Tutorial on Structural Causal Model

Yonghan Jung

Purdue University

yonghanjung.me

2022.07.11

University of Seoul 2022

Overview of Lecture Series

Outline for Lecture Series

This lecture series composes of the following topics:

1. Tutorial on Structural Causal Model (SCM)
2. Causal Effect Estimation on Any Identifiable Causal Functional.
3. Application to Interpretable Machine Learning

Introduction and Motivation

Practical Causal Query is Expressible as “What-If”

Many practical queries on causality are encoded as a “**What-If**” question.

Practical Causal Query is Expressible as "What-If"

Many practical queries on causality are encoded as a "**What-If**" question.

- Example 1. (Randomized Controlled Trials): **What** would have been Alice's headache **if** she had taken an aspirin?

Practical Causal Query is Expressible as “What-If”

Many practical queries on causality are encoded as a “**What-If**” question.

- Example 1. (Randomized Controlled Trials): **What** would have been Alice’s headache **if** she had taken an aspirin?
- Example 2. (A/B Test) Among two designs {A,B} for an online ad, **what** would have been the ad’s click rate **if** the design A has been chosen?

Example 1: Causality $\neq$ Correlation

Example 1: Causality $\neq$ Correlation

Number of Patients

Category	Treated	Non-Treated
[redacted]		
Combined	350	350

Example 1: Causality ≠ Correlation

Number of Patients		
Category	**Treated**	**Non-Treated**
[redacted]		
Combined	350	350

Number of Recovers		
Category	**Treated**	**Non-Treated**
[redacted]		
Combined	273	290

Example 1: Causality ≠ Correlation

Number of Patients		
Category	**Treated**	**Non-Treated**
[redacted]		
Combined	350	350

Number of Recovers		
Category	**Treated**	**Non-Treated**
[redacted]		
Combined	273	290

Survival Rate (%)		
Category	**Treated**	**Non-Treated**
[redacted]		
Combined	78	83

Example 1: Causality $\neq$ Correlation

Number of Patients			Number of Recovers			Survival Rate (%)		
Category	**Treated**	**Non-Treated**	**Category**	**Treated**	**Non-Treated**	**Category**	**Treated**	**Non-Treated**
[redacted]			[redacted]			[redacted]		
Combined	350	350	**Combined**	273	290	**Combined**	78	83

P(Survival | Non-Treated) > P(Survival | Treated).

Example 1: Causality $\neq$ Correlation

Number of Patients		
Category	**Treated**	**Non-Treated**
[redacted]		
Combined	350	350

Number of Recovers		
Category	**Treated**	**Non-Treated**
[redacted]		
Combined	273	290

Survival Rate (%)		
Category	**Treated**	**Non-Treated**
[redacted]		
Combined	78	83

P(Survival | Non-Treated) > P(Survival | Treated).

Does this mean that the drug is harmful?

Example 1: Causality $\neq$ Correlation

Example 1: Causality $\neq$ Correlation

Number of Patients

Category	Treated	Non-Treated
Male	87	270
Female	263	80
Combined	350	350

Example 1: Causality $\neq$ Correlation

Number of Patients

Category	Treated	Non-Treated
Male	87	270
Female	263	80
Combined	350	350

Number of Recovers

Category	Treated	Non-Treated
Male	81	235
Female	192	55
Combined	273	290

Example 1: Causality $\neq$ Correlation

Number of Patients		
Category	**Treated**	**Non-Treated**
Male	87	270
Female	263	80
Combined	350	350

Number of Recovers		
Category	**Treated**	**Non-Treated**
Male	81	235
Female	192	55
Combined	273	290

Survival Rate (%)		
Category	**Treated**	**Non-Treated**
Male	93	87
Female	73	69
Combined	78	83

Example 1: Causality $\neq$ Correlation

Number of Patients			Number of Recovers			Survival Rate (%)		
Category	**Treated**	**Non-Treated**	**Category**	**Treated**	**Non-Treated**	**Category**	**Treated**	**Non-Treated**
Male	87	270	**Male**	81	235	**Male**	93	87
Female	263	80	**Female**	192	55	**Female**	73	69
Combined	350	350	**Combined**	273	290	**Combined**	78	83

- P(Survival | Non-Treated, Male) < P(Survival | Treated, Male).
- P(Survival | Non-Treated, Female) < P(Survival | Treated, Female).

Example 1: Causality ≠ Correlation

Number of Patients			Number of Recovers			Survival Rate (%)		
Category	Treated	Non-Treated	Category	Treated	Non-Treated	Category	Treated	Non-Treated
Male	87	270	Male	81	235	Male	93	87
Female	263	80	Female	192	55	Female	73	69
Combined	350	350	Combined	273	290	Combined	78	83

- P(Survival | Non-Treated, Male) < P(Survival | Treated, Male).
- P(Survival | Non-Treated, Female) < P(Survival | Treated, Female).

Does this mean that the drug is beneficial?

Data Generating Process in Causal Inference

Survival Rate (%)		
Category	**Treated**	**Non-Treated**
Male	93	87
Female	73	69
Combined	78	83

Data Generating Process in Causal Inference

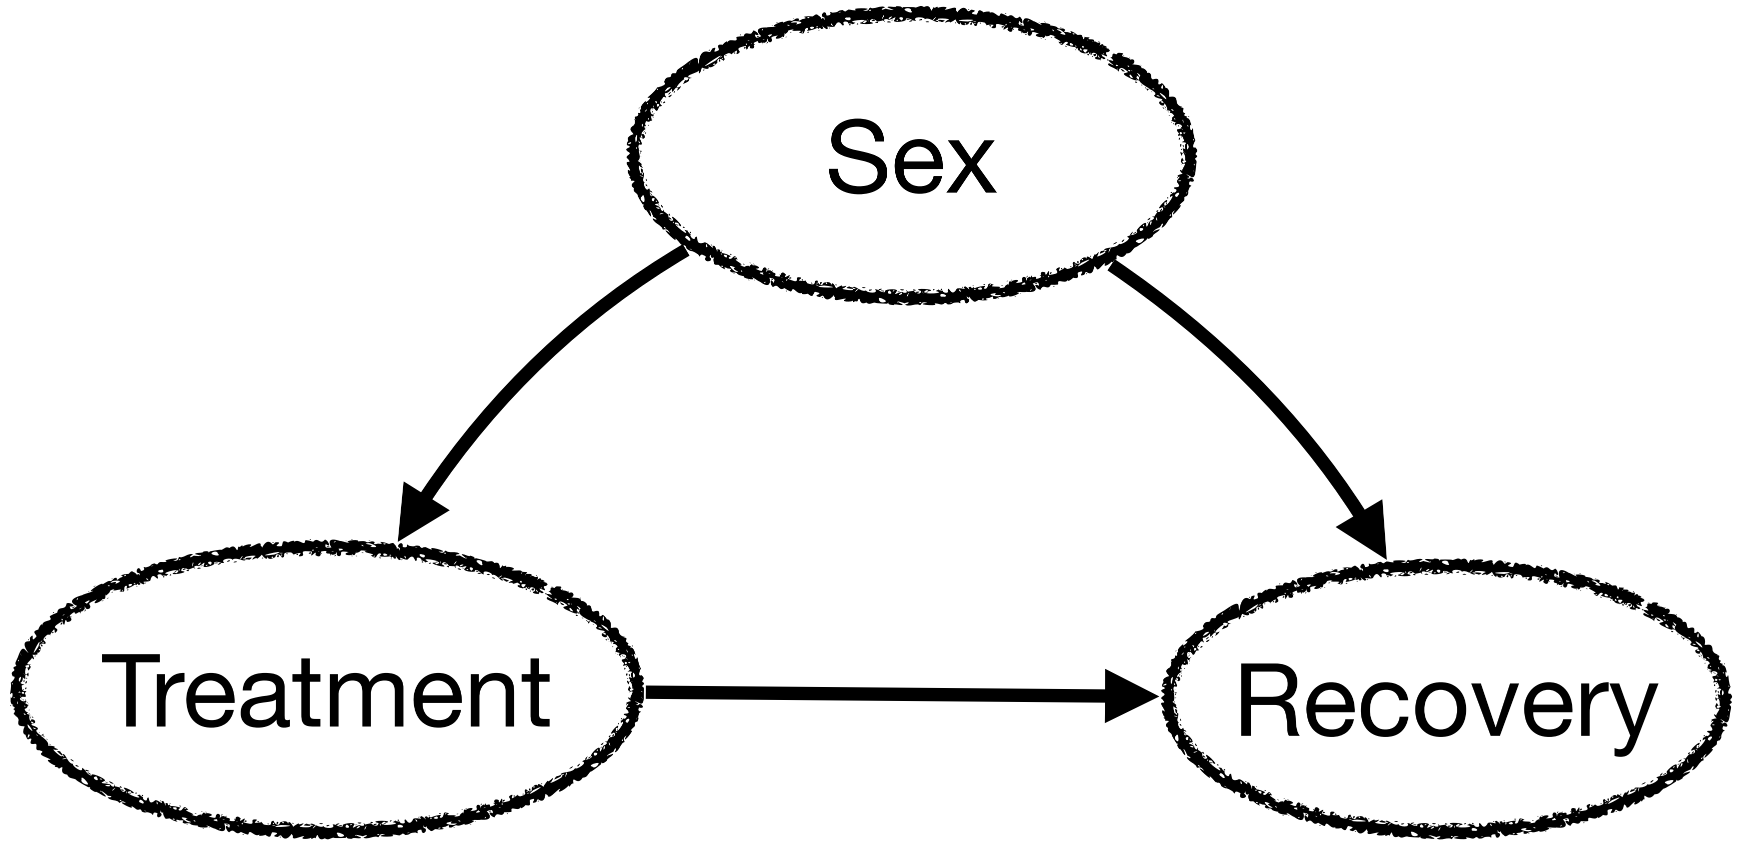

Survival Rate (%)		
Category	**Treated**	**Non-Treated**
Male	93	87
Female	73	69
Combined	78	83

Data Generating Process in Causal Inference

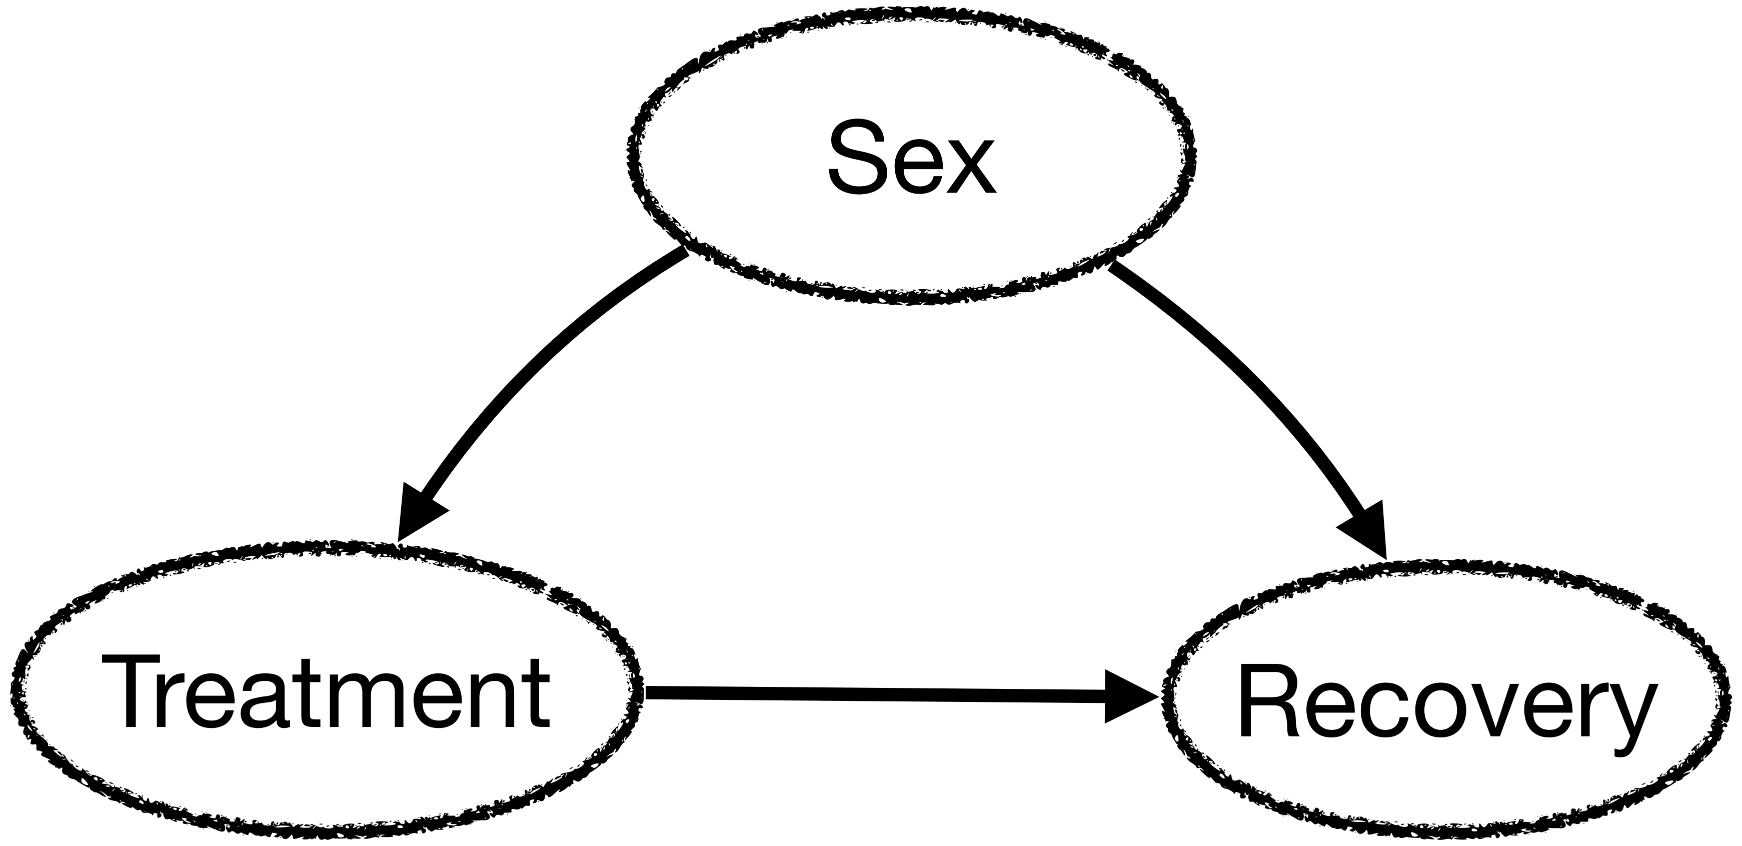

Survival Rate (%)		
Category	**Treated**	**Non-Treated**
Male	93	87
Female	73	69
Combined	78	83

If the data generating process is given as a causal diagram,

Data Generating Process in Causal Inference

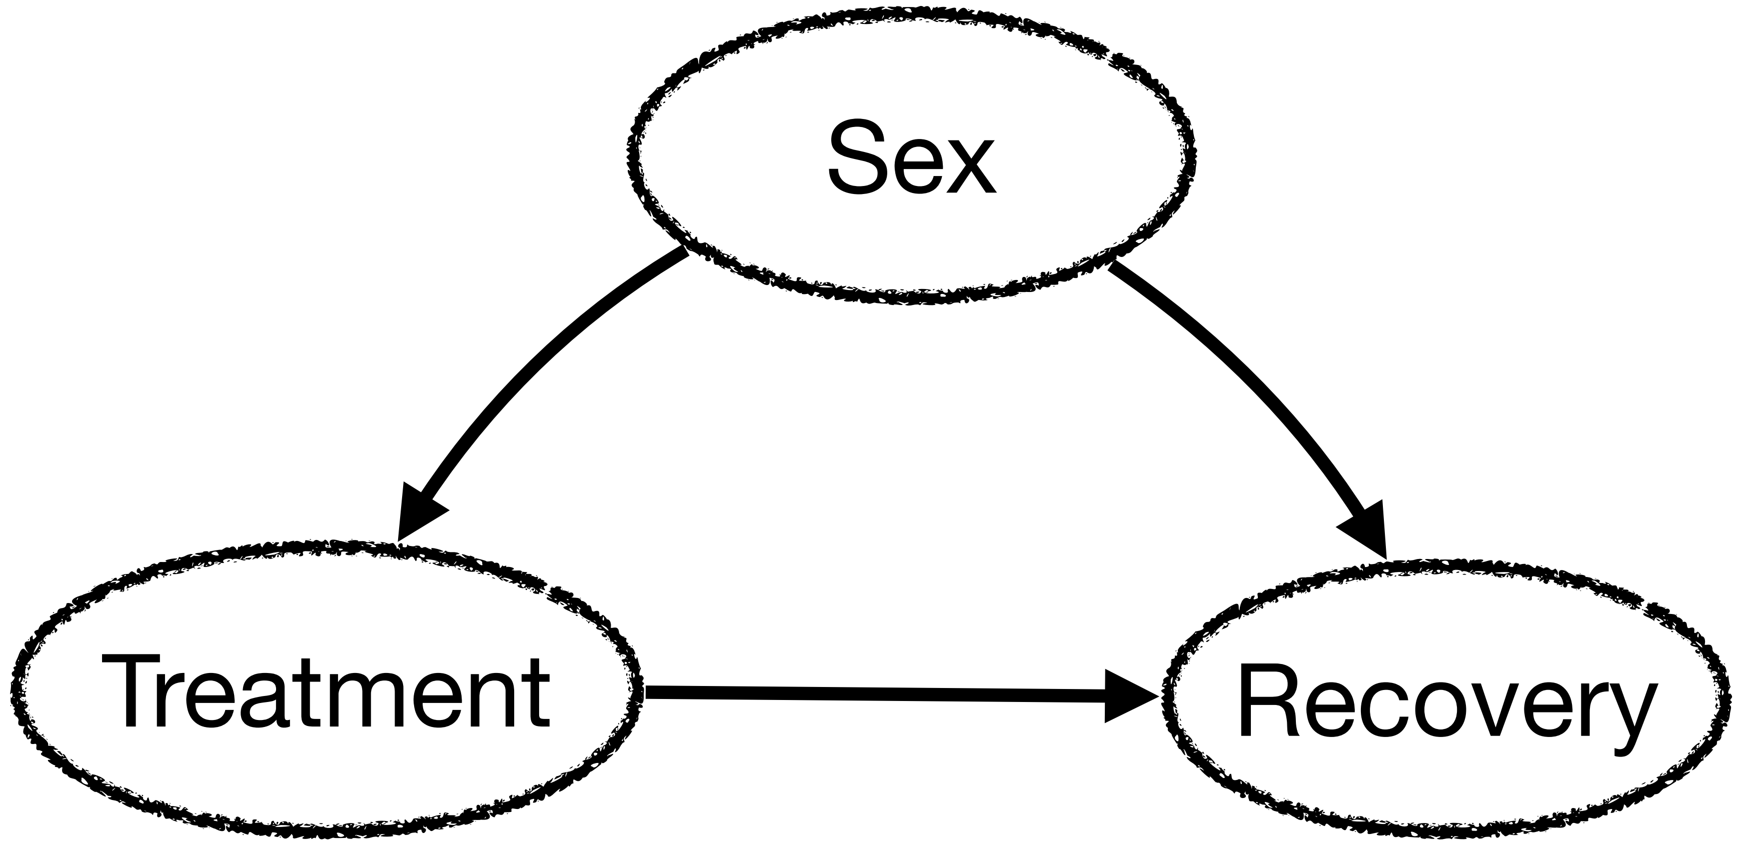

Survival Rate (%)		
Category	**Treated**	**Non-Treated**
Male	93	87
Female	73	69
Combined	78	83

If the data generating process is given as a causal diagram,

=> The treatment is beneficial.

Example 2: Causality ≠ Correlation

Number of Patients			Number of Recovers			Survival Rate (%)		
Category	**Treated**	**Non-Treated**	**Category**	**Treated**	**Non-Treated**	**Category**	**Treated**	**Non-Treated**
[redacted]			[redacted]			[redacted]		
Combined	350	350	**Combined**	273	290	**Combined**	78	83

P(Survival | Non-Treated) > P(Survival | Treated).

Does this mean that the drug is harmful?

Example 2: Causality $\neq$ Correlation

Example 2: Causality $\neq$ Correlation

Number of Patients

Category	**Treated**	**Non-Treated**
Low-BP	87	270
High-BP	263	80
Combined	350	350

Example 2: Causality $\neq$ Correlation

Number of Patients			Number of Recovers		
Category	**Treated**	**Non-Treated**	**Category**	**Treated**	**Non-Treated**
Low-BP	87	270	Low-BP	81	235
High-BP	263	80	High-BP	192	55
Combined	350	350	**Combined**	273	290

Example 2: Causality $\neq$ Correlation

Number of Patients		
Category	**Treated**	**Non-Treated**
Low-BP	87	270
High-BP	263	80
Combined	350	350

Number of Recovers		
Category	**Treated**	**Non-Treated**
Low-BP	81	235
High-BP	192	55
Combined	273	290

Survival Rate (%)		
Category	**Treated**	**Non-Treated**
Low-BP	93	87
High-BP	73	69
Combined	78	83

Example 2: Causality $\neq$ Correlation

Number of Patients		
Category	**Treated**	**Non-Treated**
Low-BP	87	270
High-BP	263	80
Combined	350	350

Number of Recovers		
Category	**Treated**	**Non-Treated**
Low-BP	81	235
High-BP	192	55
Combined	273	290

Survival Rate (%)		
Category	**Treated**	**Non-Treated**
Low-BP	93	87
High-BP	73	69
Combined	78	83

- P(Survival | Non-Treated, Low) < P(Survival | Treated, Low).
- P(Survival | Non-Treated, High) < P(Survival | Treated, Female).

Example 2: Causality ≠ Correlation

Number of Patients			Number of Recovers			Survival Rate (%)		
Category	**Treated**	**Non-Treated**	**Category**	**Treated**	**Non-Treated**	**Category**	**Treated**	**Non-Treated**
Low-BP	87	270	Low-BP	81	235	Low-BP	93	87
High-BP	263	80	High-BP	192	55	High-BP	73	69
Combined	350	350	**Combined**	273	290	**Combined**	78	83

- P(Survival | Non-Treated, Low) < P(Survival | Treated, Low).
- P(Survival | Non-Treated, High) < P(Survival | Treated, Female).

Does this mean that the drug is beneficial?

Data Generating Process in Causal Inference

Survival Rate (%)		
Category	**Treated**	**Non-Treated**
Male	93	87
Female	73	69
Combined	78	83

Data Generating Process in Causal Inference

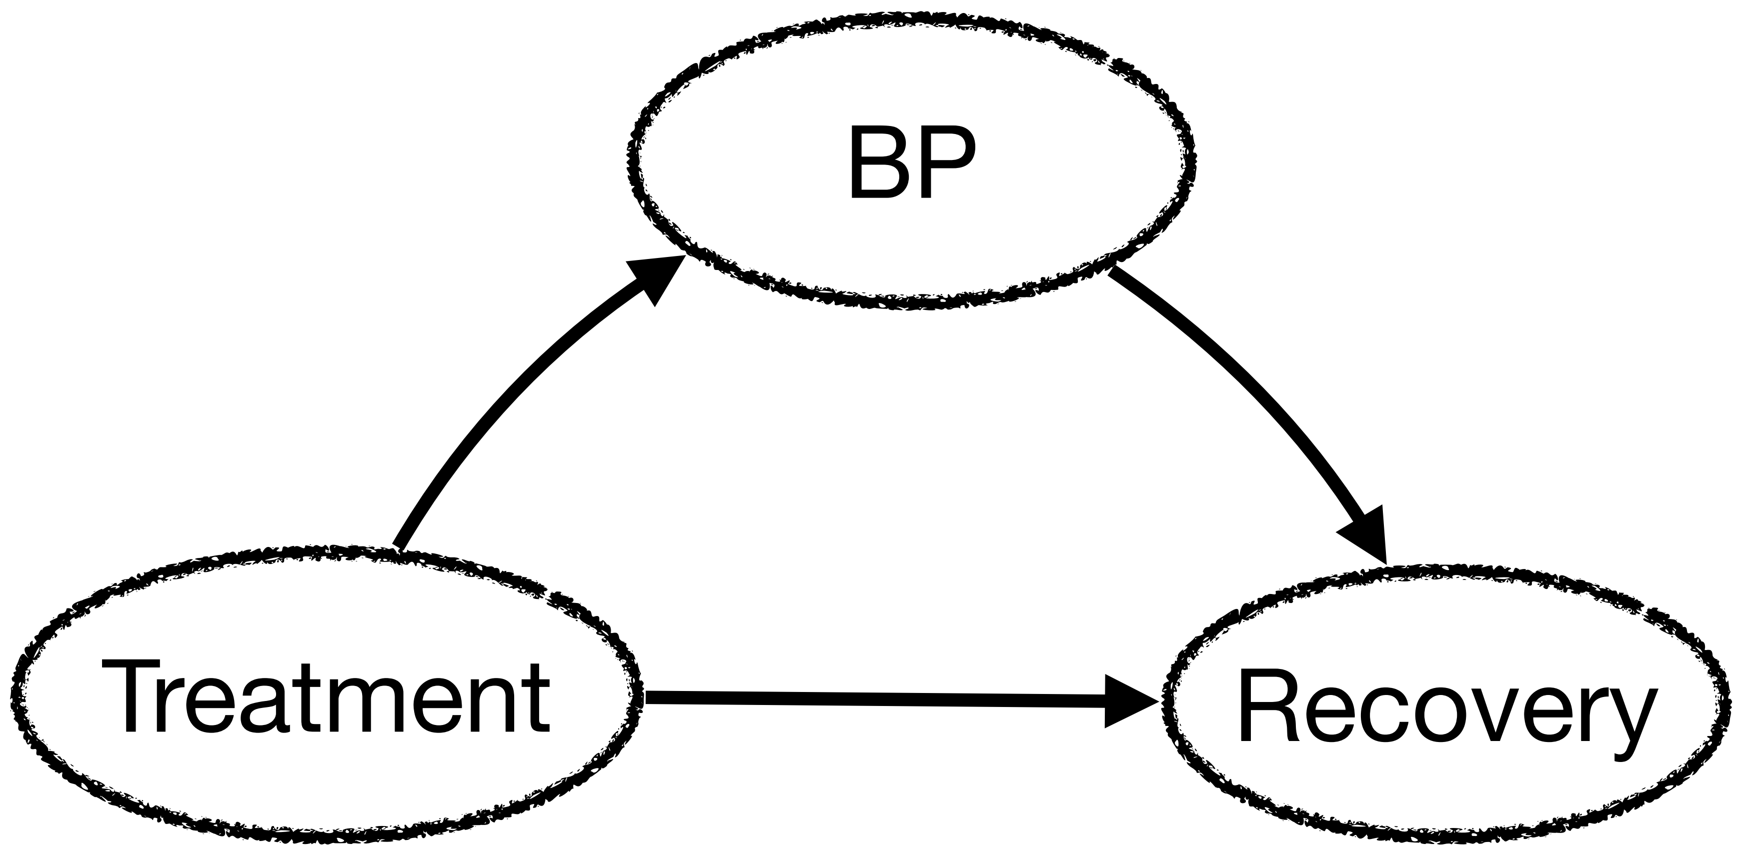

Survival Rate (%)		
Category	**Treated**	**Non-Treated**
Male	93	87
Female	73	69
Combined	78	83

Data Generating Process in Causal Inference

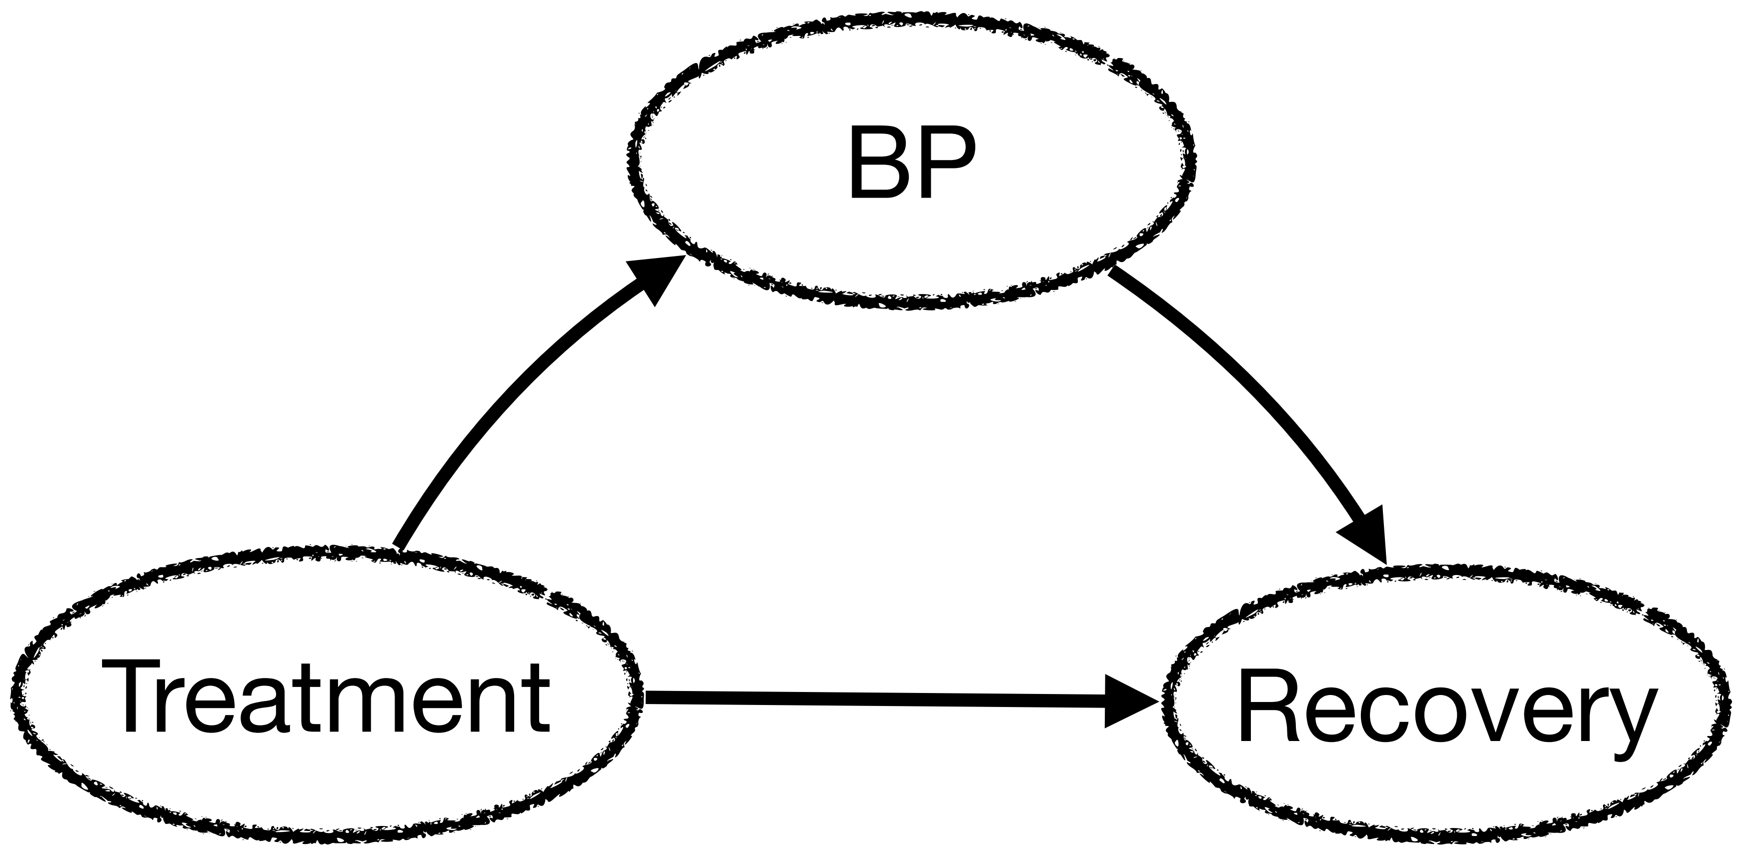

Survival Rate (%)		
Category	**Treated**	**Non-Treated**
Male	93	87
Female	73	69
Combined	78	83

If the data generating process is given as a causal diagram,

Data Generating Process in Causal Inference

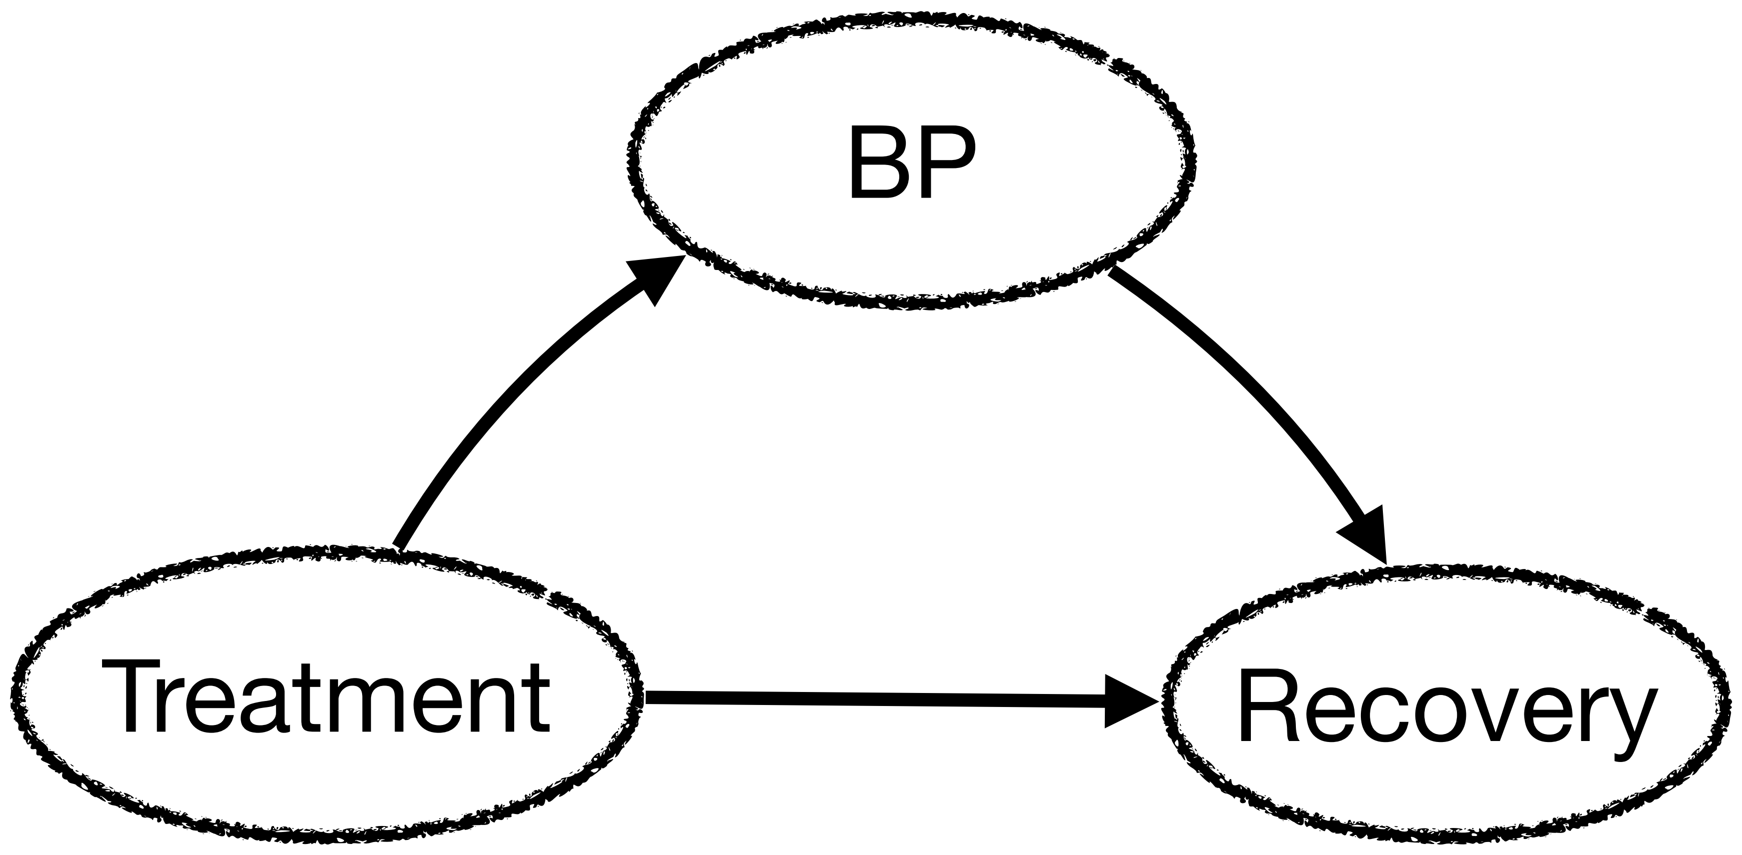

Category	Treated	Non-Treated
Survival Rate (%)		
Male	93	87
Female	73	69
Combined	78	83

If the data generating process is given as a causal diagram,

=> The treatment is harmful.

For Causal Inference, Understanding the DGP is Crucial

For Causal Inference, Understanding the DGP is Crucial

Two different DGPs have the same correlation structure but different causality structures.

For Causal Inference, Understanding the DGP is Crucial

Two different DGPs have the same correlation structure but different causality structures.

=> For causal inference, understanding the DGP is crucial.

What is Causality?

Chronicles of Causality

First Attempt: Correlation

David Hume

"We may define a ***cause*** to be an object, followed by another, and where all the objects similar to the first are followed by objects similar to the second" (1752, Hume)

First Attempt: Correlation

David Hume

"We may define a ***cause*** to be an object, followed by another, and where all the objects similar to the first are followed by objects similar to the second" (1752, Hume)

=> X is a cause of Y, if X happens and then Y happens.

First Attempt: Correlation

David Hume

“We may define a ***cause*** to be an object, followed by another, and where all the objects similar to the first are followed by objects similar to the second” (1752, Hume)

=> X is a cause of Y, if X happens and then Y happens.

=> X is a cause of Y, if X and Y has correlation.

First Attempt: Correlation

David Hume

"We may define a ***cause*** to be an object, followed by another, and where all the objects similar to the first are followed by objects similar to the second" (1752, Hume)

=> X is a cause of Y, if X happens and then Y happens.

=> X is a cause of Y, if X and Y has correlation.

Correlation $\neq$ Causation.

Second Attempt: Counterfactual / Potential-Outcome

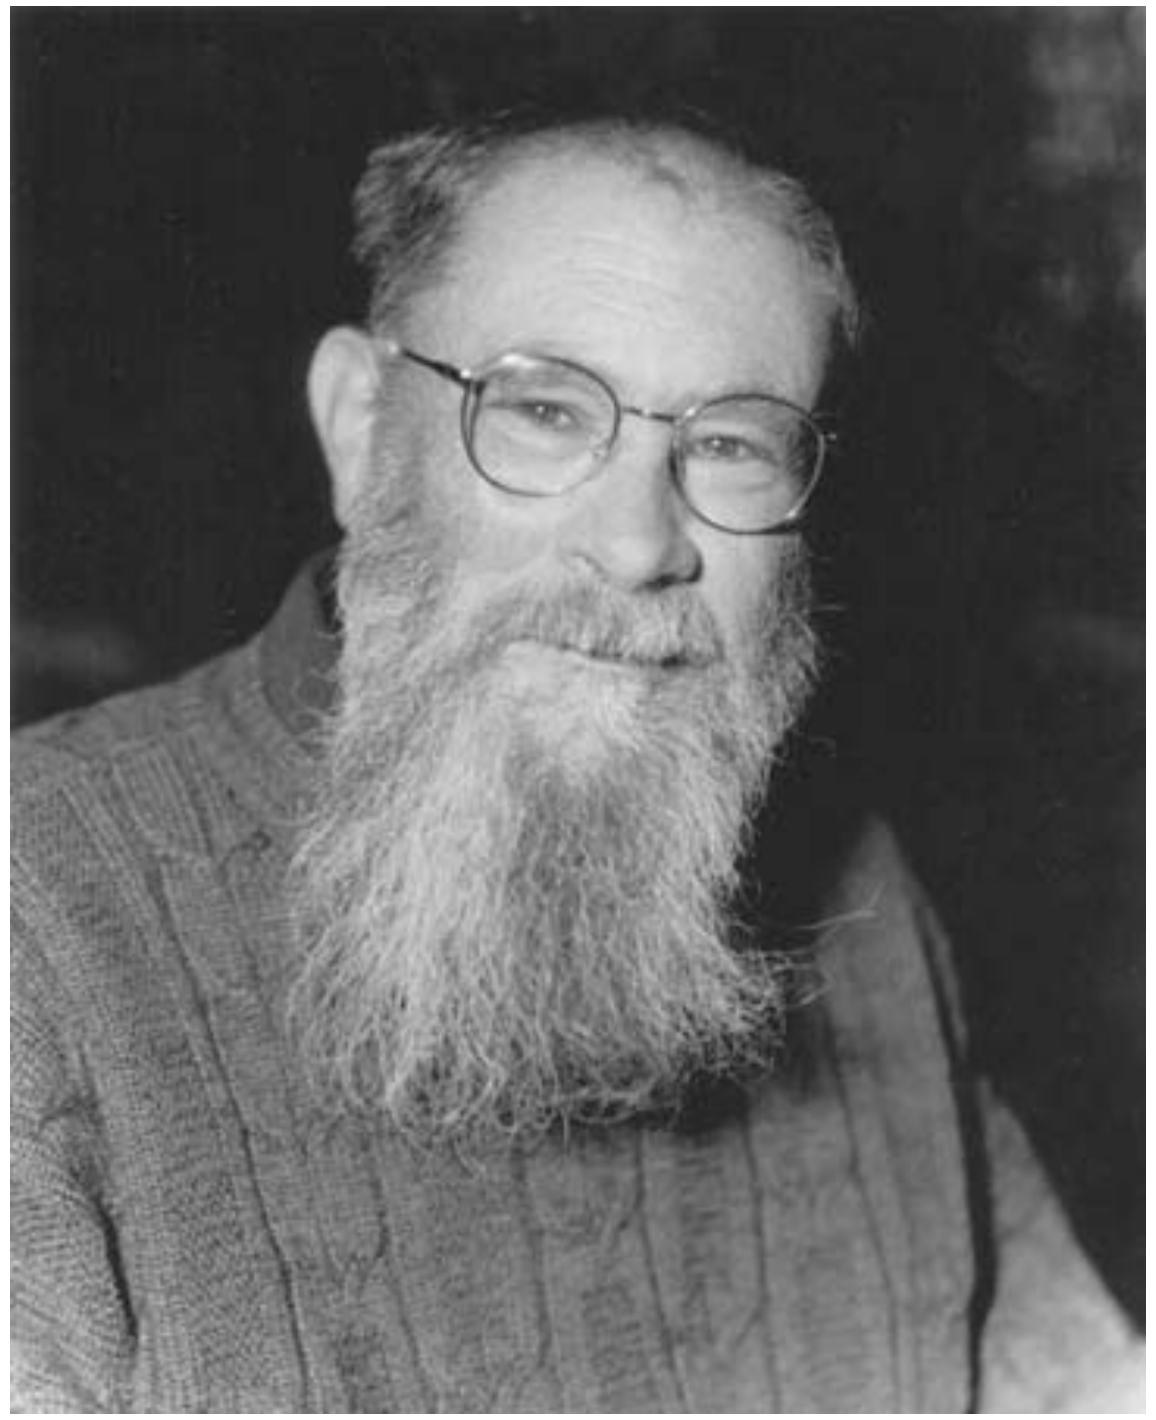

David Lewis

"We may define a ***cause*** as something that makes a difference, and the difference wouldn't happened without the **cause**" (Lewis, 1973).

Second Attempt: Counterfactual / Potential-Outcome

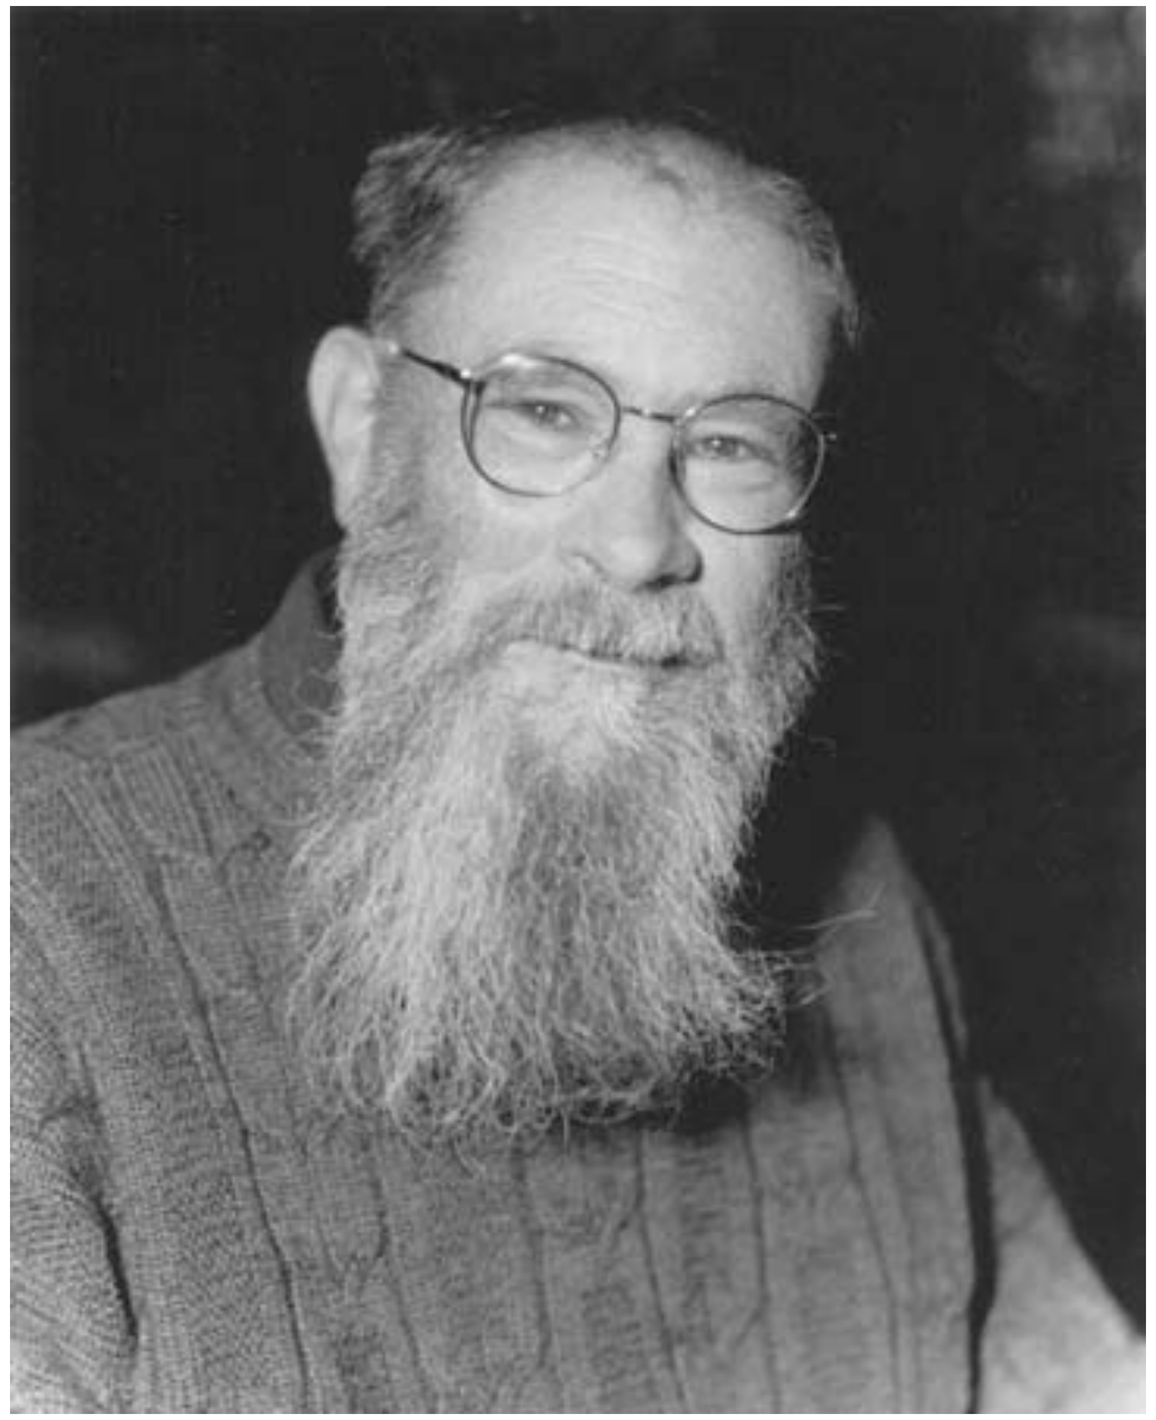

David Lewis

"We may define a ***cause*** as something that makes a difference, and the difference wouldn't happened without the **cause**" (Lewis, 1973).

X is a cause of Y, if

- Y would happened if X had been happened.
- Y wouldn't happened if X hadn't been happened.

Second Attempt: Counterfactual / Potential-Outcome

X is a cause of Y, if

- Y would happened if X had been happened.
- Y wouldn't happened if X hadn't been happened.

- **Example**: $Y(X = 1)$ is the recovery status (Y) if all patients in the population had taken the drug $(X = 1)$.
- X is a cause of Y, if $Y(X = 1) \neq Y(X = 0)$.

Second Attempt: Counterfactual / Potential-Outcome

X is a cause of Y, if

- Y would happened if X had been happened.
- Y wouldn't happened if X hadn't been happened.

Counterfactual (Potential-Outcome): $Y(X = x)$ is Y when values of X is set to x in their DGP (or population).

- **Example**: $Y(X = 1)$ is the recovery status (Y) if all patients in the population had taken the drug $(X = 1)$.
- X is a cause of Y, if $Y(X = 1) \neq Y(X = 0)$.

Counterfactual / Potential-Outcome: Example

X is a cause of Y, if

- Y would happened if X had been happened
- Y wouldn't happened if X hadn't been happened

Counterfactual / Potential-Outcome: Example

X is a cause of Y, if
- Y would happened if X had been happened
- Y wouldn't happened if X hadn't been happened

$X = 1$

$Y(X = 1) = 1$

Counterfactual / Potential-Outcome: Example

X is a cause of Y, if

- Y would happened if X had been happened
- Y wouldn't happened if X hadn't been happened

$X = 1$ $\quad$ $Y(X = 1) = 1$

$X = 0$ $\quad$ $Y(X = 0) = 0$

Counterfactual itself isn't enough: Example (1)

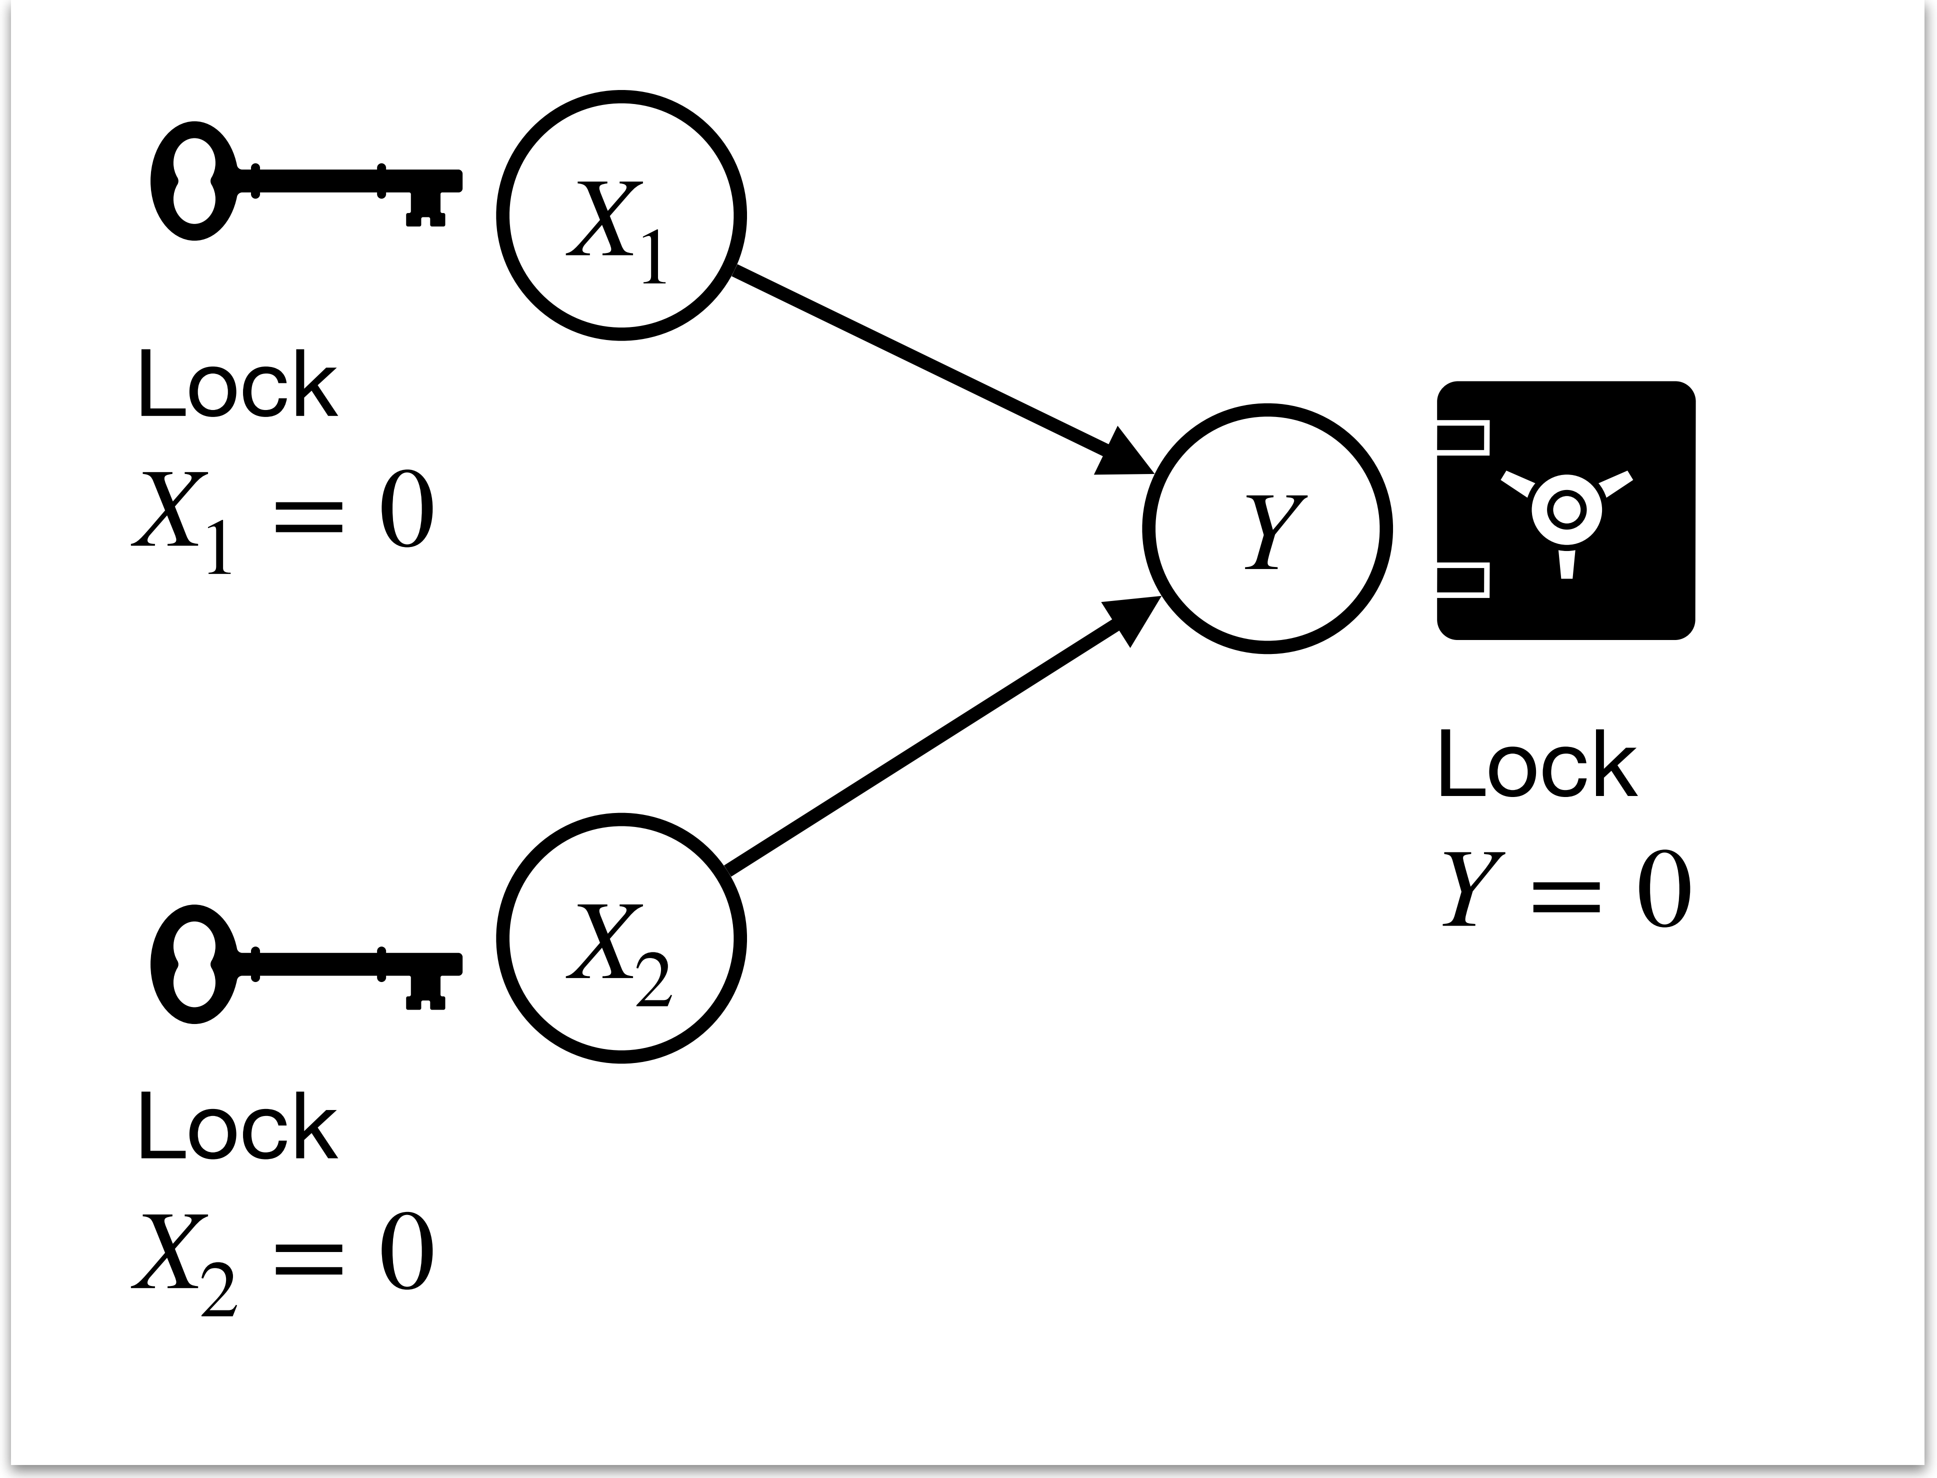

Counterfactual itself isn't enough: Example (1)

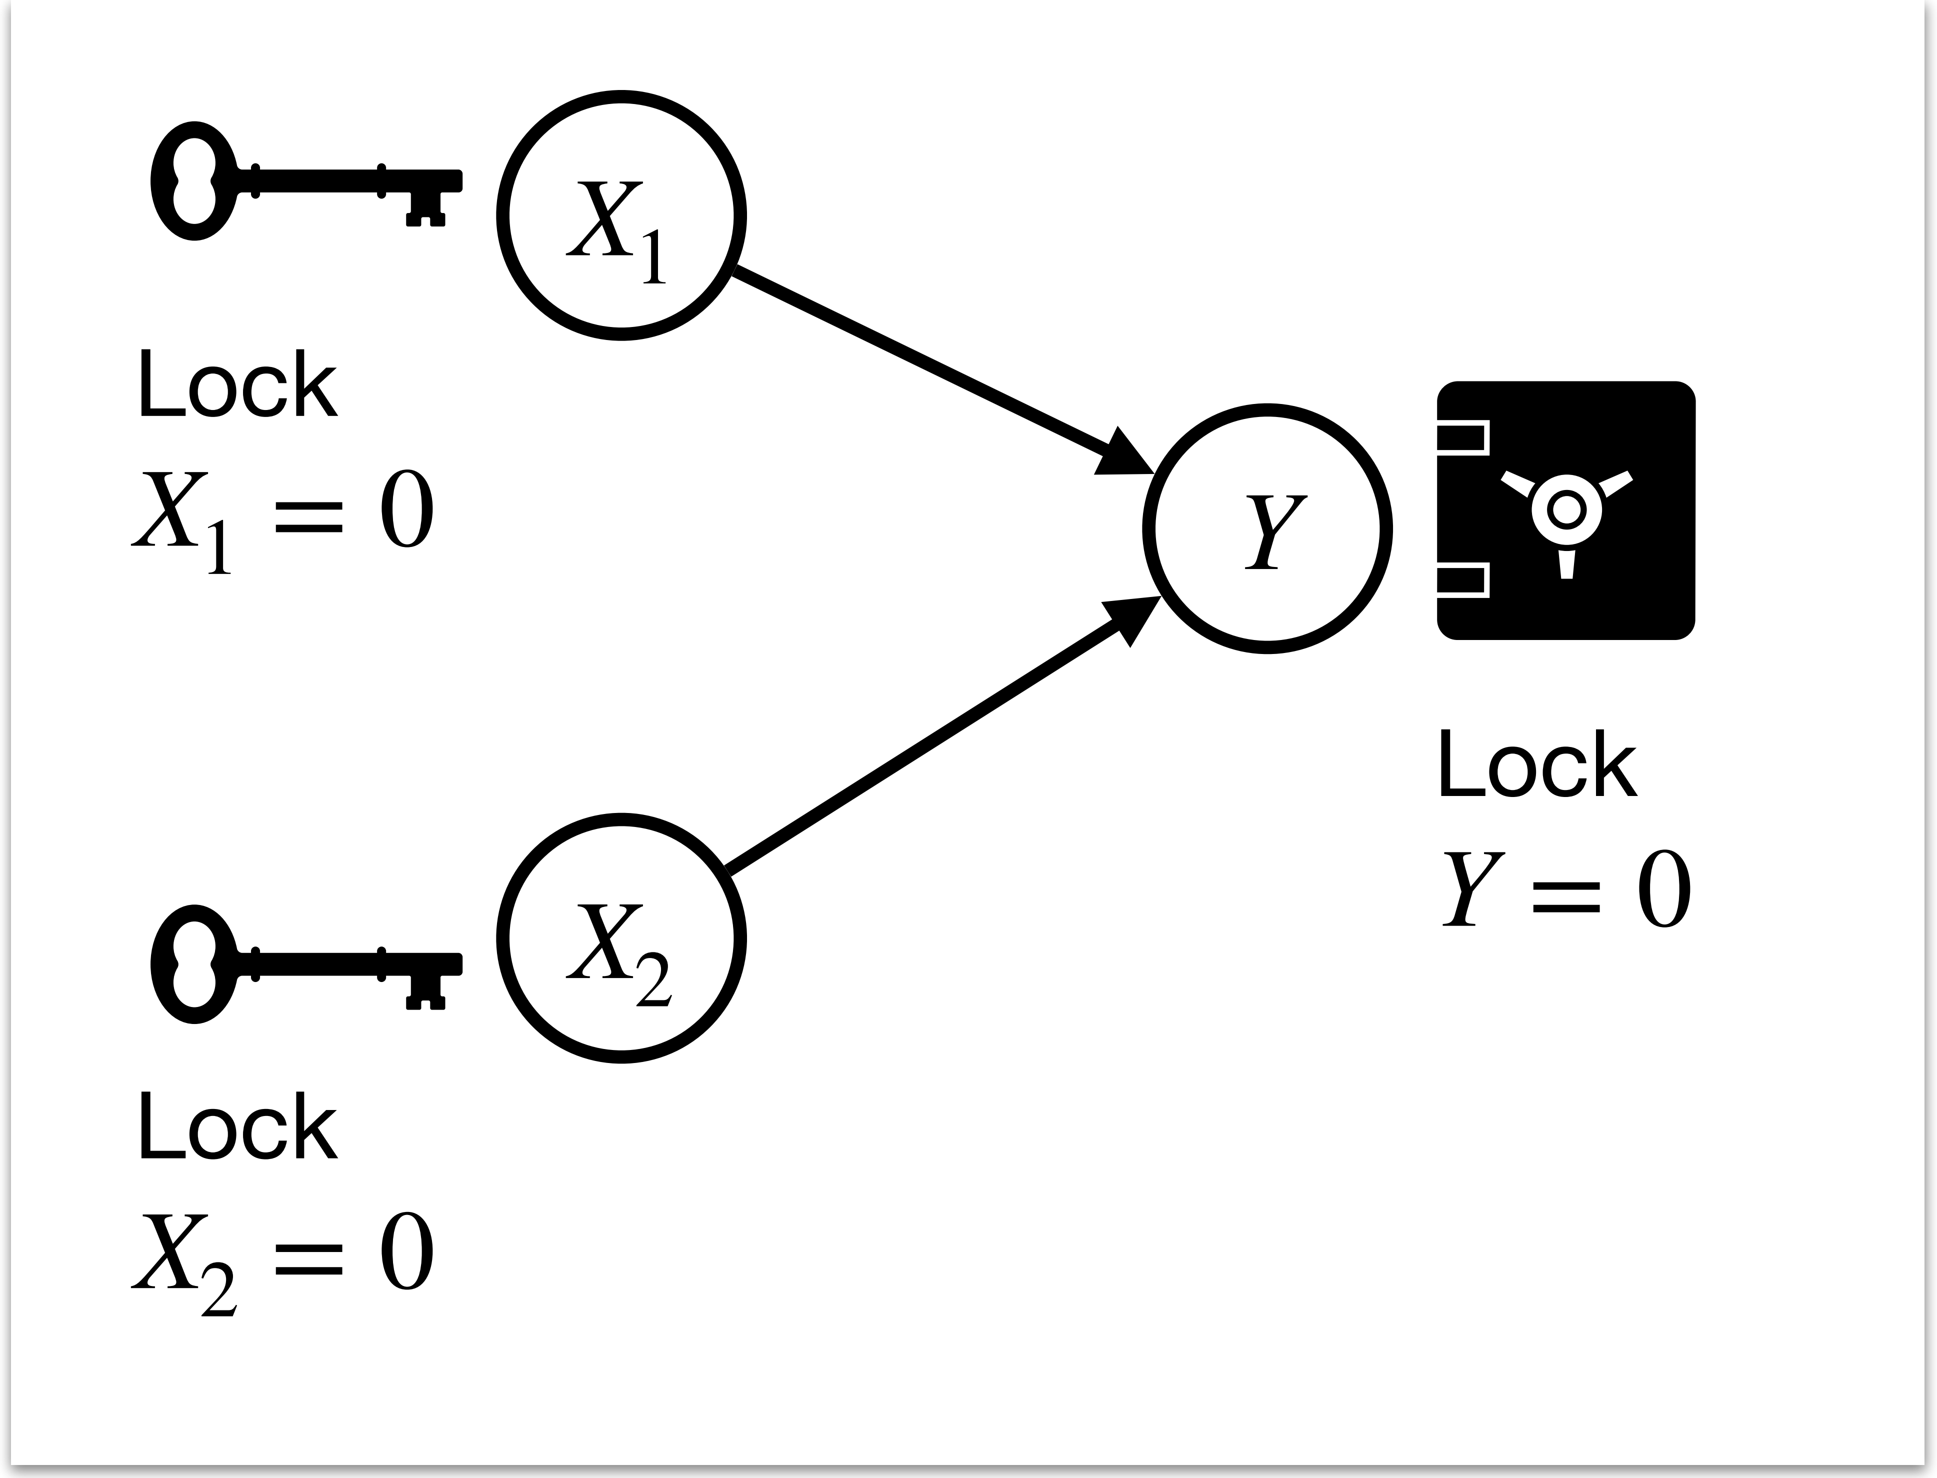

The door becomes unlocked ($Y = 1$) only when two locks are simultaneously unlocked ($X_1 = X_2 = 1$); i.e., $Y(X_1 = 1, X_2 = 1) = 1$.

Counterfactual itself isn't enough: Example (1)

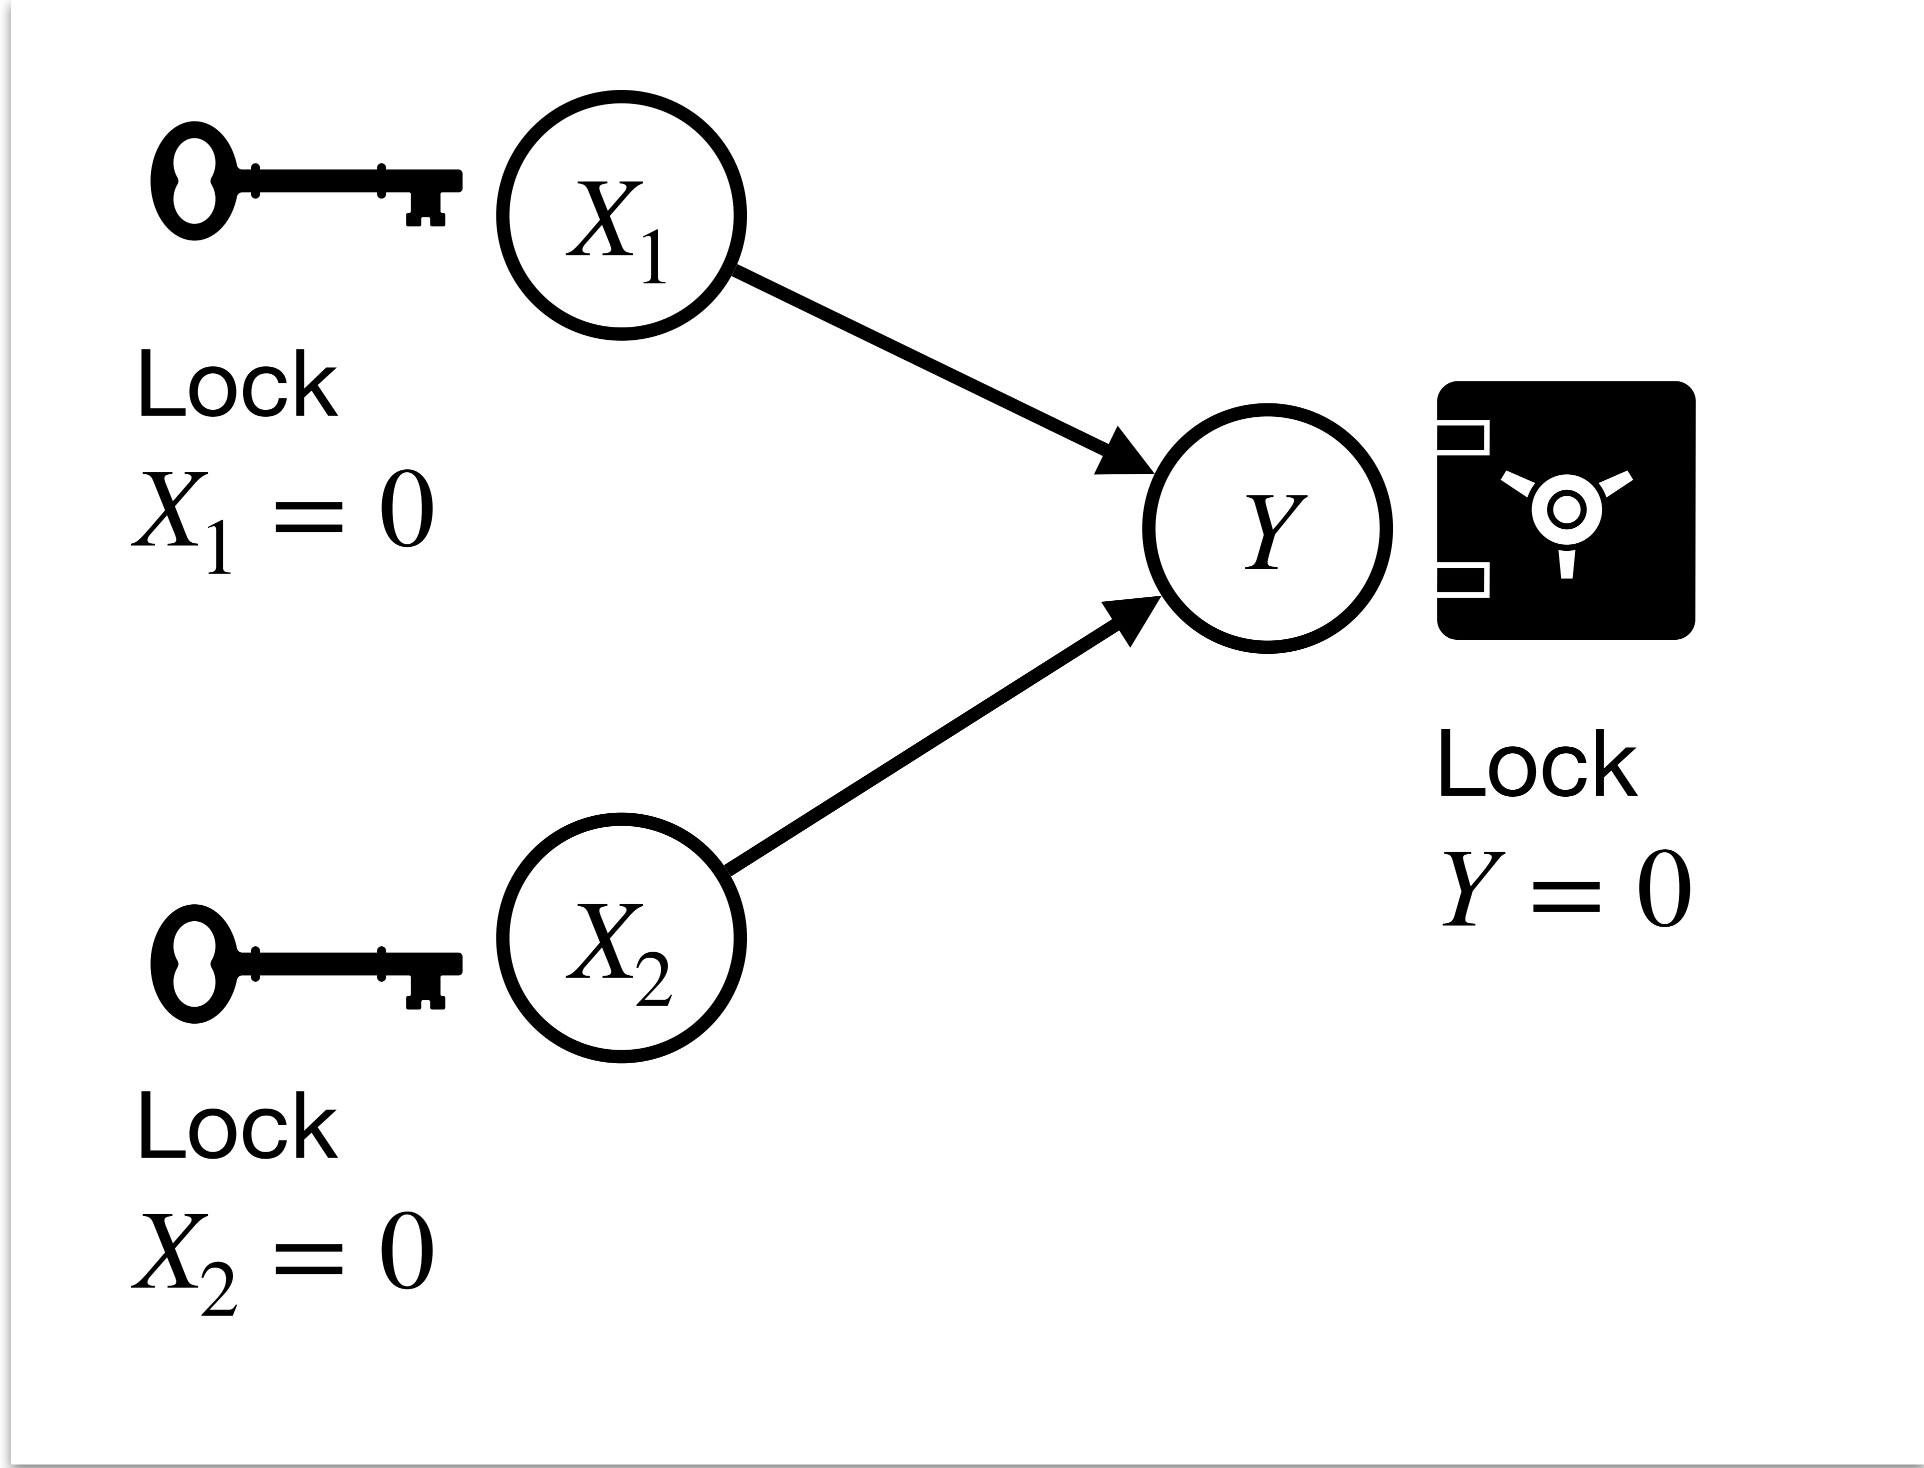

The door becomes unlocked ($Y = 1$) only when two locks are simultaneously unlocked ($X_1 = X_2 = 1$); i.e., $Y(X_1 = 1, X_2 = 1) = 1$.

A default state is that two locks are locked ($X_1 = X_2 = 0$), and the door is also locked ($Y = 0$) as a result.

Counterfactual itself isn't enough: Example (2)

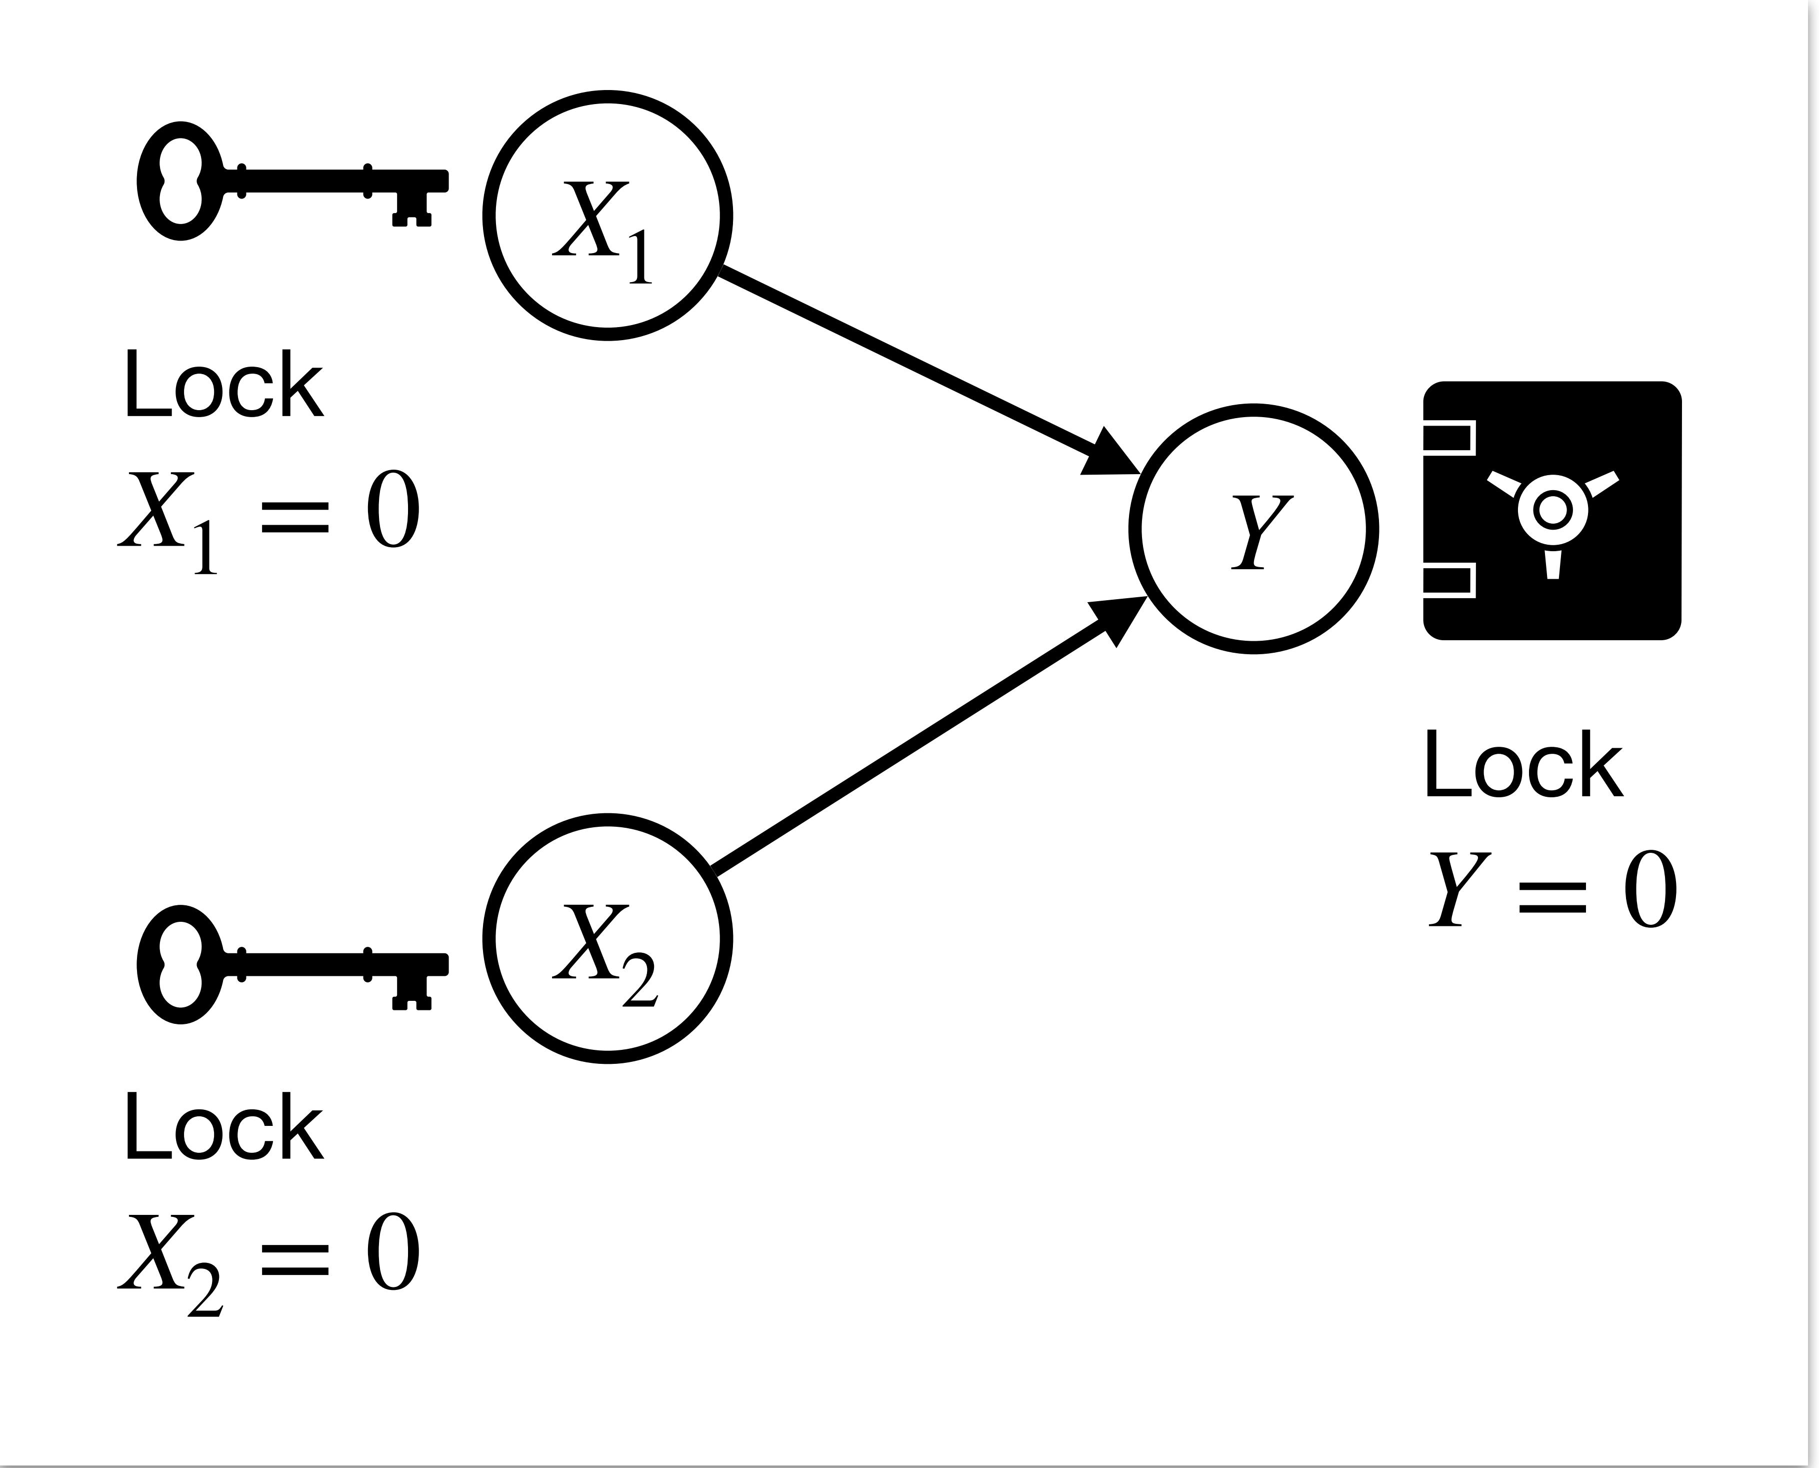

Counterfactual itself isn't enough: Example (2)

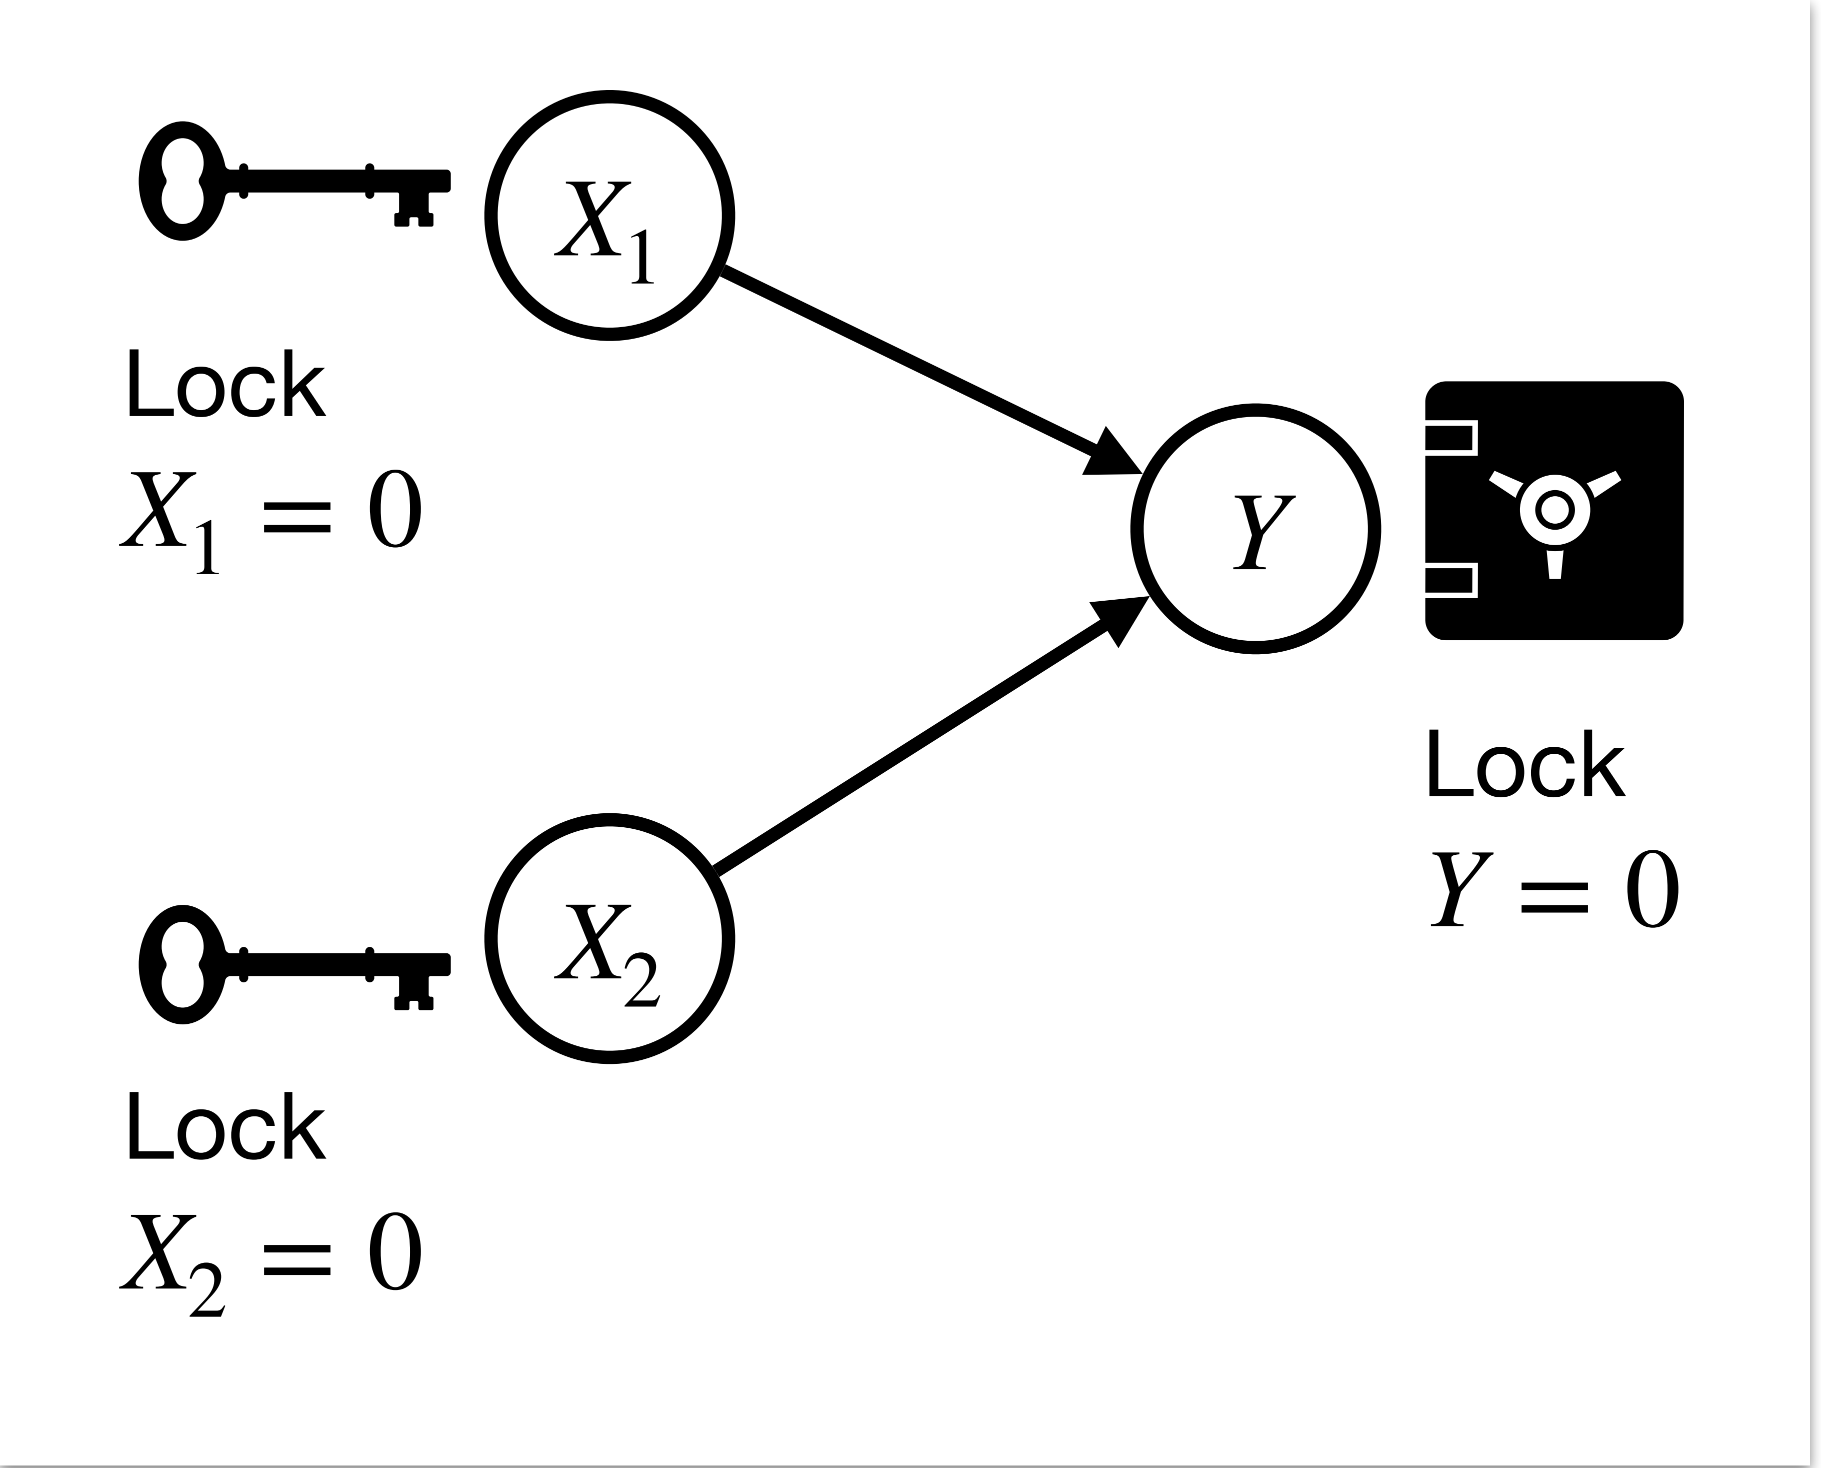

- If X_1 had been locked ($X_1 = 0$), then Y would be locked ($Y = 0$); $Y(X_1 = 0, X_2 = 0) = 0$
- If X_1 had been unlocked ($X_1 = 1$), Y would be still locked ($Y = 0$), because X_2 is set to be locked; i.e., $Y(X_1 = 1, X_2 = 0) = 0$

Counterfactual itself isn't enough: Example (2)

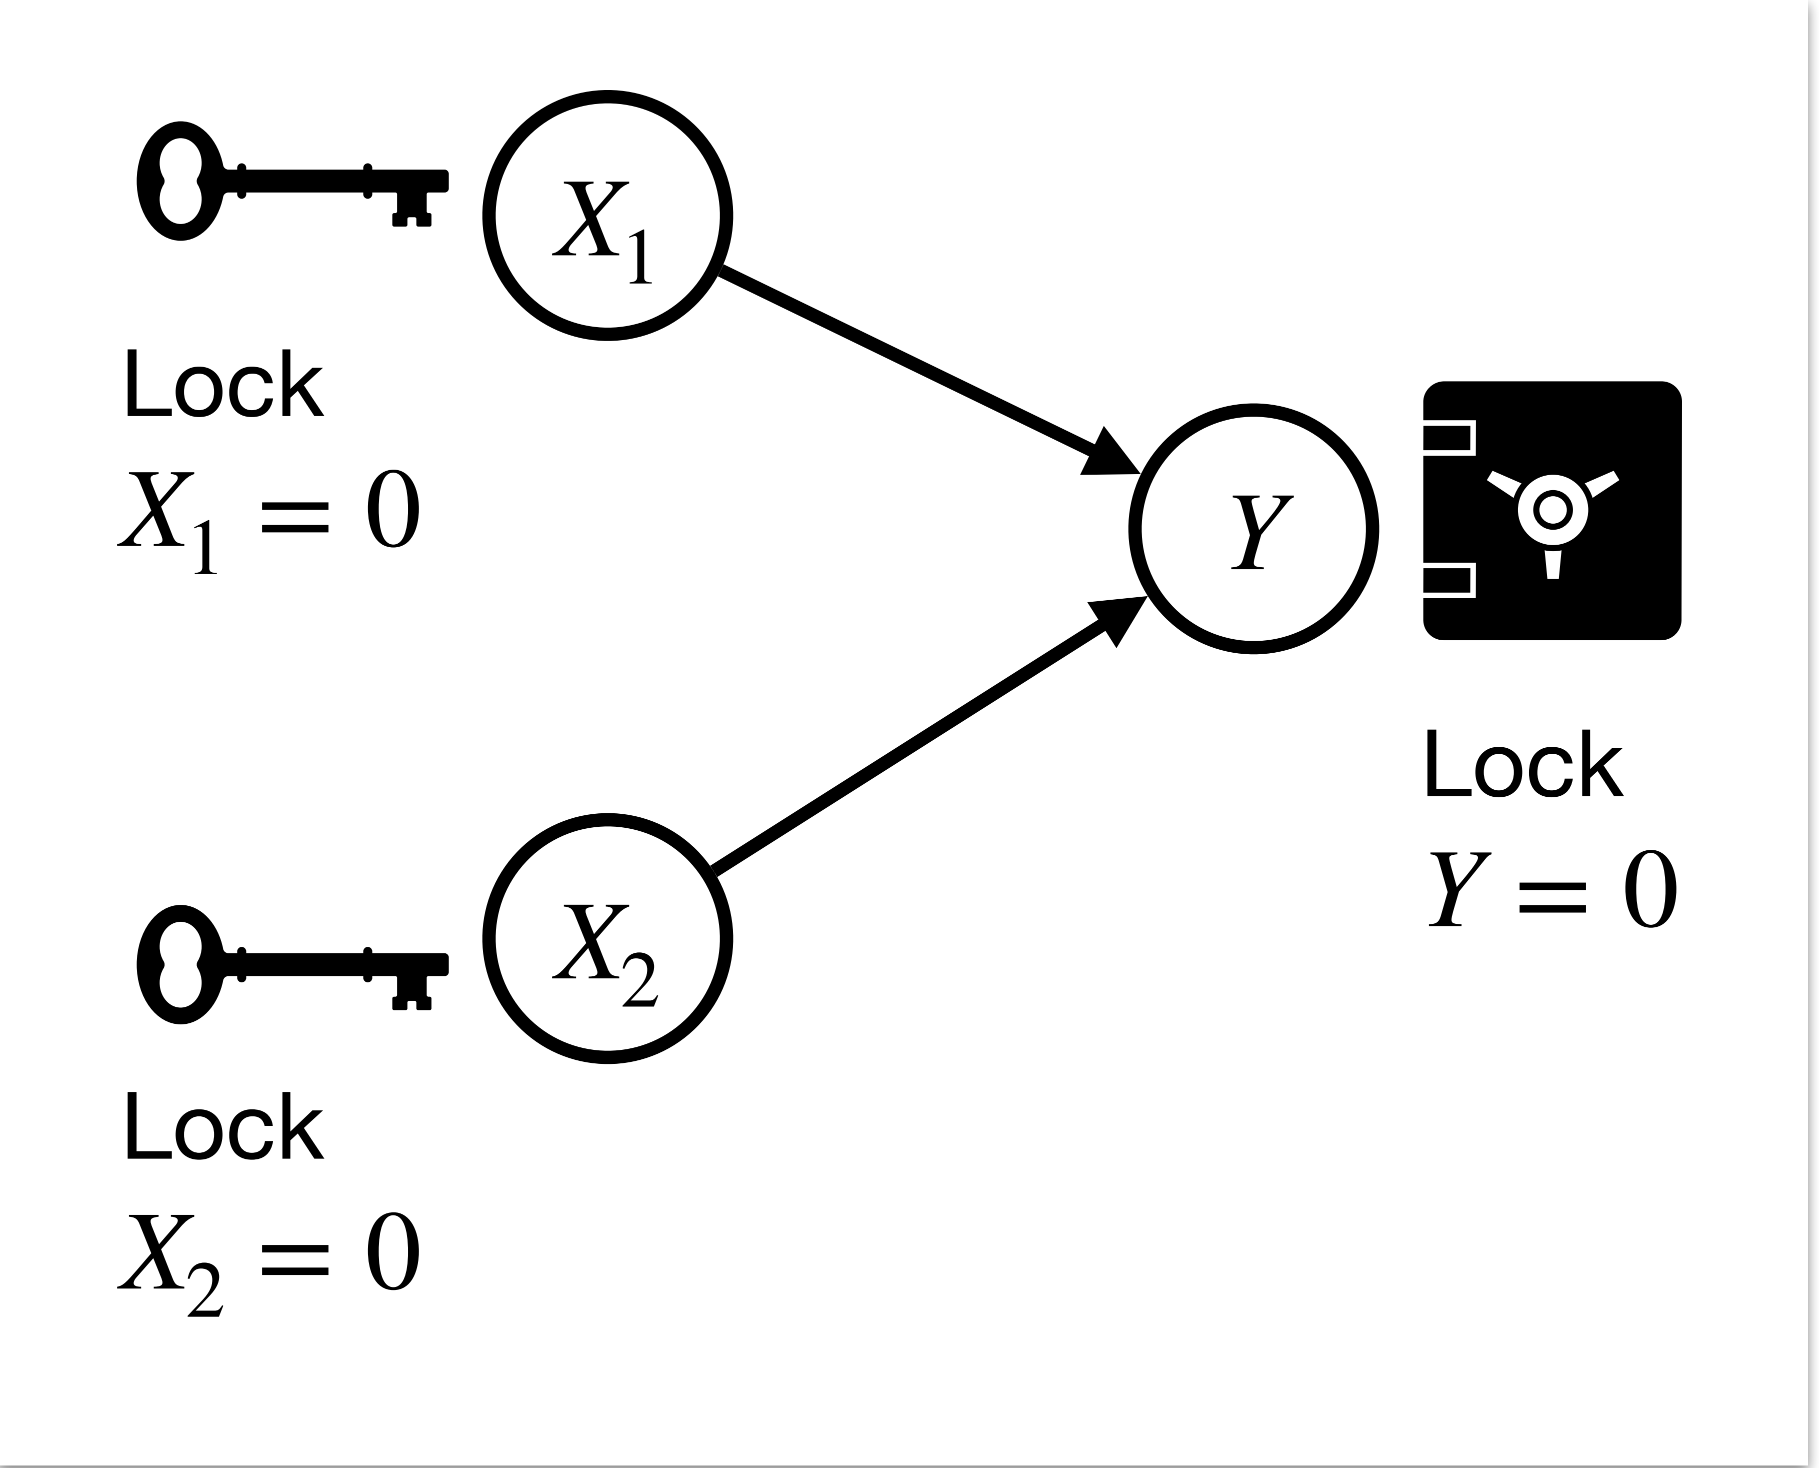

- If X_1 had been locked ($X_1 = 0$), then Y would be locked ($Y = 0$); $Y(X_1 = 0, X_2 = 0) = 0$
- If X_1 had been unlocked ($X_1 = 1$), Y would be still locked ($Y = 0$), because X_2 is set to be locked; i.e., $Y(X_1 = 1, X_2 = 0) = 0$

The event X_1 doesn't make difference of the counterfactual (potential) outcome of Y.

Counterfactual itself isn't enough: Example (2)

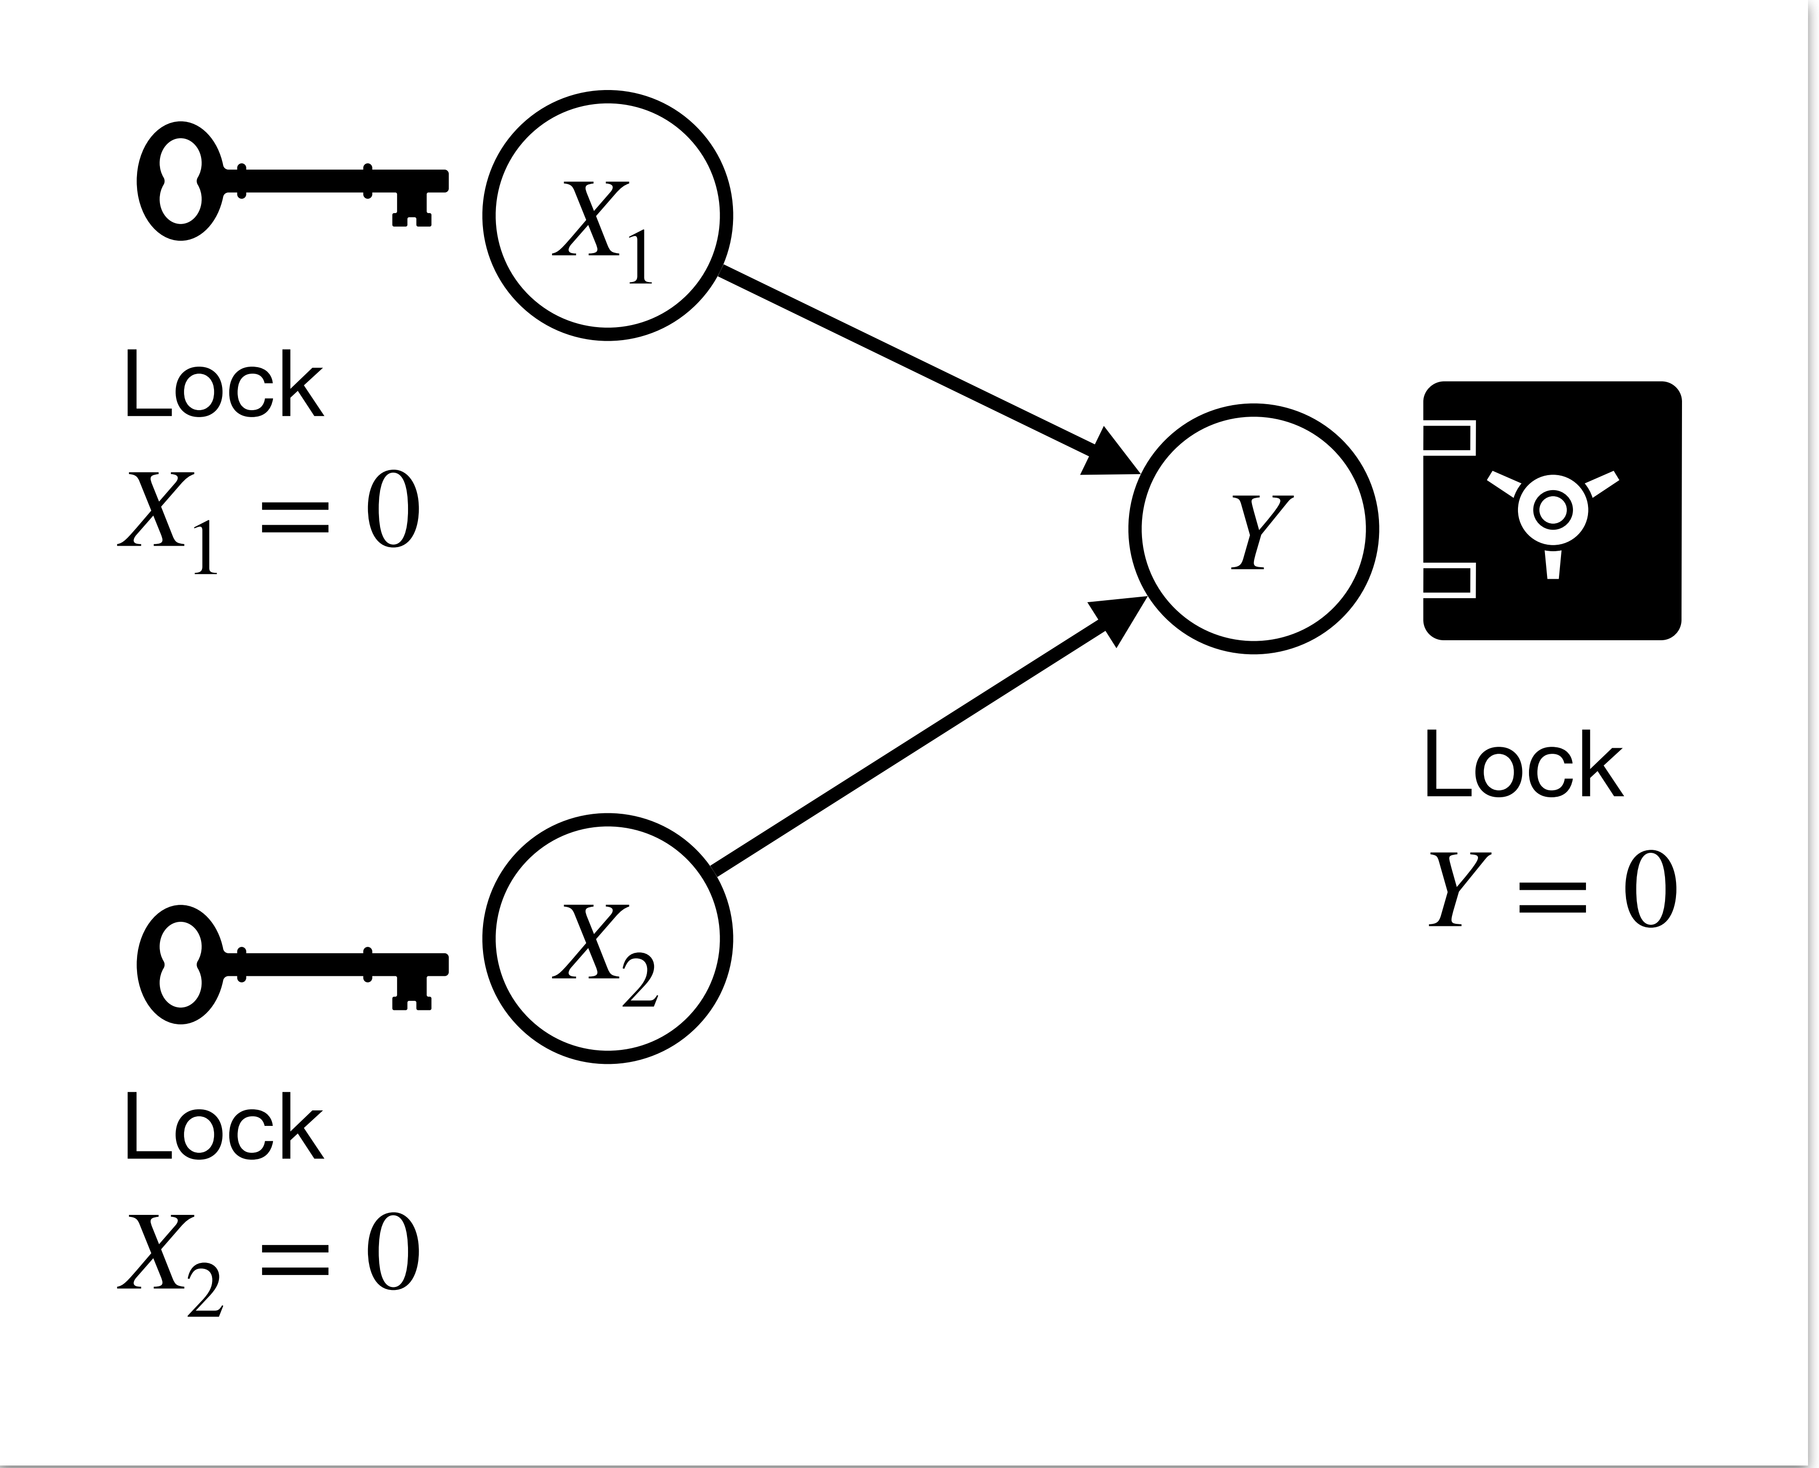

- If X_1 had been locked ($X_1 = 0$), then Y would be locked ($Y = 0$); $Y(X_1 = 0, X_2 = 0) = 0$
- If X_1 had been unlocked ($X_1 = 1$), Y would be still locked ($Y = 0$), because X_2 is set to be locked; i.e., $Y(X_1 = 1, X_2 = 0) = 0$

The event X_1 doesn't make difference of the counterfactual (potential) outcome of Y.

=> X_1 is not a cause of Y??

Counterfactual itself isn't enough: Example (3)

$do(X_2 = 1)$

Counterfactual itself isn't enough: Example (3)

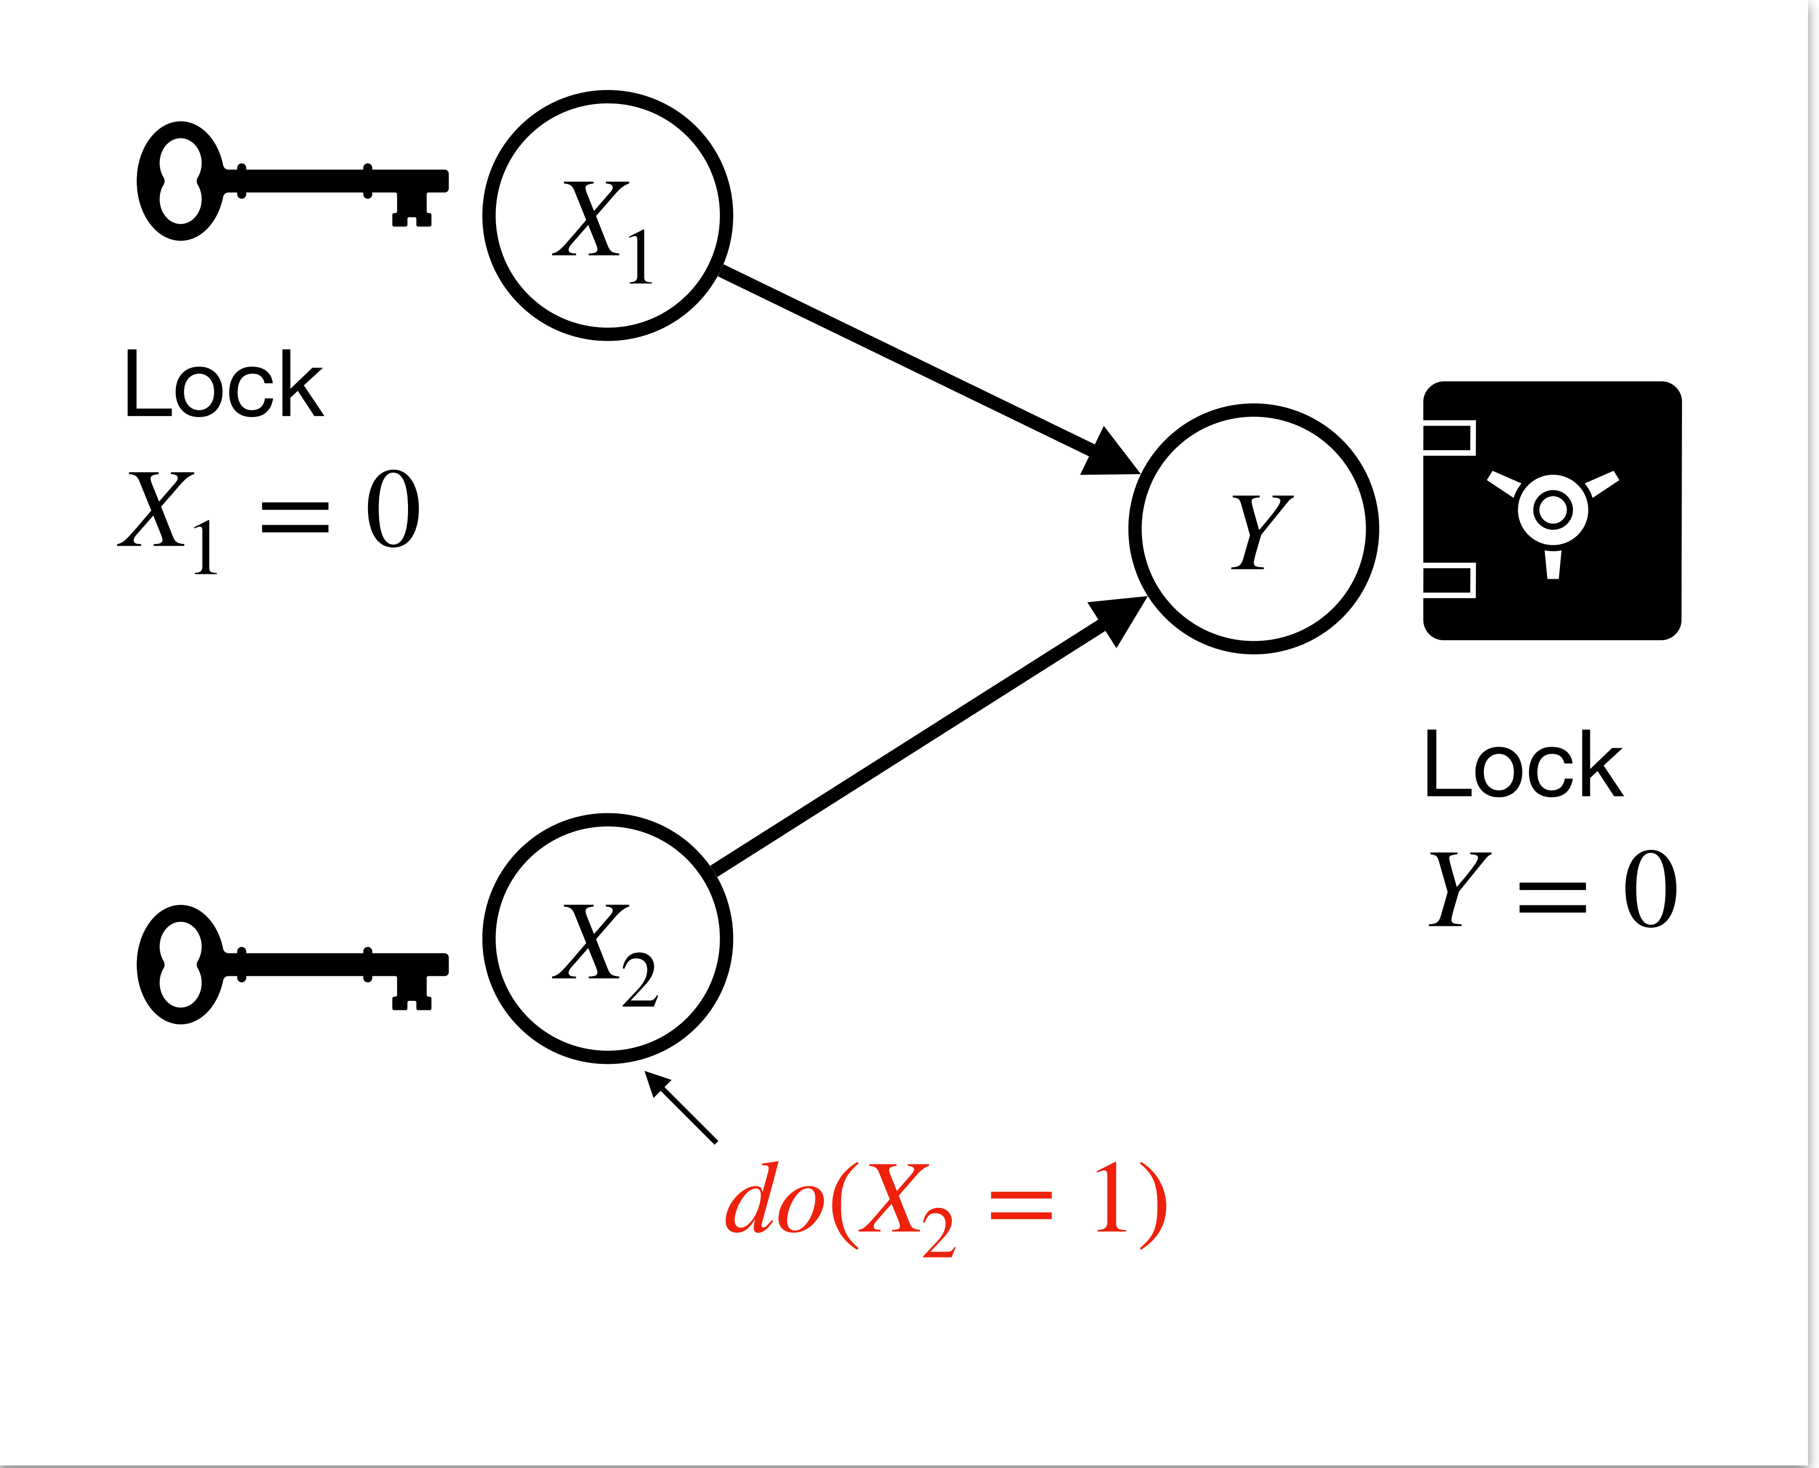

Consider X_1 in which $X_2 = 1$ had been set (i.e., X_2 is unlocked)

- If X_1 had been locked ($X_1 = 0$), then Y would be locked ($Y = 0$); $Y(X_1 = 0, X_2 = 1) = 0$
- If X_1 had been unlocked ($X_1 = 1$), Y would be unlocked ($Y = 1$), because X_2 is set to be unlocked; i.e., $Y(X_1 = 1, X_2 = 1) = 1$

Counterfactual itself isn't enough: Example (3)

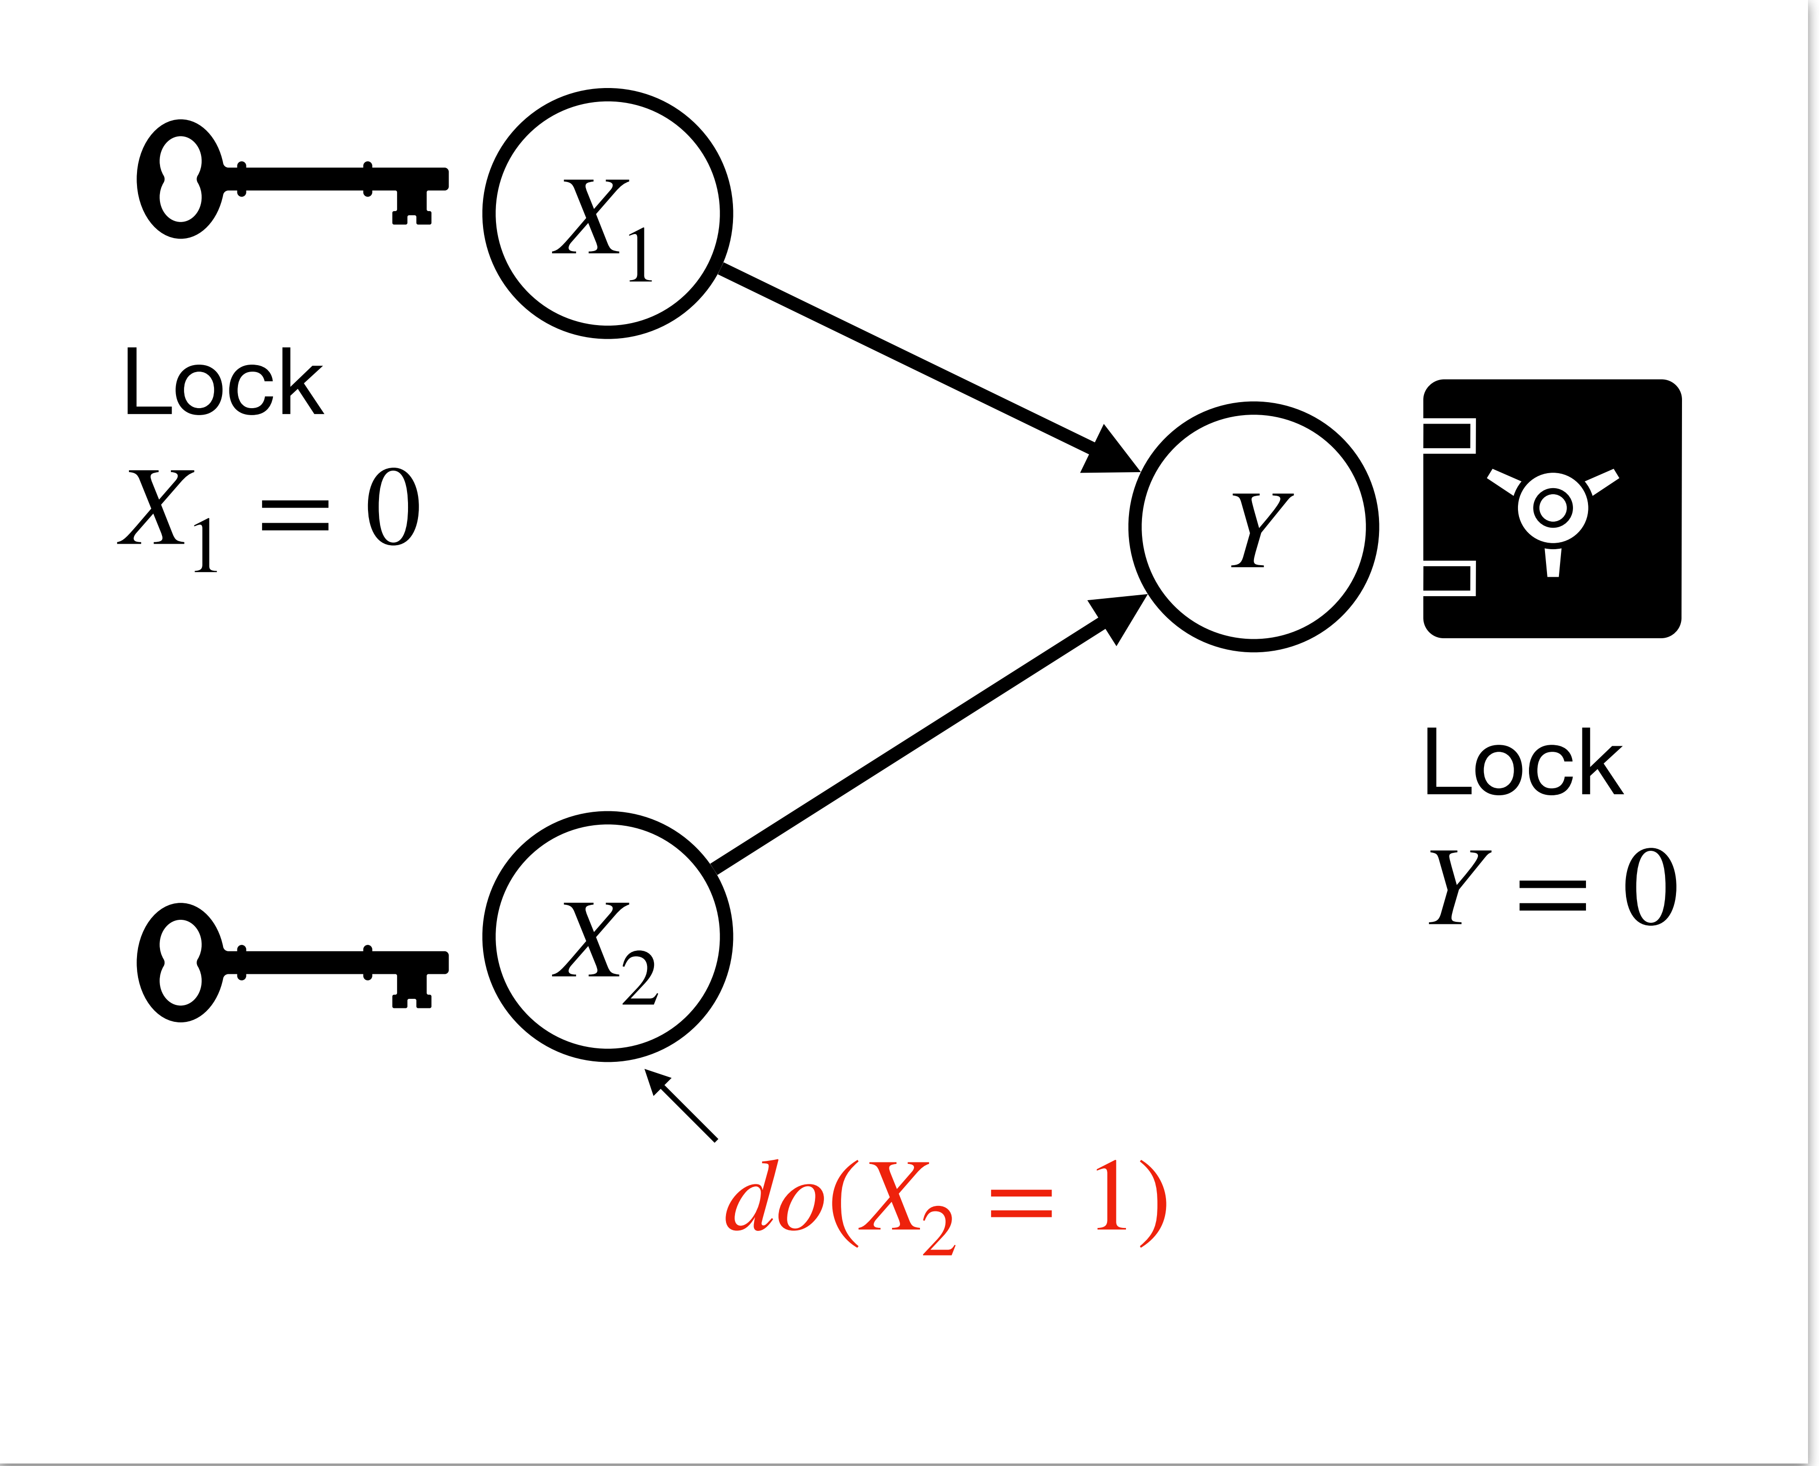

Consider X_1 in which $X_2 = 1$ had been set (i.e., X_2 is unlocked)

- If X_1 had been locked ($X_1 = 0$), then Y would be locked ($Y = 0$); $Y(X_1 = 0, X_2 = 1) = 0$
- If X_1 had been unlocked ($X_1 = 1$), Y would be unlocked ($Y = 1$), because X_2 is set to be unlocked; i.e., $Y(X_1 = 1, X_2 = 1) = 1$

The event X_1 makes difference of the counterfactual (potential) outcome of Y.

Counterfactual itself isn't enough: Example (3)

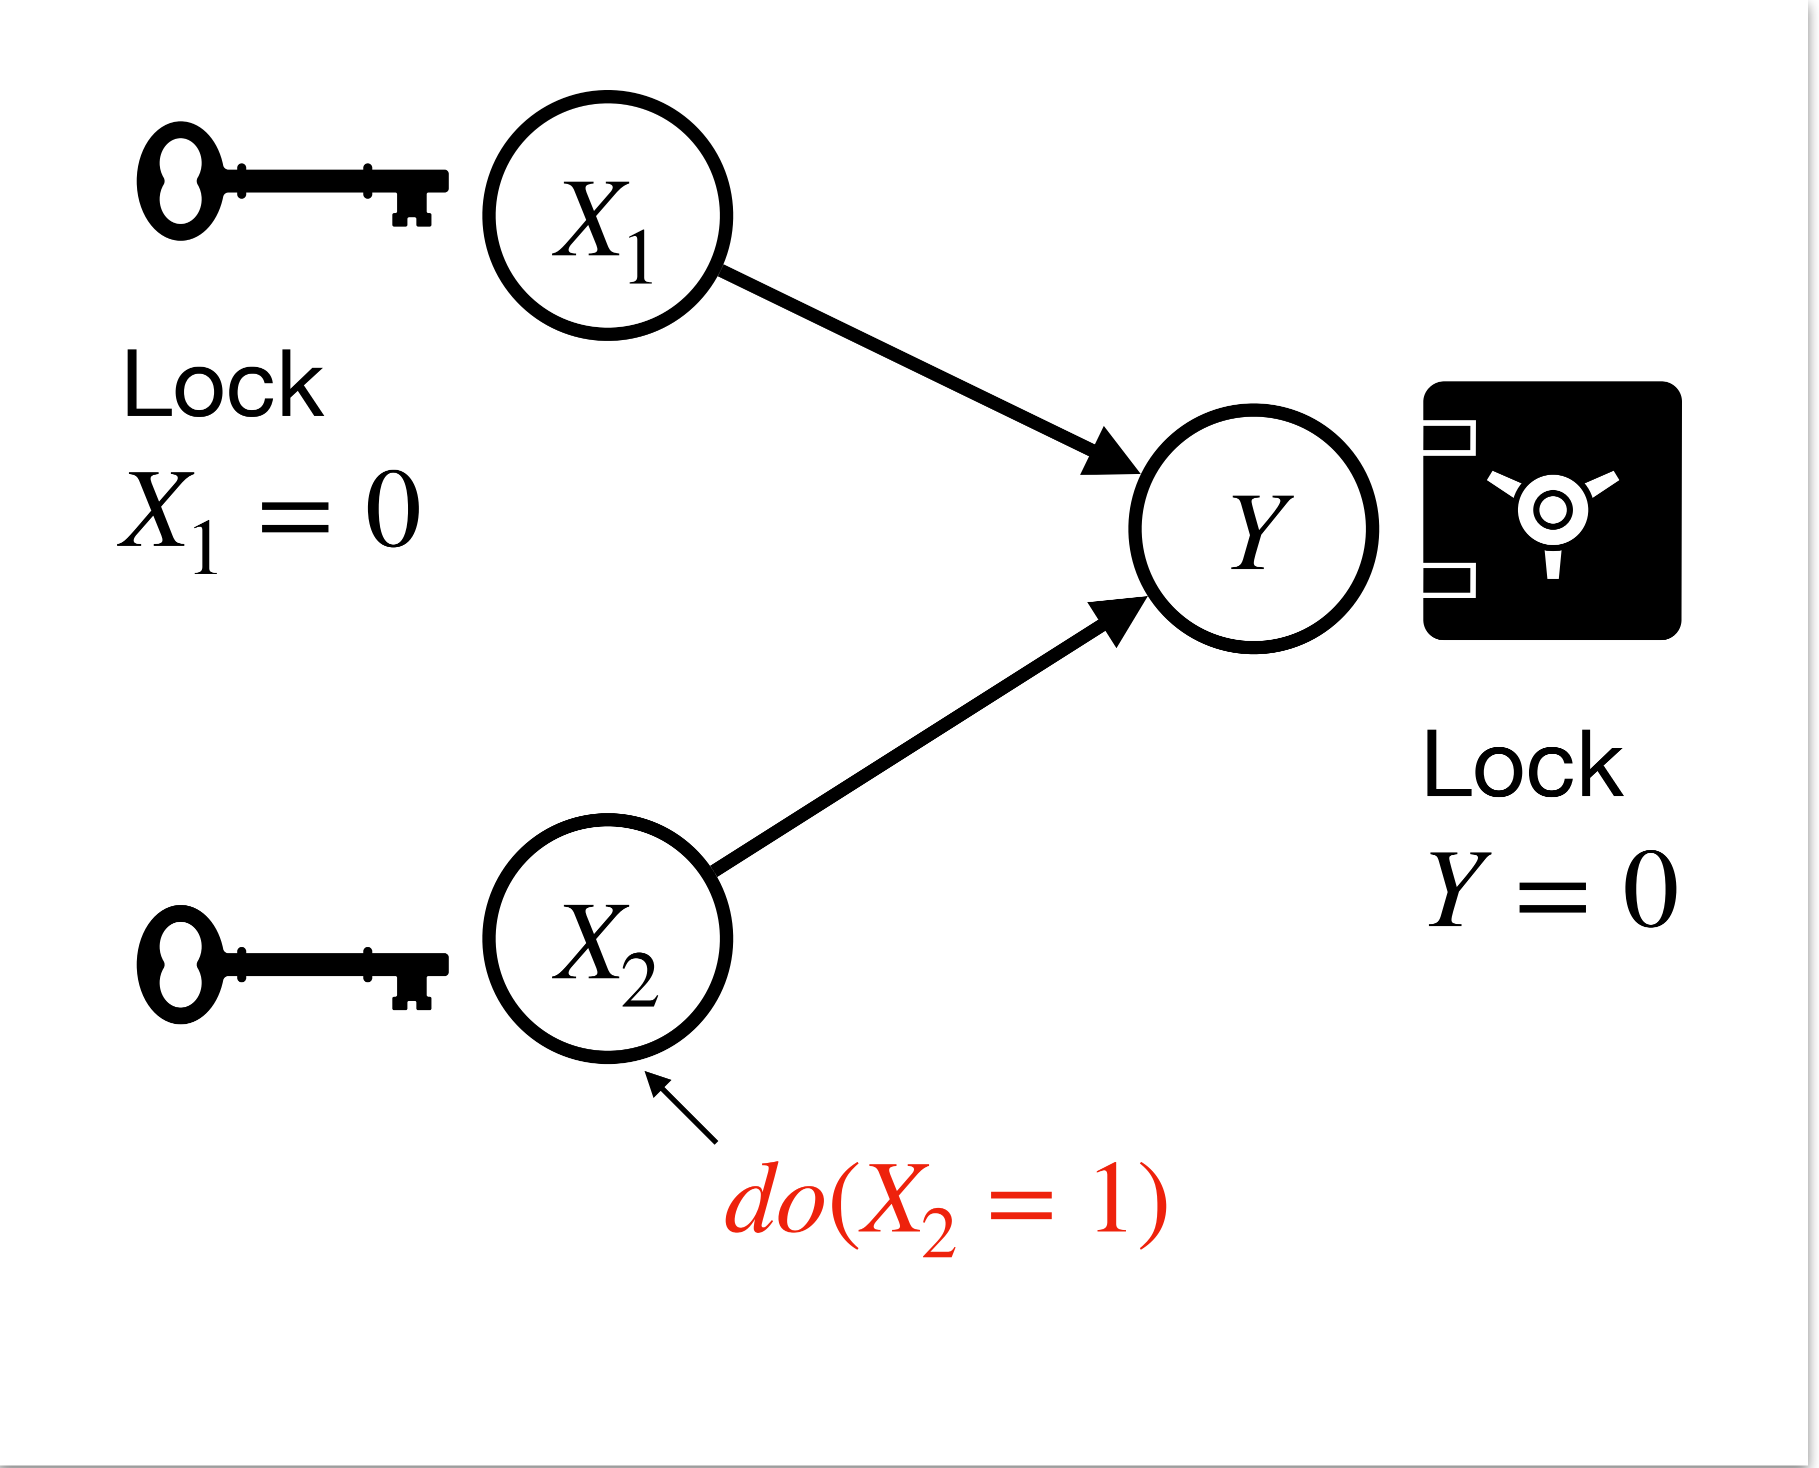

Consider X_1 in which $X_2 = 1$ had been set (i.e., X_2 is unlocked)

- If X_1 had been locked ($X_1 = 0$), then Y would be locked ($Y = 0$); $Y(X_1 = 0, X_2 = 1) = 0$
- If X_1 had been unlocked ($X_1 = 1$), Y would be unlocked ($Y = 1$), because X_2 is set to be unlocked; i.e., $Y(X_1 = 1, X_2 = 1) = 1$

The event X_1 makes difference of the counterfactual (potential) outcome of Y.

=> X_1 is a cause of Y.

Counterfactual itself isn't enough: Takeaway

Counterfactual itself isn’t enough: Takeaway

- The counterfactual itself cannot reveal the causality without considering the corresponding DGP.

Counterfactual itself isn’t enough: Takeaway

- The counterfactual itself cannot reveal the causality without considering the corresponding DGP.

- The causality can be revealed by considering ***relations between variables*** in the DGP.

Structural Causal Model

DGP of the counterfactuals (i.e., DGPs taking account of causality).

Structural Causal Model

Structural Causal Model

The structural Causal Model (SCM) can represent the DGP considering the relation of variables.

Structural Causal Model

The structural Causal Model (SCM) can represent the DGP considering the relation of variables.

Structural Causal Model $M := \langle \mathbf{V}, \mathbf{U}, \mathbf{F}, P(\mathbf{u}) \rangle$

- $\mathbf{V}$: A set of endogenous (observable) variables.
- $\mathbf{U}$: A set of exogenous (latent) variables.
- $\mathbf{F}$: A set of structural equations $\{f_{V_i}\}_{V_i \in \mathbf{V}}$ determining the value of $V_i \in \mathbf{V}$, where $V_i \leftarrow f_{v_i}(PA_{V_i}, U_{V_i})$ for some $PA_{V_i} \subseteq \mathbf{V}$ and $U_{V_i} \subseteq \mathbf{U}$.
- $P(\mathbf{u})$: A probability measure for $\mathbf{U}$.

Example of the SCM: Encoding the DGP

Example of the SCM: Encoding the DGP

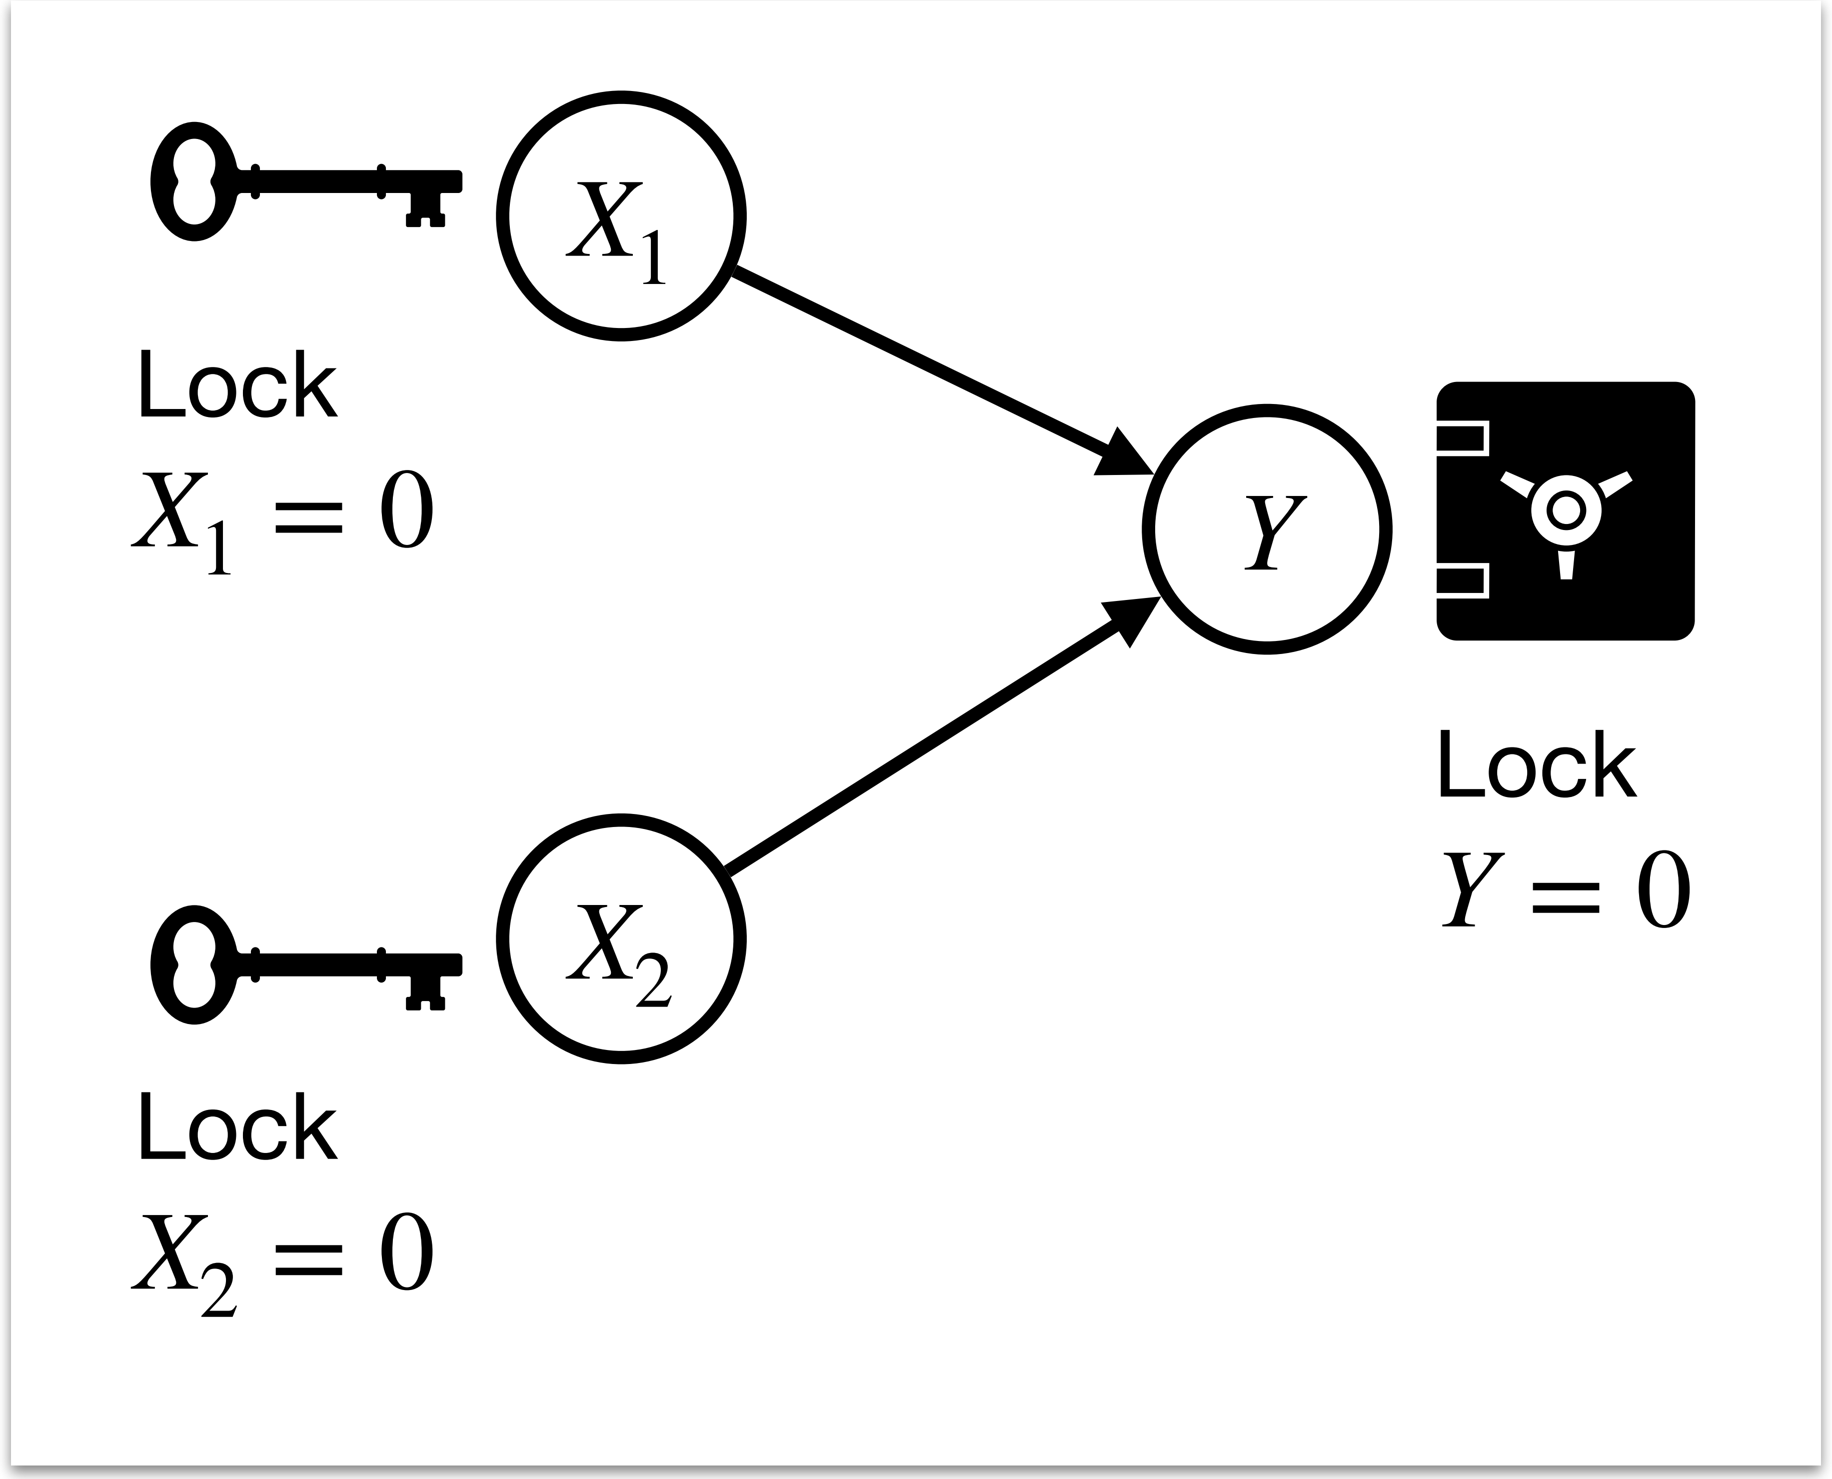

Example of the SCM: Encoding the DGP

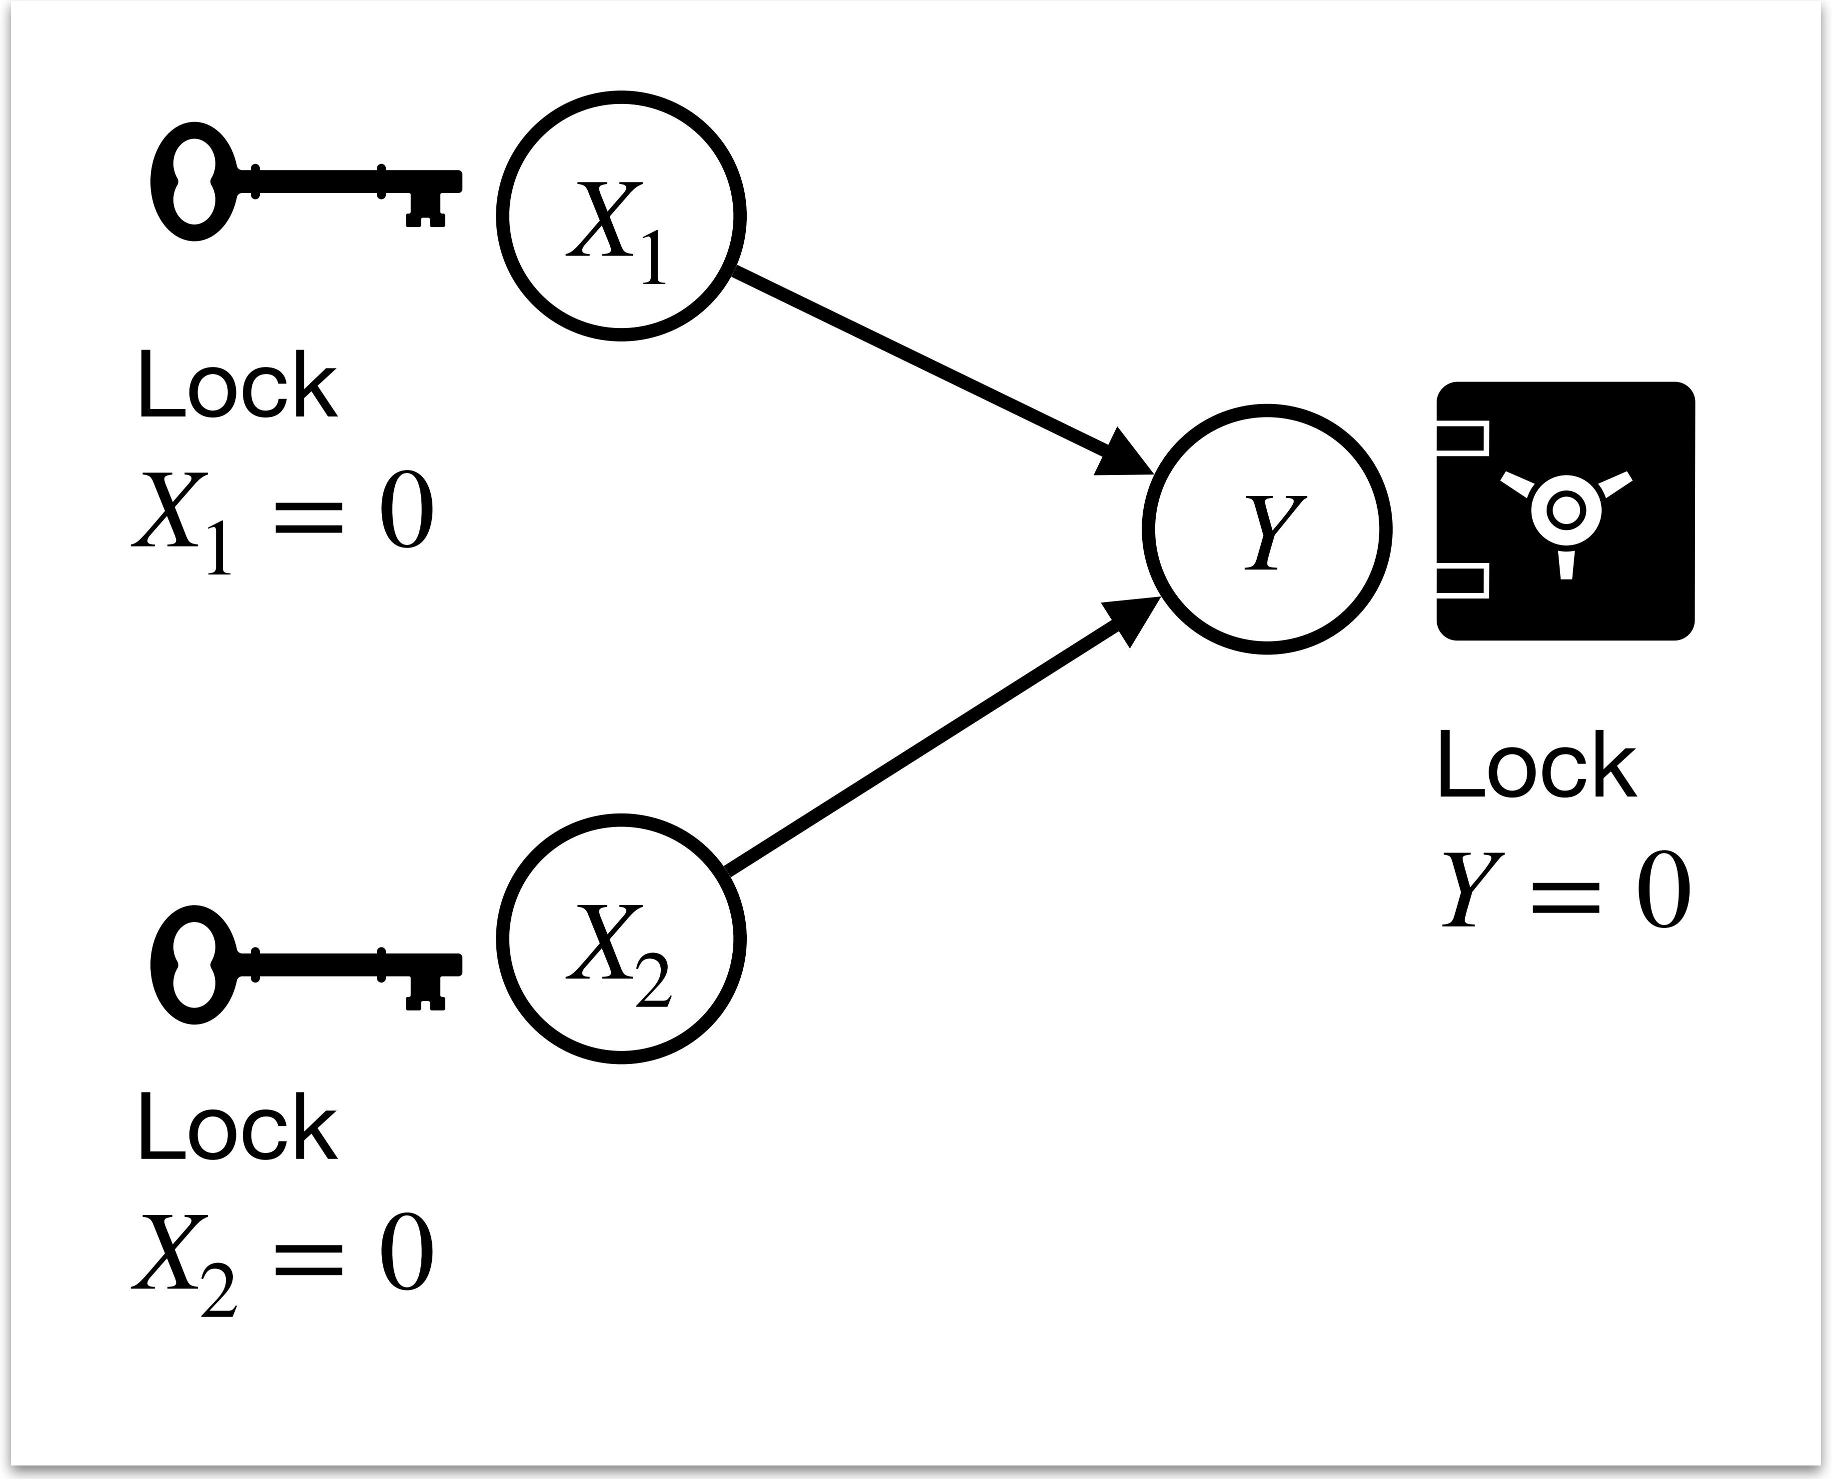

Original SCM

$$X_1 \leftarrow f_{X_1}(U_{X_1})$$
$$X_2 \leftarrow f_{X_2}(U_{X_2})$$
$$Y \leftarrow f_Y(X_1, X_2, U_Y)$$

Example of the SCM: Encoding the "What-If $X = x$"

$do(X_2 = 1)$

Example of the SCM: Encoding the "What-If $X = x$"

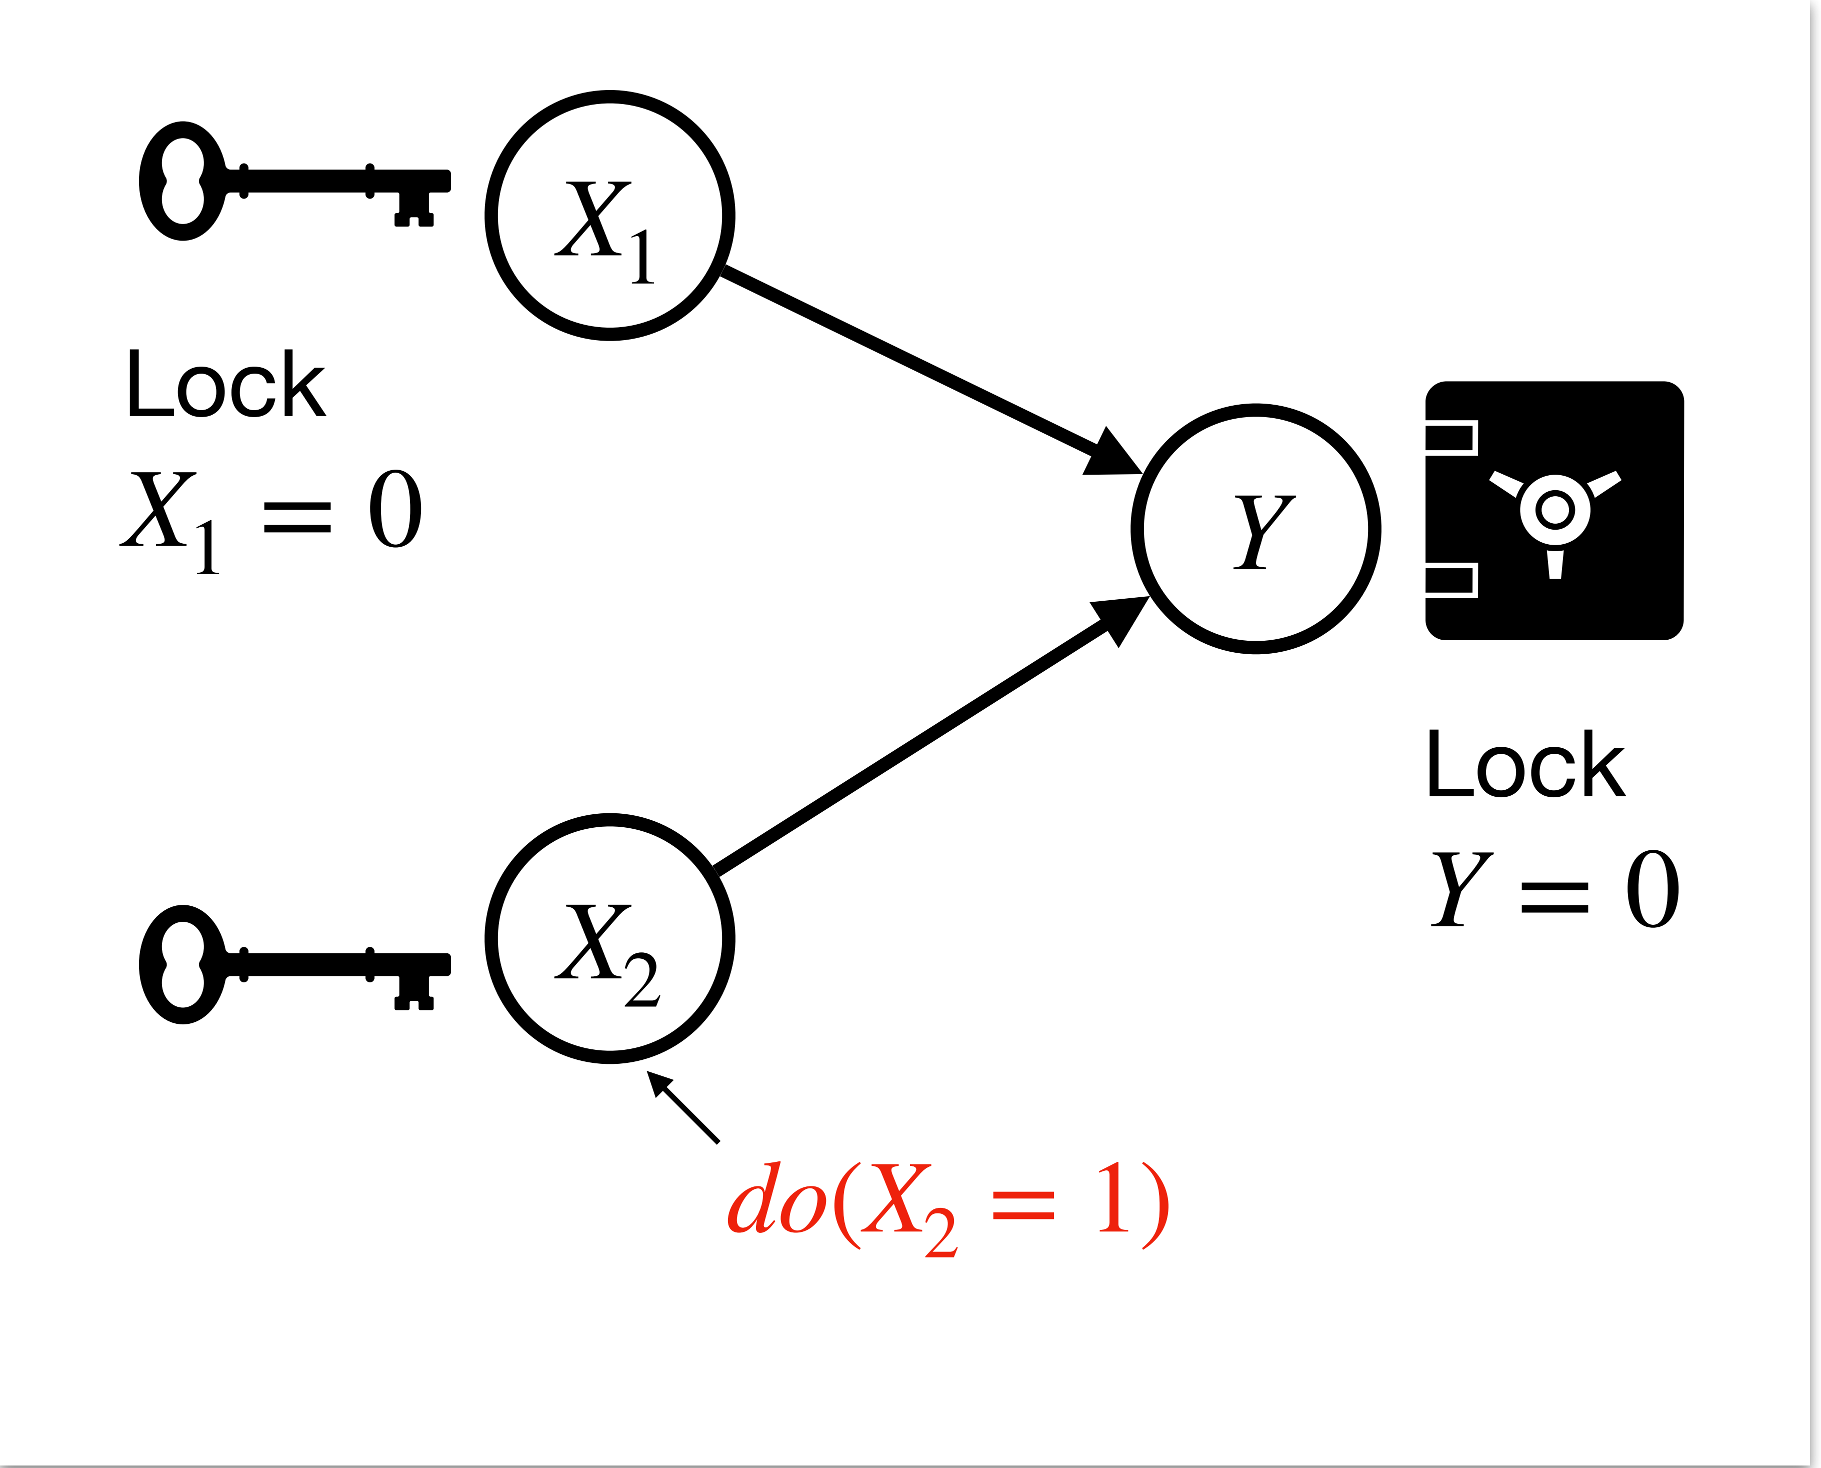

Example of the SCM: Encoding the "What-If $X = x$"

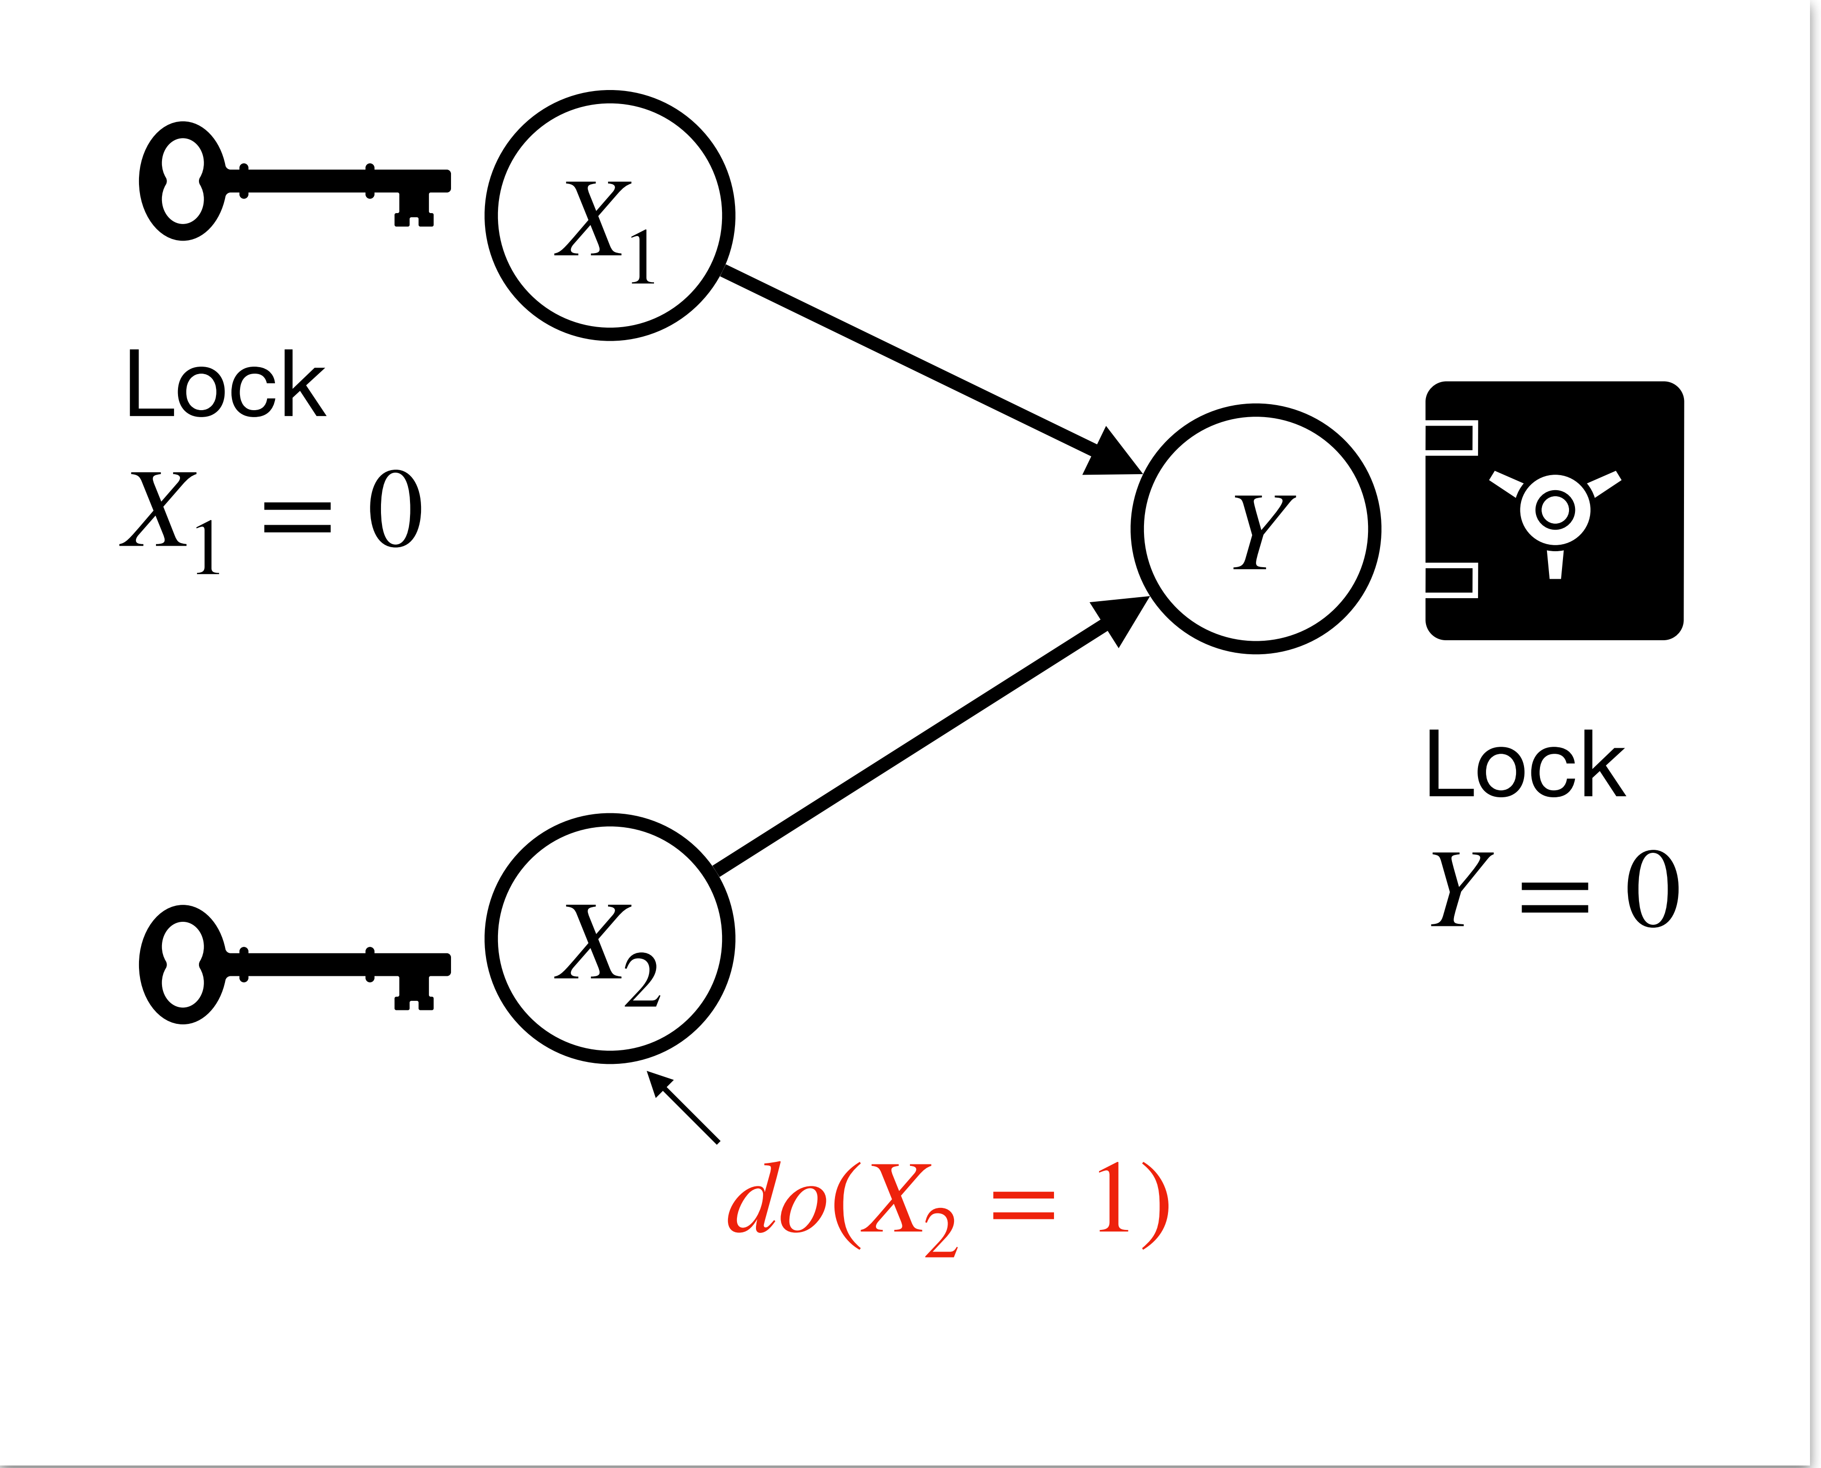

"What-If" $X_2 = 1$

$$X_1 \leftarrow f_{X_1}(U_{X_1})$$
$$X_2 \leftarrow 1$$
$$Y \leftarrow f_Y(X_1, X_2 = 1, U_Y)$$

Submodel: SCMs Induced by Fixing

Submodel: SCMs Induced by Fixing

For the original SCM M, "What if X had been fixed to x" can be encoded by replacing the function $X \leftarrow f_X(\cdot)$ to $X = x$.

Submodel: SCMs Induced by Fixing

For the original SCM M, "What if X had been fixed to x" can be encoded by replacing the function $X \leftarrow f_X(\cdot)$ to $X = x$.

Submodel of the SCM: The SCM after fixing $X = x$ is called the "submodel of the SCM" and denoted $M_{X=x}$.

Submodel: SCMs Induced by Fixing

For the original SCM M, “What if X had been fixed to x” can be encoded by replacing the function $X \leftarrow f_X(\cdot)$ to $X = x$.

Submodel of the SCM: The SCM after fixing $X = x$ is called the “submodel of the SCM” and denoted $M_{X=x}$.

Original SCM M

$$X_1 \leftarrow f_{X_1}(U_{X_1})$$
$$X_2 \leftarrow f_{X_2}(U_{X_2})$$
$$Y \leftarrow f_Y(X_1, X_2, U_Y)$$

Submodel: SCMs Induced by Fixing

For the original SCM M, "What if X had been fixed to x" can be encoded by replacing the function $X \leftarrow f_X(\cdot)$ to $X = x$.

Submodel of the SCM: The SCM after fixing $X = x$ is called the "submodel of the SCM" and denoted $M_{X=x}$.

Original SCM M

$$X_1 \leftarrow f_{X_1}(U_{X_1})$$
$$X_2 \leftarrow f_{X_2}(U_{X_2})$$
$$Y \leftarrow f_Y(X_1, X_2, U_Y)$$

Submodel $M_{X_2=1}$

$$X_1 \leftarrow f_{X_1}(U_{X_1})$$
$$X_2 \leftarrow 1$$
$$Y \leftarrow f_Y(X_1, X_2 = 1, U_Y)$$

Counterfactuals from Submodels

Counterfactuals from Submodels

The counterfactual $Y(X = x)$ can be generated by the submodel.

Counterfactuals from Submodels

The counterfactual $Y(X = x)$ can be generated by the submodel.

Counterfactual w.r.t SCM

Counterfactuals from Submodels

The counterfactual $Y(X = x)$ can be generated by the submodel.

Counterfactual w.r.t SCM

- $Y(X = x)$ is Y when values of X is set to x in their DGP (or population).

Counterfactuals from Submodels

The counterfactual $Y(X = x)$ can be generated by the submodel.

Counterfactual w.r.t SCM

- $Y(X = x)$ is Y when values of X is set to x in their DGP (or population).
- $Y(X = x)$ is Y in the submodel $M_{X=x}$.

Power of the SCM: SCM Can Reveal True Causality

Power of the SCM: SCM Can Reveal True Causality

In the previous example, we showed that the true causality of X_1 on Y was identified after fixing other variables $X_2 = 1$ in the DGP.

Power of the SCM: SCM Can Reveal True Causality

In the previous example, we showed that the true causality of X_1 on Y was identified after fixing other variables $X_2 = 1$ in the DGP.

Indeed, the state-of-the-art notion of causality is a counterfactual with fixing other variables, i.e., the counterfactual theories taking account of the relation of variables.

Power of the SCM: SCM Can Reveal True Causality

In the previous example, we showed that the true causality of X_1 on Y was identified after fixing other variables $X_2 = 1$ in the DGP.

Indeed, the state-of-the-art notion of causality is a counterfactual with fixing other variables, i.e., the counterfactual theories taking account of the relation of variables.

(Actual) Causality [Halpern and Pearl]

X is a cause of Y, if, for some $W \subseteq \mathbf{V}$, $Y(X = x) \neq Y(X = x')$ under some intervention $W = w$.

Power of the SCM: SCM Can Reveal True Causality

Actual Causality [Halpern and Pearl]

X is a cause of Y, if, for some $W \subseteq \mathbf{V}$, $Y(X = x) \neq Y(X = x')$ under some intervention $W = w$.

Power of the SCM: SCM Can Reveal True Causality

Actual Causality [Halpern and Pearl]

X is a cause of Y, if, for some $W \subseteq \mathbf{V}$, $Y(X=x) \neq Y(X=x')$ under some intervention $W = w$.

The formal definition of the actual causality is written w.r.t. the SCM.

Power of the SCM: SCM Can Reveal True Causality

Actual Causality [Halpern and Pearl]

X is a cause of Y, if, for some $W \subseteq \mathbf{V}$, $Y(X=x) \neq Y(X=x')$ under some intervention $W = w$.

The formal definition of the actual causality is written w.r.t. the SCM.

=> The SCM is a formal language that can describe the counterfactuals taking account of the relation of variables.

Power of the SCM: Unified Language

Power of the SCM: Unified Language

We just stated that the SCM is a language that can express the counterfactuals and take account of the DGP.

Power of the SCM: Unified Language

We just stated that the SCM is a language that can express the counterfactuals and take account of the DGP.

Indeed, the SCM is the **only** language that can express the counterfactuals and take account of the DGP.

Power of the SCM: Unified Language

We just stated that the SCM is a language that can express the counterfactuals and take account of the DGP.

Indeed, the SCM is the **only** language that can express the counterfactuals and take account of the DGP.

Any DGP of counterfactuals such that fixing $X = x$ (i.e., "what-if $X = x$") induces the counterfactual $W(x)$ of W (where W is any variable), this DGP is equivalent to the SCM.

Power of the SCM: Unified Language

We just stated that the SCM is a language that can express the counterfactuals and take account of the DGP.

Indeed, the SCM is the **only** language that can express the counterfactuals and take account of the DGP.

Any DGP of counterfactuals such that fixing $X = x$ (i.e., “what-if $X = x$”) induces the counterfactual $W(x)$ of W (where W is any variable), this DGP is equivalent to the SCM.

Example: In the hypothetical population where all patients in the population took the drug ($X = 1$). Suppose we measure patients’ blood pressure (W). Then, in this population, $W = W(X = 1)$.

Power of the SCM: Unified Language

Power of the SCM: Unified Language

Any DGP of counterfactuals such that fixing $X = x$ (i.e., “what-if $X = x$”) induces the counterfactual $W(x)$ of W (where W is any variable), this DGP is equivalent to the SCM.

Power of the SCM: Unified Language

Any DGP of counterfactuals such that fixing $X = x$ (i.e., “what-if $X = x$”) induces the counterfactual $W(x)$ of W (where W is any variable), this DGP is equivalent to the SCM.

Many materials teach two types of causal inference frameworks: (1) SCM and (2) Potential Outcome (PO) Framework. We now know that

Power of the SCM: Unified Language

Any DGP of counterfactuals such that fixing $X = x$ (i.e., “what-if $X = x$”) induces the counterfactual $W(x)$ of W (where W is any variable), this DGP is equivalent to the SCM.

Many materials teach two types of causal inference frameworks: (1) SCM and (2) Potential Outcome (PO) Framework. We now know that

1. Any causal inference theories taking account of the DGP are equivalent to the SCM.

Power of the SCM: Unified Language

Any DGP of counterfactuals such that fixing $X = x$ (i.e., "what-if $X = x$") induces the counterfactual $W(x)$ of W (where W is any variable), this DGP is equivalent to the SCM.

Many materials teach two types of causal inference frameworks: (1) SCM and (2) Potential Outcome (PO) Framework. We now know that

1. Any causal inference theories taking account of the DGP are equivalent to the SCM.

2. Therefore, the SCM and the PO frameworks are equivalent.

Connecting Human's Cognition and AI/ML

Hierarchical Layer	Quantity	Task	Question

Connecting Human's Cognition and AI/ML

Hierarchical Layer	Quantity	Task	Question
L1 (Association)	$P(y \mid x)$	• Classification • Regression	What does the symptom tells about my headache?

Connecting Human's Cognition and AI/ML

Hierarchical Layer	Quantity	Task	Question
L1 (Association)	$P(y \mid x)$	• Classification • Regression	What does the symptom tells about my headache?
L2 (Intervention)	$P(Y(X = x) = y)$	• Reinforcement Learning • Randomized Trial	What if I took the aspirin, will my headache be cured?

Connecting Human's Cognition and AI/ML

Hierarchical Layer	Quantity	Task	Question
L1 (Association)	$P(y \mid x)$	• Classification • Regression	What does the symptom tells about my headache?
L2 (Intervention)	$P(Y(X=x)=y)$	• Reinforcement Learning • Randomized Trial	What if I took the aspirin, will my headache be cured?
L3 (Counterfactual)	$P(Y(X=x)=y \mid x')$	• Counterfactual Thinking	Given that I didn't take the aspirin and didn't get cured, what if I did?

Connecting Human's Cognition and AI/ML: Example

Original SCM M

$$X_1 \leftarrow f_{X_1}(U_{X_1})$$
$$X_2 \leftarrow f_{X_2}(U_{X_2})$$
$$Y \leftarrow f_Y(X_1, X_2, U_Y)$$

Submodel $M_{X_2=1}$

$$X_1 \leftarrow f_{X_1}(U_{X_1})$$
$$X_2 \leftarrow 1$$
$$Y \leftarrow f_Y(X_1, X_2 = 1, U_Y)$$

Connecting Human's Cognition and AI/ML: Example

Original SCM M

$X_1 \leftarrow f_{X_1}(U_{X_1})$

$X_2 \leftarrow f_{X_2}(U_{X_2})$

$Y \leftarrow f_Y(X_1, X_2, U_Y)$

L1 (Association)

$P(Y = y \mid X_2 = 0)$

Submodel $M_{X_2=1}$

$X_1 \leftarrow f_{X_1}(U_{X_1})$

$X_2 \leftarrow 1$

$Y \leftarrow f_Y(X_1, X_2 = 1, U_Y)$

Connecting Human's Cognition and AI/ML: Example

Original SCM M

$X_1 \leftarrow f_{X_1}(U_{X_1})$

$X_2 \leftarrow f_{X_2}(U_{X_2})$

$Y \leftarrow f_Y(X_1, X_2, U_Y)$

Submodel $M_{X_2=1}$

$X_1 \leftarrow f_{X_1}(U_{X_1})$

$X_2 \leftarrow 1$

$Y \leftarrow f_Y(X_1, X_2 = 1, U_Y)$

L1 (Association)

$P(Y = y \mid X_2 = 0)$

L2 (Intervention)

$P(Y(X_2 = 1) = y)$

Connecting Human's Cognition and AI/ML: Example

Original SCM M

$X_1 \leftarrow f_{X_1}(U_{X_1})$

$X_2 \leftarrow f_{X_2}(U_{X_2})$

$Y \leftarrow f_Y(X_1, X_2, U_Y)$

Submodel $M_{X_2=1}$

$X_1 \leftarrow f_{X_1}(U_{X_1})$

$X_2 \leftarrow 1$

$Y \leftarrow f_Y(X_1, X_2 = 1, U_Y)$

L1 (Association)

$P(Y = y \mid X_2 = 0)$

L2 (Intervention)

$P(Y(X_2 = 1) = y)$

L3 (Intervention)

$P(Y(X_2 = 1) = y \mid X_2 = 0)$

Big Picture in Causal Inference

Important Problems in Causal Inference

Big Picture for Causal Inference

Big Picture for Causal Inference

SCM

$Z \leftarrow f_Z(U_Z)$
$X \leftarrow f_X(Z, U_X)$
$Y \leftarrow f_Y(X, Z, U_Y)$

Big Picture for Causal Inference

SCM

$Z \leftarrow f_Z(U_Z)$
$X \leftarrow f_X(Z, U_X)$
$Y \leftarrow f_Y(X, Z, U_Y)$

Assumption on the SCM

Big Picture for Causal Inference

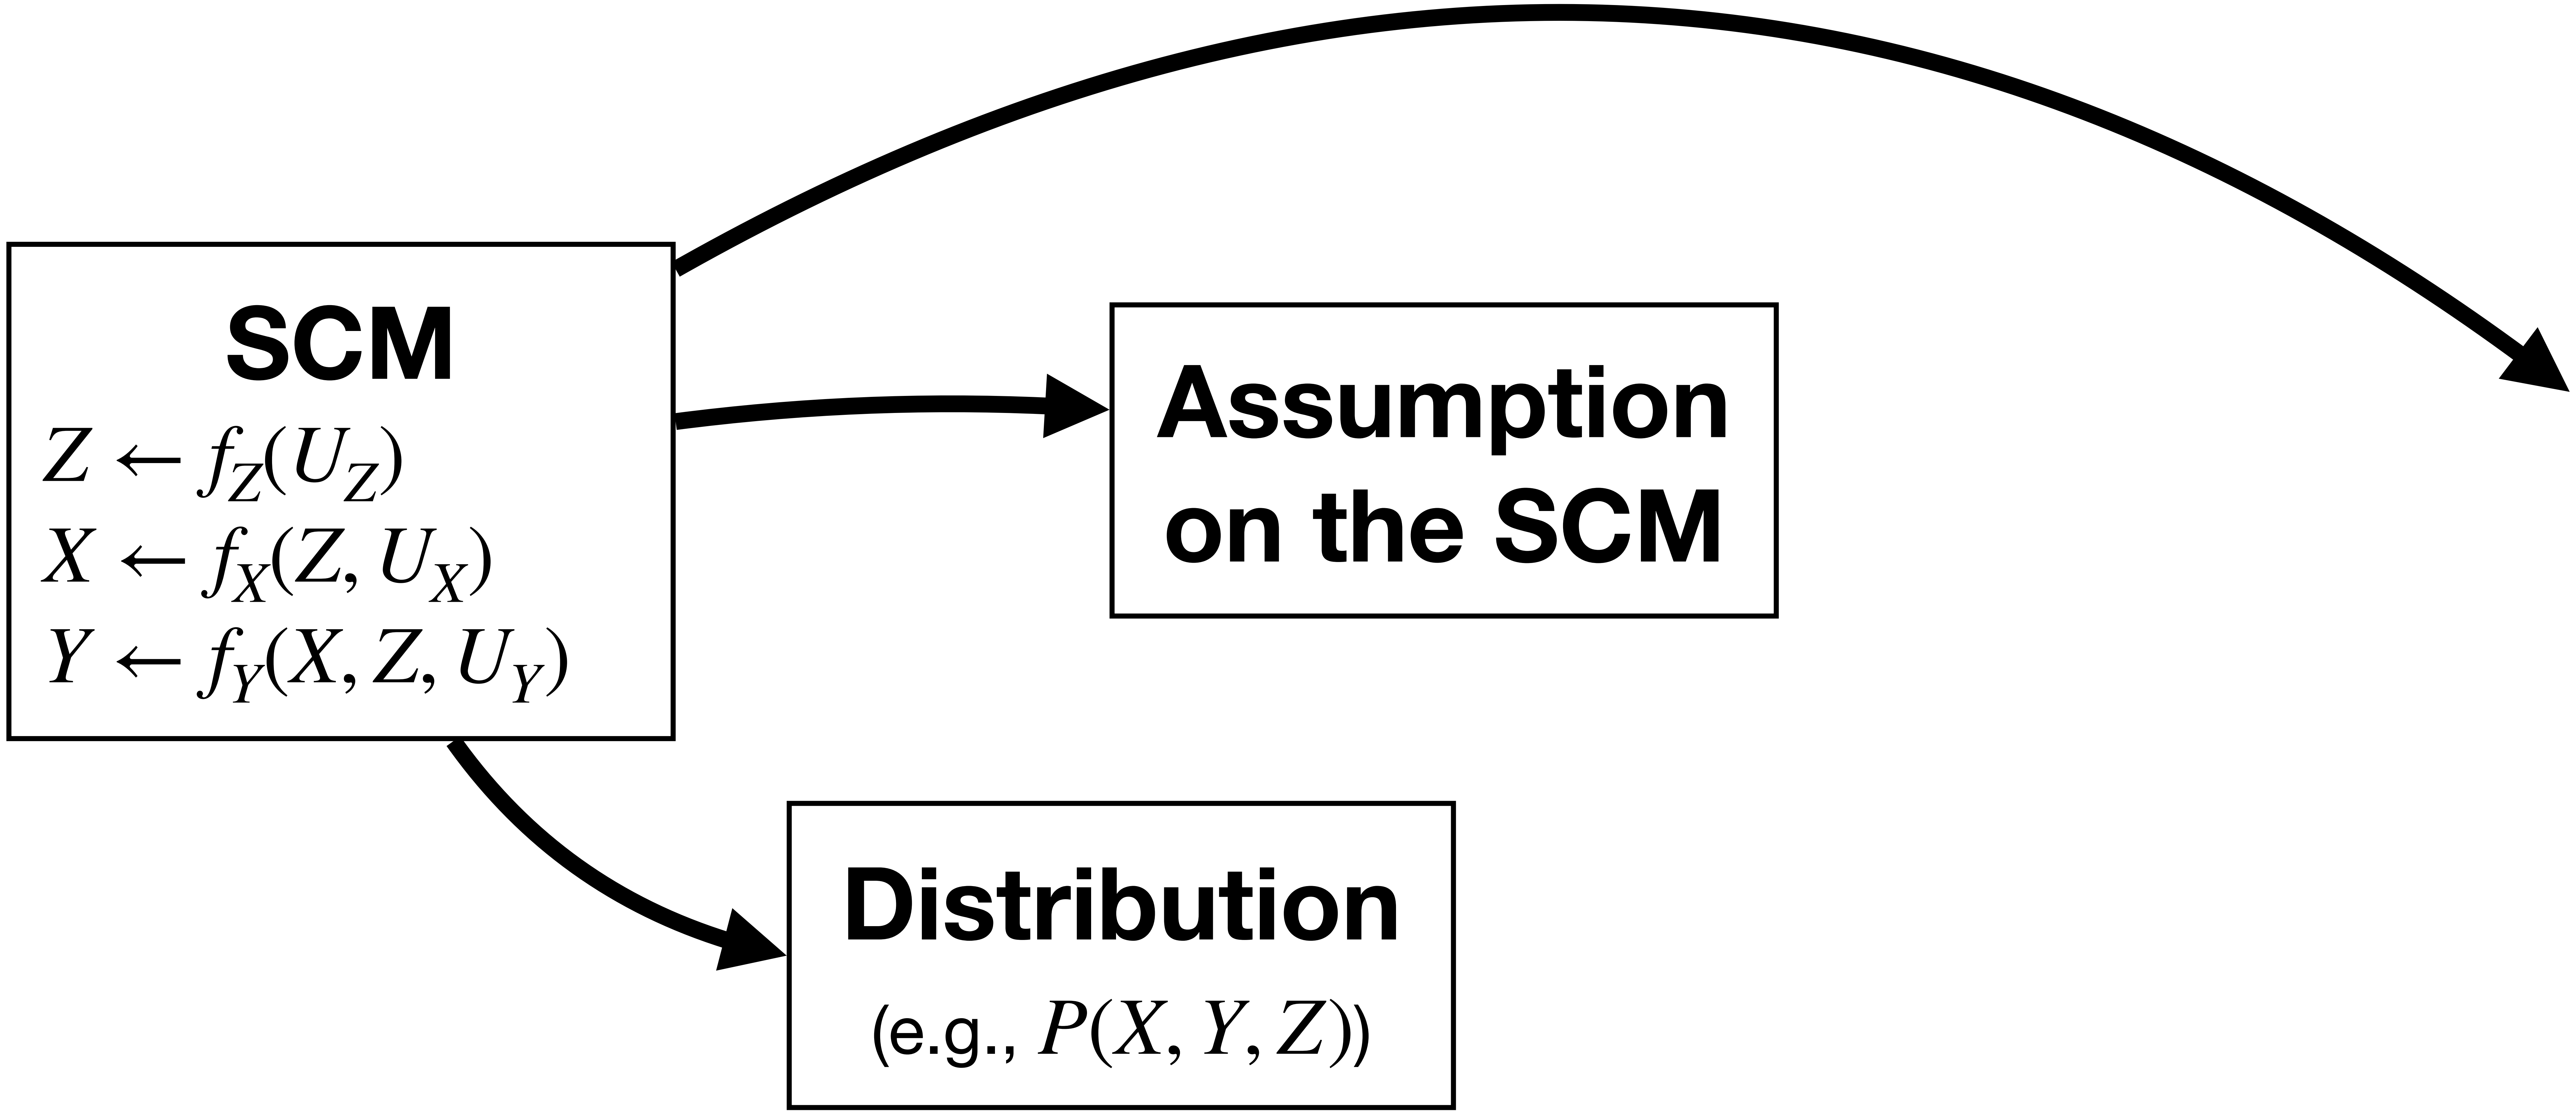

Big Picture for Causal Inference

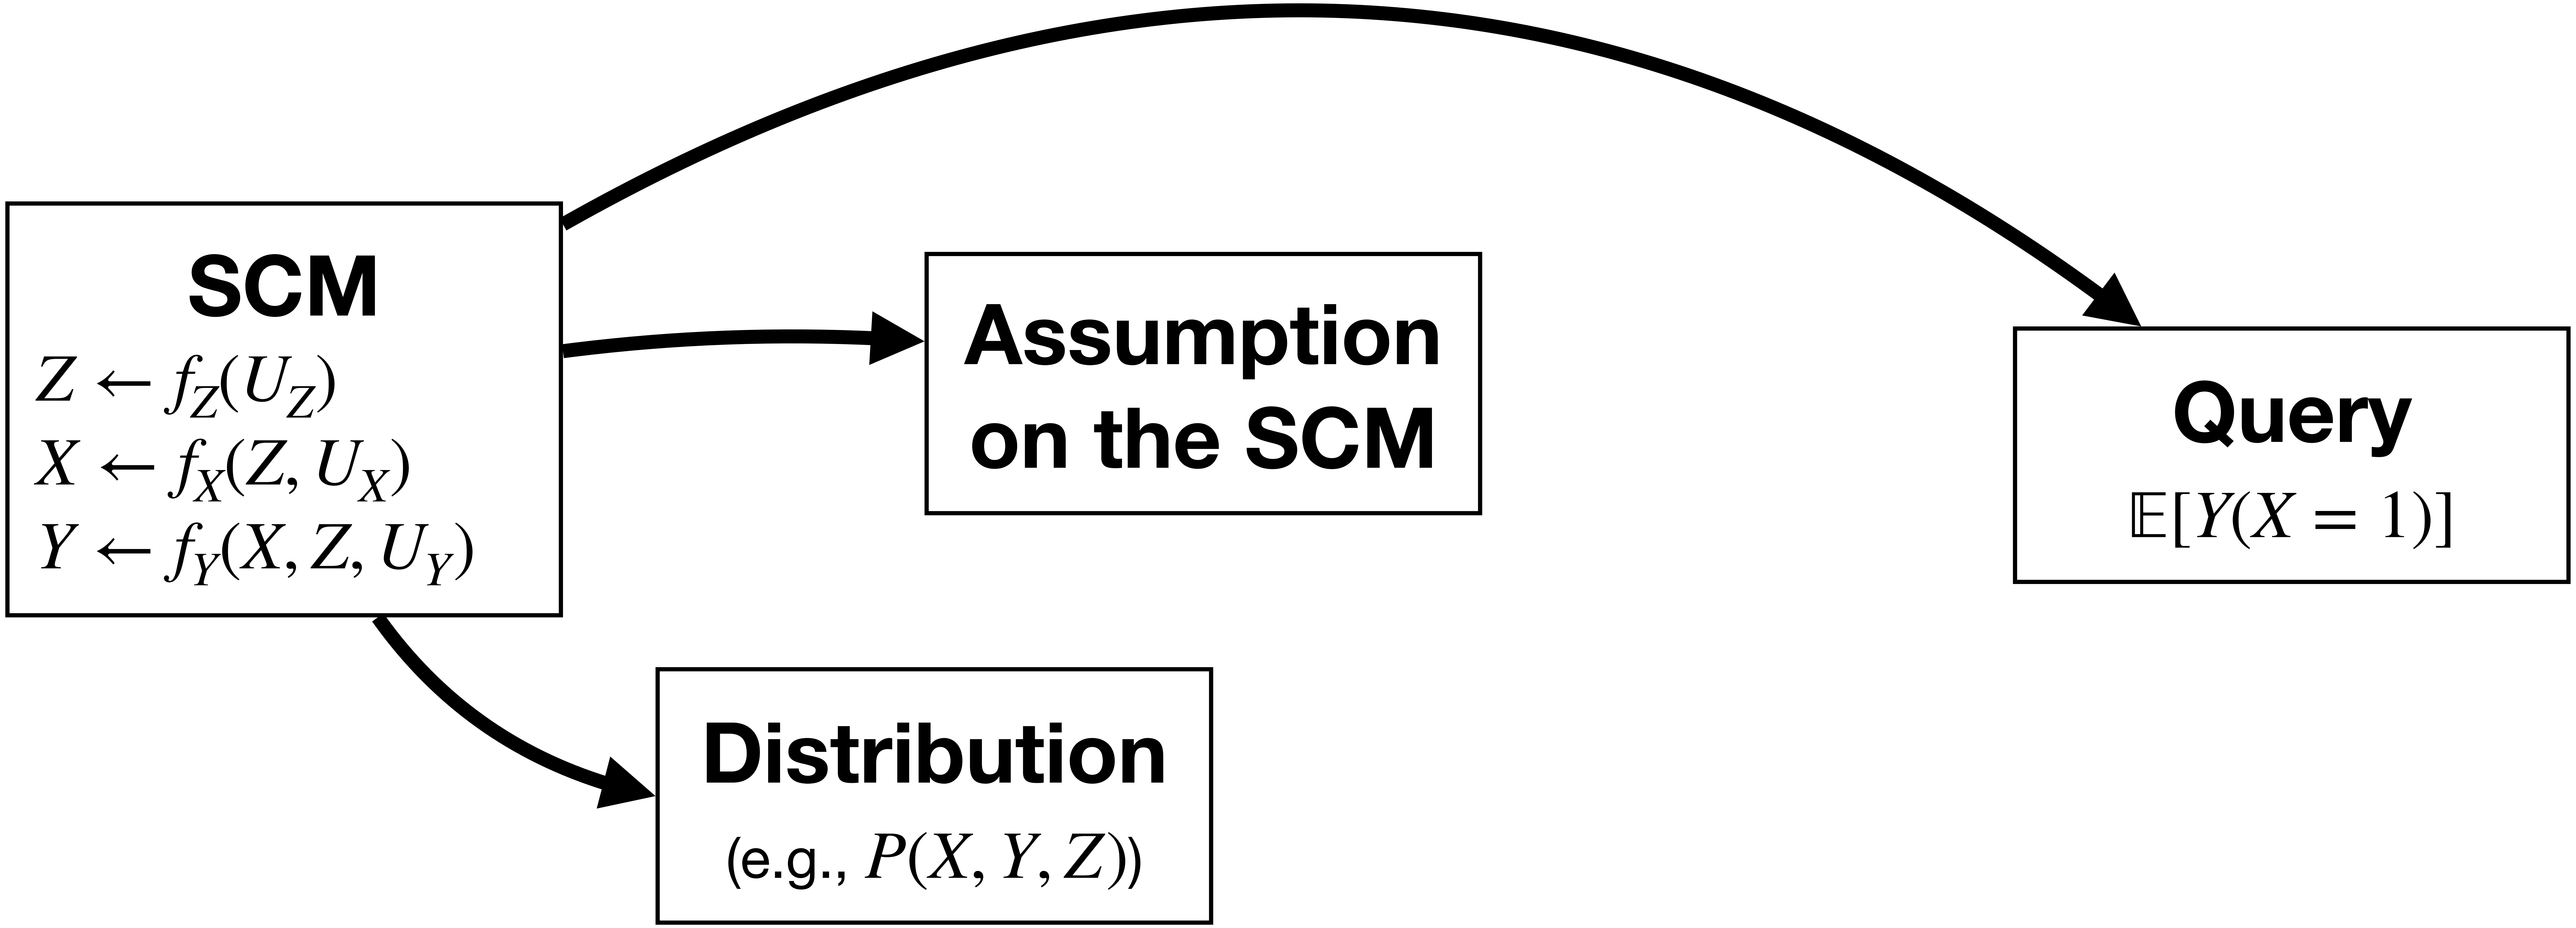

Big Picture for Causal Inference

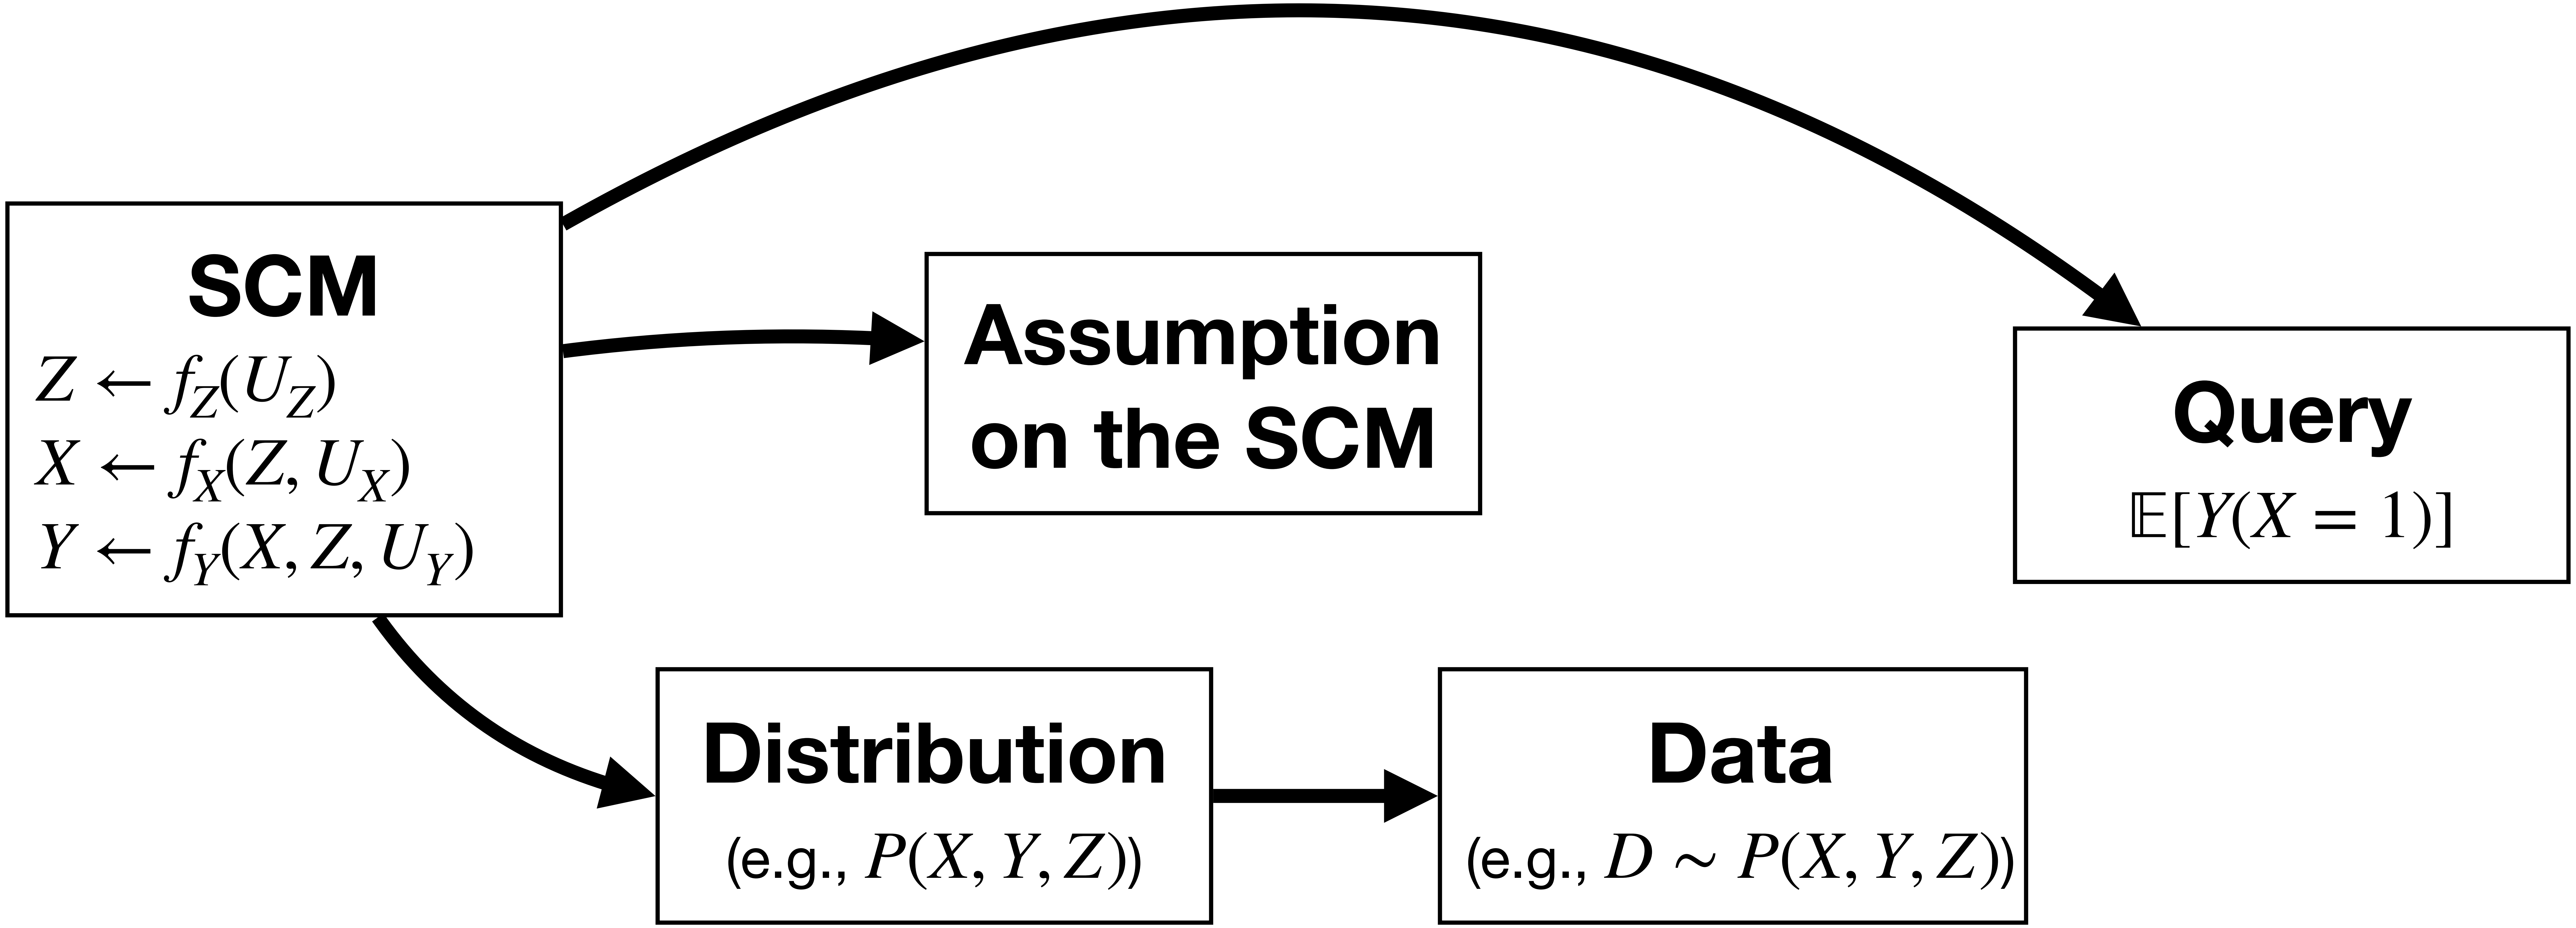

Big Picture for Causal Inference

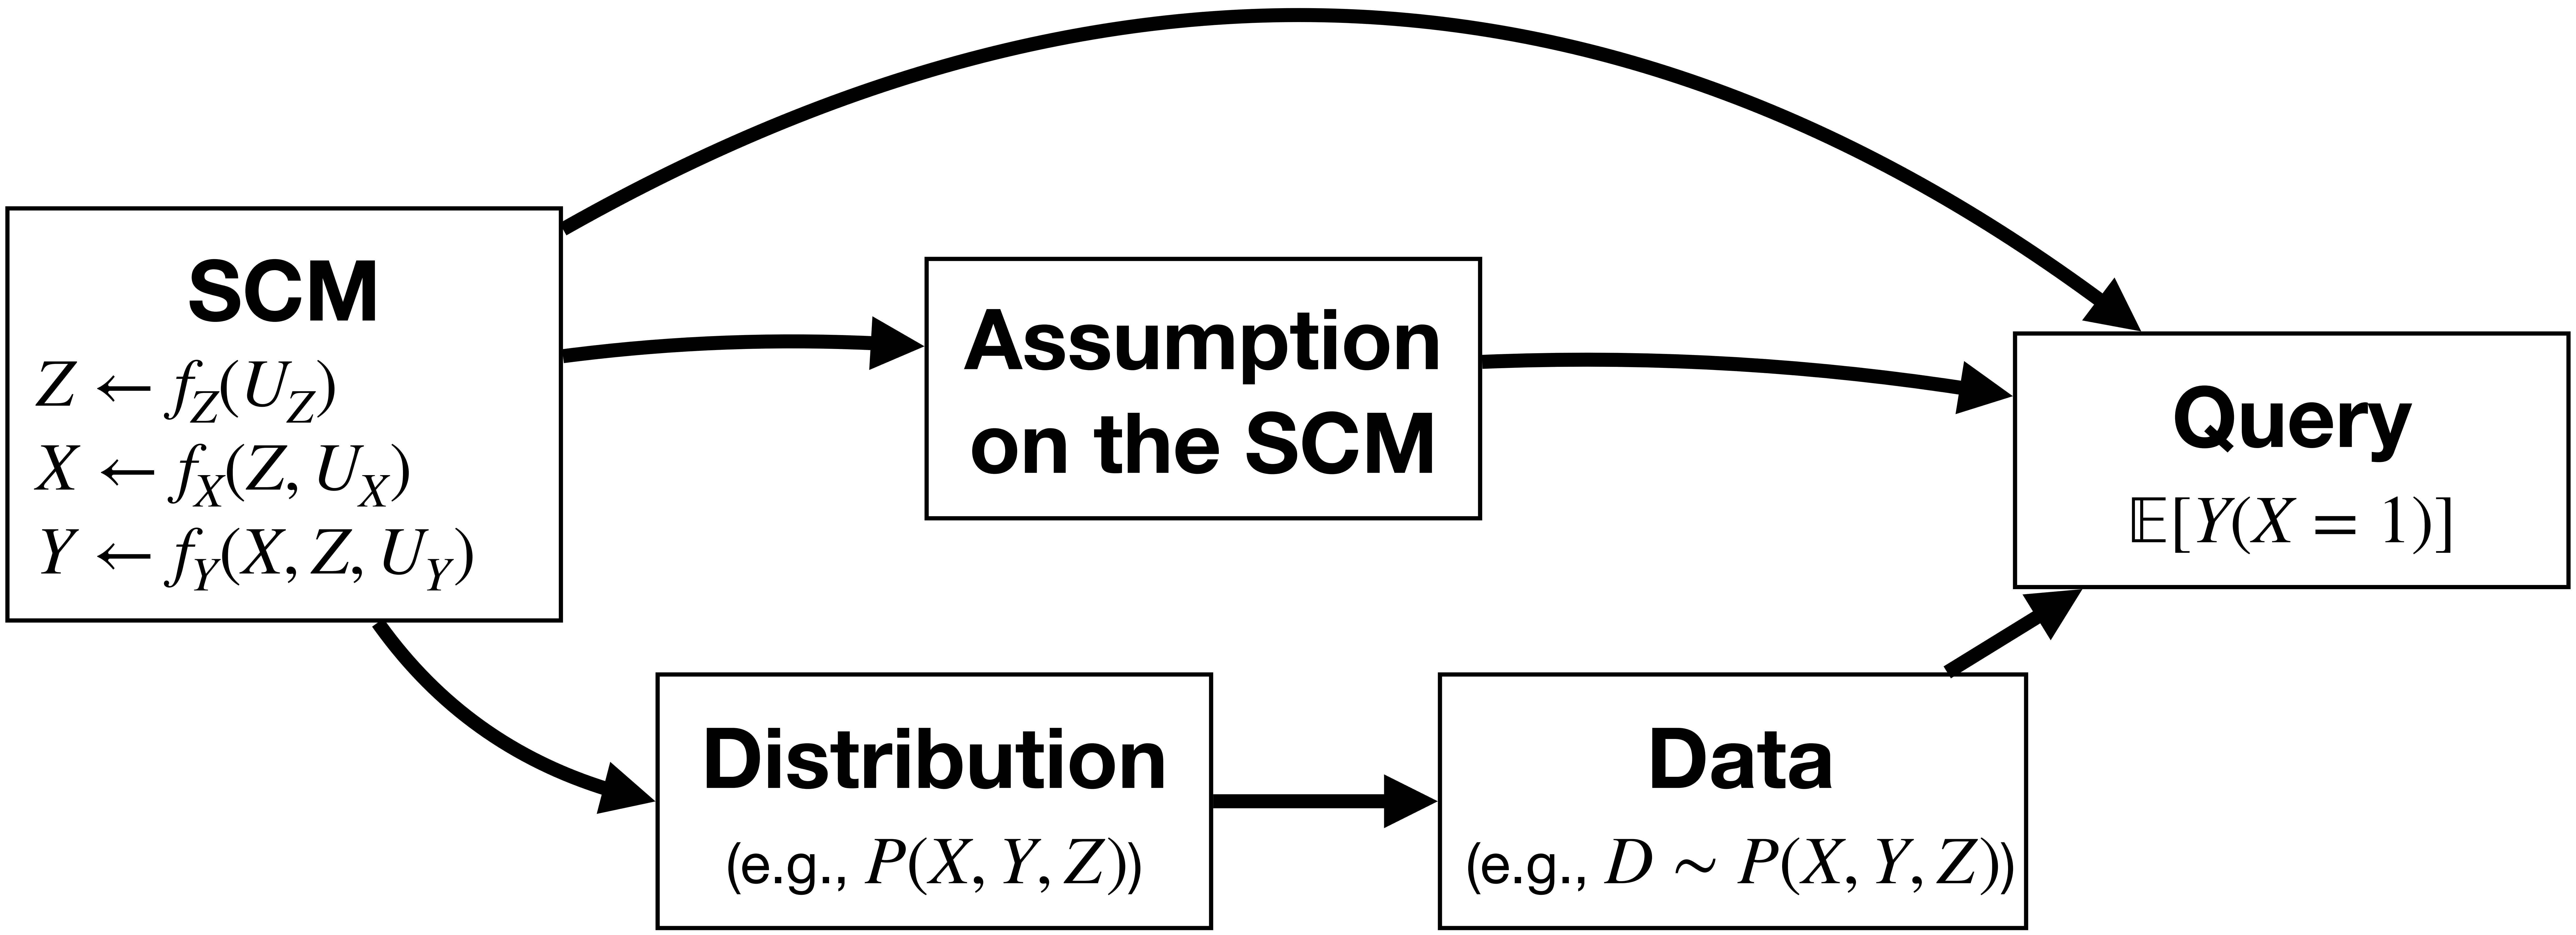

Importance of Assumptions

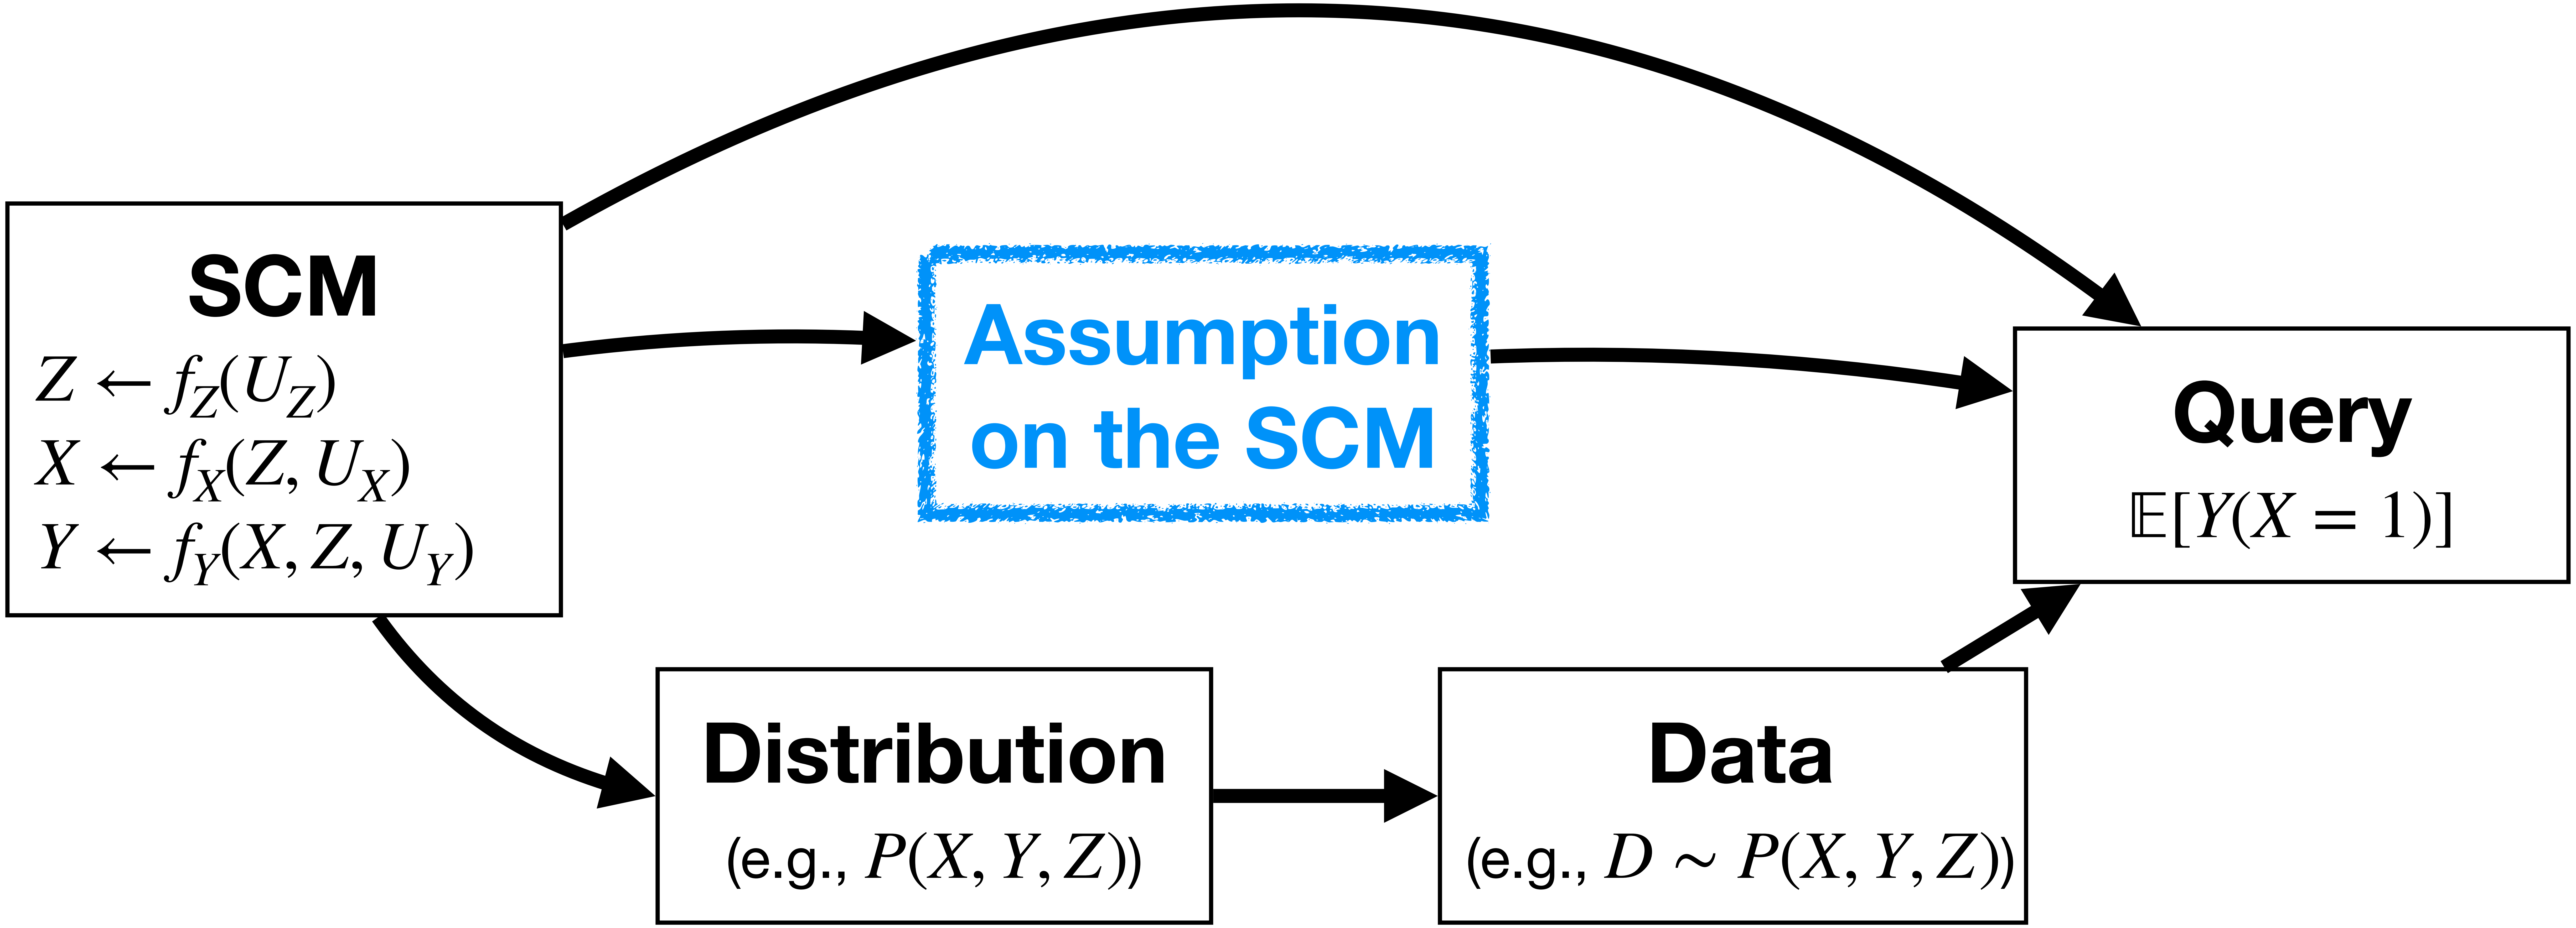

Importance of Assumptions: Causal Hierarchy Theorem (CHT)

Importance of Assumptions: Causal Hierarchy Theorem (CHT)

Why do we need assumptions on the SCM (or the DGP of counterfactuals) in the big picture?

Importance of Assumptions: Causal Hierarchy Theorem (CHT)

Why do we need assumptions on the SCM (or the DGP of counterfactuals) in the big picture?

Pearl's Causal Hierarchy Theorem [Bareinboim et al., 2020]

Importance of Assumptions: Causal Hierarchy Theorem (CHT)

Why do we need assumptions on the SCM (or the DGP of counterfactuals) in the big picture?

Pearl's Causal Hierarchy Theorem [Bareinboim et al., 2020]

- Causal inference is impossible without making any assumptions on the DGP of the counterfactuals (i.e., the SCM).

Importance of Assumptions: Causal Hierarchy Theorem (CHT)

Why do we need assumptions on the SCM (or the DGP of counterfactuals) in the big picture?

Pearl's Causal Hierarchy Theorem [Bareinboim et al., 2020]

- Causal inference is impossible without making any assumptions on the DGP of the counterfactuals (i.e., the SCM).
- Equivalently, given L_i's information (e.g., associational information L_1), the higher layer information (e.g., the causal information L_2) is not inferable without making any assumptions.

Recap:
Big Picture for Causal Inference

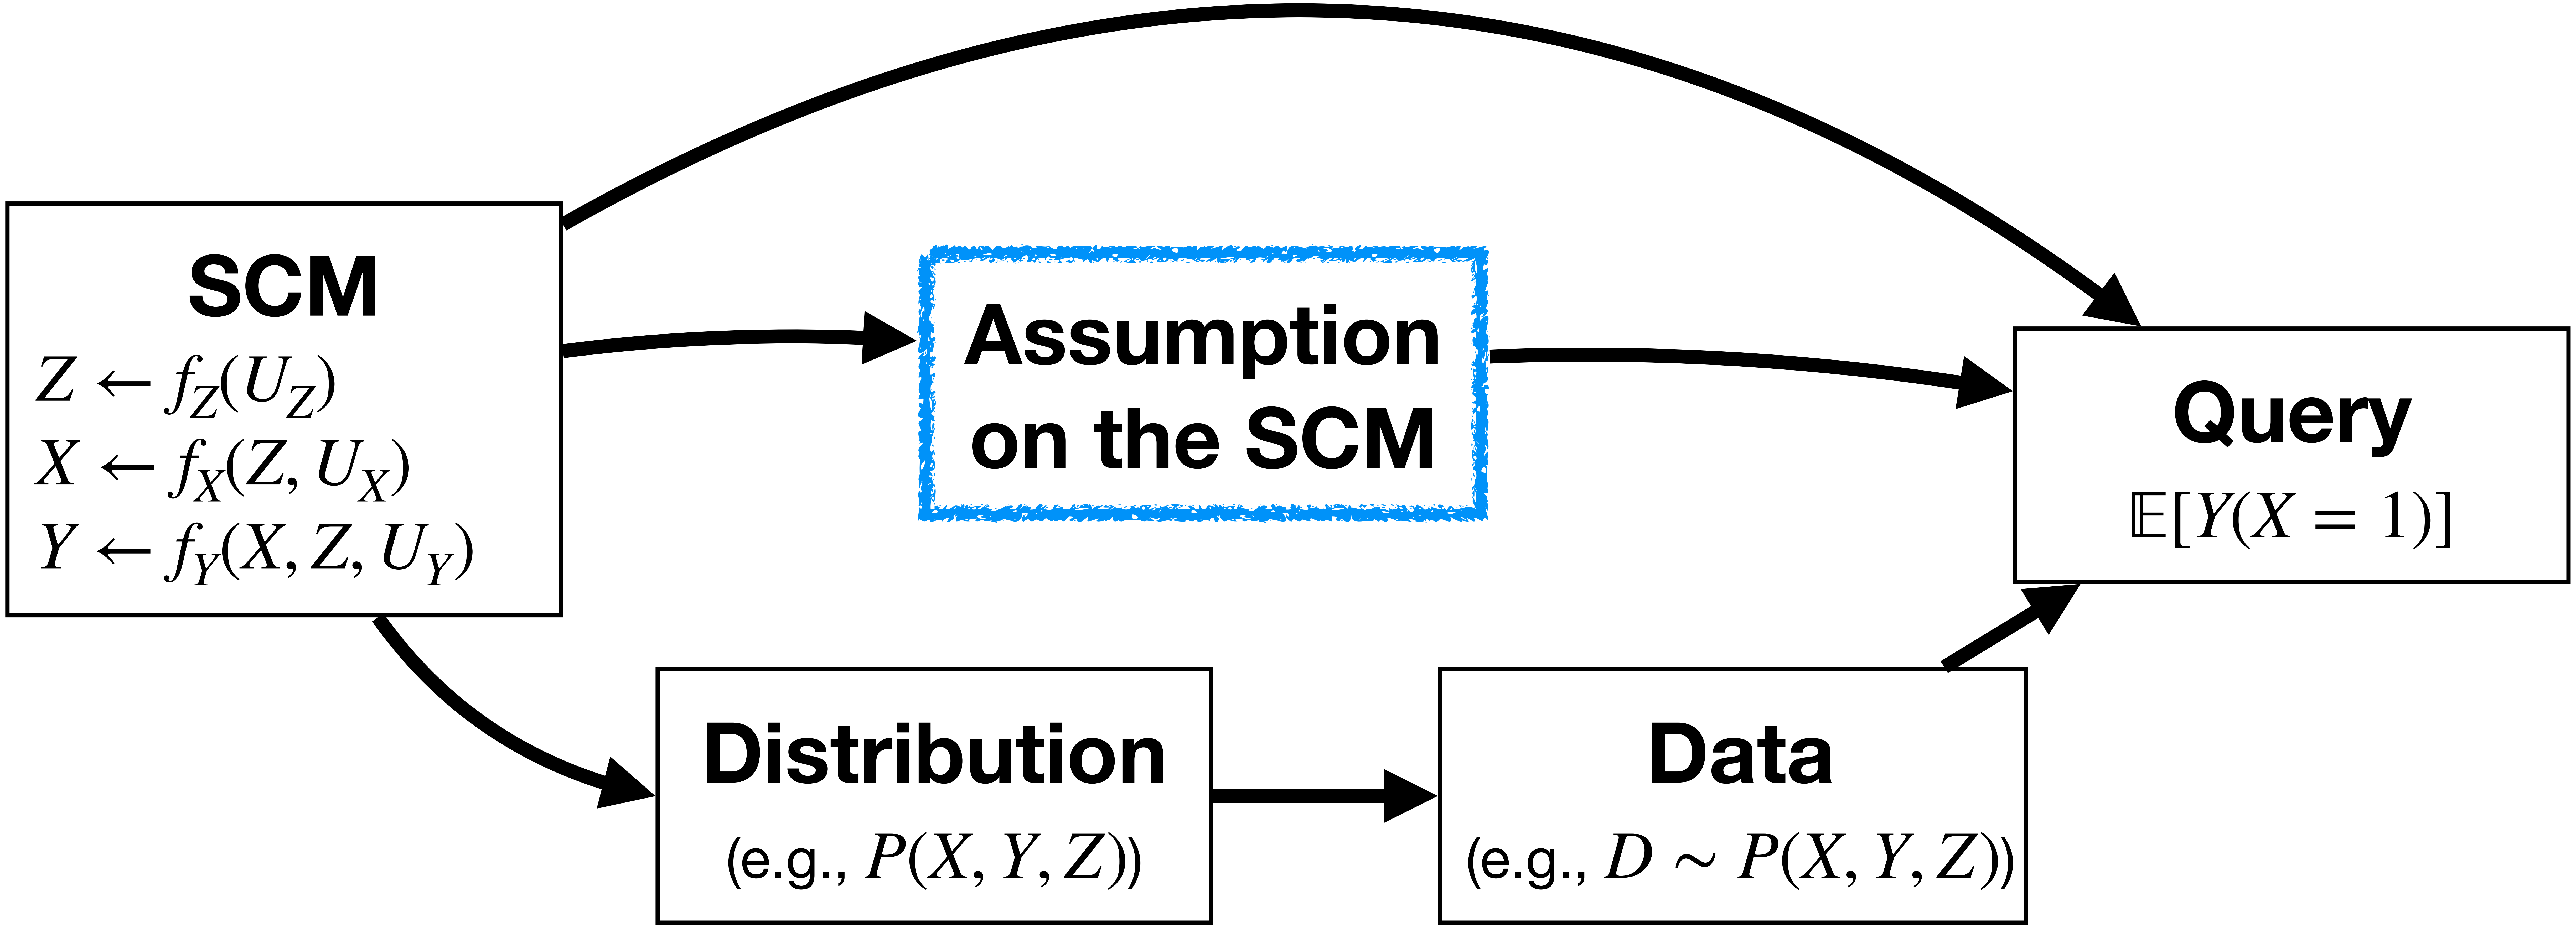

Big Picture for Causal Inference: Encoding Assumptions Thr. Graphs

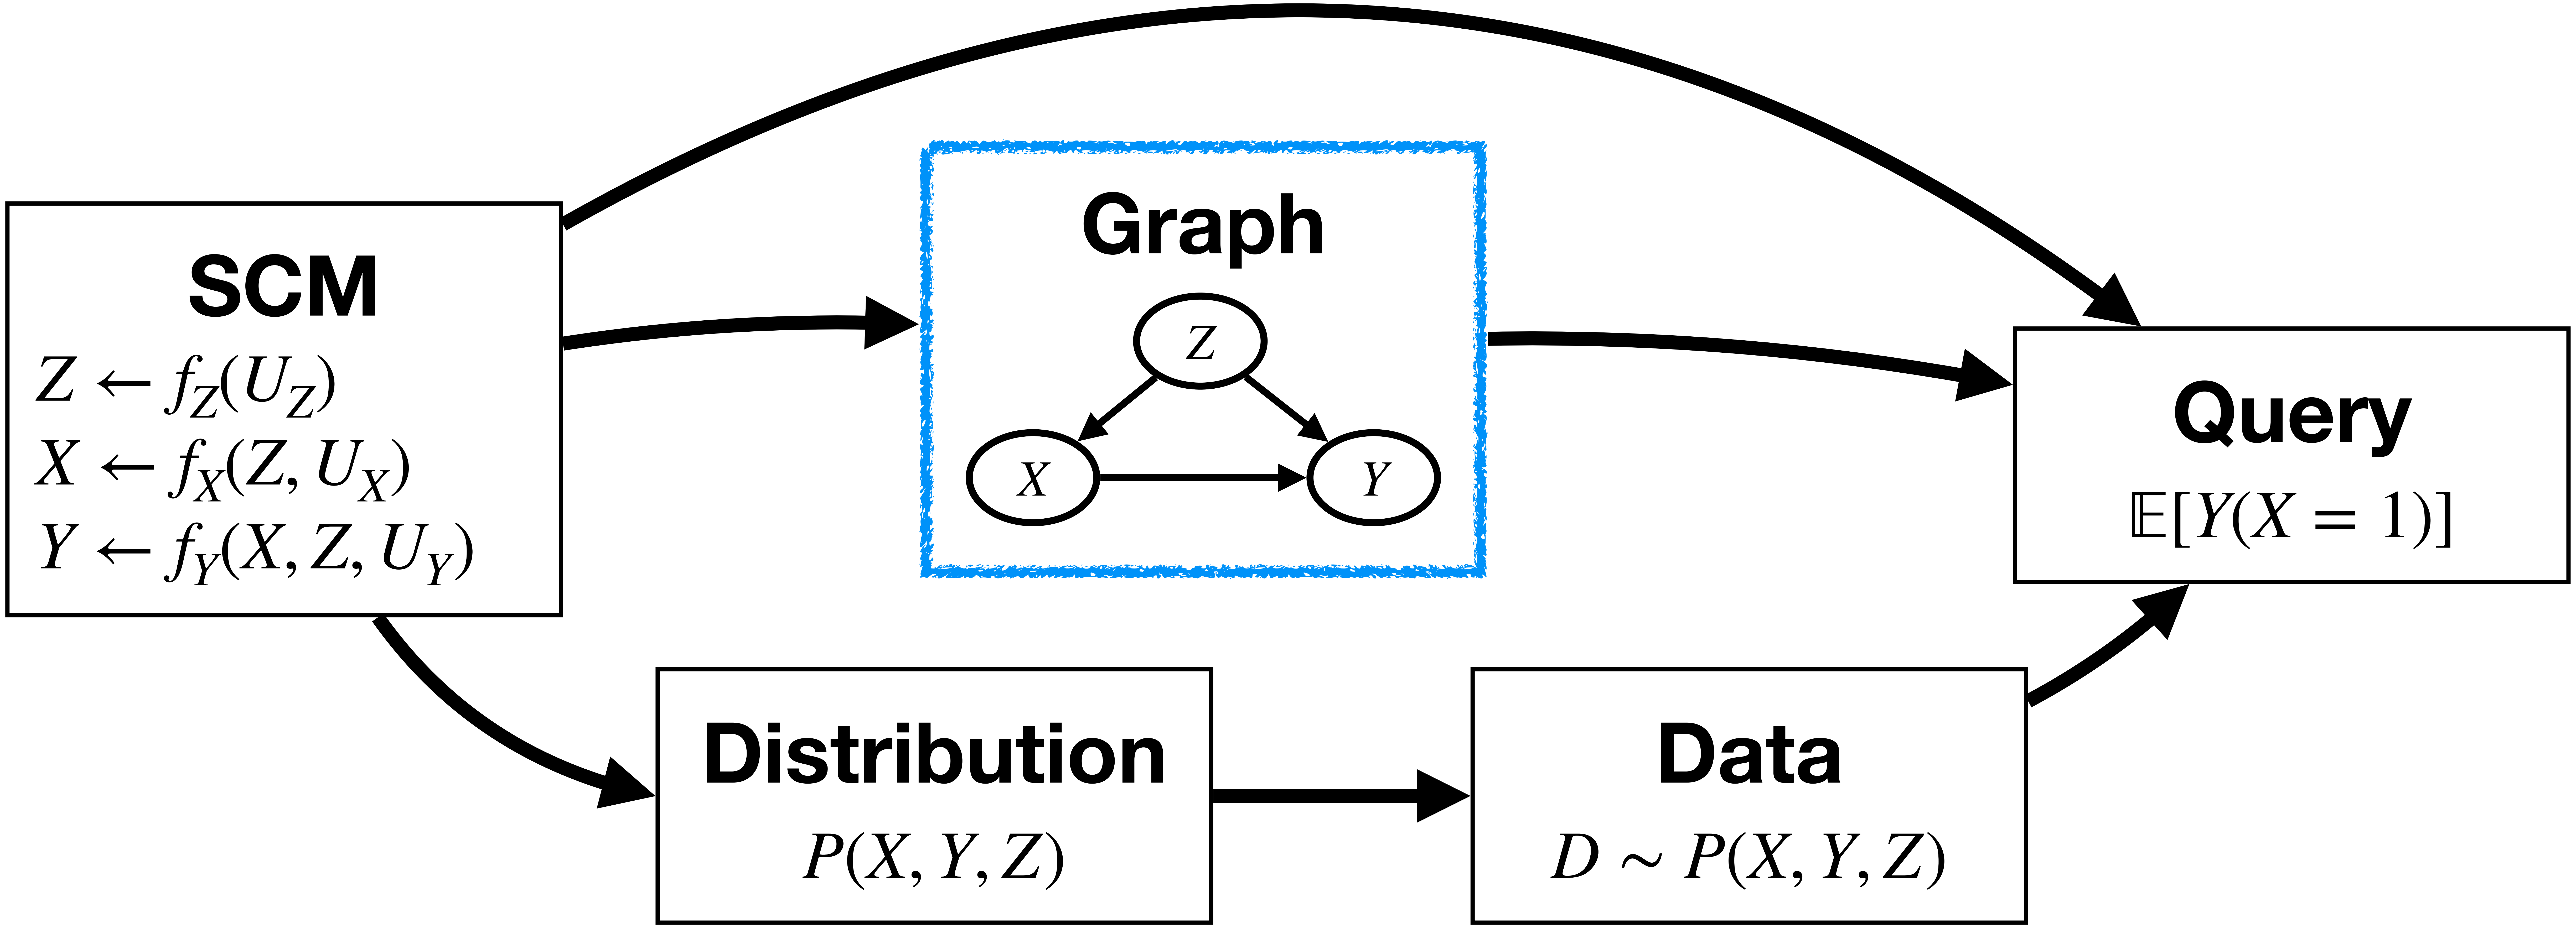

Big Picture for Causal Inference: Inaccessibility to SCMs

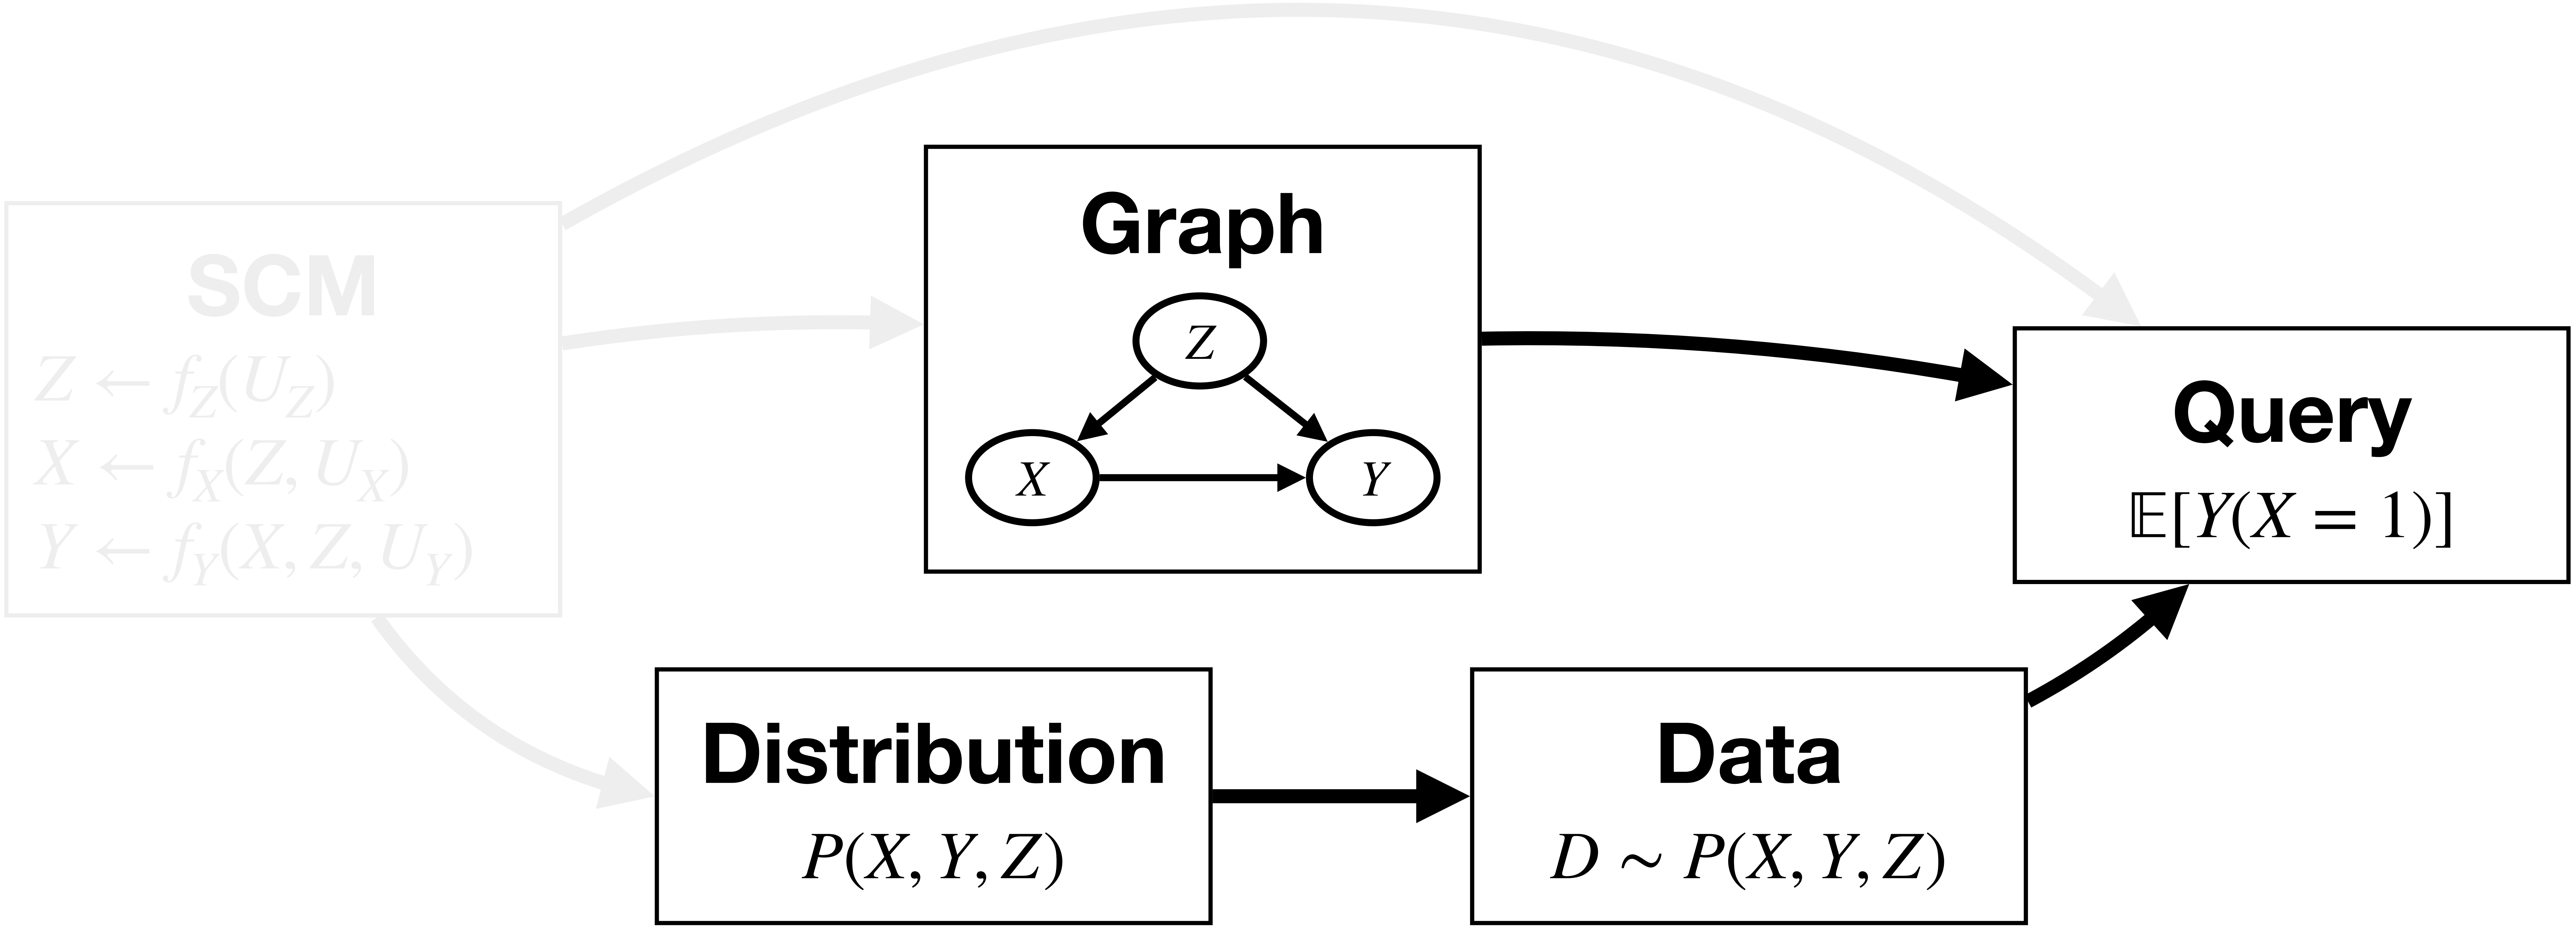

Conclusion

Conclusion

Conclusion

We studied the SCM, the unified causal inference framework.

Conclusion

We studied the SCM, the unified causal inference framework.

- SCM is a unique framework that can represent the DGP of counterfactuals.

Conclusion

We studied the SCM, the unified causal inference framework.

- SCM is a unique framework that can represent the DGP of counterfactuals.

We overviewed important data science problems in the SCM.

Conclusion

We studied the SCM, the unified causal inference framework.

- SCM is a unique framework that can represent the DGP of counterfactuals.

We overviewed important data science problems in the SCM.

- We overviewed important causal inference problems under the rubric of the SCM.

Conclusion

We studied the SCM, the unified causal inference framework.

- SCM is a unique framework that can represent the DGP of counterfactuals.

We overviewed important data science problems in the SCM.

- We overviewed important causal inference problems under the rubric of the SCM.
- We studied that practical data science problems where the DGP can be expressed as a SCM can be reduced to the causal inference problem.