Structural Causal Models

Causal Models and Meanings – Class II

Dean McHugh

University of Edinburgh
Department of Linguistics

Yerevan Academy for Linguistics and Philosophy (YALP)
June 2026

Plan

1. Why "correlation" is not enough.
2. The two basic ingredients: **variables** and **mechanisms**.
3. The formal object: a **structural causal model**.
4. **Interventions**: the difference between *seeing* and *doing*.
5. A first taste of **counterfactuals**: "what would have happened if. . . ?"

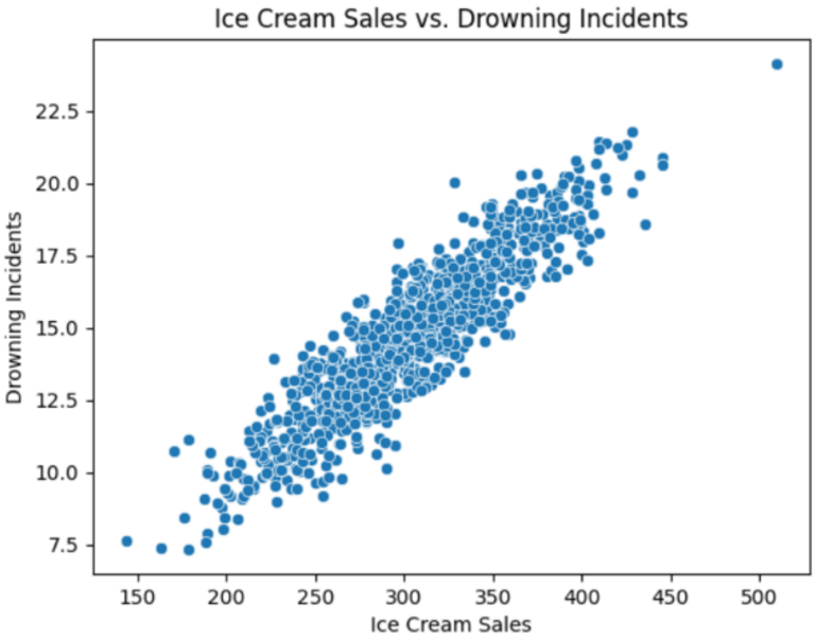
Ice Cream Sales vs. Drowning Incidents
Drowning Incidents
Ice Cream Sales
22.5
20.0
17.5
15.0
12.5
10.0
7.5
150
200
250
300
350
400
450
500

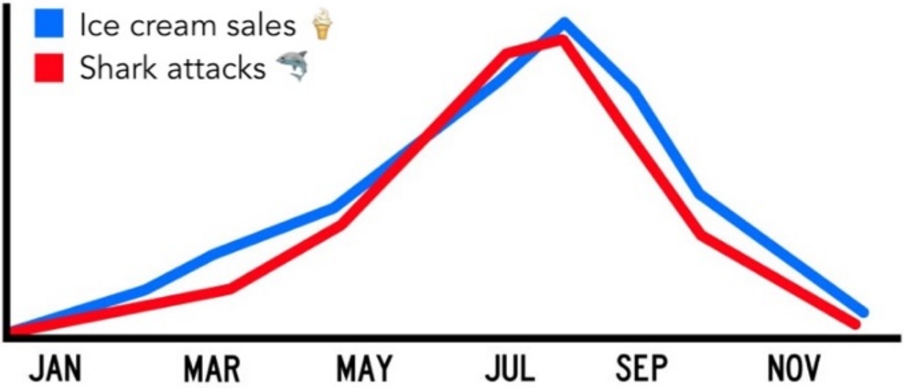
Ice cream sales
Shark attacks
JAN
MAR
MAY
JUL
SEP
NOV

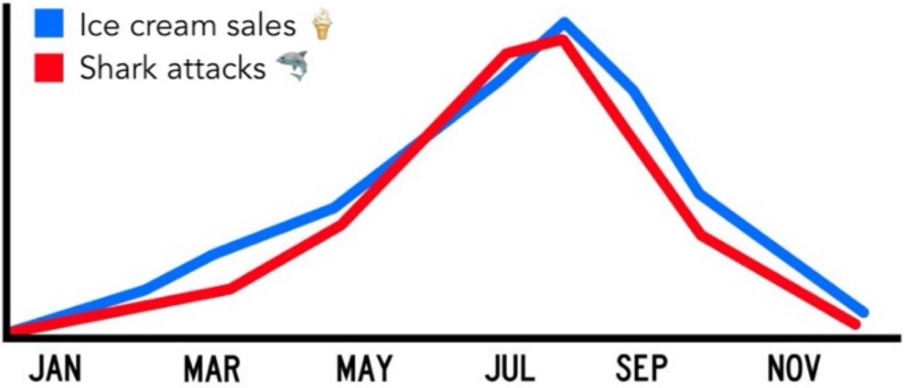

"Shark attacks are associated with ice cream sales."

The trouble with “correlation”

Two things rise and fall together:

$$\text{ice-cream sales} \xleftrightarrow{\textit{correlation}} \text{drownings}$$

Does ice cream cause drowning?

The trouble with “correlation”

Two things rise and fall together:

ice-cream sales $\xleftrightarrow{\textit{correlation}}$ drownings

Does ice cream cause drowning?
There is a *hidden common cause*:

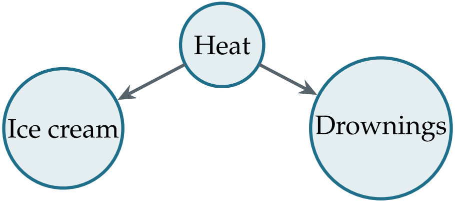

The moral

Correlation tells us *that* two things move together. It cannot, by itself, tell us *which causes which* — or whether a third thing drives both.[a]

[a] Pearl & Mackenzie (2018), *The Book of Why*, ch. 1.

The New York Times

THURSDAY, JANUARY 15, 1998

Georgia's Governor Seeks Musical Start for Babies

By KEVIN SACK

ATLANTA, Jan. 14 — As it is, newborn children in Georgia often come home from the hospital with a bag of free goodies: baby wipes, diapers, instructions about breast feeding and immunizations. Now Gov. Zell Miller wants to throw in a little something extra: a cassette tape or compact disc of classical music.

The music would not be intended to soothe the frayed nerves of parents getting their first doses of sleep-deprivation from their colicky babies. Rather, Mr. Miller, a devoted fan of

fascination came naturally.

He said that he had a stack of research on the subject, but also that his experiences growing up in the mountains of north Georgia had proved convincing.

"Musicians were folks that not only could play a fiddle but they also were good mechanics," he said. "They could fix your car."

During his budget address on Tuesday, Mr. Miller played a bit of Beethoven's "Ode to Joy" on a tape recorder and then asked the lawmakers, "Now don't you feel smart-

Florida's 'Beethoven's Babies' law

Tampa Bay Times NEWS SPORTS VIEWPOINTS CULTURE FOOD | Obituaries Classifieds

ARCHIVE

Legislature mandating classical music for tots

By **Adam C. Smith** *Former Times Staffer*

Published April 29, 1998 | Updated Sept. 13, 2005

Move over Barney, and make way for Beethoven.

The Legislature has passed a bill requiring classical music in all state-funded child care and educational programs. Mandating Mozart at day care, lawmakers hope, will stimulate brain development.

"I hadn't paid much attention to it, because I thought it was a joke," said Mary Tingiris of Tampa, president of the Florida Family Child Care Home Association. "I believe in classical music, I use classical music, but how do

What we actually want from a model

A good causal model should let us answer three kinds of question:

The "ladder of causation"[a]

[a] Pearl & Mackenzie (2018), *The Book of Why*, Introduction.

Rung	**Question**	**Example**
Seeing	What if I *observe*?	"Given wet grass, did it rain?"
Doing	What if I *act*?	"If I turn the sprinkler on. . . ?"
Imagining	What if things had been *otherwise*?	"Would it be wet had it not rained?"

Plain statistics lives on the first rung. Structural causal models climb all three.[1]

[1] See the *Pearl Causal Hierarchy* in Bareinboim et al. (2022).

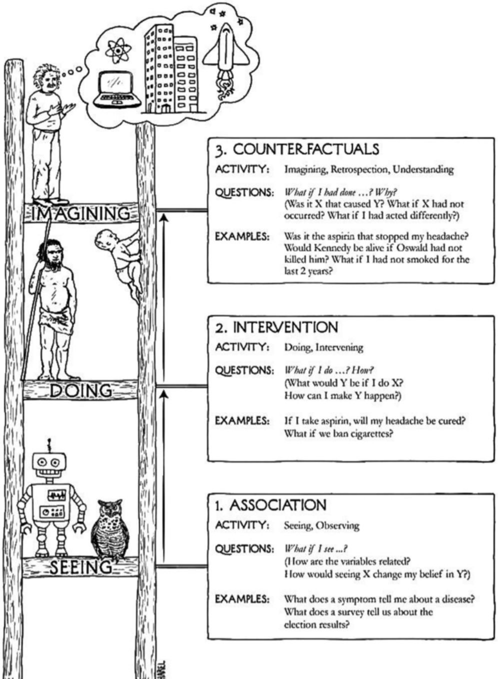
IMAGINING
DOING
SEEING
3. COUNTERFACTUALS
ACTIVITY: Imagining, Retrospection, Understanding
QUESTIONS: *What if I had done …? Why?*
(Was it X that caused Y? What if X had not occurred? What if I had acted differently?)
EXAMPLES: Was it the aspirin that stopped my headache? Would Kennedy be alive if Oswald had not killed him? What if I had not smoked for the last 2 years?
2. INTERVENTION
ACTIVITY: Doing, Intervening
QUESTIONS: *What if I do …? How?*
(What would Y be if I do X? How can I make Y happen?)
EXAMPLES: If I take aspirin, will my headache be cured? What if we ban cigarettes?
1. ASSOCIATION
ACTIVITY: Seeing, Observing
QUESTIONS: *What if I see …?*
(How are the variables related? How would seeing X change my belief in Y?)
EXAMPLES: What does a symptom tell me about a disease? What does a survey tell us about the election results?

Plan

Ingredient one: variables

A **variable** is just a question about the world with a fixed range of answers.

- *Rain* — did it rain? answers: $\{\text{yes}, \text{no}\}$
- *Wet* — is the pavement wet? answers: $\{\text{yes}, \text{no}\}$
- *Temp* — temperature, in °C answers: a number

We split variables into two families:

Endogenous (V)

Things the model *explains* — determined by other variables inside the model.
(drawn as solid circles)

Exogenous (U)

Background conditions the model takes as *given* — everything we do not model in detail.
(drawn dashed)

Ingredient two: a mechanism is an equation

For each endogenous variable we write *one* equation saying how its value is *produced* from its direct causes. Start tiny:

A one-mechanism world

$$Wet \;:=\; Rain$$

"Whether it is wet is *set by* whether it rained."

Notice the symbol := ("becomes / is set to"), not plain =.

Why the arrow-like " := " matters

$=$ is symmetric: $a = b$ is the same statement as $b = a$.
A *mechanism* is **not** symmetric. Rain makes the ground wet; drying the ground does not stop the rain. The assignment := records that direction.[a]

[a] Pearl, Glymour & Jewell (2016), *Causal Inference in Statistics: A Primer*, §1.

Where does the rest of the world go? Into U

Real outcomes have causes we do not list. We bundle all of them into an exogenous "noise" or "background" term U:

$$Wet := f_{Wet}(Rain, U_{Wet})$$

- f_{Wet} is the *mechanism* — a rule turning inputs into an output.
- U_{Wet} stands for everything else (a burst pipe, morning dew, a passing street-cleaner) lumped into one "unexplained" input.

Read it aloud

"The wetness is whatever the mechanism f_{Wet} outputs, given the rain and whatever else is going on."
Fix the inputs $\Rightarrow$ the output is fixed. That determinism is the whole idea; randomness enters only through the U's.

Plan

A small world: sprinkler, rain, and a slippery path

Let us model a garden path. Five yes/no variables:

Season	is it the dry season?
Sprinkler	is the sprinkler on?
Rain	is it raining?
Wet	is the path wet?
Slip	is the path slippery?

One mechanism per endogenous variable (suppressing the U's for readability):

$$\begin{aligned} \mathit{Sprinkler} &:= f_1(\mathit{Season}) \qquad & \mathit{Wet} &:= f_3(\mathit{Sprinkler}, \mathit{Rain}) \\ \mathit{Rain} &:= f_2(\mathit{Season}) \qquad & \mathit{Slip} &:= f_4(\mathit{Wet}) \end{aligned}$$

(Season is exogenous here — the model does not explain the weather's source.)

2

[2] The sprinkler world is the running example of Pearl (2009), *Causality*, §1.

From equations to a picture

Draw one arrow from each *input* to the variable it helps produce. The equations *become* a diagram:

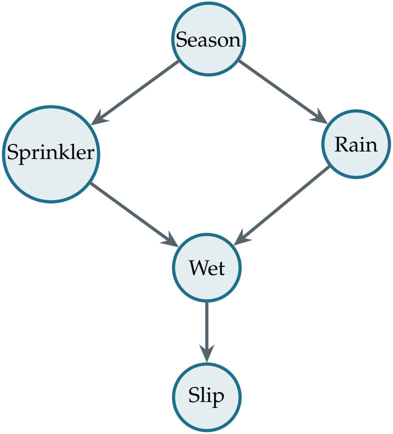

This is the **causal graph**. Arrow $X \to Y$ means "X appears on the right-hand side of Y's equation" — i.e. X is a *direct cause* of Y.

Reading the graph

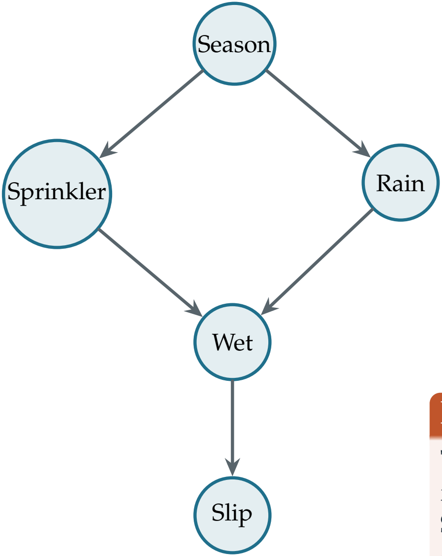

- **Parents** of *Wet* = its direct causes = $\{Sprinkler, Rain\}$.
- **Children** of *Season* = $\{Sprinkler, Rain\}$.
- A **path**: $Season \to Rain \to Wet \to Slip$.
- **Roots** (no parents) are exactly the things the model leaves unexplained.

No loops allowed

The arrows never cycle back. You cannot follow arrows and return to your start. Such a graph is a *directed acyclic graph* (DAG): nothing is its own ancestor.

Plan

A structural causal model, assembled

We have met every piece. Collect them:

Definition (informal but complete)

A **structural causal model** is a triple $M = \langle U, V, F \rangle$ where

- U — the *exogenous* variables (background, set outside the model);
- $V = \{V_1, \ldots, V_n\}$ — the *endogenous* variables (what we explain);
- $F = \{f_1, \ldots, f_n\}$ — one *structural equation* per endogenous variable,

$$V_i \; := \; f_i(pa_i,\, U_i),$$

with $pa_i \subseteq V$ the *parents* (direct causes) of V_i.

Adding a probability distribution $P(U)$ over the background gives us a *probabilistic* model: uncertainty about U flows through the equations to give probabilities for everything in V.[3]

[3] Pearl (2009), *Causality*, §1.4; Halpern (2016), *Actual Causality*, ch. 2.

The same definition, in plain English

What the triple says

1. List the things you want to explain (V) and the things you take for granted (U).
2. For *each* explained thing, write the single rule (f_i) that produces it from its direct causes and its own background noise.
3. Fix the background U, and *everything else follows* by computing the equations in order, parents before children.

Terminology: The outer variables are called the 'exogenous' variables, the inner variables the 'endogenous' variables.

A solved world

Choosing values for the outer variables (*Season* and all the U's) determines $Sprinkler, Rain, Wet, Slip$ uniquely.
Each full setting of U is one possible "run" of the world.

Why the arrows say more than statistics

Two models can predict the *same* correlations yet disagree about causation:

rain $\to$ wet　　　　wet $\to$ rain

Both fit “rain and wet occur together” equally well. They differ only in what they predict when we *intervene*. The extra content of a structural model — beyond the numbers — is precisely **the direction of the arrows**.[4]

[4] On the gap between statistical and causal models, see Pearl (2009), *Causality*, Preface; Peters, Janzing & Schölkopf (2017), ch. 1.

Plan

Doing is not the same as seeing

Compare two questions about the sprinkler:

- **Seeing:** "I *notice* the sprinkler is on. What does that tell me?" It also tells me the season is probably dry — because the season is what turned it on.
- **Doing:** "I *walk over and switch* the sprinkler on. Now what?" This tells me *nothing* about the season — I overrode it.

The do-operator

We write $\mathrm{do}(\mathit{Sprinkler} = \mathrm{on})$ for the action "force the sprinkler on, by hand." It *changes* the model, rather than observing it.[a]

[a] Pearl (2009), *Causality*, §1.3–1.4; Pearl, Glymour & Jewell (2016), §3.

An intervention is graph surgery

To compute do($Sprinkler$ = on): *delete the equation* for $Sprinkler$, fix its value by hand, and leave every other equation untouched.

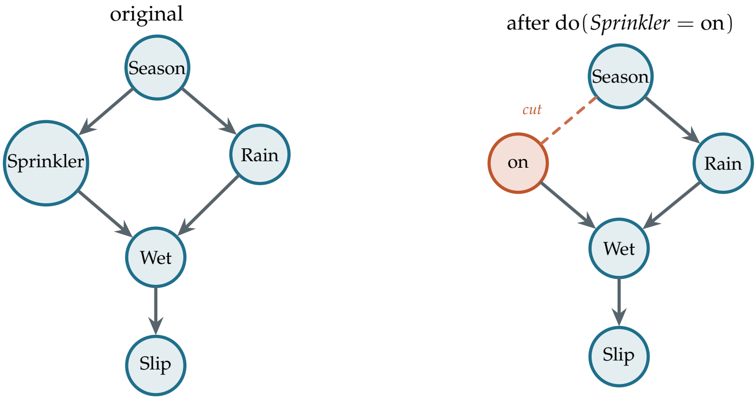

The arrow *into* the intervened variable is severed: its parents no longer matter — we, not the model, set its value. Everything downstream still runs as before.

Worked example: what does turning it on change?

Suppose the mechanisms are the obvious ones:

$$\mathit{Wet} := (\mathit{Sprinkler} \text{ or } \mathit{Rain}), \qquad \mathit{Slip} := \mathit{Wet}.$$

Compute do($\mathit{Sprinkler}$ = on)

- $\mathit{Sprinkler} = \text{on}$ (forced by us)
- $\mathit{Wet} = (\text{on or } \mathit{Rain}) = \text{yes}$ (the path is wet for sure)
- $\mathit{Slip} = \text{yes}$ (so it is slippery)
- *Rain*, *Season*: *unchanged in distribution* — the intervention told us nothing about them.

Contrast with merely *seeing* the sprinkler on, which would also make us revise our belief about the season. **Same value, different inference.**

Plan

The hardest rung: counterfactuals

A counterfactual asks about a world that did *not* happen, while holding on to the one that did:

Our question

"The path *is* slippery, and it *did* rain.
Would it *still* have been slippery, *had it not rained*?"

This mixes facts ("it did rain") with a contrary supposition ("had it not"). Plain observation and even plain intervention cannot express it — we need the full model, including the background U that produced the actual world.[5]

[5] The structural-equation treatment of counterfactuals: Halpern & Pearl (2005a,b); Halpern (2016), ch. 2. Cf. the possible-worlds tradition of Lewis (1973).

Three steps: abduction, action, prediction

Given what actually happened, evaluate a counterfactual in three moves:[6]

1. **Abduction.** Use the *evidence* (slippery, rained) to pin down the background U — figure out which "run of the world" we are in.
2. **Action.** Apply the supposed change as an intervention: $\mathrm{do}(\mathit{Rain} = \mathrm{no})$, surgically, leaving U fixed.
3. **Prediction.** Re-run the (now modified) equations with that *same* U and read off the outcome for *Slip*.

Why keep U fixed?

Holding the background constant is what makes it the *same situation*, merely with one thing changed. That is exactly the meaning of "had things been otherwise."

[6] Pearl (2009), *Causality*, §7.1 (the "three-step" recipe).

The recipe at work (sketch)

Question: *slippery & rained — still slippery had it not rained?*

1. **Abduction.** It was wet and slippery; it rained. Was the sprinkler on? That depends on U — so we *ask the model* what the background must have been.
   - If the sprinkler was *also* on, the path stays wet without rain.
   - If the sprinkler was *off*, removing rain dries the path.
2. **Action.** Set $\mathrm{do}(Rain = \mathrm{no})$, keep that same background.
3. **Prediction.** Recompute *Wet*, then *Slip*.

Payoff

The answer is *not* automatic — it depends on the hidden state we recovered in step 1. Counterfactuals require all three layers of the model at once, which is why they sit at the top of the ladder.

What to remember

The whole picture in five lines

1. A **variable** is a question with a range of answers; split into endogenous V and exogenous U.
2. A **mechanism** is one directed assignment $V_i := f_i(pa_i, U_i)$.
3. Collect them: a **structural causal model** $\langle U, V, F \rangle$, optionally with $P(U)$.
4. The equations *are* a **causal graph** (a DAG); arrows = direct causes.
5. **Intervene** by graph surgery (do); **counterfactuals** = abduction, action, prediction.

Same equations, three questions — seeing, doing, imagining.

Mini cheat-sheet

Symbol / term	Means
V (solid circle)	endogenous variable — explained by the model
U (dashed circle)	exogenous variable — background, taken as given
$V_i := f_i(pa_i, U_i)$	the mechanism: V_i is *set by* its parents and noise
pa_i	parents = direct causes of V_i
$X \to Y$	X is a direct cause of Y
DAG	directed graph with no cycles
$\mathrm{do}(X = x)$	force X to x; cut the arrows into X
counterfactual	abduction $\to$ action $\to$ prediction

Further reading

- **Gentlest start.** Pearl & Mackenzie (2018), *The Book of Why* — the ladder of causation, almost no formulas.
- **First textbook.** Pearl, Glymour & Jewell (2016), *Causal Inference in Statistics: A Primer* — short, worked exercises.
- **The reference.** Pearl (2009), *Causality* — full treatment of do, identification, and counterfactuals.
- **Causation, philosophically.** Halpern (2016), *Actual Causality*; Woodward (2003), *Making Things Happen*; the structural account of Halpern & Pearl (2005a,b); the counterfactual tradition of Lewis (1973).
- **Learning structure from data.** Spirtes, Glymour & Scheines (2000); Peters, Janzing & Schölkopf (2017).

References I

Bareinboim, E., Correa, J. D., Ibeling, D. & Icard, T. (2022). On Pearl's Hierarchy and the Foundations of Causal Inference. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, pp. 507–556. ACM. doi:10.1145/3501714.3501743.

Halpern, J. Y. (2016). *Actual Causality*. MIT Press.

Halpern, J. Y. & Pearl, J. (2005a). Causes and Explanations: A Structural-Model Approach. Part I: Causes. *The British Journal for the Philosophy of Science* 56(4), 843–887. doi:10.1093/bjps/axi147.

Halpern, J. Y. & Pearl, J. (2005b). Causes and Explanations: A Structural-Model Approach. Part II: Explanations. *The British Journal for the Philosophy of Science* 56(4), 889–911. doi:10.1093/bjps/axi148.

Lewis, D. (1973). Causation. *The Journal of Philosophy* 70(17), 556–567. doi:10.2307/2025310.

Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*, 2nd ed. Cambridge University Press.

References II

- Pearl, J., Glymour, M. & Jewell, N. P. (2016). *Causal Inference in Statistics: A Primer*. John Wiley & Sons.

- Pearl, J. & Mackenzie, D. (2018). *The Book of Why: The New Science of Cause and Effect*. Basic Books.

- Peters, J., Janzing, D. & Schölkopf, B. (2017). *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press.

- Spirtes, P., Glymour, C. & Scheines, R. (2000). *Causation, Prediction, and Search*, 2nd ed. MIT Press.

- Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.