A Calculus for Stochastic Interventions: Causal Effect Identification and Surrogate Experiments

Juan D. Correa and Elias Bareinboim
{jdcorrea, eb}@cs.columbia.edu

COLUMBIA UNIVERSITY

February, 2020, New York

AAAI-20, New York

Outline

Outline

- Hard/atomic interventions vs. Soft/non-atomic interventions
- Graphical representation
- Inferences rules for soft interventions (σ-calculus)
- Imperfect surrogate experiments
- Conclusions

Motivating example

Motivating example

- Consider a tutoring program in place at a certain school.

Motivating example

- Consider a tutoring program in place at a certain school.
- For each student, we observe the GPA at the beginning of the term, their motivation (low, high), whether they got tutoring or not, and their GPA at the end.

Motivating example

- Consider a tutoring program in place at a certain school.
- For each student, we observe the GPA at the beginning of the term, their motivation (low, high), whether they got tutoring or not, and their GPA at the end.

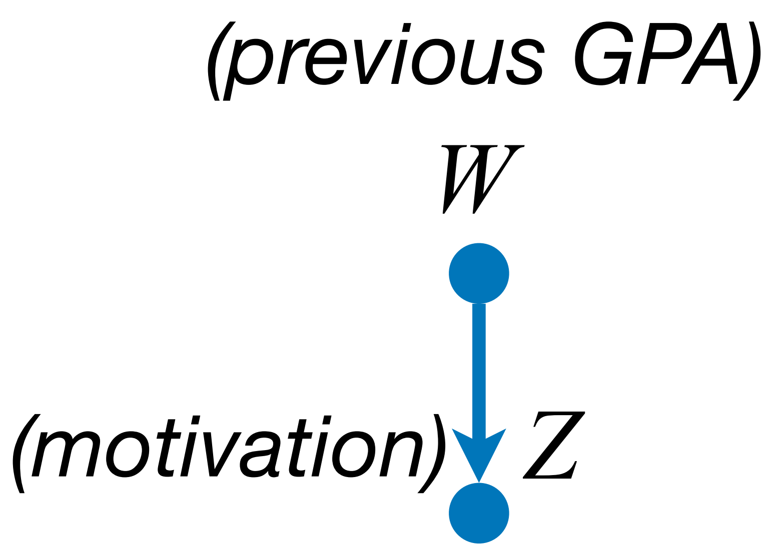

- Motivation depends (among other not observed factors) on the previous GPA.

Motivating example

- Consider a tutoring program in place at a certain school.
- For each student, we observe the GPA at the beginning of the term, their motivation (low, high), whether they got tutoring or not, and their GPA at the end.

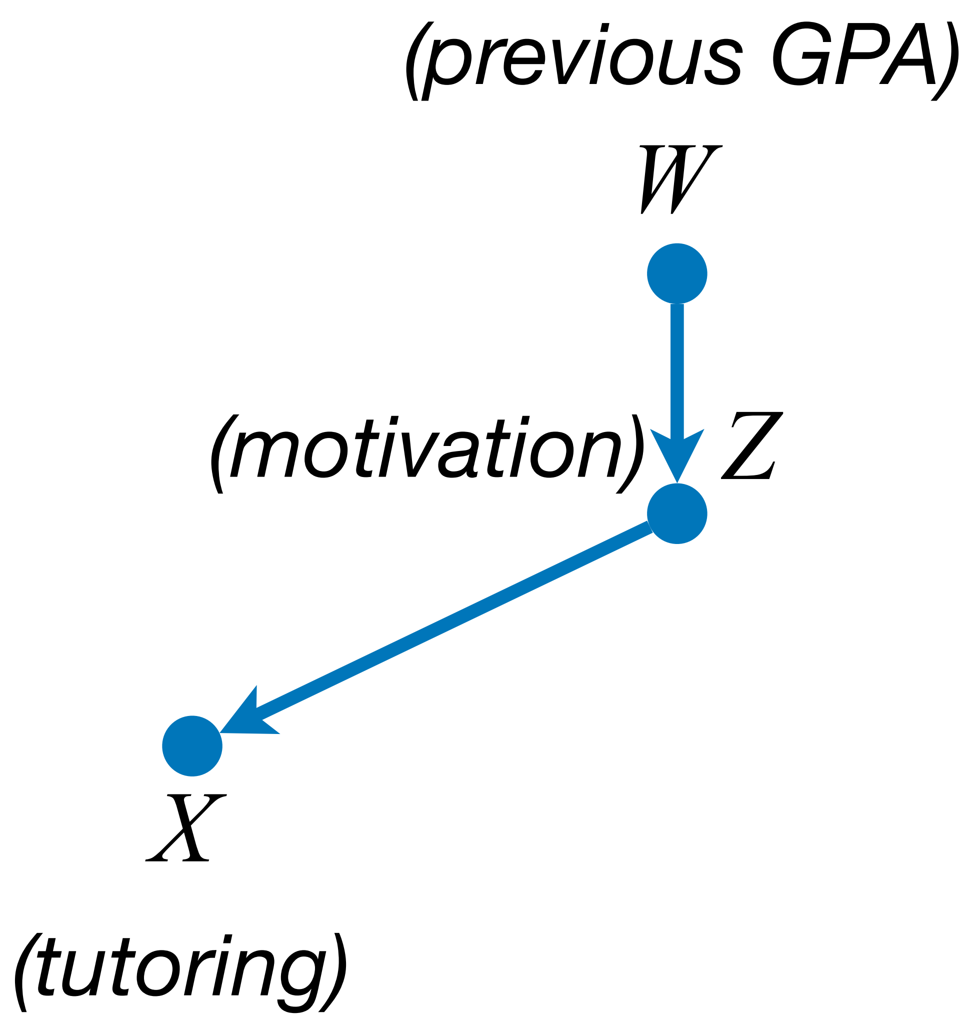

- Motivation depends (among other not observed factors) on the previous GPA.
- Students get tutoring depending on their motivation.

Motivating example

- Consider a tutoring program in place at a certain school.
- For each student, we observe the GPA at the beginning of the term, their motivation (low, high), whether they got tutoring or not, and their GPA at the end.

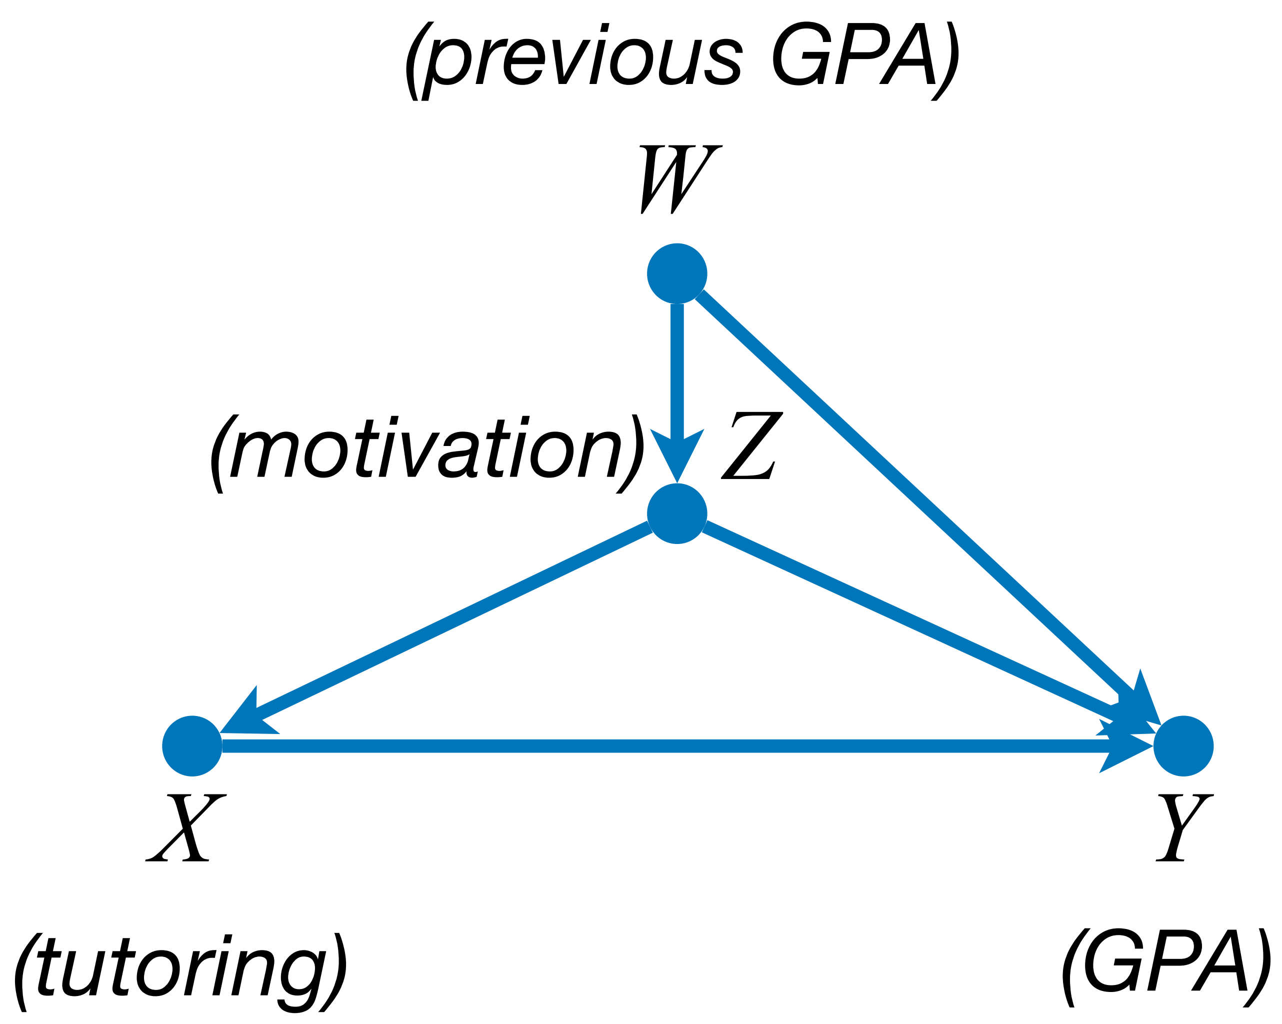

- Motivation depends (among other not observed factors) on the previous GPA.
- Students get tutoring depending on their motivation.
- The GPA at the end of the term is a function of the previous GPA, student's motivation and getting tutoring or not.

Motivating example

- Consider a tutoring program in place at a certain school.
- For each student, we observe the GPA at the beginning of the term, their motivation (low, high), whether they got tutoring or not, and their GPA at the end.

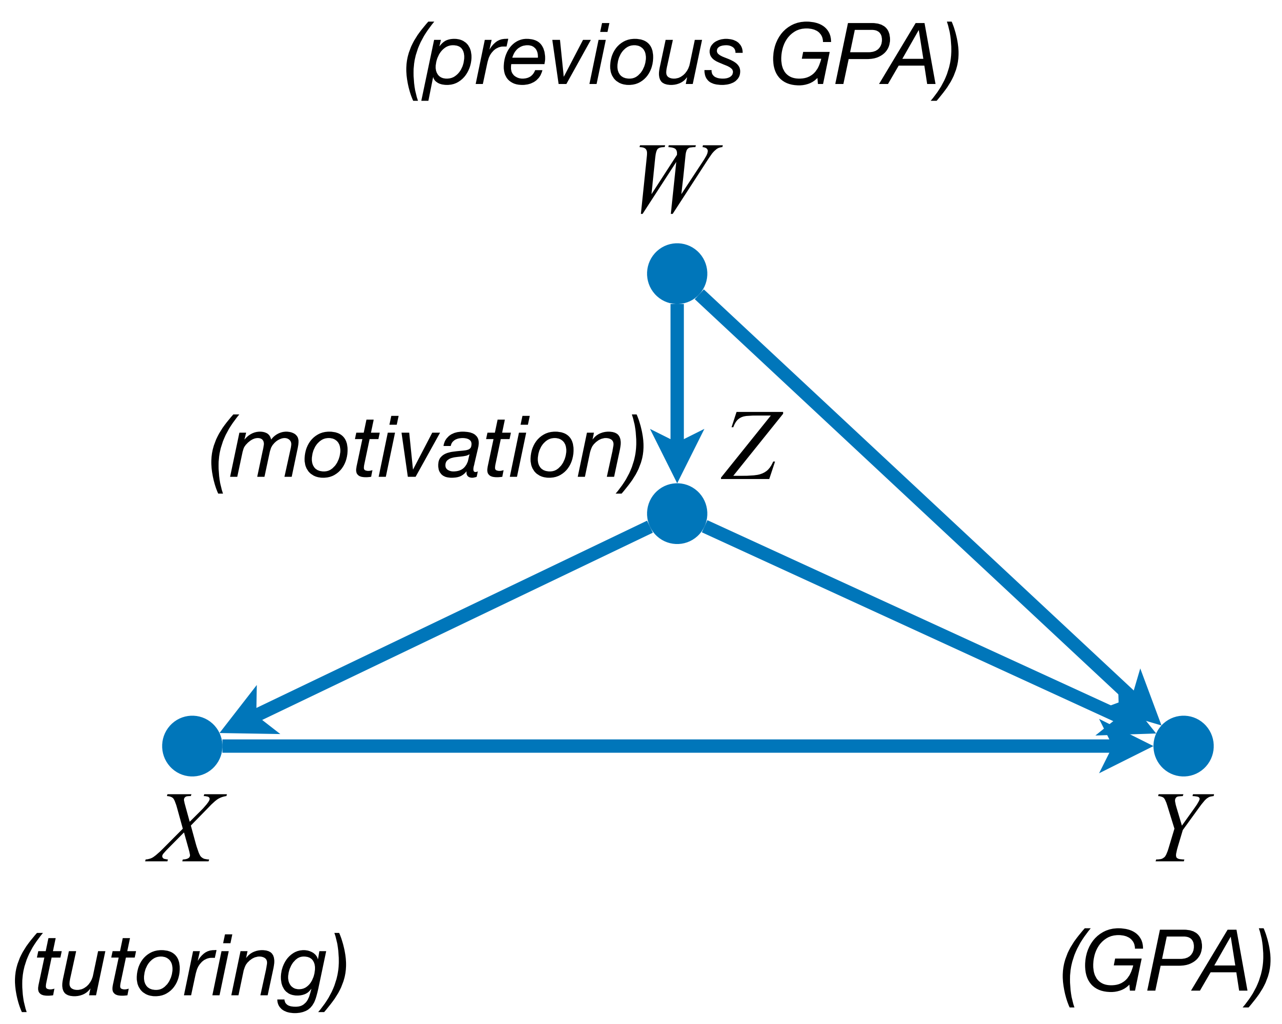

G **Natural (current) Regime**

- Motivation depends (among other not observed factors) on the previous GPA.
- Students get tutoring depending on their motivation.
- The GPA at the end of the term is a function of the previous GPA, student's motivation and getting tutoring or not.

Motivating example

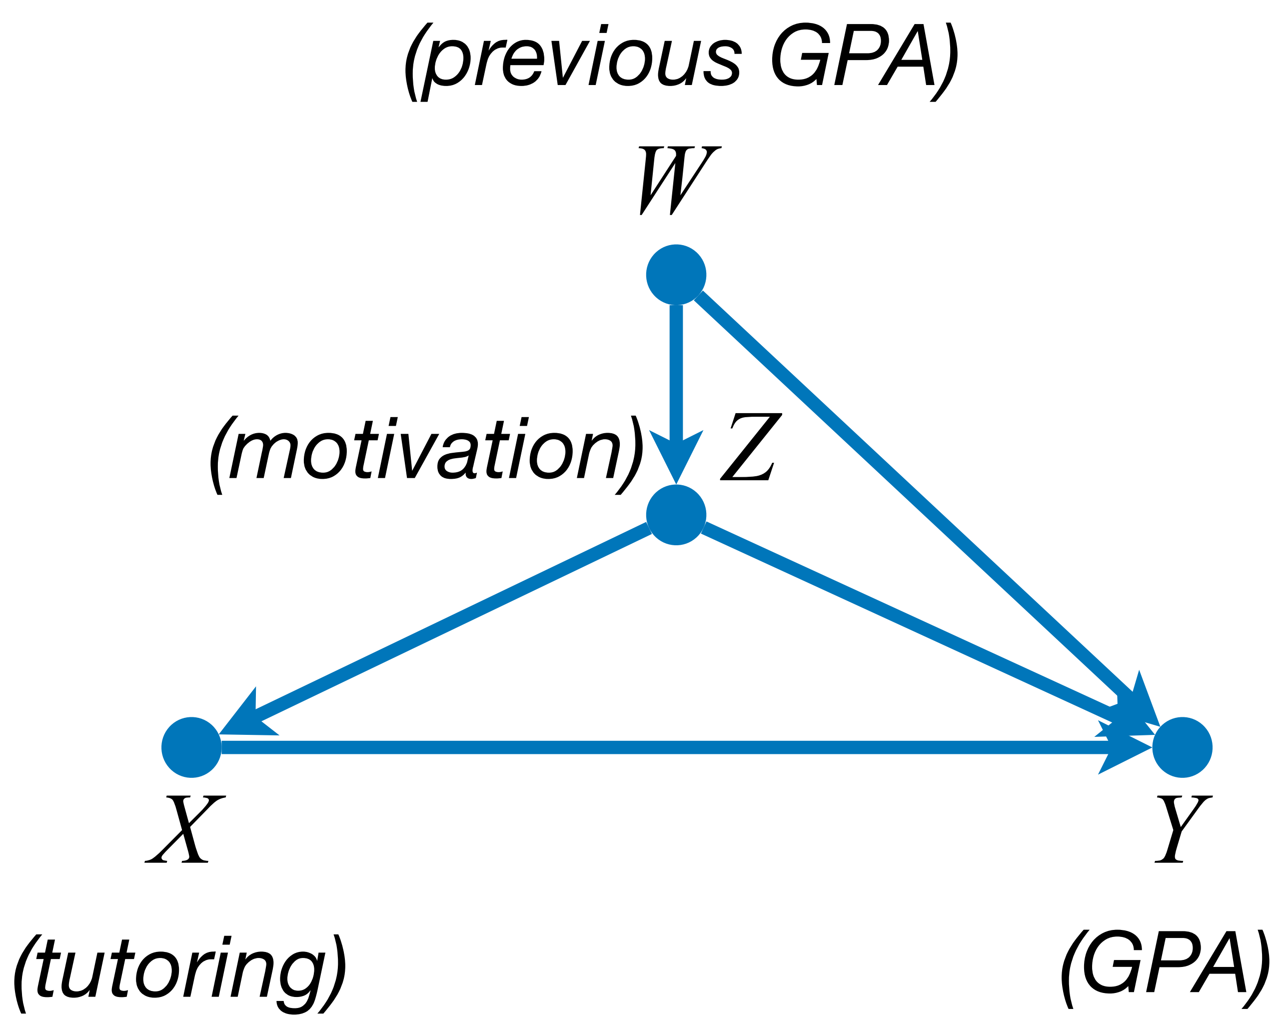

$\mathcal{G}$ **Natural (current) Regime**

Motivating example

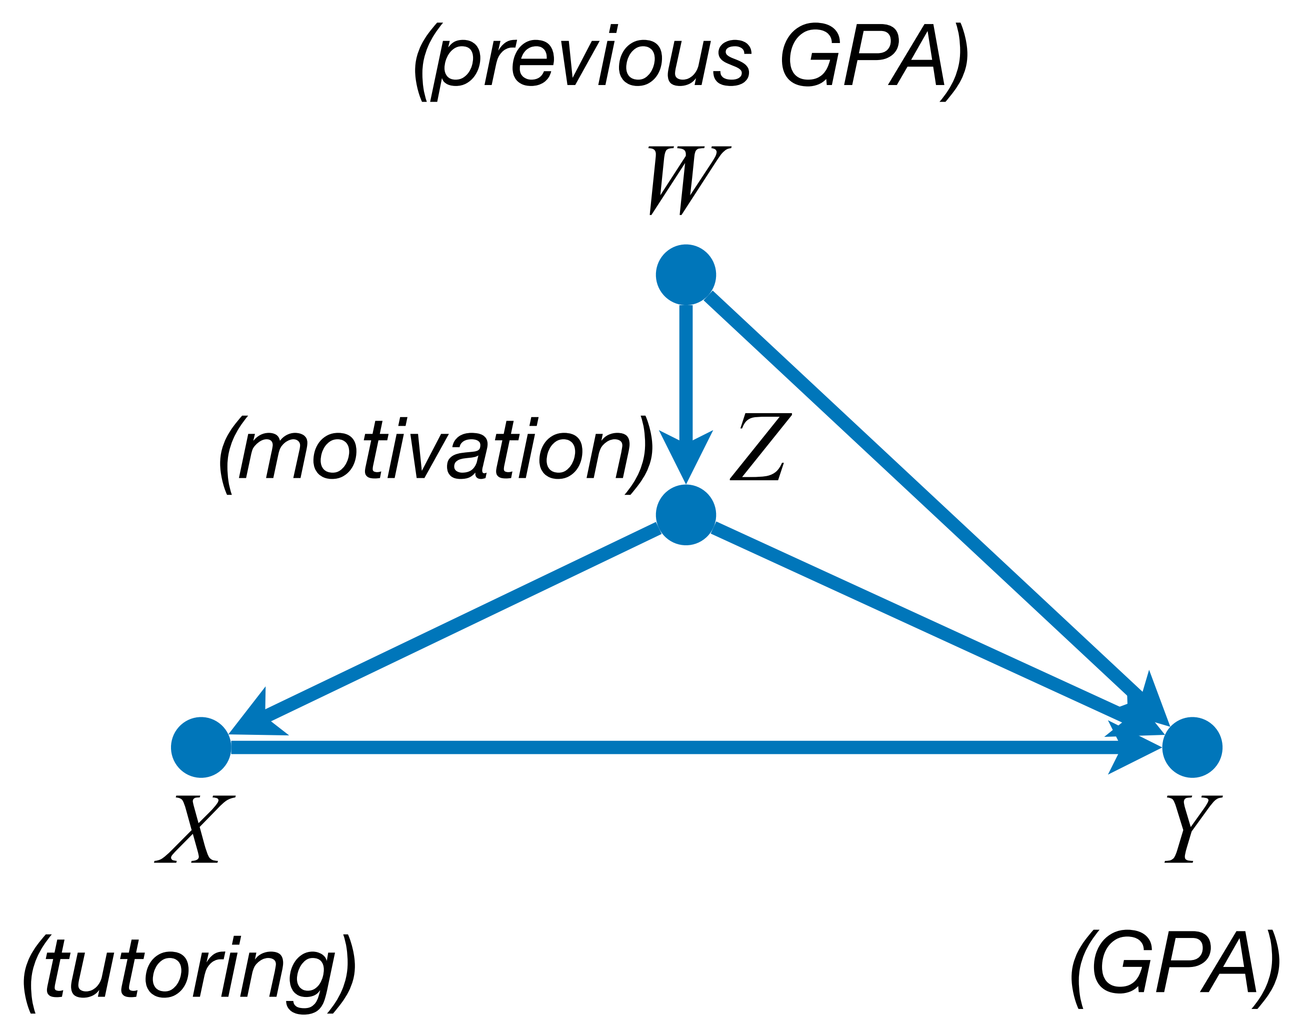

$\mathcal{G}$ **Natural (current) Regime**

- Using machine learning, and with enough data, a students GPA can be predicted with small error given other features i.e., $P(y \mid w, z, x)$.

Motivating example

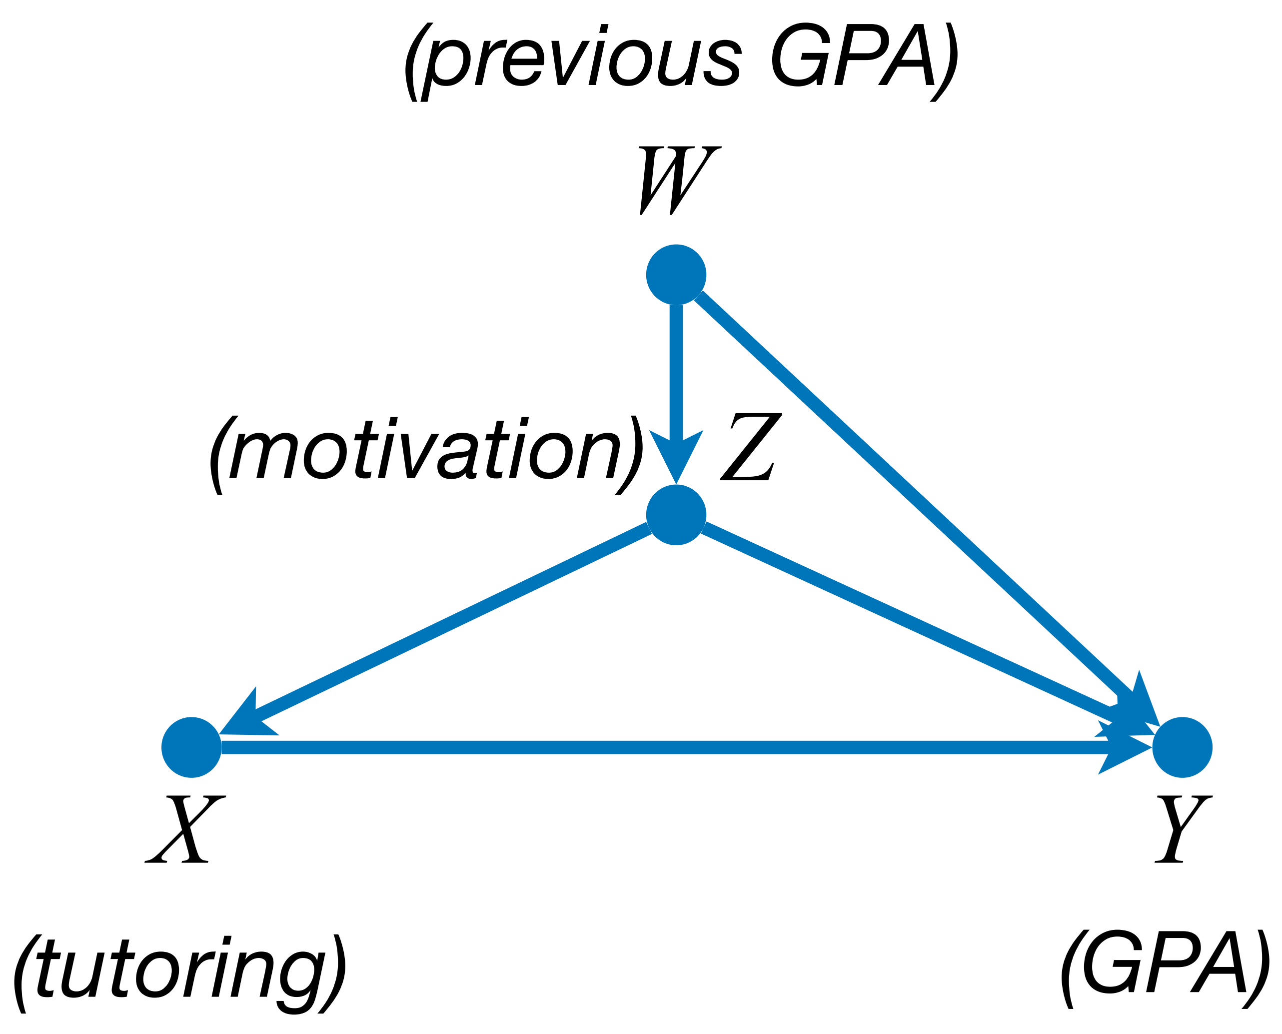

$\mathcal{G}$ **Natural (current) Regime**

- Using machine learning, and with enough data, a students GPA can be predicted with small error given other features i.e., $P(y \mid w, z, x)$.
- This distribution is a model that reflects the current/natural regime, but we are interested in taking decisions to improve the students GPA.

Motivating example

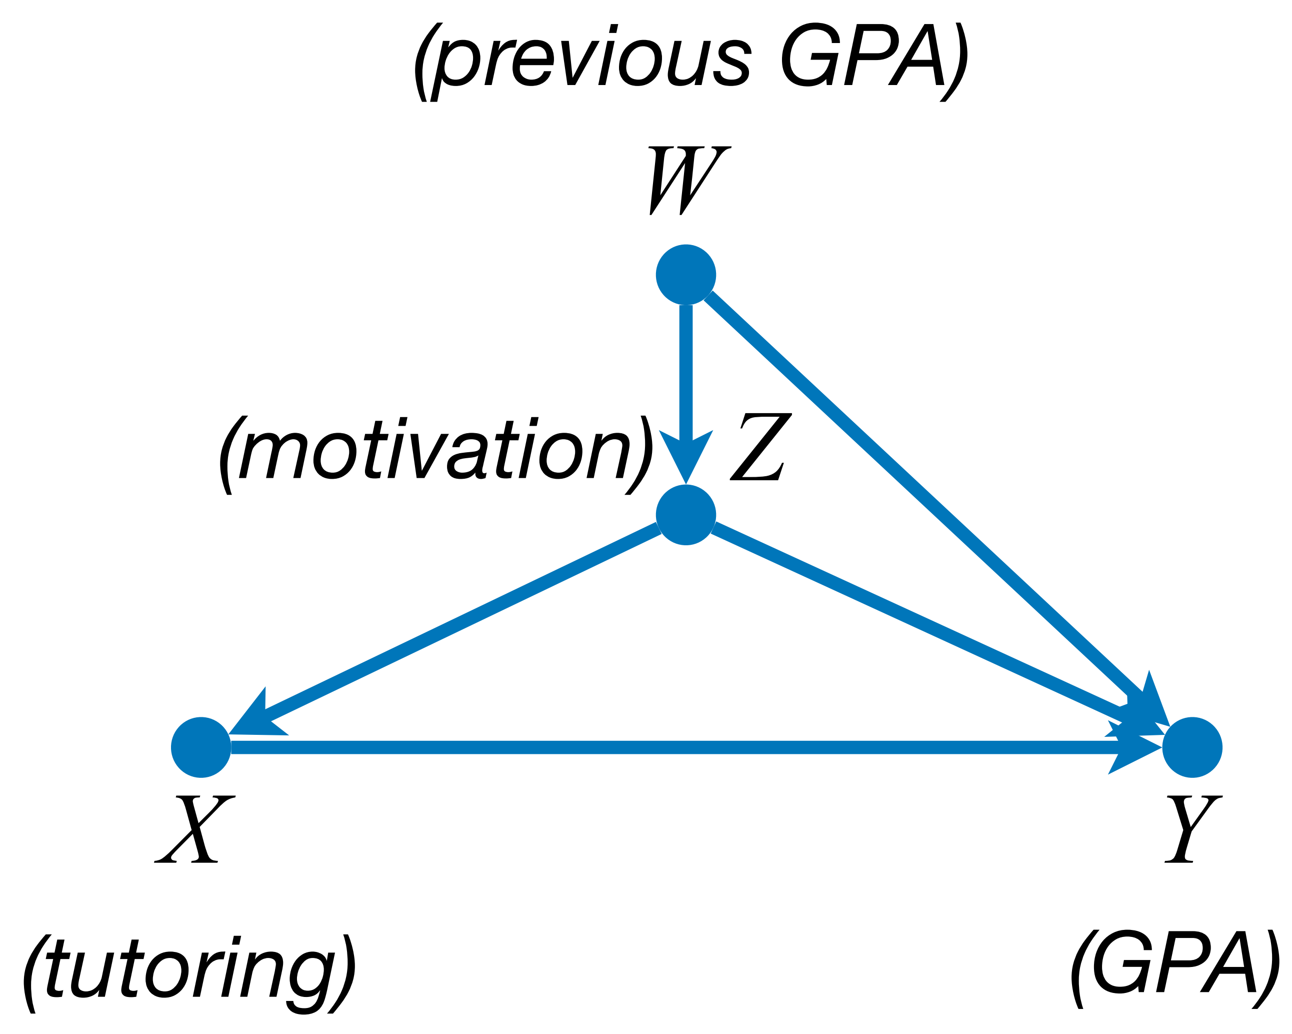

$\mathcal{G}$ **Natural (current) Regime**

- Using machine learning, and with enough data, a students GPA can be predicted with small error given other features i.e., $P(y \mid w, z, x)$.
- This distribution is a model that reflects the current/natural regime, but we are interested in taking decisions to improve the students GPA.
- Taking decisions amount to intervening the current regime. Hence, we are interested in predicting student's GPA receiving tutoring in a hypothetical (unrealized) reality.

Motivating example

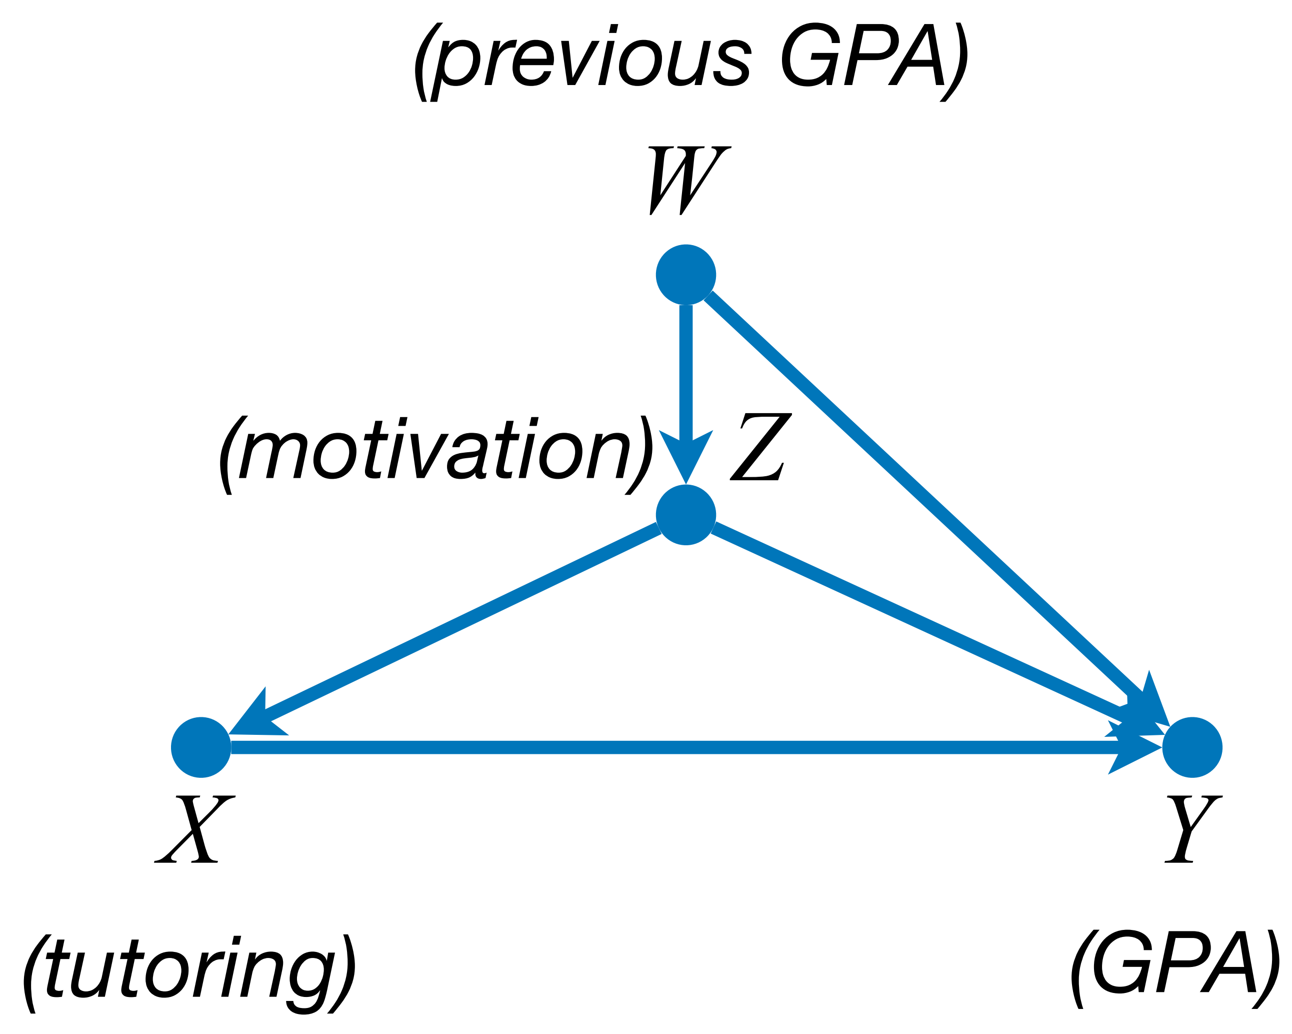

$\mathcal{G}$ **Natural (current) Regime**

- Using machine learning, and with enough data, a students GPA can be predicted with small error given other features i.e., $P(y \mid w, z, x)$.
- This distribution is a model that reflects the current/natural regime, but we are interested in taking decisions to improve the students GPA.
- Taking decisions amount to intervening the current regime. Hence, we are interested in predicting student's GPA receiving tutoring in a hypothetical (unrealized) reality.
- This is a causal inference question!

Some types of Interventions

Some types of Interventions

- **Hard/atomic**: $\sigma_X = do(X=x)$ set variable X to a constant value x.
  (Pearl's original treatment considered mostly this intervention.)

Some types of Interventions

- **Hard/atomic**: $\sigma_X = do(X=x)$ set variable X to a constant value x.
  (Pearl's original treatment considered mostly this intervention.)
  - Every student gets tutoring.

Some types of Interventions

- **Hard/atomic**: $\sigma_X = do(X=x)$ set variable X to a constant value x.
  (Pearl's original treatment considered mostly this intervention.)
  - Every student gets tutoring.

- **Conditional**: $\sigma_X = g(\boldsymbol{w})$ sets the variable X to output the result of a function g that depends on a set of observable variables $\boldsymbol{W}$.

Some types of Interventions

- **Hard/atomic**: $\sigma_X = do(X = x)$ set variable X to a constant value x.
 (Pearl's original treatment considered mostly this intervention.)
  - Every student gets tutoring.

- **Conditional**: $\sigma_X = g(\mathbf{w})$ sets the variable X to output the result of a function g that depends on a set of observable variables $\mathbf{W}$.
  - Students get tutoring if and only if they have a low GPA.

Some types of Interventions

- **Hard/atomic**: $\sigma_X = do(X=x)$ set variable X to a constant value x.
  (Pearl's original treatment considered mostly this intervention.)
  - Every student gets tutoring.

- **Conditional**: $\sigma_X = g(\mathbf{w})$ sets the variable X to output the result of a function g that depends on a set of observable variables $\mathbf{W}$.
  - Students get tutoring if and only if they have a low GPA.

- **Stochastic**: $\sigma_X = P^*(x|\mathbf{w})$ sets the variable X to follow a given probability distribution conditional on a set of variables $\mathbf{W}$.

Some types of Interventions

- **Hard/atomic**: $\sigma_X = do(X=x)$ set variable X to a constant value x.
  (Pearl's original treatment considered mostly this intervention.)
  - Every student gets tutoring.

- **Conditional**: $\sigma_X = g(\mathbf{w})$ sets the variable X to output the result of a function g that depends on a set of observable variables $\mathbf{W}$.
  - Students get tutoring if and only if they have a low GPA.

- **Stochastic**: $\sigma_X = P^*(x|\mathbf{w})$ sets the variable X to follow a given probability distribution conditional on a set of variables $\mathbf{W}$.
  - Students with low GPA enter a raffle for 80% of the spots, other interested students enter for the remaining 20%.

Hard/Atomic Interventions

Hard/Atomic Interventions

- What if we make tutoring mandatory for every student?

Hard/Atomic Interventions

- What if we make tutoring mandatory for every student?

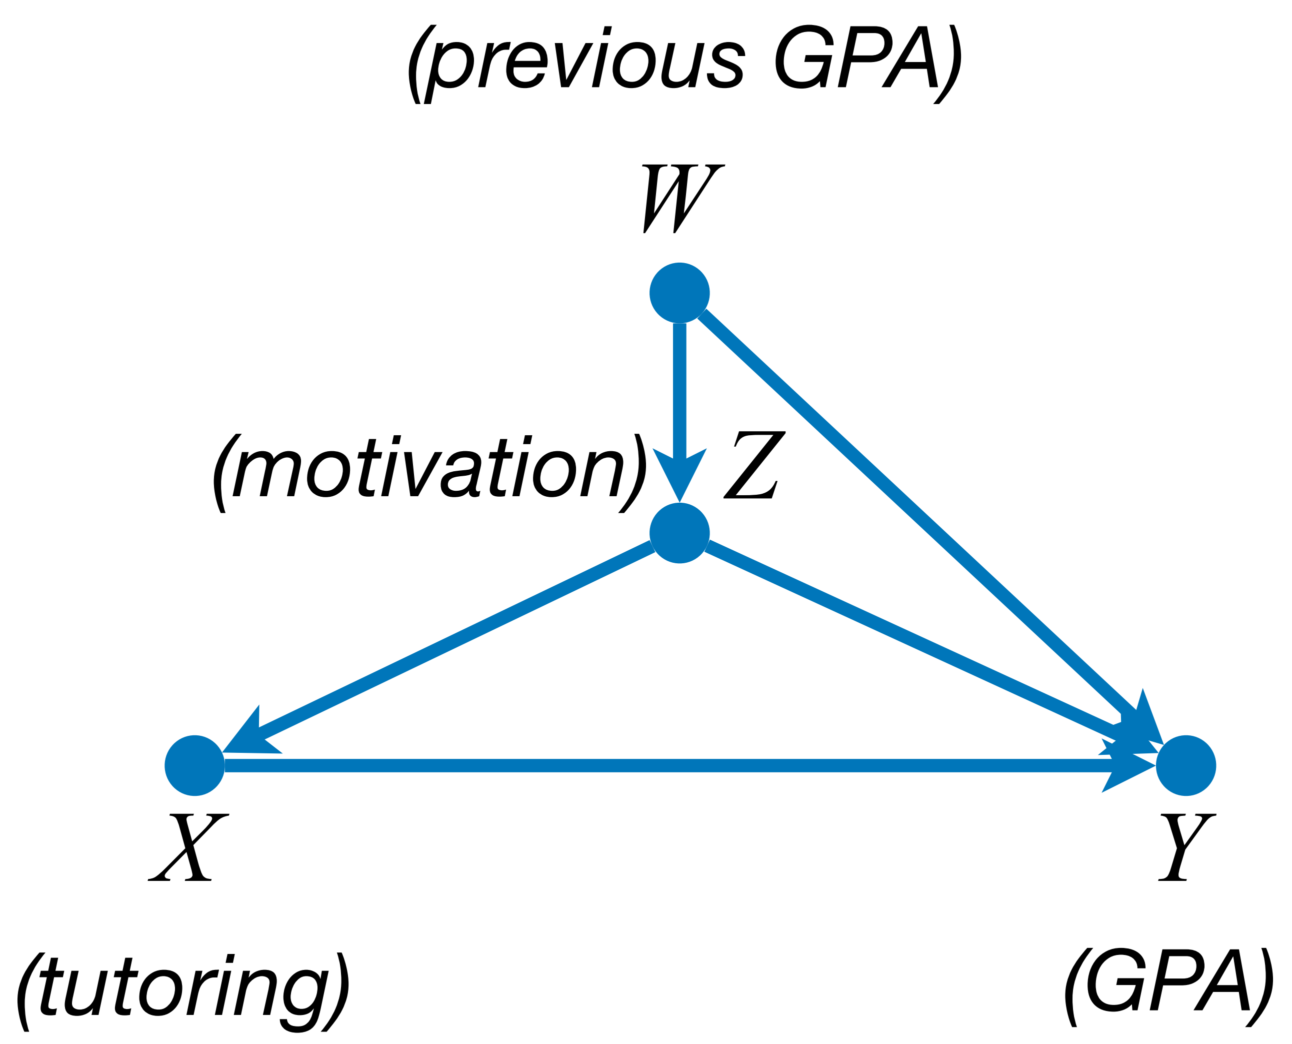

Hard/Atomic Interventions

- What if we make tutoring mandatory for every student?

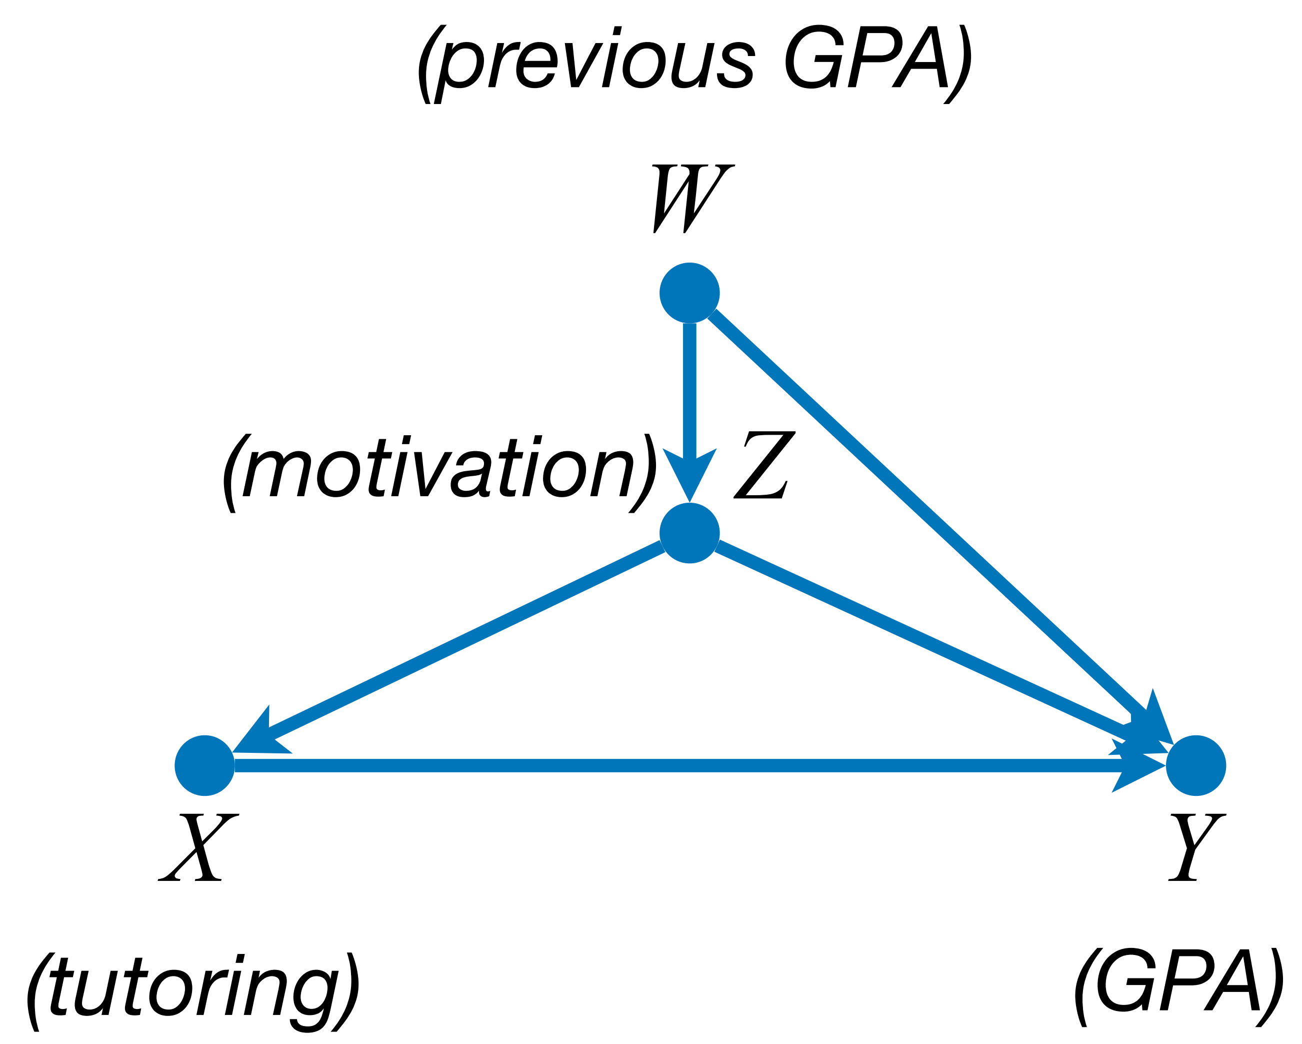

$\mathscr{G}$ **Natural (current) Regime**

Hard/Atomic Interventions

- What if we make tutoring mandatory for every student?

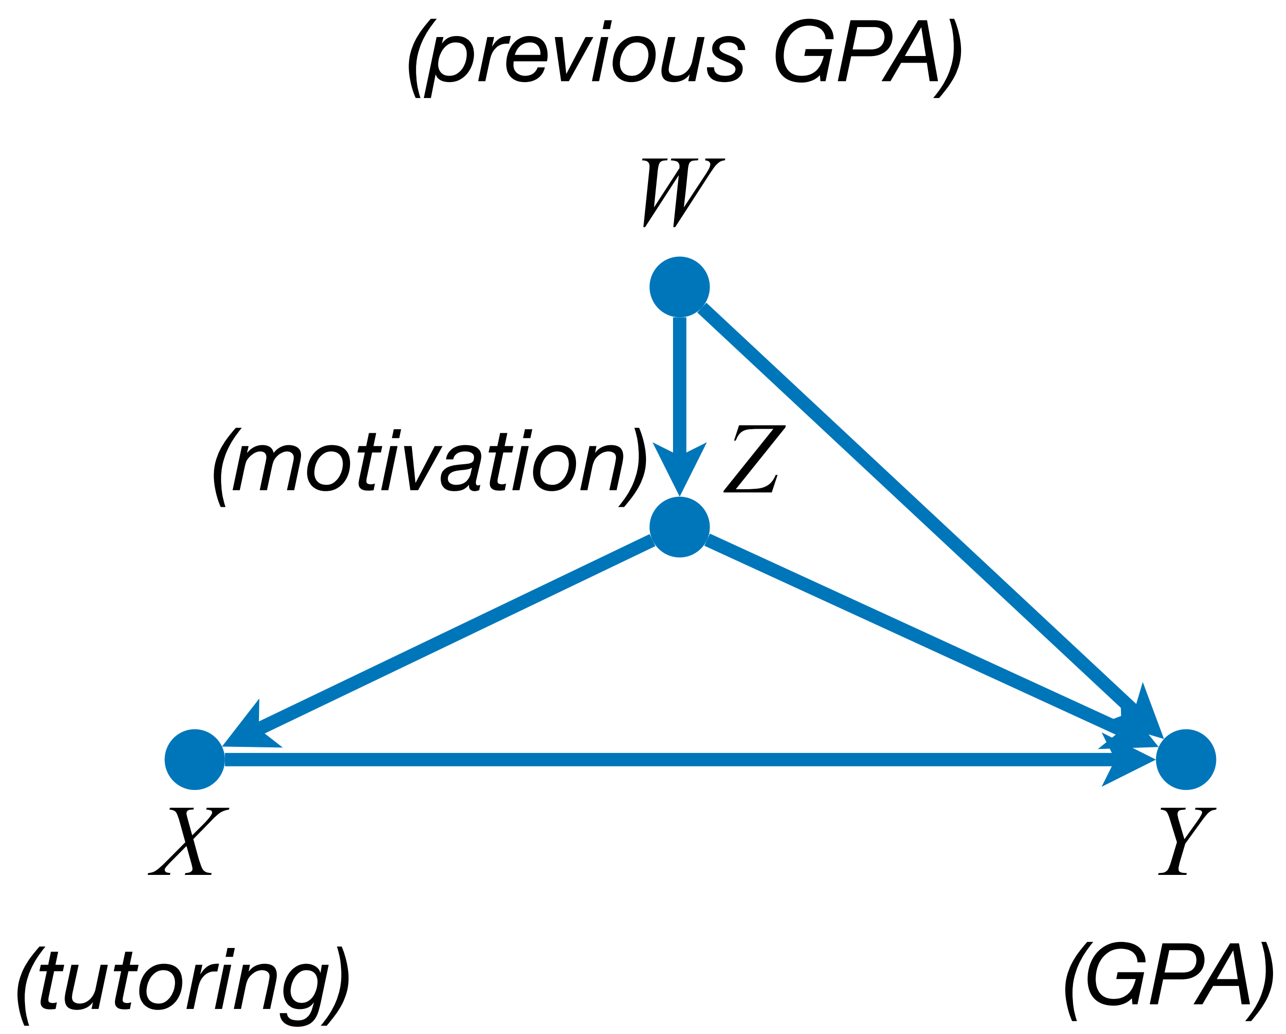

$\mathcal{G}$ **Natural (current) Regime**

Intervention $do(X = 1)$

Make tutoring mandatory for ***all*** *students.*

Hard/Atomic Interventions

- What if we make tutoring mandatory for every student?

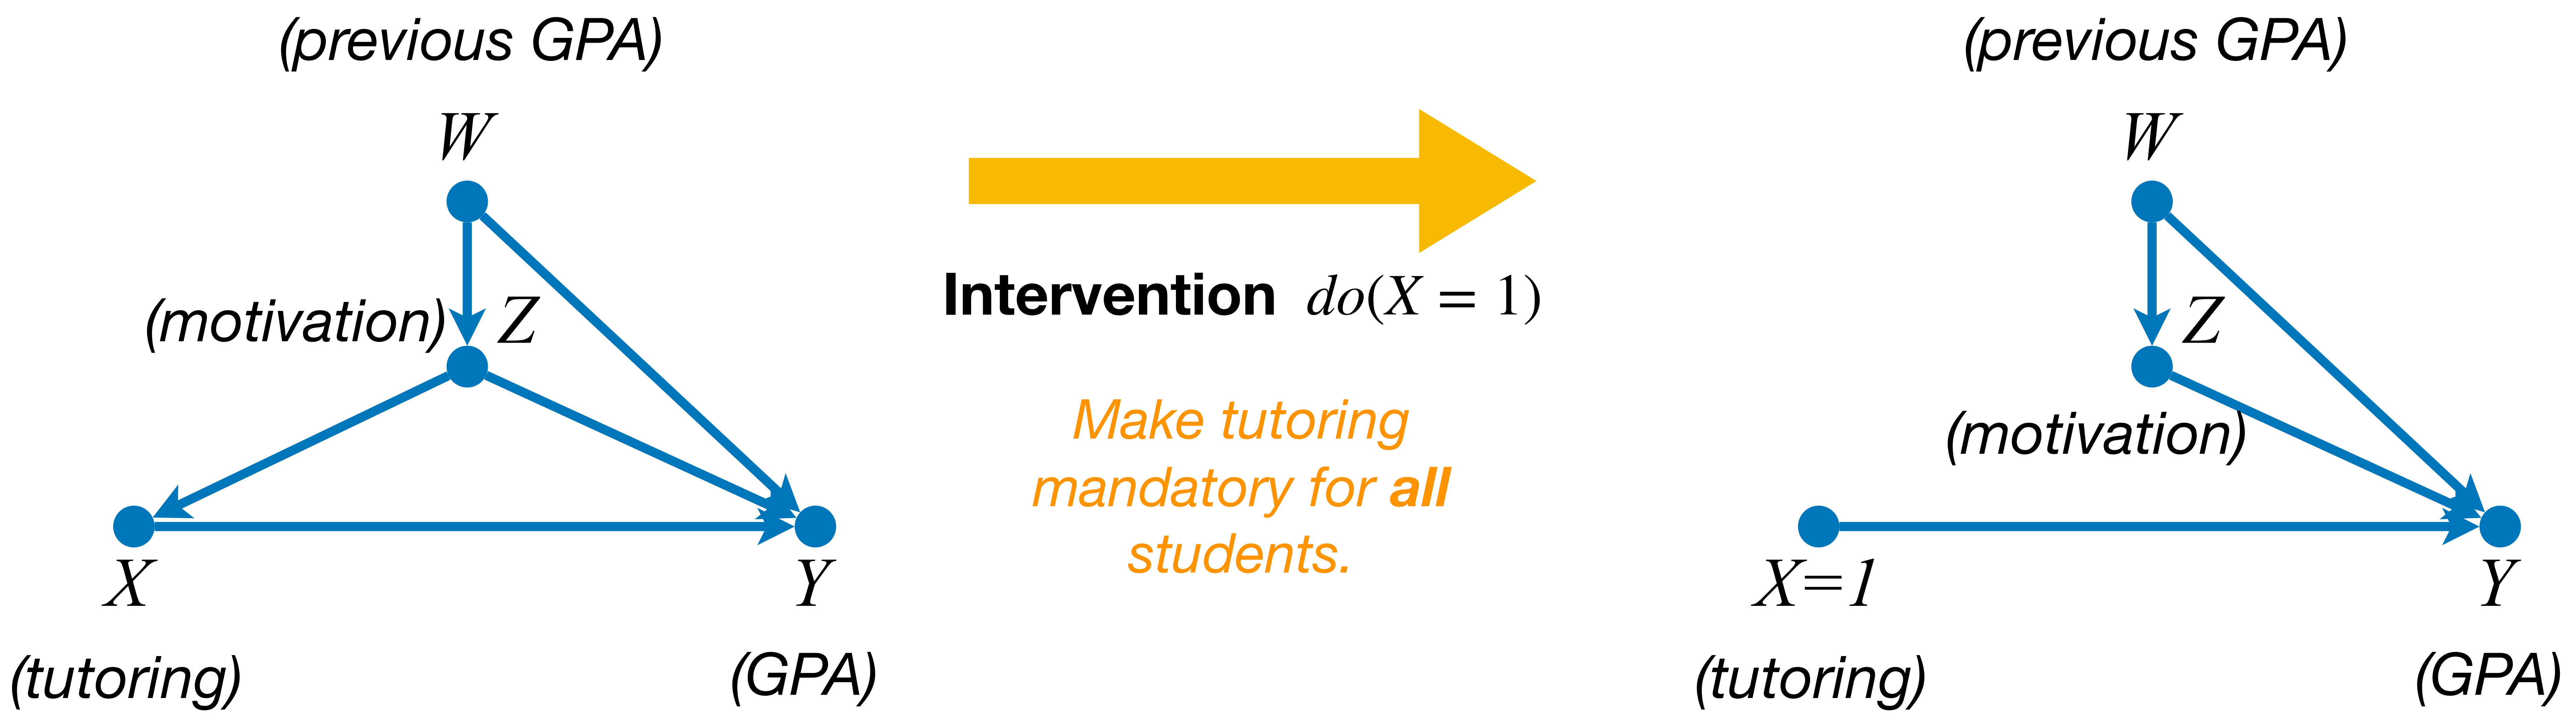

$\mathcal{G}$ **Natural (current) Regime**

Hard/Atomic Interventions

- What if we make tutoring mandatory for every student?

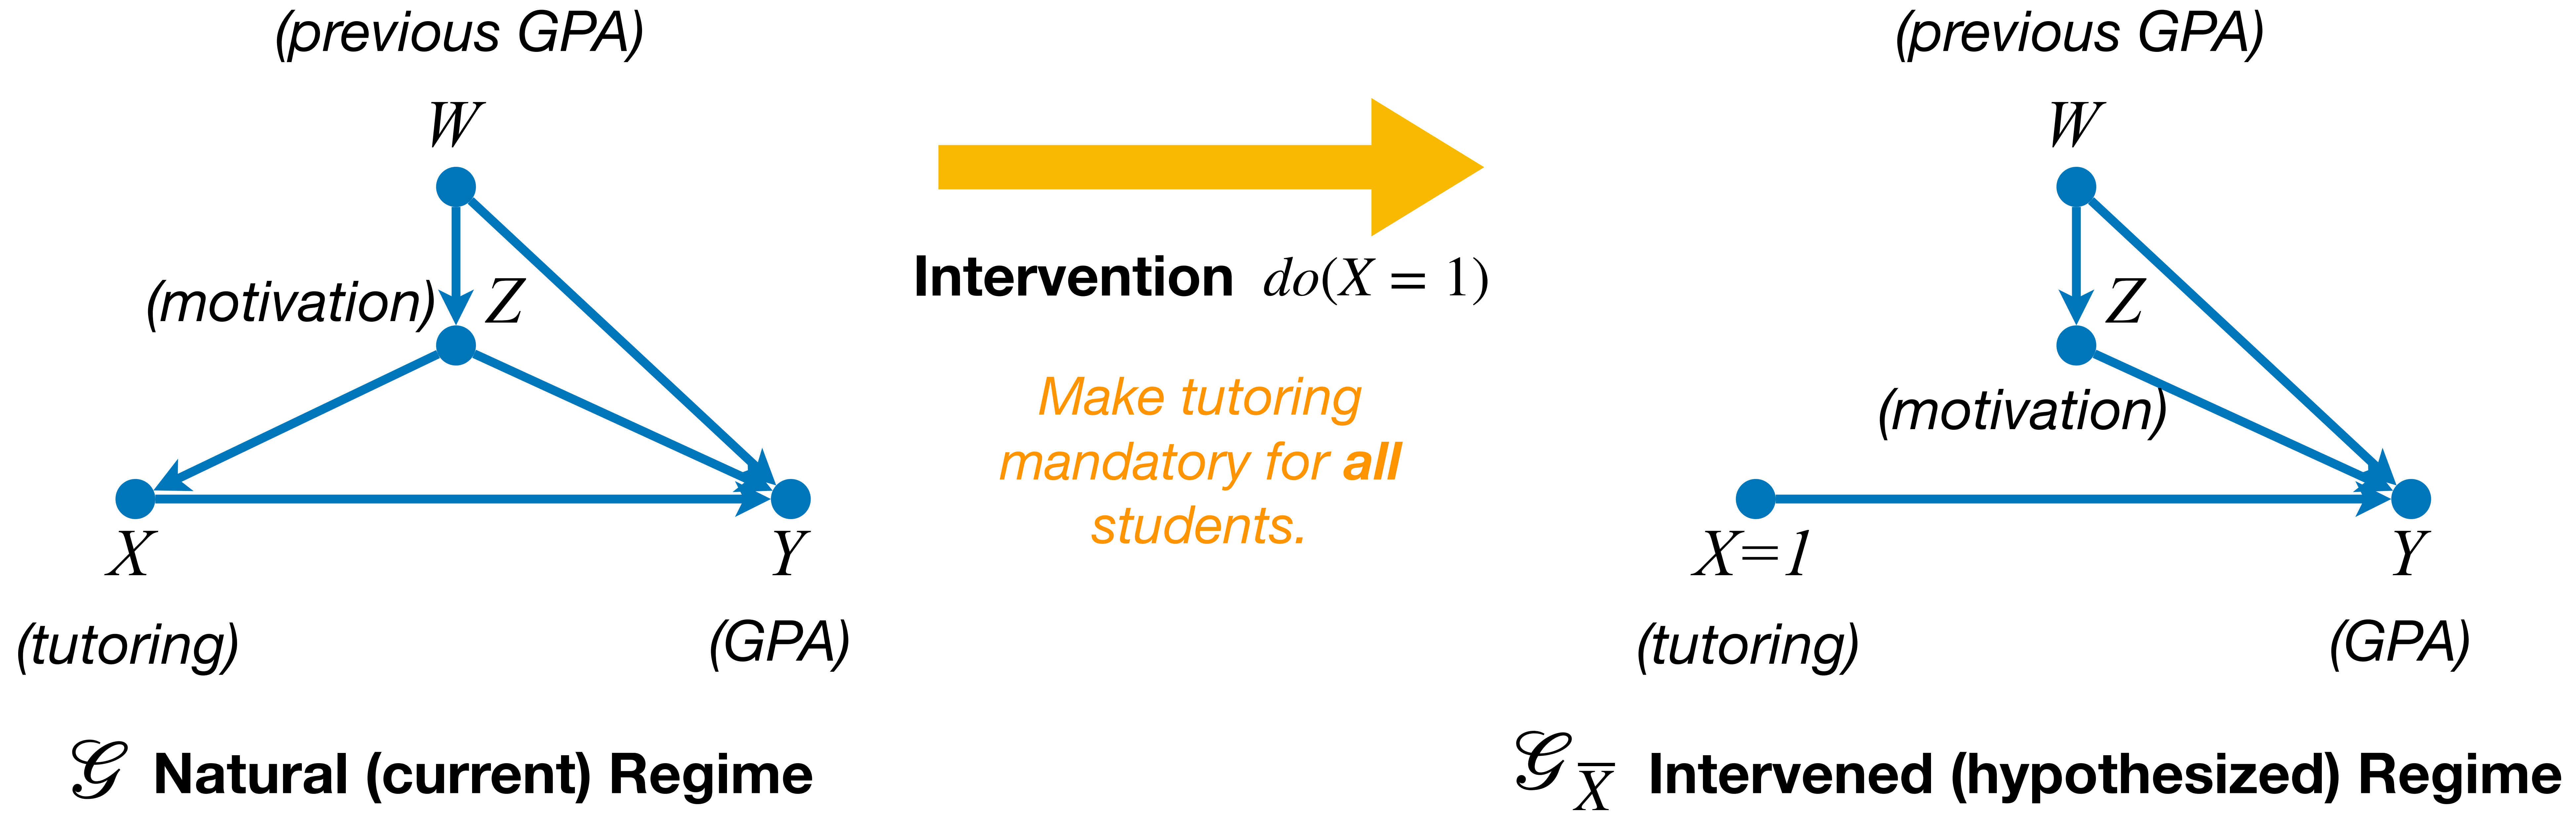

Hard/Atomic Interventions

- What if we make tutoring mandatory for every student?

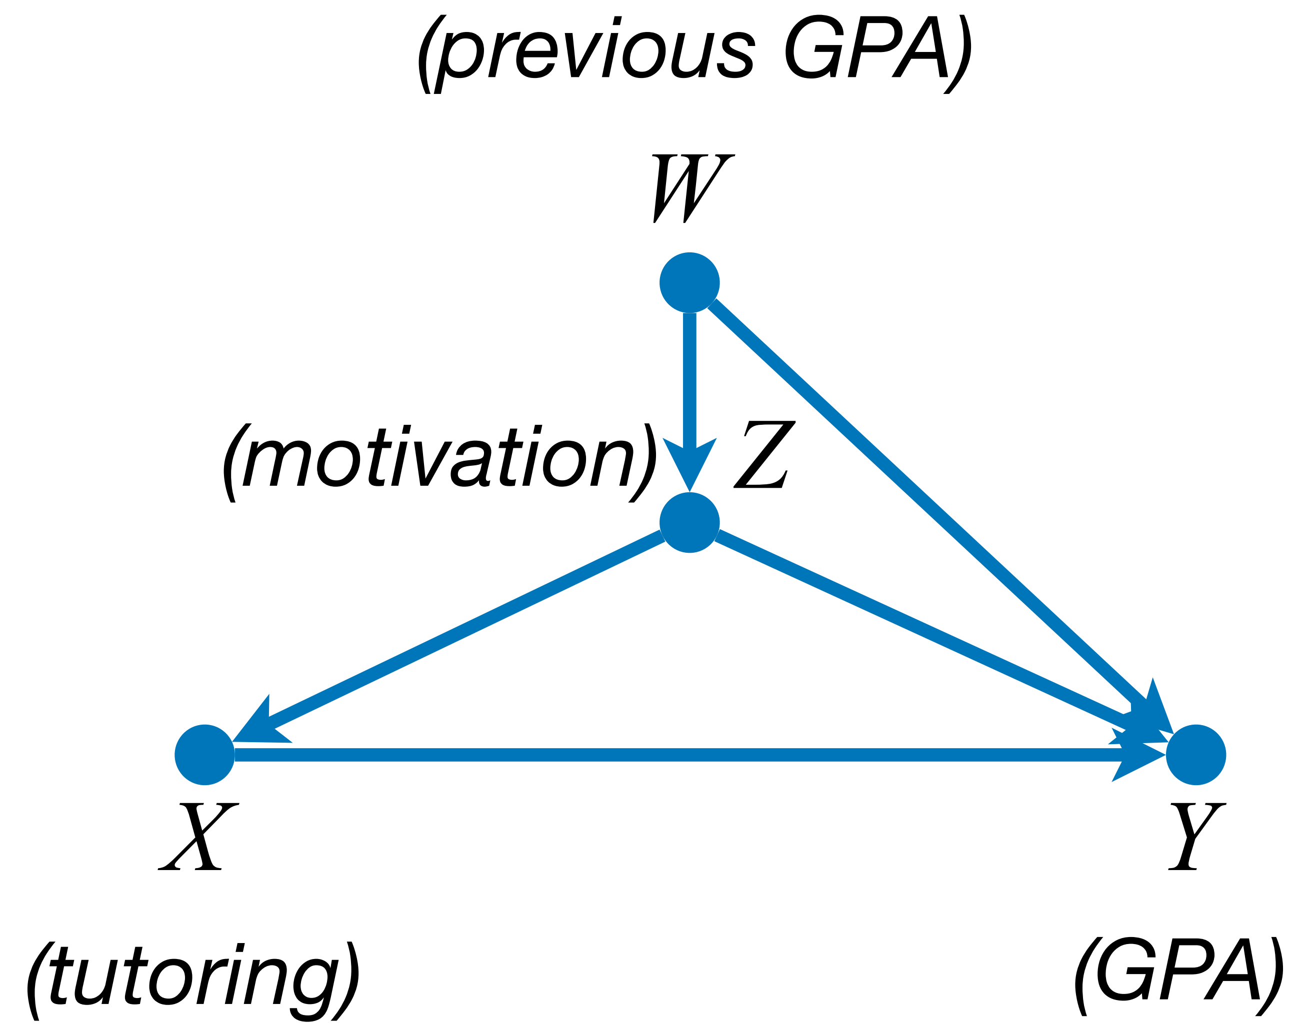

$\mathcal{G}$ **Natural (current) Regime**

Intervention $do(X = 1)$

Make tutoring mandatory for ***all*** *students.*

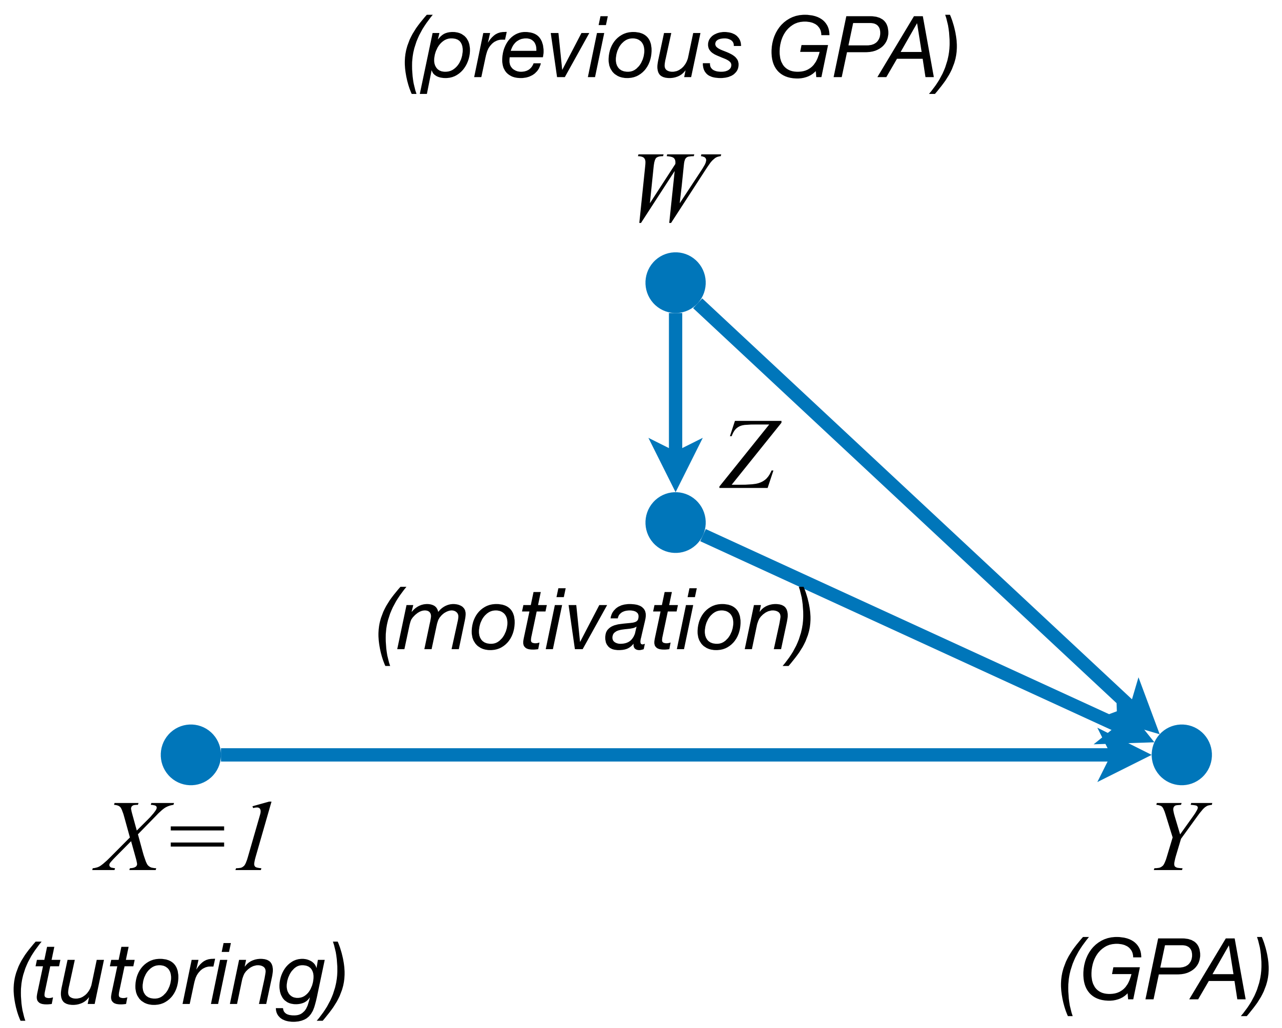

$\mathcal{G}_{\overline{X}}$ **Intervened (hypothesized) Regime**

Instead of $P(y \mid X=1)$ we are reasoning about $P(y \mid do(X=1))$, or, more generally, $P(y; \sigma_X = do(X=1))$

Soft Interventions

Soft Interventions

- A more realistic, or interesting, type of intervention is, for example, to consider the effect of making tutoring mandatory for students with historically low GPA and only to them, on their current GPA.

Soft Interventions

- A more realistic, or interesting, type of intervention is, for example, to consider the effect of making tutoring mandatory for students with historically low GPA and only to them, on their current GPA.

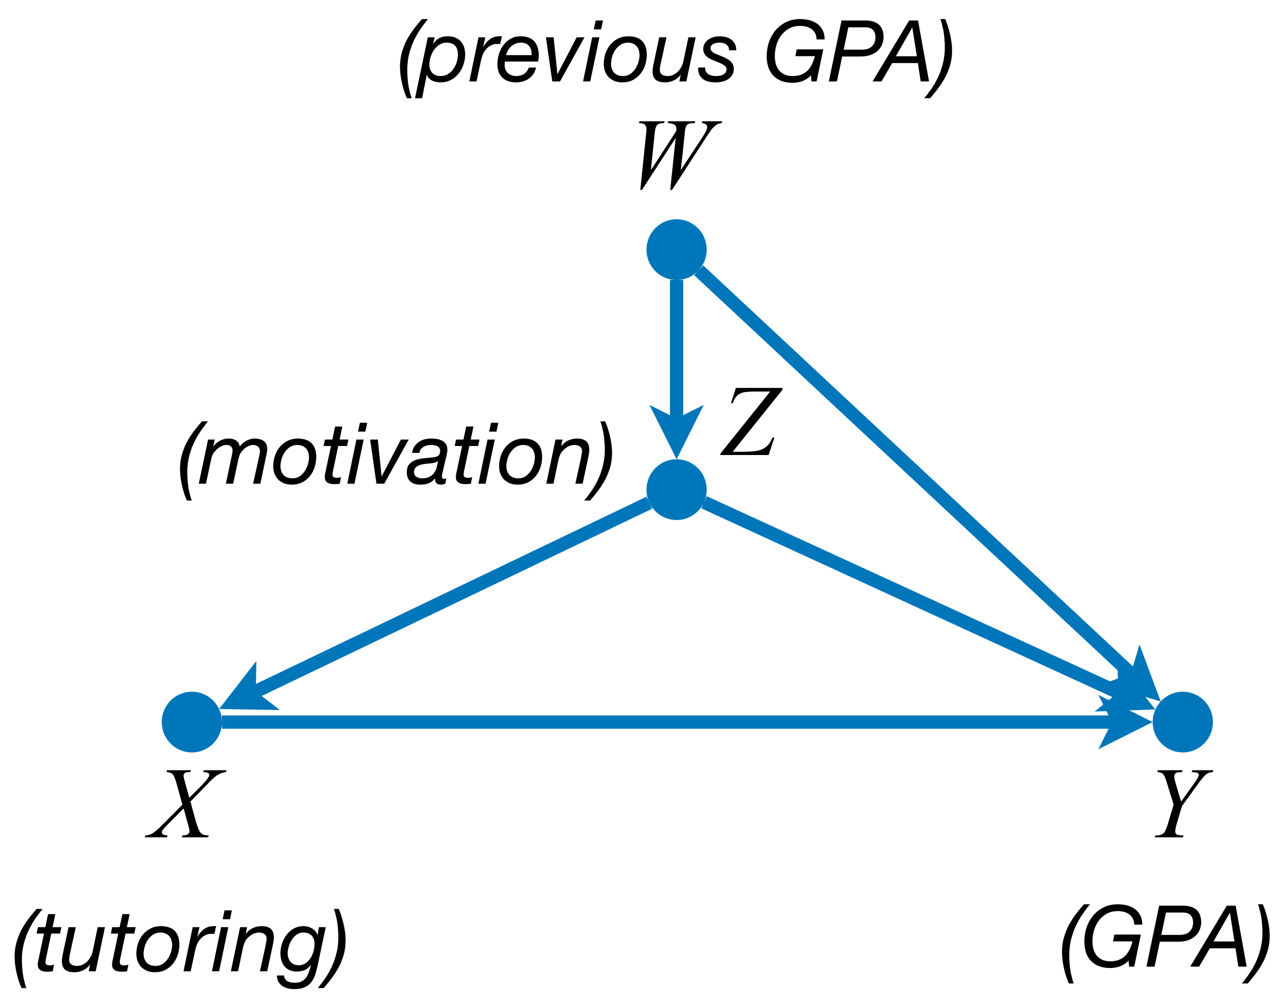

Soft Interventions

- A more realistic, or interesting, type of intervention is, for example, to consider the effect of making tutoring mandatory for students with historically low GPA and only to them, on their current GPA.

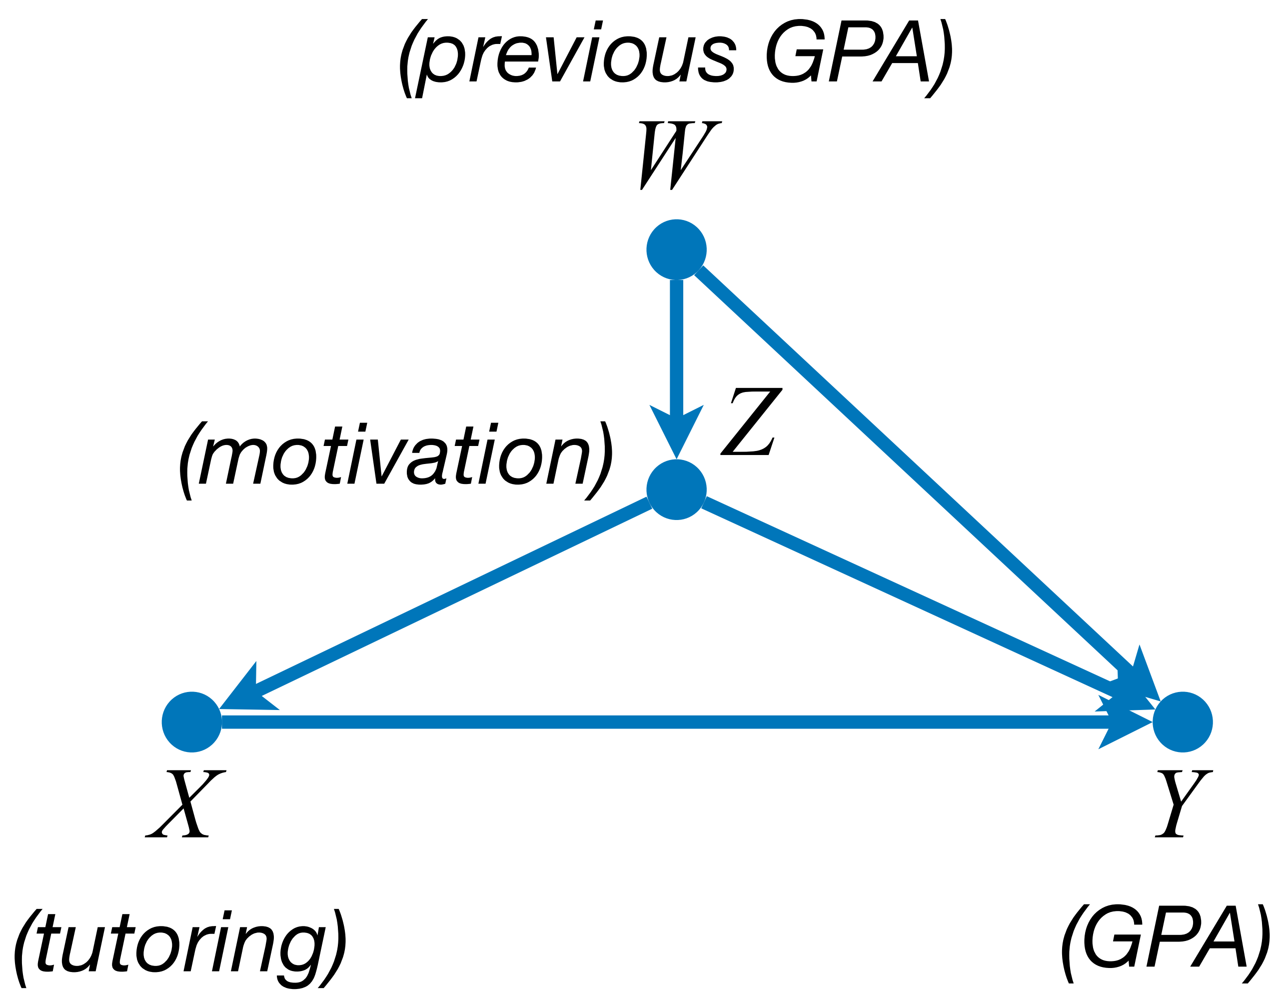

$\mathcal{G}$ **Natural (current) Regime**

Soft Interventions

- A more realistic, or interesting, type of intervention is, for example, to consider the effect of making tutoring mandatory for students with historically low GPA and only to them, on their current GPA.

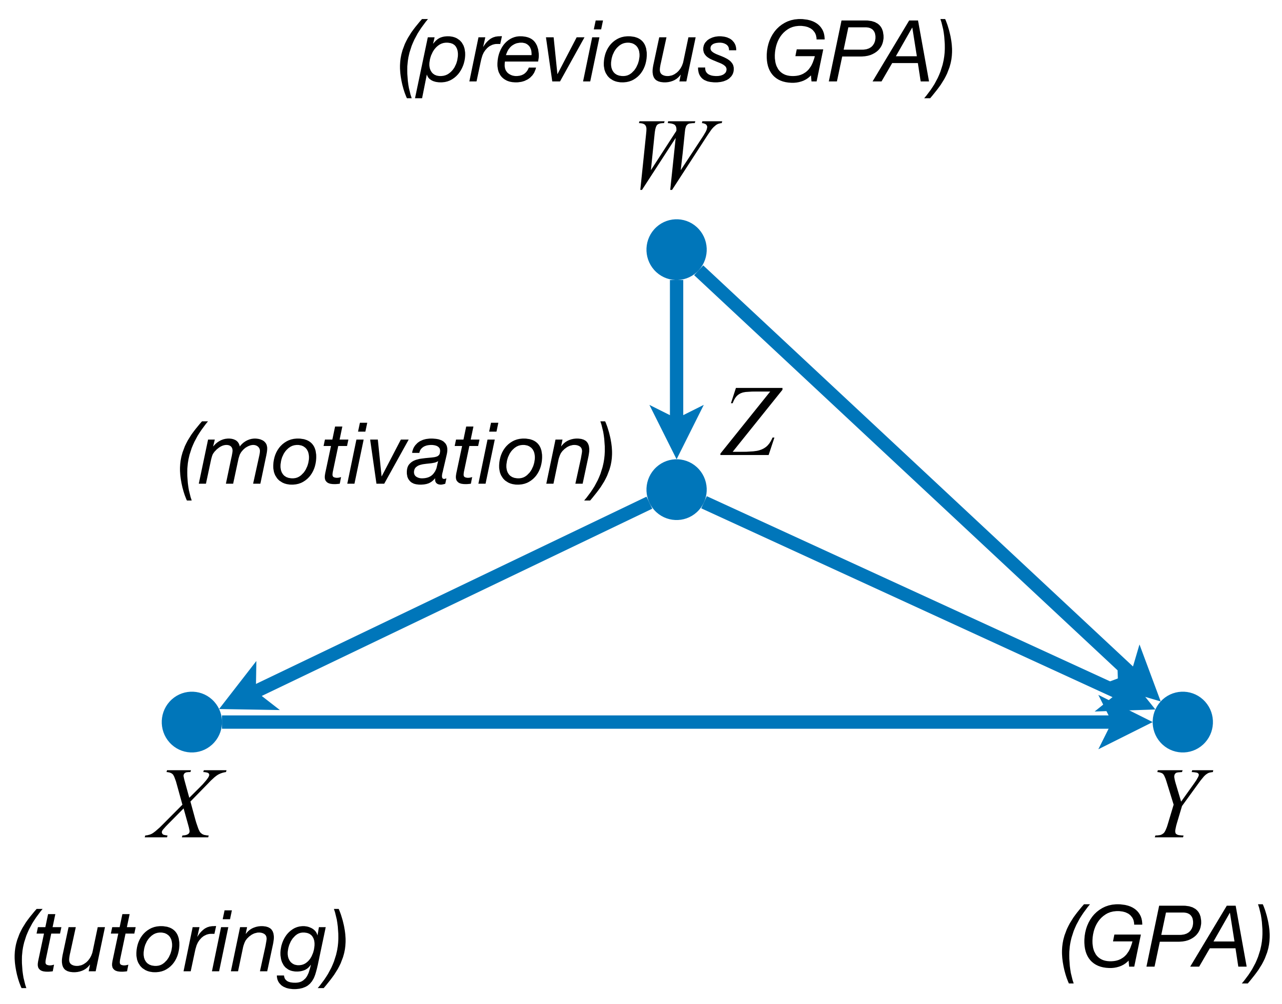

$\mathcal{G}$ **Natural (current) Regime**

Intervention $\sigma_X = 1[W = 1]$

Assign tutoring only to students with low GPA.

Soft Interventions

- A more realistic, or interesting, type of intervention is, for example, to consider the effect of making tutoring mandatory for students with historically low GPA and only to them, on their current GPA.

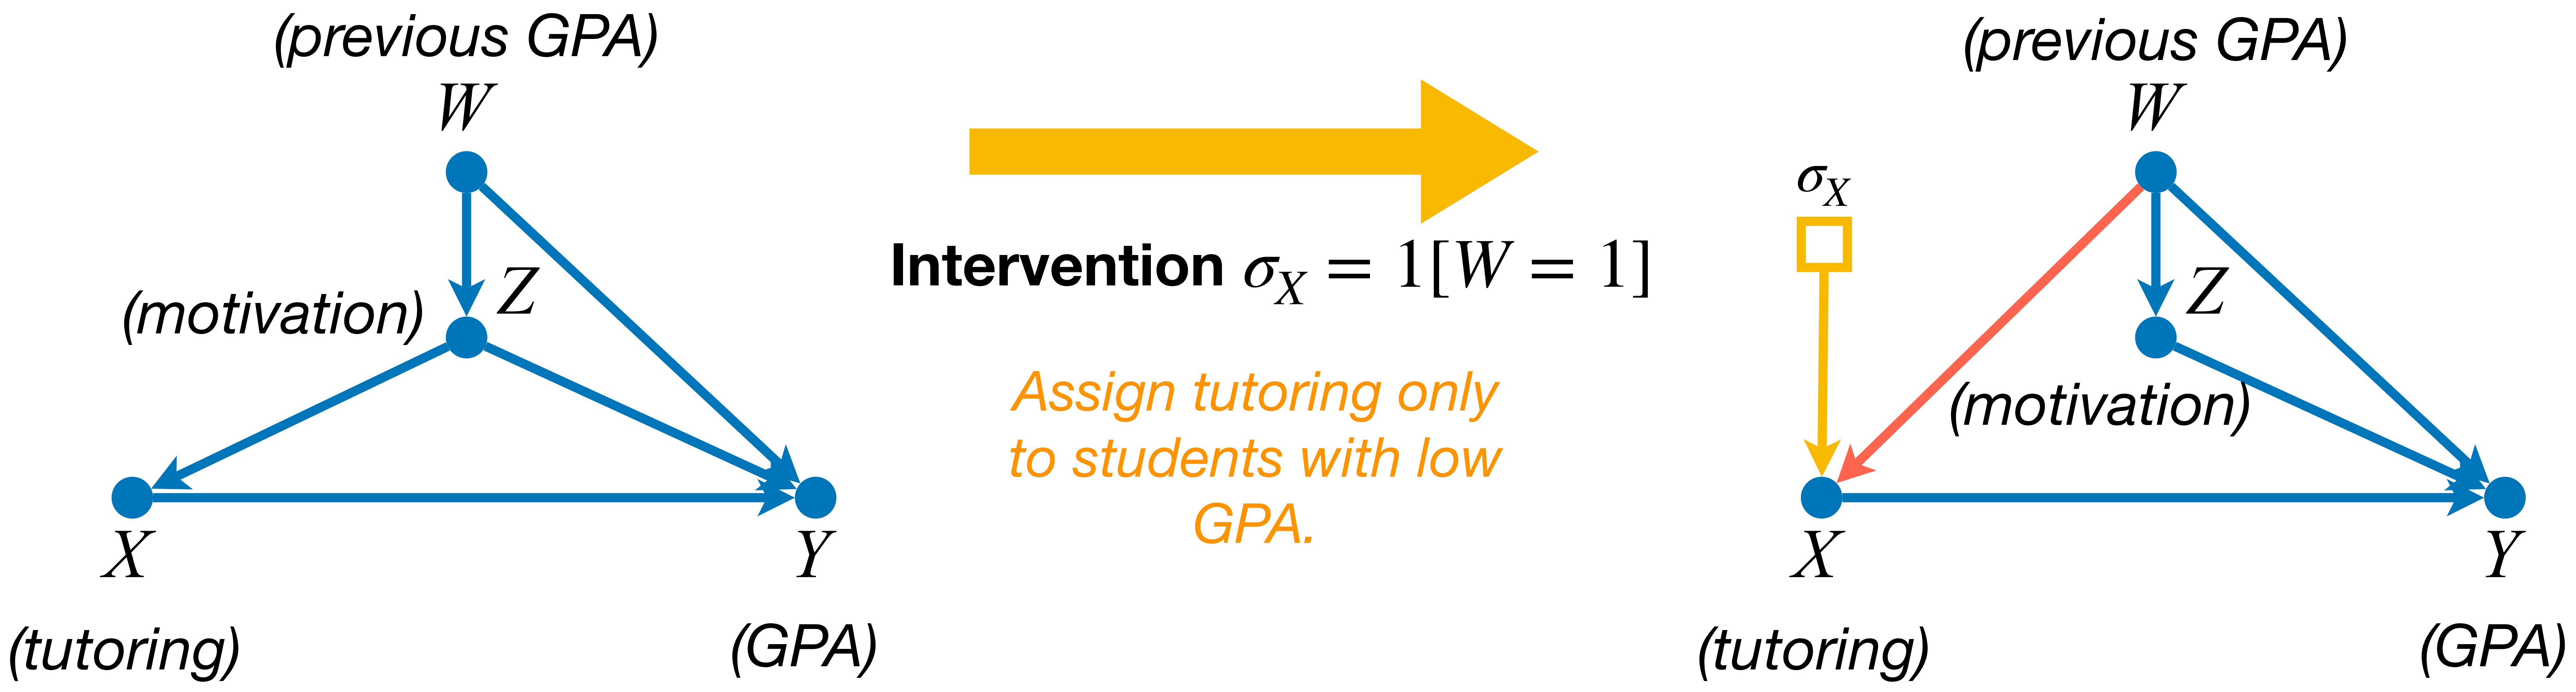

$\mathcal{G}$ **Natural (current) Regime**

Soft Interventions

- A more realistic, or interesting, type of intervention is, for example, to consider the effect of making tutoring mandatory for students with historically low GPA and only to them, on their current GPA.

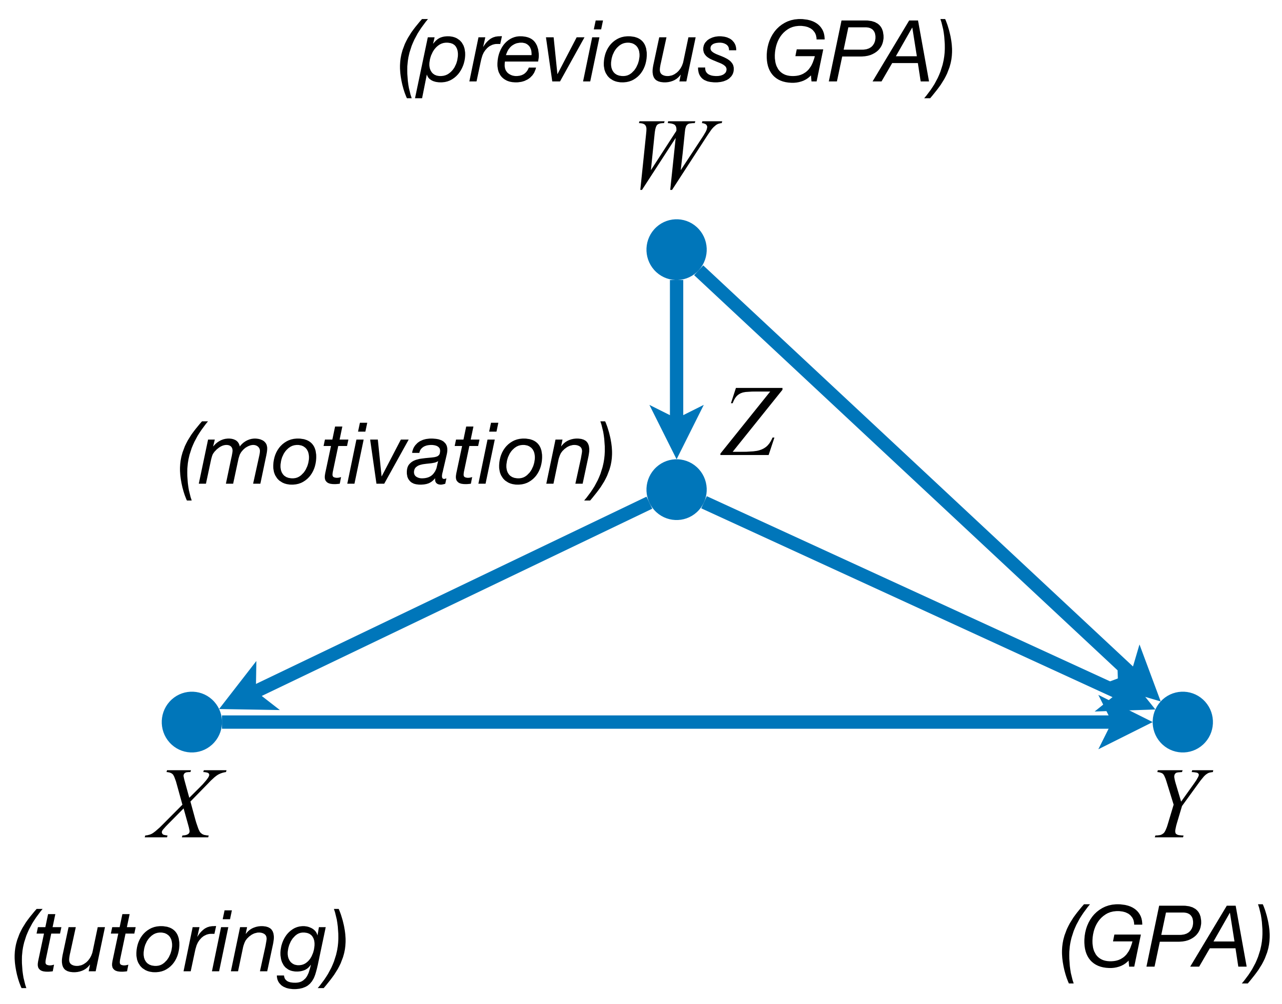

$\mathcal{G}$ **Natural (current) Regime**

Intervention $\sigma_X = 1[W = 1]$

Assign tutoring only to students with low GPA.

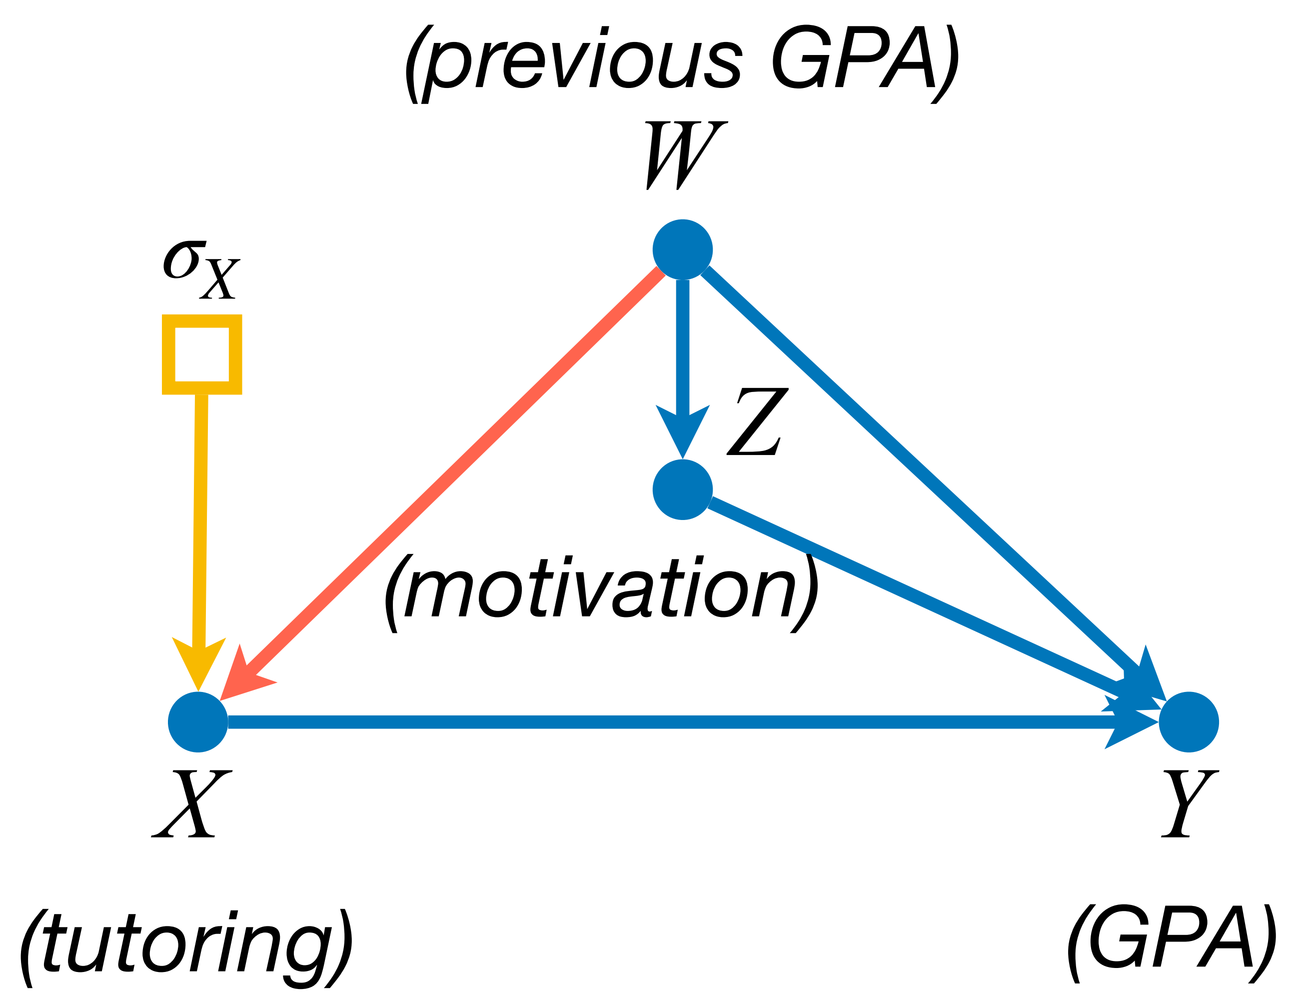

$\mathcal{G}_{\sigma_X}$ **Intervened (hypothesized) Regime**

Soft Interventions

- A more realistic, or interesting, type of intervention is, for example, to consider the effect of making tutoring mandatory for students with historically low GPA and only to them, on their current GPA.

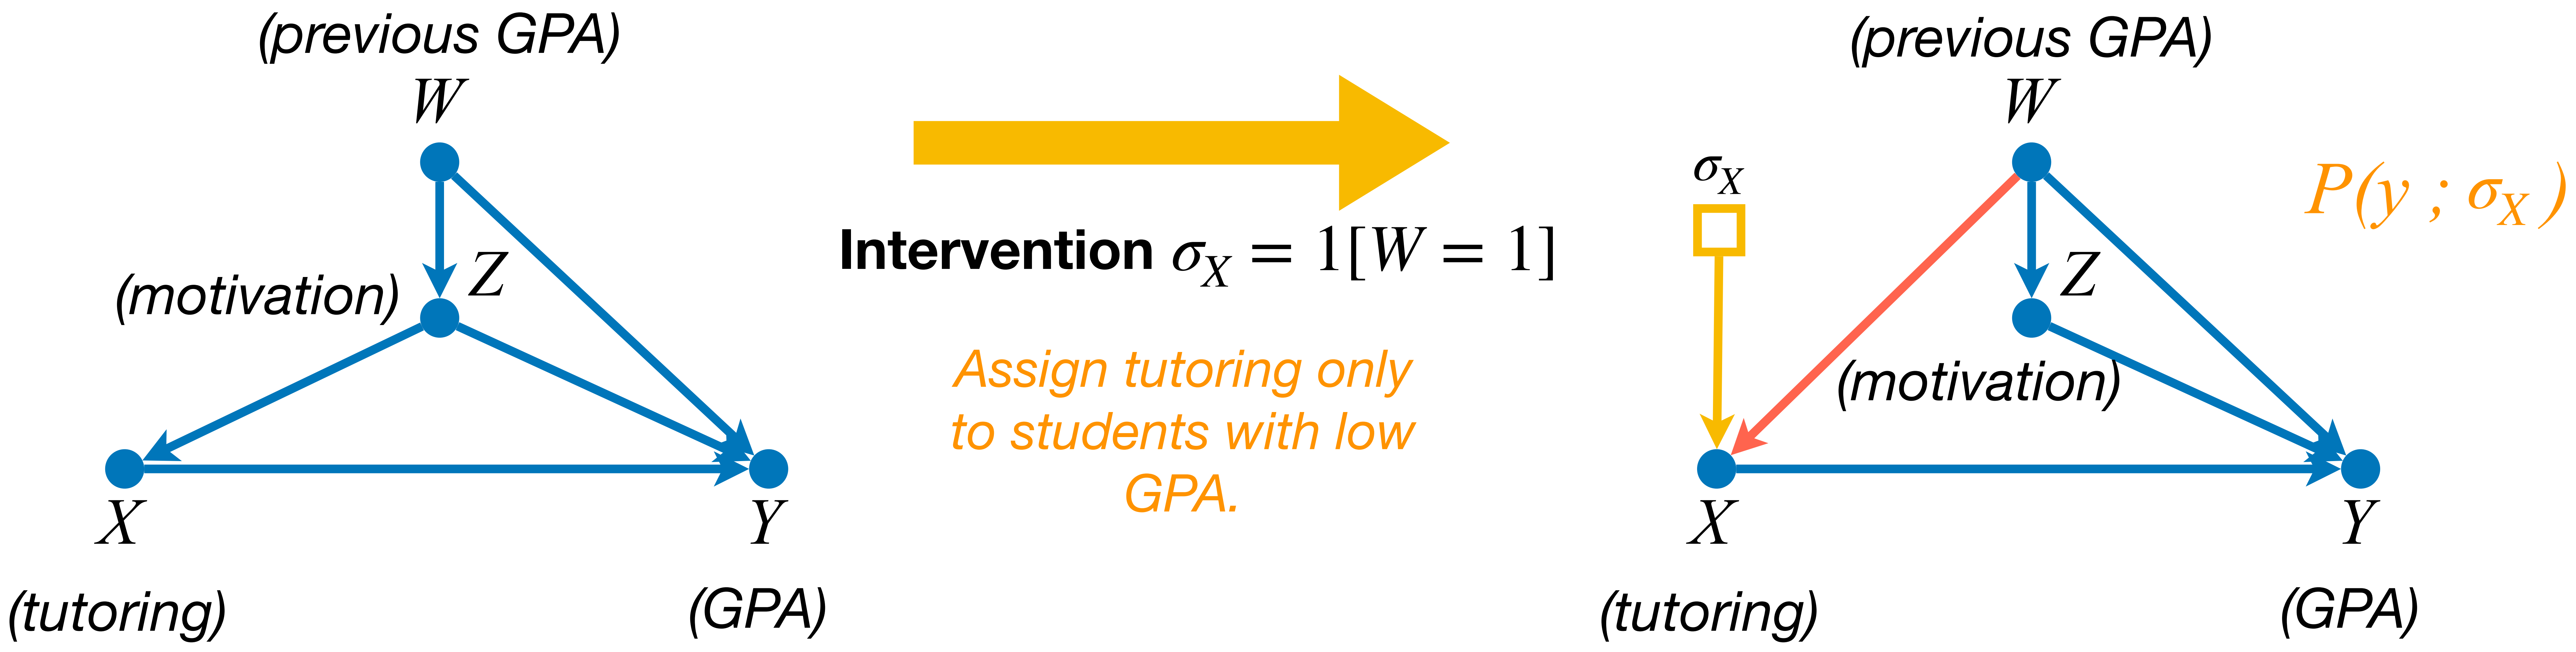

$\mathcal{G}$ **Natural (current) Regime**

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

σ-calculus (simplified)

σ-calculus (simplified)

- *Insertion/deletion of observations:*

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

σ-calculus (simplified)

- *Insertion/deletion of observations:*

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

- *Change of regimes under observation:*

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

σ-calculus (simplified)

- *Insertion/deletion of observations:*

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

- *Change of regimes under observation:*

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

- *Change of regimes without observations:*

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

Using σ-calculus

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

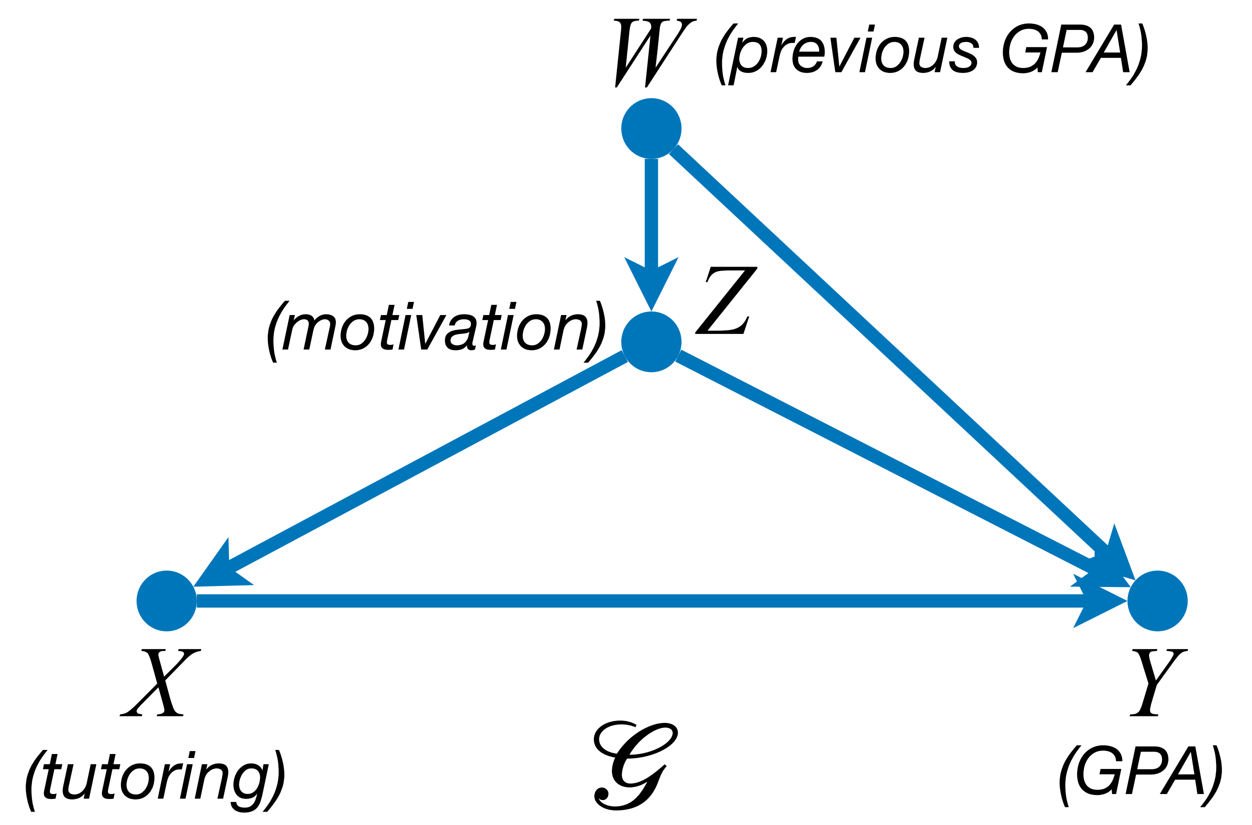

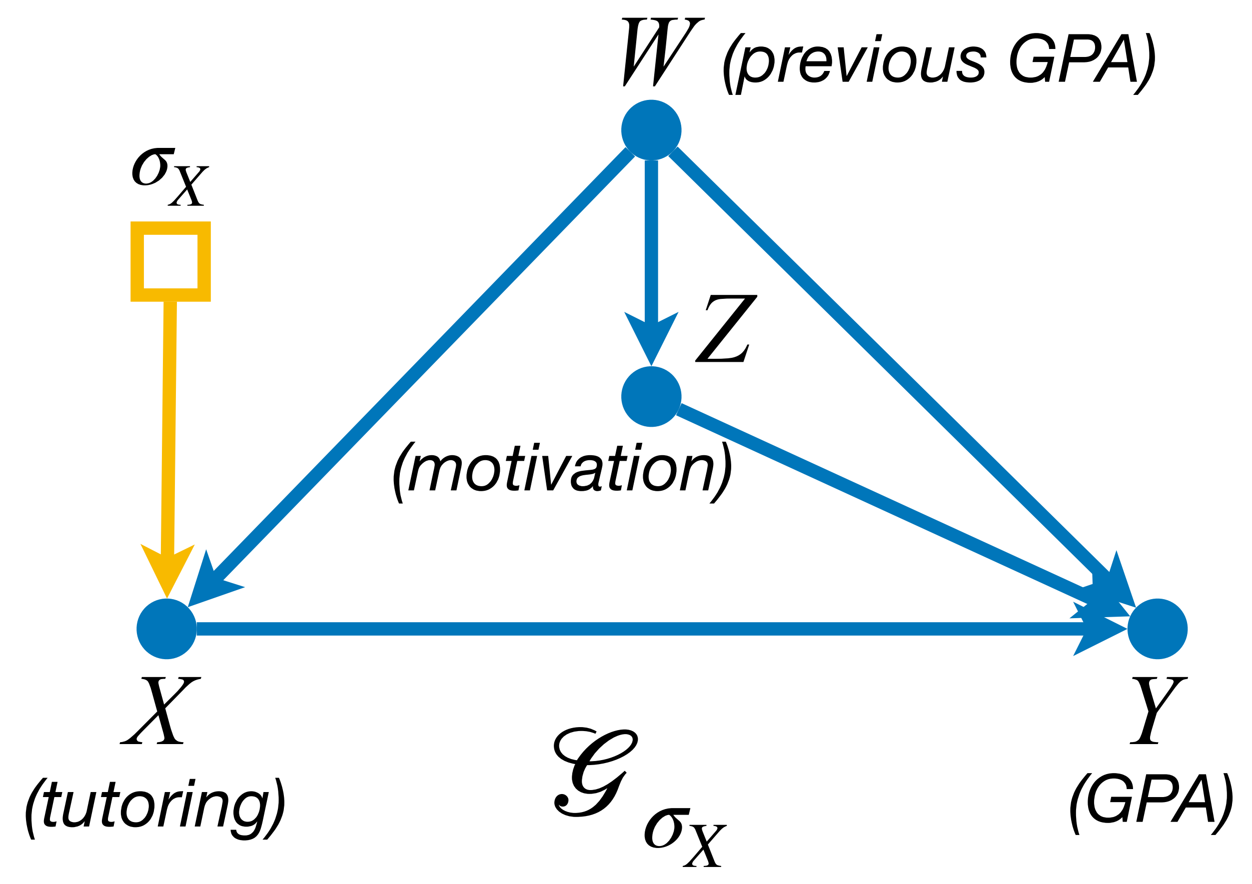

Using σ-calculus

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

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

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

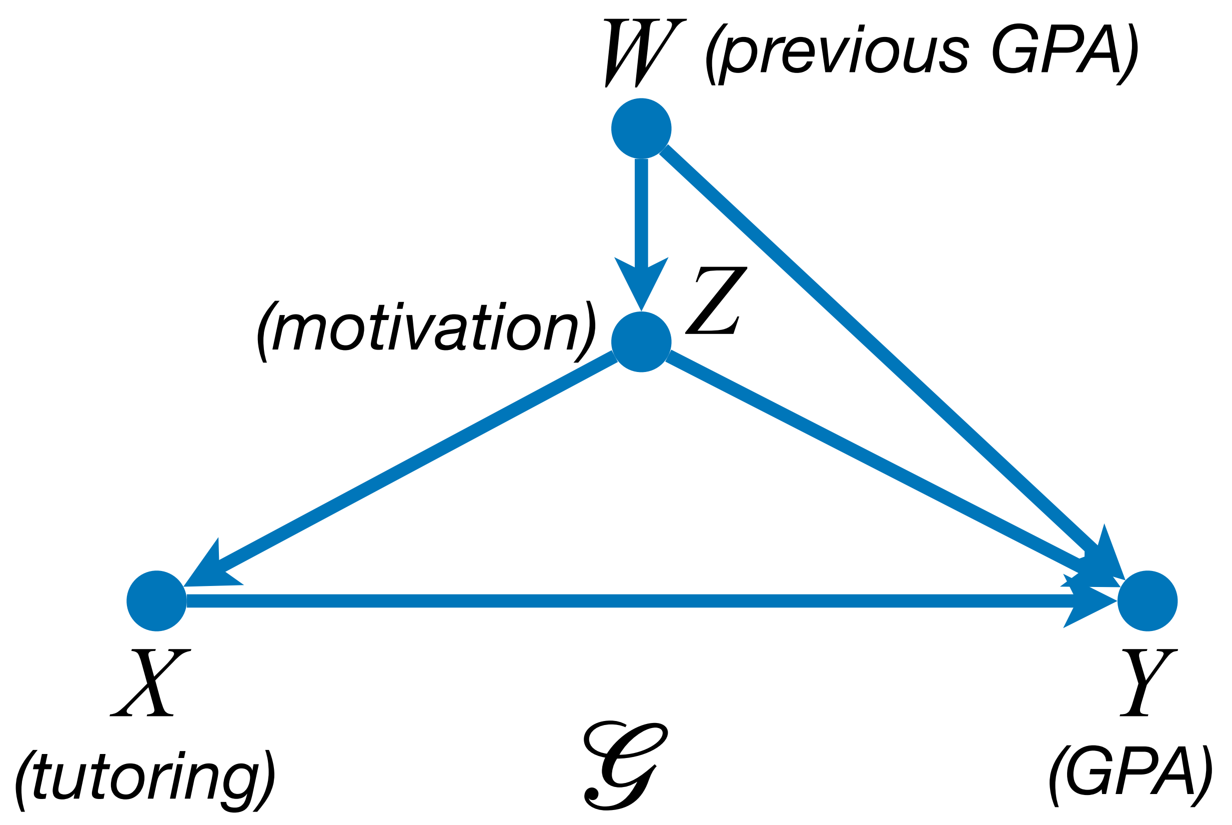

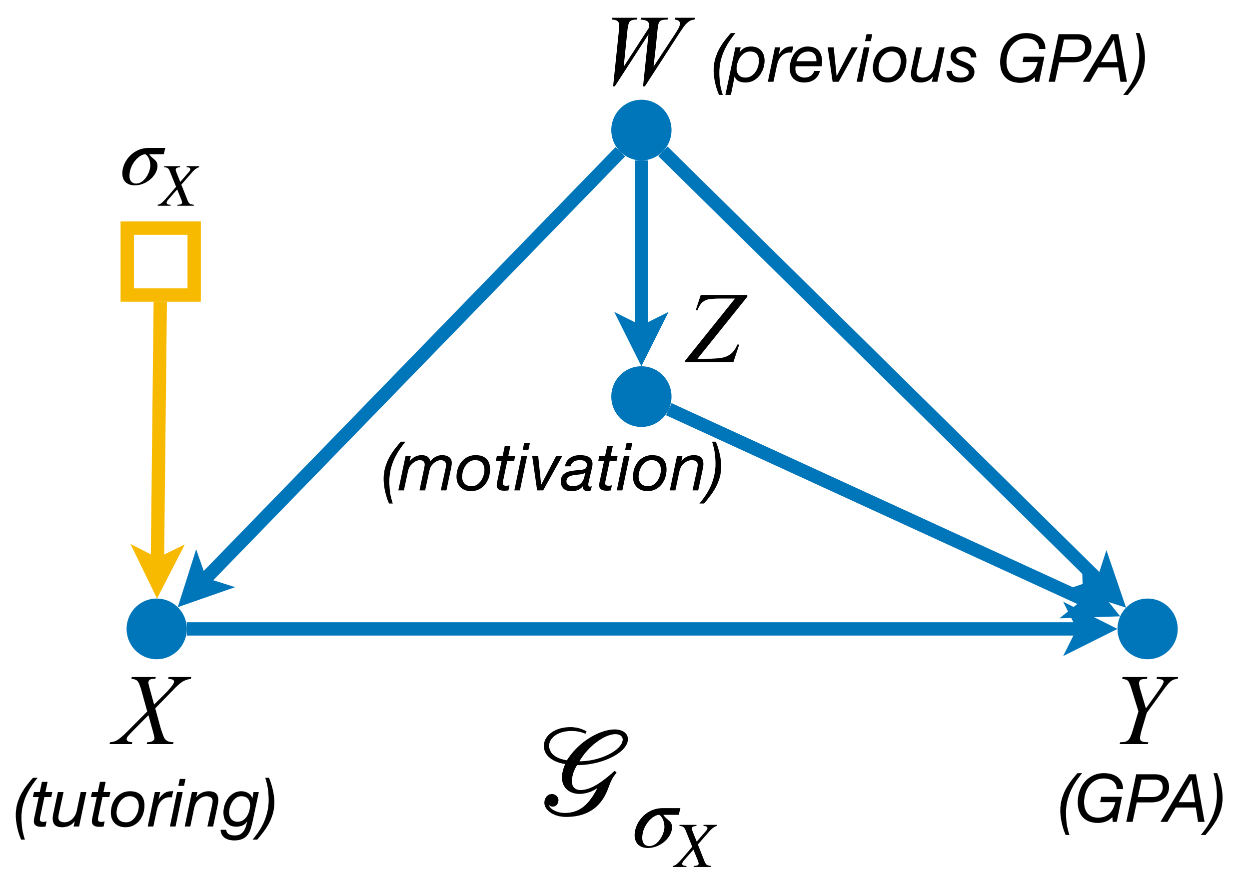

Using σ-calculus

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

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

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

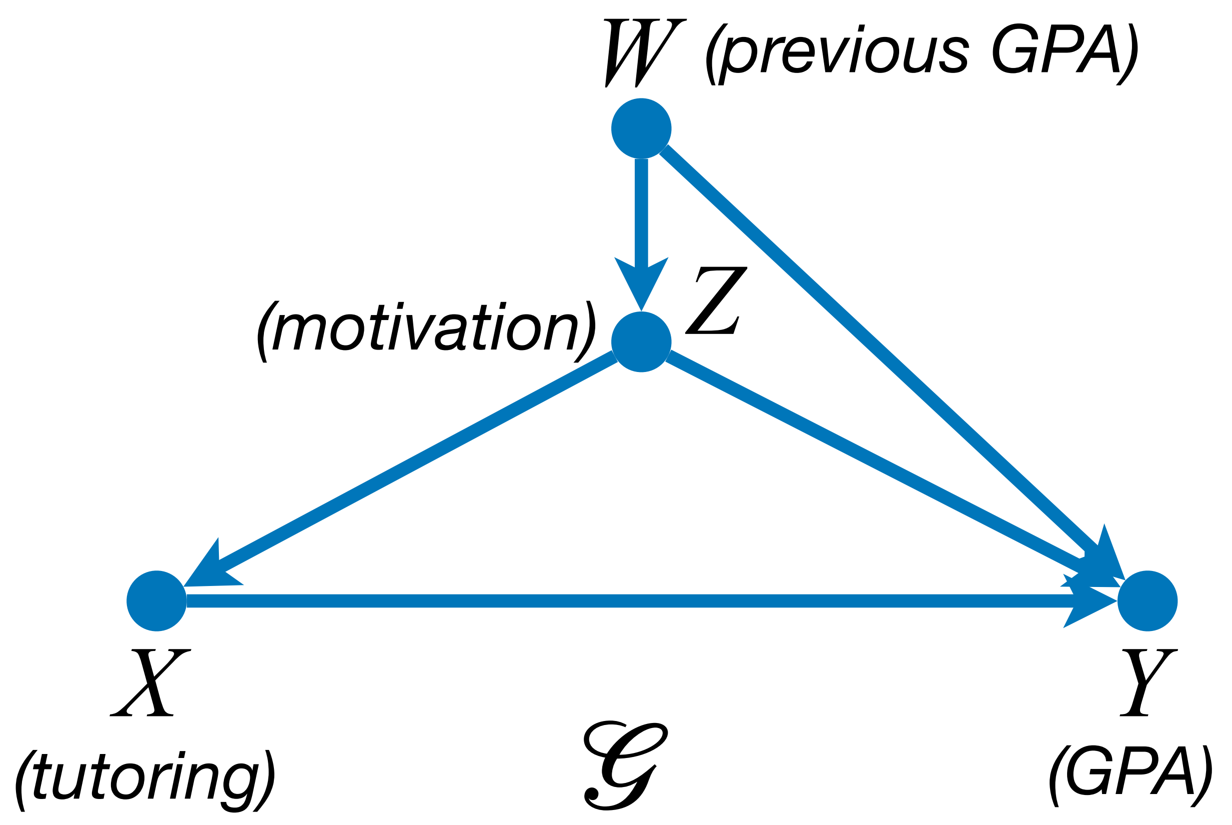

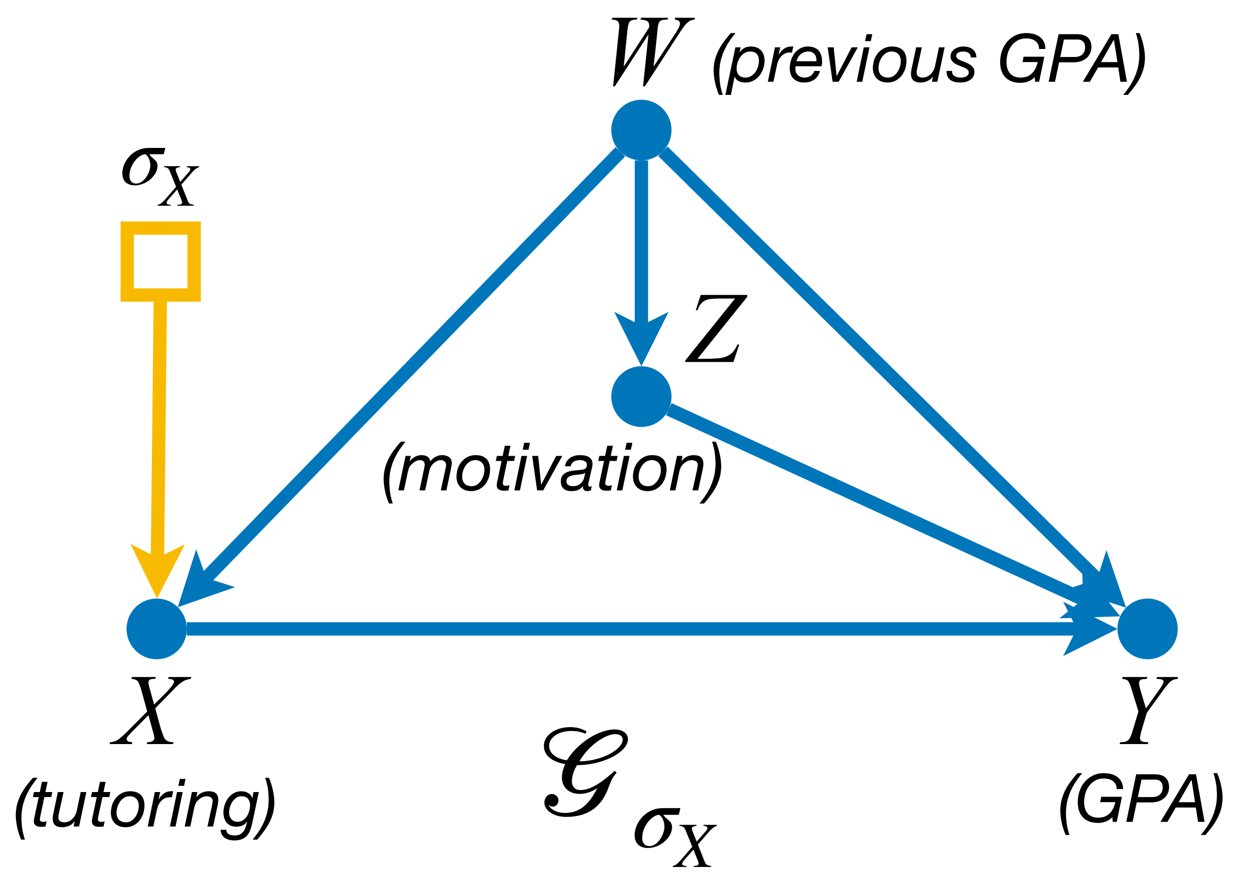

Using σ-calculus

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

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

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

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

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

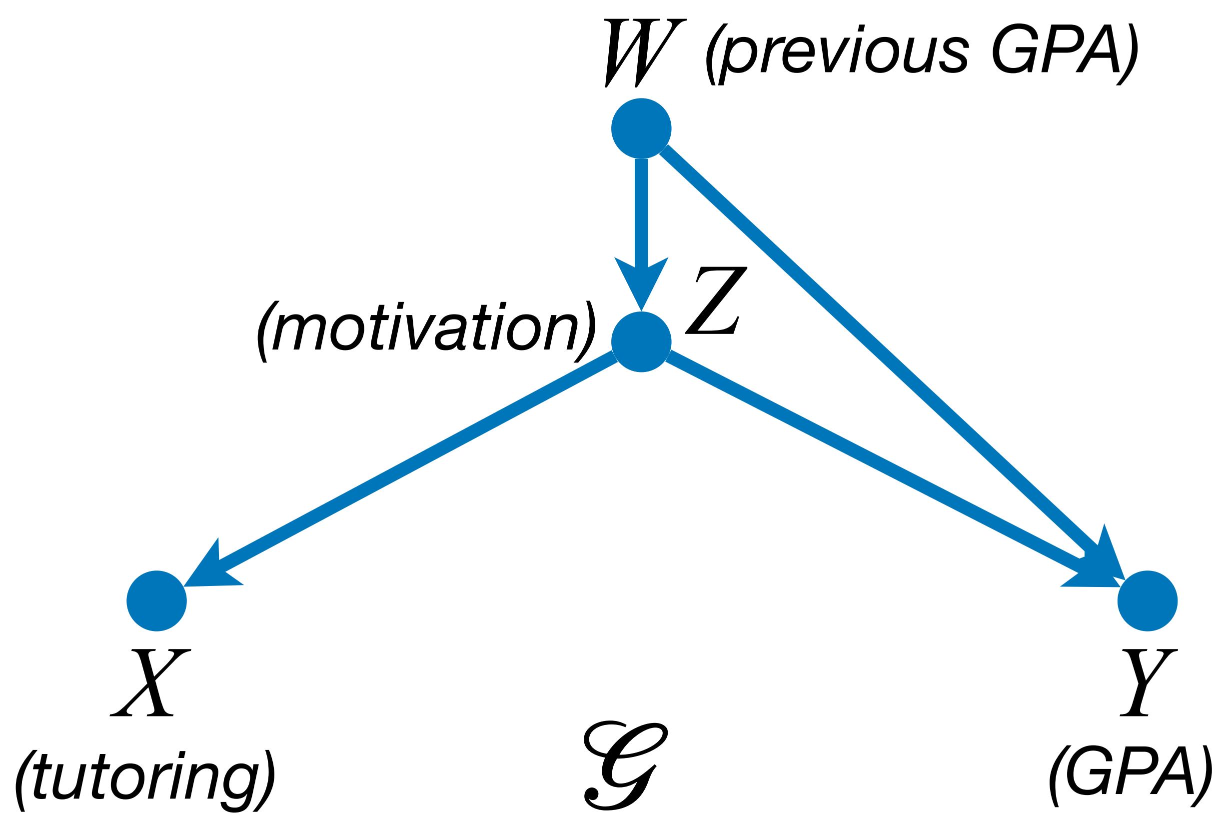

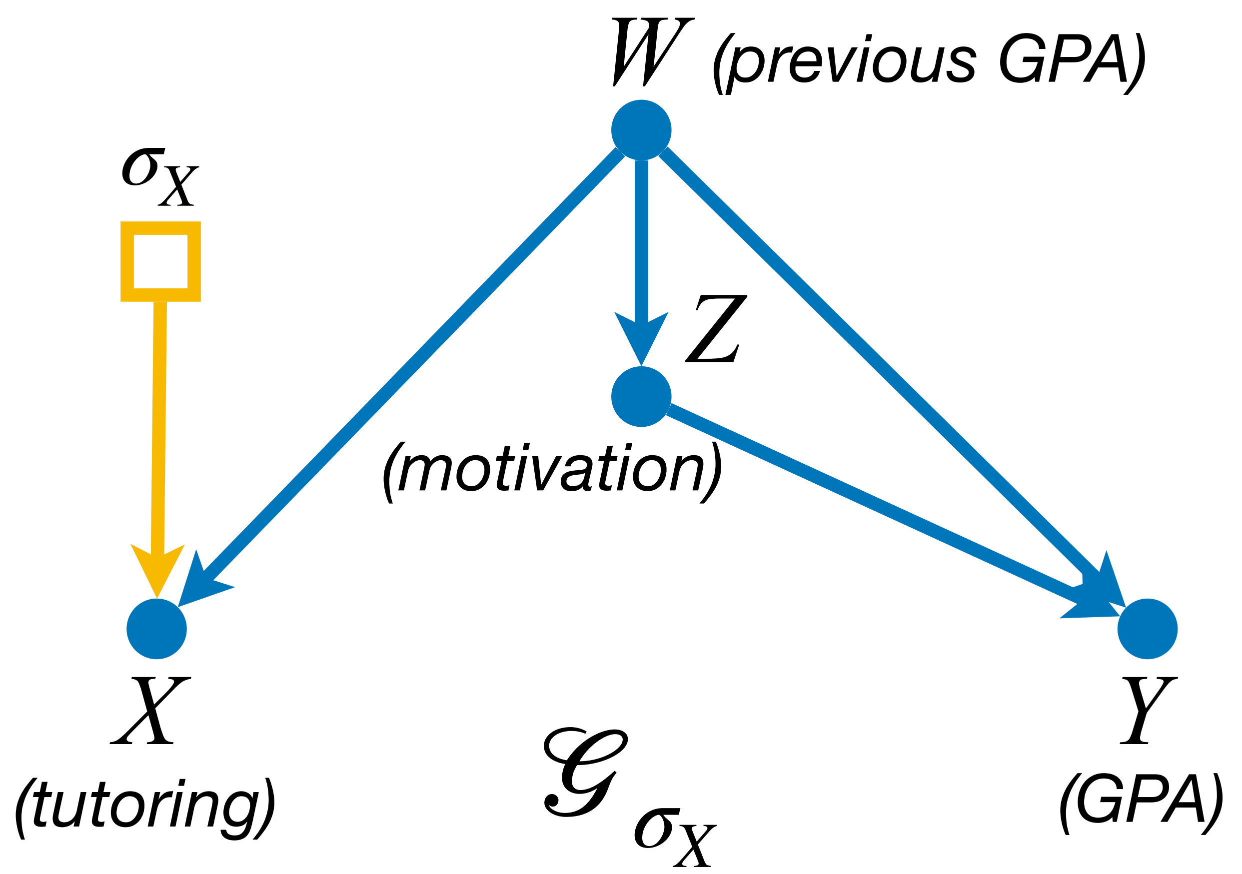

Using σ-calculus

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

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

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

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

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

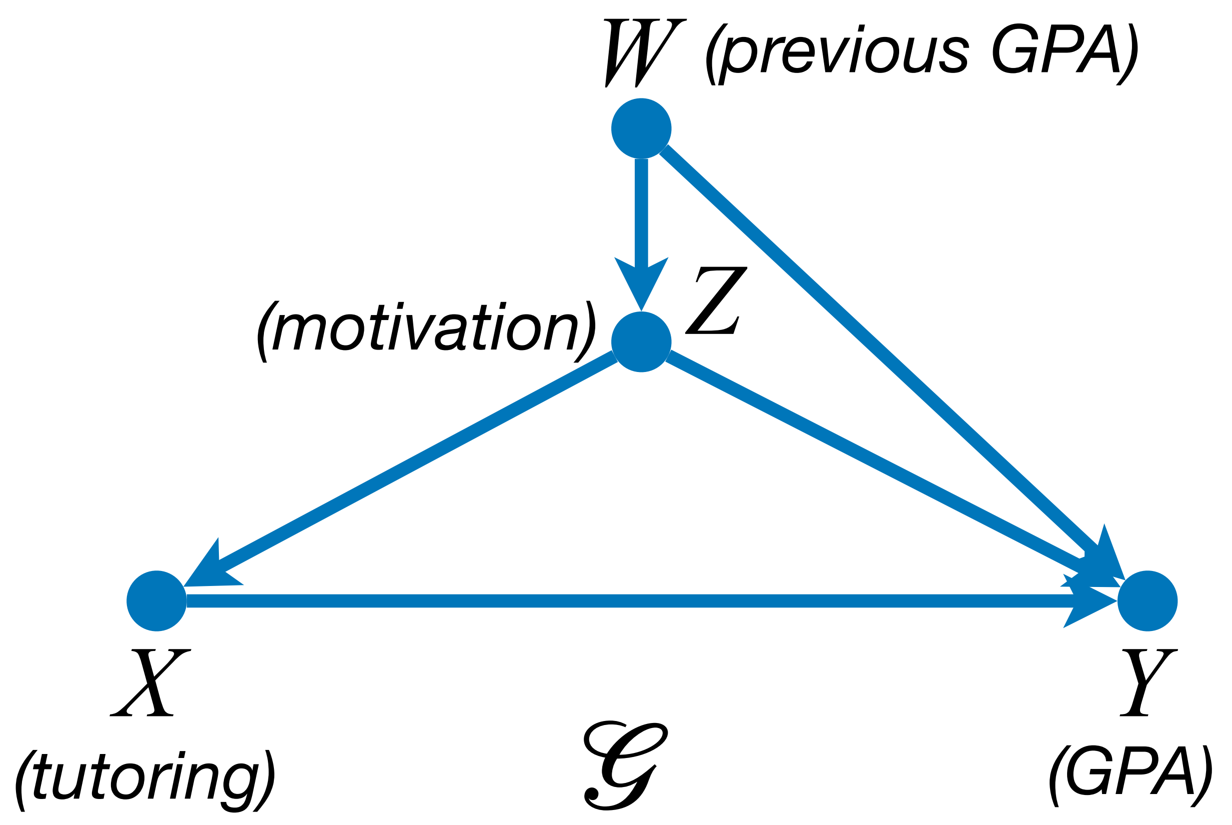

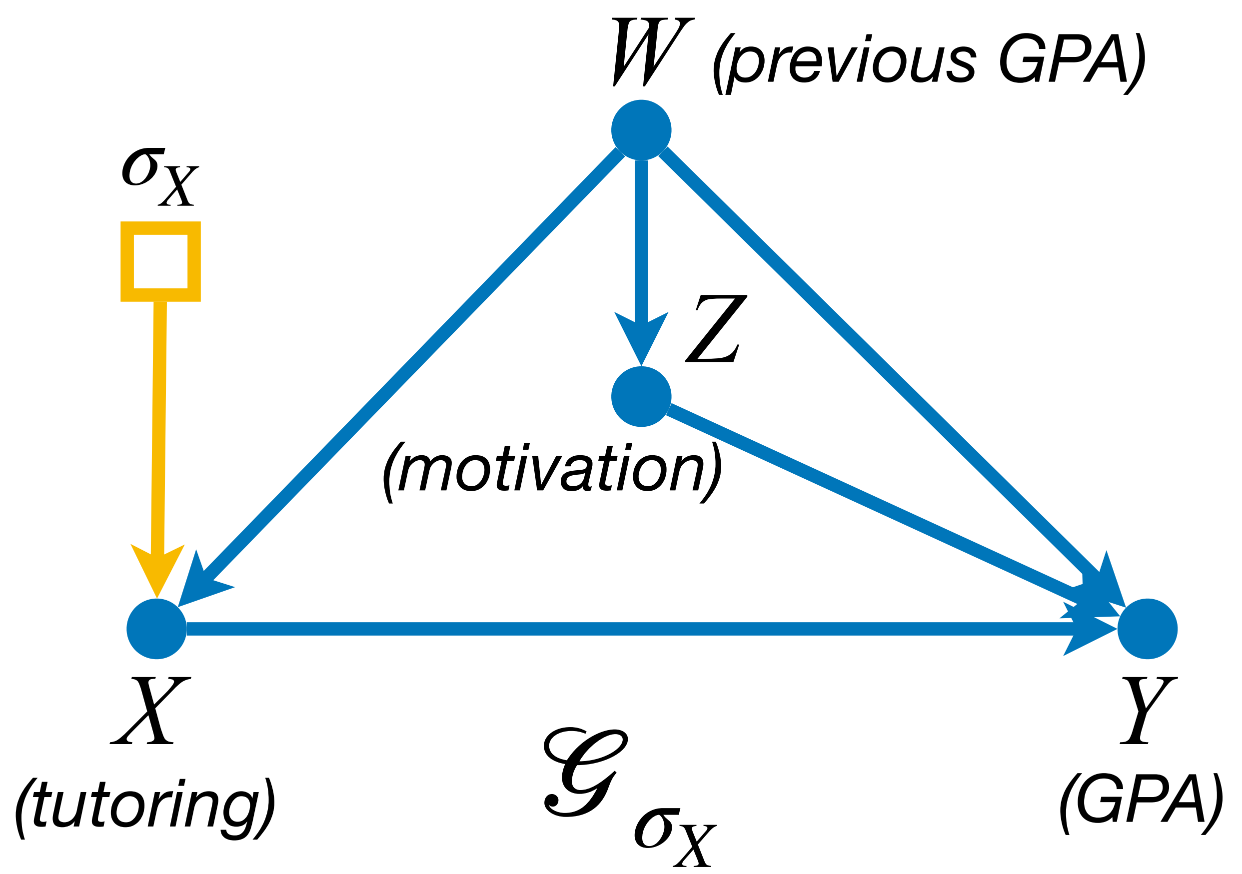

Using σ-calculus

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

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

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

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

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

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

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

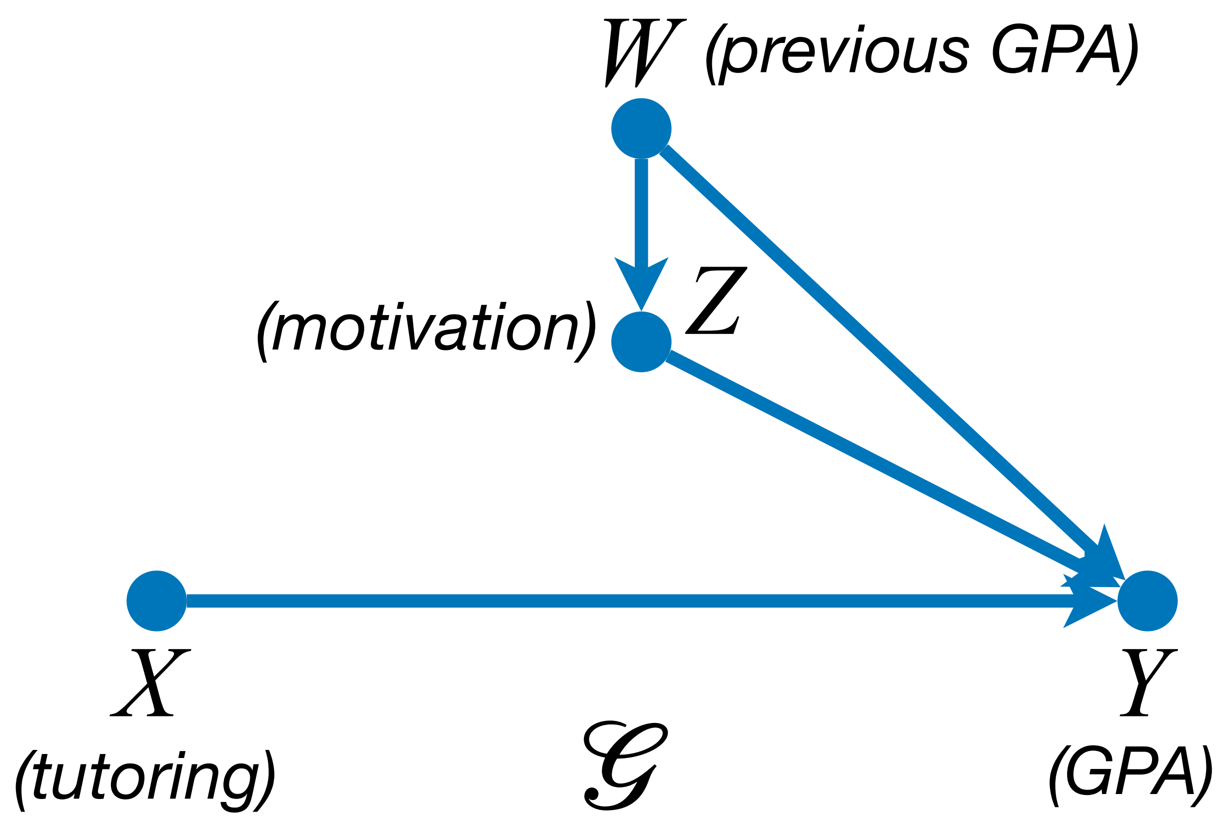

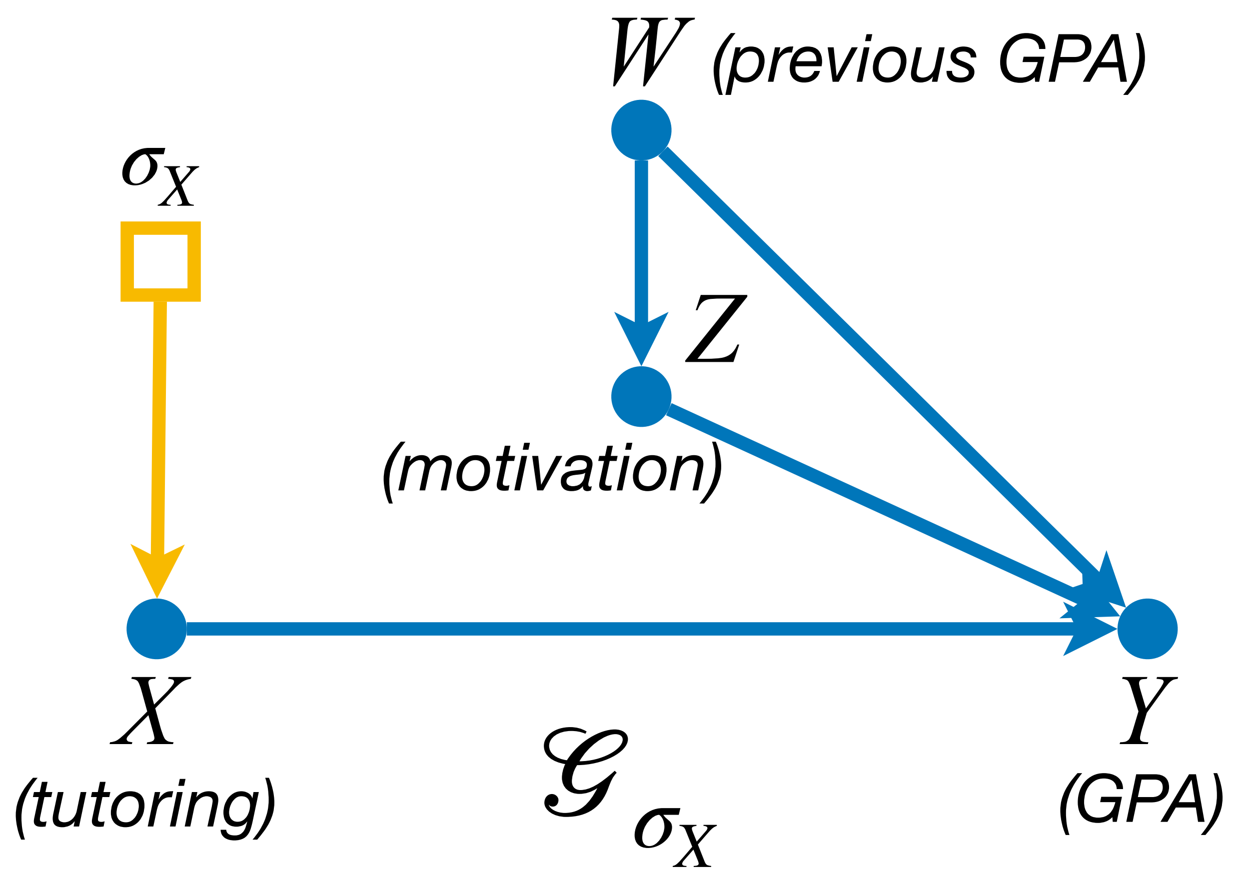

Using σ-calculus

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

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

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

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

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

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

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

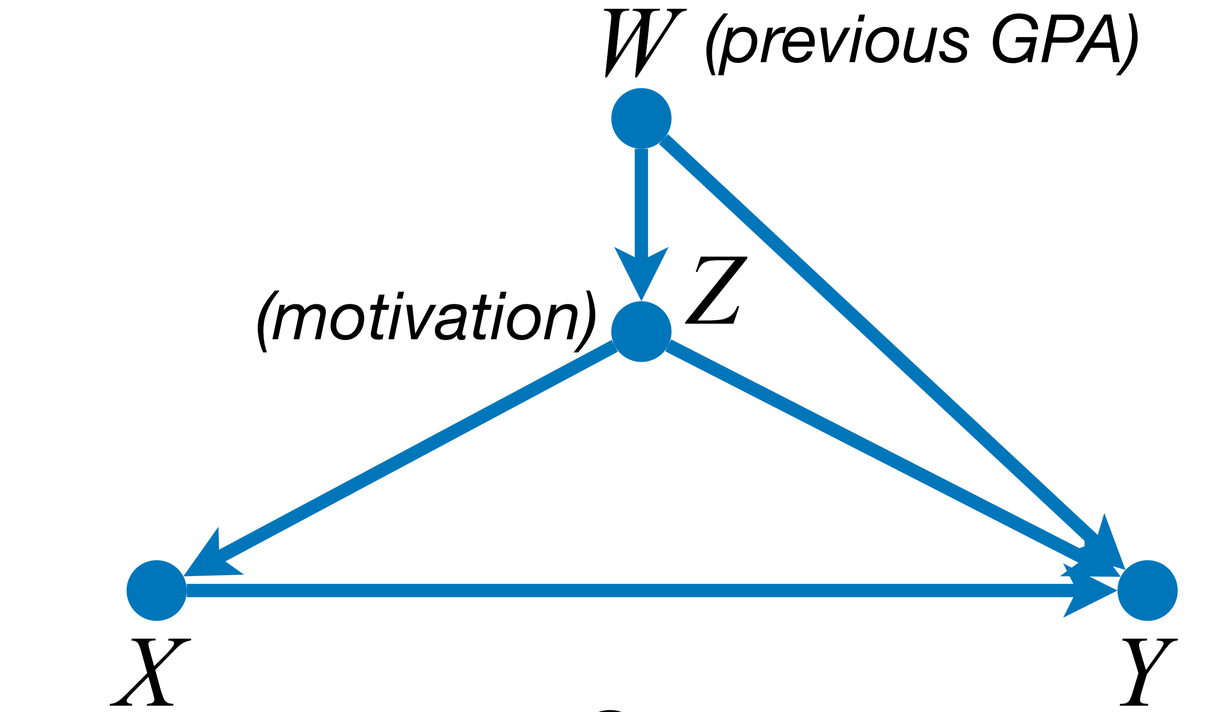

$\mathcal{G}$

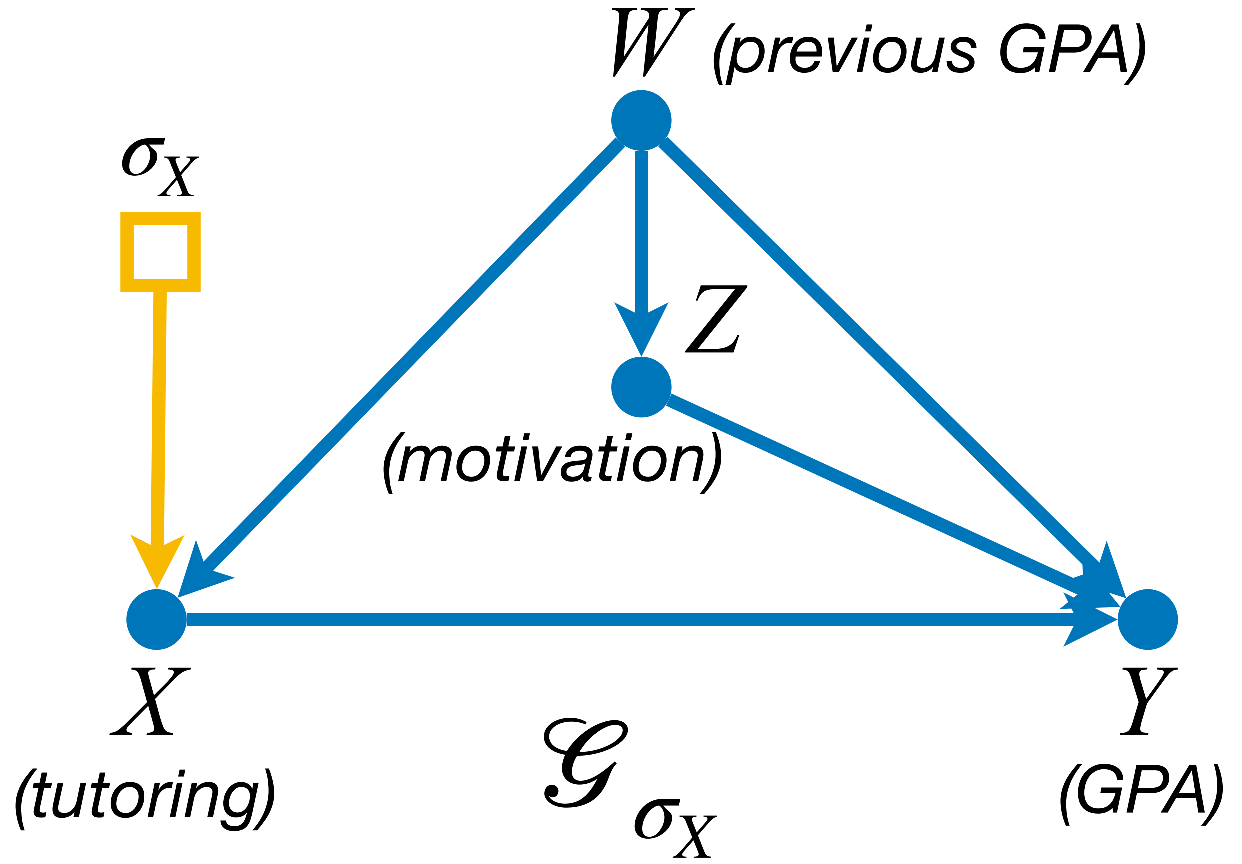

$\mathcal{G}_{\sigma_X}$

Using σ-calculus

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

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

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

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

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

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

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

Defined by σ_X

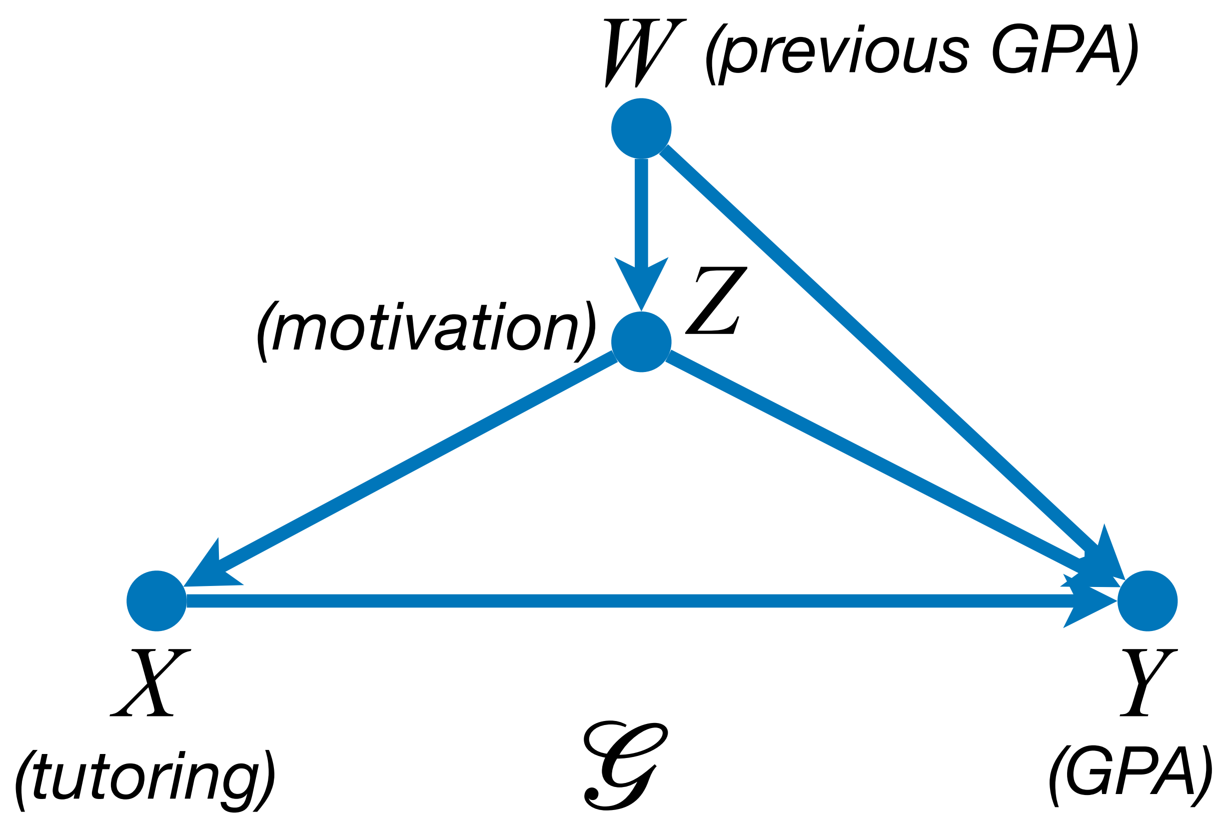

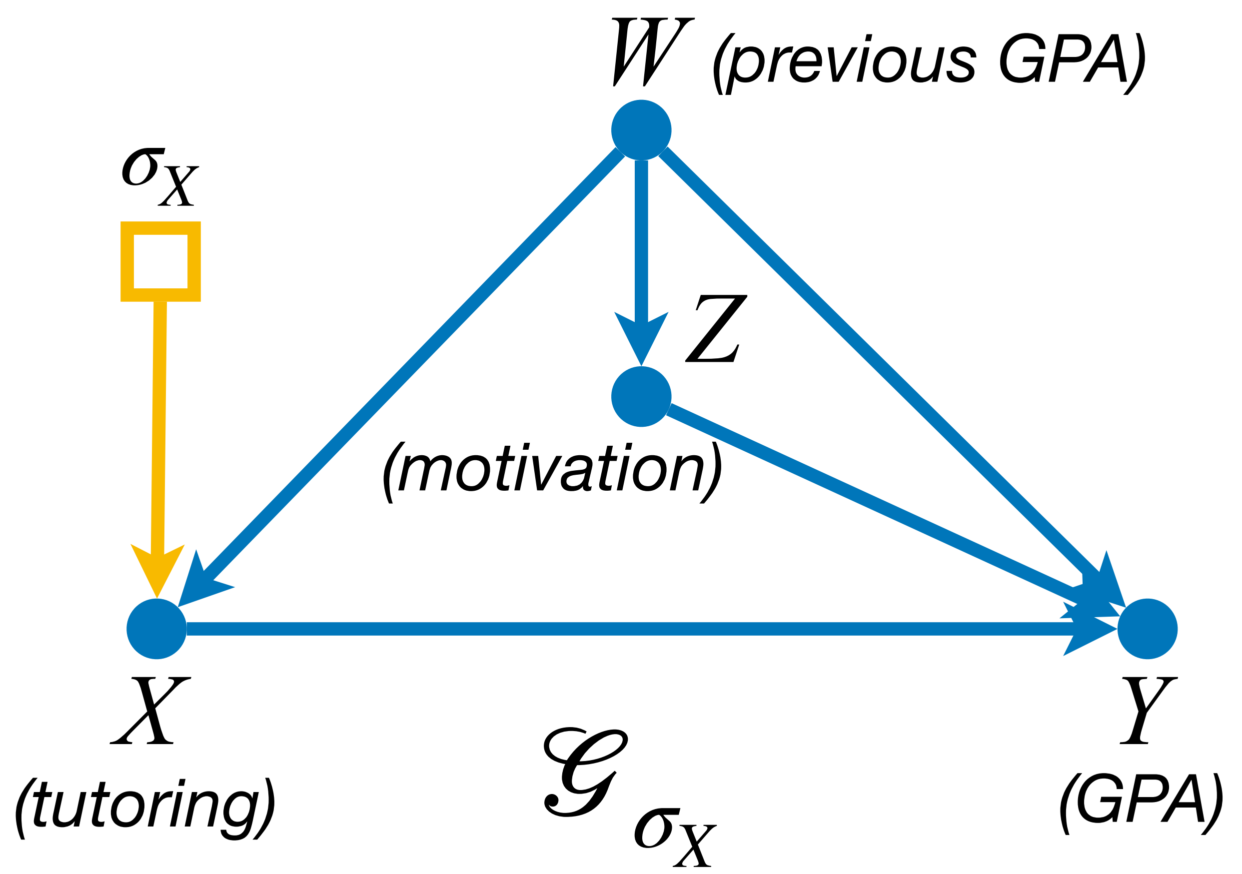

Using σ-calculus

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

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

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

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

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

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

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

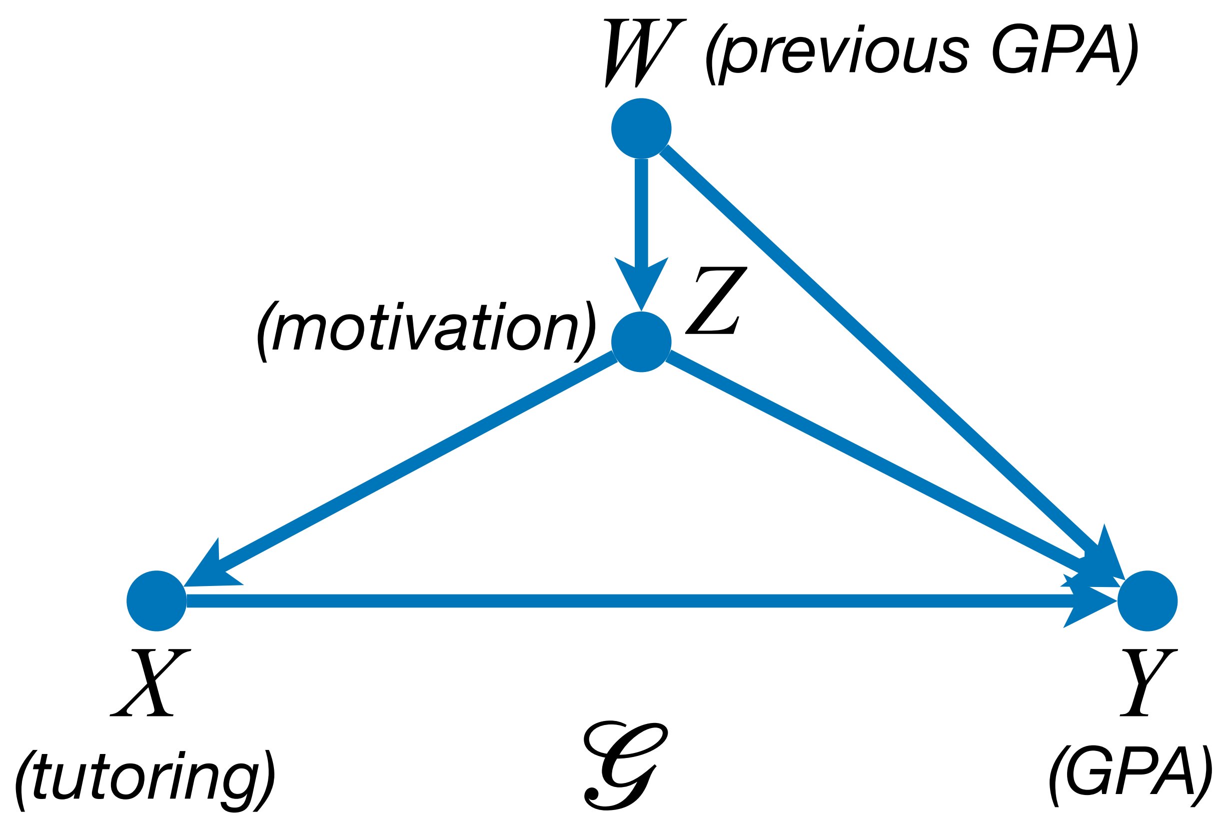

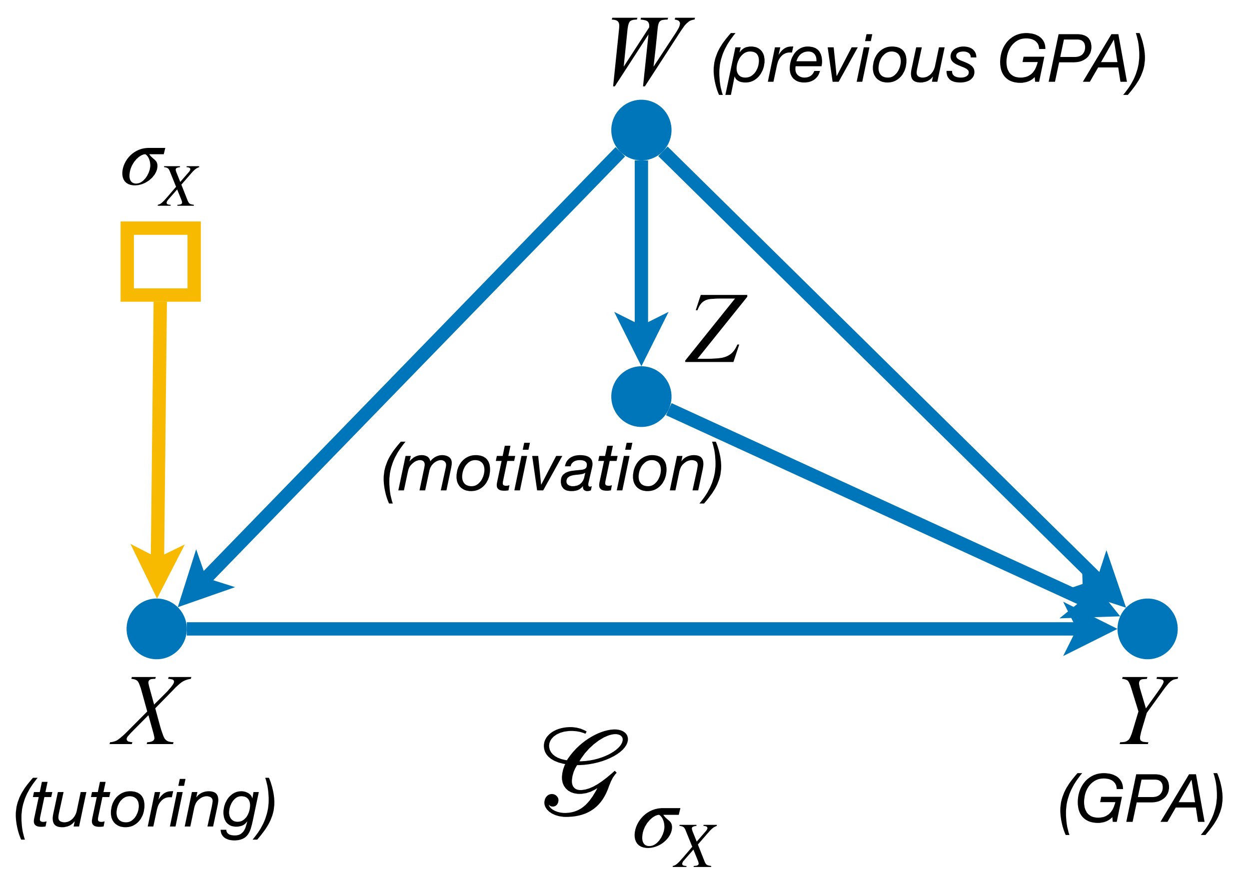

Surrogate Experiments

Surrogate Experiments

- It's not uncommon that the effect of a certain intervention is not identifiable (not uniquely computable) from observational data alone whenever unobserved confounders are present.

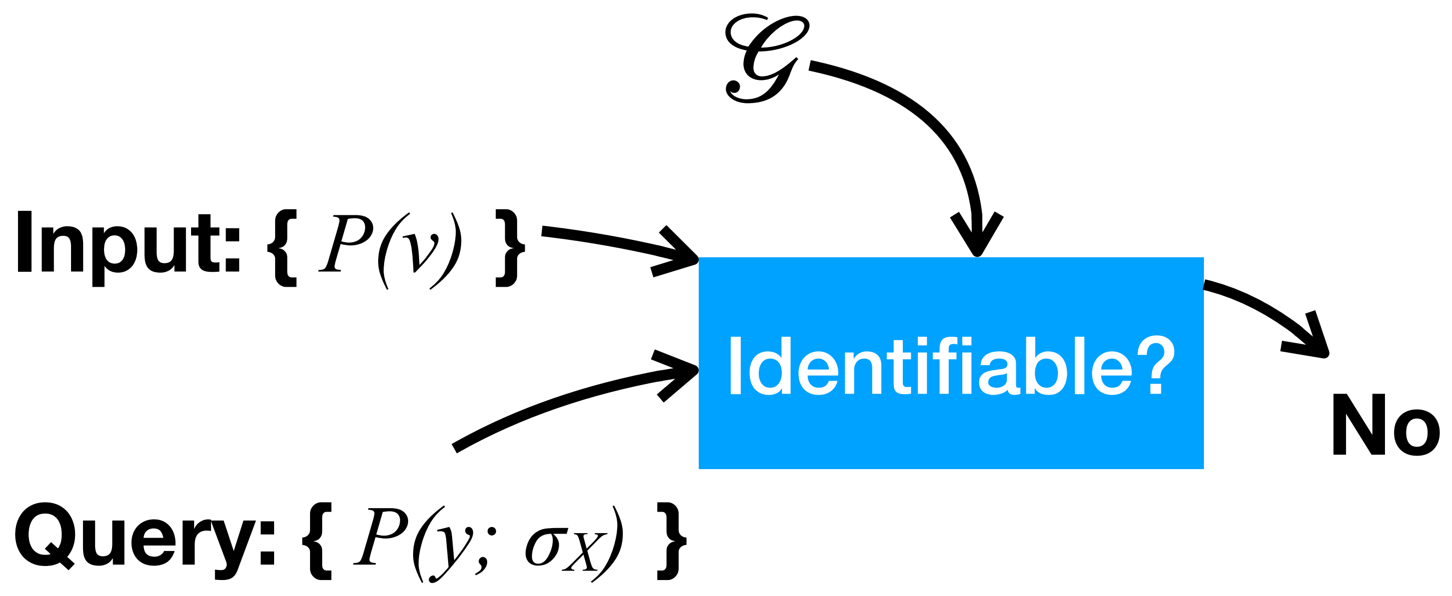

Surrogate Experiments

- It's not uncommon that the effect of a certain intervention is not identifiable (not uniquely computable) from observational data alone whenever unobserved confounders are present.

- Experiments over a set of surrogate variables $\boldsymbol{Z}$ may be more accessible to manipulation than the target effect σ_X, e.g., randomizing diet vs randomizing cholesterol.

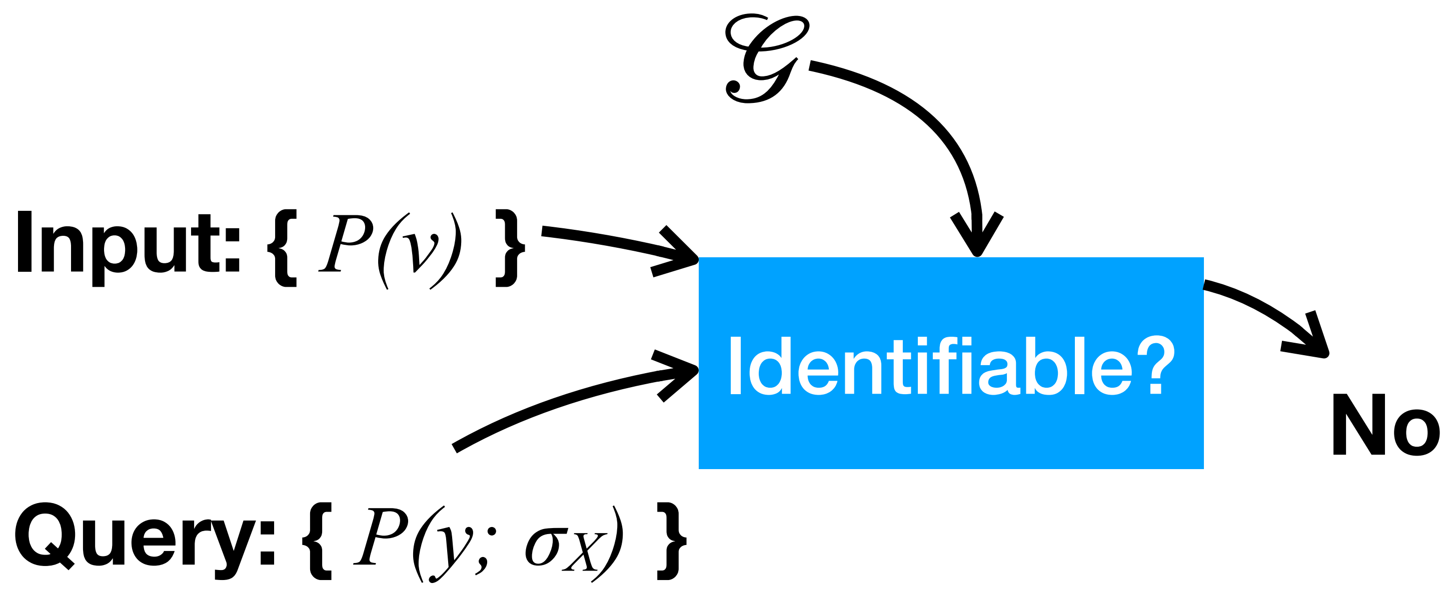

Surrogate Experiments

- It's not uncommon that the effect of a certain intervention is not identifiable (not uniquely computable) from observational data alone whenever unobserved confounders are present.

- Experiments over a set of surrogate variables $\boldsymbol{Z}$ may be more accessible to manipulation than the target effect σ_X, e.g., randomizing diet vs randomizing cholesterol.

- Those surrogate experiments can be leveraged to identify the effect of the interventions of interest.

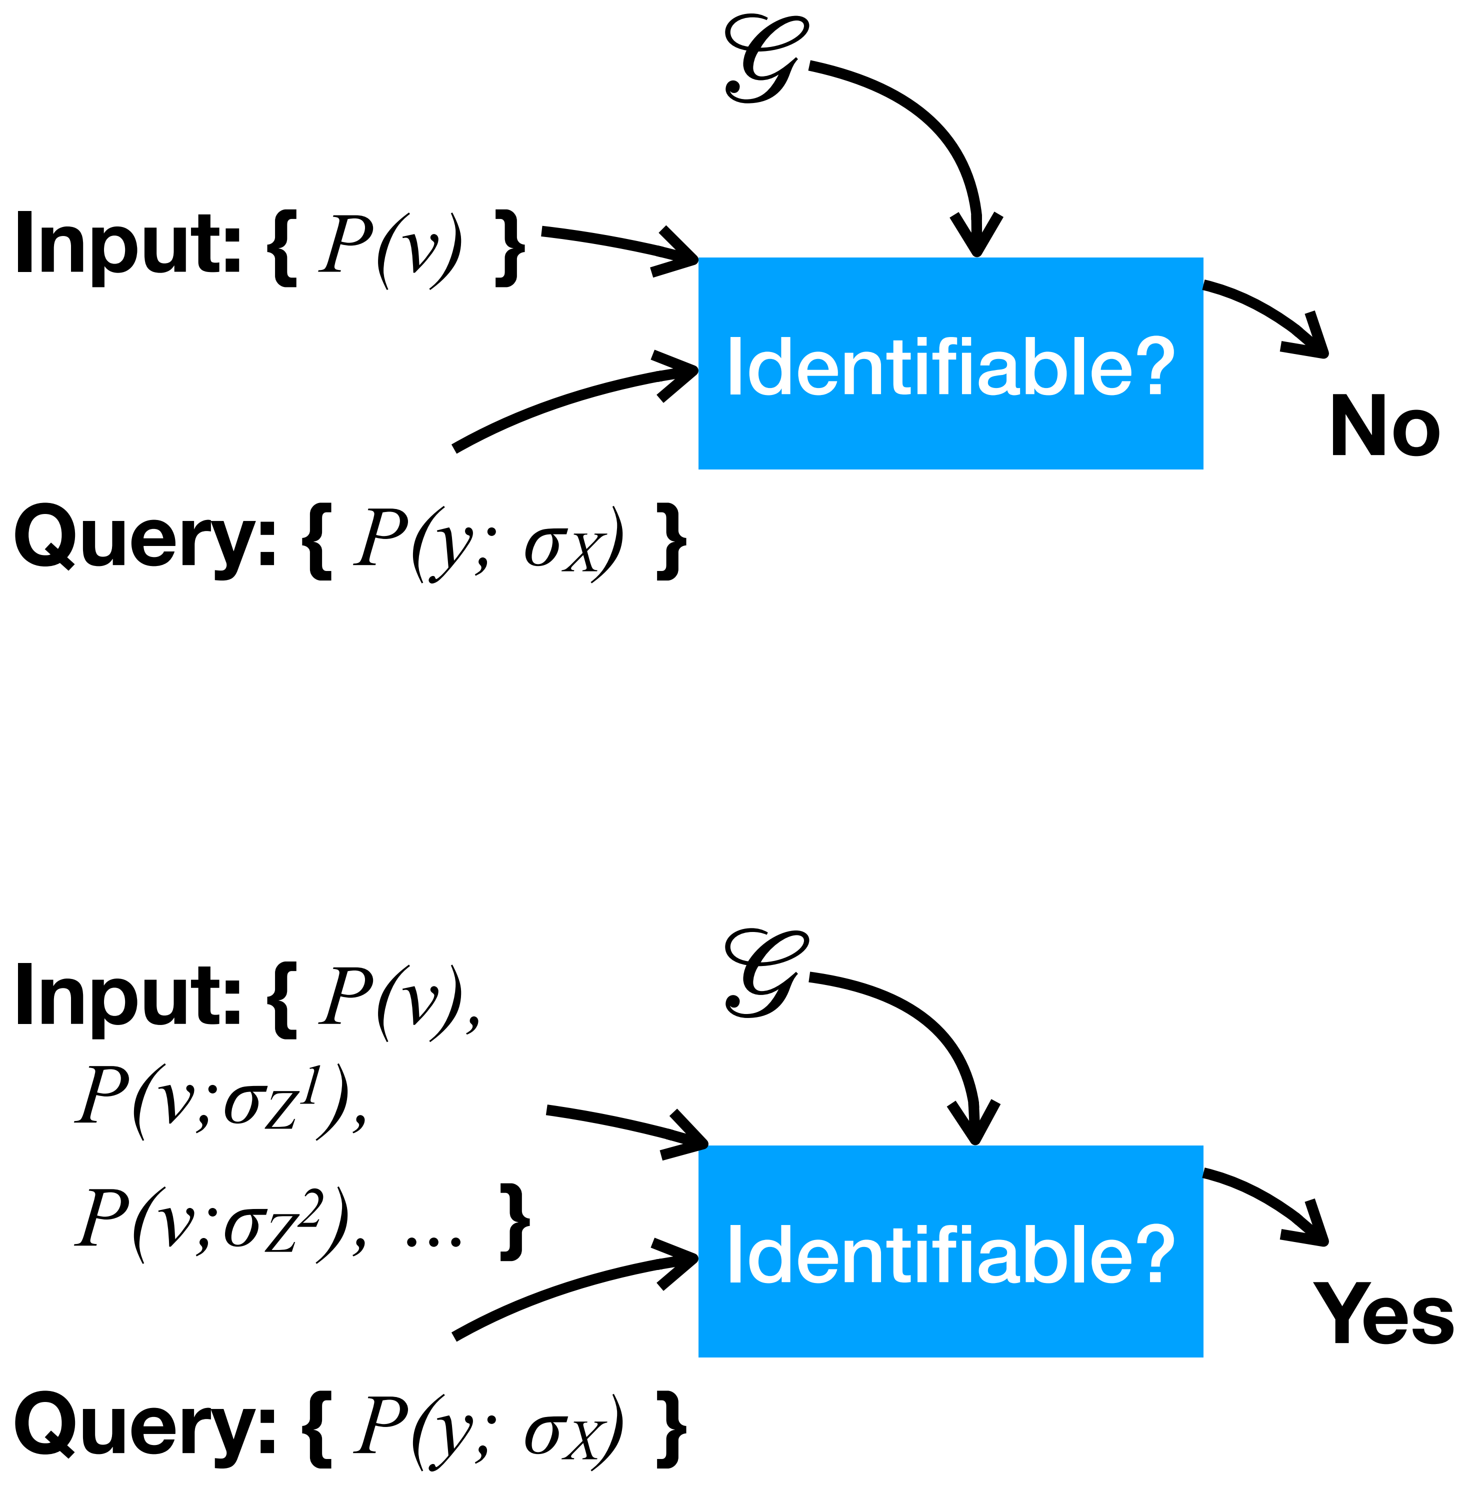

Surrogate Experiments

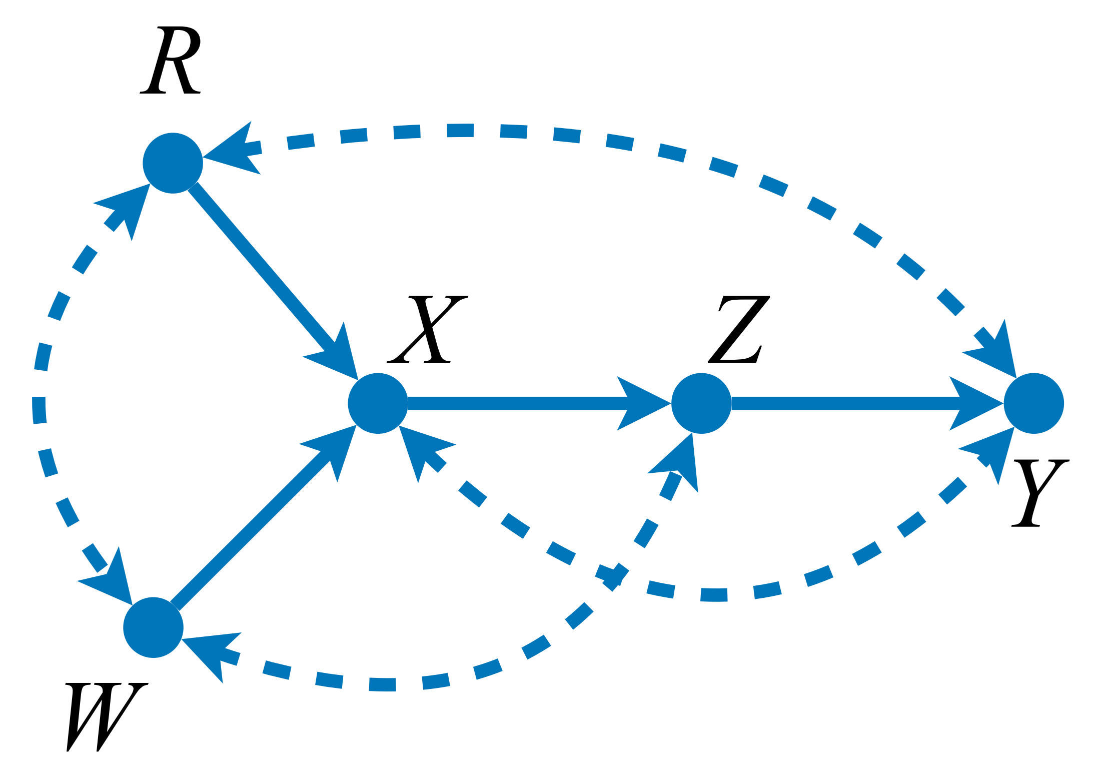

Surrogate Experiments

- **Input:** $\{P(\mathbf{v}), P(\mathbf{v} \mid \sigma_Z{=}P^*(Z|X))\}$

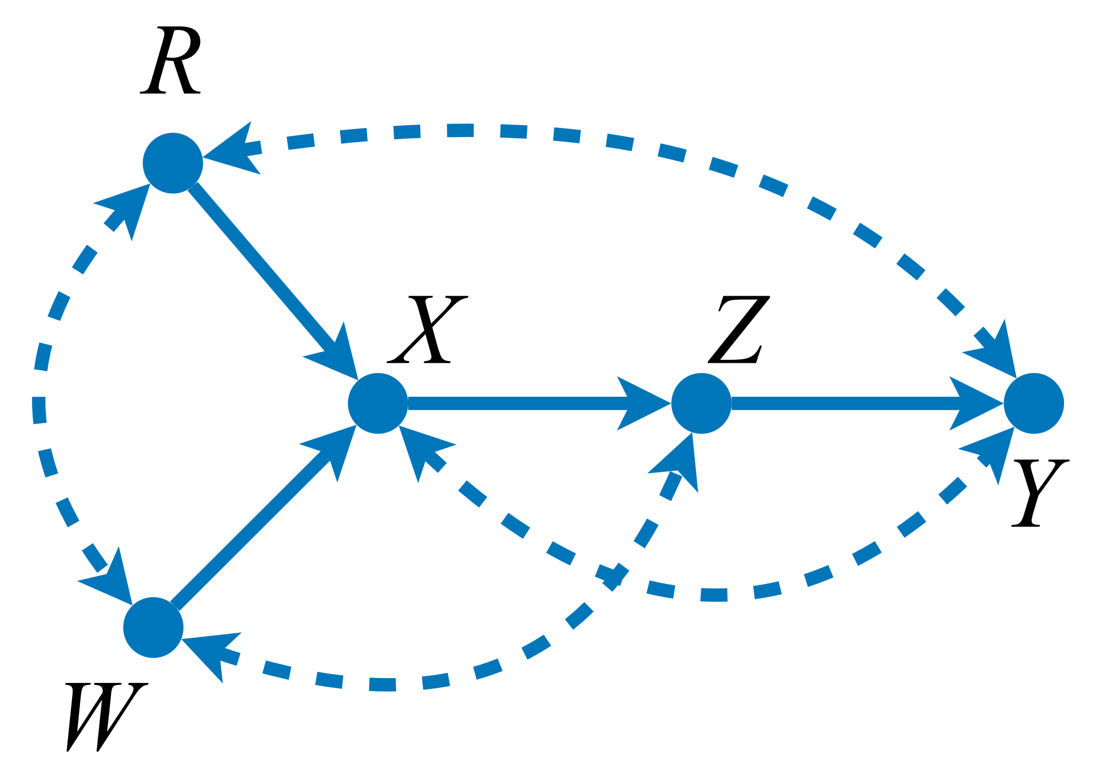

Surrogate Experiments

- **Input:** $\{P(\mathbf{v}), P(\mathbf{v} \mid \sigma_Z = P^*(Z|X))\}$
- **Query:** $P(y \mid \sigma_X = P^*(X|R))$

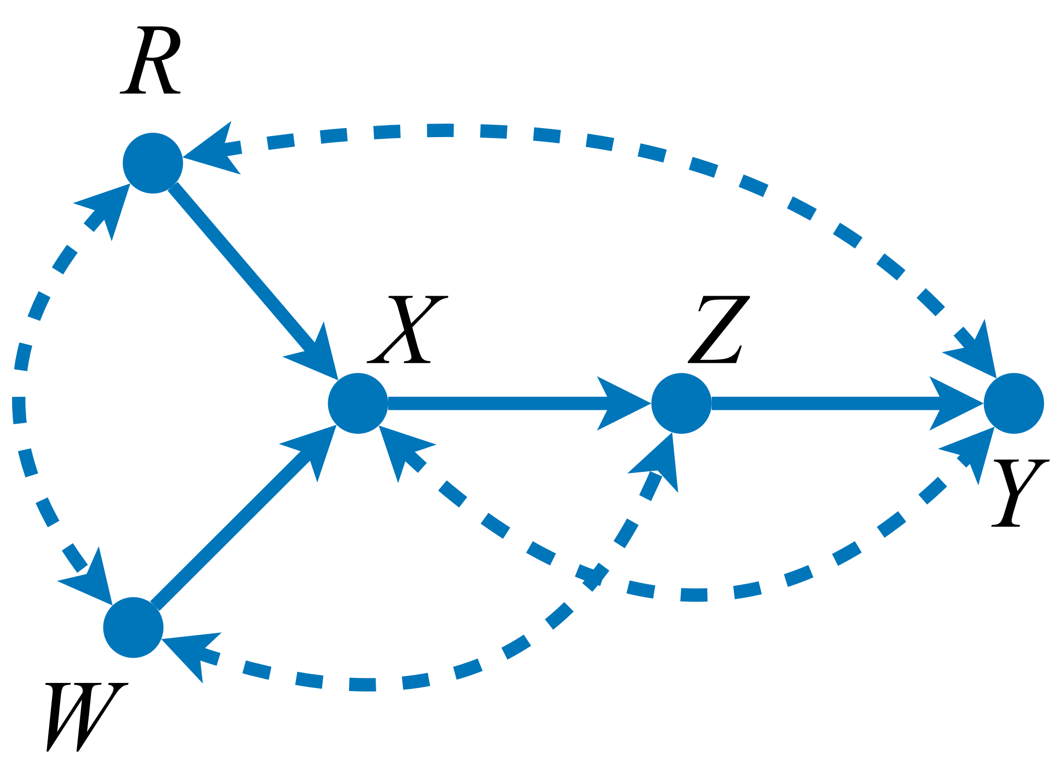

Surrogate Experiments

- **Input:** $\{P(\mathbf{v}), P(\mathbf{v} \mid \sigma_Z = P^*(Z|X))\}$
- **Query:** $P(y \mid \sigma_X = P^*(X|R))$

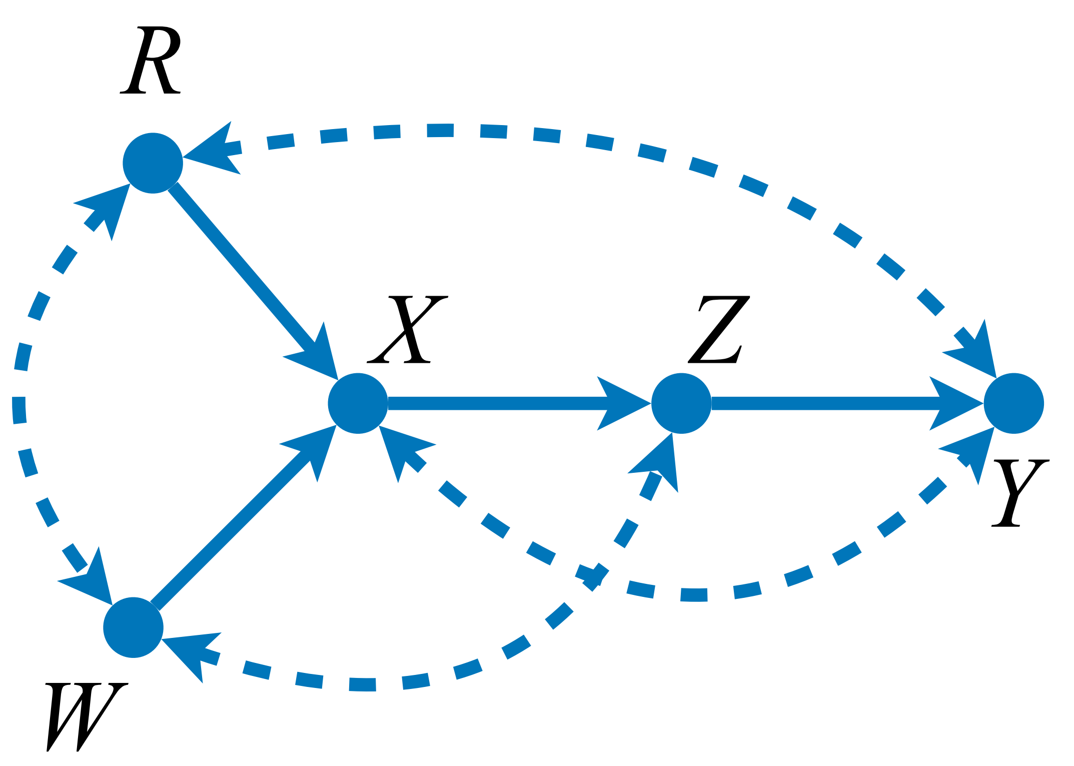

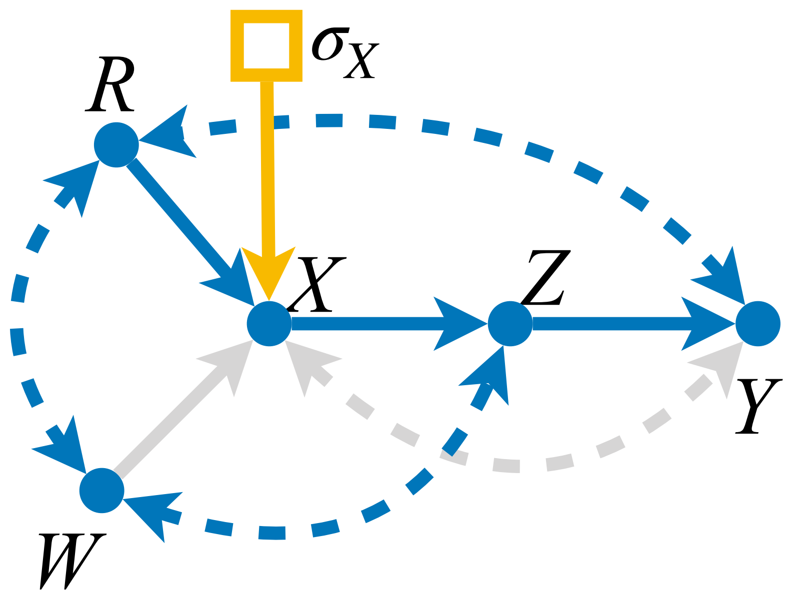

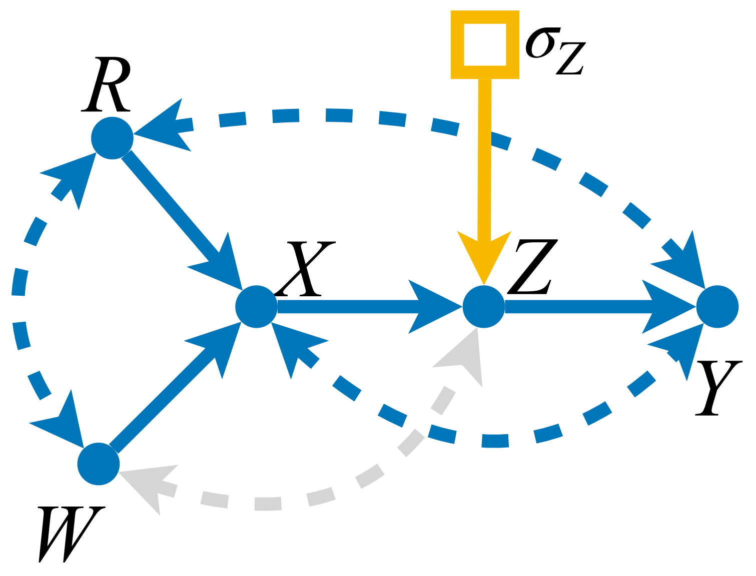

Surrogate Experiments

- **Input**: $\{P(\mathbf{v}), P(\mathbf{v} \mid \sigma_Z = P^*(Z|X))\}$
- **Query**: $P(y \mid \sigma_X = P^*(X|R))$
- Not identifiable from $P(\mathbf{v})$ alone, but

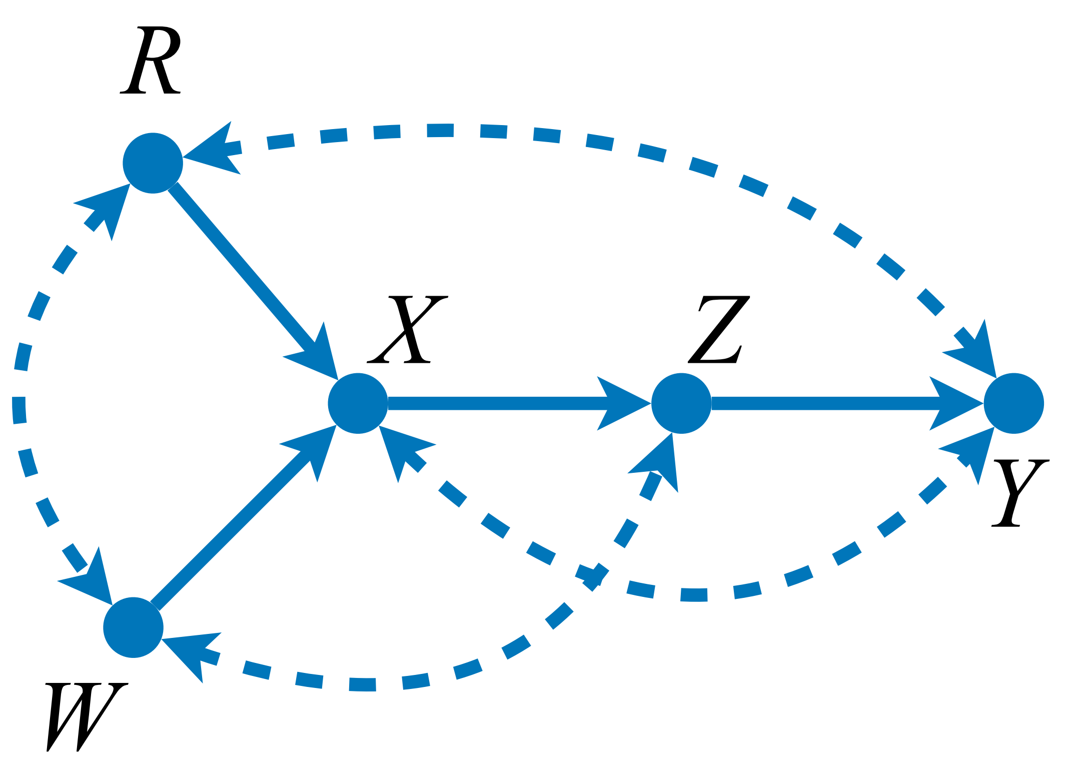

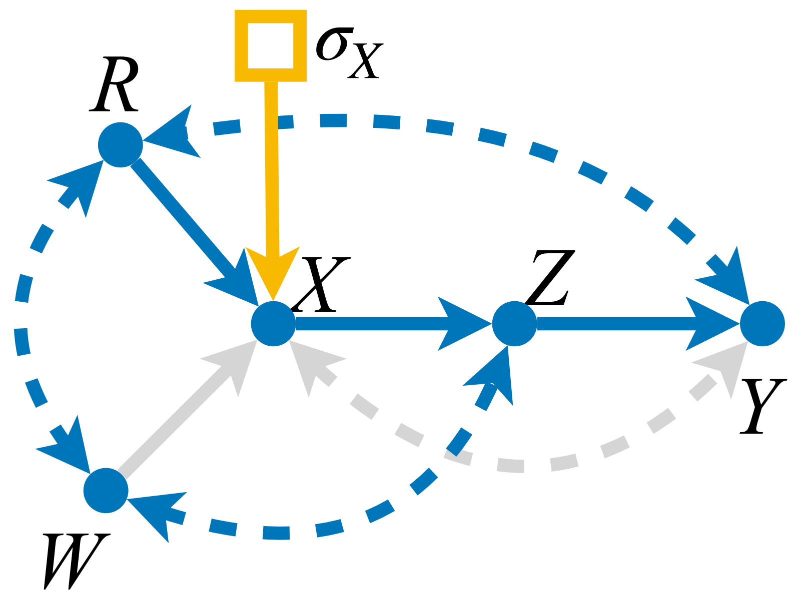

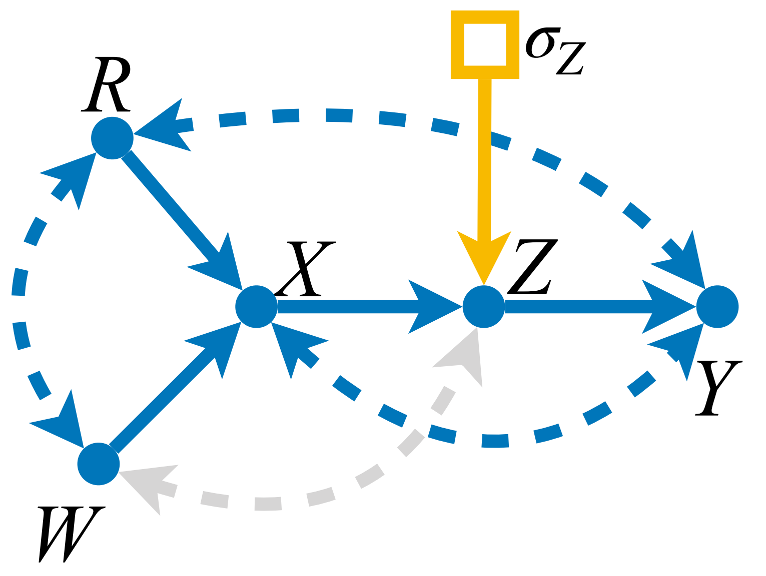

Surrogate Experiments

- **Input**: $\{P(\mathbf{v}), P(\mathbf{v} \mid \sigma_Z{=}P^*(Z|X))\}$
- **Query**: $P(y \mid \sigma_X{=}P^*(X|R))$
- Not identifiable from $P(\mathbf{v})$ alone, but

$P(y; \sigma_X)$

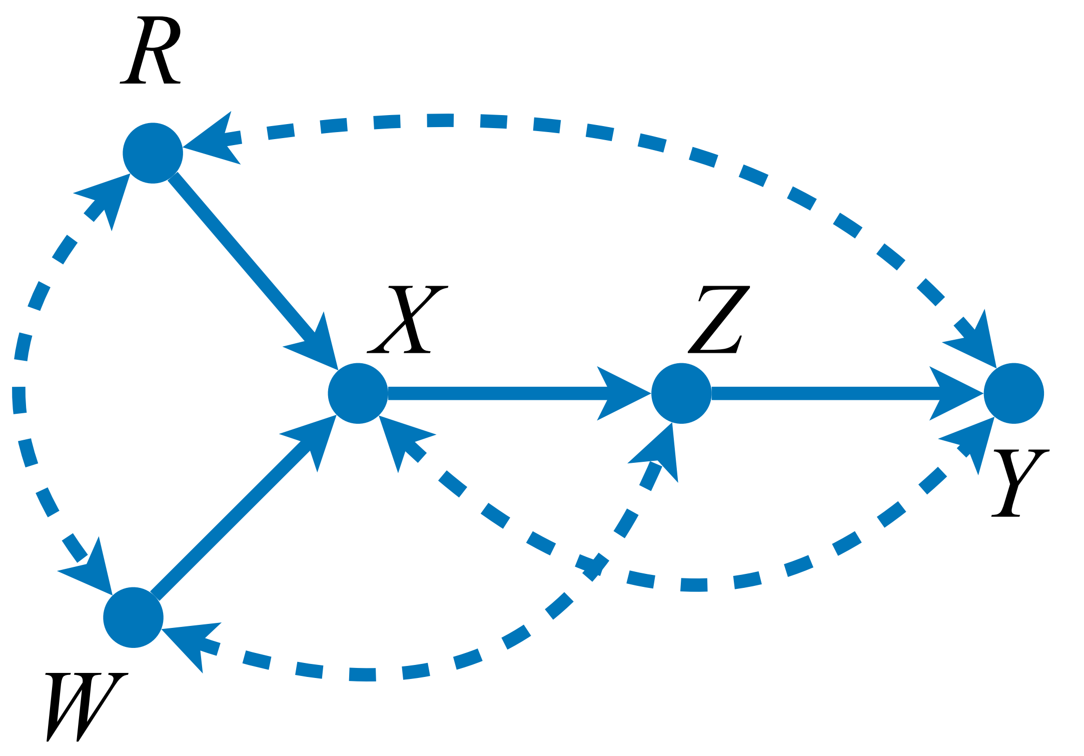

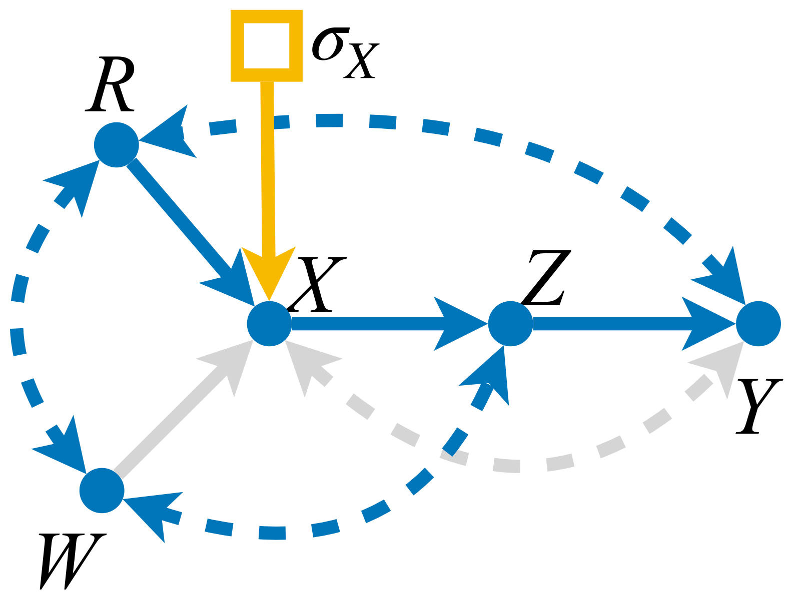

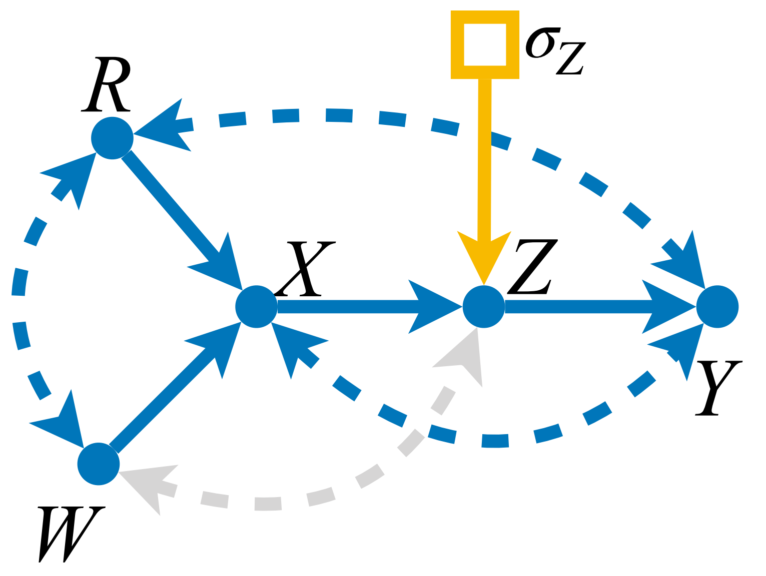

Surrogate Experiments

- **Input:** $\{P(\mathbf{v}), P(\mathbf{v} \mid \sigma_Z{=}P^*(Z|X))\}$
- **Query:** $P(y \mid \sigma_X{=}P^*(X|R))$
- Not identifiable from $P(\mathbf{v})$ alone, but

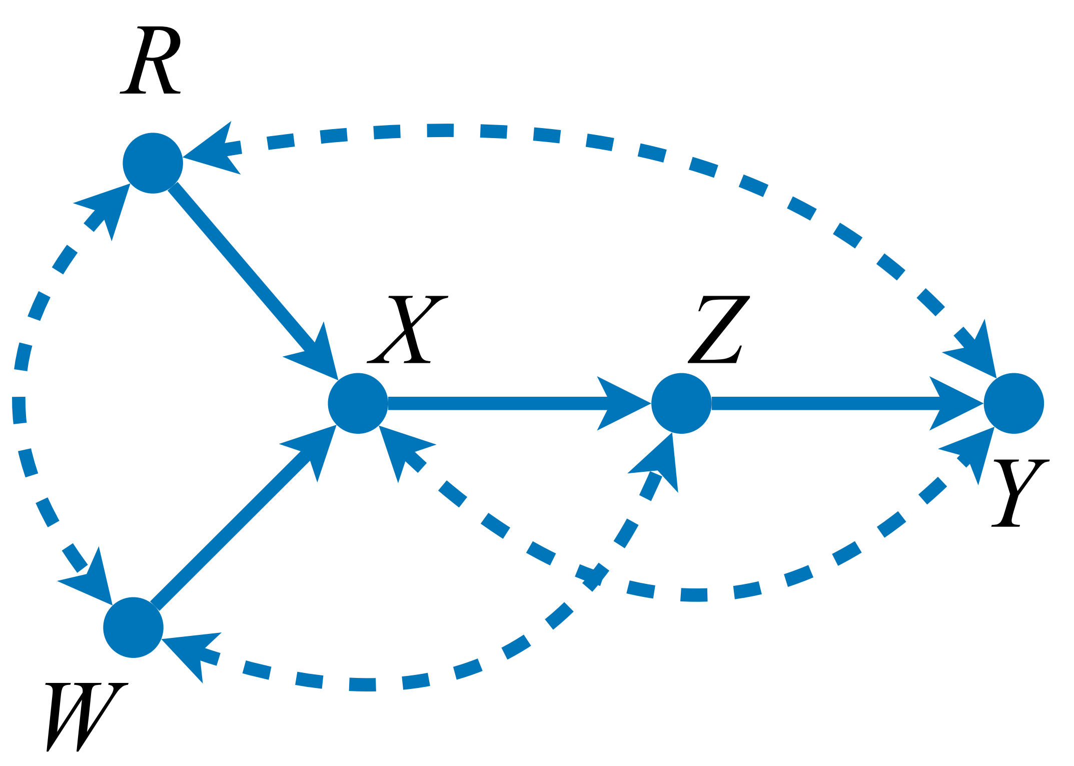

$$P(y; \sigma_X) = \sum_{r,w,x,z} P(r)P(x \mid r; \sigma_X)P(z \mid r, x, w)P(w \mid r) \sum_{x'} P(y \mid r, x', z; \sigma_Z)P(x' \mid r)$$

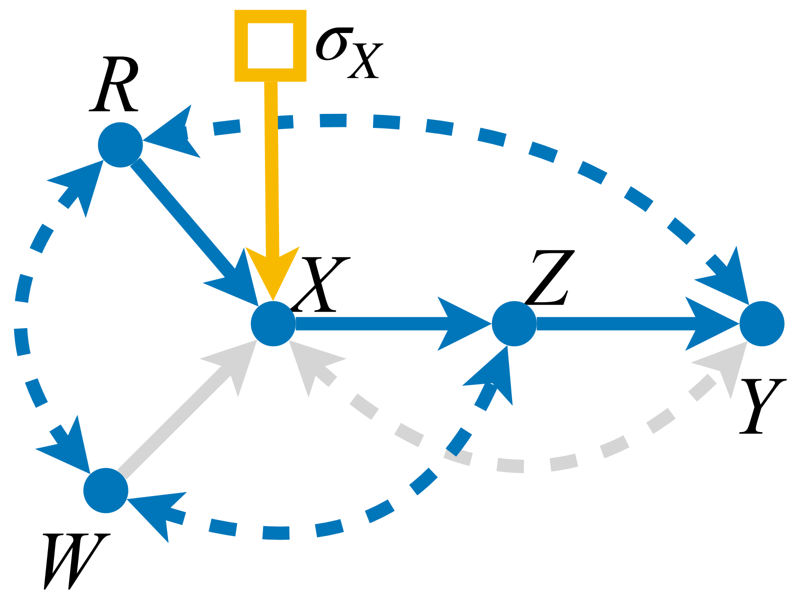

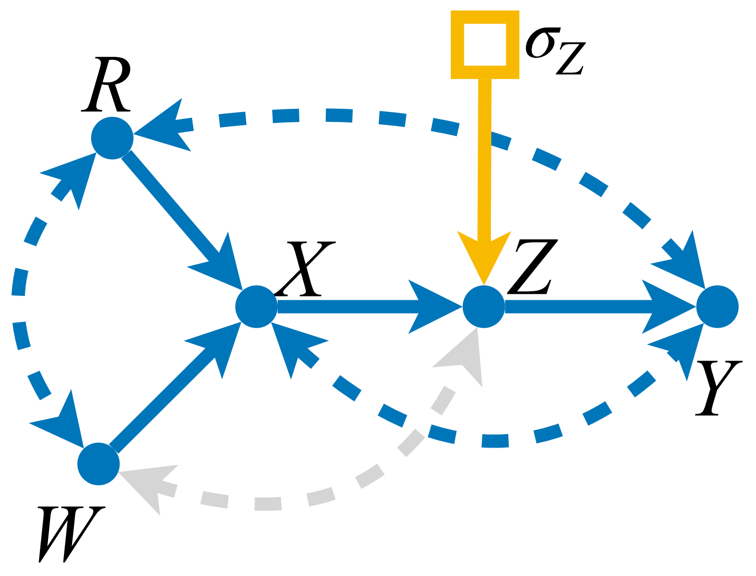

Surrogate Experiments

- **Input**: $\{P(\mathbf{v}), P(\mathbf{v} \mid \sigma_Z{=}P^*(Z|X))\}$
- **Query**: $P(y \mid \sigma_X{=}P^*(X|R))$
- Not identifiable from $P(\mathbf{v})$ alone, but

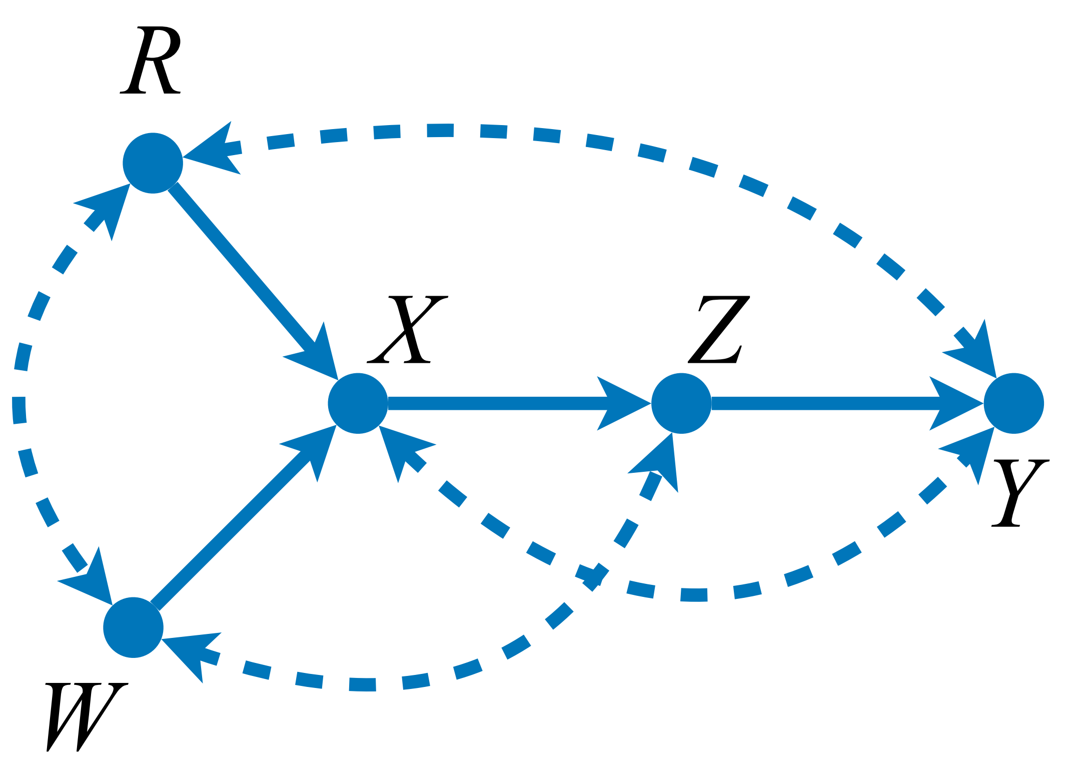

$$P(y; \sigma_X)$$
$$= \sum_{r,w,x,z} P(r)P(x \mid r; \sigma_X)P(z \mid r, x, w)P(w \mid r) \sum_{x'} P(y \mid r, x', z; \sigma_Z)P(x' \mid r)$$

From surrogate experiment

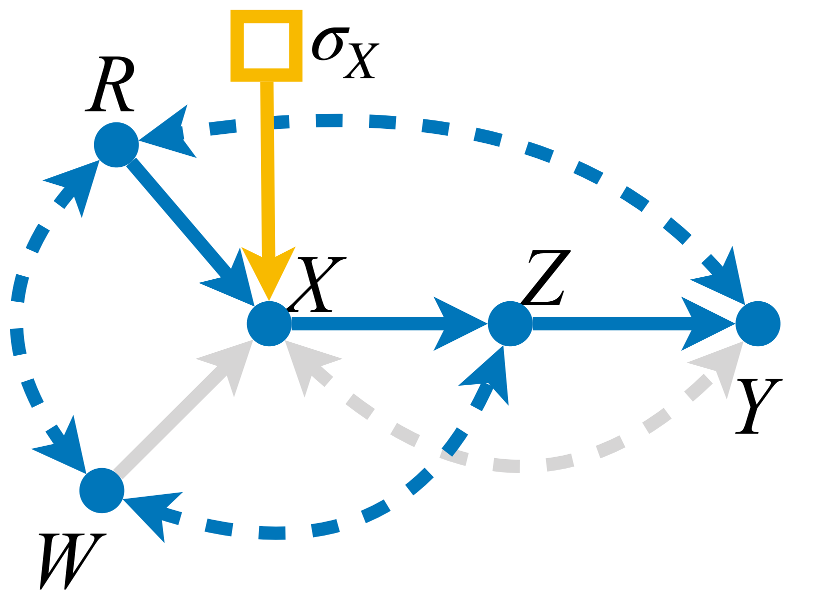

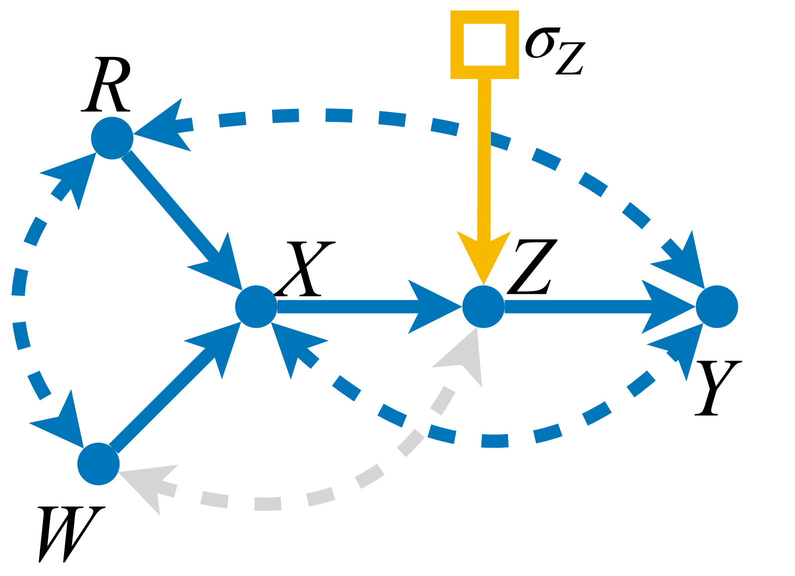

Surrogate Experiments

- **Input:** $\{P(\mathbf{v}), P(\mathbf{v} \mid \sigma_Z{=}P^*(Z|X))\}$
- **Query:** $P(y \mid \sigma_X{=}P^*(X|R))$
- Not identifiable from $P(\mathbf{v})$ alone, but

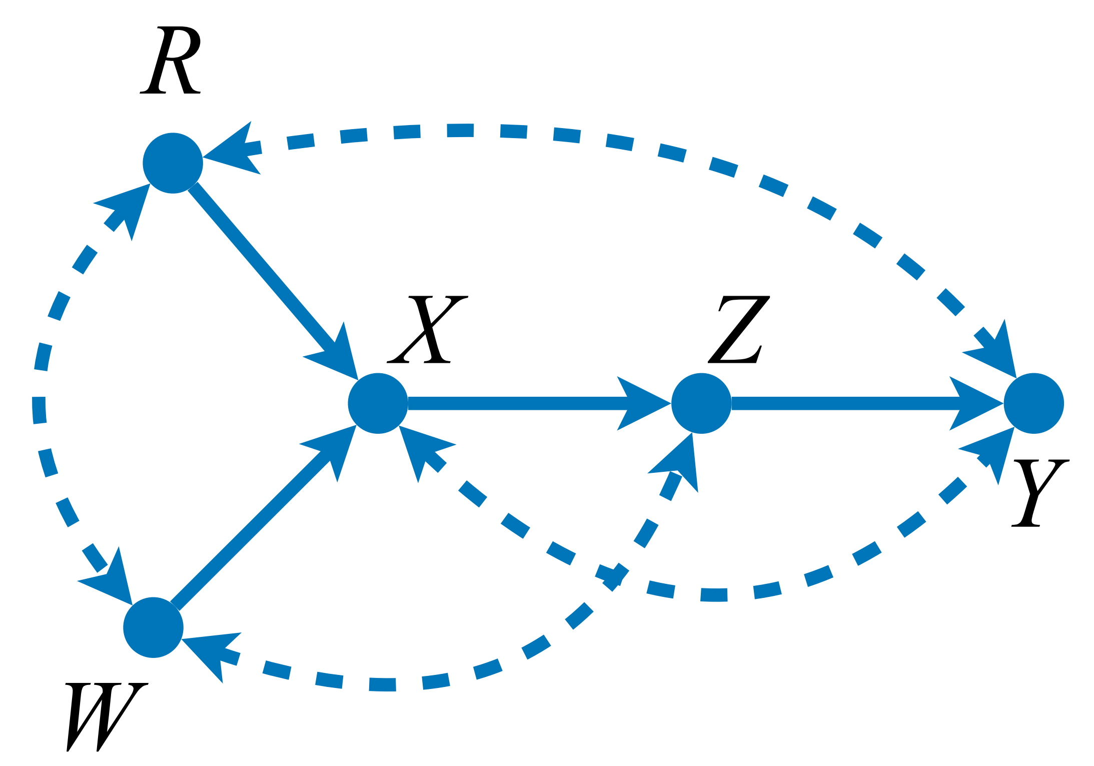

$$P(y;\sigma_X)$$

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

Natural regime

From surrogate experiment

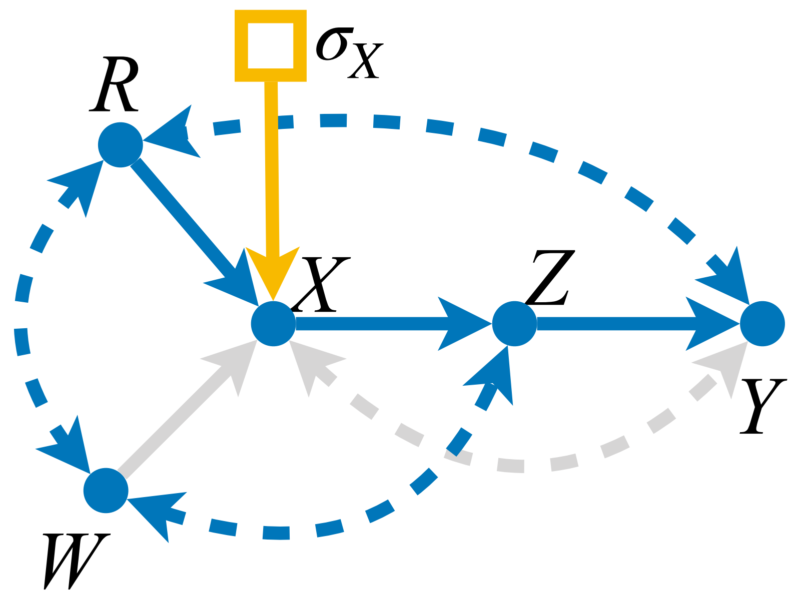

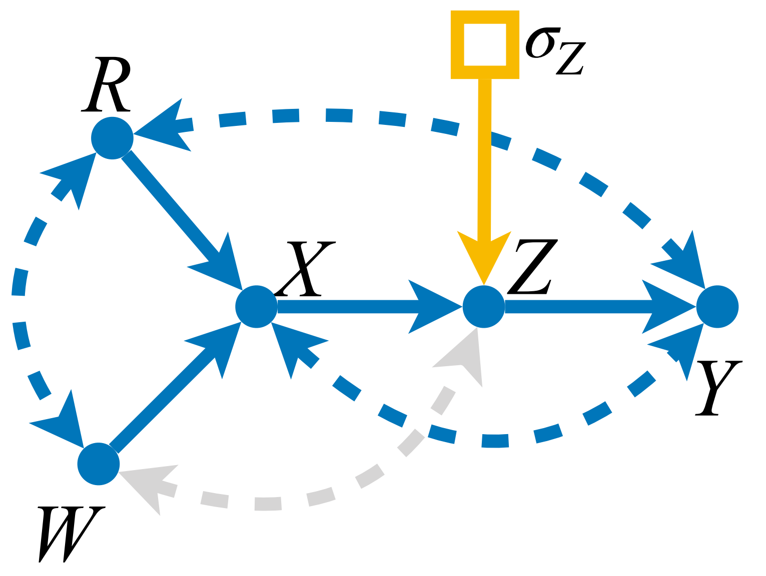

Surrogate Experiments

- **Input:** $\{P(\mathbf{v}), P(\mathbf{v} \mid \sigma_Z{=}P^*(Z|X))\}$
- **Query:** $P(y \mid \sigma_X{=}P^*(X|R))$
- Not identifiable from $P(\mathbf{v})$ alone, but

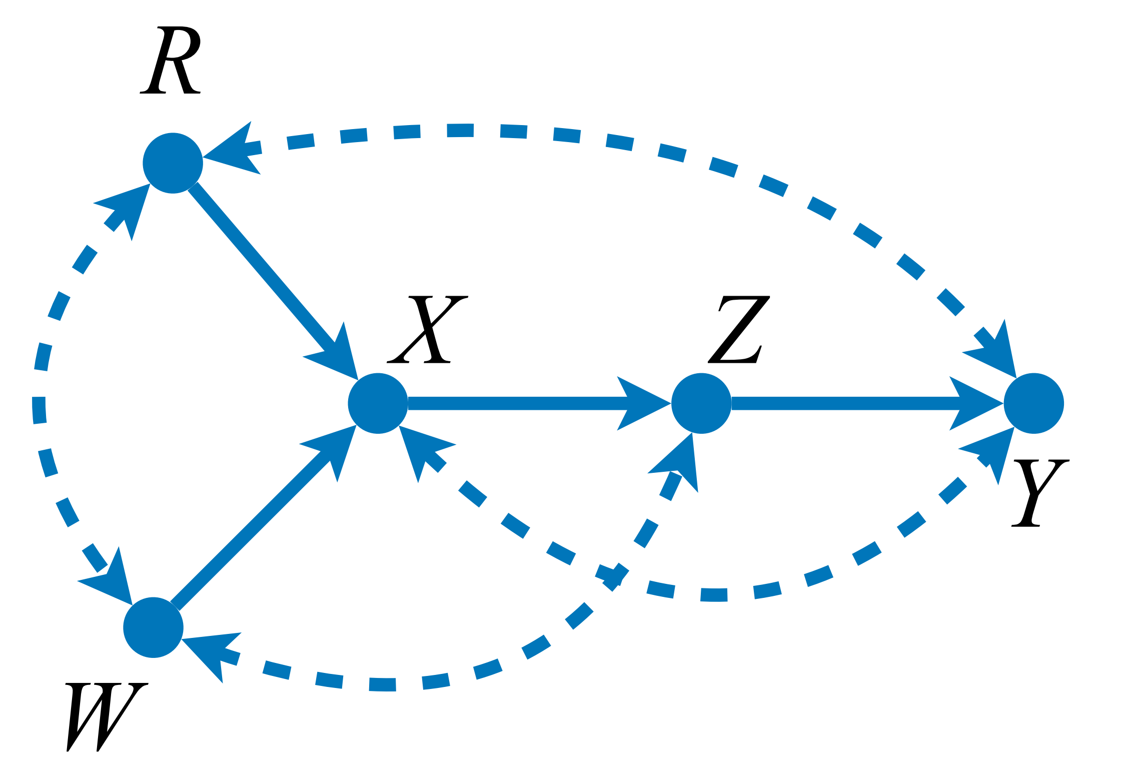

$$P(y; \sigma_X) = \sum_{r,w,x,z} P(r)P(x \mid r; \sigma_X)P(z \mid r, x, w)P(w \mid r) \sum_{x'} P(y \mid r, x', z; \sigma_Z)P(x' \mid r)$$

Defined by intervention

Natural regime

From surrogate experiment

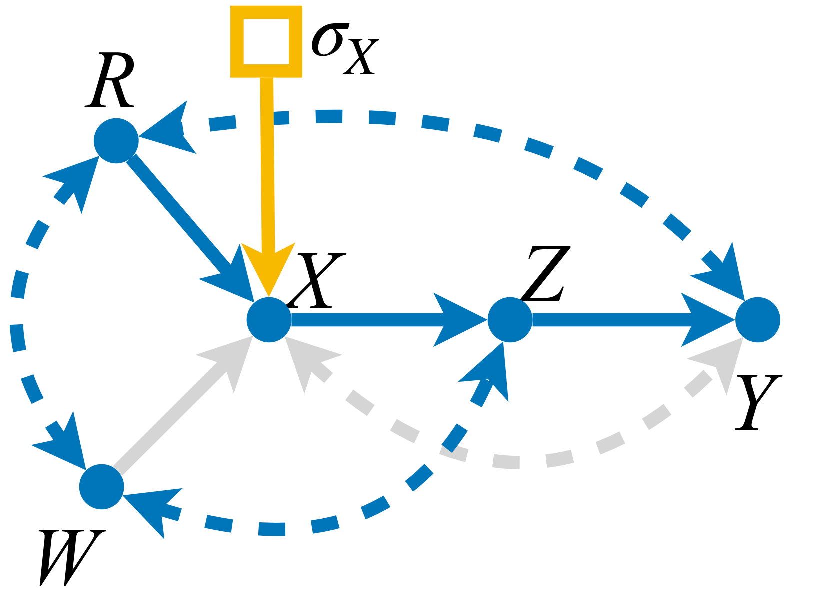

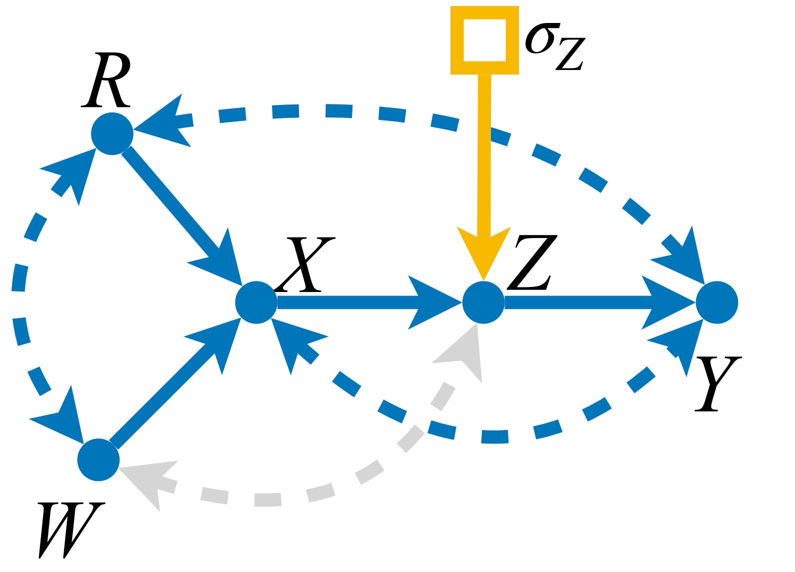

Summary of the Results

Summary of the Results

1. 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. Further, we provide a syntactical method for deriving and verifying claims about such interventions given a causal graph.

Summary of the Results

1. 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. Further, we provide a syntactical method for deriving and verifying claims about such interventions given a causal graph.

2. We develop an efficient procedure to determine the identifiability of the (conditional) effect of non-atomic interventions from a combination of observational and experimental data given a causal diagram.

Proposed Strategy

Proposed Strategy

1. Encode qualitative assumptions natural and intervened domain graphically.

Proposed Strategy

1. Encode qualitative assumptions natural and intervened domain graphically. → Diagrams annotated with □ σ_X nodes.

Proposed Strategy

1. Encode qualitative assumptions natural and intervened domain graphically. → Diagrams annotated with □ σ_X nodes.

2. Find the mechanisms composing the effect of intervention.

Proposed Strategy

1. Encode qualitative assumptions natural and intervened domain graphically. → Diagrams annotated with □ σ_X nodes.

2. Find the mechanisms composing the effect of intervention.

3. Derive the needed mechanisms from the given distributions.

Proposed Strategy

1. Encode qualitative assumptions natural and intervened domain graphically. → Diagrams annotated with □ σ_X nodes.

2. Find the mechanisms composing the effect of intervention.

3. Derive the needed mechanisms from the given distributions.

4. Construct an estimator from the available data.

Proposed Strategy

1. Encode qualitative assumptions natural and intervened domain graphically. → Diagrams annotated with □ σ_X nodes.

2. Find the mechanisms composing the effect of intervention.

3. Derive the needed mechanisms from the given distributions.

4. Construct an estimator from the available data.

(Steps 2–4:) Use σ-calculus or equivalent algorithmic procedure.

Conclusions

Conclusions

- σ-calculus allows one to discover and verify, from a causal graph, logical statements about general interventions suitable to capture real-world situations.

Conclusions

- σ-calculus allows one to discover and verify, from a causal graph, logical statements about general interventions suitable to capture real-world situations.
- These rules can be used to identify the effect of interventions from a combination of observational and experimental data.

Conclusions

- σ-calculus allows one to discover and verify, from a causal graph, logical statements about general interventions suitable to capture real-world situations.
- These rules can be used to identify the effect of interventions from a combination of observational and experimental data.
- Our algorithm searches for a reduction of the effect of interest to the set of observed distributions (observational and experimental); if found, it returns a corresponding mapping expression.

Thank you!