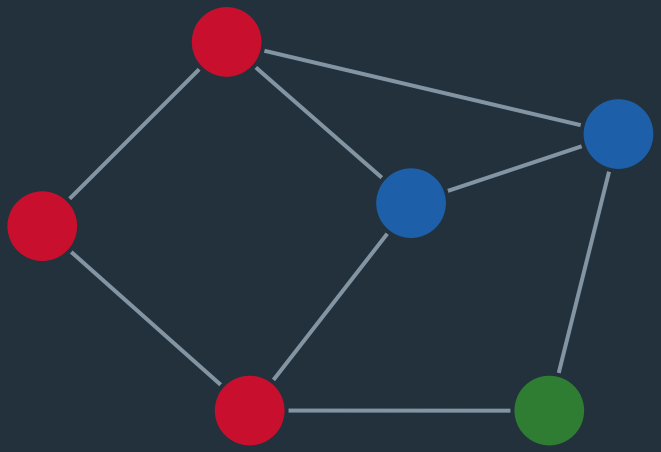
Stochastic actor-oriented models



Course B · Week 9

Chapter 25 of Network Analysis with Stata (Stata Press), from building up a model on Glasgow to the end · one lecture and one lab, about 90 minutes each

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

Today: a lecture, a break and a lab

0:00	1	**Building up a model on Glasgow**	20 min
0:20	2	**Convergence is not significance**	20 min
0:40	3	**Does the model fit?**	20 min
1:00	4	**Selection or influence**	30 min
1:30		Break	10 min
1:40	5	**Lab: fit, restart, check**	45 min
2:25	6	**Lab: co-evolution**	45 min

Week 8 ended with a first model: a rate, two effects, a converged fit

```
. nwwebuse glasgow, nwclear
. nwsaom, wave1(glasgow1) wave2(glasgow2) ///
>         outdegree reciprocity seed(1)
Actors: 50                    Estimated rate:  5.552 (0.856)
   outdegree |   -2.26413    .127996   -17.69   0.000
 reciprocity |    2.50883   .2328995    10.77   0.000
Overall maximum convergence ratio: 0.178 (criterion: below 0.25; met)
Convergence t-ratios (criterion: all below 0.1 in absolute value): met
```

Today. Add effects, check that the estimation settled, check the fit, and let drinking change too.

Rate

How often an actor may change a tie: 5.55 opportunities per pupil between the waves.

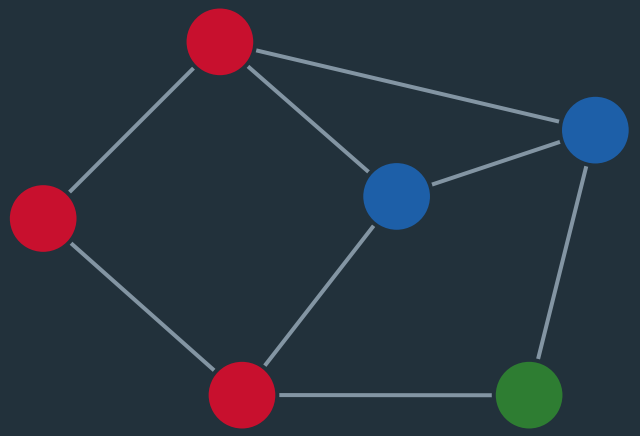
Objective function

Which change the actor makes: strongly negative outdegree, strongly positive reciprocity.

1

Building up a model on Glasgow

20 min · Chapter 25: transitive closure, homophily on smoking, two waves and both periods



Does friendship sort by smoking, net of the network's own tendencies?

Effect	A tie is more attractive when it …	Names the tendency
`outdegree`	is simply one more friend (a negative baseline keeps the network sparse)	density
`reciprocity`	returns a friendship the pupil received	reciprocation
`transtrip`	completes a triangle: the friend of my friend	closure
`nodematch(smoke1)`	joins two pupils with the same smoking status	homophily

The model, not the data, separates these. A tendency to cluster by smoking can be present descriptively (Chapter 11) and still be attributed to reciprocity and closure once those are in the model.

A first real fit: reciprocity and closure are strong, smoking homophily is not

```
. nwsaom, wave1(glasgow1) wave2(glasgow2) outdegree ///
>         reciprocity transtrip nodematch(smoke1) seed(12345)

Actors: 50                    Estimated rate:  6.047 (1.003)
           SAOM | Coefficient  Std. err.      z    P>|z|
----------------+----------------------------------------
       outdegree |  -2.710452   .2222379   -12.20   0.000
     reciprocity |   2.143375   .2537855     8.45   0.000
nodematch_smoke1 |   .2934813   .2079884     1.41   0.158
       transtrip |   .5790748   .1130786     5.12   0.000
Overall maximum convergence ratio: 0.115 (criterion: ...; met)
```

Chapter 25, building up a real model on Glasgow. The rate is 6.05 per pupil in this period.

Strong

Reciprocity (z = 8.45) and closure (z = 5.12): choices favor returning ties and completing triangles.

Not distinguishable

Smoking homophily: .29 with p = 0.158, in a group of 50 pupils.

Each coefficient is a log-odds in one actor's choice

Effect	Coefficient	Odds factor	p
outdegree	-2.710	0.07	0.000
reciprocity	2.143	**8.5**	0.000
transtrip	.579	**1.8**	0.000
nodematch_smoke1	.293	**1.3**	0.158

Odds factor = exp(coefficient). It compares two alternatives of an activated pupil that differ by one unit of the effect statistic.

It says

Other things equal, the alternative that returns a tie is about 8.5 times as likely to be chosen as one that does not.

It does not say

How often pupils get to choose. That is the rate. And a coefficient of 0.29 is not a verdict on smoking homophily: with 50 pupils and few smokers the data carry little information.

Both periods at once: shared coefficients, one rate per period

```
. nwsaom, waves(glasgow1 glasgow2 glasgow3) outdegree ///
>         reciprocity transtrip nodematch(smoke1) seed(12345)
Actors: 50                              Periods: 2
            SAOM | Coefficient  Std. err.      z    P>|z|
-----------------+------------------------------------------
       outdegree |  -2.761655   .1525161   -18.11   0.000
     reciprocity |   2.443945   .1847386    13.23   0.000
nodematch_smoke1 |   .1341893   .1517608     0.88   0.377
       transtrip |   .6221807   .0796816     7.81   0.000
Rate parameters (estimated, one per inter-wave period):
              |   period1    period2
         rate |    6.4343     5.2125
           se |    1.1104     0.8909
Overall maximum convergence ratio: 0.092 (criterion: ...; met)
```

Coefficients shared

One objective function for both periods: an assumption of the pooled model, not a finding.

Rates differ by period

6.43 then 5.21, with standard errors near 1: no evidence that opportunity for change slowed.

More data, a smaller and more precise homophily estimate

Coefficient (standard error)	Period 1 only	Both periods
reciprocity	2.143 (.254)	2.444 (.185)
transtrip	.579 (.113)	.622 (.080)
nodematch_smoke1	**.293 (.208)**	**.134 (.152)**
p of nodematch_smoke1	0.158	0.377
rate (standard error)	6.047 (1.003)	6.434 and 5.213

Pooling shrinks the standard error but moves the estimate toward zero: the more precise estimate is the smaller one.

It says

In this panel, smoking homophily is small and imprecise once reciprocity and closure are in the model.

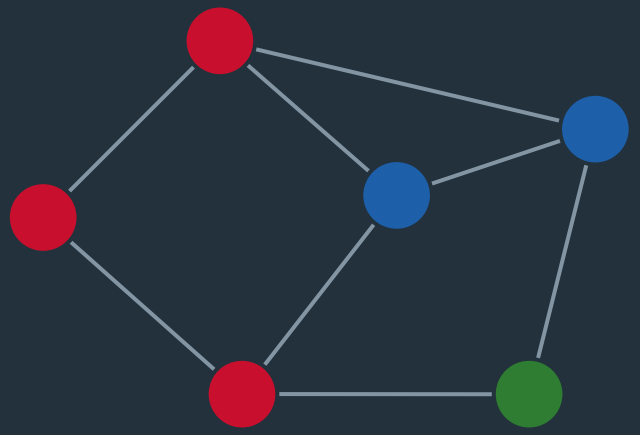
It does not say

That smoking plays no role in who befriends whom. The standard error is as large as the estimate, so the data cannot tell a small effect from none.

2

Convergence is not significance

20 min · Chapter 25: Method of Moments, t-ratios, the restart



Two different kinds of number sit in the same output

z and P>|z|

In the coefficient table, one pair per effect.

A Wald test: is this tendency different from zero?

Read after the estimation has settled.

Convergence t-ratios

Below the table, one per estimated parameter.

Do simulated networks reproduce the observed target statistics?

Read first, before any coefficient.

```
nodematch_smoke1 |   .2934813   .2079884     1.41   0.158      <- z: is the effect zero?
Overall maximum convergence ratio: 0.115 (criterion: below 0.25; met)   <- has it settled?
Convergence t-ratios (criterion: all below 0.1 in absolute value): met
```

The Method of Moments matches simulated and observed statistics

1 Simulate

At trial values of the rate and the coefficients, run the ministep process forward from wave 1.

2 Compare

Compute the target statistics of the simulated wave 2 and compare them with the observed ones.

3 Adjust

Move the trial values to shrink the gap (Robbins-Monro). Repeat until the average deviation is near zero.

No likelihood is computed. Wave 2 given wave 1 would need a sum over every sequence of ministeps. The moment condition avoids it: one condition per effect.

Two criteria, both must be met

< 0.25

overall maximum convergence ratio

< 0.1

every t-ratio, in absolute value

```
. matrix list e(tconv)
          outdegree   reciprocity  nodematch_~1     transtrip
tconv     .06210207     .03757292     .05874033     .05812226
```

Why two criteria

The overall ratio is the largest t-ratio over all linear combinations of the deviations. It is at least as large as any single t-ratio, so its threshold is looser.

Stored

The individual ratios are in e(tconv).

When a criterion fails, or to check stability: restart from the estimates

```
. nwsaom, waves(glasgow1 glasgow2 glasgow3) outdegree ///
>         reciprocity transtrip seed(12345)
Overall maximum convergence ratio: 0.075 (criterion: ...; met)
. nwsaom refit, seed(12346)
Overall maximum convergence ratio: 0.173 (criterion: ...; met)
nwsaom refit: 1 run; largest |t| 0.076, overall ratio 0.173
(converged)
```

Estimate	First run	After refit
outdegree	-2.683	-2.682
reciprocity	2.472	2.477
transtrip	.621	.617
overall ratio	0.075	0.173

Options of refit

seed() for a new random stream; k3() for a longer final phase; maxruns() to repeat until both criteria are met.

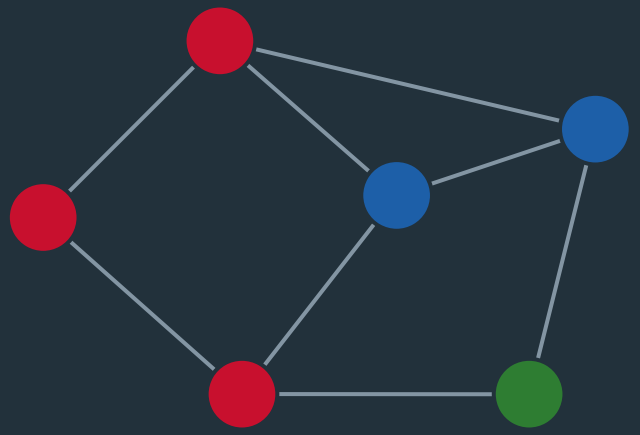
Stable estimates

Estimates that hold across restarts are the stronger reassurance.

3

Does the model fit?

20 min · Chapter 25: estat gof, Mahalanobis distances, where a model fails



Simulate from the fitted model and see whether the data look typical

Simulate

From the observed first wave of each period, run the fitted model many times (nsim).

Summarize

Outdegrees, indegrees and geodesic distances: statistics the model does not contain.

Compare

Mahalanobis distance of the observed distribution from the simulated ones; p is the share of simulated distributions as far away.

Small p, below 0.05: the observed network is unusual under the model on that statistic. A large p is only the absence of evidence against the model.

Degrees are reproduced, the geodesic distances are not

```
. estat gof, nsim(100) seed(1)
Goodness of fit: Mahalanobis distance test
100 simulations at the estimates, 2 periods pooled
           Statistic |   Mahalanobis  P-value      Fit
---------------------+------------------------------
           Outdegree |         9.922    0.380       ok
            Indegree |         8.373    0.480       ok
            Geodesic |        28.119    0.000 rejected
Violin plots: gof_outdegree gof_indegree gof_geodesic
```

After outdegree, reciprocity and transtrip on three waves, after the restart.

Degrees: ok

p = 0.38 and 0.48. Not surprising for such a flat degree distribution.

Geodesics: rejected

p = 0.000: none of the 100 simulated networks was as far from the center as the observed one.

Adding reciprocity and closure brings the distance down from 220 to 28

Statistic	outdegree only	+ reciprocity, transtrip
Outdegree: distance (p)	5.408 (0.560)	9.922 (0.380)
Indegree: distance (p)	6.728 (0.560)	8.373 (0.480)
Geodesic: distance (p)	**219.780 (0.000)**	**28.119 (0.000)**

220 to 28

geodesic distance, outdegree only against three effects

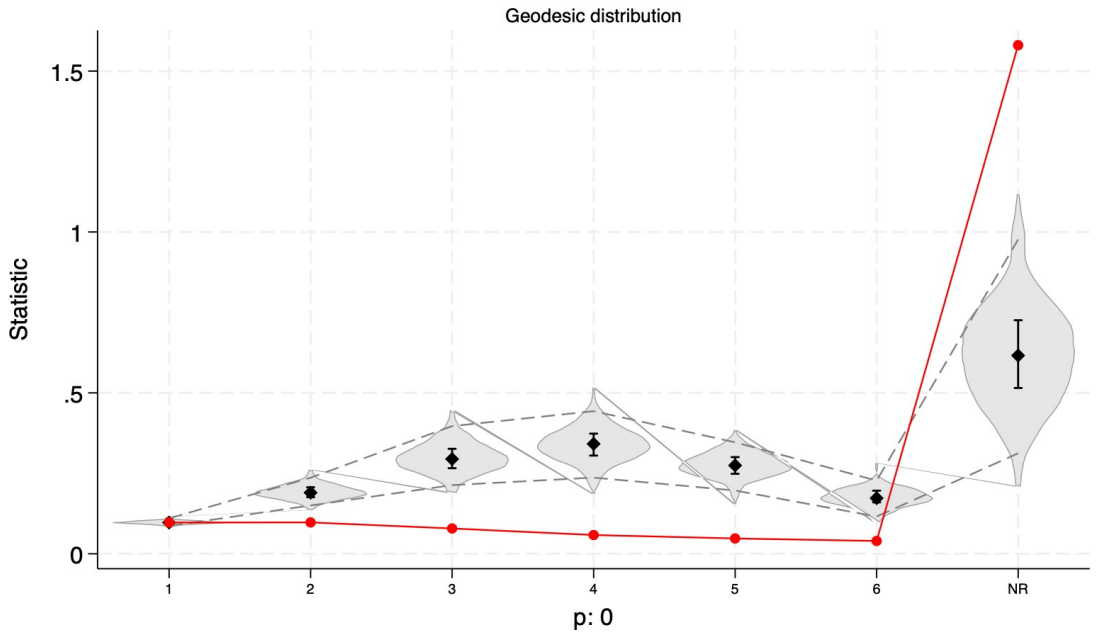
It says

Reciprocity and closure explain much of the shape of the network, and the degrees are reproduced by both models.

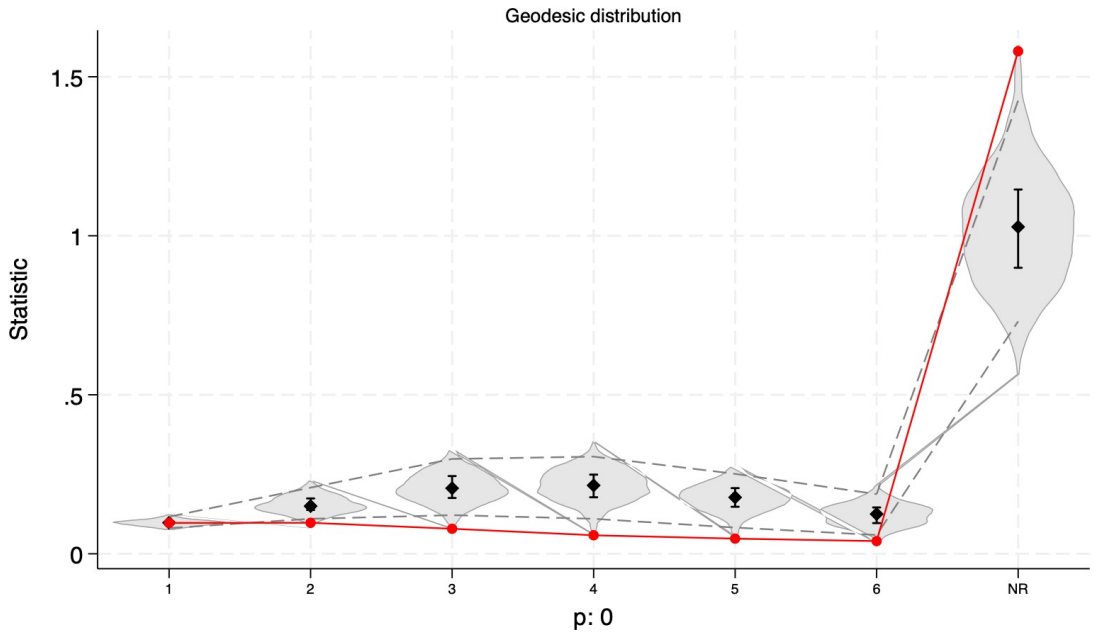
It does not say

That the model fits. The geodesic p-value is 0.000 for both: the model is useful and incomplete. Compare distances within a statistic across models, not across statistics.

The violin plot shows where: far more unreachable pairs than the model produces



Outdegree only



Outdegree, reciprocity, transtrip

Red: the observed network. Violins: 100 simulated ones. NR is the pairs of pupils with no path. The observed network has far more of them than either model produces.

estat gof, nsim(100) seed(1), graph display gof_geodesic; nwsaom runs of 17 Nov 2026.

Useful and incomplete: report the limit, and look where it fails

It says

Where the model reproduces the data (degrees) and where it does not (distances): a map of what its coefficients can be trusted to describe.

It does not say

What to add, or that the effects are wrong. A passing statistic is the absence of evidence against the model, not evidence for it.

- **Rerun near 0.05**

  The reference distribution is simulated: try another seed().

- **Each period**

  join(off) tests the periods separately.

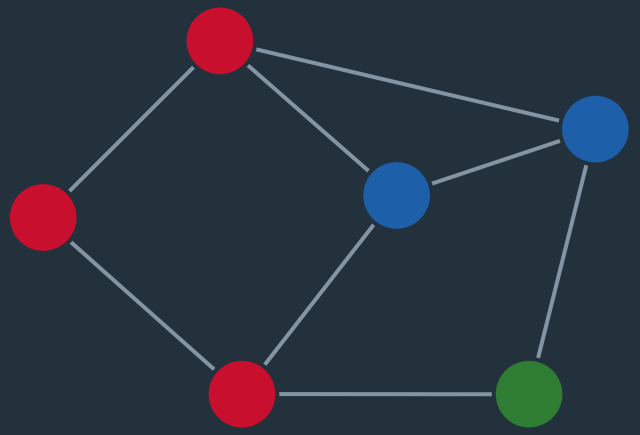
- **Report the limit**

  If the distances still fail, say so.

4

Selection or influence

30 min · Chapter 25: co-evolution of friendship and drinking, behsim and avalt



Friends are alike: two stories leave the same trace in a cross-section

Selection

Pupils choose friends who drink like them. The network changes toward similarity.

Influence

Pupils come to drink like their friends. The behavior changes toward the friends' level.

A common cause

A new teacher, a club, a neighborhood: the same pattern with no choice and no influence.

A cross-section shows the pattern, not its origin. A panel shows which came first, and a model with both processes can give each its own coefficient. With one wave only, an autologistic actor attribute model (Chapter 24, Section 24.4; the optional block of Week 7) describes how smoking and ties line up, and cannot separate the two.

Co-evolution: two processes, each with its own rate and objective function

	Network part	Behavior part
Ministep	an actor changes one tie	an actor moves its drinking up or down by one step
Rate	`network_rate`	`behavior_rate`
Objective function	`outdegree reciprocity transtrip behsim`	`linear quadratic avalt`
Question it answers	**Selection: do similar pupils choose each other? (behsim)**	**Influence: do pupils move toward their friends? (avalt)**

Both effects are evaluated on the current, simulated state. Each is estimated net of the other process, with the structural tendencies of the network held in the model.

Look at the behavior first: alcohol moves, smoking hardly does

```
. nwsaom describe, waves(glasgow1 glasgow2 glasgow3) ///
>         behavior(alcohol1 alcohol2 alcohol3)
    Behavior |        Mean   Std. dev.     Missing
      Wave 1 |       2.880       1.118           0
      Wave 2 |       3.100       1.199           0
      Wave 3 |       3.360       0.985           0
      Change |          Up        Down        Same
        1->2 |          16           6          28
        2->3 |          18           9          23
```

Alcohol use has five levels, from not drinking to drinking more than once a week.

Alcohol: enough change

About half of the pupils move in each period: 22 and 27 of 50.

Smoking: too little

Three levels, rarely changing: the behavior estimates would be very imprecise.

The co-evolution fit: selection is clear, influence is weaker

```
. nwsaom, waves(glasgow1 glasgow2 glasgow3) outdegree ///
>    reciprocity transtrip behsim linear quadratic avalt ///
>    behavior(alcohol1 alcohol2 alcohol3) seed(12345)
         SAOM | Coefficient  Std. err.      z    P>|z|
--------------+----------------------------------------
    outdegree |   -2.76126   .1460683   -18.90   0.000
  reciprocity |    2.35587   .1982928    11.88   0.000
    transtrip |    .621357   .0803867     7.73   0.000
       behsim |   1.460717   .6218014     2.35   0.019
   beh_linear |   .3948453   .2135712     1.85   0.064
beh_quadratic |  -.6002813    .320045    -1.88   0.061
    beh_avalt |   1.305018   .7780128     1.68   0.093
Overall maximum convergence ratio: 0.138 (criterion: ...; met)
```

p = 0.019

selection (behsim): similar pupils choose each other

p = 0.093

influence (beh_avalt): weaker evidence

Drinking rates: 1.31 and 1.80, against 6.49 and 5.17 for friendship.

What the two coefficients say, and what they do not

It says

Pupils prefer friends whose drinking resembles their own (selection, p = 0.019), and pupils whose friends drink more tend to drink more (influence, p = 0.093): clear for selection, weaker for influence.

It does not say

That selection is stronger because 1.46 exceeds 1.31: the two coefficients are on different scales. The difference in evidence is not itself significant. And an unmodeled common cause, such as shared evenings out, could produce both patterns.

Write it carefully

Report both coefficients with their standard errors. Avoid causal wording such as the effect of friends on drinking.

The fit

Both convergence criteria are met (overall ratio 0.138). estat gof adds a behavior test: distance 0.848, p = 0.950.

What the separation delivers, and what it cannot deliver

Delivers

A test of each mechanism net of the other, given the model and the data.

Does not deliver

A causal effect. Homophily and contagion are in general confounded in observational data (Shalizi and Thomas, 2011).

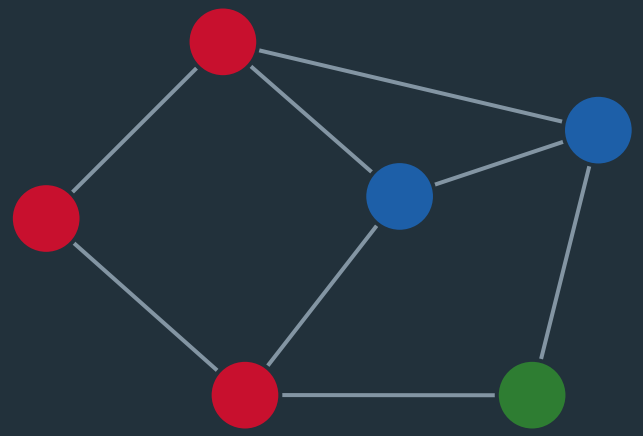
Fixed in advance

The form of influence: avalt pulls toward the average of the friends, not toward the closest one or one heavy drinker.

Remedies: design (matching on dynamic data), a sensitivity analysis of how strong an unobserved cause would have to be, and cautious wording.

Break

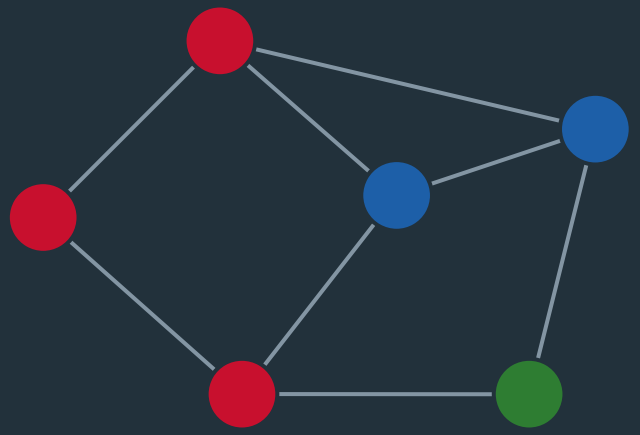
10 minutes



5

Lab: fit, restart, check

45 min · Run along the chapter's models; Ch. 25, Ex. 2 and Ex. 3



Run along: the three-effect model, its restart and the fit

```
. nwwebuse glasgow, nwclear
. nwsaom, wave1(glasgow1) wave2(glasgow2) outdegree ///
>         reciprocity transtrip nodematch(smoke1) seed(12345)
. nwsaom, waves(glasgow1 glasgow2 glasgow3) outdegree ///
>         reciprocity transtrip seed(12345)
. nwsaom refit, seed(12346)
. estat gof, nsim(100) seed(1)
```

If a number differs

Is seed() typed as printed? Does refit directly follow the fit, and estat gof directly follow the refit? Both use the last fit, so any other nwsaom command in between changes what they work on.

Checkpoints

Two waves: reciprocity 2.143, transtrip .579, nodematch_smoke1 .293 (p 0.158)

Three waves: reciprocity 2.472, transtrip .621

Refit: largest |t| 0.076

gof: 9.922, 8.373, 28.119 (p 0.000)

Hands-on (25 min)

Ch. 25, Ex. 2

Fit the first model (outdegree and reciprocity) on the second period: wave1(glasgow2) wave2(glasgow3), seed(1). Restart with nwsaom refit, seed(2). Report the rate and the two coefficients of both runs and whether each run meets both convergence criteria. Does the restart move any coefficient by more than a tenth of its standard error, and is reciprocity in the second period higher or lower than the 2.51 of the first?

Hint nwsaom refit starts from the estimates of the last fit, so it must follow the period-2 fit directly. Divide each change between the two runs by the coefficient's standard error. Whether the difference between the periods is more than chance is the question of estat timetest in Chapter 26.

Ch. 25, Ex. 3

Refit the two-wave model (outdegree reciprocity transtrip nodematch(smoke1), seed(12345)) with nodematch(sport1) in place of smoking. Report coefficient and p-value of nodematch_sport1 next to smoking (0.29, p = 0.16): is either kind of homophily distinguishable from zero? Then run estat gof, nsim(100) seed(1) on the sport model, and again with seed(2). Does the geodesic distribution fail as clearly as before, and what do the two runs together tell you?

Hint Only the variable in nodematch() changes. tabulate sport1 shows the two groups (37 of the 50 do sport regularly); check the two convergence lines before reading the table. The two estat gof runs differ only in the seed: ask whether the difference is large enough to change your reading.

Debrief Ex. 2: reciprocity is stronger in the second period, and a restart moves it a little

Period 2 (wave 2 to 3)	First run, seed(1)	After refit, seed(2)
rate (standard error)	4.667 (0.709)	4.656 (0.713)
outdegree	-2.578 (.183)	-2.548 (.190)
reciprocity	**3.352 (.299)**	**3.298 (.324)**
overall ratio	0.139, met	0.099, met
largest |t|	below 0.1	0.096

Period 1 for comparison: rate 5.552, outdegree -2.264, reciprocity 2.509.

It says

The restart moves outdegree by 0.16 and reciprocity by 0.18 of a standard error: more than a tenth, far less than one. Reciprocity (3.35) is higher than the 2.51 of period 1.

It does not say

That reciprocity changed between the periods. Two fits with standard errors near 0.3 do not show it. The test is estat timetest (Chapter 26).

Debrief Ex. 3: sport homophily is no clearer than smoking, and the gof verdict flips with the seed

Two-wave model	smoke1	sport1
nodematch, coefficient (se)	.293 (.208)	.108 (.186)
nodematch, p	0.158	0.561
overall ratio	0.115	0.117

Geodesic, nsim(100)	Distance	p	Verdict
seed(1)	13.266	0.080	ok
seed(2)	14.290	**0.030**	**rejected**

13 of 50 pupils do not do sport regularly, 37 do (tabulate sport1).

It says

Neither homophily is distinguishable from zero. The sport model's geodesic distance (13 to 14) is far below the 28 of the three-wave model, with a p-value on the line.

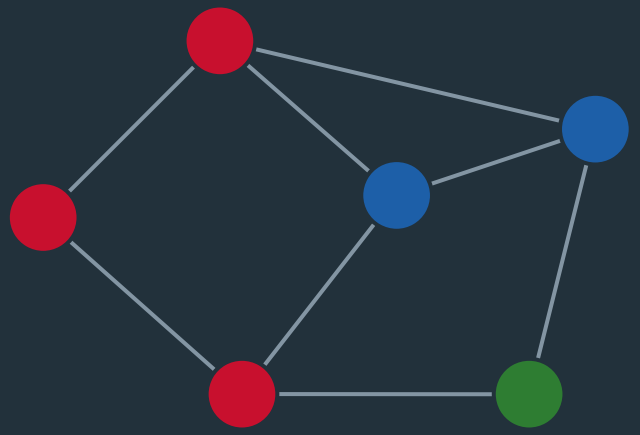
It does not say

A pass at seed(1) and a rejection at seed(2) are not two verdicts: the p-value sits near 0.05. Rerun with another seed and report the result as borderline.

6

Lab: co-evolution

45 min · Run along the drinking model; Ch. 25, Ex. 5 and Ex. 8



Run along: the co-evolution model and its goodness of fit

```
. nwsaom, waves(glasgow1 glasgow2 glasgow3) outdegree ///
>    reciprocity transtrip behsim linear quadratic avalt ///
>    behavior(alcohol1 alcohol2 alcohol3) seed(12345)
. estat gof, nsim(100) seed(1)
           Statistic |   Mahalanobis  P-value      Fit
           Outdegree |         8.775    0.440       ok
            Indegree |         9.950    0.370       ok
            Geodesic |        30.554    0.000 rejected
            Behavior |         0.848    0.950       ok
```

The new line

After a co-evolution fit, estat gof adds the distribution of the behavior: the drinking distribution is reproduced, the geodesic distribution fails as before.

Checkpoints

behsim 1.461 (p 0.019)

beh_avalt 1.305 (p 0.093)

overall ratio 0.138

behavior gof: 0.848, p 0.950

Hands-on (25 min)

Ch. 25, Ex. 5

Refit the co-evolution model with avsim (average similarity) in place of avalt, keeping everything else, including seed(12345). If the convergence line says that a criterion is not met, restart with nwsaom refit, seed(12346). Report the coefficients and p-values of behsim and beh_avsim and compare them with the avalt model (selection 1.46, p = 0.019; influence 1.31, p = 0.093). Is the conclusion about selection and influence the same?

Hint Compare signs and p-values, not sizes: avsim and avalt measure influence on different scales, so their coefficients cannot be compared directly.

Ch. 25, Ex. 8

Suppose a new teacher seats pupils who drink alike next to each other in the second year, and they then become friends. Which of the two coefficients of the co-evolution model, behsim or avalt, would the seating plan inflate, and why? Name one variable you would add to the model, or one feature of the study design, that would separate the teacher's seating from the pupils' own choice.

Hint The seating changes which ties form, not how much anyone drinks: ask which effect belongs to the network part of the model and which to the behavior part. For the second part, ask what record the seating plan leaves, and whether it describes single pupils or pairs of pupils.

Debrief Ex. 5: the same conclusion, and a first run that needs its restart

Coefficient (p)	avalt (the chapter)	avsim, first run	avsim, after refit
behsim: selection	1.461 (0.019)	1.410 (0.027)	**1.402 (0.018)**
influence	1.305 (0.093)	3.993 (0.051)	**3.954 (0.070)**
overall ratio	0.138, met	0.179	0.112, met
t-ratios	met	not met: rate_period1 -0.115	met (largest 0.075)

It says
Selection clear (p near 0.02), influence positive and weaker (p between 0.05 and 0.1): the conclusion is the same under both specifications.

It does not say
That avsim finds stronger influence because 3.95 exceeds 1.31: avsim and avalt are on different scales. And the first run of avsim is not usable until it has been restarted.

Debrief Ex. 8: the seating plan inflates behsim, and a dyadic covariate separates it

Part of the answer	What a good answer contains
Which coefficient	behsim: the seating changes which ties form, not how much anyone drinks. It belongs to the network part of the model.
Why	Pupils who drink alike become friends for a reason outside their own choice, so similar drinkers look like selection.
Separating variable	The seating plan as a pair-level record: who sits next to whom, entered as a dyadic covariate in the network part.
Or a design feature	Seating assigned at random, or classes with and without a seating plan.

A sketch, not the book's answer

The book gives hints for the open exercises, not solutions. Grade the reasoning: which part of the model the seating feeds, and why the fix is pair-level.

What to take away: check the estimation, check the fit, and word the separation carefully

Check the estimation

Convergence t-ratios, not z, say whether the algorithm settled. Restart with nwsaom refit; stable estimates are the reassurance.

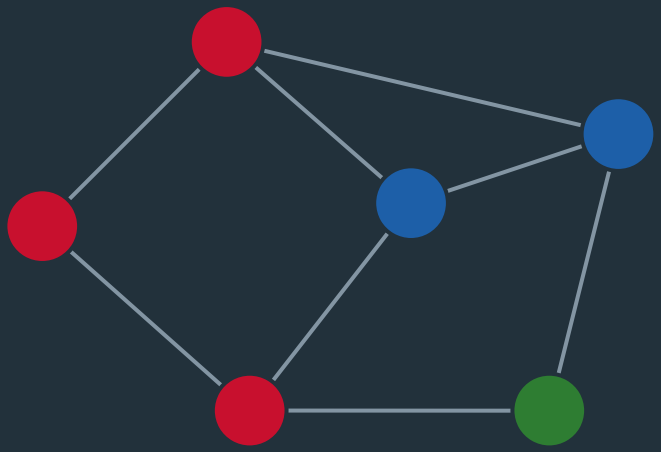
Check the fit

estat gof shows where a model fails. A p-value near 0.05 is not a verdict: change the seed.

Word the separation

Selection and influence in one model: a test of each, net of the other, given the assumptions. Not a causal effect.

To read after the session



Read Chapter 26, Advanced stochastic actor-oriented models

Left over from Chapter 25: Ex. 4, 6 and 9 (and Ex. 1 and 7 if not done in Week 8)

Fit the three-effect model with a restart: Week 10 builds on it

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

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