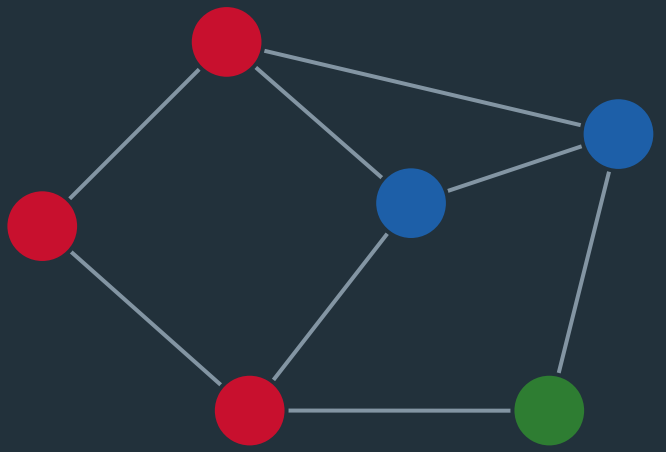
A first network model

Course C · Day 5, morning

Chapter 21: exponential random graph models · about three hours

Instructors are free to use, adapt and share these slides.



Today, morning: four blocks, one break

0:00	**1 Ties that depend on ties**	30 min
0:30	**2 The advice model in four steps**	55 min
1:25	Break	15 min
1:40	**3 Checking the simulation and the fit**	60 min
2:40	**4 What the model says, and the afternoon**	15 min

Day 4 left the tendencies tangled together

What Day 4 showed

The advice network's reciprocity (0.274) and transitivity (0.329) are far above random networks of the same density. Advice and friendship correlate at 0.418, and friends and colleagues in one office ask each other for advice.

What it could not do

Weigh the tendencies against each other. QAP has no term for a lawyer who asks a colleague because that colleague asked her the day before: her request depends on another advice tie.

Today

An ERGM puts the dependence among ties into the model, so that reciprocity, closure, homophily and status are estimated together, each net of the others.

The question: which local tendencies are associated with the observed advice network, and how strongly, once the others are taken into account?

Before we start

```
. which nwergm
. nwwebuse lazega, nwclear
. nwsummarize lazega_adv
```

nwergm is a separate package, installed after the core (Getting started).

Check at the start that every participant can run nwergm: a missing package costs ten minutes later.

Law firm

71 lawyers; advice network lazega_adv with 892 arcs; office, practice and status.

Every fit takes seed()

Results then repeat exactly, on every machine.

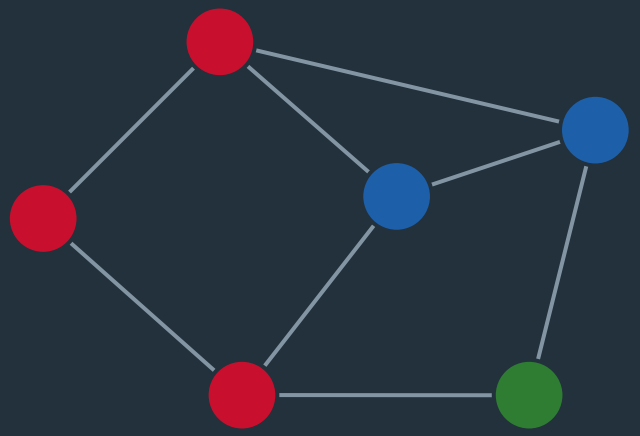
Open the do-file

The book's do-file for Chapter 21 holds every command of this morning.

1

Ties that depend on ties

30 min · Chapter 21: from dyads to dependence; the change statistic; two ways of estimating



Four tendencies compete for the same ties

Tendency	Claim about a tie from i to j	ERGM term
Reciprocity	likelier when the tie from j to i exists	`mutual`
Closure	likelier when i and j share a partner	`gwesp(.5)`
Homophily	likelier when i and j share an attribute	`nodematch(office practice)`
Status	likelier toward, and from, the senior	`nodeofactor(status) nodeifactor(status)`

Each predicts a local pattern; the observed network shows the sum of all of them. Lawyers of one office share many colleagues, so a tie within an office also tends to be mutual and to close a triangle.

Where QAP stops and an ERGM begins

Reciprocity

The tie i to j depends on one specific other tie, j to i.

Transitivity

The tie i to j depends on the joint presence of two others, through a shared partner k.

Popularity

A tie to j is likelier the more ties j already has: it depends on the whole configuration at j.

Homophily

Reduces to the pair and its attributes: QAP already handles it, with no ERGM needed.

The first three cannot be written as a function of i, j and their attributes alone.

The network is the outcome

$$P(Y = y) = \exp\{ \theta_1 s_1(y) + \theta_2 s_2(y) + \dots \} / \kappa(\theta)$$

s(y): counts of patterns

Edges, mutual pairs, closed triads, ties within an office. A count of closed triads is a joint property of several ties: that is where dependence enters.

θ: one weight each

A positive weight makes networks with more of that pattern more probable than networks with less.

κ(θ): the normalizer

A sum over every possible network on the same nodes. For 71 lawyers there are 2 to the power 4,970 directed networks: it cannot be computed.

A coefficient is read through a change statistic

$$\text{logit } P(Y_{ij} = 1 \mid \text{all other ties}) = \theta_1 \delta_1(y)_{ij} + \theta_2 \delta_2(y)_{ij} + \ldots$$

Open path B - A - C

Lawyer A is tied to B and to C, B and C are not tied. Toggling B-C closes the triad: the triangle count goes from 0 to 1, so $\delta = 1$ for this pair. A tie from B to a lawyer D with no shared contact closes nothing: $\delta = 0$.

Path A→B→C→D

Adding B→A, the reverse of the present A→B, turns one asymmetric pair into a mutual one: the change statistic of mutual is 1 for this toggle.

δ is how much a statistic changes if one tie is switched on, every other tie held fixed. Every term works this way; they differ only in which configurations make δ nonzero.

What an ERGM describes, and what it does not

It says
Which local configurations (mutual pairs, closed triads, ties within offices) are over-represented in the observed network, once the others are taken into account.

It does not say
Why they arise, or how the network changed. A reciprocity coefficient from one snapshot cannot say whether lawyers came to return each other's requests, or whether pairs who return requests are the ones who stayed in the data.

Tie-oriented

A probability over whole networks; each tie given the rest follows from it.

The SAOM differs

Actors change their own ties, step by step, between waves: this afternoon's option for panel data.

This is the standard form of each model: temporal ERGMs and single-network SAOMs exist, and are advanced.

Two ways of estimating

	MPLE	MCMLE
Idea	each dyad's conditional probability is a logistic observation	draw networks from the model with a Markov chain; move θ until simulated statistics match the observed ones
When	exact if the model has only edges and dyadic covariates	needed once mutual or gwesp make ties depend on ties
Cost	seconds for a thousand nodes	each iteration simulates networks; seconds to minutes here
Check	standard errors, deviance	estat mcmcdiag: autocorrelation, effective sample size

The coefficients are only as trustworthy as the simulation, which is why the diagnostics come before the interpretation.

Build a specification in stages

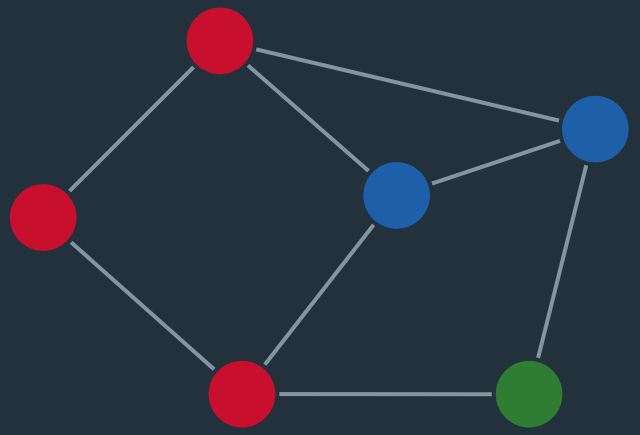
Stage	Adds	Estimation
1	`edges`	MPLE
2	`nodematch(), nodeofactor(), nodeifactor()`	MPLE (dyad-independent)
3	`mutual`	MCMLE
4	`gwesp(.5)`	MCMLE
5	diagnostics: estat mcmcdiag	
6	comparison and fit: estimates stats, lrtest, estat gof	

One dependence term at a time, so that a problem can be traced to the term that caused it. Edges alone is never the model of interest, but it confirms that the network and the software behave.

2

The advice model in four steps

55 min · Chapter 21: edges, attributes, reciprocity, closure



Who asks whom for advice in the firm

892

advice arcs among 71 lawyers: an arc from i to j means i names j as a source of advice

0.179

density; of 2,485 pairs 192 are mutual, 508 asymmetric and 1,785 null

0.274

reciprocity

Two questions. Do lawyers seek advice from colleagues in their own office and practice, and do partners attract and give advice more than associates? And once these are controlled, is advice reciprocated and does it cluster into closed triads beyond what the attributes imply?

```
. nwwebuse lazega, nwclear
```

String attributes office, practice and status enter as they are; levels in alphabetical order.

Step 1: edges alone is the density

```
. nwergm lazega_adv, edges seed(12345)

Estimation:            =  MPLE

------------------------------------------------
  lazega_adv | Coefficient  Std. err.      z    P>|z|
-------------+----------------------------------
       edges |  -1.519896   .0369634   -41.12   0.000
------------------------------------------------
```

−1.52

= ln(0.1795 / 0.8205), the log-odds of the density

With edges alone every possible arc is independent and equally likely: all 71 lawyers are alike and advice ties are scattered at random. This is the null model that the later steps improve on.

- **Log-likelihood −2338.88**

  AIC 4679.8; the model of later comparisons starts here.

- **Estimation: MPLE**

  exact for a model without tie dependence.

Step 2: homophily and status

```
. nwergm lazega_adv, edges nodematch(office practice) ///
>     nodeofactor(status) nodeifactor(status) seed(12345)

--------------------+-------------------------------------
              edges |  -4.095448   .1309373   -31.28   0.000
   nodematch_office |   1.422373   .0887433    16.03   0.000
 nodematch_practice |   1.037945   .0837376    12.40   0.000
nodeofactor_status_2 |   1.125832   .0842818    13.36   0.000
nodeifactor_status_2 |   .7831546   .0819723     9.55   0.000
```

Term	Odds ratio
same office	4.1
same practice	2.8
partner sends	3.1
partner is asked	2.2

exp(1.42), exp(1.04), exp(1.13), exp(0.78)

All four are positive. The model is still dyad-independent, so MPLE is exact and a logistic regression of the 4,970 ordered pairs would give the same results: these are effects QAP could already give.

- **nodematch(office practice)**
  same office, same practice
- **nodeofactor(status)**
  partners send more
- **nodeifactor(status)**
  partners are asked more

Step 3: reciprocity, and the switch to MCMLE

```
. nwergm lazega_adv, edges mutual nodematch(office practice) ///
>     nodeofactor(status) nodeifactor(status) seed(12345)

Estimation:            =  MCMLE

               edges |  -3.961329   .1247155   -31.76   0.000
              mutual |   1.083464   .1279072     8.47   0.000
    nodematch_office |   1.198955   .0882039    13.59   0.000
(...)
MCMLE converged after 7 iteration(s).
```

3.0

odds ratio of mutual, exp(1.083)

The z column can be read here: it comes from a likelihood that contains the dependence terms, so it is valid as far as the model is adequate.

It says

Compared with an otherwise identical pair, an advice arc i to j has about 3.0 times the odds when the arc j to i exists (mutual 1.08).

It does not say

That an arc is three times as probable: near a probability of 0.2 the ratio of probabilities is nearer 2.

Step 4: closure with gwesp

```
. nwergm lazega_adv, edges mutual nodematch(office practice) ///
>     nodeofactor(status) nodeifactor(status) gwesp(.5) seed(12345)

                edges |  -4.662469   .1821854   -25.59   0.000
               mutual |   1.031625   .1244439     8.29   0.000
    nodematch_office |   1.016895   .0839039    12.12   0.000
  nodematch_practice |   .7571778   .0784148     9.66   0.000
nodeofactor_status_2 |   .8242996   .0907388     9.08   0.000
nodeifactor_status_2 |   .4019506   .0892556     4.50   0.000
           gwesp_.5 |   .5597618   .1091626     5.13   0.000
```

gwesp(.5)

Counts closed two-paths with diminishing returns: each further shared contact adds less. It replaces the raw triangle count, which tends to give degenerate models.

0.56

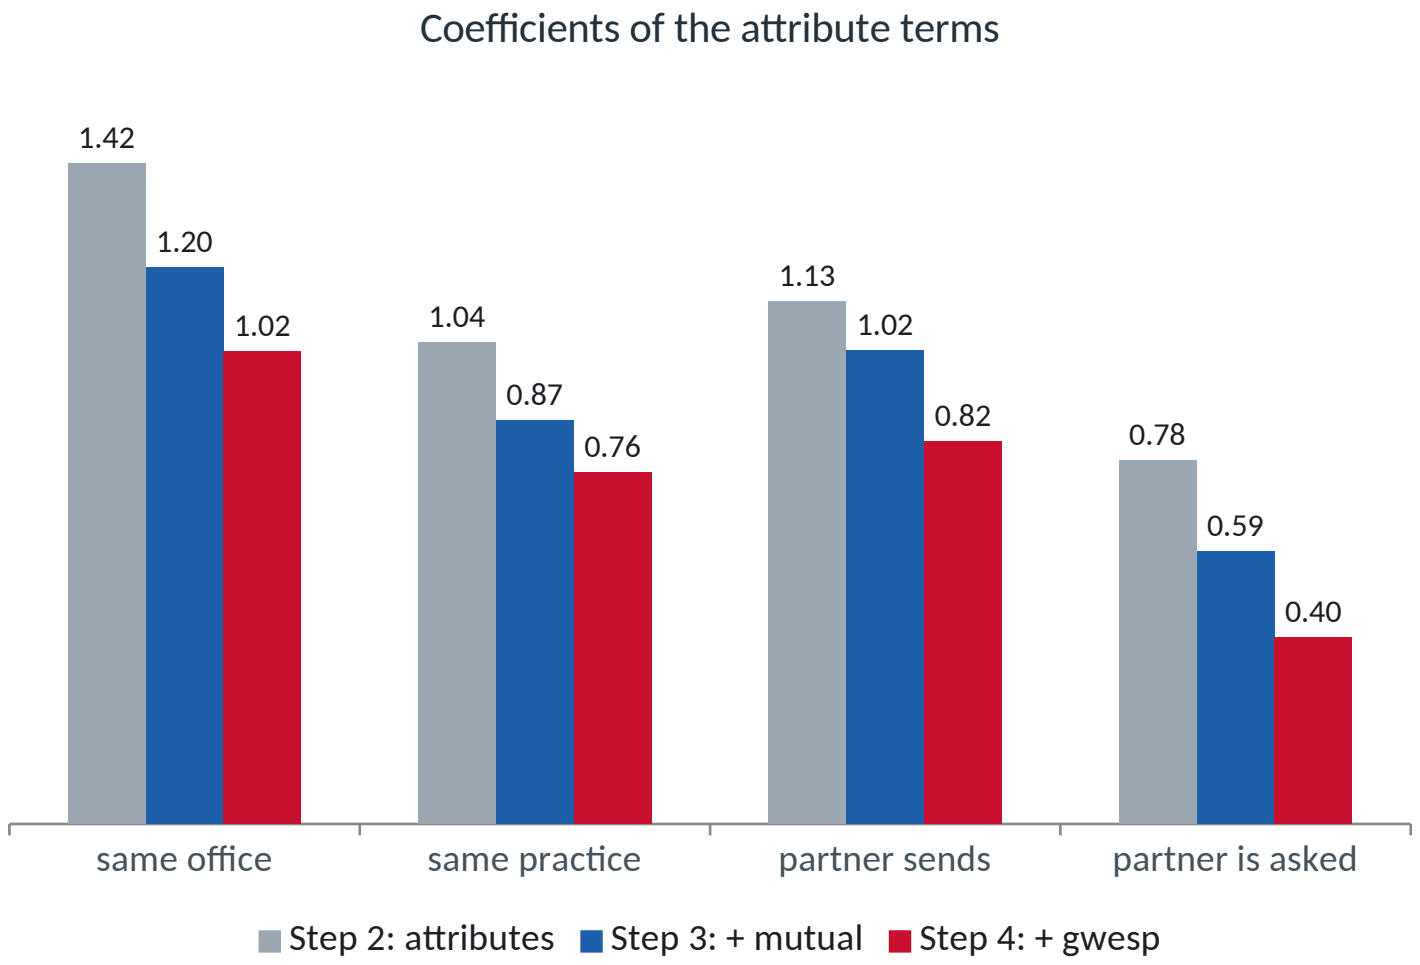
gwesp coefficient (standard error 0.11)

−1954.3

log-likelihood, from −1971.8 without gwesp

An advice arc is likelier when it closes paths through shared contacts. MCMLE converged after 5 iterations.

Coefficients move as structure enters



Ties within an office also tend to be mutual and to close triangles. **Without those terms, the attribute coefficients absorb the overlap.**

Coefficients of differently specified models are not on a common scale: the drop is not a share of homophily that turned out to be closure.

Reading the final model as odds ratios

```
. foreach t in mutual nodematch_office nodematch_practice ///
>     nodeofactor_status_2 nodeifactor_status_2 {
  2.     display %-24s "`t'" %8.2f exp(_b[`t'])
  3. }
mutual                    2.81
nodematch_office          2.76
nodematch_practice        2.13
nodeofactor_status_2      2.28
nodeifactor_status_2      1.49
```

- **gwesp has no single odds ratio**
  Its change statistic depends on how many shared contacts the arc would close.

- **How to report it**

  All three tendencies with their odds ratios, and in the same sentence that the model does not tell us why: the coefficients cannot distinguish cheap access through a shared office from a preference for similar colleagues.

Comparing the stages: likelihood, AIC, BIC

Model	df	ll(model)	AIC	BIC
covariates (step 2)	5	-2004.493	4018.985	4051.541
reciprocity (step 3)	6	-1971.793	3955.585	3994.652
closure (step 4)	7	-1954.322	3922.643	3968.222

```
. estimates stats covariates reciprocity closure
. lrtest reciprocity closure
Likelihood-ratio test
Assumption: reciprocity nested within closure
 LR chi2(1) =  34.94
Prob > chi2 = 0.0000
```

Each term earns its place

Each added term raises the log-likelihood and both criteria prefer each larger model. The chi-square reference is only approximate for one network and an estimated log-likelihood, but a statistic this large leaves no doubt.

Hands-on (10 min)

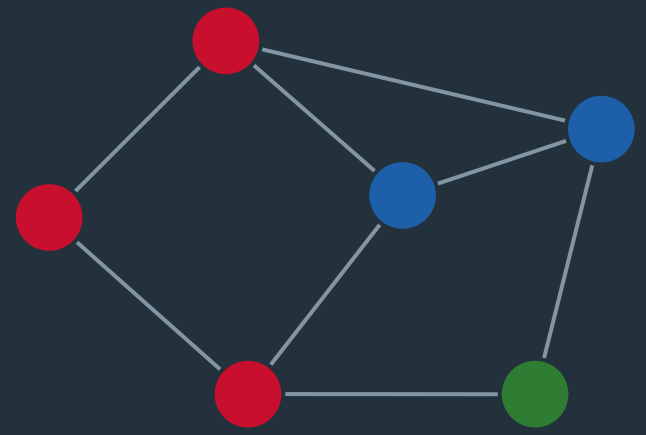
Ch. 21, Ex. 1

In the final advice model, the coefficient of nodematch_office is 1.017, an odds ratio of 2.76. Using the conditional log-odds form of this chapter, state precisely what this says about an advice arc from one lawyer to another of the same office. Then name two things it does not say, for example about two lawyers of the same office chosen at random from the firm, or about why lawyers of one office ask each other for advice.

Hint Write the conditional log-odds of the arc i to j for a pair in the same office and for an otherwise identical pair in different offices, with every other arc held fixed. The two differ by the coefficient times the change statistic of nodematch_office, which is 1 for a pair in the same office.

Break

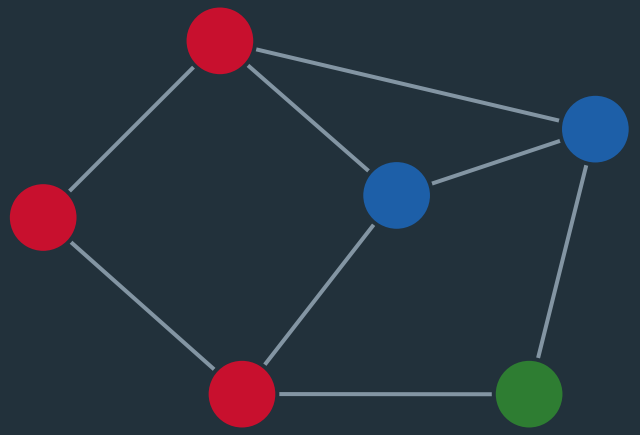
15 minutes



3

Checking the simulation and the fit

60 min · Chapter 21: estat mcmcdiag, estat gof



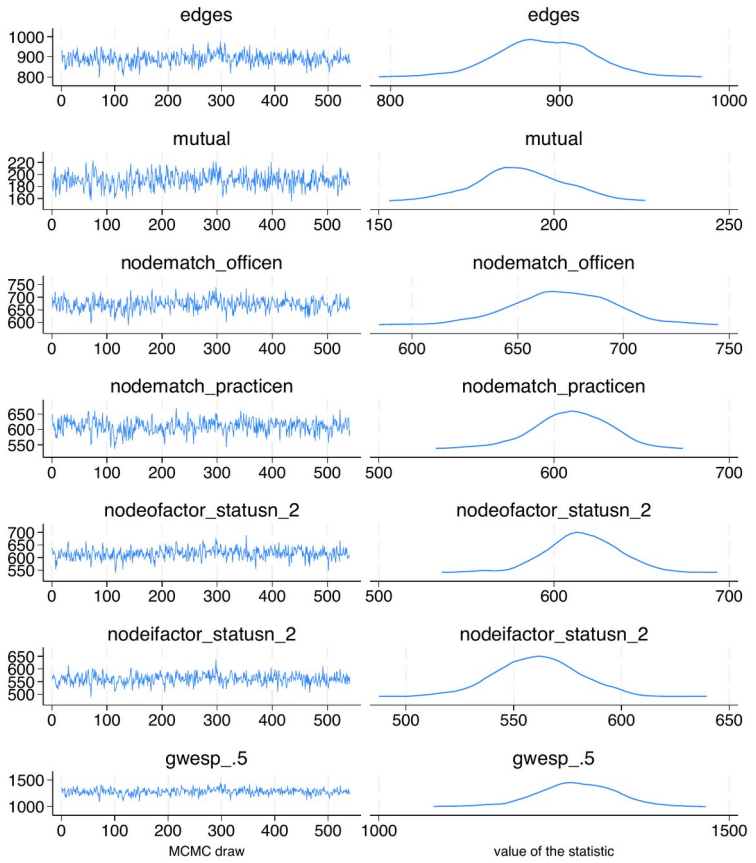
Converged is not a verdict: read the sample

```
. estat mcmcdiag, plot

           Statistic |  Autocorr       ESS   Geweke z
---------------------+-----------------------------
               edges |     0.218     348.4     -1.11
              mutual |     0.219     347.1     -1.09
    nodematch_office |     0.259     319.5     -0.66
  nodematch_practice |     0.301     292.2      0.39
nodeofactor_status_2 |     0.236     335.4     -2.25*
nodeifactor_status_2 |     0.231     338.7     -0.76
            gwesp_.5 |     0.207     356.3     -1.09
```

Final model. Columns Mean and SD omitted; sample size 541, interval 4,096, acceptance rate 0.4998.

Effective sample sizes of 292 to 356 meet the working rule of 200 or more. One Geweke z of seven beyond 1.96 is what chance gives; the trace plots decide.



estat mcmcdiag, plot

A slow chain gives itself away in Autocorr and ESS

```
. nwergm lazega_adv, edges mutual nodematch(office practice) ///
>     nodeofactor(status) nodeifactor(status) seed(12345) ///
>     mcmcinterval(50) mcmcburnin(3000) mcmleeffectivesize(0)
```

0.98

autocorrelation of every statistic

22 to 28

effective draws among 3,000

Interval	Smallest ESS
50	30
250	109
1,024	345
4,096	1,427

Burn-in 16,384, 3,000 draws (Chapter 21).

The fit reports convergence and every Geweke and Heidelberger-Welch test passes. The coefficients happen to lie close to the default fit (edges −3.948 against −3.961), but that is luck: only the diagnostics warn.

Goodness of fit: the model's own statistics fit, as they must

```
. estimates store closure
. estat gof, gof(model distance odegree ///
>     idegree espartners) seed(12345) plot
```

Statistic	Observed	Mean	MC p
edges	892	893.20	0.96
mutual	192	192.55	0.98
nodematch_office	674	679.34	0.80
nodematch_practice	610	611.17	0.98
nodeofactor_status_2	616	615.94	0.92
nodeifactor_status_2	562	561.23	0.94
gwesp_.5	1284.86	1288.29	0.94

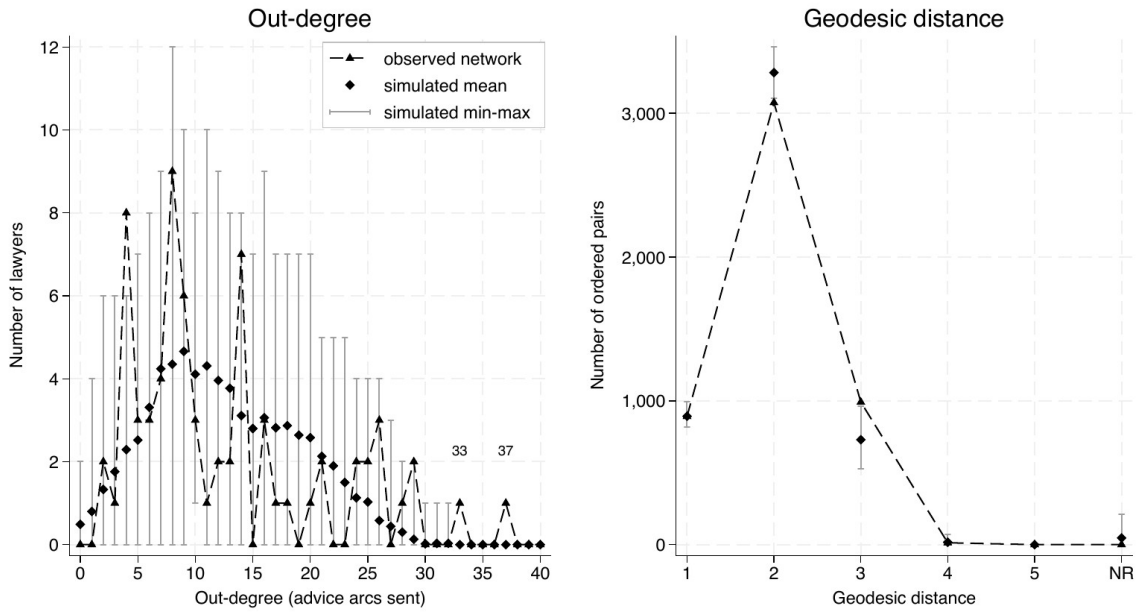
estat gof simulates 100 networks from the fitted coefficients and compares them with the observed network, first on the statistics in the model. The Monte Carlo p-value is twice the smaller tail share.

All p-values are 0.80 or higher: guaranteed, since the estimation matched these statistics.

Statistics the model does not contain do not fit

Distance	Observed	Mean	Min–max	MC p
1	892	893.2	819–993	0.96
2	**3,074**	**3,283**	**3,104–3,461**	**0.00**
3	**991**	**730**	**527–963**	**0.00**
Inf	0	46.8	0–209	1.00

Fewer pairs at distance 2 and more at distance 3 than in any draw. In the simulated networks 47 pairs on average cannot reach each other, while in the observed network every lawyer reaches every other. Two lawyers who name 33 and 37 colleagues lie beyond every simulated network.



Out-degree and distance panels of estat gof (Chapter 21)

Shared partners: the clearest failure

34 | 8.7

arcs with 10 shared partners: observed | simulated mean

17 | 1.9

arcs with 12 shared partners

95 | 155

arcs with 2 shared partners: fewer than simulated

The firm is more clustered than the model

Denser local clusters than the model reproduces: a closure term with another decay, or degree terms for the very active lawyers, may close part of the gap (Exercise 6).

Borderline p-values move

Rows with p between 0.02 and 0.06 can change with another seed. The large gaps do not.

Useful and incomplete

It says
The coefficients describe the over-represented patterns: reciprocity, closure, homophily and status, each net of the others.

It does not say
That the model is correct. Agreement on the statistics it contains is guaranteed. It still misses the spread of the out-degrees and the number of pairs at distance 2, so two heavy advice-seekers and a core of tightly clustered colleagues are not explained.

In the write-up

Say both. Add what the gof output points to (a sender term for the heavy advice-seekers, a closure term with another decay), refit, run estat gof again. Where the gaps cannot be closed, report them.

Fit on statistics a model does not contain is the main evidence that it is more than a re-description of the data.

What can go wrong

Reporting a coefficient because MCMLE converged

The message reads like a verdict on the model, but it only records that the algorithm stopped. A fit can report convergence while the effective sample sizes are a few dozen among thousands of draws. Read estat mcmcdiag after every MCMLE fit.

Treating agreement on the fitted statistics as proof

Agreement is guaranteed, since the estimation matched those statistics by construction. The final advice model reproduces every statistic it contains (p of 0.80 or higher) and still misses the spread of the out-degrees and the number of pairs at distance 2.

Both were shown in this block: the 22 to 28 effective draws behind a fit that reported convergence, and the distance 2 and 3 rows of a model that fits every statistic it contains.

Hands-on (25 min)

Ch. 21, Ex. 3

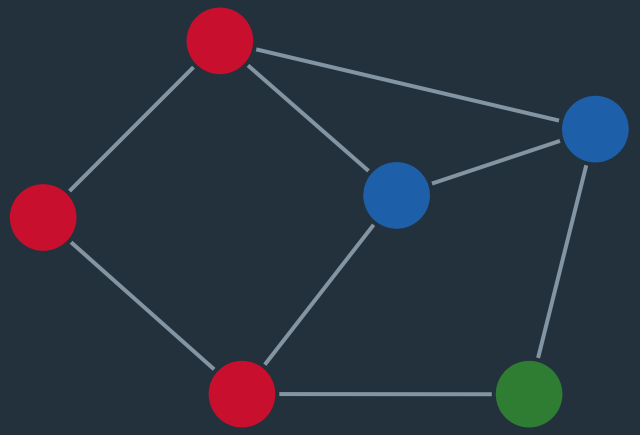
Extend the final advice model with homophily on gender (gender, a string variable like office): nwergm lazega_adv, edges mutual nodematch(office practice gender) nodeofactor(status) nodeifactor(status) gwesp(.5) seed(12345). Report the coefficient of nodematch_gender, its odds ratio and its p-value, and say whether the other coefficients change when it is added. Then run estat mcmcdiag and report the smallest effective sample size and any Geweke test that is flagged.

Hint Store the final model with estimates store closure as in Step 4 and the new fit under another name, for example gender; estimates table closure gender puts the two side by side. A flagged Geweke test among several statistics is read as in the Step 4 diagnostics: look at the trace plots (estat mcmcdiag, plot).

4

What the model says, and the afternoon

15 min · Chapter 21, with a look at Chapters 22 to 29



What the advice model shows

2.81

odds ratio of mutual: advice is returned

2.76

same office (4.1 before reciprocity and closure)

2.28

partner sends; partner is asked 1.49

Advice follows the formal structure (office, practice, status), is reciprocated and clustered, each net of the others. The size of a homophily coefficient depends on whether the local patterns are in the model: office falls from 4.1 to 2.8 and practice from 2.8 to 2.1 in odds ratios.

The model is incomplete

Too many pairs at distance 2 are produced, and two lawyers with 33 and 37 named colleagues lie beyond anything it simulates.

The morning in four commands

Question	Command	Look at
Which local tendencies go with the ties, net of each other?	`nwergm`	edges, nodematch(), gwesp(), mutual
Did the simulation behind the estimates mix?	`estat mcmcdiag`	effective sample size, autocorrelation, Geweke
Does the model reproduce the observed network?	`estat gof`	degree, shared partners, distances
Which of two models is better?	`estimates stats, lrtest`	AIC, BIC, likelihood ratio

Take away. An ERGM makes the dependence among ties the object of the model. Build it in stages, check the simulation, then compare simulated and observed networks on statistics the model does not contain. It describes which local configurations are over-represented, not why they arise.

This afternoon: choose one of two options

Option A: panel data

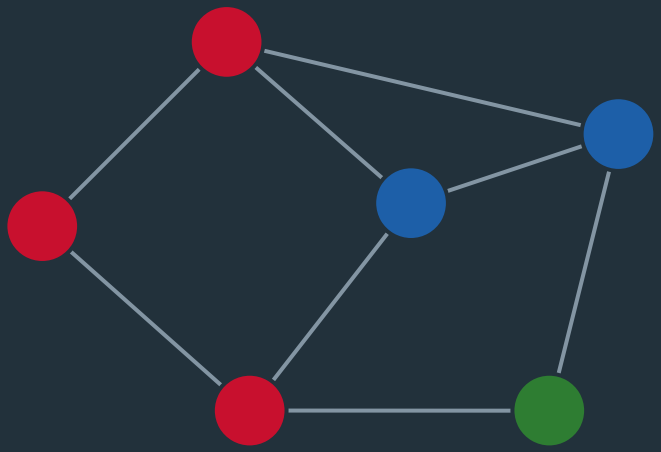
Chapters 14 and 25. How much does a network change between waves (nwturnover), and what do actors' choices look like in a model of change (nwsaom describe, nwsaom on the three Glasgow waves)? Exercise: Ch. 25, Ex. 3, homophily in sport.

Option B: the capstone

Chapter 29. One question about the law firm (do lawyers seek advice from those above them?) carried from theory to conclusion with the QAP and ERGM models of this week. Exercise: Ch. 29, Ex. 4, a step or a slope.

Choose A if your own data are panel waves of the same actors; choose B to bring the week together on one worked analysis.

To read after the session



Chapter 21 · Exponential random graph models

Next: Chapter 22 (simulation, degeneracy) and Chapters 23 to 29

Afternoon: panel data (Ch. 14, 25) or the capstone (Ch. 29)

Do-files and teaching material: nwcommands.org/book

Instructors are free to use, adapt and share these slides.