EPSE 596: Correlational Designs and Analysis

Ed Kroc

University of British Columbia

ed.kroc@ubc.ca

"Mediators" and "moderators"

- There is a very common terminology used throughout the social sciences, and throughout a significant amount of the health sciences, that we should clarify: mediation/mediators and moderation/moderators.
- A **mediator** in a regression model is a predictor M that "explains" (accounts for) the relationship between another predictor X and the response Y.
  - I.e., when you include M in the model, the effect of X on Y goes away (or is seriously mitigated).
- A **moderator** in a regression model is a predictor M that "affects the strength" of the relationship between another predictor X and the response Y.
  - I.e., there is a significant interaction between X and M in the model.

“Mediators” and “moderators”

- There is a long history of the mediator/moderator lingo in the behavioural sciences.
- The terminology is often inextricably linked with ideas/claims of *causal inference*.
- Note, in particular, that “mediation/moderation” both implicitly rely on some kind of *directionality* between the relationships between the model variables.
- In fact, the terminology of “mediation/moderation” came out of theoretical models for causality in psychology in the early part of the 20th Century.
- Notably, these theoretical models were *not* mathematical/statistical in nature; they relied on then-fasionable philosophies of causality in psychology.

"Mediators" and "moderators"

- From a purely statistical point of view, the language of "mediation/moderation" offers no benefits; in fact, it is mathematically *inconsistent and ill-defined*, largely because of the *directionality* bit:
  - Realize: nowhere in our regression framework (or in the math) do we have any notion of a predictor X *casually affecting* (or even occuring before) we observe a response Y; i.e., you can just as easily specify a regression model as
  $$X \sim N(\beta_0 + \beta_1 Y, \sigma_1^2)$$
  as you can specify
  $$Y \sim N(\gamma_0 + \gamma_1 X, \sigma_2^2).$$
- There are several mathematical/statistical models/formalizations of *causality*, but the classical "mediator/moderator" lingo that is so prevalent in the behavioural sciences today rarely use these ideas.

3 semi-coherent causal frameworks in use

- Fisher: Classical experimental design:
  - Experimental *control* and *manipulation* of key variables
  - Randomization of treatment
  - Exchangeability of sample units across treatments
- Neyman-Rubin: Counterfactual causality:
  - Counterfactual probability (i.e., we observe outcome "survival" but could have observed outcome "death")
  - Missing data problems
- Pearl: "do"–calculus:
  - Counterfactual probability augmented
  - Directed acyclic graphs (DAGs)

Causal inference and pseudo-causal inference

- There are a variety of methods out there that claim to be useful for causal inference.
- Some methods are statistically valid (when properly done):
  - Regression discontinuity-designs (local causal inference)
  - Propensity score methods (partial causal inference)
  - Instrumental variable methods (full causal inference, theoretically possible)
- Some methods are **not** statistically valid (without assuming a bunch of untestable and unlikely structure to the problems you are studying):
  - Mediation modelling
  - Naive structural equation modelling (SEM)
  - Naive network analysis
  - Imputation and other missing data techniques

Causal inference and pseudo-causal inference

- Causality doesn't come from specifying some kind of model; it comes from the structure of your design (control, manipulation, randomization, etc.).
- At best, all causal inference in an observational setting (even reasonable approaches like PS or IV) is model-dependent. Mediation/moderator analysis doesn't even try to achieve "balance" like PSs, or take advantage of incidental randomization (like IVs), or of a natural experiment creating exchangeability at a point in time (like DDs).
- So in particular, everything model-based is nonsense unless that model is valid to begin with. Mediation analysis requires 4 different models to be valid simultaneously. Without this, you cannot even begin to think about if the model's estimated correlations are actually reflecting anything causal. ✓✓

Mediation modelling

- "The triangle of mediation" (AKA *the triangle of broken dreams*):

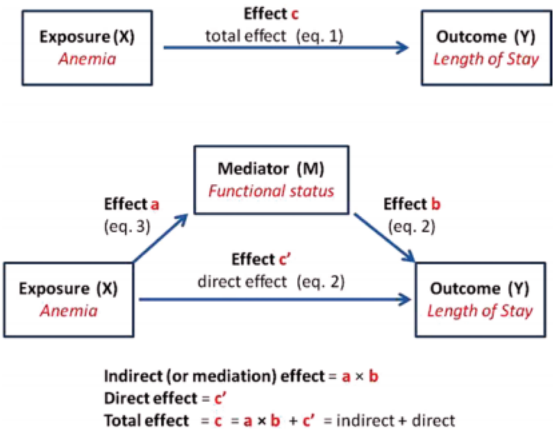

Mediation modelling

Classical mediation analysis recipe comes from Baron & Kenny (1986):

(1) Test that the "total effect" of X on Y by fitting $Y = \beta_0 + \beta_1 X + \varepsilon_1$. If $\beta_1 \not\approx 0$, then proceed.

(2) Test that X predicts M by fitting $M = \alpha_0 + \alpha_1 X + \varepsilon_2$. If $\alpha_1 \not\approx 0$, then proceed.

(3) Test that M predicts Y by fitting $Y = \gamma_0 + \gamma_1 M + \varepsilon_3$. If $\gamma_1 \not\approx 0$, then proceed.

(4) Fit the composite model $Y = \lambda_0 + \lambda_1 X + \lambda_2 M + \varepsilon_4$ and calculate the "direct effect" of X on Y, λ_1 and the "indirect effect" of X on Y (mediated by M), $\beta_1 - \lambda_1 = \alpha_1 \lambda_2$.

If $\lambda_2 \not\approx 0$, M is said to "(partially) mediate" the relationship between X and Y. If also $\lambda_1 \approx 0$, then M is said to "fully mediate" the relationship between X and Y.

Mediation modelling

- Creates the standard "triangle of mediation":

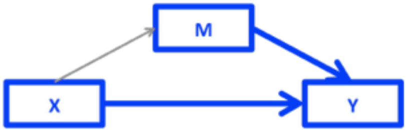

- Notice the arrows (i.e., directed edges) in the above path diagram, implying a *directional* (i.e., causal) relationships between the variables.
- But regression modelling (in fact, all of probability theory) *has no concept of directionality* encoded in anything we do. This diagram is an artifice.

Mediation modelling

(1) Test that the "total effect" of X on Y by fitting $Y = \beta_0 + \beta_1 X + \varepsilon_1$. If $\beta_1 \not\approx 0$, then proceed.

(2) Test that X predicts M by fitting $M = \alpha_0 + \alpha_1 X + \varepsilon_2$. If $\alpha_1 \not\approx 0$, then proceed.

(3) Test that M predicts Y by fitting $Y = \gamma_0 + \gamma_1 M + \varepsilon_3$. If $\gamma_1 \not\approx 0$, then proceed.

(4) Fit the composite model $Y = \lambda_0 + \lambda_1 X + \lambda_2 M + \varepsilon_4$ and calculate the "direct effect" of X on Y, λ_1 and the "indirect effect" of X on Y (mediated by M), $\beta_1 - \lambda_1 = \alpha_1 \lambda_2$.

In practice, people often skip steps (2) and/or (3). Statistically immaterial, as neither approach does what it is trying to do; i.e., understand if the correlation between X and Y is spurious and attributable to some other variable M.

Mediation modelling

- Notice that nowhere in the mediation analysis tradition is there any mention of *considering if your regression models are reasonable and actually appropriately describe the data relationships*.
- Indeed, implicit in mediation modelling working is the assumption that *all* models in steps (1)–(4) are *correctly specified*, something that is virtually guaranteed to never hold in practice.
- Still, mediation analysis is ubiquitous in the behavioural sciences. So let's see an example to illustrate its vacuity: ✓✓

Mediation modelling

- Have 100 randomly generated (Y, X_1, M_1) observations.
- Want to check if M_1 "mediates" the relationship between X_1 and the response Y at all.
  - Step (1): Fit the first-order model where X_1 predicts Y:

```
Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)  -18.289      7.152  -2.557   0.0121 *
x1            -7.100      1.371  -5.180 1.19e-06 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' '

Residual standard error: 13.82 on 98 degrees of freedom
Multiple R-squared:  0.215,     Adjusted R-squared:  0.2069
F-statistic: 26.83 on 1 and 98 DF,  p-value: 1.187e-06
```

Mediation modelling

- Have 100 randomly generated (Y, X_1, M_1) observations.
- Want to check if M_1 "mediates" the relationship between X_1 and the response Y at all.
    - Step (2): Fit the first-order model where X_1 predicts M_1:

```
Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept) -10.0806     1.0873  -9.272 4.61e-15 ***
x1            2.0056     0.2084   9.626 7.86e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 '

Residual standard error: 2.101 on 98 degrees of freedom
Multiple R-squared:  0.486,     Adjusted R-squared:  0.4807
F-statistic: 92.65 on 1 and 98 DF,  p-value: 7.859e-16
```

Mediation modelling

- Have 100 randomly generated (Y, X_1, M_1) observations.
- Want to check if M_1 "mediates" the relationship between X_1 and the response Y at all.
  - Step (3): Fit the first-order model where M_1 predicts Y:

```
Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept) -54.9266     1.4969 -36.693  < 2e-16 ***
m1            1.5334     0.5149   2.978  0.00365 **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' '

Residual standard error: 14.94 on 98 degrees of freedom
Multiple R-squared:  0.08299,	Adjusted R-squared:  0.07364
F-statistic:  8.87 on 1 and 98 DF,  p-value: 0.003655
```

Mediation modelling

- Have 100 randomly generated (Y, X_1, M_1) observations.
- Want to check if M_1 "mediates" the relationship between X_1 and the response Y at all.
- Sufficient to just perform steps (1) and (4) of the standard recipe:
  - Step (4): Fit the first-order model where X_1 and M_1 both predict Y:

```
Coefficients:
             Estimate Std. Error t value Pr(>|t|)    
(Intercept)   45.5227     2.6788   16.99   <2e-16 ***
x1           -19.7956     0.5226  -37.88   <2e-16 ***
m1             6.3301     0.1817   34.85   <2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 

Residual standard error: 3.778 on 97 degrees of freedom
Multiple R-squared:  0.9419,	Adjusted R-squared:  0.9407 
F-statistic: 786.7 on 2 and 97 DF,  p-value: < 2.2e-16
```

So M_1 seems to "partially mediate" the effect of X_1 on Y.

Mediation modelling

- But how reasonable were the regression models to begin with?
- Remember that the model estimates cannot be trusted unless all model assumptions are satisfied; in particular, only if the functional form of the model is correctly specified..
- Res.v.fit plot for $Y = \beta_0 + \beta_1 X_1 + \varepsilon_1$ model looks good:

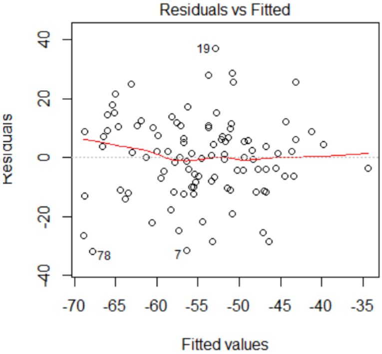

Mediation modelling

- But how reasonable were the regression models to begin with?
- Remember that the model estimates cannot be trusted unless all model assumptions are satisfied; in particular, only if the functional form of the model is correctly specified..
- Res.v.fit plot for $M_1 = \alpha_0 + \alpha_1 X_1 + \varepsilon_2$ model looks good:

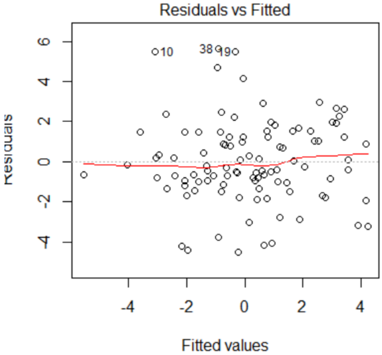

Mediation modelling

- But how reasonable were the regression models to begin with?
- Remember that the model estimates cannot be trusted unless all model assumptions are satisfied; in particular, only if the functional form of the model is correctly specified..
- Res.v.fit plot for $Y = \gamma_0 + \gamma_1 M_1 + \varepsilon_3$ model looks good:

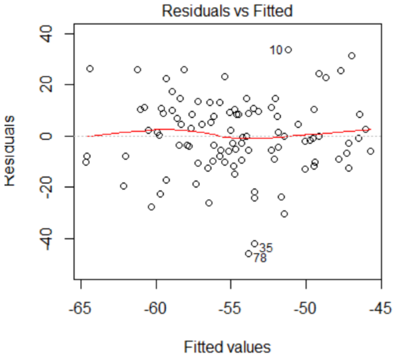

Mediation modelling

- But res.v.fit plot for composite $Y = \lambda_0 + \lambda_1 X_1 + \lambda_2 M_1 + \varepsilon_4$ model looks bad: obvious model misspecification, perhaps missing curvature?

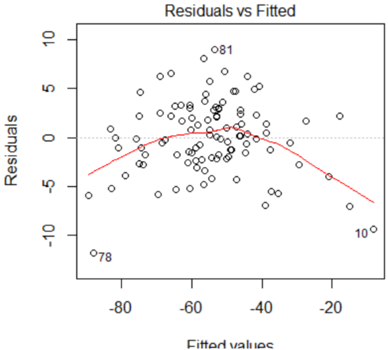

Mediation modelling

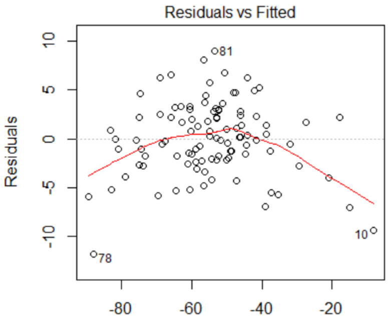

- Actually here, the model misspecification is due to an *omitted confounding variable* X_2 that is not included in the model.

Mediation modelling

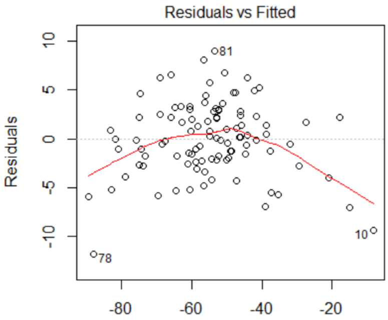

- Note: patterns in the residuals are usually due to missing/omitted variables; sometimes these are missing higher order terms of already-included predictors, but sometimes they are totally separate variables (perhaps not even measured/collected with the sample data).

Mediation modelling

- Let's see what happens to our "mediator" and "direct/indirect/mediated effects" if we include the missing predictor X_2:

```
Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept) 52.85022    1.60761  32.875  < 2e-16 ***
x1          -6.96627    0.94638  -7.361 6.23e-11 ***
m1          -0.07645    0.46042  -0.166    0.868
x2          14.21345    0.99543  14.279  < 2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' '

Residual standard error: 2.149 on 96 degrees of freedom
Multiple R-squared:  0.9814,    Adjusted R-squared:  0.9808
F-statistic:  1689 on 3 and 96 DF,  p-value: < 2.2e-16
```

Now M_1 does not seem to "mediate" the effect of X_1 on Y at all. In fact, one could say that X_2 "mediates" the relationship instead.

Mediation modelling

- But how reasonable is this new regression model?
- Res.v.fit plot for revised model including X_2 does not look appreciably better:

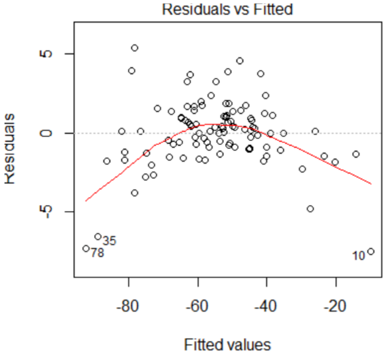

Mediation modelling

- In fact, there were *two* important missing predictors in the model: X_2 and the *interaction term* $X_1 \times X_2$.

```
Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)  1.31411    2.60889   0.504    0.616
x1           3.71582    0.66233   5.610 1.99e-07 ***
m1          -0.27825    0.19904  -1.398    0.165
x2           4.38673    0.64395   6.812 8.70e-10 ***
x1:x2        2.03994    0.09954  20.493  < 2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' '

Residual standard error: 0.9278 on 95 degrees of freedom
Multiple R-squared:  0.9966,    Adjusted R-squared:  0.9964
F-statistic:  6902 on 4 and 95 DF,  p-value: < 2.2e-16
```

- Again, M_1 does not seem to "mediate" the effect of X_1 on Y at all. But one could say that X_2 now simultaneously "mediates" and "moderates" the X_1 and Y relationship.

- And, as always, we check the res.v.fit plot for the new revised model; now things look quite good:

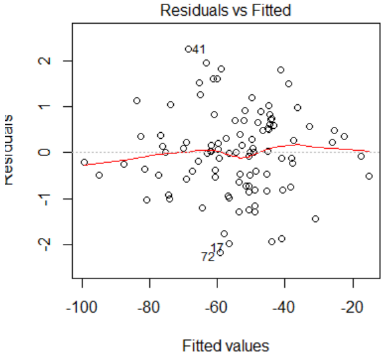

Mediation modelling

- But notice too that, throughout all of this analysis, we were mostly concerned with quantifying the effect of X_1 on the response Y.
- The first $Y \sim N(\beta_0 + \beta_1 X_1, \sigma^2)$ model yielded $\widehat{\beta}_1 = -7.10$.
- The composite $Y \sim N(\beta_0 + \beta_1 X_1 + \beta_2 M_1, \sigma^2)$ model yielded $\widehat{\beta}_1 = -19.79$.
- The revised $Y \sim N(\beta_0 + \beta_1 X_1 + \beta_2 M_1 + \beta_3 X_2, \sigma^2)$ model yielded $\widehat{\beta}_1 = -6.96$.
- The final revised $Y \sim N(\beta_0 + \beta_1 X_1 + \beta_2 M_1 + \beta_3 X_2 + \beta_4 X_1 X_2, \sigma^2)$ model yielded an estimated effect of X_1 on Y that *depended on the value of X_2*: $\widehat{\beta}_1 + \widehat{\beta}_4 X_2 = 3.71 + 2.04 X_2$. Here, the X_2 data ranged between about $[-7, -2]$, so the (local linear) effect of X_1 on Y ranges between about $[-10, 0]$.
- Notice just how poorly the "mediation analysis" did at quantifying and understanding this effect.

Mediation modelling

In summary:

- The language of "mediation" and "moderation" is not very useful, but often abused and a major source of bad inference/modelling/science.
- Mediation analysis is **not** causal modelling.
- Mediation analysis is extremely naive regression modelling, unlikely to lead to precise, generalizable, or sometimes even mildly accurate predictions or explanations of real world (complex) phenomena.
- Mediation analysis bears a superficial resemblance to a valid kind of causal modelling, i.e., to the *instrumental variables* technique.
- But mediation analysis ignores a critical component called *instrumental validity*. Would learn more about this in a causal inference/modelling course (EPSE 581C?).

Graduated mediation modelling: SEMs

- Structural equation models - SEMs (can take entire classes on this subject).
- Very useful for modelling latent variables (constructs) with measurement error (usually via a factor analysis on multiple *items*).
- However, SEMs are **not** causal modelling on their own.
- There is a *long* history of conflating SEMs with causal modelling (again, mostly in the behavioural sciences).
- The vast majority of SEM work is purely correlational, not causal.
- But see the work of J. Pearl et al. for a more rigorous, causally flavoured approach to SEMs that is made mathematically/statistically coherent by fusing it with actual causal inference techniques (e.g., instrumental variables).✓✓