

Day 12 Presenter Notes

Shahryar Minhas

DAY 12 PRESENTER NOTES Day 12: Stochastic Actor-Oriented Models Network change, co-evolution, and what the model can claim

HOW TO USE THESE NOTES

Keep the teaching deck visible and use this file as the spoken route. Each section follows one slide in the current 33-slide deck, gives a conversational explanation rather than a transcription, and ends with a transition to the next slide. Do not run the models during class. The rendered walkthrough already contains the executable code, full output, and technical detail students need.

SLIDE 1: Day 12: Stochastic Actor-Oriented Models

What to say

Yesterday, the TERGM treated change as a transition from one observed network to the next. Today, the actor takes

the modeled micro-step. A SAOM represents the interval between waves as a sequence of opportunities in which one actor can reconsider one outgoing tie, or later, one behavior value.

The payoff is a sharper language for network change and co-evolution, but the sharper language comes with stronger process assumptions. The micro-steps are not observed decisions. They are the model's account of how the observed waves could be connected.

Transition

Before opening RSiena, start with the model-choice question: is an actor-oriented process defensible for this relation?

SLIDE 2: Is This Actually a SAOM Question?

What to say

A longitudinal network is necessary for a SAOM, but it is not sufficient. The substantive process should plausibly involve actors controlling outgoing relations and receiving opportunities to make relatively small changes between meaningful observation waves.

For friendship nominations, the actor-oriented framing is plausible because a student names outgoing friends. For relations that are imposed, jointly controlled, or changed only through large coordinated events, the framing may be much weaker.

Use the prompt as a genuine model-selection test: can the

process be approximated as actors making small tie changes while the network, and possibly an actor behavior, evolves? If the answer is no, extra waves do not rescue the model choice.

Transition

Once the actor-oriented framing is plausible, move from the software question to the social science question that motivates co-evolution.

SLIDE 3: Did Similarity Produce the Relationship?

What to say

The recurring question is whether similarity preceded the relationship or the relationship preceded greater similarity. States may ally because their security policies already align, legislators may collaborate because their positions already match, and armed groups may cooperate because their strategies are already compatible. The reverse sequence is also possible in every example.

Repeated observations help establish which measured state came earlier, but they do not reveal every change inside an interval. A SAOM can show patterns consistent with selection or influence under its continuous-time micro-step process. It does not prove either mechanism merely because both variables were measured more than once.

Keep the fitted application in view: the data follow friendship nominations and drinking among 50 Scottish adoles-

cents across three waves. Every substantive result later in the deck should return to those actors, outcomes, and intervals.

Transition

Now distinguish the applied questions answered by a network SAOM, a co-evolution SAOM, and the TERGM comparator from Day 11.

SLIDE 4: The Specifications Answer Different Applied Questions

What to say

The network SAOM asks which candidate tie changes receive more weight in an actor's objective function. The co-evolution SAOM adds a behavior equation so that selection into similar friendships and movement toward friends' behavior are represented separately. The TERGM instead specifies a distribution for the next network conditional on the past network and covariates.

These are not three interchangeable routes to the same coefficient. A SAOM result should name a candidate micro-step or a simulation-based outcome contrast. A TERGM result should name a conditional network transition. A positive SAOM evaluation parameter is not an unconditional probability that a tie forms.

Transition

With the estimand and process question clear, inspect whether the panel actually defines the actor set and risk set the model needs.

SLIDE 5: Pre-Flight Checks Come Before Fitting

What to say

Begin with actor composition, structural zeros, network change, and measurement. Ask who was present and eligible in each interval, which particular ties were truly impossible, whether the network has both continuity and change, and whether missingness, behavior ranges, direction, and tie definitions are defensible.

Keep two data operations separate. A structural zero removes a particular dyad from the actor's possible choices. A composition-change object specifies when an actor belongs to the modeled population. Joining or leaving is not a control variable.

In this application, **netify** verifies a common actor order, direction, and matrix dimensions before the three friendship waves become the RSiena array. The behavior matrices remain separate because drinking and smoking are actor attributes rather than network outcomes.

Transition

After the observed panel is defensible, explain the unobserved path that the model inserts between adjacent waves.

SLIDE 6: The Model Fills the Space Between Waves

What to say

The data show a friendship network at one interview and another network at the next interview. The model connects those endpoints with many latent micro-steps. At each step, one actor receives an opportunity to add one outgoing tie, remove one outgoing tie, or make no change.

The sequence on the slide is not a reconstruction of events that were recorded. RSiena integrates over many possible paths that could connect the observed endpoints under the specified rate and evaluation functions.

Use a tiny mental example: if Alice receives an opportunity, list her eligible additions, deletions, and no-change option. That is the choice set the evaluation function will score.

Transition

The next distinction is essential because the model separates how often an opportunity occurs from which option receives greater weight.

SLIDE 7: Rate and Evaluation Parameters Do Different Jobs

What to say

The rate function describes the pace of the latent process: how often actors receive opportunities between waves. The evaluation function describes direction: which available next state is more attractive when an opportunity occurs.

More opportunities do not imply a preference for more ties. An actor can receive many opportunities and repeatedly choose no change, or can add and later remove the same tie. Rate parameters therefore should not be read as log odds or expected degree.

Transition

Now open the evaluation function and show how familiar network statistics become contributions to a candidate state's score.

SLIDE 8: The Evaluation Function Scores Candidate States

What to say

The evaluation function adds contributions from statistics such as outdegree, reciprocity, transitive triplets, and covariate-related selection. For one candidate state, each statistic is multiplied by its fitted coefficient and the pieces are summed.

The useful question is not, “What is the probability of friendship from this coefficient?” The useful local question is, “If actor i considers the tie from i to j , which statistics change, and how does each change alter the candidate state’s score?”

Connect this to Day 11 carefully. ERGMs and SAOMs both use changes in network statistics, but the probability model differs. An ERGM change statistic enters a conditional tie comparison in a graph distribution. A SAOM statistic enters an actor’s choice among available micro-steps.

Transition

Because the micro-step path is missing from the observations, estimation must simulate it rather than treat the observed transitions as a conventional regression table.

SLIDE 9: Estimation Simulates the Missing Path

What to say

Walk through the loop in order. The rate function generates opportunities, the evaluation function assigns probabilities across available one-step changes, and RSiena simulates a complete path to the next wave. It then compares statistics from the simulated path with the corresponding statistics in the observed panel.

If the simulated process produces too little reciprocity, too little closure, or the wrong amount of change, the parameters are revised and the simulation is repeated. Estimation learns parameter values under which the observed changes

are typical for the specified micro-step process.

Transition

The next slide writes the target of that simulation loop explicitly and separates method of moments from maximum likelihood.

SLIDE 10: What RSiena Is Solving

What to say

Let the observed target statistics be one vector and the average simulated statistics under the current parameters be another. RSiena seeks parameters for which their difference is approximately zero.

Robbins-Monro stochastic approximation updates the parameters using noisy simulation-based discrepancies. In the final phase, the estimate is held fixed while additional simulations support standard errors and convergence checks.

Default SAOM estimation here is simulation-based method of moments, not likelihood maximization. Convergence means the targeted simulated moments are sufficiently close to their observed counterparts within Monte Carlo error. It does not establish that actors consciously followed the modeled objective function.

Transition

That estimation burden is one reason to begin with the smallest specification that answers the substantive question.

SLIDE 11: Fit the Smallest Model That Answers the Question

What to say

Start with the basic density or outdegree tendency, add reciprocity when it is substantively plausible, choose one defensible closure term, and include the focal covariate or behavior mechanism. Each effect should have a clear job before it enters the specification.

Do not treat the RSiena effect catalog as a menu from which every available term should be selected. Each added statistic changes the actor's choice surface, creates new dependencies among parameters, and makes convergence harder to diagnose.

The code record should remain readable: create the effects object, add effects deliberately with `includeEffects()`, and retain a direct mapping between each term and the question it is meant to answer.

Transition

Before reading any coefficient from that first fit, establish that the simulated moments actually converged toward their targets.

SLIDE 12: Convergence Is About Simulated Moments Matching Targets

What to say

Inspect the overall maximum convergence ratio, the effect-specific t-ratios for deviations, and stability across repeated or continued runs. These are diagnostics of whether estimation finished, not tests of the substantive hypotheses.

If a run is close but unfinished, continuing from the previous estimate can be appropriate. Restarting blindly wastes the information already learned, while adding more effects to a nonconverged model usually makes the source of the problem less visible.

A printed coefficient is not ready for interpretation merely because RSiena returned a table. The simulated targets must first align well enough with the observed change statistics.

Transition

With convergence established, name exactly what the first network-only fit represents before showing its evaluation coefficients.

SLIDE 13: What the First Fit Is Describing

What to say

This first fit models friendship change only. Between waves, students receive opportunities to reconsider one outgoing nomination. Smoking is a fixed wave-1 covariate in this model; smoking itself does not evolve.

The fitted rates are about 6.52 opportunities per student

from waves 1 to 2 and 5.21 from waves 2 to 3. Describe that as somewhat more modeled activity in the first interval, not as more friendships or stronger preferences.

An opportunity can end in no change, and the same tie can be reconsidered more than once. The rate parameters therefore describe the latent pace of the fitted process rather than observed counts of changed nominations.

Transition

Now read the evaluation effects together, beginning with signs and the candidate changes they score.

SLIDE 14: The First Fit in One Table

What to say

One more outgoing nomination has a negative contribution of -2.78, so additional nominations are less attractive in an otherwise comparable candidate state. Returning a nomination contributes 2.48, and closing one friend-of-friend path contributes 0.66.

The three-cycle estimate is -0.09 with a standard error of 0.29, and same baseline smoking is 0.14 with a standard error of 0.15. Those data do not clearly distinguish either pattern from zero after the other modeled features are included.

Use estimate divided by standard error only as a rough uncertainty check. The more important interpretive habit is to state the feature, the proposed change, and whether that feature makes the candidate state more or less attractive

while the other modeled contributions are held fixed.

Transition

The next two slides slow down on the two clearest structural signals, beginning with reciprocity.

SLIDE 15: Returned Friendships Carry the Strongest Signal

What to say

Compare two proposed nominations from i to j . In one case, j does not nominate i . In the other, j already nominates i , so adding the outgoing tie returns that nomination and creates reciprocity.

The coefficient difference is 2.48 in the objective-function score. Exponentiating gives about 11.9, so the reciprocating candidate receives roughly twelve times the choice weight of an otherwise comparable candidate that does not reciprocate.

Do not say that the students are twelve times more likely to be friends. The actual action probability depends on every eligible addition, deletion, and no-change option, plus every statistic that changes for each option.

Transition

Use the same candidate-change logic for closure, where one proposed nomination can close one or several friend-of-friend paths.

SLIDE 16: Friends of Friends Are More Attractive Candidates

What to say

Alice nominates Ben, and Ben nominates Cara. If Alice adds a nomination to Cara, the proposed change closes one transitive path. The fitted contribution is 0.66 for each path closed.

Exponentiating 0.66 gives about 1.94, so closing one path nearly doubles the choice weight of the proposed nomination, holding the other modeled features constant. Closing three paths contributes three times 0.66 to the score because this specification treats the transitive-triplet effect as linear.

This is evidence about the fitted closure pattern, not proof that an adolescent consciously searched through friends' nomination lists.

Transition

Balance those positive results by stating clearly what the same specification does not establish.

SLIDE 17: What the First Fit Does Not Show

What to say

The three-cycle estimate does not reveal a clear circular exchange pattern after reciprocity and transitivity are included. The baseline same-smoking estimate likewise does

not reveal a clear smoking-similarity pattern after the network terms are included.

Avoid the stronger sentence that smoking and friendship are unrelated. Smoking is frozen at wave 1 in this model, so the specification cannot represent later smoking changes or peer influence on smoking.

The defensible summary is narrower: friendships tend to be returned, friends of friends tend to become connected, and baseline smoking similarity contributes little in this particular network-only specification.

Transition

Before moving to co-evolution, combine several effects in one proposed nomination to show why coefficients cannot be interpreted in isolation.

SLIDE 18: One Proposed Nomination Changes Several Statistics

What to say

Alice's proposed nomination to Cara simultaneously adds one outgoing tie, returns Cara's nomination, and closes two friend-of-friend paths. The contributions are -2.78 for out-degree, +2.48 for reciprocity, and +1.32 for two transitive paths.

The combined score change is +1.02. Relative to making no change, that is about 2.77 times the choice weight. It remains a comparison inside the modeled choice set, not a

tie probability and not an account of Alice's private reason for the nomination.

This example also explains why correlated statistics can be difficult to estimate separately. If reciprocating ties almost always close paths, the data offer limited independent variation for separating reciprocity from transitivity.

Transition

The next slide interprets the same-smoking indicator as a direct 0-to-1 predictor change.

SLIDE 19: Interpret the Same-Smoking Coefficient

What to say

Start with the variable. The same-smoking indicator equals zero when the two students report different smoking categories and one when they report the same category. Moving from zero to one is therefore the familiar one-unit change in a binary predictor.

The fitted coefficient is 0.14. Exponentiating gives 1.15. In a multinomial choice model, this is a conditional probability ratio between two otherwise identical possible actions.

Use this sentence: at a modeled friendship-change opportunity, a nomination to a student with the same smoking status is about 15 percent more likely to be chosen than an otherwise identical nomination to a student with a different smoking status.

Be precise about the unit. This is 15 percent higher conditional probability relative to the comparison nomination, not a 15 percentage-point increase. The coefficient has a standard error of 0.15, so the fitted difference is uncertain.

Transition

Now make drinking endogenous and separate the two histories that can produce similar friends at the final wave.

SLIDE 20: One Pattern, Two Possible Histories

What to say

Alice and Cara end as friends with similar drinking, but the endpoint does not reveal the sequence. Under selection, their drinking was already similar and the friendship appeared later. Under influence, the friendship existed first and drinking became more similar afterward.

Several waves provide information about ordering, but they still do not reveal every tie and behavior change inside an interval. The co-evolution model supplies a continuous-time sequence in which one actor changes one tie or one behavior step at a time.

The central identification point is modest: the model parameterizes selection and influence separately. Whether the data and assumptions separate them convincingly is a further question.

Transition

The next slide names the dependent variable and focal effect in each of the two connected equations.

SLIDE 21: The Co-Evolution Model Uses Two Equations

What to say

Make the break from Slide 19 explicit. Both friendship and drinking now change, so this is a different model with two dependent outcomes, two evaluation equations, and separate change rates for each process and interval.

In the friendship equation, the complete specification includes outdegree, reciprocity, transitive triplets, three-cycles, drinking ego, drinking alter, and drinking similarity. At a friendship opportunity, a student may add one outgoing nomination, delete one, or make no change. The focal drinking-similarity selection estimate is 1.44 with a standard error of 0.61.

In the drinking equation, the complete specification includes two baseline terms plus average friends' drinking. RSiena calls the baseline terms linear shape and quadratic shape. The linear term describes whether drinking generally tends to move upward or downward. The quadratic term describes whether movement tends toward the middle of the scale or toward its extremes. These are not time trends. At a drinking opportunity, a student may move one response category up, move one down, or make no change. The focal influence estimate is 1.36 with a standard error of 0.95.

The fitted friendship process favors students with more similar drinking as possible friends. The influence point estimate is positive but too uncertain to distinguish from no influence.

Transition

Translate the selection coefficient by asking what happens when a possible friend is one drinking category closer.

SLIDE 22: Interpret the Drinking-Similarity Coefficient

What to say

Start by naming both sides of the interpretation. The predictor is similarity in reported drinking between a student and a possible friend. The immediate outcome is how attractive that possible friendship nomination is in the model.

Drinking runs from 1 for none to 5 for more than once a week. The middle categories are once or twice a year, once a month, and once a week. Two students in the same category receive a similarity score of 1. Students one category apart receive 0.75, students two categories apart receive 0.50, and so on. Moving one category closer therefore raises similarity by 0.25.

Multiply that change by the fitted drinking-similarity coefficient: 0.25 times 1.44 equals 0.36. Exponentiating gives 1.43.

Give students the substantive sentence before discussing the technical scale: at a modeled friendship-change opportunity,

a student who is one drinking category closer is about 43 percent more likely to be chosen than an otherwise comparable student.

Then explain what “43 percent more likely” means. It is a conditional probability ratio of 1.43 between otherwise comparable possible nominations. It is not a 43 percentage-point increase because the student is choosing among all eligible additions, deletions, and no change. The probability assigned to this nomination depends on that whole choice set.

Transition

Use an equally careful outcome for influence, where movement toward friends may be upward or downward on the drinking scale.

SLIDE 23: Put the Influence Result on the Drinking Scale

What to say

Compare the complete co-evolution model with its average-friend influence coefficient at 1.36 against the same complete model with only that coefficient set to zero. Both scenarios retain the fitted rates, the entire friendship equation including selection, and the drinking linear and quadratic shape terms.

Then restrict attention to actor-intervals in which an adolescent begins with at least one friend and the adolescent's drinking differs from the starting average among

those friends. Measure the absolute distance between the simulated next-wave drinking value and that fixed starting friend average.

With the influence term set to zero, the average next-wave distance is 0.88 response categories. In the fitted model, it is 0.82 categories. The term-ablation average first difference is -0.06 categories, so including the fitted influence term reduces the simulated distance to the friends' starting average by about six hundredths of a category.

Do not summarize influence as the probability that drinking increases. Movement toward friends can require an increase for one adolescent and a decrease for another. The distance outcome respects both directions.

Again, this is a point-estimate term-ablation average first difference. It is not an inverse-logit transformation, not a causal average marginal effect, and not a causal intervention. The influence coefficient is 1.36 with a standard error of 0.95, and parameter uncertainty is not included in the displayed 0.06-category contrast.

Transition

After the two outcome-scale summaries, return immediately to the assumptions that prevent separate equations from becoming automatic causal proof.

SLIDE 24: Separate Equations Are Not Automatic Proof

What to say

An omitted common cause can still affect both friendship and drinking. A new peer group, a family change, or another unmeasured event can generate a pattern that the fitted model allocates to selection or influence.

The panel also leaves the exact ordering within each interval unknown, nominations and self-reports can be measured with error, and the meaningful process may operate faster or slower than the observation schedule assumes.

The two equations make the mechanisms explicit and separately parameterized. A causal claim still needs a defensible intervention, confounding control, timing assumptions, measurement model, network boundary, and process specification.

Transition

With the causal guardrail in place, ask whether the fitted process reproduces features of the observed network and behavior that estimation did not directly force it to match.

SLIDE 25: Goodness-of-Fit Asks What the Model Missed

What to say

Convergence and goodness-of-fit answer different questions. Convergence asks whether the estimator matched its selected target moments. Goodness-of-fit asks whether simulations from the fitted model reproduce other features

that matter for the substantive argument.

The model reproduces the broad distribution of nominations received with $p = .589$ and the drinking distribution with $p = .912$. The triad census has $p = .007$, indicating that the fitted process misses important local three-person structure.

Do not read a goodness-of-fit p-value as a conventional test that proves or rejects the entire model. Use the pattern of discrepancy to identify which claim-relevant structure the current specification fails to reproduce.

Transition

The failed triad check motivates a revision, but the next slide shows why revision must be disciplined rather than driven by a search for a larger p-value.

SLIDE 26: A Failed Check Suggests a Revision, Not a Fishing Expedition

What to say

Degree-based popularity is a plausible candidate because unequal popularity can create local clustering. Adding that effect, however, raises the overall convergence measure from 0.163 to 0.266 and leaves the triad-census p-value at .004.

That expanded fit is not a successful repair. It did not converge adequately, and its triad discrepancy did not improve. Do not interpret its coefficients or present it as evidence that popularity failed.

The next step is to obtain convergence, rerun the goodness-

of-fit check, and consider other theoretically motivated specifications. The goal is not to keep adding effects until a p-value crosses .05.

Transition

Before paying for every expanded fit, a score test can screen whether a proposed effect appears to strain against a zero restriction.

SLIDE 27: A Score Test Temporarily Pins the New Effect at Zero

What to say

The score test includes the indegree-popularity statistic in the calculations while fixing its coefficient at zero. `initialValue = 0` states the restricted value, `fix = TRUE` holds it there, and `test = TRUE` asks whether the observed statistic is poorly reproduced under that restriction.

The printed estimate remains 0.000 and the standard error is NA because the coefficient was not freely estimated. The test examines local pressure away from zero rather than estimating a new effect.

Transition

Now interpret the score statistic narrowly and list the questions it cannot answer.

SLIDE 28: What the Score Test Tells Us Here

What to say

For indegree popularity, the score test gives chi-squared 5.27 and $p = .022$. The observed friendship changes contain more popularity-related structure than the restricted model reproduces, so a freely estimated popularity effect is worth investigating.

The score test does not provide the coefficient's direction, magnitude, or standard error after estimation. It does not tell us how the existing effects will move, whether the expanded model will converge, or whether goodness-of-fit will improve.

Do not interpret the sign of RSiena's one-sided Z substantively because it depends on the score convention. Fit the expanded model to learn direction and magnitude, then repeat convergence and goodness-of-fit checks.

Transition

Even a well-specified average process may be inadequate if the evaluation parameters differ across intervals.

SLIDE 29: Time Heterogeneity Challenges One-Process Stories

What to say

Ask whether one set of evaluation parameters is plausible for every transition. Institutions and opportunities may change, observation intervals may differ, a shock may alter the network process, or the behavior scale and actor composition may shift.

A single average parameter can conceal substantively different regimes. Time-varying effects should be motivated by both the observed context and statistical evidence, not added merely because the software offers time dummies.

Transition

This returns to the model-family choice from the opening: SAOM and TERGM can describe the same panel through different temporal stories.

SLIDE 30: SAOM and TERGM Encode Different Temporal Stories

What to say

The SAOM is actor-oriented. It represents opportunities, objective-function scores, and choices over outgoing ties through a latent sequence of micro-steps. The TERGM is graph-oriented. It specifies a probability distribution for the next network conditional on prior networks and covariates.

Neither family observes the actual sequence of changes between widely spaced waves. Both approximate that missing

process, and either can be badly specified. The choice should follow actor control, temporal resolution, the research question, and the interpretation the analysis needs.

Avoid claims that one family is universally more realistic. A relation may be well represented by actor opportunities, by a network transition distribution, by a relational event model with observed events, or by none of these without further design work.

Transition

Turn those distinctions into a compact checklist for reading and reviewing SAOM applications.

SLIDE 31: The Referee Checklist

What to say

Check the application in six passes. First, ask why actor-oriented change is defensible. Second, inspect composition, missingness, structural zeros, direction, and the risk set. Third, verify convergence across effects and repeated runs.

Fourth, check whether simulations reproduce relevant network and behavior features. Fifth, match selection, influence, and causal language to the assumptions actually defended. Sixth, ask whether a TERGM or relational event model would answer a different and possibly better question.

The checklist is not a demand that every model be perfect. It is a demand that the reported claim remain inside the

support supplied by the data, model, diagnostics, and design.

Transition

Use the penultimate slide to assemble those pieces into one defensible applied result.

SLIDE 32: After the Model Runs, Describe What Changed

What to say

Begin the result with the actors, relation, behavior, and observation schedule: 50 Scottish adolescents, friendship nominations, drinking, and three waves. Then separate the rate function's change opportunities from the evaluation function's preferences over candidate states.

Report that the fitted friendship process gives greater weight to reciprocal and transitive states. The drinking-similarity selection estimate is 1.44 with a standard error of 0.61, consistent with selection into friendships with similar peers. The peer-drinking influence estimate is 1.36 with a standard error of 0.95, so this panel does not distinguish the influence pattern from zero.

If an outcome-scale magnitude is useful, add the defined simulation-based first difference and its risk set. Do not substitute that number for convergence, goodness-of-fit, uncertainty, or the remaining confounding assumptions.

Transition

Close by naming the five ideas that should transfer to the causal network material on Day 13.

SLIDE 33: What to Carry Into Day 13

What to say

A SAOM fills intervals between panel waves with actor-oriented micro-steps. Rate and evaluation parameters describe opportunities and preferences, so they are not interchangeable. Co-evolution represents selection and influence jointly but does not identify them causally without strong assumptions.

Convergence and goodness-of-fit remain separate requirements. SAOM and TERGM encode different temporal questions, and probability-scale summaries must be defined through a coherent risk set and contrast rather than an inverse-logit transformation of a SAOM coefficient.

The bridge to Day 13 is a change in inferential focus. Network change models describe an evolving relation. Causal network analysis asks how treatments and outcomes interact through a network, how interference changes the estimand, and how dyadic dependence changes uncertainty.

Closing line

Tomorrow we separate two problems that are often blurred together: interference changes the estimand, while dyadic dependence changes uncertainty.