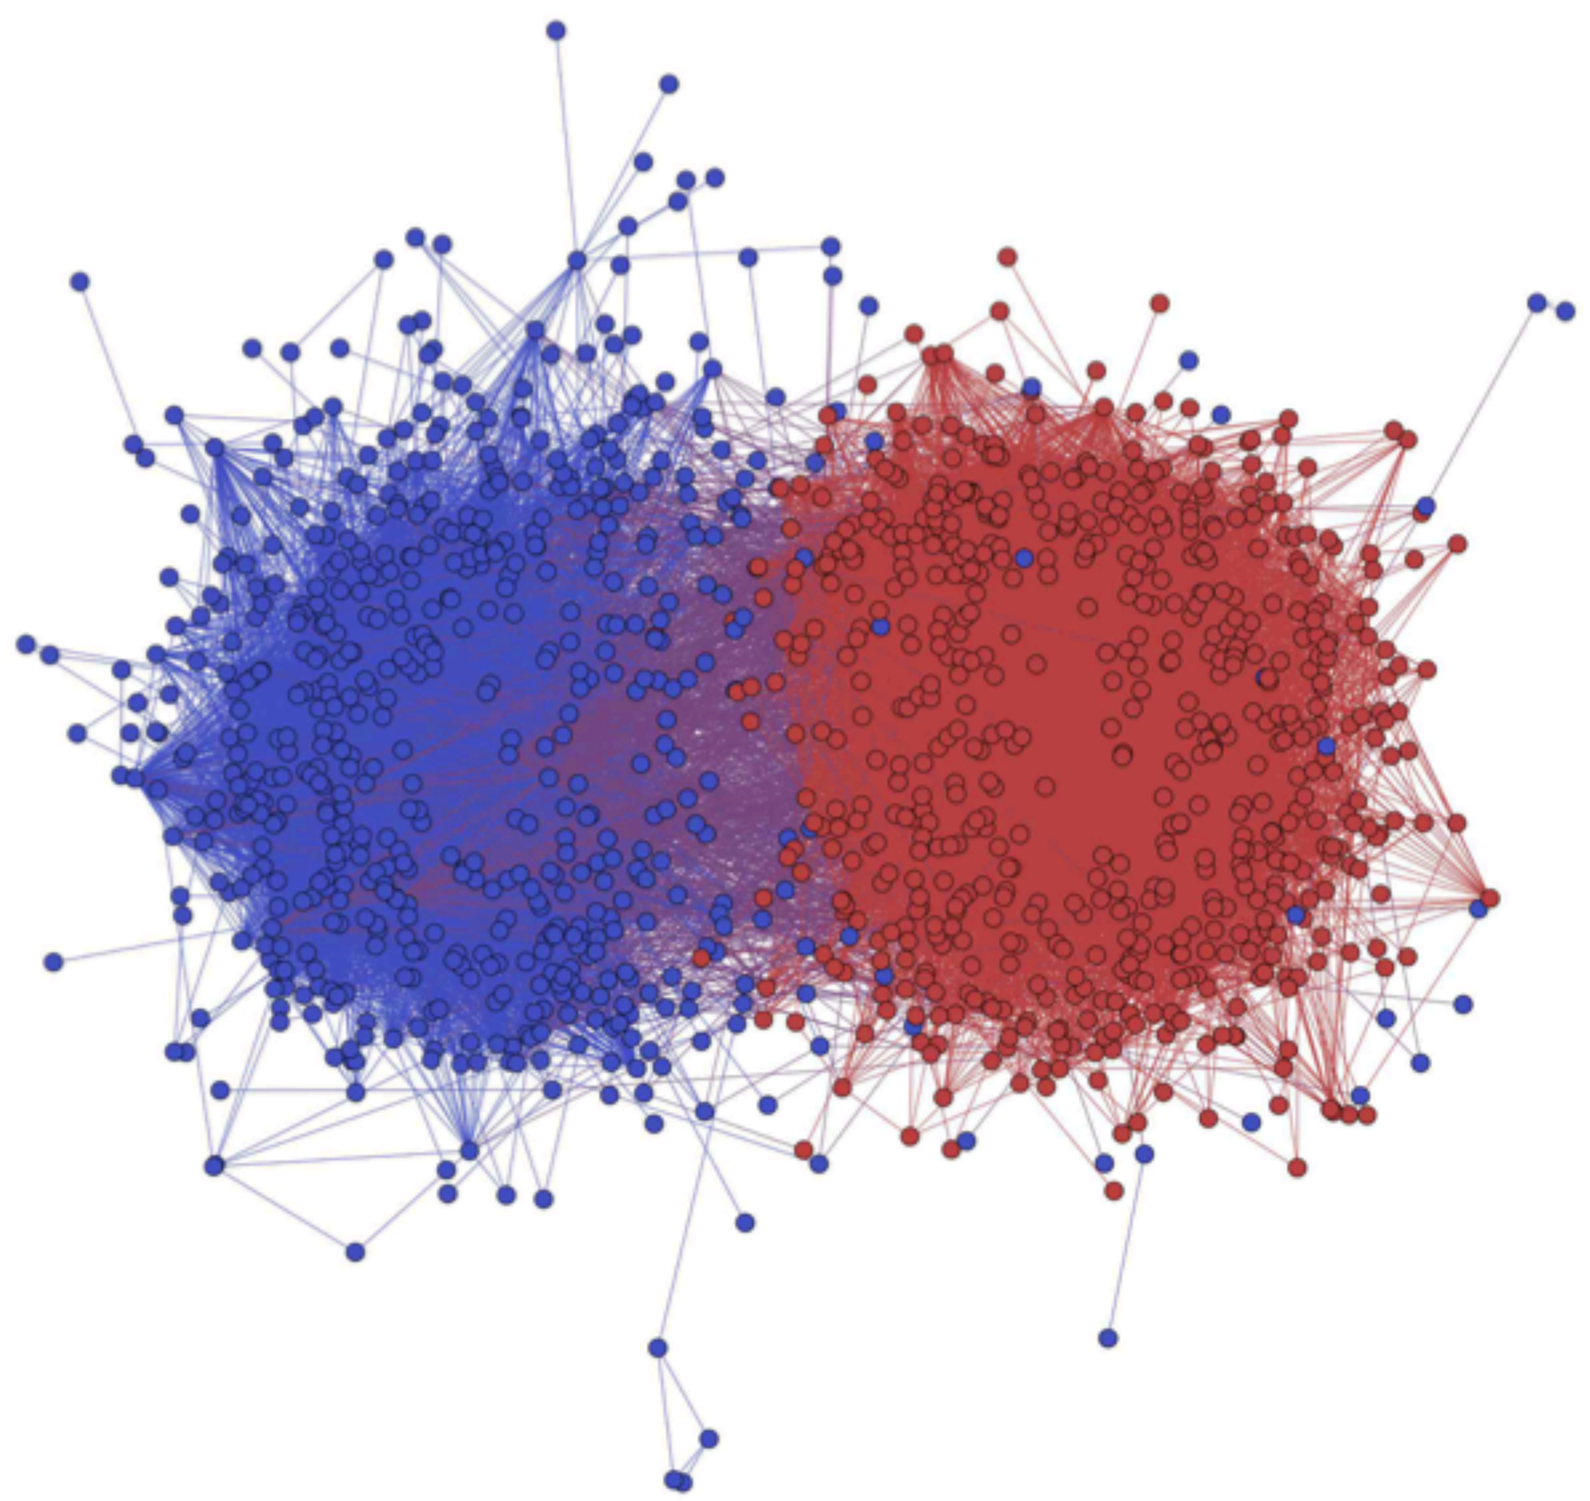

Abbe, "Community Detection and Stochastic Block Models", JMLR 2018

SBM Laplacian in Expectation

- We analyze the Laplacian $\bar{L}_G = \mathbb{E}\left[L_G\right]$ of an SBM graph **in expectation.**
- Then $\bar{L}_G$ takes the following form, where p is the in-group connection probability, and q is the out-group connection probability.
- Next, consider the discrete innate opinion vector $\vec{s}$, with length *n*.
  - $\vec{s}$ is *the second eigenvector of* $\bar{L}_G$ ***!***

$$\bar{L}_G = \begin{bmatrix} d(1) & \dots & -p & -q & \dots & -q \\ \vdots & d(2) & \vdots & \vdots & \ddots & \\ -p & \dots & \ddots & -q & & -q \\ -q & \dots & -q & \ddots & \dots & -p \\ \vdots & \ddots & & \vdots & \ddots & \vdots \\ -q & & -q & -p & \dots & d(n) \end{bmatrix}$$

$$\vec{s} = [1, \quad \dots \quad 1, \quad -1, \quad \dots \quad -1]$$

Expected SBM Laplacian w/ Fixed Edges

- For *uniform* fixed edges, set γ, which is some small connection probability independent of p *and* q.

- We analyze the expected Laplacian $\hat{L}_G = \mathbb{E}\left[L_G + L_F\right]$ of an SBM graph with probabilities $(p + \gamma)$ and $(q + \gamma)$.

- Note that if $\gamma > 0$, neither of the above quantities can equal 0…

$$\hat{L}_G = \begin{bmatrix} d(1) & \dots & -p-\gamma & -q-\gamma & \dots & -q-\gamma \\ \vdots & d(2) & \vdots & \vdots & \ddots & \\ -p-\gamma & \dots & \ddots & -q-\gamma & & -q-\gamma \\ -q-\gamma & \dots & -q-\gamma & \ddots & \dots & -p-\gamma \\ \vdots & \ddots & & \vdots & \ddots & \vdots \\ -q-\gamma & & -q-\gamma & -p-\gamma & \dots & d(n) \end{bmatrix}$$

$$\vec{s} = [1, \quad \dots \quad 1, \quad -1, \quad \dots \quad -1]$$

Deriving P(L) and D(L) using $\bar{L}_G$

$$\lambda_{n-1}(\bar{L}_G) = \frac{n(p+q)}{2} - \frac{n(p-q)}{2} = qn \qquad \lambda_2((\bar{L}_G + \mathbf{I})^{-1}) = \frac{1}{qn+1} \leftrightarrow \overrightarrow{v_2} = \vec{s}$$

$$P(L, \vec{s}) = ||\vec{z}||_2^2 = \vec{s}^T(\mathbf{I} + L)^{-2}\vec{s} \qquad D(L, s) = \vec{s}^T(\mathbf{I} + L)^{-1}L(\mathbf{I} + L)^{-1}\vec{s}$$

$$(\bar{L}_G + \mathbf{I})^{-1}\vec{s} = \frac{1}{qn+1} \cdot \vec{s} \qquad \bar{L}_G(\bar{L}_G + \mathbf{I})^{-1}\vec{s} = \frac{qn}{qn+1} \cdot \vec{s}$$

$$P(\bar{L}_G, s) = \frac{1}{(qn+1)^2} \cdot \vec{s}^T\vec{s} = \frac{n}{(qn+1)^2} \qquad D(\bar{L}_G, s) = \frac{qn}{(qn+1)^2} \cdot \vec{s}^T\vec{s} = \frac{qn^2}{(qn+1)^2}$$

Checking these Derivations

- We generate SBM graphs with two blocks — discrete innate opinions $\in \{-1,1\}$.
- We change p and q over time, and show that reducing q appears similar to the behavior of our model.

$$P(L_G, s) = \vec{s}^T(\mathbf{I} + L)^{-2}\vec{s}$$

$$D(L_G, s) = \vec{s}^T(\mathbf{I} + L)^{-1}L(\mathbf{I} + L)^{-1}\vec{s}$$

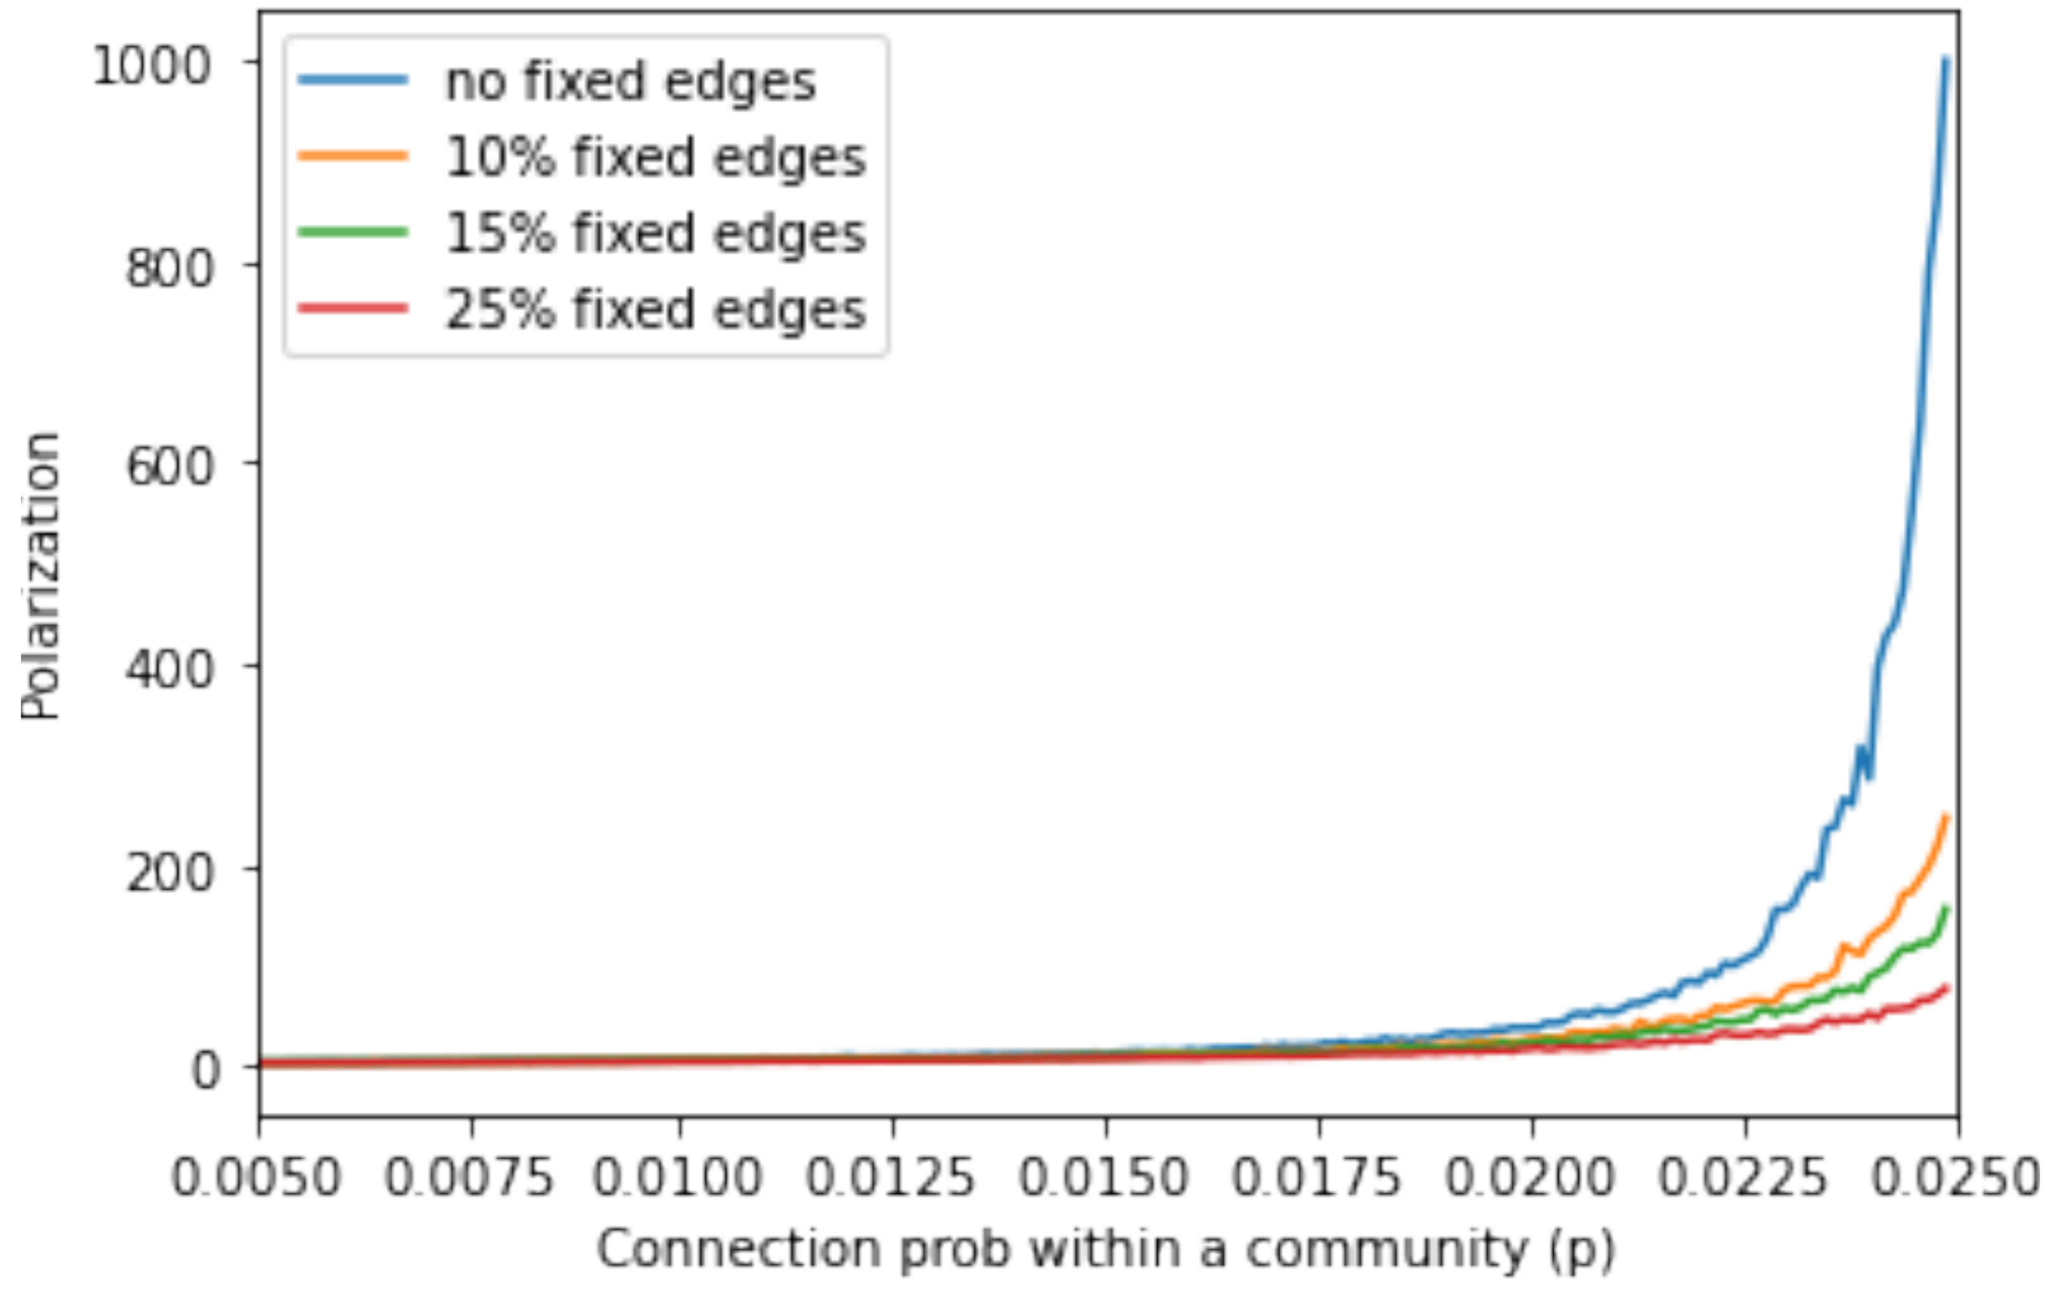

As q gets smaller, polarization increases, peaking when the two groups are fully disconnected without fixed edges.

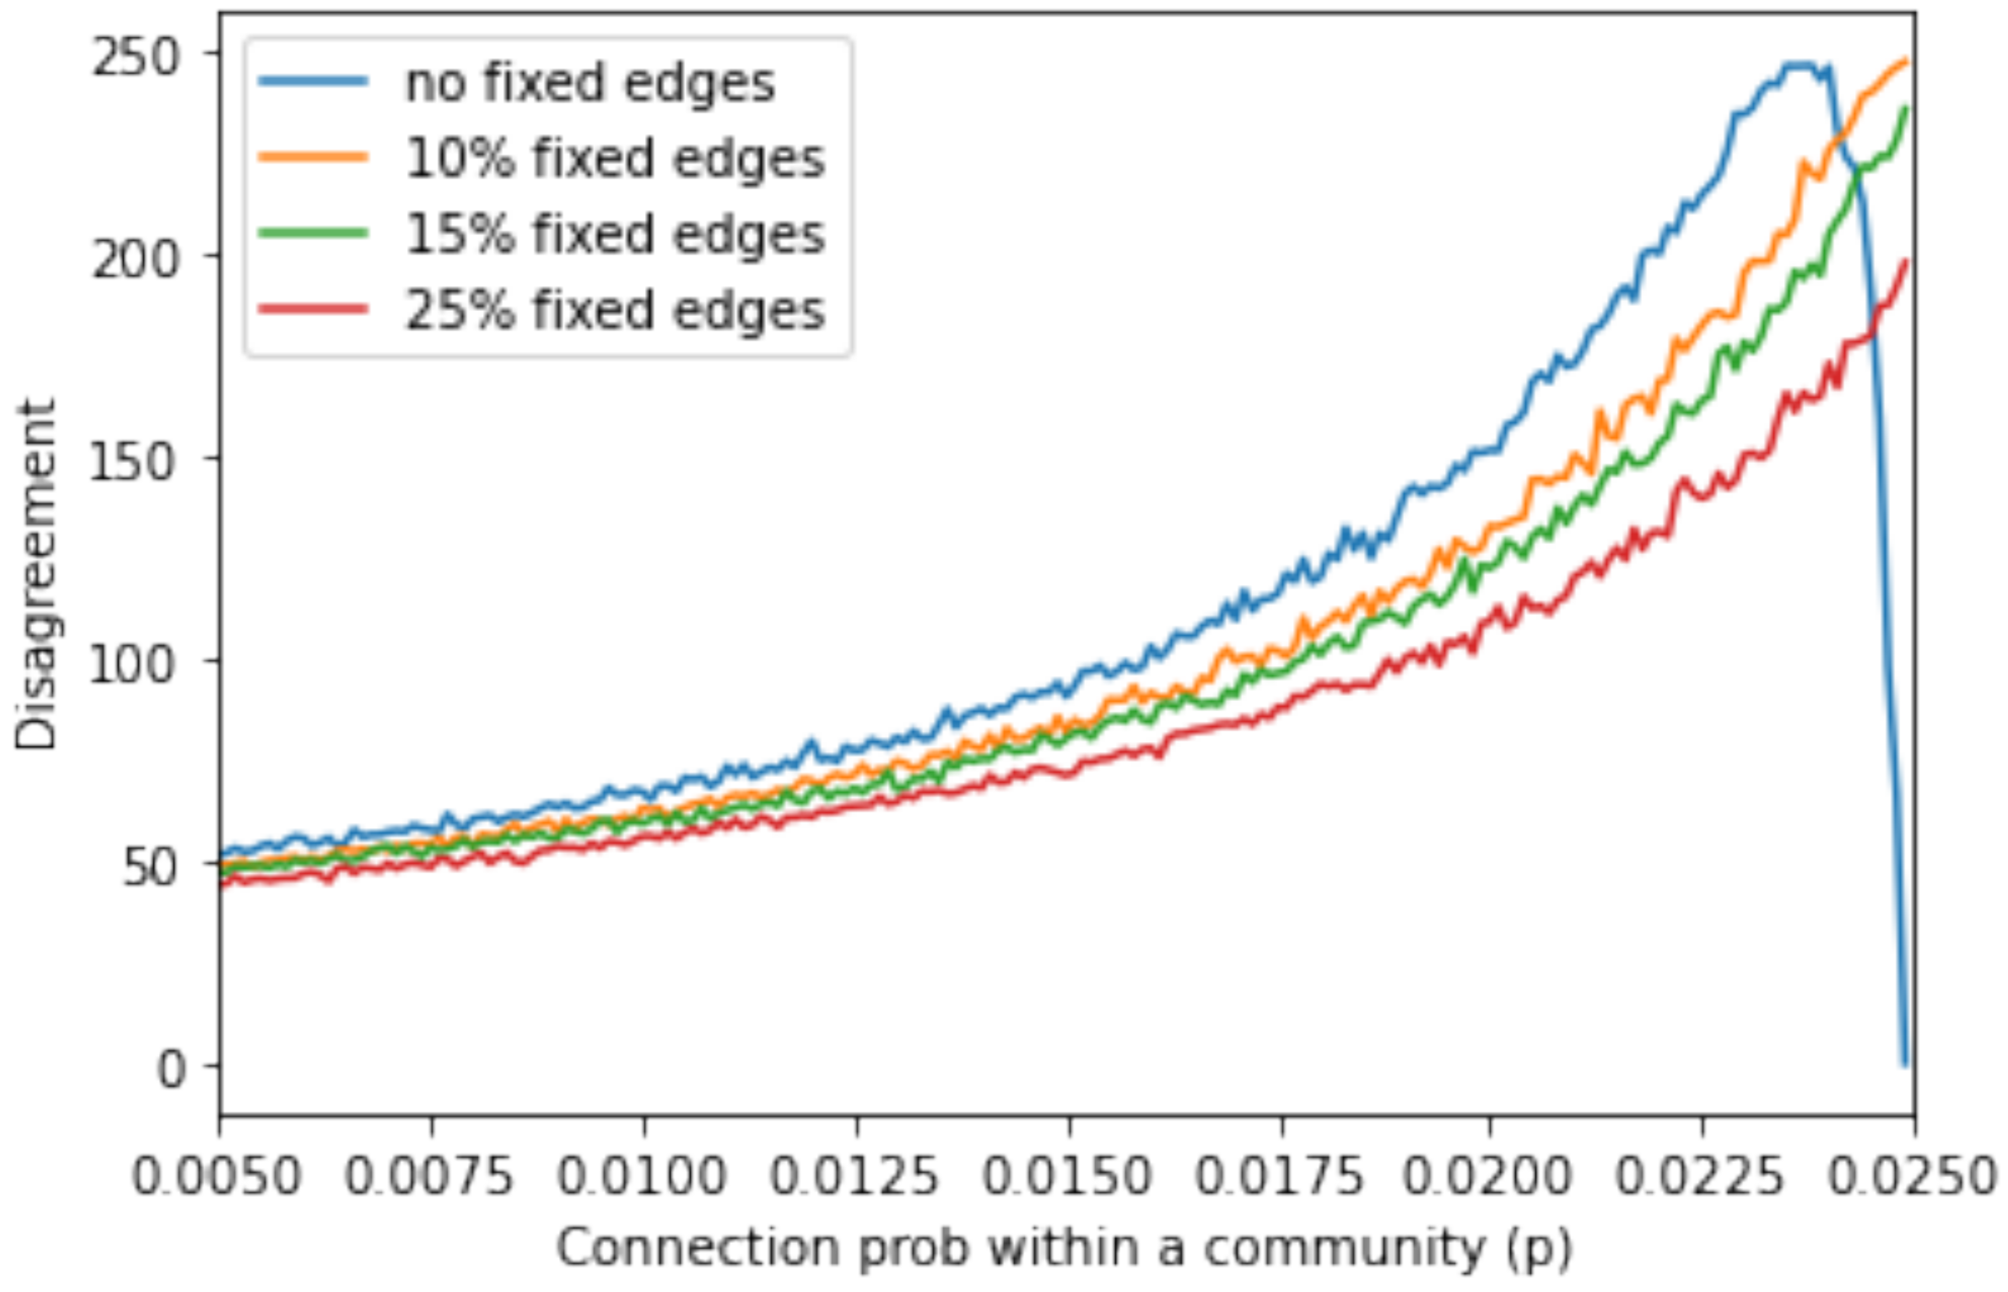

As q gets smaller, disagreement spikes initially and then falls when the blocks fully disconnect without fixed edges.

Checking these Derivations

- We generate SBM graphs with two blocks — discrete innate opinions $\in \{-1,1\}$.
- We change p and q over time, and show that reducing q appears similar to the behavior of our model.

$$P(\hat{L}_G, \bar{s}) = \frac{n}{((q+\gamma)n+1)^2} \qquad D(\hat{L}_G, \bar{s}) = \frac{(q+\gamma)n^2}{((q+\gamma)n+1)^2}$$

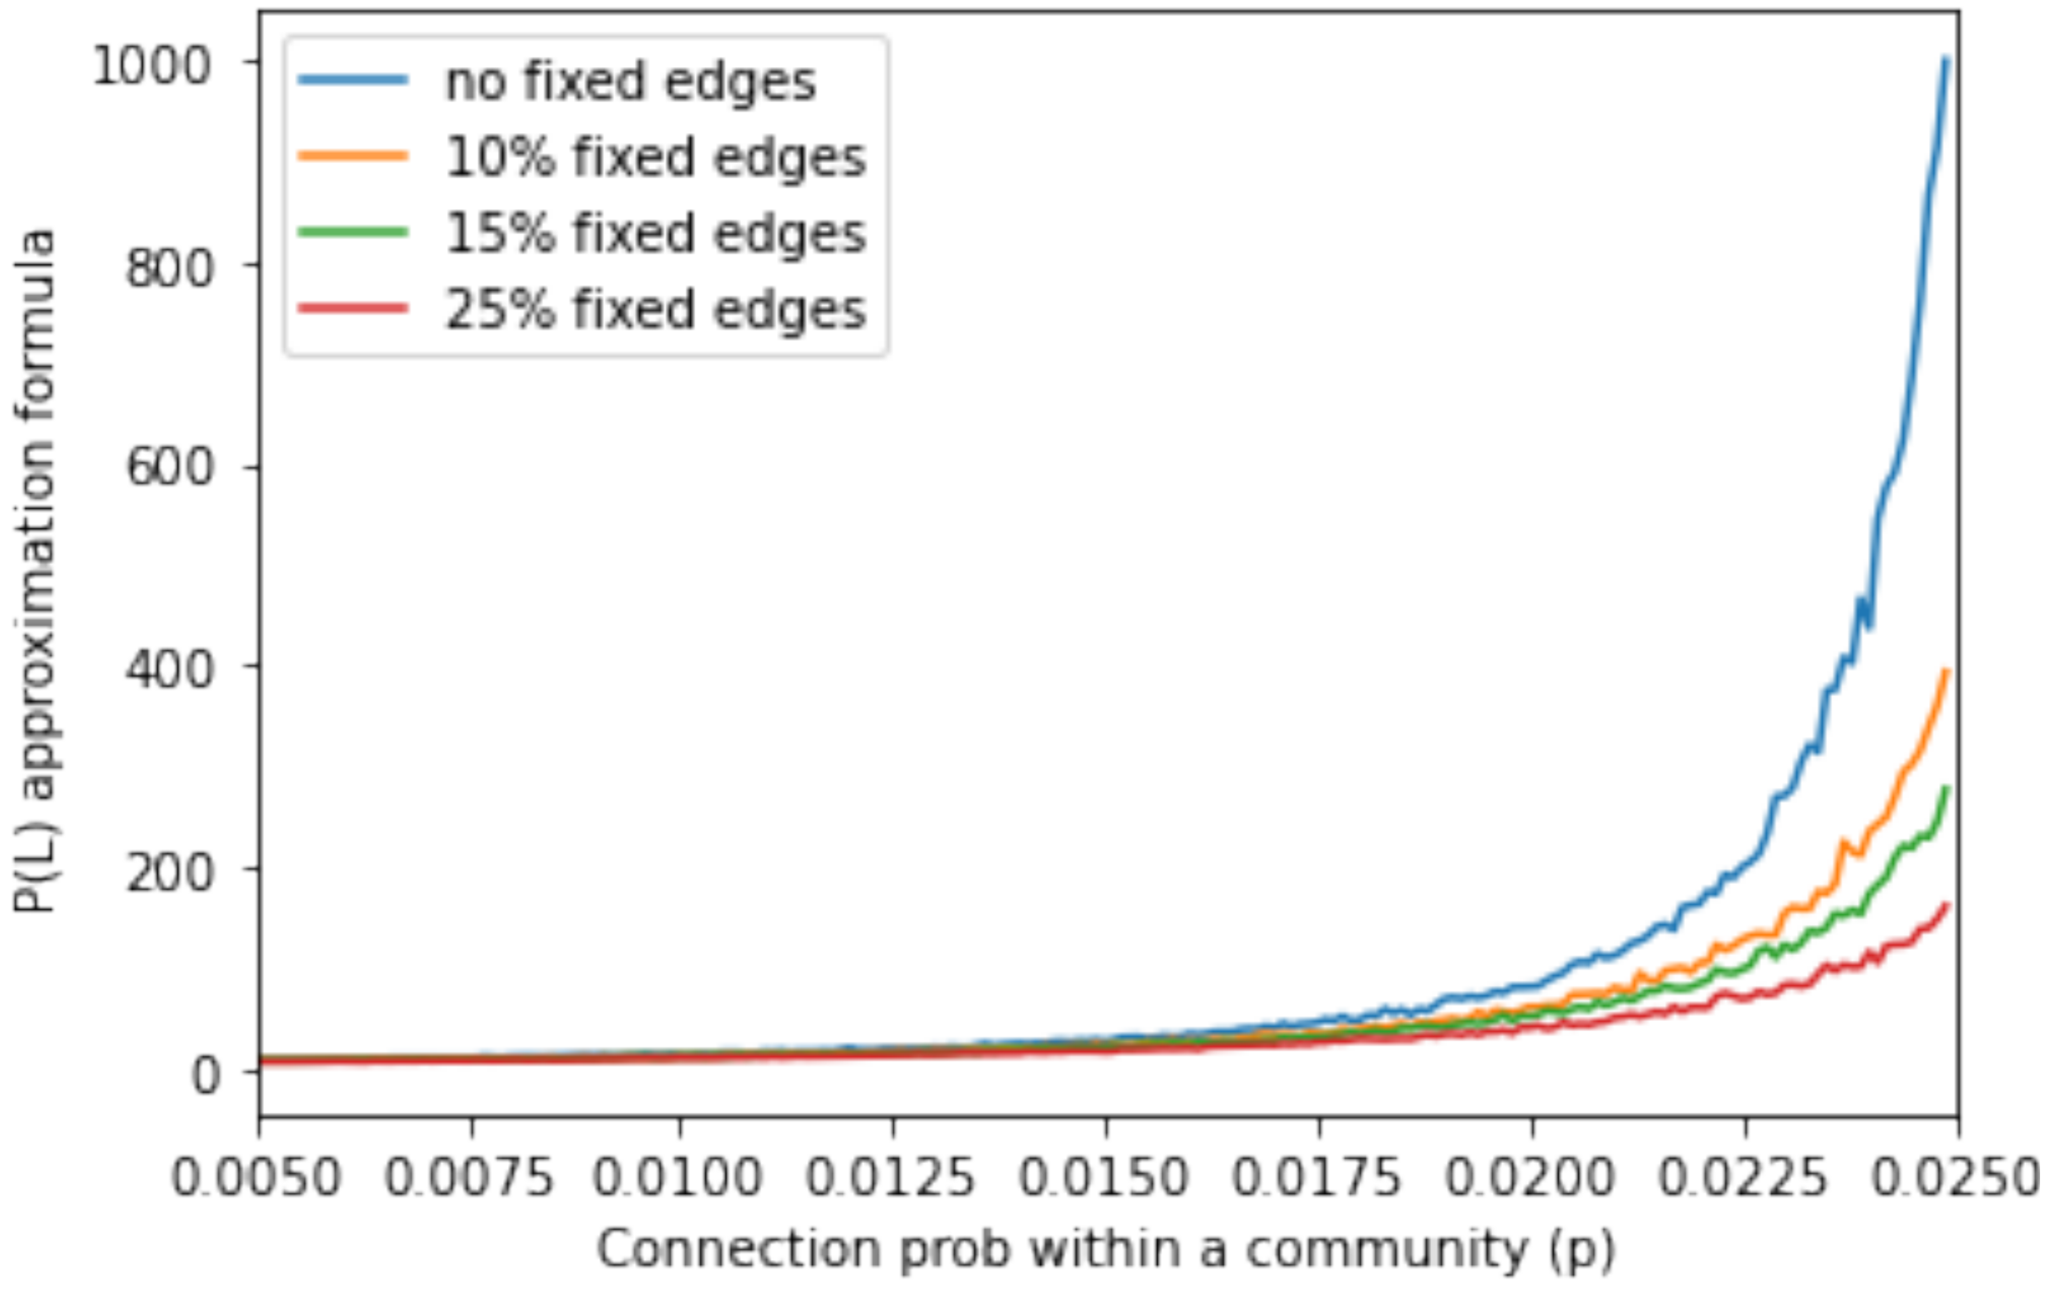

Polarization approximation — Tracks closely with the actual value for polarization.

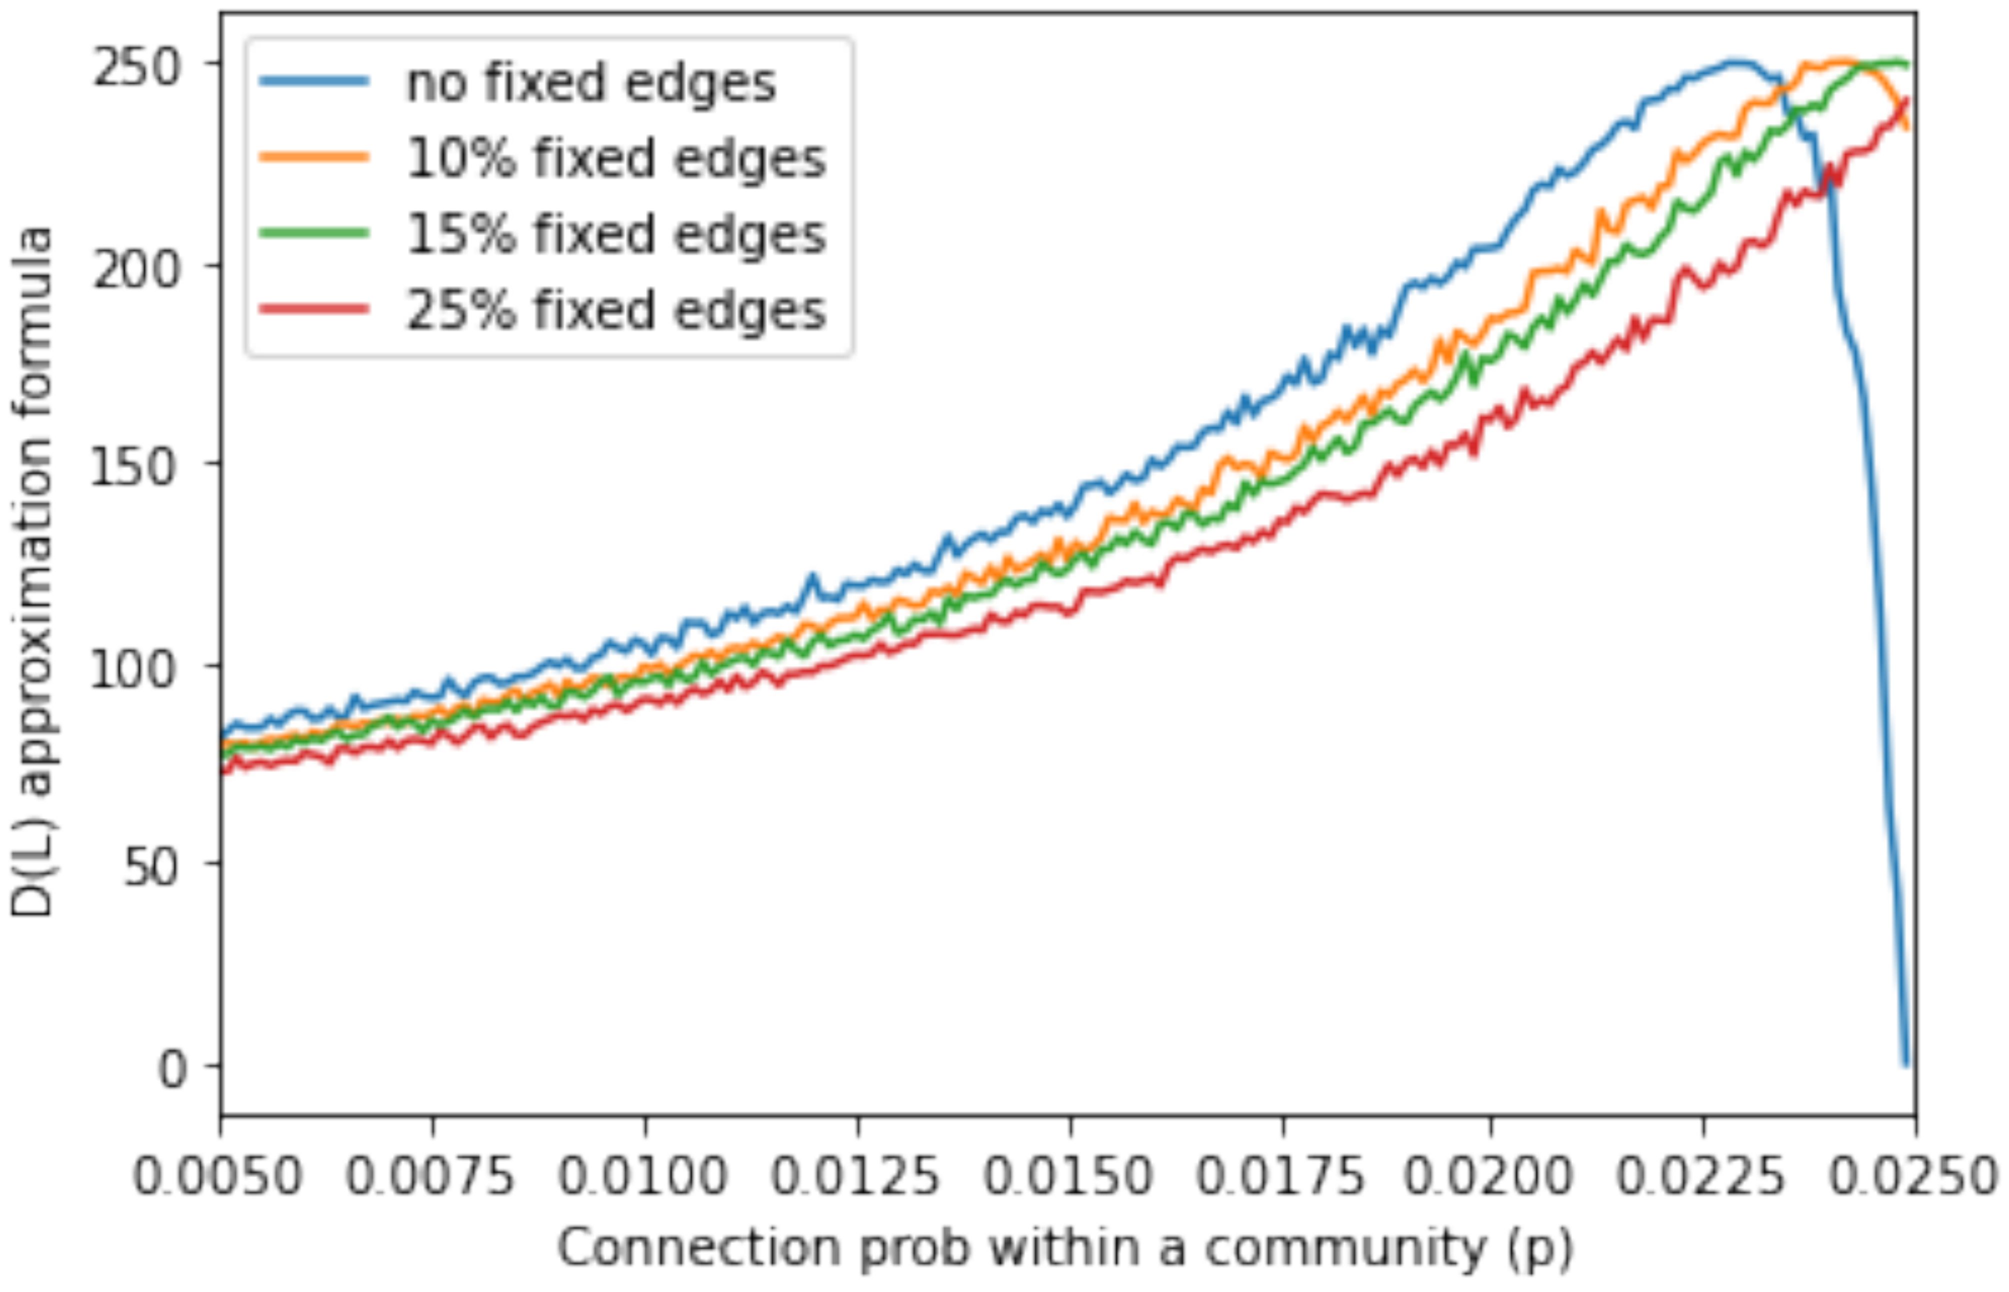

Disagreement approximation — Tracks closely with the actual value for disagreement.

Real-World Verification

- We have this model and understand it theoretically; we want to get some idea of how *realistic it is* — verify our model against *real world data*.
- Related works attempt to do this by quantifying graph structure, through measures such as the *global clustering coefficient, degree distribution, and small world quotient.*
- We explore two methods:
  - The first method leverages graph measures to show that our model creates realistic structures.
  - The second constructs a real-world graph based on temporal data, so we can analyze how it changes over time.

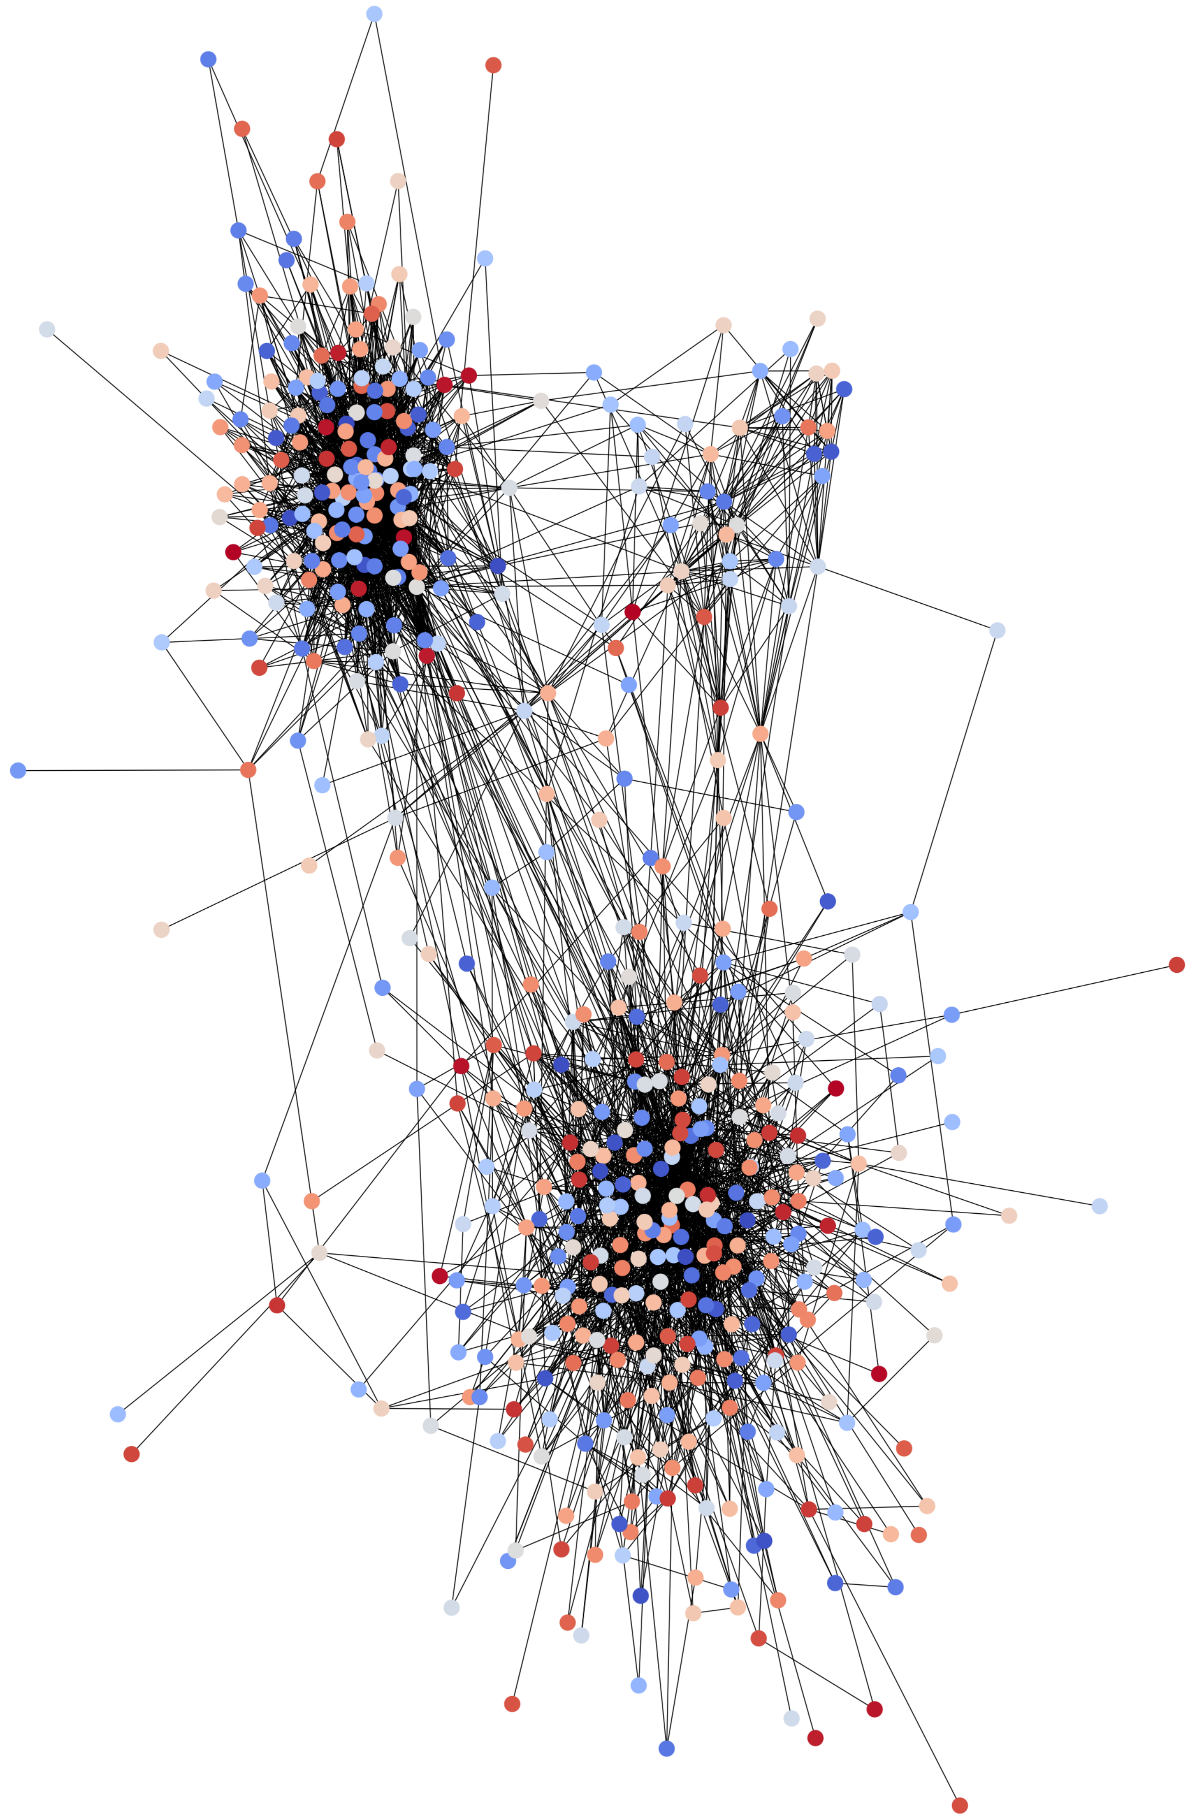

Twitter (Delhi 2013) Dataset
De et al., Learning Linear Influence Models in Social Networks from Transient Opinion Dynamics. 2019.

Verification through Graph Measures

- In this section, we study a real-world snapshot from Twitter, which captured interactions surrounding a 2013 Delhi election [De et al. '19]
    - $n = 531$ nodes $|E| = 3621$ edges
- We generate a graph, which has 531 nodes, ~3600 edges, and fixed edges in $\mathbf{p}_f = \{15\%, 25\%, 35\%\}$
    - Both Erdős–Rényi and Barabási–Albert graphs
- On these generated graphs, we simulate edge dynamics with both friend-of-friend recommendations and confirmation bias edge deletion

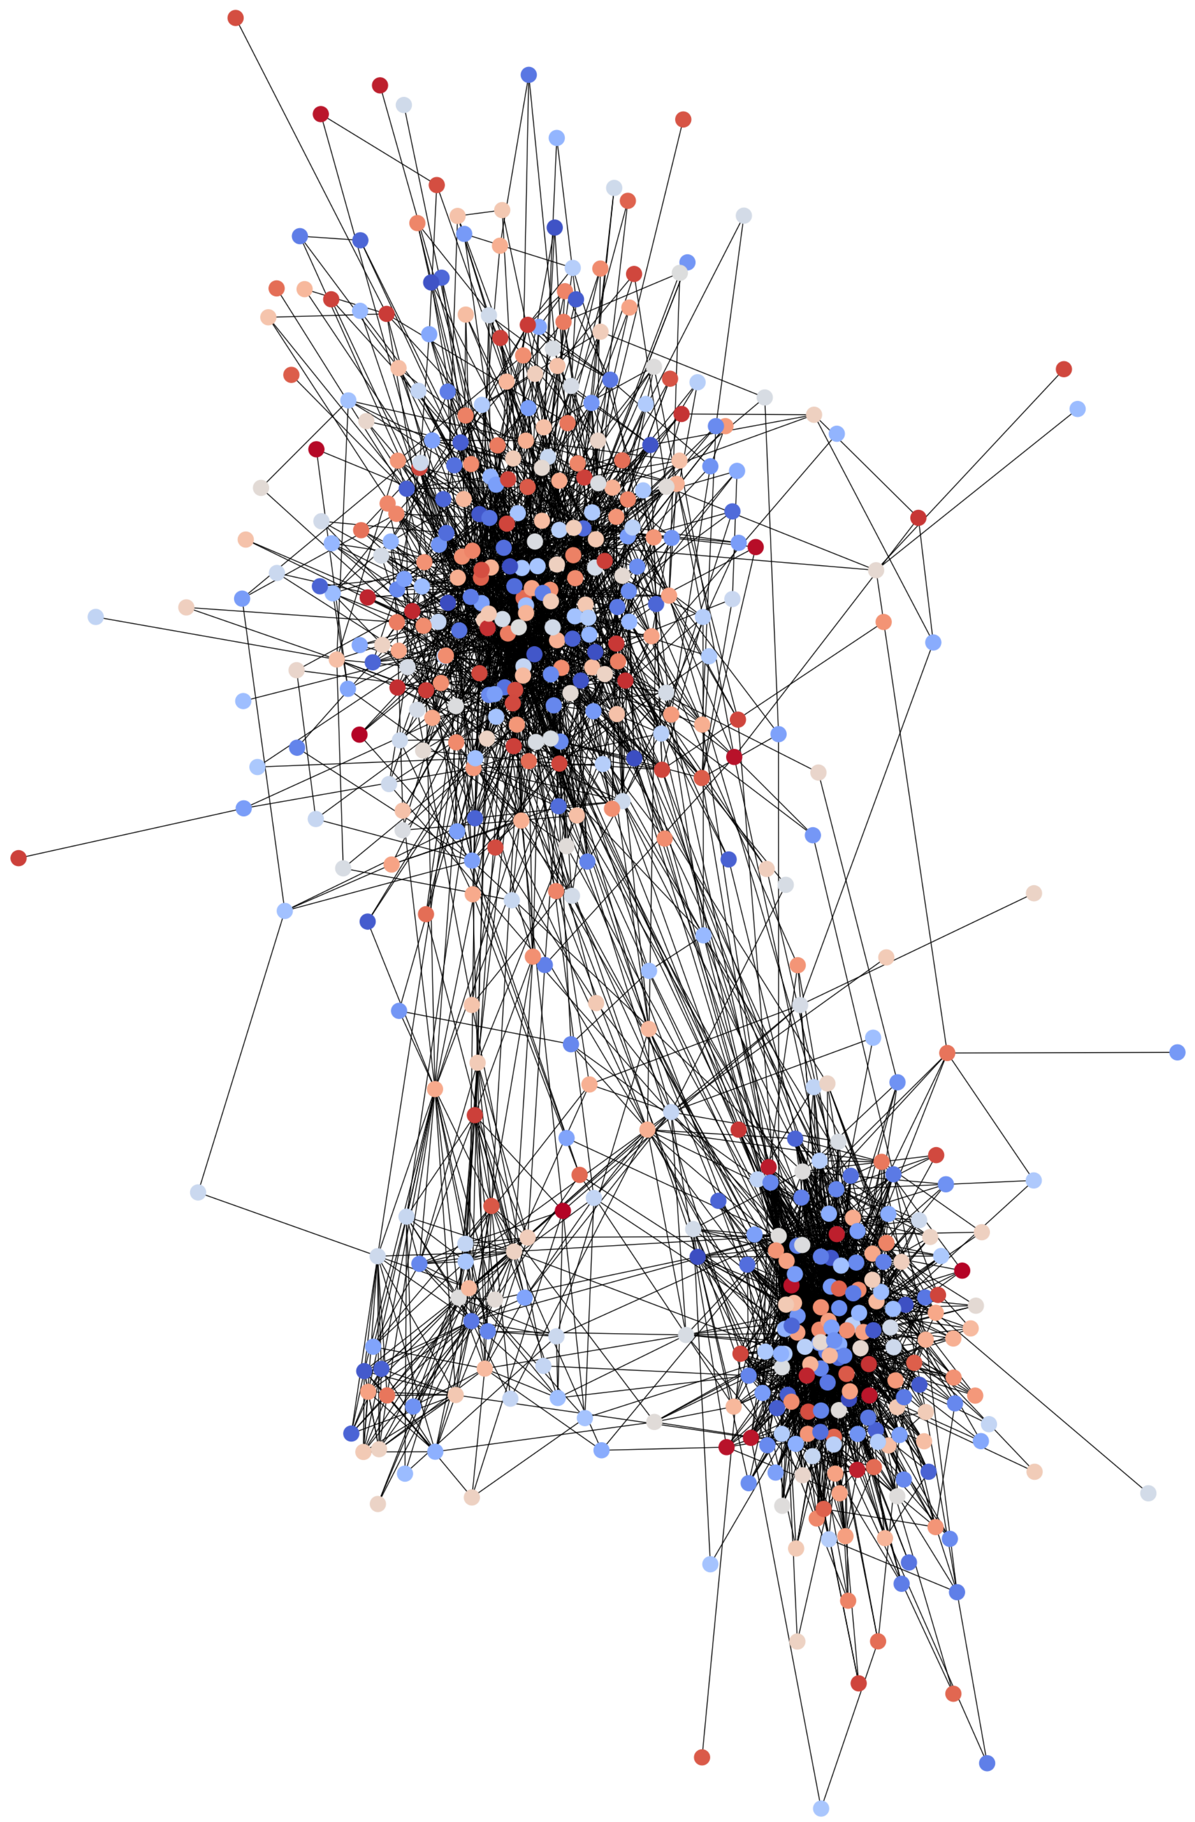

Twitter (Delhi 2013) Dataset
[De et al. '19]

Verification through Graph Measures — Clustering

$$\mathrm{CC}_{Glo} = \frac{3 \cdot \#\text{ of triangles} \in G}{\#\text{ of open \& closed triads} \in G}$$

- Different fixed edges yield different values for the global clustering coefficient — **use this value to *tune* against real data.**
  - *25% fixed edges is closest*
- Next we'll compare other measures between the steady-state snapshots and Twitter graph.
- *Twitter global clustering coefficient ~* ***0.227***

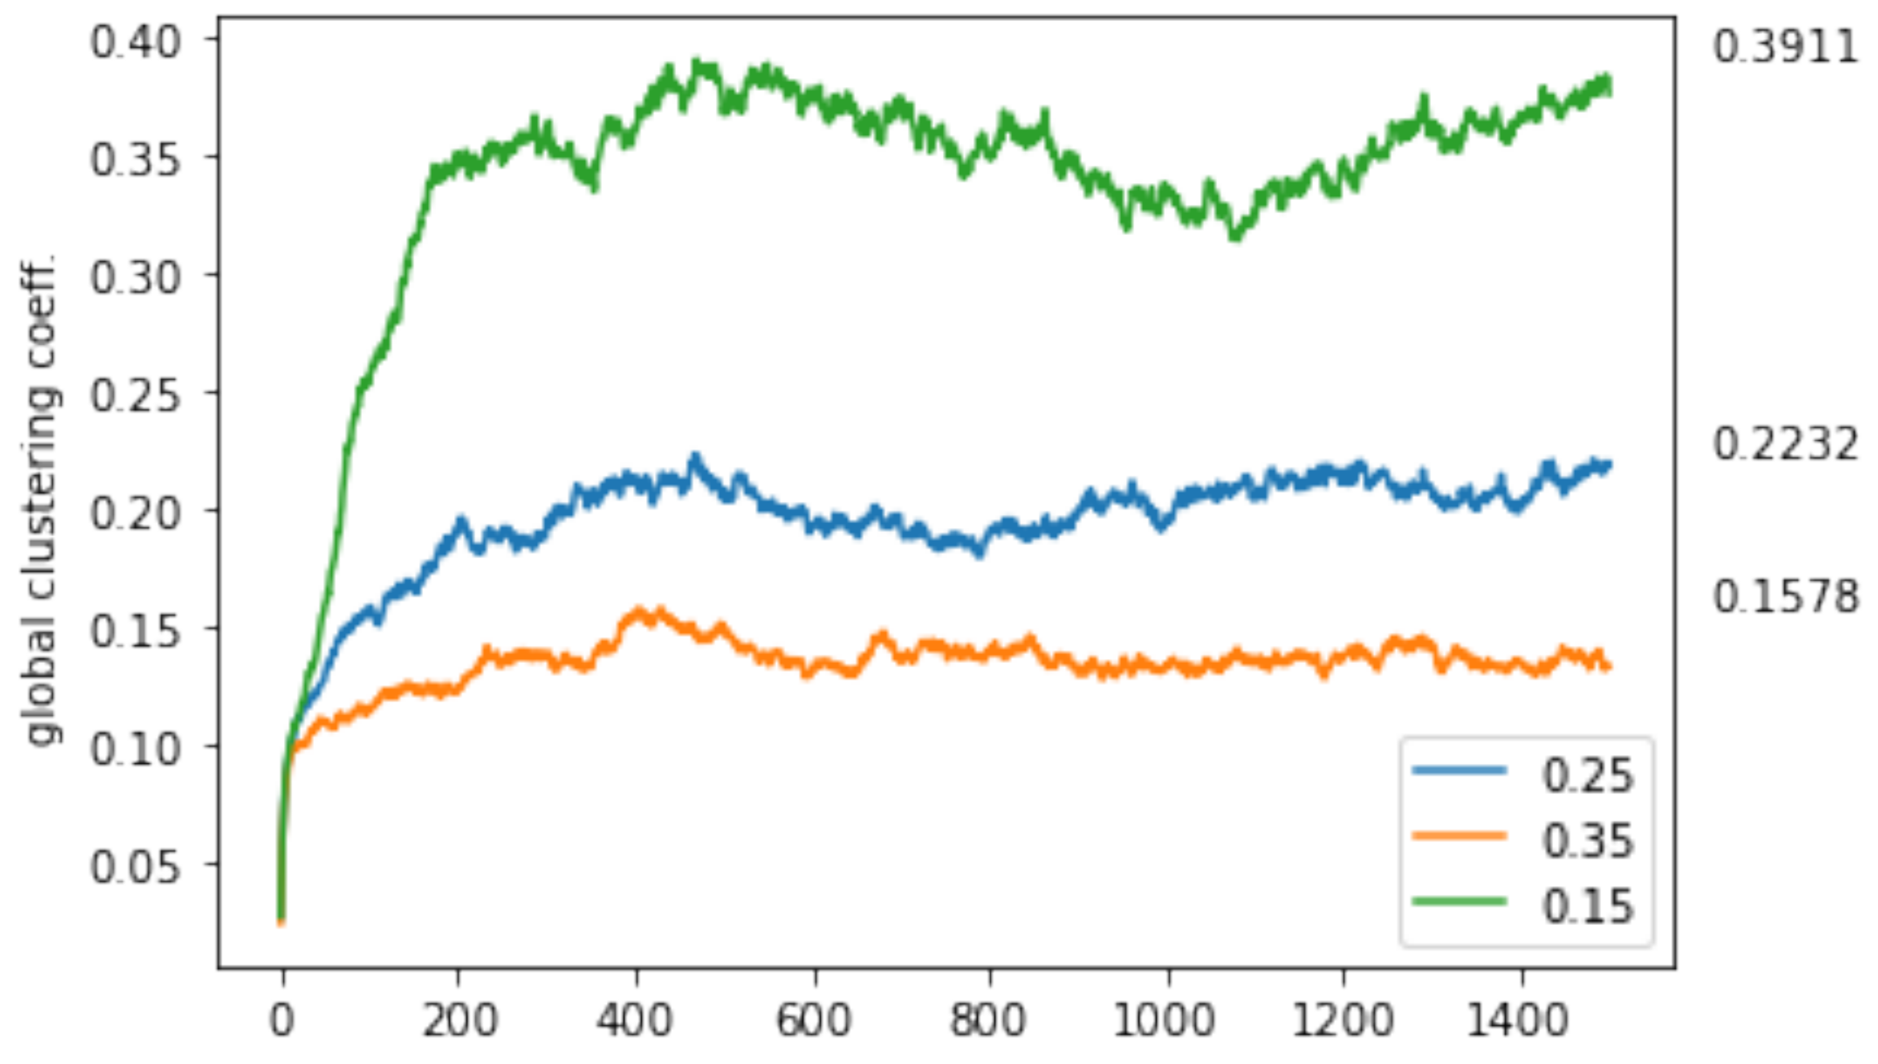

Global clustering coefficient over time for Erdős–Rényi graphs

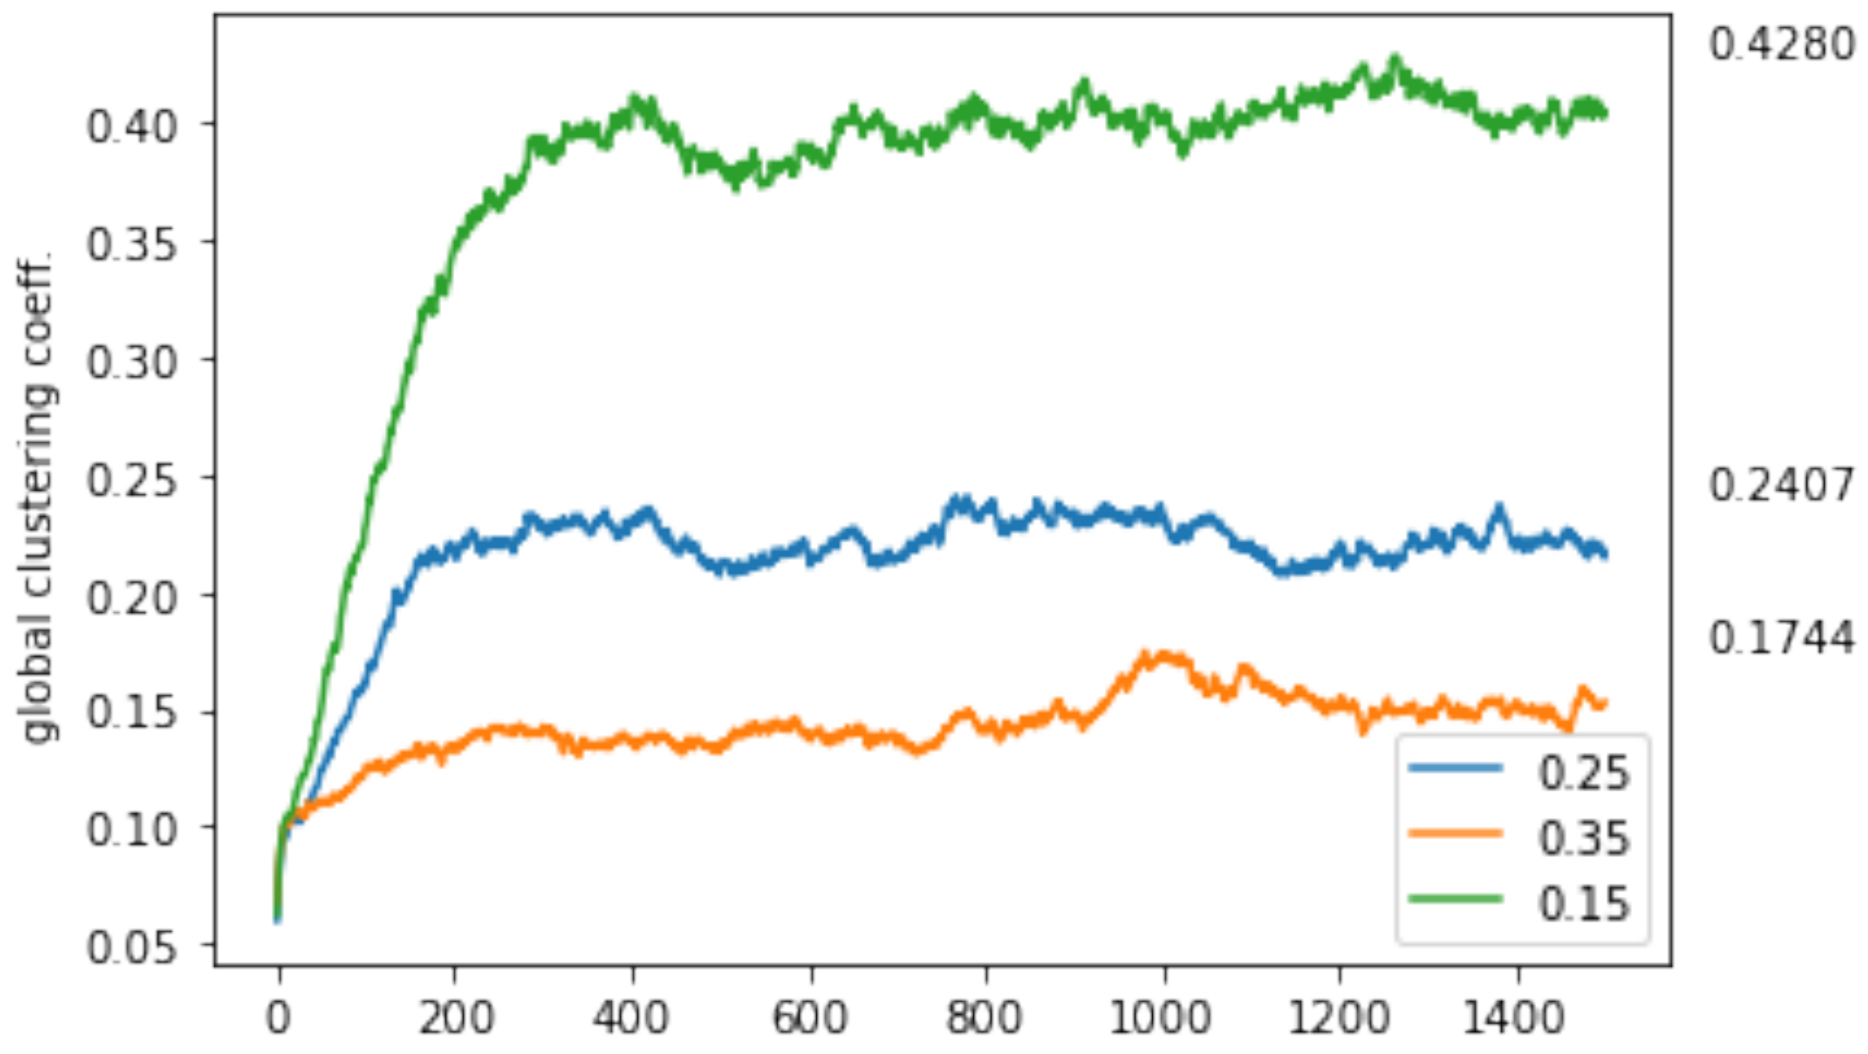

Global clustering coefficient over time for Barabási–Albert graphs

Verification through Graph Measures — Visually

- Taking a look at the *initial generated graphs,* compared against the Twitter graph.

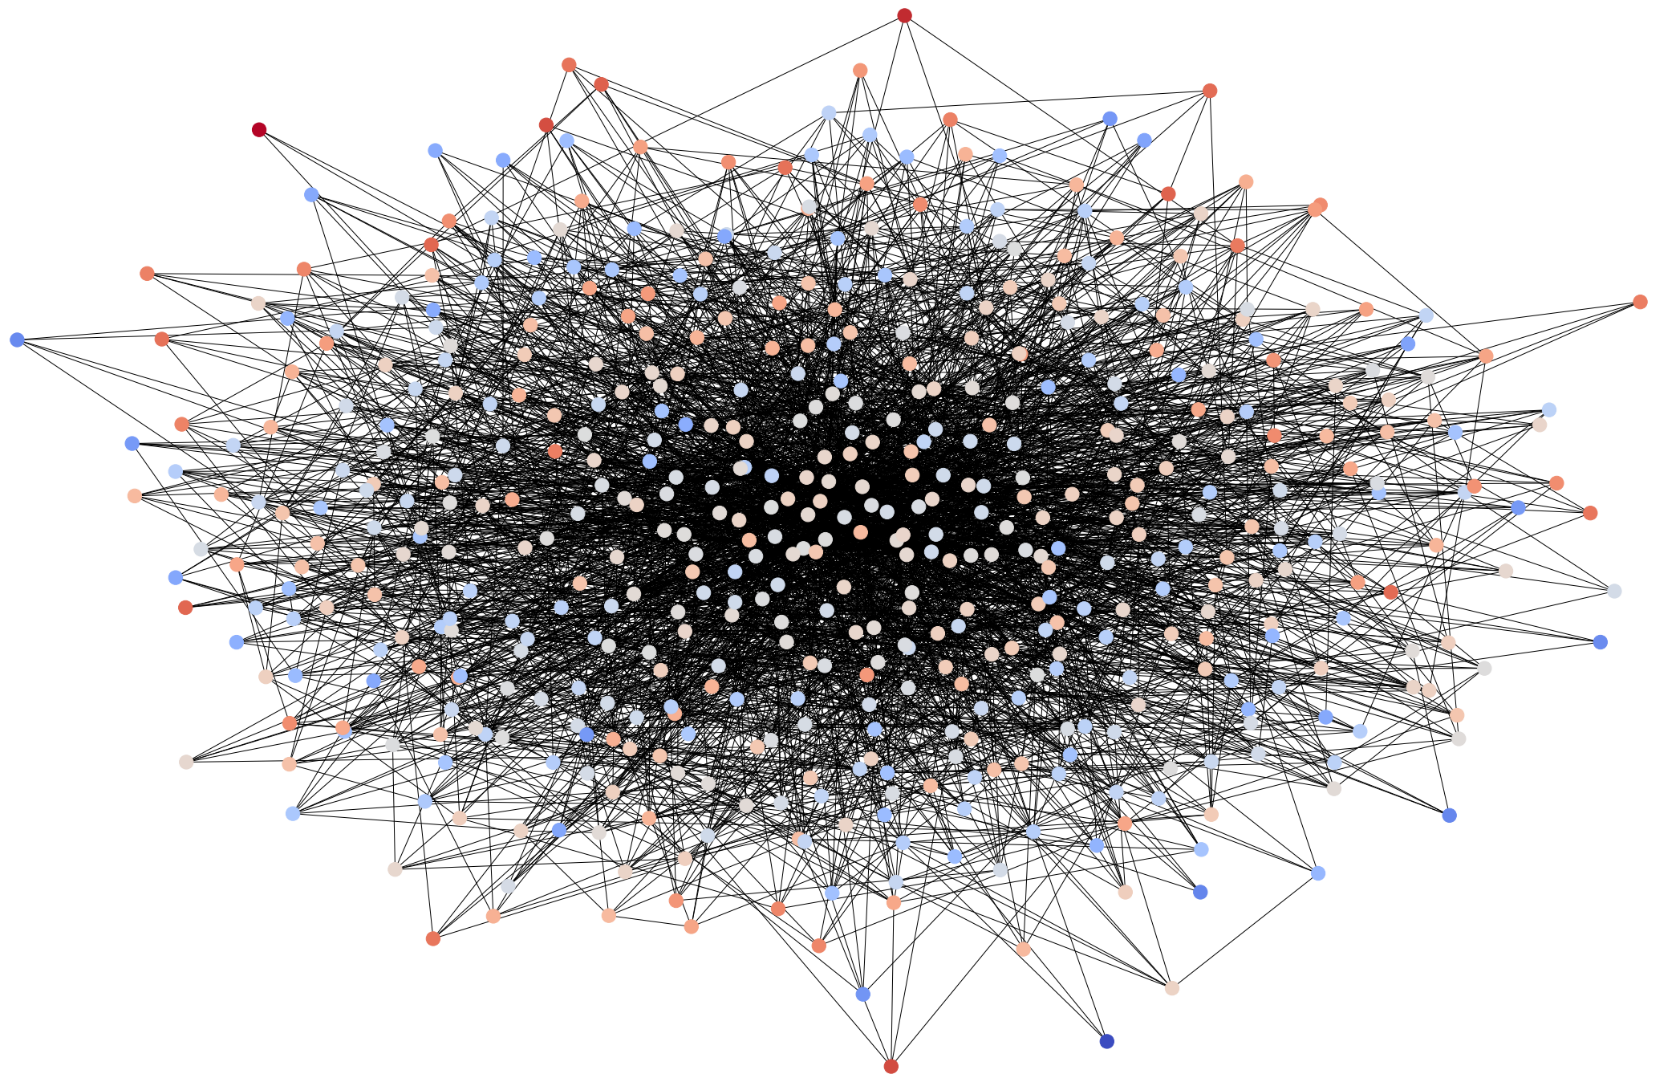

Barabási–Albert graph with 25% fixed edges, before edge dynamics

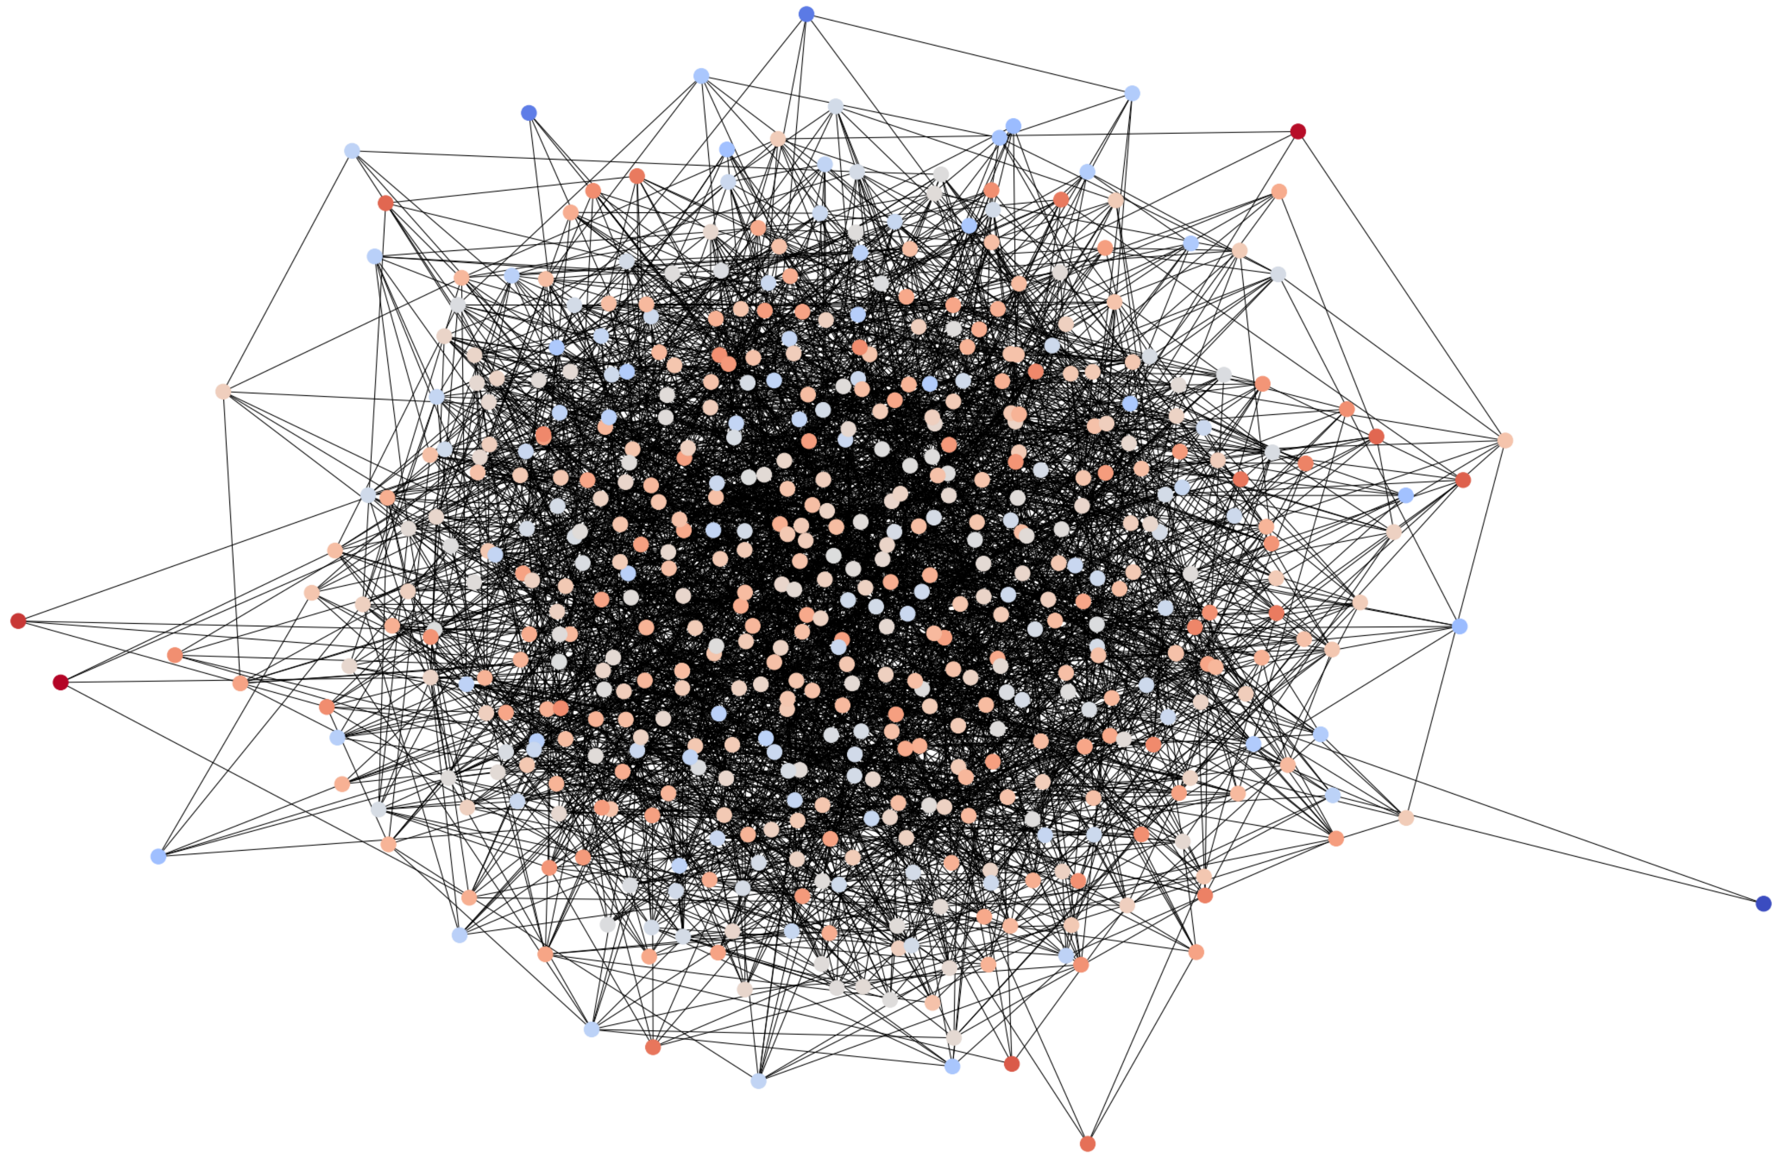

Erdős–Rényi graph with 25% fixed edges, before edge dynamics

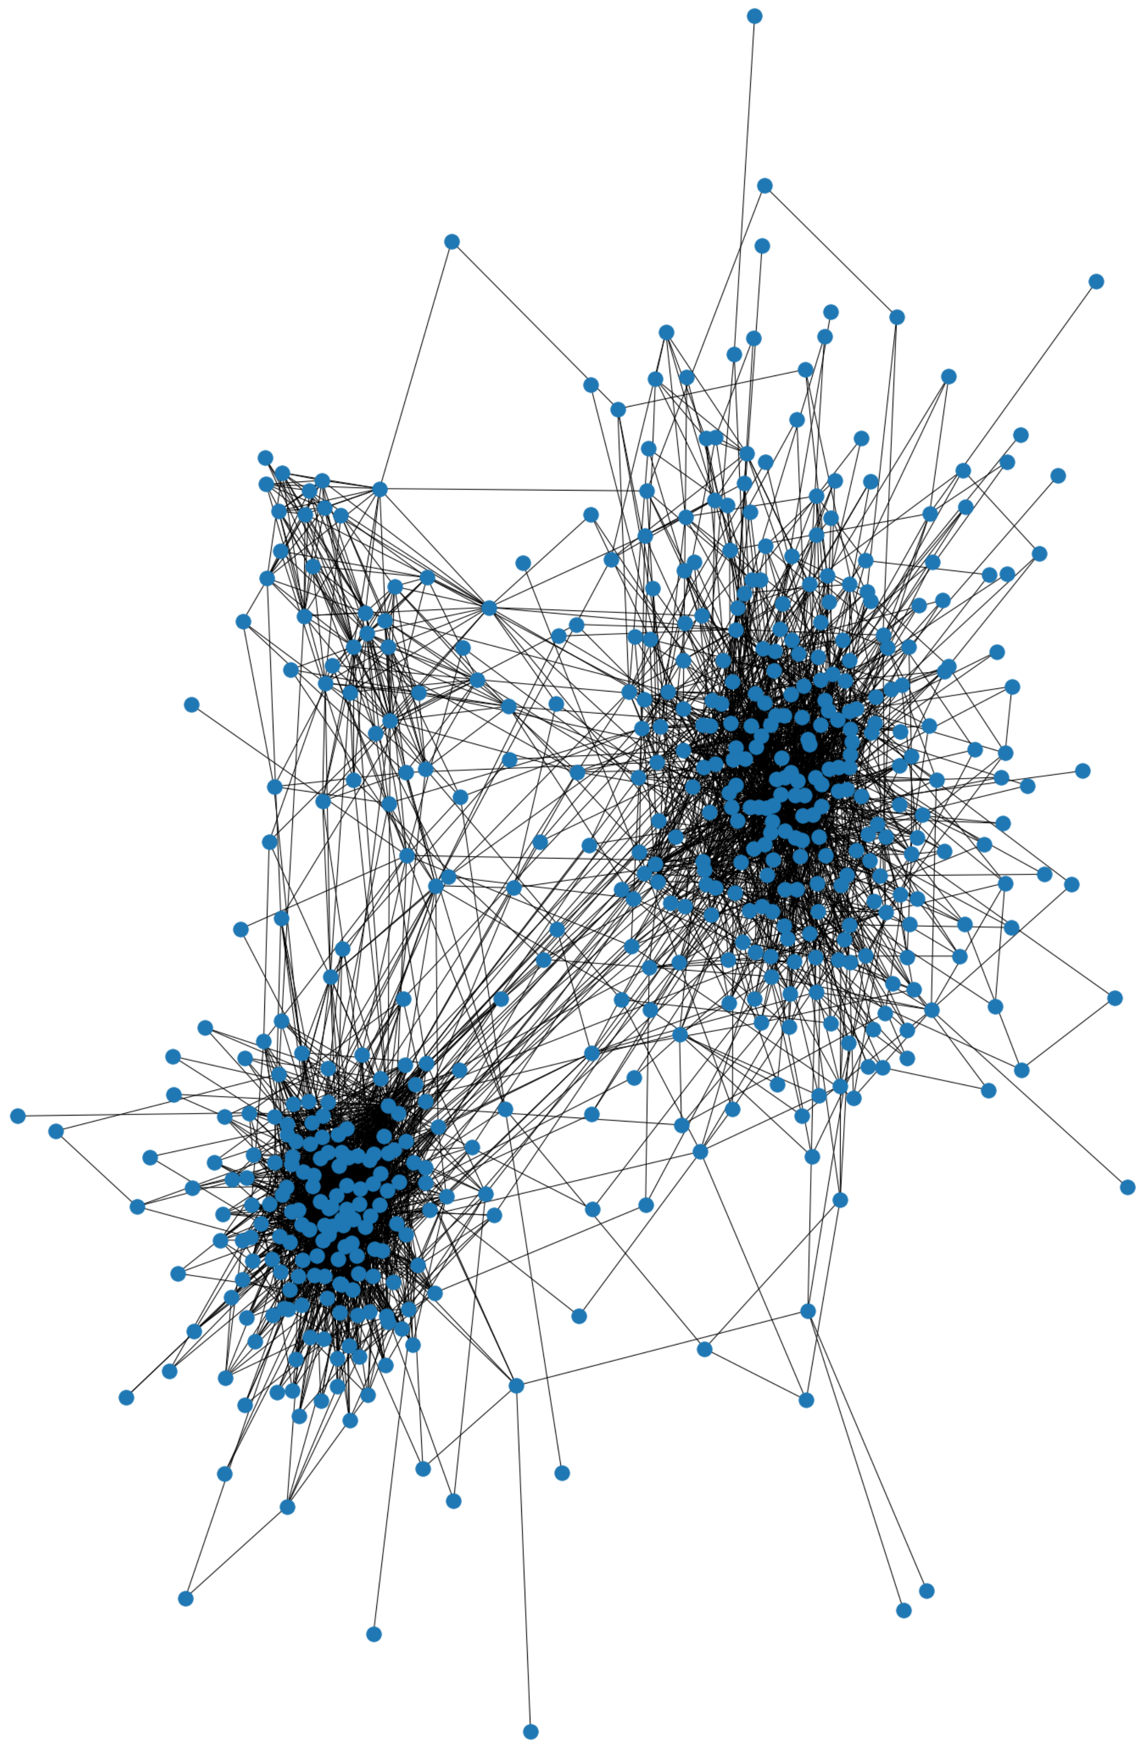

Twitter (Delhi 2013) snapshot visualized

Verification through Graph Measures — Visually

- Looking at the *simulated steady-state graphs*, compared against the Twitter graph.

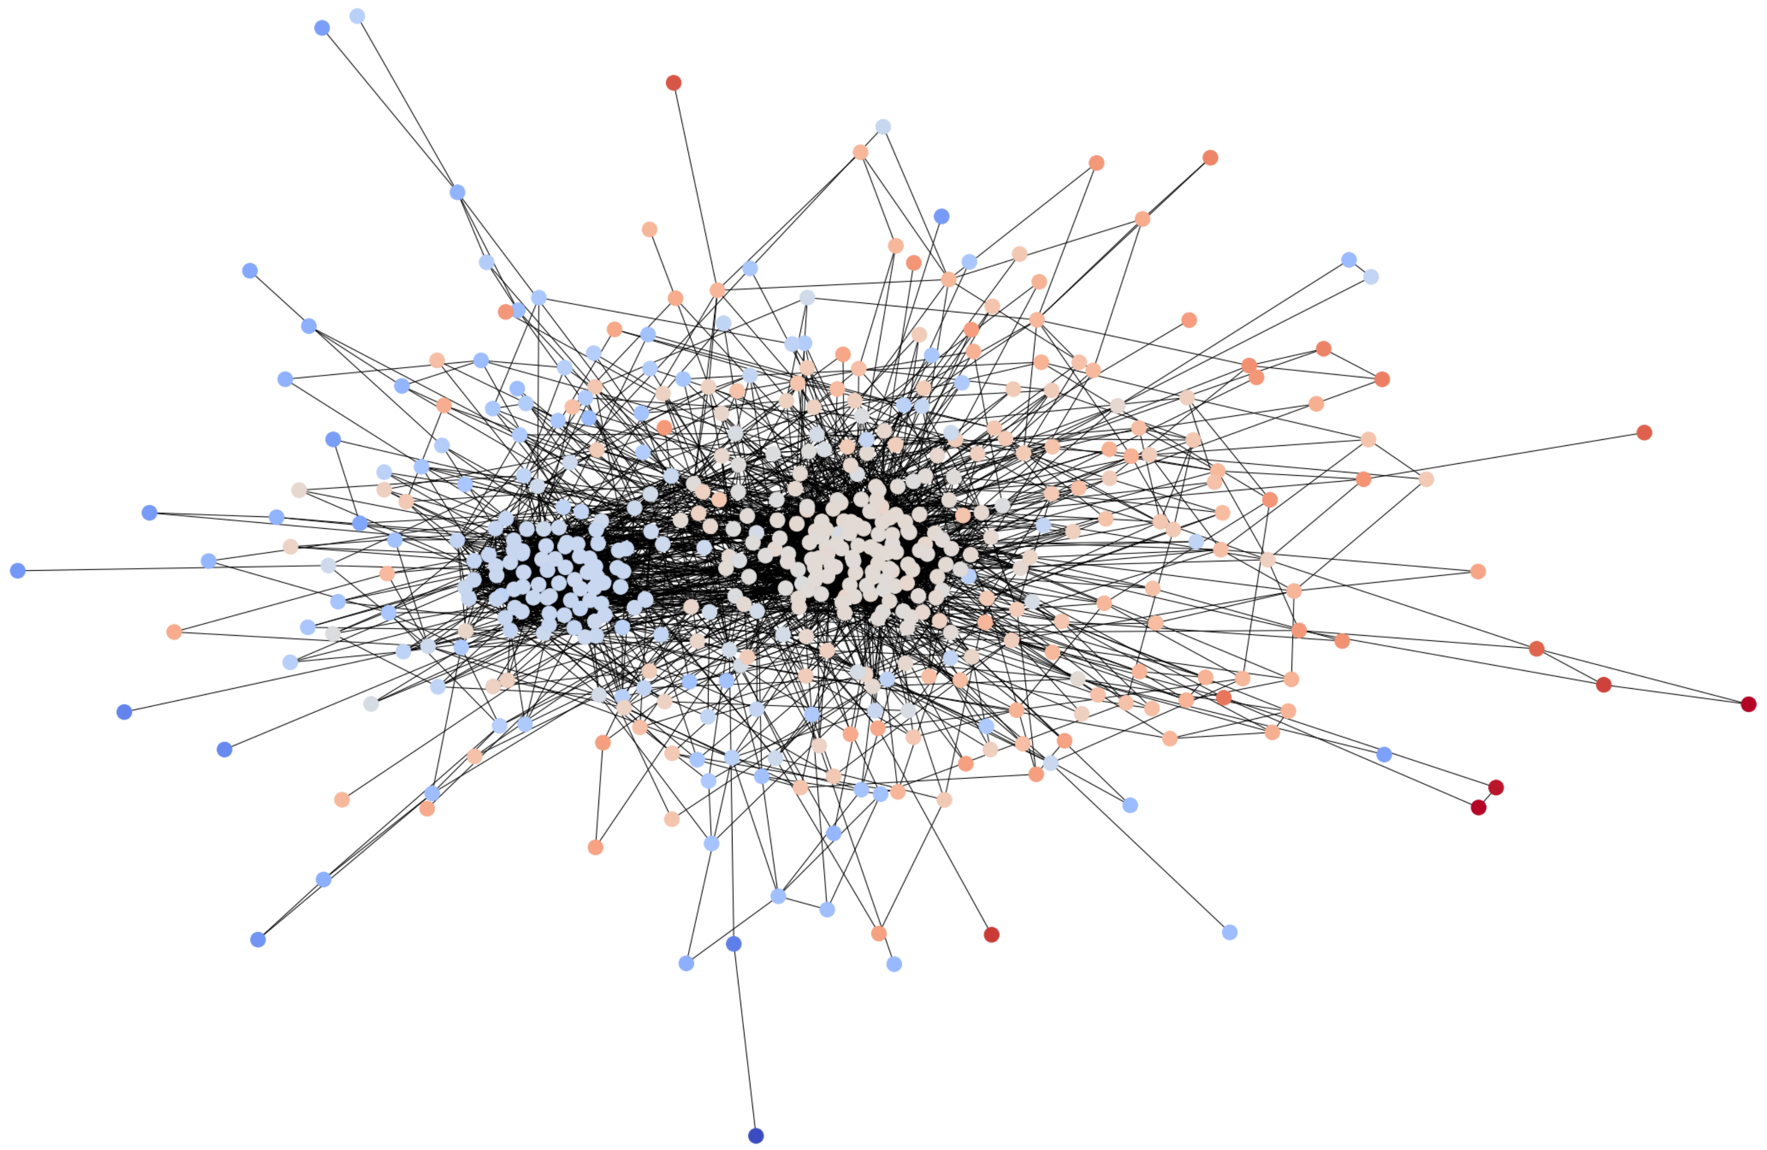

Barabási–Albert graph w/ 25% fixed edges, **AFTER** edge dynamics

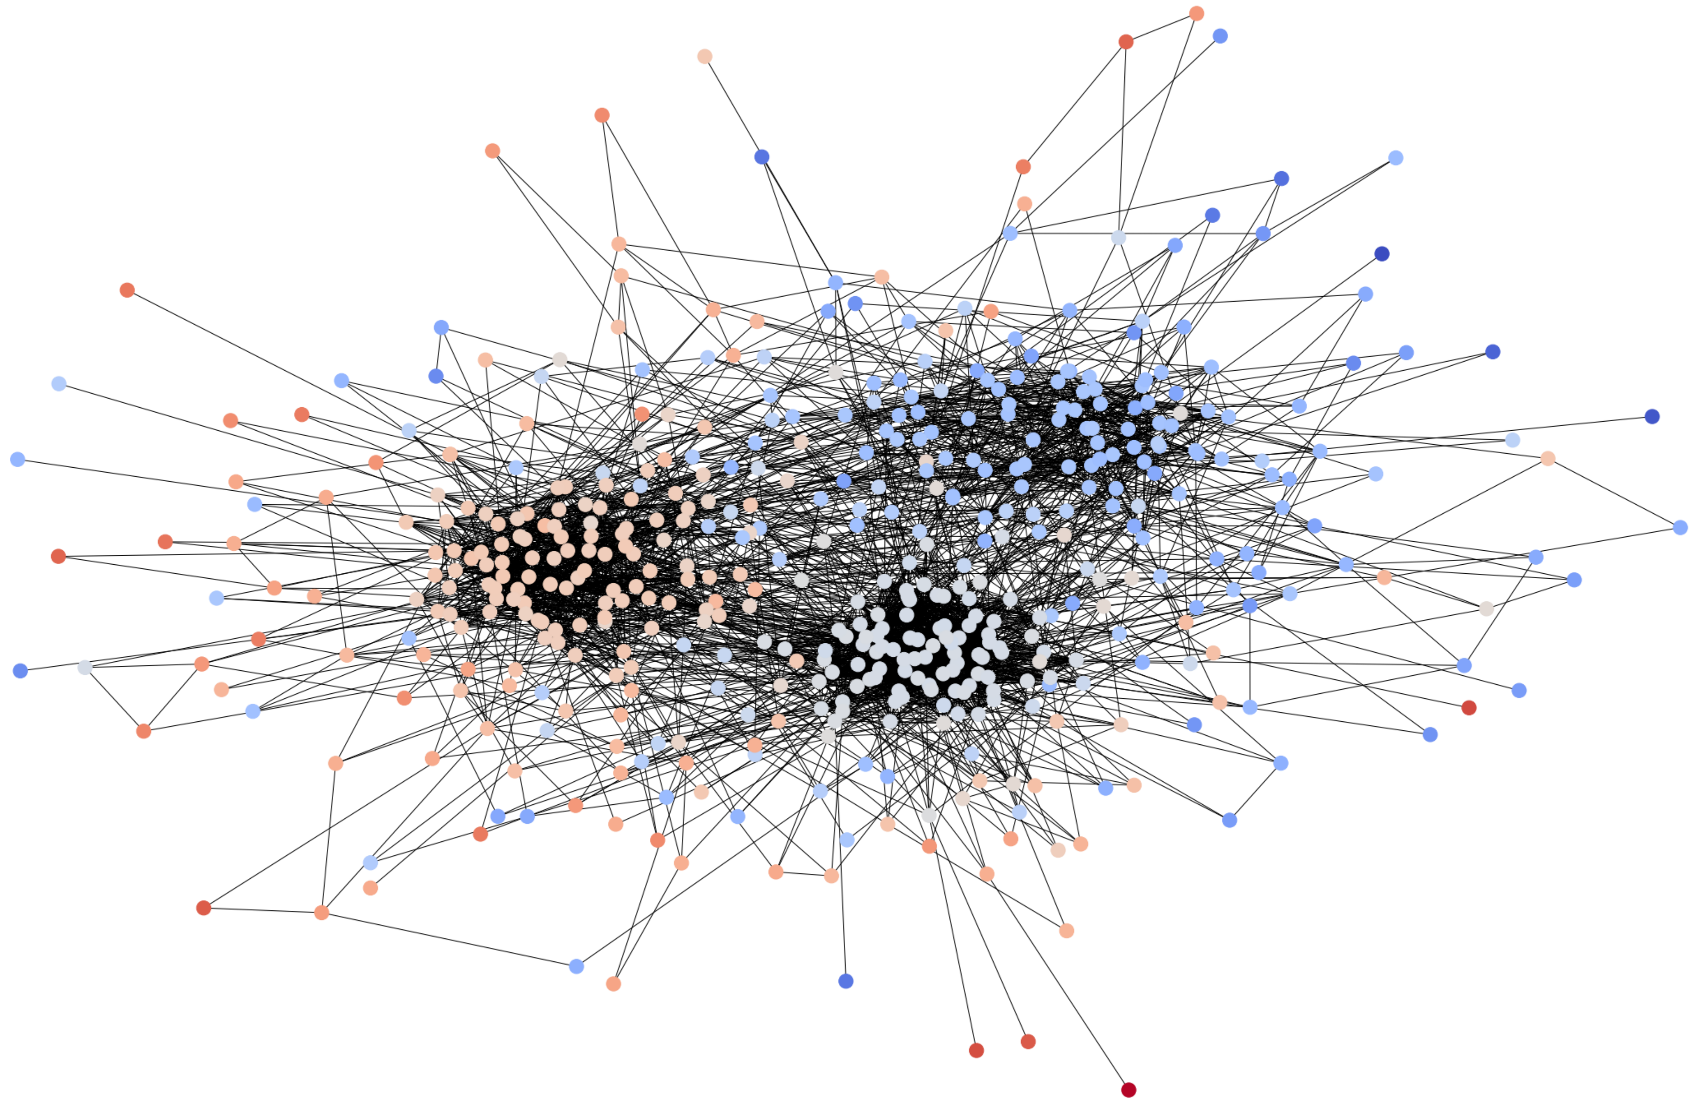

Erdős–Rényi graph w/ 25% fixed edges, **AFTER** edge dynamics

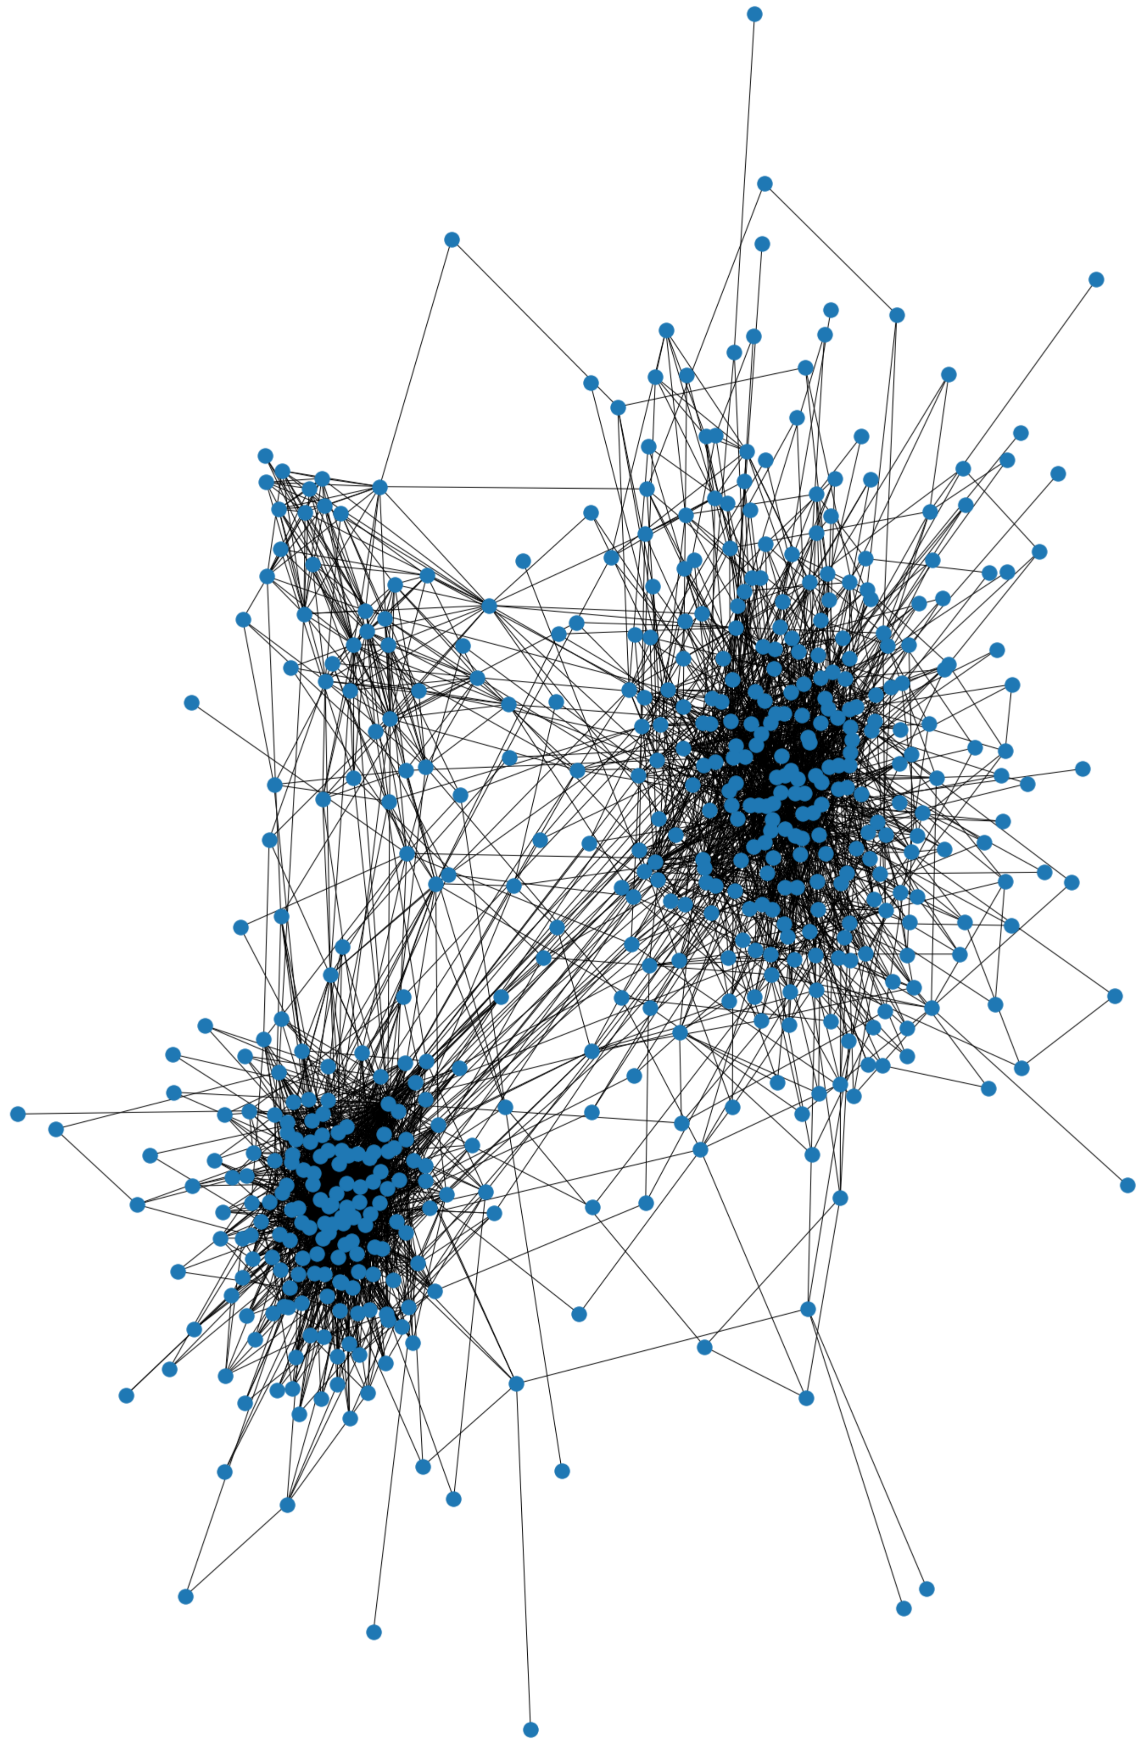

Twitter (Delhi 2013) snapshot visualized

Verification through Graph Measures — Degree Distribution

- Degree distribution of the *initial generated graphs,* compared against the Twitter graph.
- Known that real-world networks exhibit this type of power law degree dist.   [Muchnik et al. '13]

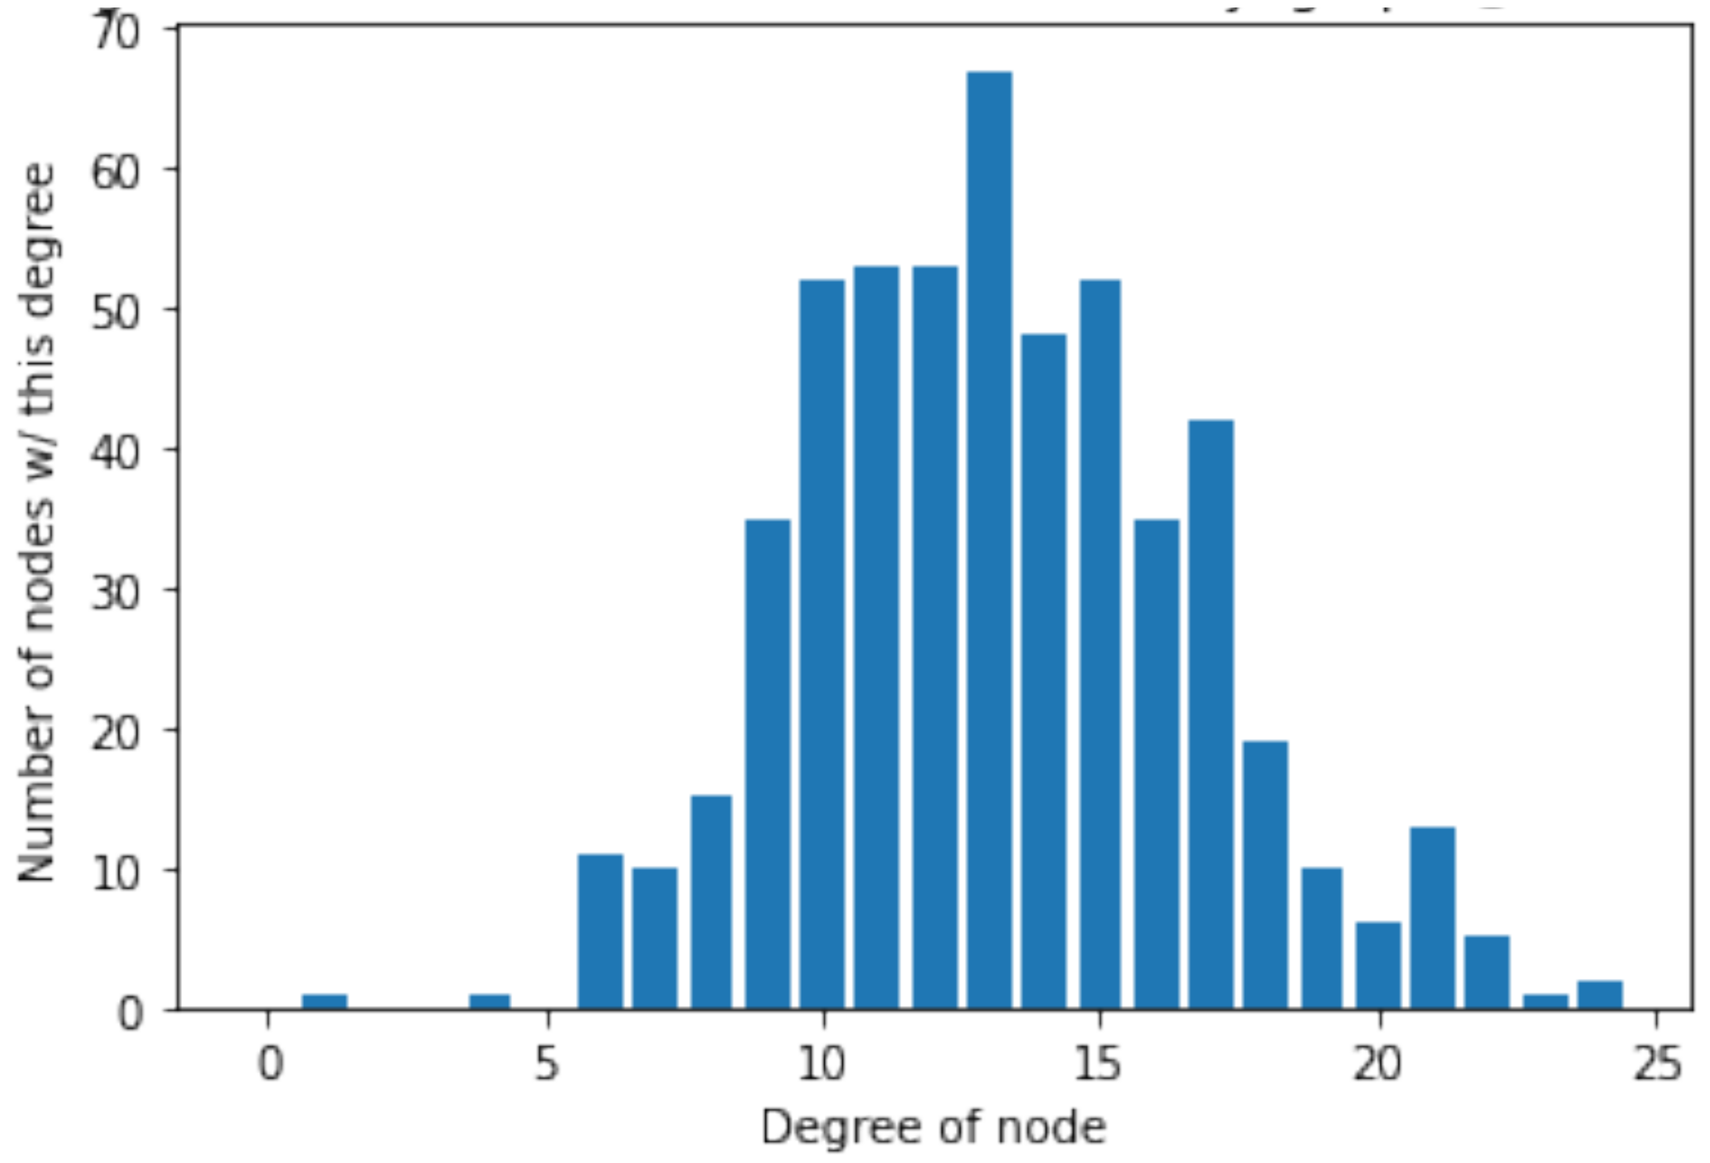

Erdős–Rényi graph with 25% fixed edges, degree distribution before edge dynamics

Twitter (Delhi 2013) degree distribution

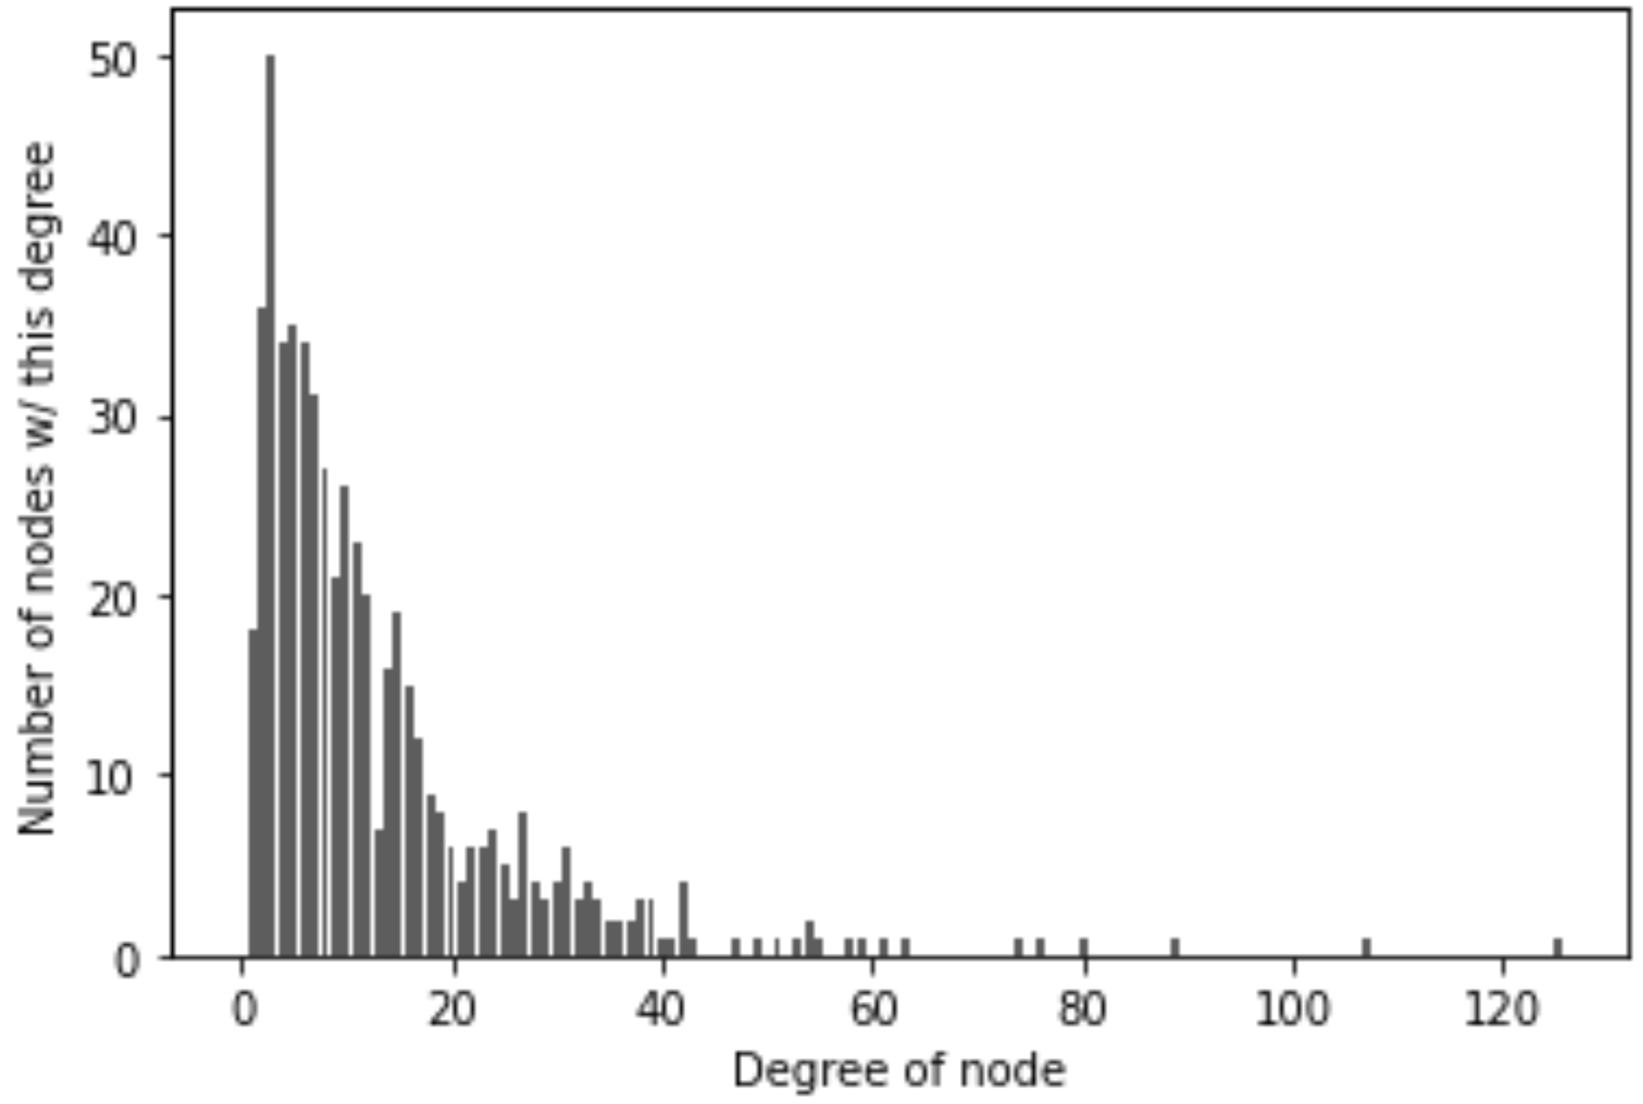

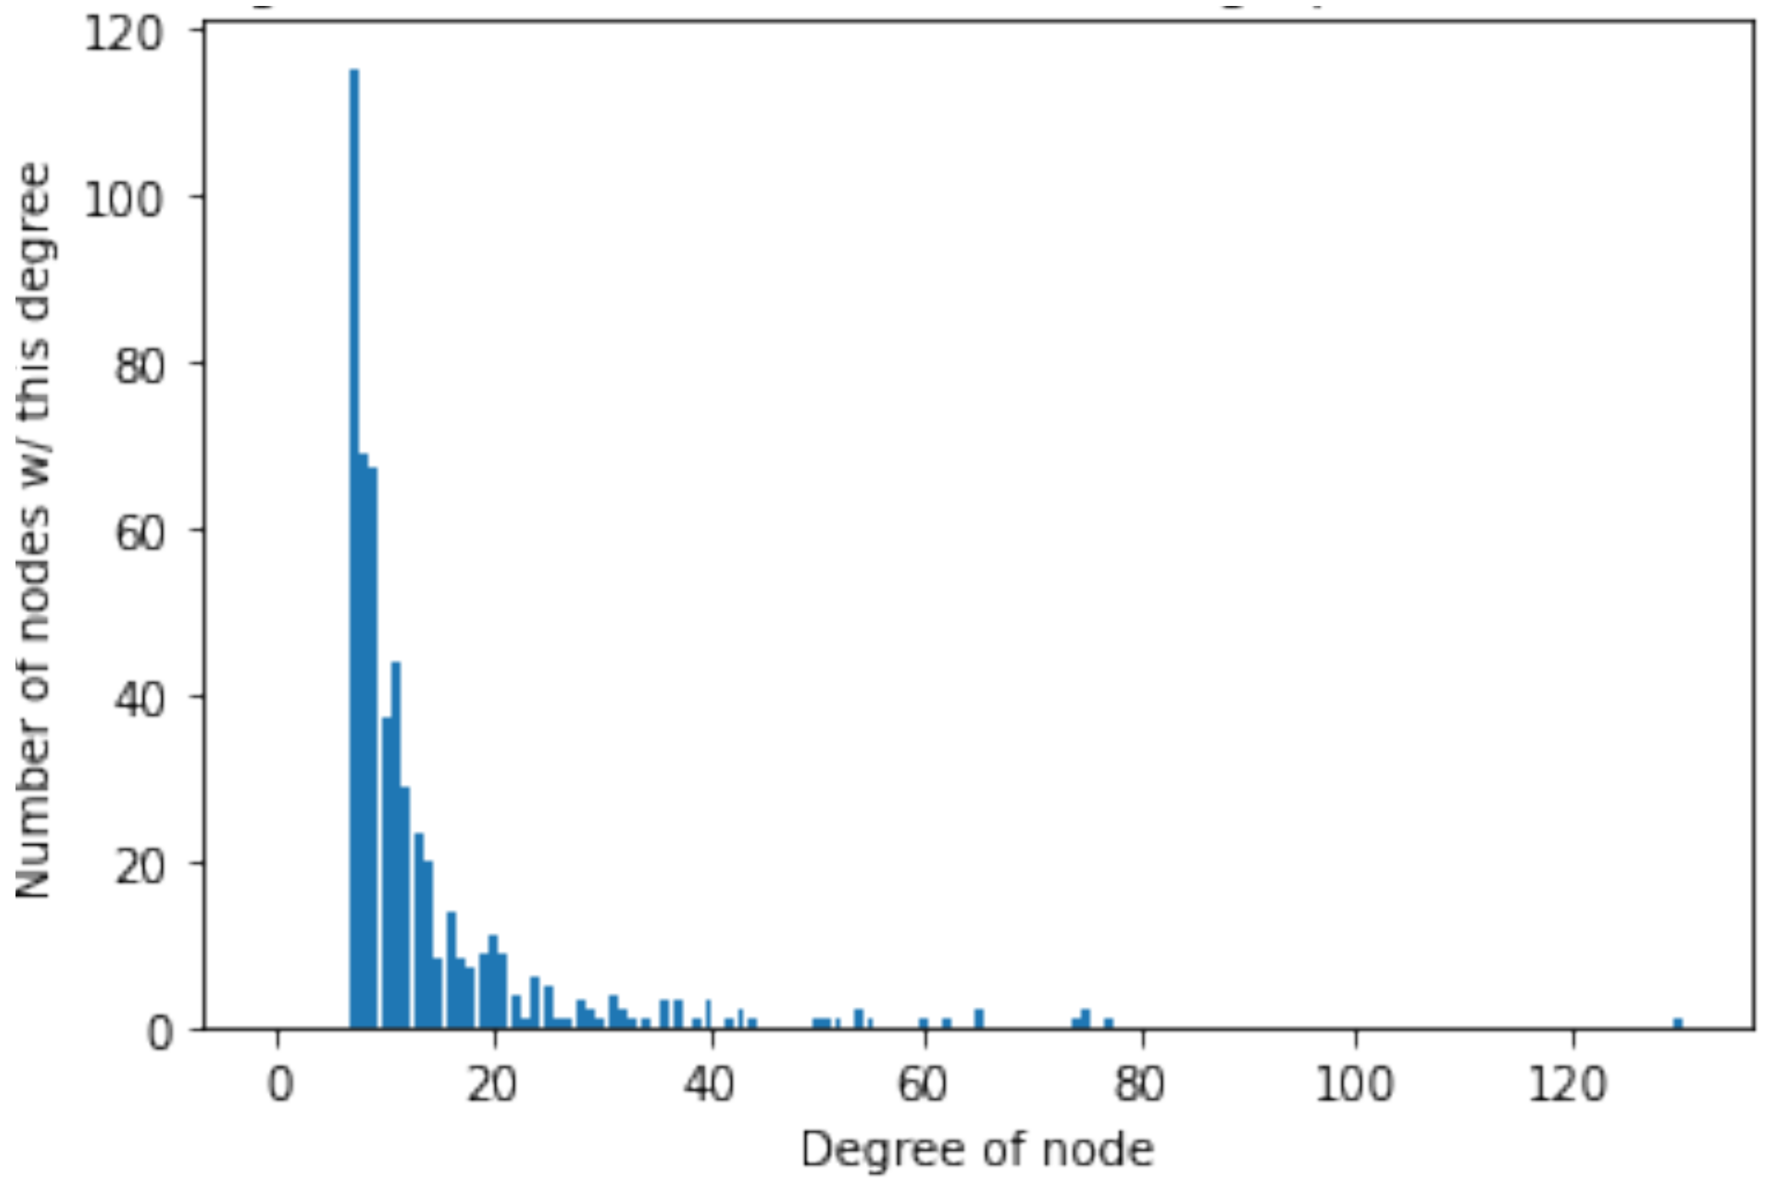

Barabási–Albert graph with 25% fixed edges, degree distribution before edge dynamics

Verification through Graph Measures — Degree Distribution

- Degree distribution of the *simulated steady-state graphs,* compared against the Twitter graph.
- Interesting result — prior work [Sasahara et al. '20] has not been able to *alter* the degree distribution of a network through a synthetic model.

Twitter (Delhi 2013) degree distribution

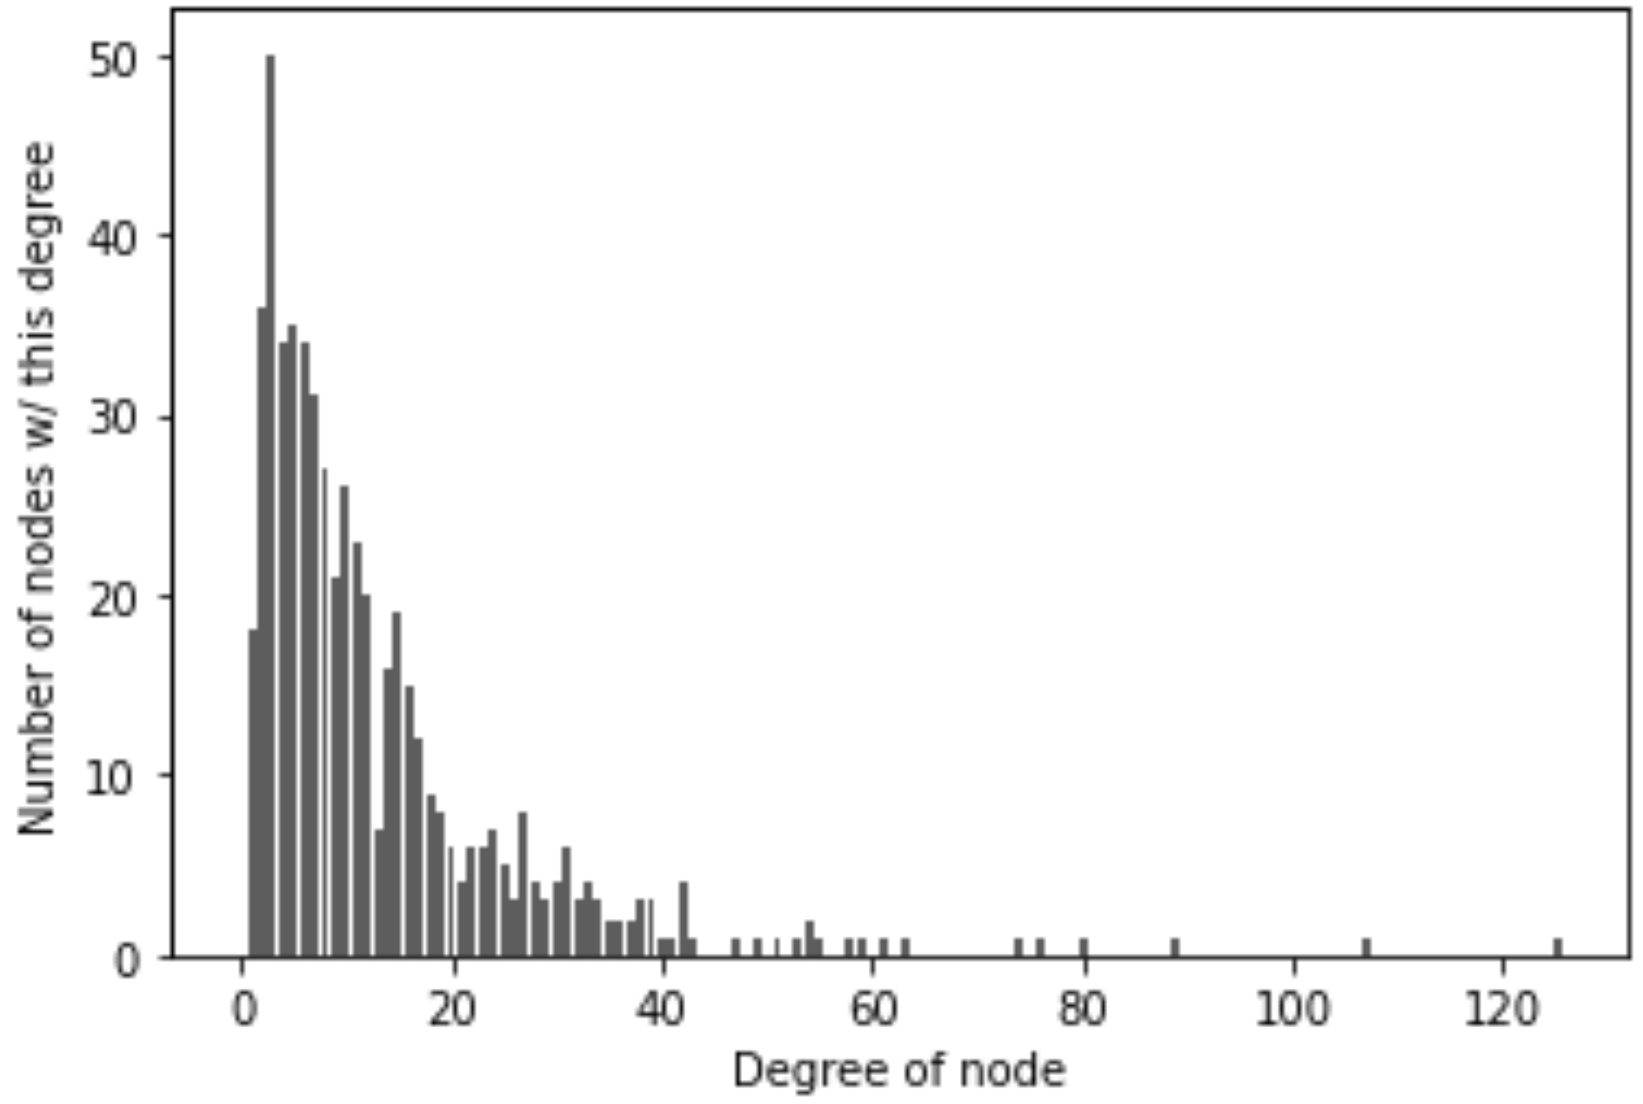

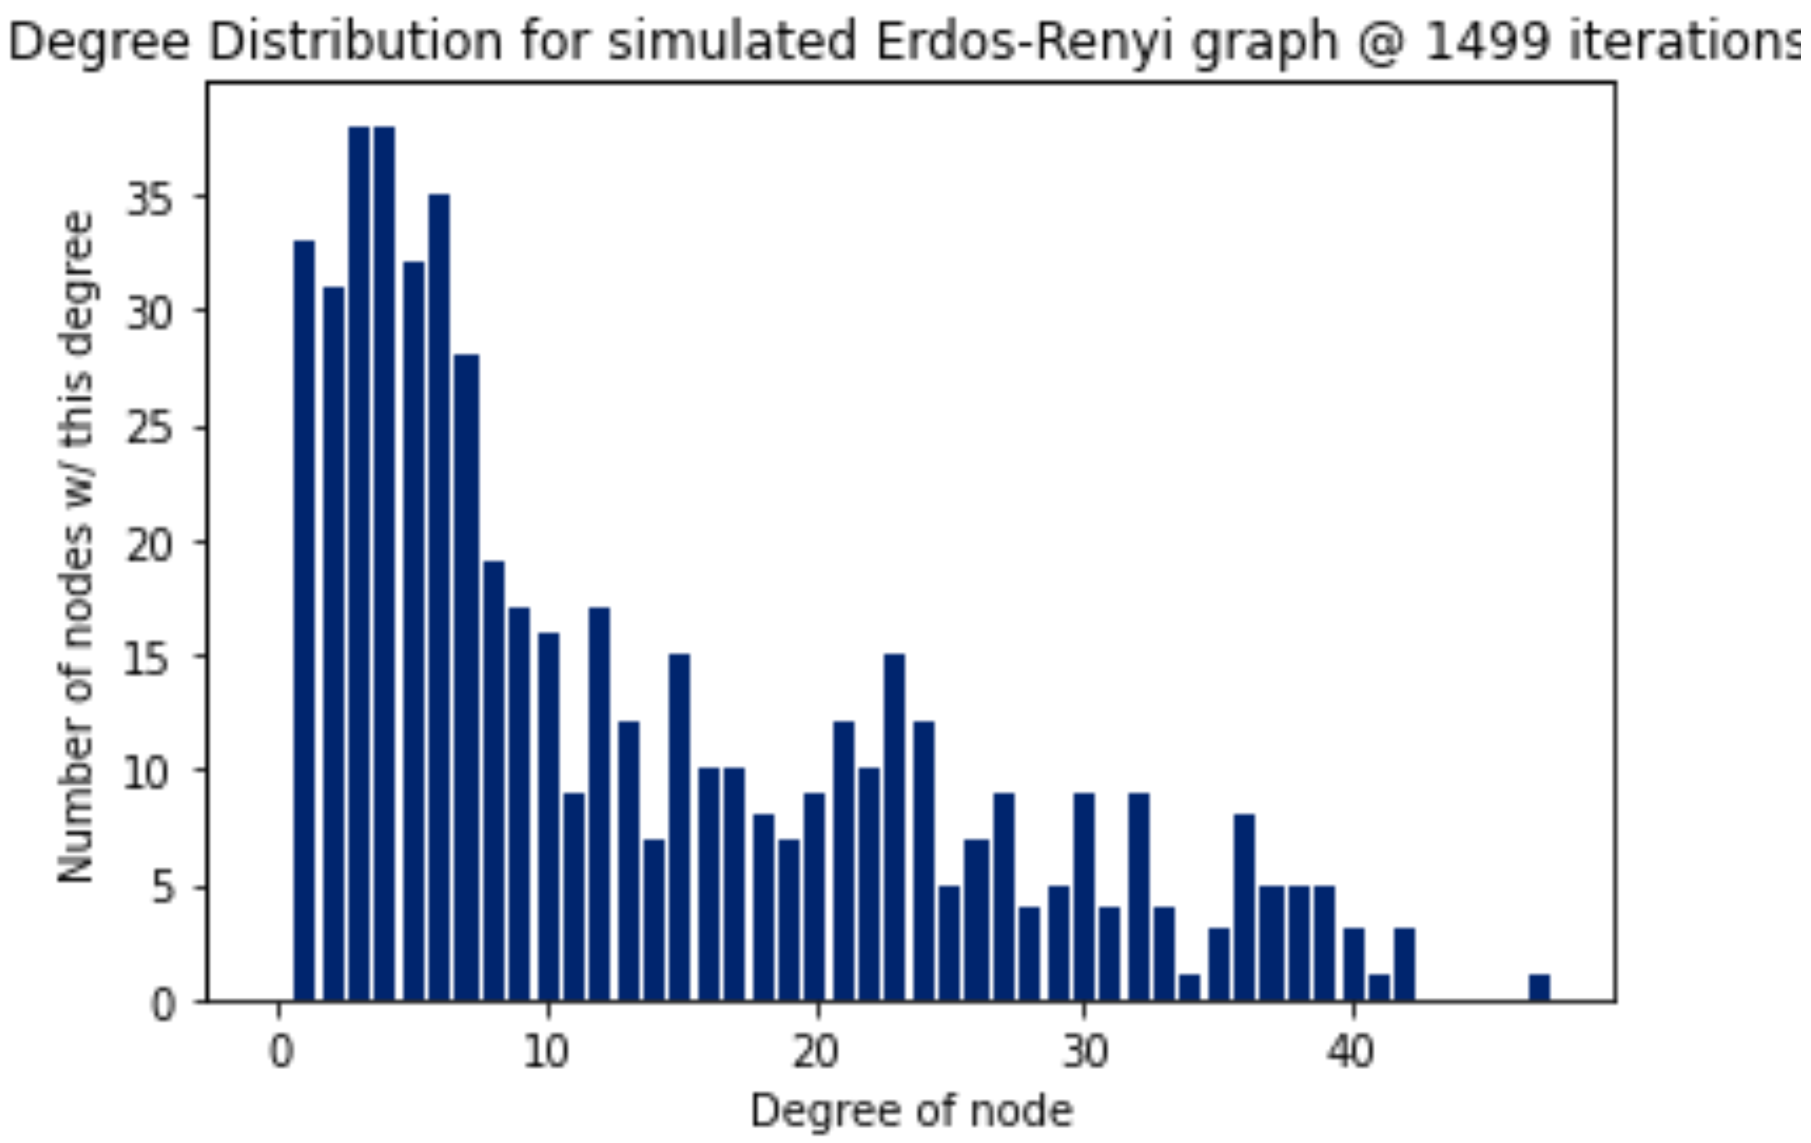

Erdős–Rényi graph with 25% fixed edges, degree distribution **AFTER** edge dynamics

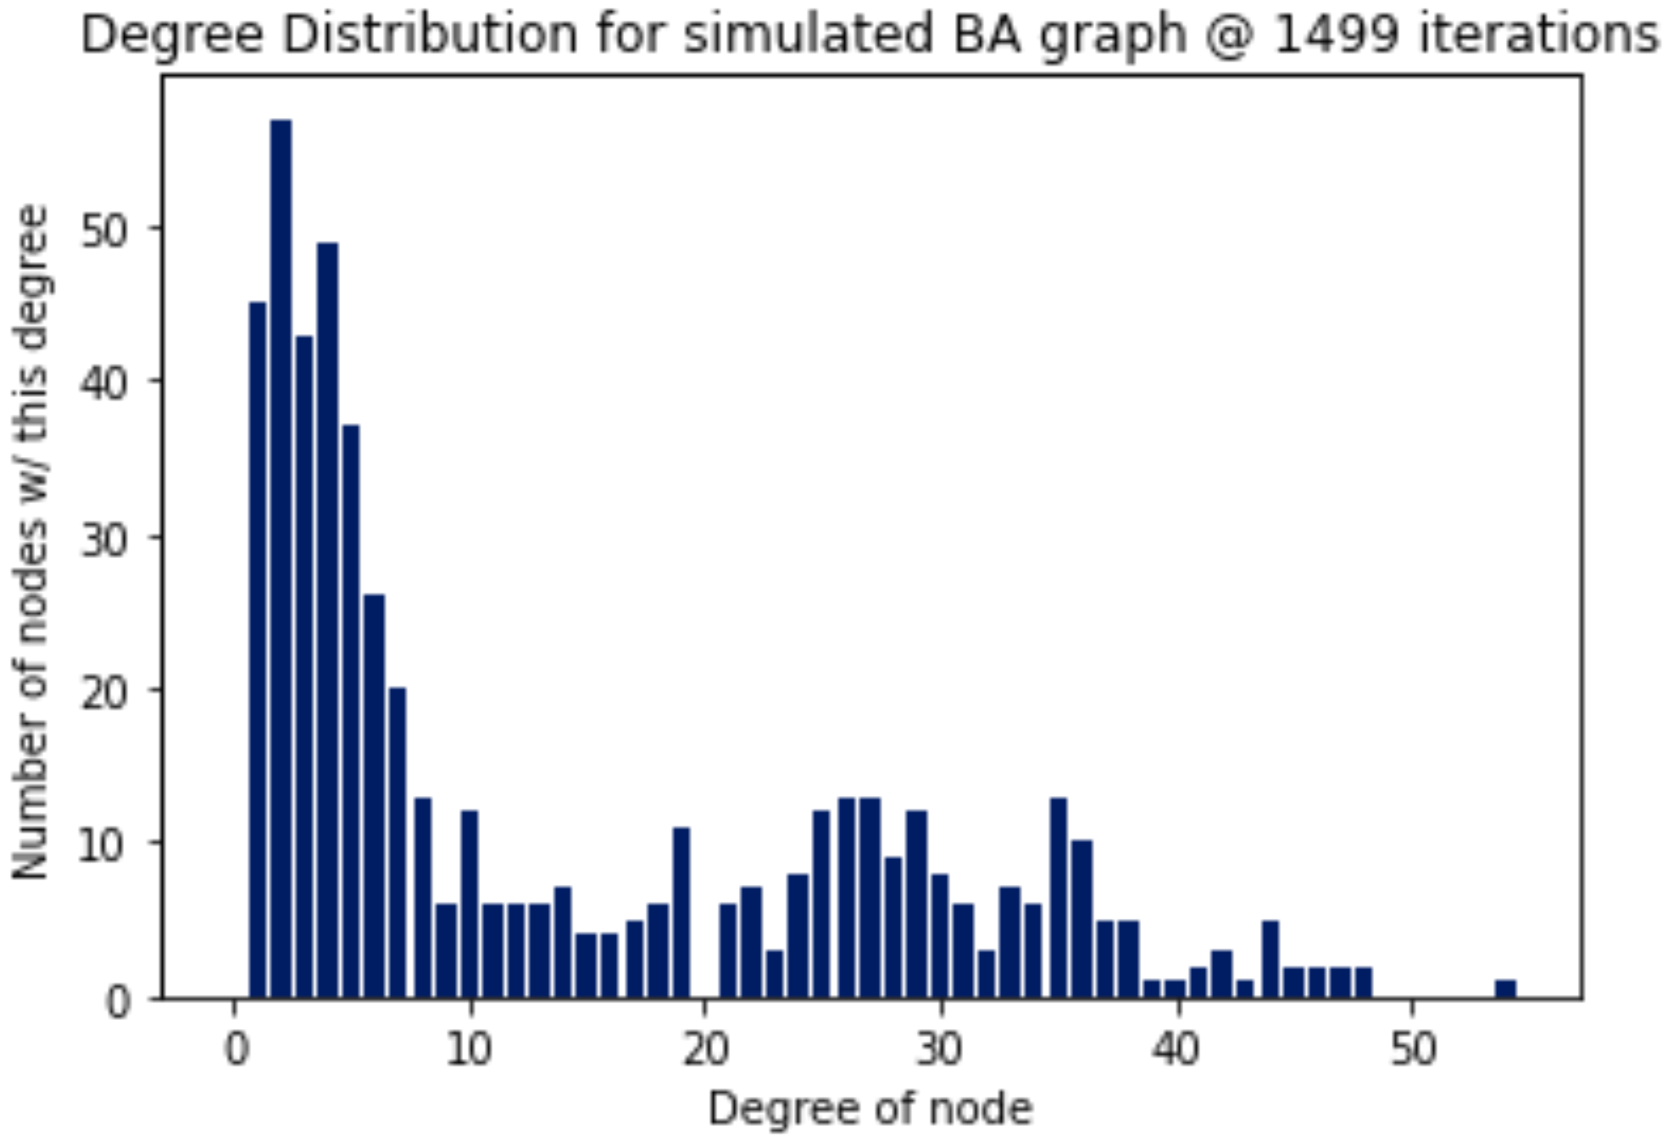

Barabási–Albert graph with 25% fixed edges, degree distribution **AFTER** edge dynamics

Real-World Verification through Temporal Data

- In this section, we study a data set containing U.S. Congress co-sponsorship records over time. [Benson et al. '18, Fowler '06]
  - $n = 1{,}718$ nodes, $260{,}851$ temporal co-sponsorship records
- Each record represents a bill and a list of people (nodes) who co-sponsored it.
- We also recovered partisan IDs for each node, so we have "innate opinions".
- **High-Level Intuition:** We iterate through the records, sampling from each record to *add timestamped edges.*
  - With each added edge, we delete the *oldest* edge in the graph to remove the influence of connections that have no recent activity.
  - Initially, edge deletion is disabled, so we can build up density in the graph.

Temporal Verification — Clustering & Triangles

- Both the global clustering coefficient and number of triangles rising over time.
- Vertical lines indicate the time step when edge deletion is turned on and experiment starts.

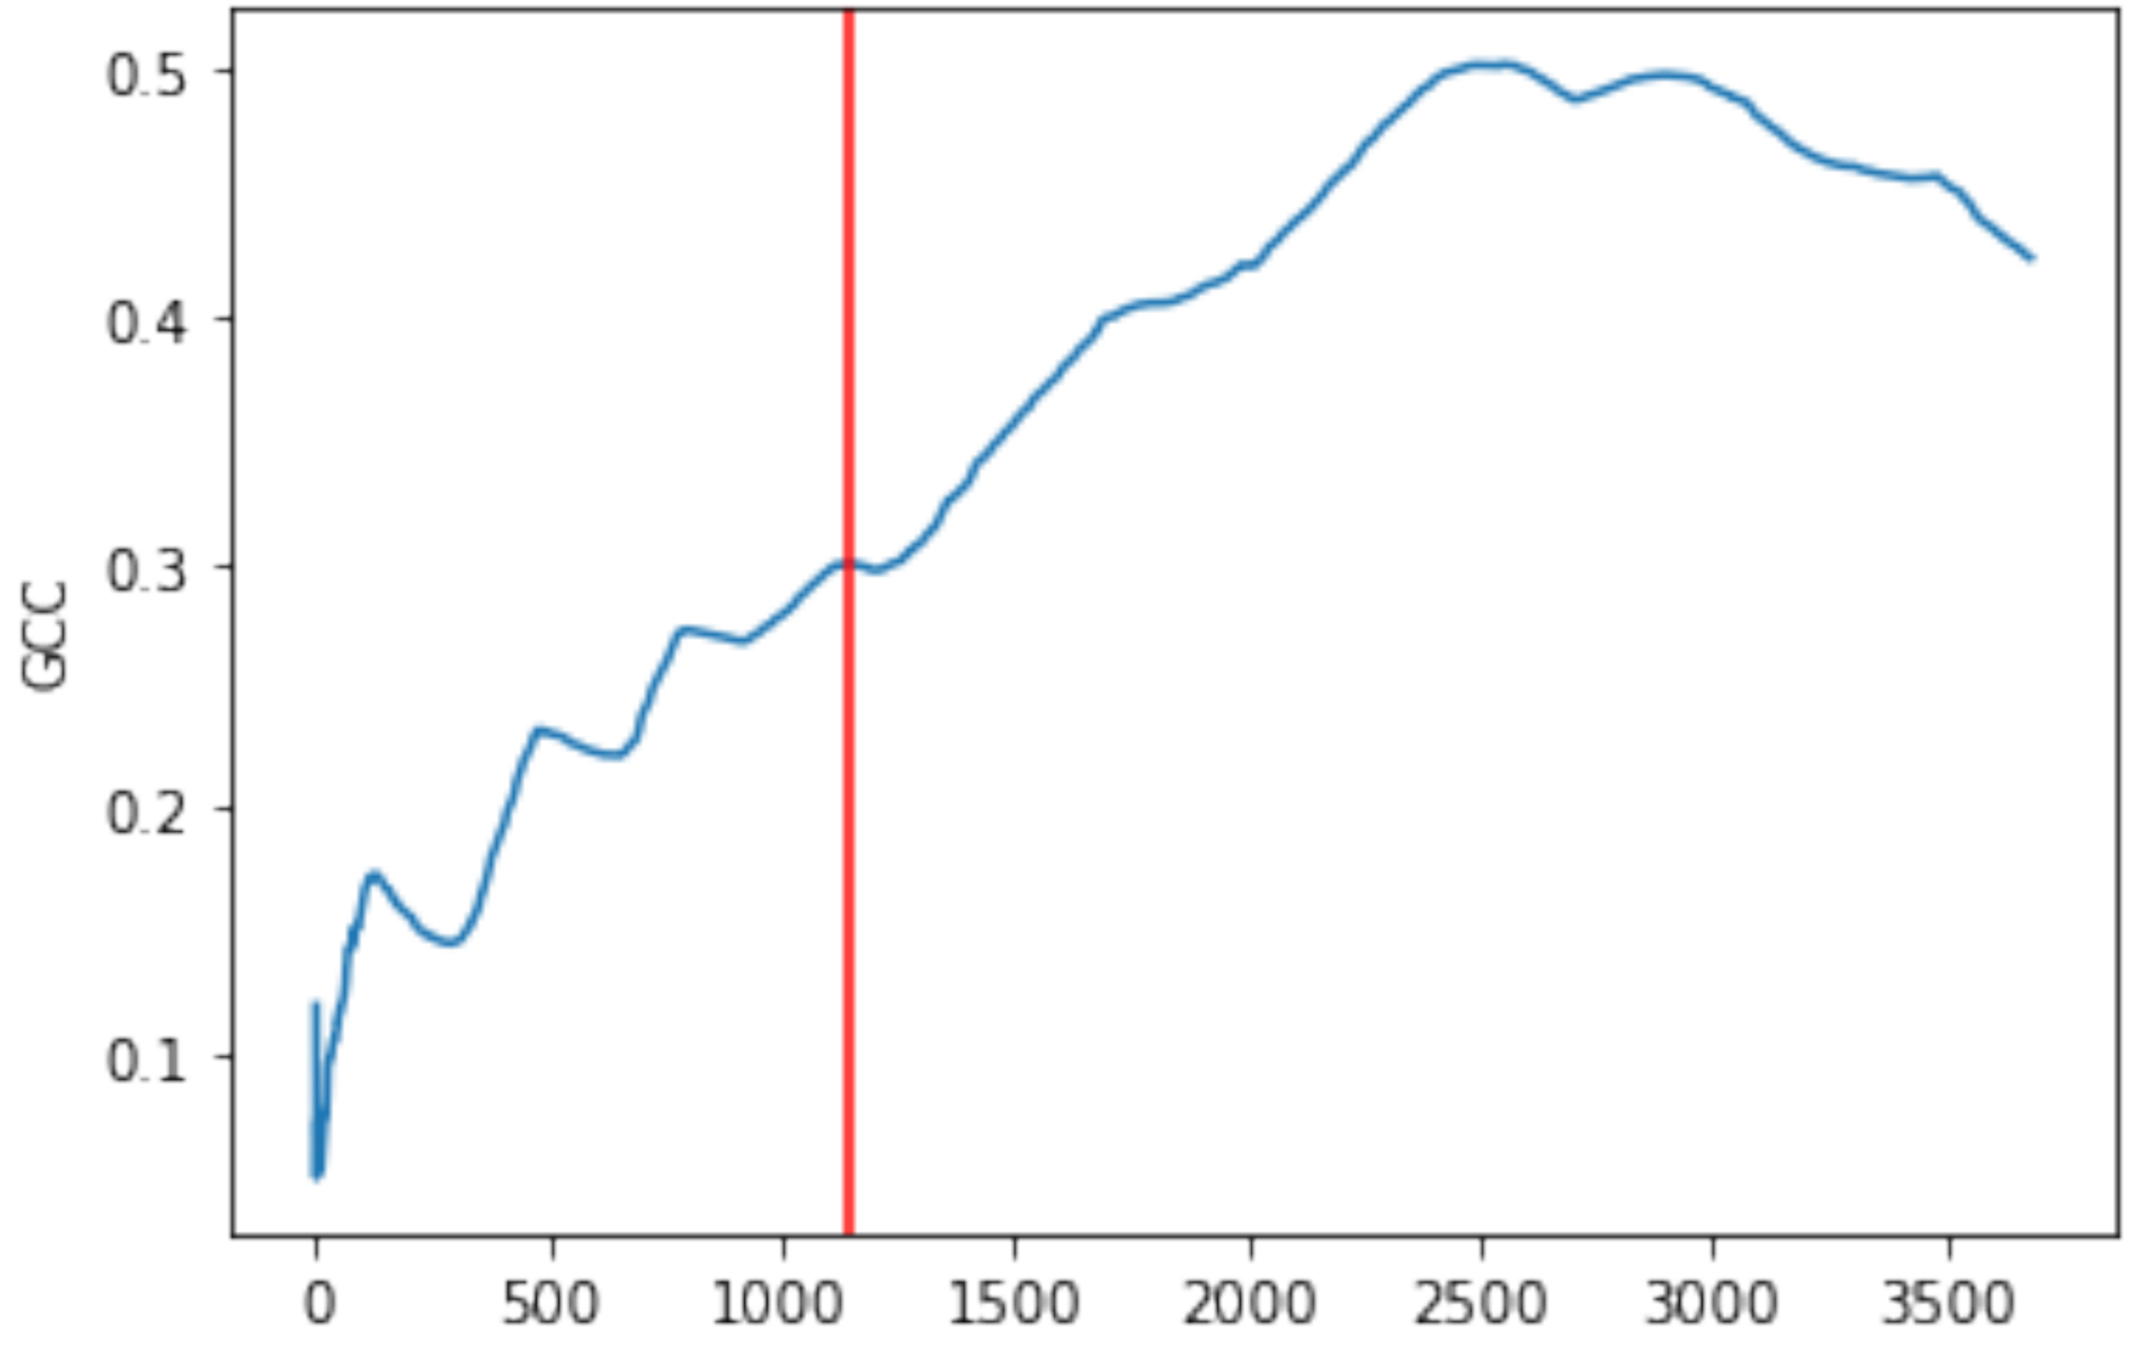

Global clustering coefficient over time for temporal data set. Vertical line indicates timestep where density threshold was met.

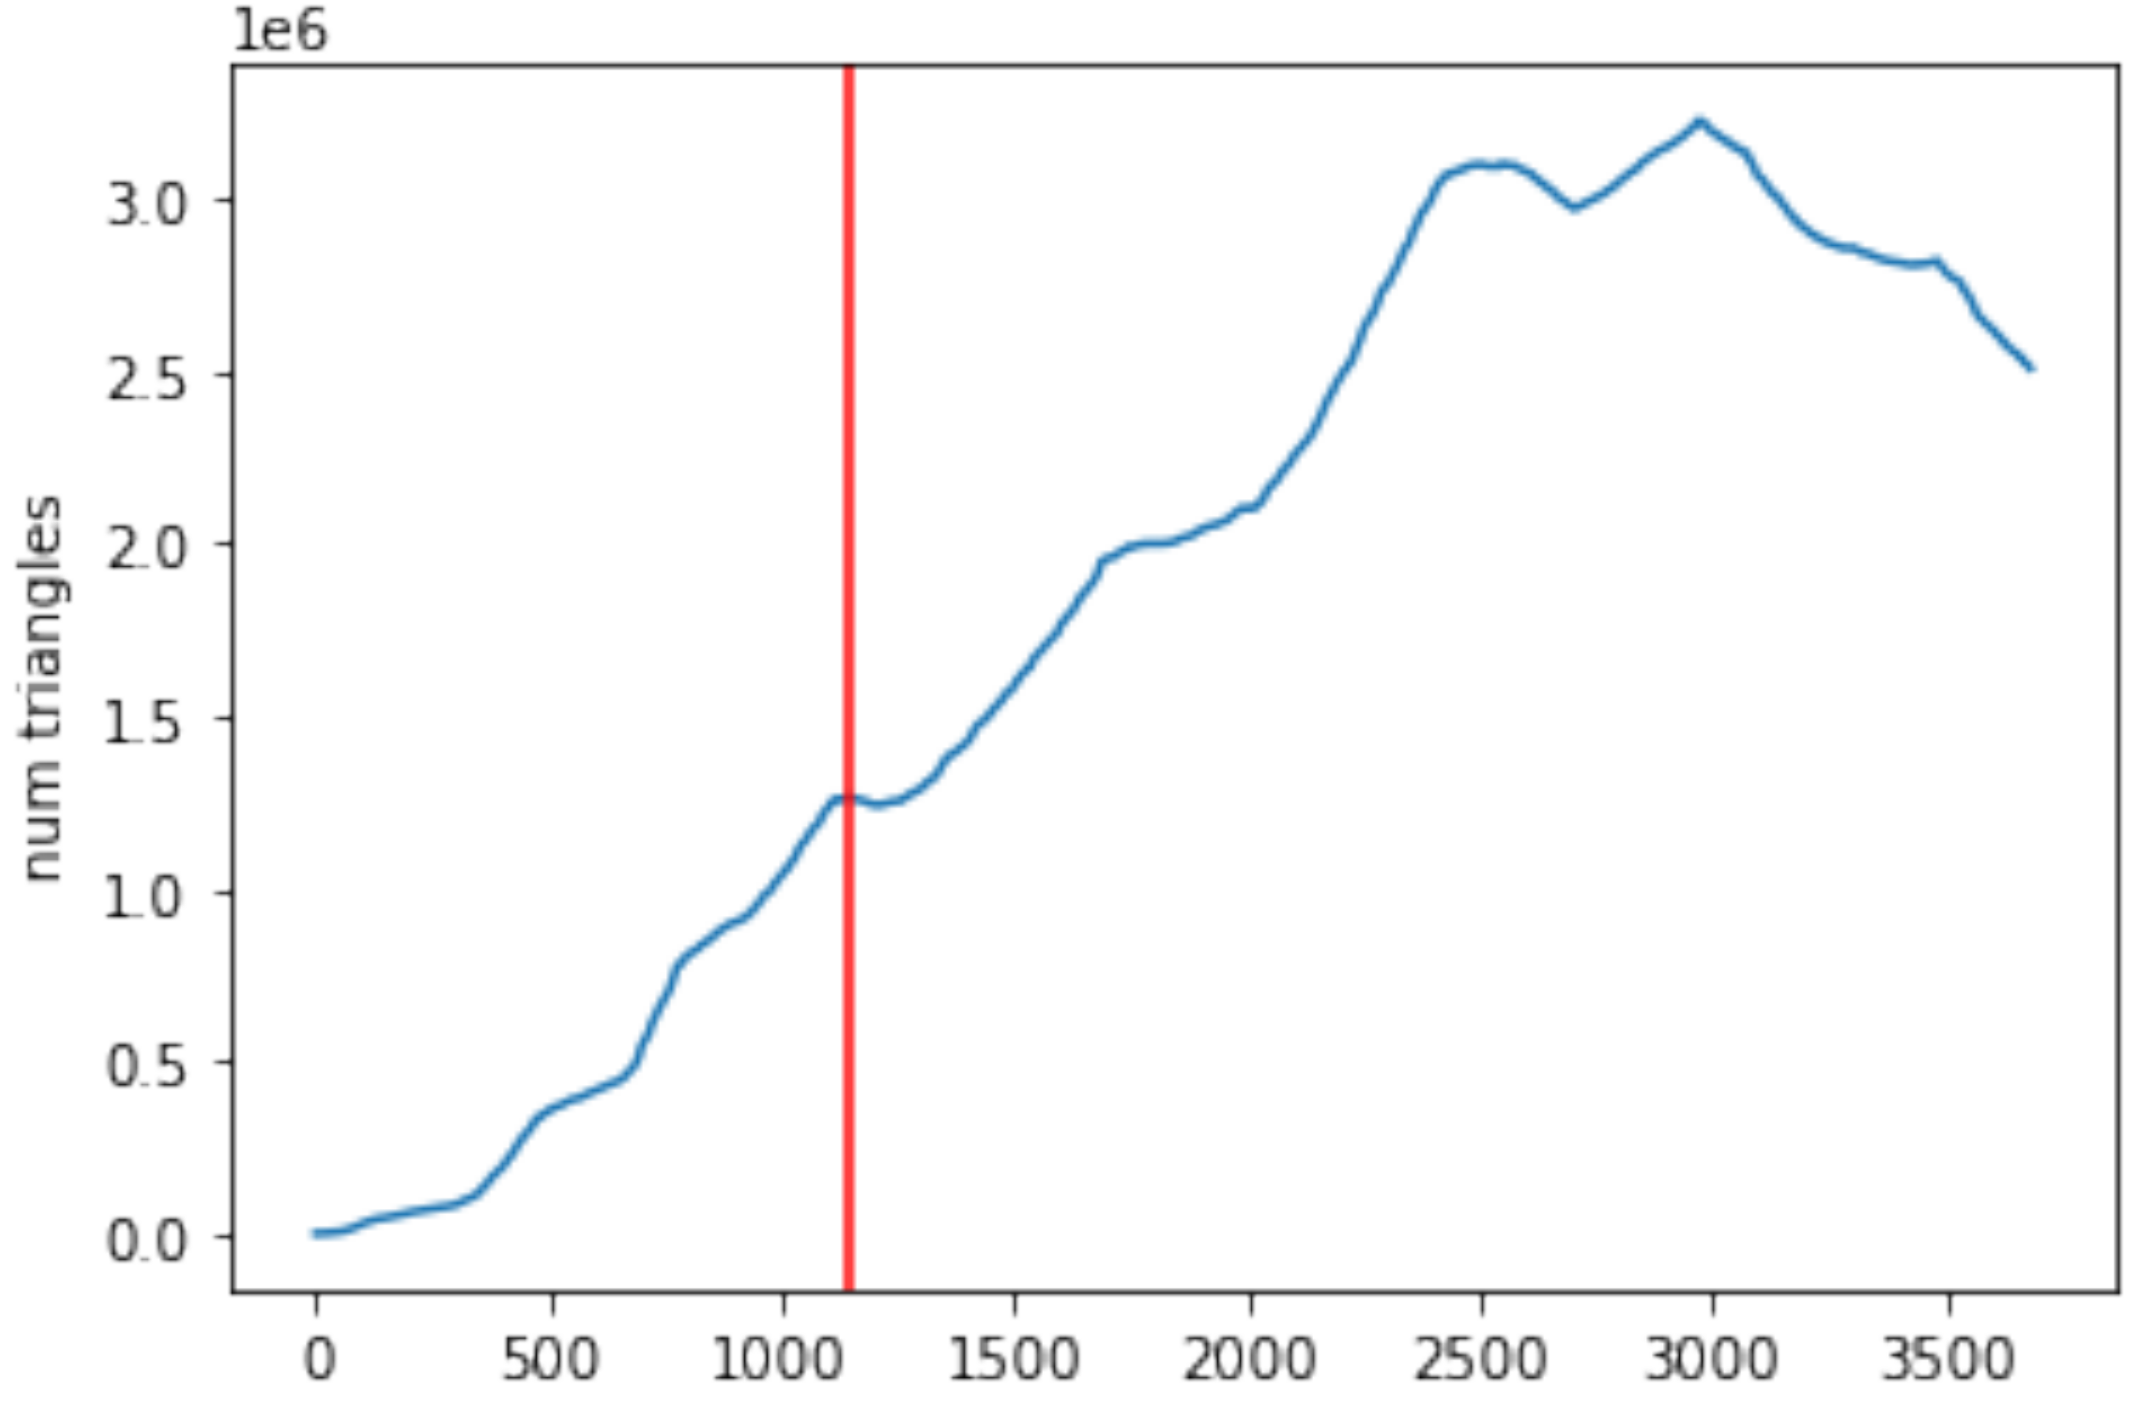

Triangles over time for temporal data set. Vertical line indicates timestep where density threshold was met

Visualizing the temporal graph evolving over time

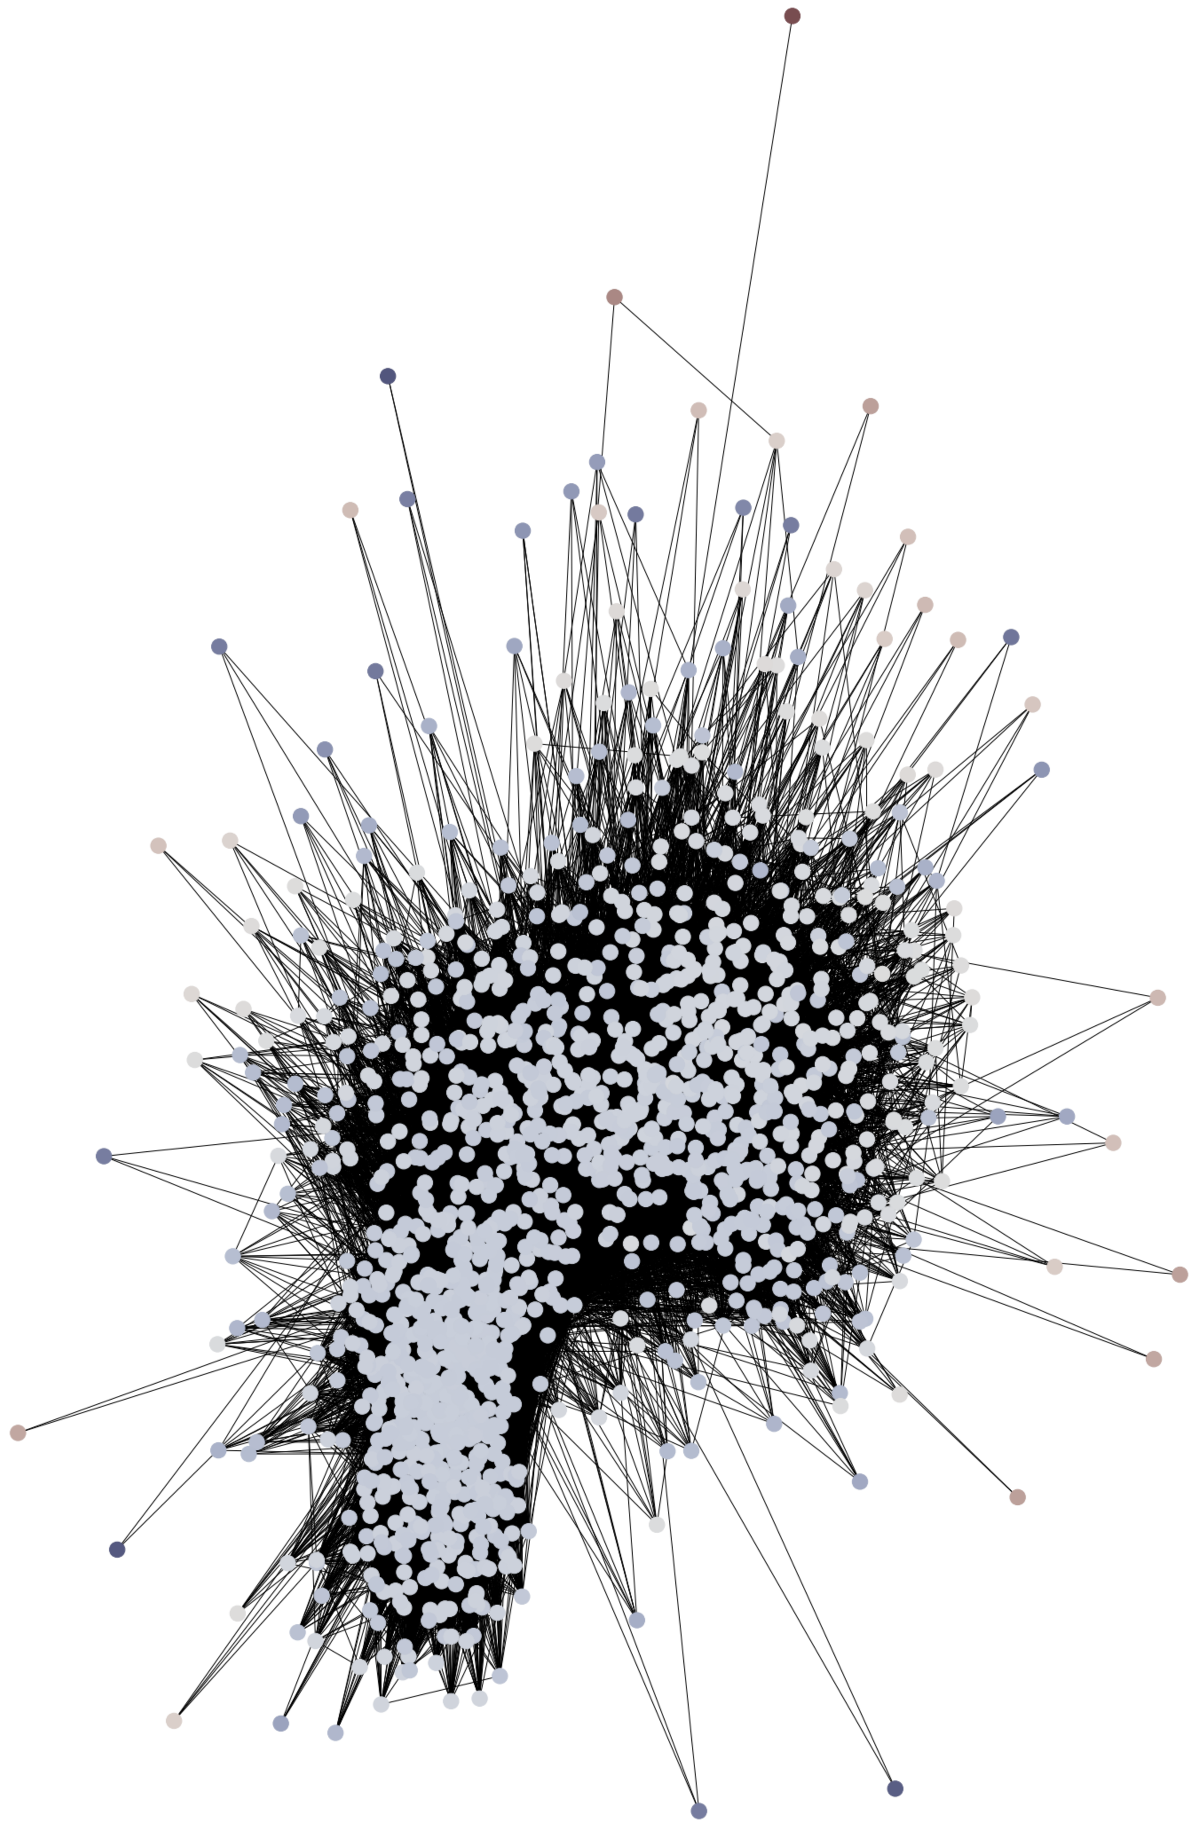

Temporal graph at $t = 0$

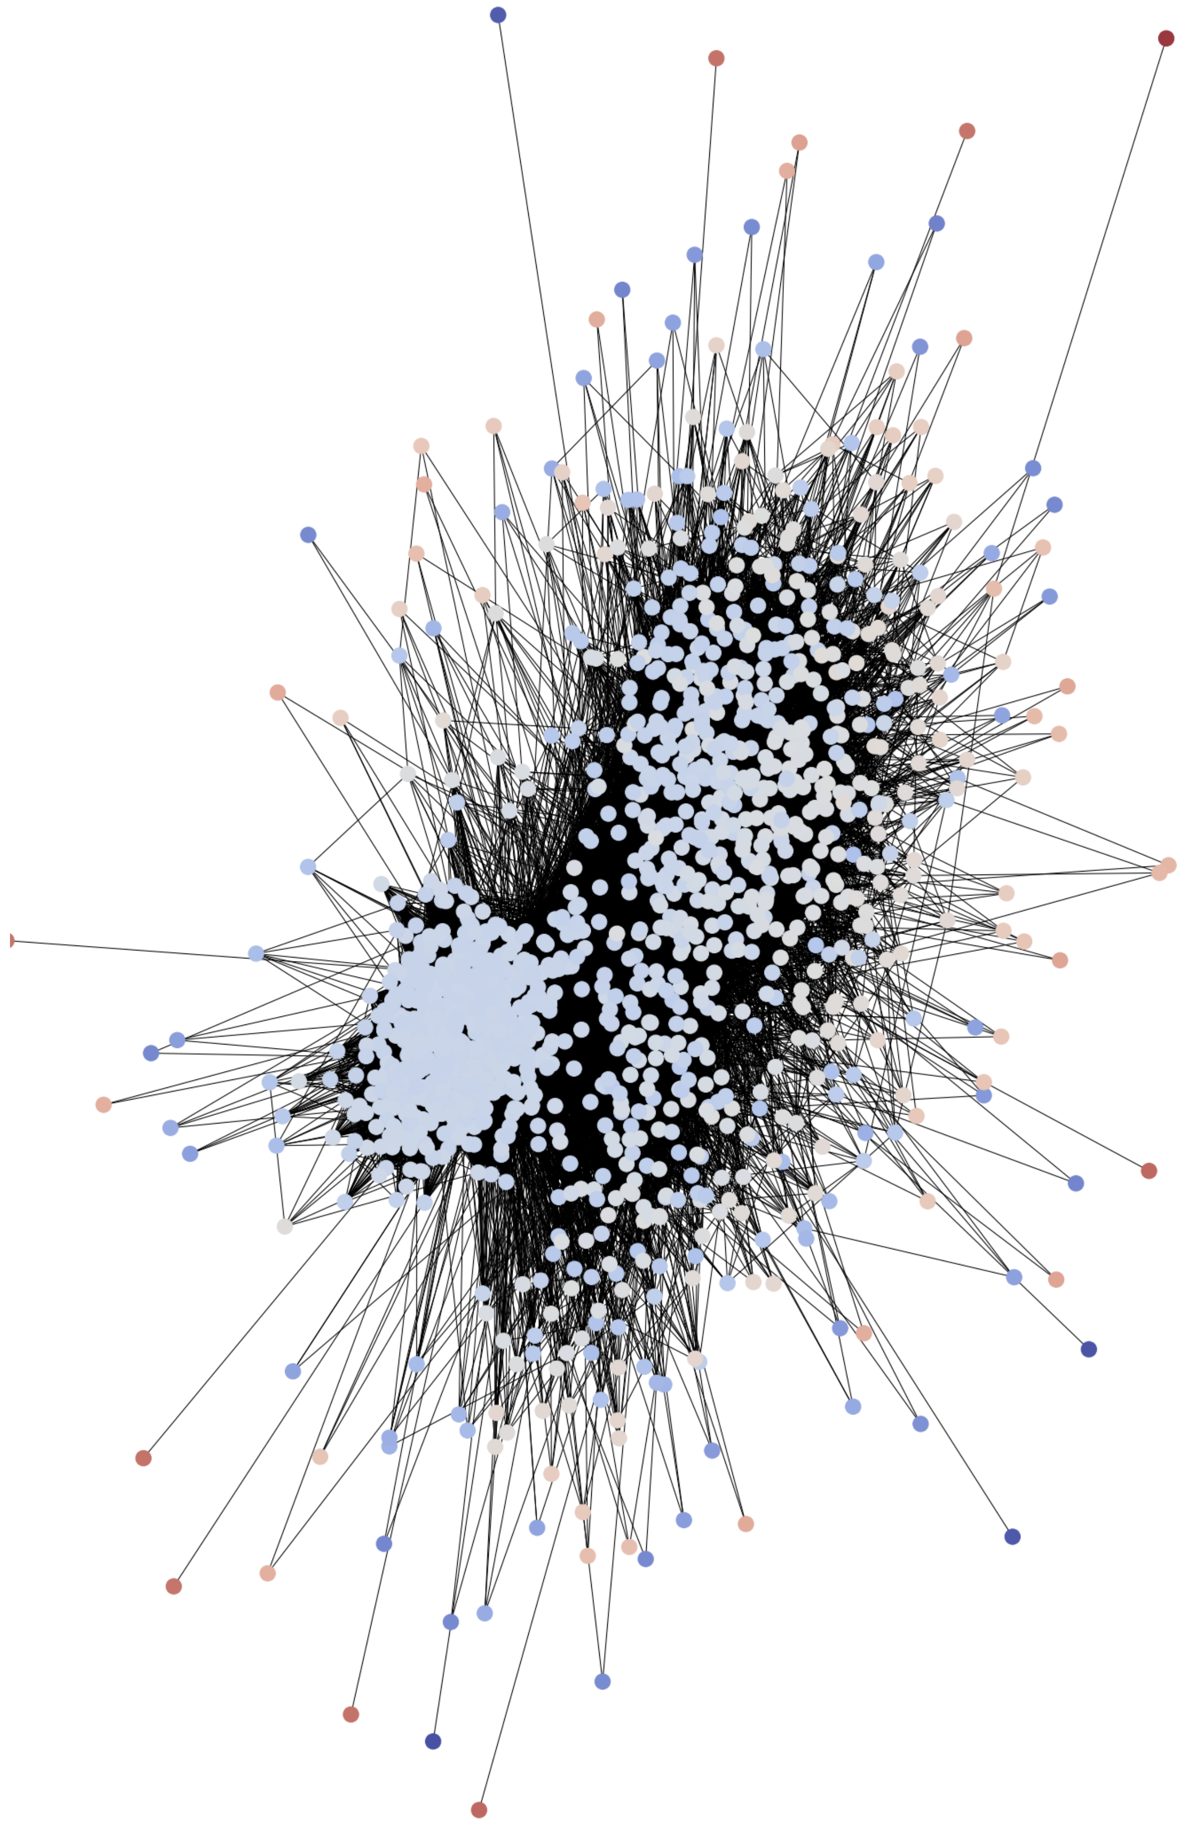

Temporal graph at $t = 2200$

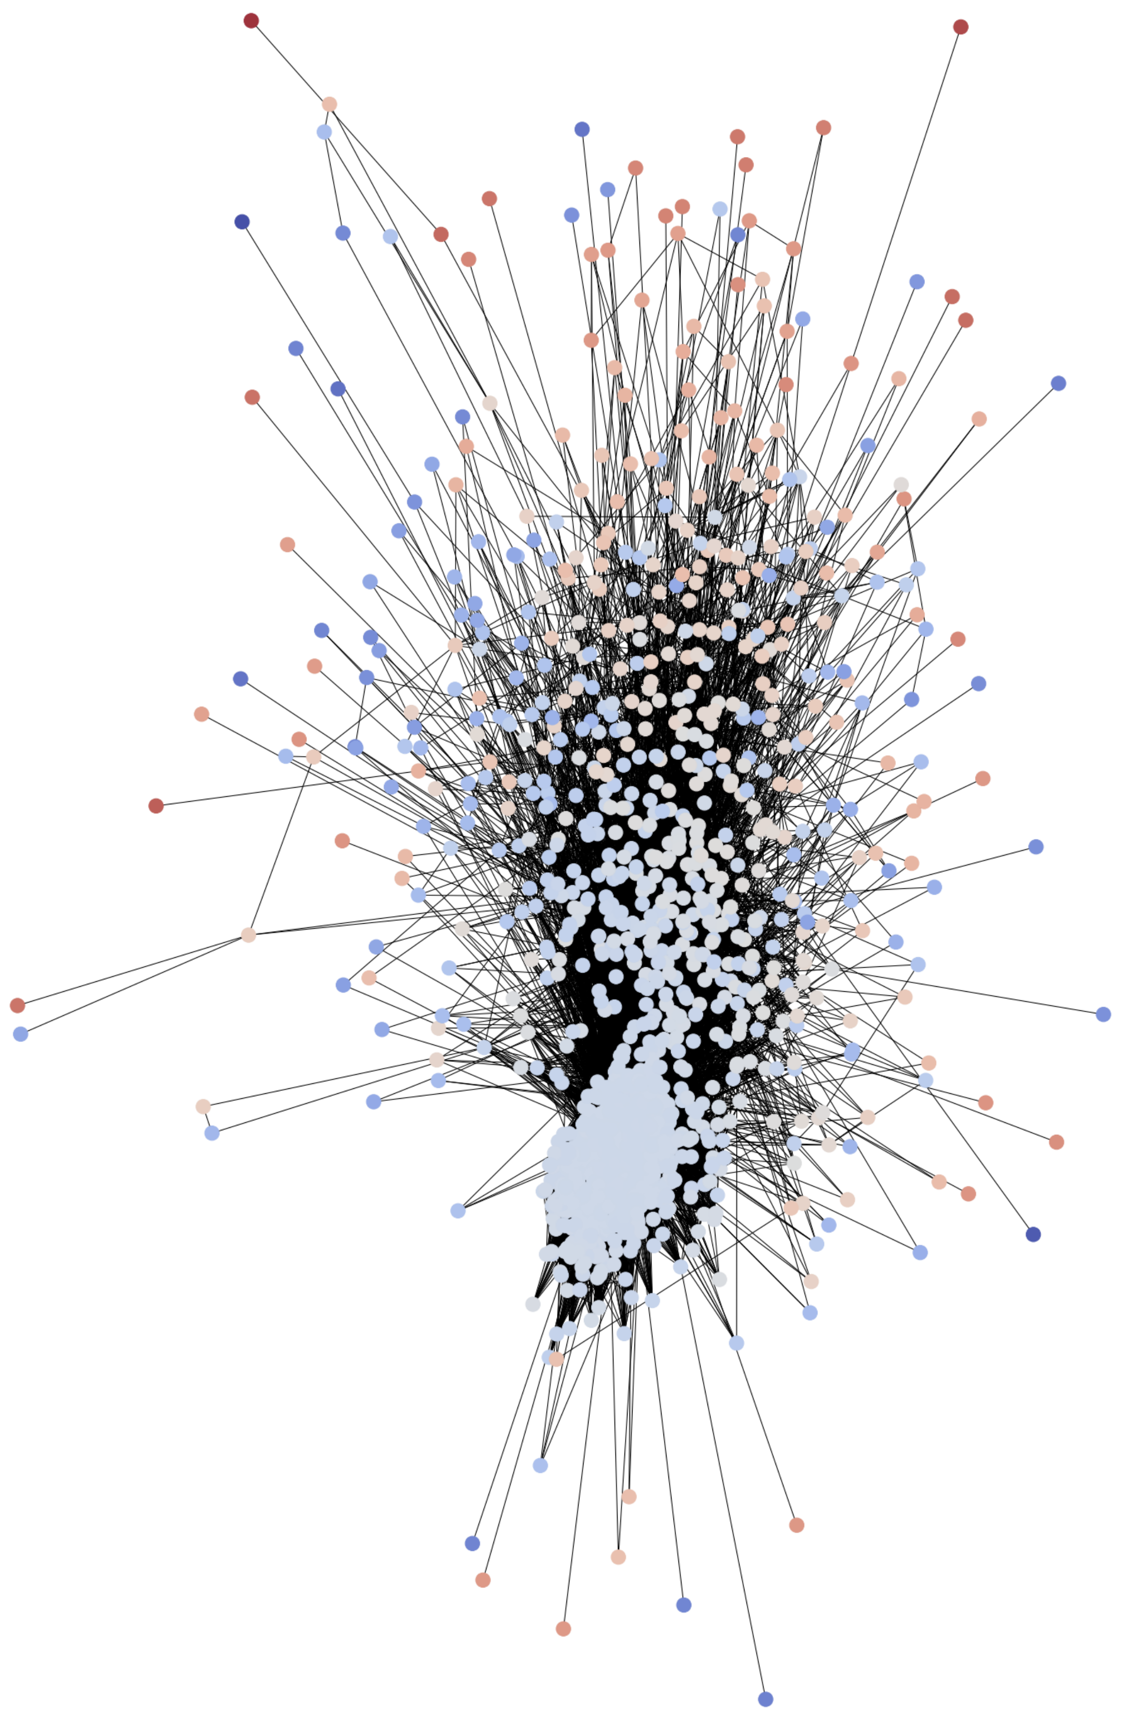

Temporal graph at $t = 2800$

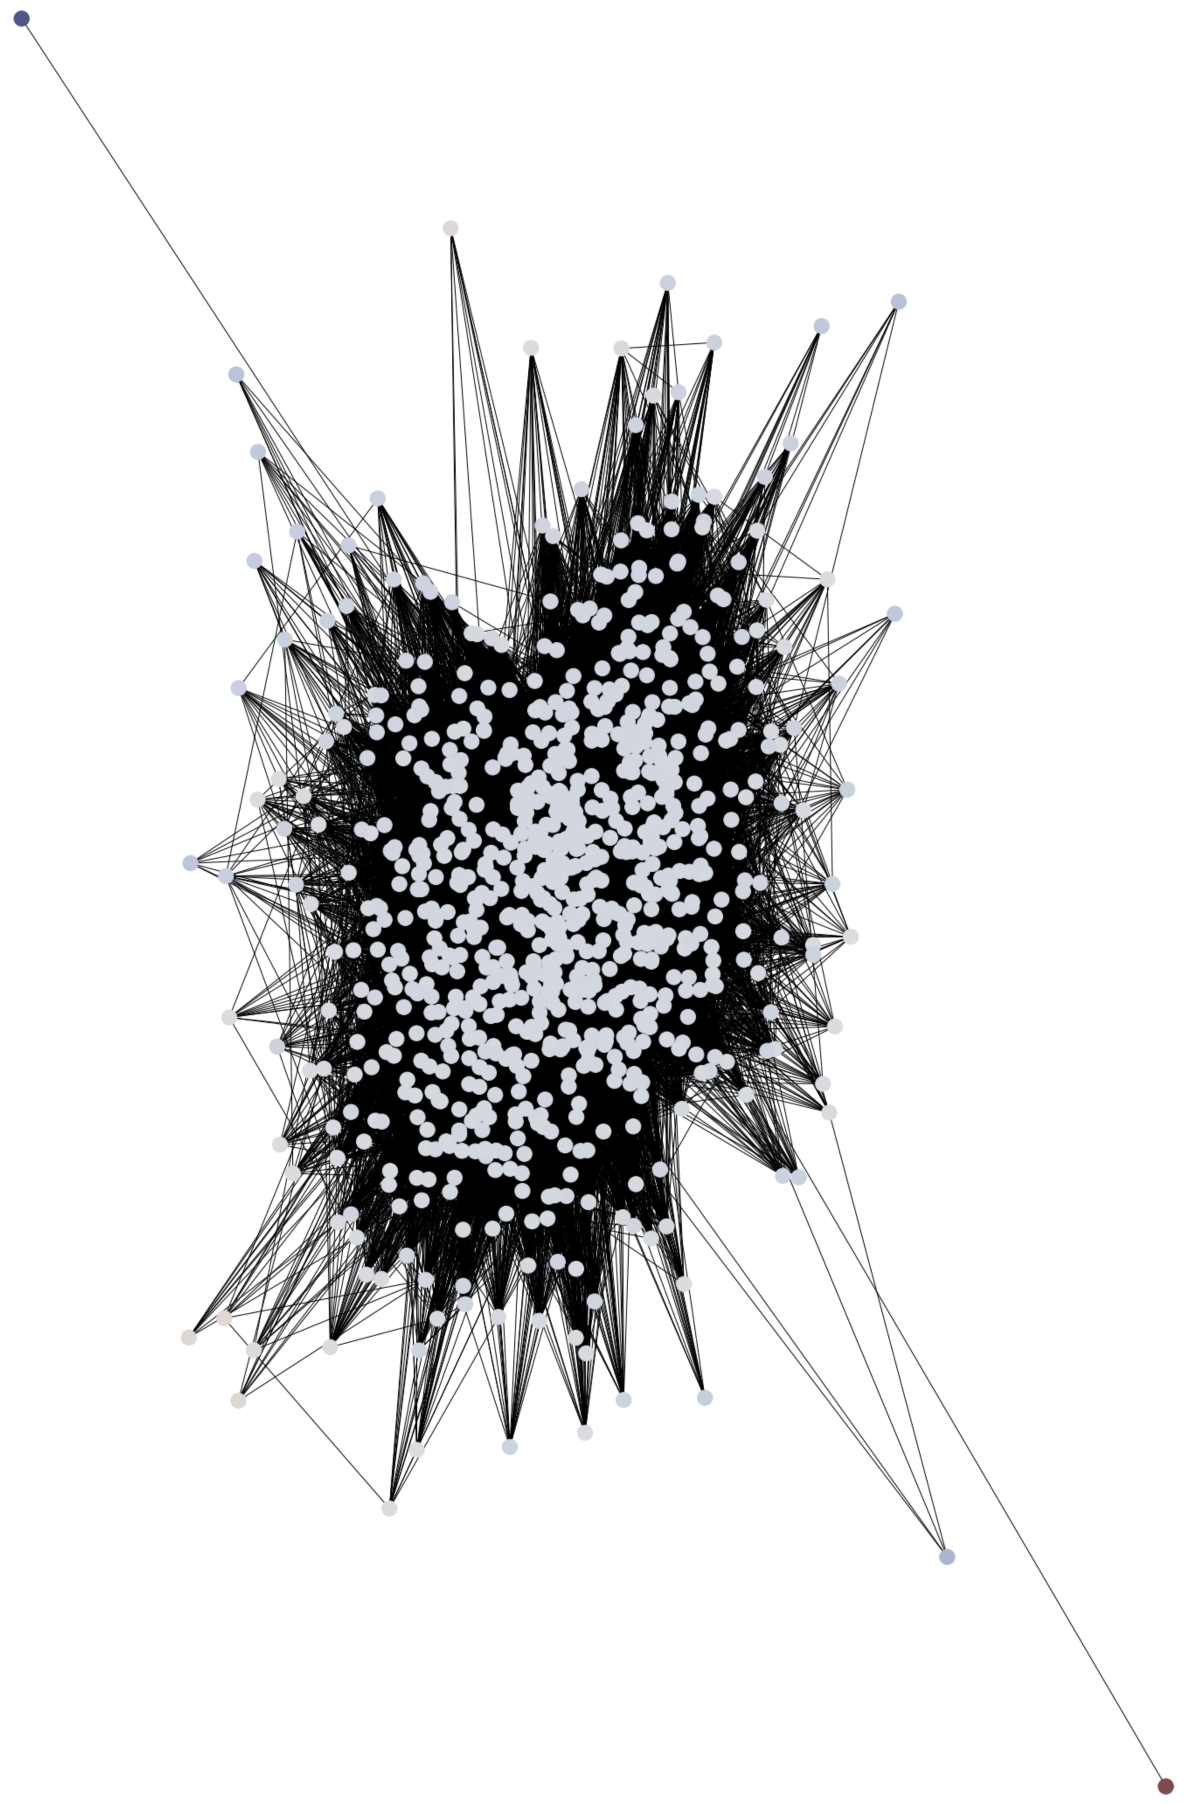

Temporal graph at $t = 3500$

Temporal Verification — Polarization & Disagreement

- Using the party IDs as innate opinions, we see polarization rising over time, and disagreement rises before falling — does looks similar to our model at surface level
- Vertical line indicates the time step when edge deletion is turned on.

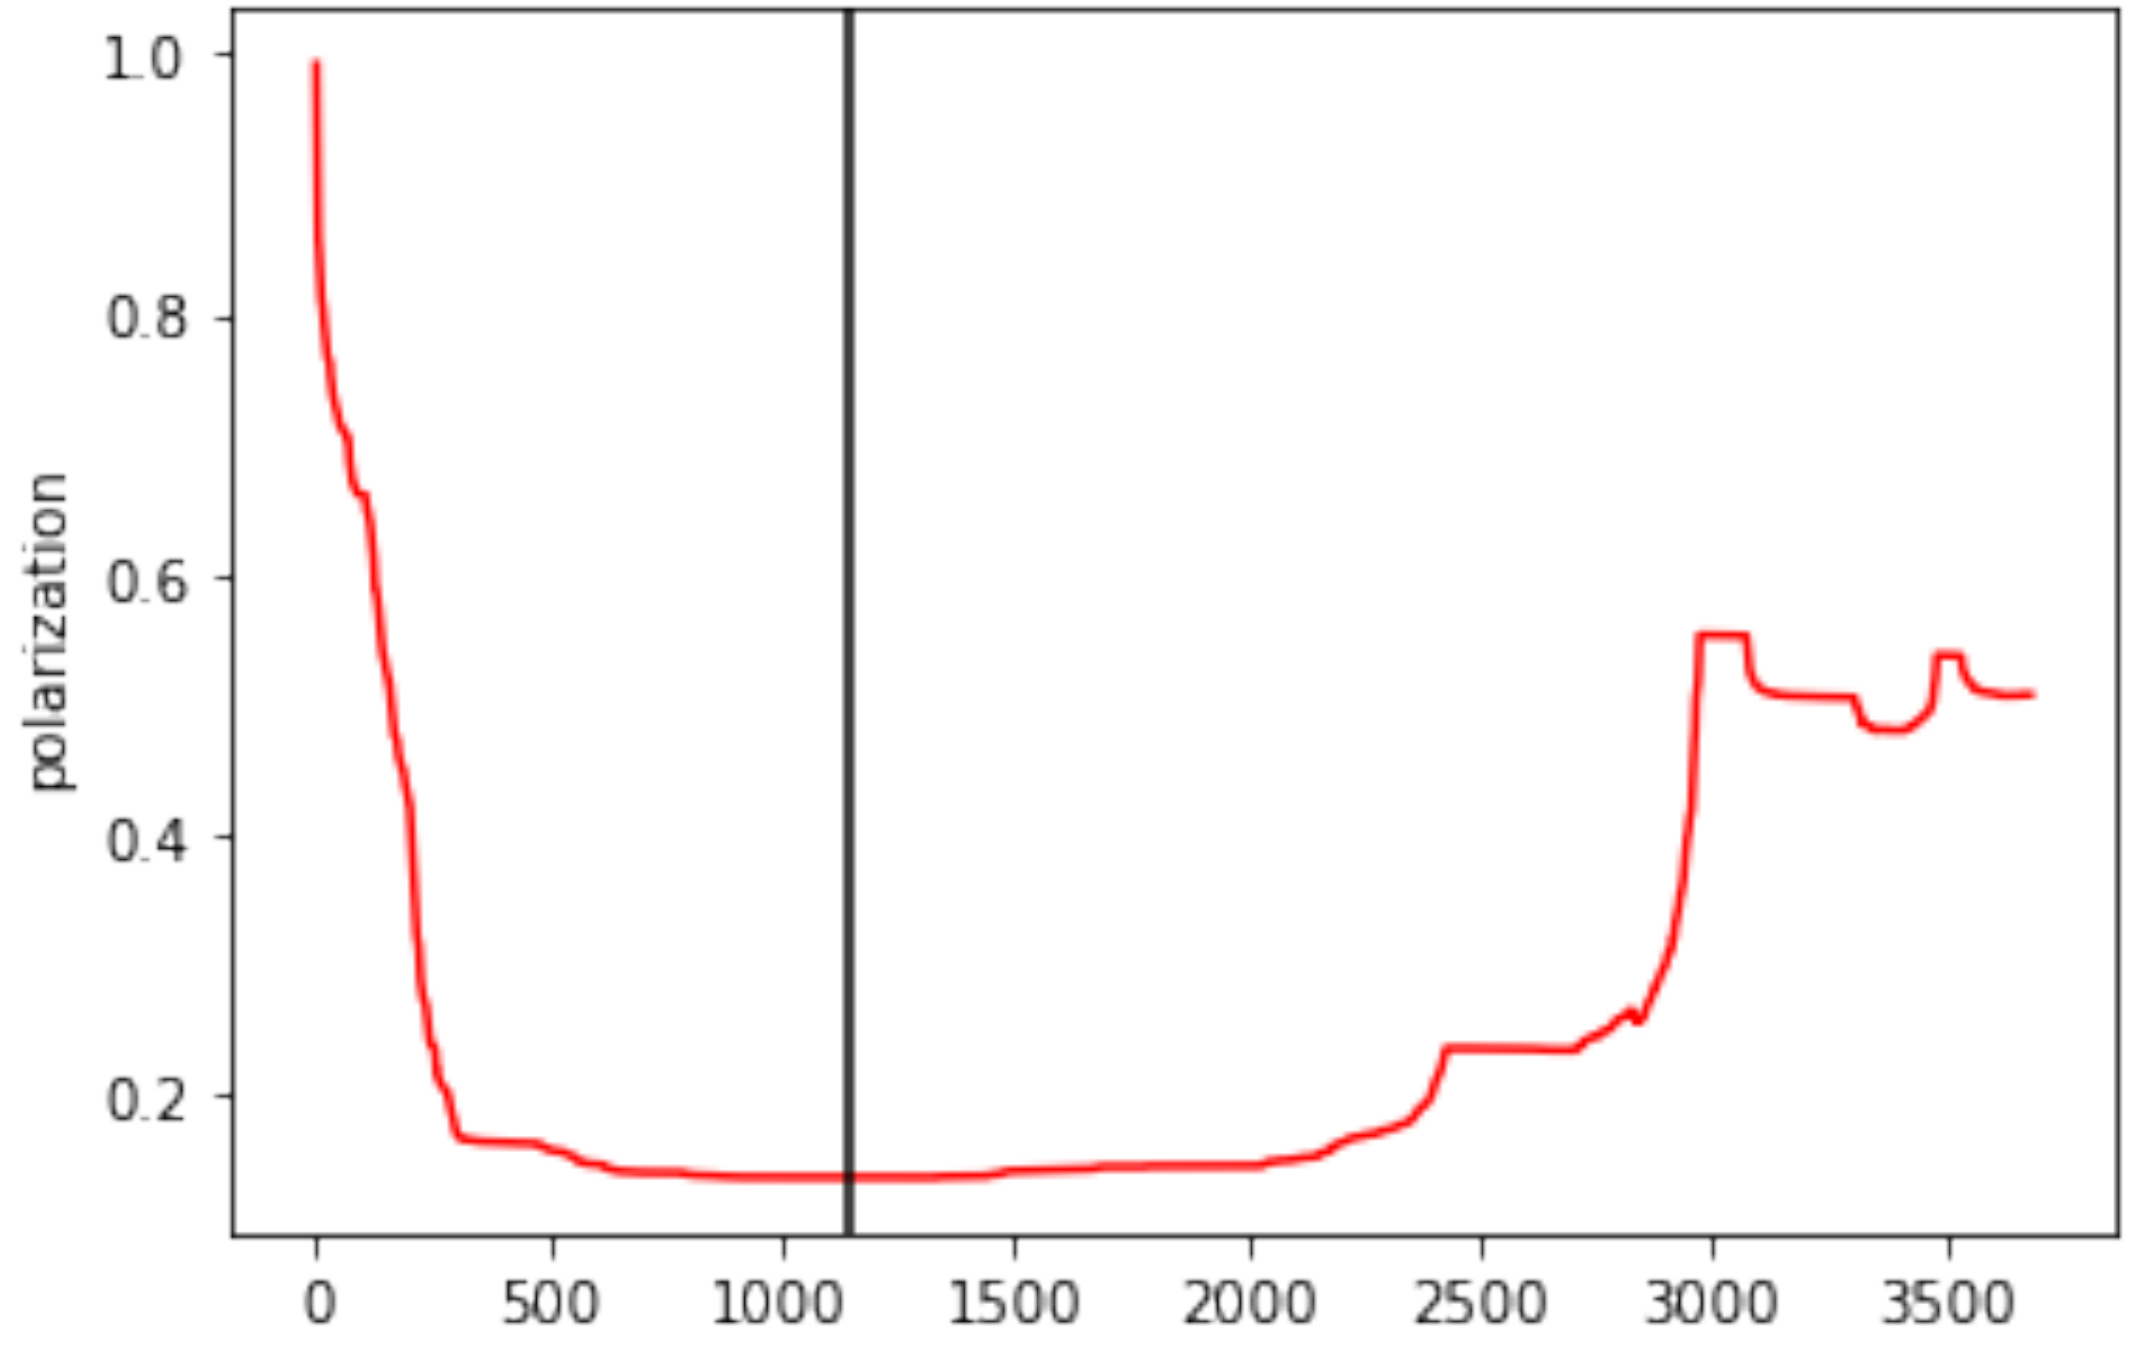

Polarization over time for temporal data set. Vertical line indicates timestep where density threshold was met.

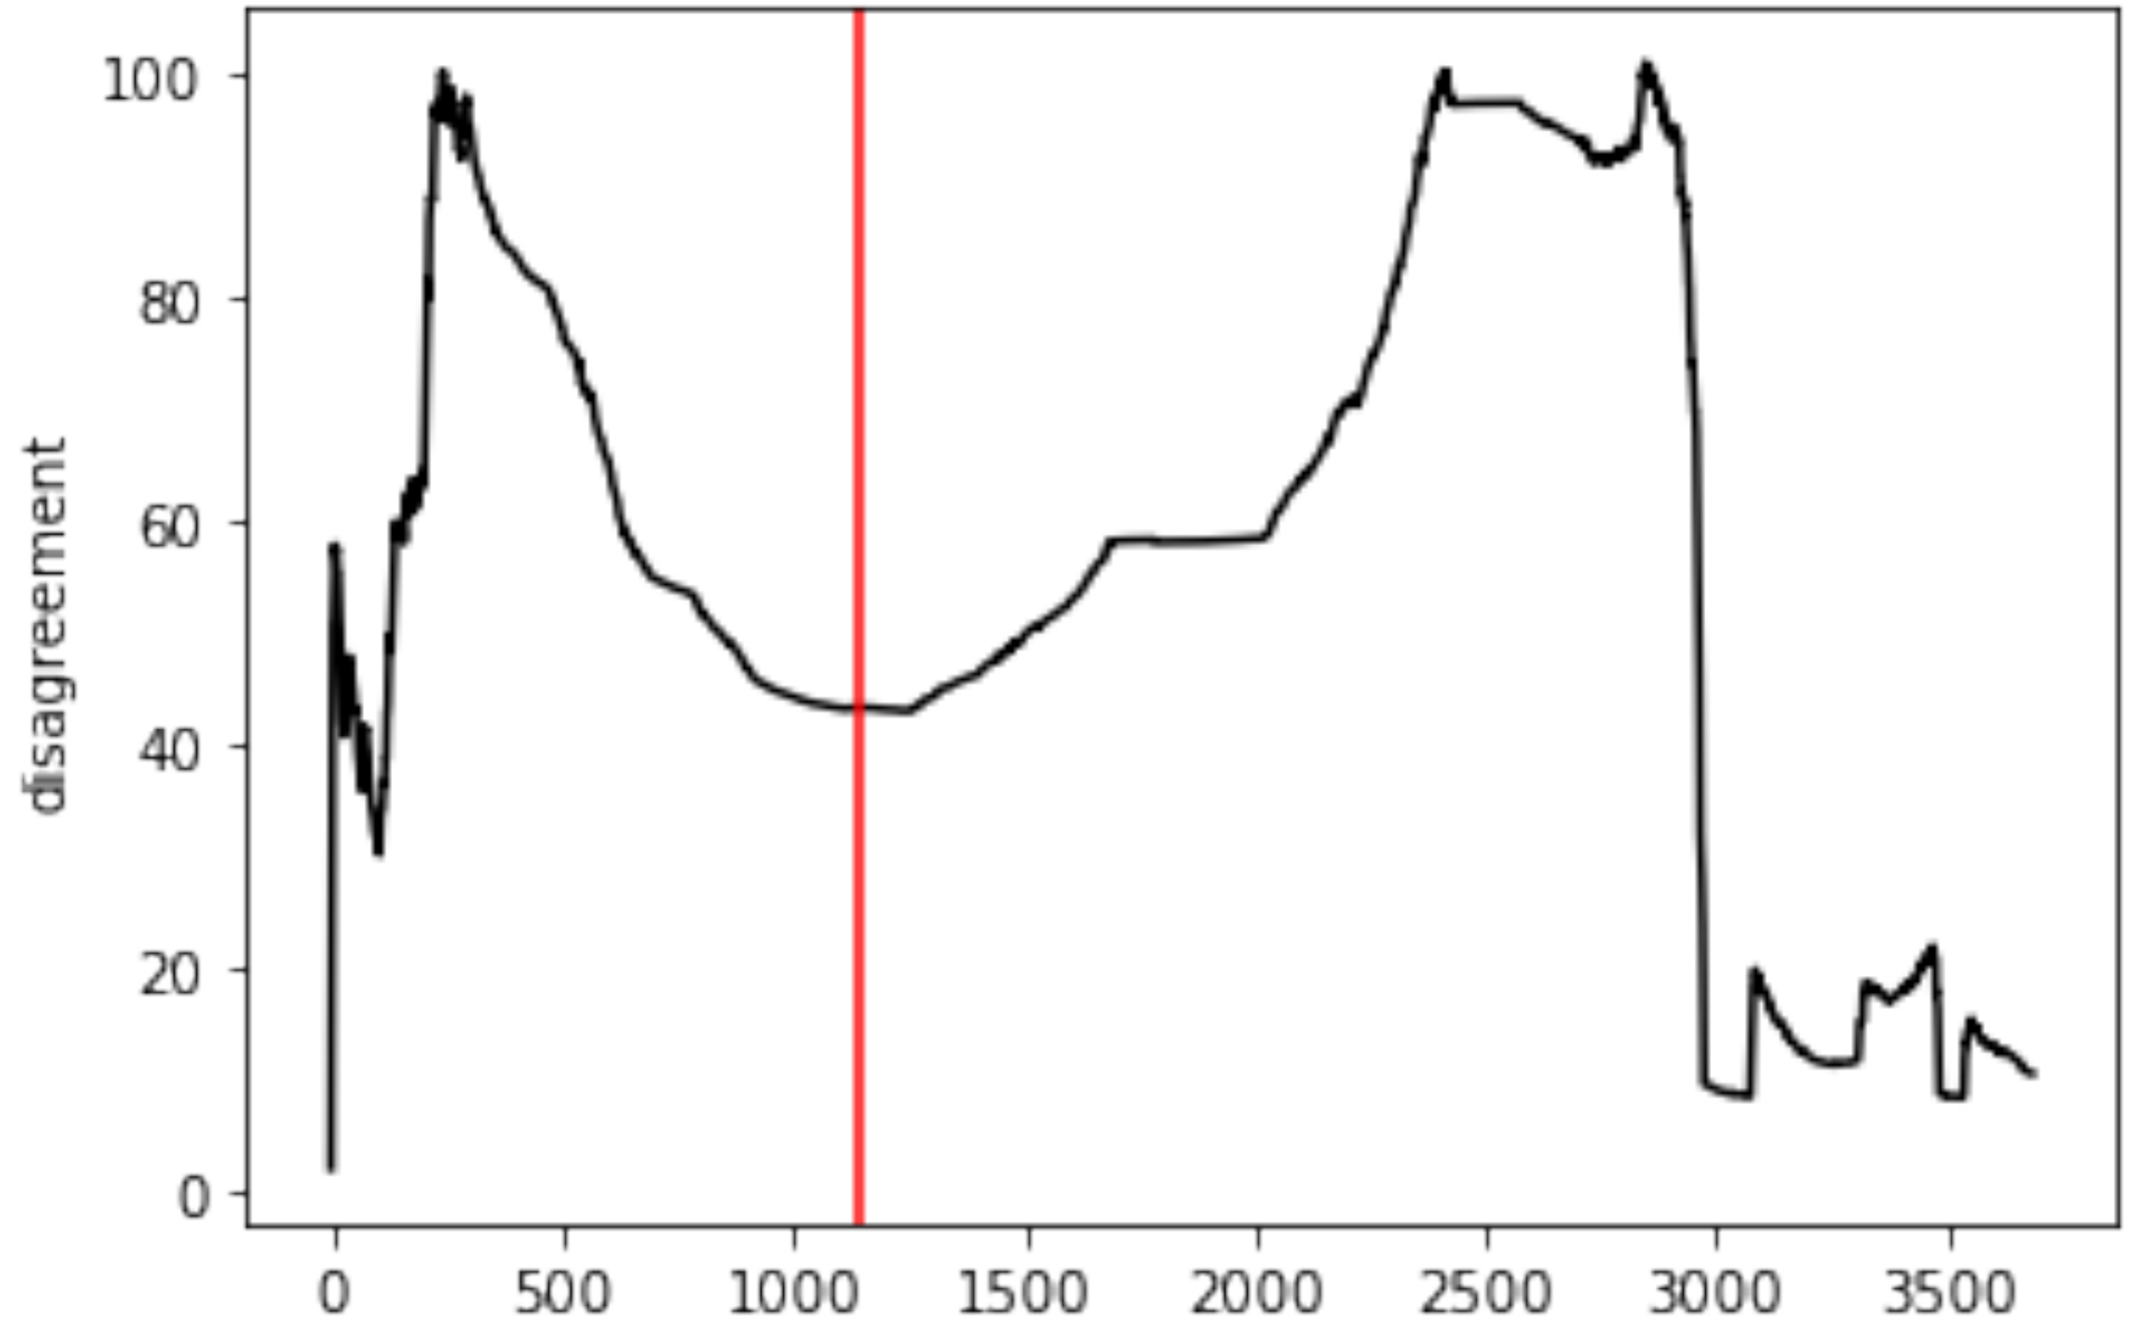

Disagreement for temporal data set. Vertical line indicates timestep where density threshold was met

Temporal Verification — Polarization & Disagreement

- Using the party IDs as innate opinions, we see polarization rising over time, and disagreement rises before falling — does looks similar to our model at surface level
- Vertical line indicates the time step when edge deletion is turned on.

- Quantities of polarization and disagreement mirror the **size of the largest connected component.**
- After edge deletion is turned on, the graph starts disconnecting around $t = 3000$
- Inconclusive result, but there is evidence that nodes are being sorted into clusters

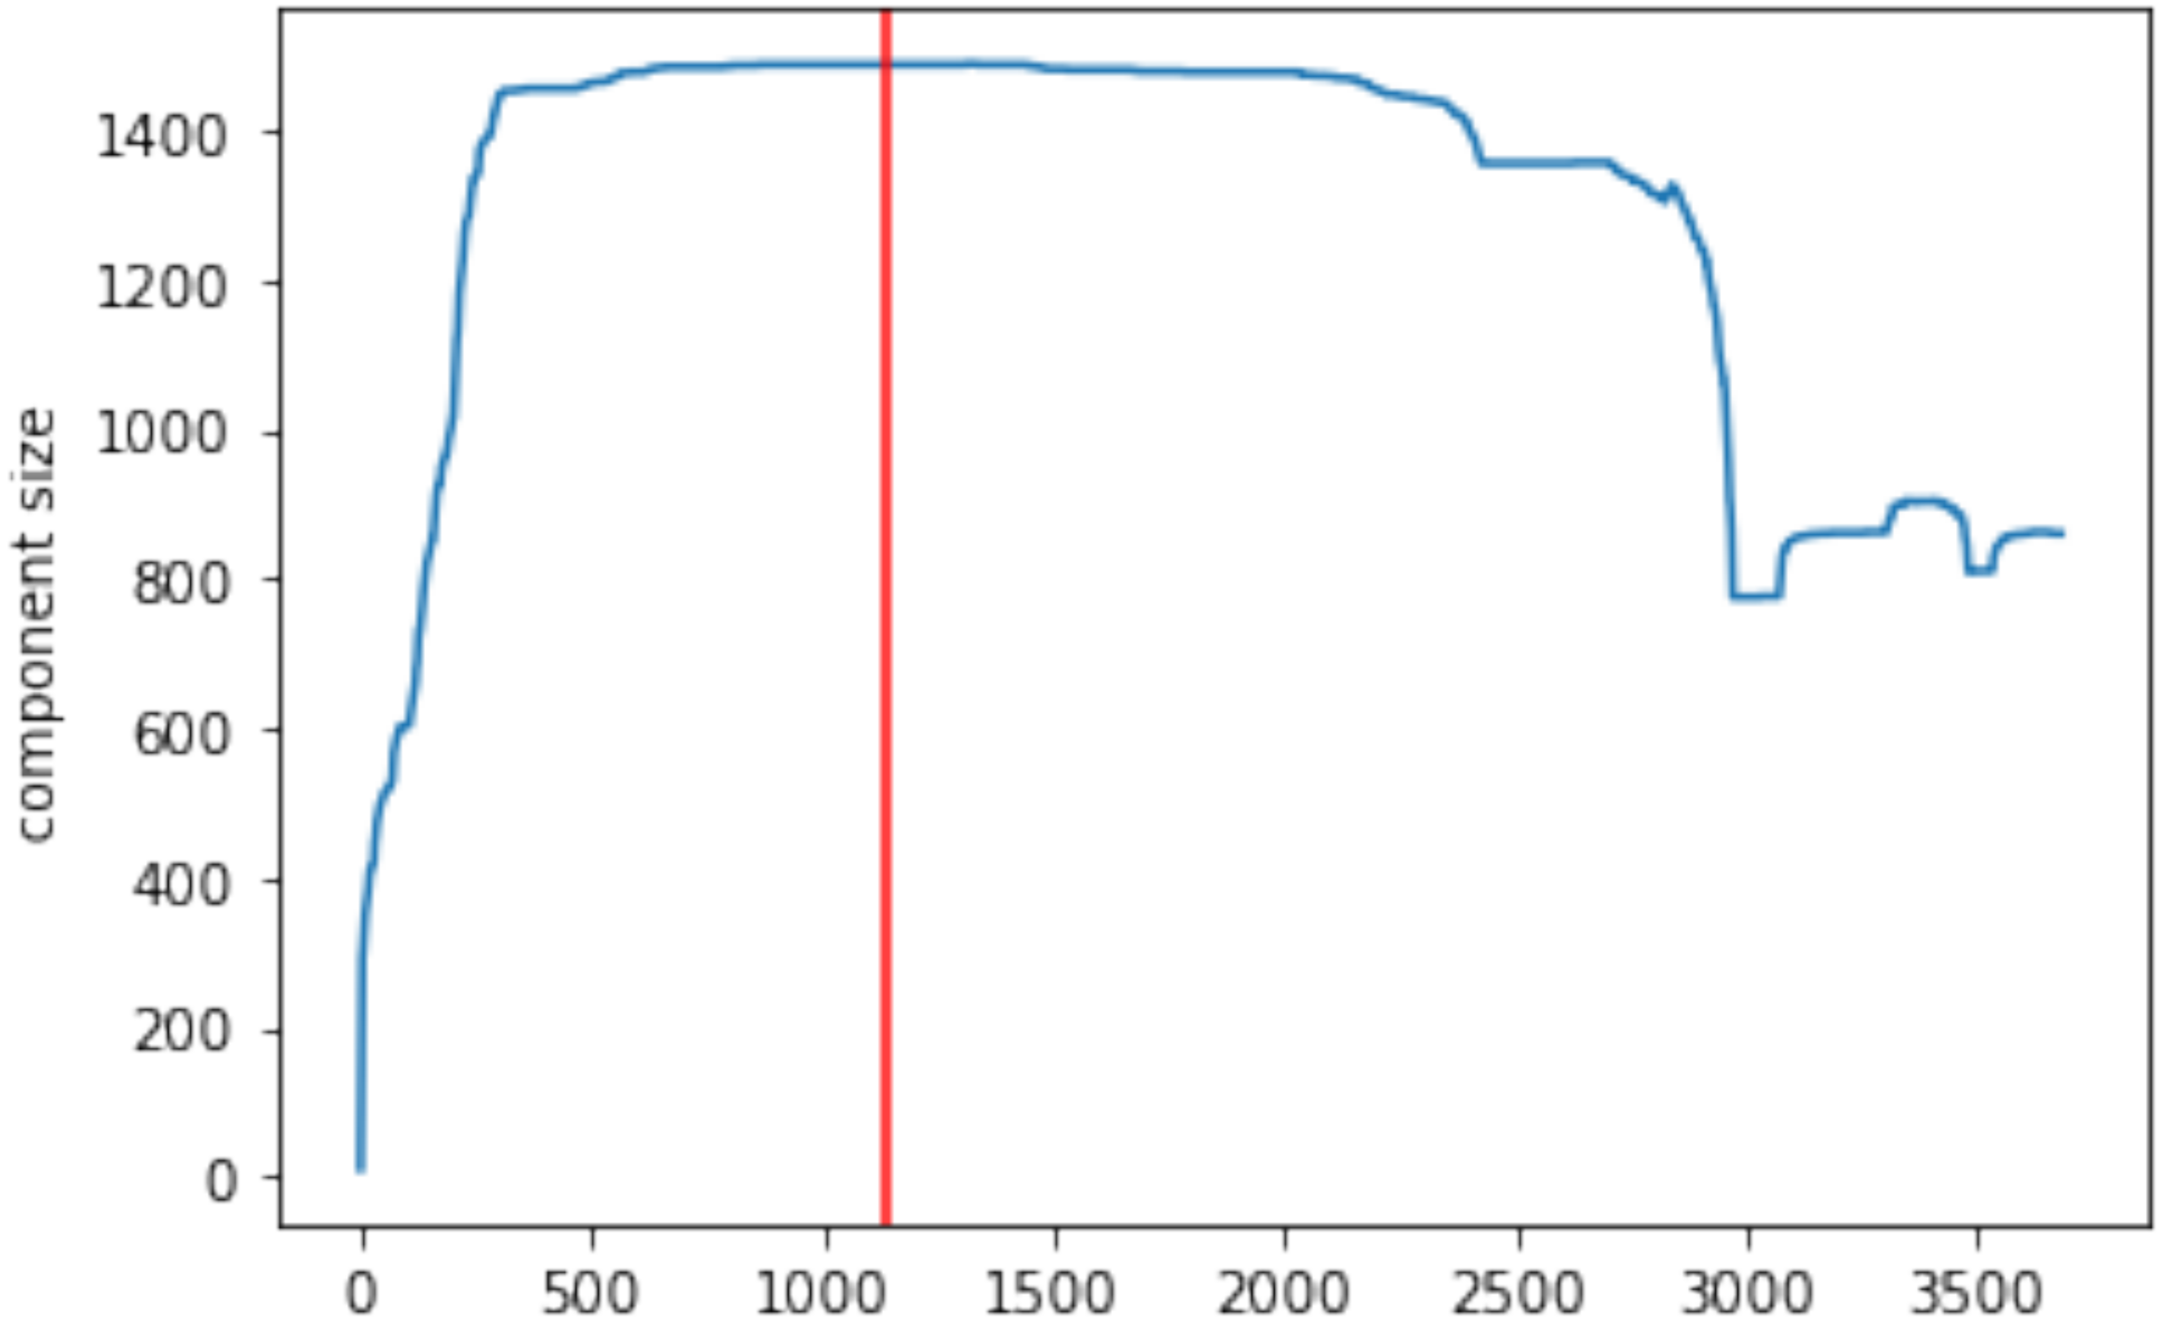

Size of largest connected component for temporal data set. Vertical line indicates timestep where density threshold was met

Conclusions / Future Steps

- We introduce a model of social networks that combines F-J opinion dynamics with edge dynamics; friend-of-friend recommendations and confirmation bias edge deletion.

- We think we have a fairly strong theoretical understanding of our model.

  - The techniques used in the theoretical study of our model could be extended further — leverage relationships to approximate different quantities, generalize to more diverse opinion settings, etc.

  - Potential to extend past the simple F-J model shown here for more realistic opinion dynamics.

Conclusions / Future Steps

- Interesting preliminary results for real-world verification.

  - For the snapshot method, generalization to other measures such as *small world quotient*, *closeness centrality* is something to explore further — also, fixed edges aren't the only way that we could constrain edge dynamics.

  - Initial temporal results are inconclusive, but worth further inquiry — having access to an evolving real-world network opens up a number of possibilities for further comparison

Q & A / Discussion