Causal Inference

CMPUT 366: Intelligent Systems

Bar §3.4

Lecture Outline

1. Recap & Logistics
2. Causal Queries
3. Identifiability

Labs & Assignment #1

- Assignment #1 was due **Feb 4 (today)** before lecture
- Today's lab is from **5:00pm to 7:50pm** in **CAB 235**
    - Last-chance lab for late assignments
    - Not mandatory
    - Opportunity to get help from the TAs

Recap: Independence in a Belief Network

Belief Network Semantics:
Every node is **independent** of its **non-descendants**, **conditional only** on its **parents**

Patterns of dependence:

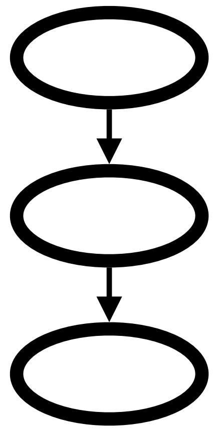

1. **Chain**: Ends are not marginally independent, but conditionally independent given middle

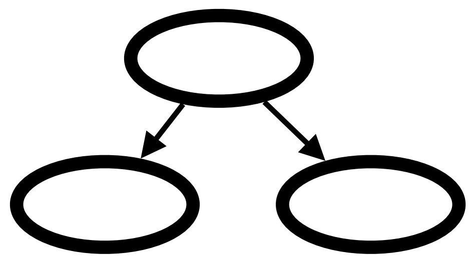

2. **Common ancestor**: descendants are not marginally independent, but conditionally independent given ancestor

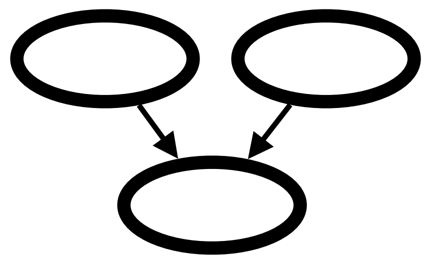

3. **Common descendant**: Ancestors are marginally independent, but not conditionally independent given descendant

Recap: Simpson's Paradox

- The joint distribution factors as
  P(G,D,R) = P(R | D, G) x P(D | G) x P(G)
- Per-gender queries seem **sensible**:
  - Is the drug effective for males?
    P(R | D=true, G=male) = 0.60
    P(R | D=false, G=male) = 0.70
  - Is the drug effective for females?
    P(R | D=true, G=female) = 0.20
    P(R | D=false, G=female) = 0.30
- Marginal query seems **wrong**:
  - Is the drug effective?
    P(R | D=true) = 0.50
    P(R | D=false) = 0.40

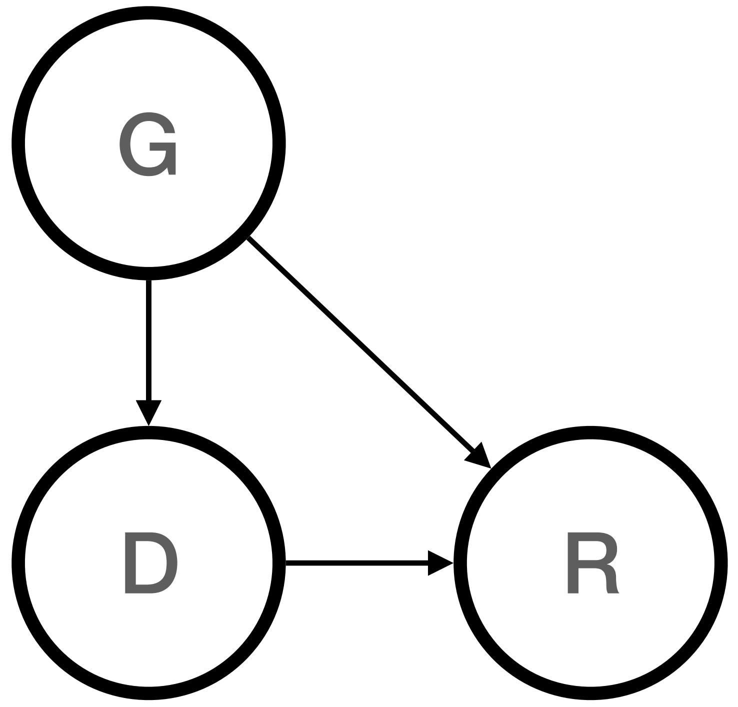

Recap: Selection Bias

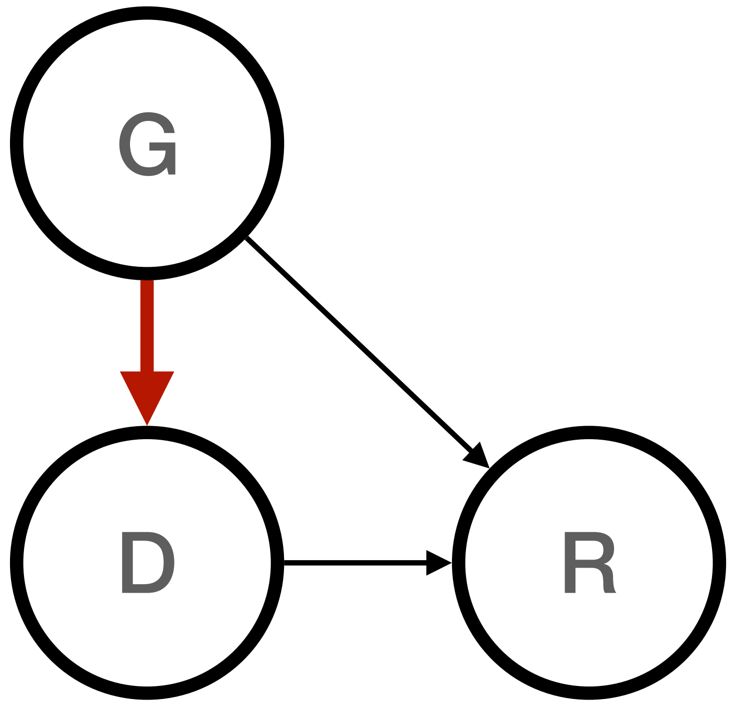

- Simpson's paradox is an example of **selection bias**
- Whether subjects received treatment is **systematically related** to their **response** to the treatment
- **Observational query** is computed as

$$P(R \mid D) = \frac{P(R, D)}{P(D)} = \frac{\sum_G P(G, D, R)}{\sum_{G,R} P(G, D, R)} = \frac{\sum_G P(R \mid D, G) \boxed{P(D \mid G)} P(G)}{\sum_{G,R} P(R \mid D, G) \boxed{P(D \mid G)} P(G)}$$

  - This is the **correct answer** for the **observational** query
- For the causal question, we don't want to condition on P(D | G), because our query is about **forcing** D=true

Post-Intervention Distribution

- The causal query is really a query on a **different distribution** in which we have **forced** *D=true*
    - We will refer to the two distributions as the **observational** distribution and the **post-intervention** distribution
- With a post-intervention distribution, we can compute the answers to causal queries using existing techniques (e.g., variable elimination)

Post-Intervention Distribution for Simpson's Paradox

- **Observational distribution:**
  P(G,D,R) = P(R | D, G) × P(D | G) × P(G)

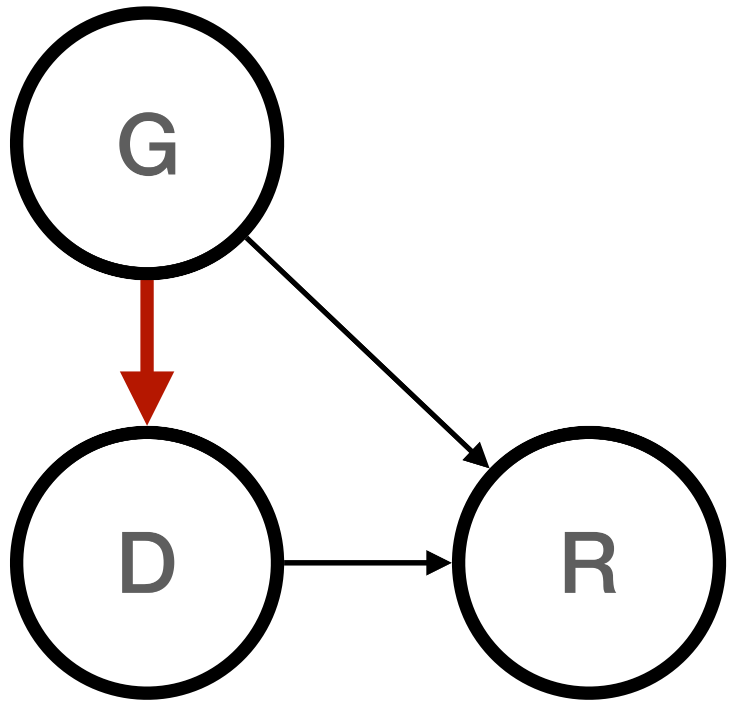

- **Question:** What is the post-intervention distribution for Simpson's Paradox?
  - We're forcing D=true, so P(D=true | G) = 1 for all g∈dom(G)
  - That's the same as just omitting the P(D | G) factor

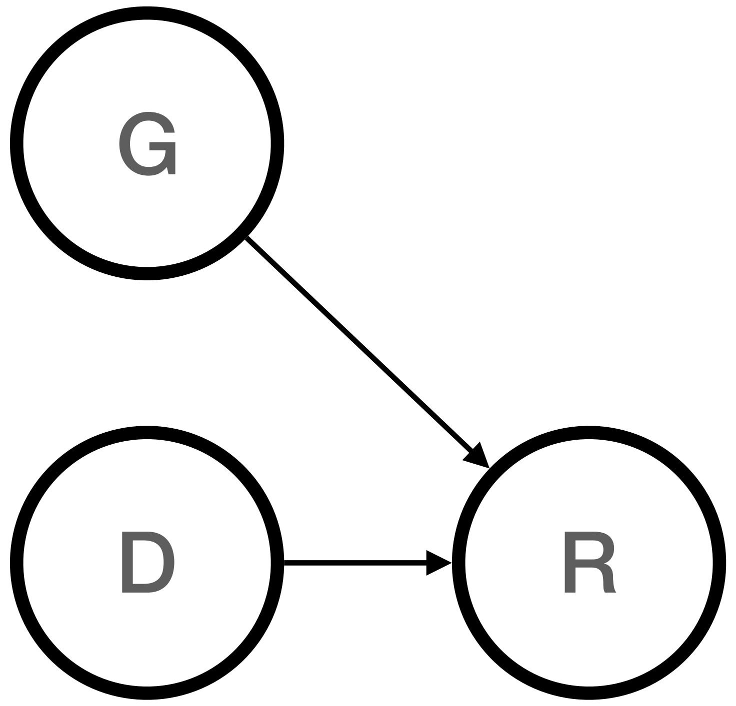

- **Post-intervention distribution:**
  P(G,D,R) = P(R | D, G) × P(G)

The Do-Calculus

- How should we **express** causal queries?
- One approach: The do-calculus
- Condition on **observations**:
  $P(Y \mid X = x)$
- Express **interventions** with special **do** operator:
  $P(Y \mid do(X=x))$
- Allows us to mix observational and interventional information:
  $P(Y \mid Z=z, do(X=x))$

Evaluating Causal Queries With the Do-Calculus

Given a query $P(Y \mid do(X=x), Z=z)$:

1. Construct post-intervention distribution $\hat{P}$ by **removing** all links from X's direct parents to X

2. Evaluate the **observational** query $\hat{P}(Y \mid X=x, Z=z)$ in the **post-intervention distribution**

Example: Simpson's Paradox

- **Observational distribution:**
  P(G,D,R) = P(R | D, G) × P(D | G) × P(G)

- **Observational query:**

$$P(R \mid D) = \frac{P(R, D)}{P(D)} = \frac{\sum_G P(G, D, R)}{\sum_{G,R} P(G, D, R)} = \frac{\sum_G P(R \mid D, G) P(D \mid G) P(G)}{\sum_{G,R} P(R \mid D, G) P(D \mid G) P(G)}$$

- Observational query values:
  P(R | D=true) = 0.50
  P(R | D=false) = 0.40

- **Post-intervention distribution** for causal query P(R | do(D=true)):
  P̂(G,D,R) = P(R | D, G) × P(G)

- **Causal query:**

$$P(R \mid do(D = true)) = \hat{P}(R \mid D = true) = \frac{\sum_G P(R \mid D, G) P(G)}{\sum_{G,R} P(R \mid D, G) P(G)}$$

- Causal query values:
  P(R | do(D=true)) = 0.40
  P(R | do(D=false)) = 0.50

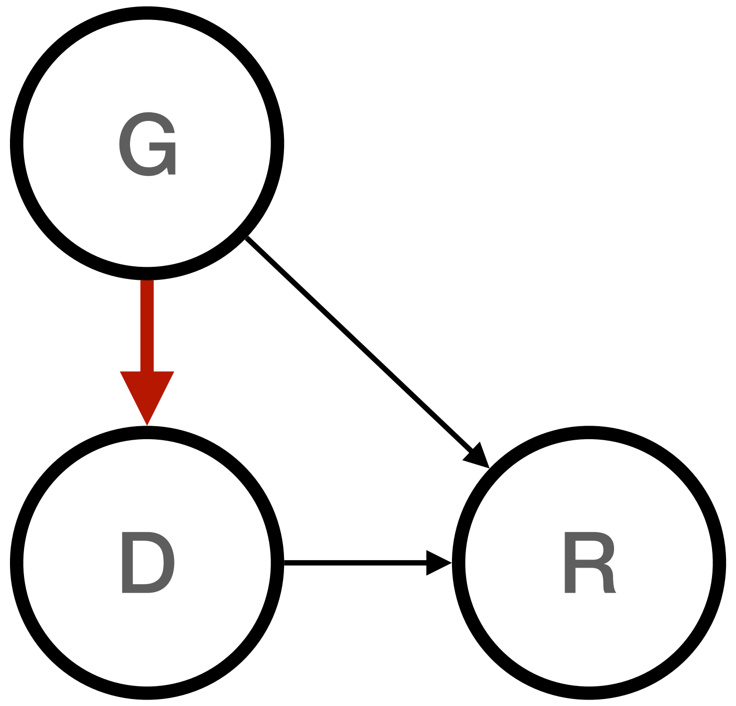

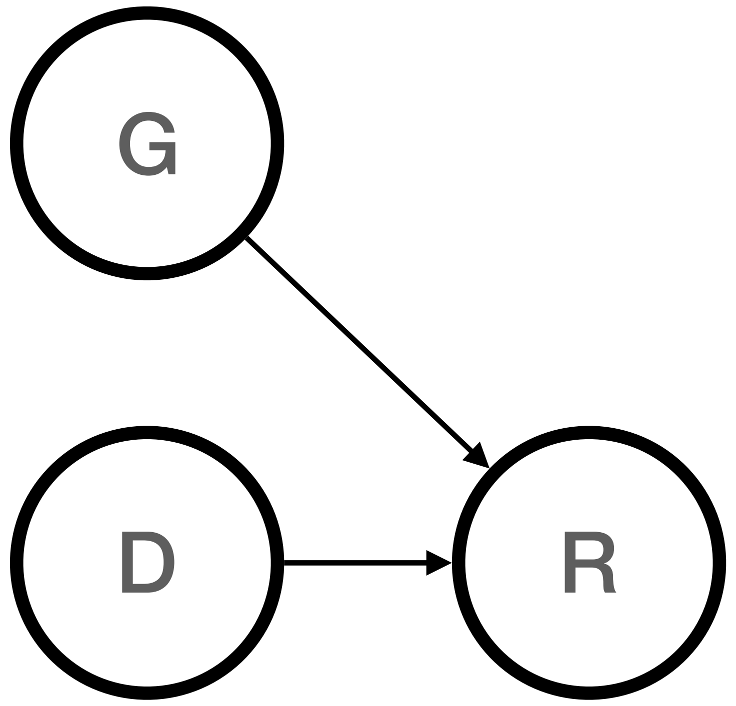

Example: Rainy Sidewalk

Query: P(Rain | do(Wet=true)

Natural network:

- Observational distribution:
  P(Wet, Rain) = P(Wet|Rain)P(Rain)
- Post intervention distribution:
  $\hat{P}$(Wet=true, Rain) = P(Rain)P(Wet)
- P(Rain | do(Wet=true)) = **.50**

Inverted network:

- Observational distribution:
  P(Wet, Rain) = P(Rain | Wet)P(Rain)
- Post intervention distribution:
  $\hat{P}$(Wet=true, Rain) = P(Rain | Wet)P(Wet)
- P(Rain | do(Wet=true)) = **.78**

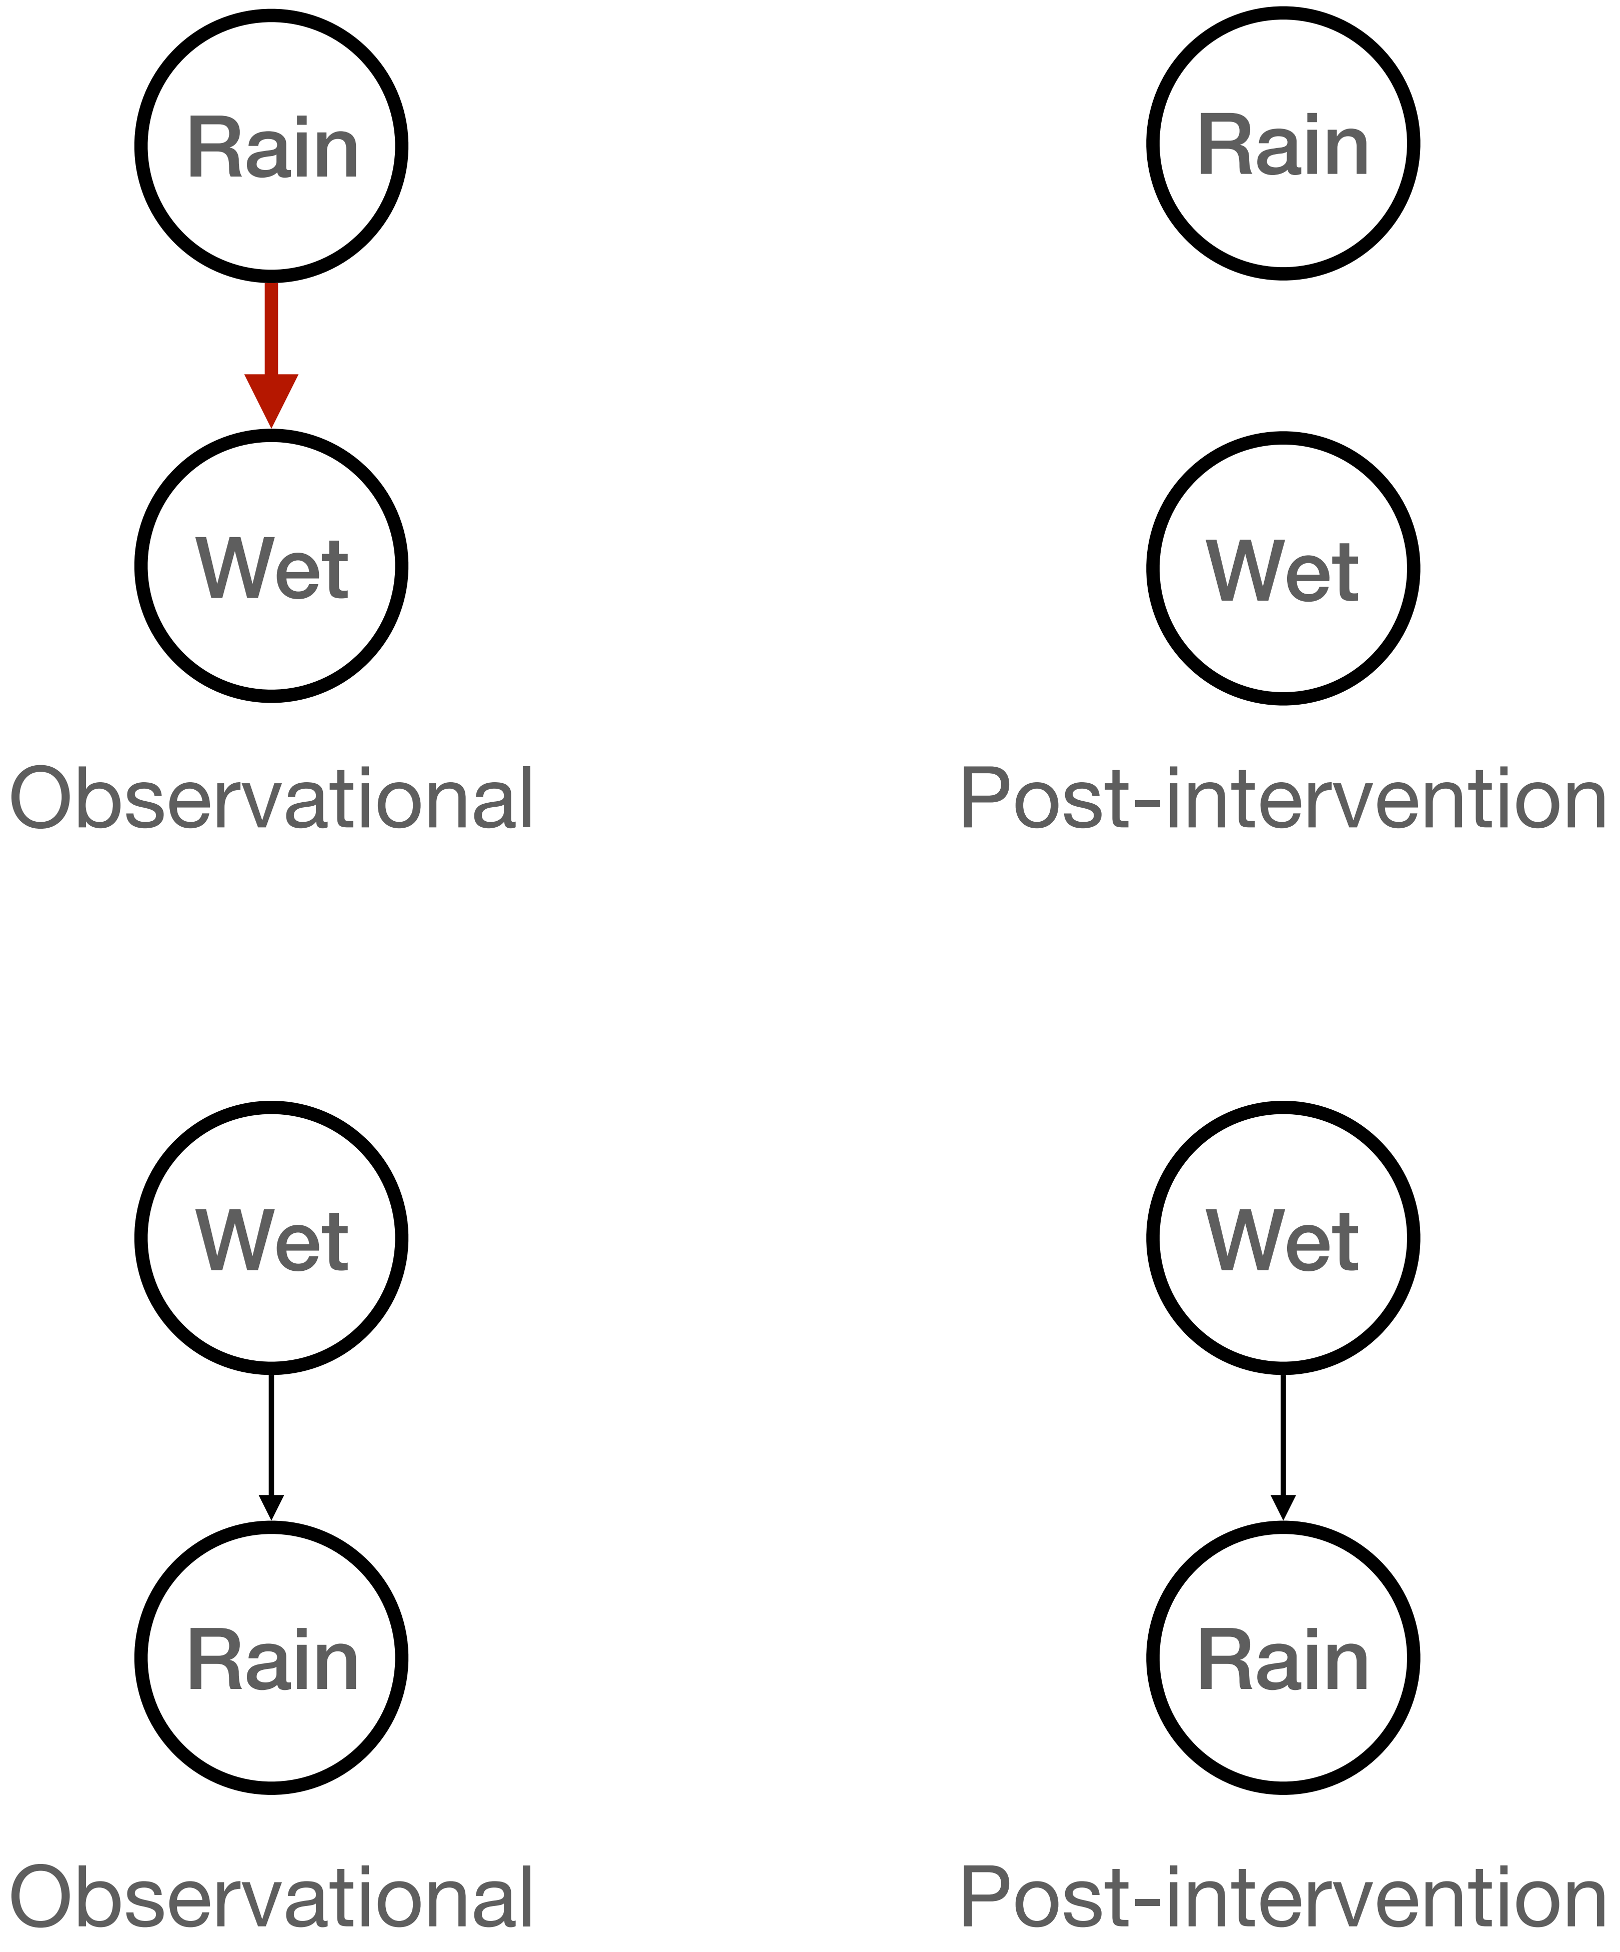

Causal Models

A: Both networks encode valid factorings of the observational distribution, but the inverted network does not encode the correct causal structure.

- The **natural network** gives the correct answer to our causal query, but the **inverted network** does not (**Why**?)

- Not every factoring of a joint distribution is a valid causal model

Definition:
A **causal model** is a directed acyclic graph of random variables such that for every edge $X \to Y$, the value of random variable X is **realized before** the value of random variable Y.

Alternative Representation: Influence Diagrams

Instead of adding a new operator, we can instead represent causal queries by **augmenting** the causal model with **decision variables** F_D for each potential intervention target D.

$\mathrm{dom}(F_D) = \mathrm{dom}(D) \cup \{idle\}$

$$P(D \mid pa(D), F_D) = \begin{cases} P(D \mid pa(D)) & \text{if } F_D = idle, \\ 1 & \text{if } F_D \neq idle \wedge D = F_D, \\ 0 & \text{otherwise.} \end{cases}$$

Influence Diagrams Examples

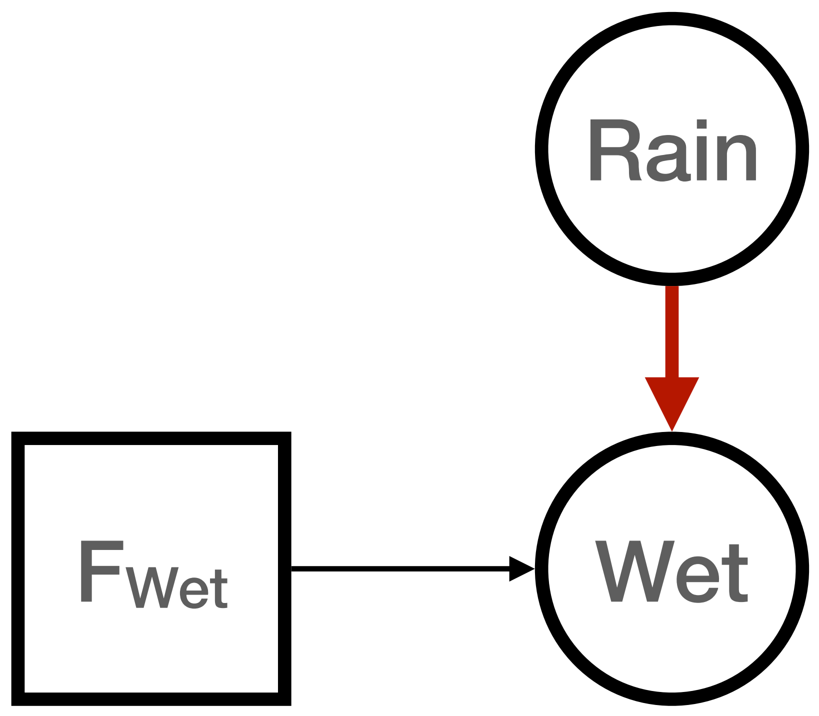

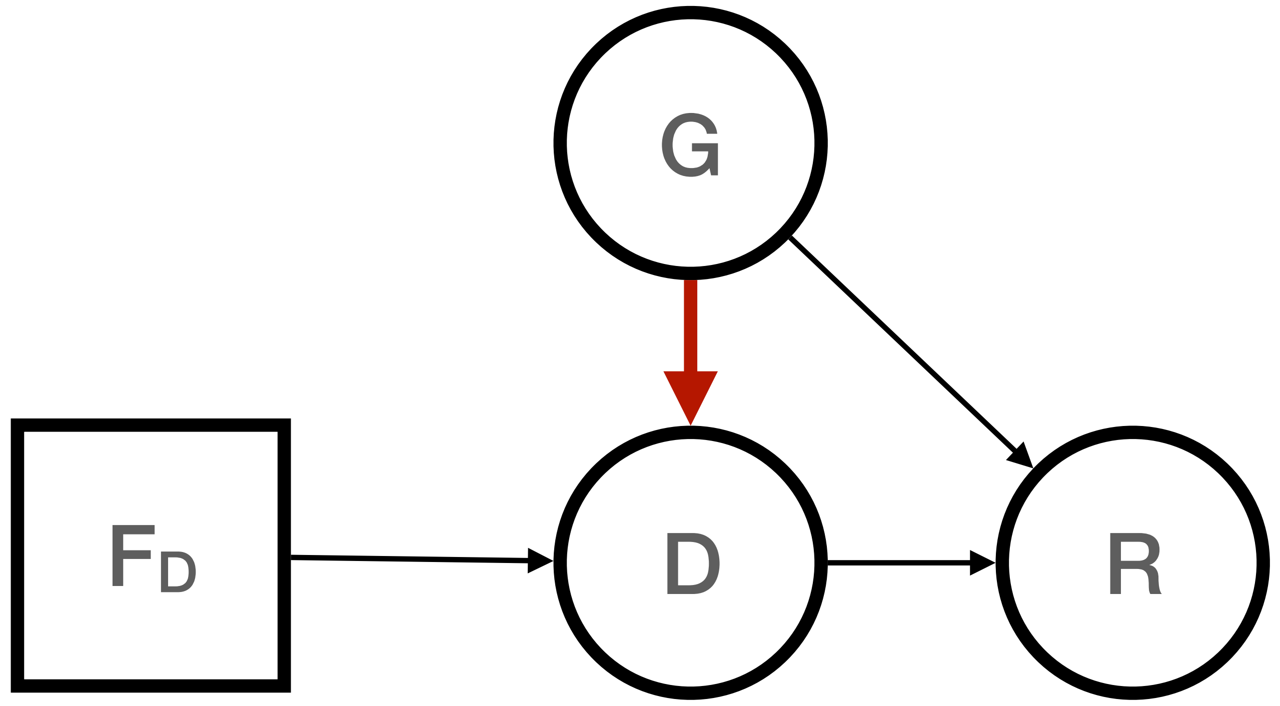

Partially Observable Models

- Sometimes we will have a **causal model** (i.e., graph), but not all of the **conditional distributions**
    - This is the case in most experiments!
    - **Question:** Why/how could this happen?
        - Observational data that didn't include all variables of interest
        - Some causal variables might be unobservable even in principle
- **Question:** Can we still answer **observational** questions?
- **Question:** Can we still answer **causal** questions?

Simpson's Paradox Variations

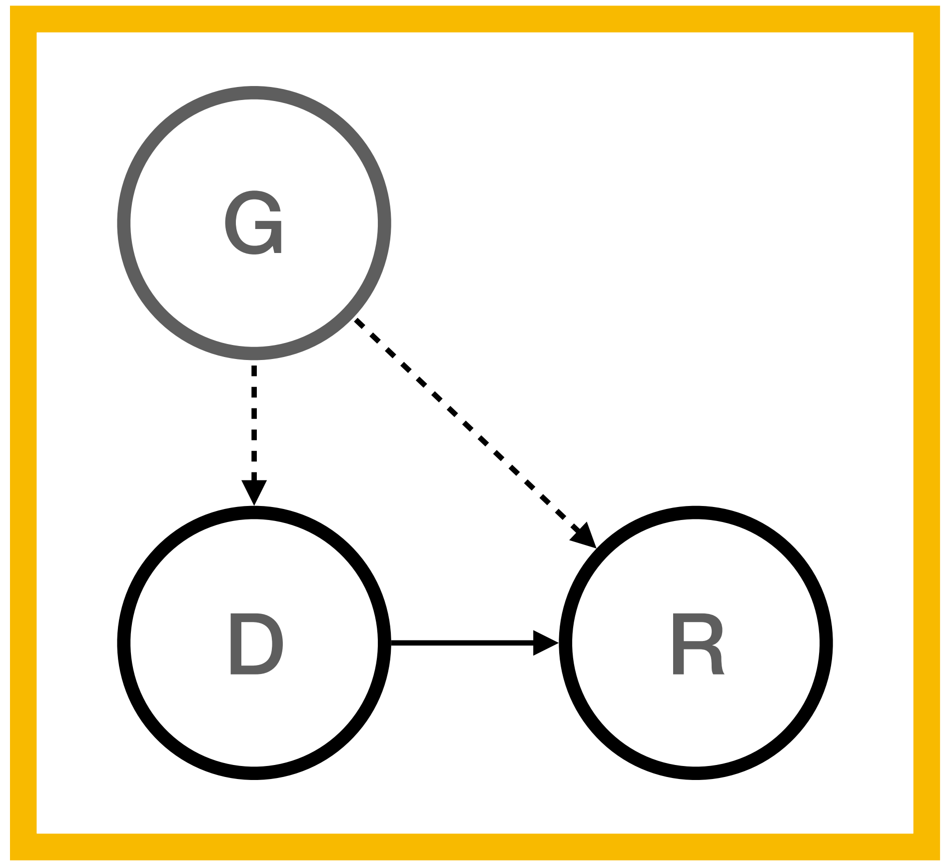

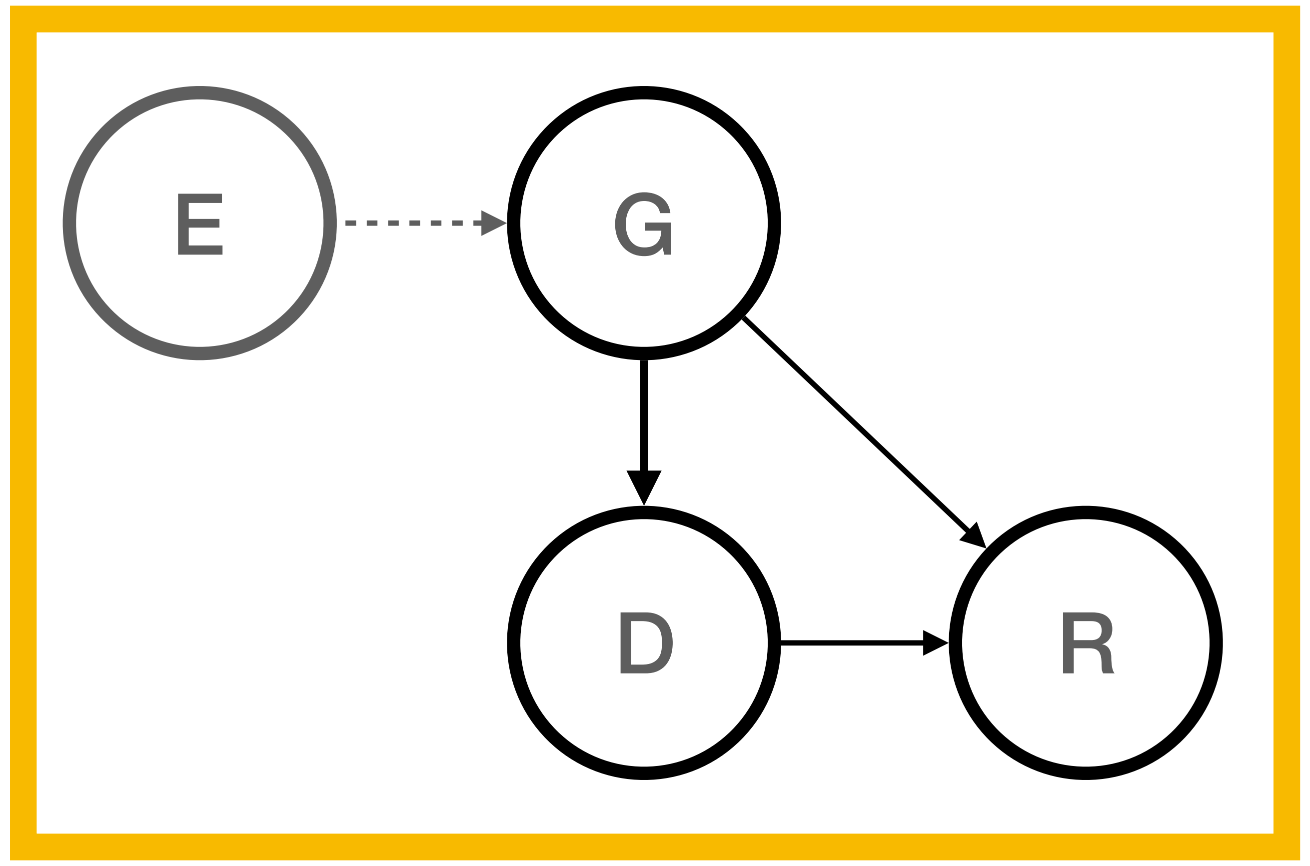

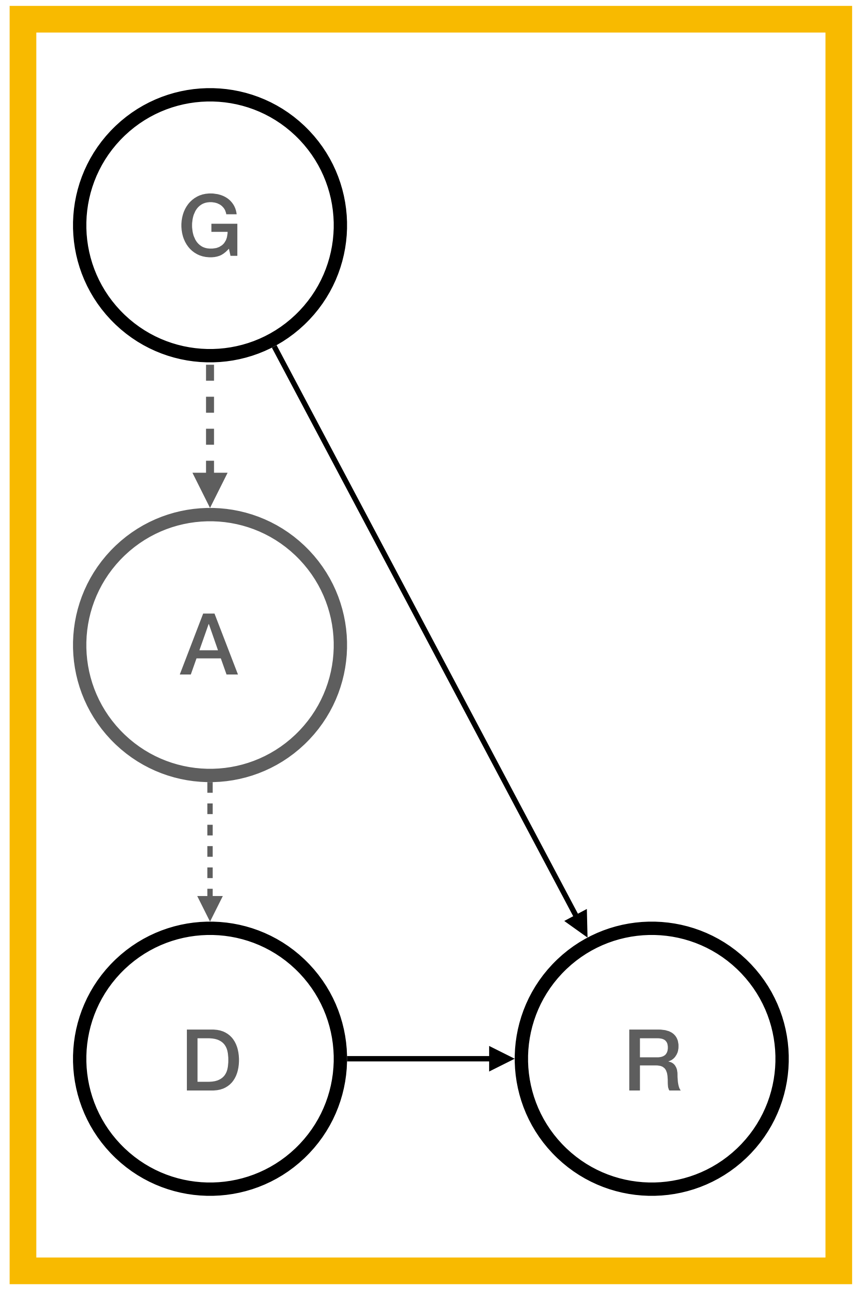

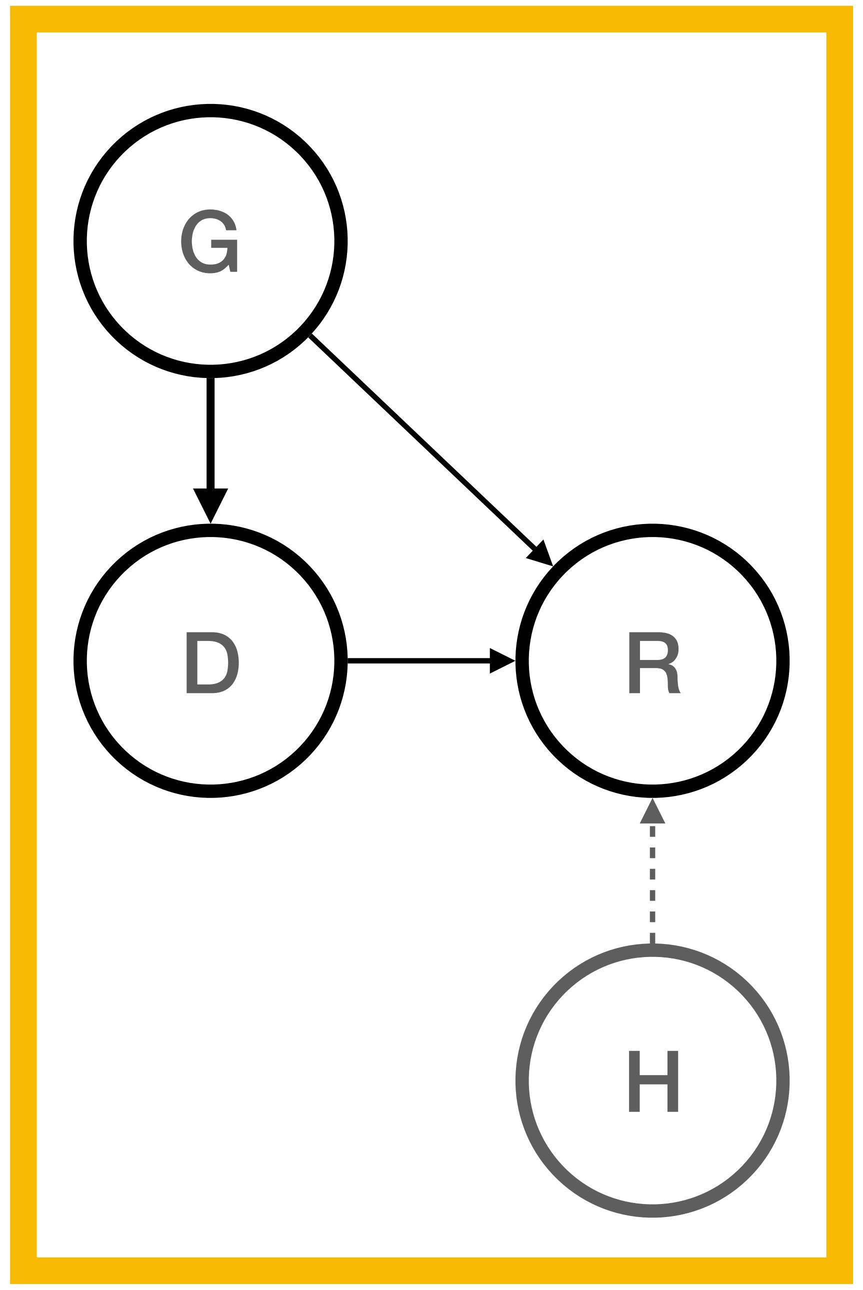

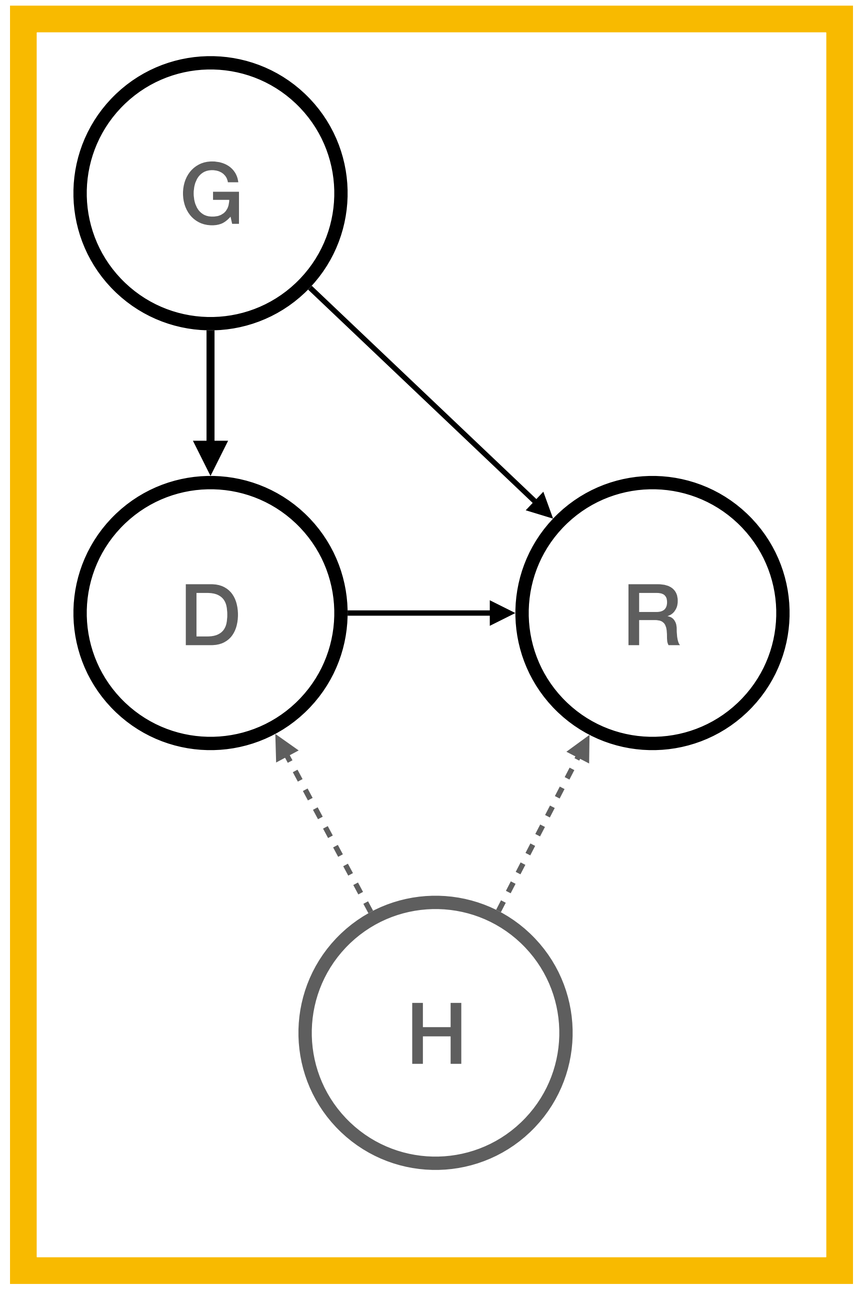

Question: Can we answer the query **P(R | do(D))** in these causal models?

(answers in subsequent slides)

Identifiability

- Many different distributions can be consistent with a given causal model
- A causal query is **identifiable** if it is the same in every distribution that is consistent with the **observed variables** and the **causal model**

Definition: (Pearl, 2000)
The causal effect of X on Y is **identifiable** from a graph *G* if the quantity P(Y | do(X=x)) can be computed uniquely from any positive probability of the observed variables.

I.e., if $P_{M1}(Y \mid do(X=x)) = P_{M2}(Y \mid do(X=x))$ for every pair of models $M1$,$M2$ such that

1. The **causal graph** of both $M1$ and $M2$ is G
2. The **joint distributions** on the **observed** variables v are equal: $P_{M1}(v) = P_{M2}(v)$

Direct Causes Criterion

Theorem: (Pearl, 2000)
Given a causal graph G of any **Markovian** model in which a subset of variables V are observed, the causal effect $P(Y \mid do(X=x))$ is **identifiable** whenever $\{X \cup Y \cup pa(X)\}$ are observable.
That is, whenever X, Y, and **all parents** of X are observable.

Simpson's Paradox Revisited #1

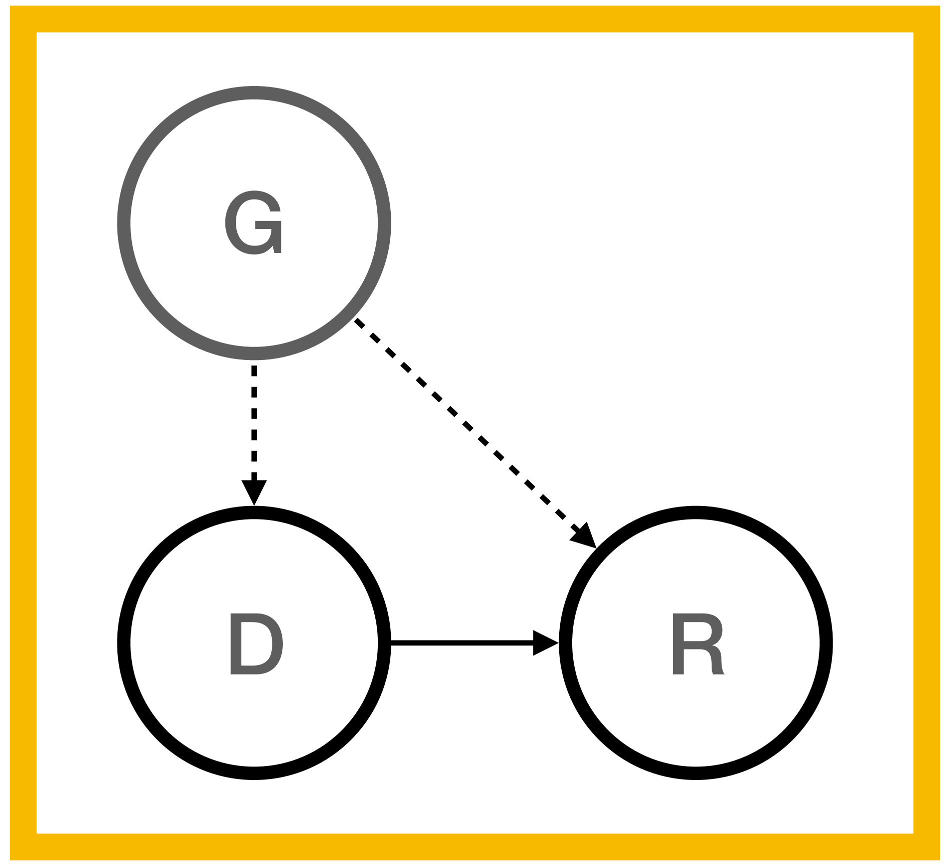

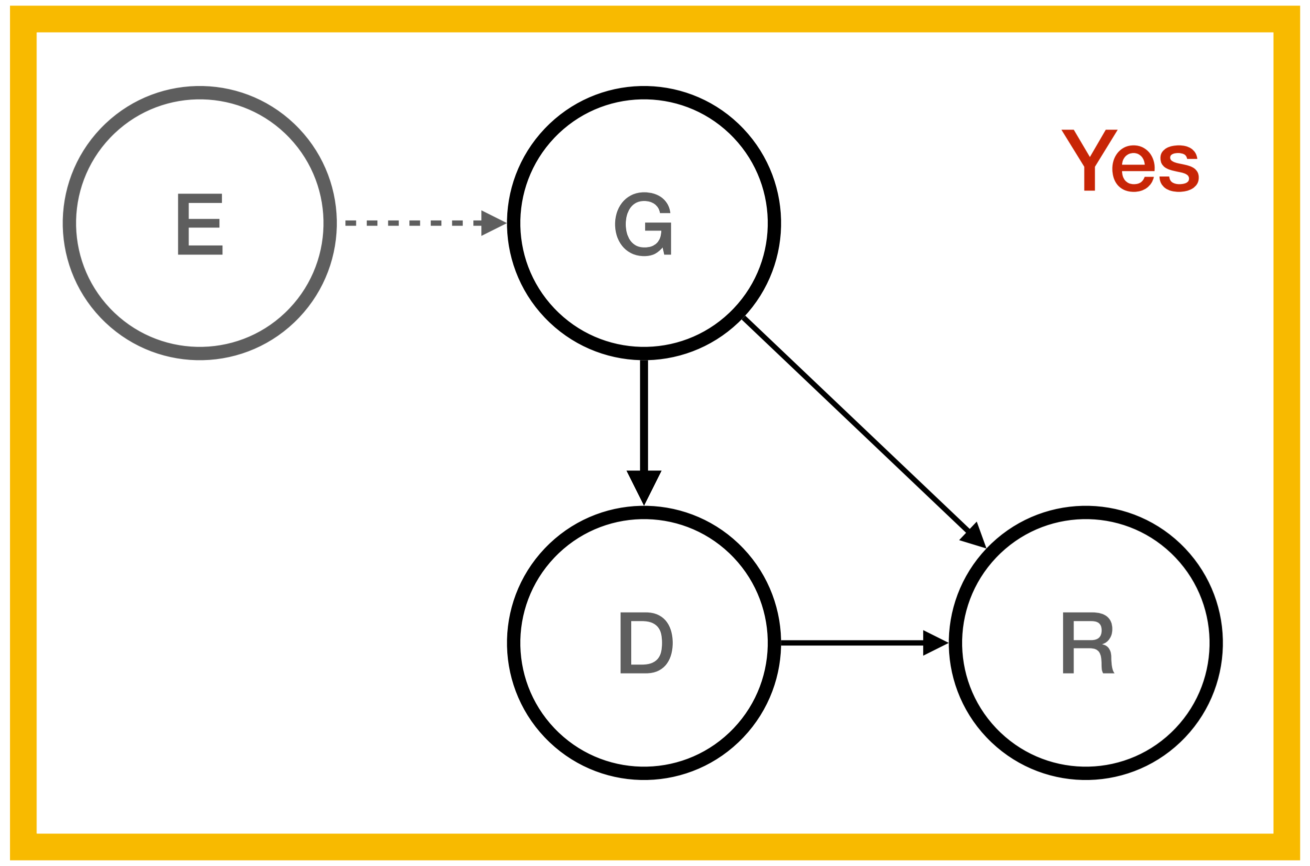

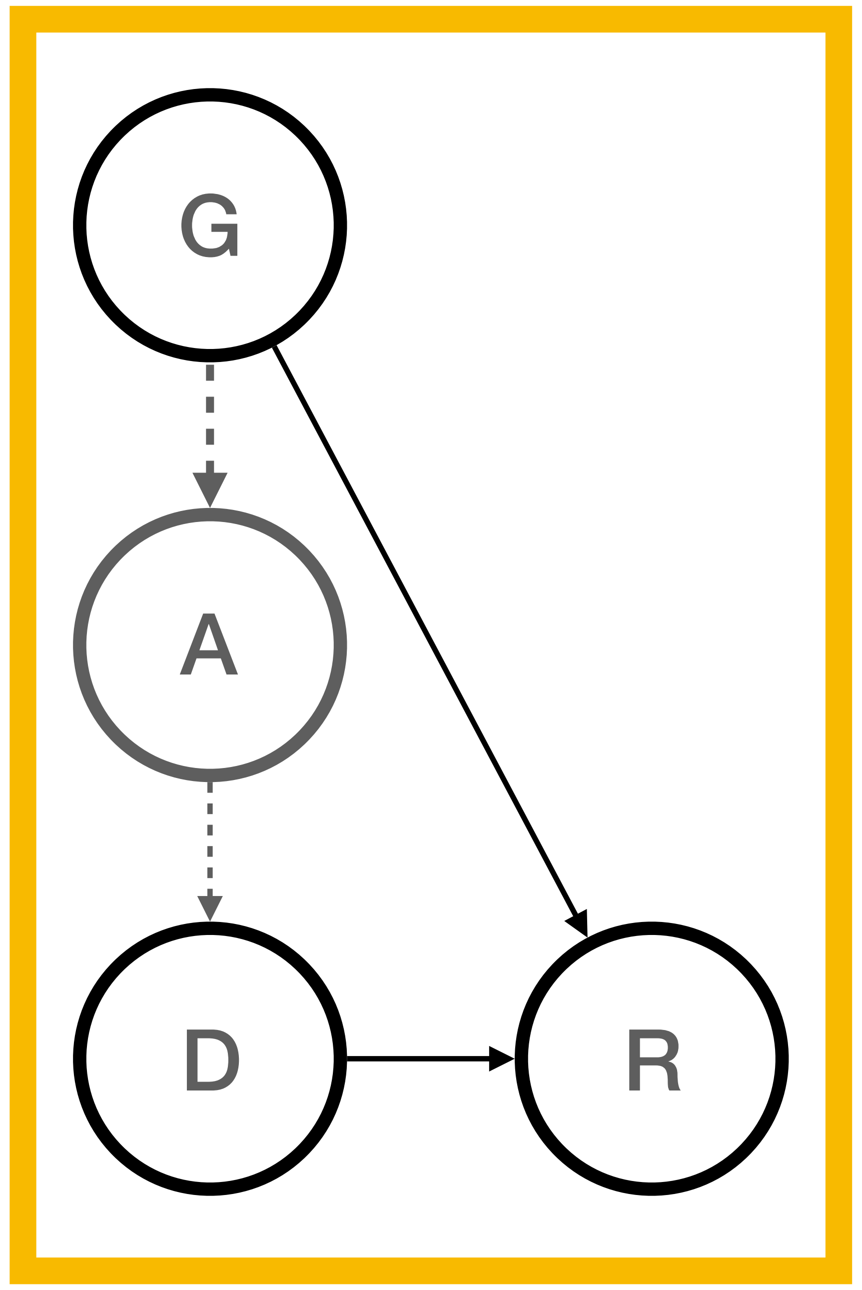

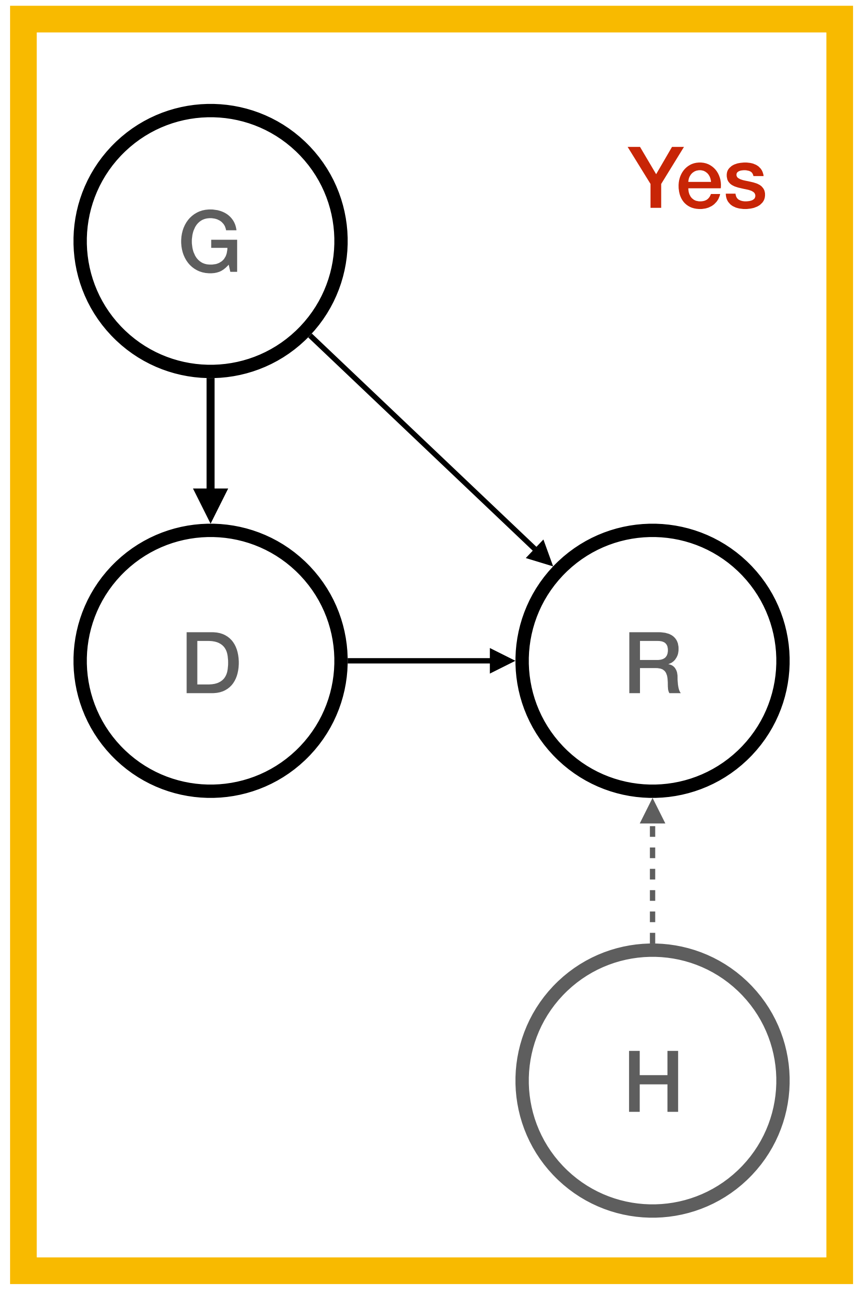

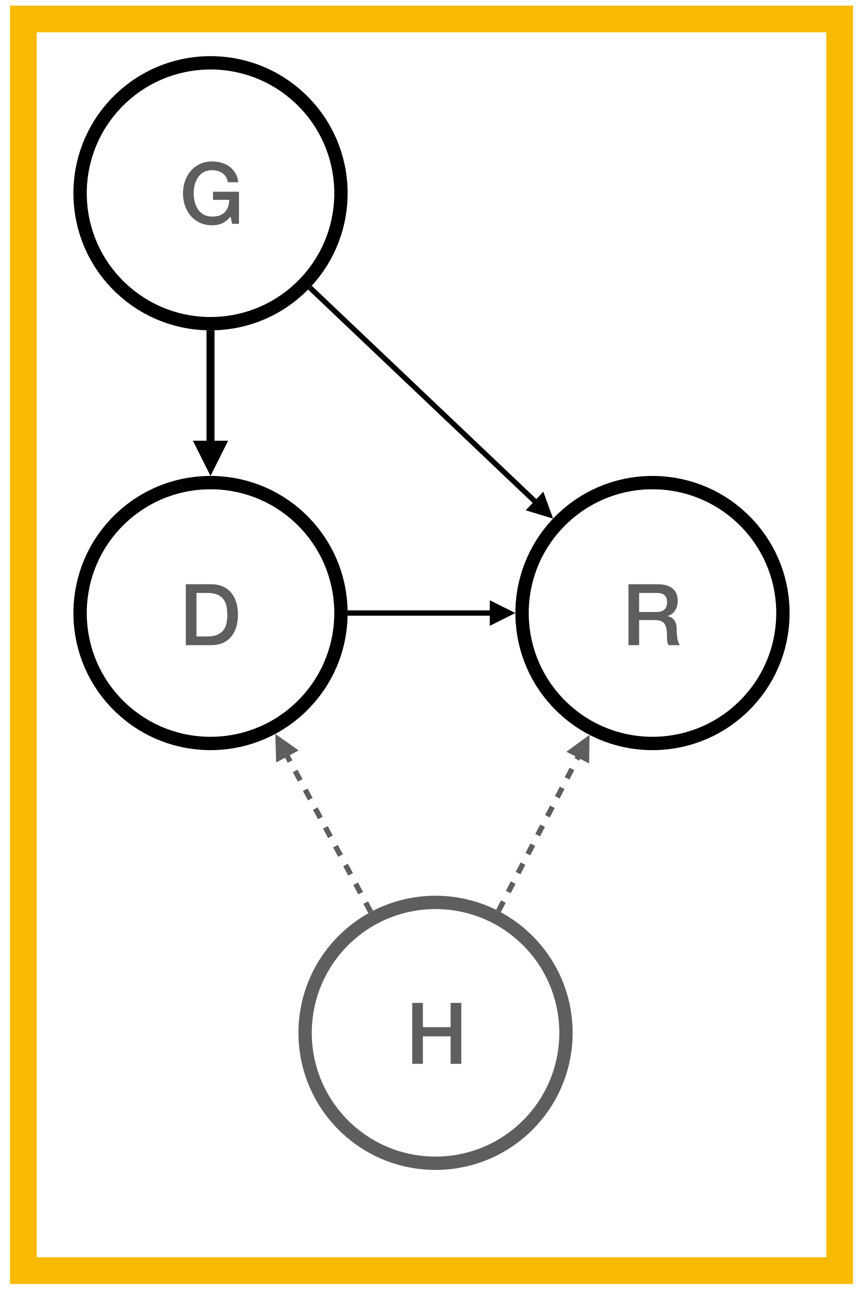

Question: Can we answer the query **P(R | do(D))** in these causal models?

(answers in subsequent slides)

Back Door Paths

- An **undirected path** is a path that ignores edge directions
  - **Examples:** *X,Y,Z* and *A,B,C* above
- A **back-door path** from *S* to *T* is an undirected path from *S* to *T* where the **first arc** enters *S*
  - **Examples:**
    - *A,B,C* is a back-door path
    - *Y,Z* is a back-door path
    - *X,Y,Z* is **not** a back-door path

Back Door Criterion

Definition:
A set Z of variables satisfies the **back-door criterion** with respect to a pair of variables X, Y if

1. No node in Z is a **descendant** of X, and

2. Z blocks every **back-door path** from X to Y

Theorem: (Pearl 2000)
If a set of **observed** variables Z satisfies the back-door criterion with respect to X,Y, then the causal effect of X on Y is identifiable and is given by the formula $P(Y \mid \mathrm{do}(X = x)) = \sum_{z \in dom(Z)} P(Y \mid X = x, Z = z) P(Z = z)$.

Simpson's Paradox Revisited #2

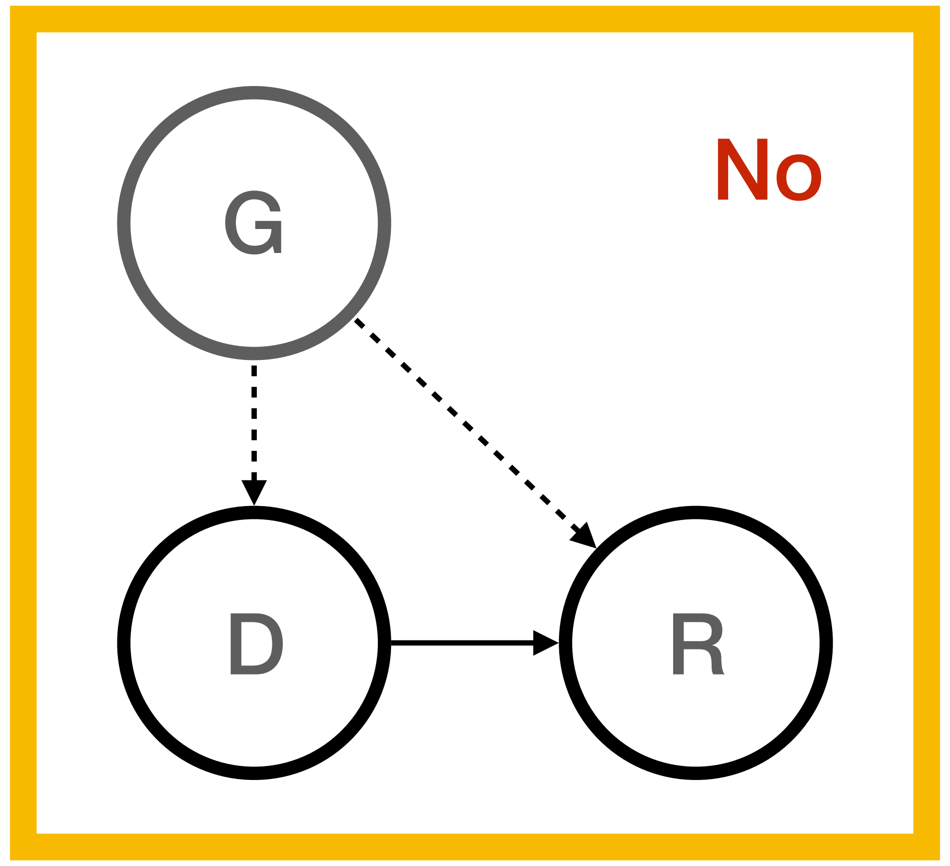

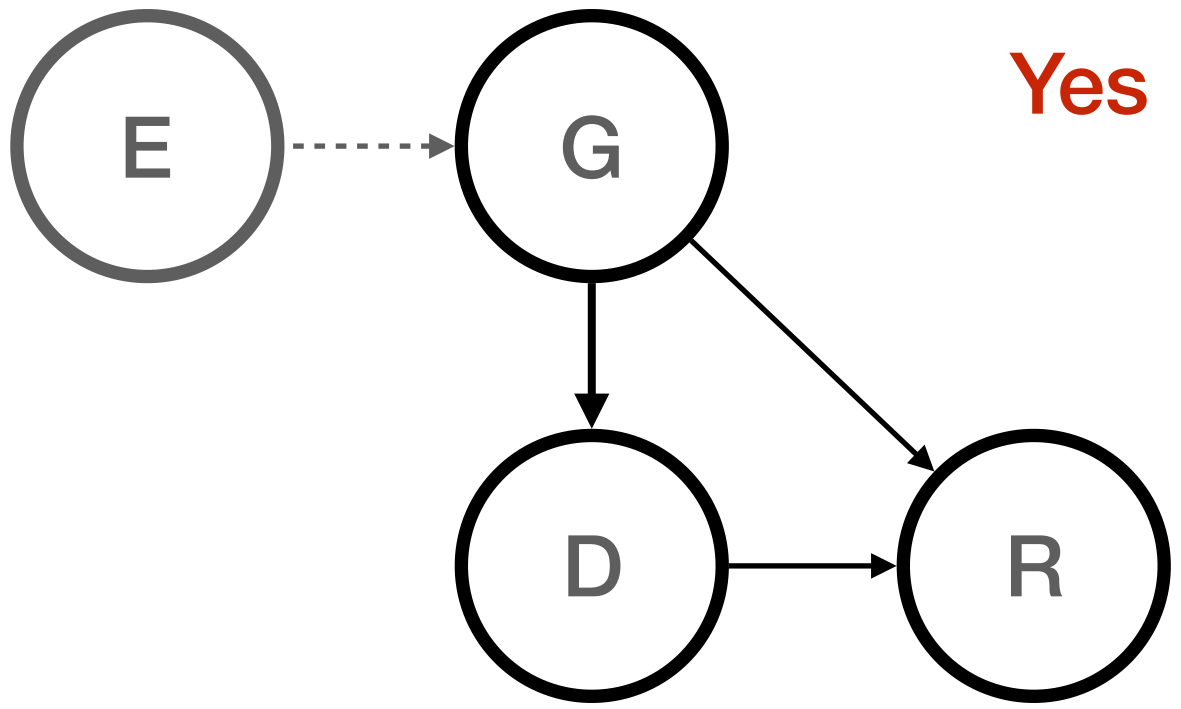

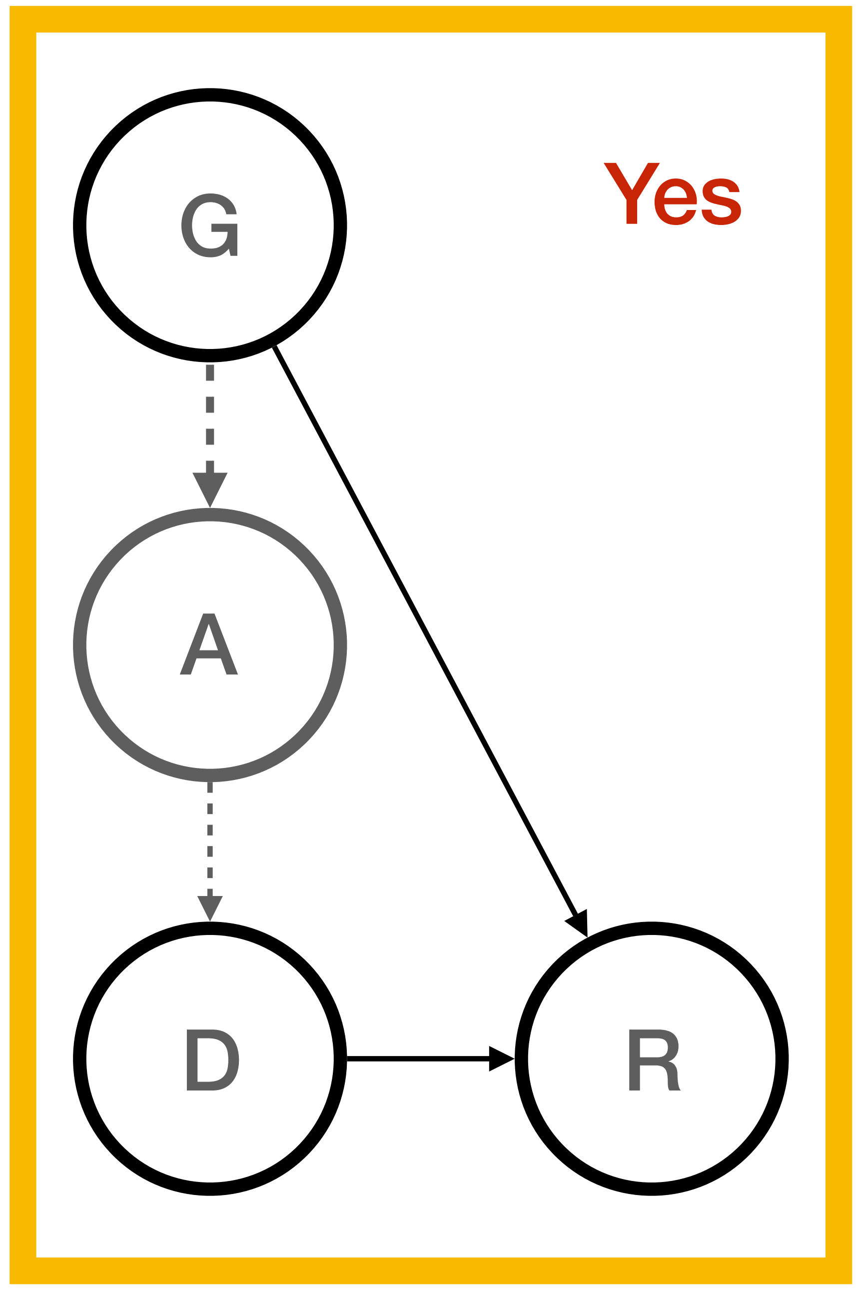

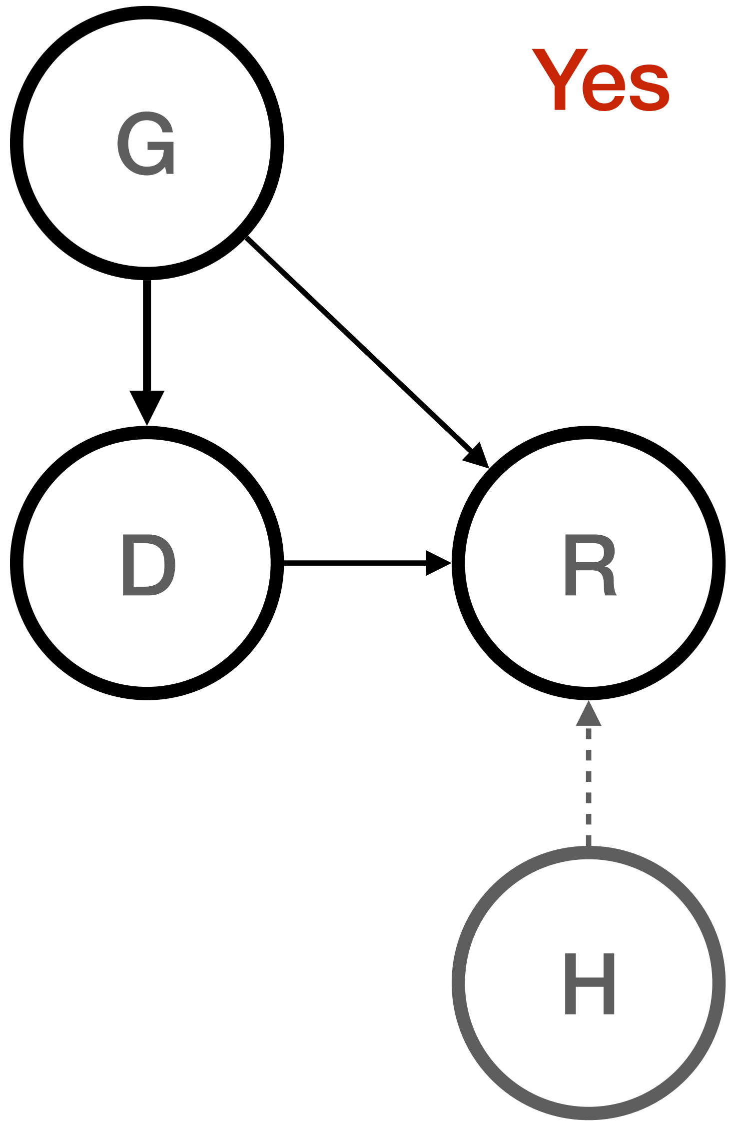

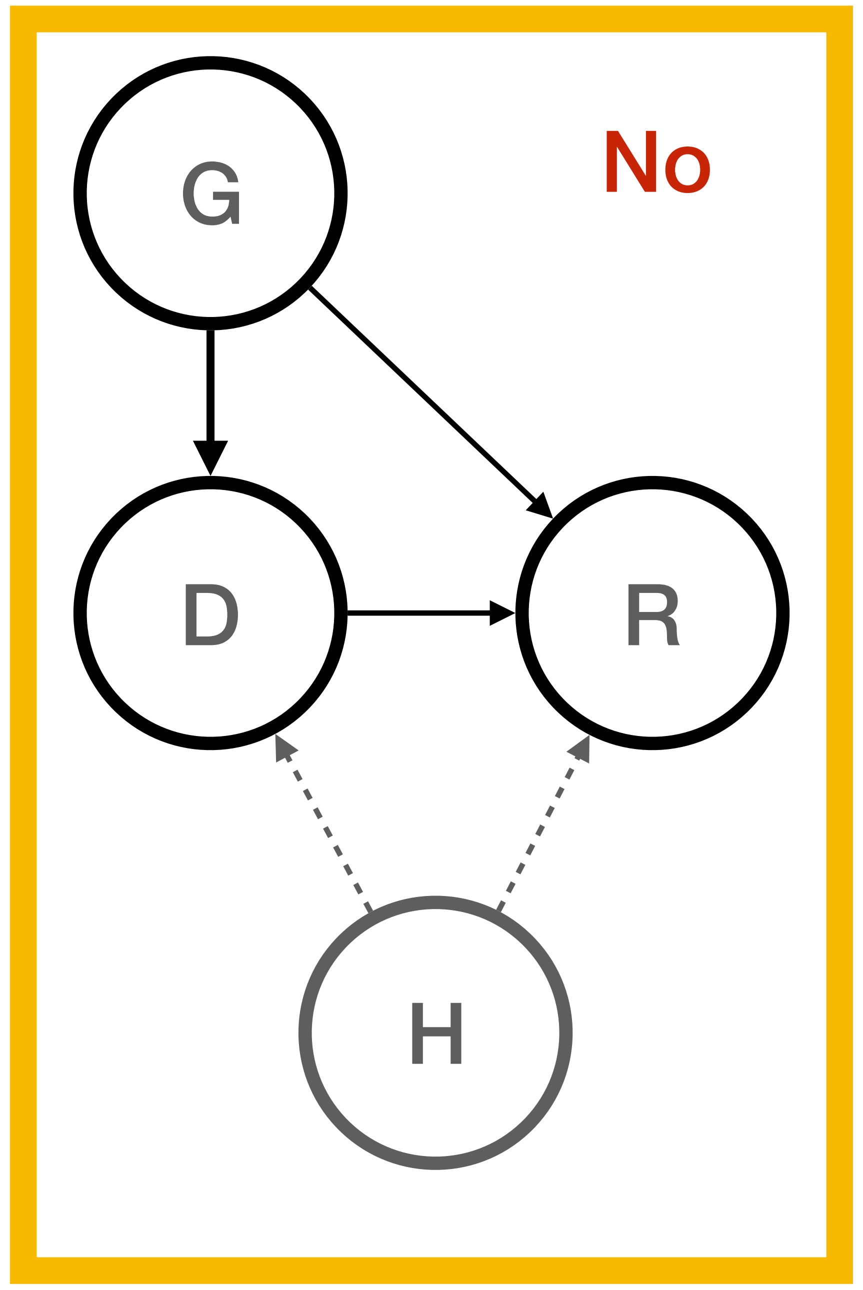

Question: Can we answer the query **P(R | do(D))** in these causal models?

Summary

- **Observational** queries $P(Y \mid X{=}x)$ are different from **causal** queries $P(Y \mid do(X{=}x))$
- To evaluate causal query $P(Y \mid do(X{=}x))$:
  1. Construct post-intervention distribution $\hat{P}$ by **removing** all links from X's direct parents to X
  2. Evaluate the **observational** query $\hat{P}(Y \mid X{=}x, Z{=}z)$ in the **post-intervention distribution**
- Not every correct Bayesian network is a valid causal model
- Causal effects can **sometimes** be **identified** in a partially-observable model:
  - Direct causes criterion
  - Back-door criterion