

Multilevel analysis of network dynamics

Tom A.B. Snijders



University of Groningen
University of Oxford



March, 2023

Traditionally, network analysis tended to consist of cases studies of single networks.

However, it is preferable to generalize to a **population** of networks.

Traditionally, network analysis tended to consist of cases studies of single networks.

However, it is preferable to generalize to a **population** of networks.

This can be achieved, in principle, by multilevel network analysis in the sense of analyzing multiple similar networks, mutually independent.

This was proposed by Snijders & Baerveldt (*J. Math. Soc.* 2003).

Also see Entwisle, Faust, Rindfuss, & Kaneda (*AJS*, 2007) who gave an overview of empirical work until then involving multiple networks.

Sample from Population of Networks

Suppose we have a sample indexed by $j = 1, \dots, N$ from a population of networks on disjoint node sets, where the networks are 'replications' of each other in the following sense:

Sample from Population of Networks

Suppose we have a sample indexed by $j = 1, \dots, N$ from a population of networks on disjoint node sets, where the networks are 'replications' of each other in the following sense:

they all are regarded as realizations of processes obeying the same model, but having different parameters $\theta_1, \dots, \theta_j, \dots, \theta_N$.

Each disjoint network is called a **group**.

We assume we have network panel data for each group, and wish to analyze these by fitting Stochastic Actor-Oriented Models ('SAOMs', **RSiena**).

Several approaches are possible for combining such data:

Several approaches are possible for combining such data:

- 1 Multi-group analysis:
assume all parameters are identical.

Several approaches are possible for combining such data:

- 1 Multi-group analysis:
assume all parameters are identical.
- 2 Integrated hierarchical approach:
Assumption: population of networks, normal distribution.

Several approaches are possible for combining such data:

- 1 Multi-group analysis:
assume all parameters are identical.
- 2 Integrated hierarchical approach:
Assumption: population of networks, normal distribution.
- 3 Two-step approach like in meta analysis:
Assumption: population of networks,
no distributional assumptions:
two-stage 'meta analysis'.

Several approaches are possible for combining such data:

- 1 Multi-group analysis:
assume all parameters are identical.
- 2 Integrated hierarchical approach:
Assumption: population of networks, normal distribution.
- 3 Two-step approach like in meta analysis:
Assumption: population of networks,
no distributional assumptions:
two-stage 'meta analysis'.
- 4 Two-step approach without population assumption:
Fisher combination of independent tests.

Such data sets have multilevel structure;
all caveats and considerations from usual multilevel analysis apply!

See T.A.B. Snijders & R.J. Bosker,
'Multilevel Analysis' (2nd ed., Sage, 2012);

T.A.B. Snijders, 'The Multiple Flavours of Multilevel Issues for Networks',
in E. Lazega & T.A.B. Snijders, 'Multilevel Network Analysis' (Springer, 2016).

Multilevel designs, generalization, and the risks of having few groups

In 'standard' multilevel analysis, we have several populations.


Consider a data set with individuals nested in groups: then we wish to generalize to the population of individuals as well as the population of groups. Often, we only have small number of groups, which may jeopardize generalization.

For some basic intuitive understanding of the relevant issues, consider the simplest case: estimation of the population mean μ of some variable.

We have $N = 5$ randomly drawn groups each of size $n = 20$. Data range from -2.5 to $+3.5$, as in the histogram (next page).

Is there evidence that the population mean is positive?


Test $H_0 : \mu = 0$ against $H_1 : \mu > 0$ (one-sided).



mean = 0.36, s.d. = 1.18, $t_{99}=3.40$, $p = 0.0003$

linear model (OLS): yes

Frequency histogram of five groups




multilevel $t_4 = 1.49$, $p = 0.12$

multilevel model: no

but by group $p = 0.11, 0.84, 0.034, 0.021, 0.00002$

Frequency histogram of five groups




multilevel $t_4 = 1.49, p = 0.12$

multilevel model: no

but by group $p = 0.11, 0.84, 0.034, 0.021, 0.00002$

Frequency histogram of five groups



An even more extreme example:

OLS $t_{99} = 7.9$, $p \approx 0$; multilevel $t_4 = 1.50$, $p = 0.12$

by group $p = 0.64, 0.002, 0.000012, < 0.0001, < 0.0001$

The conclusion for both data sets is not only, that the multilevel model tells us μ is not significantly positive.

There is also significant heterogeneity between the groups: μ is not zero in all groups (this can be concluded already from an OLS analysis by group).

And one of the groups has such a small p , that the significance survives multiple testing with Bonferroni.

If we wish to generalize to the population of all groups, there is no evidence of a positive mean; but there is evidence that at least one of these groups has a positive mean, and that the mean of these five groups is positive.

After this sensitization, on with the four approaches:

- 1 Multi-group analysis:
assume all parameters are identical.
- 2 Integrated hierarchical approach:
Assumption: population of networks, normal distribution.
- 3 Two-stage approach:
Assumption: population of networks,
no distributional assumptions:
first analysis by group, then combination 'meta-analysis'.
- 4 Two-stage approach without population assumption:
*first analysis by group,
then Fisher combination of independent tests.*


We treat how this can be handled with SAOMs,
using the package **RSiena**.

Multi-group analysis

Multi-group analysis (`sienaGroupCreate`, `siena07`)

assumes all parameters are the same.

The groups are treated just like subsequent waves, strung in a sequence, where (of course) transitions between groups are not analyzed.



Example for 3 groups, with 3 / 2 / 3 waves;

the analyzed 2+1+2 periods are indicated as bold lines.

Groups can have different sizes,

must have same variables and a common model specification.

(Earlier approaches combined all networks in one big network with structural zeros between them; this is now superseded by multi-group analysis.)

The multi-group analysis assumes all parameters are identical.

The multi-group analysis assumes all parameters are identical.

possible but risky

Gives a global impression, risk of underestimated standard-errors.

Try to find relatively homogeneous sets of groups for doing this.

The multi-group analysis assumes all parameters are identical.

possible but risky

Gives a global impression, risk of underestimated standard-errors.

Try to find relatively homogeneous sets of groups for doing this.

`sienaTimeTest()` can be used to test homogeneity assumption.

Random effects: hierarchical multilevel structure

- 1 On the **tie level** there is a dynamic process governed by the SAOM.
- 2 On the **network level** there is a Stochastic Actor-Oriented Model (SAOM) with parameter vector $(\theta_j^{(1)}, \eta)'$ for group j .
Here $\theta_j^{(1)} \sim \mathcal{N}(\mu, \Sigma)$.
- 3 On the **global level** there is a population of networks with either a multivariate normal distribution $\mathcal{N}(\mu, \Sigma)$ for the randomly varying parameters $\theta_j^{(1)}$, and a common parameter η for the rest; or without distribution assumption.

Hierarchical multilevel analysis

Assumption: population of networks;

$\theta_j \sim$ multivariate normal distribution,
perhaps conditionally on network-level covariates.

integrated hierarchical approach:

some parameters varying, others constant across groups.

(cf. random slopes and fixed effects

in Hierarchical Linear Model of standard multilevel analysis.)

multiSiena: *multi-level option*, `sienaBayes()`

Integrated procedure:

Estimate the distribution of θ_j and consider the 'posterior' distribution of θ_j given the data.

Integrated procedure:

Estimate the distribution of θ_j and consider the 'posterior' distribution of θ_j given the data.

Advantage:

The analysis of the separate networks draws strength from the total sample of networks by regression to the mean.

Useful especially for many rather small networks.

The method implemented in `sienaBayes` is in the paradigm of *Bayesian* as opposed to *frequentist* statistics.

Further explanation is in the set of slides

'Multilevel analysis of network dynamics using sienaBayes' .

(Frequentist methods would also be possible, but have not been elaborated.)

Two-stage approach ~ Random Effects Model

$\theta_1, \dots, \theta_j, \dots, \theta_N$ are drawn randomly from a *population* $\mathcal{P}[\text{net}]$ of networks, no further distributional assumptions are made.

Parameter θ_j can be written as

$$\theta_j = (\theta_{1j}, \theta_{2j}, \dots, \theta_{Kj}).$$

Two-stage approach ~ Random Effects Model

$\theta_1, \dots, \theta_j, \dots, \theta_N$ are drawn randomly from a *population* $\mathcal{P}[\text{net}]$ of networks, no further distributional assumptions are made.

Parameter θ_j can be written as

$$\theta_j = (\theta_{1j}, \theta_{2j}, \dots, \theta_{Kj}).$$

Two-stage procedure:

estimate each θ_j separately,

combine the results in a meta-analysis (Cochran 1954),

(‘*V-known problem*’ in multilevel analysis)

which allows testing hypotheses about $\mathcal{P}[\text{net}]$

such as, for a coordinate k ,

$$\begin{aligned} H_0^{\text{total}} &: \text{all } \theta_{kj} = 0; \\ H_0^{\text{mean}} &: E\{\theta_{kj}\} = 0; \\ H_0^{\text{spread}} &: \text{var}\{\theta_{kj}\} = 0. \end{aligned}$$

The meta-analysis is directed at one particular coordinate k (e.g., average alter).

The input for the meta-analysis consists of estimates $\hat{\theta}_{kj}$ and their standard errors s.e. $_{kj}$.

The meta analysis is constructed based on the model

$$\hat{\theta}_{kj} = \theta_{kj} + E_{kj} = \mu_k + U_{kj} + E_{kj} ,$$

where μ_k is the population mean for parameter coordinate k , θ_{kj} is the true coordinate k of the parameter in group j , U_{kj} is the true effect of group j as deviation from μ_k , and E_{kj} is the statistical error of estimation.

The meta-analysis is directed at one particular coordinate k (e.g., average alter).

The input for the meta-analysis consists of estimates $\hat{\theta}_{kj}$ and their standard errors s.e._{kj} .

The meta analysis is constructed based on the model

$$\hat{\theta}_{kj} = \theta_{kj} + E_{kj} = \mu_k + U_{kj} + E_{kj} ,$$

where μ_k is the population mean for parameter coordinate k , θ_{kj} is the true coordinate k of the parameter in group j , U_{kj} is the true effect of group j as deviation from μ_k , and E_{kj} is the statistical error of estimation.

U_{kj} and E_{kj} are independent residuals with mean 0, the U_{kj} are *i.i.d.* with unknown variance σ^2 , and $\text{var}(E_{kj}) = \text{s.e.}_{kj}^2$ ('V-known').

Implemented in MLwiN, HLM, R package *Metafor*, **RSiena** function `siena08()`; Snijders & Baerveldt (2003).

Meta-Analysis ~ Random Effects Model (contd.)

This has been applied in quite many studies. Examples:

Lubbers (2003):

homophily in 57 classrooms with 1466 students
(also with random coefficient p^* approach);

Baerveldt, van Duijn, Vermeij, van Hemert (2004):
ethnic homophily in 20 schools, 1317 students;

Valente, Fujimoto, Chou, and Spruit-Metz (2009):
friendship & obesity, 17 classrooms with 617 students;

Mercken, Snijders, Steglich, & de Vries (2009):
fr. & smoking, 7704 adolescents, 70 schools, 6 countries;
also other studies by Liesbeth Mercken.

Kruse & Kroneberg (2022):

liking & school grades, 40 schools, 4293 students.

There are a variety of R packages for meta-analysis.
For example: `metafor`, `mvmeta`.

`RSiena` also includes the function `siena08` for meta-analysis.
This carries out meta-analysis according to S & B (2003),
but with an iterated weighted least squares ('WLS') method. `metafor` calls
this *WLS analysis* of the *random effects* model.

`metafor` also has more options than `siena08`.

`siena08` can be used for what is not provided by `metafor`:
overall test of H_0^{total} ; Fisher combination (see below).

Wolfgang Viechtbauer, 'Conducting meta-analyses in R with the `metafor` package', *Journal of Statistical Software*, 2010.

Some terminology:

The method carried out by `siena08` is, in the language of metafor, the *WLS analysis* of the *random effects* model.

Fixed effects: conclusion only about the average of these N groups (usually weighted inversely by variances);
random effects: generalization to population of groups.

Viechtbauer (2005): REML (one of the methods implemented in metafor) is somewhat better than WLS (implemented in `siena08`).

Example: Chris Baerveldt's data

Houtzager and Baerveldt (1999): *Dutch Social Behavior study*.

19 classrooms in secondary schools,
evolution of friendship network and delinquent behavior.

Also used in Snijders & Baerveldt (*J. Math. Soc.*, 2003) and
Baerveldt, Völker & Van Rossem (*Can. J. Crim. Crim. Just.*, 2008).

3 classrooms had difficulties for convergence and were dropped;
in 2 further classrooms,
one parameter had to be fixed to 0 for stable convergence.

Number of pupils per group varied 33–91.

Average degrees per wave 1.0–3.2 (low!)

Jaccards 0.24–0.38.

The following pages show a table that was produced as a hybrid from metafor and siena08 results.

It has results for the global parameters, μ as well as σ .

Note: it is important for a meta-analysis to present results for mean tendency as well as between-group variability!!!

(Many researchers unfortunately present only the former.)

All results in the table are from metafor, except p_T , the p -value for H_0^{total} from siena08.

	N	$\hat{\mu}$	$se(\hat{\mu})$	ρ_{μ}	$\hat{\sigma}$	ρ_{σ}	ρ_T
<i>Friendship dynamics</i>							
outdegree (density)	16	-2.955	0.149	0.00	0.00	0.65	0.00
reciprocity	16	3.588	0.253	0.00	0.44	0.52	0.00
GWESP	16	2.377	0.120	0.00	0.18	0.27	0.00
indegree - popularity	16	-0.118	0.024	0.00	0.02	0.42	0.00
outdegree - activity	16	0.121	0.024	0.00	0.00	0.91	0.00
rec.degree - activity	16	-0.204	0.067	0.00	0.08	0.63	0.10
indegree - activity	16	-0.204	0.072	0.00	0.05	0.70	0.17
same ethnicity	16	-0.005	0.089	0.96	0.19	0.21	0.20
sex (M) alter	16	-0.329	0.086	0.00	0.00	0.64	0.03
sex (M) ego	16	-0.169	0.173	0.33	0.42	0.08	0.05
same sex	16	0.554	0.080	0.00	0.14	0.39	0.00
crimi alter	16	-0.053	0.046	0.24	0.00	0.83	0.74
crimi ego	16	0.074	0.080	0.35	0.16	0.22	0.15
crimi ego × alter	16	0.006	0.057	0.92	0.11	0.22	0.21
reciprocity × GWESP	14	-0.568	0.272	0.04	0.00	0.96	0.68

	N	$\hat{\mu}$	$se(\hat{\mu})$	p_{μ}	$\hat{\sigma}$	p_{σ}	p_T
<i>Delinquency dynamics</i>							
linear shape	16	0.015	0.108	0.89	0.00	0.93	0.93
quadratic shape	16	-0.295	0.048	0.00	0.05	0.70	0.00
outdegree	16	0.042	0.047	0.37	0.00	0.99	0.99
total alter	16	0.076	0.064	0.23	0.00	1.00	1.00
sex (M)	16	0.610	0.139	0.00	0.00	0.99	0.09

- N : number of groups combined;
- $\hat{\mu}$, $se(\hat{\mu})$, p_{μ} : inference for μ ;
- $\hat{\sigma}$, p_{σ} : inference for σ ;
- p_T : p -value for H_0^{total} .

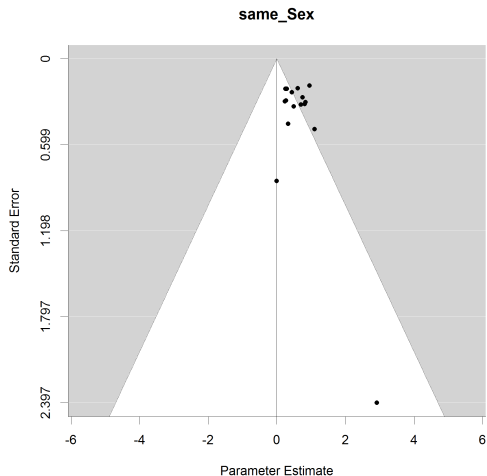
It may be noted that σ^2 never is significantly positive.

However, `sienaTimeTest`, which uses a more powerful test (multivariate, and using also information used for estimation) concludes that the parameters are significantly different.

Some example conclusions:

- reciprocity: strongly significant, but also quite variable across groups;
- reciprocal degree-activity: clearly significant, variability across groups might exist; if so, no or hardly any groups with positive parameter;
- same sex: strongly significant, variable across groups, will be positive in all groups;
- total alter: not significant; if non-zero, then more likely to be positive; note that $p_\sigma = p_\tau = 1$ but $p_\mu = 0.23$, which can be explained by the fact that tests do not produce certainty, and that the test for μ is more focused.

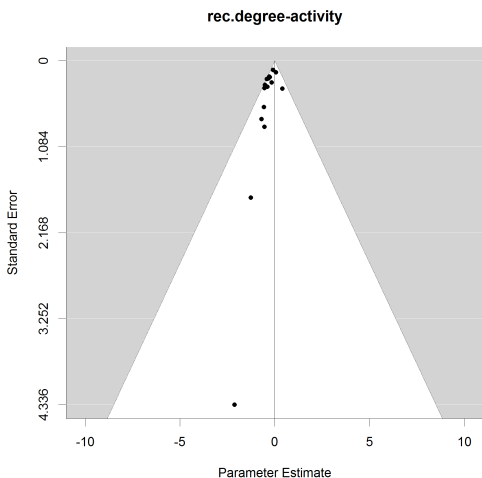
For 'positive in all groups', think of $\hat{\mu} \pm 2\hat{\sigma}$.



The *funnel plot* produced by metafor shows groupwise estimates and standard errors.


Important to get a feeling for what's going on!

RSiena also has a function `<funnelPlot>`.



Note that for the reciprocated degree-activity effect, none of the individual estimates is significant, but jointly the effect is strongly significant.

Crimi_total_alter



Meta-Analysis without a population assumption

$\theta_1, \dots, \theta_j, \dots, \theta_N$ are arbitrary values,
no assumption about a population is made.

Meta-Analysis without a population assumption

$\theta_1, \dots, \theta_j, \dots, \theta_N$ are arbitrary values,
no assumption about a population is made.

two-stage procedure:

estimate each θ_j separately, combine the results by
Fisher's procedure for combining independent tests:

'is there any evidence for a hypothesized effect?'

See, e.g., Snijders & Bosker, "Multilevel Analysis", Section 3.7.

Meta-Analysis ~ without population assumption (contd.):

For coordinate k of the parameter, test null hypothesis

$$H_0 : \theta_{kj} = 0 \text{ for all } j$$

against alternative hypothesis

$$H_1 : \theta_{kj} \neq 0 \text{ for at least one } j .$$

RSiena: first estimate for all groups, then use `siena08()`.
Useful when there are few groups.

Mercken, Snijders, Steglich, & de Vries (2009)
applied this in a study of smoking initiation:
7704 adolescents in 70 schools in 6 countries.

This principle here is applied in a double test for θ_{kj}
(for each given effect k)

- 1 for detecting if any groups j have $\theta_{kj} > 0$,
the null hypothesis tested is
 H_0^r : For all groups j we have $\theta_{kj} \leq 0$
with the alternative right-sided hypothesis;
 H_1^r : For at least one group j , $\theta_{kj} > 0$;
- 2 for detecting if any groups j have $\theta_{kj} < 0$,
the null hypothesis tested is
 H_0^l : For all groups j we have $\theta_{kj} \geq 0$
with the alternative left-sided hypothesis;
 H_1^l : For at least one group j , $\theta_{kj} < 0$;

For each effect, `siena08` gives the right-sided p -value p_r ,
and the left-sided p -value p_l .

To account for the double testing, the significance level α should be divided by 2; corresponding to doubling the p -values.

It is very well possible that p_r as well as p_ℓ are significant (e.g., for a combined $\alpha = 0.05$, that p_r as well as p_ℓ are < 0).

This means that some groups have positive, others negative parameter values; a perfectly possible conclusion.

To account for the double testing, the significance level α should be divided by 2; corresponding to doubling the p -values.

It is very well possible that p_r as well as p_ℓ are significant (e.g., for a combined $\alpha = 0.05$, that p_r as well as p_ℓ are < 0).

This means that some groups have positive, others negative parameter values; a perfectly possible conclusion.

The following pages present the two one-sided p -values for the Baerveldt data for comparison combined with the earlier results for μ .

	N	$\hat{\mu}$	$se(\hat{\mu})$	ρ_{μ}	ρ_l	ρ_r
<i>Friendship dynamics</i>						
outdegree (density)	16	-2.955	0.149	0.00	1.00	0.00
reciprocity	16	3.588	0.253	0.00	0.00	1.00
GWESP	16	2.377	0.120	0.00	0.00	1.00
indegree - popularity	16	-0.118	0.024	0.00	1.00	0.00
outdegree - activity	16	0.121	0.024	0.00	0.00	1.00
rec.degree - activity	16	-0.204	0.067	0.00	1.00	0.00
indegree - activity	16	-0.204	0.072	0.00	1.00	0.00
same ethnicity	16	-0.005	0.089	0.96	0.36	0.42
sex (M) alter	16	-0.329	0.086	0.00	1.00	0.00
sex (M) ego	16	-0.169	0.173	0.33	0.66	0.04
same sex	16	0.554	0.080	0.00	0.00	1.00
crimi alter	16	-0.053	0.046	0.24	0.88	0.31
crimi ego	16	0.074	0.080	0.35	0.09	0.73
crimi ego × alter	16	0.006	0.057	0.92	0.25	0.52
reciprocity × GWESP	14	-0.568	0.272	0.04	0.98	0.13

	N	$\hat{\mu}$	$se(\hat{\mu})$	ρ_{μ}	ρ_l	ρ_r
<i>Delinquency dynamics</i>						
linear shape	16	0.015	0.108	0.89	0.65	0.77
quadratic shape	16	-0.295	0.048	0.00	1.00	0.00
outdegree	16	0.042	0.047	0.37	0.55	0.93
total alter	16	0.076	0.064	0.23	0.38	0.99
sex (M)	16	0.610	0.139	0.00	0.00	1.00

In this case, the results of Fisher's combination of tests are in line with the random-effects meta-analysis; there is a difference in 'significance' only for reciprocity \times GWESP.

This is related to the small variability between the groups, and the fact that no parameters have ambiguous signs.

In other cases, it may well be different!
Statistical reasoning gives no certainty...

See the manual for further explanation.

Literature

The example using Chris Baerveldt's data is in the script *RscriptMultipleGroups_meta.R* on the Siena website.

References

RSiena manual: Chapter 11.

Johan H. Koskinen and Tom A. B. Snijders (2022), Multilevel Longitudinal Analysis of Social Networks. *arXiv preprint arXiv:2201.12713*.

DOI: <https://doi.org/10.48550/arXiv.2201.12713>
> To be published (in revised form) in *Journal of the Royal Statistical Society, vol. A*.

Tom A.B. Snijders (2016). The Multiple Flavours of Multilevel Issues for Networks. Chapter 2 in Emmanuel Lazega and Tom A.B. Snijders (eds.), *Multilevel Network Analysis for the Social Sciences*. Cham: Springer.

Tom A.B. Snijders and Chris Baerveldt (2003). A Multilevel Network Study of the Effects of Delinquent Behavior on Friendship Evolution. *Journal of Mathematical Sociology* 27, 123–151.

Tom A.B. Snijders and Roel J. Bosker (2012). *Multilevel Analysis: An Introduction to Basic and Advanced Multilevel Modeling*, 2nd edition. London: Sage.

Wolfgang Viechtbauer (2010). 'Conducting meta-analyses in R with the metafor package', *Journal of Statistical Software*.