

Day 13 Live Presenter Script

Shahryar Minhas

Day 13 Live Presenter Script

Why Causal Claims Are Hard in Networks

Pacing Guide

Slides 1 to 5, opening and causal foundations: minutes 0 to 15.

Slides 6 to 9, the household voting study: minutes 15 to 38.

Slides 10 to 14, exposure, support, and weighting: minutes 38 to 80.

Break: minutes 80 to 90.

Slides 15 to 20, the middle-school program: minutes 90 to 120.

Slides 21 to 25, election observers in Ghana: minutes 120 to 145.

Slides 26 to 32, friends' GPA, negative controls, and latent variables: minutes 145 to 170.

Slides 33 to 36, synthesis and closing: minutes 170 to 180.

Slide 1: Day 13: Why Causal Claims Are Hard in Networks

Today we are changing the kind of question we ask. Earlier in the course, we learned how to describe ties, find groups, represent hidden social structure, and model network change. Today the question is: what would have happened if an intervention had been different?

That question becomes especially hard in a network. An action aimed at one person may affect a friend, a household member, or a nearby community. At the same time, connected people often looked alike before any influence occurred.

We will use four published studies rather than begin with equations. One study sends voting messages to households. One works with an anti-conflict program in middle schools. One places election observers at voter-registration centers in Ghana. One asks whether friends' grades affect a student's later grades.

The goal is not to memorize four methods. The goal is to learn how to slow down a causal sentence and ask what the evidence truly supports.

By the end, a sentence such as “friends improve grades” should immediately make you ask: What changed, who could be affected, what is the comparison, and why should we trust it?

Let us begin with those four questions.

Slide 2: What Are We Trying to Learn?

Causal inference begins with a simple idea: what would change if we took one action rather than another?

Suppose a campaign sends a voting message. We want to compare turnout after that message with turnout under a meaningful alternative, such as a recycling message. The difference is the effect of changing the message.

For every case today, we will ask four questions. What action changed? Whose outcome could change? What two situations are being compared? Why should we trust that comparison?

Notice that the statistical method is not the first question. A sophisticated model cannot help if we are vague about the action or if we compare people whose experiences were fundamentally different.

Take thirty seconds and think of a causal claim you have heard recently. It might concern social media, education, public health, policing, or elections. What was the action, and what was the outcome?

Keep that example in mind as we define the comparison more carefully.

Slide 3: A Causal Effect Compares Two Possible Outcomes

Imagine one particular voter. In one possible world, that voter hears a message encouraging turnout. In another possible world, the same voter hears a recycling message.

The causal effect for that voter is the difference between what the voter would do in those two worlds. The difficulty is that we observe only one world. Once the voter hears one message, we cannot rewind time and give the same voter the other message under exactly the same conditions.

This missing second outcome is why causal inference needs a comparison group. We use other voters to estimate what would probably have happened under the message that our focal voter did not receive.

Random assignment is valuable because, before the messages are delivered, the two groups should be similar on average. The comparison message group can then stand in for the missing possible outcome of the voting-message group.

Randomization does not let us see both worlds for one person. It makes the groups credible substitutes for one another.

Now we can see what a network adds to this already difficult problem.

Slide 4: Why Networks Make This Harder

Without spillovers, we might describe my experience using only my own assignment. I received the program, or I did not.

In a network, a second route appears. My outcome may change because someone connected to me received the program. A friend may repeat a message, a household member may discuss it, or a nearby organization may change its be-

havior.

This means that two people labeled “untreated” may have had very different experiences. One untreated student may have several friends in a program. Another may know no one in it.

The technical word for one unit’s assignment affecting another unit’s outcome is interference. The word does not tell us how the effect travels. It could travel through conversation, observation, competition, movement, or shared resources.

Networks also create a second difficulty. People choose friends, households share circumstances, and nearby places may share politics or markets. Similar outcomes among connected units can therefore reflect influence, prior similarity, or both.

Turn to someone near you and name one setting in which an intervention could reach an unassigned person. What is the connection, and what might travel along it?

The four studies give us four concrete versions of this problem.

Slide 5: Four Applications, Four Versions of the Problem

The first case is a door-to-door voting experiment in two-voter households. It gives us a very small network: the person who answers the door and the other registered voter in the home.

The second case is a student-led anti-conflict program in

New Jersey middle schools. Some students were invited to help lead the program, but their friends and the rest of the school could also see its activities.

The third case concerns voter registration before Ghana's 2008 national election. Observers were stationed at some local registration centers. The concern was that monitoring might stop questionable activity at one center but push it to another center nearby.

The fourth case uses a large US adolescent friendship study. It asks whether friends' earlier grades affect a student's later grades when friendship itself was not randomly assigned.

These studies measure different outcomes and use different designs. We should not rank them by comparing the sizes of their estimates. We should compare the logic of their claims.

As we move through them, keep returning to the same four questions: what changed, who could respond, what is compared, and why is the comparison credible?

We begin with the smallest network, a household with two registered voters.

Slide 6: Case 1: A Voting Message Enters a Household

Nickerson studied door-to-door campaigning during the 2002 congressional primary elections in Denver and Minneapolis. The study focused on homes with two registered voters.

A canvasser knocked on the door and spoke with the person

who answered. Households were assigned to receive either a message encouraging voting or a message encouraging recycling. The recycling message created a useful comparison because both groups experienced a canvassing contact, but the content of the message differed.

The researchers later used administrative voting records. They did not have to rely on people remembering or admitting whether they voted. They could see whether the person who heard the script voted and whether the other registered voter in the household voted.

Why care about the second voter? Campaigns rarely reach people in isolation. Household members talk, share plans, remind one another, and often travel to vote together. A message delivered to one person may therefore reach two voters.

The action was the assigned script. The two outcomes were turnout for the person at the door and turnout for the other voter.

That one action creates two distinct comparisons.

Slide 7: The Household Creates Two Comparisons

The first comparison is direct. Among people who answered the door, was turnout higher when the canvasser delivered the voting message rather than the recycling message?

The second comparison is the network comparison. Among the other registered voters, was turnout higher when the person at their door heard the voting message rather than

the recycling message?

The second voter never heard the canvasser directly. A difference for that voter is therefore evidence that the assigned message had an effect beyond the person who received it.

Conversation is one possible pathway, but it is not the only one. The person at the door might place a reminder on the table, make a voting plan, or simply change the household's schedule.

Here is the important boundary. The experiment changed the message. It did not randomly force the first person to vote. A partner effect is therefore a spillover from the message, not proof that one person's act of voting caused the other person's act of voting.

In pairs, finish this sentence: "The randomized action was blank, and the partner outcome was blank."

With the action and outcomes clear, the result is easy to read.

Slide 8: Turnout Was Higher for Both People

The horizontal axis reports percentage-point differences from the recycling-message group. A value of ten percentage points would mean that turnout was ten points higher, such as 40 percent rather than 30 percent.

For the people who heard the voting message, turnout was about 9.7 percentage points higher than for people who heard the recycling message.

For the other registered voters in those households, turnout

was about 5.9 percentage points higher. The second estimate is smaller, but it still points to a meaningful within-household spillover from the assigned message.

The dots are the best estimates from these data. The horizontal lines show uncertainty. They remind us that a sample estimate is not an exact population value.

The partner increase is roughly three fifths as large as the increase for the person who heard the message. That is a comparison between two estimated increases. It does not mean that 60 percent of partners were persuaded, and it does not mean turnout rose by 60 points.

If we looked only at the person who answered the door, we would miss part of what the campaign changed.

Before moving on, we need to keep the causal statement at the level of the actual design.

Slide 9: What the Household Study Teaches Us

The careful conclusion is that assignment to the voting message changed turnout for both people in successfully contacted two-voter households, compared with assignment to the recycling message.

The phrase “successfully contacted” matters. Households were assigned before the canvasser knew whether anyone would answer, but this analysis focuses on homes where contact occurred. The causal reading therefore also needs the contacted voting-message and recycling-message households to remain comparable.

The study does not show that every attempted household would respond in the same way. It also does not reveal the exact household conversation or action that carried the message.

Most importantly, it does not identify the effect of the first person's actual vote on the partner's vote. The first person's vote was an outcome that happened after the message, not the randomized intervention.

Take one minute and rewrite this overclaim: "Voting is contagious because one partner's vote made the other partner vote." A better sentence should name the voting message, the partner's turnout, and the contacted households.

The household was simple because each person had one obvious connection. A larger network requires a way to describe many possible forms of exposure.

Slide 10: "Untreated" Can Hide Different Experiences

Imagine that a school invites some students to join a program. Alex is not invited and has no invited friend. Ben is not invited but does have an invited friend.

If we label both Alex and Ben "untreated," we erase the very network pathway we want to study. Ben may hear about the program, see its symbols, or join activities through a friend. Alex may have none of those experiences.

Cara is invited but has no invited friend. Dani is invited and also has an invited friend. Their direct assignment is the same, but their social environments differ.

These four students show why a single treatment label may be too crude. We may need to record both a person's own assignment and what happened to connected people.

Take a minute to choose one outcome, such as wearing a program wristband, reporting less conflict, or changing an opinion. Which of the four students would you expect to differ, and why?

To compare these experiences systematically, researchers define exposure conditions.

Slide 11: An Exposure Condition Is Just a Useful Description

An exposure condition is a compact description of the parts of an assignment pattern that might matter for one person's outcome.

For a student, it might record whether the student was invited and whether at least one friend was invited. For a neighborhood, it might record whether an observer was assigned there and whether an observer was assigned nearby.

Why do we need this shortcut? In a large network, every person's outcome could in principle depend on the assignment of everyone else. There are far too many complete assignment patterns to compare one by one.

The exposure condition compresses that enormous pattern into a few situations that have a substantive meaning. It turns "everything that happened in the network" into a comparison we can actually describe.

This compression is also an assumption. If we group one

invited friend and six invited friends together, we are saying that the difference between them is not central to the causal question.

Random assignment can tell us how likely each declared exposure was. It cannot tell us whether we declared the right exposure.

The next slide shows how different definitions can produce different questions.

Slide 12: The Exposure Definition Can Change the Answer

Suppose a student has invited friends. We could use a simple yes-or-no definition: no invited friends versus at least one invited friend.

We could instead count the number of invited friends. One, two, and six would then be different levels of exposure.

We could also focus only on close friends. A message from one trusted friend may matter more than messages from several weak acquaintances.

None of these definitions is always correct. The right choice depends on how the program could travel. If one public invitation is enough to make the program visible, “at least one” may be sensible. If repeated encouragement matters, a count or share may fit better.

Measurement matters too. A friendship survey may miss ties, treat a one-way nomination as mutual, or fail to distinguish a close friendship from a casual one.

In small groups, imagine a vaccination campaign, a study-

skills program, or a protest message. Would you use any treated neighbor, the number treated, the share treated, or only strong ties? Give one sentence explaining the mechanism.

Even a sensible exposure definition is useful only when the design makes the desired comparison possible.

Slide 13: Some Comparisons Are Not Available

Before comparing two exposure conditions, ask a practical question: could the people in this analysis realistically have experienced either condition under the study design?

A student with no recorded friends can never have an invited friend. That student cannot tell us what would happen under peer exposure.

A student who was not eligible to receive an invitation can never experience direct assignment. That student cannot contribute to a direct invitation comparison.

Some school-level designs also make combinations impossible. A student in a control school cannot have a program invitation if invitations occur only in program schools.

This requirement is often called positivity or support. The name is less important than the idea: both sides of a causal comparison must be genuinely available to the people the claim describes.

No statistical model can create an experience that the design made impossible. The honest response is often to narrow the population, such as focusing on students who were eligible

for all relevant conditions.

Look back at Alex, Ben, Cara, and Dani. Which comparisons could an isolated student never enter? Which comparisons could an ineligible student never enter?

When conditions are possible but occur with unequal chances, weighting can help.

Slide 14: Why Weighting Sometimes Appears

Consider two students. One has six friends who were eligible for invitation. The other has only one eligible friend. Even under fair random assignment, the first student has more ways to end up with at least one invited friend.

If we simply average everyone observed with an invited friend, highly connected students may dominate that group. The exposed group may then describe the network opportunities of those students more than the population we intended to study.

Design-based weighting uses the known randomization to correct that imbalance. An observation that was very likely to appear in its condition receives less weight. An observation that was unlikely to appear there receives more weight. The intuition is similar to giving a rare experience more voice so that the final comparison represents the target group rather than only the people who most easily entered each condition.

Weighting does not repair a bad exposure definition, and it cannot help someone whose probability of reaching a con-

dition is zero. It does one narrower job: it accounts for unequal chances created by the design.

The exact formula is optional today. The question to remember is: did some people have many more ways than others to reach this exposure?

We now have the tools needed for a full school network. Let us take a ten-minute break, then apply them to a student-led anti-conflict program.

Break: Ten Minutes

Slide 15: Case 2: An Anti-Conflict Program in Middle Schools

Paluck, Shepherd, and Aronow studied 56 public middle schools in New Jersey during the 2012 to 2013 school year. The larger project involved more than 24,000 students from ages 11 to 15.

The Roots program was designed to change how students handled everyday peer conflict and how acceptable constructive responses seemed within the school. Half of the schools were assigned to host the program.

Within program schools, some eligible students were randomly invited to become seed students. These students met regularly, identified common sources of conflict in their own schools, and designed slogans, posters, public events, and other activities.

The program was student-led for a reason. A message about peer behavior may carry more weight when visible students

make it part of ordinary school life than when adults simply deliver another rule.

The broader study found fewer disciplinary incidents involving peer conflict in program schools. Our narrower network question is how signs of the program reached students through friendship ties.

To answer that question, the analysis uses an outcome that may look unusual at first: an orange wristband.

Slide 16: What the Orange Wristbands Meant

The wristbands were orange and carried the Roots tree logo. Seed students handed them out when they saw friendly or conflict-reducing behavior, and students could wear them visibly at school.

For the program, the wristband was a small public signal. Wearing one could communicate support for the anti-conflict message and make that message visible to classmates.

For the researchers, reported wristband wearing offered a concrete measure of program reach. They could ask whether the symbol appeared only among invited seed students or also among their friends and other students.

The wristband is not a direct measure of reduced conflict. A student might wear one without changing behavior. Another student might behave more constructively without wearing one.

That distinction is why the wristband matters. It helps answer a diffusion question: did the campaign's visible symbol

travel beyond the students who received invitations?

Think of another campaign symbol, such as a ribbon, badge, profile frame, sticker, or yard sign. What can observing the symbol tell us, and what can it not tell us?

Now we need to separate receiving an invitation from choosing to participate.

Slide 17: Assignment and Participation Are Different

What was randomized was whether an eligible student received an invitation to join Roots. Participation happened later and involved a choice.

About 24 percent of invited seed students did not accept the invitation. That is not a small detail. Students who accept may be more motivated, more available, or already more supportive of the program than students who decline.

If we compare participants with nonparticipants, those prior differences may become mixed with the program effect. Random invitation gives us a cleaner comparison because invitation did not depend on a student's later decision.

The results therefore describe invitation assignment and exposure to friends' invitation assignments. They do not automatically describe the effect of actual participation.

This is often called an intention-to-treat interpretation, but the plain-language version is enough: compare people according to what the randomized design offered them, not according to the choice they made afterward.

Complete this sentence together: "The study randomized

blank, while blank was a later choice.”

Combining school assignment, student invitation, and friendship produces several distinct experiences.

Slide 18: The School Design Creates Several Experiences

The first experience is no program exposure. The student is in a control school, is not invited, and has no invited friend.

The second is program-school context only. The student attends a Roots school but is not invited and has no invited friend. Posters, events, or a changing school climate could still reach this student.

The third is friend assignment. The student is not invited, but at least one friend is invited. This is the clearest place to look for reach through the friendship network.

The fourth is direct assignment without an invited friend. The fifth combines the student’s own invitation with at least one invited friend.

The friendship network was measured before the program. That timing matters because the intervention itself could have changed whom students spent time with or nominated as friends.

The table is not a ladder from “no dose” to “more dose.” It lists qualitatively different situations created by school assignment, student assignment, and prior friendship ties.

Choose two rows and state exactly what changes between them. Then state what remains fixed. This habit prevents us from attaching a simple label to a comparison that

changes several things at once.

With those experiences defined, we can read where the wristband appeared.

Slide 19: The Program Reached Students Through More Than One Route

Every point on this figure compares an exposure condition with eligible students in control schools who had no program exposure.

Students who were invited themselves were about 30 percentage points more likely to report wearing the wristband. This is the largest and clearest difference, which makes sense because invitations connected students directly to the program.

Students who were not invited but had at least one invited friend were 15.4 percentage points more likely to report wearing the wristband than the control-school reference group.

The program-school-only estimate was 5.7 percentage points, but its uncertainty interval includes zero. The data do not give a precise positive difference for students with only the broader school context.

The 15.4-point comparison changes two features together: the student moves from a control school to a program school and also has an invited friend. It is therefore safest to say that the program's visible reach extended beyond directly invited students, with friendship assignment helping define that exposure.

The dots summarize the estimated differences. The lines

remind us that the direct and peer-exposed groups have different levels of precision.

If the only people wearing wristbands had been invited seed students, what would that have suggested about the campaign? What does the uninvited friend group add to the story?

The result is meaningful only if we keep the outcome and exposure definition in view.

Slide 20: What the School Study Does and Does Not Tell Us

The network-based conclusion is that assignment to Roots reached students who were not invited themselves but had an invited friend.

Here, “reached” means a higher probability of reporting wristband wearing. It does not by itself mean that those students had fewer conflicts, treated classmates better, or changed the whole school climate.

The exposure definition also groups together students with one invited friend and students with several invited friends. It does not distinguish a best friend from a casual peer or a highly visible student from a less visible one.

Finally, the comparison concerns randomized invitations. It is not automatically the effect of choosing to attend the seed meetings.

These limits do not make the finding unimportant. They tell us what the finding is about: the movement of a visible program symbol across assignment and friendship conditions.

Take two minutes and write one careful headline for the result and one headline that would overclaim. A careful headline should include the invitation, the friendship route, and wristband wearing.

So far, spillovers have involved messages and visible behavior. The next case shows that interference can also arise because strategic actors move activity to another place.

Slide 21: Case 3: Election Monitoring May Move Behavior Nearby

Ichino and Schündeln studied Ghana's 2008 voter-registration period before a closely contested national election. This was not a separate local election. It was local registration for a national vote.

Registration lists matter because they determine who is eligible to cast a ballot. If people are registered improperly or in the wrong place, the problem can affect representation and confidence in the election.

Nearly two million new names were approaching the rolls, while roughly 800,000 people were expected to have become newly eligible simply by age. Some of that gap could have innocent explanations, but it intensified accusations that parties were moving supporters or registering people improperly.

The process happened at local registration centers. That local scale matters because a national voter roll is built from many local decisions, and monitoring resources are limited.

Imagine that an observer reduces unusual registration at one

center, but organizers send people to an unmonitored center five kilometers away. The monitored center looks better, yet the broader problem may simply have changed location.

That is why we care about outcomes both where observers were sent and at nearby places where they were not.

The study used the placement of domestic election observers to examine exactly that possibility.

Slide 22: How the Observer Study Worked

Ghana's Coalition of Domestic Election Observers, known as CODEO, sent observers to registration centers during the 13-day registration period.

The observers were meant to make the process visible and discourage questionable practices where they were stationed. Their presence could raise the chance that irregular behavior was noticed or challenged.

The assignment occurred in stages. Some constituencies entered the observer program, and within those constituencies roughly one quarter of electoral areas were assigned an observer. This created monitored sites, nearby unmonitored sites, and more distant sites.

The researchers then examined growth in the number of registered voters from 2004 to 2008. They asked whether growth changed at the focal center and whether it changed when observer-assigned centers were nearby.

Randomization helps because political organizers did not simply choose the hardest or easiest sites for observers. But the network exposure still needs to include geography be-

cause people and strategies can move across nearby locations.

Think of the observer as changing the cost of questionable activity at one point on a map. What might a strategic organizer do next?

That response produces spillover without requiring any friendship influence.

Slide 23: This Spillover Does Not Require Social Influence

The first possible response is deterrence. Registration growth may fall at a center where an observer is present because questionable activity becomes riskier.

The second possible response is displacement. Registration growth may rise at a nearby unmonitored center if activity shifts away from the observer.

No one has to copy a friend's behavior for this spillover to occur. The relevant connection is geographic and strategic. One location becomes less attractive, so another becomes more attractive.

The same logic appears in many policies. Police patrols may move crime to a nearby block. Environmental rules at one plant may move production elsewhere. Closing one clinic may increase demand at another clinic. Monitoring one border crossing may redirect traffic.

This broadens our definition of a network. Units are connected whenever an action at one unit can change opportunities or behavior at another.

In small groups, name one policy that might displace an outcome. What are the connected units, and where would the outcome move?

The Ghana estimates show why measuring nearby outcomes changes the policy story.

Slide 24: Registration Growth Fell Locally and Rose Nearby

At a registration center assigned an observer, fitted registration growth was about 3.5 percentage points lower than at the relevant comparison center.

For an unmonitored center, one additional observer-assigned center within five kilometers was associated with about 2.7 percentage points more registration growth, holding the other parts of the fitted comparison fixed.

The signs tell a coherent story. Growth went down where monitoring occurred and went up at nearby places without an observer.

If we examined only monitored centers, we might conclude that observers simply reduced the outcome. Looking nearby reveals that part of the response may have shifted across space.

The nearby estimate does not tell us that exactly the same people moved from one center to another. It shows a change in the distribution of registration growth that is consistent with displacement.

Notice again how the exposure definition matters. “Within five kilometers” is the study’s way of describing a nearby

treatment environment. A different distance could produce a different question.

Now we must separate the measured outcome from the larger political interpretation.

Slide 25: The Outcome Is Registration Growth

The study measured growth in registered voters from 2004 to 2008. It did not observe a label on each registration saying “valid” or “fraudulent.”

Registration can grow for legitimate reasons, including population movement, delayed registration, or newly eligible citizens. Unusually high growth may be evidence relevant to irregularities, but it is not identical to directly observed fraud.

The randomized observer assignment gives strong evidence that monitoring changed registration growth in the studied areas under the proposed geographic exposure model.

The local decline and nearby increase are consistent with deterrence and displacement. They do not reveal who moved, which registrations were improper, or whether the pattern would look the same across all of Ghana.

A careful conclusion is therefore narrower than “observers stopped fraud” or “parties moved fraudulent voters.” It says that observer assignment reduced registration growth locally and coincided with greater growth nearby.

Rewrite this sentence in careful language: “Observers proved that fraud moved down the road.” Keep the

observer assignment, the registration-growth outcome, and the nearby pattern.

The first three cases began with randomized actions. Our final application asks what can be learned when friends and their grades were never randomized.

Slide 26: Case 4: Friends' GPA and Later Academic Performance

Egami and Tchetgen Tchetgen study more than 10,000 adolescents in Add Health, a large US school-based study from 1994 and 1995.

Students were in grades 7 through 12. They could name up to five male and five female friends, which gave researchers a friendship network. The study also recorded grades in English, mathematics, history or social studies, and science, then summarized them as GPA.

The question is easy to understand: would a student's later GPA change if the student's friends had a higher average GPA at baseline?

We care because the answer changes how we think about educational policy. If friends causally shape performance, study groups, mentoring, classroom composition, or peer-based programs might generate benefits beyond the students directly served.

But schools do not randomly hand out friends. Students choose one another, share classes and neighborhoods, and may become friends because they already have similar interests, resources, or goals.

The outcome is later GPA, the proposed exposure is friends' earlier average GPA, and the central problem is hidden selection into friendship.

Two students with similar grades can therefore fit two very different stories.

Slide 27: Similar Friends Do Not Automatically Show Influence

Suppose Maya and Jordan are friends and both earn high grades. One story is influence: Jordan's study habits, advice, or expectations help Maya improve.

Another story is selection: Maya and Jordan became friends because both were already academically motivated. Their similar grades would then appear even if neither changed the other.

Shared context can produce the same pattern. They may take the same advanced class, live in similar households, have access to the same teachers, or receive similar family support.

A regression can adjust for every measured variable in the data and still miss a common cause that was not measured well. Motivation, course placement, family resources, and academic identity are difficult to capture completely.

This problem is called hidden homophily. An unmeasured factor helps create the friendship and also helps explain the outcome.

Look at the two cards and ask what data pattern would distinguish the stories. The uncomfortable answer is that

similar grades among friends alone cannot do it.

The researchers use negative controls as clues about the hidden selection process.

Slide 28: Negative Controls Look for Clues About Hidden Selection

The word “negative” does not mean that the estimated effect should be below zero. It means the control variable is chosen because it should not lie on the causal pathway being studied, under the proposed assumptions.

One clue comes from the student’s own baseline GPA. It reflects earlier academic orientation and other hidden forces that may also shape later GPA. Under the timing argument, friends’ baseline GPA should not reach backward and cause this earlier measure.

Another clue comes from the network side. Peers-of-peers’ baseline GPA may reflect the same school and friendship-selection process that shaped friends’ GPA, while being assumed not to directly change the focal student’s later GPA after the stated adjustments.

The researchers also examined peers’ headache levels as another network-side clue. It may carry information about shared selection or context, but its usefulness depends on it being informative enough and not directly affecting the student’s GPA through another route.

The basic idea is that hidden selection may leave fingerprints in several observed variables. If those fingerprints satisfy the required restrictions, they can help correct the friends’ GPA

comparison.

These restrictions are not facts that the method discovers. Researchers must defend why the controls are related to hidden selection and why they are not themselves part of the peer-effect pathway.

With that caution in place, compare the ordinary regression with the negative-control results.

Slide 29: The Peer-GPA Estimate Became Much Smaller

The conventional adjusted regression estimated that a one-point increase in friends' average baseline GPA was associated with a 0.176-point increase in the student's later GPA.

On a four-point GPA scale, a full one-point difference among friends is substantial. The conventional estimate therefore describes a noticeable academic difference and has a narrow uncertainty interval.

Using peers-of-peers' GPA as the network-side negative control reduced the estimate to 0.033. Using peers' headaches reduced it to 0.078.

Both negative-control intervals include zero and are much wider. The headache-based result is especially uncertain because that clue was not very informative about the relevant hidden process in these data.

The key visual pattern is not merely that the dots moved. A large, precise association became smaller and less precise when the analysis took hidden friendship selection more seriously.

This does not prove the corrected value is exactly zero. It shows that the original 0.176 estimate was not stable across plausible approaches to hidden confounding.

That difference changes how carefully we should describe the evidence.

Slide 30: How to Read the Peer-GPA Result

The study does not prove that friends have no influence. Students clearly can study together, exchange information, and shape expectations.

The result says something more specific. An ordinary adjusted regression produced a large positive association, but two proposed negative-control strategies produced much smaller and more uncertain estimates.

The causal conclusion therefore depends heavily on whether the negative controls truly capture useful traces of hidden selection and whether their exclusion assumptions are credible.

Wider uncertainty is not a methodological failure. It is an honest sign that the observational data contain less causal information than the conventional regression appeared to provide.

The main lesson is sensitivity. A persuasive-looking peer association can change substantially when we address the reasons similar students may have become friends in the first place.

Take two minutes and compare these two headlines: “High-achieving friends raise GPA” and “The estimated peer ef-

fect shrank after adjustment for hidden friendship selection.” Which one matches the evidence, and what uncertainty does it preserve?

Negative controls use selected observed clues. A different strategy asks whether the network itself contains clues about the missing variables.

Slide 31: Can the Network Help Reveal Hidden Selection?

This question connects today’s causal problem to the latent-variable models we studied earlier.

Suppose academic orientation was not measured well. If academically oriented students tend to form ties with one another, that hidden trait may leave a footprint in the larger pattern of friendships.

A latent-space model might place students with similar tie patterns near one another in an inferred social space. A blockmodel might place them in a shared group. A latent-factor model might represent several hidden dimensions that help explain who connects to whom.

Researchers can use those estimated positions or groups as proxies for unobserved social structure, then compare students who are similar in observed characteristics and in the inferred network structure.

McFowland and Shalizi have used latent community or latent-space information in this way to adjust peer-influence estimates under explicit assumptions. So this is not only a hypothetical connection. It is an existing methodological

strategy for latent homophily.

The appeal is intuitive: the network may contain fingerprints of variables the survey failed to measure. A student's position in the full network can reveal more than a short list of observed covariates.

But an inferred position is a model-based clue, not a direct observation of ambition, family resources, or academic orientation.

That is why latent adjustment can help without automatically solving the causal problem.

Slide 32: Latent Adjustment Can Help, but It Is Not Magic

Latent adjustment is most promising when the hidden cause of friendship creates a strong and recoverable pattern in the network.

The latent model must recover the part of that pattern that also affects the outcome. A beautifully fitted social position is not useful for causal adjustment if it captures popularity or community structure but misses the hidden academic factor causing the bias.

The outcome analysis must then use the recovered structure appropriately. Researchers also need to carry uncertainty about the estimated latent positions into the final causal estimate.

A latent position is not a measured confounder. We should not rename one dimension “ambition” simply because it predicts ties. It remains an inferred summary whose meaning

depends on the model and the network data.

This creates an important distinction between prediction and identification. Adding latent positions may improve predictions of friendship or GPA while leaving the causal comparison biased.

Negative controls and latent adjustment follow a related intuition: hidden selection may leave observable traces. They use different traces and require different assumptions, but neither creates random friendship assignment.

In one sentence, explain the boundary: “Latent network adjustment can help if blank, but it does not prove blank.”

We can now collect the four difficulties that appeared across the applications.

Slide 33: Four Difficulties Keep Returning

The first difficulty is spillovers. Can another person’s or place’s assignment affect my outcome? We saw this inside households, through school friendships, and across nearby registration centers.

The second is exposure definition. What exactly counts as being reached? At least one invited friend, the number of invited friends, or an observer within five kilometers are different definitions.

The third is missing comparisons. Could the people being described realistically have experienced both situations? An isolated student cannot have an invited friend, and an ineligible student cannot receive an invitation.

The fourth is hidden selection. Did connected people already

resemble one another before any influence occurred? This was the central threat in the friendship and GPA study.

Different studies use different tools because they face different combinations of these difficulties. There is no single network causal model that solves them all.

Match each case to its most visible difficulty, then name a second difficulty that also matters. The household case is not only about spillover because successful contact matters. The school case is not only about exposure because support and outcomes matter too.

The next distinction helps us see what each tool can and cannot do.

Slide 34: Design, Modeling, and Uncertainty Do Different Jobs

Research design explains why the groups being compared are credible stand-ins for one another. Random assignment played this role in the household, school, and Ghana studies.

A statistical model summarizes relationships, adjusts for measured differences, or represents a proposed process. Exposure mappings, regression adjustments, negative-control bridges, and latent positions all belong somewhere in this modeling work.

An uncertainty calculation tells us how much the estimate might vary. In network data, connected observations may move together, so ordinary uncertainty formulas may be too optimistic.

These jobs are related but not interchangeable. A flexible

model cannot quietly replace a missing comparison. A wider or more appropriate standard error cannot remove hidden confounding. Randomization does not rescue an exposure definition that ignores a real spillover route.

When reading a paper, ask three separate questions. What made the comparison credible? What did the model add? What sources of variation did the uncertainty calculation allow?

Take one minute and classify each item: random invitation, latent position, and a 95 percent interval. Which job is each item doing?

The safest final claim should reflect the scale of the design and the outcome actually measured.

Slide 35: Read the Claim at the Scale of the Evidence

For the household study, the careful conclusion is that the voting message increased turnout for the contacted voter and for the other registered voter in the household comparison. The claim is about message assignment, not one person's vote causing another person's vote.

For the school study, invitation assignments reached uninvited students who had invited friends. The network outcome shown here is reported wristband wearing, not conflict reduction itself.

For the Ghana study, monitoring reduced registration growth locally and was associated with greater growth nearby. The outcome is registration growth, not a direct

count of fraudulent registrations.

For the GPA study, the large conventional association became much smaller under two proposed corrections for hidden selection. That is not proof of no peer influence, but it weakens a confident causal reading of the ordinary regression.

Every careful conclusion names three things: the action or exposure, the measured outcome, and the group or setting supported by the evidence.

Here is a final practice. Choose one row and make the claim too large on purpose. Then repair it by returning to what was assigned, what was measured, and who was studied.

We can close with four questions that travel well beyond these particular applications.

Slide 36: Four Questions to Take With You

First, what exactly changed? Name the voting message, the invitation, the observer assignment, or the friends' GPA exposure rather than using the vague word "treatment."

Second, who could be affected, including through connections? Look beyond the person or place assigned directly.

Third, what two situations are being compared? State both sides clearly and notice what else changes between them.

Fourth, what assumptions make that comparison credible? Look for randomization, support, an exposure definition, timing, and defenses against hidden selection.

If you can answer these four questions, you can read a net-

work causal study intelligently even when you would not reproduce its estimator yourself.

If one answer is vague, slow down before accepting the causal language. The technical method comes after the intervention, the possible routes of exposure, and the comparison are clear.

Networks make causal claims harder because outcomes can travel and because ties are not random. They also make the questions more realistic, because interventions rarely stay inside the boundary where they were first delivered.

The final takeaway is simple: follow the action through the network, compare only experiences the design can support, and keep the conclusion at the scale of the evidence.