A CALCULUS FOR SOFT INTERVENTIONS

Juan D. Correa <jdcorrea@cs.columbia.edu>
joint work with **Elias Bareinboim** <eb@cs.columbia.edu>

based on our paper in AAAI-20

COLUMBIA UNIVERSITY
IN THE CITY OF NEW YORK

Causal Data Science Meeting
Nov, 2020

Motivating Example - Tutoring Program

- For the students we observe their GPA at the beginning of the term, their motivation (low, high), whether they get tutoring or not, and their GPA at the end of the semester.

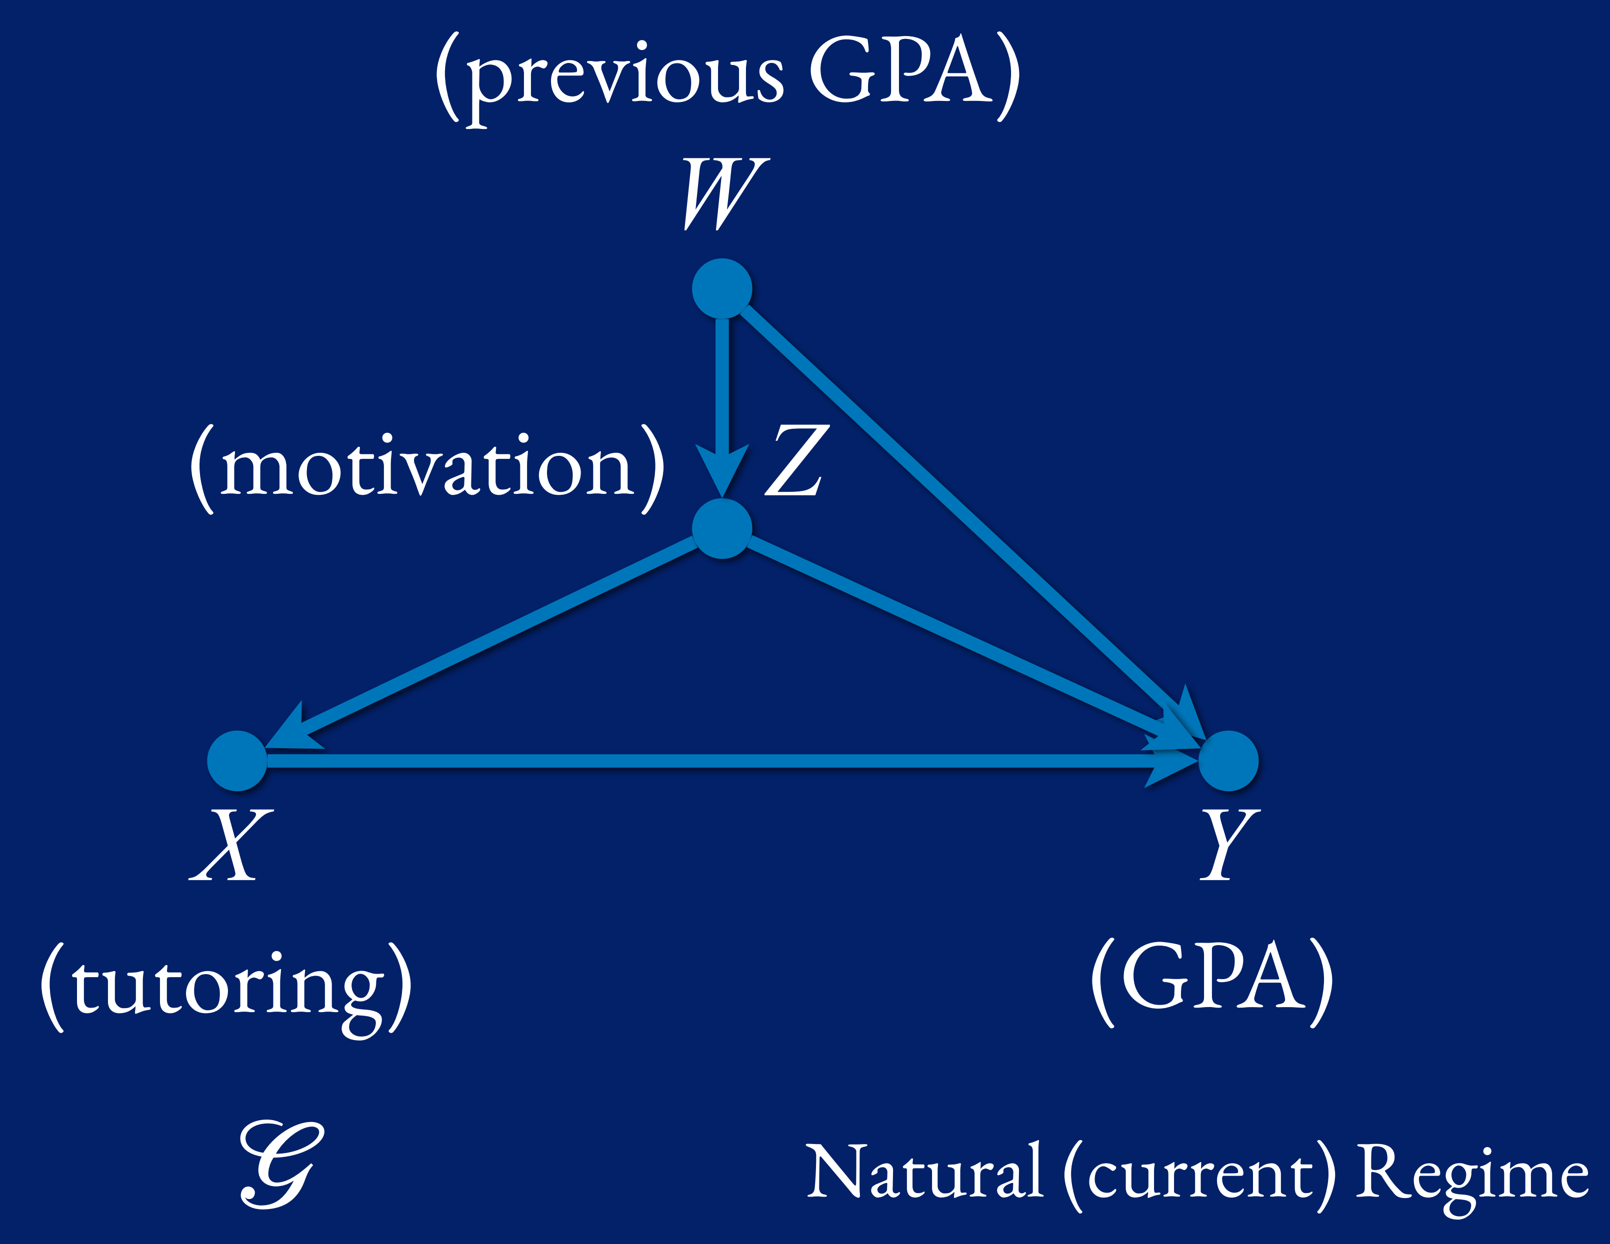

- Using machine learning, and with enough data, a students GPA can be predicted with small error given other features i.e., $P(y \mid w, z, x)$.
- However, this data reflects the current/natural regime, yet we aim to assess the impact of a new unobserved policy (intervention) on the students GPA.

Motivating Example - Tutoring Program

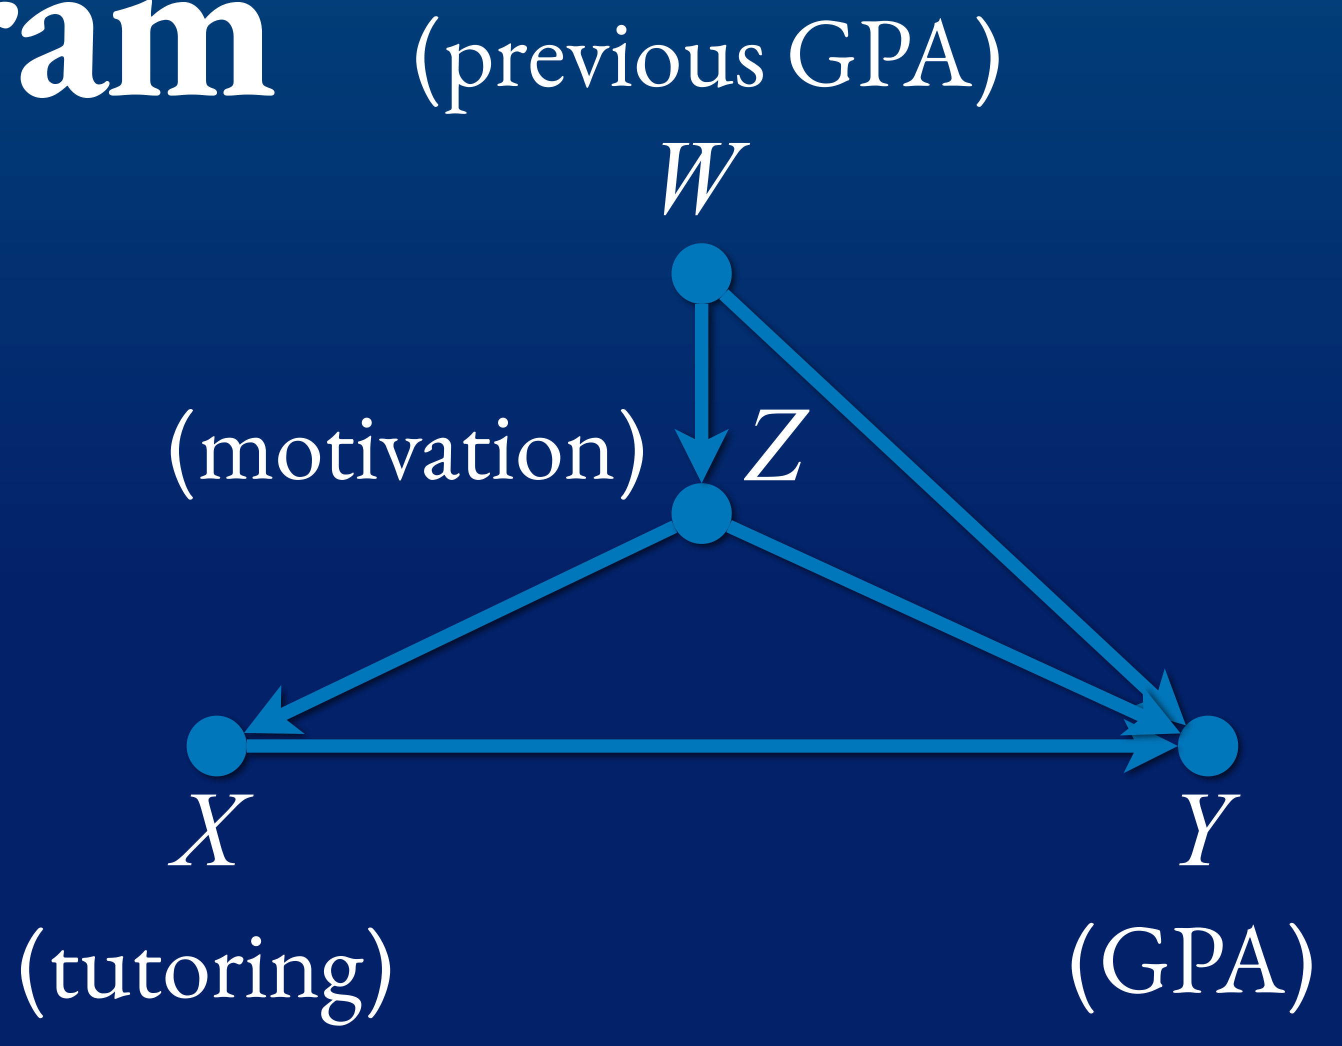

$\mathcal{G}$ Natural (current) Regime

- What can be inferred from $P(W, Z, X, Y)$ and the causal graph in terms of causal effects?
- Possibly, the causal effect of X on Y, that is:

$$P(y \mid do(x)) = \sum_z P(y \mid do(x), z) P(z \mid do(x))$$

Condition on Z

$$= \sum_z P(y \mid do(x), z) P(z)$$

$$= \sum_z P(y \mid x, z) P(z)$$

Or, simply note that Z is backdoor admissible relative to (X, Y).

X

is the same for Y

Motivating Example - Tutoring Program

- What does $do(X = 1)$ mean in the real-world?
  - Make tutoring mandatory for every student

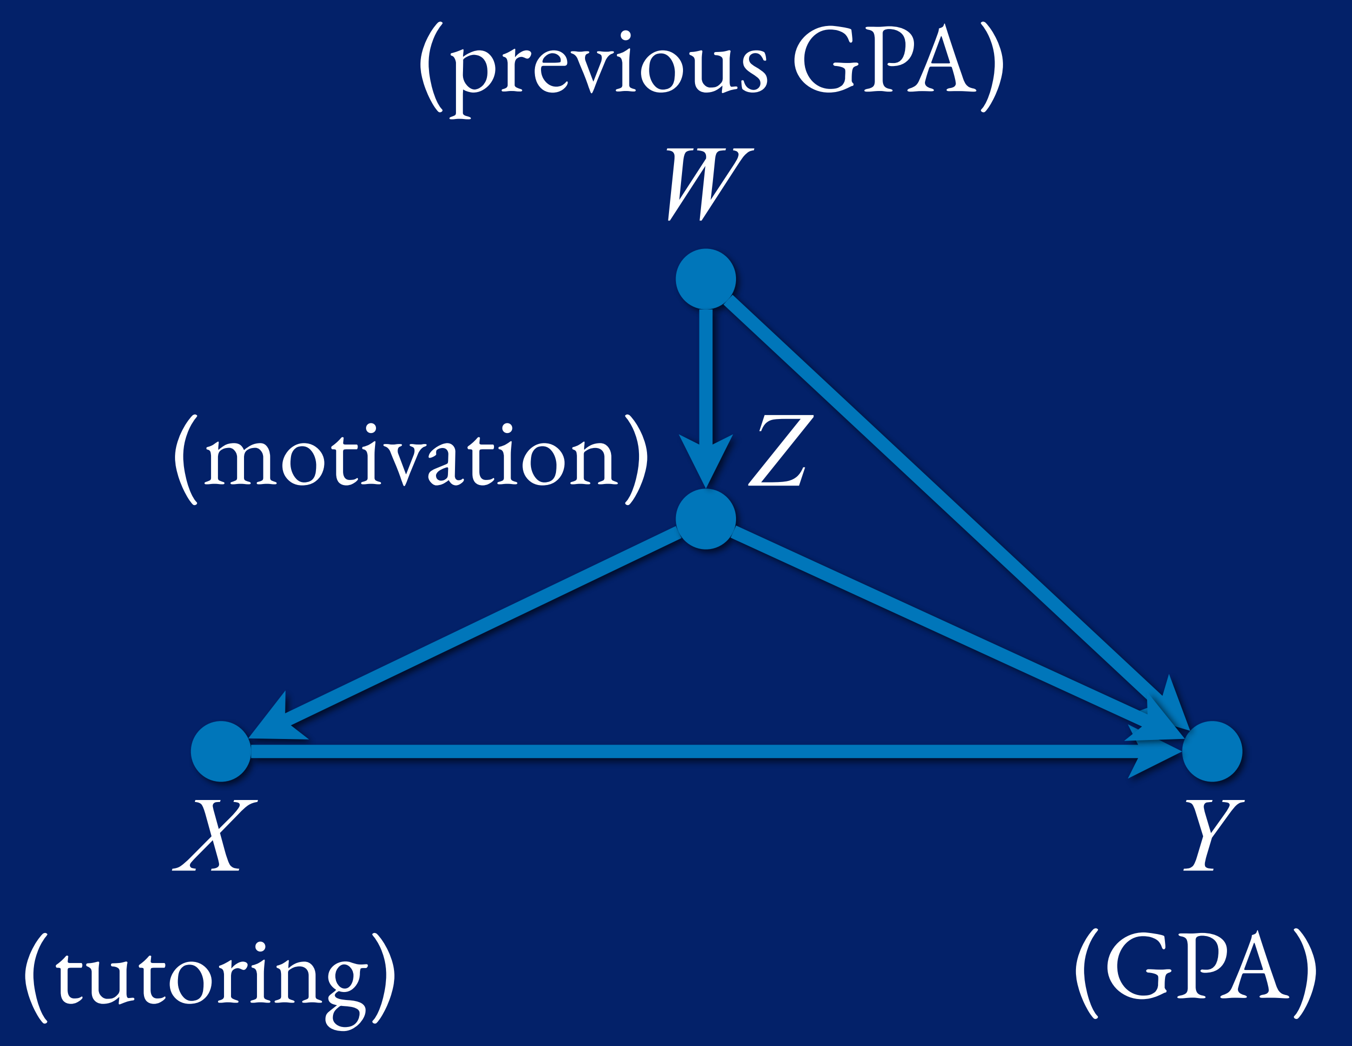

$\mathcal{G}$ Natural (current) Regime

Intervention

$do(X = 1)$

Make tutoring mandatory for all students.

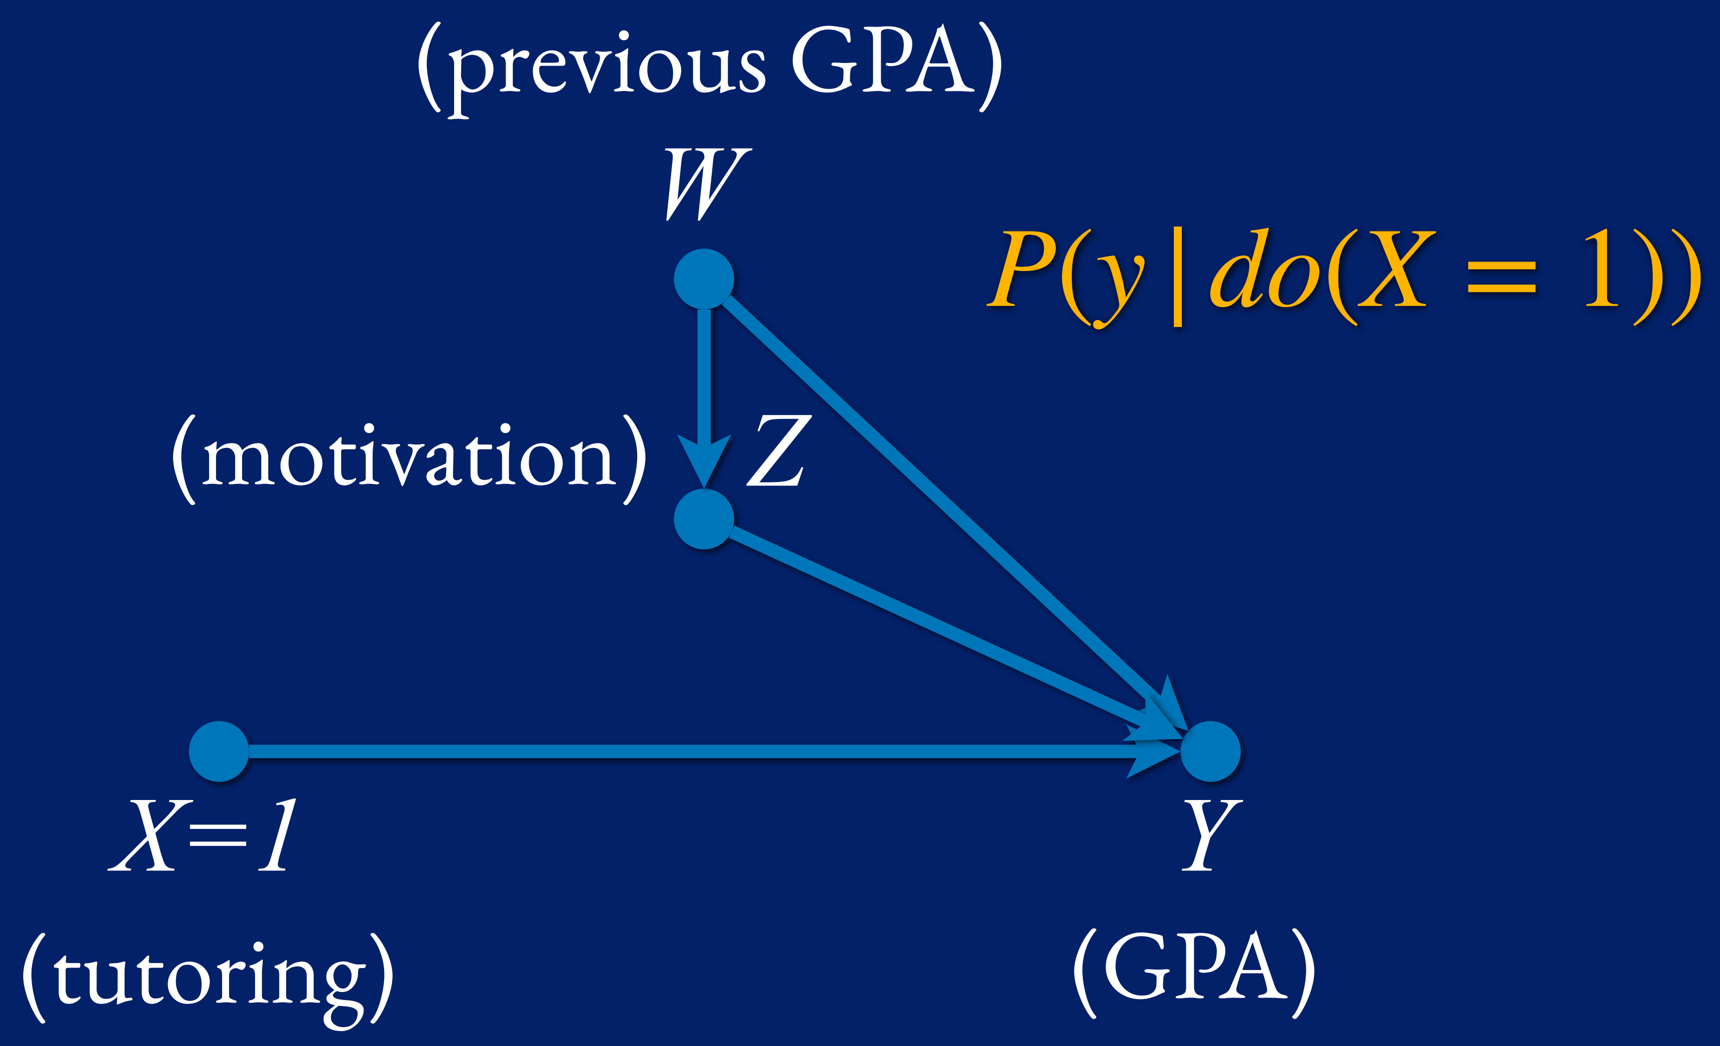

$\mathcal{G}_{\overline{X}}$ Intervened (hypothesized) Regime

Implementation of do()-interventions

In decision making scenarios, even if the effect of a do() intervention is identifiable …

- Available resources may be insufficient to implement the corresponding policy.
  - There are no enough teachers to cover all the hours of tutoring needed for every single student in the school.

- Effectiveness of the intervention cannot be guaranteed:
  - Patients assigned treatment may not follow it.

Implementation of do()-interventions (cont)

- For practical purposes, one may care about the effect of realizable interventions.

Do-like Intervention	Realistic Intervention
Make sure no one smokes	Reduce tabaco consumption to 20% of current consumption
Provide treatment to all patients	Administer the treatment if and only if patient is in a critical condition
Move a robotic arm exactly to coordinates (X, Y, Z)	Move arm to (X, Y, Z) w/ normally dist. error (considering physical constraints)
Make all applicants male	Mark all applicants as males (on paper)

Some Canonical types of Interventions [Dawid 02, Tian 08]

- **Hard/atomic**: $\sigma_X = do(X = x)$ set variable X to a constant value x.
  (Do-calculus original treatment considered mostly this type of intervention.)
  - Every student gets tutoring.
- **Conditional**: $\sigma_X = g(w)$ sets the variable X to output the result of a function g that depends on a set of observable variables W.
  - Students get tutoring if and only if they have a low GPA.
- **Stochastic**: $\sigma_X = P^*(x \mid w)$ sets the variable X to follow a given probability distribution conditional on a set of variables W.
  - Students with low GPA enter a raffle for 80% of the spots, other interested students enter for the remaining 20%.

A more realistic intervention

- Suppose $P(X = 1 \mid Z = 1) = 0.3$, that is, under normal conditions, only 30% of motivated students get tutoring.
- What is the effect of making it 60%? That is, $P^*(X$

Regime node used to encode the fact that X has been intervened on.

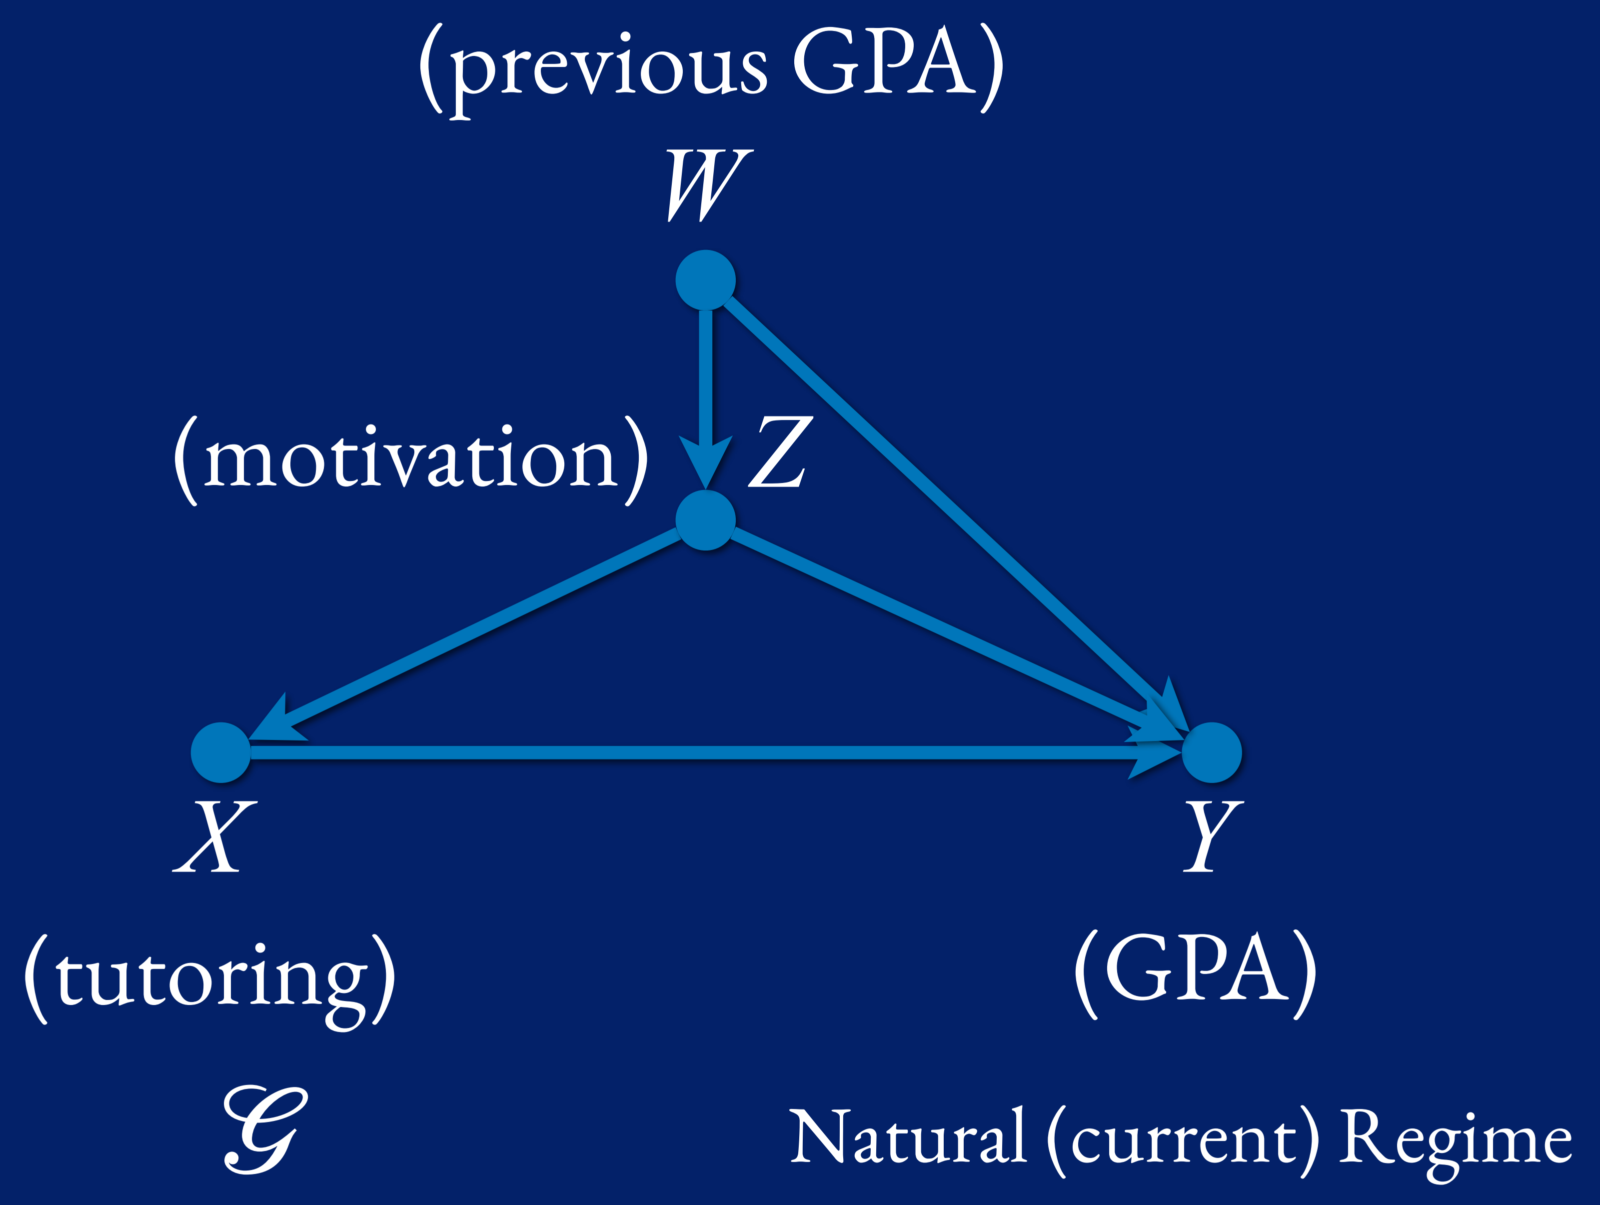

Intervention

$\sigma_X = P^*(X \mid Z)$

Double participation rate in tutoring program for motivated students.

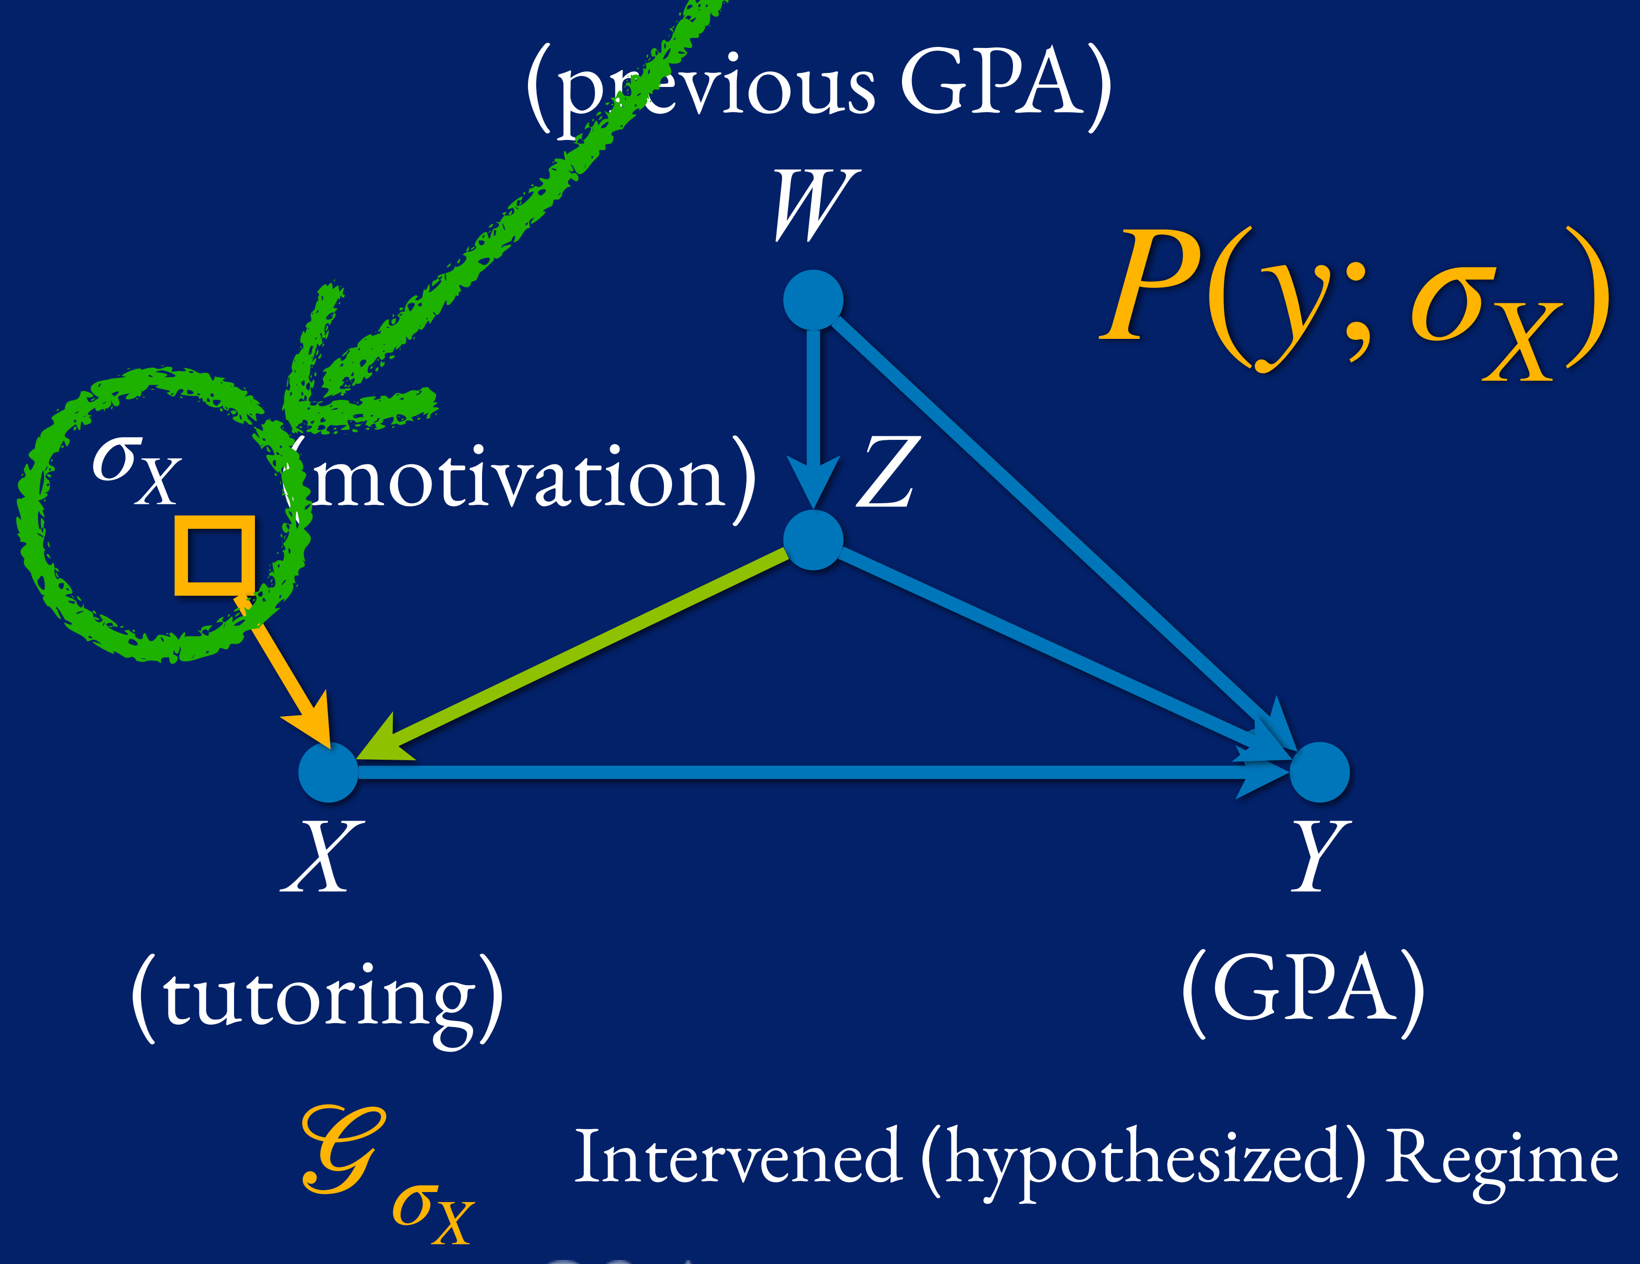

Using do-calculus intuition

$$P(y;\sigma_X) = \sum_z P(y \mid z;\sigma_X)P(z;\sigma_X)$$ Condition on Z

$$= \sum_z P(y \mid z;\sigma_X)P(z)$$ Rule 3 ($Z \perp\!\!\!\perp X$) in $\mathcal{G}_{\overline{X}}$

$$= \sum_z \sum_x P(y \mid z, x;\sigma_X)P(x \mid z;\sigma_X)P(z)$$ Condition on X

$$= \sum_z \sum_x P(y \mid z, x;\sigma_X)P^*(x \mid z)P(z)$$ Definition of

$$= \sum_z \sum_x P(y \mid z, x)P^*(x \mid z)P(z)$$ Rule 2 ($Y \perp\!\!\!\perp X \mid Z$) in $\mathcal{G}_{\underline{X}}$

Estimable from current regime

Defined by σ_X

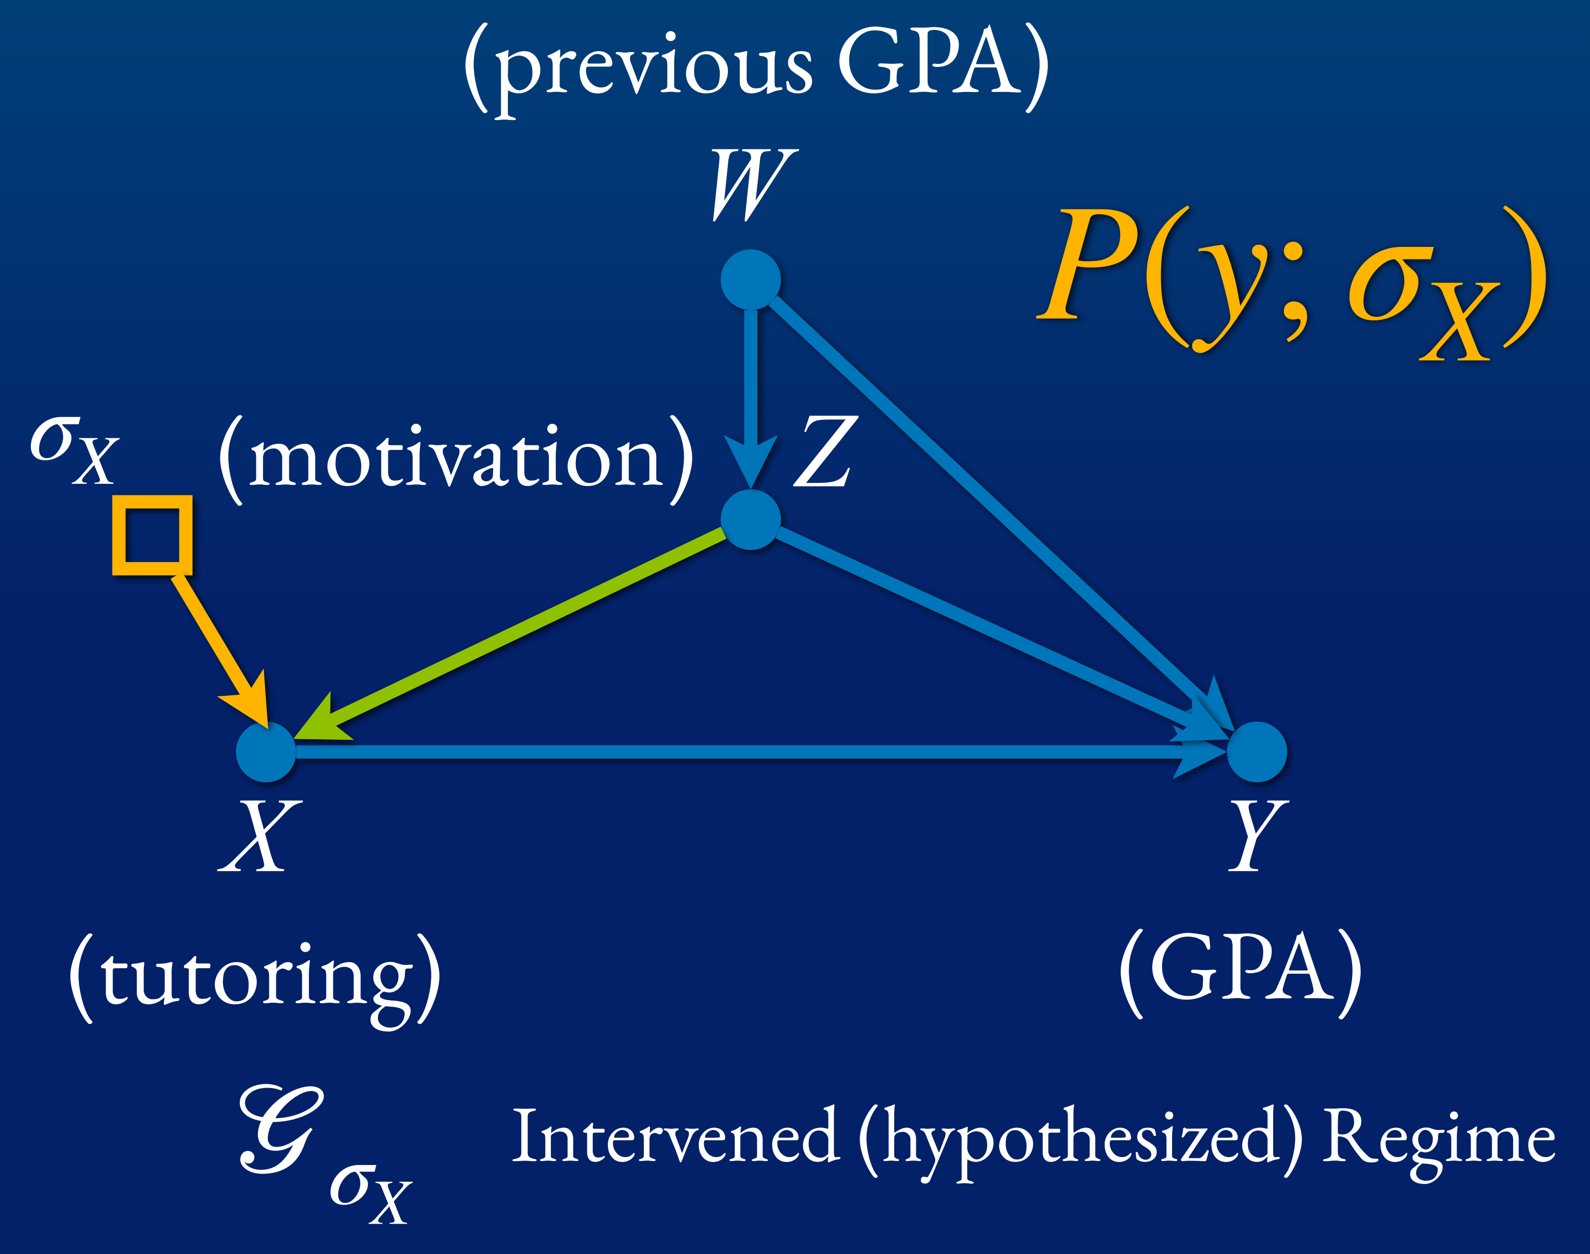

Is this derivation strategy (do-calculus-like) sufficient to solve the problem?

Identifying the effect of soft interventions

- Although the identification of the effect of soft interventions can be reduced to identification of atomic interventions [Pearl 2000, Tian 2008], there is a gap in terms of end-to-end derivations, based on rules akin to do-calculus.
- Such rules allow for a better understanding not only of the assumptions entailed by the graphical model, but also for building intuition of the identification procedure.
- Next, we will see an example where this conceptual gap leads to an incorrect conclusion.

A classical example [Pearl and Robins, 95]

- Consider a situation in which a patient receives a sequence of treatments (for now, say two times).

- After the first treatment X_1, a second physician checks the patient (and observes $Z = z$), and then decide on a second treatment X_2.

- Finally the patient may survive or not; $Y = 1$ or $Y = 0$, respectively.

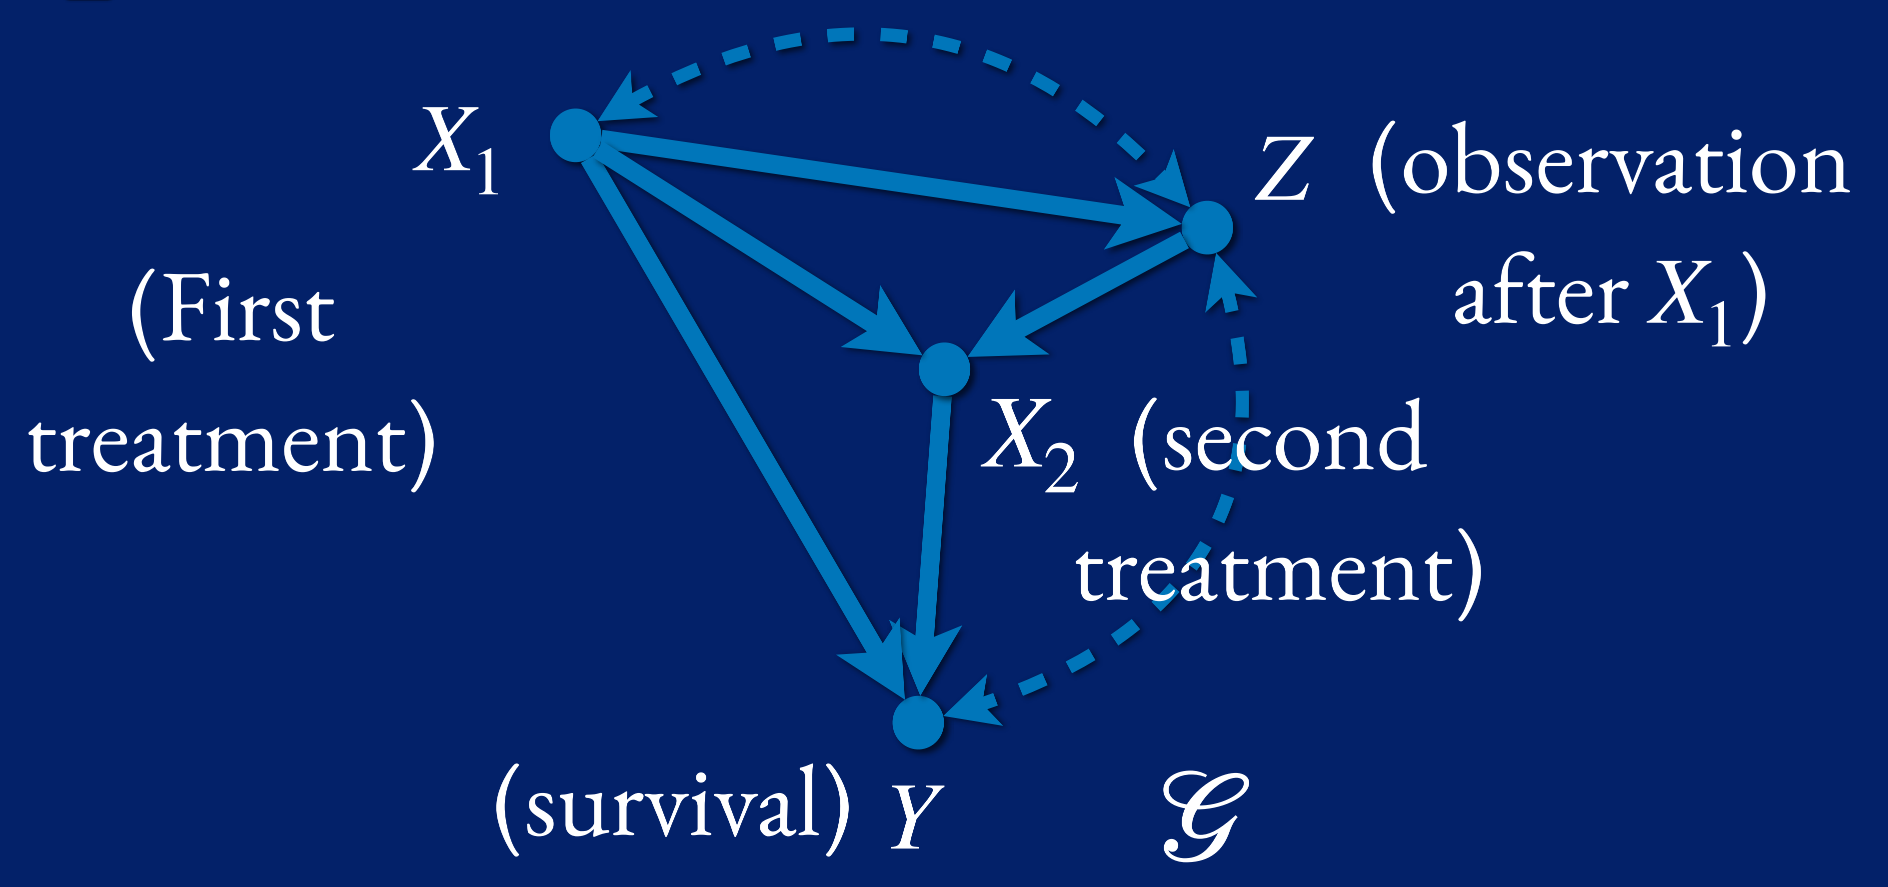

A classical example [Pearl and Robins, 95]

- What is the effect of an intervention when we fix $X_1 = x_1$ but $X_2 = g(x_1, z)$, that is, X_2 is prescribed depending on what x_1 was and the observation Z.

- This can be written as $do(x_1), do(x_2 = g(x_1, z))$ or $\sigma_X = \{X_1 = x_1, X_2 = g(x_1, z)\}$.

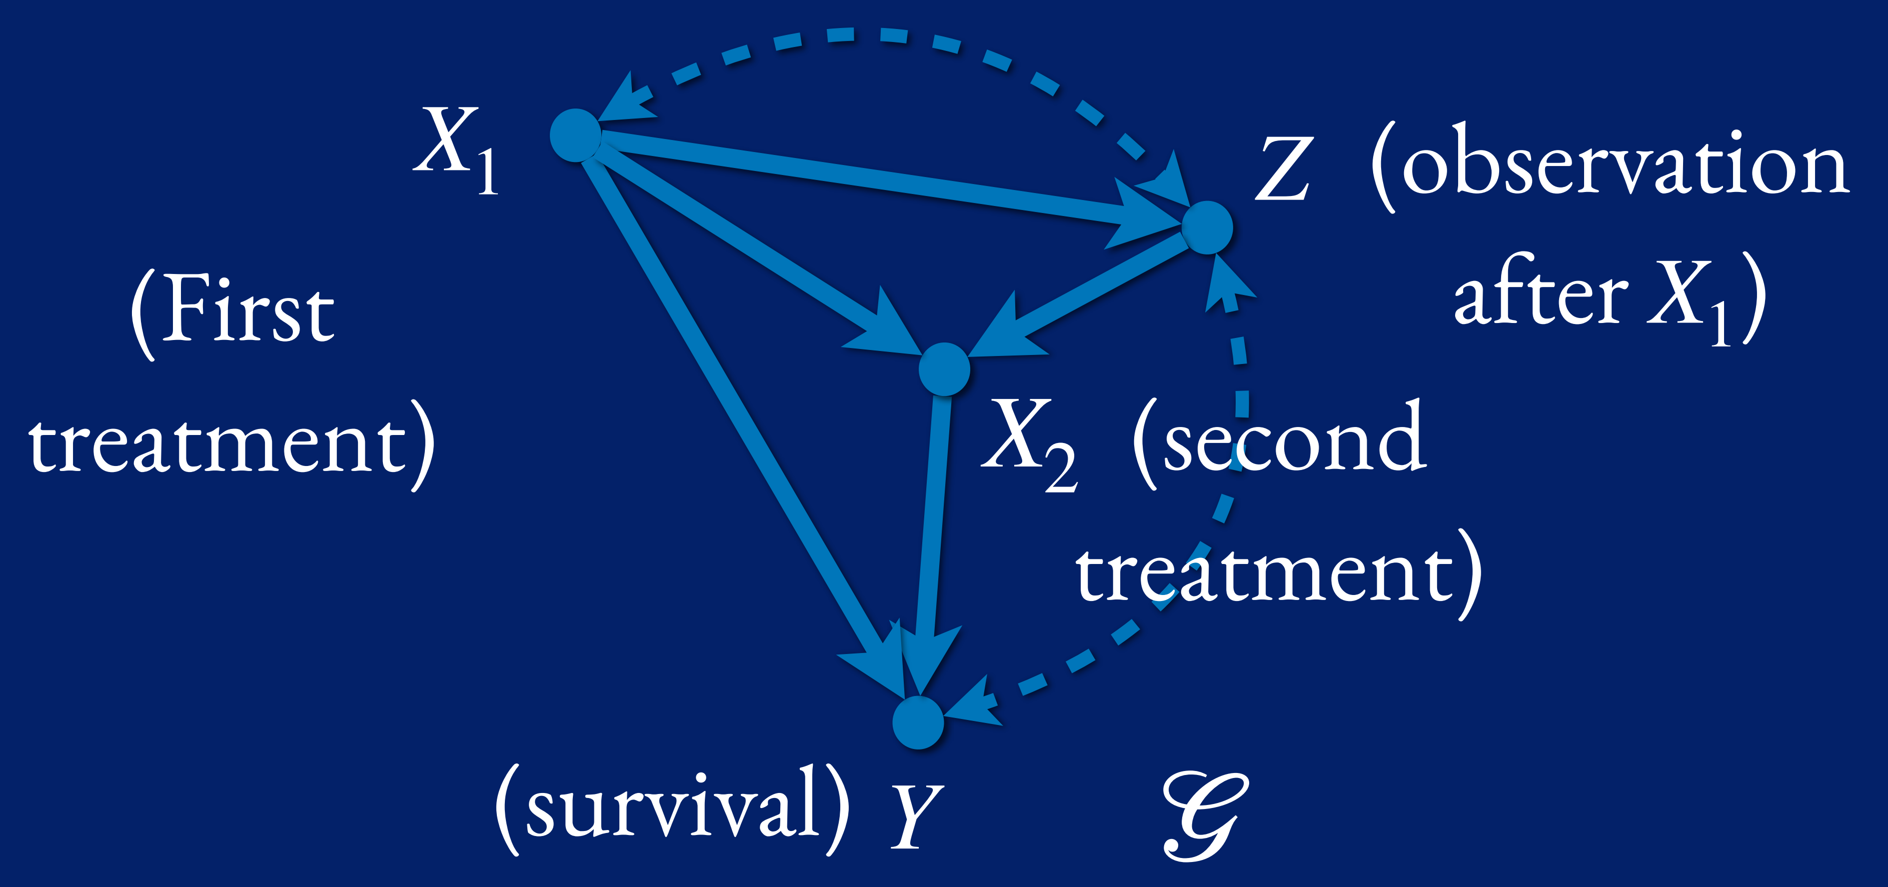

A classical example [Pearl and Robins, 95]

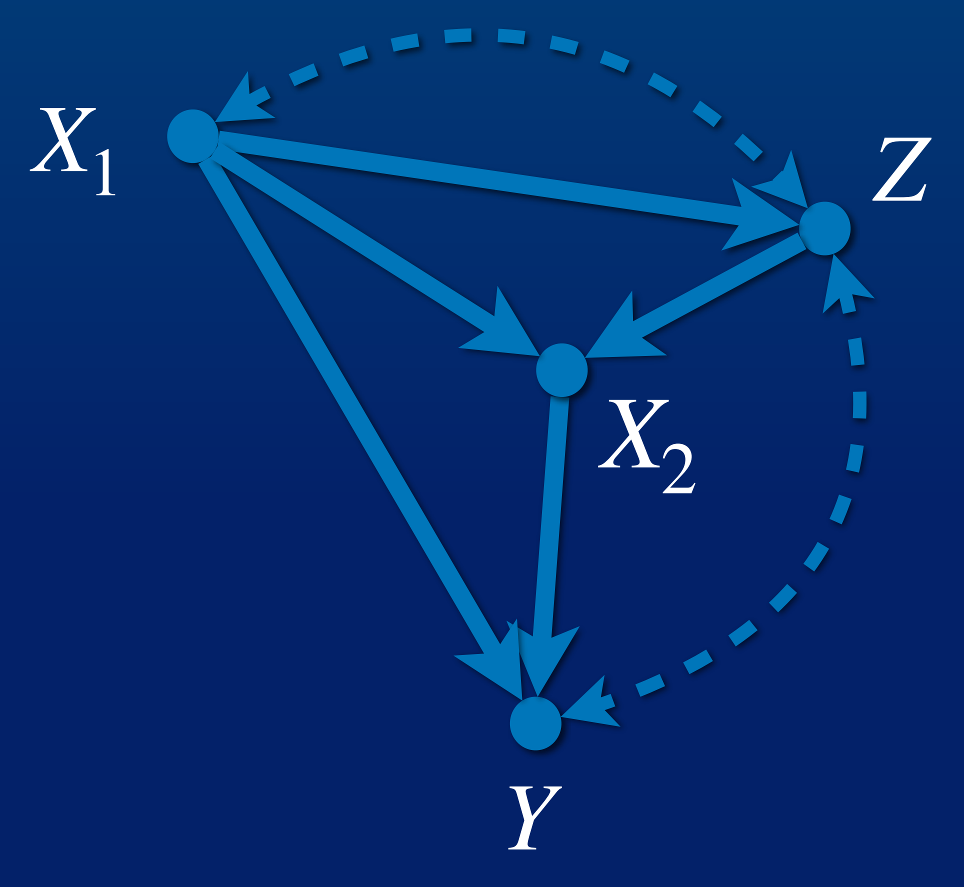

- We could try to identify this as with do-interventions:

$$P(y \mid do(x_1), do(X_2 = g(x_1, z)))$$

$$= P(y \mid x_1, do(X_2 = g(x_1, z)))$$ Rule 2 ($Y \perp\!\!\!\perp X_1 \mid X_2$) in $\mathcal{G}_{\overline{X_2}\underline{X_1}}$

$$= \sum_z P(y \mid x_1, do(X_2 = g(x_1, z)), z) P(z \mid x_1, do(X_2 = g(x_1, z)))$$ Condition on Z

$$= \sum_z P(y \mid x_1, do(X_2 = g(x_1, z)), z) P(z$$

Turns out this effect is not identifiable.
What went wrong with the derivation?

$$= \sum_z P(y \mid x_1, x_2, z)\big|_{x_2 = g(x_1, z)} P(z \mid x_1)$$ Definition of σ_X

A classical example [Pearl and Robins, 95]

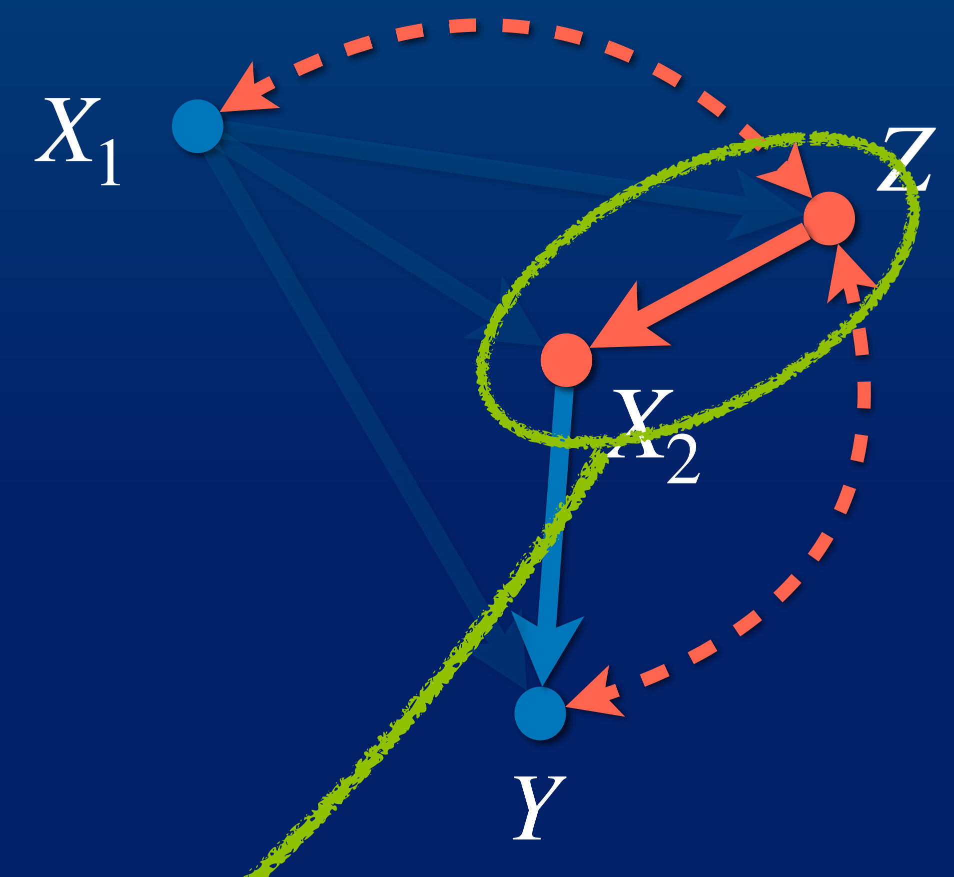

- We could try to identify this as with do-interventions:

$$P(y \mid do(x_1), do(X_2 = g(x_1, z)))$$

$$= P(y \mid x_1, do(X_2 = g(x_1, z))) \quad \text{Rule 2 } (Y \perp\!\!\!\perp X1 \mid X2) \text{ in } \mathcal{G}_{\overline{X_2}\underline{X_1}}$$

- Under $do(x_2 = g(x_1, z))$, the edges incoming to X_2 are still active, hence $\mathcal{G}_{\overline{X_2}\underline{X_1}}$ is not the right graph to look at.
- This rule application needs to be considered w.r.t. $\mathcal{G}_{\underline{X_1}}$, where the separation does not hold due to the fact that, given the intervention on X_2, Z becomes an active collider opening a bidirected path from X_1 to Y.

Where do-calculus intuition breaks

For do-interventions, any dependence between the intervened variable and its parents disappear, but for soft interventions may:

- keep all or some dependences with parents,
- change the distribution of the variable given its parents, or
- even add new dependences (new parents)

Rules of σ- calculus

- Rule 1
  - do-calculus

$$P(\mathbf{y} \mid do(\mathbf{x}), \mathbf{w}, \mathbf{t}) = P(\mathbf{y} \mid do(\mathbf{x}), \mathbf{w}) \quad \text{if } (\mathbf{Y} \perp \mathbf{T} \mid \mathbf{W}, \mathbf{X}) \quad \text{in} \quad \mathcal{G}_{\overline{\mathbf{X}}}$$

  - σ-calculus

$$P(\mathbf{y} \mid \mathbf{w}, \mathbf{t}; \sigma_{\mathbf{X}}) = P(\mathbf{y} \mid \mathbf{w}; \sigma_{\mathbf{X}}) \quad \text{if } (\mathbf{Y} \perp \mathbf{T} \mid \mathbf{W}) \quad \text{in} \quad \mathcal{G}_{\sigma_{\mathbf{X}}}$$

    - Graph depends on the specification of the intervention

Rules of σ- calculus

- Rule 2

  - do-calculus

$$P(\mathbf{y} \mid do(\mathbf{x}), \mathbf{w}) = P(\mathbf{y} \mid \mathbf{x}, \mathbf{w}) \qquad \text{if } (\mathbf{Y} \perp \mathbf{X} \mid \mathbf{W}) \quad \text{in} \quad \mathcal{G}_{\underline{\mathbf{X}}}$$

X is observed

  - σ-calculus

$$P(\mathbf{y} \mid \mathbf{x}, \mathbf{w}; \sigma_{\mathbf{X}}) = P(\mathbf{y} \mid \mathbf{x}, \mathbf{w}) \qquad \text{if } (\mathbf{Y} \perp \mathbf{X} \mid \mathbf{W}) \quad \text{in} \quad \mathcal{G}_{\sigma_{\mathbf{X}}\underline{X}} \text{ and } \mathcal{G}_{\underline{X}}$$

- Separation statement needs to hold in the pre-interventional and post-interventional graphs

σ-calculus will not allow this derivation

- Recall the sequential treatment example from before

$$P(y \mid do(x_1), do(X_2 = g(x_1, z)))$$
$$= P(y; \sigma_{X_1, X_2})$$

- If we try rule 2, the required separation is $(Y \perp X_1)$ in the two graphs to the right.
- It holds only on the second one, so the rule is not applicable.

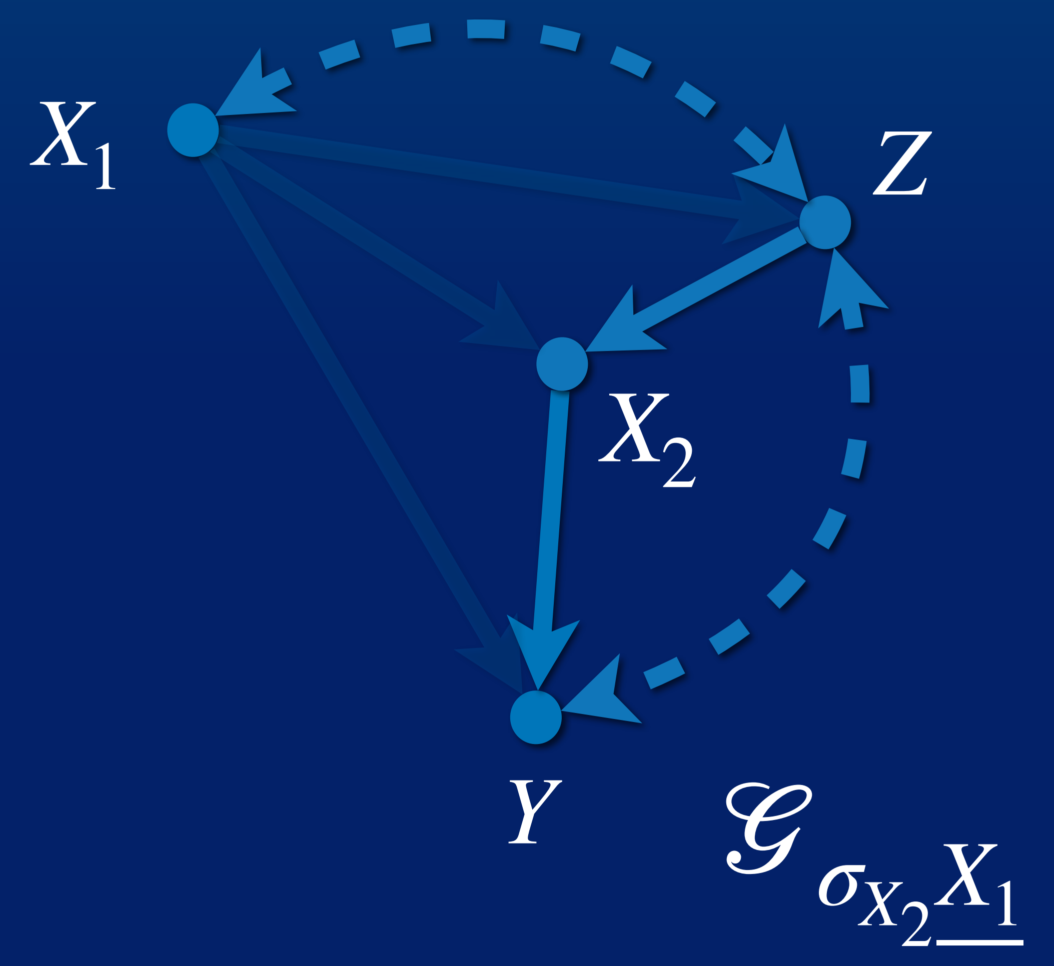

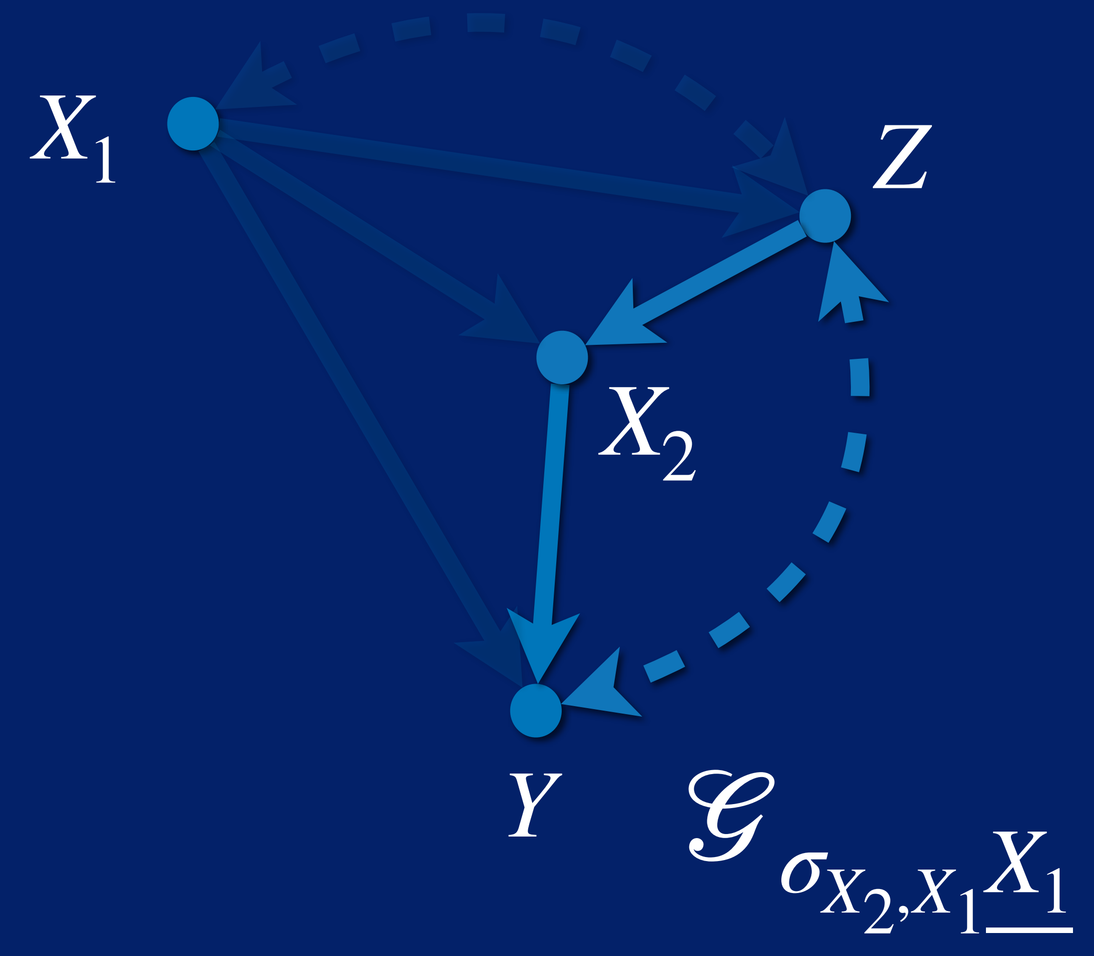

Rules of σ- calculus

- Rule 3
  - do-calculus

$$P(\mathbf{y} \mid do(\mathbf{x}), \mathbf{w}) = P(\mathbf{y} \mid \mathbf{w}) \qquad \text{if } (\mathbf{Y} \perp \mathbf{X} \mid \mathbf{W}) \quad \text{in} \quad \mathcal{G}_{\overline{\mathbf{X}(\mathbf{W})}}$$

X is not observed

  - σ-calculus

$$P(\mathbf{y} \mid \mathbf{w}; \sigma_{\mathbf{X}}) = P(\mathbf{y} \mid \mathbf{w}) \qquad \text{if } (\mathbf{Y} \perp \mathbf{X} \mid \mathbf{W}) \quad \text{in} \quad \mathcal{G}_{\sigma_{\mathbf{X}}\overline{X(W)}} \text{ and } \mathcal{G}_{\overline{X(W)}}$$

- Separation statement needs to hold in the pre-interventional and post-interventional graphs

Another intervention

- Resources are limited so we want to focus on students that need tutoring the most.
- From now on, students with low GPA have to get tutoring and the service will only be available to them. That is: $P^*(X = 1 \mid W = 0) = 1$, $P^*(X = 1 \mid W = 0) = 0$.

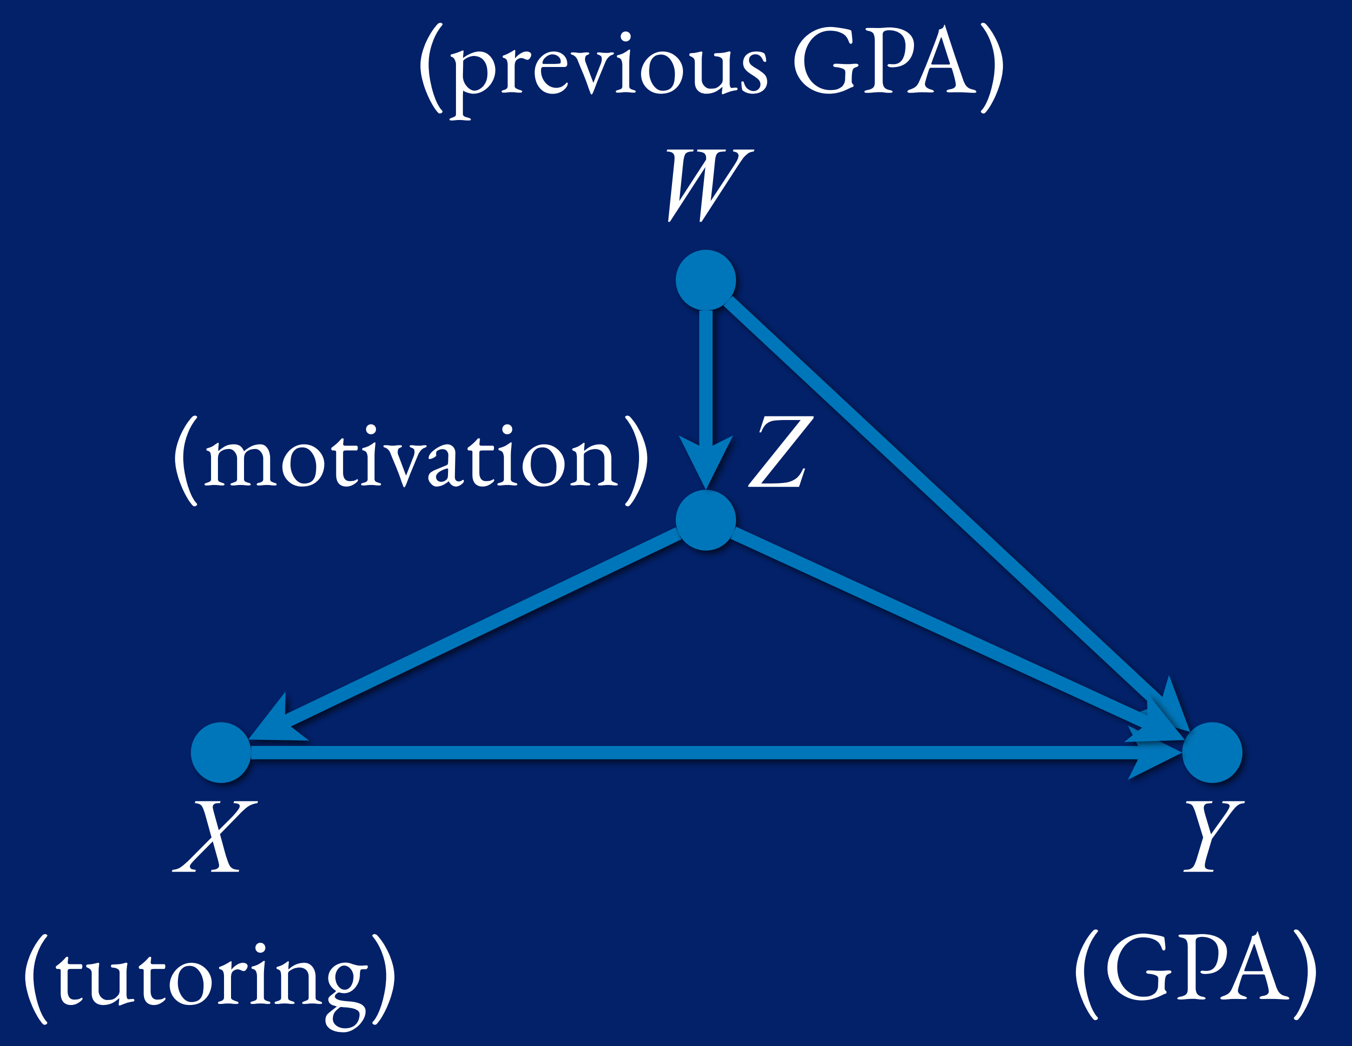

$\mathcal{G}$ Natural (current) Regime

Intervention

$\sigma_X = 1[W = 1]$

Assign tutoring only to students with low GPA.

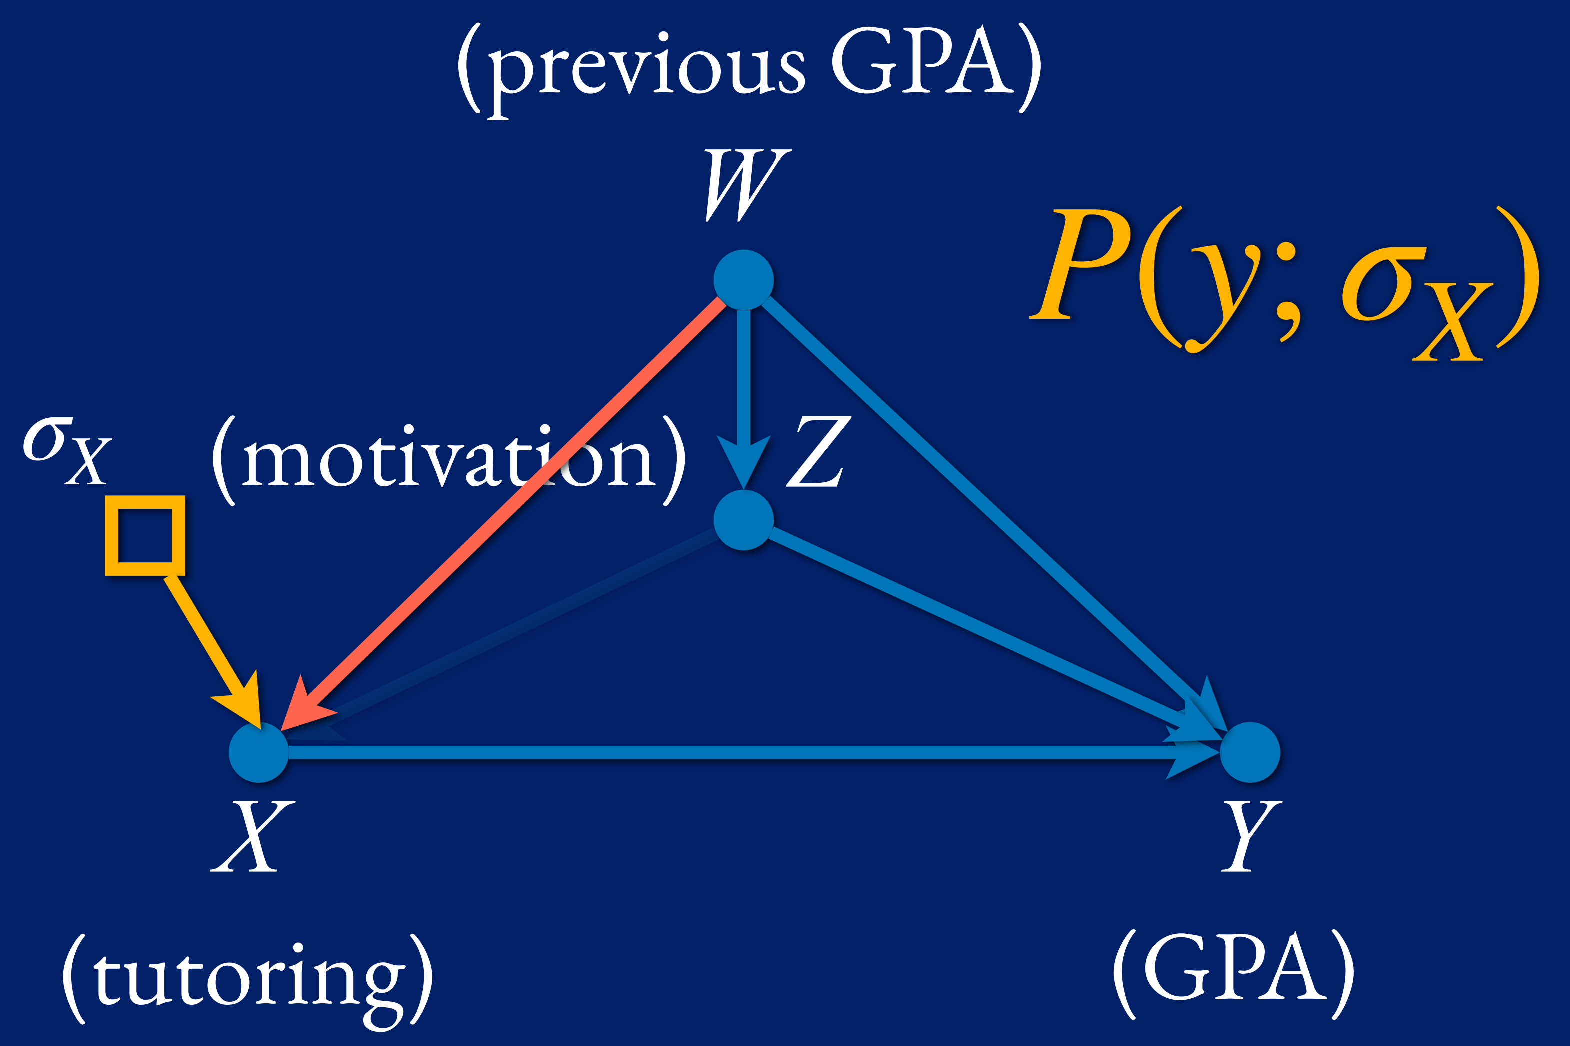

$\mathcal{G}_{\sigma_X}$ Intervened (hypothesized) Regime

Using σ-calculus

$$P(y;\sigma_X)$$

$$= \sum_{w,z,x} P(y \mid x, w, z; \sigma_X) P(x \mid w, z; \sigma_X) P(w, z; \sigma_X)$$

$$= \sum_{w,z,x} P(y \mid x, w, z; \sigma_X) P(x \mid w; \sigma_X) P(w, z; \sigma_X)$$

Rule 1 $(X \perp Z \mid W)$ in $\mathcal{G}_{\sigma_X}$

$$= \sum_{w,z,x} P(y \mid x, w, z) P(x \mid w; \sigma_X) P(w, z; \sigma_X)$$

Rule 2 $(Y \perp X \mid W, Z)$ in $\mathcal{G}_{\sigma_X \underline{X}}$ and $\mathcal{G}_{\underline{X}}$

$$= \sum_{w,z,x} P(y \mid x, w, z) P(x \mid w; \sigma_X) P(w, z)$$

Rule 3 $(W, Z \perp X)$ in $\mathcal{G}_{\sigma_X \overline{X}}$ and $\mathcal{G}_{\overline{X}}$

Estimable from current regime

Defined by σ_X

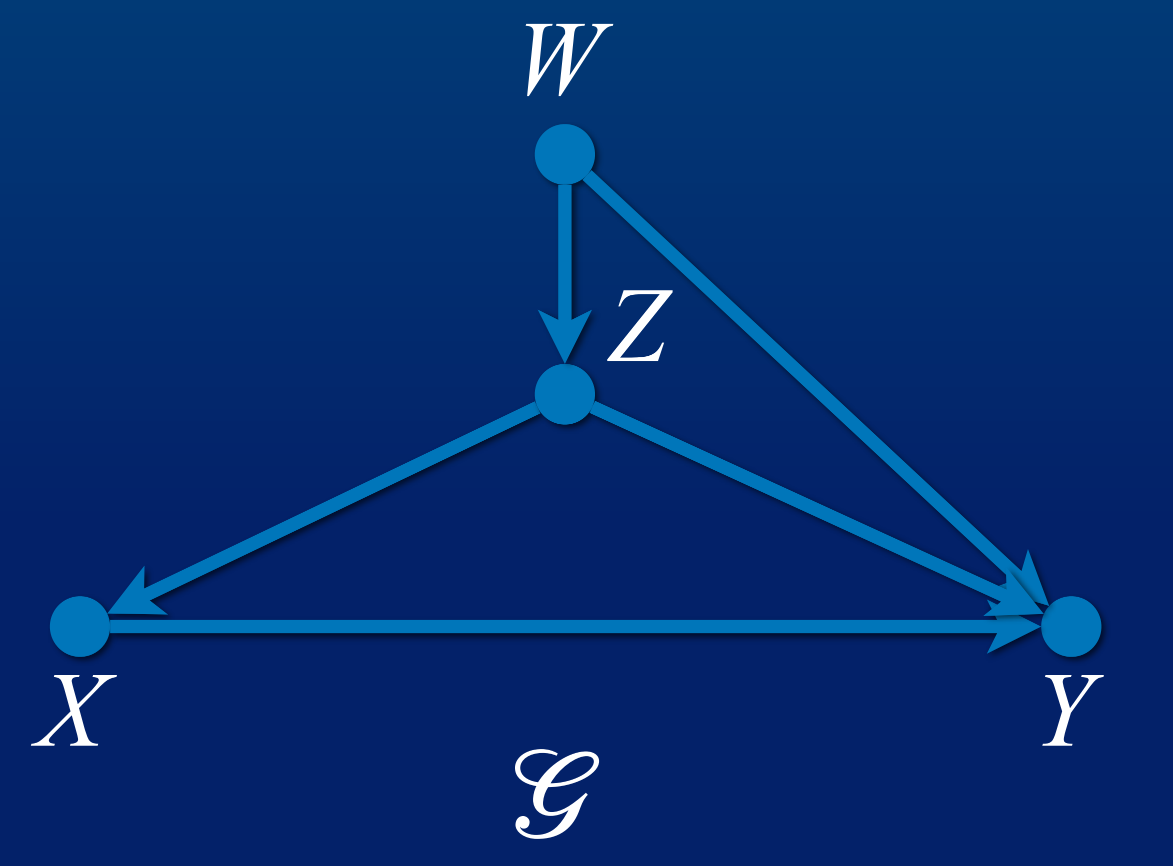

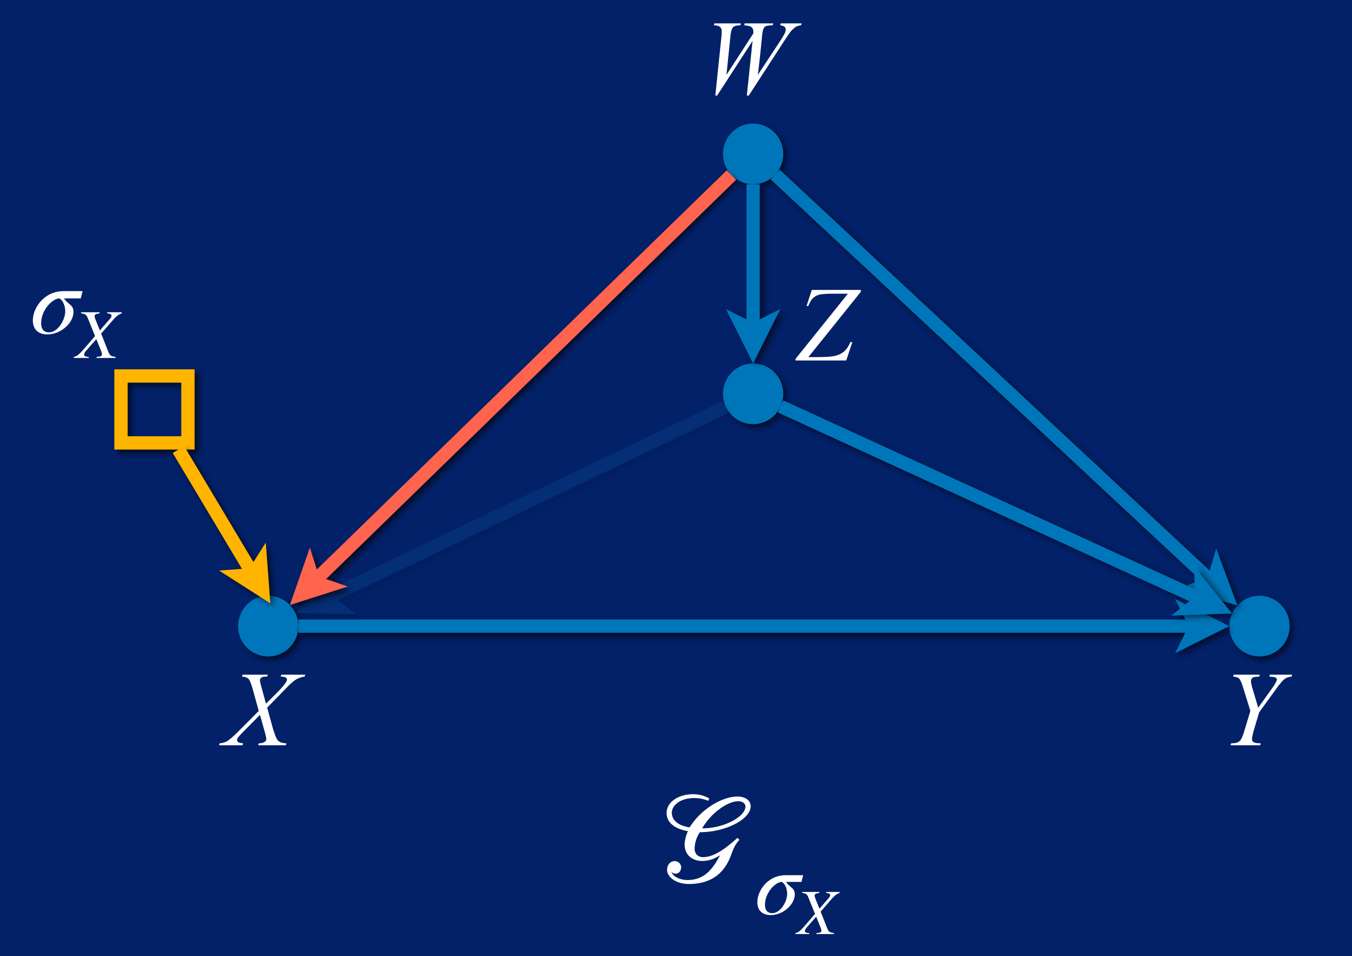

Summary

- For many realistic situations, soft interventions are more suitable for representing plans and policies that can actually be implemented.
- We introduce a set of inference rules called σ-calculus, which generalizes Pearl's do-calculus, to reason about the effect of general types of interventions.
- These rules provide a syntactical method for deriving and verifying claims about soft interventions given a causal graph.

Thank you! Questions?