Causality

Part I: Interventions

Adnan Darwiche

UCLA

Why Causality?

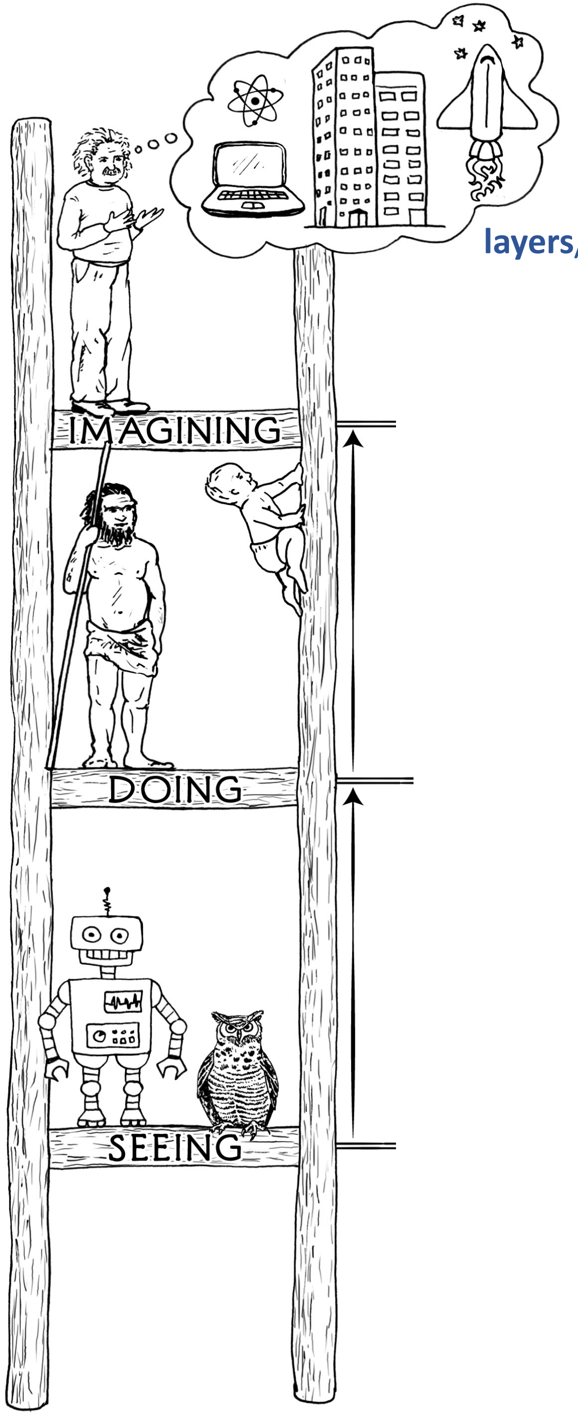

layers/rungs of the causal hierarchy

next lecture

this lecture

class so far

Themes

- Why do we need it?
- What additional information do we need so we can do it?
- How can we do it in an idealized setting?
  (complete model: causal graph + parameters)
- How can we do it in a practical setting?
  (partial model: causal graph + data)

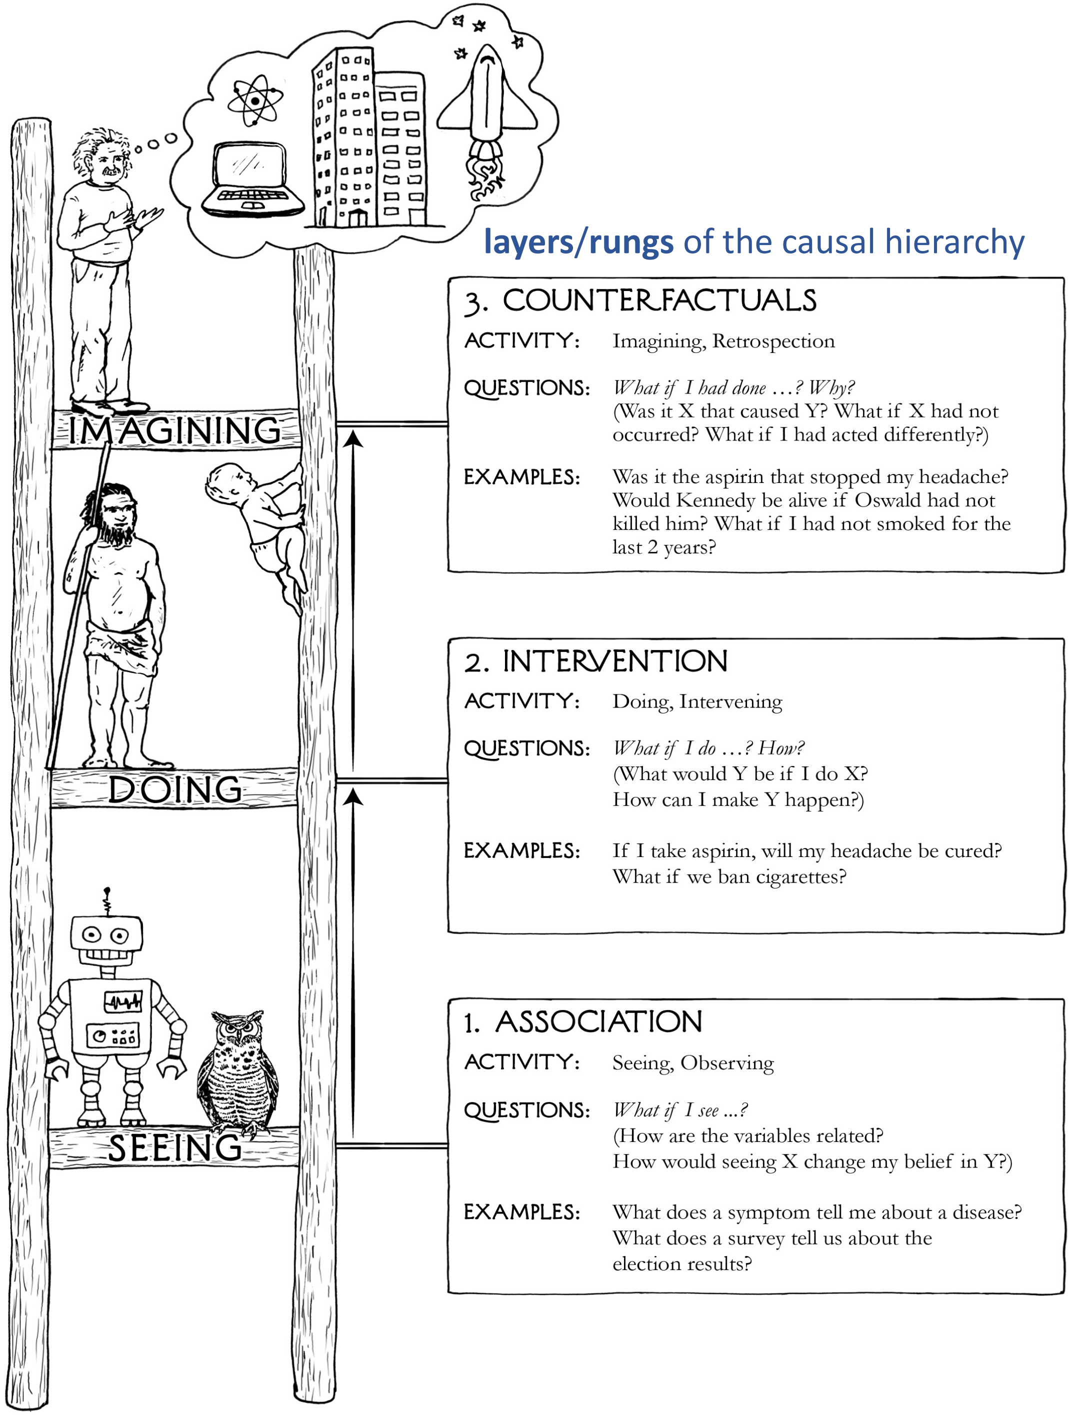

next lecture

this lecture

class so far

1987

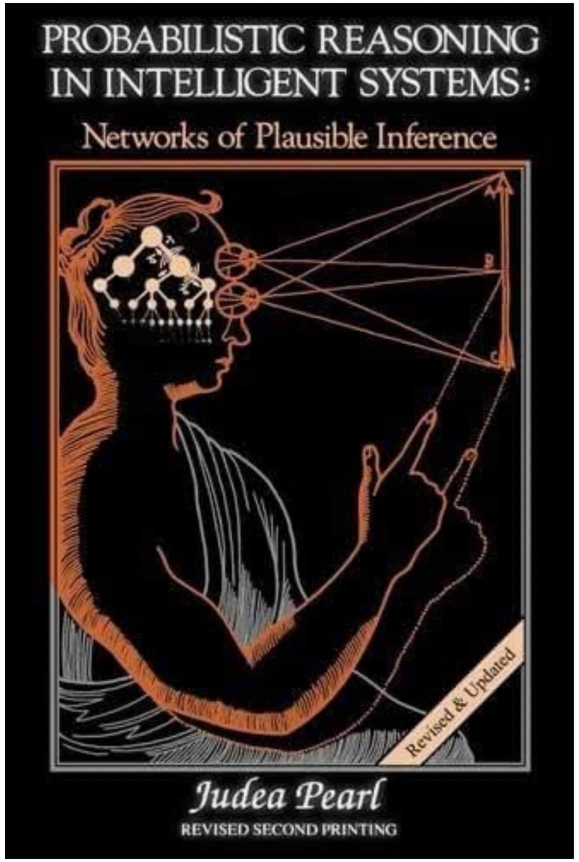

2009

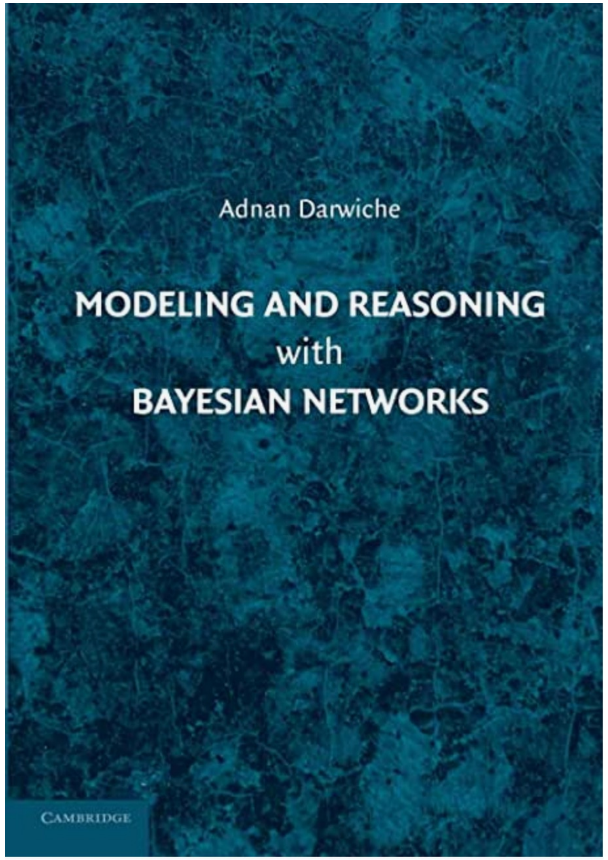

2009

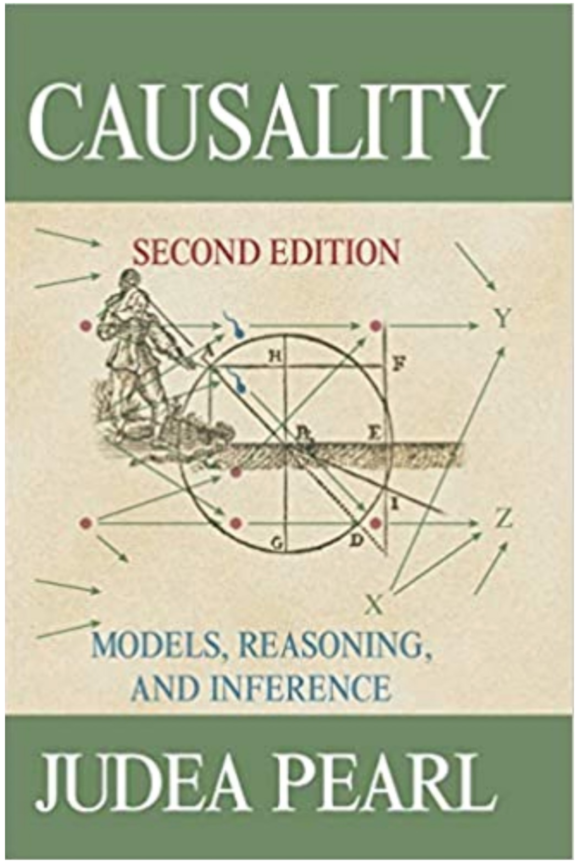

2018

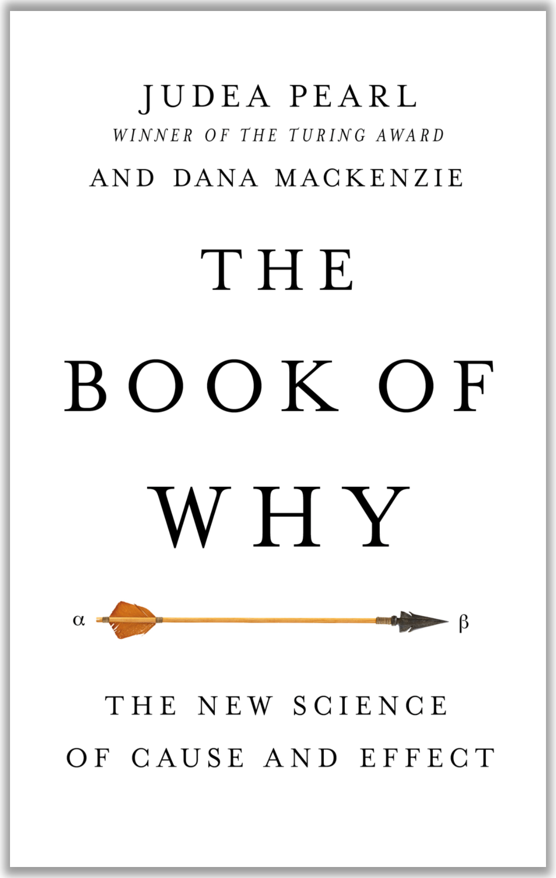

association

intervention & counterfactuals

2016

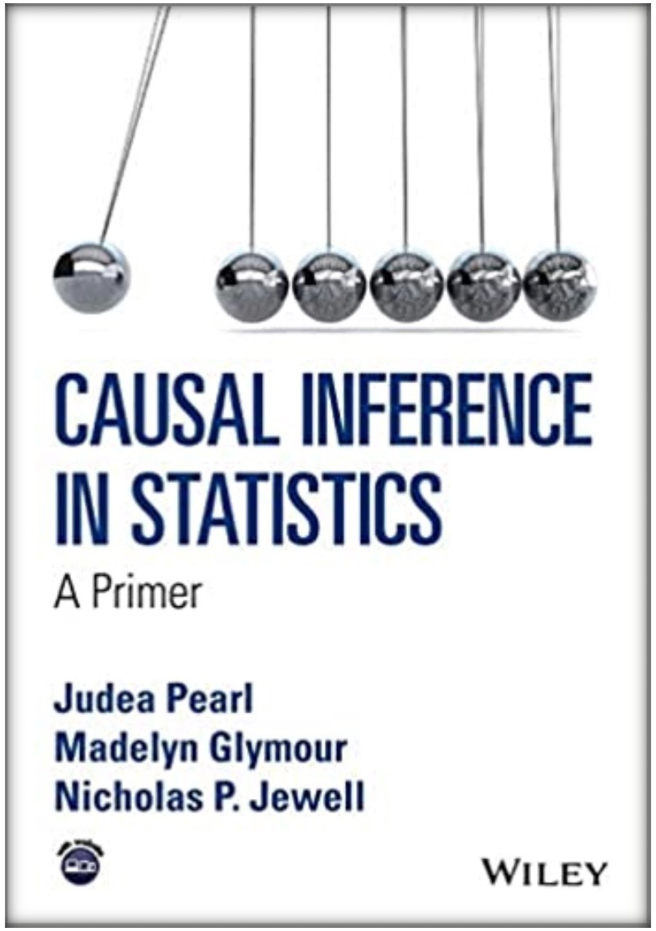

http://bayes.cs.ucla.edu/PRIMER/

2009

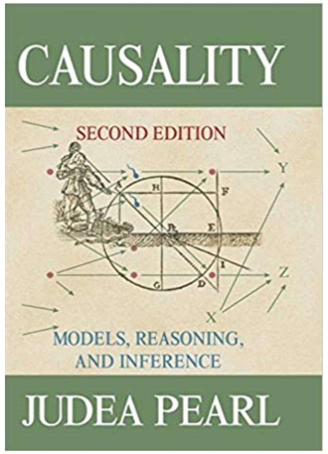

Why Do We Need It? causal effect

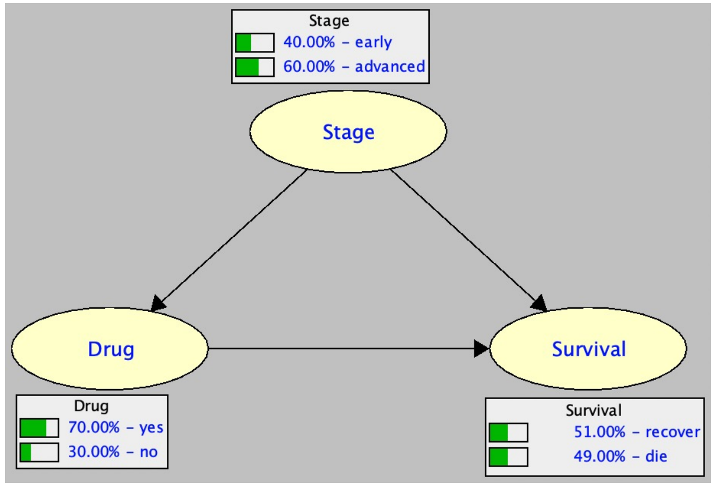

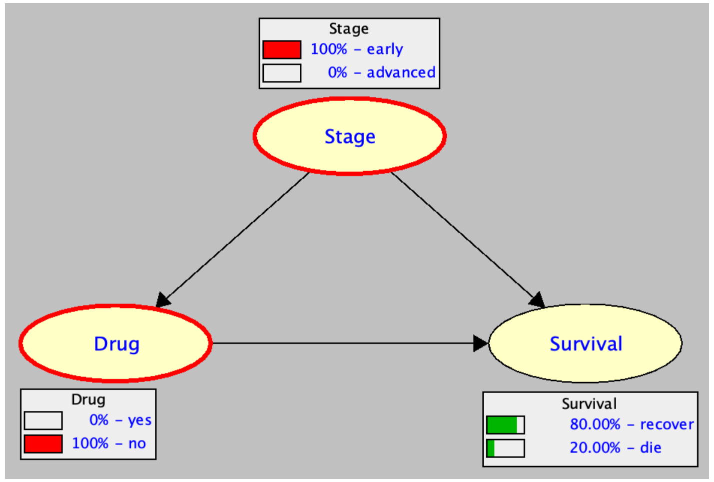

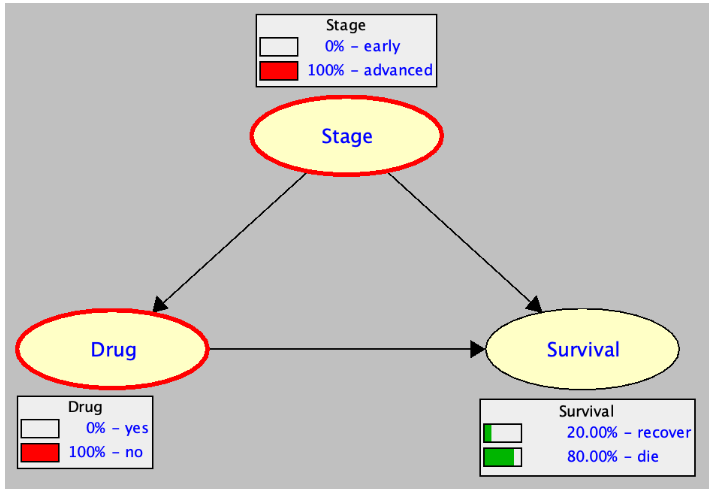

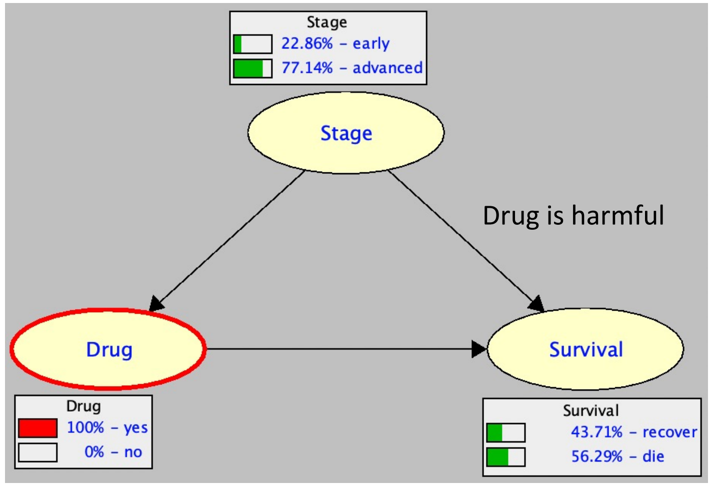

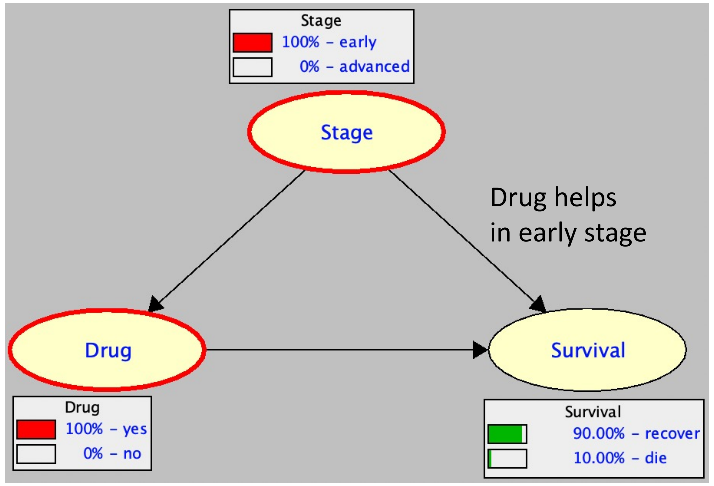

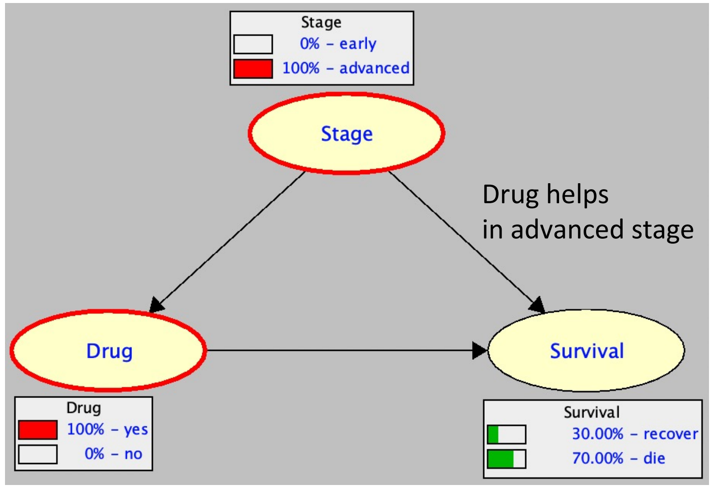

Simpson's Paradox

http://reasoning.cs.ucla.edu/samiam/

How Can We Do It? causal effect

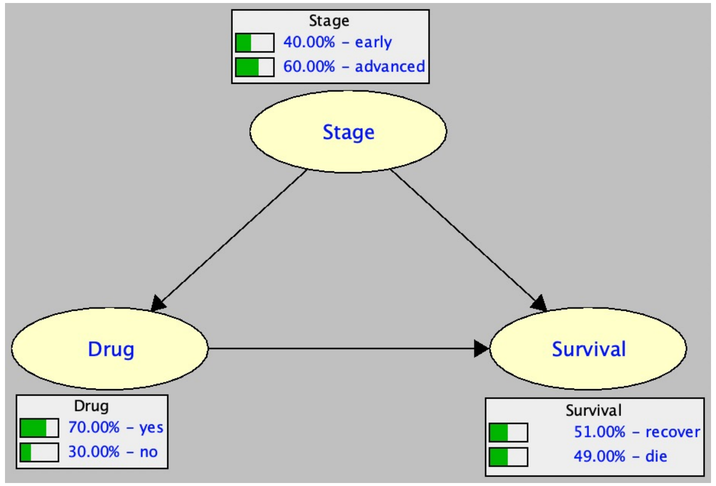

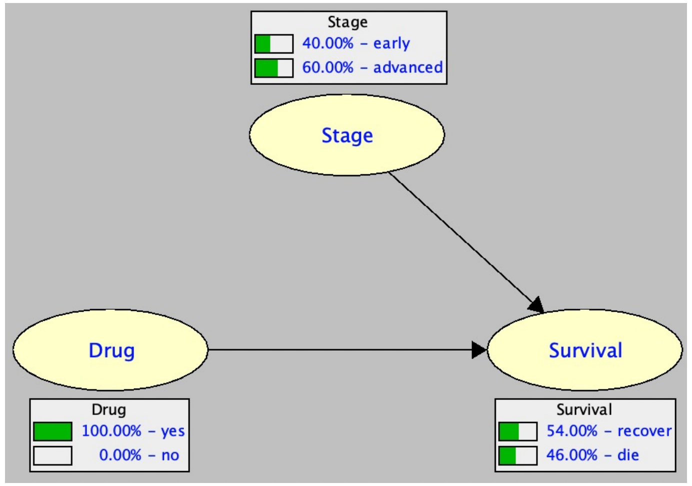

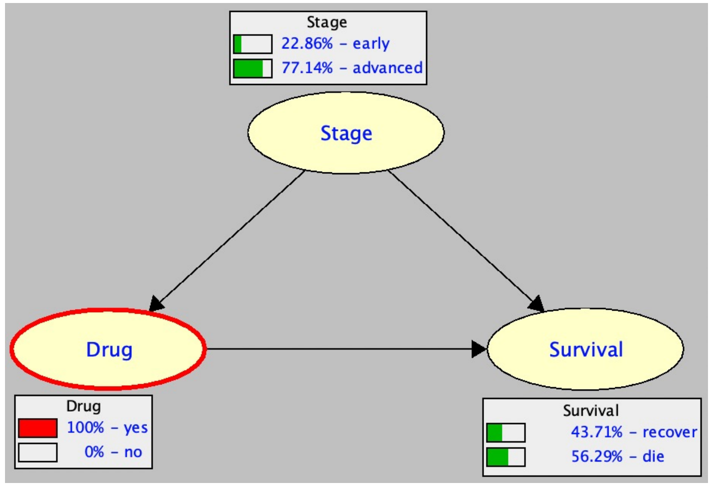

Original CPT: $Pr(Drug|Stage)$. New CPT: $Pr'(Drug)$.

$$
\begin{aligned}
Pr'(Drug{=}yes) &= 1 \\
Pr'(Drug{=}no) &= 0
\end{aligned}
$$

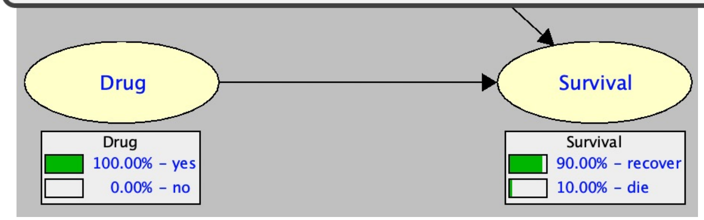

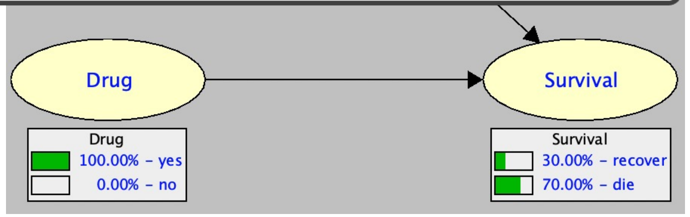

What Additional Information Do We Need?

model 1

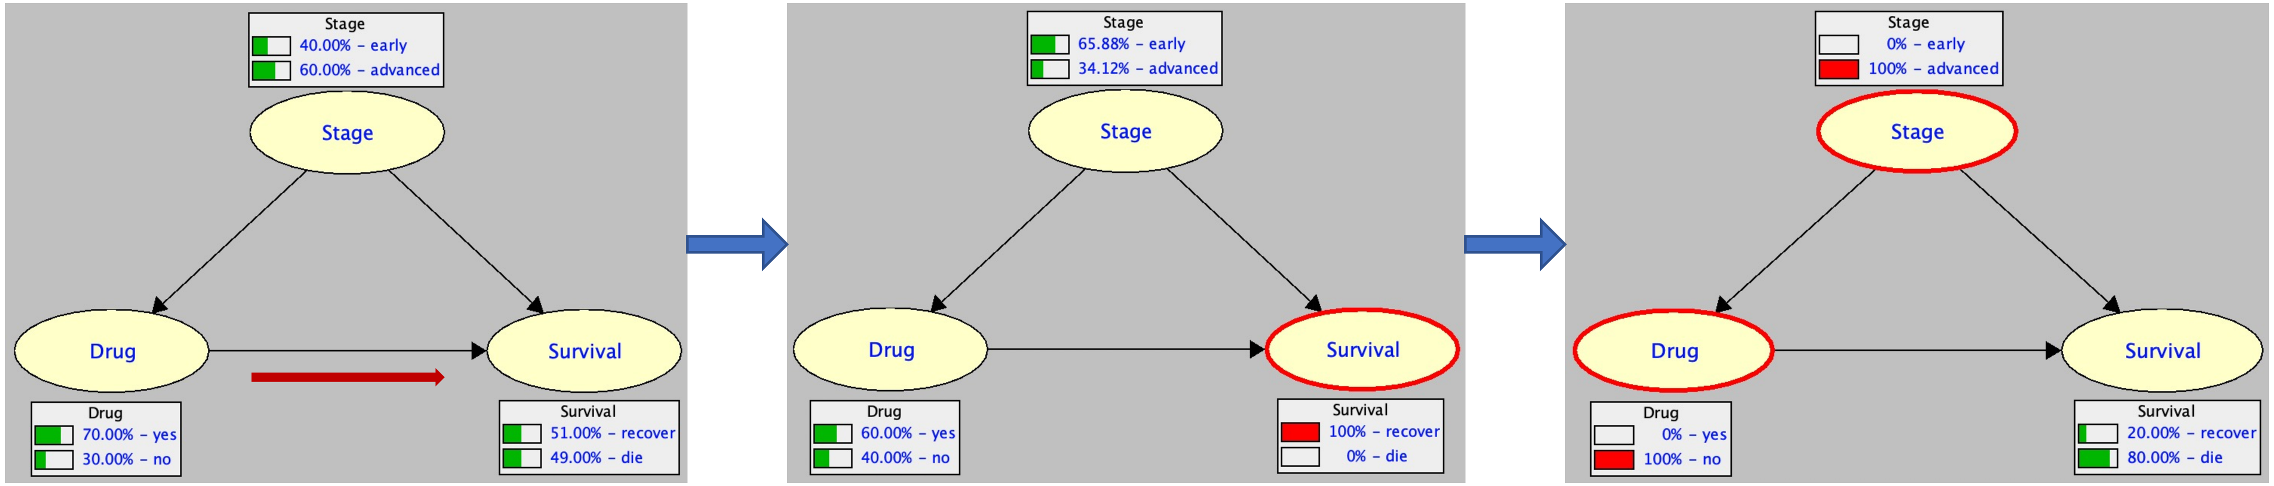

model 2

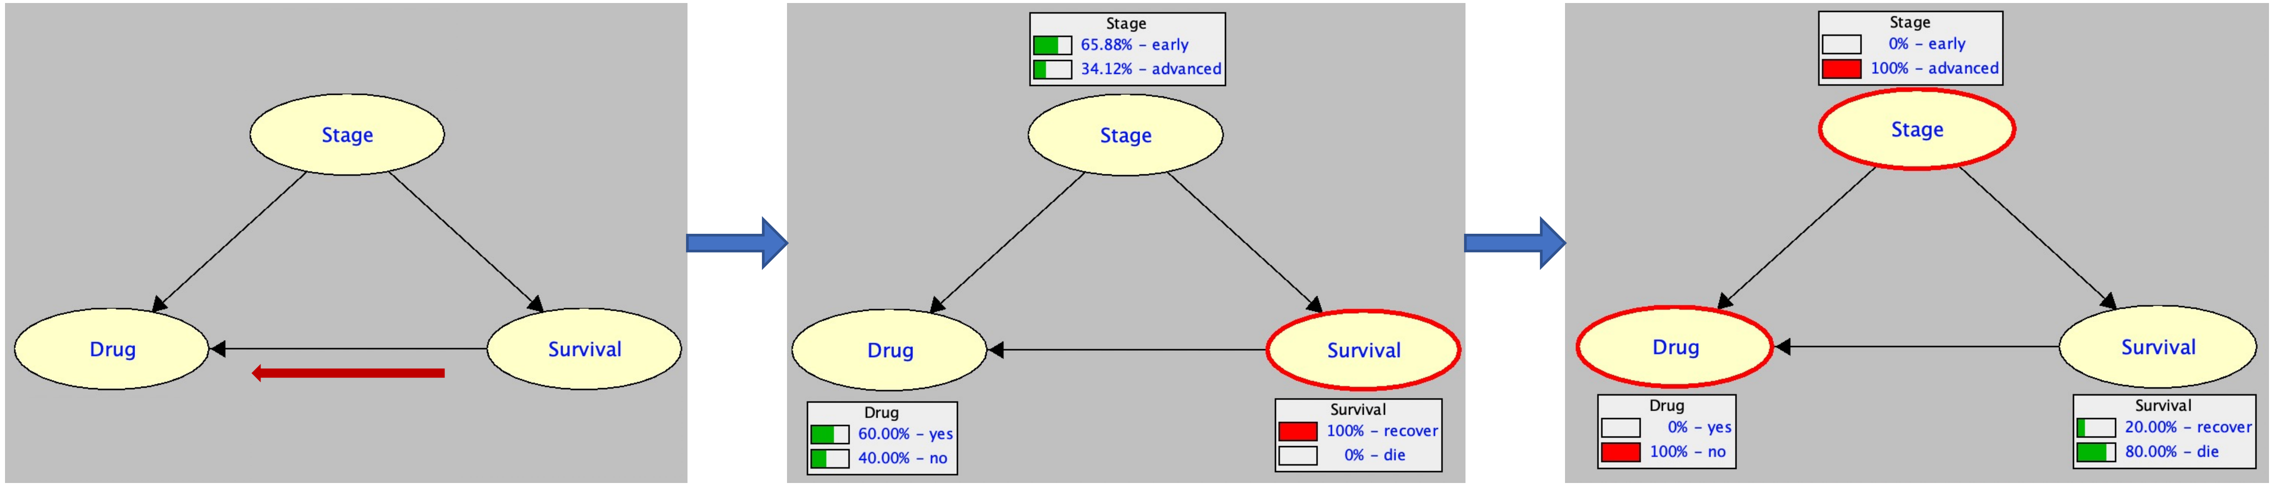

What Additional Information Do We Need?

model 1

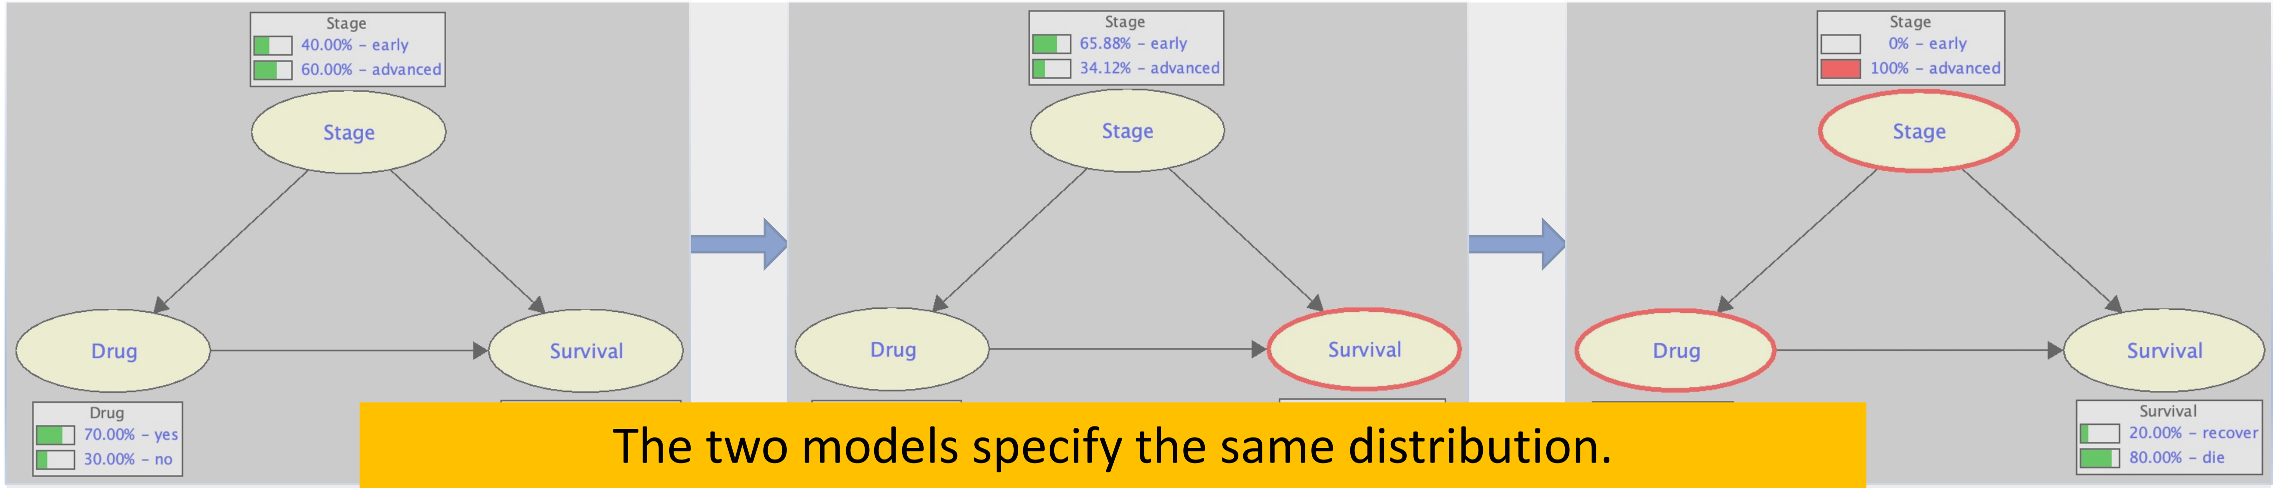

The two models specify the same distribution.

They are **equivalent** as far as answering associational queries.

They are **not equivalent** as far as answering interventional queries.

model 2

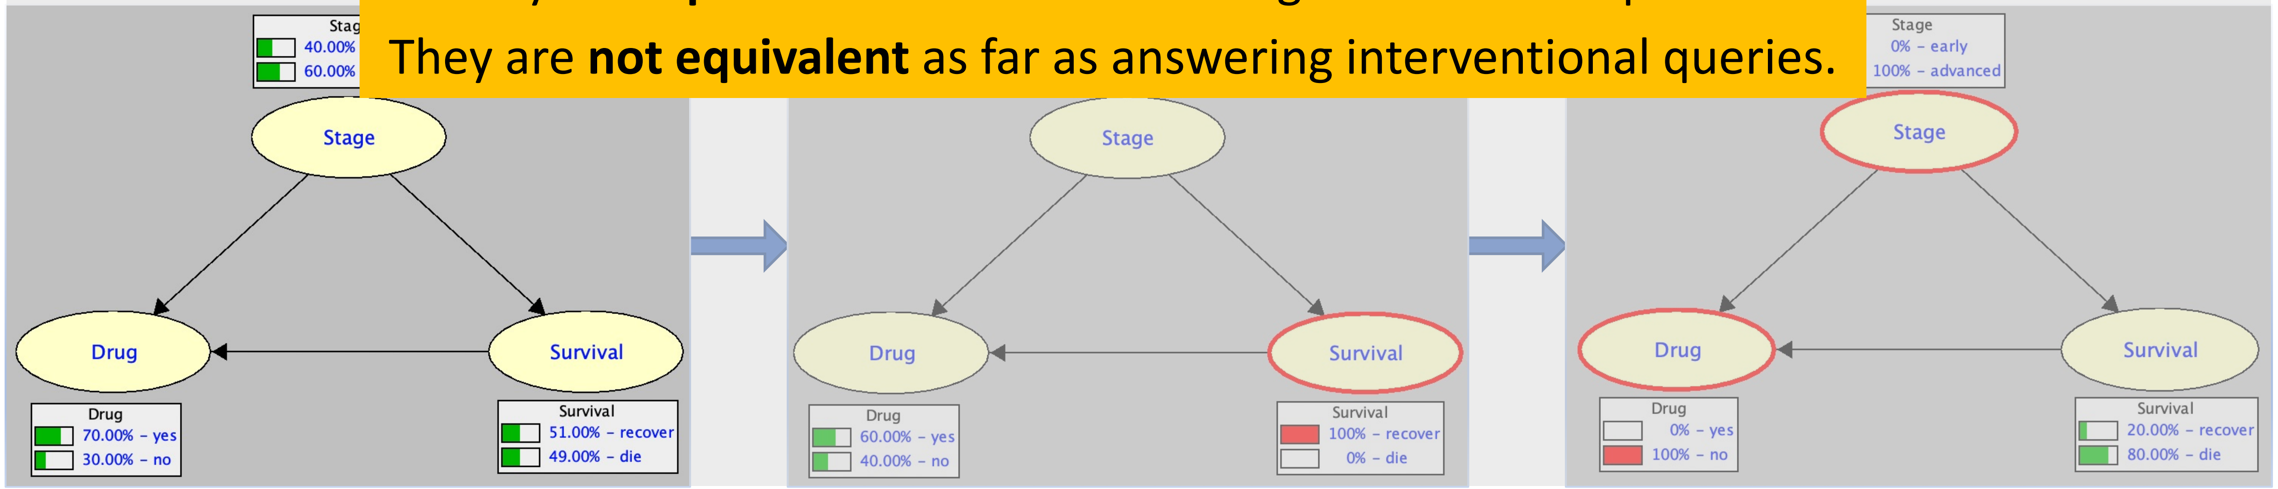

Data Is Not Enough!

- The previous two models specify the same distribution
- They will generate the same data
- We cannot tell which model generated the data
- Using data only, we cannot answer interventional queries
- We need further information: The causal graph!

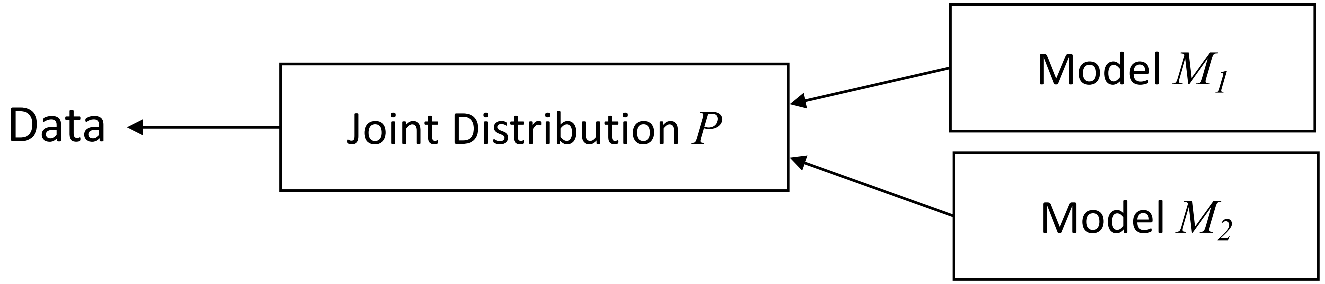

How Else Can We Do It? causal effect

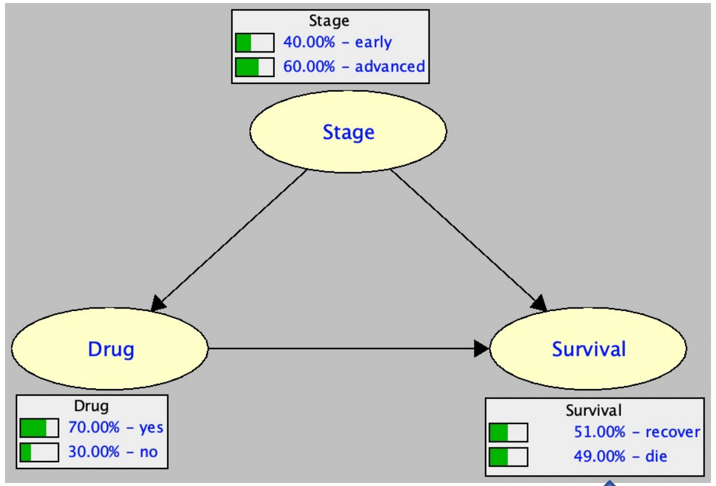

method 1

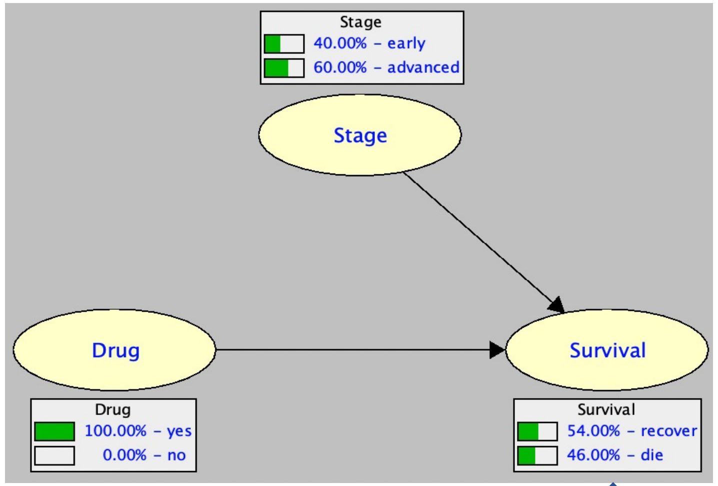

method 2

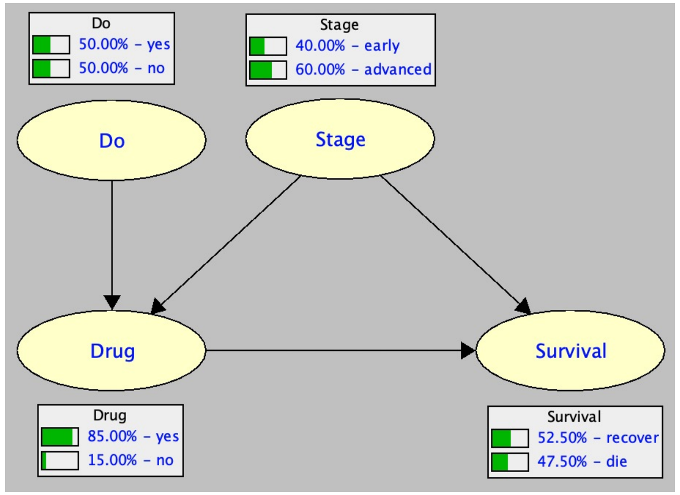

Original CPT: $Pr(Drug|Stage)$. New CPT: $Pr'(Drug|Stage, Do)$.

$$
\begin{aligned}
Pr'(Drug{=}d \mid Stage{=}s, Do{=}no) &= Pr(Drug{=}d \mid Stage{=}s) \\
Pr'(Drug{=}yes \mid Stage{=}s, Do{=}yes) &= 1 \\
Pr'(Drug{=}no \mid Stage{=}s, Do{=}yes) &= 0
\end{aligned}
$$

How Else Can We Do It? an algebraic view

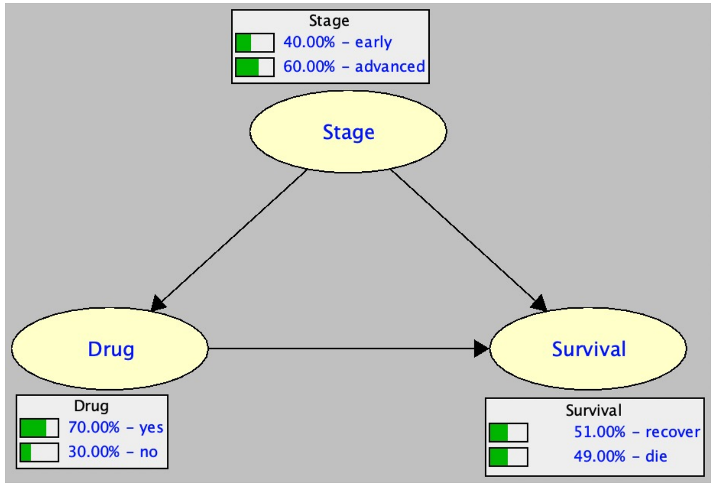

method 1

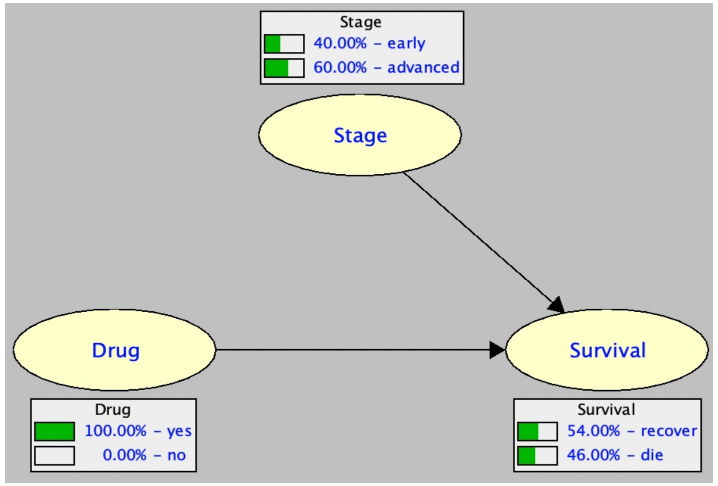

Truncated Formula for **Interventional Distribution**

$$Pr(Survival, Drug, Stage) = Pr(Survival|Drug, Stage)Pr(Drug|Stage)Pr(Stage)$$
$$Pr_{Drug=yes}(Survival, Stage) = Pr(Survival|Drug{=}yes, Stage)Pr(Stage)$$

How Can We Do It? complete model

- Surgery
- Auxiliary do-node
- Truncated formula

Surgery

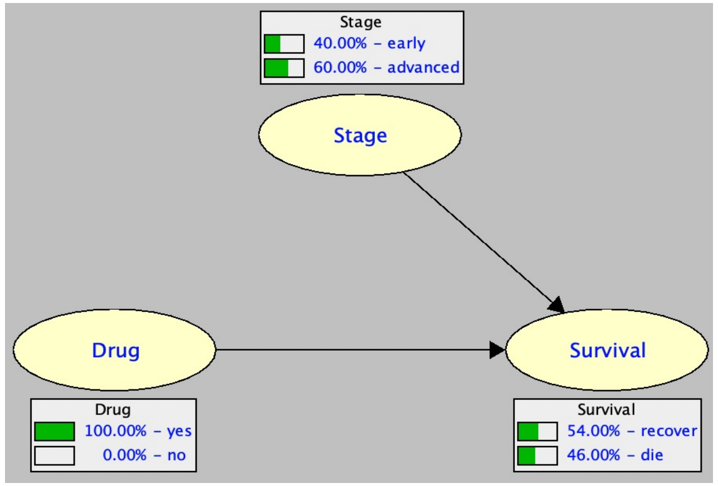

Auxiliary do-node

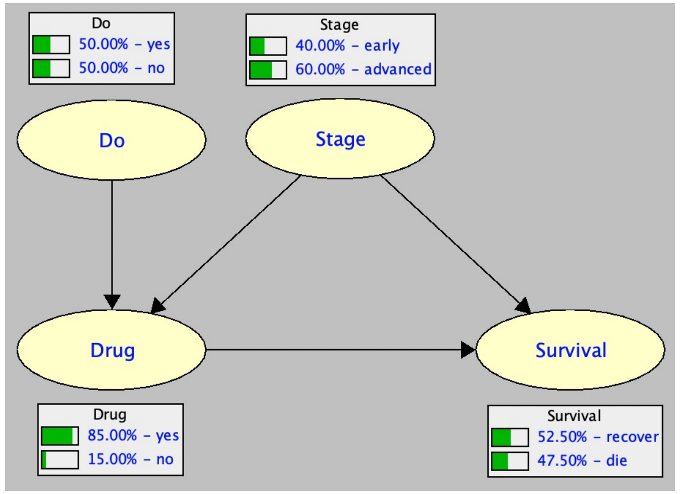

Truncated Formula for **Interventional Distribution**

$$Pr(Survival, Drug, Stage) = Pr(Survival|Drug, Stage)Pr(Drug|Stage)Pr(Stage)$$

$$Pr_{Drug=yes}(Survival, Stage) = Pr(Survival|Drug=yes, Stage)Pr(Stage)$$

Notation

Causal Effect (CE) of $X = x$ on $Y = y$

$$Pr(Y = y | do(X = x))$$

$$Pr(y | do(x))$$

$$Pr(y_x)$$

Distribution

$$Pr(X, Y, Z)$$

Interventional Distribution for $do(X = x)$

$$Pr_{X=x}(Y, Z)$$

How Else Can We Do It?

Causal Effect Rule

$$Pr(y|do(x)) = Pr(y_x) =$$

expressing interventional probabilities using associational probabilities

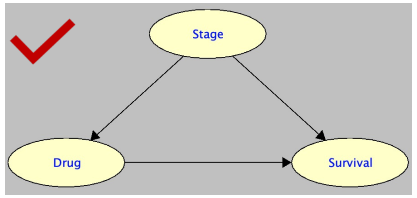

catch: you need to know the parents of X

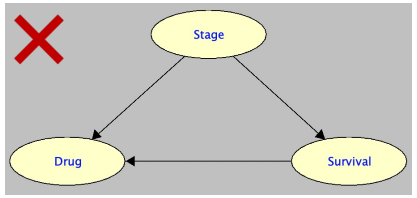

How Can We Do It? partial model + data

Input:

- Causal graph (no parameters)
- Data (observational)

Output:

- Causal effect

Key observations:

- This may not always be possible
- When it is possible, we say the causal effect is **identifiable**
- Identifiability depends on **type of causal graph** and **available data**
- Several criteria for deciding identifiability: some are **complete**, some are not

Example Input-Output

input

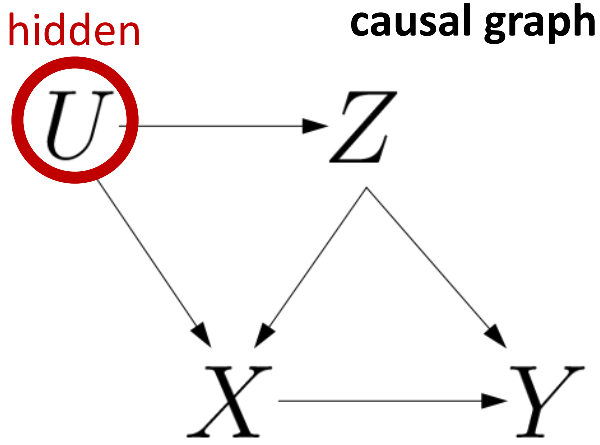

causal graph

data

X	Y	Z	% of population
Yes	Yes	Male	0.116
Yes	Yes	Female	0.274
Yes	No	Male	0.009
Yes	No	Female	0.101
No	Yes	Male	0.334
No	Yes	Female	0.079
No	No	Male	0.051
No	No	Female	0.036

output

$Pr(y_x)$ identifiable

$Pr(x_z)$ not identifiable

notation

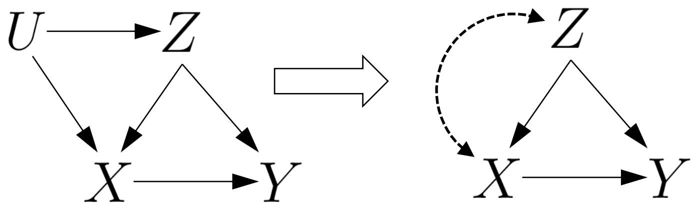

variables U and U_i usually denote hidden variables

Types of Data

- **Observational** data
  - Variables are either hidden or fully observed
  - Variables are either hidden, fully observed or partially observed
- **Interventional** data

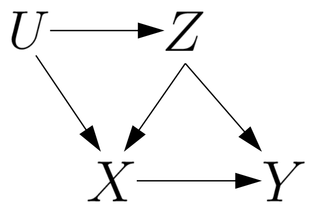

data variable U is hidden

X	Y	Z	% of population
Yes	Yes	Male	0.116
Yes	Yes	Female	0.274
Yes	No	Male	0.009
Yes	No	Female	0.101
No	Yes	Male	0.334
No	Yes	Female	0.079
No	No	Male	0.051
No	No	Female	0.036

Why CE May Not Be Identifiable?

Consider the model

If X and Y are off, cannot tell if X turned off Y or if U turned off both

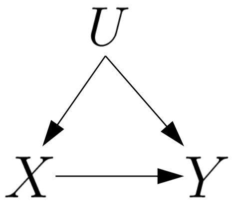

- $Pr(y|x,u) = Pr(\bar{y}|\bar{x},u) = 1$ and $Pr(x|u) = q$
- $Pr(\bar{x}|\bar{u}) = Pr(\bar{y}|X,\bar{u}) = 1$
- $Pr(u) = p$

Causal effect (using truncated formula of interventional distribution)

$$\begin{aligned} Pr(y|do(x)) &= Pr_x(y) \\ &= Pr_x(u,y) + Pr_x(\bar{u},y) \\ &= Pr(u)Pr(y|x,u) + Pr(\bar{u})Pr(y|x,\bar{u}) \\ &= p(1) + (1-p)(0) = p \end{aligned}$$

Data generated by the model

U	X	Y	frequency
u	x	y	pq
u	$\bar{x}$	$\bar{y}$	$p(1-q)$
$\bar{u}$	$\bar{x}$	$\bar{y}$	$1-p$

If *U* is hidden, we see this data

X	Y	frequency
x	y	pq
$\bar{x}$	$\bar{y}$	$1-pq$

Cannot recover the causal effect…

Example: if *pq=.14 then p=.7, q=.2* and *p=.2, q=.7* are solutions, but with different causal effects

both *p=.7, q=.2* and *p=.2, q=.7* are maximum-likelihood parameters

Types of Causal Graphs hidden variables are roots

Markovian Model

each hidden variable has **at most** one child

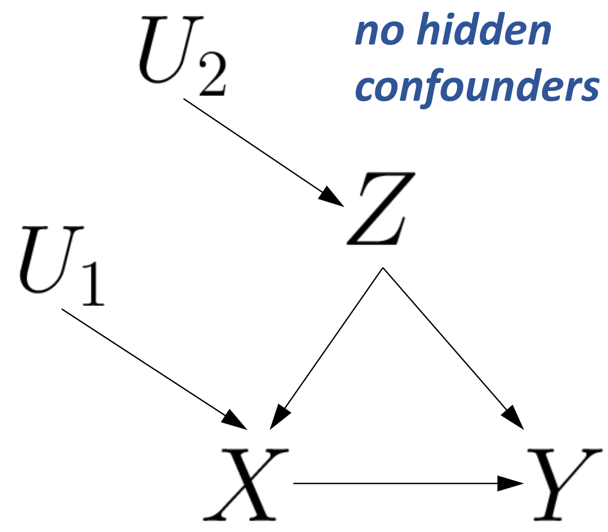

causal effect **always** identifiable

Semi-Markovian Model

some hidden variable has **more than** one child

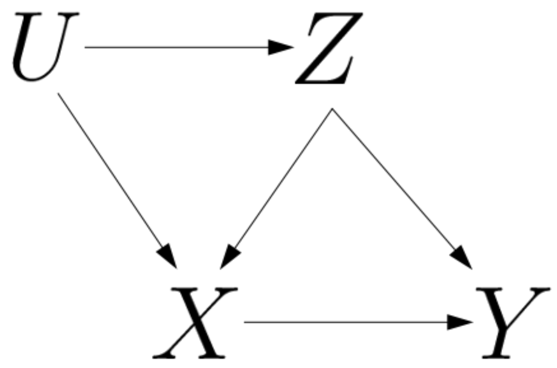

$Pr(y_x)$ identifiable $Pr(x_z)$ not identifiable

causal effect **not always** identifiable

Identifiability Criteria identifiability tests

- Causal Effect Rule (incomplete)
- Backdoor Criteria (incomplete)
- Frontdoor Criteria (incomplete)
- Do-Calculus (complete)

- There are other criteria…
  some use additional information like context-specific independence, or functional dependencies

Causal Effect Rule

Let G be a causal graph which contains variables X and Y and let $\mathbf{Z}$ be the parents of X in G. Then

$$Pr(y|do(x)) = \sum_{\mathbf{z}} Pr(y|x, \mathbf{z}) Pr(\mathbf{z})$$

causal effect is identifiable if we have $Pr(X, Y, \mathbf{Z})$

causal effect is identifiable if X, Y, $\mathbf{Z}$ are observed

Example: Causal Effect Rule

Let G be a causal graph which contains variables X and Y and let $\mathbf{Z}$ be the parents of X in G. Then

$$Pr(y|do(x)) = \sum_{\mathbf{z}} Pr(y|x,\mathbf{z})Pr(\mathbf{z})$$

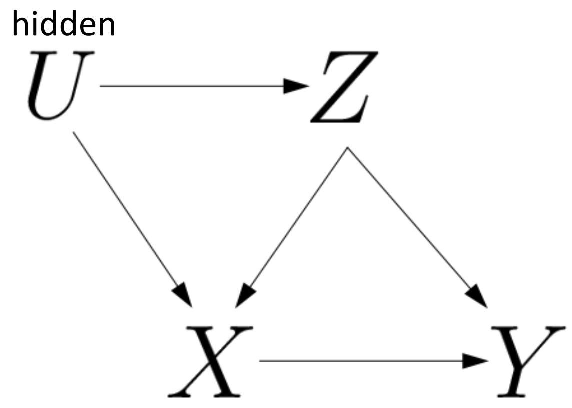

$Pr(x_z)$	cannot use CER	**not identifiable**
$Pr(y_x)$	cannot use CER	**identifiable**

Backdoor Criteria d-separation review

Is the path blocked by variable $\boldsymbol{Z}$?

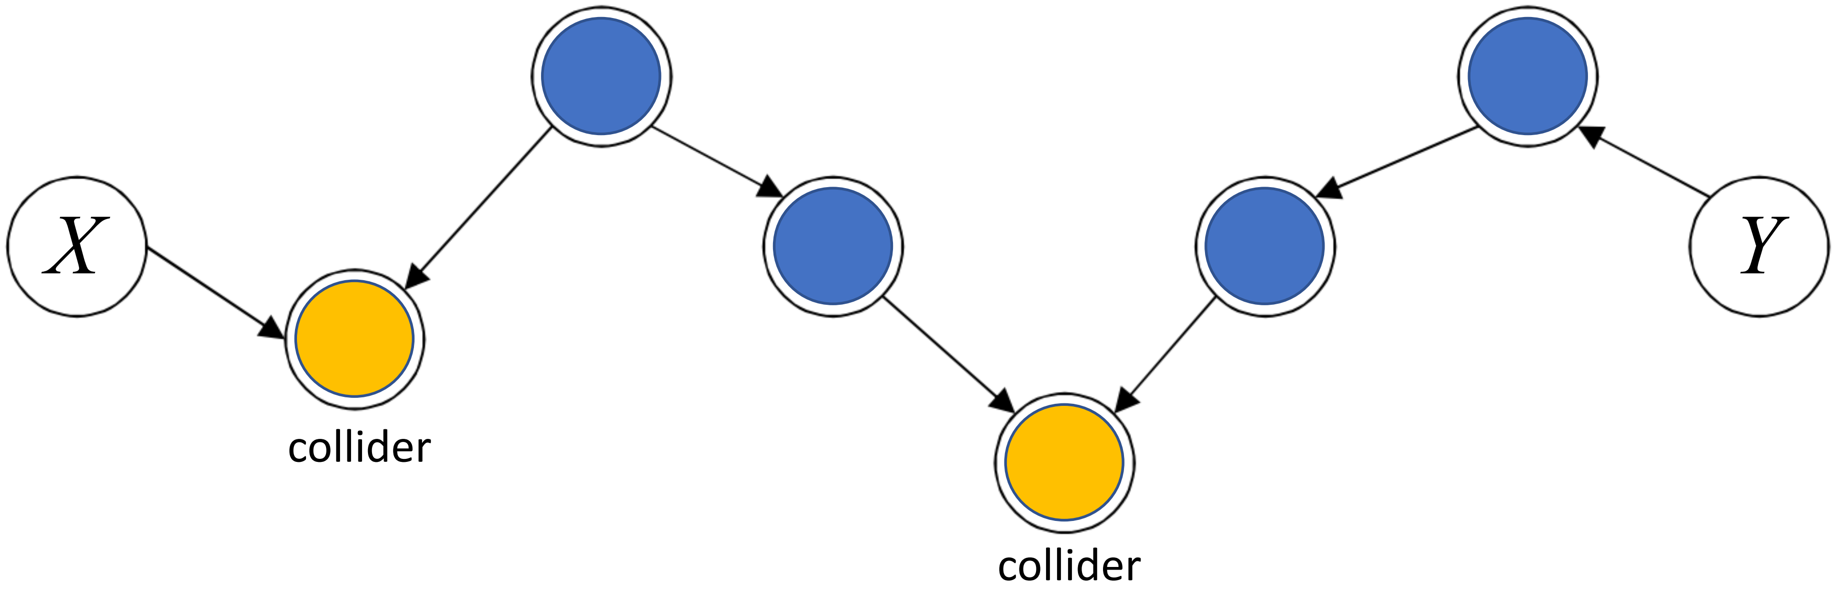

Path between X and Y is **blocked** by $\boldsymbol{Z}$ iff:

- some non-collider is in $\boldsymbol{Z}$, or
- some collider and none of its descendants are not in $\boldsymbol{Z}$

Backdoor Criteria

Consider a causal graph G and causal effect $Pr(y_x)$. A set of variables $\mathbf{Z}$ satisfies the backdoor criteria iff
(1) no node in $\mathbf{Z}$ is a descendant of X
(2) $\mathbf{Z}$ blocks every path between X and Y that contains an arrow into X

If $\mathbf{Z}$ is a backdoor, then

$$Pr(y_x) = \sum_{\mathbf{z}} Pr(y|x, \mathbf{z}) Pr(\mathbf{z})$$

interventional   associational

Example: Backdoor

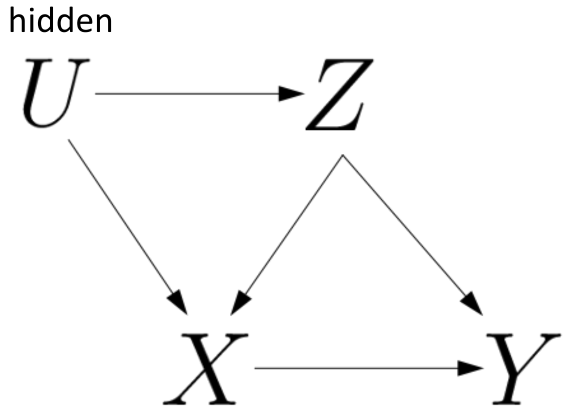

$Pr(y_x)$ cannot use CER **identifiable**

$Pr(x_z)$ cannot use CER **not identifiable**

$Pr(y_x)$ Z is a backdoor Paths $Y \leftarrow Z \rightarrow X$ and $Y \leftarrow Z \leftarrow U \rightarrow X$ are blocked by Z

$Pr(x_z)$ no backdoor Path $Z \leftarrow U \rightarrow X$ cannot be blocked by observed variables

Example: Backdoor

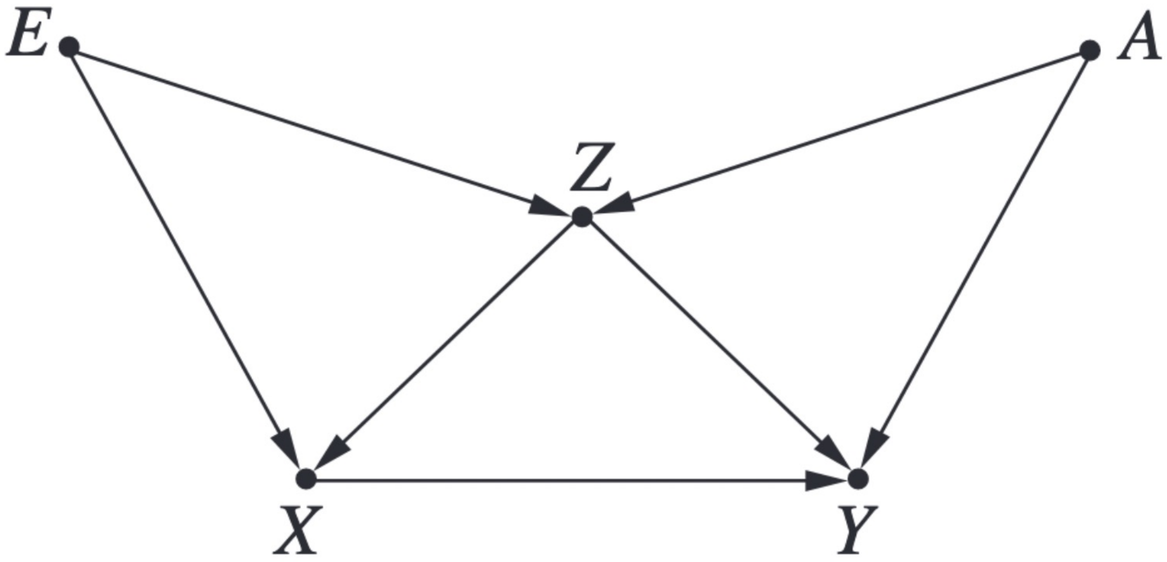

find backdoor for causal effect of X on Y

all variables are observed (no hidden variables)

Example: Backdoor

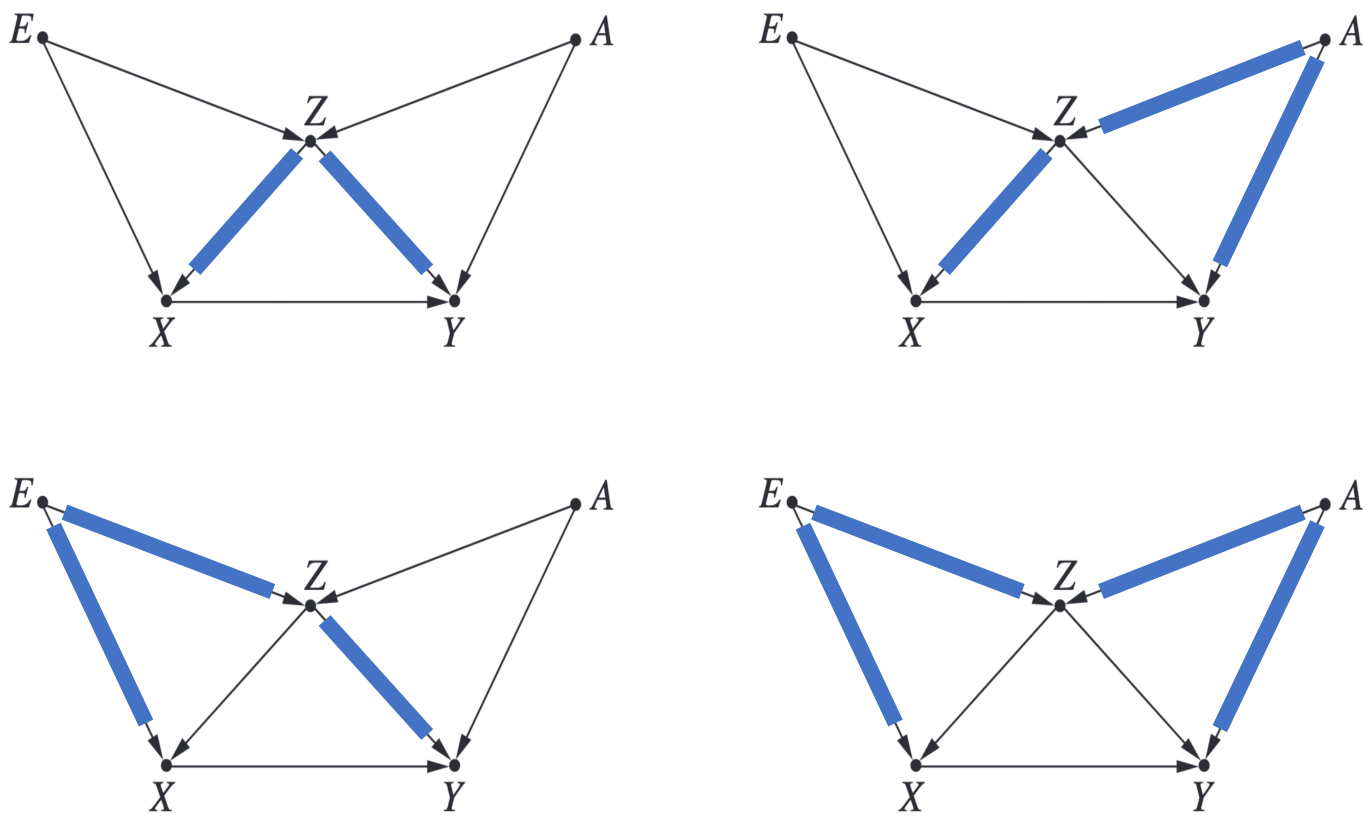

backdoor 1: *A, Z*

backdoor 2: *E, Z*

backdoor 3: *A, E, Z*

enumerating backdoor paths

Frontdoor Criteria

Consider a causal graph G and causal effect $Pr(y_x)$. A set of variables $\mathbf{Z}$ satisfies the frontdoor criteria iff

⋮

If $\mathbf{Z}$ is a frontdoor, then

$$Pr(y_x) = \sum_{\mathbf{z}} Pr(\mathbf{z}|x) \sum_{x'} Pr(y|x', \mathbf{z}) Pr(x')$$

interventional associational

Incompleteness of Backdoor Criterion

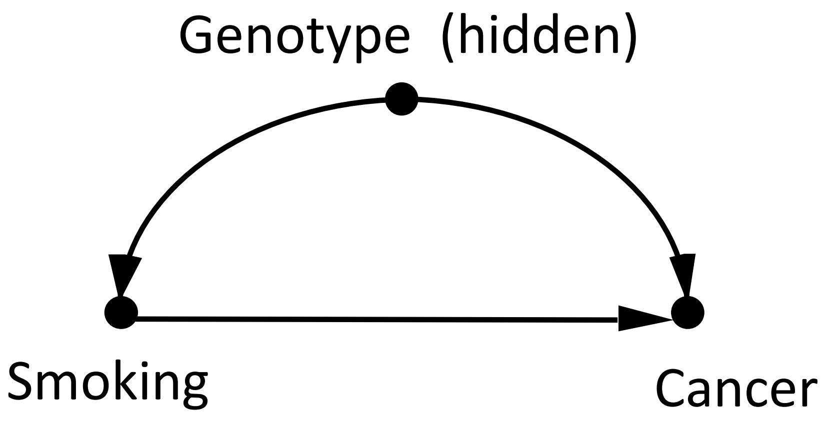

no backdoor

causal effect **is not identifiable**

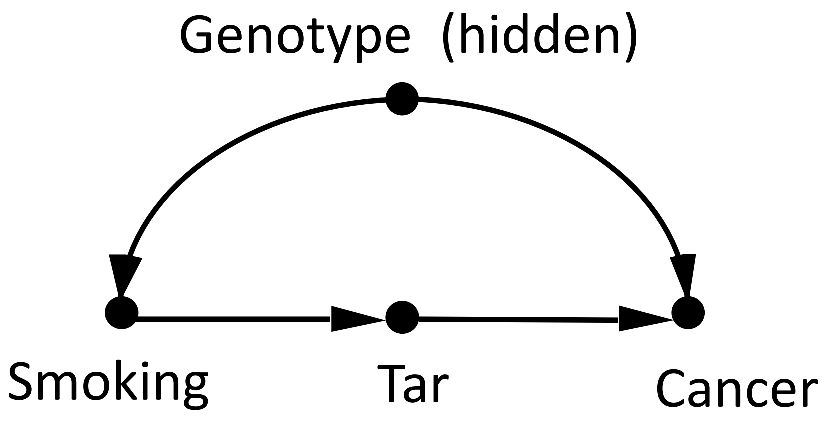

no backdoor

causal effect **is identifiable!**

causal effect of smoking on cancer $Pr(c|do(s))$

The Do-Calculus positivity assumption

$\mathbf{X}$, $\mathbf{Y}$, $\mathbf{Z}$, $\mathbf{W}$ are disjoint sets of variables

Rule 1: Ignoring observations.

$$Pr(\mathbf{y}|do(\mathbf{x}), \mathbf{z}, \mathbf{w}) = Pr(\mathbf{y}|do(\mathbf{x}), \mathbf{w}) \text{ if } dsep(\mathbf{Y}, \mathbf{XW}, \mathbf{Z})_{G_{\overline{\mathbf{X}}}}$$

Rule 2: Action/observation exchange.

$$Pr(\mathbf{y}|do(\mathbf{x}), do(\mathbf{z}), \mathbf{w}) = Pr(\mathbf{y}|do(\mathbf{x}), \mathbf{z}, \mathbf{w}) \text{ if } dsep(\mathbf{Y}, \mathbf{XW}, \mathbf{Z})_{G_{\overline{\mathbf{X}}\underline{\mathbf{Z}}}}$$

Rule 3: Ignoring actions.

$$Pr(\mathbf{y}|do(\mathbf{x}), do(\mathbf{z}), \mathbf{w}) = Pr(\mathbf{y}|do(\mathbf{x}), \mathbf{w}) \text{ if } dsep(\mathbf{Y}, \mathbf{XW}, \mathbf{Z})_{G_{\overline{\mathbf{X}}\overline{\mathbf{Z}(\mathbf{W})}}}$$

$G_{\overline{\mathbf{X}}}$: graph G after removing edges pointing into variables $\mathbf{X}$

$G_{\underline{\mathbf{X}}}$: graph G after removing edges outgoing from variables $\mathbf{X}$

The Do-Calculus: Example Derivation

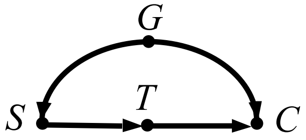

$Pr(c|do(s))$ is identifiable **there is a polytime algorithm (ID) for the do-calculus**

$$P(c \mid do(s)) = \sum_t P(c \mid do(s), t) P(t \mid do(s))$$

Probability Axioms

$$= \sum_t P(c \mid do(s), do(t)) P(t \mid do(s))$$

Rule 2 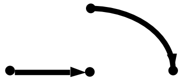

$$= \sum_t P(c \mid do(s), do(t)) P(t \mid s)$$

Rule 2 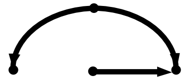

$$= \sum_t P(c \mid do(t) P(t \mid s)$$

Rule 3 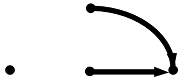

$$= \sum_{s'} \sum_t P(c \mid do(t), s') P(s' \mid do(t)) P(t \mid s)$$

Probability Axioms

$$= \sum_{s'} \sum_t P(c \mid t, s') P(s' \mid do(t)) P(t \mid s)$$

Rule 2 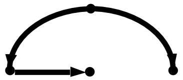

$$= \sum_{s'} \sum_t P(c \mid t, s') P(s') P(t \mid s)$$

Rule 3 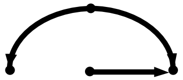

Concluding Remarks

- Causal hierarchy has three layers
  1) Associational reasoning
  2) Interventional reasoning
  3) Counterfactual reasoning
- A distribution is not enough for layers 2 and 3
  (we need a causal graph)
- Data is not enough for layers 2 and 3
- **Cross-layer inference** is sometimes feasible
  (e.g., computing interventional probabilities using (partial) data + causal graph)

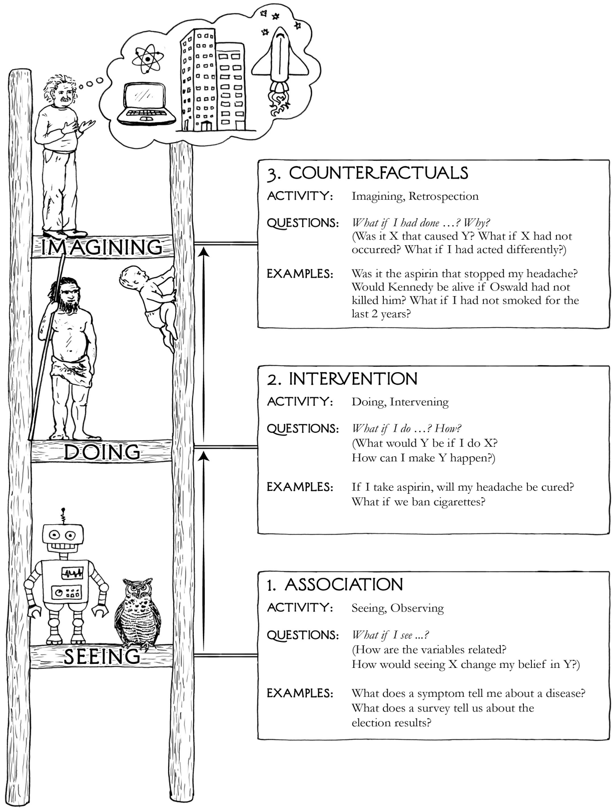

Thank You!