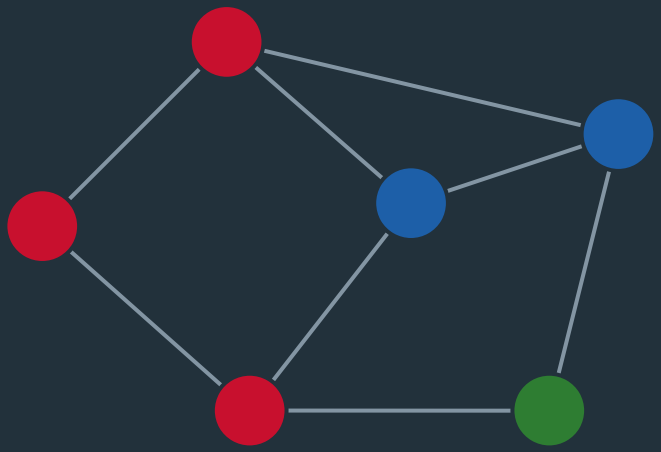
SAOM: adequacy, estimators and several groups



Course B, Week 10

Chapter 26 · a lecture and a lab

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

Where Weeks 8 and 9 left the Glasgow model

Week 8: change

Ties are kept, formed or dissolved. An SAOM models actors' small choices: a rate function (how often) and an objective function (which change).

Week 9: a model, checked

Convergence t-ratios, z tests, goodness of fit, co-evolution. The three-effect Glasgow model converged and failed on geodesic distances.

Week 10: is it adequate?

Which effect is missing? Do effects stay constant? Does the estimator matter? Can the groups of a larger study be pooled?

28.1

Mahalanobis distance of the geodesic distances, p = 0.000: the model fails there

The failure on the geodesic distances makes the first worry concrete: the model has been checked on the statistics it was built to reproduce and on little else.

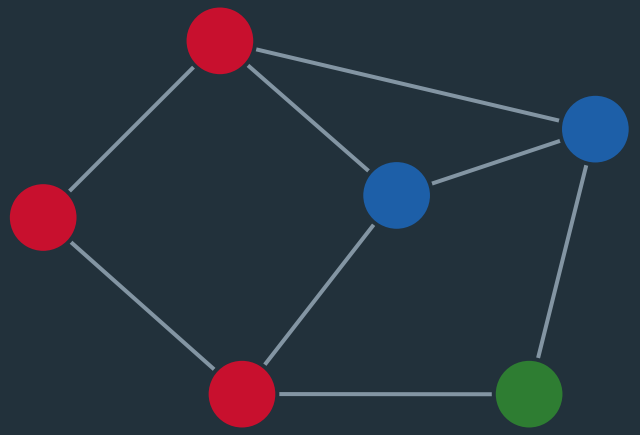
Today: a lecture, a break and a lab

0:00	**1 Is the specified process adequate?**	25 min
0:25	**2 What does the fitted model imply?**	15 min
0:40	**3 Which estimator?**	25 min
1:05	**4 Several groups: pool or not?**	25 min
1:30	Break	10 min
1:40	**5 Lab: checks, estimators and groups**	60 min

1

Is the specified process adequate?

25 min · Chapter 26: estat score, estat timetest, estat importance



A coefficient table can be wrong in ways it does not show

A tendency left out

Actors follow a rule the model lacks; the effects it has absorb it and their coefficients are biased. Tool: estat score.

Constant rules

One reciprocity coefficient averages over periods that differ. Tool: estat timetest.

A less precise estimator

Another estimator fits the same model with smaller standard errors. Tool: ml and bayes.

One group, many groups

A model for one class says nothing about whether other classes behave alike. Tool: groups() and meta.

Simulation shows what a fitted model implies; none of these tools replaces the goodness-of-fit test of Week 9.

estat score: fit the restricted model, test what it leaves out

```
. nwwebuse glasgow, nwclear
. quietly nwsaom, waves(glasgow1 glasgow2 glasgow3) outdegree reciprocity ///
>         transtrip seed(12345)
. quietly nwsaom refit, seed(12346)
. matrix b = e(b)
. quietly nwsaom, waves(glasgow1 glasgow2 glasgow3) outdegree reciprocity ///
>         transtrip indegpopularity outactivity cycle3 ///
>         theta0(`=b[1,1]' `=b[1,2]' 0 0 `=b[1,3]' 0) ///
>         fix(indegpopularity outactivity cycle3) seed(1)
```

The zeros in theta0() follow the order of the table: outdegree, reciprocity, indegpopularity, outactivity, transtrip, cycle3.

0.041

overall maximum convergence ratio of the restricted fit (criterion: below 0.25)

- **fix()**

  holds an effect at its starting value, here 0, so it is not estimated.

Score test: outdegree activity and indegree popularity

```
. estat score

         Effect |  Fixed at      chi2   df  Prob > chi2  One-sided z
----------------+-----------------------------------------------------
indegpopularity |    0.0000     6.554    1       0.0105       -2.560
    outactivity |    0.0000    11.835    1       0.0006       -3.440
         cycle3 |    0.0000     0.164    1       0.6853       -0.405
----------------+-----------------------------------------------------
          Joint |             13.900    3       0.0030
```

.0006

outdegree activity: the strongest candidate

.69

three-cycles: not flagged, with 50 pupils

chi2

With one degree of freedom, a value above 3.84 is significant at 5%.

One-sided z

The sign says in which direction the coefficient would move if freed: negative, below its fixed value.

Joint

All three fixed effects together: chi2 13.9 on 3 degrees of freedom, p = .0030.

Add the strongest candidate, refit, test again

It says

Freeing indegree popularity or outdegree activity would improve the fit (p = .0105 and .0006; joint p = .0030), and both coefficients would move down: ties are spread more evenly than the model produces.

It does not say

That cycles are absent (p = .69 with 50 pupils), how large the coefficients would be (one-step estimates are linear extrapolations: outdegree would move from −2.68 to −1.56), or that all candidates can be freed at once.

The next step

Free outactivity, the strongest candidate.

Refit.

Test the remaining two again.

Run estat gof: do the geodesic distances improve?

Each test is conditional on the effects in the model at that moment.

The time test: no change detected is not constancy

```
. quietly nwsaom, waves(glasgow1 glasgow2 glasgow3) outdegree reciprocity ///
>         transtrip seed(12345)
. quietly nwsaom refit, seed(12346)
. estat timetest, fixrates

      Effect |          chi2        df        Prob > chi2
-------------+----------------------------------------------
   outdegree |         0.152         1             0.6963
 reciprocity |         1.319         1             0.2508
   transtrip |         0.649         1             0.4205
       Joint |         4.267         3             0.2340
```

fixrates treats the rates as known, as RSiena's time test does.

Joint p = .23

No effect shows time heterogeneity. That is not evidence that the rules stayed the same.

+0.71

The second-period reciprocity coefficient is 0.71 above the first (p = .25), over a quarter of the pooled 2.48.

Importance: outdegree and reciprocity carry the choices

```
. estat importance

      Effect | Estimate  Std. err.    Sigma  Semi-std.     RI p1     RI p2
-------------+-----------------------------------------------------------
   outdegree |  -2.6819     0.1183   0.4384    -1.1757    0.4434    0.5139
 reciprocity |   2.4768     0.1973   0.2140     0.5300    0.3907    0.3216
   transtrip |   0.6166     0.0790   0.6514     0.4017    0.1659    0.1645
-------------+-----------------------------------------------------------
   Certainty |                                           0.2683    0.3487
```

Semi-standardized

Estimate times sigma puts the effects on one scale: per standard deviation, outdegree weighs about twice as much as reciprocity, in the opposite direction.

Relative importance

Reciprocity and outdegree carry most of the choices (RI .44 and .39 in period 1); transitive triplets about a sixth, though clearly nonzero.

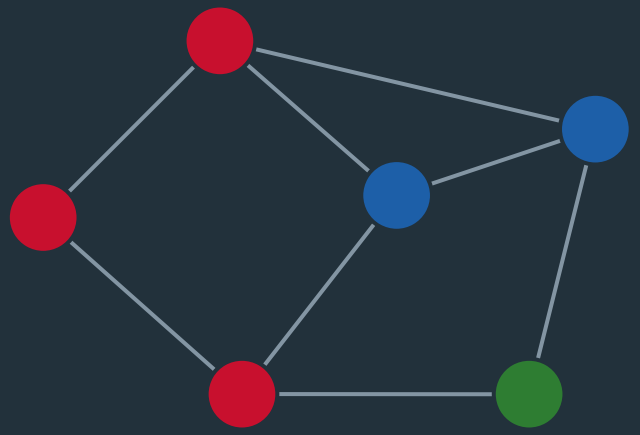
Certainty R_H

0.27 and 0.35: low. The model leaves most of each pupil's choice open.

2

What does the fitted model imply?

15 min · Chapter 26: estat simulate, estat mems, nwsaom simulate



estat simulate: the model run forward from the observed wave

```
. quietly nwsaom, waves(glasgow1 glasgow2) outdegree reciprocity ///
>         transtrip seed(4)
. estat simulate, nsim(100) seed(5)

   Statistic |       Sim       Obs
-------------+--------------------
Period 1     |
        Ties |     115.9       116
      Mutual |      35.5        35
  Transitive |      92.7        88
```

A wrong model has wrong implications that look just as precise. The adequacy check remains estat gof and the tests of block 1.

Sim next to Obs

Sim is the mean of 100 simulated end waves; Obs is the observed second wave.

Close, as expected

The estimation chose the coefficients to match these statistics, so agreement mostly confirms convergence.

estat mems: the network without reciprocity

```
. program mutualties, rclass
  1.     args netname
  2.     nwtomata `netname', mat(MM)
  3.     mata: st_numscalar("__mut", sum(MM :* MM') / 2)
  4.     return scalar stat = scalar(__mut)
  5. end
. estat mems, effect(reciprocity) macro(mutualties) nsim(100) seed(1)
```

```
             |  Estimate  Std. dev.    [95% conf. interval]  MC p-value
-------------+------------------------------------------------------------
        MEMS |   37.1000    11.8973     17.0000     62.6250      0.0000
Prop. change |    0.9716
```

37.1

fewer mutual dyads at the end of the period without the reciprocity effect

17 to 63

95% interval; the size is rough, the sign is not in doubt

MC p-value 0.0000 means below 1/101: none of the 100 draws had the opposite sign.

nwsaom simulate: what do chosen coefficients mean?

```
. nwsaom simulate, wave1(glasgow1) wave2(glasgow2) outdegree reciprocity ///
>         transtrip theta(-2.3 2 0.5) rates(5) nsim(100) seed(7)
```

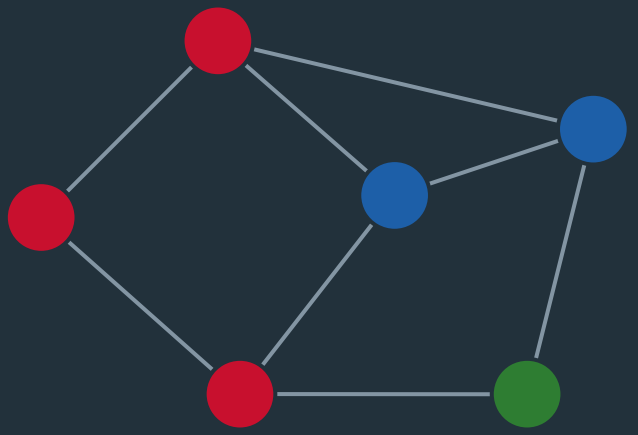
Coefficients (outdegree, reciprocity, transtrip) and rate	Ties	Mutual	Transitive
Observed second wave	116	35	88
theta(-2.3 2 0.5), rates(5)	129.5	37.2	81.3
theta(-2.3 0 0.5), rates(5): no reciprocity	64.3	2.4	6.3
theta(-2.3 2 0.5), rates(1): slower	121.0	37.8	77.5

Without the reciprocity term the ties halve too. A coefficient does not move only its own statistic: new ties are no longer held in place.

3

Which estimator?

25 min · Chapter 26: Method of Moments, maximum likelihood (ml), Bayesian estimation (nwsaom bayes)



Both likelihood-based estimators sample the hidden path

Method of Moments

The default and the fastest. Solves the moment equations. Convergence: t-ratios of the statistics (Weeks 8 and 9).

Maximum likelihood

Option ml. Starts from the Method-of-Moments fit, then solves the likelihood equations of the augmented data. Convergence: t-ratios of the scores.

Bayesian

nwsaom bayes. Samples the posterior of the same model with the same chain. Diagnostic: the effective sample size of the draws.

Both likelihood-based estimators use the same chain over hidden paths (data augmentation); they differ in what they return.

Maximum likelihood: the same model, a different criterion

```
. nwsaom, waves(glasgow1 glasgow2 glasgow3) outdegree reciprocity ///
>         transtrip ml seed(12345) mloptions(n3(3000))

Maximum likelihood estimates (Robbins-Monro on the score of the augmented data)
        SAOM | Coefficient  Std. err.      z    P>|z|
-------------+----------------------------------------
   outdegree |  -2.655079   .1063517   -24.97   0.000
 reciprocity |   2.258475    .180527    12.51   0.000
   transtrip |   .6932367   .0676438    10.25   0.000
...
Overall maximum t-ratio of the scores: 0.099 (good; below 0.1)
```

0.099

maximum t-ratio of the scores with n3(3000); the default run gave 0.129

77 s

run time of this call on a laptop, much longer than the Method of Moments

The longer last phase left the estimates unchanged and brought every score below 0.1. If it had not, the next step is a longer run again: mloptions(n3(#)), nstart2(#).

Bayesian estimation returns draws, not one estimate

```
. quietly nwsaom, waves(glasgow1 glasgow2 glasgow3) outdegree reciprocity ///
>         transtrip seed(12345)
. quietly nwsaom refit, seed(12346)
. nwsaom bayes, nmain(250) seed(1)
. frame nwsaom_bayes_draws: summarize b1 b2 b3
```

Prior and run length	Setting
Effects	flat
Rates	normal, sd 1 (ratepriorsd())
Draws	50 + 50 + 250 main iterations
Run time	about 110 seconds

The posterior mean

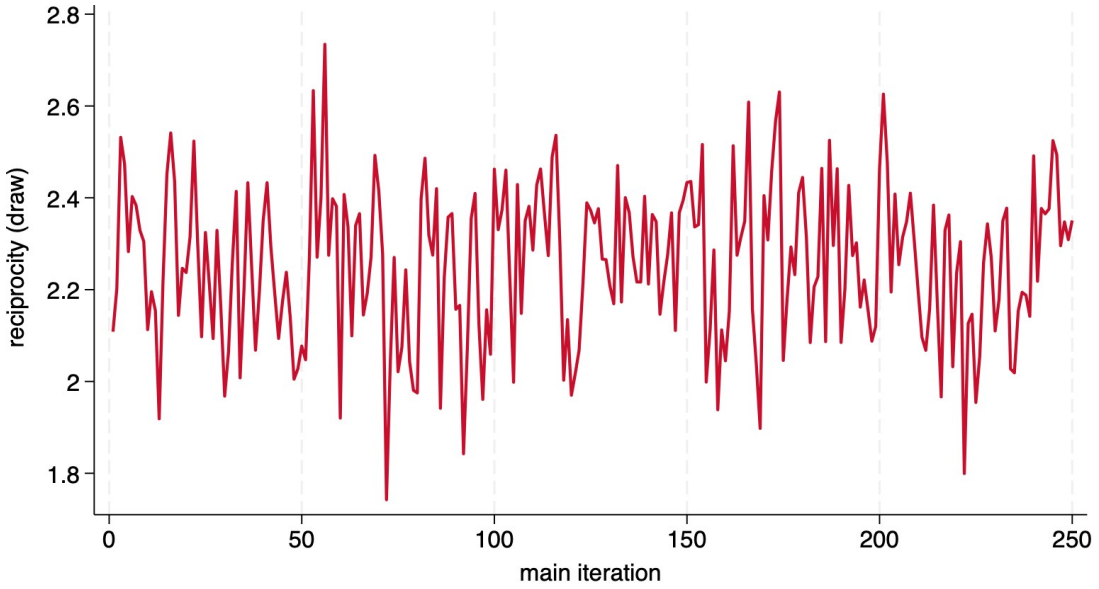
is the point estimate; the posterior standard deviation plays the role of the standard error.

A prior

is flat for the effects; the rates get a normal prior with standard deviation 1, so their posterior sd is partly the prior's.

250 draws are worth 86 to 151 independent ones

Effect	Mean	Std. dev.	[95% cred. interval]		Median	ESS
outdegree	-2.6567	0.1106	-2.8632	-2.4501	-2.6566	137
reciprocity	2.2560	0.1687	1.9391	2.5401	2.2710	151
transtrip	0.6963	0.0634	0.5773	0.8212	0.6945	141
rate_period1	6.4376	0.6929	5.1809	7.7413	6.4611	86
rate_period2	5.2104	0.6245	4.0034	6.5777	5.2214	137



Draws of the reciprocity coefficient, seed(1): consecutive draws are similar.

- ESS is the number of independent draws that would carry the same information.
- 86 to 151 of 250: enough for means and sds to a first decimal, not for the ends of the intervals.
- For a paper increase nmain() until every ESS is in the hundreds.

Three estimators, one substantive answer

	Method of moments	Maximum likelihood	Bayesian (posterior mean)
outdegree	−2.682 (0.118)	−2.655 (0.106)	−2.657 (0.111)
reciprocity	2.477 (0.197)	2.258 (0.181)	2.256 (0.169)
transtrip	0.617 (0.079)	0.693 (0.068)	0.696 (0.063)
rate, period 1	6.49 (1.07)	6.47 (1.00)	6.44 (0.69)
rate, period 2	5.19 (0.89)	5.30 (0.90)	5.21 (0.62)

Likelihood-based pair

Maximum likelihood and Bayes are within a tenth of a standard error on every effect and both rates.

The moment estimate sits apart

Its reciprocity estimate is 0.22 higher, over one standard error of the likelihood-based ones.

Same substance

Strong reciprocity, clear transitivity, a negative baseline for tie formation, under all three.

Which estimator to use

Method of Moments

Keep it when its convergence check is met and the question needs no likelihood-based quantity. The default and the fastest.

Maximum likelihood

When restarts leave a poor t-ratio, or a likelihood-based interval matters for the argument.

Bayesian

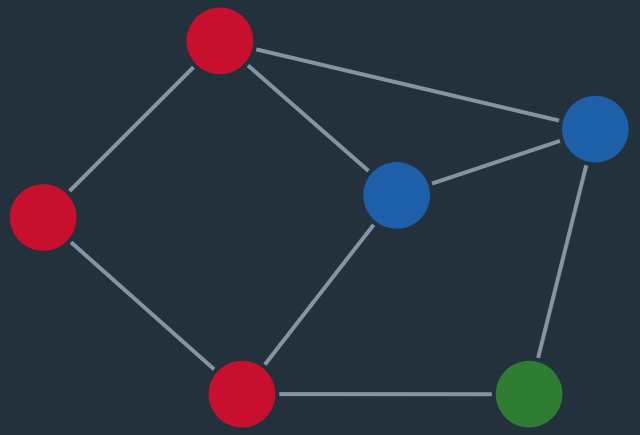
When we want the posterior distribution itself and the effective sample sizes as a diagnostic of how far the chain can be trusted.

All three inherit what the specification leaves out. Agreement among them is not a test of adequacy.

4

Several groups: pool or not?

25 min · Chapter 26: nwstack, groups(), nwsaom meta, multinet; 19 Dutch schools (baerveldt19)



Joint model or meta-analysis: two assumptions

Joint model: groups(school)

Same rules at different speeds. All effects are shared, every period of every group has its own rate. Works with small groups and a behavior. Small standard errors, because it does not estimate how much the groups differ.

Meta-analysis: nwsaom meta

Each group fitted alone, the estimates combined: a mean, a between-group sd and a heterogeneity test Q. No assumption that groups share their effects, but every group must converge on its own.

When the heterogeneity is significant, the joint standard errors, which assumed equal effects, are too optimistic: the variation between groups is part of the uncertainty about the average effect.

19 schools, 990 pupils: stack the networks into one

```
. nwwebuse baerveldt19, nwclear
. clear
. nwstack school*_w1, name(w1) group(school)
. nwstack school*_w2, name(w2)
. nwstack school*_ethn, name(ethn)
. bysort school: generate pupil = _n
. merge 1:1 school pupil using "$nwwebpath/baerveldt19_pupils.dta", nogenerate
. nwtomata ethn, mat(E)
. mata: st_matrix("Eth", E)
. nwsaom describe, waves(w1 w2)
```

990

pupils in 19 schools of 31 to 91 pupils

0.288

Jaccard index of the two waves: many ties change, few stay

One block-diagonal network per wave, no ties between schools. Ethnicity must become a Stata matrix for the joint model; the Jaccard index below 0.3 calls for careful convergence checks.

The joint model converges, with rates 2.3 to 12.0

```
. quietly nwsaom, waves(w1 w2) groups(school) outdegree reciprocity ///
>         transtrip seed(1)
. nwsaom refit, seed(2)

        SAOM | Coefficient  Std. err.      z    P>|z|
-------------+----------------------------------------
   outdegree |  -2.406625   .0293444   -82.01   0.000
 reciprocity |   2.129775   .0582958    36.53   0.000
   transtrip |   .5098972   .0276183    18.46   0.000
Overall maximum convergence ratio: 0.060 (criterion: below 0.25; met)
```

2.3 to 12.0

the 19 school rates: the schools differ in speed by a factor of about five

Small standard errors: every school contributes to one set of coefficients. They describe one common process; whether it is common is what the meta-analysis tests.

Fit each school on its own, twice, and store every fit

```
. matrix C = J(19, 3, .)
. local fits ""
. forvalues s = 1/19 {
.     quietly nwsaom, waves(school`s'_w1 school`s'_w2) outdegree reciprocity ///
>         transtrip seed(1)
.     quietly nwsaom refit, seed(2)
.     estimates store school`s'
.     matrix b = e(b)
.     matrix C[`s', 1] = b[1, 1], b[1, 2], b[1, 3]
.     local fits `fits' school`s'
. }
```

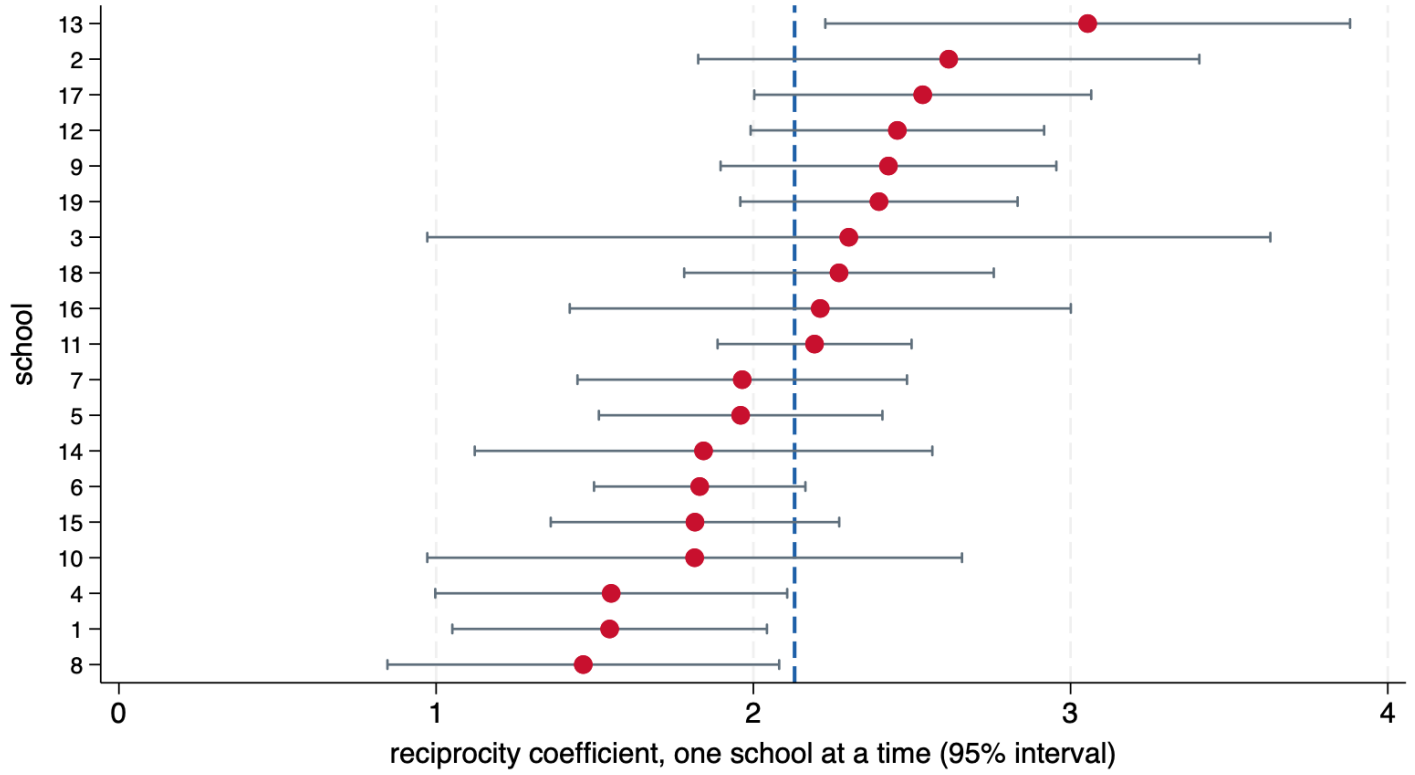
Run twice

With seed 1 alone, five of the 19 schools missed the 0.1 criterion for at least one effect.

Converged first

A fit that did not converge adds an error that the heterogeneity statistic counts as a difference between schools.

The schools differ: reciprocity from 1.46 to 3.05



Separate fits, 95% interval; the vertical line is the joint estimate 2.13.

1.46 to 3.05

reciprocity in the 19 separate fits

0.34 to 1.41

transitive triplets (transtrip)

−3.36 to −1.97

outdegree

nwsaom meta: the heterogeneity test rejects equal effects

```
. nwsaom meta `fits'

      Effect |  N      Mean Std. err.        t   P>|t|      SD       Q     P>Q
-------------+-------------------------------------------------------------
   outdegree | 19   -2.4144    0.0653   -36.99   0.000  0.2845   53.42   0.000
 reciprocity | 19    2.0902    0.0897    23.31   0.000  0.3909   33.82   0.013
   transtrip | 19    0.5790    0.0465    12.46   0.000  0.2025   36.15   0.007
```

Q is significant for all three

Equal effects are rejected: outdegree p < .001, reciprocity p = .013, transtrip p = .007 (18 degrees of freedom).

Small against the means

Between-school sd 0.39 for reciprocity against a mean of 2.09; direction and size agree with the joint model.

A warning, not a measurement

Q says the groups are not alike; with 19 schools the between-school sd is a rough guide.

Joint standard errors are too optimistic by a factor of 1.6 to 2.2

Standard error of the mean effect	outdegree	reciprocity	transtrip
Joint model, groups(school)	.029	.058	.028
Meta-analysis, nwsaom meta (IWLS)	.065	.090	.046
Ratio meta to joint	2.2	1.6	1.6

It says

The joint standard errors describe one common process; with significant heterogeneity they are too optimistic as a statement about schools in general. The meta-analytic standard error includes the variation between schools.

It does not say

That the joint model is wrong: the means agree (reciprocity 2.13 against 2.09) and the direction is the same. A non-significant Q would not show equal rules either: with few groups the test has little power.

A significant Q: let the effect depend on the group

Fit, check, report

Joint model first, then the meta-analysis on the same effects. If Q is significant, report the meta-analytic mean with its standard error next to the joint one.

Let the rules differ

Interact the effect with a group indicator or a characteristic of the group (size, school): egox() adds the main effect, interact() multiplies it with the effect.

Two networks: multinet

Two networks on the same actors that evolve together, with cross-network effects such as crprod. The second network must be a relation measured in its own right.

```
nwsaom multinet, net1(f1 f2 f3) net2(a1 a2 a3) behavior(b1 b2 b3) eff1(outdegree reciprocity transtrip) ///
        eff2(outdegree reciprocity) cross1(crprod) beheff1(linear quadratic avalt)
```

Syntax as in Chapter 26; the Glasgow data have no second relation, so it is not run.

What to take away from the lecture

Adequacy

The score test says which effect to try next, from the restricted model alone. The time test with few periods and actors shows little when it fails to reject.

Implications

Simulation and the importance table show what a model implies and which effects matter. Neither tests adequacy; both inherit the model's errors.

Estimators

Compare the three on one model; read each one's convergence diagnostic before the table. ESS in the hundreds before trusting a Bayesian interval.

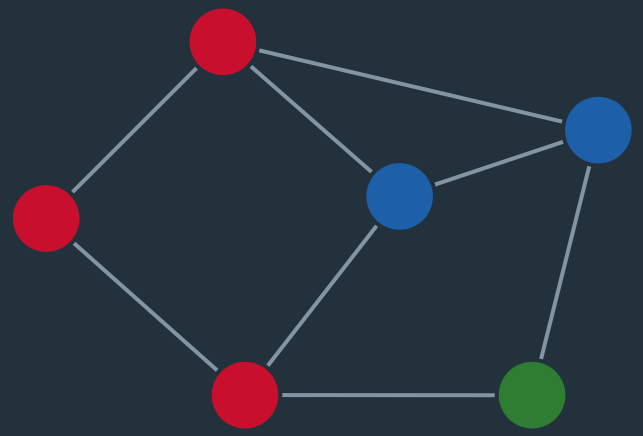
Groups

Joint model first, meta-analysis as the check. A significant Q: interact the effect with the group, report both standard errors.

Four mistakes to avoid: freeing every flagged effect at once; reading a short Bayesian run to the last digit; a cross-network effect for nearly nested networks; joint standard errors without a heterogeneity check.

Break

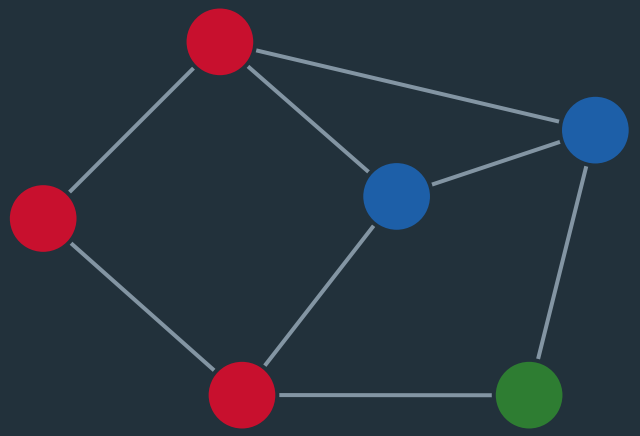
10 minutes



5

Lab: checks, estimators and groups

60 min · Run along: restart, time test, importance, joint model; Ch. 26, Ex. 1, 2, 3 and 4



Run along 1: restart, time test, importance

```
. nwwebuse glasgow, nwclear
. quietly nwsaom, waves(glasgow1 glasgow2 glasgow3) outdegree reciprocity ///
>         transtrip seed(12345)
. quietly nwsaom refit, seed(12346)
. estat timetest, fixrates
. estat importance
```

Predict first. What will the joint p-value of the time test be: below .05 or above?

With 50 pupils and two periods the test cannot detect a deviation of a quarter of the pooled coefficient.

Checkpoint	Value
time test, joint chi2 (3 df)	4.267
time test, joint p	0.2340
reciprocity deviation, period 2	0.7056
relative importance p1: outdegree	0.4434
certainty R_H, p1 and p2	0.2683, 0.3487

Run along 2: stack the schools, fit the joint model

```
. nwwebuse baerveldt19, nwclear
. clear
. nwstack school*_w1, name(w1) group(school)
. nwstack school*_w2, name(w2)
. nwstack school*_ethn, name(ethn)
. bysort school: generate pupil = _n
. merge 1:1 school pupil using "$nwwebpath/baerveldt19_pupils.dta", nogenerate
. nwtomata ethn, mat(E)
. mata: st_matrix("Eth", E)
```

```
. nwsaom describe, waves(w1 w2)
. quietly nwsaom, waves(w1 w2) groups(school) outdegree reciprocity ///
>         transtrip seed(1)
. nwsaom refit, seed(2)
. estimates store joint
```

Checkpoint	Value
Jaccard, period 1->2	0.288
outdegree	−2.4066
reciprocity	2.1298
transtrip	.5099
overall ratio after refit	0.060

Start from an empty data set: nwstack creates one observation for each pupil, and the merge comes before the joint model. The 19 separate fits and nwsaom meta take minutes: the instructor starts them at the beginning of the break.

Hands-on (20 min)

Ch. 26, Ex. 3

The score test flagged outdegree activity most clearly (p = .0006). Free outactivity and keep the other two candidates at 0: nwsaom, waves(glasgow1 glasgow2 glasgow3) outdegree reciprocity transtrip indegpopularity outactivity cycle3 fix(indegpopularity cycle3) seed(1). Report its coefficient and p-value, then run estat score and estat gof, nsim(100) seed(1). Is indegree popularity still significant at 5%, and has the geodesic distribution improved on Chapter 25's model (Mahalanobis 28.1, p = 0.000)?

Hint Effects held at 0 need no theta0(). A fixed effect is not estimated, and its t-ratio is the deviation the score test measures, so judge convergence by the overall ratio and the freed effects.

Ch. 26, Ex. 1

In the table of three estimators the Method-of-Moments estimate of reciprocity is 2.477 and the maximum likelihood estimate 2.258, with a standard error of 0.181. Express the difference in standard errors of the maximum likelihood estimate. Then write the one or two sentences you would put in a paper: which estimate you report, and what you say about the other.

Hint Divide the difference by 0.181. The paragraph on whether the three estimators agree says which estimate to report when they differ by more than a standard error, and on what condition.

Ch. 26, Ex. 3 and 1: the answers

Ex. 3: the next step	Chapter 25's model	outactivity freed
coefficient of outactivity	(not in model)	−.170, p < .001
indegree popularity, score test	p = .0105	p = .075: not at 5%
geodesic distances, Mahalanobis	28.1, p = .000	15.1, p = .010
overall convergence ratio		0.132 (met)

Ex. 3: improved, not adequate

The geodesic distance halves but is still rejected (p = .010). Test again, add the next effect, run estat gof.

Ex. 1: 1.21 standard errors

(2.477 − 2.258) / 0.181. Report the maximum likelihood estimate if its run converged (it did, 0.099); say the moment estimate is higher, by 1.2 se.

Hands-on (20 min)

Ch. 26, Ex. 2

Rerun the Bayesian estimation twice, with seed(2) and seed(3) in place of seed(1), each time after the same restarted three-effect fit. Put the posterior means, posterior standard deviations and ESS of the three effects next to those of the chapter's run. Which of these changes more from run to run, measured against the posterior standard deviation? What would you change in the call to make the less stable ones more reliable?

Hint Each run: the quiet fit with seed(12345), the quiet nwsaom refit, seed(12346), then nwsaom bayes, nmain(250) seed(2); one run takes one to two minutes. Divide the difference between two posterior means by the posterior standard deviation. The paragraph after the table of draws names the option to increase.

Ch. 26, Ex. 4

A research team has two waves of friendship data on 40 school classes of 20 to 30 pupils each. Say whether you would fit a joint model with groups(), a meta-analysis with nwsaom meta, or both, and in which order. List two things you would check before reporting a single pooled coefficient for reciprocity.

Hint The section on pooling gives the condition each procedure needs: the meta-analysis needs every class to converge on its own, the joint model assumes the classes share their effects. Ask which output would show that a class cannot be estimated alone, and which that classes do not share one reciprocity effect.

Ch. 26, Ex. 2 and 4: the answers

Ex. 2: seeds 1, 2, 3	outdegree	reciprocity	transtrip
posterior mean	−2.657, −2.653, −2.641	2.256, 2.230, 2.223	.696, .706, .697
posterior sd	.111, .102, .101	.169, .187, .175	.063, .062, .065
ESS	137, 93, 129	151, 38, 127	141, 91, 93

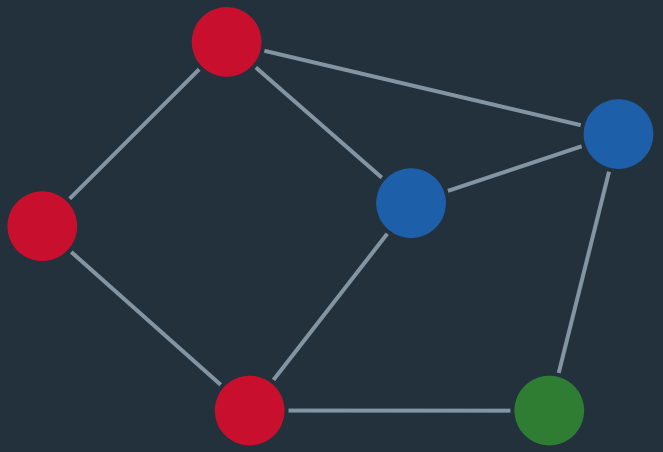
Ex. 2: the ESS moves most

The means differ by 0.14 to 0.19 posterior sd, the sds by up to 11%, the ESS of reciprocity by a factor of four (38 to 151). Increase nmain() until every ESS is in the hundreds.

Ex. 4: joint first, meta as the check

Check that each class converges alone (and has enough change), and the heterogeneity test Q with the ratio of the meta-analytic to the joint standard error. Report both if Q is significant.

To read after the session



Week 11: Chapters 27 and 28 (relational event models), the first part of the path Change and events

Left over from Chapter 26: Ex. 5 (answering a reviewer about a significant heterogeneity test)

Re-read the section on pooling the groups before the capstone (Week 12)

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

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