An example of multilevel analysis of network dynamics using `sienaBayes`

Tom A.B. Snijders

University of Groningen
University of Oxford

June, 2024

Example: data Andrea Knecht

This is an example of function `sienaBayes` in package `multiSiena` for the estimation of multilevel longitudinal network models.

The theory is in
Johan H. Koskinen and Tom A. B. Snijders (2023), 'Multilevel Longitudinal Analysis of Social Networks'. *Journal of the Royal Statistical Society, Series A*, 186, 376–400.
DOI: `https://doi.org/10.1093/jrsssa/qnac009`

The script used is `RscriptsienaBayes_5.r`.

The data set used is about friendship networks in 21 school classes from the study by Andrea Knecht (PhD thesis Utrecht, 2008).

See Knecht, Snijders, Baerveldt, Steglich, & Raub,
'Friendship and Delinquency: Selection and Influence Processes in Early Adolescence',
Social Development, 2010.

We consider a model for a longitudinal study with 2 waves, dependent variables friendship and delinquent behavior.

Data description 1

Group sizes range from 13 to 33, with mean 26.

Except for 1 group, all sizes range from 21 to 33.

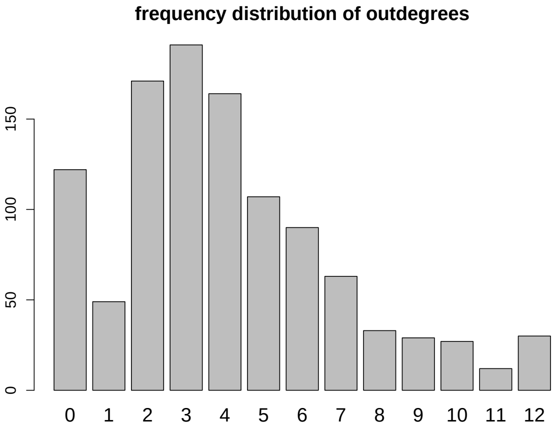

Data description 2

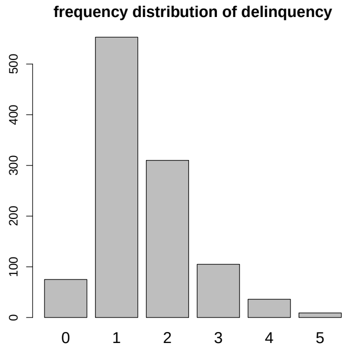

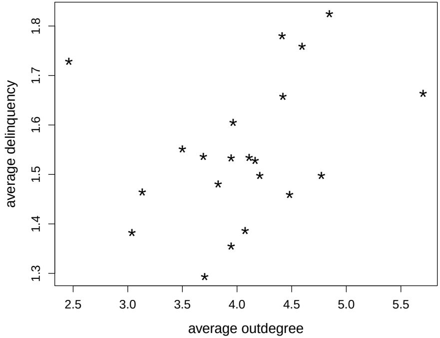

Model specification

In addition to the regular effects, for multilevel models
we should think about group-level effects.

1. The groups may have different numbers n of actors.
   Snijders (2005; Section 11.13(B)) derives that for the empty model the outdegree parameter will have a component approximately $-\frac{1}{2}\log(n)$.
   Therefore it is recommendable to include $\log(n)$ as a covariate;
   the expected regression coefficient is something like -0.5,
   but for non-empty models the value will differ and is unknown.
   An alternative is to use a creation effect of $\log(n)$;
   then the expected coefficient is something like –1.
2. Furthermore, it may be advisable to include interactions of $\log(n)$ with reciprocity and transitivity.
3. Other group-level variables may also be relevant.

However, the number of group-level variables should not be too large!
The same considerations apply as for the number of covariates,
given sample size, in linear regression models;
sample size here is number of groups, which usually is small.

Within- and between-group regressions

Similar to the Hierarchical Linear Model of multilevel analysis,
we should be aware that within-group regression coefficients
may differ from between-group coefficients.
The group mean of covariates, or dependent behavior variables,
may be included in the model to account for this.

By implication, cross-level interactions may be included.

This represents here that social processes may be different in groups
depending on their average delinquency;
or, perhaps, represents regression to the mean.

In this case, we include the group mean of delinquency (ego)
and the interaction of this group mean with delinquency alter.

Group-level variables

Log group size is a variable defined at the group level
and gets an `egoX` effect.

Delinquency is an individual-level variable.
The effect of the current group average of delinquency
can be represented by effects `avGroupEgoX` (on friendship)
and `avGroup` (on delinquency itself).

('Current' refers to the simulation process.)

Effects of delinquency on network evolution

From Snijders & Lomi (*Network Science*, 2019),
we know that actor variables (here: delinquency)
may have a variety of effects on networks,
because such effects imply a level transition monadic ⇒ dyadic.

In a first analysis, the five-parameter model was used:
$V(\text{ego})$, $V(\text{alter})$, $V^2(\text{ego})$, $V^2(\text{alter})$, $\big(V(\text{ego}) - V(\text{alter})\big)^2$.

Script `RscriptsienaBayes_5.r` reports this 'first analysis',
and goes on to 'A slightly different model'.

From the first analysis it seemed that delinquency alter,
for given ego, has approximately a linear effect.
Therefore the model was reduced to four parameters:
$V(\text{ego})$, $V(\text{alter})$, $V^2(\text{ego})$, $V(\text{ego}) \times V(\text{alter})$.

Summary of model specification

Network dynamics:

outdegree; reciprocity; transitive triplets; transitive reciprocated triplets;
indegree popularity; outdegree activity; reciprocal degree-activity;
old friends; same sex; $\log(n)$;

dependence on delinquency V:
V(ego), V(alter); V^2(ego); V(ego) $\times$ V(alter);
V(group mean ego); V(group mean ego) $\times$ V(alter);

Delinquency dynamics:

linear shape; quadratic shape; sex; average alter; group mean delinquency.

For 21 groups, 2 is a rather high number of group-level variables.
Therefore a prior distribution
(mean –0.5, variance 0.25, corresponding to s.d.=0.5)
was assumed for the evaluation effect (egoX) of $\log(n)$.

Random / fixed

The choice for which parameters to define as random was based on a preliminary multi-group analysis by MoM (`siena07`)
where all parameters were assumed fixed,
followed by `sienaTimeTest` to test parameter homogeneity;
parameters with the largest test statistics were defined as random.

The following effects were considered to vary randomly:

Network dynamics:
outdegree; reciprocity; transitive triplets; indegree popularity;
reciprocal degree-activity; delinquency ego;

Delinquency dynamics:
linear shape; quadratic shape.

The number of randomly varying effects should be smaller
when there are fewer groups in the data set.

Initially, also the effects of outdegree activity and same sex were considered to be randomly varying, but an MDS plot suggested that this number of randomly varying effects was too high.

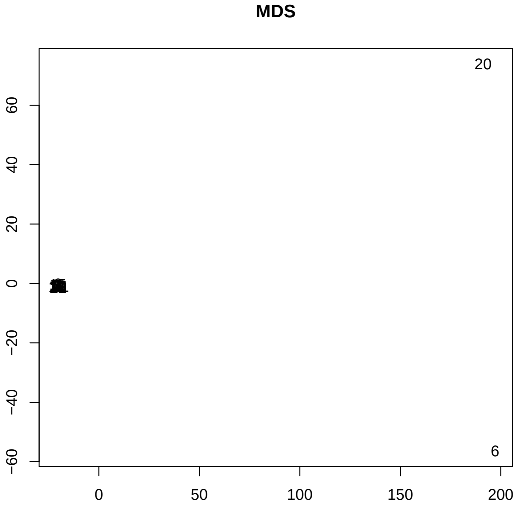

Prior distributions

Prior for rates: data dependent, calculated internally; for the rest:
prior means for μ:
outdegree –2, recip 1, others 0;
prior variances for μ: outdegree and linear shape 0.1, others 0.01;
prior covariances: all 0;

group-level variables get priors for η:
prior means for η: for $\log(n)$ -0.5,
for V(group mean) and V(group mean) $\times$ V(alter) 0;
prior variances for η: infinite, but for the group variables 0.25;
prior κ: 0.01.

This means that the between-group differences of parameters $\theta_j^{(1)}$
are thought to be in the order of magnitude of $\sqrt{0.01} = 0.1$,
and the uncertainty about the values of the prior means μ of $\theta_j^{(1)}$
is of the order of $\sqrt{0.01/0.01} = 1$.

(Recall that the above are variances, which are on the scale of squares.)

First a multi-group was estimated, and used as `prevAns` for `sienaBayes`.

The parameter `mult` was determined to keep all autocorrelations below 0.40.

Then a first estimation (**A**) was done,
using the following parameters for the MCMC algorithm:

1. groupwise number of MH iterations for sampling mini-steps varies between 600–1600 depending on distance between observed networks (determined by the multiplication factor);
2. ⇒ 500 iterations sampling $\theta_j^{(1)}, \eta, \mu, \Sigma$ for warm-up
   ⇒ 1000 iterations sampling $\theta_j^{(1)}, \eta, \mu, \Sigma$ for estimation,
   with a thinning ratio of 1:20.

This was studied, and all seemed OK.

Next, 5 independent estimation runs were carried out (**B1**)
(using different random number seeds)
with 1000 main iterations and a thinning ratio of 40
all starting from the results of the first estimation **A** (using `prevBayes`)
and the results analyzed for the R-hat as computed by package `rstan`.

The highest R-hat was 1.13, not adequate.

The procedure was repeated (**B2**),
now each run starting from the end of its value in **B1**.
This led to a highest R-hat for **B1+B2** of 1.21, worse than before.

The procedure was repeated again (**B3**),
now each run starting from the end of its value in **B2**.
This led to a highest R-hat for **B1+B2+B3** of 1.10.
This was found adequate (but see further on).

Combining the **B** runs was done using `glueBayes`.

In the following slides, trace plots are given for one run out of **B1+B2+B3**.

Trace plots for posterior μ_k and σ_k for rate parameters:

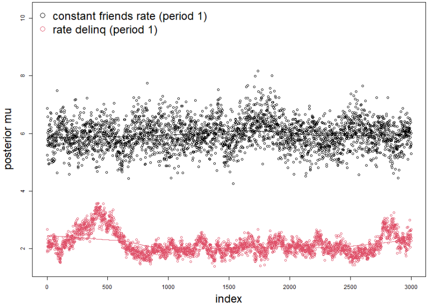

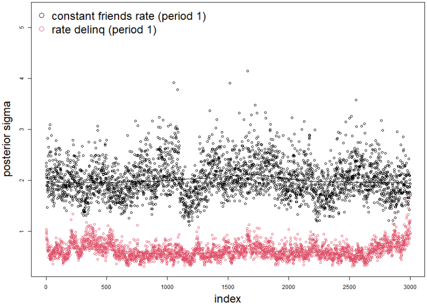

Trace plots for μ_k for structural parameters for friendship and η_k for parameters for delinquency:

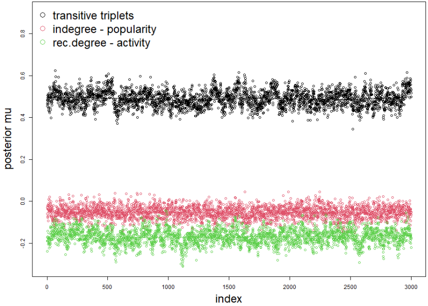

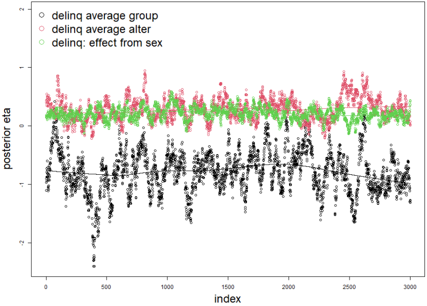

The trace plots are stationary, but do show waves;
it helps that there are 5 independent sets of runs.

The following pages show posterior means and standard deviations
with 95 % credibility interval for μ_k, η_k, and σ_k
computed from the union of the 5 posterior samples in **B1+B2+B3**,
using `glueBayes`.

	$\hat{\mu}_k, \hat{\eta}_k$	(mFrom	mTo)	$\hat{\sigma}_k$	(sFrom	sTo)
Network dynamics: friendship						
outdegree (density)	–1.937	(–2.283	–1.579)	0.480	(0.308	0.694)
reciprocity	2.113	(1.815	2.410)	0.215	(0.108	0.375)
transitive triplets	0.497	(0.421	0.577)	0.116	(0.081	0.162)
transitive recipr. triplets	–0.207	(–0.299	–0.115)	.	.	.
indegree - popularity	–0.055	(–0.118	0.005)	0.118	(0.082	0.166)
outdegree - activity	–0.002	(–0.029	0.023)	.	.	.
rec.degree - activity	–0.164	(–0.239	–0.091)	0.115	(0.081	0.160)
oldties	0.375	(0.201	0.550)	.	.	.
del. alter	0.036	(–0.050	0.129)	.	.	.
del. ego	0.081	(–0.149	0.425)	0.306	(0.157	0.538)
del. squared ego	–0.068	(–0.253	0.065)	.	.	.
del. ego x del. alter	0.048	(–0.050	0.152)	.	.	.
del. group-av. ego	–0.184	(–1.069	0.729)	.	.	.
del. group-av. ego x del. alter	–1.436	(–3.310	1.167)	.	.	.
same sex	0.510	(0.374	0.643)	.	.	.
nn ego	–0.413	(–1.182	0.376)	.	.	.

$\hat{\mu}_k, \hat{\eta}_k$: posterior mean for μ_k or η_k, as the case may be,
with 95% credibility interval from `mFrom` to `mTo`;

$\hat{\sigma}_k$: posterior between-groups s.d.,
with 95% credibility interval from `sFrom` to `sTo`.

	$\hat{\mu}_k, \hat{\eta}_k$	(`mFrom`	`mTo`)	$\hat{\sigma}_k$	(`sFrom`	`sTo`)
Behavior dynamics: delinquency						
delinq linear shape	0.347	(–0.097	0.792)	0.342	(0.240	0.477)
delinq quadratic shape	–0.266	(–0.385	–0.154)	0.168	(0.102	0.254)
delinq average group	–0.758	(–1.467	–0.020)	.	.	.
delinq average alter	0.348	(–0.013	0.765)	.	.	.
delinq: effect from sex	0.199	(–0.012	0.417)	.	.	.

$\hat{\mu}_k, \hat{\eta}_k$: posterior mean for μ_k or η_k, as the case may be,
with 95% credibility interval from `mFrom` to `mTo`;

$\hat{\sigma}_k$: posterior between-groups s.d.,
with 95% credibility interval from `sFrom` to `sTo`.

The hierarchical multilevel approach gives much more information:
for each group, the posterior distribution of the parameters
(which are constant across groups for the fixed parameters).

The following pages show some density plots,
again computed from the union of the 5 posterior samples in **B1+B2+B3**.

For the varying parameters, note the difference between
the posterior density of μ and that of average θ_j;

the average θ_j has a bearing on this sample only;
for μ, there is the extra uncertainty
due to generalisation from sample to population.

Density plot for outdegree (density) effect for friendship.

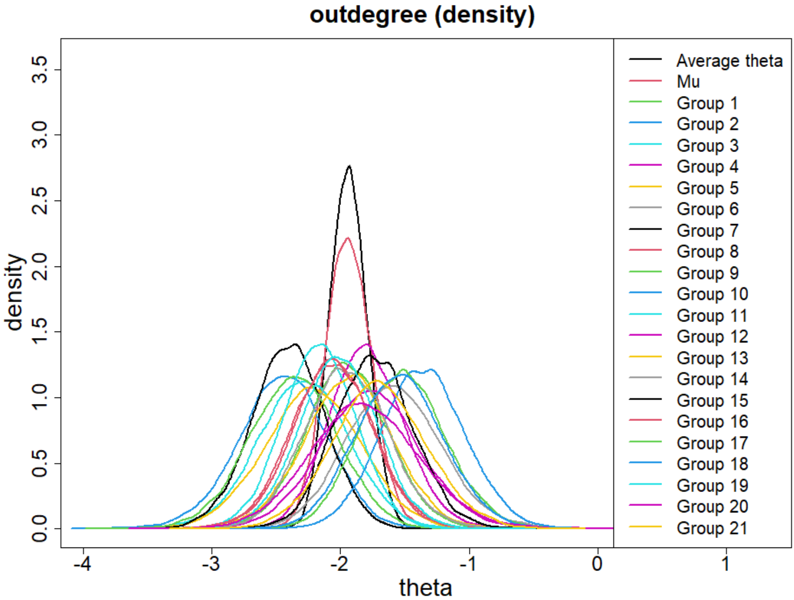

Density plot for reciprocity effect for friendship.

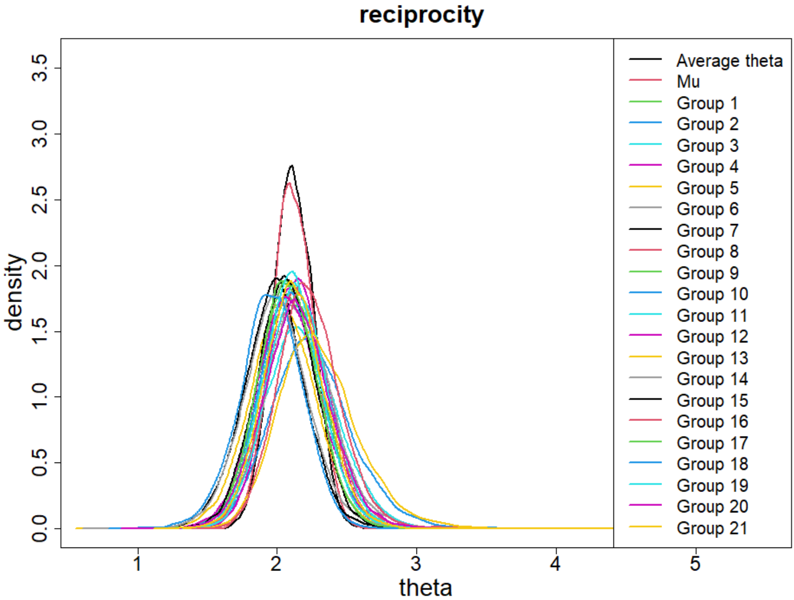

Density plot for reciprocal degree effect for friendship.

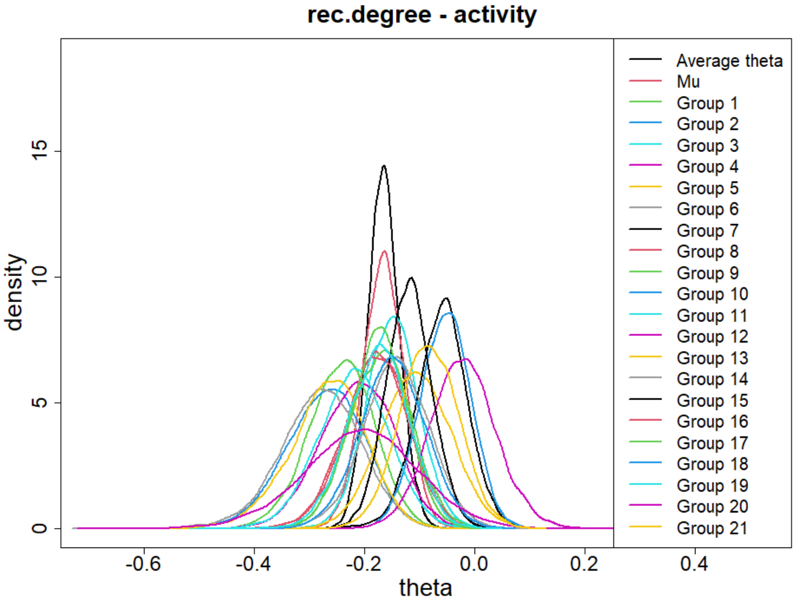

Density plot for delinquency ego × alter effect for friendship.

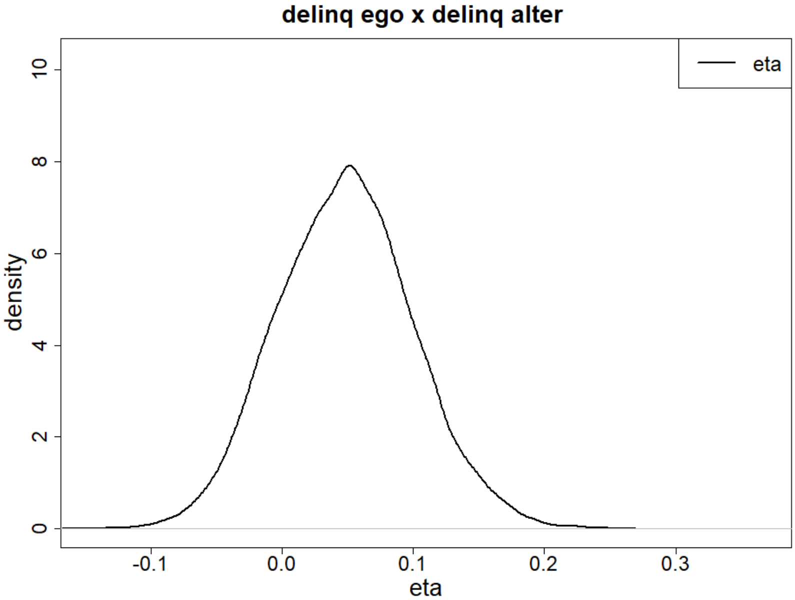

Density plot for same sex effect for friendship.

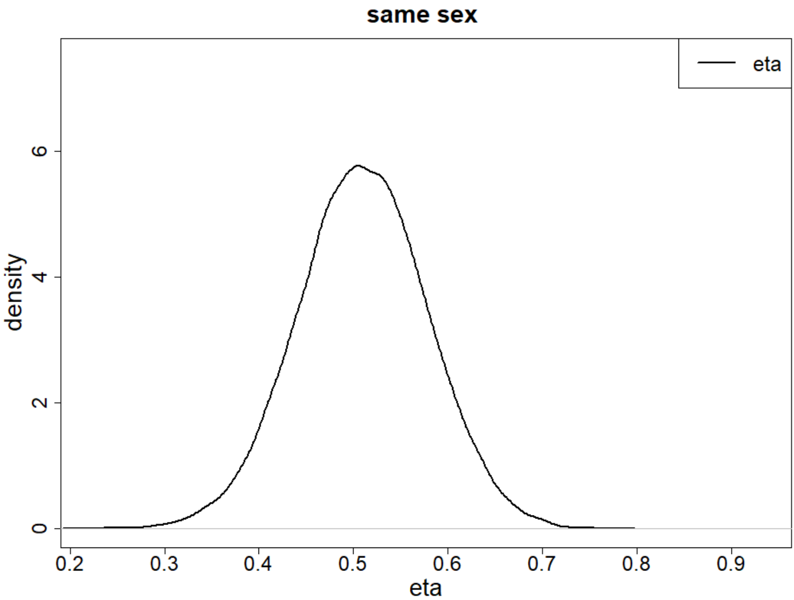

Density plot for $\log(n)$ effect for friendship.

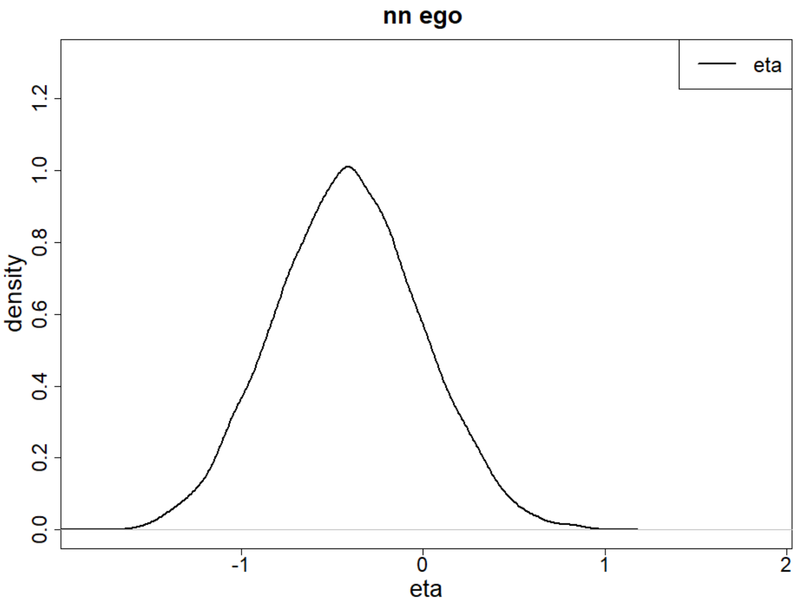

Scatter plot for the effects on friendship of delinquency(group mean) and delinquency (group mean) × alter.

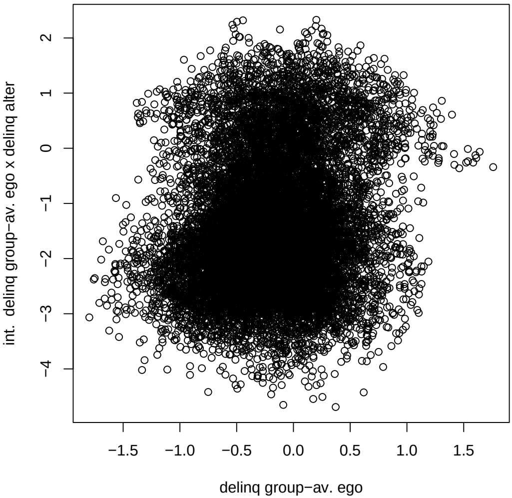

Density plot for quadratic shape effect for delinquency.

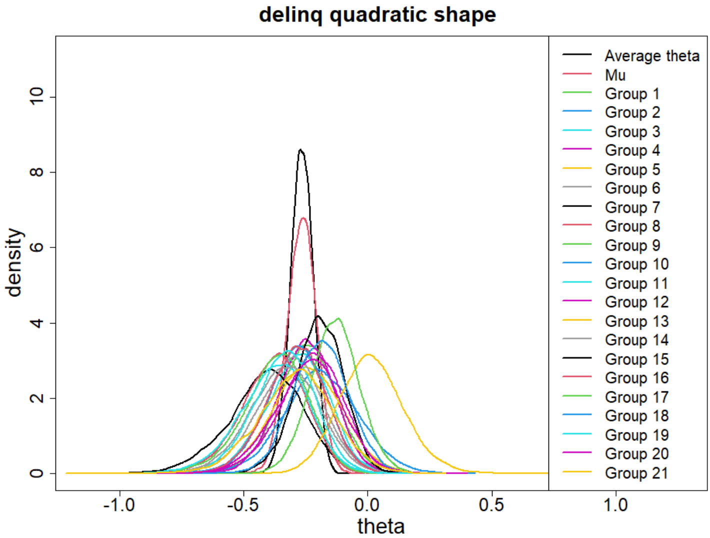

It would be interesting to look more closely at group 13,

the only group that has a positive posterior mean for quadratic shape.

Density plot for effect of sex on delinquency.

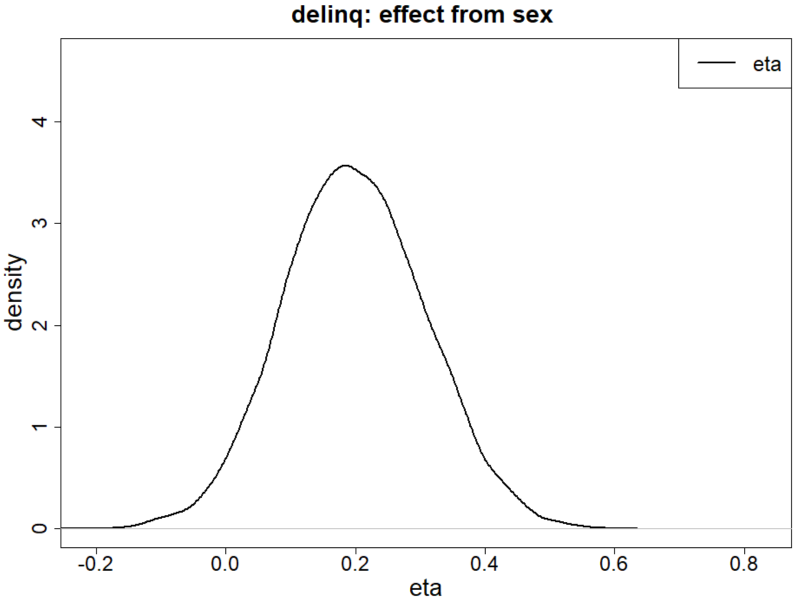

Density plot for effect of average alter (friendship) on delinquency.

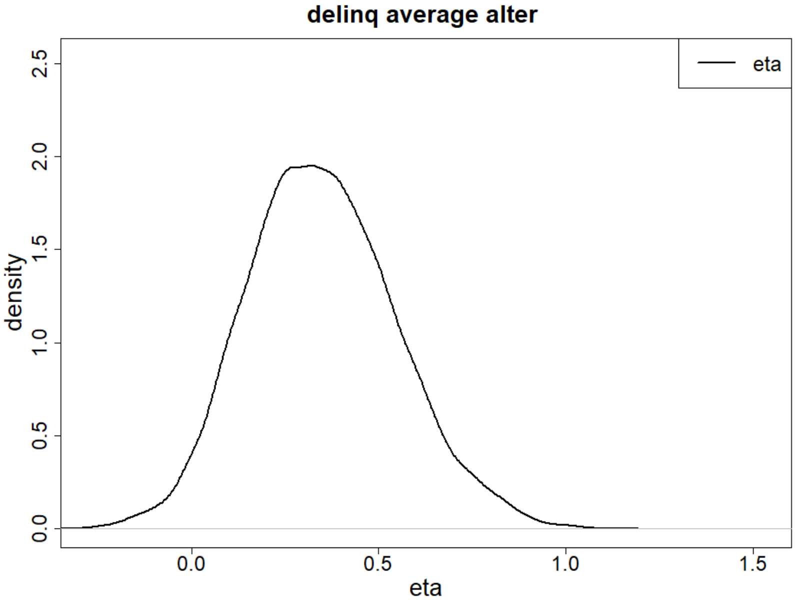

The within-group effect of delinquency on friendship selection can be shown by a selection plot, made by `SelectionTables.r`.

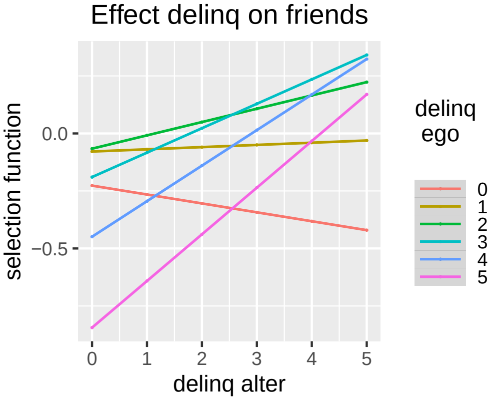

The delinquency group mean is not included; it was centered about 1.5 (close to average), so this plot refers to groups with a mean delinquency of 1.5.

But the effects of group mean delinquency on friendship selection are not important.

Although the value of 1.10 for R-hat is satisfactory, convergence is not perfect.

The collection of trace plots produced by the `bayesplot` package (see script `RscriptsienaBayes_5.r`) shows that convergence still is not very good.

It is clear that the group-level variables give the most problems, and further analysis with a reduced model might be good.

For the effects on friendship of $\log(n)$, V(group mean ego), and V(group mean ego) $\times$ V(alter), the posterior distributions do not signal that they are different from 0.

Therefore, further analysis could be done for a model without the effects involving the group mean of delinquency.

The density plots suggest that, if an effect would have to be chosen for dropping the assumption of variability between groups, the reciprocity effect would be a good candidate.

A better analysis of friendship and delinquency

The paper by Koskinen and Snijders (2023) contains an analysis of a larger part of this data set (3 waves) with some selection (82 classrooms with Jaccard ≥ 0.20, and few missing respondents), where delinquency was represented as a two-mode network (4 items).

It was found there that there is social influence w.r.t. delinquency at the level of the individual items (stealing, vandalism, graffiti, fighting).

The script is at

`https://www.stats.ox.ac.uk/~snijders/siena/sienaBayes_Script.R` .