

Path Analysis and Mediation

PSY3100 – From Single Equations to Systems of Equations

Frederick Anyan

Associate Professor of Quantitative Research Methods

[NTNU](#) | [Personal website](#)

NTNU, Department of Psychology

Fall 2026

Today's plan

- Types of research questions, variables, relationships, and parameters
- What is a path model, and why do we need one?
- Mediation: simple, parallel, and serial
- A full path analysis model: two predictors, one mediator, one outcome
- Free, fixed, and constrained parameters
- Model fit and identification, introduced for the first time
- Estimating and interpreting path models in `lavaan`
- Common mistakes, and a bonus look at the `manymome` package

Types of Research Questions

Mediation asks: why, or how?

Does provocation increase aggression *because* it increases hostile thinking? Mediation asks about the **mechanism** that transmits an effect.

Moderation asks: for whom, or when?

Does the effect of provocation on aggression depend on some third variable – for instance, is it stronger for one group of people than another? Moderation asks about **boundary conditions**.

Path analysis asks: what is the whole system?

Path analysis considers several regression equations **simultaneously**, including cases where a variable is an outcome in one equation and a predictor in another.

Types of Variables and Relationships

Observed / manifest

A variable that is directly measured – for example, a provocation score from a questionnaire. Drawn as a **rectangle**.

Latent / factor

An unmeasured construct inferred from multiple observed indicators – for example, a “hostility” factor measured by several items. Drawn as an **ellipse** or **circle**. Today’s models use only observed variables; latent variables arrive with common factor models (EFA/CFA).




Exogenous

Not predicted by anything else in the model. No arrow points into it. Its variance is simply estimated from the data.

Endogenous

Predicted by at least one other variable in the model – has at least one arrow pointing into it. Carries a **disturbance** (residual), representing everything about it that is not explained by its predictors.

Third variables: mediator, moderator, confound



Direct, indirect, total, common cause, common outcome

- **Direct effect:** the association between a predictor and an outcome that is not transmitted through any other variable in the model
- **Indirect effect:** the association transmitted *through* one or more intervening variables – the product of the paths along the way
- **Total effect:** direct effect plus all indirect effects; equivalently, the effect of X on Y *ignoring* any mediators
- **Common cause:** two variables are related because a third variable causes both of them
- **Common outcome:** two variables are both caused by two other, potentially unrelated, causes (an “inverted fork”)

Types of Parameters

Free

A parameter estimated by the model to best reproduce the observed data – for example, every regression coefficient we have fit all semester is a free parameter.

Fixed

A parameter set to a specific constant by the researcher, not estimated at all – most commonly fixed to zero, meaning “this path does not exist.” Every path you *omit* from a diagram is implicitly fixed to zero.

Constrained

A parameter required to bear some relationship to another parameter, rather than fixed to a specific number. The most common case is an **equality constraint**: two paths are forced to take the same value, so only one is actually estimated.

From regression to path analysis

Briefly, multiple regression

$$Y = B_0 + B_1X_1 + B_2X_2 + e \quad (1)$$

B_0 is the intercept, B_1 and B_2 are unstandardized regression coefficients, and e is the residual. One outcome, one equation: every X is a predictor, Y is the only outcome, and that role never changes.

- every model this semester so far, including logistic regression, has had exactly this shape: one or more predictors, one outcome, one equation
- today asks what happens once that restriction is lifted

The one new idea for today


What if a variable is an outcome in one equation, and then itself a *predictor* in a second equation?

That is the entire conceptual jump from regression to path analysis – nothing else changes.

$$M = i_M + aX + e_M, \quad Y = i_Y + c'X + bM + e_Y$$

- M (hostile cognition) is predicted by X (provocation) in the first equation
- M then predicts Y (aggression) in the second equation
- this is the conceptual bridge from regression into path analysis: a path model is what you get when you let several regression equations like this share variables and estimate them together

Path diagram notation



- **rectangle** = observed variable
- **one-headed arrow** ($\rightarrow$) = a hypothesized directional (regression) relationship
- **two-headed arrow** ($\leftrightarrow$) = a covariance or correlation, no direction implied
- **e** (or a disturbance arrow with no source box) = residual, unexplained variance in an endogenous variable

Path analysis as a system of regression equations

The diagram on the previous slide, written as equations

$$M = i_M + aX + e_M \quad (2)$$

$$Y = i_Y + c'X + bM + e_Y \quad (3)$$

- M is an **outcome** in equation (2) and a **predictor** in equation (3)
- i_M, i_Y are intercepts; e_M, e_Y are residuals – exactly as in every regression this semester
- the two equations are fit *simultaneously*, not one after the other, so that every path is estimated net of the whole system

Path analysis is not causality

A path diagram can look causal. A path model does not, by itself, establish causality.

An arrow in a path diagram represents a *hypothesized* directional relationship – not proof of causation.

- cross-sectional data cannot establish temporal precedence
- confounding: an unmeasured common cause can produce exactly the same pattern of associations
- alternative models: provocation $\rightarrow$ hostile cognition $\rightarrow$ aggression is statistically indistinguishable, in cross-sectional data, from hostile cognition $\rightarrow$ provocation $\rightarrow$ aggression, or from a model with all three paths reversed
- experimental manipulation of X , with M measured afterward and Y measured after that, is what actually licenses causal-mechanism language – the statistical model alone does not

What this section covers

Does provocation increase aggression *because* it increases hostile cognition? We build the same worked example three ways – simple, parallel, and serial mediation – each answering a slightly different research question about the mechanism.

- Where mediation analysis came from, and why time order matters
- A first look at lavaan syntax
- Model 1: simple mediation (one mediator)
- Model 2: parallel mediation (two mediators, no ordering assumed between them)
- Model 3: serial mediation (two mediators, one feeding into the other)

Where mediation analysis came from

Hyman (1955): elaboration

Kenny (2008) credits Hyman with the earliest systematic discussion of what we now call mediation, under the name *elaboration* – explaining an association between two variables by introducing a third.

The sequential time-order criterion

Hyman's original formulation required a **time-ordered** relationship: the mediator must occur, causally and temporally, between the predictor and the outcome. Modern treatments (Baron & Kenny, 1986; Preacher & Hayes, 2004) often understate this requirement.

- Hyman–Tate conceptual time criterion: even without longitudinal data, the *theoretical* ordering of predictor, mediator, and outcome must be defensible
- Tate (2015) argues that a large share of the overuse and misuse of mediation analysis traces back to skipping this question

A first look at lavaan syntax: the general pattern

Three operators cover almost everything in this lecture

<code>Y ~ X</code>	<code># regression: "Y is predicted by X"</code>
<code>X ~~ W</code>	<code># covariance: "X and W covary, undirected"</code>
<code>a*X</code>	<code># labels the path a, so it can be reused below</code>
<code>indirect := a*b</code>	<code># a defined parameter: a function of labeled paths</code>

- every model fit in this lecture, from the simplest mediation triangle to the full multi-group analysis near the end, is built entirely out of these operators
- labeling a path (`a*X`) costs nothing and changes no estimate – it only gives that path a name we can refer to later
- `:=` is how indirect and total effects are obtained directly from the fitted model, with correct standard errors, rather than computed by hand afterward

A first look at lavaan syntax: a worked example

The classic three-variable mediation setup

```
model <- '  
  Y ~ c*X           # direct effect  
  M ~ a*X           # X predicts the mediator  
  Y ~ b*M           # M predicts the outcome  
  ab := a*b         # indirect effect, defined  
  total := c + (a*b) # total effect, defined  
,  
fit <- sem(model, data = Data)
```

- $\sim$ specifies a regression: $Y \sim X$ means “regress Y on X ”
- $a*X$ *labels* the path a – labels let us define new quantities from the paths
- $:=$ *defines* a new parameter as a function of labeled paths – this is how indirect and total effects are obtained directly from the fitted model, with correct standard errors

Getting bootstrap confidence intervals in R

Why bootstrap, and not the delta method


An indirect effect (ab) is a *product* of two estimates, so its sampling distribution is generally skewed – the symmetric CI a delta-method standard error would imply is not trustworthy here. Every indirect and total-indirect CI in this lecture is a nonparametric bootstrap CI instead.

The R syntax, once, for any model in this lecture

```
fit <- sem(model, data = Data, se = "bootstrap", bootstrap = 5000)
parameterEstimates(fit, boot.ci.type = "perc", level = .95)
```


- `se = "bootstrap"` tells lavaan to compute standard errors, and confidence intervals, from 5,000 nonparametric case-resampling replications instead of the delta method
- `boot.ci.type = "perc"` requests the percentile bootstrap CI – the 2.5th and 97.5th percentiles of the 5,000 resampled estimates – matching Preacher and Hayes's (2004) standard recommendation for indirect effects
- this single argument change is what produces every 95% CI reported for the rest of this lecture, for `:=` defined parameters exactly as for ordinary regression paths

The theory behind today's running example: the General Aggression Model



- **Inputs:** features of the *person* (stable traits, e.g. irritability) and of the *situation* (e.g. being provoked) that meet in a social encounter
- **Routes:** inputs affect behavior by changing the *present internal state* – hostile thoughts (cognition), hostile feelings (affect), and physiological arousal – which also influence one another
- **Outcomes:** the internal state shapes how the situation is appraised and which response is chosen, impulsive or considered; aggression is one possible response
- shaded boxes are the variables measured in our data ($N = 300$); each model today tests one specific piece of this theory

Model 1: does provocation increase aggression through hostile cognition?



Model 1: results

lavaan syntax

```
m1 <- 'hostile_cognition ~ a*provocation
      aggression ~ cp*provocation + b*hostile_cognition
      indirect := a*b      total := cp + (a*b)'
```

Path	<i>B</i> (unstd.)	<i>SE</i>	<i>z</i>	<i>p</i>	<i>β</i> (std.)
<i>a</i> : provocation → hostile cognition	0.416	0.043	9.58	< .001	0.501
<i>b</i> : hostile cognition → aggression	1.022	0.151	6.76	< .001	0.342
<i>c'</i> : provocation → aggression (direct)	0.793	0.124	6.40	< .001	0.319
<i>c</i> : provocation → aggression (total)	1.218	0.115	10.56	< .001	0.490

Indirect effect

$ab = 0.416 \times 1.022 = 0.425$, 95% bootstrap CI [0.274, 0.600] (5,000 resamples) – excludes 0, so the indirect effect is supported.

Results output: Model 1


Regressions:

	Estimate	Std.Err	z-value	P(> z)	Std.lv	Std.all
hostile_cognition ~						
provocatn (a)	0.416	0.043	9.577	0.000	0.416	0.501
aggression ~						
provocatn (cp)	0.793	0.124	6.400	0.000	0.793	0.319
hstl_cgnt (b)	1.022	0.151	6.763	0.000	1.022	0.342

Defined Parameters:

	Estimate	Std.Err	z-value	P(> z)	Std.lv	Std.all
indirect	0.425	0.082	5.196	0.000	0.425	0.171
total	1.218	0.115	10.561	0.000	1.218	0.490

Model 2: does provocation also work through negative affect?



Model 2: results

Path	<i>B</i> (unstd.)	<i>SE</i>	<i>z</i>	<i>p</i>	β (std.)
a_1 : provocation $\rightarrow$ hostile cognition	0.416	0.043	9.57	< .001	0.501
a_2 : provocation $\rightarrow$ negative affect	0.486	0.058	8.34	< .001	0.433
b_1 : hostile cognition $\rightarrow$ aggression	0.743	0.161	4.62	< .001	0.248
b_2 : negative affect $\rightarrow$ aggression	0.638	0.122	5.21	< .001	0.289
c' : provocation $\rightarrow$ aggression (direct)	0.599	0.129	4.65	< .001	0.241

Correlating the mediators' residuals – why?

$\text{cor}(e_{M_1}, e_{M_2}) = 0.31$. Hostile cognition and negative affect share unmeasured causes (e.g., general arousal) beyond provocation; leaving this covariance out would silently force it to zero and bias the other paths.

Two indirect effects, one total

- via hostile cognition: $a_1 b_1 = 0.309$ [95% CI 0.163, 0.471]
- via negative affect: $a_2 b_2 = 0.310$ [95% CI 0.176, 0.458]
- total indirect: $a_1 b_1 + a_2 b_2 = 0.619$ [95% CI 0.442, 0.811]
- total effect: $c' + \text{total indirect} = 0.599 + 0.619 = 1.218$, matching Model 1's total effect exactly

Results output: Model 2

Regressions:

	Estimate	Std.Err	z-value	P(> z)	Std.lv	Std.all
hostile_cognition ~						
provocatn (a1)	0.416	0.043	9.570	0.000	0.416	0.501
neg_affect ~						
provocatn (a2)	0.486	0.058	8.338	0.000	0.486	0.433
aggression ~						
provocatn (cp)	0.599	0.129	4.649	0.000	0.599	0.241
hstl_cgnt (b1)	0.743	0.161	4.620	0.000	0.743	0.248
neg_affct (b2)	0.638	0.122	5.212	0.000	0.638	0.289


Covariances:

	Estimate	Std.Err	z-value	P(> z)	Std.lv	Std.all
.hostile_cognition ~~						
.neg_affect	8.140	1.635	4.980	0.000	8.140	0.311

Defined Parameters:

	Estimate	Std.Err	z-value	P(> z)	Std.lv	Std.all
ind1	0.309	0.078	3.949	0.000	0.309	0.124
ind2	0.310	0.073	4.263	0.000	0.310	0.125
total_ind	0.619	0.095	6.500	0.000	0.619	0.249

Model 3: does hostile cognition itself intensify negative affect?



Model 3: results

Path	<i>B</i> (unstd.)	<i>SE</i>	<i>z</i>	<i>p</i>	β (std.)
a_1 : provocation $\rightarrow$ hostile cognition	0.416	0.043	9.57	< .001	0.501
a_2 : provocation $\rightarrow$ negative affect (direct)	0.304	0.061	5.01	< .001	0.271
d_1 : hostile cognition $\rightarrow$ negative affect	0.438	0.077	5.68	< .001	0.324
b_1 : hostile cognition $\rightarrow$ aggression	0.743	0.161	4.61	< .001	0.248
b_2 : negative affect $\rightarrow$ aggression	0.638	0.122	5.21	< .001	0.289
c' : provocation $\rightarrow$ aggression (direct)	0.599	0.129	4.64	< .001	0.241

Three indirect effects now, not two

via M_1 only: $a_1 b_1 = 0.309$ [0.163, 0.471]

via M_2 only: $a_2 b_2 = 0.194$ [0.095, 0.311]

serial, $M_1 \rightarrow M_2$: $a_1 d_1 b_2 = 0.116$ [0.057, 0.190]

total indirect: $a_1 b_1 + a_2 b_2 + a_1 d_1 b_2 = 0.619$ [0.441, 0.812] (4)

Results output: Model 3

Regressions:

	Estimate	Std.Err	z-value	P(> z)	Std.lv	Std.all
hostile_cognition ~						
provocatn (a1)	0.416	0.043	9.570	0.000	0.416	0.501
neg_affect ~						
provocatn (a2)	0.304	0.061	5.012	0.000	0.304	0.271
hstl_cgnt (d1)	0.438	0.077	5.684	0.000	0.438	0.324
aggression ~						
provocatn (cp)	0.599	0.129	4.643	0.000	0.599	0.241
hstl_cgnt (b1)	0.743	0.161	4.608	0.000	0.743	0.248
neg_affct (b2)	0.638	0.122	5.209	0.000	0.638	0.289

Defined Parameters:

	Estimate	Std.Err	z-value	P(> z)	Std.lv	Std.all
ind1	0.309	0.078	3.938	0.000	0.309	0.124
ind2	0.194	0.055	3.504	0.000	0.194	0.078
ind3	0.116	0.034	3.431	0.001	0.116	0.047
total_ind	0.619	0.095	6.487	0.000	0.619	0.249


A Full Path Analysis Model

What this section covers

Does a stable trait – irritability – also shape hostile cognition and aggression, alongside momentary provocation? This section builds a model with two exogenous inputs – one situational, one a stable trait – and then uses it to introduce model fit, identification, and constraint testing for the first time.

- Model 4: two exogenous predictors, one mediator, one outcome
- Model fit – an introduction
- Identification – an intuitive introduction
- Free, fixed, and constrained parameters, tested with χ^2 difference tests

Model 4: does trait irritability matter too?



Model 4: results

lavaan syntax

```
m4 <- 'hostile_cognition ~ a1*provocation + a2*trait_irritability
      aggression ~ c1*provocation + c2*trait_irritability + b1*hostile_cognition
      provocation ~~ trait_irritability
      ind_x1 := a1*b1      ind_x2 := a2*b1'
```

Path	<i>B</i> (unstd.)	<i>SE</i>	<i>z</i>	<i>p</i>	β (std.)
a_1 : provocation $\rightarrow$ hostile cognition	0.367	0.043	8.48	< .001	0.441
a_2 : trait irritability $\rightarrow$ hostile cognition	0.153	0.032	4.82	< .001	0.215
b_1 : hostile cognition $\rightarrow$ aggression	0.922	0.156	5.91	< .001	0.308
c_1 : provocation $\rightarrow$ aggression (direct)	0.748	0.125	5.99	< .001	0.301
c_2 : trait irritability $\rightarrow$ aggression (direct)	0.270	0.105	2.56	.010	0.126
cor(provocation, trait irritability) = 0.276, $p < .001$					

Two indirect effects, one for each predictor

- via provocation: $a_1 b_1 = 0.338$ [95% CI 0.204, 0.490]
- via trait irritability: $a_2 b_1 = 0.141$ [95% CI 0.074, 0.219]
- both exclude 0

Results output: Model 4

Regressions:

	Estimate	Std.Err	z-value	P(> z)	Std.lv	Std.all
hostile_cognition ~						
provocatn (a1)	0.367	0.043	8.478	0.000	0.367	0.441
trt_rrtbl (a2)	0.153	0.032	4.815	0.000	0.153	0.215
aggression ~						
provocatn (c1)	0.748	0.125	5.991	0.000	0.748	0.301
trt_rrtbl (c2)	0.270	0.105	2.562	0.010	0.270	0.126
hstl_cgnt (b1)	0.922	0.156	5.906	0.000	0.922	0.308

Covariances:

	Estimate	Std.Err	z-value	P(> z)	Std.lv	Std.all
provocation ~~						
trait_irrtblty	11.540	2.313	4.989	0.000	11.540	0.276

Defined Parameters:

	Estimate	Std.Err	z-value	P(> z)	Std.lv	Std.all
ind_x1	0.338	0.073	4.626	0.000	0.338	0.136
ind_x2	0.141	0.037	3.797	0.000	0.141	0.066

Two ways to ask how well a model fits

Beyond prediction

So far we have only asked how well a model *predicts* the outcome (R^2). A path model lets us additionally ask: how well does the proposed model reproduce *all* of the observed relationships among the variables, not just the ones pointing at Y ?

- χ^2 : tests whether the model-implied covariance matrix differs from the observed one (smaller, non-significant is better – unusual among test statistics)
- CFI, TLI: compare the fitted model to a “worst case” baseline model with no relationships at all (closer to 1 is better)
- RMSEA: a badness-of-fit measure that penalizes model complexity (closer to 0 is better)
- SRMR: the average discrepancy between observed and model-implied correlations (closer to 0 is better)

Putting fit indices together

Fit statistic	What does it tell us?	Interpretation
χ^2	Exact fit	Model–data discrepancy
RMSEA	Approximate fit	Error per degree of freedom
CFI	Comparative fit	Improvement over baseline
TLI	Comparative fit	Improvement accounting for complexity
SRMR	Residual fit	Average standardized residual

- Conventional guidelines (e.g. Hu & Bentler, 1999) are guidelines, not universal laws – they were derived from one specific simulation design
- Sample size, model complexity, misspecification type, and data characteristics all shift what a “good” value looks like
- Evaluate several indices together, alongside theory – never a single number in isolation

Common cutoff conventions

Index	Good fit	Acceptable fit
χ^2 test	$p > .05$	– (very sensitive to N ; rarely used alone)
CFI, TLI	$\geq .95$	$\geq .90$
RMSEA	$\leq .06$	$\leq .08$ (poor if $> .10$)
SRMR	$\leq .08$	–

These are conventions, not laws

Hu and Bentler's (1999) widely-cited guidelines were derived from a specific simulation study, not a universal statistical fact. Treat them as one input into a judgment, alongside sample size, model complexity, and substantive theory – never as a single automatic pass/fail rule. We return to this with more nuance in the model evaluation lecture.

Checking our own two models against these

Complete mediation ($c'_1 = c'_2 = 0$): CFI = .85, RMSEA = .25, SRMR = .088 – all three fail even the acceptable threshold. Equality constrained ($a_1 = a_2$): CFI = .96, RMSEA = .19, SRMR = .047 – CFI looks good, but RMSEA still signals poor fit and SRMR is borderline: the indices genuinely disagree here, partly because RMSEA is known to behave erratically with only 1 degree of freedom. Consistent with the χ^2 difference tests on slides 41 and 42: both constraints are rejected, but the equality constraint is the milder violation.

Identification: what a just-identified model looks like

Is there enough information to estimate every parameter?

$$p = \frac{k(k+1)}{2} \text{ (unique values in the observed covariance matrix); } q = \text{number of free parameters;}$$
$$df_M = p - q.$$

$k = 4$ variables (provocation, trait irritability, hostile cognition, aggression): $p = 4(5)/2 = 10$

Free parameters in Model 4, as drawn (all 10 listed)

$\text{Var}(X_1), \text{Var}(X_2), \text{Cov}(X_1, X_2)$	3
a_1, a_2 (paths into hostile cognition)	2
$\text{Var}(e_M)$ (residual variance of hostile cognition)	1
b_1, c'_1, c'_2 (paths into aggression)	3
$\text{Var}(e_Y)$ (residual variance of aggression)	1

Total: $q = 10$

$df_M = 10 - 10 = 0$

- **just-identified:** $df_M = 0$ – exactly enough information; the model reproduces the data perfectly, by construction, and $\chi^2 \approx 0$, exactly what we saw for Model 4 as drawn
- **over-identified:** $df_M > 0$ – fewer free parameters than unique covariances; the model can now, in principle, fit poorly
- **under-identified:** $df_M < 0$ – not enough information; the model cannot be estimated at all

Over-identified: fixing one direct effect, in the simplest model

Same counting logic, applied to Model 1

Take Model 1's simple mediation triangle (provocation $\rightarrow$ hostile cognition $\rightarrow$ aggression, plus the direct effect c) and fix c to zero, exactly as we will soon do to Model 4's two direct effects.

$k = 3$ variables (provocation, hostile cognition, aggression): $p = 3(4)/2 = 6$					
Model 1, as drawn (saturated)			Model 1, with c fixed to 0		
$\text{Var}(X), a, \text{Var}(e_M), c, b, \text{Var}(e_Y)$			$\text{Var}(X), a, \text{Var}(e_M), b, \text{Var}(e_Y)$		
$q = 6 \Rightarrow df_M = 6 - 6 = 0$			$q = 5 \Rightarrow df_M = 6 - 5 = 1$		
Model	χ^2	df	CFI	TLI	RMSEA
Saturated (Model 1, as drawn)	0.00	0	1.00	–	.000
c fixed to 0	32.28	1	.85	.54	.32

Now genuinely over-identified

$df_M = 1 > 0$: there is one more piece of observed information than free parameters, so the model *can* fit poorly – and here it does. $\Delta\chi^2(1) = 32.28 - 0.00 = 32.28$, $p < .0001$: provocation's direct effect on aggression is not actually zero, so removing it damages fit badly. This is the exact same logic as the complete-mediation test for Model 4 on slide 41, shown here first in the simplest possible model.

Results output: the over-identified model

```
> fitMeasures(fit1_over, c("chisq","df","cfi","tli","rmsea"))
  chisq    df    cfi    tli  rmsea
32.281  1.000  0.846  0.537  0.323

> anova(fit1_sat, fit1_over)
Chi-Squared Difference Test

      Df  AIC  BIC Chisq Chisq diff RMSEA Df diff Pr(>Chisq)
fit1_sat   0 4091 4109   0.0
fit1_over   1 4121 4136  32.3      32.3 0.323      1  1.3e-08 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Under-identified: when there is not enough information

What if we dropped provocation and estimated hostile cognition and aggression as **reciprocally** causing each other – each a predictor of the other, with nothing to tell the two directions apart?

$$k = 2 \text{ variables (hostile cognition, aggression): } p = 2(3)/2 = 3$$

$$\text{Attempted free parameters: } \text{Var}(e_M), \text{Var}(e_Y), b (M \rightarrow Y), b' (Y \rightarrow M)$$

$$\text{Total: } q = 4 \quad df_M = p - q = 3 - 4 = -1$$

$df_M < 0$: not enough information to estimate the model

Only 3 observed quantities (two variances, one covariance) are available for 4 free parameters, so infinitely many combinations of b and b' reproduce the data equally well. No amount of data resolves this: it is a problem of model *structure*, not sample size.

The remedy: instrumental variables

This is the classic **non-recursive** (feedback loop) problem, which can be solved by using **instrumental variables**. An instrumental variable is an observed variable that directly predicts one variable in the loop but, by theory, has no direct path to the other and is unrelated to the other's disturbance. Because it can reach that other variable only *through* the loop, it supplies the extra information needed to estimate that direction. Each reciprocal path needs its own instrument.

Results output: the under-identified model

```
Warning: lavaan->lav_model_vcov():
```

```
Could not compute standard errors! The information matrix could not be inverted. This may be a symptom that the model is not identified.
```

```
Model Test User Model:
```

Test statistic	NA
Degrees of freedom	-1
P-value (Unknown)	NA

```
Regressions:
```

	Estimate	Std.Err	z-value	P(> z)
aggression ~				
hstl_cg (b)	1.279	NA		
hostile_cognition ~				
aggrssn (bprm)	0.031	NA		

```
Variances:
```

	Estimate	Std.Err	z-value	P(> z)
.aggression	167.230	NA		
.hostile_cogntn	22.679	NA		

Free, fixed, and constrained parameters

Recall, from earlier

Every parameter in a path model is free, fixed, or constrained. Model 4 as drawn used only free parameters – that is why it was saturated. We now impose two different kinds of restriction on that same model, and watch what an over-identified χ^2 test looks like.

- **a fixed parameter:** what if we fix both direct effects, c'_1 and c'_2 , to zero? This is the classical test of **complete mediation**.
- **an equality constraint:** what if we require $a_1 = a_2$ – that provocation and trait irritability shift hostile cognition by exactly the same amount?

The complete-mediation model: fixing the direct effects to zero

lavaan syntax – the restricted model, and comparing it to the saturated one

```
m4_complete <- 'hostile_cognition ~ a1*provocation + a2*trait_irritability
               aggression ~ b1*hostile_cognition # c1, c2 omitted = fixed to 0
               provocation ~~ trait_irritability'
fit_saturated <- sem(m4, data = gam_mediation)
fit_complete <- sem(m4_complete, data = gam_mediation)
anova(fit_saturated, fit_complete) # runs exactly the test below
```

Degrees of freedom, counted explicitly

$q = 8$: the same 10 free parameters as the saturated model, minus c'_1 and c'_2 , now fixed to 0. $df_M = p - q = 10 - 8 = 2$.

Model	χ^2	df	p	CFI	RMSEA	SRMR
Saturated (Model 4, as drawn)	0.00	0	–	1.00	.000	.000
Complete mediation ($c'_1 = c'_2 = 0$)	38.57	2	< .0001	.85	.25	.088

The χ^2 difference test, computed step by step

$$\Delta\chi^2 = \chi^2_{\text{complete mediation}} - \chi^2_{\text{saturated}} = 38.57 - 0.00 = 38.57; \quad \Delta df = 2 - 0 = 2$$

Compare 38.57 to a χ^2 distribution with 2 degrees of freedom: `pchisq(38.57, df = 2, lower.tail = FALSE)`, or `anova(fit_saturated, fit_complete)`
Result: $p < .0001$. Fixing both direct effects to zero significantly worsens fit – provocation and trait irritability retain genuine direct effects on aggression.
Complete mediation is rejected in favor of partial mediation.

The equality-constrained model: constraining two paths to be equal

lavaan syntax – the same label forces equality, and comparing it to the saturated one

```
m4_equal <- 'hostile_cognition ~ a*provocation + a*trait_irritability
  aggression ~ c1*provocation + c2*trait_irritability + b1*hostile_cognition
  provocation ~~ trait_irritability'
fit_equal <- sem(m4_equal, data = gam_mediation)
anova(fit_saturated, fit_equal) # runs exactly the test below
```

Degrees of freedom, counted explicitly

$q = 9$: the same 10 free parameters as the saturated model, but a_1 and a_2 now share one label (a), so they count as a single free parameter instead of two.
 $df_M = p - q = 10 - 9 = 1$.

Model	χ^2	df	p	CFI	RMSEA	SRMR
Saturated (Model 4, as drawn)	0.00	0	–	1.00	.000	.000
Equality constrained ($a_1 = a_2 = 0.247$)	11.46	1	.0007	.96	.19	.047

The χ^2 difference test, computed step by step

$\Delta\chi^2 = 11.46 - 0.00 = 11.46$; $\Delta df = 1 - 0 = 1$; `pchisq(11.46, df = 1, lower.tail = FALSE)`

Result: $p = .0007$. The constraint is rejected: provocation's pull on hostile cognition ($a_1 = 0.367$) is not the same size as trait irritability's ($a_2 = 0.153$) – a situational input moves hostile cognition roughly $2.4\times$ more than the stable trait does. Both restrictions are rejected: CFI drops to .96 here, against .85 for the complete-mediation model.

Results output: the complete-mediation and equality-constrained models

```
> anova(fit4, fit4_complete)
Chi-Squared Difference Test

            Df  AIC  BIC Chisq Chisq diff RMSEA Df diff Pr(>Chisq)
fit4              0 7997 8034   0.0
fit4_complete    2 8031 8061  38.6      38.6 0.247      2  4.2e-09 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

> fitMeasures(fit4_complete, c("chisq","df","cfi","tli","rmsea","srmr"))
   chisq    df    cfi    tli  rmsea  srmr
38.568  2.000  0.852  0.557  0.247  0.088

> anova(fit4, fit4_equal)
Chi-Squared Difference Test

            Df  AIC  BIC Chisq Chisq diff RMSEA Df diff Pr(>Chisq)
fit4              0 7997 8034   0.0
fit4_equal        1 8006 8039  11.5      11.5 0.187      1  0.00071 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

> fitMeasures(fit4_equal, c("chisq","df","cfi","tli","rmsea","srmr"))
   chisq    df    cfi    tli  rmsea  srmr
11.458  1.000  0.958  0.746  0.187  0.047
```

Why constraints matter

- a **fixed** parameter tests whether a path is needed at all – here, whether mediation is complete or only partial
- a **constrained** parameter tests a specific quantitative hypothesis about two paths – here, whether two predictors matter equally
- both are tested the same way: fit the restricted model, compare χ^2 against a less-restricted model, with df equal to the number of restrictions

Where this returns

- multi-group analysis: are indirect effects equal across two groups?
- measurement invariance: are factor loadings equal across groups or time?
- model comparison generally: is a simpler, nested model a more parsimonious yet still adequate account of the data than a more complex one?

Moderation via Multi-Group Analysis

A new research question

So far, moderators have entered as product terms within a single sample. When the moderator is categorical, there is a second way to test moderation: fit the same model separately in each group, and ask whether the indirect effect itself differs. **Does the provocation → hostile cognition → aggression pathway operate the same way for men and for women?**

- a purpose-built dataset ($N = 300$), designed so that the indirect effect genuinely differs by gender
- the same path model, fit simultaneously in two groups, in lavaan multi-group syntax
- a full likelihood-ratio χ^2 test of whether the a path, the path that drives the group difference in the indirect effect, is equal across groups, computed step by step

A new dataset, built to show a genuine group difference

A new dataset, with a grouping variable added

A new, purpose-built simulated dataset ($N = 300$).

Group	n	a (unstd.)	SE	p	a (std.)
Men	146	0.247	0.063	< .001	0.308
Women	154	0.493	0.058	< .001	0.567

- b path (hostile cognition $\rightarrow$ aggression): men 1.532 (std. .514); women 1.595 (std. .540) – similar across groups
- indirect effect ab : men 0.379 [boot. CI 0.206, 0.571]; women 0.786 [boot. CI 0.554, 1.046]
- the a path, not the b path, is where the groups visibly diverge – provocation moves hostile cognition roughly twice as much for women as for men in this simulated sample

Fitting the same model in two groups: lavaan syntax

The `group=` argument, and where the equality constraint actually lives

```
model_free <- 'hostile_cognition ~ provocation
              aggression ~ hostile_cognition + provocation'
fit_free <- sem(model_free, data = gam_multigroup, group = "gender")

model_constrained <- 'hostile_cognition ~ a*provocation
                     aggression ~ hostile_cognition + provocation'
fit_constrained <- sem(model_constrained, data = gam_multigroup, group = "gender")
anova(fit_free, fit_constrained)
```

- `fit_free`: no path is labeled, so lavaan's default multi-group behavior applies – every path is free to take a different value in each group
- `fit_constrained`: one label (`a`) in the model string shared by both groups forces the `a` path to be equal across groups – the “same label forces equality” mechanism of the equality-constrained model, now applied *across groups*
- a common mistake: a single label added only to name a path silently constrains it across groups (lavaan warns about this). To name a path and keep it free, give one label per group:
`c(a_women, a_men)*provocation`

Results output: the free multi-group model

Group 1 [Women]:

Regressions:

	Estimate	Std.Err	z-value	P(> z)	Std.lv	Std.all
hostile_cognition ~						
provocation	0.493	0.058	8.536	0.000	0.493	0.567
aggression ~						
hostile_cogntn	1.595	0.189	8.416	0.000	1.595	0.540
provocation	0.781	0.165	4.735	0.000	0.781	0.304

Group 2 [Men]:

Regressions:

	Estimate	Std.Err	z-value	P(> z)	Std.lv	Std.all
hostile_cognition ~						
provocation	0.247	0.063	3.914	0.000	0.247	0.308
aggression ~						
hostile_cogntn	1.532	0.194	7.911	0.000	1.532	0.514
provocation	0.700	0.155	4.508	0.000	0.700	0.293

Results output: the constrained multi-group model

Group 1 [Women]:

Regressions:

	Estimate	Std.Err	z-value	P(> z)	Std.lv	Std.all
hostile_cognition ~						
provocatin (a)	0.382	0.043	8.831	0.000	0.382	0.465
aggression ~						
hstl_cgntn	1.595	0.187	8.517	0.000	1.595	0.529
provocatin	0.781	0.153	5.087	0.000	0.781	0.316

Group 2 [Men]:

Regressions:

	Estimate	Std.Err	z-value	P(> z)	Std.lv	Std.all
hostile_cognition ~						
provocatin (a)	0.382	0.043	8.831	0.000	0.382	0.442
aggression ~						
hstl_cgntn	1.532	0.191	8.032	0.000	1.532	0.529
provocatin	0.700	0.165	4.251	0.000	0.700	0.280

```
> anova(fit_free, fit_constrained)
```

Chi-Squared Difference Test

	Df	AIC	BIC	Chisq	Chisq diff	RMSEA	Df diff	Pr(>Chisq)
fit_free	0	4022	4074	0.00				
fit_constrained	1	4028	4077	8.13	8.13	0.218	1	0.0043 **

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Testing equality of the a path across groups, step by step

The same test, worked by hand as a likelihood-ratio comparison

Compare the log-likelihoods of the free model (a differs by group) and the constrained model (a shared); twice their difference is referred to a χ^2 distribution. This is the arithmetic inside `anova(fit_free, fit_constrained)`, shown on the previous slide.

Model	Log-likelihood	q	χ^2	df
Free (a differs by group)	-1997.140	14	0.00	0
Constrained (a shared)	-2001.207	13	8.13	1

The likelihood-ratio χ^2 test, computed step by step

$$LR = 2 \times (LL_{\text{free}} - LL_{\text{constrained}}) = 2 \times (-1997.140 - (-2001.207)) = 2 \times 4.067 = 8.13$$

$$\Delta df = 14 - 13 = 1 \quad (\text{equivalently, } \Delta \chi^2 = 8.13 - 0.00 = 8.13 \text{ on } \Delta df = 1 - 0 = 1)$$

Compare 8.13 to a χ^2 distribution with 1 degree of freedom: `pchisq(8.13, df = 1, lower.tail = FALSE)`

Result: $\chi^2(1) = 8.13$, $p = .0043$. The a path genuinely differs by group – this is moderated mediation, with gender as a categorical moderator.

What this tells us

- the indirect effect through hostile cognition is larger for women ($0.493 \times 1.595 = 0.786$) than for men ($0.247 \times 1.532 = 0.379$) in this simulated sample; the difference, 0.408, has a 95% bootstrap CI of [0.103, 0.722], which excludes 0
- the group difference traces specifically to the *a* path (provocation $\rightarrow$ hostile cognition), which differs reliably between groups ($\chi^2(1) = 8.13, p = .004$), not the *b* path – worth reporting which stage of the mediation is doing the moderating, not just that “the indirect effect differs”
- multi-group analysis and the product-term (continuous-moderator) approach from the previous lecture answer the same underlying question – “does this effect depend on a third variable?” – and the choice between them depends on whether the moderator is categorical (multi-group) or continuous (product term)

Common Mistakes: Things Not to Say – and What to Say Instead

Avoid

- “The path analysis *proves* that provocation causes aggression.”
- “The significant indirect effect *proves* that hostile cognition is the mechanism.”
- “Because the model fits well, the model is *true*.”
- “A non-significant direct effect means there is *no* relationship.”
- “Adding more paths always *improves* the model.”

Prefer

- “The results are *consistent with* the hypothesized indirect association.”
- “The model suggests the association *may operate partly* through hostile cognition.”

The manymome Package: Many Models, One Mediation

Why a dedicated package

`manymome` (Cheung & Cheung, 2024) computes indirect effects, total effects, and *conditional* indirect effects (moderated mediation) directly from an already-fitted `lm` or `lavaan` object – no path labels, no `:=` definitions written in advance, bootstrap done once and reused for every quantity asked of it afterward.

- works whether the mediator and outcome models were fit with `lm()` or `sem()`
- computes the indirect effect at chosen moderator values (e.g., ± 1 SD) automatically – exactly what moderated mediation requires
- self-documenting: every result can be asked to print the actual arithmetic behind it

Worked example 1: the serial mediation model, revisited

No path labels needed

```
fit <- sem(model3, data = gam_mediation) # an ordinary, unlabeled lavaan fit
out1 <- indirect_effect(x = "provocation", y = "aggression",
                       m = "hostile_cognition", fit = fit,
                       boot_ci = TRUE, R = 5000, seed = 4171)
out2 <- indirect_effect(x = "provocation", y = "aggression",
                       m = c("hostile_cognition", "neg_affect"), fit = fit,
                       boot_ci = TRUE, R = 5000, seed = 4171)
```

Path	Estimate	95% boot. CI	Matches
$X \rightarrow M_1 \rightarrow Y$ (out1)	0.309	[0.163, 0.471]	$a_1 b_1$, Model 3
$X \rightarrow M_1 \rightarrow M_2 \rightarrow Y$ (out2)	0.116	[0.057, 0.190]	$a_1 d_1 b_2$, Model 3

The point

Both numbers, CIs included, are identical to the ones computed for Model 3: both bootstraps use the same seed (4171), so they draw the same 5,000 resamples. When more than one mediator is given, `m = c(...)` lists them in the order the path passes through them: `c("hostile_cognition", "neg_affect")` means provocation $\rightarrow$ hostile cognition $\rightarrow$ negative affect $\rightarrow$ aggression ($X \rightarrow M_1 \rightarrow M_2 \rightarrow Y$).

Results output: worked example 1

```
> out1
== Indirect Effect ==
```

```
Path:          provocation -> hostile_cognition -> aggression
Indirect Effect: 0.309
95.0% Bootstrap CI: [0.163 to 0.471]
```

```
Computation Formula:
  (b.hostile_cognition-provocation)*(b.aggression-hostile_cognition)
```

```
Computation:
  (0.41584)*(0.74286)
```

Percentile confidence interval formed by nonparametric bootstrapping with 5000 bootstrap samples.

```
Coefficients of Component Paths:
              Path Coefficient
hostile_cognition-provocation    0.416
aggression-hostile_cognition     0.743
```

```
> out2
== Indirect Effect ==
```

```
Path:          provocation -> hostile_cognition -> neg_affect -> aggression
Indirect Effect: 0.116
95.0% Bootstrap CI: [0.057 to 0.190]
```

Worked example 2: conditional indirect effects, a continuous moderator

Research question: does trait irritability moderate the provocation → hostile cognition → aggression pathway?

```
d$c_provocation <- d$provocation - mean(d$provocation) # centered, as in Topic 1
d$c_trait_irr <- d$trait_irritability - mean(d$trait_irritability)
d$prov_x_irr <- d$c_provocation * d$c_trait_irr # the product term
fit_mome <- sem('hostile_cognition ~ c_provocation + c_trait_irr + prov_x_irr
               aggression ~ hostile_cognition + c_provocation', data = d)
out3 <- cond_indirect_effects(wlevels = mod_levels("c_trait_irr", fit = fit_mome),
                             x = "c_provocation", y = "aggression", m = "hostile_cognition",
                             fit = fit_mome, boot_ci = TRUE, R = 5000, seed = 4171)
```

Trait irritability	Conditional indirect effect	95% boot. CI
$M - 1SD = 23.00$	0.414	[0.259, 0.594]
Mean = 30.00	0.375	[0.239, 0.529]
$M + 1SD = 37.00$	0.335	[0.177, 0.511]

Index of moderated mediation

`index_of_mome(x = "c_provocation", y = "aggression", m = "hostile_cognition", w = "c_trait_irr", ...)`
→ $a_3b_1 = -0.006$, 95% boot. CI $[-0.018, 0.006]$: includes 0. Consistent with the interaction term itself ($a_3 = -0.0056$, $SE = 0.0061$, $p = .363$), trait irritability does *not* significantly moderate this pathway in these data – an honest non-significant result, reported the same way a significant one would be.

Results output: worked example 2

```
> out3
```

```
== Conditional indirect effects ==
```

```
Path: c_provocation -> hostile_cognition -> aggression
```

```
Conditional on moderator(s): c_trait_irr
```

```
Moderator(s) represented by: c_trait_irr
```

	[c_trait_irr]	(c_trait_irr)	ind	CI.lo	CI.hi	Sig	hostile_cognition-c_provocation	aggression-hostile_cognition
1	M+1.0SD	6.997	0.335	0.177	0.511	Sig	0.327	1.022
2	Mean	-0.000	0.375	0.239	0.529	Sig	0.366	1.022
3	M-1.0SD	-6.997	0.414	0.259	0.594	Sig	0.405	1.022

```
> index_of_mome(...)
```

```
== Index of Moderated Mediation ==
```

```
Levels compared: Row 1 - Row 2
```

	x	y	Index	CI.lo	CI.hi
Index	c_provocation	aggression	-0.006	-0.018	0.006

```
- [CI.lo, CI.hi]: 95% percentile confidence interval.
```

Advantages over hand-written lavaan syntax

- **No advance path labeling:** any indirect path through a fitted model can be requested after the fact, unlike `:=`, which must be written before fitting
- **Bootstrap once, reuse everywhere:** `do_boot(fit, R = 5000, seed = 4171)` runs the resampling a single time; every subsequent `indirect_effect()` or `cond_indirect_effects()` call reuses it (`boot_out =`) instead of re-bootstrapping – our worked examples above would ordinarily need three separate bootstrap runs, one per quantity
- **Any moderator level, on demand:** `wlevels` accepts any values, not just $\pm 1SD$ – re-running at a new level needs no model refit
- **Self-documenting:** `get_one_cond_indirect_effect()` prints the exact substituted formula and numbers behind any single conditional effect, useful for teaching and for verifying a result by hand
- **Handles missing data and multiple imputation**, and works identically for latent-variable mediation once the fitted object is a full SEM rather than a simple path model
- **Group as moderator:** the same `cond_indirect_effects()` machinery treats a categorical group variable as a moderator, giving a second route to the gender comparison from the multi-group section

Post-estimation: drawing the fitted model

Turning a fitted object back into a diagram

```
library(semPlot)
semPaths(fit, what = "std", layout = "tree", edge.label.cex = 0.9)



library(lavaanPlot)
lavaanPlot(model = fit, coefs = TRUE, stand = TRUE, sig = .05)
```

- `semPaths()` draws the path diagram directly from a fitted lavaan object, with standardized coefficients placed on the arrows automatically – the same diagrams seen throughout this lecture can be regenerated this way rather than drawn by hand
- `lavaanPlot()` produces a similar, publication-ready diagram with an alternative rendering engine, and can grey out non-significant paths
- `manyome` itself does not draw diagrams – it computes effects, not visuals; diagram production and effect computation are complementary, separate tools

A preview: when the moderator changes a specific path

Two different research questions

Today's multi-group section asked whether the indirect effect differs by group. A sharper question: *which specific path* does the moderator change – the $X \rightarrow M$ path, or the $M \rightarrow Y$ path?



The two equations

First-stage moderated mediation

$$M = i_1 + a_1X + a_2W + a_3(XW) + e_M, \quad \theta_{X \rightarrow M|W} = a_1 + a_3W$$

Second-stage moderated mediation

$$Y = i_2 + c'X + b_1M + b_2W + b_3(MW) + e_Y, \quad \theta_{M \rightarrow Y|W} = b_1 + b_3W$$

- in each case, the conditional path (θ) is itself a linear function of W – exactly the same logic as the conditional-effect equation from the interaction lecture, now embedded inside a mediation model
- the index of moderated mediation (a_3b_1 for first-stage) is the single-number test of whether that conditional path genuinely changes with W – what `index_of_mome()` computed on the previous slide, for a first-stage example
- full worked examples of both, including combined (moderated at both stages) models, are in Hayes's *Introduction to Mediation, Moderation, and Conditional Process Analysis*

Summary

- a path model is a system of regression equations, estimated together and drawn as a diagram
- path coefficients describe specific relationships; direct, indirect, and total effects answer different questions about the same predictor-outcome pair
- a statistically estimated path model does not, by itself, establish causality
- every parameter is free, fixed, or constrained – and each of those can be tested against the data with a χ^2 difference test
- `manymome` extends this toolkit to conditional and multi-step indirect effects, computed post hoc from a fitted model
- model fit and identification, introduced today for the first time, return in full starting with common factor models next