Visualizing Environmental Econometrics

Sharing DAGs to communicate causal inference logic

Pierce Donovan

Department of Economics, Connecticut College

University of Nevada, Reno
April 2023

When we lack the facilities to present causal knowledge, we end up approaching causal inference blindfolded.

Claims of causally-interpretable regression output require sharing our understanding of the underlying data generating process.

[†]Netflix, *Bird Box* (2018)

Directed Acyclic Graphs help us share causal information

identifying variation:
- variables and relationships of interest
- setting in which data are collected
- research design (variation selection)

A Directed Acyclic Graph is a visualization of an observed, structured process. It represents our best understanding of the various causal and spurious links between variables of interest.

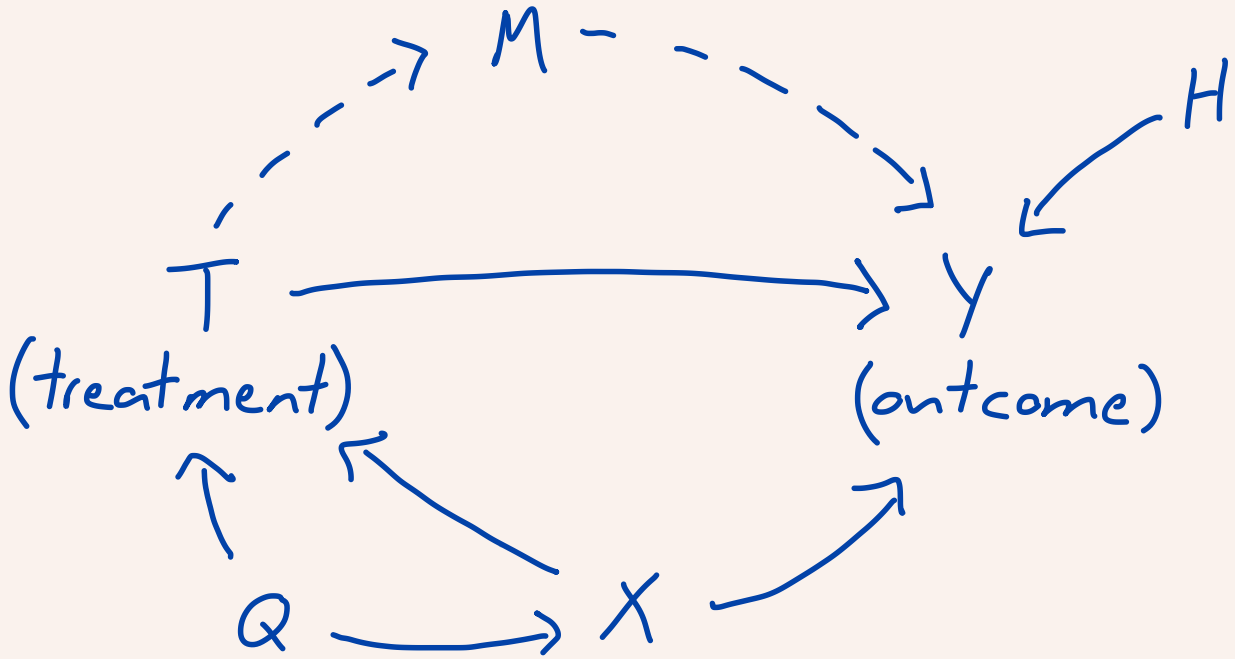

We want to focus on the *links* (edges) between nodes. They tell a story about our DGP, inform research design and presentation, and **validate claims of causal identification in regression output**.

Any directed paths that "leave" the independent variable of interest (T) are potentially-causal {green}. A path that "enters" the independent variable is a non-causal *back-door* path if it also links up with the dependent variable (Y) {purple}.

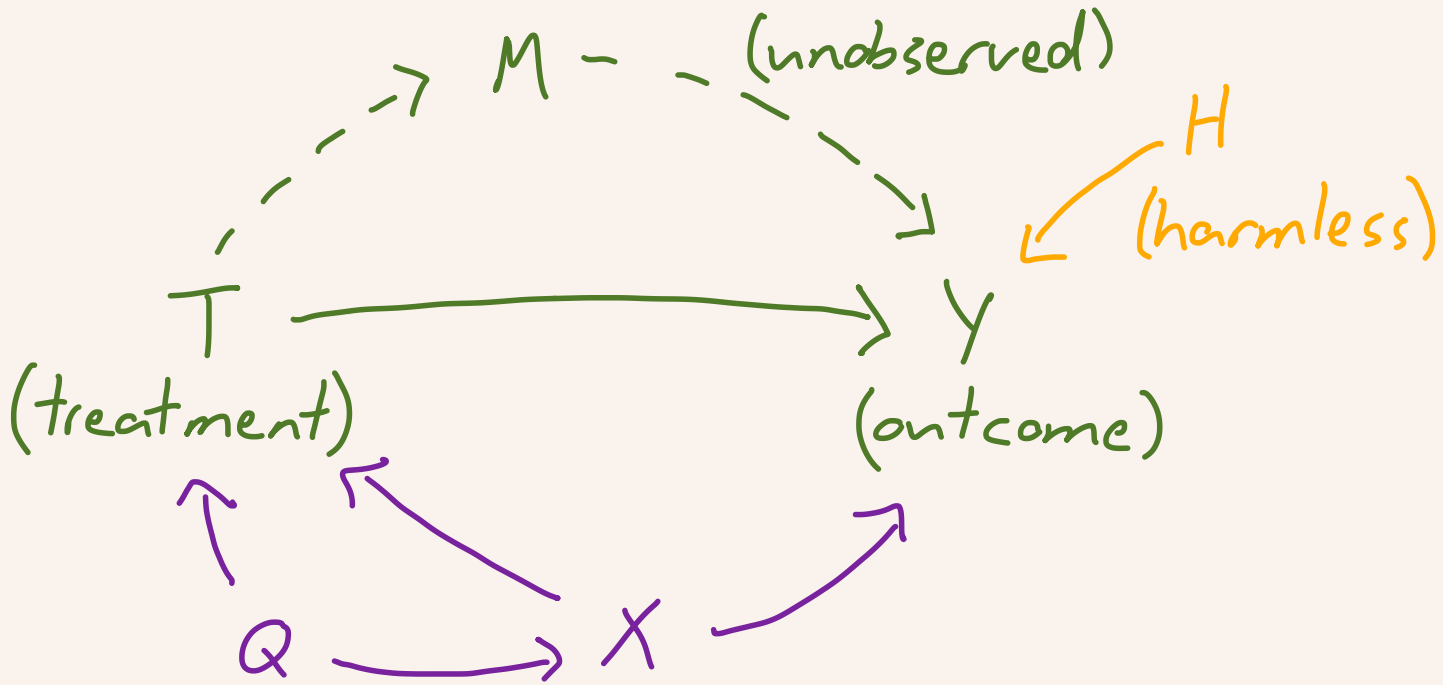

A confounding variable creates spurious correlations (non-causal paths) between two variables when omitted from a regression.

A path is *closed* once we control for/condition on/hold constant at least one variable along that path. When reporting regression output, $\hat{\beta}$ reflects all remaining *open* paths from T to Y.

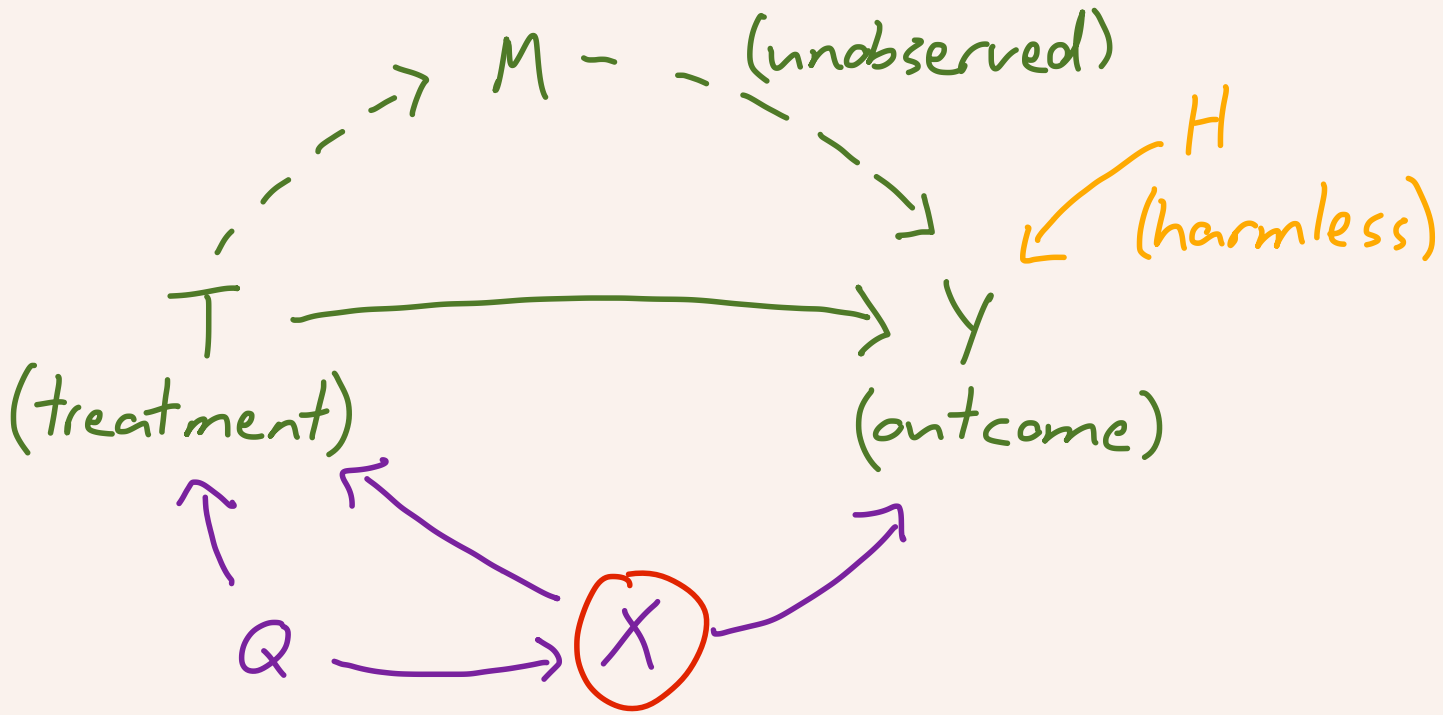

Can you determine a viable control strategy for measuring $T \rightarrow Y$?

The “Big Four” identification methods

(1) Matching Estimators: "closing your backdoors"

Regression automatically isolates the non-overlapping variation between regressors. It is our responsibility to determine whether or not this variation generates a measurement of a causal effect.

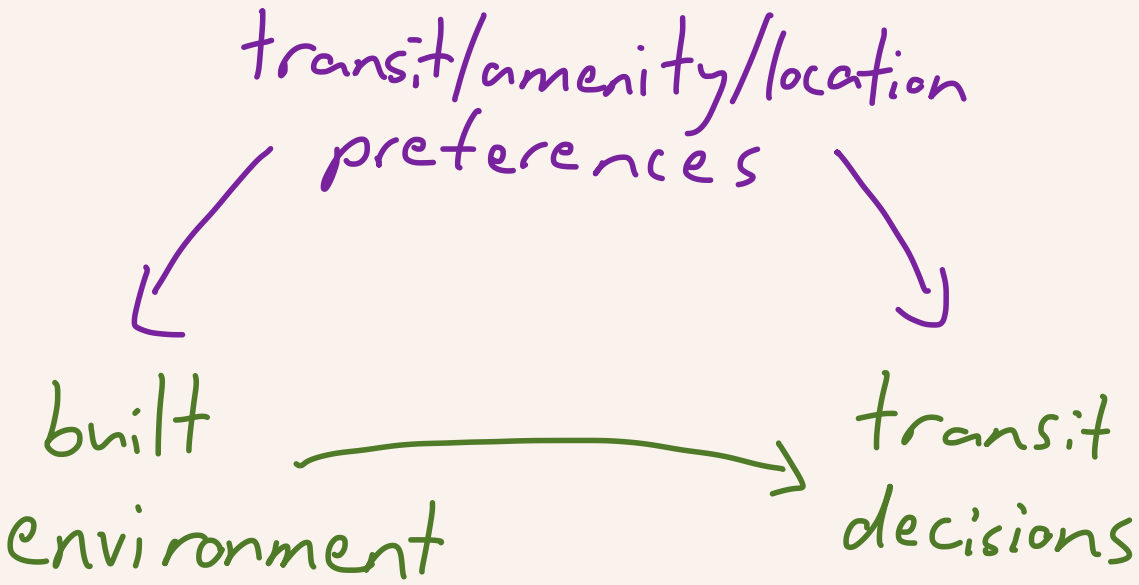

Illustrative example: Does the local built environment induce changes in transportation modality? (With a naïve regression, wouldn't transit and amenity preferences be confounders?)

Matching is more complicated than we think. Some researchers include "controls" only because they are available. Many of these decisions are made without any causal logic to back them up.

Illustrative example: *How* do droughts impact tree health?

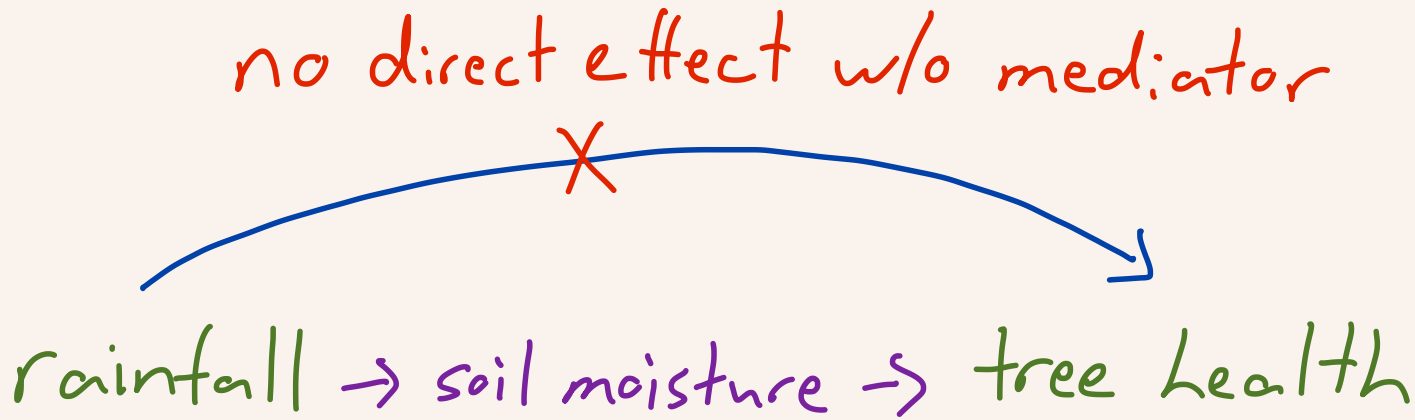

Over-conditioning: eliminating the variation that you wanted to exploit. Rainfall can't "reach" tree health without impacting soil moisture—but controlling for soil moisture closes this path.

Some control strategies can inadvertently bias estimators. For example, what if we want to estimate the impact of high PM2.5 levels on respiratory health, but use data from hospital admissions?

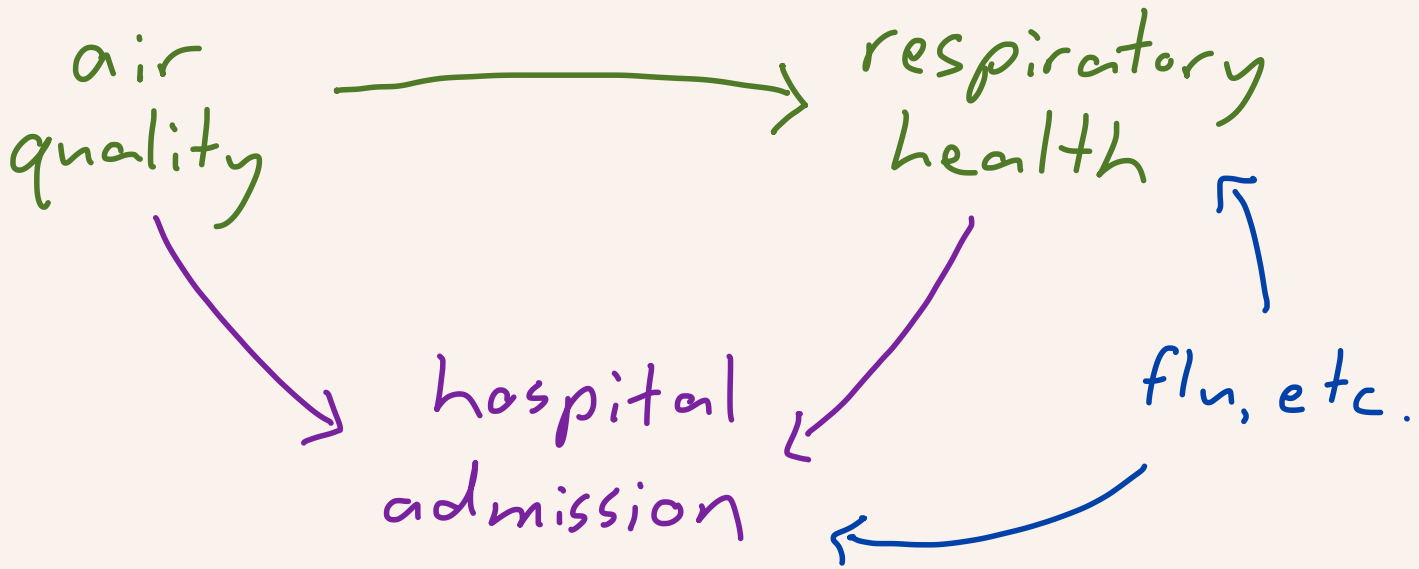

Some variables create spurious correlations between treatment and outcome *only once they are "controlled for."* Stratifying on a *collider* variable will open an otherwise closed path.

Mediators + colliders: Does adopting an electric vehicle decrease household emissions? Controlling the mediator closes a causal path (replacing a gas-powered car); controlling the collider opens a spurious one (forces a replacement). [under/over-estimate]

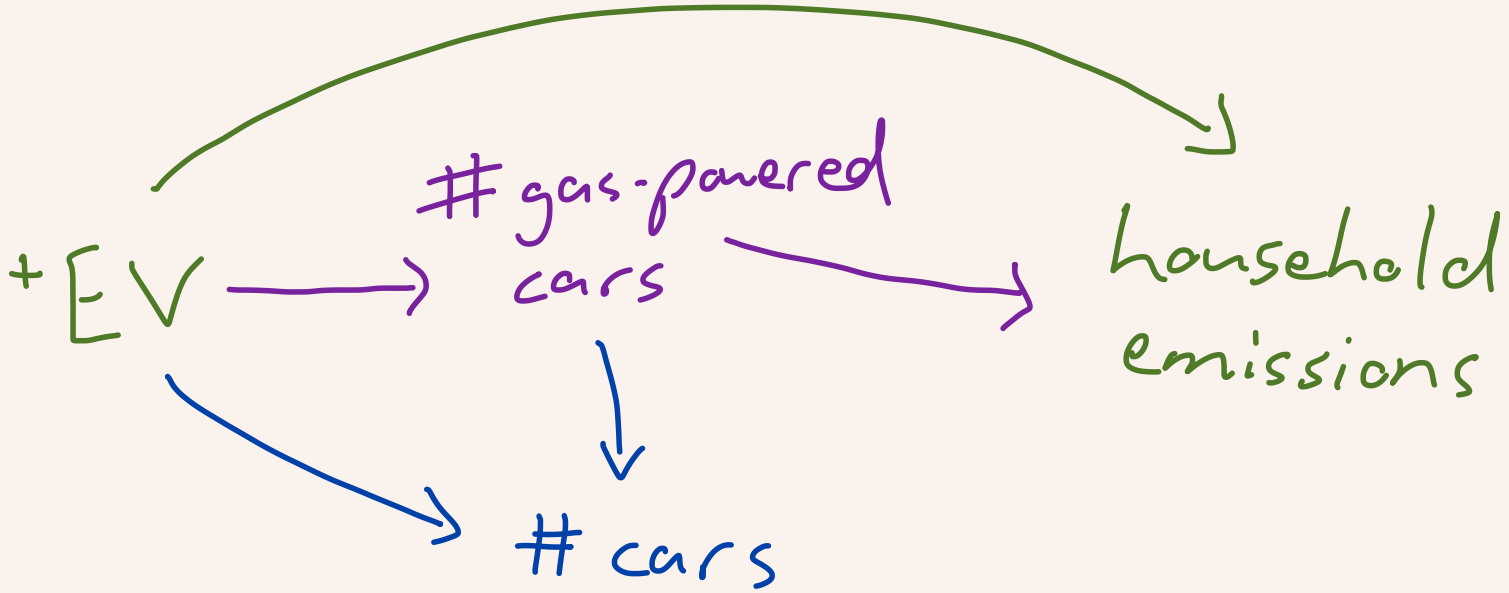

Matching is pretty hard. But the DAG alerts us to the potential pitfalls of an identification strategy. We don't need omniscience.

(2) Instrumental Variables: "colliders can be helpful, too"

How do you measure the price elasticity of demand for fish (Graddy, 2006)? You observe several [counterfactual] price-quantity pairs along a single demand curve.

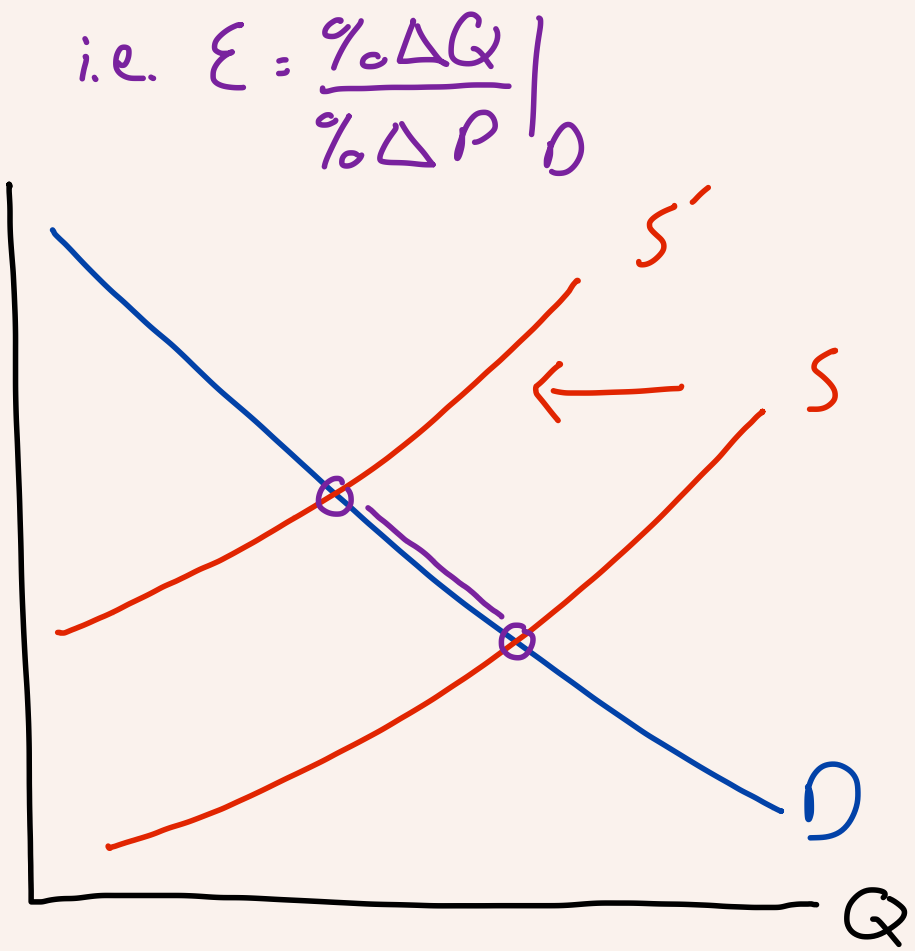

Two regressions, one fishy, one sound.

Dependent variable:	log(quantity sold)	
Independent variable	OLS	2SLS
log(price)	-0.549	-0.960
	(0.184)	(0.406)
Monday	-0.318	-0.322
	(0.227)	(0.225)
Tuesday	-0.684	-0.687
	(0.224)	(0.201)
Wednesday	-0.535	-0.520
	(0.221)	(0.219)
Thursday	0.068	0.106
	(0.251)	(0.232)
Time trend	-0.001	-0.003
	(0.003)	(0.003)
First stage for log(price)		
Wave height (feet)		0.103
		(0.022)
F statistic for IV		22.638

Graddy used the maximum wave height during recent fishing trips as an *instrument* for the price of fish at the Fulton fish market. More difficult fishing conditions drive a supply-side price shock.

$$IV: \begin{cases} price_{dt} = \pi + \underline{\lambda} \cdot wave\_height_{dt} + \gamma_d + \delta_t + \vartheta_{dt} \\ quantity\_sold_{dt} = \alpha + \underline{\eta} \cdot wave\_height_{dt} + \theta_d + \phi_t + \varepsilon_{dt} \end{cases}$$

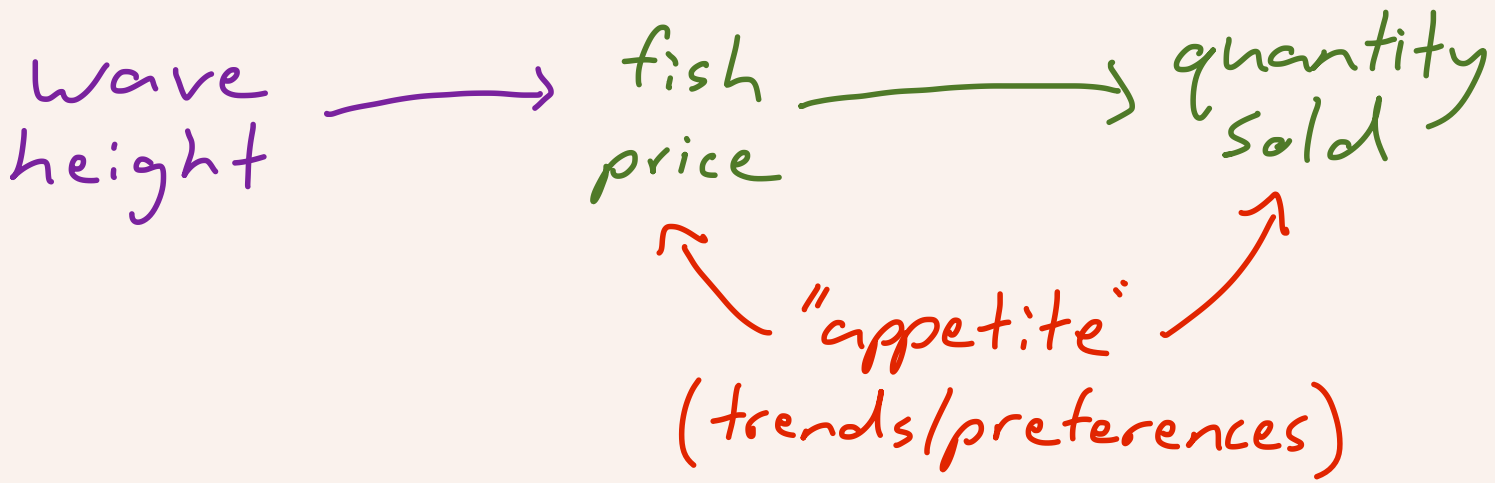

Our DAG turns price into a *collider*, which closes the spurious path between price and quantity! The coefficients $\hat{\lambda}$ and $\hat{\eta}$—and $\hat{\beta}_{IV} = \hat{\eta}/\hat{\lambda}$—now have a causal interpretation. ($\hat{\beta}_{2SLS}$ works too.)

The IV DAG reveals three identification assumptions which can be communicated transparently and evaluated with a cursory glance.

Our instrument must satisfy three assumptions:

1. *Relevancy* (i.e. have a causal impact on price)
2. *Independence* (i.e. no relation to demand shifts)
3. *Validity** (i.e. no impact on quantity sold except via price)

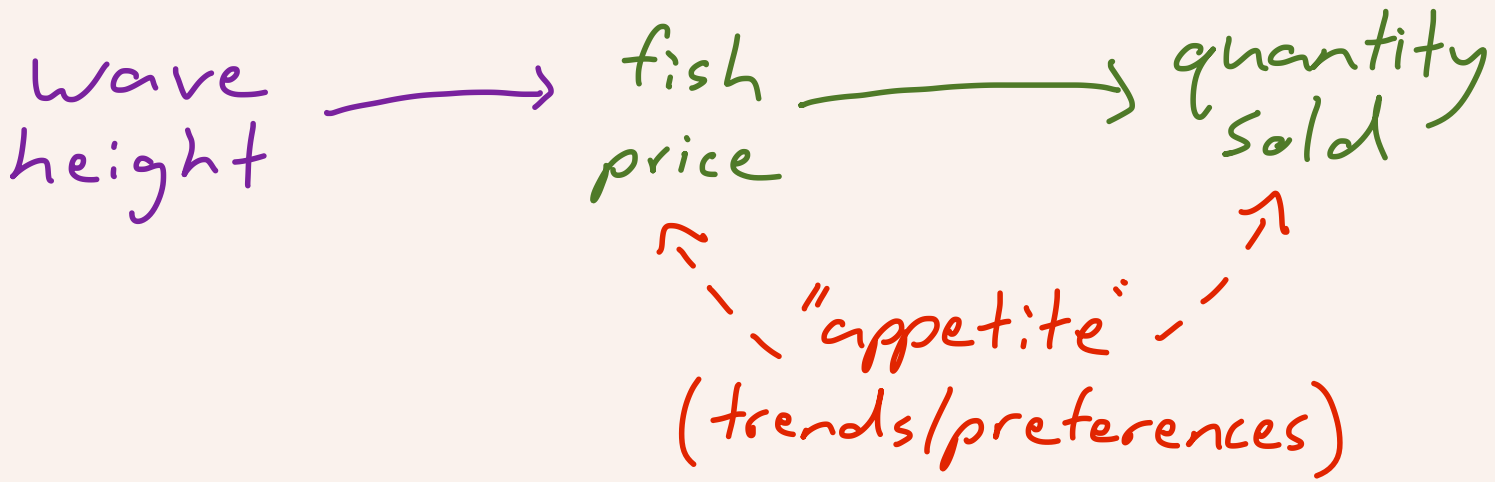

(3) Regression Discontinuity: "finding your offensive linemen"

An LA transit worker strike resulted in huge increases in congestion on roadways (Anderson, 2014). What makes us so sure?

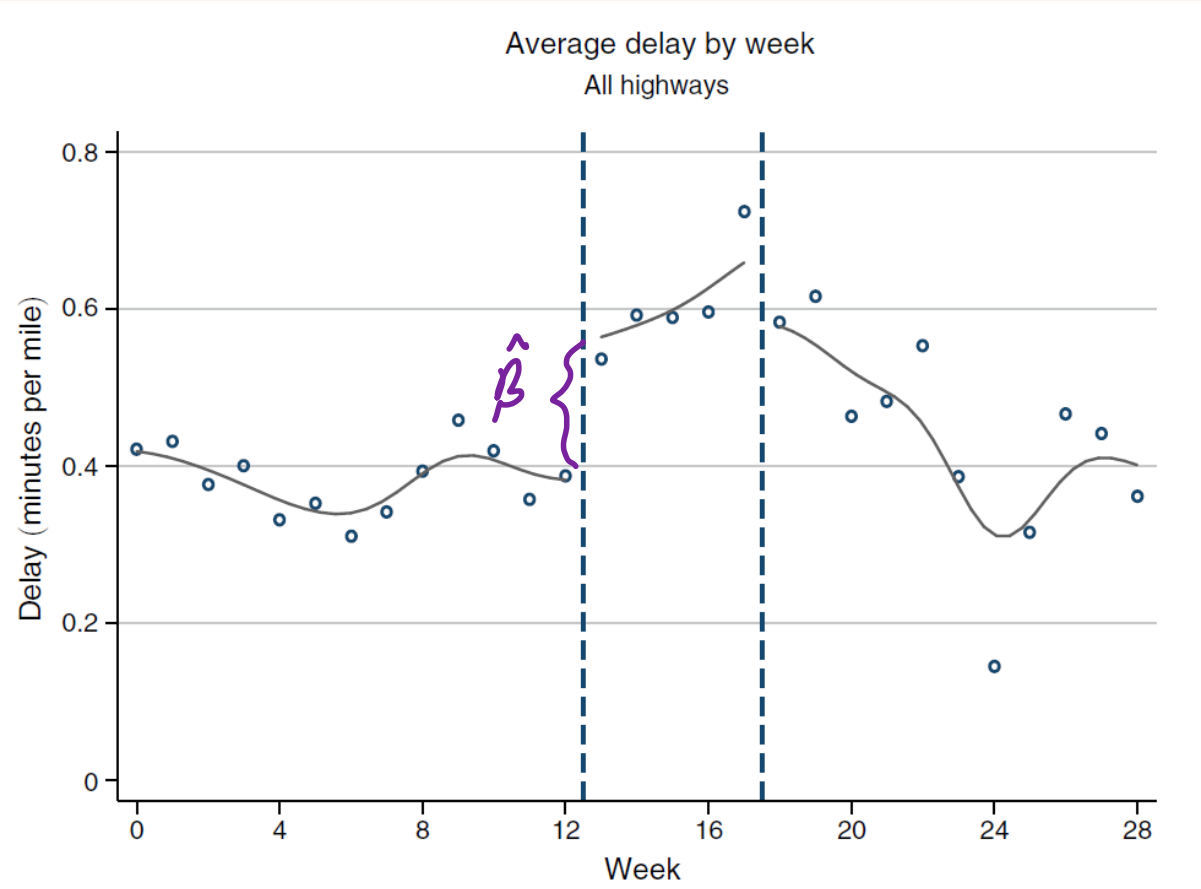

(because nothing else caused a jump in delays around the same time)

A running variable controls for the "smooth" variation in delay. Since the strike date is a function of *week*, all *continuous* [time-varying] BDPs are closed, leaving the discontinuity for $\hat{\beta}$.

OLS: $delay_t = \alpha + \beta \cdot strike_t + \gamma \cdot f(week_t) + \delta \cdot strike_t \cdot f(week_t) + \varepsilon_t$

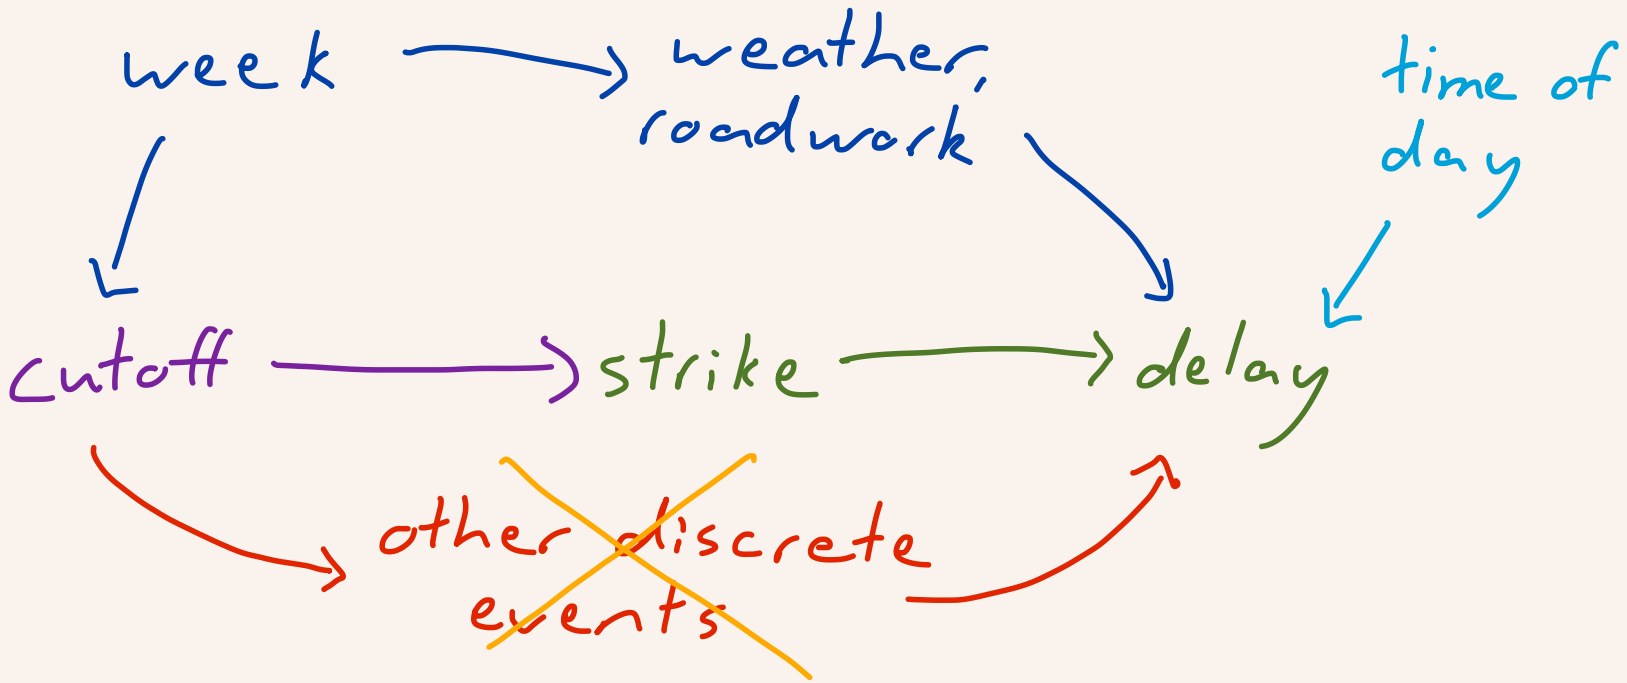

Because we want to attribute $\hat{\beta}$ solely to the strike, we assume that there are no other causal paths concurrent with the strike.

(4) Differences-in-Differences: "parallel trends and metonymy"

(fixed-effects)

The introduction of cell phones in Kerala seriously reduced price dispersion in fish markets (Jensen, 2007). How do we know?

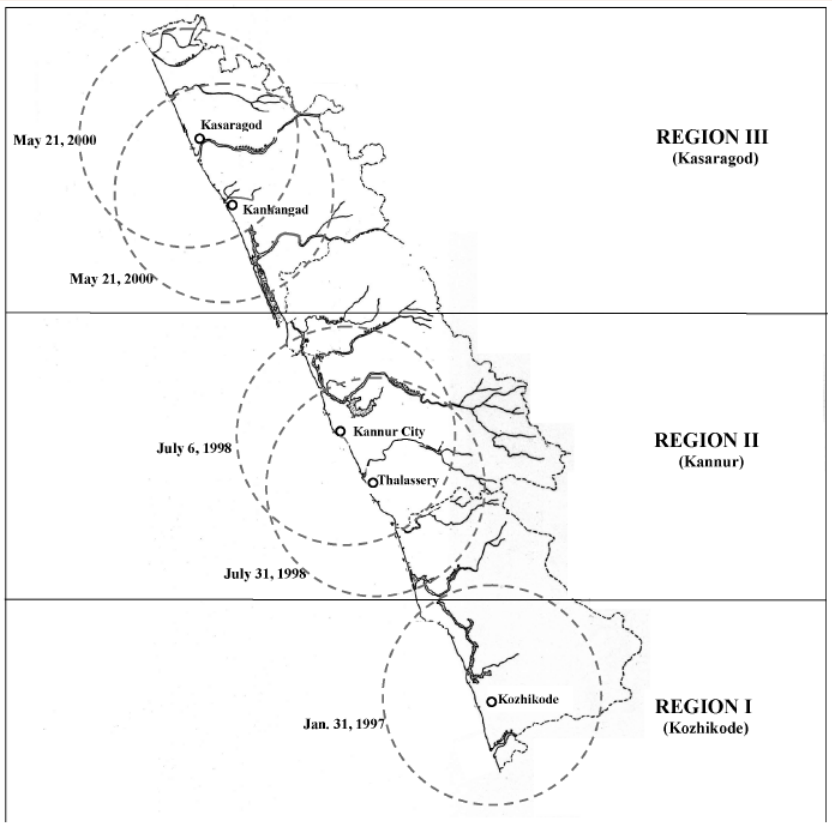

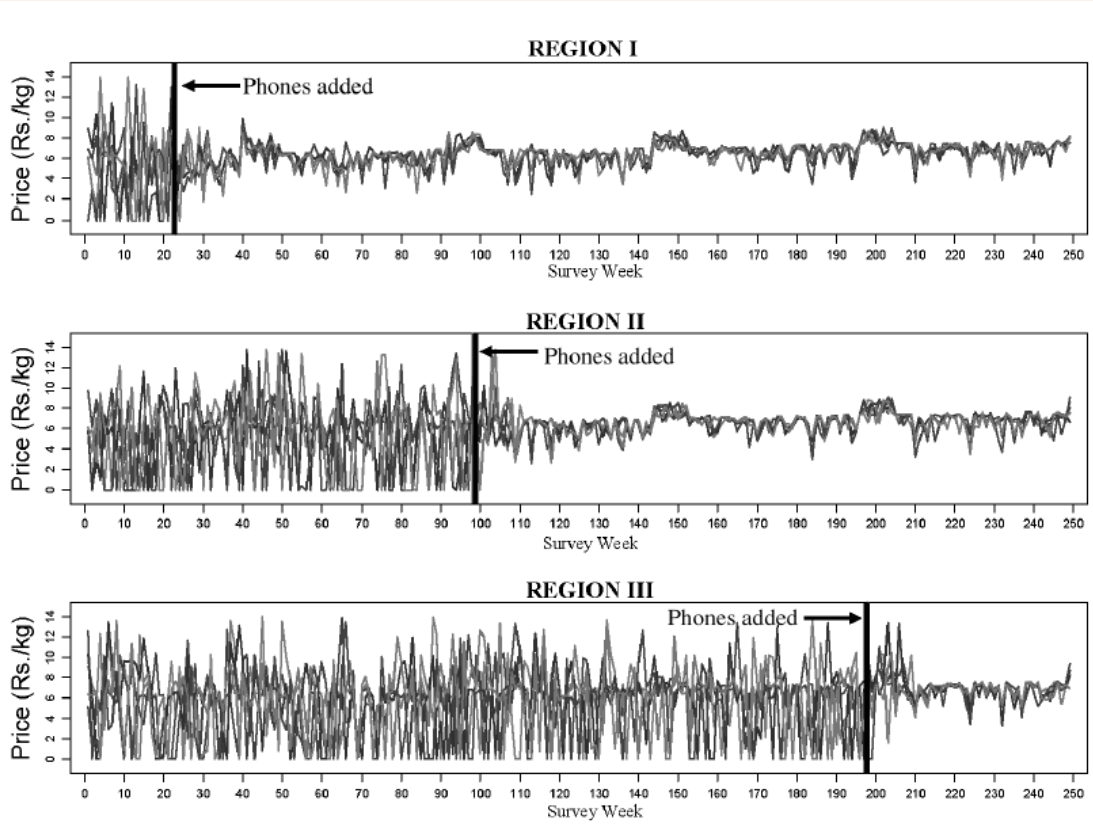

(b/c these coastal towns were on similar trajectories w/o intervention)

When treatments perfectly correlate with time, we cannot control for time without eliminating all variation in the treatment variable. Our solution? Give the treatment group a friend to compare to.

OLS: $spread_{it} = \alpha + \beta \cdot treat_i + \gamma \cdot post_t + \delta \cdot treat_i \cdot post_t + \varepsilon_{it}$

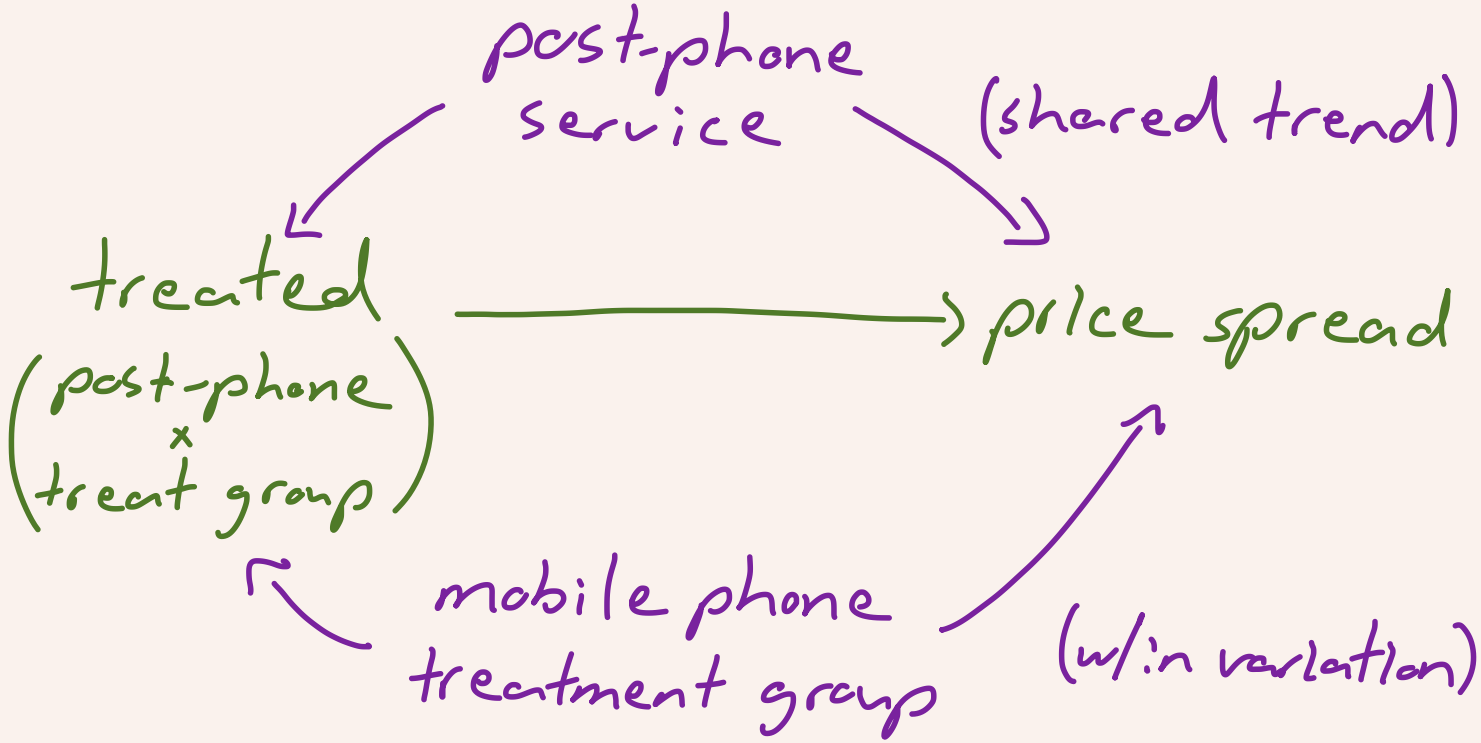

The DiD DAG suggests a *parallel trends* assumption: w/o treatment, the two groups would move in-step (effect attribution).

(5?) The Front-Door Criterion: "causality and the chain rule"

In the presence of confounders, we can decompose an unobserved relationship $T \to Y$ into two separate causally-interpretable measurements when an exogenous mediator M is available.

$$SUR: \begin{cases} M_i = \pi + \lambda \cdot T_i + \vartheta_i \\ Y_i = \alpha + \gamma \cdot M_i + \delta \cdot T_i + \varepsilon_i \end{cases}$$

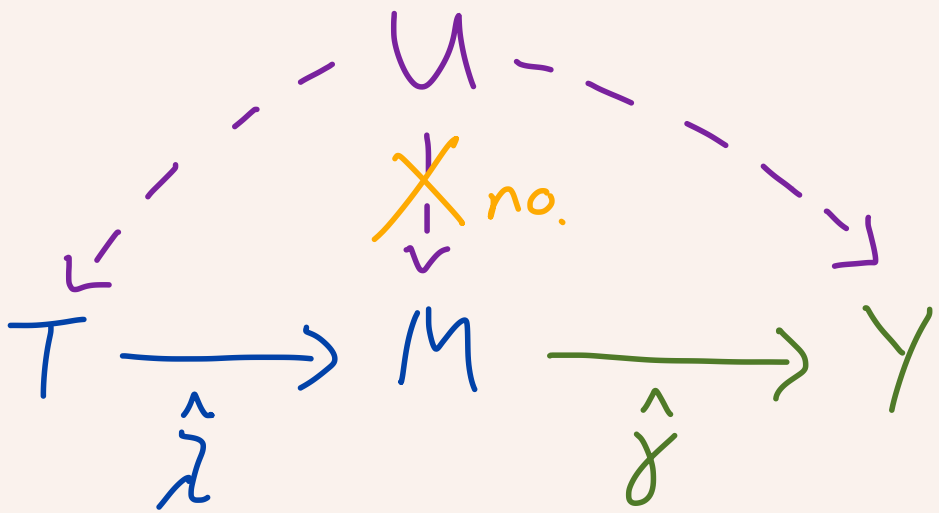

For $\hat{\beta}_{FDC} = \hat{\lambda} \cdot \hat{\gamma}$ to be causal, the FDC DAG suggests three new assumptions: (1) T only impacts Y via M, (2) T and M aren't confounded, (3) Y and M aren't confounded after controlling T.

Current FDC work: The efficacy of cloud seeding programs is not well-understood, despite the gobs of money thrown at it and the amount of data that's already been collected (bad seeding rules).

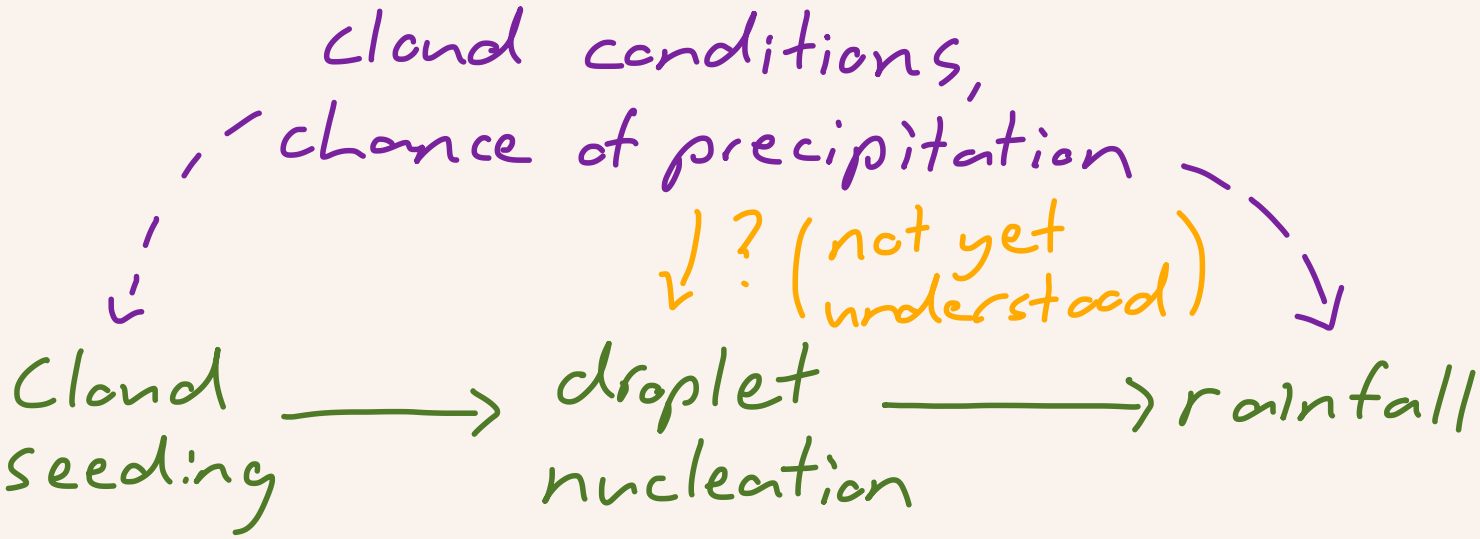

I'm working on [conditionally] verifying these three claims:

1. seeding only impacts rainfall via nucleation
2. seeding and nucleation are not confounded
3. rainfall and nucleation are not confounded after conditioning on seeding

Plausibly-exogenous?

Thank you!

pdonovan@conncoll.edu

[†] Sony Pictures, *Snatch* (2000)