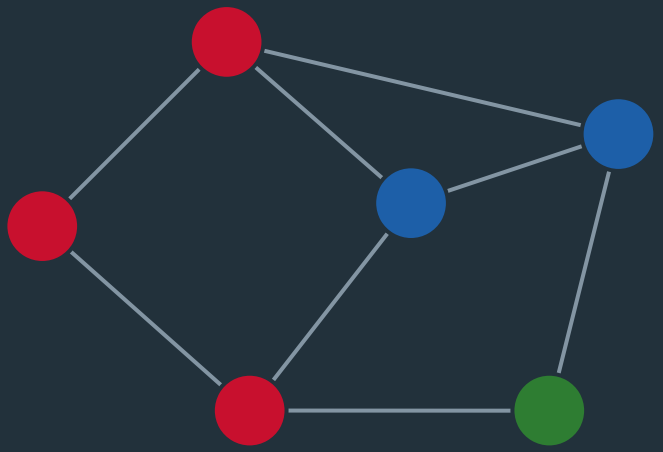
Inference for one network

Course C · Day 4

Random baselines, permutation tests and QAP regression, with the data and commands of Network Analysis with Stata (Stata Press); about six hours

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



Today: five blocks, built on one question

9:00	**1 Chance levels: density and random graphs**	75 min
10:15	Break	15 min
10:30	**2 The conditional uniform graph test**	85 min
11:55	Break	60 min
12:55	**3 Permutation inference and comparing networks**	70 min
14:05	Break	15 min
14:20	**4 QAP regression**	90 min
15:50	**5 Choosing a test, and reporting it**	30 min

Before we start: the data and the two numbers

```
. nwwebuse lazega, nwclear
. nwsummarize lazega_adv, detail
   Nodes: 71
   Arcs: 892
   Density:  .179
   Mean outdegree:  12.56
   Reciprocity:  .274
   Transitivity:  .329
```

Output trimmed to the lines used today. Chapter 15, listing in the section on random graphs.

Check the installation: nwsummarize should print 71 nodes, 892 arcs, density .179. If it does, every command of the day will run.

71 lawyers, three relations

lazega_adv (advice), lazega_cow (co-work) and lazega_fr (friendship), with office, practice, status and gender as variables.

Two numbers to explain

Reciprocity .274 and transitivity .329. Is that a lot, for a network of this size and density?

Seeds

Every simulation today is preceded by set seed or has a seed() option, so your numbers match the slides.

Four tools, four answers to “more than chance?”

Tool	What is randomized	Question it answers	Command
Random baseline (Ch. 15)	A new random network of the same size and density	What would chance give?	nwrandom
CUG test (Ch. 18)	Random networks that share density, dyad census or degrees	Is a statistic of one network unusual?	nwcug
Permutation (Ch. 19)	The labels of the nodes of the observed network	Do two networks, or a network and an attribute, line up?	nwcorrelate, nwmixing
QAP regression (Ch. 20)	The labels of the outcome network, then the model is refitted	Which predictors go with ties, taken together?	nwqap

Each rests on its own assumption about what may be exchanged. The day compares them, then asks which to use when.

Repetitions: the book's numbers, and a lab's numbers

Test	Book	Run time here	Lab with many students	Smallest p
nwcug, density	reps(1000)	14 s	reps(200)	2/(R+1): .002; .010
nwcug, degrees	reps(1000)	6 s	reps(200)	2/(R+1): .002; .010
nwcorrelate	permutations(1000)	under 1 s	keep 1000	2/(R+1): .002
nwqap	permutations(500)	9 s	permutations(100)	1/(R+1): .002; .0099

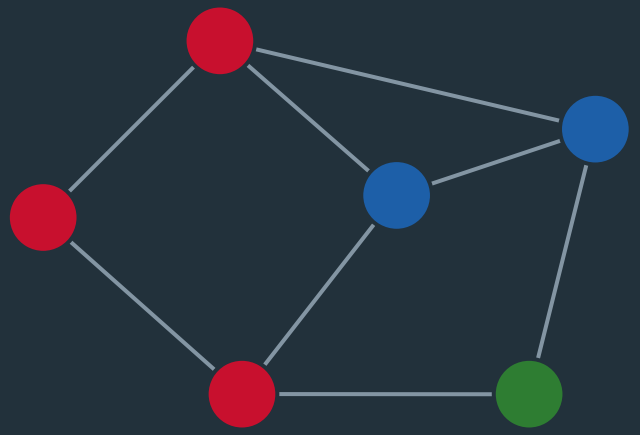
A smaller number gives the same lesson in less time, with a coarser p-value. Say the number of repetitions in every result; the smallest p-value depends on it.

Run times: one Mac laptop, Stata/BE 18, one run at a time, 17 November 2026. Appendix A of the book advises setting reps() or permutations() deliberately.

1

Chance levels: density and random graphs

75 min · Chapter 15 (density and random graphs) · includes ten minutes of setup



Is a transitivity of .33 a lot?

.329

transitivity of the advice network

.274

reciprocity of the advice network

The question for the first hour: how much of the clustering and of the returned requests would appear from chance alone?

Not large or small by itself

It depends on what a network of 71 lawyers and 892 ties would show if no one closed a triangle or returned a request on purpose.

A generator is a theory, run

State a mechanism, run it, and see what that mechanism alone produces.

The features of a network are tied to one another

More ties, more reachability

With nothing else changed, more pairs can reach each other.

More ties, shorter paths

Paths shorten by sheer opportunity.

More ties, more closed triangles

More two-paths happen to close.

More ties, more returned ties

More requests happen to be answered.

None of this needs a tendency toward closure or reciprocity; it is arithmetic. Size and density set the chance level.

One command draws the baseline

```
. nwwebuse lazega, nwclear
. set seed 12345
. nwrandom 71, density(.179) name(rand71)
. nwsummarize rand71, detail
   Arcs: 889
   Density:  .179
   Reciprocity:  .1
   Transitivity:  .174
```

Output trimmed. Chapter 15, section on random graphs.

density(): fixed number of ties

The number of ties is fixed in advance and placed at random among the possible pairs.

prob(): the Bernoulli graph

Every tie is present independently with probability p; the number of ties varies from draw to draw.

The firm against one random graph of its size and density

	Advice network	Random draw
Arcs	892	889
Reciprocity	.274	.100
Transitivity	.329	.174
In-degree centralization	.195	.137
Out-degree centralization	.354	.166
Average shortest path	2.025	1.898

Clustered and reciprocal

Transitivity almost twice, and reciprocity nearly three times, the chance level of the draw.

Short paths prove little

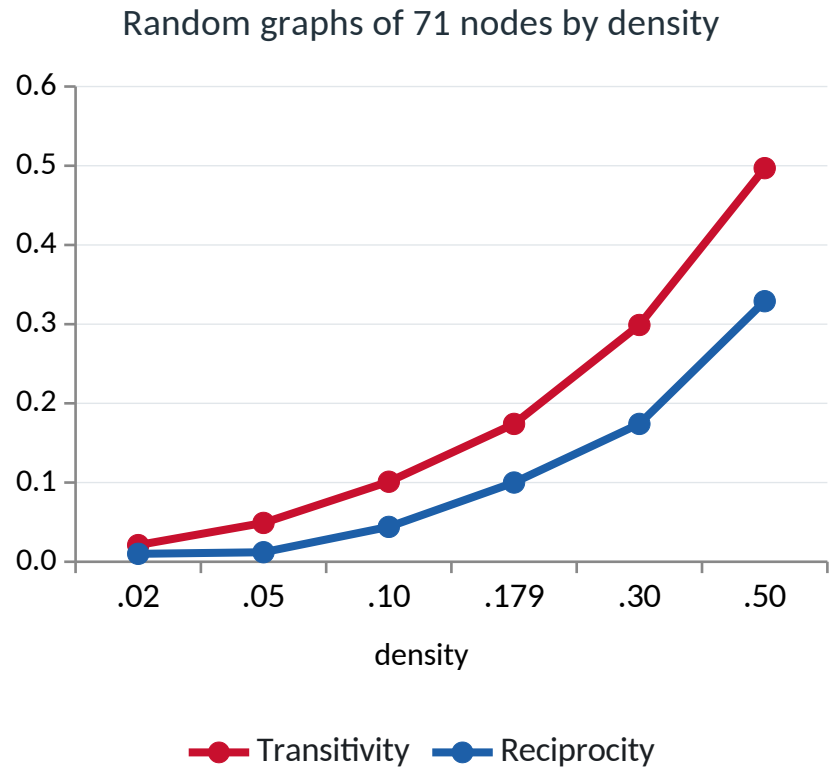
At this density a random graph has paths of 1.9 steps; the firm has 2.0.

One draw is one network from a distribution. Block 2 turns the comparison into a test.

Six densities, one rule

```
. foreach d in .02 .05 .1 .179 .3 .5 {
>     set seed 12345
>     quietly nwrandom 71, density(`d') name(r)
>     quietly nwsummarize r, detail
>     display %6.3f `d' %10.3f r(reciprocity) %10.3f r(transitivity)
>     nwdrop r
> }
 0.020     0.010     0.021
 0.050     0.012     0.049
 0.100     0.044     0.101
 0.179     0.100     0.174
 0.300     0.174     0.299
 0.500     0.329     0.497
```

Columns: density, reciprocity, transitivity. The book stores the rows in a matrix (Chapter 15).



Density sets the chance level of everything

density	reciprocity	transitivity	reachable	avgpath
.020	.010	.021	.275	.
.050	.012	.049	.891	.
.100	.044	.101	1.000	2.377
.179	.100	.174	1.000	1.898
.300	.174	.299	1.000	1.701
.500	.329	.497	1.000	1.500

The row at .179 is the firm's density: a network with no tendency of any kind would show a transitivity of about .18, a reciprocity of about .10 and paths of about 1.9 steps.

Transitivity is the density

Two contacts of one person are tied with the probability of any tie.

Reciprocity is p/(2−p)

The share of mutual pairs among connected pairs when each direction is drawn independently.

Paths are free from .1 on

Below .05 many pairs are unconnected and avgpath is undefined (dot).

Size matters too, separately from density

nodes	density	transitivity	avgpath
71	.179	.175	1.897
142	.089	.087	2.207
284	.044	.044	2.508
710	.018	.018	2.854

```
. local d = 12.56 / (`n' - 1)
. quietly nwrandom `n', density(`d') name(r)
```

Body of the loop over `n' = 71, 142, 284, 710 (Chapter 15): 12.56 contacts per node throughout.

2×

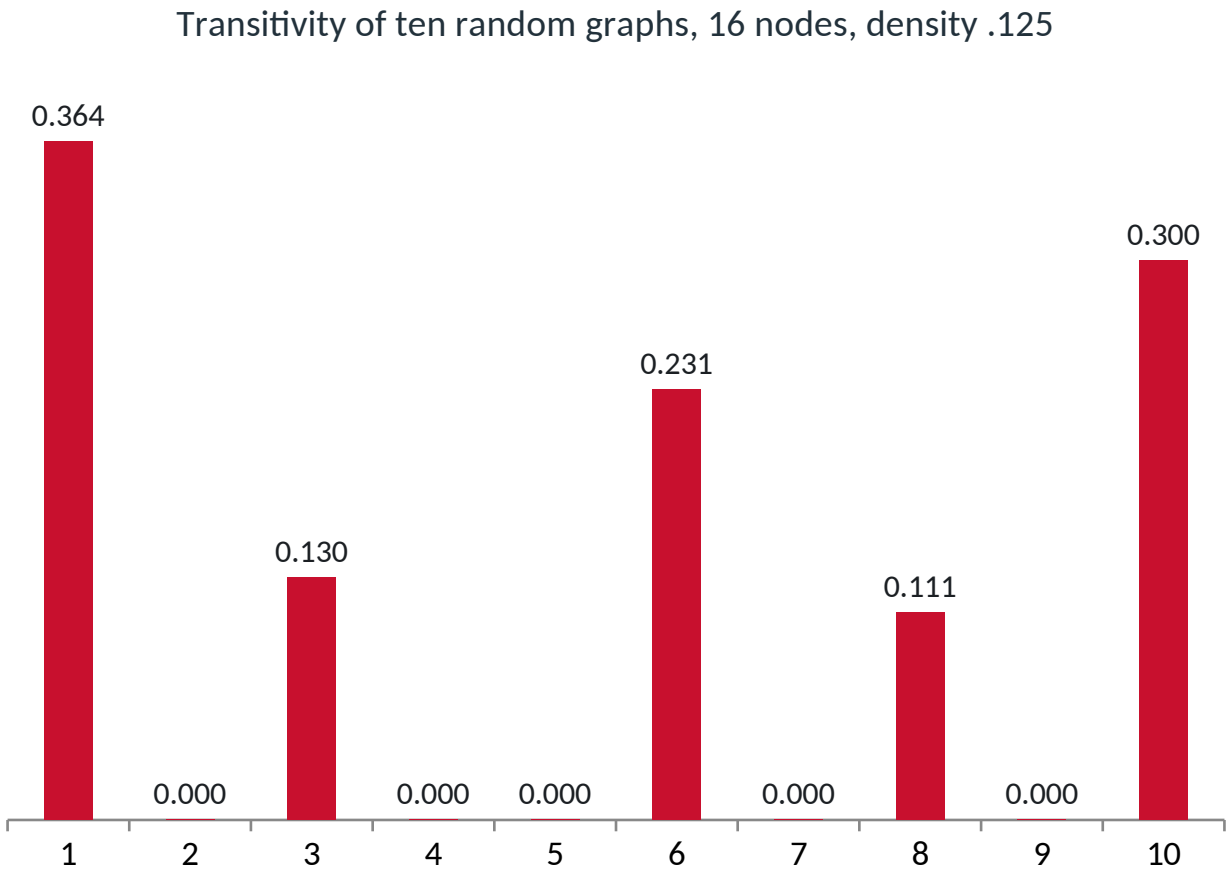
the chance level that .329 is in a firm of 71

18×

the chance level that .329 would be at 710 nodes, with the same contacts per person

Paths grow only slowly with size, roughly with its logarithm. A raw statistic cannot be compared across networks of different size or density.

One draw is one draw: small networks vary a lot




```
. set seed 1
. quietly nwrandom 16, density(.125) ///
>     undirected ntimes(10) name(small)
```

71 nodes: .175 to .181

Five draws of the firm's random graph hardly move.

16 nodes: 0 to .364

Five of ten draws have no closed triad at all.

“Random” is a choice of what to hold fixed

Size and density

Fixes nodes and the number of ties. Chance level of transitivity: .179 for the firm. nwrandom, and nwcug with condition(density).

Plus the dyad census

Also fixes the counts of mutual, asymmetric and null dyads. Directed networks only. condition(census).

Plus every degree

Also fixes each lawyer’s number of requests sent and received. A few busy lawyers close more triangles by chance. condition(degrees).

Each choice answers a narrower question. A comparison with the firm is only as informative as what the baseline holds fixed.

Hands-on (20 min)

Ch. 15, Ex. 2

Before generating it, predict the transitivity and the reciprocity of a directed random graph with 71 nodes and density 0.25, using the density sweep. Then check your prediction with one draw: nwrandom 71, density(.25), followed by nwsummarize, detail.

Hint In the sweep the transitivity of each random graph is close to its density. For reciprocity, use p/(2−p); check the formula against the printed rows for 0.179 and 0.3 before using it for 0.25.

Ch. 15, Ex. 7

Five of the ten 16-node random graphs had a transitivity of zero. Repeat the ten draws with 30 and with 60 nodes at the same density of 0.125 (set seed 1). For each size, report how many draws have a transitivity of zero and the standard deviation across the ten draws. What does this imply for reading a single draw of a small network?

Hint Change only the number of nodes in the nwrandom line and give each run its own name() and save() file. In the saved file, count if transitivity == 0 counts the zeros and summarize transitivity gives the standard deviation.

What to take away from the first block

Chance is not zero

At the firm's density a random network has transitivity .18, reciprocity .10 and paths of 1.9 steps.

Size and density matter

The same contacts per person give a tenth of the chance transitivity in a network ten times larger.

One draw says little

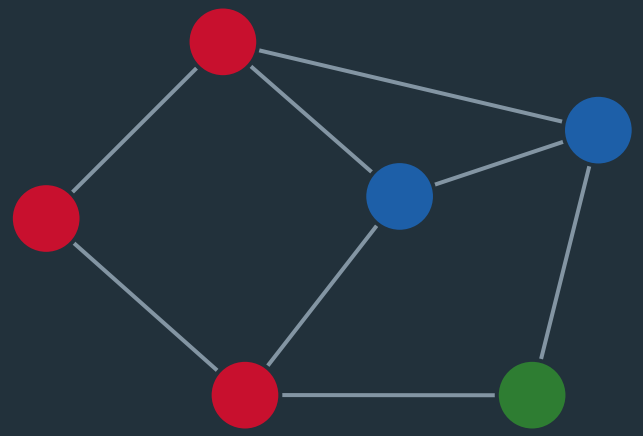
Small networks vary a lot between draws. A comparison rests on many.

Next: make the comparison a test, with 1,000 random networks and a stated list of what they share with the firm.

2

The conditional uniform graph test

85 min · Chapter 18



Dyads are not independent observations

4,970

ordered pairs among the 71 lawyers

71

lawyers who generate them

- A tie from A to B and one from A to C both involve A, with her own need for advice, office and standing.
- Pairs that share a lawyer tend to be more alike than independent draws would be.
- Ordinary tests assume independent observations, so they cannot be applied as they stand.

Also, a statistic needs a baseline

More ties close more triangles and return more requests by sheer opportunity. A value is high or low only relative to a baseline that holds size and density fixed.

The question is how ties depend

Whether they do is rarely in doubt. Part IV responds with tests (this day) and models (Day 5).

A test built from random networks

1 Compute the statistic

On the observed network: the transitivity of lazega_adv is .329.

2 Draw random networks

1,000 networks drawn uniformly from all those that share what you hold fixed; compute the statistic on each.

3 Locate the observed value

Near the middle: chance produces it easily. Far in a tail: it does not.

Conditional uniform graph (CUG) test. Uniform: every network that meets the conditions is equally likely to be drawn. Conditional: the conditions, such as size and density, are held fixed (Anderson, Wasserman and Crouch 1999).

First test: is the advice clustering more than chance?

```
. nwcug lazega_adv, stat(nwclustering) rname(cluster_global) reps(1000) ///
>         seed(12345) condition(density) plot

Conditional uniform graph test                  Number of draws =      1,000
Network: lazega_adv                             Condition       =    density
---------------+--------------------------------------------------------------
               |  Observed Null mean   Null sd  P >= obs  P <= obs Two-sided
cluster_global |  .3293109  .1791272  .0040803    0.0010    1.0000    0.0020
```

Output trimmed to the header and the table row. Chapter 18, section on the conditional uniform graph test.

.329

observed transitivity

.179

null mean, equal to the density (sd .004)

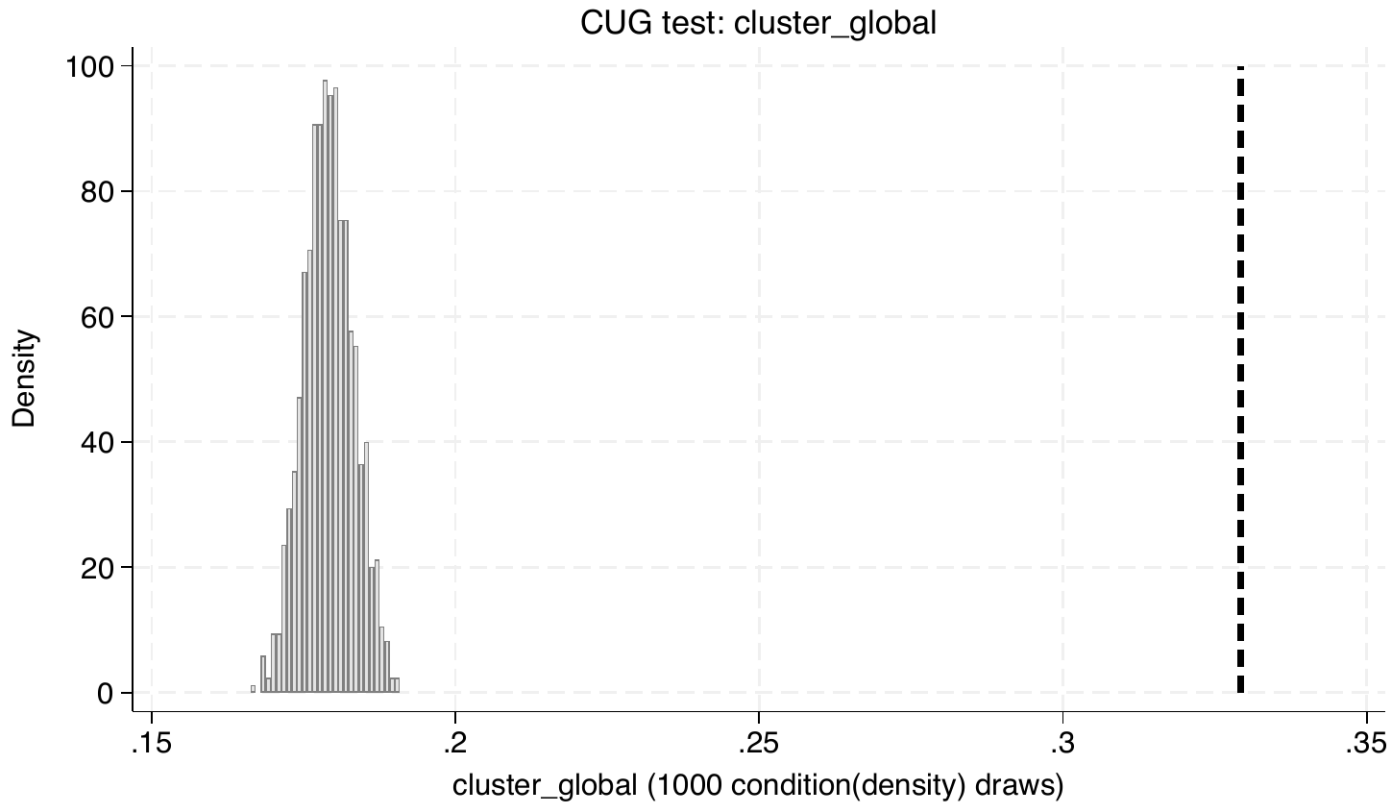
None of the 1,000 random networks reaches .329. Almost 37 null standard deviations above the mean: (.3293 − .1791)/.0041 = 36.8.

Reading the nwcug table

Column	Advice transitivity	What it is
Observed	.329	The statistic on the real network
Null mean	.179	Average over the 1,000 random networks; equal to the density
Null sd	.004	Spread of the draws; the observed value is almost 37 sd above the mean
P >= obs	0.0010	Share of draws at least as large; the observed network is one more draw
P <= obs	1.0000	Share of draws at most as large
Two-sided	0.0020	Twice the smaller tail, capped at 1; the definition used throughout the book

The claim “larger than chance” concerns one tail: the P >= obs column. Choose the column before you run the test.

The histogram shows what the p-value compresses



The 1,000 density-conditioned null draws and the observed .329 (dashed). Chapter 18.

A tight null

The draws cluster between about .17 and .19.

A wide gap

The dashed line sits far off the histogram: the picture of an overwhelming result. A marginal result would fall at the edge of the bulk.

p = .001 is a floor, not a zero

1/1001

smallest upper-tail p with 1,000 draws

The observed network is a draw

The plus-one correction (Phipson and Smyth 2010): a Monte Carlo p-value is never reported as exactly zero.

The two-sided value doubles

The smaller tail times two: 2/1001 = .002, the floor for 1,000 draws.

Write "p ≤ 2/1001 with 1,000 draws", not "p = 0". A value at the floor means no draw reached the observed value; more draws would give a finer value.

Any number left in r() can be tested

```
. nwcug lazega_adv, stat(nwdyads) rname(reciprocity) reps(1000) ///
>         seed(12345) condition(density)

-------------+------------------------------------------------------------
             |  Observed Null mean   Null sd  P >= obs  P <= obs Two-sided
  reciprocity |  .2742857  .0990426  .0101065    0.0010    1.0000    0.0020
```

.099

null mean of the reciprocity, again the value that density alone gives

Placeholder and your own programs

If the network is not the first argument, ##net## marks where its name goes: stat(nwcorrelate lazega_fr ##net##) rname(corr). A program of your own that leaves a number in r() can be tested the same way.

Observed .274: more than ten null standard deviations above the mean.

What a significant result does not say

It says

Advice ties close more triangles, and are returned more often, than a random network of the same size and density produces. Density alone does not account for the pattern.

It does not say

That lawyers close triangles on purpose. Office sorting, shared settings and a few very busy lawyers can all raise transitivity, and a density null leaves every one of them in play.

Imagine a firm sorted by office

If people ask for advice almost only inside their own office, transitivity is high with nobody closing a triangle on purpose: two colleagues of one office share most of their contacts. The next slides hold more fixed to separate some of these rivals.

The conditioning choice is part of the question

Option	Held fixed	The question it answers
`condition(density)`	Number of nodes and of ties	Is this unusual for a network this sparse?
`condition(census)`	Also the counts of mutual, asymmetric and null dyads (directed networks only)	What comes on top of the observed reciprocity?
`condition(degrees)`	Also every node's out- and in-degree; ties are swapped, trades() sets how many	Is this unusual given the busy lawyers?

The calls differ by one word, and the output differs only in the condition named in its header. State the choice in every result.

Census conditioning: transitivity yes, reciprocity no

```
. nwcug lazega_adv, stat(nwclustering) rname(cluster_global) reps(1000) ///
>         seed(12345) condition(census)
               |  Observed Null mean   Null sd  P >= obs  P <= obs Two-sided
cluster_global |  .3293109  .1793014  .0047774    0.0010    1.0000    0.0020

. nwcug lazega_adv, stat(nwdyads) rname(reciprocity) reps(1000) ///
>         seed(12345) condition(census)
Note: the null standard deviation is 0 under condition(census); this statistic
is fixed by the dyad census, so the test is uninformative.
             |  Observed Null mean   Null sd  P >= obs  P <= obs Two-sided
 reciprocity |  .2742857  .2742857         0    1.0000    1.0000    1.0000
```

Rule of thumb: density conditioning as the default; census conditioning when the question is what comes on top of the observed reciprocity.

Transitivity

Null mean .179, sd .005: the answer does not change.

Reciprocity

The census counts the mutual pairs, so every draw has exactly .274: sd 0, two-sided p = 1.

Degrees alone raise the chance level

```
. nwcug lazega_adv, stat(nwclustering) rname(cluster_global) reps(1000) ///
>         seed(12345) condition(degrees)
               |  Observed Null mean   Null sd  P >= obs  P <= obs Two-sided
cluster_global |  .3293109  .2719183  .0043436    0.0010    1.0000    0.0020

 reciprocity |  .2742857  .1363848  .0099951    0.0010    1.0000    0.0020
(same call with stat(nwdyads) rname(reciprocity))
```

.272

chance transitivity once every lawyer's degrees are fixed (.179 under density)

Three fifths of the way: (.272 − .179)/(.329 − .179) = .62. Unequal degrees raise the chance level, but no draw reaches .329 or .274.

Entangled

Closure can itself make some lawyers busier: the degrees say how much clustering they would bring, not how much they cause.

Six tests side by side

Network and statistic	Conditioning	Observed	Null mean	Null sd
Advice, transitivity	density	.329	.179	.004
Advice, transitivity	census	.329	.179	.005
Advice, reciprocity	density	.274	.099	.010
Advice, reciprocity	census	.274	.274	0
Advice, transitivity	degrees	.329	.272	.004
Advice, reciprocity	degrees	.274	.136	.010

The same two observed values, three different questions. Choose, then say which.

Only degrees raise the baseline

For transitivity: the census leaves it at .179.

Only census can void a test

Sd 0 for reciprocity: nothing to test.

A result that is not unusual

```
. set seed 2024
. nwrandom 71, density(.179) name(draw)
. nwcug draw, stat(nwclustering) rname(cluster_global) reps(1000) ///
>         seed(12345) condition(density)

               |  Observed Null mean   Null sd  P >= obs  P <= obs Two-sided
cluster_global |  .1737822  .1784492  .0040069    0.8831    0.1179    0.2358
```

.236

two-sided p of a draw from the null itself

A calibration. A random graph tested against its own null sits well inside the distribution: about one sd below the mean.

Large p is not proof

It is consistent with chance. A small network gives a wide null and little power.

Hands-on (25 min)

Ch. 18, Ex. 3

lazega_fr has density 0.116 and transitivity 0.186. Before running anything, say what null mean you expect under condition(density), and why. Then run the test with reps(1000) seed(12345) plot. Report the null mean, the null sd and how many sd 0.186 lies above the mean. Why could 0.186 alone not show closure?

Hint Ask with what probability two neighbors of a node are tied in a random graph of this density. The command is the first nwcug call of the block with lazega_fr in place of lazega_adv.

Ch. 18, Ex. 5

Within one office, sorting by office has nothing to work with. Keep the Boston lawyers: nwsubset lazega_adv if office == "Boston", name(adv_boston). Test adv_boston with condition(density) and condition(degrees), reps(1000) seed(12345). Is the clustering above both baselines? What does that say about office sorting as the whole explanation?

Hint The Boston network is denser than the whole firm (0.239): compare its transitivity with its own null means, not with 0.179.

Ch. 18, Ex. 1

For the reciprocity tests of lazega_adv, compute for density and for degree conditioning how many null sd the observed 0.274 lies above the null mean. Which of the three p-value columns answers the claim that advice is more reciprocal than chance, and what is its value?

Hint Use the unrounded numbers: (observed − null mean)/null sd, for density (0.2743 − 0.0990)/0.0101. A claim that a statistic is larger than chance concerns one tail.

Hands-on (further practice, 10 min)

Practice 1

Is the friendship network more reciprocal than chance? Test lazega_fr with stat(nwdyads) rname(reciprocity), condition(density), reps(1000) seed(12345). Compare the observed reciprocity with the null mean.

Hint nwdyads lazega_fr prints the observed reciprocity. Read the P >= obs column.

Practice 2

What do the degrees do to the chance level of co-work? Test the transitivity of lazega_cow (observed .356) with condition(density) and with condition(degrees), reps(1000) seed(12345). By what share of the distance to the observed value do the degrees raise the null mean?

Hint Same call as the advice test with lazega_cow. Share = (null mean under degrees − null mean under density) / (observed − null mean under density).

Practice 3

Does advice line up with friendship more than a random network of the same density does? Run nwcug lazega_adv, stat(nwcorrelate lazega_fr ##net##) rname(corr) reps(200) seed(12345) condition(density). Compare with the permutation test of Block 3.

Hint ##net## marks where nwcug puts the name of each random network. With reps(200) the two-sided floor is 2/201 = .01.

What to take away from the second block

Build the reference

Ties are not independent, so the comparison comes from random networks that share some features with the observed one.

The conditions matter

Density, census or degrees: each fixes more, and each answers a narrower question. State it.

Significant is not explained

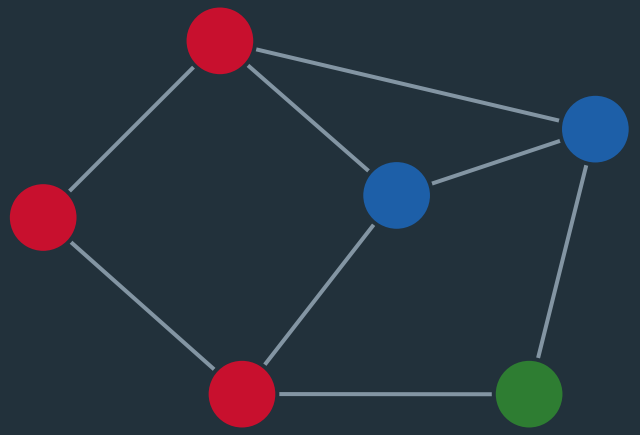
More than chance at one level of structure leaves closure, sorting and busy lawyers in play.

The firm: transitivity .329 against .179 (density) and .272 (degrees); reciprocity .274 against .099 and .136. Report the number of draws and the seed.

3

Permutation inference and comparing networks

70 min · Chapter 19



Two claims about association

Advice goes with friendship

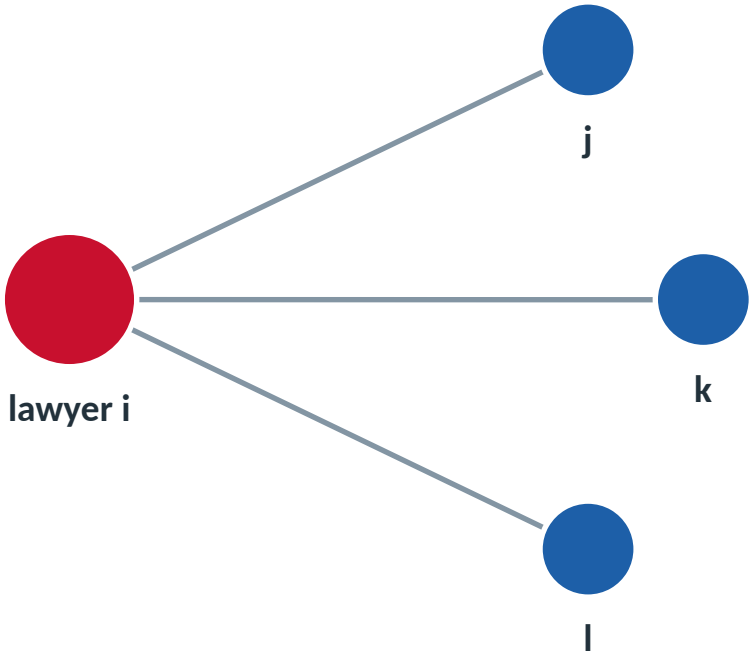
Lazega's cooperation among peers: advice is embedded in friendship. Compares two networks on the same 71 lawyers.

Office predicts who asks whom

The homophily argument. Compares a network with a node attribute. Three quarters of the advice arcs lie inside an office.

Each comparison must say what it preserves and what it randomizes. Preserved: the observed network, with all its internal dependence. Randomized: which lawyer is which.

4,970 rows are not 4,970 pieces of evidence



The dyads (i,j), (i,k) and (i,l) share the lawyer i.

- A sociable or senior lawyer raises the probability in every dyad that involves her at once.
- Fed into a correlation or a regression, 4,970 ordered pairs count as 4,970 separate pieces of evidence, when 71 actors generate them.
- The effective amount of information is smaller, so the uncertainty is understated, sometimes severely.

Relabeling keeps the structure and breaks the link to identity

Property	Random-graph CUG test	Node-label permutation
Number of nodes	preserved	preserved
Number of ties	preserved under density conditioning	preserved
Degree sequence	generally changes	preserved up to relabeling
Reciprocity	may change (fixed under census)	preserved
Triadic structure	may change	preserved
Alignment with another network or attribute	randomized by generating new networks	randomized by relabeling

Use the CUG test for a feature of one network. Use relabeling when the claim is about how two observed structures line up and each should keep its own shape.

One relabeling, by hand

```
. set seed 12345
. nwpermute lazega_fr, generate(fr_perm)
. nwcorrelate lazega_adv lazega_fr
   .41767433
. nwcorrelate lazega_adv fr_perm
   .00295191
```

Correlations as printed in a one-column table. Chapter 19, section on node-label permutation.

.418

advice with friendship

.003

advice with the relabeled friendship

The relabeled network is as clustered and as reciprocal as the original. Only its alignment with the lawyers' other relations is gone.

Everything else is unchanged

fr_perm has the same 575 arcs, density .116, reciprocity .148 and transitivity .186 as lazega_fr; only the link between identity and position is broken.

Correlating two networks: nwcorrelate with permutations()

```
. set seed 12345
. nwcorrelate lazega_adv lazega_fr, permutations(1000)

 Correlation     P-value   [95% permutation interval]
-------------------------------------------------------------
   .41767433      0.0020      -.03147169   .03327747
```

.418

observed correlation

±.03

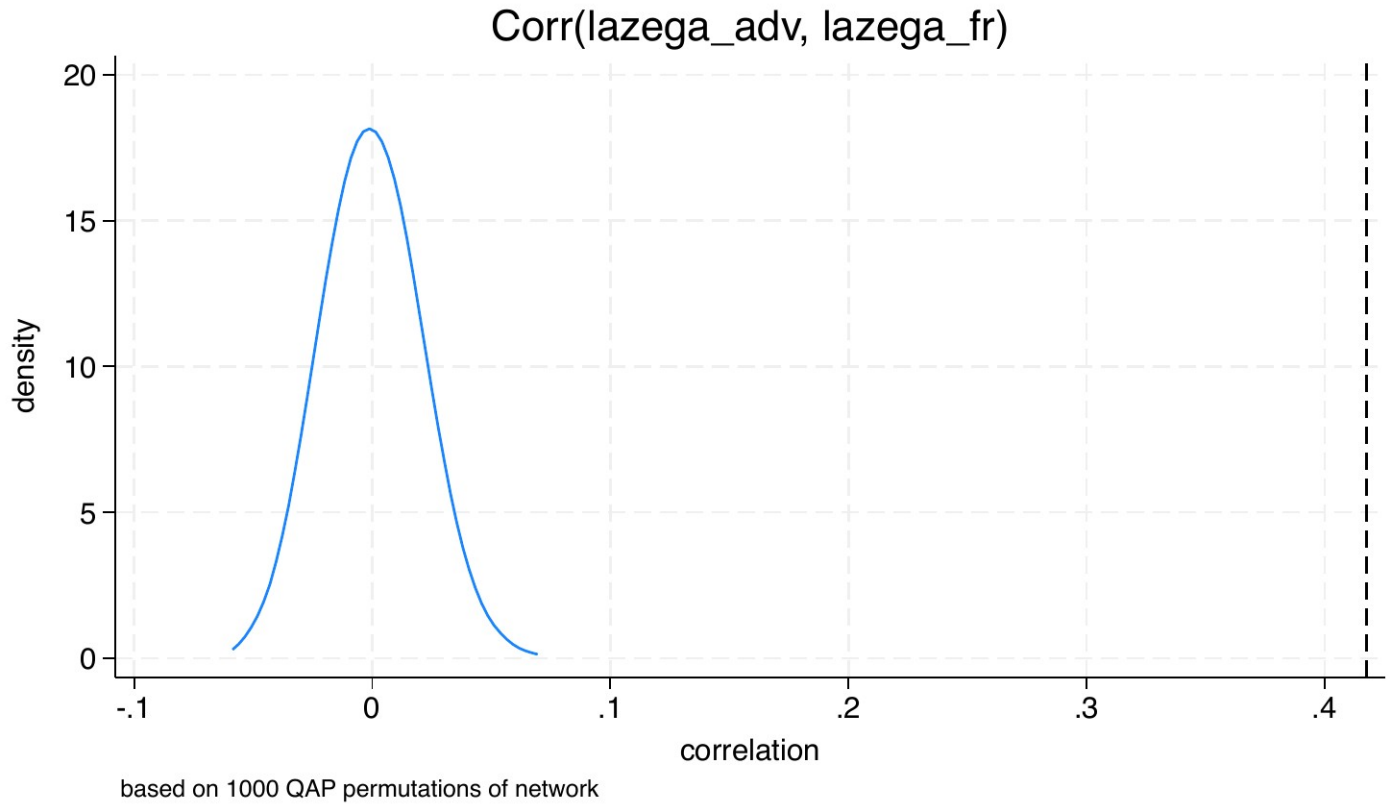
95% of 1,000 relabelings lie between −.031 and .033

p = .002 is the floor for 1,000 relabelings. The correlation is computed over all eligible cells, including the many pairs with neither relation, so .42 does not mean 42% of advice ties are friendship ties.

Size and surprise are different

The correlation gives the size of the association in the observed networks; the p-value says how unusual it is among relabelings.

The whole permutation distribution



1,000 relabelings of the friendship network against advice; dashed line at the observed .418. Chapter 19.

Centered on zero

Relabeling destroys the alignment, so the permuted correlations fall between −.03 and .03.

The line is far outside

None of the 1,000 relabelings reached .418: $p \leq 2/1001$.

What the correlation says, and what it does not

It says

Lawyers linked by friendship were much more often linked by advice too than a random alignment of the two networks would give. None of the 1,000 relabelings matched .418; the observed arrangement counts as one more draw, with the smaller tail doubled (2/1001).

It does not say

That the probability is zero: 2/1001 is the floor for 1,000 draws. Nor that friendship caused advice or the reverse.

A bivariate association

It cannot tell a direct link between friendship and advice from a shared cause. Lawyers of one office may both befriend and consult each other because they share a floor. Partialling out a third network or attribute is the subject of Block 4.

A network against an attribute: nwmixing

```
. nwmixing lazega_adv, attribute(office) permutations(1000)
   E-I Index: -.511   p-value: .002
. nwmixing lazega_adv, attribute(status) permutations(1000)
   E-I Index: -.152   p-value: .002
(output omitted)
```

The p-value of the E-I index is this permutation test. With the default 100 relabelings the floor was .02; with 1,000 it is .002.

Attribute	Arcs inside	Relabeling	E-I
Office	76%	52.5%	−.511
Status	58%	49%	−.152

The reference point is what relabeling gives, not zero: an E-I of about −.05 for office, +.014 for status. Office sorting is strong, status sorting modest.

What may be exchanged is part of the design

Exchangeability

Relabeling treats every lawyer as interchangeable with every other. Office, status and practice make them differ, and relabeling also destroys the alignment with those attributes.

A statement about this firm

The 71 lawyers are the whole firm at one time, not a sample. The p-value says how unusual the alignment is among relabelings of this very network (Freedman and Lane 1983).

Too few permutations

The two-sided p cannot be smaller than 2/(R+1): .002 for 1,000, about .02 for 100. A value at the floor means more are needed if a finer value matters.

The CUG test draws new networks; permutation relabels the same one. They are not interchangeable.

Hands-on (20 min)

Ch. 19, Ex. 4

Test the correlation of the advice network with the co-work network lazega_cow: set seed 12345, then nwcorrelate lazega_adv lazega_cow, permutations(1000). Report the correlation, the 95% permutation interval and the p-value as an inequality. Larger or smaller than advice and friendship (0.418)? Which would you expect to overlap more with advice, and why?

Hint When no relabeling reaches the observed value, write p ≤ 2/1001 and name the number of permutations. For the last question, think about what the lawyers ask one another for advice about.

Ch. 19, Ex. 5

Is advice sorted by gender? From the group sizes (53 men, 18 women; tabulate gender) compute the E-I index that relabeling alone would give. Then set seed 12345 and run nwmixing lazega_adv, attribute(gender) permutations(1000). Report the observed E-I index and its p-value: homophily, heterophily or neither?

Hint The share of ordered pairs within a group is (53·52 + 18·17)/(71·70); the E-I index expected under relabeling is one minus twice that share. Compare with this value, not with zero.

Ch. 19, Ex. 2

With permutations(1000), what is the smallest two-sided p-value that nwcorrelate can report when the observed correlation is more extreme than every relabeling? Compare it with the printed p-value for the advice–friendship correlation of 0.418 and say what that value does and does not tell you.

Hint Count the observed arrangement as one more draw, 1,001 in all, and remember that nwcorrelate doubles the smaller tail, as the CUG test of Block 2 does.

Hands-on (further practice, 10 min)

Practice 1

Is advice sorted by practice? Run set seed 12345, then nwmixing lazega_adv, attribute(practice) permutations(1000). Report the E-I index and its p-value and compare with the E-I index for office (−.511).

Hint tabulate practice gives the group sizes; compute the expected E-I as in Ex. 5.

Practice 2

Which pair of relations overlaps most? Correlate co-work with friendship: set seed 12345, then nwcorrelate lazega_cow lazega_fr, permutations(1000). Put the three correlations (advice–friendship .418, advice–co-work, co-work–friendship) in order.

Hint Report each p-value as an inequality with the number of permutations.

Practice 3

Run the advice–friendship test again with permutations(100) after set seed 12345. What happens to the p-value and to the permutation interval, and why?

Hint The smallest two-sided p-value with R = 100 is 2/101.

What to take away from the third block

Relabel, do not regenerate

Permuting node labels preserves the observed structure and destroys the link between identity and position.

The data in hand

Not other firms; not a cause. A value at the floor says only that no relabeling matched the observed value.

Bivariate, not yet a model

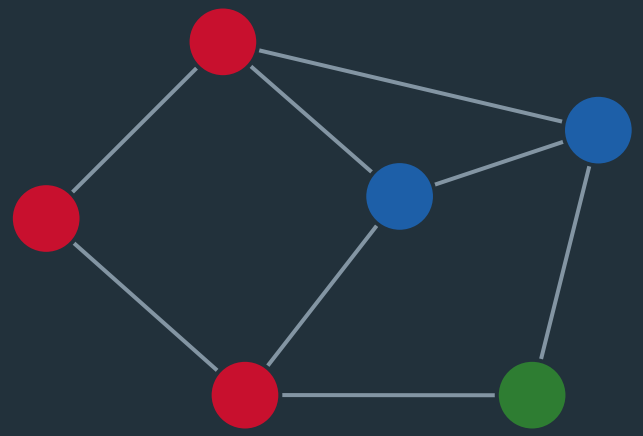
Advice goes with friendship (.418), and with office. Do they go together net of each other?

The firm: friendship and advice correlate at .418 while 95% of 1,000 relabelings lie between −.031 and .033; 76% of advice arcs are inside an office against 52.5% under relabeling.

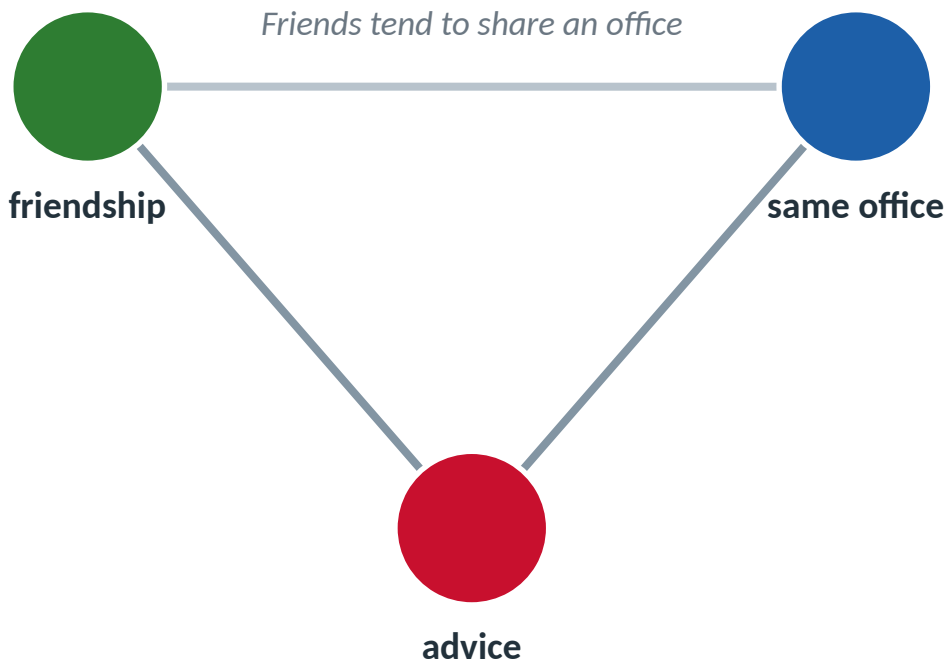
4

QAP regression

90 min · Chapter 20



Friendship, or the office in disguise?



- Lawyers who are friends also tend to share an office, and lawyers who share an office tend to ask each other for advice.
- A bivariate correlation cannot separate the two: friendship may be the office in disguise.
- So: several predictors in one model, each coefficient read holding the others fixed.

A dyadic regression, and why its standard errors fail

$$\text{logit } \Pr(Y_{ij} = 1) = b0 + b1\, X_{ij} + b2\, Z_{ij}$$

Rows share lawyers

The 70 rows in which a given lawyer is the sender move together: a generally sociable lawyer raises the probability in all of them at once. The logit counts 70 independent pieces of evidence.

Predictors correlate

Friends tend to work in the same office, so a coefficient can borrow strength from, or be confused with, a neighbor (Dekker, Krackhardt and Snijders 2007).

Each dyad is one observation; Y is 1 if there is a tie, X and Z are dyadic predictors. The point estimates of an ordinary logit are fine as summaries; the standard errors are not.

QAP keeps the estimates and replaces the inference

1 Fit the logit

Advice on friendship and three same-value indicators. The point estimates stay.

2 Relabel the outcome

Permute the rows and columns of the dependent network together, keep the predictors, refit the whole model on each permutation.

3 Compare

Each coefficient with its own permutation distribution. The p-value counts both tails around its mean; the observed value is one more draw.

The smallest p-value is 1/(R+1) for R permutations. The dependence between dyads that share a lawyer stays in the reference distribution, because whole networks are relabeled (Krackhardt 1987, 1988).

Predicting advice

```
. set seed 12345
. nwqap lazega_adv lazega_fr office practice status, ///
>         permutations(500) plot

  Regression            =  logit
  Permutations          =  500
  Number of vertices    =  71
  Number of arcs        =  892

            lazega_adv |      Coef.    P-value
-----------------------+-------------------------
             lazega_fr |      2.198     .002
           same_office |      1.062     .002
         same_practice |       .752     .002
           same_status |       .407     .002
                 _cons |     -3.228
```

Attributes become dyads

office, practice and status turn into same_office, same_practice and same_status by the default mode(same).

4,970 rows

The 71-node directed network is reshaped into dyad-level observations.

Odds ratios from the printed coefficients

Predictor	Coefficient	Odds ratio	p
Friendship tie (lazega_fr)	2.198	9.0	.002
Same office	1.062	2.9	.002
Same practice	.752	2.1	.002
Same status	.407	1.5	.002

exp(2.198) ≈ 9.0, exp(1.062) ≈ 2.9, exp(.752) ≈ 2.1, exp(.407) ≈ 1.5. Arithmetic from the printed coefficients: in Stata, display exp(2.198).

9.0

times the odds of an advice tie: friendship, holding office, practice and status fixed

Crude odds ratio: 11.9

From the cross-tabulation (358 × 3,861)/(217 × 534). Odds ratios from tables and from models with other predictors are not on a common scale.

What the coefficients say, and what they do not

It says

A friendship tie from i to j multiplies the odds of an advice tie from i to j by about 9.0, holding office, practice and status fixed (p = .002). The three shared attributes multiply the odds by about 2.9, 2.1 and 1.5.

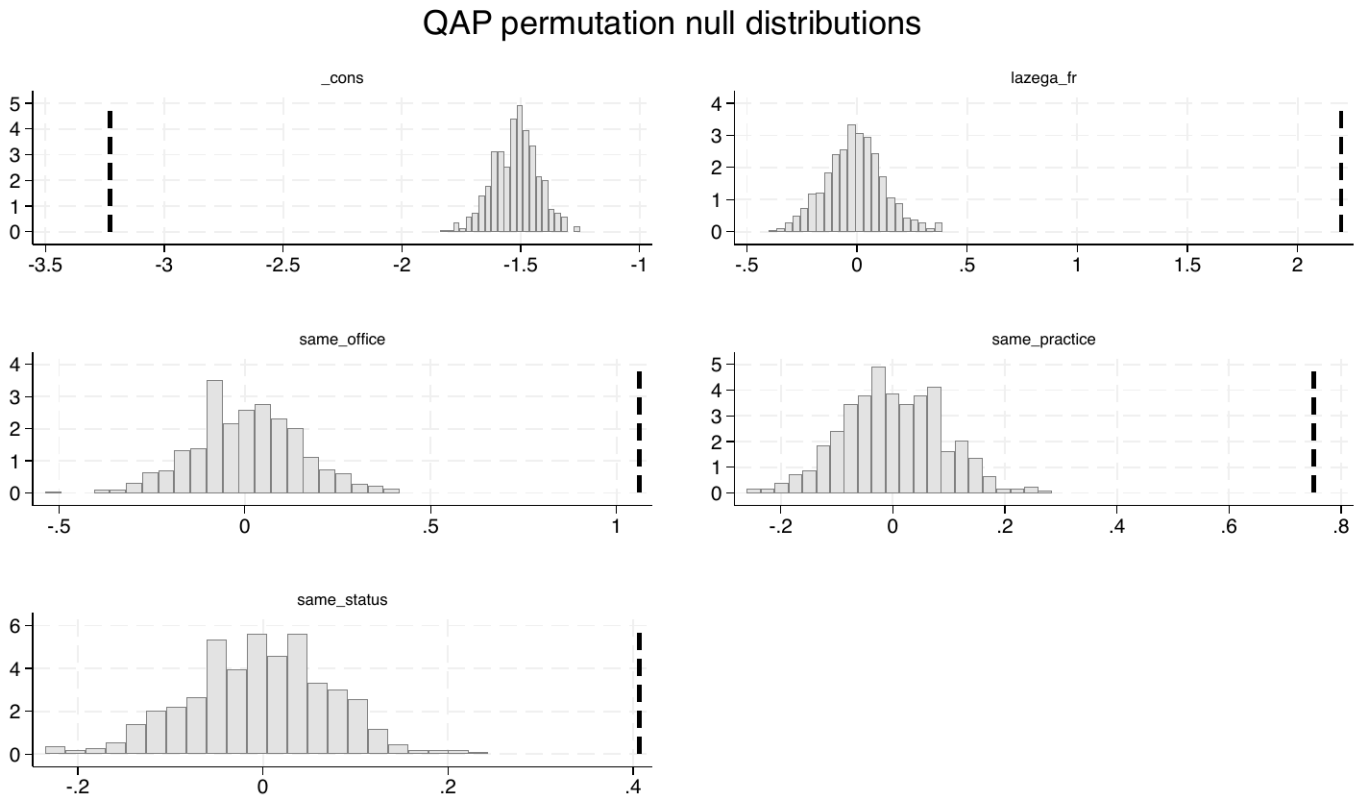
It does not say

How the effects add up for a specific pair: these are odds ratios, not probability ratios. Nor that friendship causes advice.

A bivariate test would borrow

Separate bivariate tests would answer “is this predictor related to advice?” without holding the others fixed; because friends share offices, each association would be partly borrowed from another predictor.

The permutation distributions behind the p-values



One panel per coefficient, 500 permutations, dashed lines at the observed values. Chapter 20.

Outside every histogram

The dashed line sits outside each panel: what p = .002 looks like.

Status is the closest call

The permuted coefficients reach about .25; the observed .41 lies only a little beyond.

Correlated predictors: double semi-partialling

```
. set seed 12345
. nwqap lazega_adv lazega_fr office practice status, ///
>         permutations(500) qapspp

  Estimation            =  QAP, double semi-partialling

             lazega_fr |      2.198     .002
           same_office |      1.062     .002
         same_practice |       .752     .002
           same_status |       .407     .002
```

Coefficient rows only; the header is as before. Chapter 20.

Same coefficients, same conclusion. Only the test differs: for these data the result does not depend on the permutation scheme.

Plain QAP

Permutes the whole outcome network: some association, not strictly net of the others.

qapspp

Permutes each predictor's residuals from the others (Dekker et al. 2007): much less sensitive to correlated predictors.

From a nodal attribute to a dyadic predictor: mode()

mode()	Dyadic variable	Use for
`same (default)`	1 if v_i = v_j	A categorical attribute: pooled homophily (office, practice)
`dist`	v_i − v_j	A signed difference
`absdist`	|v_i − v_j|	A continuous or ranked attribute where only the size of the gap matters
`sender / receiver`	v_i / v_j	Directed main effects: does the lawyer who asks, or is asked, have a characteristic?

The choice is part of the research question. A preference for partners and a preference for similar seniority are different ideas; the same indicator measures neither. With several attributes, mode() takes one entry per attribute, in order.

Partners as receivers, and seniority gaps

```
. set seed 12345
. nwqap lazega_adv lazega_fr partner seniority, ///
>         mode(receiver absdist) permutations(500)

                lazega_fr |      2.472     .002
           receiver_partner |       .647     .002
          absdist_seniority |      -.017     .002
                      _cons |     -1.955
```

Friendship: 2.47 here, 2.20 before. A coefficient controls only for what is in the model. In this model the seniority gap also stands in for office and practice, which are left out.

Partners as receivers

exp(.647) ≈ 1.9: a partner has about 1.9 times the odds of an associate of being asked, whoever asks.

Seniority gap

exp(−.017) ≈ .98 per rank step; over a gap of 50 ranks exp(−.85) ≈ .43.

Predicted tie probabilities

```
. nwqap lazega_adv lazega_fr office practice status, ///
>         permutations(500) predict(adv_fitted)
. nwtoedge adv_fitted lazega_adv
. gsort -adv_fitted _ego _alter
. tabulate lazega_adv if adv_fitted == adv_fitted[1]

 lazega_adv |      Freq.     Percent        Cum.
------------+-----------------------------------
          0 |         54       33.54       33.54
          1 |        107       66.46      100.00
      Total |        161      100.00
```

predict() stores the fitted values as a new valued network; nwtoedge turns networks into an edge list, one row per ordered pair, which can be sorted and tabulated.

.767

highest fitted probability, shared by 161 dyads

By hand

−3.228 + 2.198 + 1.062 + .752 + .407 = 1.191, and invlogit(1.191) ≈ .767.

Significant is not the same as accurate

.767

predicted probability of the top group of 161 dyads

66%

actual advice ties in that group: 107 of 161

.038

predicted probability with no friendship and a different office, practice and status

- All four coefficients are at the floor, yet the model overstates the probability for the top group by about ten percentage points.
- It is an additive logit: the same effects are added for every combination, and the data need not agree in each cell.
- This shows miscalibration in the fitted data; it says nothing about performance on new data.

An earlier wave as a predictor: the Glasgow girls

```
. nwwebuse glasgow, nwclear
. set seed 12345
. nwqap glasgow2 glasgow1 smoke1 sport1, permutations(200)

              glasgow2 |      Coef.    P-value
-----------------------+-------------------------
              glasgow1 |      3.653     .005
           same_smoke1 |       .514      .03
           same_sport1 |       .217     .378
                 _cons |     -4.125
```

Prior friendship is by far the strongest predictor. The model with smoking alone was not fitted, so how much its coefficient moved is left to an exercise of the chapter.

Predictor	Odds ratio
Wave-1 friendship	38.6
Same smoking	1.67
Same sport	1.24

Only 200 permutations: the p-values are coarse. One wave of attributes cannot separate selection from influence. With a dominant correlated predictor, use qapspp for the net association of smoking.

What QAP does not model

Ties that depend on ties

There is no term for “this tie is more likely because two other ties in the same network exist”: a returned request, a triangle closed.

Left in the error

Reciprocity and closure may load onto predictors that correlate with them. Friendship is the likeliest candidate: friends also return each other’s requests.

The next step

Exponential random graph models put reciprocity, closure and degree into the model as terms to estimate. Day 5.

If the question is only whether friendship and the formal organization go with advice, QAP with double semi-partialling is enough, and the write-up should say that it takes no position on how ties depend on one another.

Hands-on (25 min)

Ch. 20, Ex. 5

Fit the advice model without lazega_fr: set seed 12345, then nwqap lazega_adv office practice status, permutations(500). Report the three coefficients next to those of the four-predictor model, and say in which direction they moved and why that direction is expected. Why can the size of the change not be read as the share of the friendship–advice association that the attributes carry?

Hint Ask how friends are spread over offices, practices and status groups, and what an attribute coefficient can pick up when friendship is no longer in the model. For the last question, recall the crude odds ratio of 11.9 against 9.0: odds ratios across models are not on a common scale.

Ch. 20, Ex. 7

The coefficient of same_office, 1.06, is read as homophily, a preference for colleagues who are like oneself. Name one other process that would give the same positive coefficient, and say what data on the lawyers would tell the two apart.

Hint Ask what sharing an office does for two lawyers on an ordinary working day, apart from making them alike. For the data, ask what each process predicts for a lawyer who moves to another office.

Hands-on (optional extra, 10 min)

Ch. 20, Ex. 3 (optional extra)

(a) Why are the five highest fitted probabilities in the predict() example all 0.767, and how many different fitted probabilities can this model give at most? (b) Rerun it with nwtoedge adv_fitted lazega_adv lazega_fr, egovars(office practice status) altervars(office practice status) and check your count with codebook adv_fitted, compact. (c) Show with summarize adv_fitted lazega_adv that the mean fitted probability equals the density of advice, and explain why. What does this say, and not say, about the top group of 161 dyads, of which only 66% are advice ties? (d) Generate an indicator for each attribute that is 1 when ego and alter agree (generate same_office = office_ego == office_alter) and collapse to the cells of friendship and the three indicators with collapse (mean) fitted=adv_fitted observed=lazega_adv (count) dyads=lazega_adv, by(lazega_fr same_office same_practice same_status). Among the cells with at least 50 dyads, which has the largest gap between fitted and observed share, and which assumption of an additive logit does that gap point to?

Hint (a) Each of the four predictors takes two values. (c) The likelihood equations of a logit with a constant force the fitted probabilities to add up to the observed number of ties; a total says nothing about single cells. (d) Compare the friendship cells with the same attribute pattern without friendship, and ask whether friendship adds the same log-odds in every row; a cell with fewer than about 50 dyads is too thin to read.

Hands-on (further practice, 10 min)

Practice 1

Turn the model around: friendship as the outcome. After set seed 12345, run nwqap lazega_fr lazega_adv office practice status, permutations(500). Is advice a predictor of friendship in the same way? Compare the coefficients and p-values with the advice model.

Hint Compare lazega_adv here with lazega_fr there, and look at the p-value of same_status.

Practice 2

Add gender to the advice model: set seed 12345, then nwqap lazega_adv lazega_fr office practice status gender, permutations(500). Which coefficients change, and by how much?

Hint gender is expanded to same_gender by the default mode(same).

Practice 3

Rerun the advice model with permutations(100) after set seed 12345. What are the p-values, why, and how long did the run take compared with permutations(500)?

Hint The floor of nwqap is 1/(R+1). Use timer on 1 and timer off 1, then timer list.

What to take away from the fourth block

Same estimates

QAP keeps the logit's point estimates and replaces the p-values with permutation-based ones.

Report the routine

Plain QAP or qapspp, the number of permutations and mode(). With correlated predictors the method matters.

Significant is not accurate

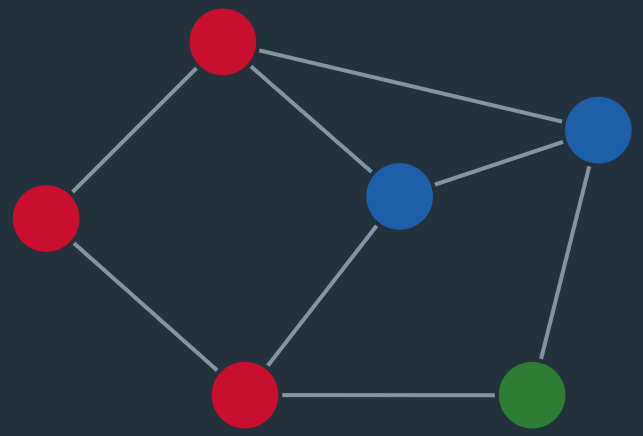
A model can have every coefficient at the floor and still miscalibrate a group of pairs. predict() shows it.

The firm: net of office, practice and status, a friendship tie goes with about 9.0 times the odds of an advice tie; the shared attributes with 2.9, 2.1 and 1.5. Advice overlaps with the formal organization without being reducible to it.

5

Choosing a test, and reporting it

30 min · Chapters 18 to 20



Three tests, three assumptions about what may be exchanged

Test	What is exchanged	Assumption	Command
CUG test (Ch. 18)	The observed network is replaced by random networks sharing density, census or degrees	Ties are exchangeable once the stated features are fixed	nwcug
Permutation (Ch. 19)	The labels of the nodes of the observed network	Lawyers are interchangeable; the structure is kept	nwcorrelate, nwmixing
QAP regression (Ch. 20)	The labels of the outcome network, then the model is refitted	As permutation; predictors keep their correlations unless qapspp is used	nwqap

Generating random networks (Chapter 18), relabeling an observed one (Chapter 19) and permuting within a regression (Chapter 20) are three different claims about what could have been otherwise.

Which command for which question

Question	Command	Key options
Is a statistic unusual for networks of this size and density?	nwcug	condition(density), reps(), seed()
Is it unusual given the dyad census (directed networks only)?	nwcug	condition(census)
Is it unusual given every node's degrees?	nwcug	condition(degrees), trades()
Is the relation between two networks, or a network and an attribute, stronger than chance?	nwcorrelate, nwmixing	permutations()
Same, with several predictors at once?	nwqap	permutations(), mode(), qapspp

Ties that depend on other ties? That is the ERGM of Day 5.

The law firm: what the day found

.329

transitivity of advice: null .179 (density), .272 (degrees)

.274

reciprocity of advice: null .099 (density), .136 (degrees)

.418

correlation of advice with friendship; relabelings: −.031 to .033

9.0

odds ratio of friendship, net of office, practice, status

None of 1,000 draws or 1,000 or 500 relabelings reached the observed values. Size and density do not account for the clustering and reciprocity; degrees account for part of the clustering; friendship goes with advice even net of the formal organization.

What the tests leave open: whether lawyers close triangles on purpose, and whether friends become advisors or advisors friends.

What to write in a results paragraph

The null

What the random networks hold fixed: condition(density), census or degrees, or relabeling of the nodes.

The draws

The number of repetitions, the seed, and the p-value as an inequality when it is at the floor: $p \leq 2/1001$ with 1,000 draws.

The routine

For a QAP: plain or qapspp, mode() for attributes, and which predictors are in the model.

The claim

About this network, not a population; more than chance at that level, not evidence for a mechanism.

Take the numbers from r() and e(), never retype them (Chapter 30).

Hands-on (10 min, in pairs)

Ch. 18, Ex. 6

A colleague reports that the clustering of a network of trade between countries is “significant” against a density null. Write three questions you would ask before accepting that this shows a social mechanism such as closure, and for each say which test of Chapter 18, or which further data, would answer it.

Hint Start from what a density null leaves free to vary. Then ask what else, besides closure, could make two trading partners of a country trade with each other: countries that trade with very many others, and attributes that sort ties.

Ch. 20, Ex. 9

A colleague fits an ordinary logit to all 4,970 dyads of the advice network and reports its standard errors as they come. Write a short paragraph that explains what is wrong with those standard errors, and say which output of nwqap, and of its option plot, you would show a skeptical reader to make the point.

Hint The rows are not independent observations (the dyadic regression section of Chapter 20); the permutation distributions that replace the standard errors are shown in the section on visualizing them.

Next: ties that depend on other ties

What today could not do

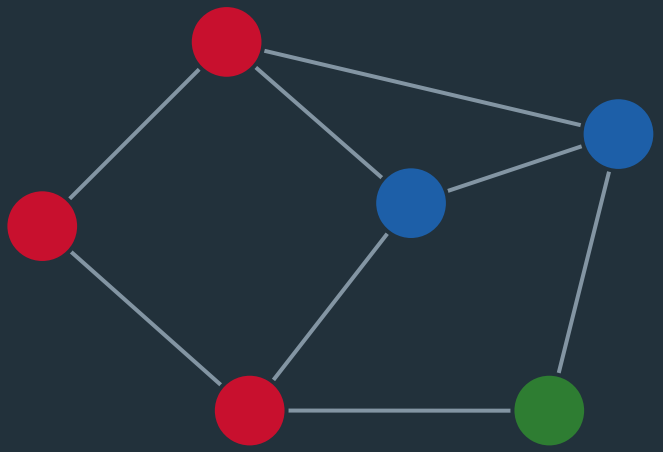
QAP leaves reciprocity and closure in the error. The CUG test says clustering is more than density gives, not why.

Day 5: exponential random graph models

Make the dependence among ties the object of the model. Coefficients are read as conditional log-odds through change statistics; fit is checked on statistics the model does not contain.

Overnight (optional): rerun a test of your own network with set seed and reps(), and state in one line what the null holds fixed.

To read after the session



Chapter 15: sections on density and random graphs

Chapters 18, 19 and 20: random networks and null models; permutation inference; QAP regression

Exercises: Ch. 15 Ex. 2 and 7; Ch. 18 Ex. 1, 3, 5 and 6; Ch. 19 Ex. 2, 4 and 5; Ch. 20 Ex. 3, 5, 7 and 9

Tomorrow: Chapter 21, exponential random graph models

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

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