GENEVA GRADUATE INSTITUTE

INSTITUT DE HAUTES ÉTUDES INTERNATIONALES ET DU DÉVELOPPEMENT

GRADUATE INSTITUTE OF INTERNATIONAL AND DEVELOPMENT STUDIES

ERGM

Introduction to Social Networks

James Hollway

Bay area network of actors and water management institutions

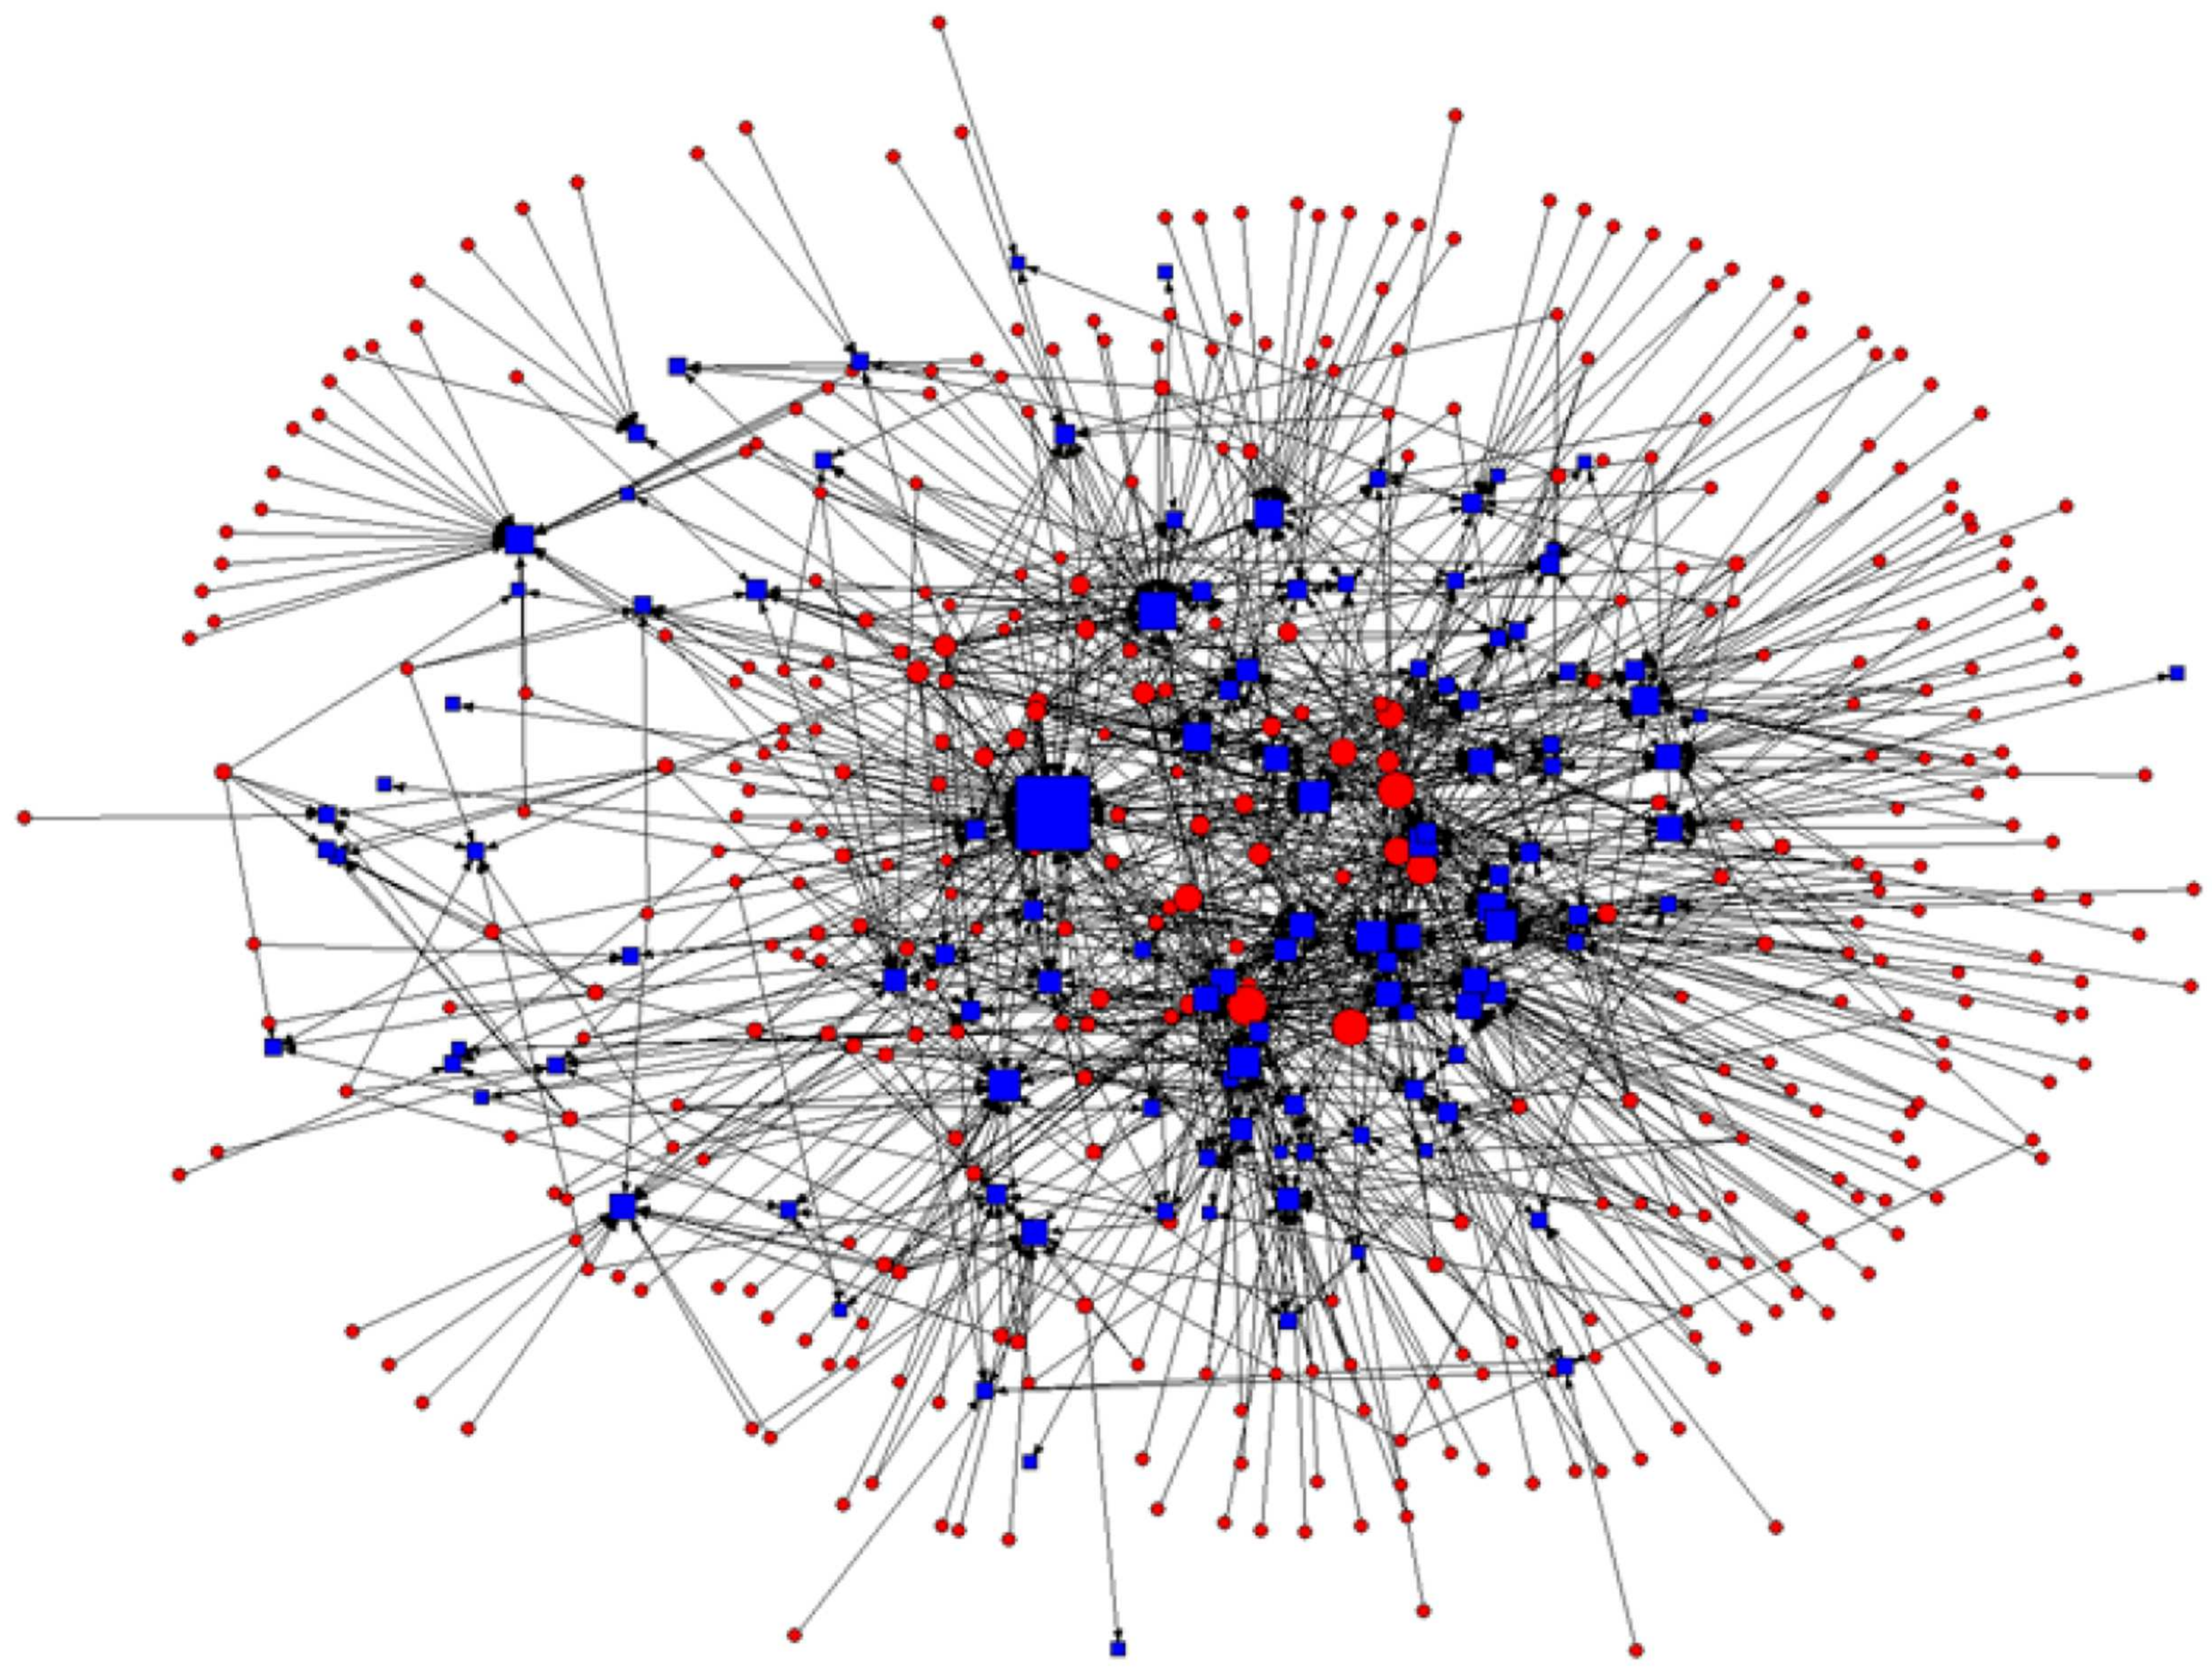

Questions?

Why some institutions more popular than others?

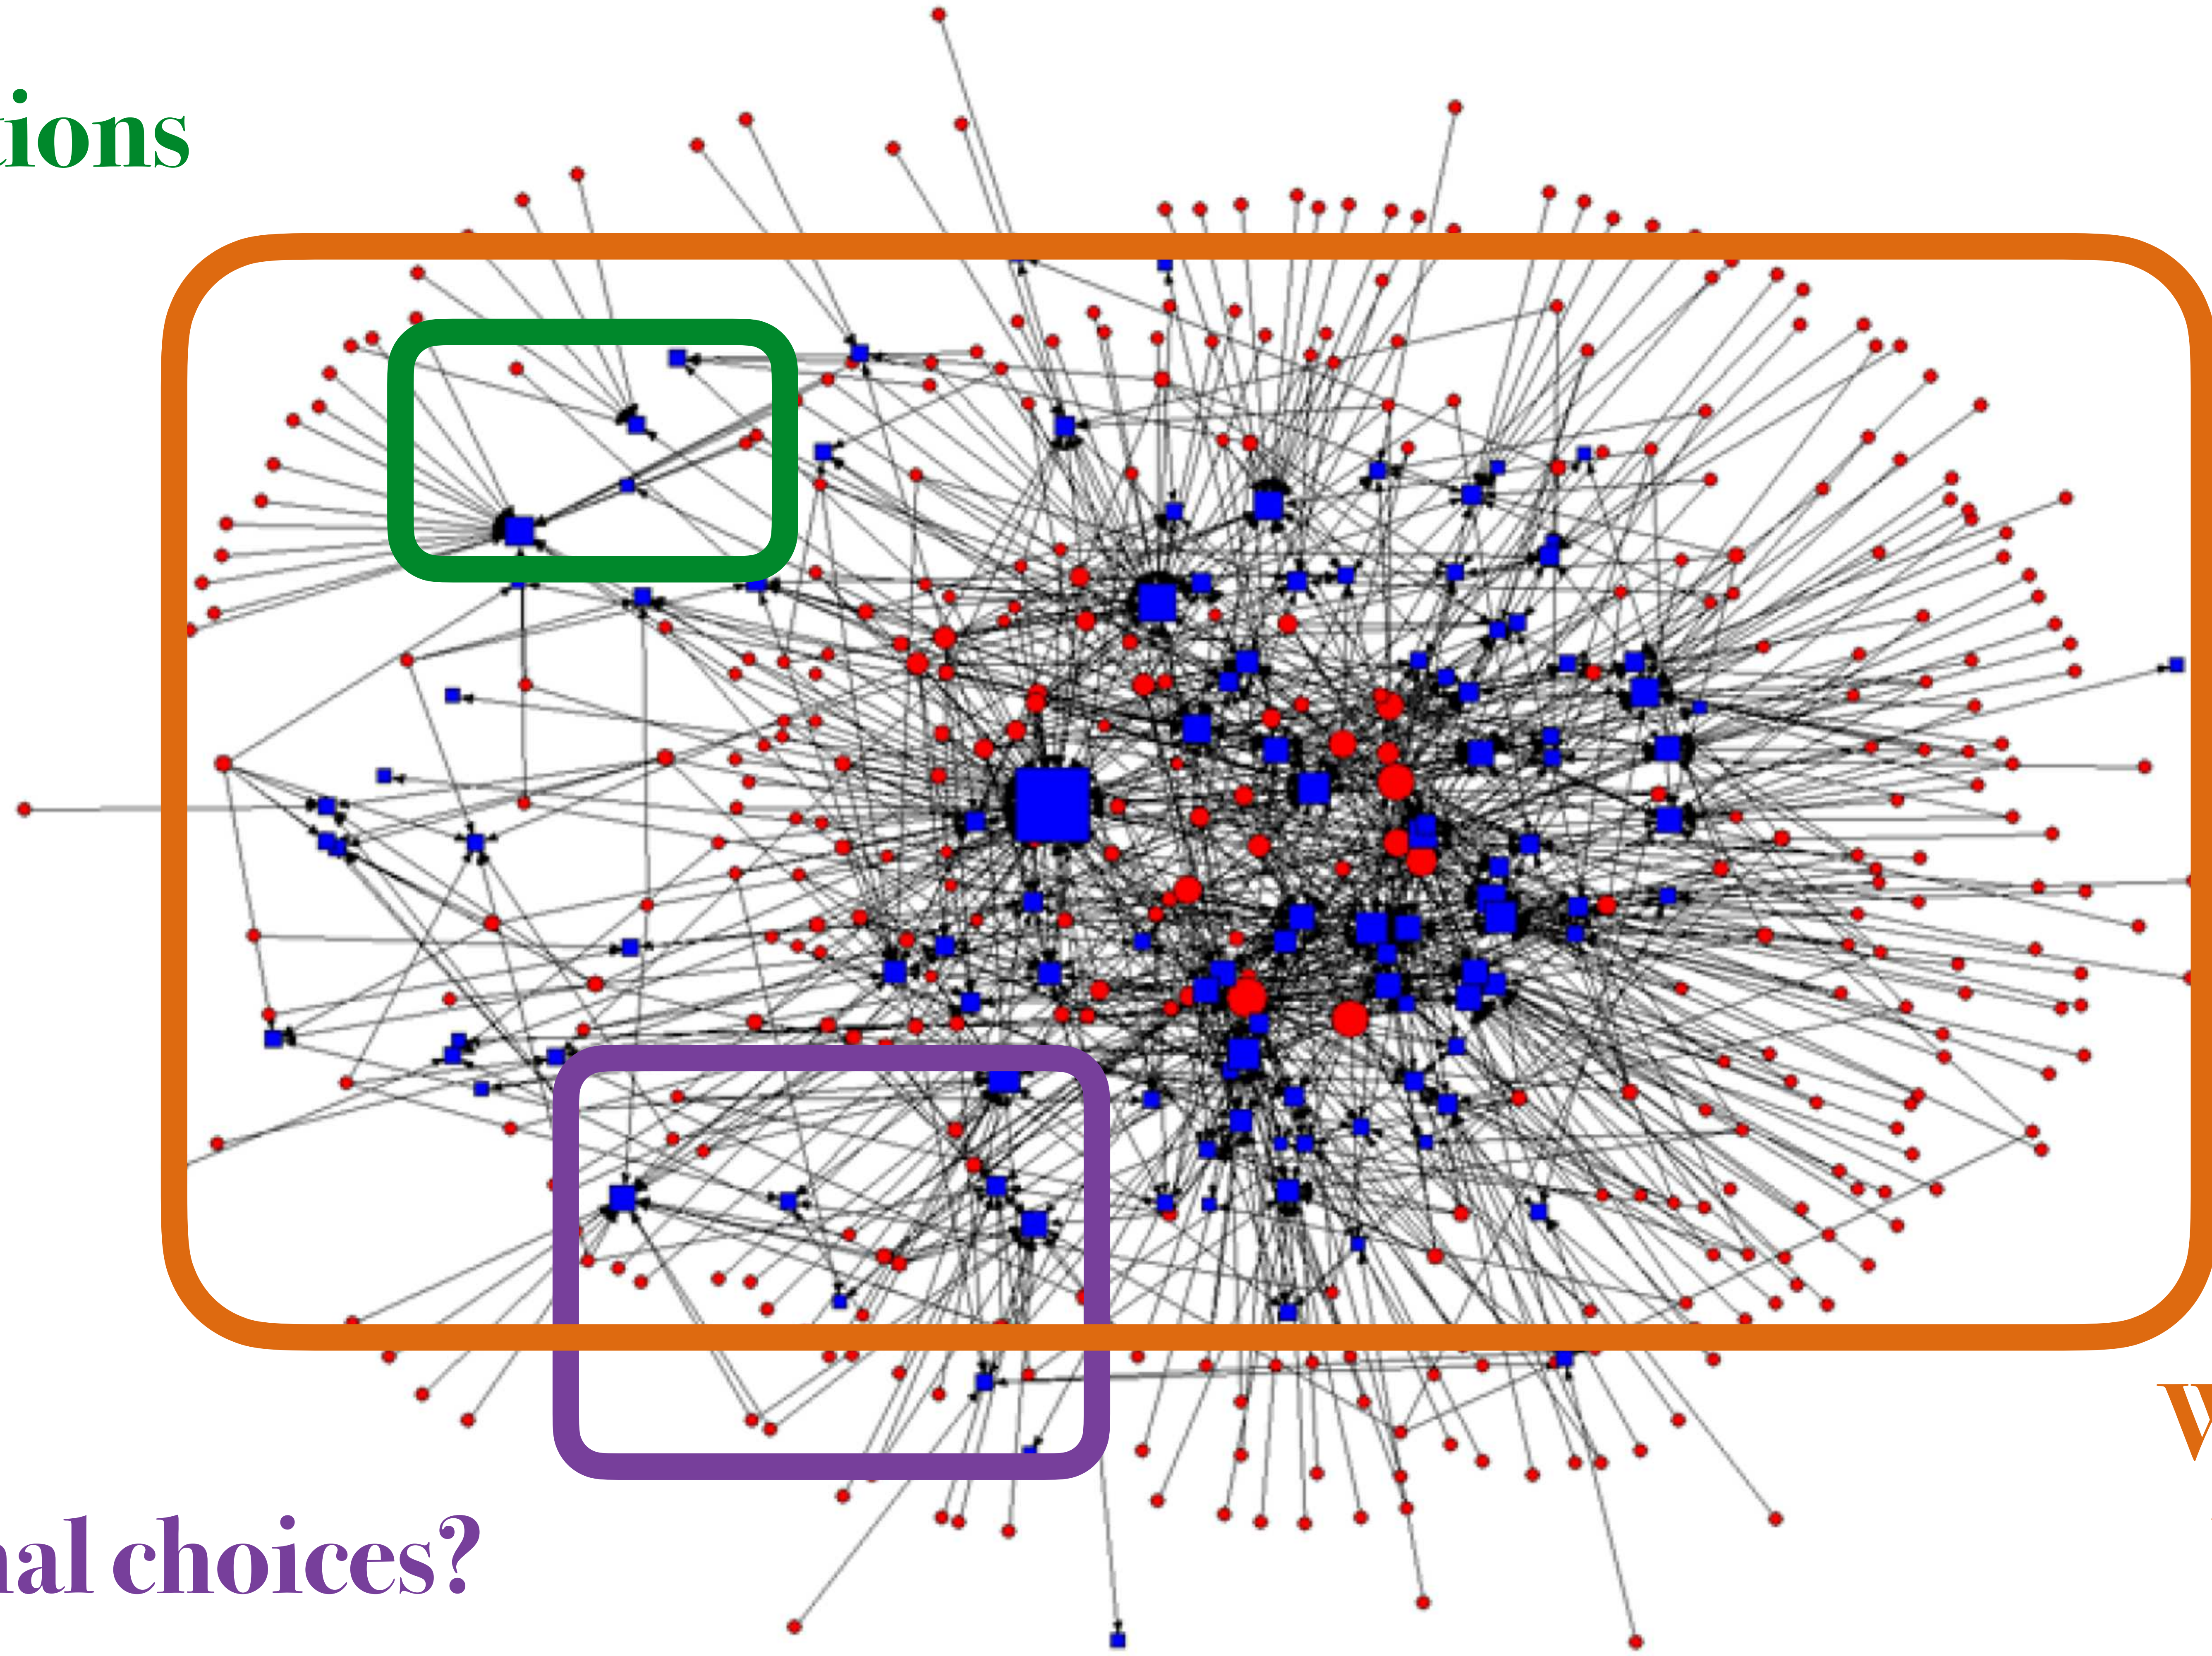

Do actors overlap in their institutional choices?

Why does the network have this structure?

Answers?

Design (Koremenos et al 2001)

Homophily (McPherson et al 2001)

Popularity (Merton 1968)

Transitivity (Simmel 1902)

Structural holes (Burt 1994)

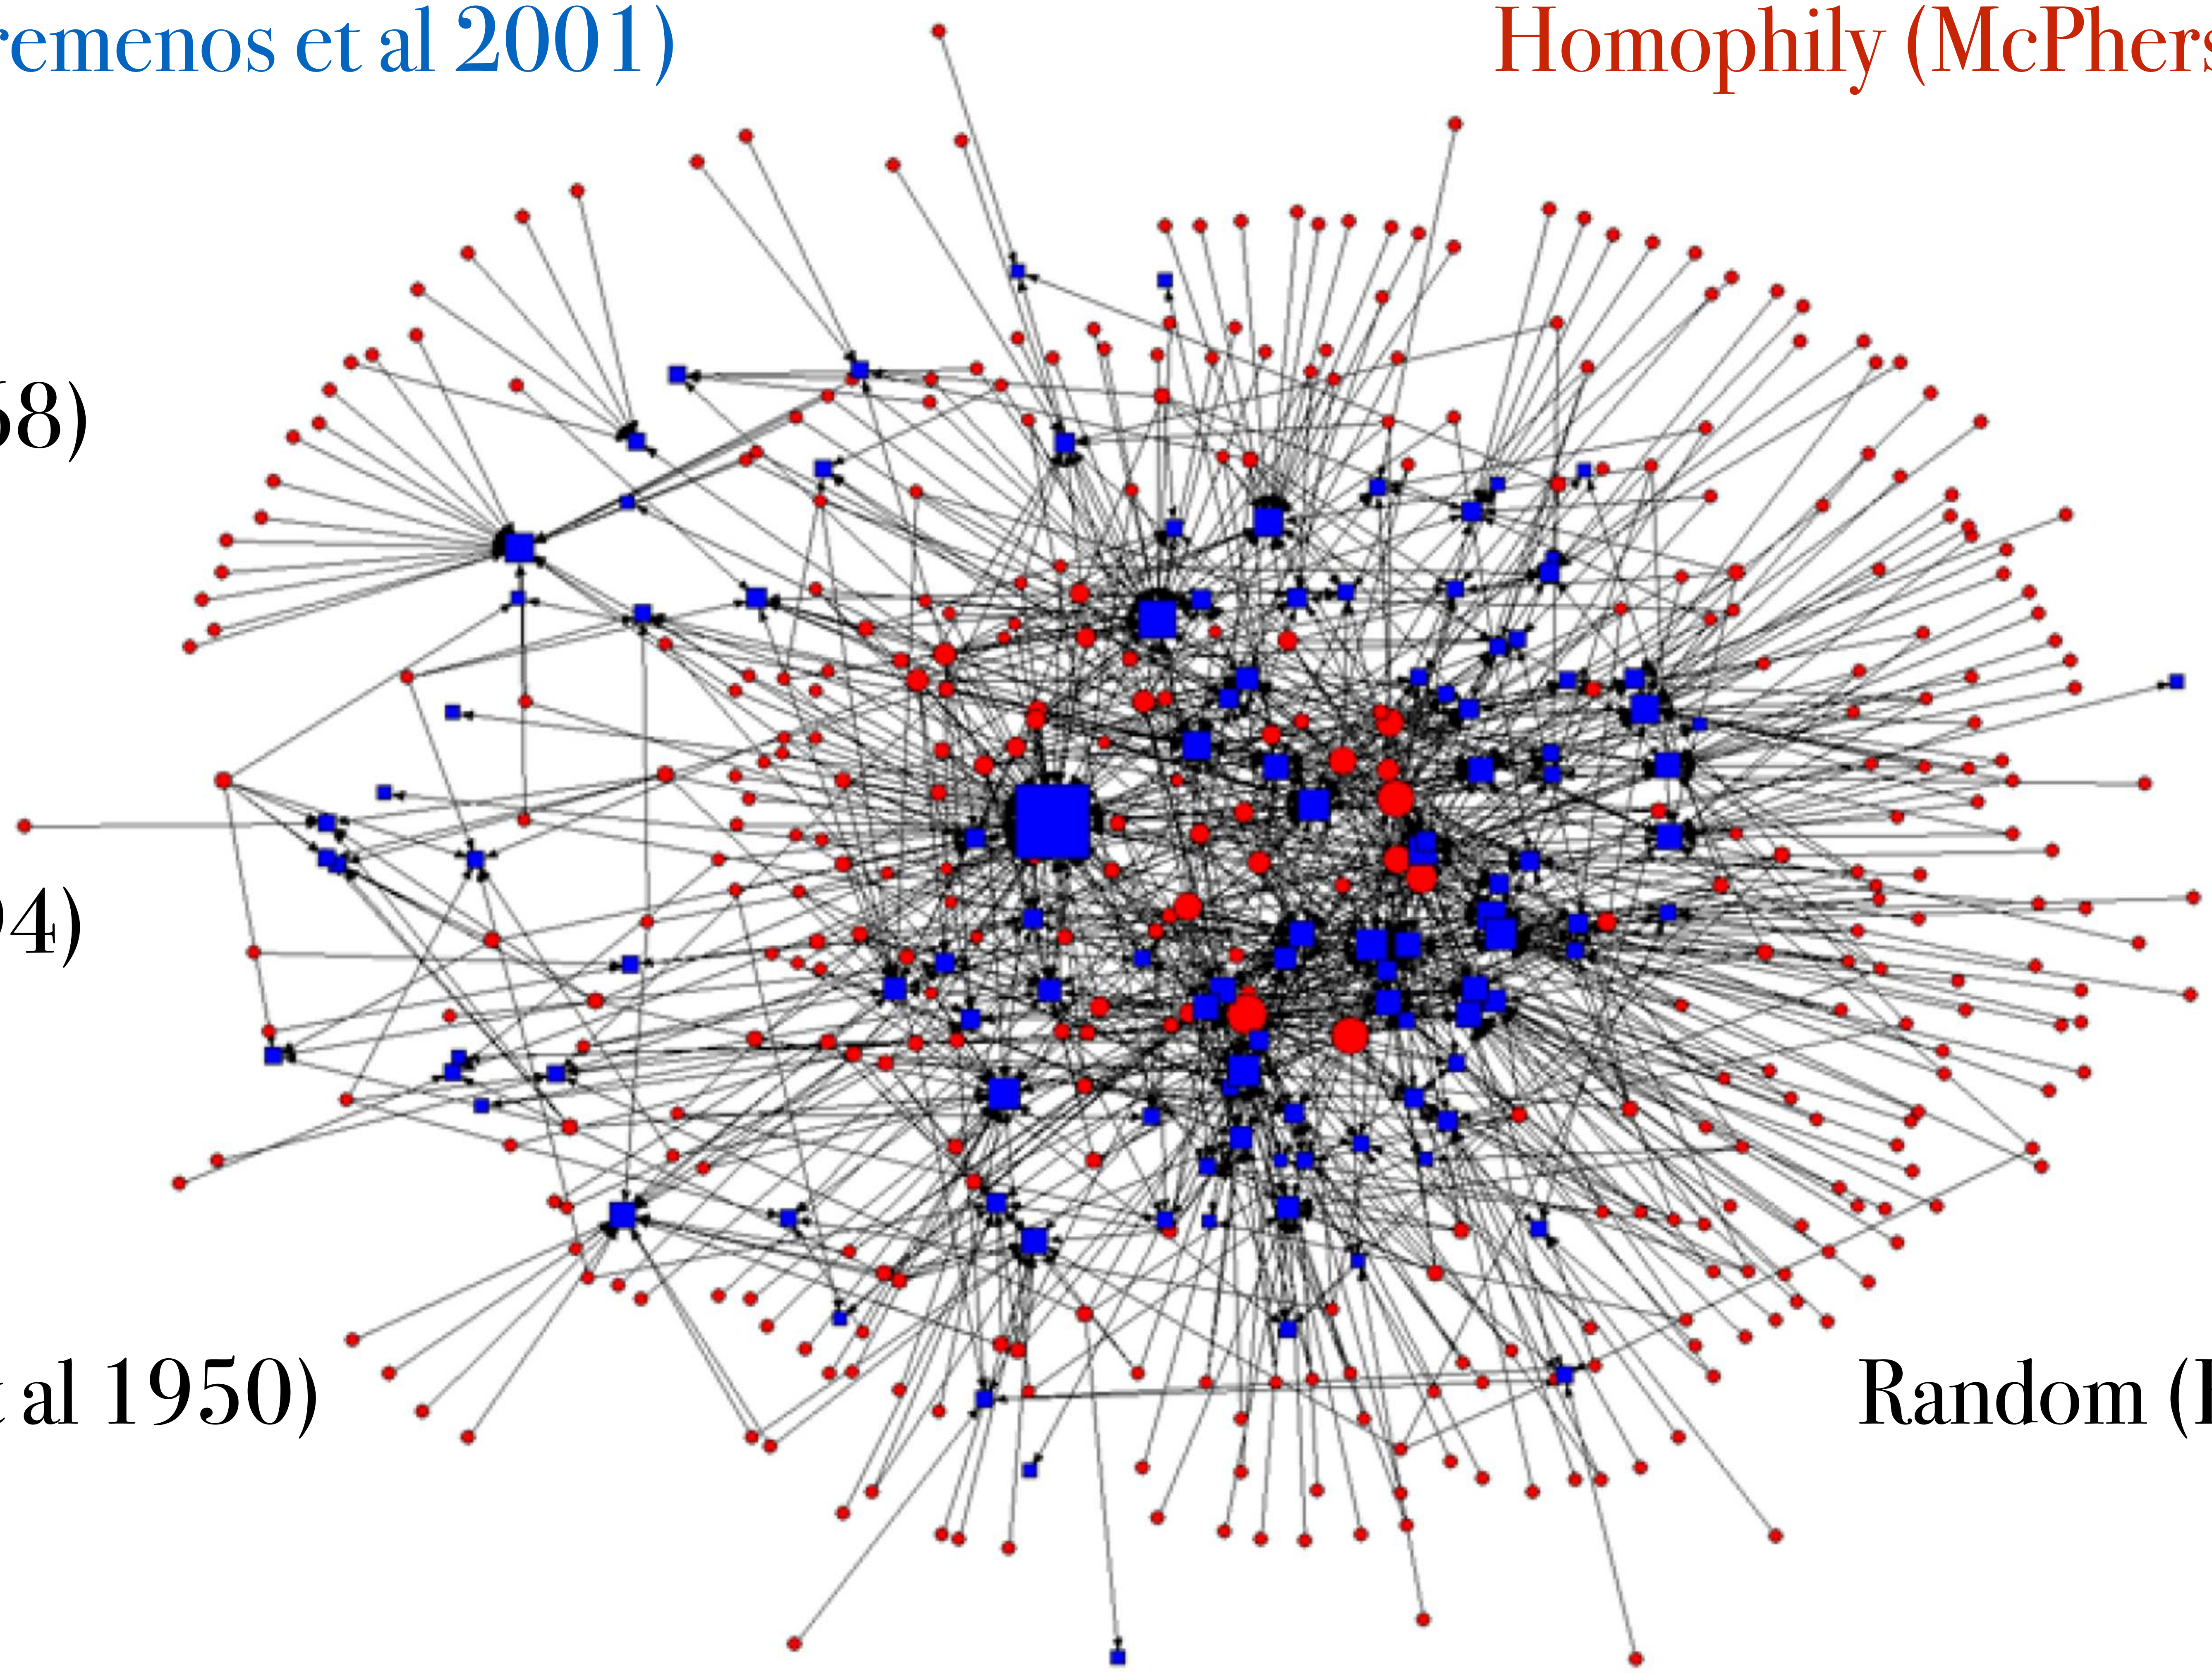

Propinquity (Festinger et al 1950)

Random (Erdős and Rényi 1960)

Could it be a bit of all of these? Maybe some more than others?

Design (Koremenos et al 2001)

Homophily (McPherson et al 2001)

Popularity (Merton 1968)

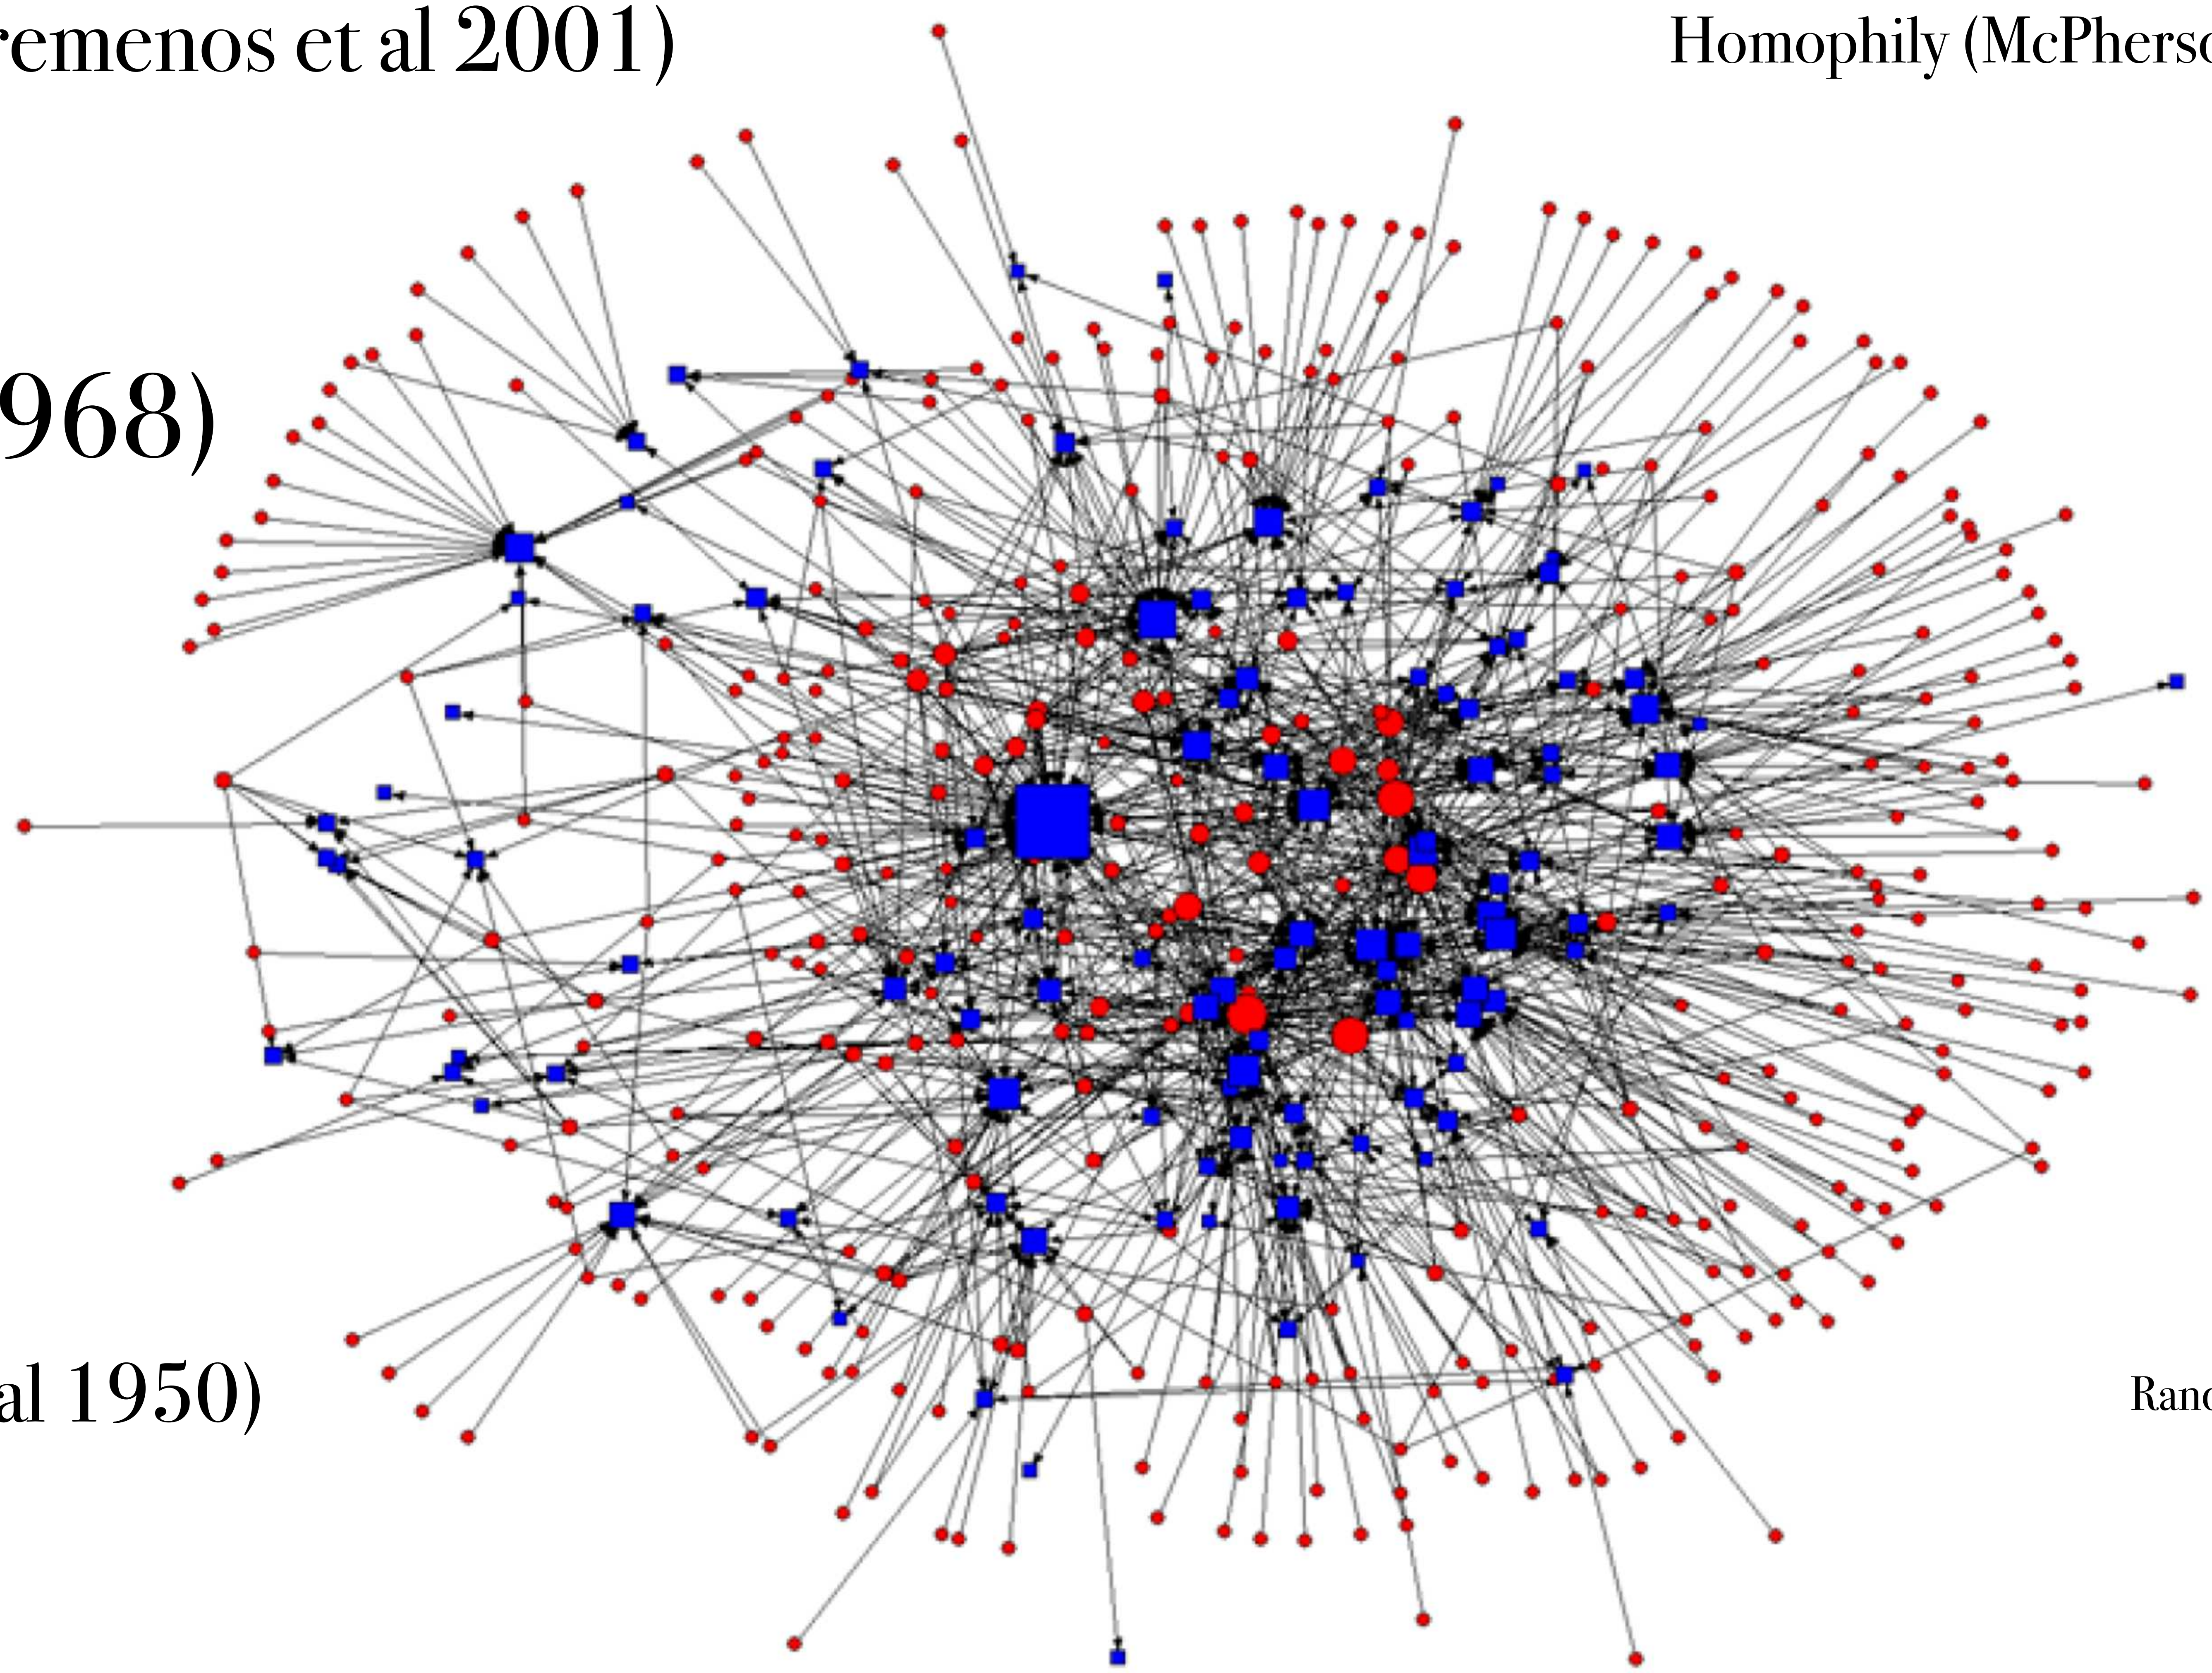

Transitivity (Simmel 1902)

Structural holes (Burt 1994)

Propinquity (Festinger et al 1950)

Random (Erdős and Rényi 1960)

Strategy…

- 1. Get out all your **ingredients** (Effects)
- 2. Use special **oven** (ERGM)
- 3. Test look and **taste** (Diagnostics)

ERGM

Ingredients

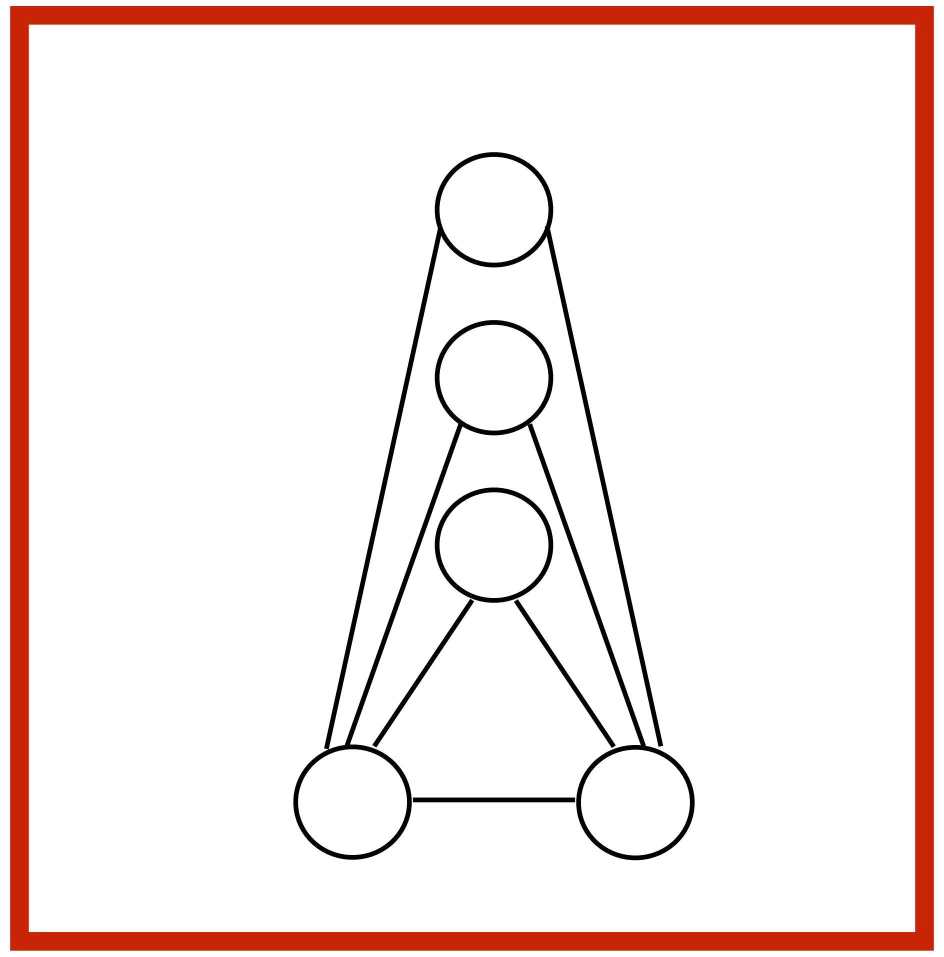

Explore various
effects available

Oven

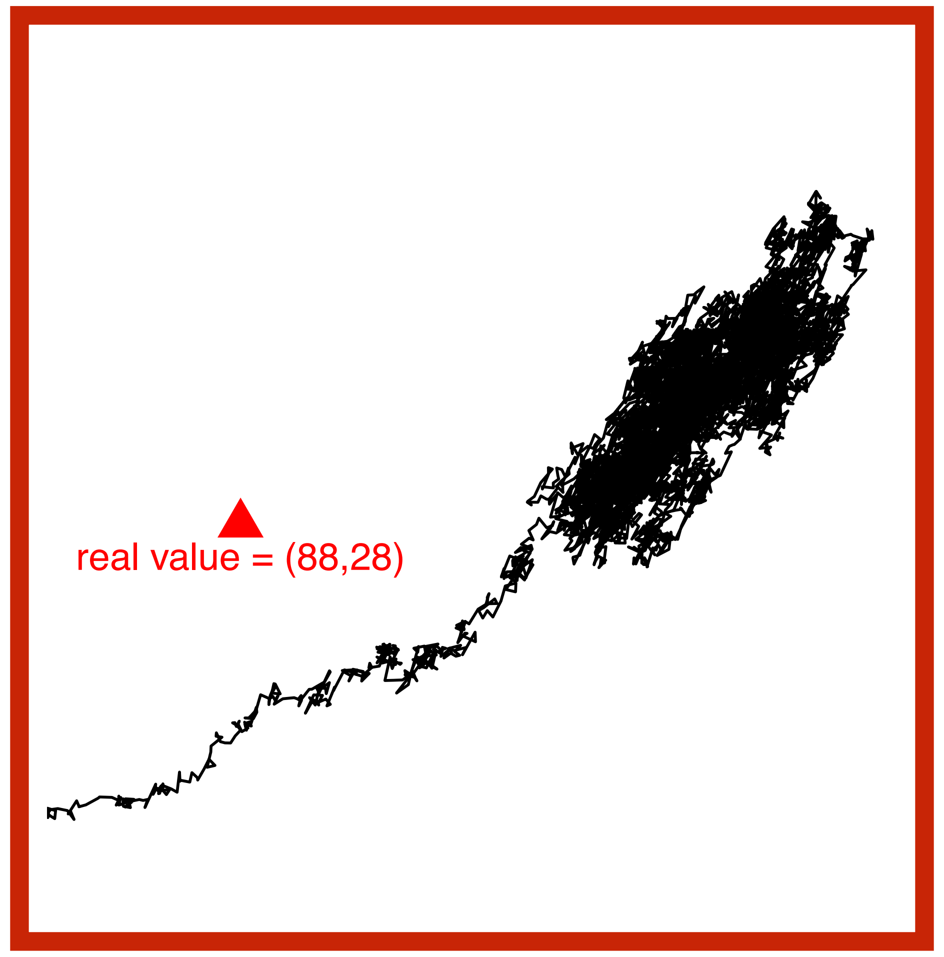

Understand model
intuition and estimation

Taste Test

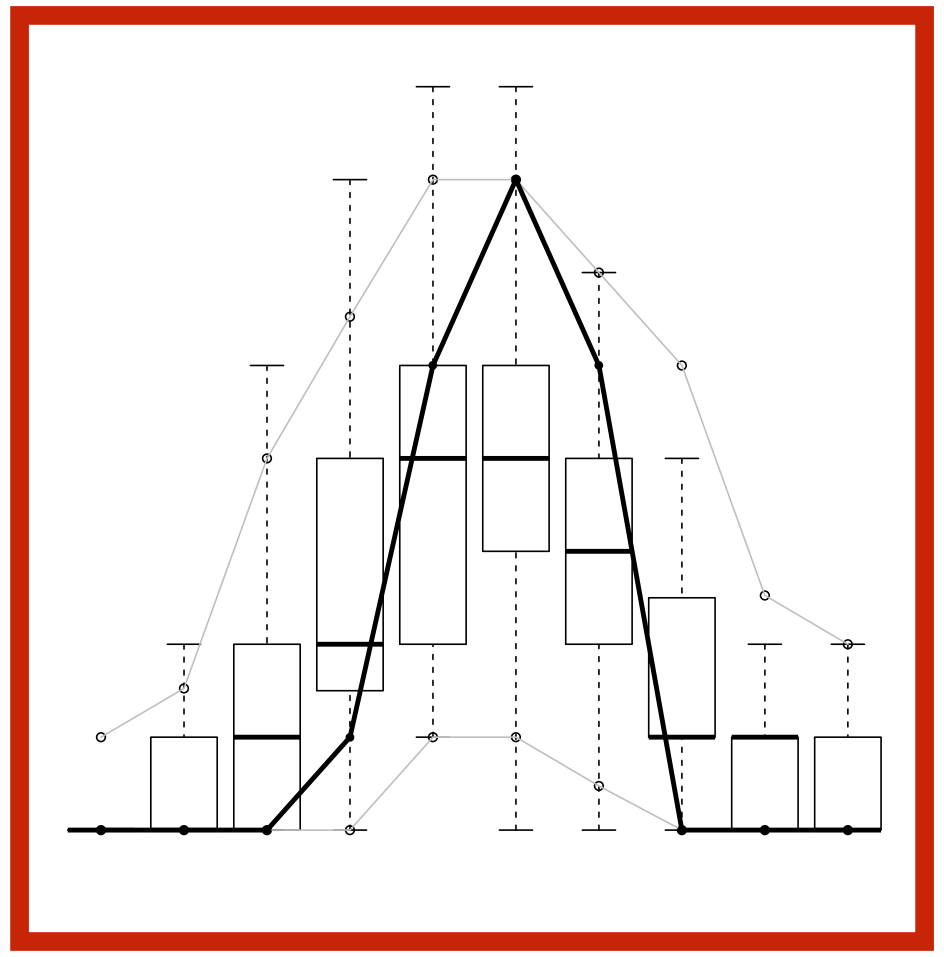

Recognise when a model
converges and fits

ERGM

Effects

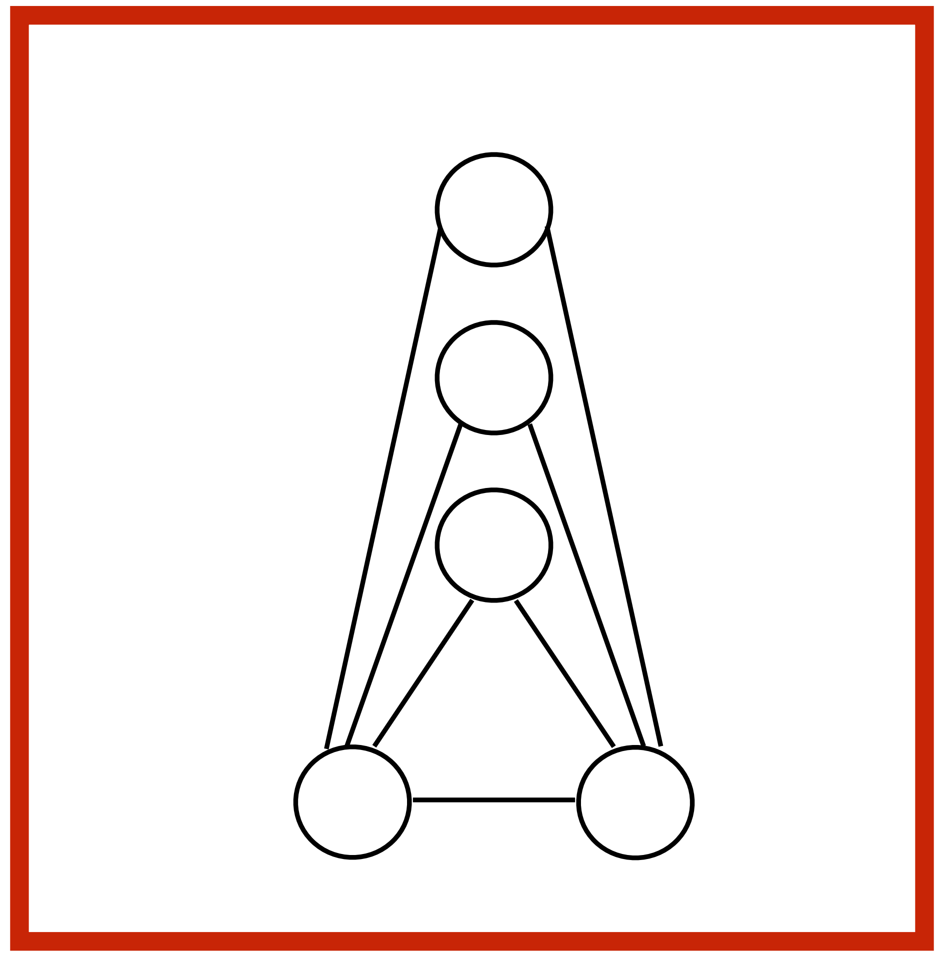

Model

Diagnostics

How & Why Ties Form?

- Randomness
- Covariates (nodal attributes)
  - Monadic
    - Sender
    - Receiver
  - Dyadic effects
    - Matching
    - Similarity
- Exogenous contexts
  - Spatial factors
  - Other networks

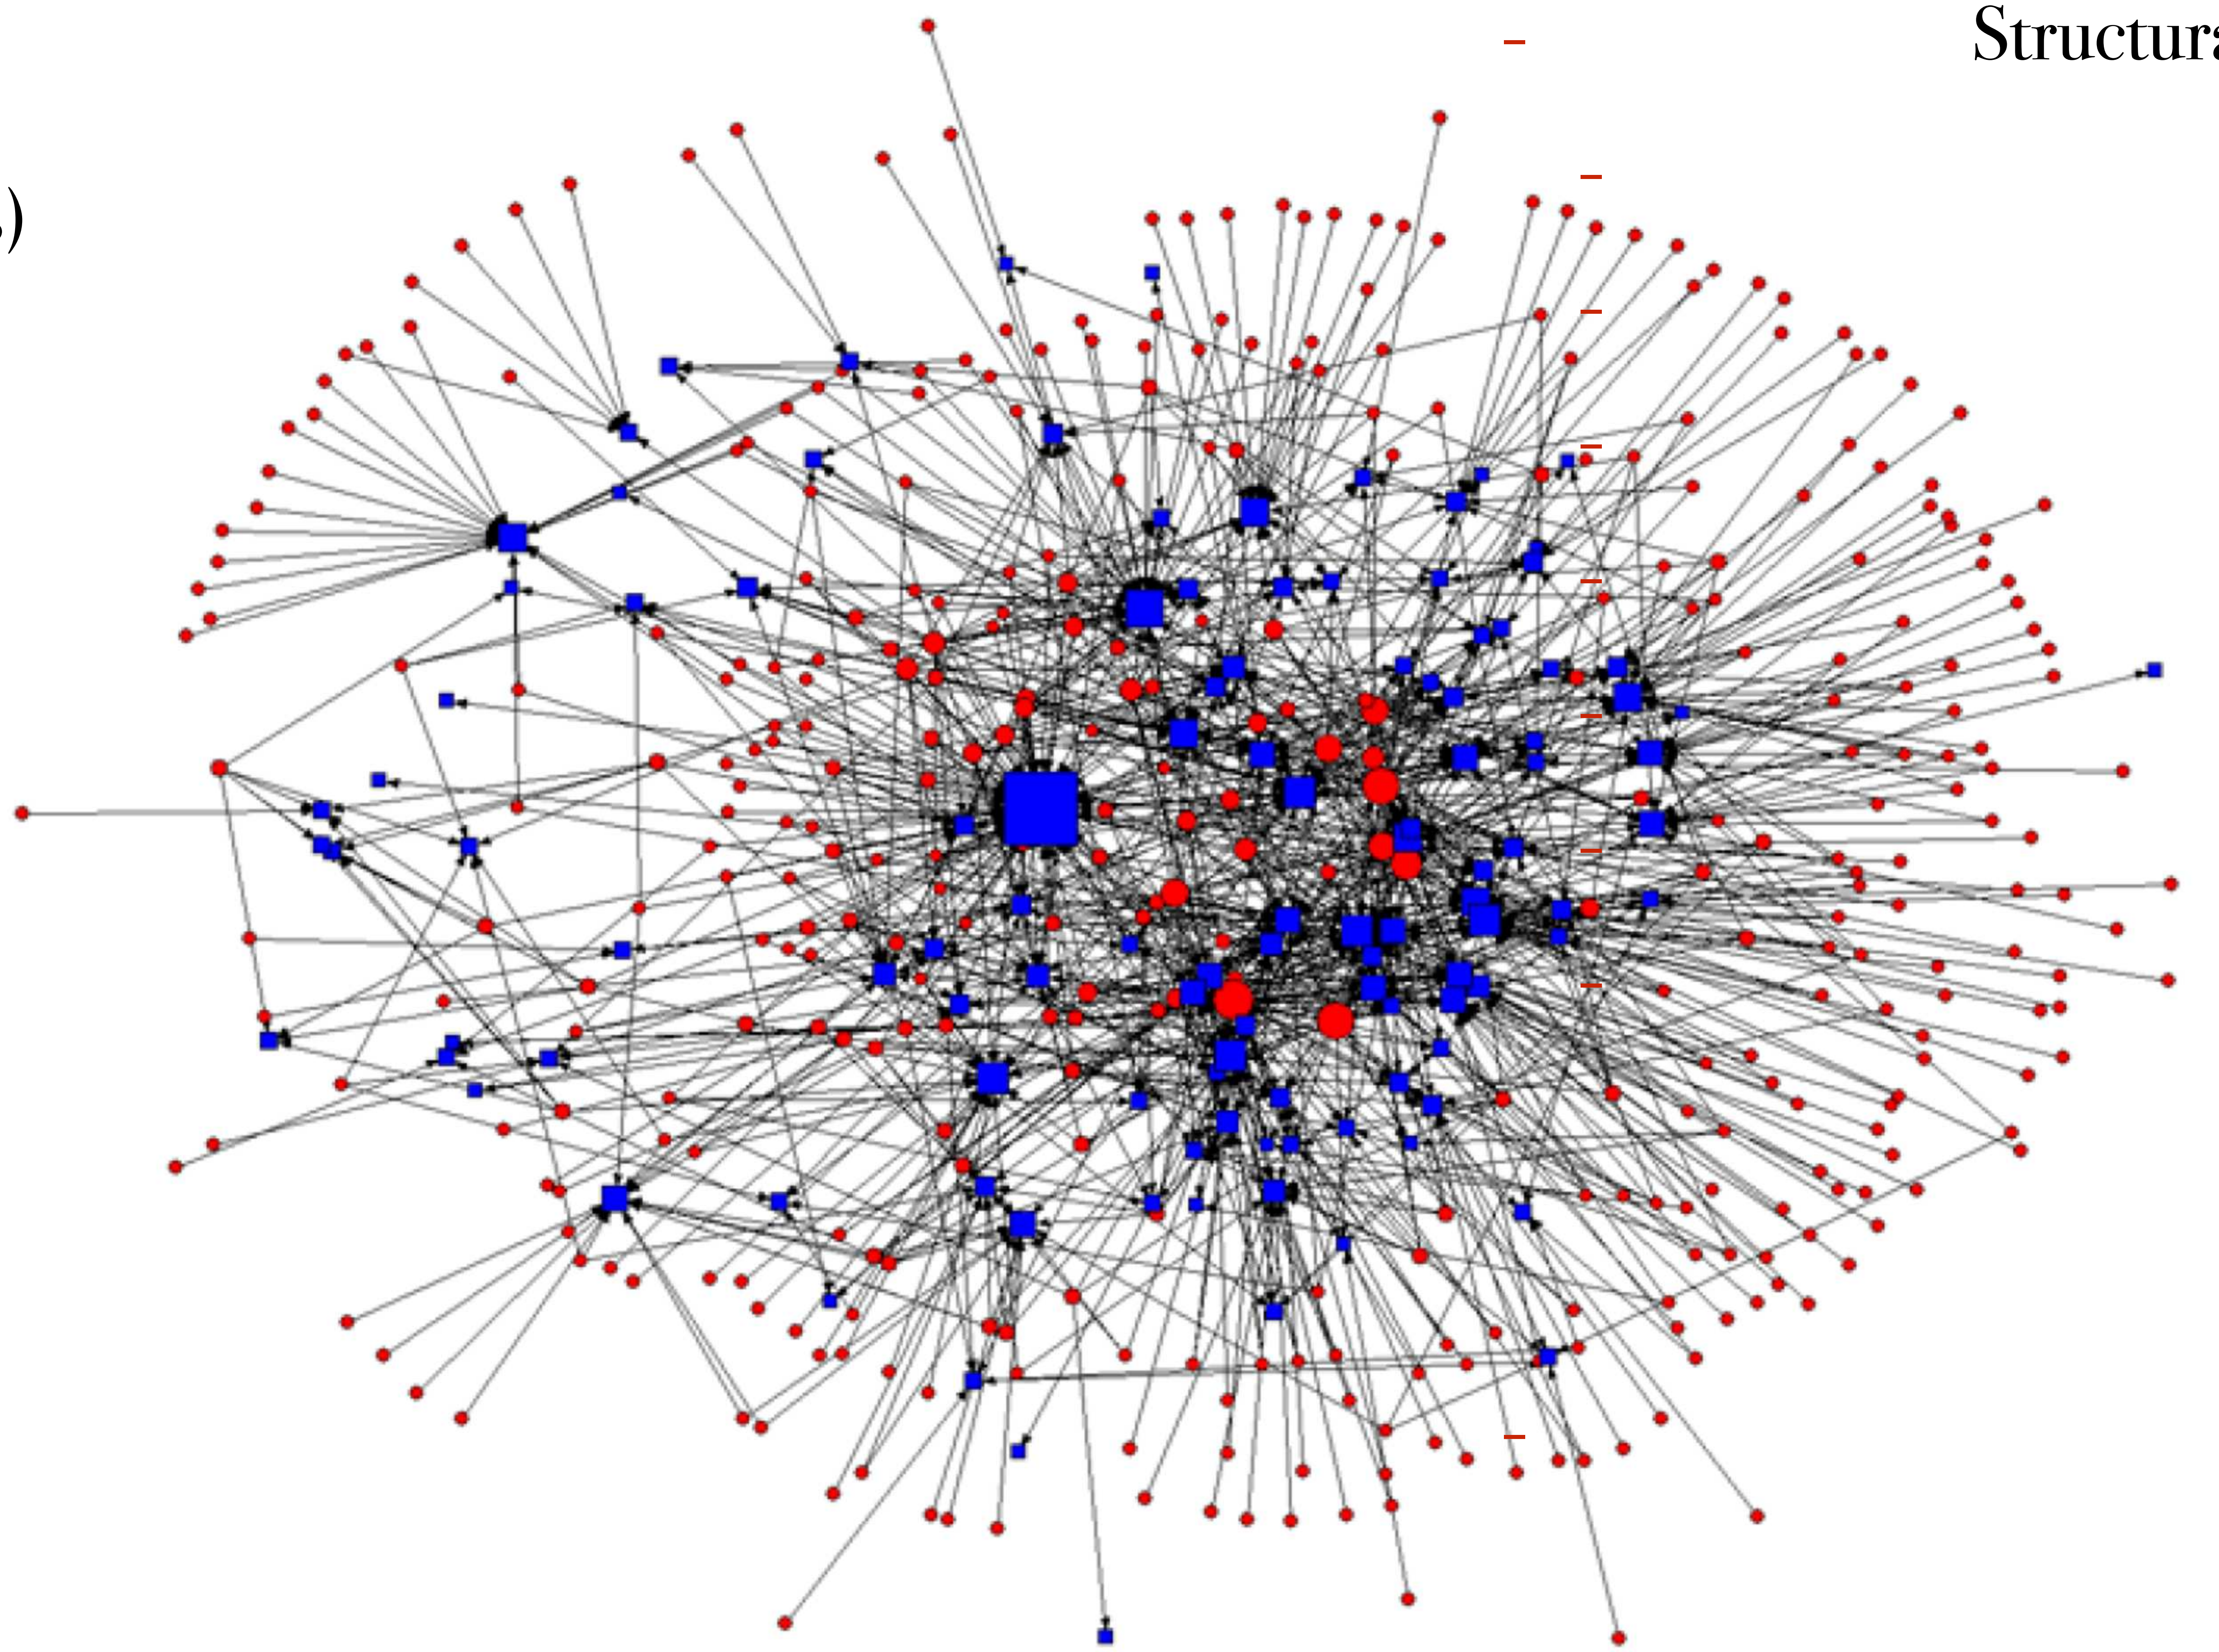

- Structural (Network self-organization)
  - Activity
  - Popularity
  - Reciprocity
  - Transitivity
  - Three-Cycles
  - Four-Cycles
  - Brokerage

Robins & Lusher 2012: 23

Levels of Dependence

- Independence
  - *Logistic regression*
  - Density, attributes

If we believe that particular attributes are responsible for ties, then include counts of

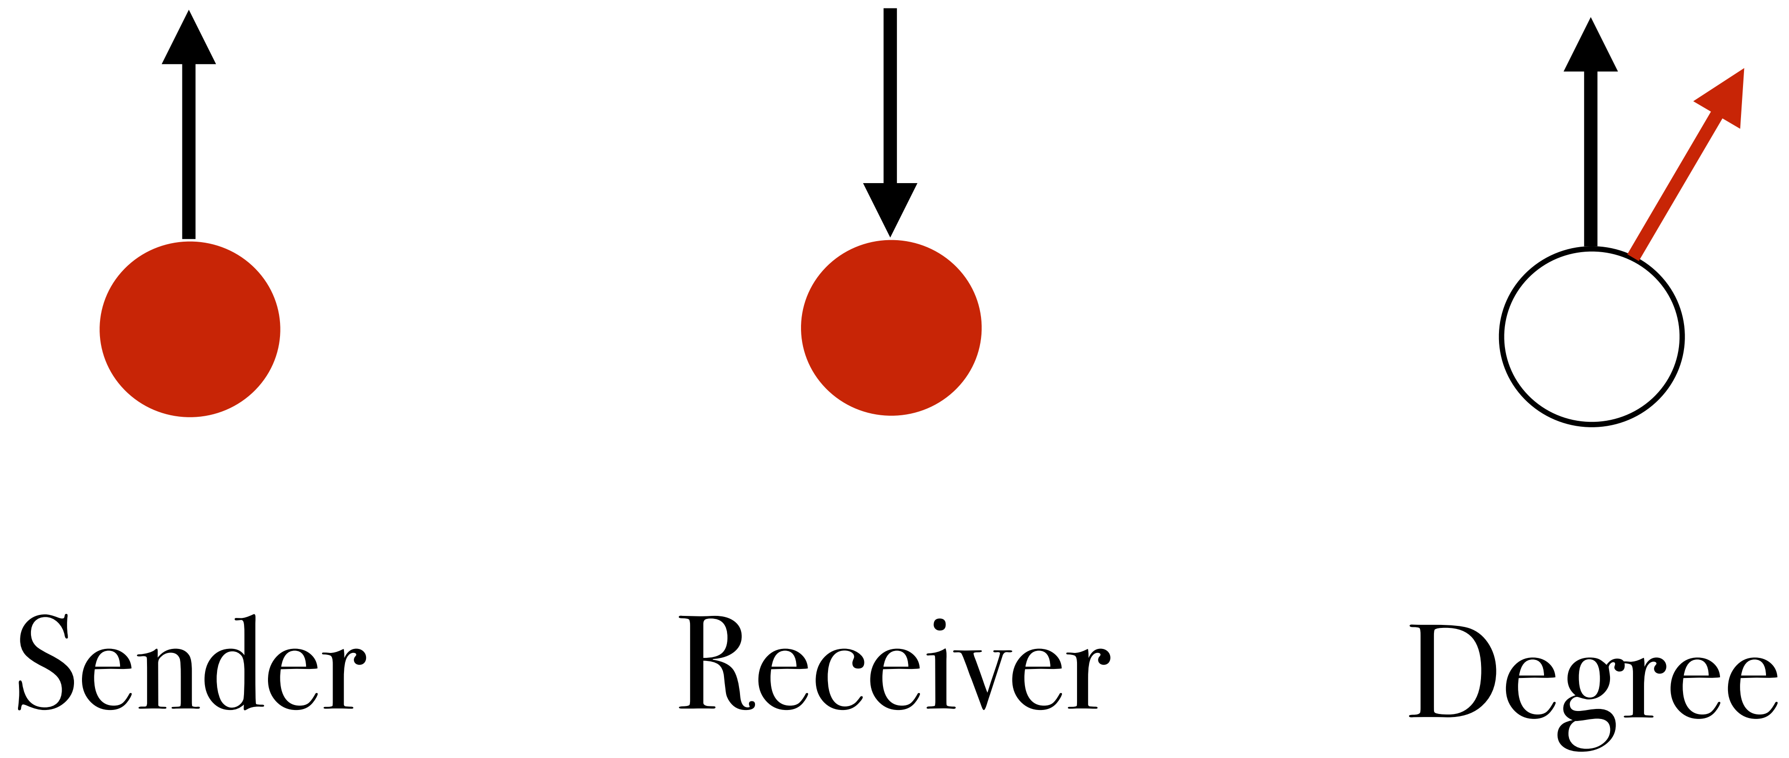

Levels of Dependence

- Independence
  - *Logistic regression*
  - Density, attributes
- Dyad-independence
  - *p1*
  - Reciprocity, homophily

If we believe that reciprocity or homophily
are responsible for ties,
then include counts of

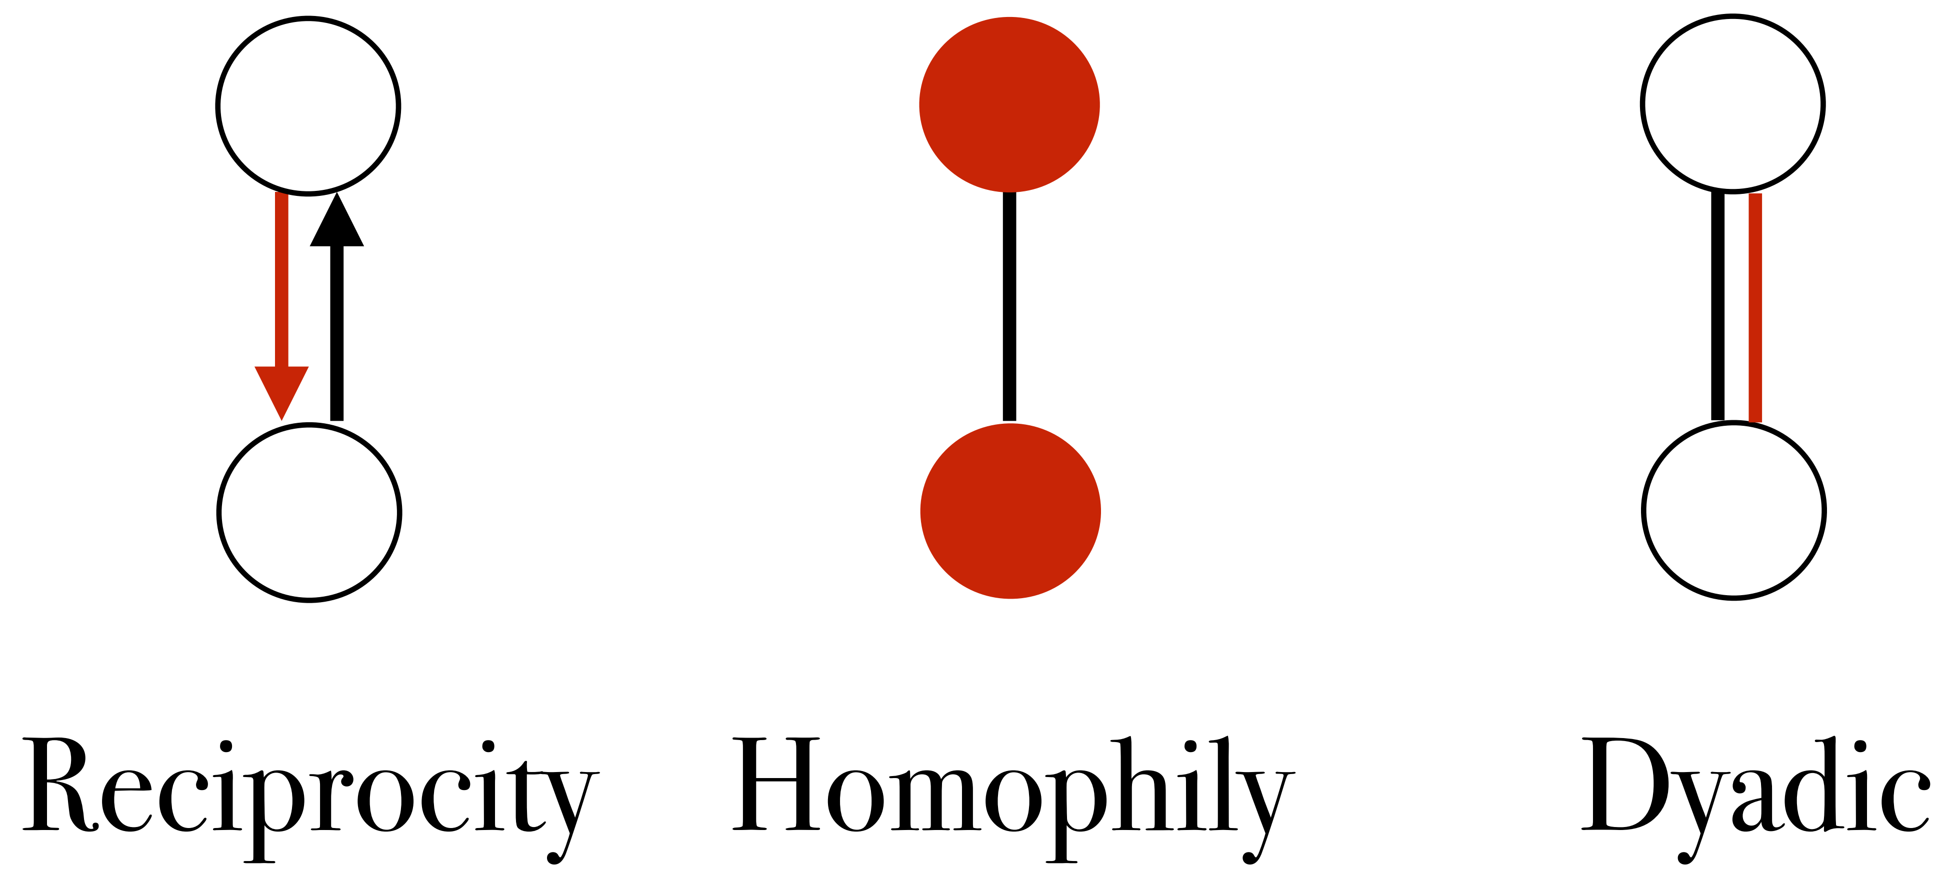

Holland & Leinhardt 1981, Fienberg & Wasserman 1979, 1981

Levels of Dependence

- Independence
  - *Logistic regression*
  - Density, attributes
- Dyad-independence
  - *p1*
  - Reciprocity, homophily
- Markov-dependence
  - *p*/ERGM*
  - Transitivity, popularity

may depend on
one step removed…

If we believe that popularity or transitivity
are responsible for ties,
then include counts of

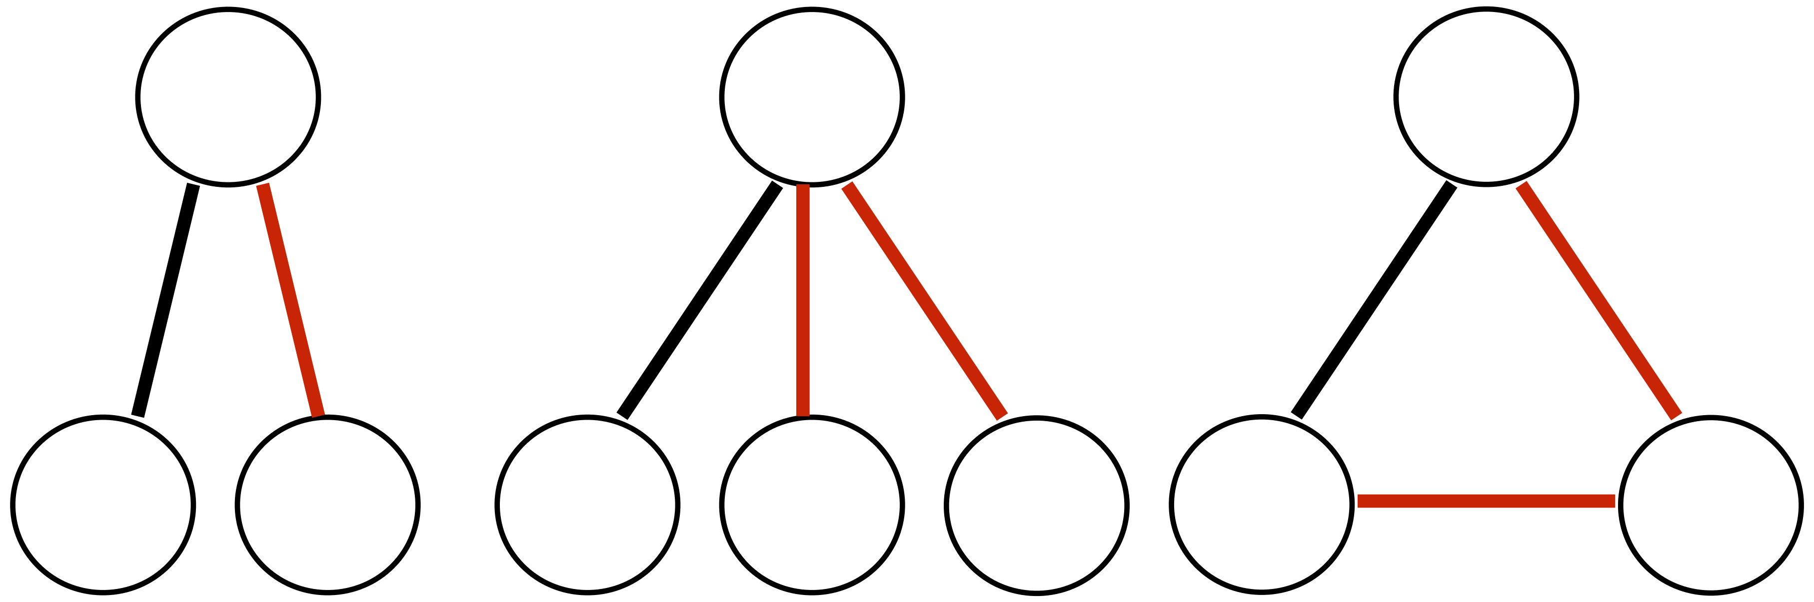

k-Stars

Triangles

Levels of Dependence

- Independence
  - *Logistic regression*
  - Density, attributes
- Dyad-independence
  - *p1*
  - Reciprocity, homophily
- Markov-dependence
  - *p*/ERGM*
  - Transitivity, popularity
- Social circuit dependence
  - New specifications
  - Geometrically weighted edgewise shared partners (GWESP), four-cycles

If we believe that ties are coordinated or that clustering aggregates, then include counts of

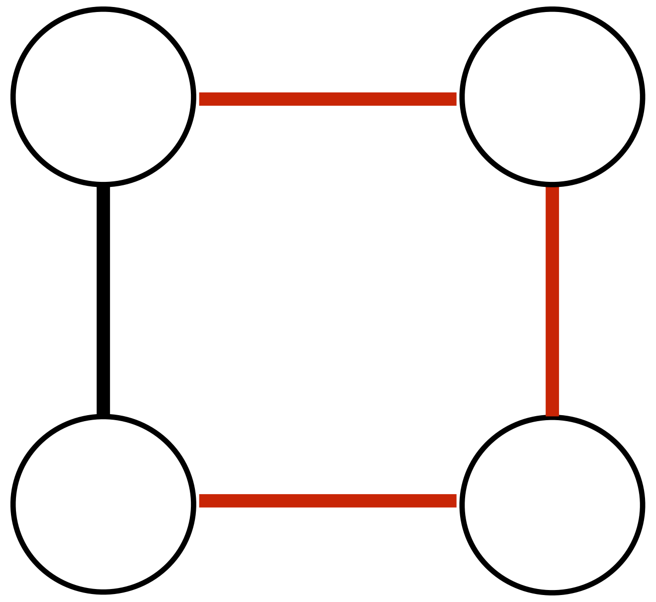

Four Cycles

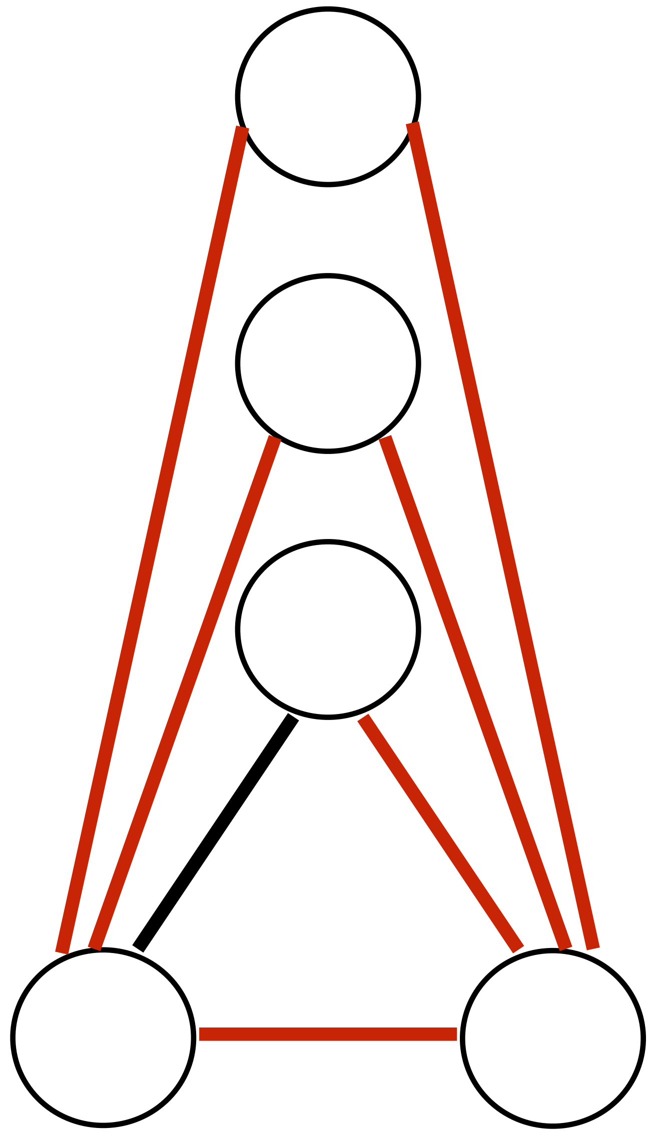

GWESP

Snijders et al 2006, Hunter & Handcock 2006

Dependence assumption and corresponding model	Dependence graph D	Cliques of D
Independence $X_{ij} \perp\!\!\!\perp X_{hk},\ \forall\ i,j,h,k \in \mathcal{N}$ Bernoulli random graph models	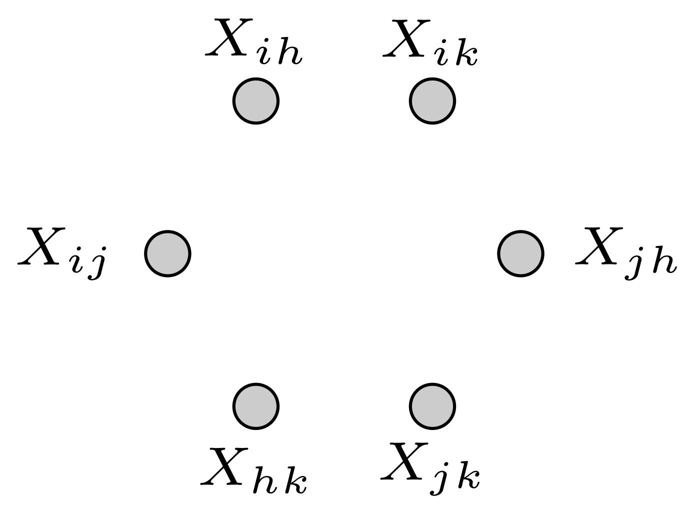	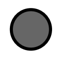
Dyadic Dependence $X_{ij} \not\perp\!\!\!\perp X_{hk},\ \forall\ \{i,j\} = \{h,k\}$ Dyadic dependence models	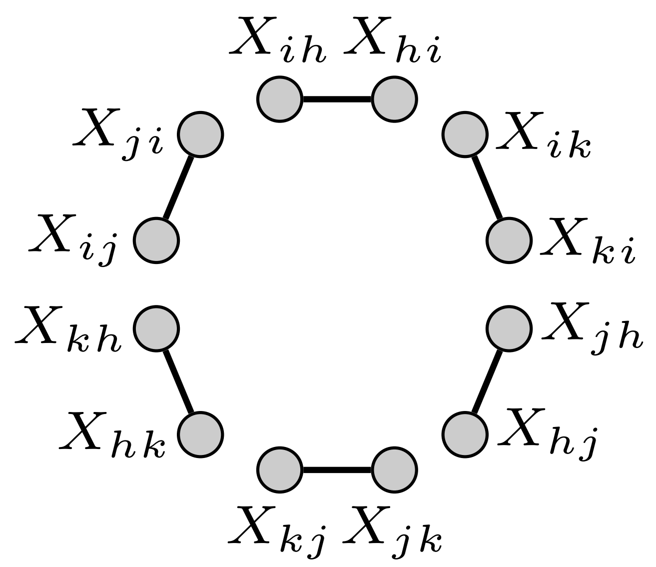	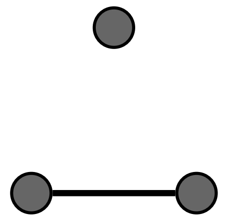
Markov Dependence $X_{ij} \not\perp\!\!\!\perp X_{hk}\ if\ \{i,j\} \cap \{h,k\} \neq \emptyset$ Markov graphs	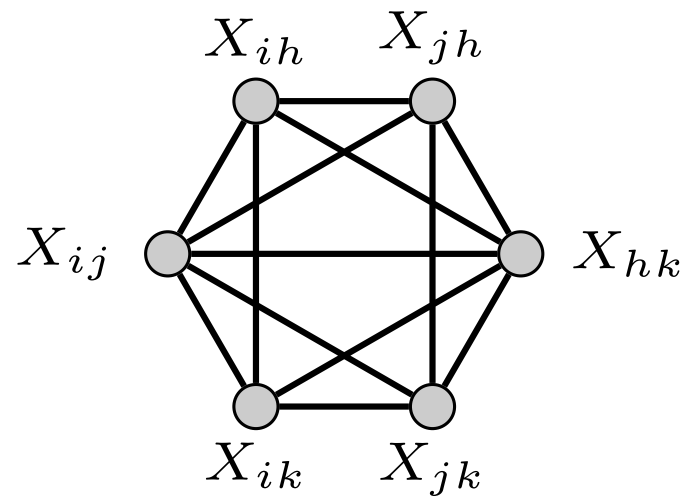	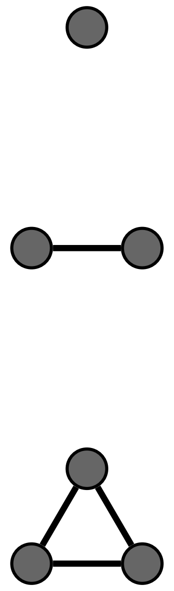
Partial conditional dependence E.g., social circuit dependence $X_{ij} \not\perp\!\!\!\perp X_{hk}\ if\ X_{ih} = X_{jk} = 1\ or$ $X_{ik} = X_{jh} = 1$ Exponential random graph models	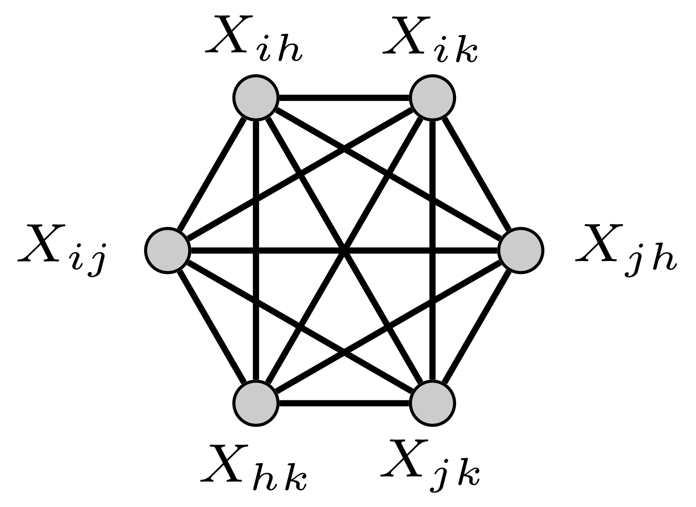	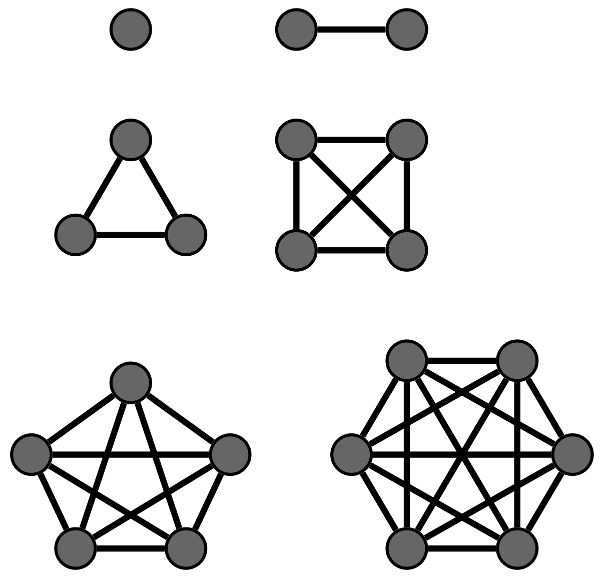

Network property	Statistics		
	Undirected	Directed	
Density	 Edges	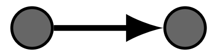 Arcs	
Reciprocity		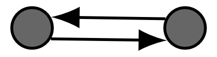 Mutual dyads	
Degree distribution	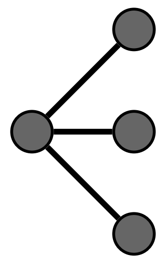 Stars	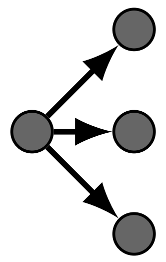 Out-stars	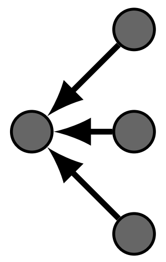 In-stars
Connectivity	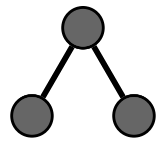 Two-paths	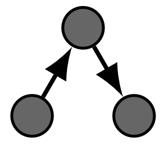 Two-paths	
Closure	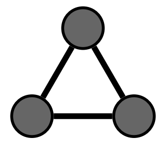 Triangles	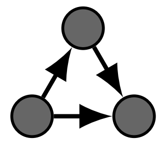 Transitive triads	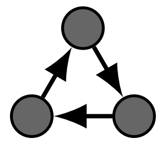 3-cycles
Clustering of triangles	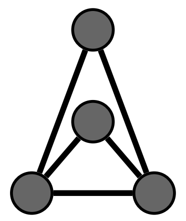 Alternating-k-triangles	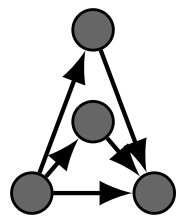 Alternating-transitive-triangles	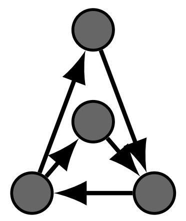 Alternating-3-cycles
Clustering of 2-paths	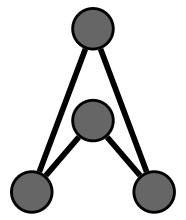 Alternating-k-paths	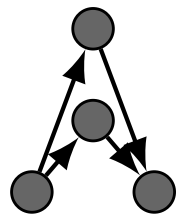 Alternating-k-twopaths	

Amati, Lomi, and Mira. 2018. “Social Network Modeling.” *Annual Review of Statistics and Its Application* 5 (1): 343–69.

Relevant statistics in the water management case

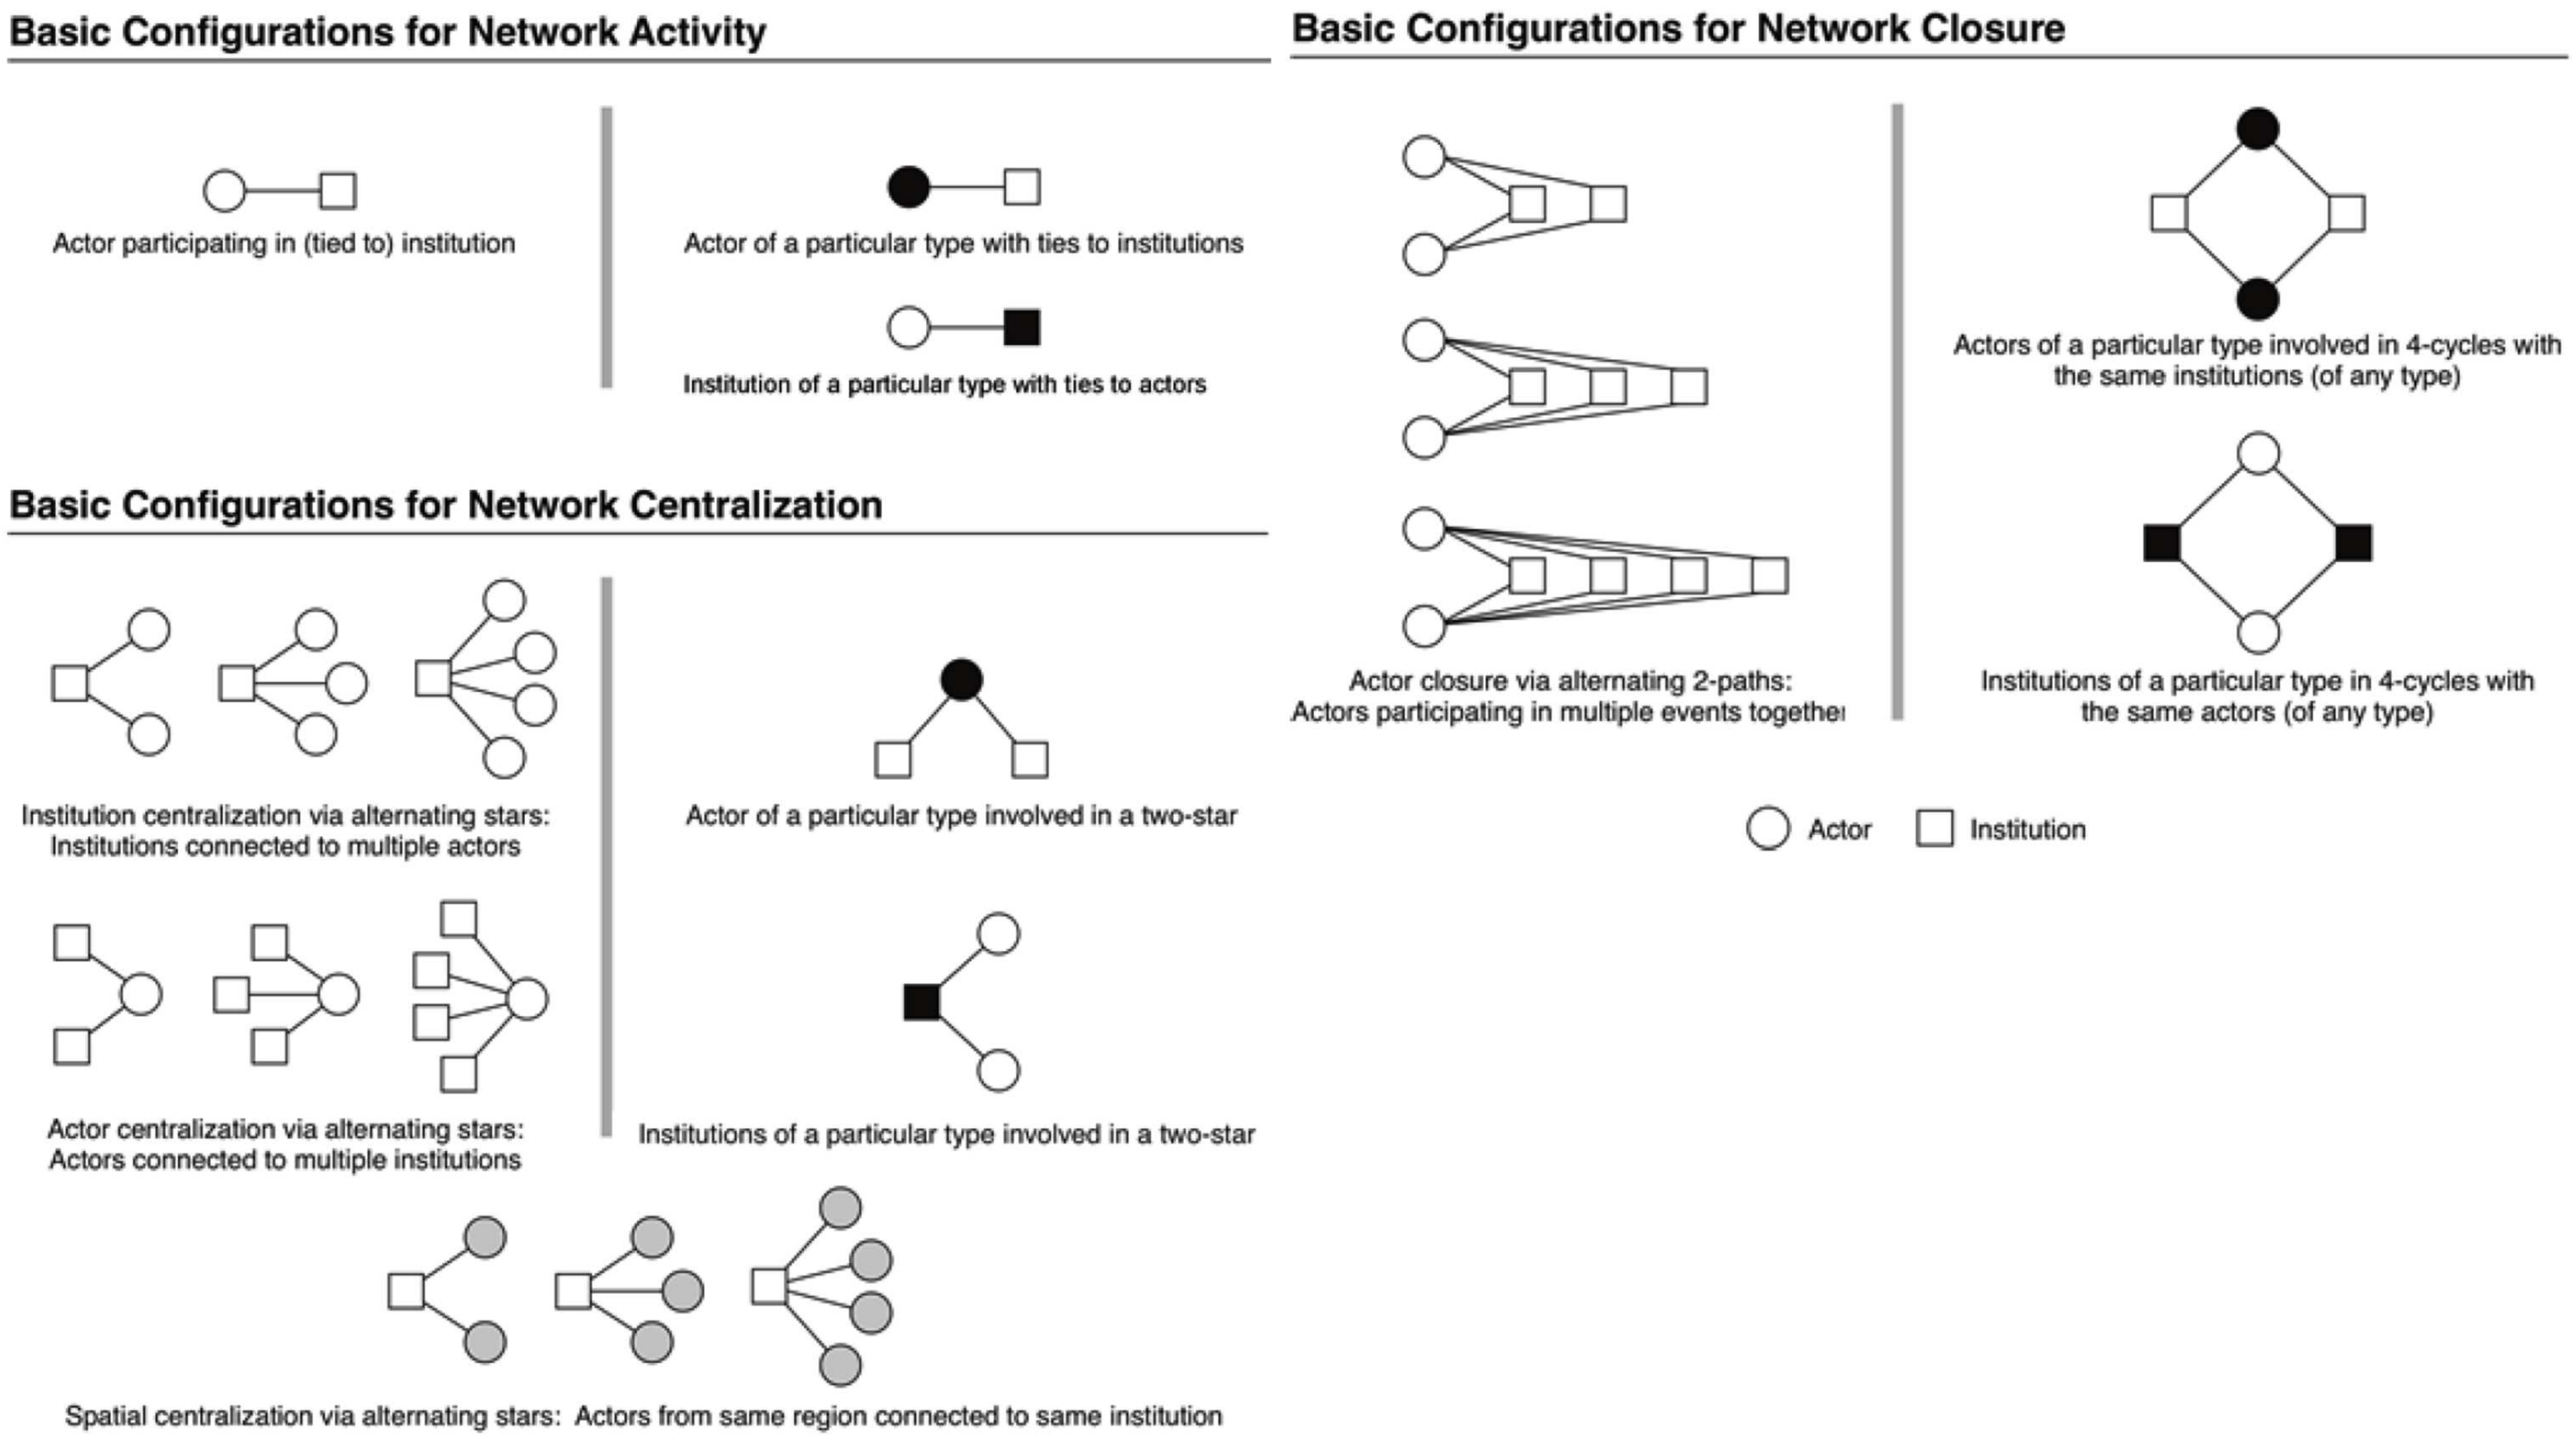

Notes on ingredients

- Start with most basic effects (e.g. density)
- Add effects from increasing levels of dependence (e.g. Markov, social circuit)
- Always include more fundamental forms from within more complex configurations (e.g. monadic before homophily, degree before closure)
- Often useful to contrast structural-only models with with covariates-added models

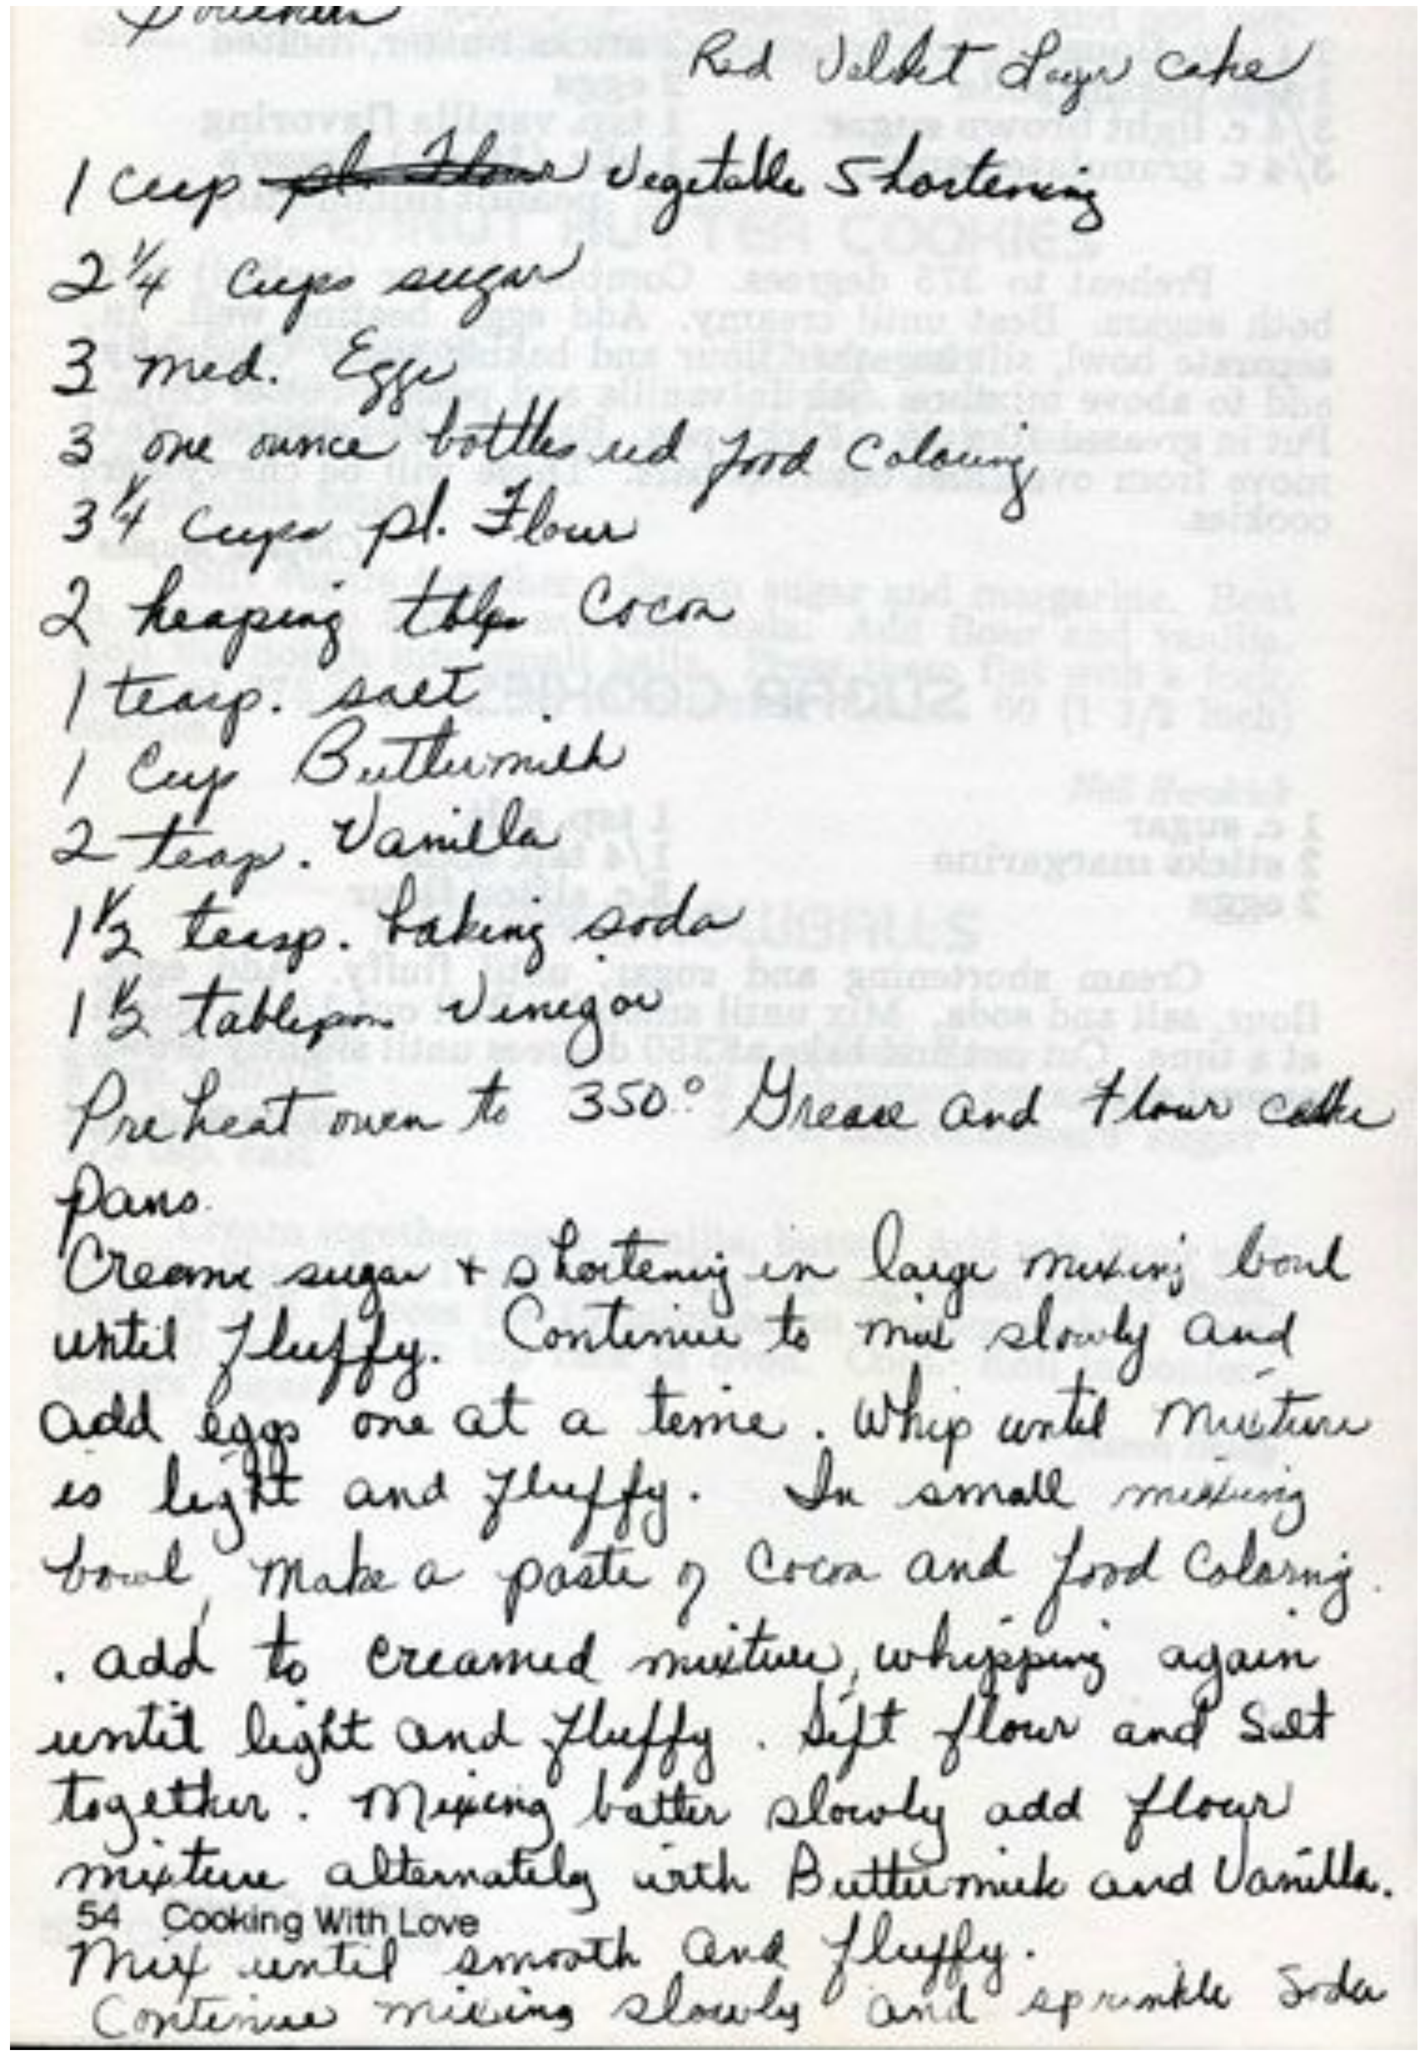

Red Velvet Layer cake

1 cup ~~pl. flour~~ Vegetable Shortening
2 1/4 Cups sugar
3 med. Eggs
3 one ounce bottles red food Coloring
3 1/4 Cups pl. Flour
2 heaping tbls Cocoa
1 teasp. salt
1 Cup Buttermilk
2 teasp. Vanilla
1 1/2 teasp. baking soda
1 1/2 tablesp. Vinegar

Preheat oven to 350° Grease and flour cake pans.
Cream sugar + shortening in large mixing bowl until fluffy. Continue to mix slowly and add eggs one at a time. Whip until mixture is light and fluffy. In small mixing bowl make a paste of Cocoa and food Coloring. add to creamed mixture, whipping again until light and fluffy. Sift flour and salt together. Mixing batter slowly add flour mixture alternately with Buttermilk and Vanilla. Mix until smooth and fluffy. Continue mixing slowly and sprinkle Soda

54 Cooking With Love

ERGM

Effects

Model

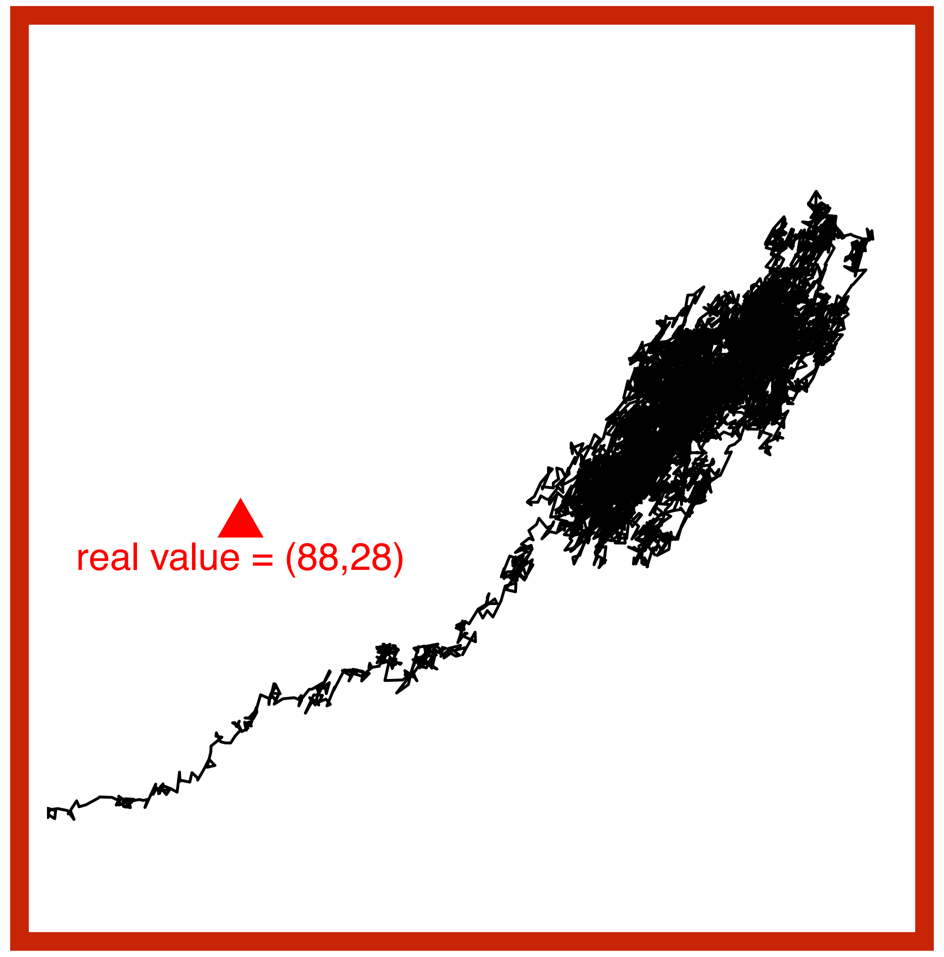

Diagnostics

Nota bene..

- We'll cover a lot of ground here
  - Some vocabulary may be unfamiliar
- Don't worry if you don't understand everything
  - Focus on getting the big picture
  - `statnet` puts a lot of this behind the curtain, so you often don't have to deal with it (except the details matter when the modelling breaks down)
- So: don't be afraid to ask questions! (today, tutorial, office hours, Moodle, consultancies)

What are ERGMs?

- ERGMs (pronounced *örgums* - this is important)
  - “are statistical models for network structure, permitting inferences about how network ties are patterned” (Robins & Lusher 2012)

- Since the random graphs in our model form an exponential family, we call the model an *exponential (family) random graph model* (ERGM… EFRGM would be *too* much of a tongue-twister!)

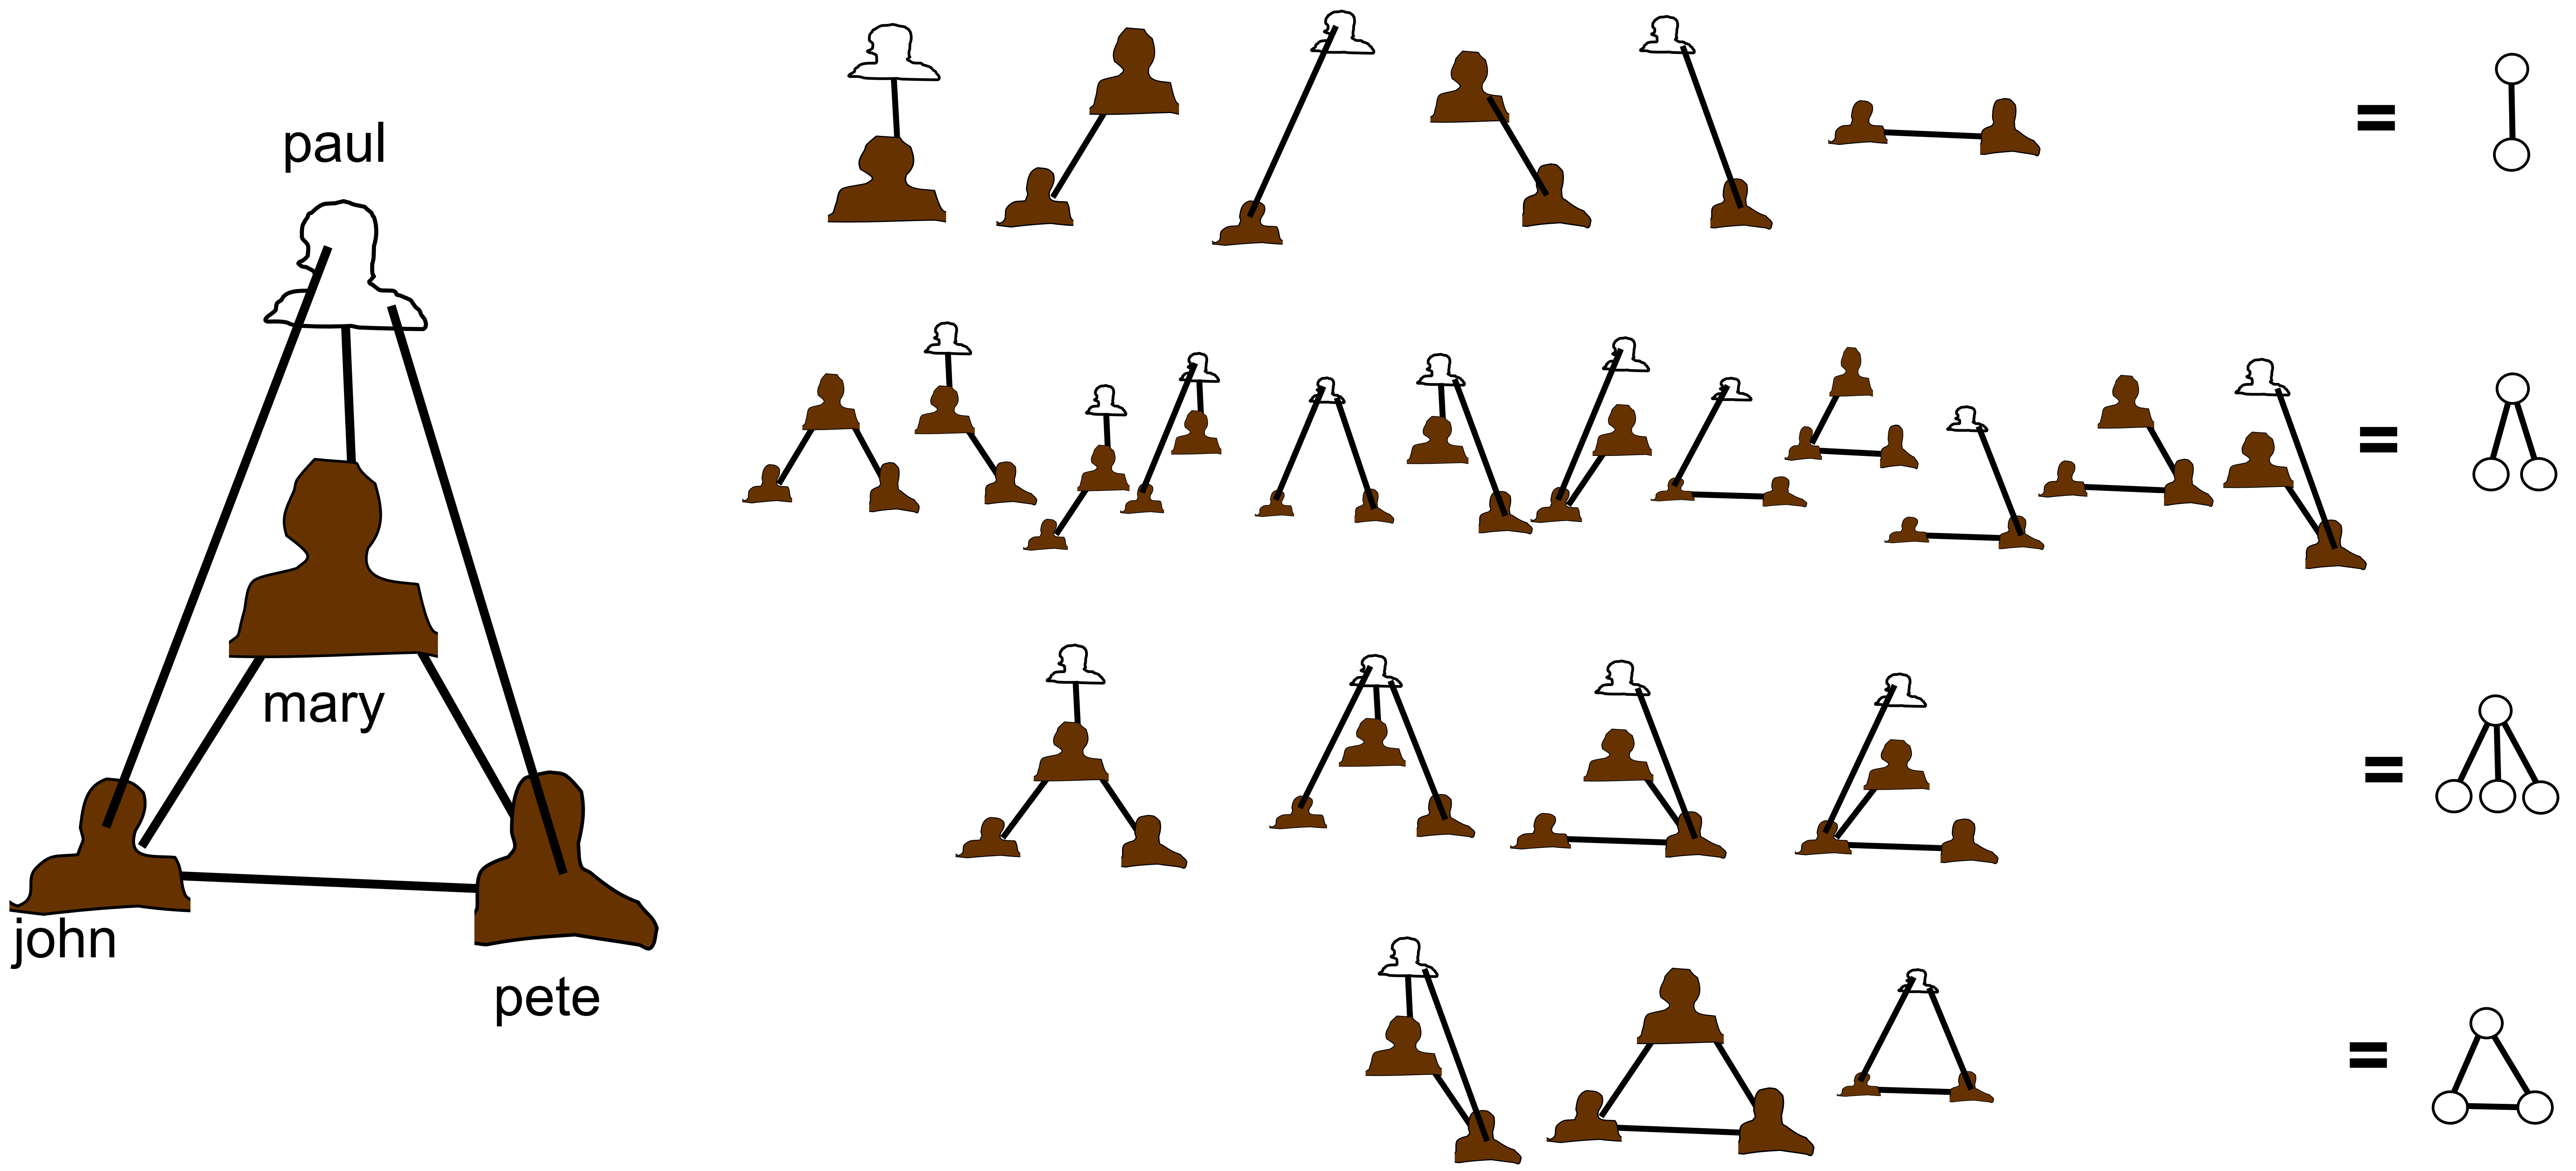

- Aim to explain observed network ties or structure as function of "ingredients" you put in to it
  - these ingredients can be exogenous (monadic and dyadic covariates), or
  - endogenous (structural effects, like activity/popularity or transitivity)

- Once you have a model of how much of each "ingredient" to put in, we can use this model to:
  - predict ties (e.g. how likely is it that Paul and Ringo have a tie?)
  - simulate networks (e.g. how would the cake look if I added more butter?)

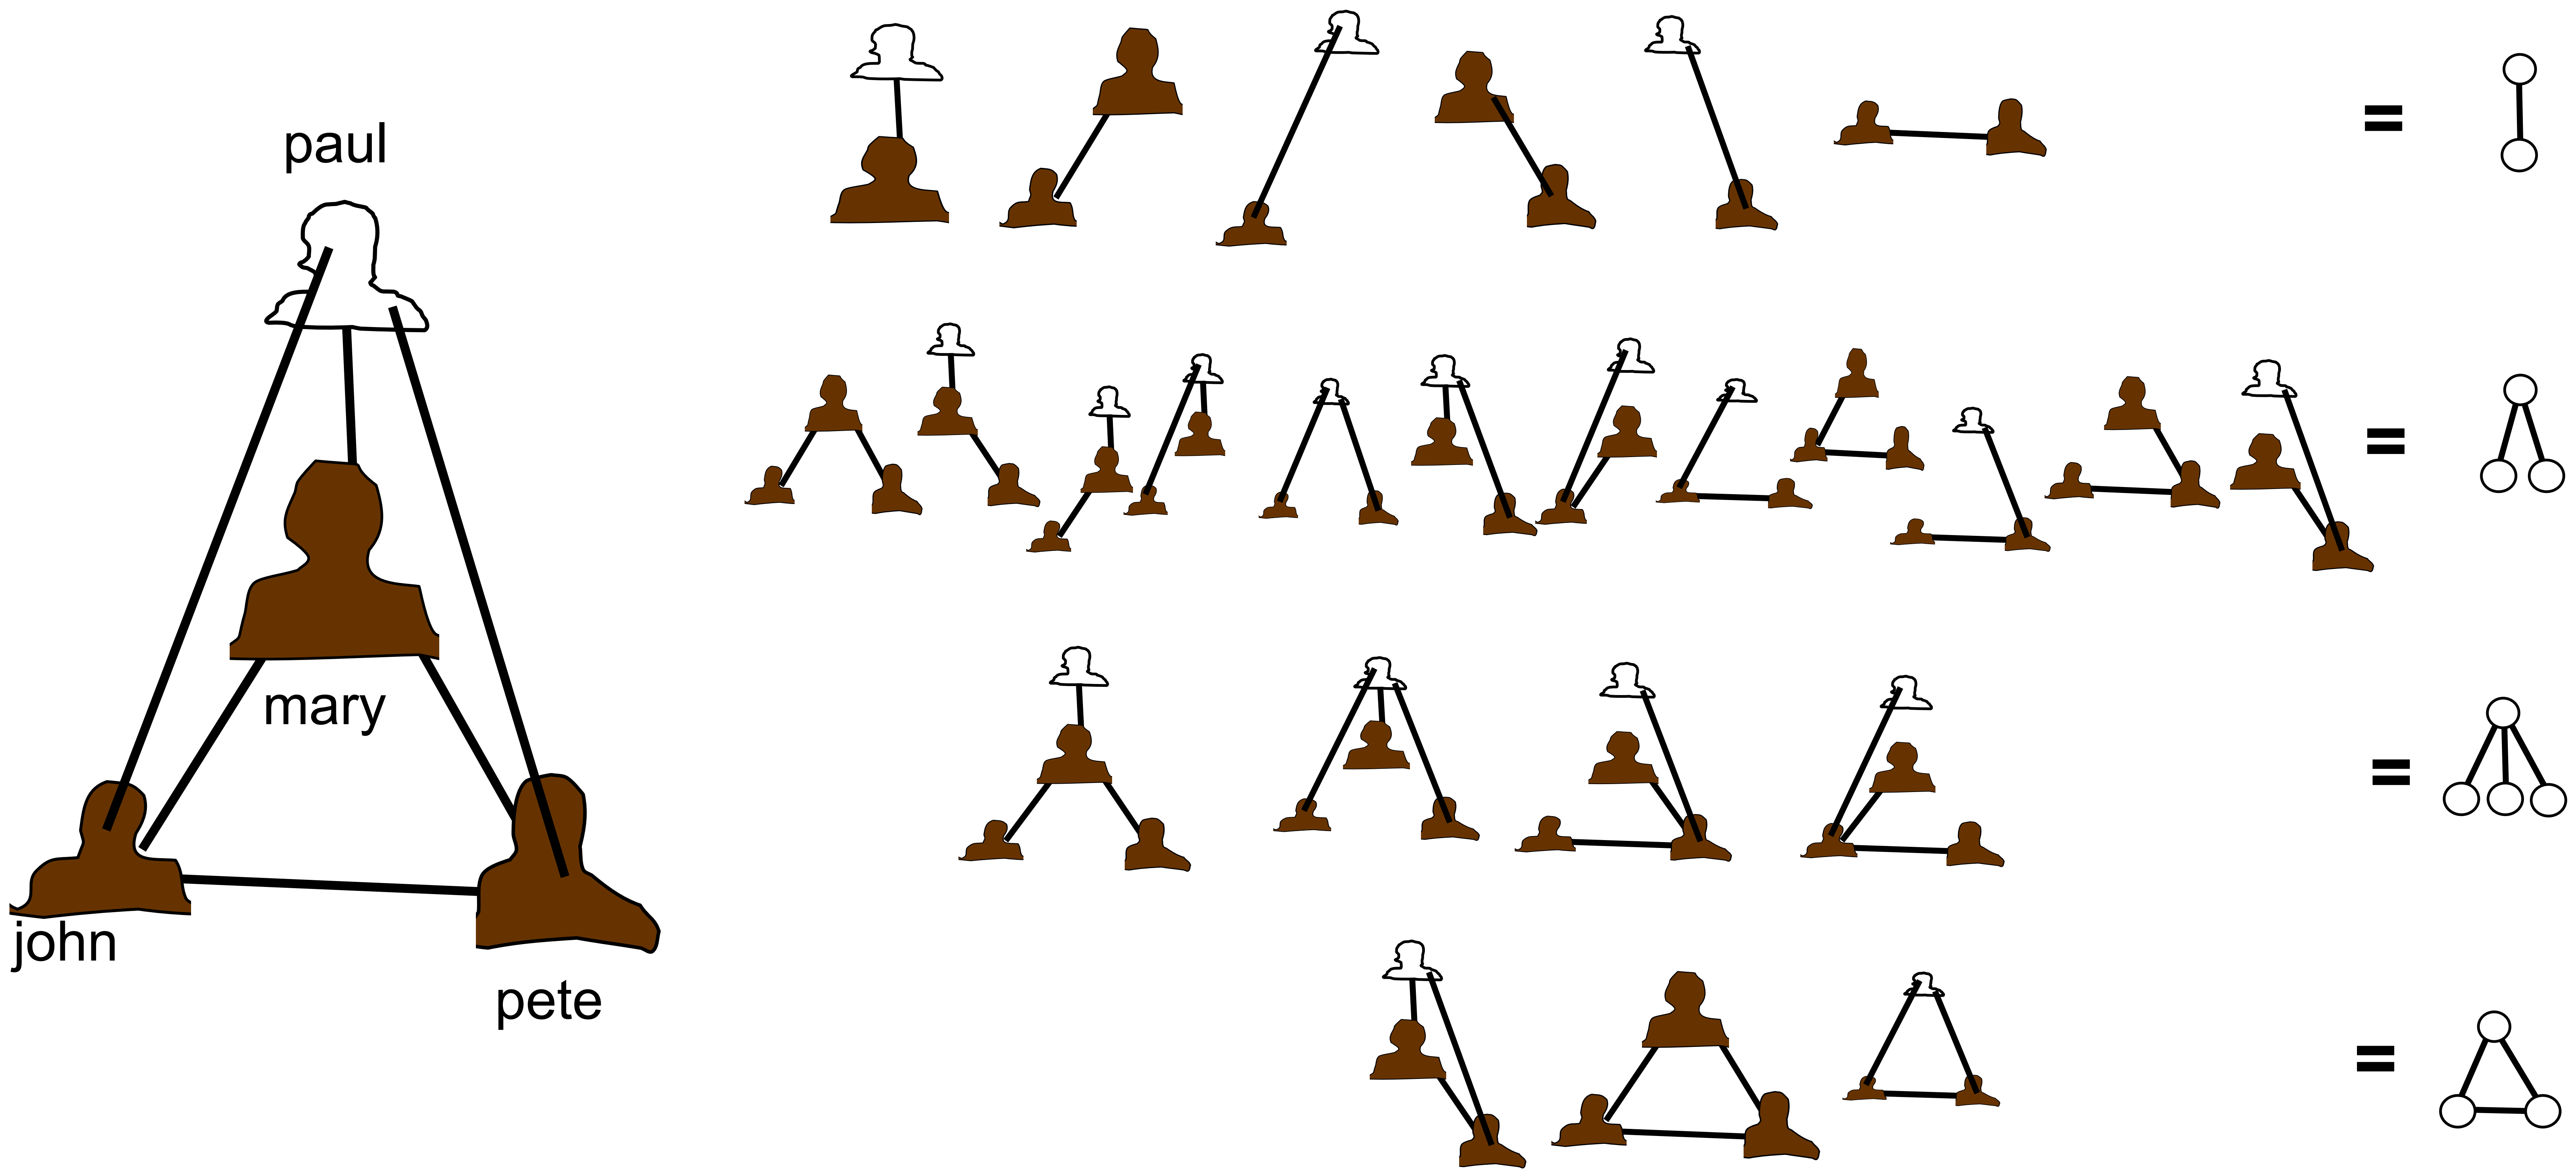

- Social networks are locally emergent; structured, yet stochastic

- Local configurations homogenous and those that appear more often than by chance and over attribute explanations evince endogenous mechanisms and multiple processes can operate simultaneously

$$P(X = x|\theta) = \frac{\exp(\sum_k \hat{\theta}_k z_k(x))}{\kappa}$$

The Secret Sauce

- Probability of network x is given by
  - a sum of network statistics (z)
    - expresses counts of network configurations (e.g. counts of reciprocal, transitive, or homophilic subgraphs)
  - that is weighted (θ)
    - expresses the importance of each configuration
  - inside an exponential (e)
    - this is an exponential-*family* random graph model, so that probabilities [0,1]
  - and is normalised (κ)
    - over all possible graphs of the same size (x' in X)

$$P(X = x|\theta) = \frac{\exp(\sum_k \hat{\theta}_k z_k(x))}{\kappa}$$

Problem: Oh κ!

- Ideally use maximum likelihood estimation, $L(\theta \mid x)$, directly, to find estimates of θ that make x most likely
- But remember κ?

$$\kappa = \sum_{x' \in X} \exp\left(\sum_k \theta_k z_k(x')\right)$$

  - Directed, binary network of n nodes has $2^{n(n-1)}$ states
  - Really, really large, making κ not computable except for very small graphs
- How large?...

How large?

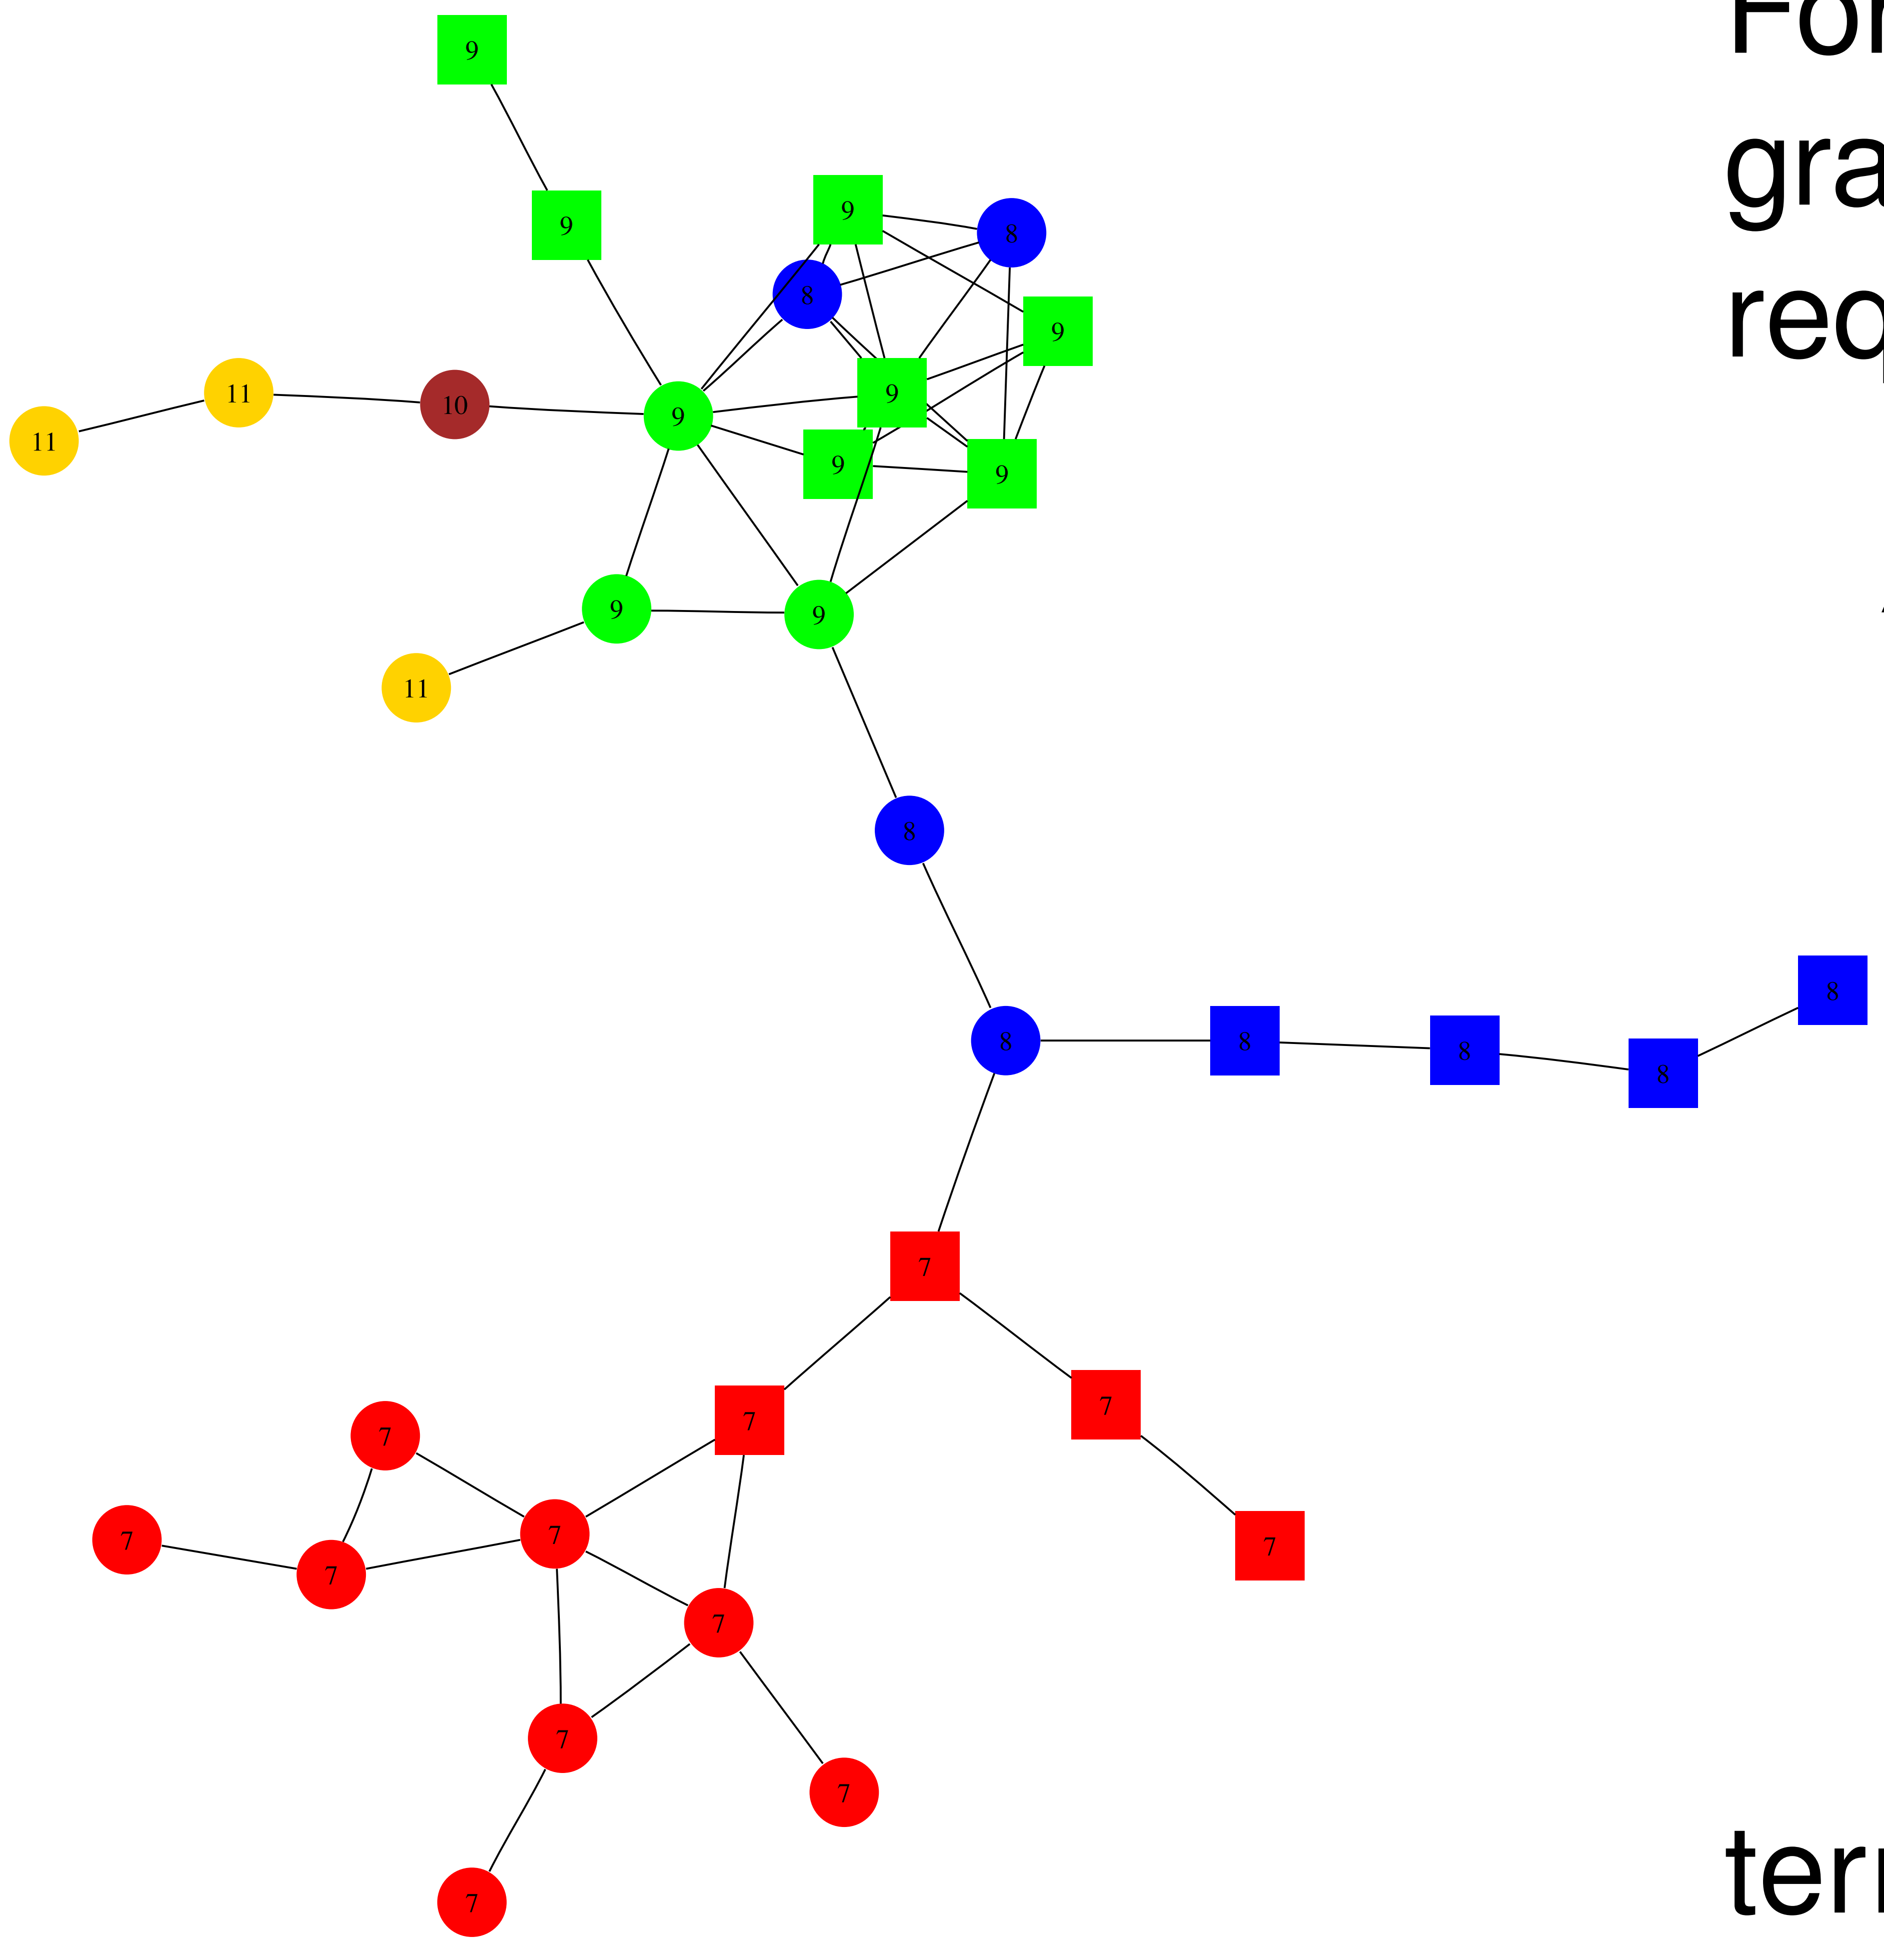

For this undirected, 34-node graph, computing $c(\theta)$ directly requires summation of

7,547,924,849,643,082,704,483,
109,161,976,537,781,833,842,
440,832,880,856,752,412,600,
491,248,324,784,297,704,172,
253,450,355,317,535,082,936,
750,061,527,689,799,541,169,
259,849,585,265,122,868,502,
865,392,087,298,790,653,952

terms.

PENNSTATE 1855

Oκ, how do we get κ?

- Markov Chain Monte Carlo (MCMC)
  - Different variations available (Gibbs, Metropolis-Hastings)
- Main idea: Simulate a discrete-time Markov chain whose stationary distribution is the distribution we want to sample from

Markov Chain Monte Carlo (MCMC)

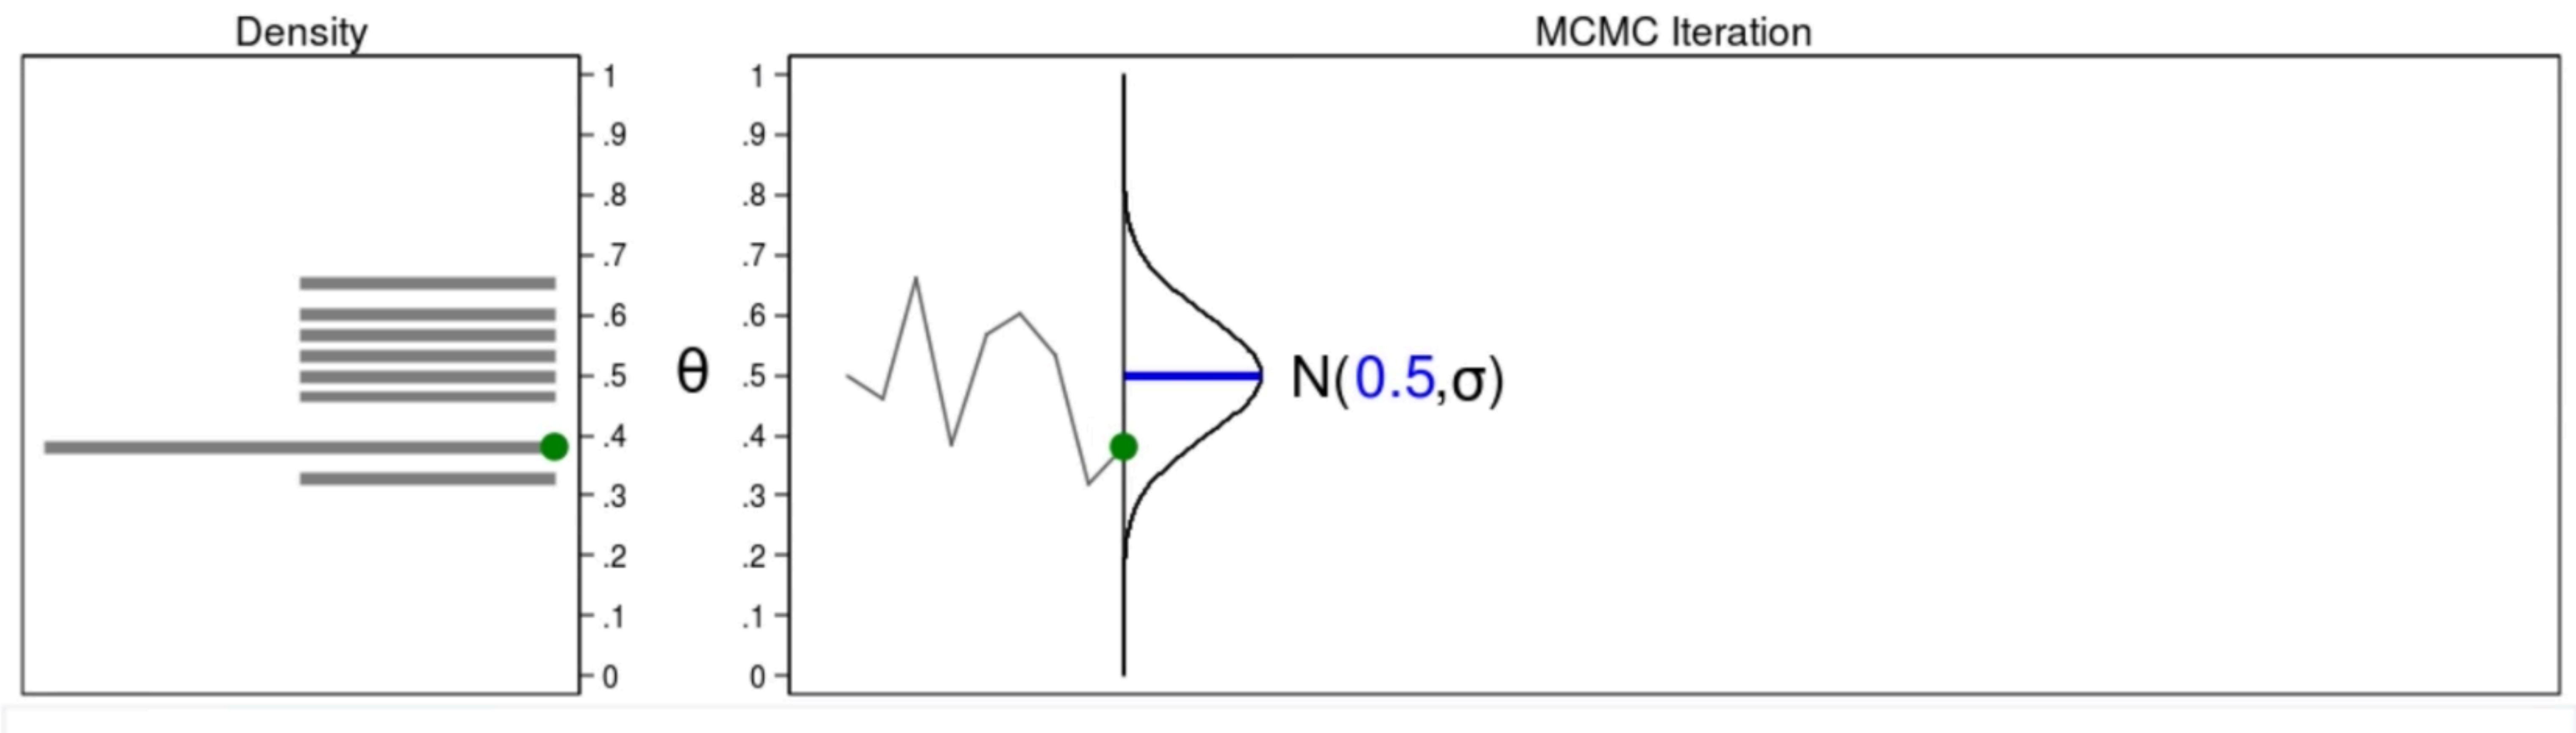

- refers to the part that relies on the generation of random numbers
- note that the distribution on the left resembles the distribution we are drawing from and that the proposal distribution does not move

Markov Chain Monte Carlo (MCMC)

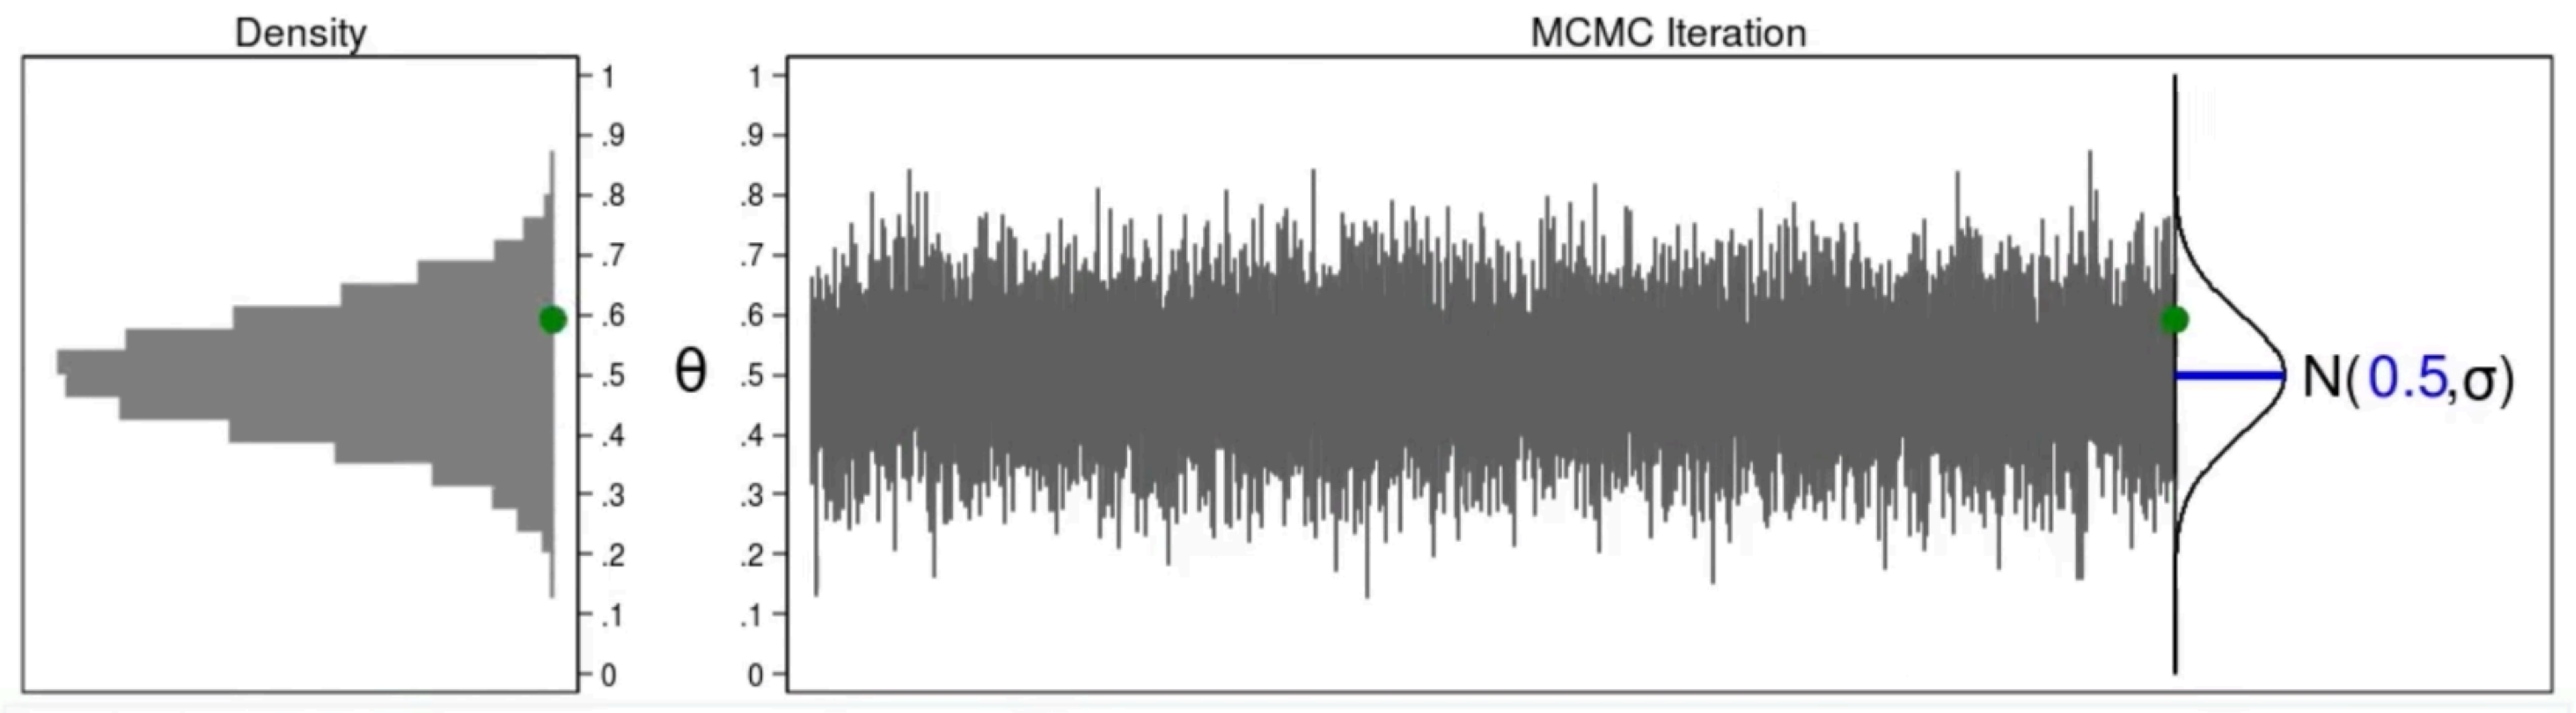

- refers to the part that relies on the generation of random numbers
- note that the distribution on the left resembles the distribution we are drawing from and that the proposal distribution does not move

Markov Chain Monte Carlo (MCMC)

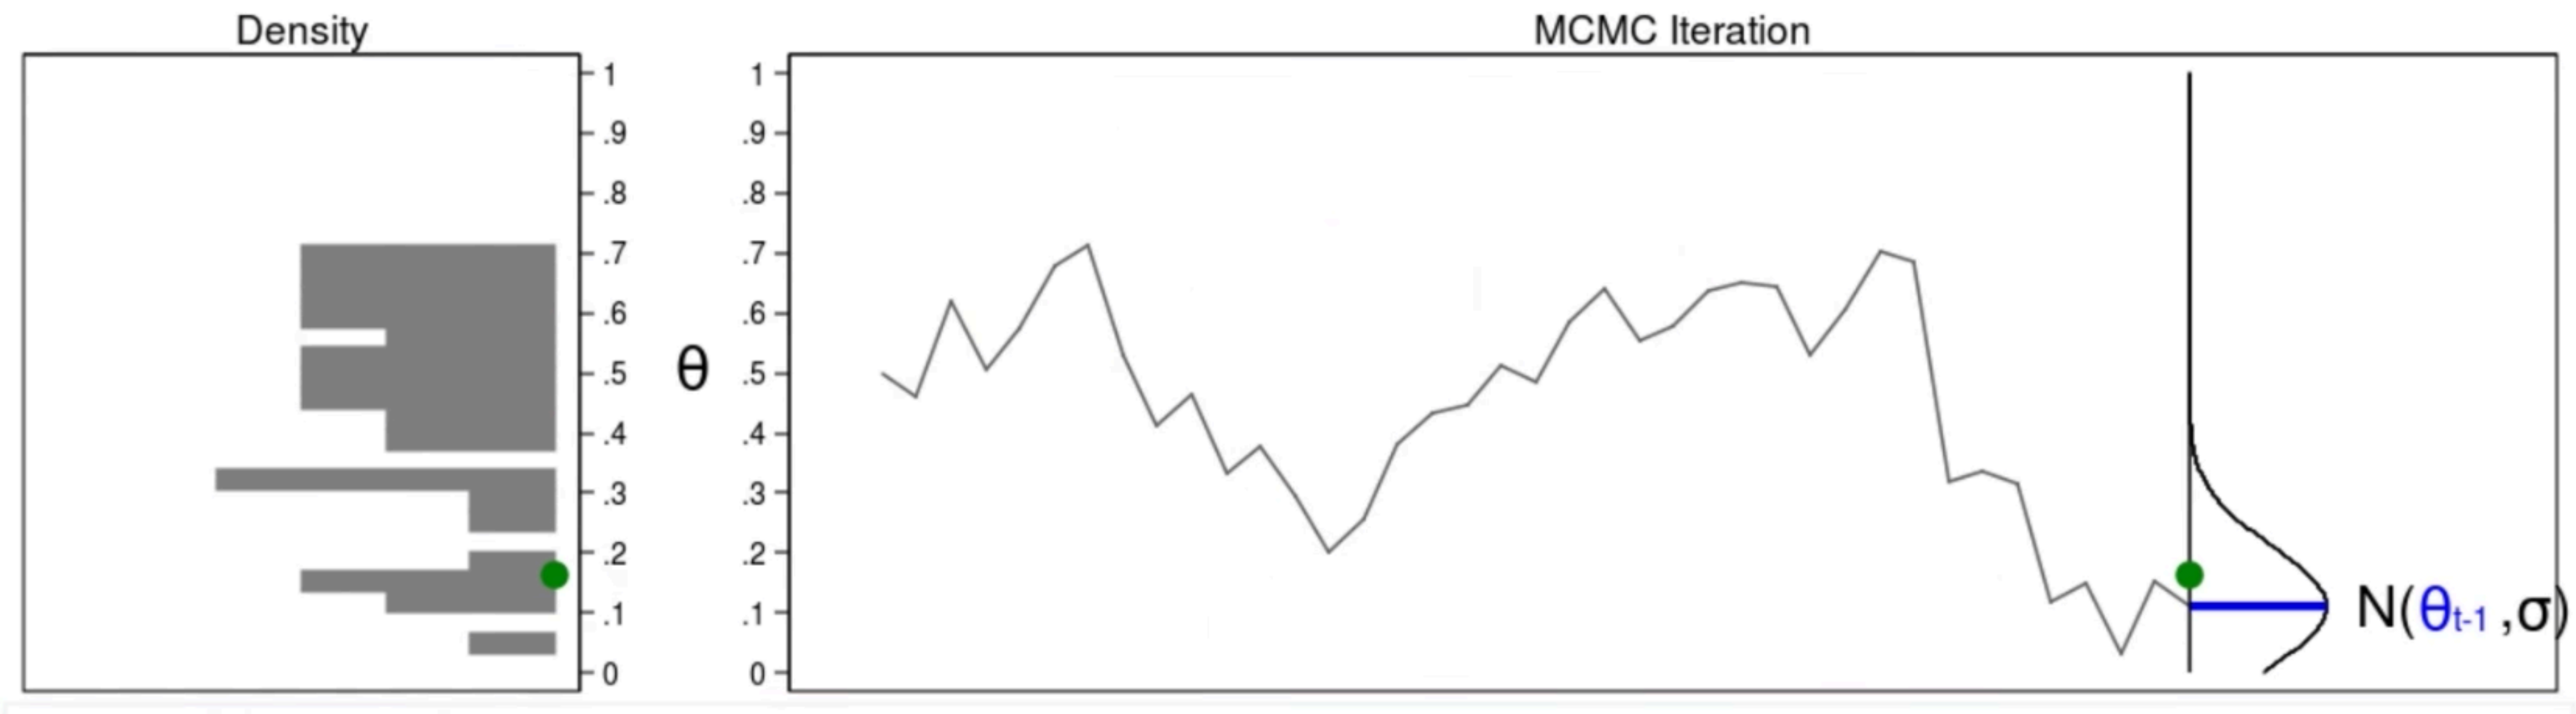

- is a sequence of numbers in which each number is dependent (only) on the previous number

- traceplot seems to wander like in a random walk

Markov Chain Monte Carlo (MCMC)

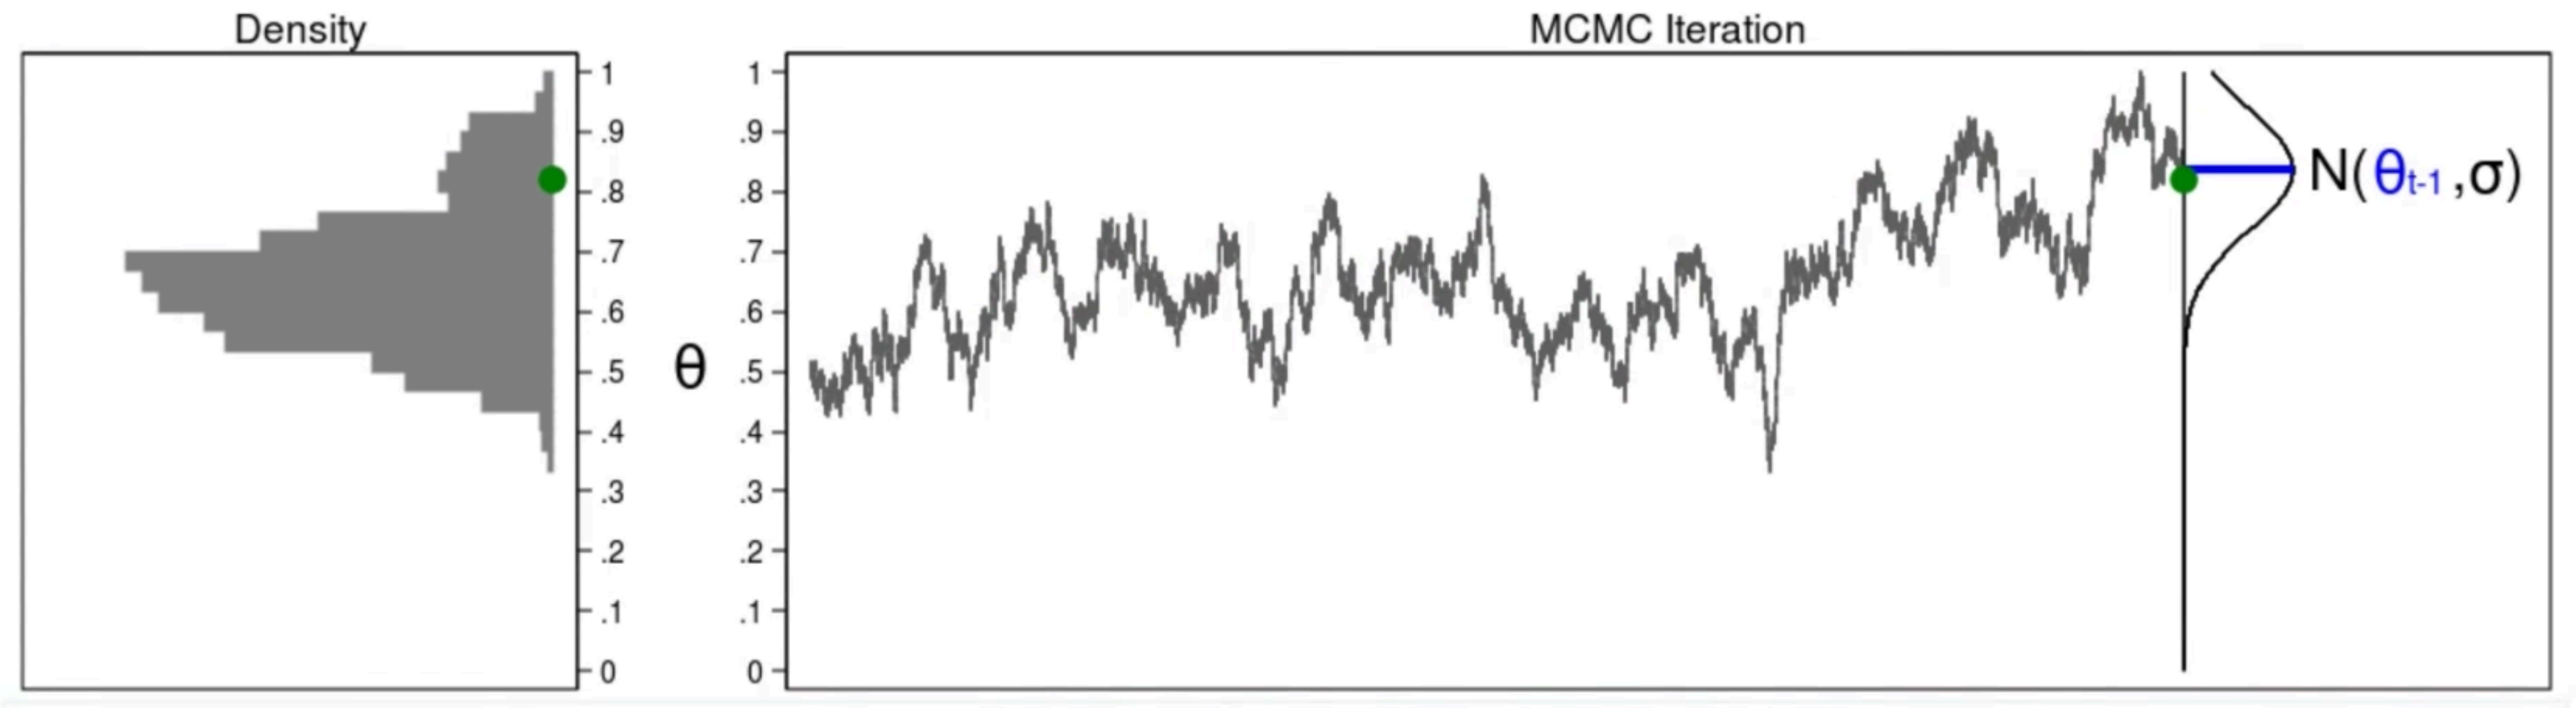

- is a sequence of numbers in which each number is dependent (only) on the previous number

- traceplot seems to wander like in a random walk

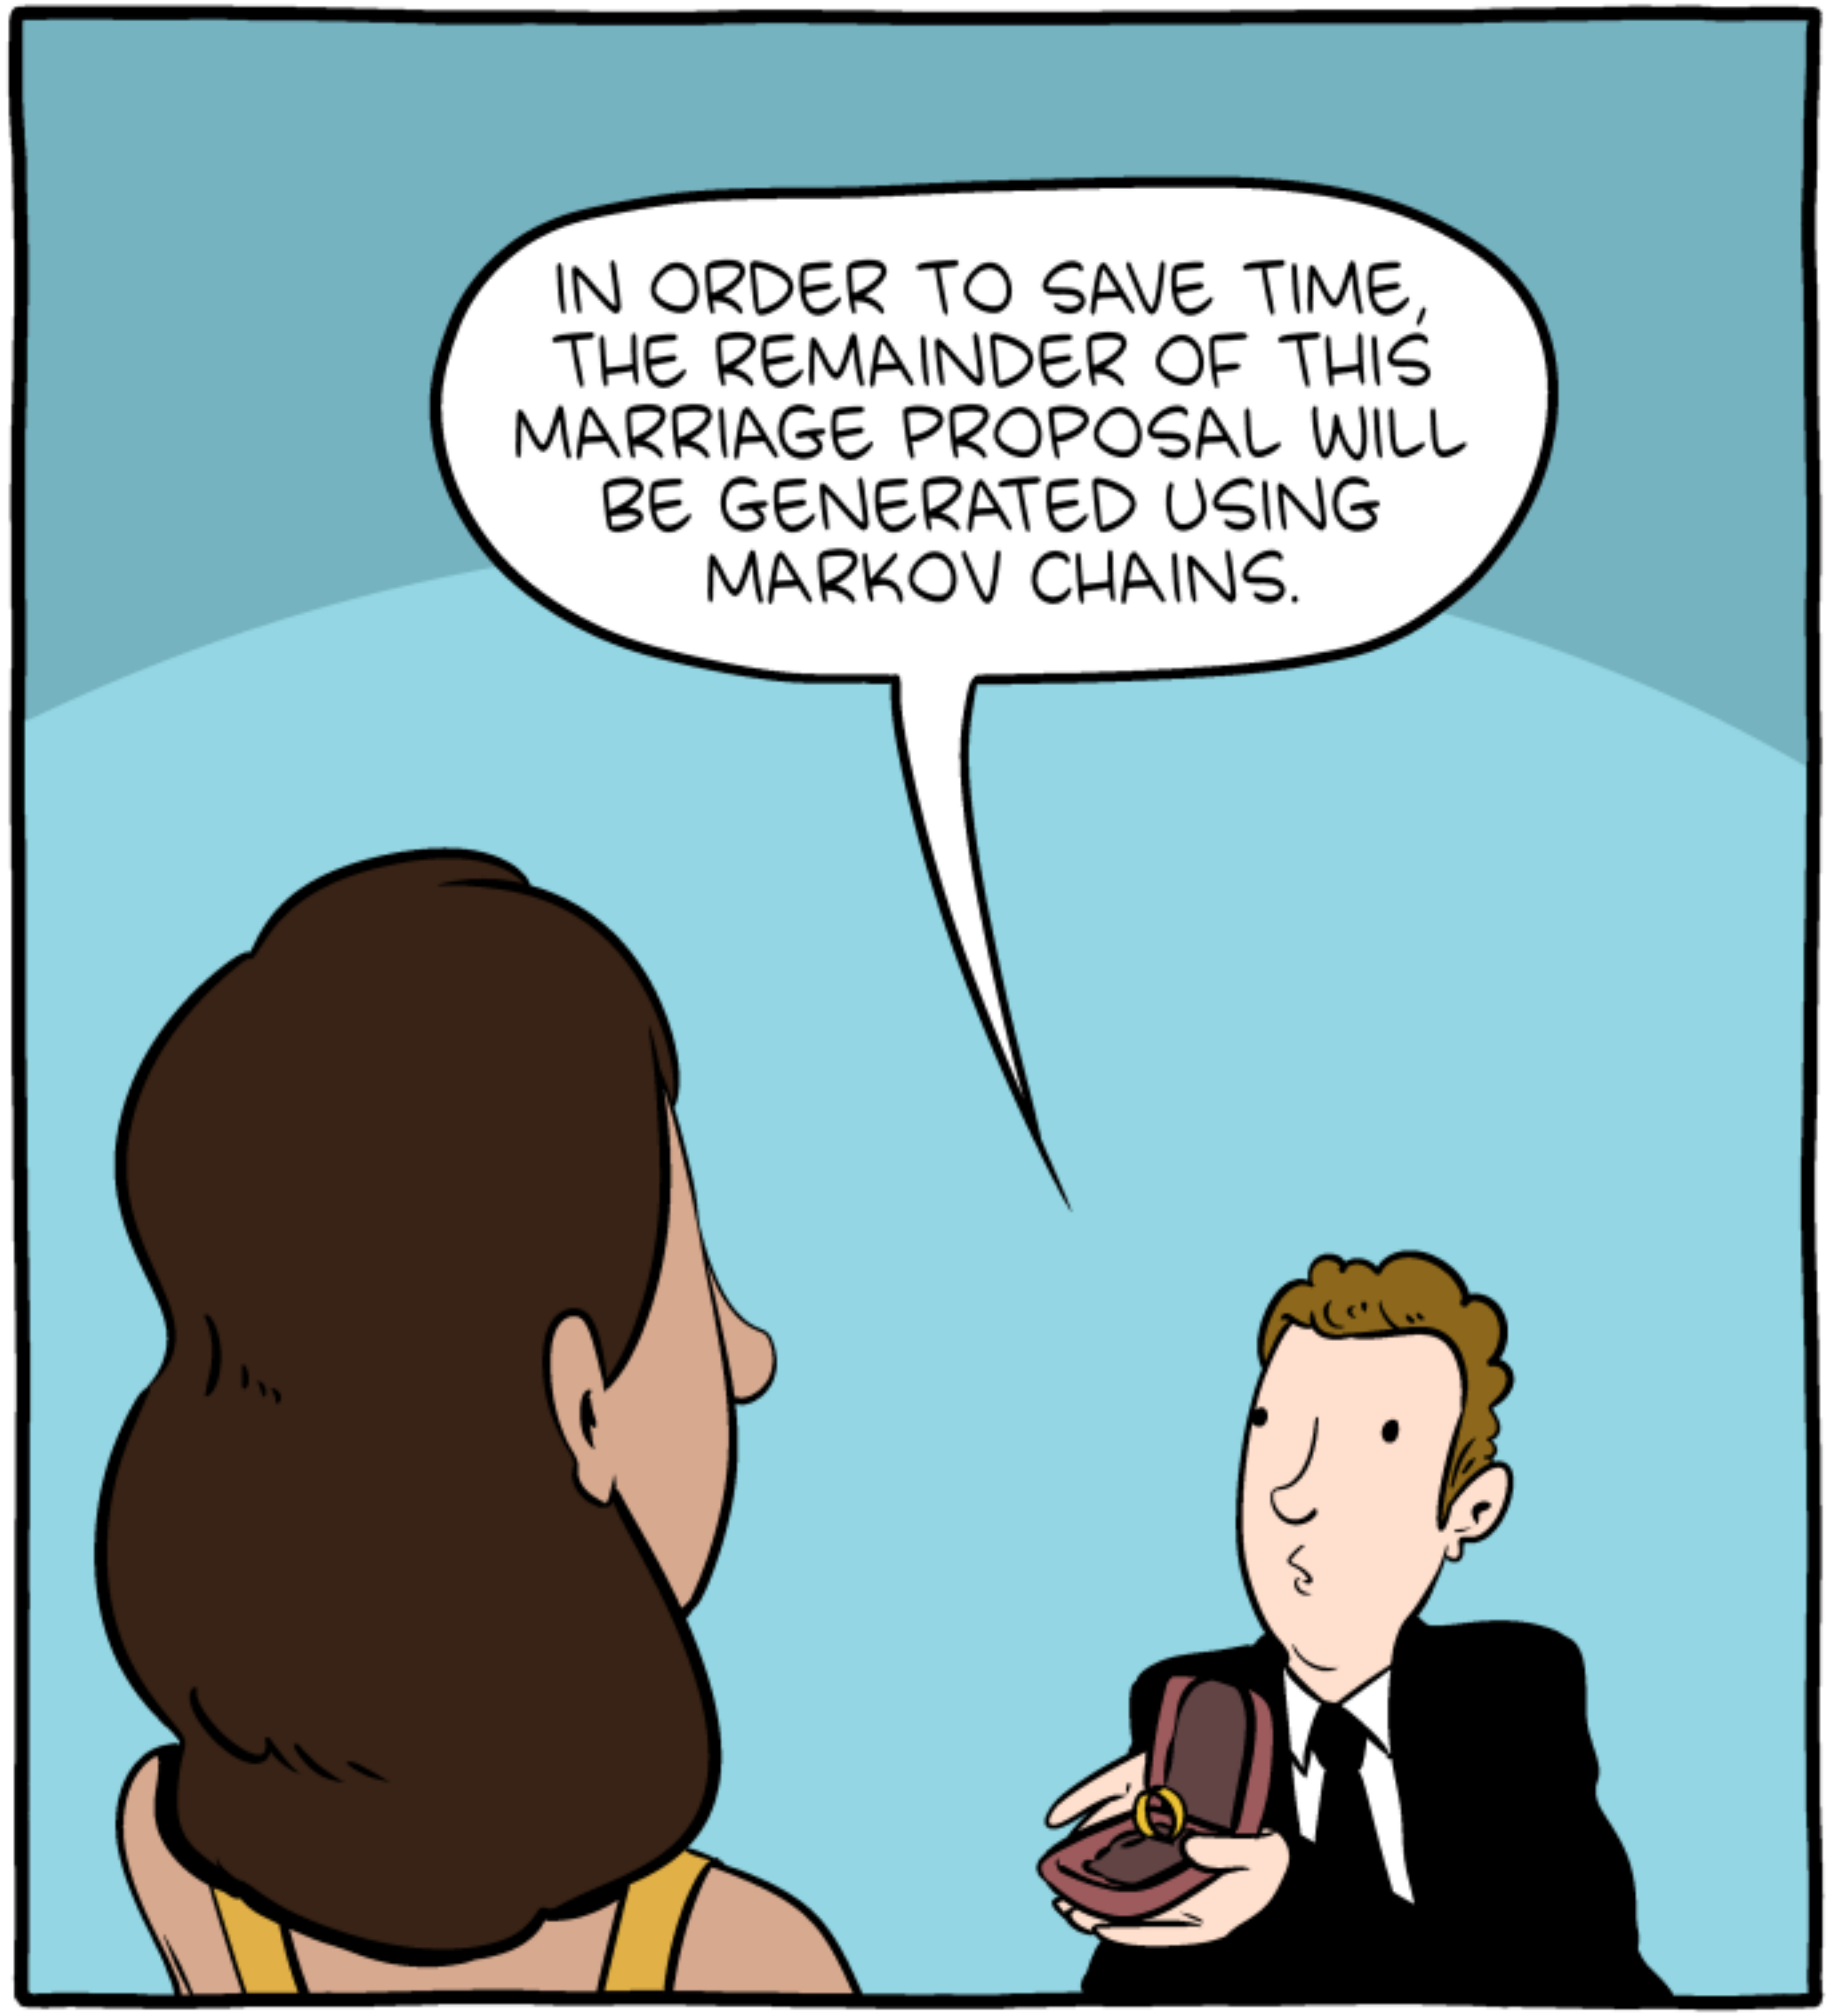
IN ORDER TO SAVE TIME, THE REMAINDER OF THIS MARRIAGE PROPOSAL WILL BE GENERATED USING MARKOV CHAINS.

An underlying Markov chain

- The ERGM is also the stationary distribution of a Markov random walk with transition probabilities

$$p(x \to x^{i \rightsquigarrow j}; \theta) = \frac{1}{N(N-1)} \cdot \frac{\exp\left(\sum_k \theta_k z_k(x^{i \rightsquigarrow j}) \right)}{\exp\left(\sum_k \theta_k z_k(x) \right) + \exp\left(\sum_k \theta_k z_k(x^{i \rightsquigarrow j}) \right)}$$

- In theory, if we just let this random walk run long enough, it will approximate the stationary distribution and thus the ERGM for a given parameter θ

- In practice, this problem is again intractable

Sampling from the Markov chain to estimate

- However, we can use the Markov chain to simulate networks $x^{(1)}, x^{(2)}, \ldots, x^{(M)}$ that are a good sample of the space of all networks

- Just need to make sure that these simulated networks have a *low autocorrelation* and are *representative of the sample space*

- Calculate the sample equivalent of $E_{\hat{\theta}}(z(X))$

$$\bar{z}_\theta = \frac{1}{M}\left(z(x^{(1)}) + z(x^{(2)}) + \ldots + z(x^{(M)})\right)$$

- Check whether $\bar{z}_\theta - z(x_{obs}) = 0$
  - If yes, $\theta = \hat{\theta}$
  - If no, update $\hat{\theta}$

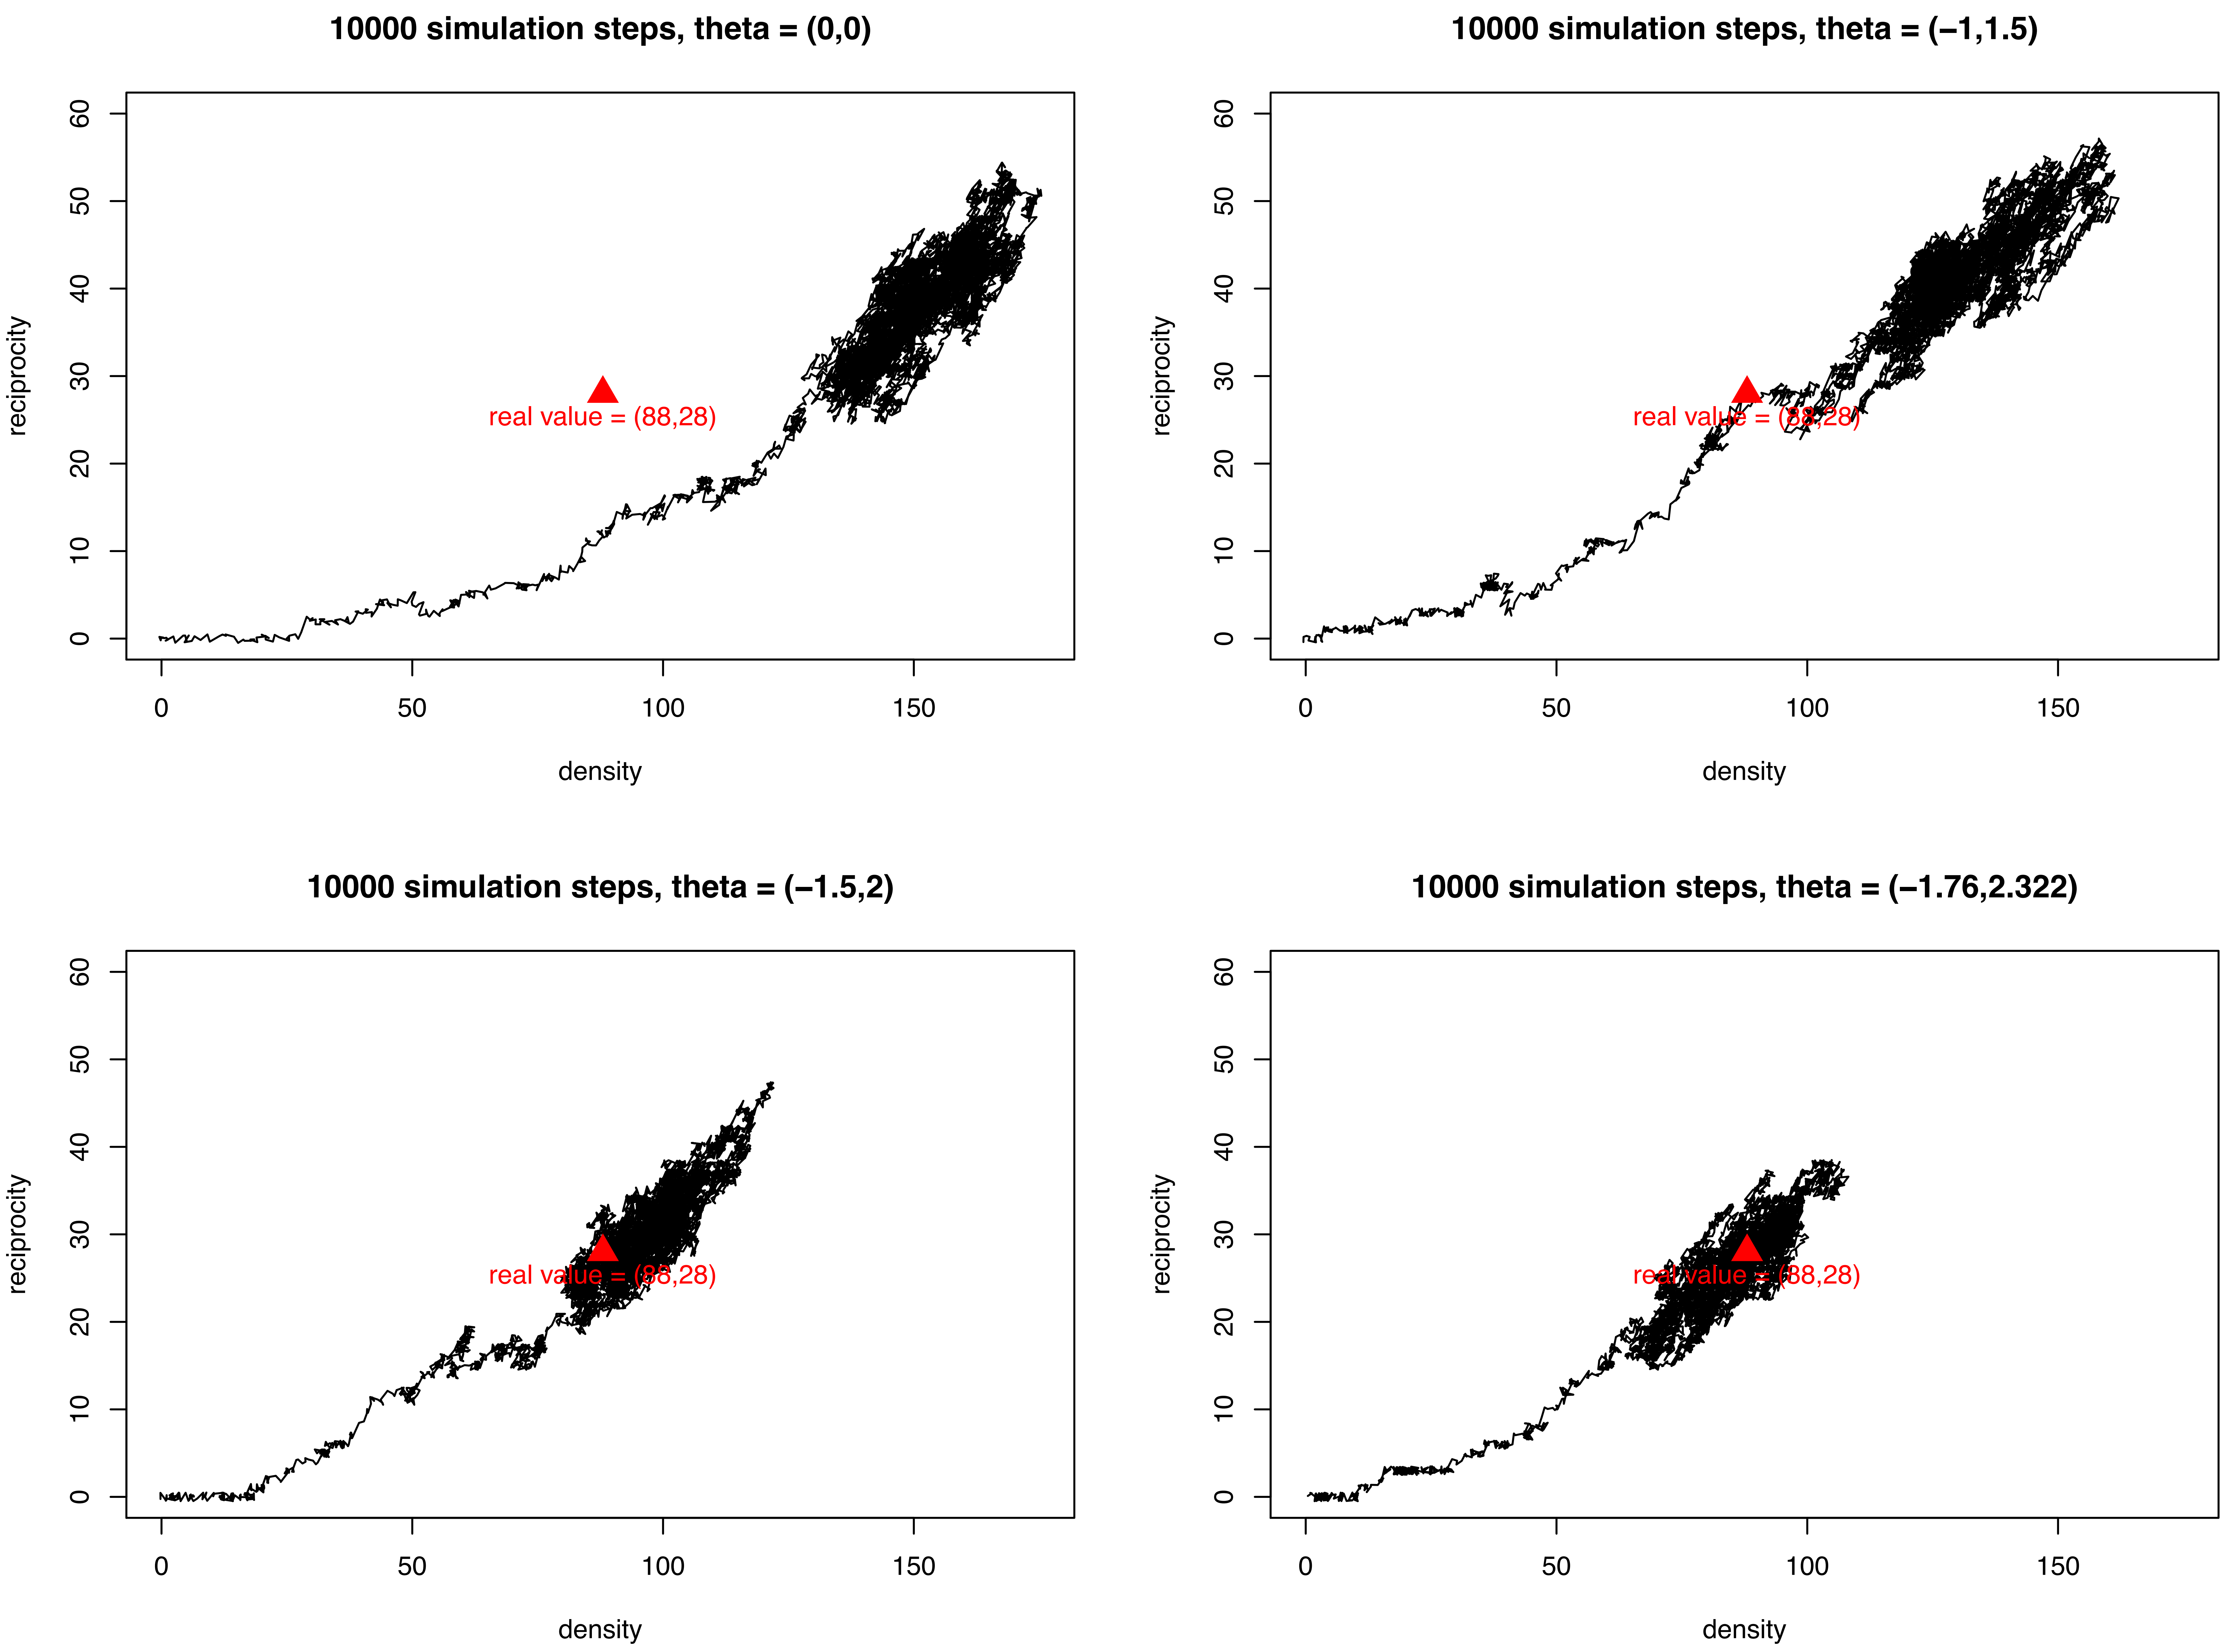
10000 simulation steps, theta = (0,0)
10000 simulation steps, theta = (−1,1.5)
10000 simulation steps, theta = (−1.5,2)
10000 simulation steps, theta = (−1.76,2.322)
real value = (88,28)
reciprocity
density
0
10
20
30
40
50
60
0
50
100
150

Convergence

- The ERGM tries to produce a combination (vector) of parameter estimates that together generate simulated networks that don't differ (much) from the observed network on the salient statistics

- When it has settled on estimates (any updates are very small and tend to oscillate around a particular point estimate) we can say that the model has converged

So far, the empirical approach to Zeno's Paradox has been inconclusive.

3 Main Issues

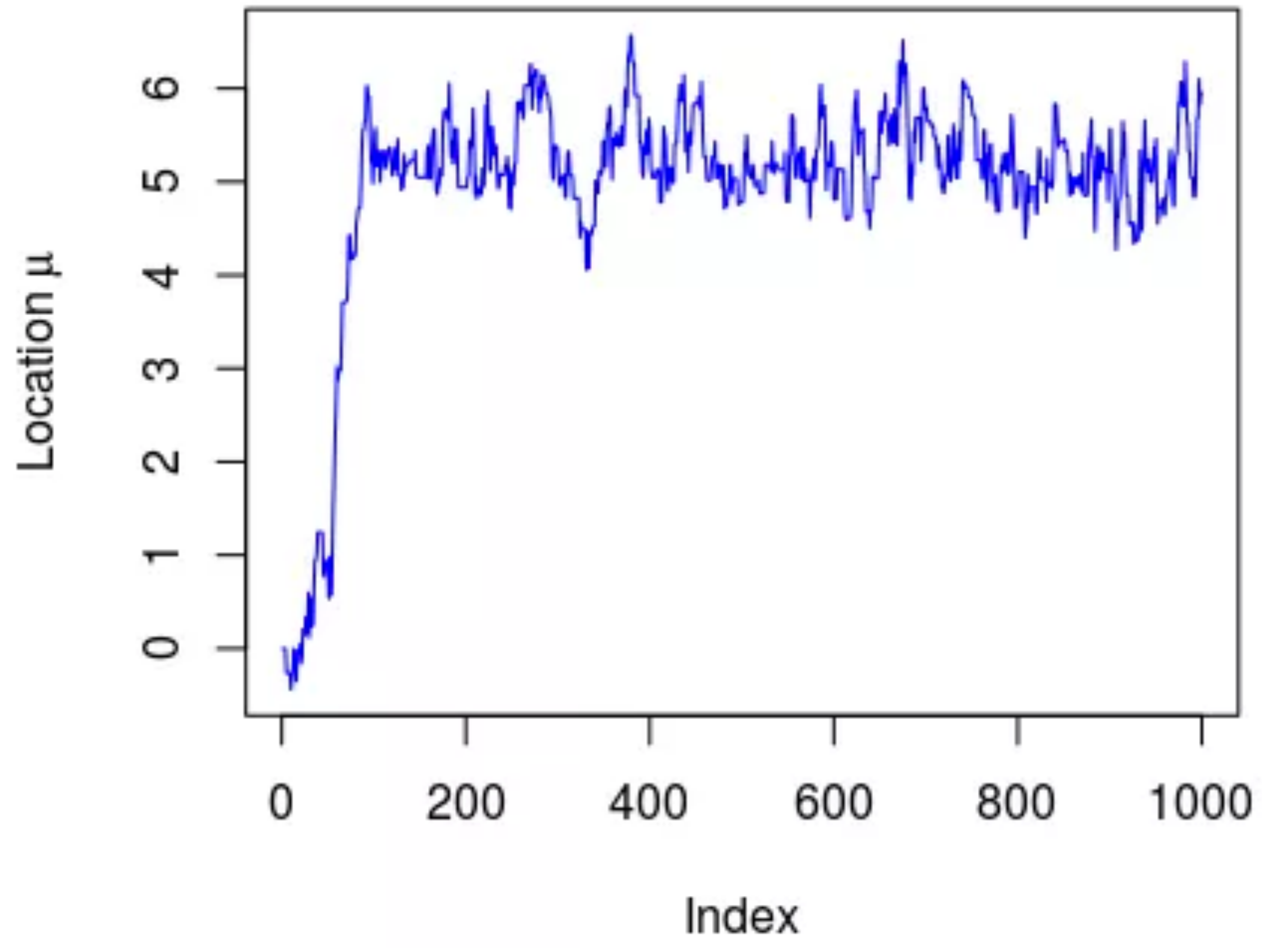

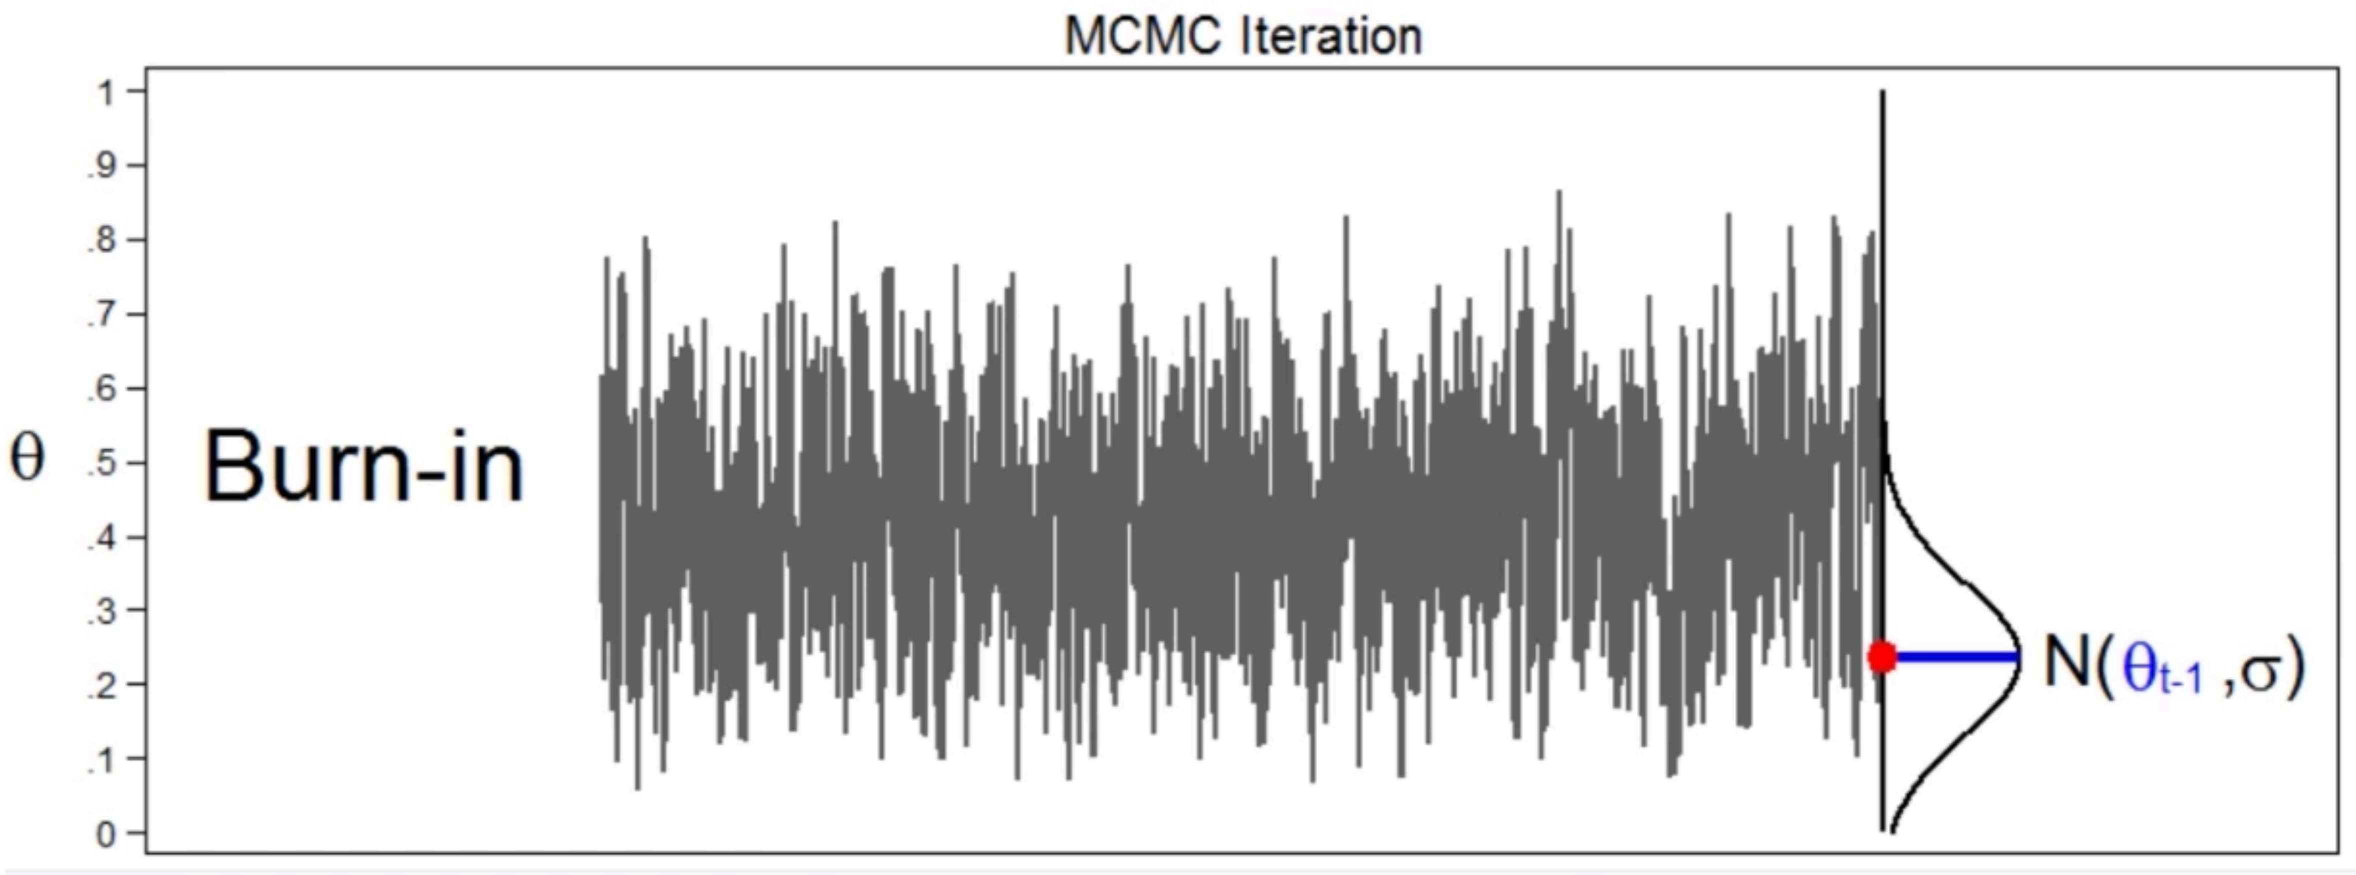

- 1. Dependence on starting values
  - Problem: Some starting values (e.g. 0) may be biased
  - Solution: Increase *burnin* period to discard first samples from Markov chain to give it time to stabilise or restart with new values

3 Main Issues

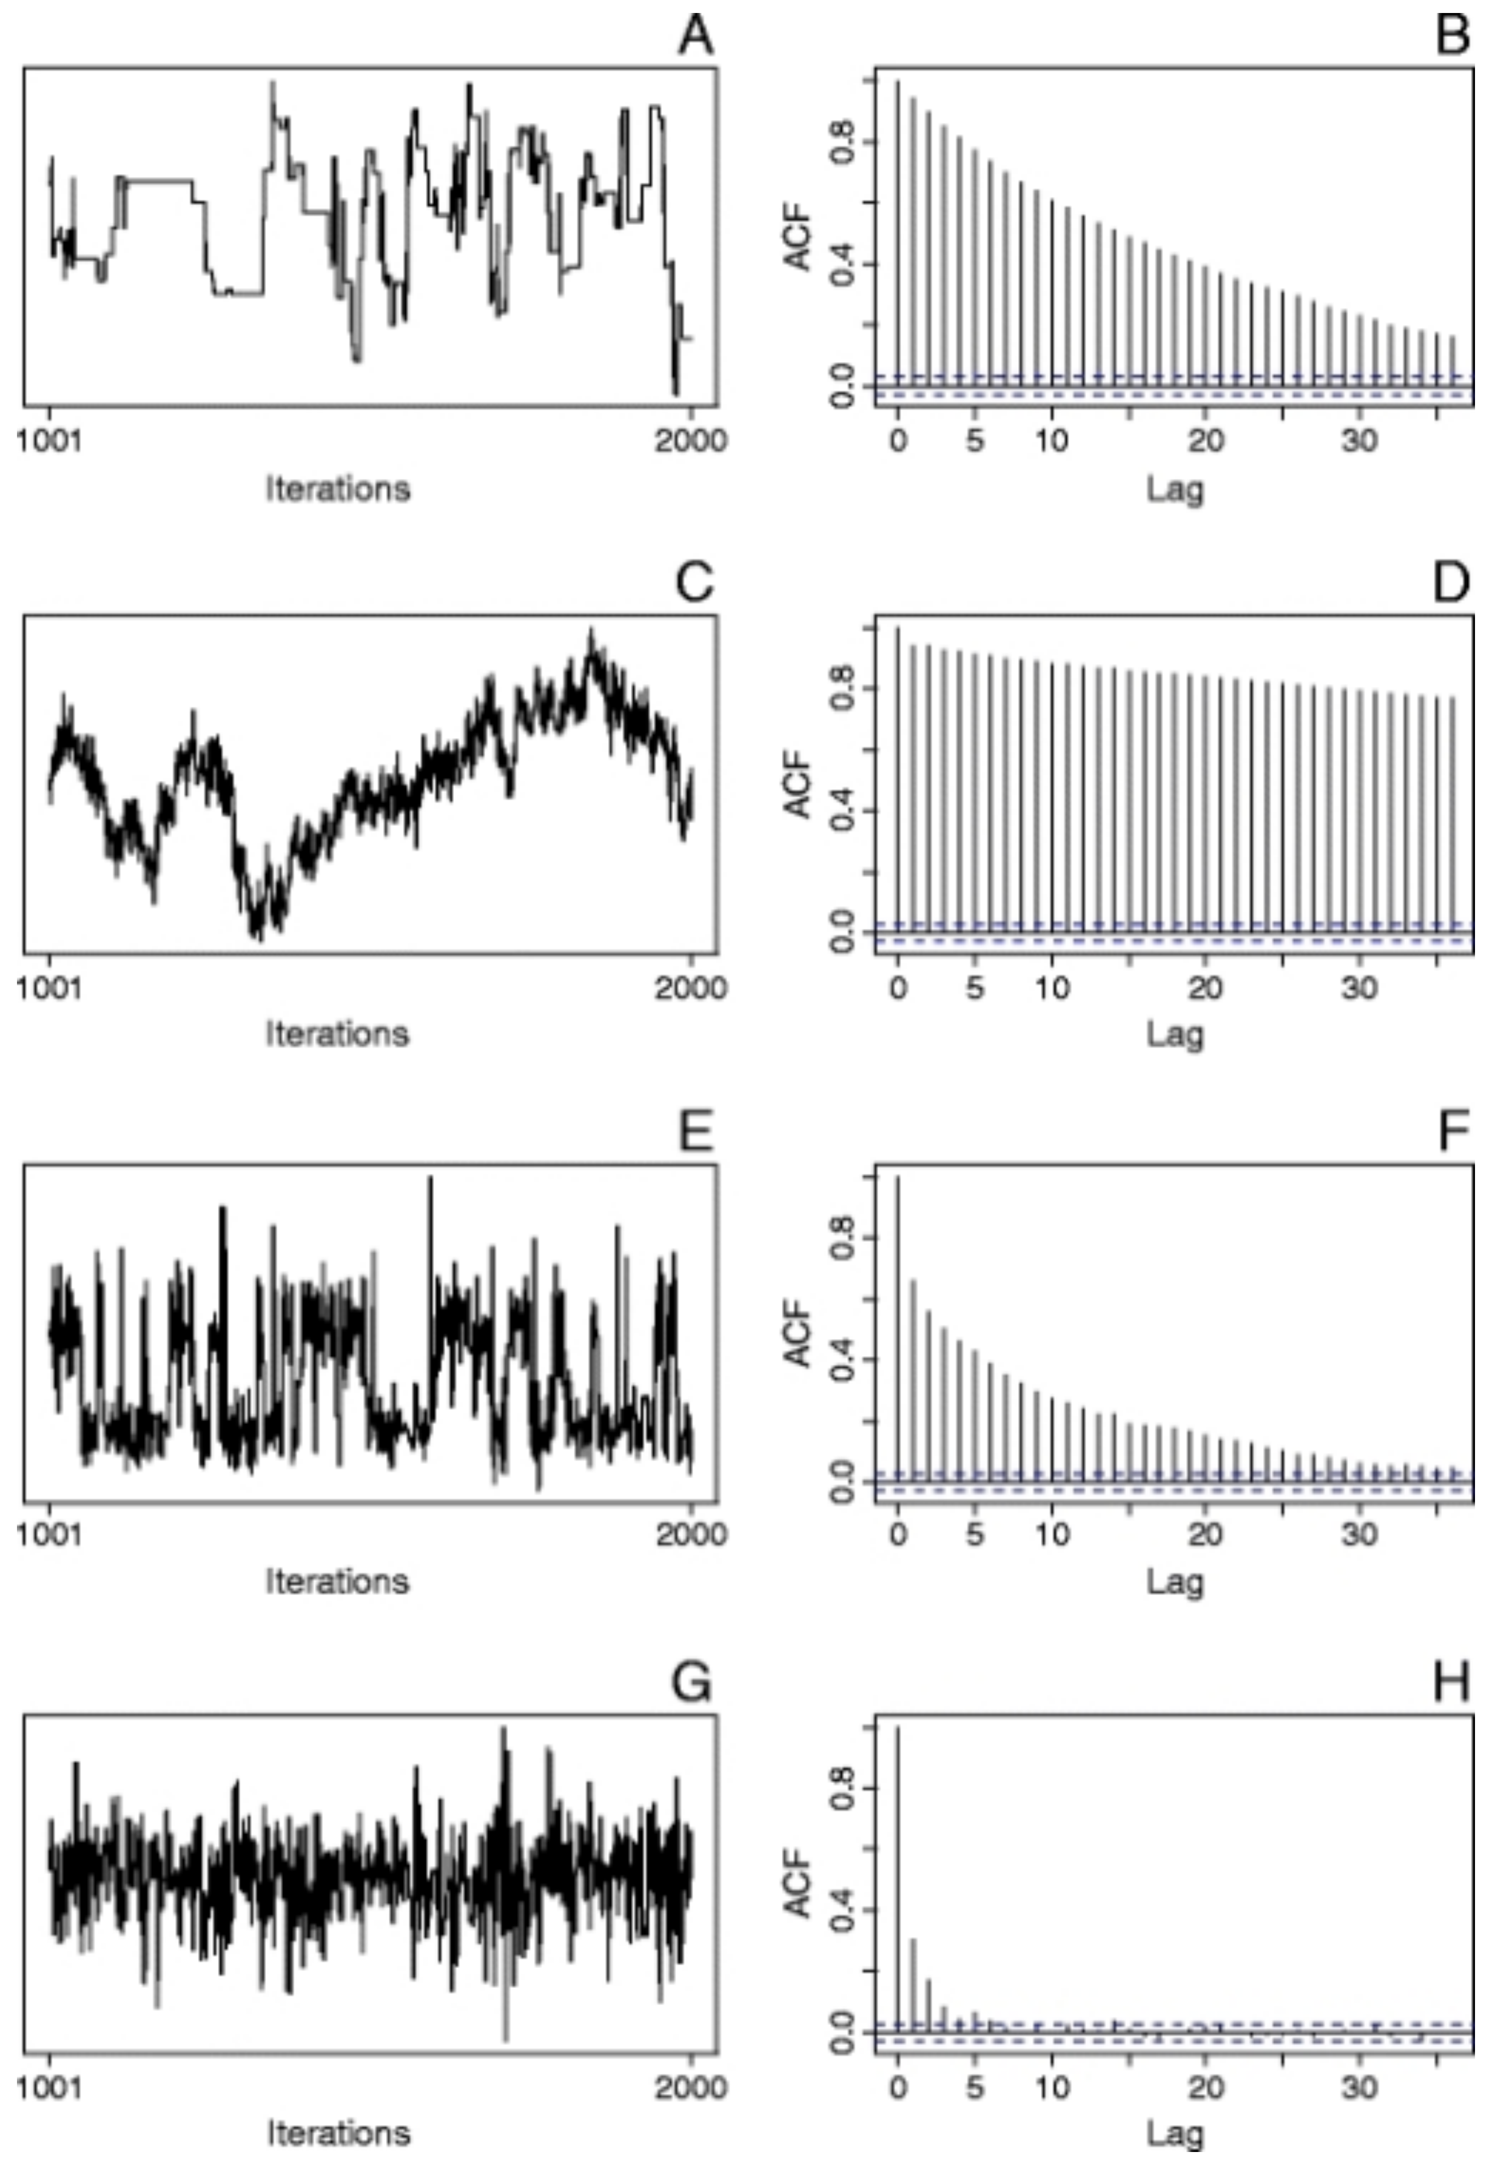

- 2. Autocorrelation due to Markov chain
  - Problem: Some normal, but should drop down and waver around 0 quite quickly or is considered excessive (not mixing well)
  - Solution: Increase *thinning* or change model

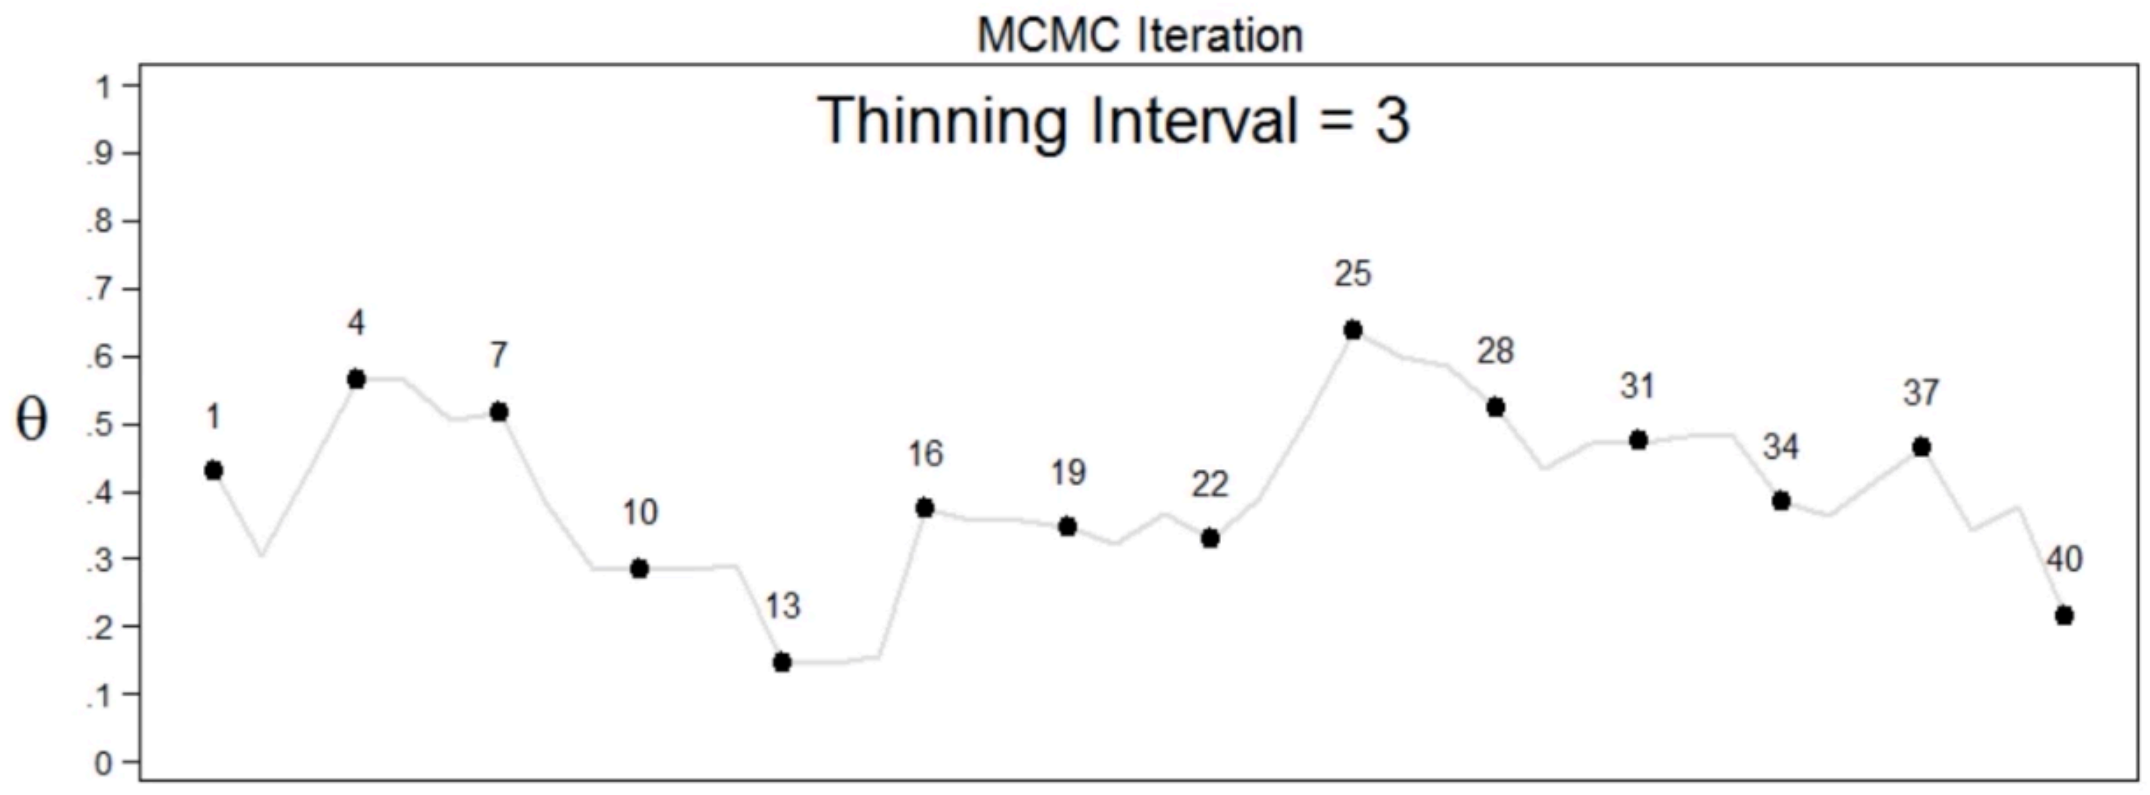

3 Main Issues

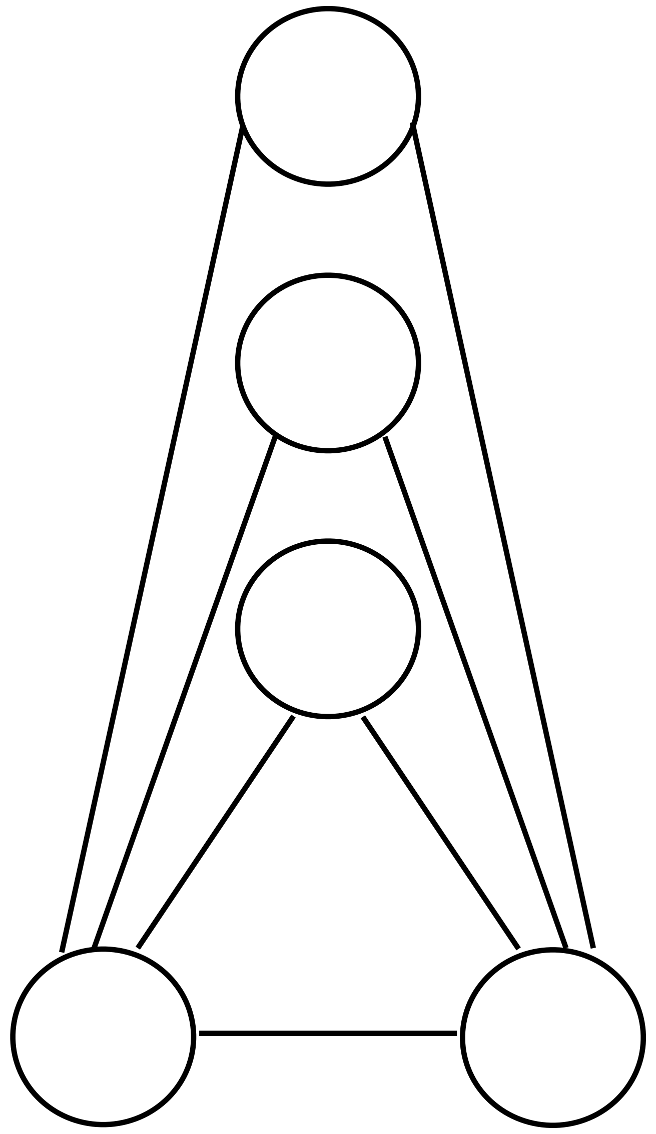

- 3. A major problem with ERGMs is that some ingredients shouldn't be scaled linearly:
  - Having a friend in common obviously makes our friendship more likely
  - But should each additional friend contribute the same? Three friends = thrice as likely? Four friends = …*fource* as likely? Same "info" in each?
  - If scaled linearly, then simulated networks would end up *degenerate*: impossibly dense, sparse, etc.

GWESP WTF?

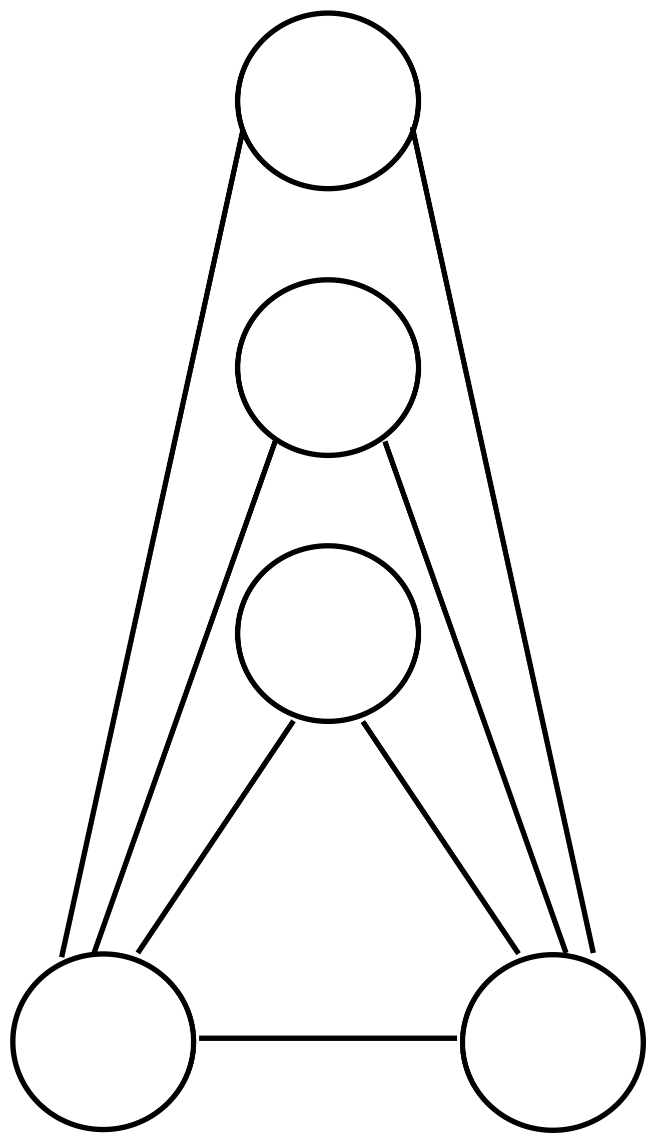

- Maybe better to discount additional friends?
  - *Alternating* k-stars and triangles effectively alternate the contributions of successive ties positively and negatively
  - *Geometrically-weighted* degrees and edgewise-shared partners discounts additional contributions by α
- Basically the same:
  - $\alpha = 0$, then GWESP statistic = number of edges in at least one triangle
  - $\alpha \to \infty$, then GWESP statistic -> 3x number of triangles
  - so as $\alpha \to 0$, subsequent ties/partners discounted more
- The lower α, model less likely to be degenerate, so start by fixing α low, say 0.25 or so (possible to estimate together with the coefficient, but slooooow)

Snijders et al 2006, Hunter 2007

A results table

	Naïve Actor Model	Political Capacity Model	Strategic Decision Model	Strategic Geography Model
General Parameters				
Density	-3.88 (0.03)*	-3.75 (0.07)*	-7.01 (0.35)*	-5.77(0.36)*
Centralization (actors)	–	–	0.61 (0.11)*	-0.21(0.11)
Centralization (institutions)	–	–	1.36 (0.18)*	0.56(0.18)*
Closure (actors)	–	–	-0.19(0.05)*	-0.06(0.04)
Geographic Centralization	–	–	–	1.57(0.05)*
Actor Type Activity Parameters (Local Government is Excluded Category)				
Federal Government	–	0.45 (0.15)*	0.43 (0.16)*	1.82(0.18)*
State Government	–	0.19 (0.14)	0.16 (0.13)	1.35(0.16)*
Water Special District	–	0.13 (0.09)	0.12 (0.09)	0.42(0.10)*
Environmental Special District	–	0.29 (0.17)	0.26 (0.17)	0.46(0.19)*
Environmental Group	–	-0.18 (0.10)	-0.16 (0.09)	-0.01(0.10)
Industry Group	–	-0.59 (0.26)*	-0.50 (0.23)*	0.05(0.29)
Education/Consulting	–	-0.40 (0.18)*	-0.32 (0.17)	-0.06(0.19)
Actor Coalition	–	-0.03 (0.34)	-0.03 (0.33)	0.44(0.38)
Other Activity	–	0.07 (0.48)	0.11 (0.43)	1.33(0.54)*
Institution Type Activity Parameters (Collaborative Partnership is Excluded Category)				
Interest Group Association Activity	–	-0.22 (0.10)*	-0.09 (0.09)	-0.04(0.06)
Advisory Committee Activity	–	-0.16 (0.12)	-0.10 (0.11)	-0.03(0.06)
Regulatory Process Activity	–	-0.78 (0.16)*	-0.61(0.15)*	-0.36(0.12)*
Actor as Venue Activity	–	-0.70 (0.19)*	-0.47 (0.16)*	-0.26(0.13)*
Joint Powers Authority Activity	–	0.16 (0.16)	0.15 (0.15)	0.06(0.10)
Mahalanobis distance as an indicator of model fit (smaller values indicate greater fit)	46,208	15,541	4,173	638

Note: Cell entries are ERGM parameter estimates with standard errors in parentheses. All models are estimated with "exogenous hubs," with fixed degree distributions for nodes with more than 20 edges. *Reject null hypothesis of parameter = 0, $p < 0.05$.

- Much like a logit
  - Coefficients represent change in (log-odds) likelihood of a tie for a unit change in predictor
- Predictors are network-level statistics that represent Markovian processes, so we can think about their changes locally
- Practical script goes into this in more detail, but we see here that there is geographic centralisation, and that this effect flips actor centralisation and mutes institutional centraliation and actor closure

Lubell, Robins, and Wang 2014

Interpretation

- Parameters of social network models (ERGM, SAOM) notoriously difficult to interpret because:
  - No uni-dimensional dependent variable
  - No single 'error term' (endogeneity)
  - Nonlinearity of the model (cf. logistic regression)
  - Substantive effects sometimes/often represented by multiple effects (model terms) in the model
  - Fundamentally N=1 models – model-based inference
  - Between-network comparisons face difficulties related to: different numbers of nodes, different average degrees, for the SAOM, different time lengths between waves

ERGM

Effects

Model

Diagnostics

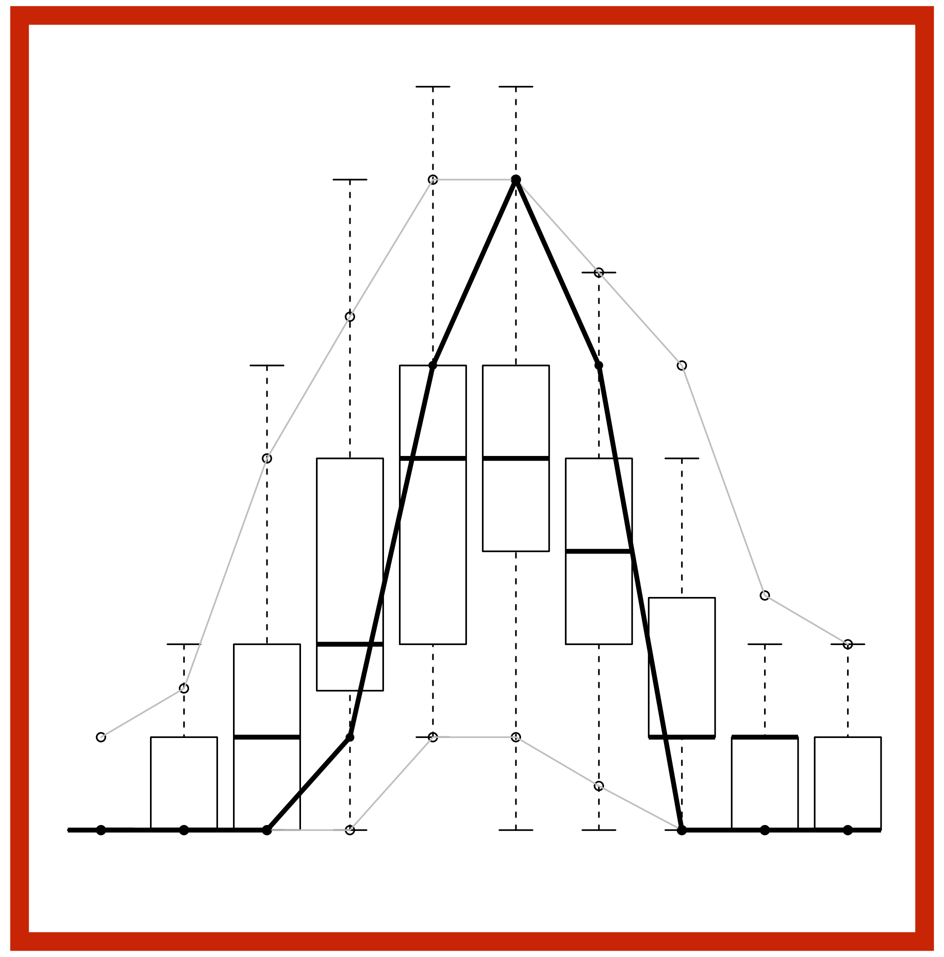

Diagnostics

In Silico

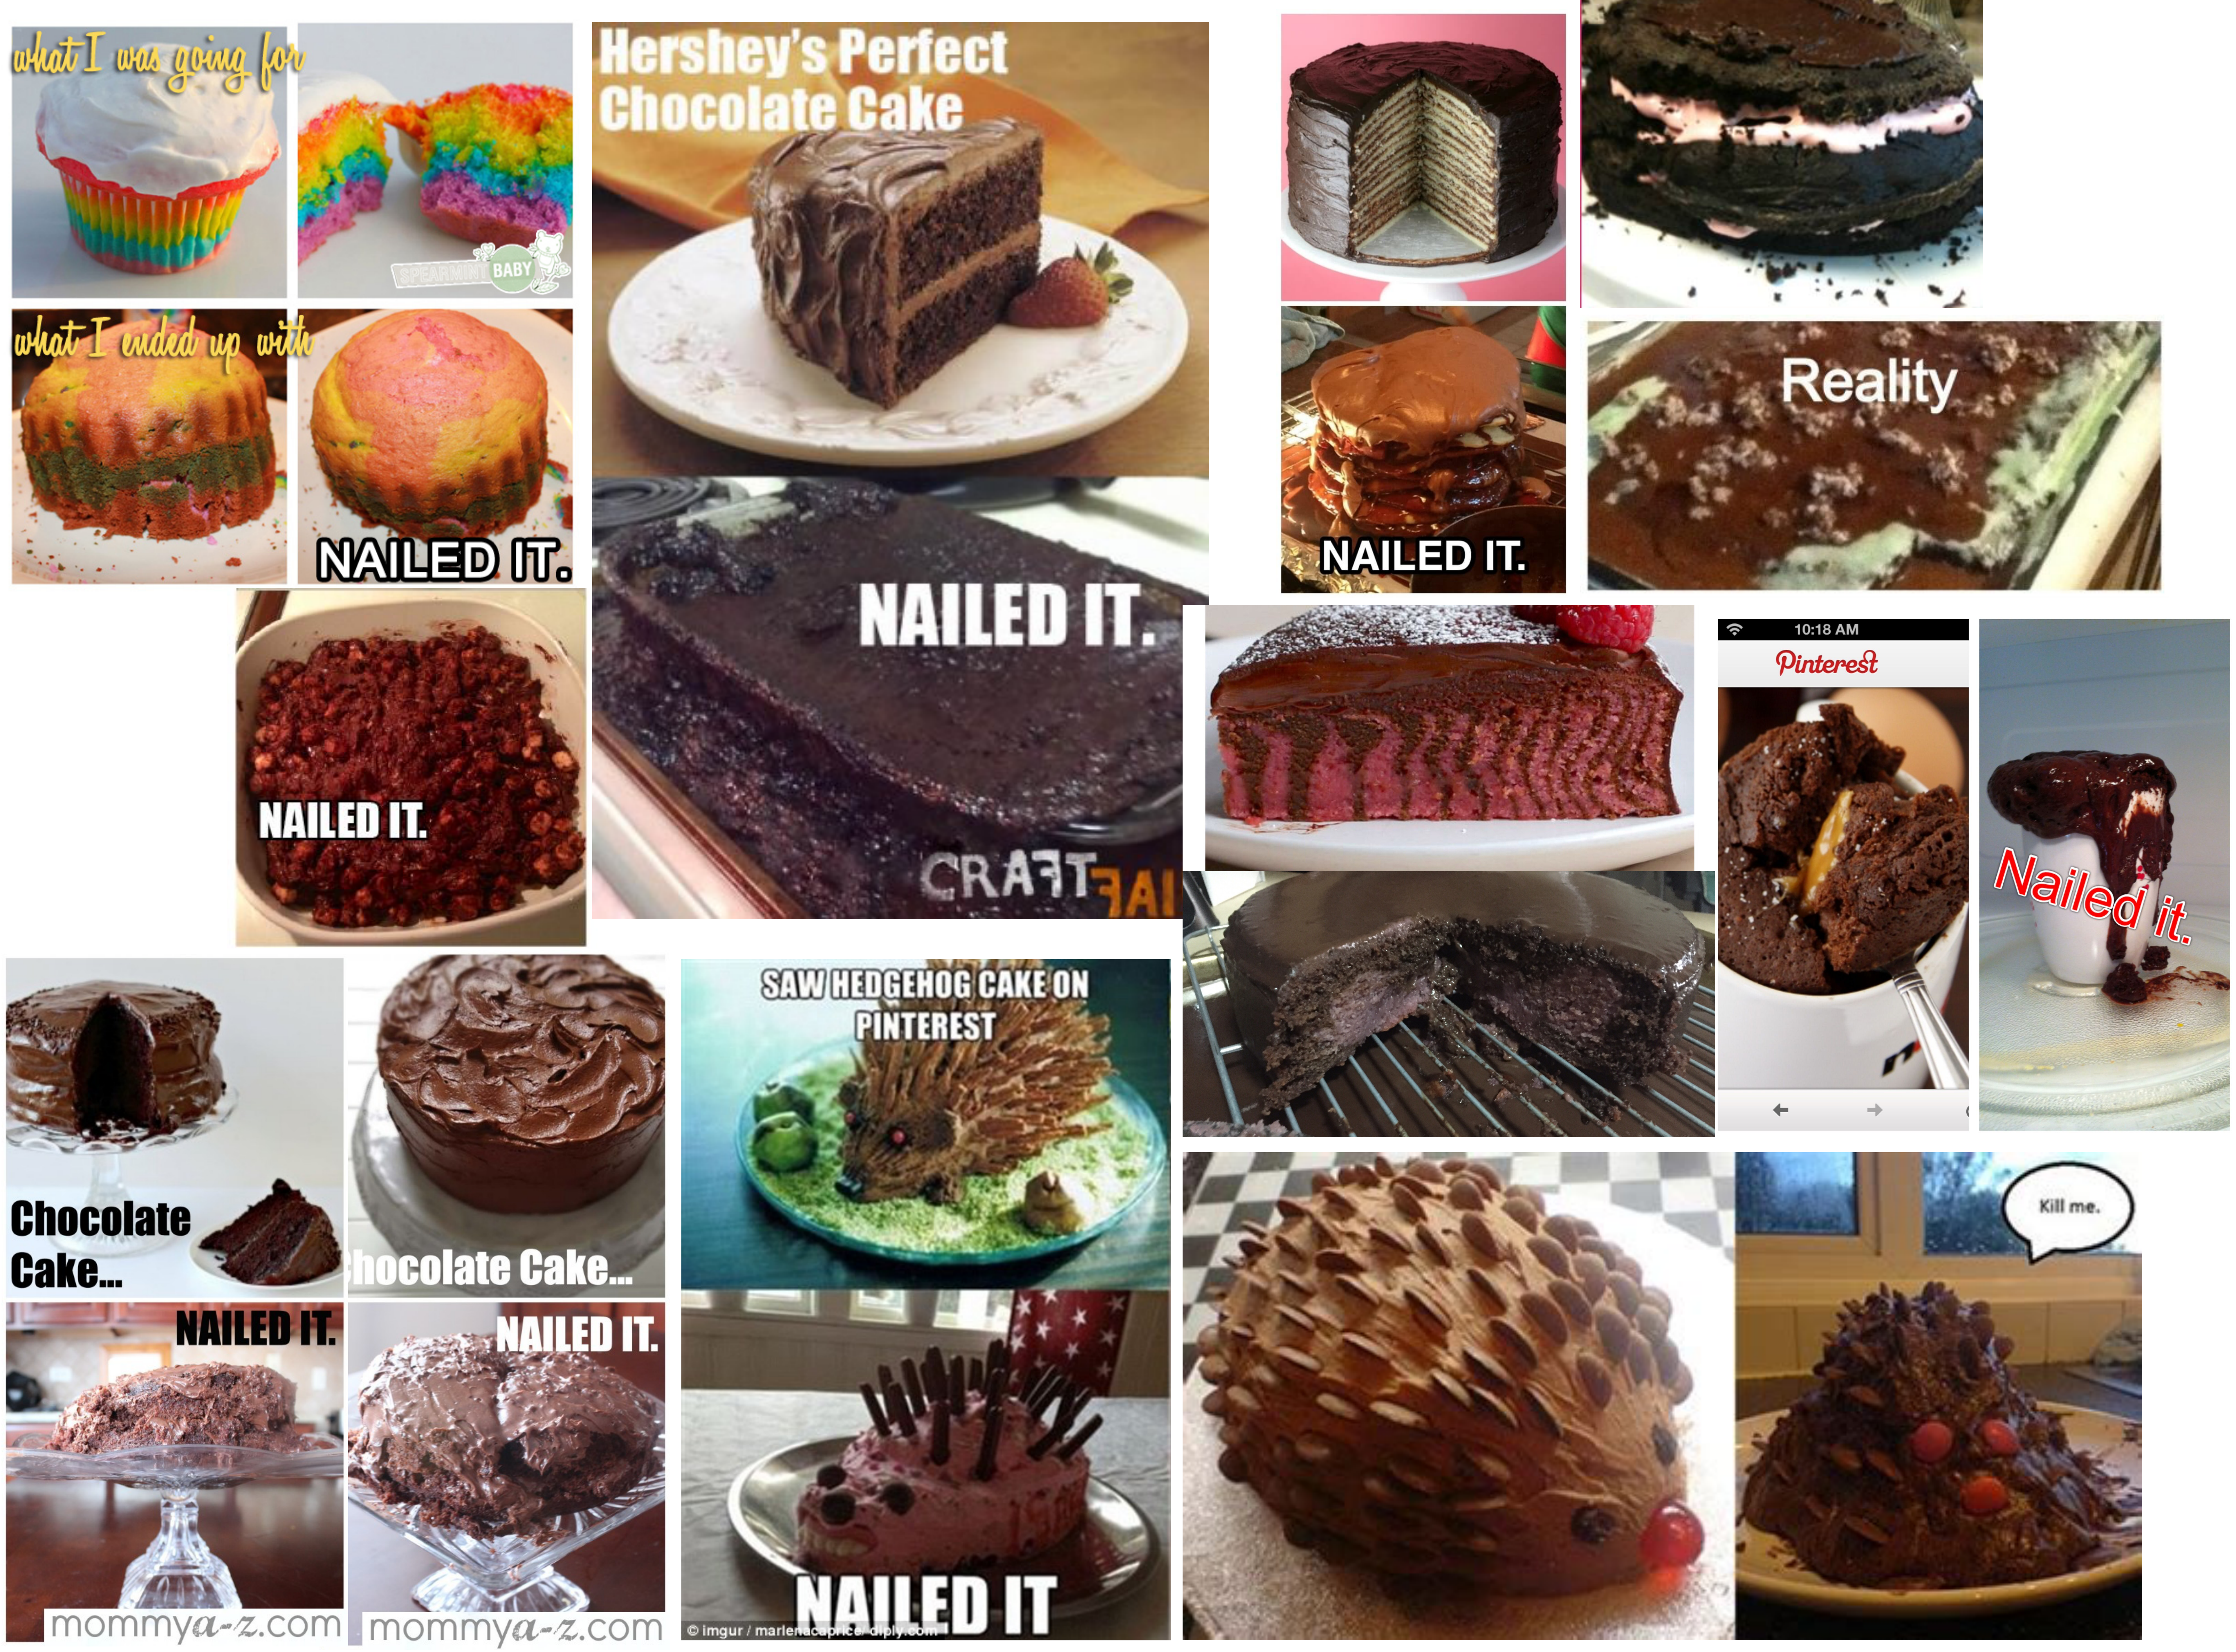
what I was going for
SPEARMINT BABY
what I ended up with
NAILED IT.
Hershey's Perfect Chocolate Cake
NAILED IT.
CRAFT
NAILED IT.
NAILED IT.
Reality
10:18 AM
Pinterest
Nailed it.
Chocolate Cake...
hocolate Cake...
NAILED IT.
NAILED IT.
mommya-z.com
mommya-z.com
SAW HEDGEHOG CAKE ON PINTEREST
NAILED IT
Kill me.

But is the converged model a *good* one?

- Goodness-of-fit (GOF) evaluates whether the simulated networks are similar to the observed one…
- In terms of statistics that are *not* explicitly modelled
  - degree distribution
  - triad census
  - geodesic distances
- Why does it have to be *other* statistics?
- GOFs can be considered equivalent to an R^2 statistic in regression models, though χ^2 and F tests are not available

Goodness-of-fit diagnostics

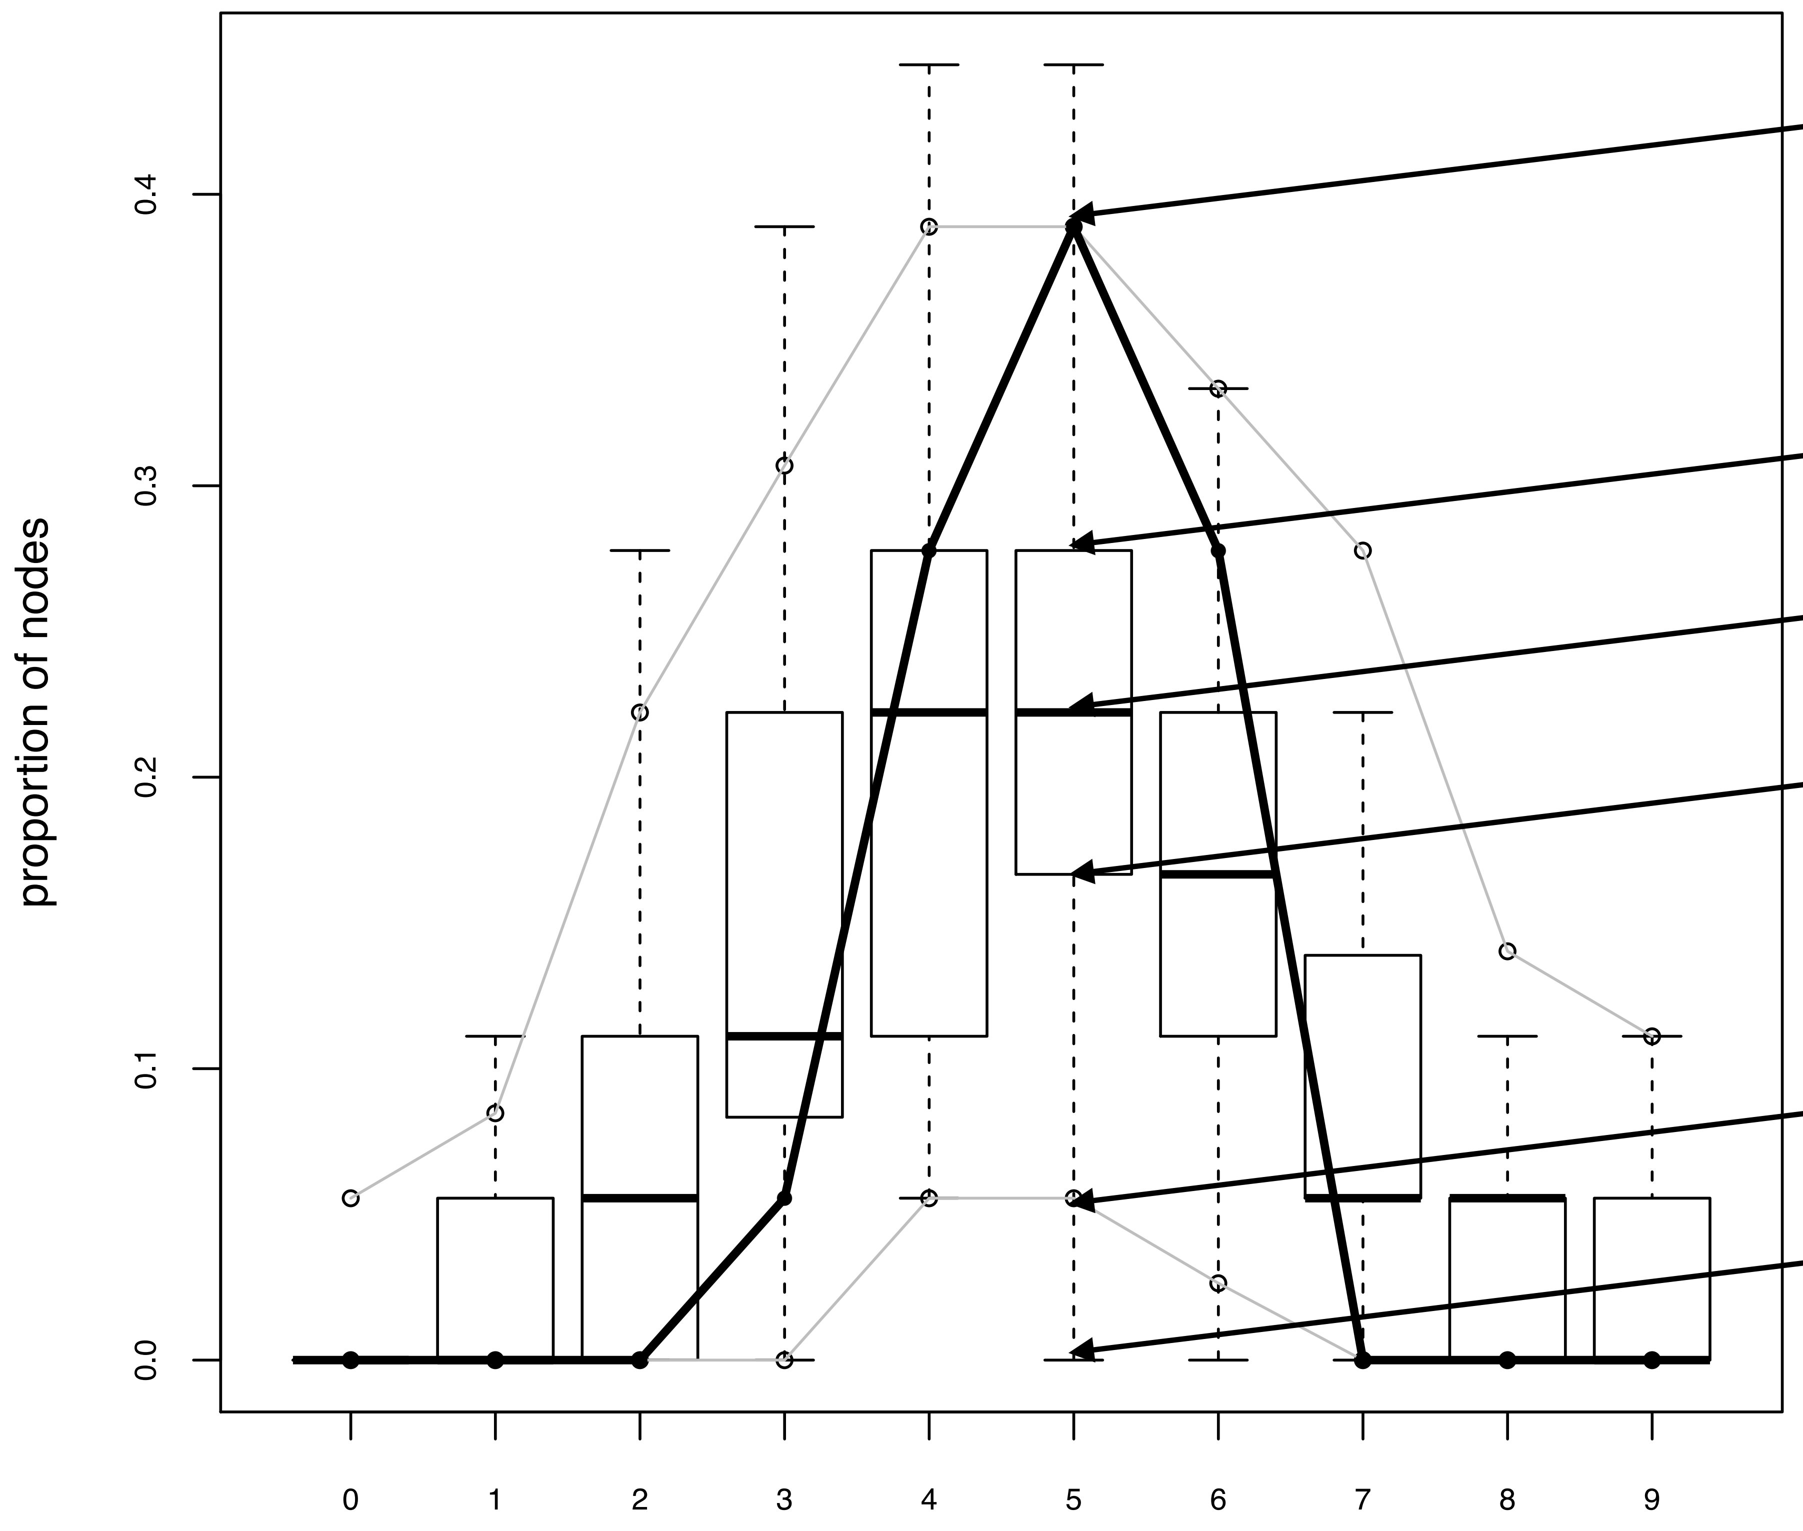

Thick line observed statistics

Box plots show distribution of statistic for simulated networks

Median

25th Percentile

95%

IQR

Goodness-of-fit diagnostics

Overestimating low degrees

Overestimating high degrees

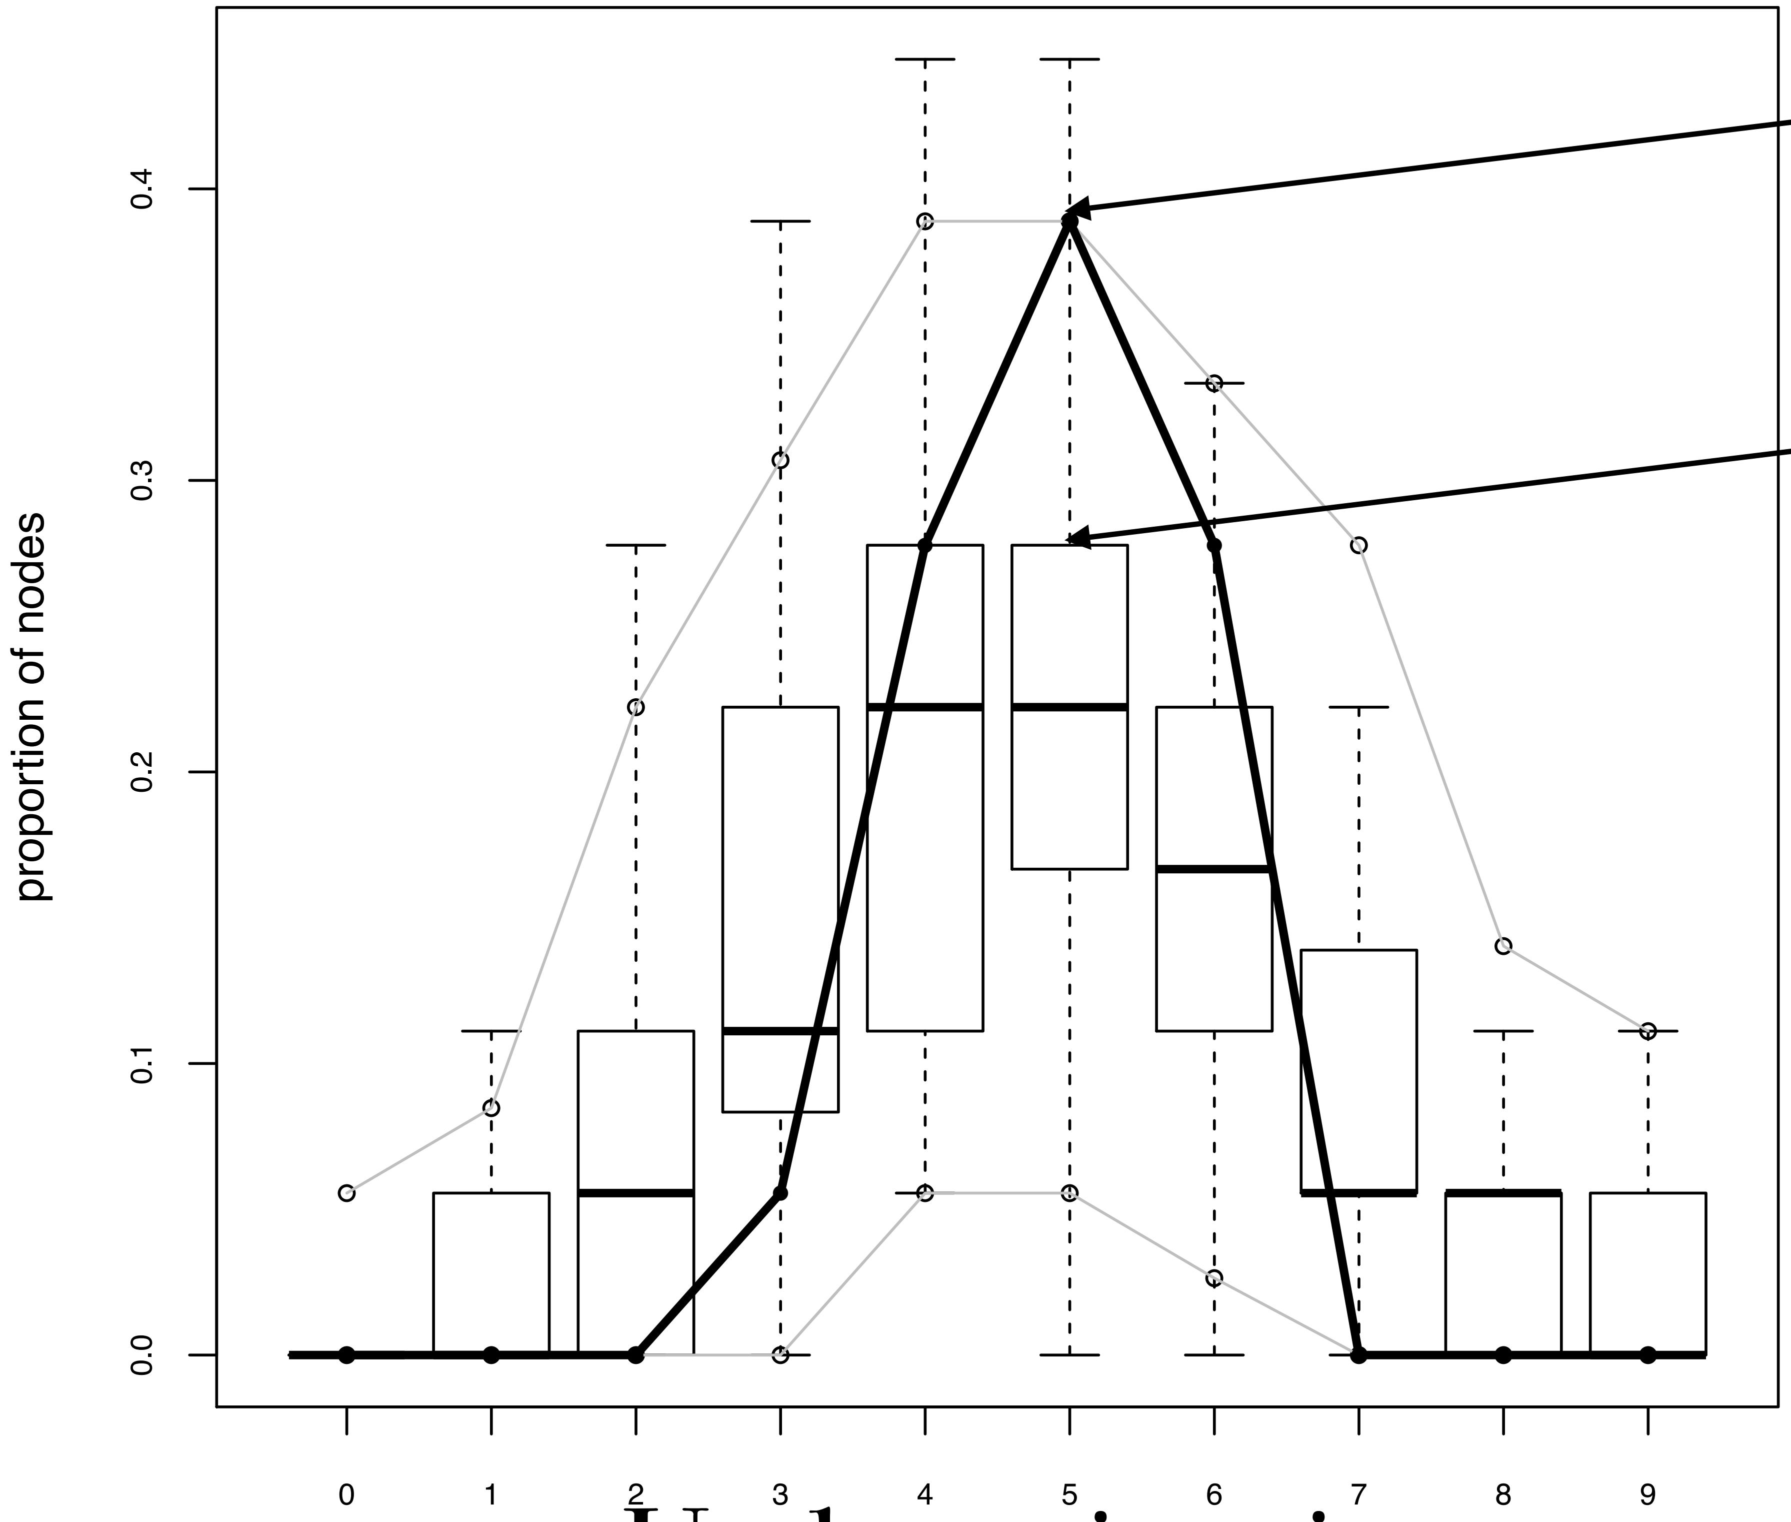

Underestimating medium degrees

Thick line observed statistics

Box plots show distribution of statistic for simulated networks

Goodness-of-fit diagnostics

Overestimating low degrees

Overestimating high degrees

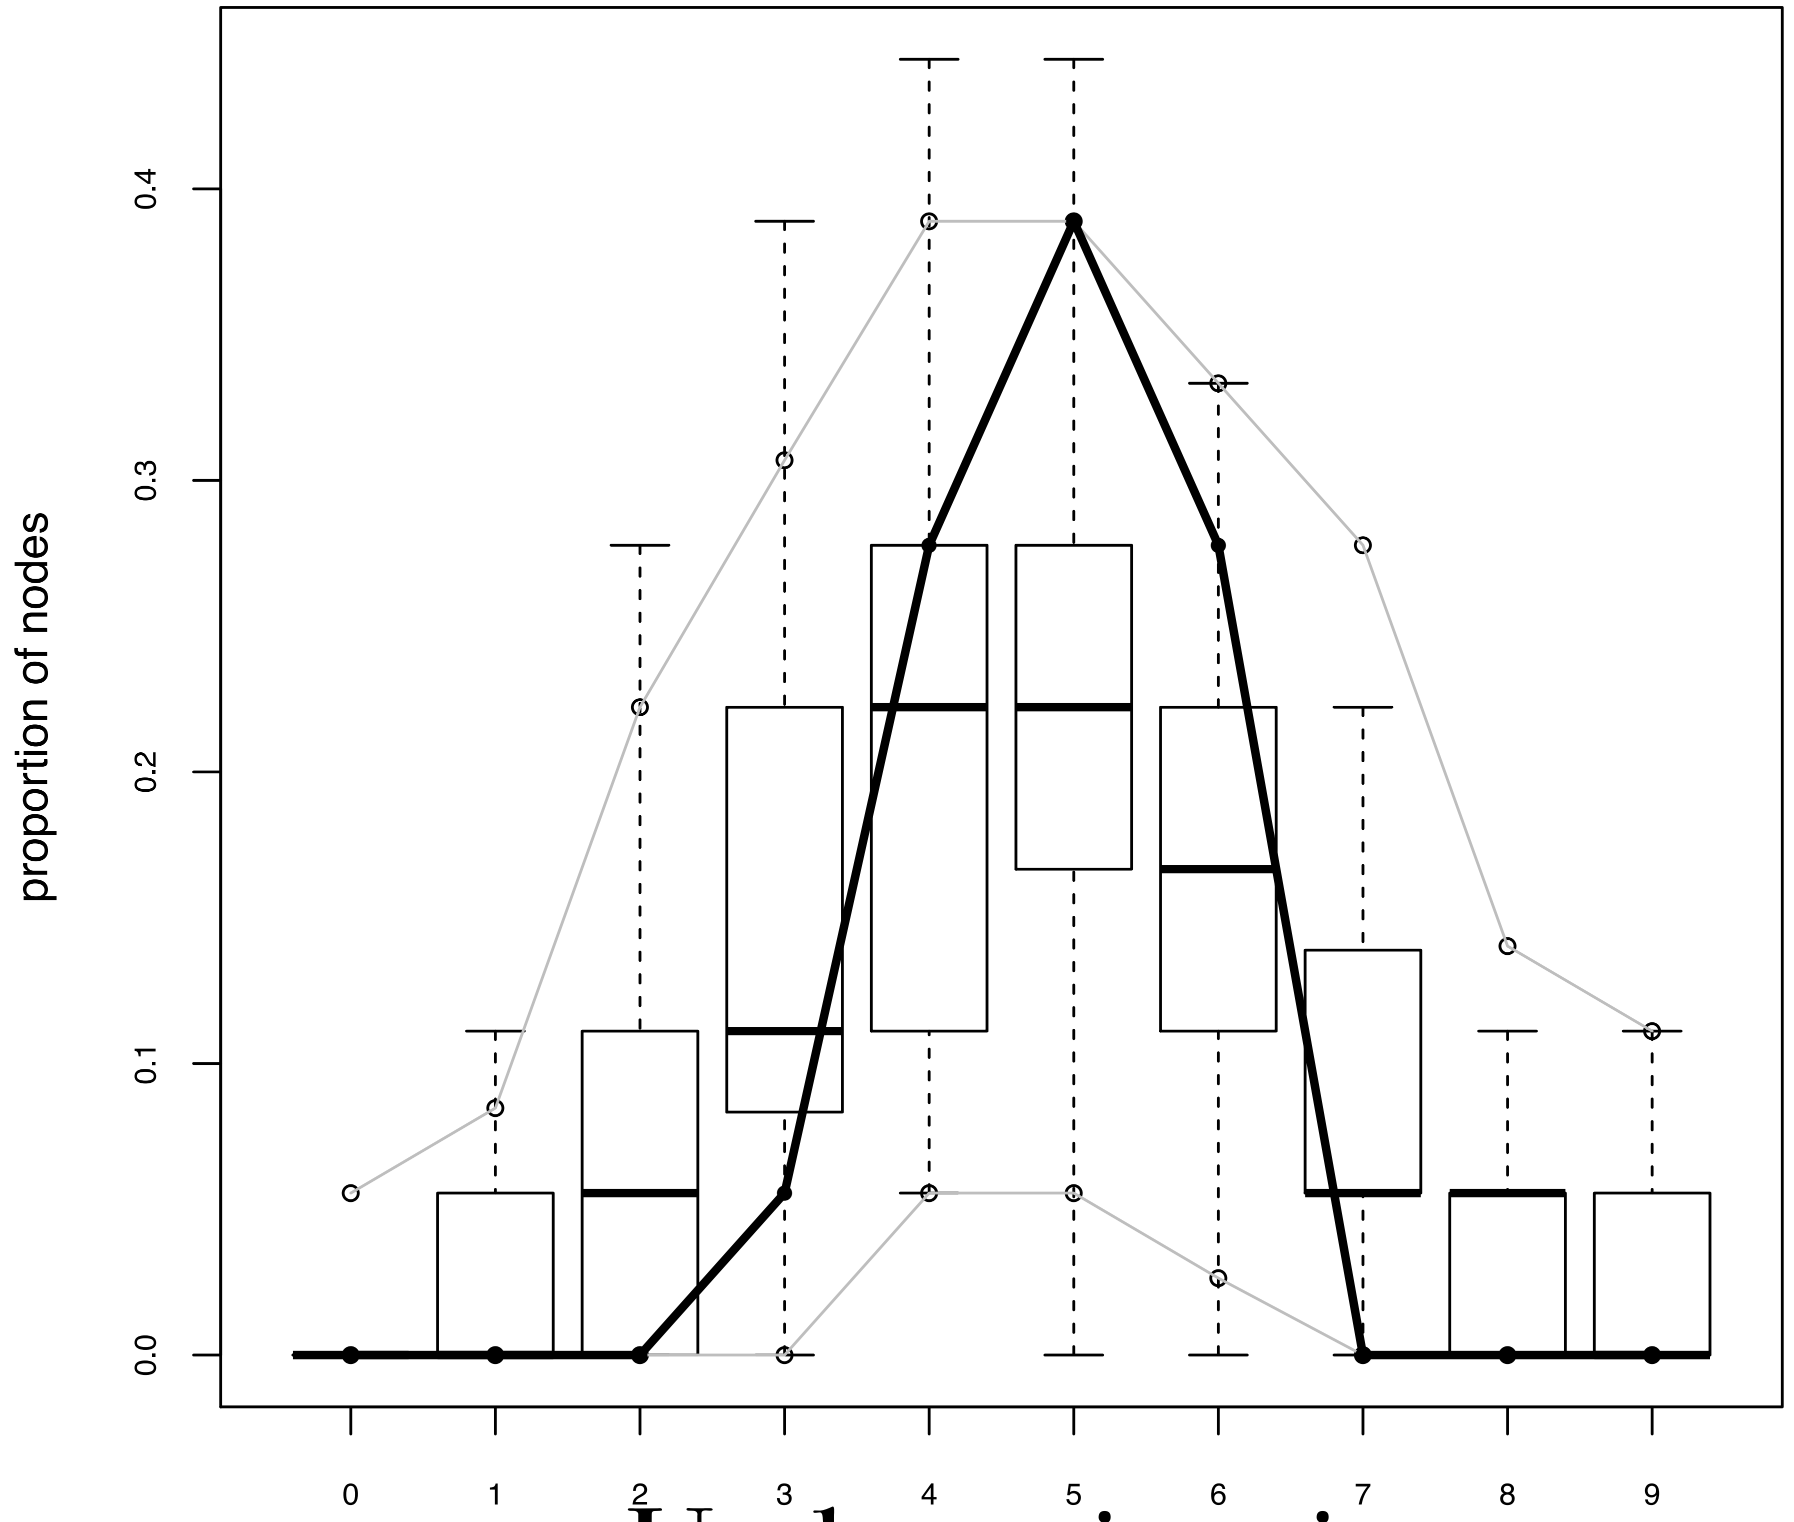

Underestimating medium degrees

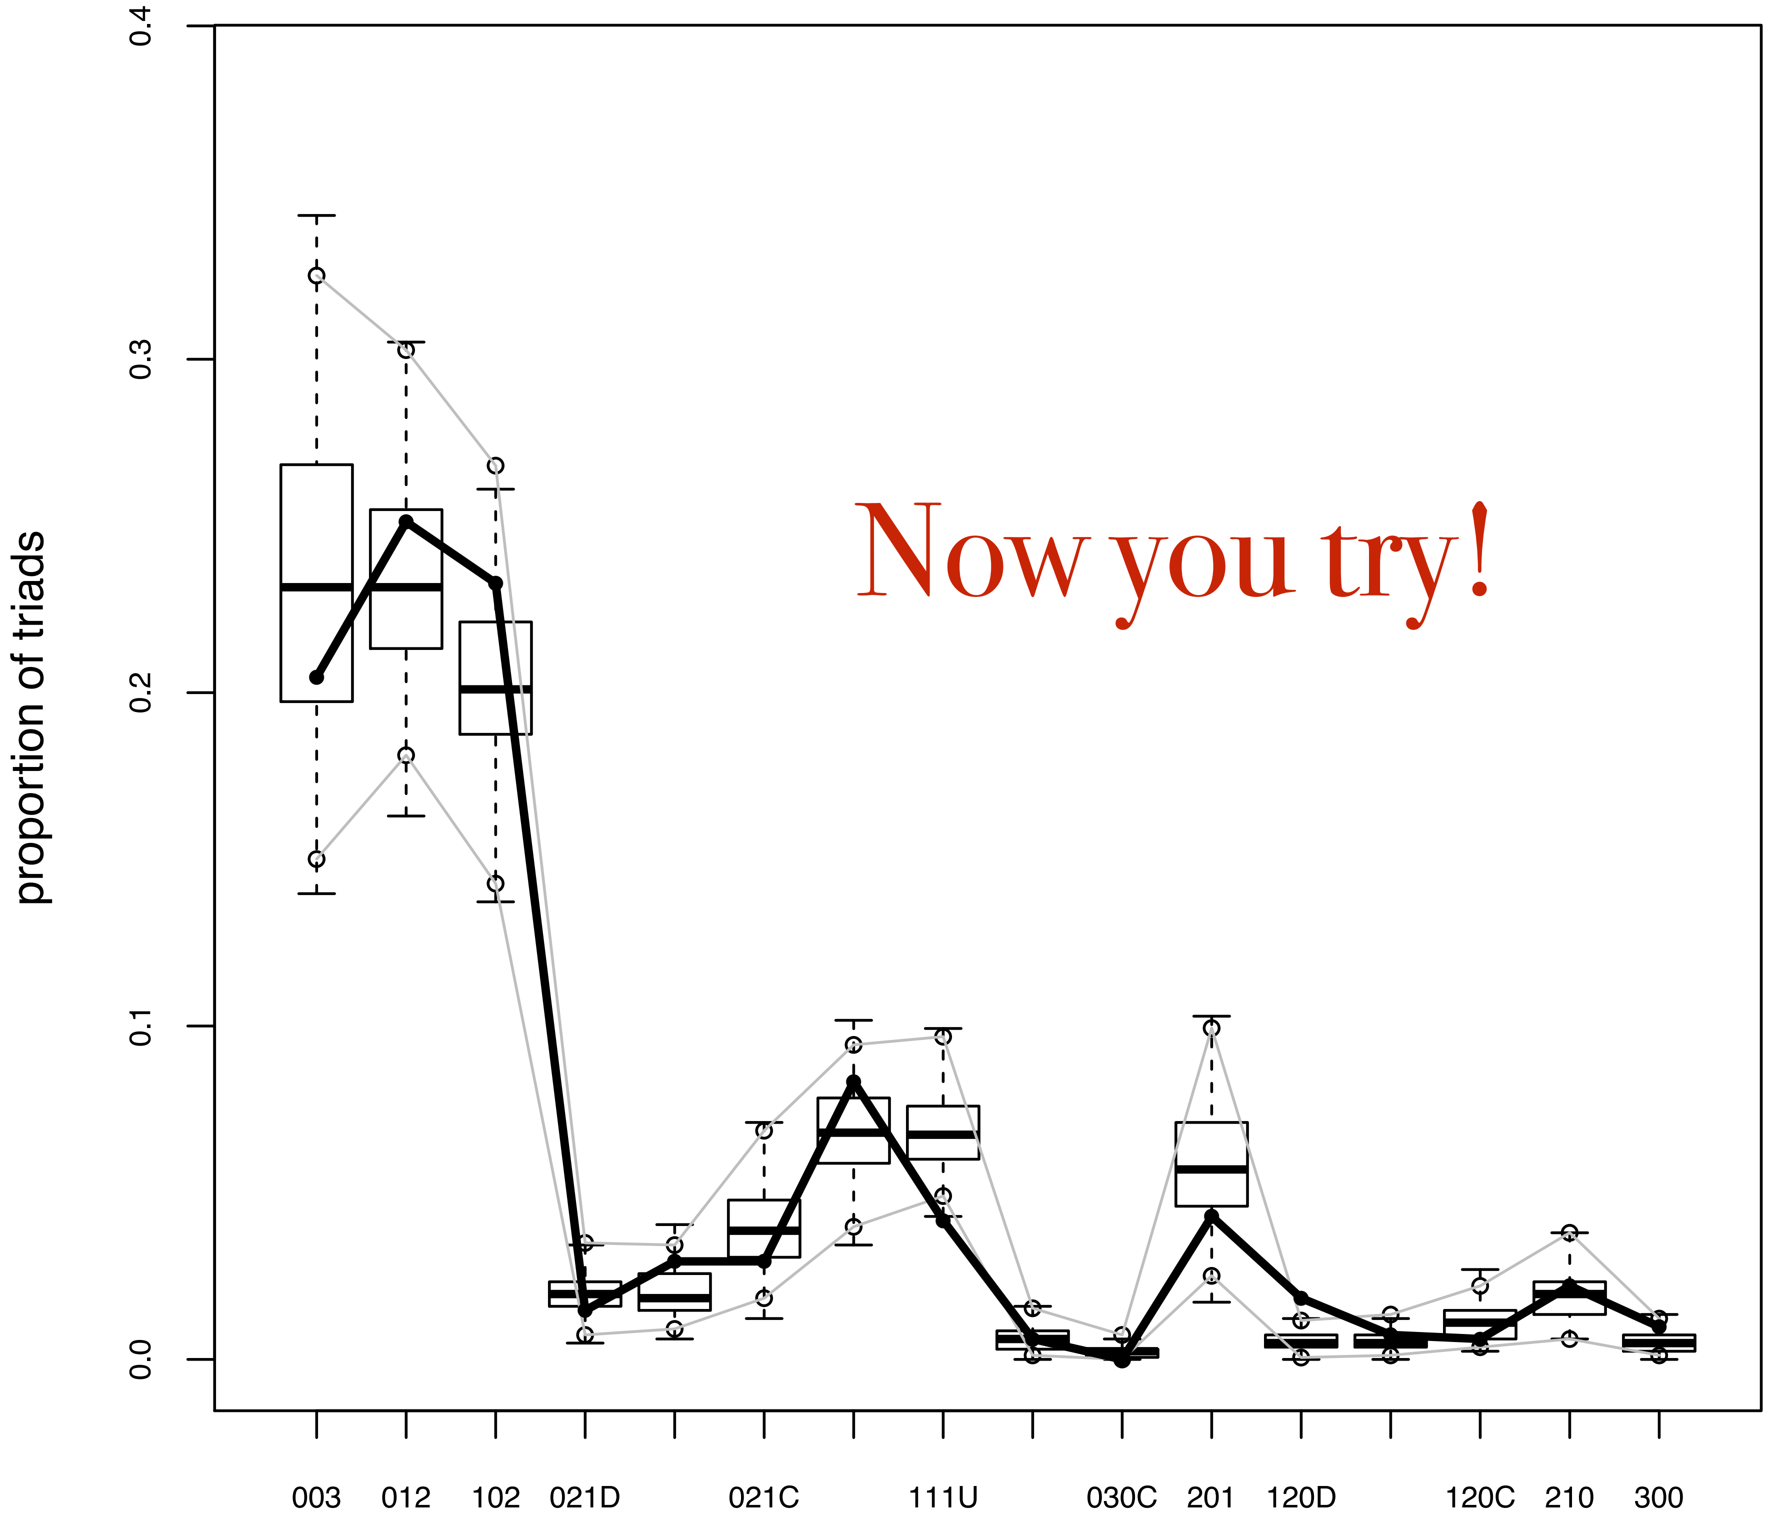


Fail
Cake

Summary: What ERGMs do

$$P(x;\theta) = \frac{\exp\left(\sum_k \theta_k z_k(x)\right)}{\kappa}$$

- Explains the probability of observing a specific network formation/tie in a network
- Dependence between observations taken into account through statistic functions z that represent local patterns
  - Density
  - Reciprocity
  - Homophily
  - Transitivity
  - Similar institutional portfolios, ...

Why ERGMs?

- ERGMs increasingly **understood** (sociology, political science, economics)
- ERGMs increasingly **used** (sociology, political science, economics)
- ERGMs increasingly **useful** (directed, bipartite, multilevel, valued, longitudinal, actor attributes, missing data, snowball designs)

The real weapon…?

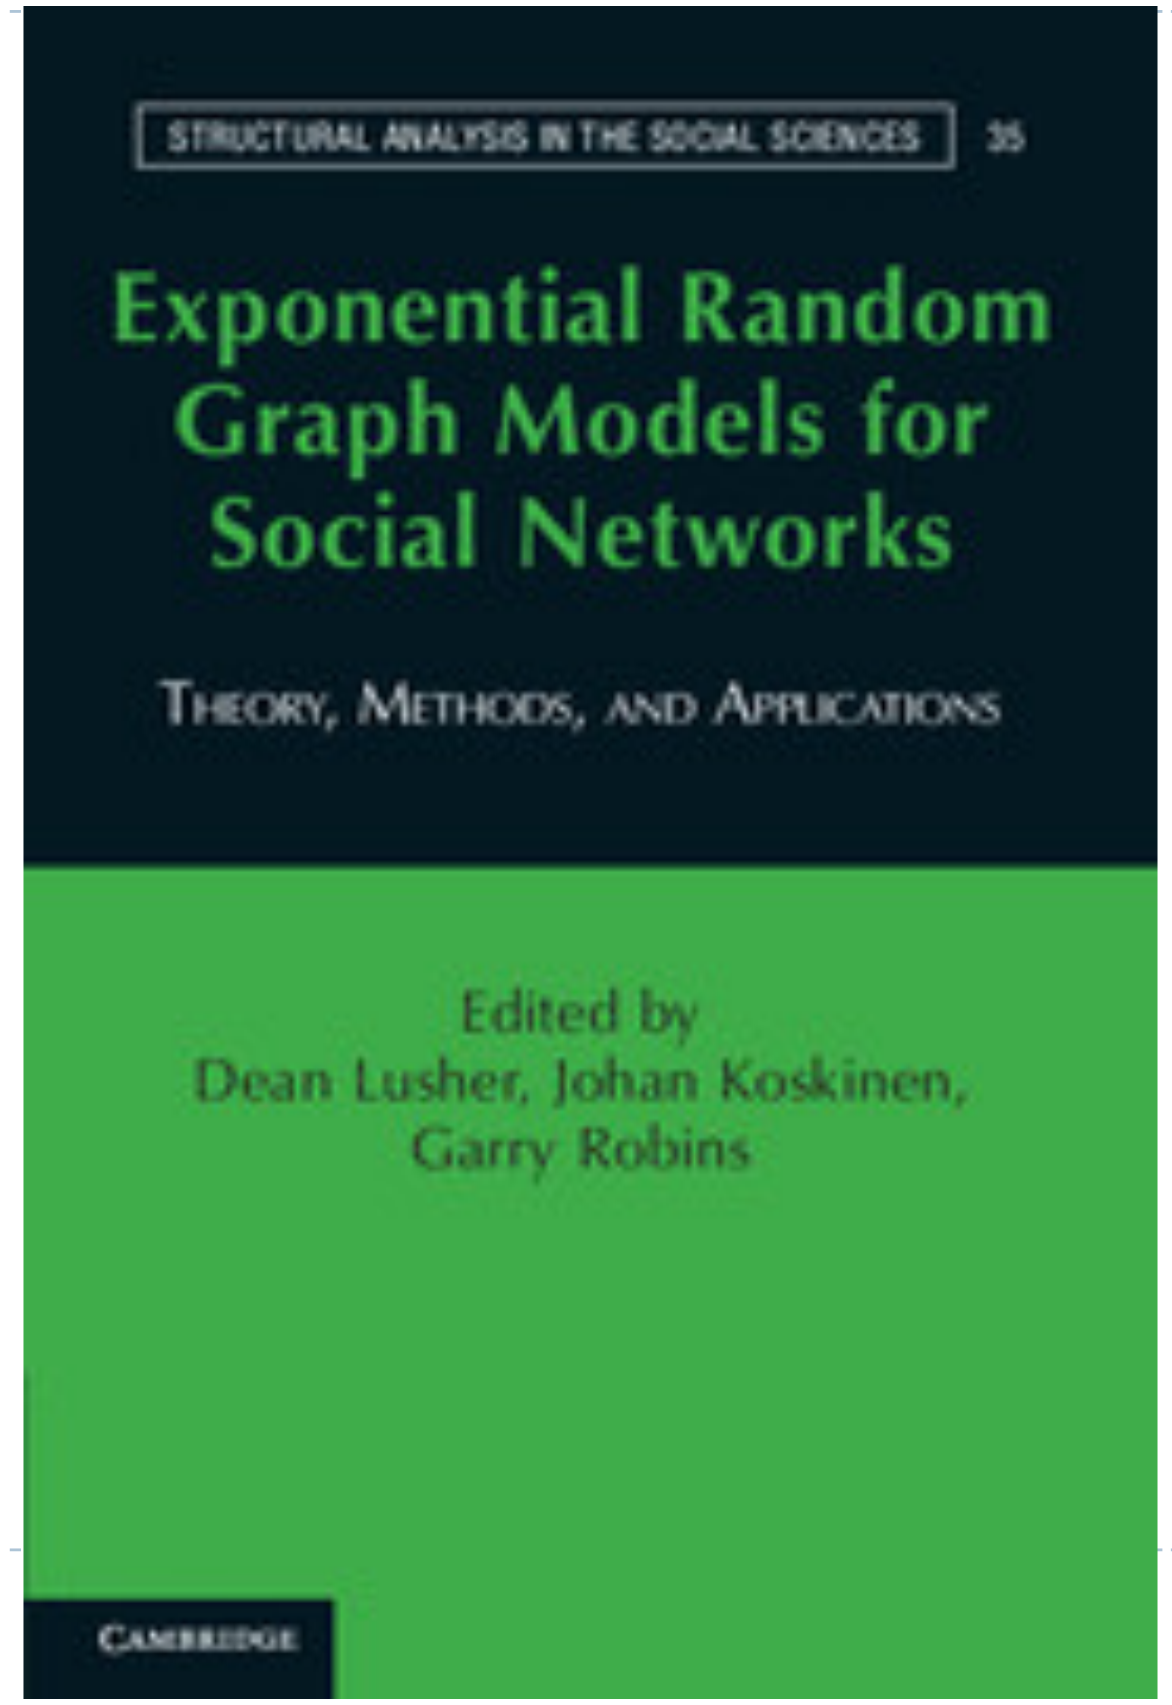

ERGM

Effects

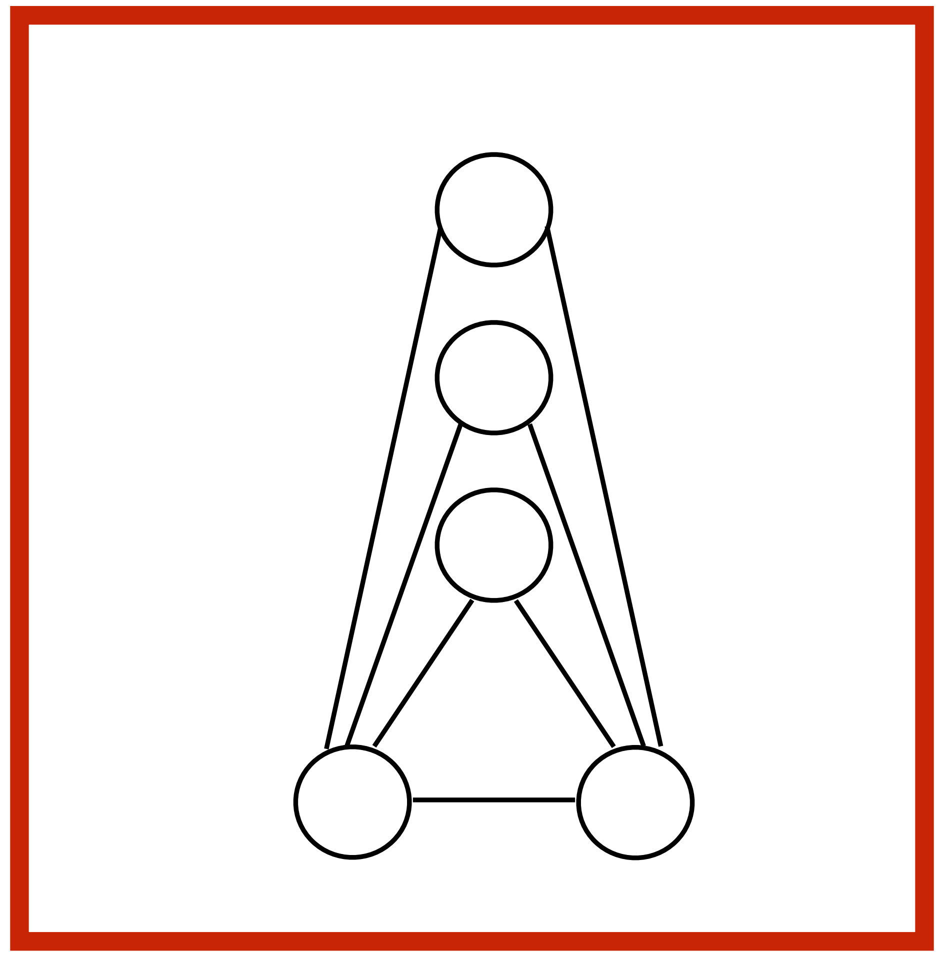

Explore various
effects available

Model

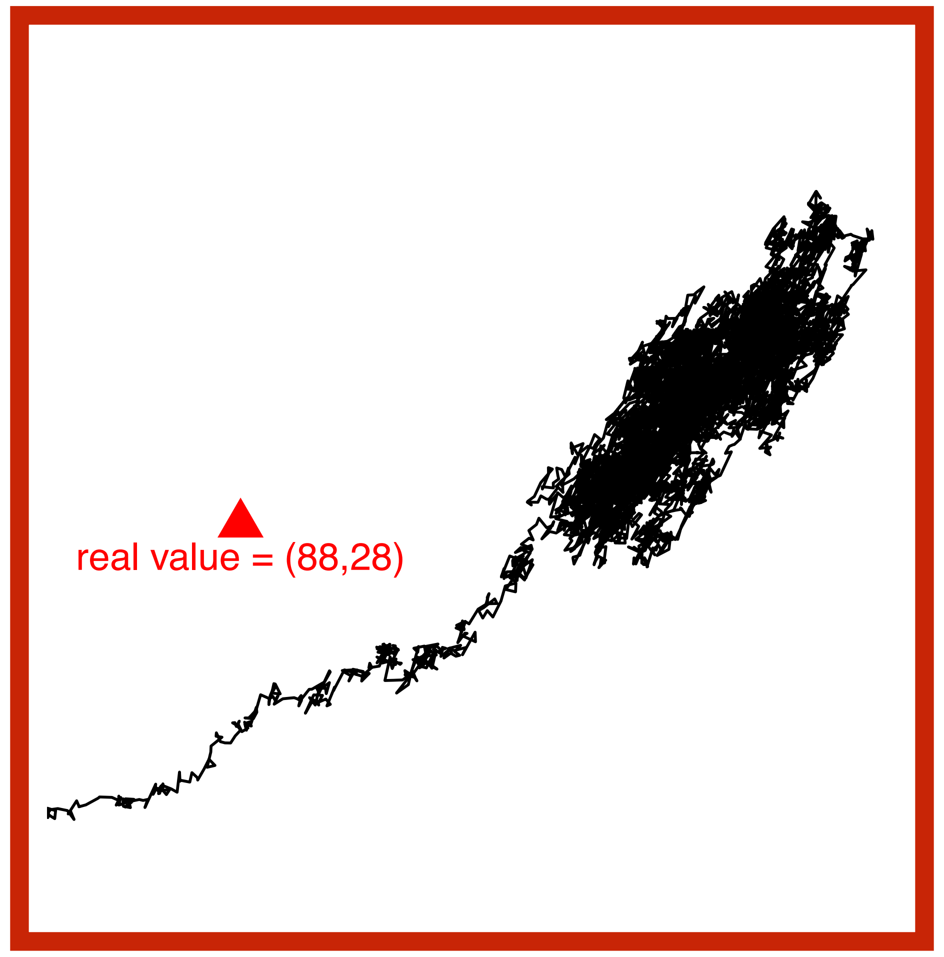

Understand model
intuition and estimation

Diagnostics

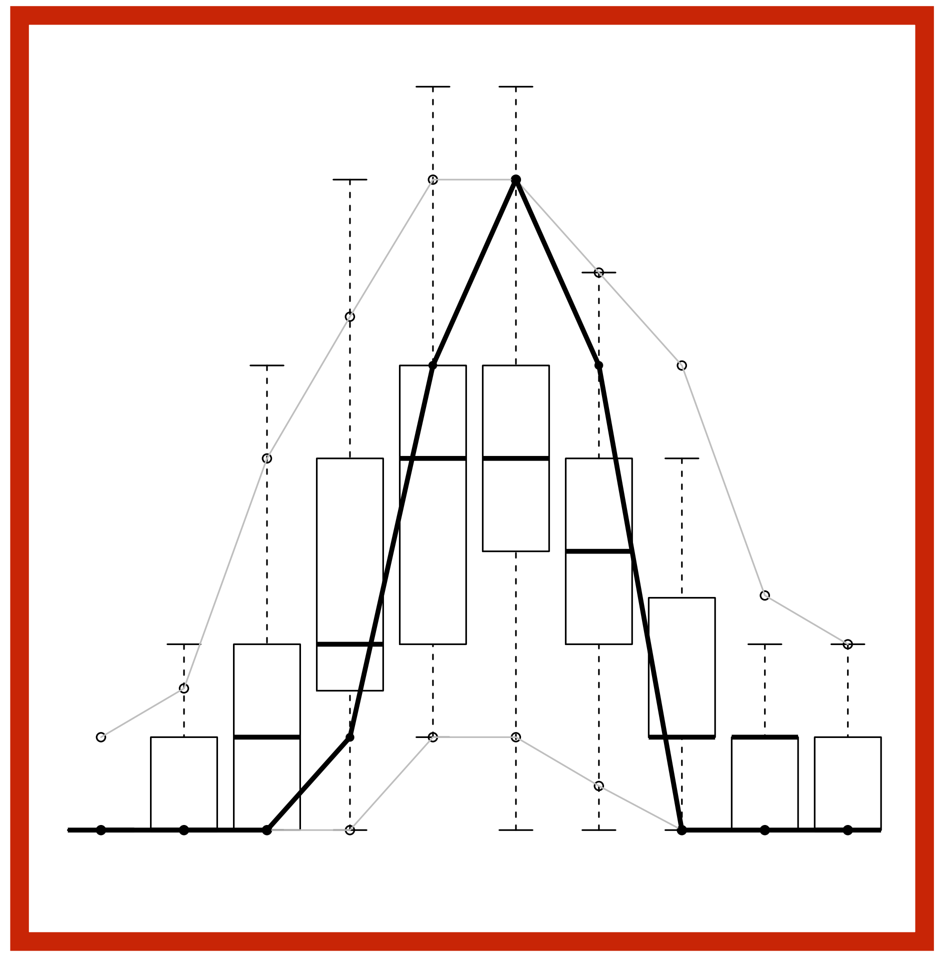

Recognise when a model
converges and fits


Dig in!