Introduction to Graphical Causal Modelling

Tom Claassen

Institute for Computer and Information Science,

Radboud University Nijmegen

SIKS course on Causal Modelling – 30 May, 2023

Radboud University Nijmegen

Outline

1. **Introduction to causality**
2. Prediction vs. causation
3. Causal graphs and how to read them
4. Cause-effect estimation
5. The missing link & conclusion

Many important research questions are rooted in causality

benefits of exercise and healthy nutrition

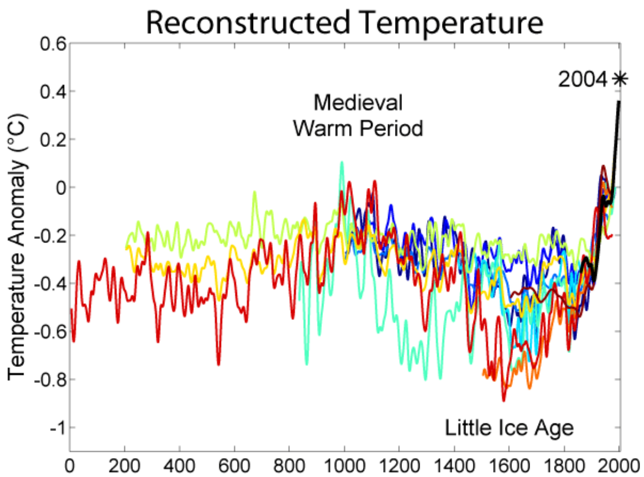

climate change

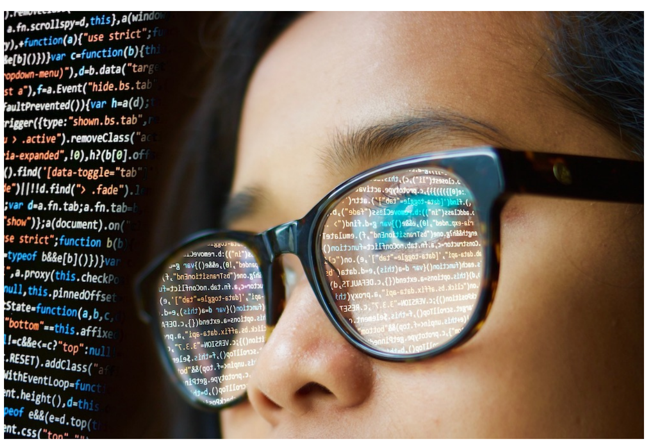

racial and gender bias in AI

Covid vaccine efficacy

Many important research questions are rooted in causality

- Answers typically involve '*what causes it*?' and '*how*?'

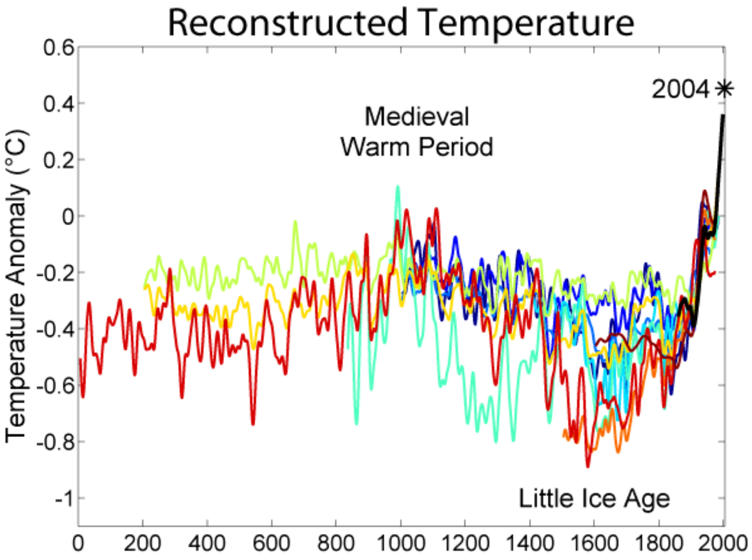

"does human activity cause climate change?"

Many important research questions are rooted in causality

- ‘*Finding a connection*’ does not imply we know what causes what …

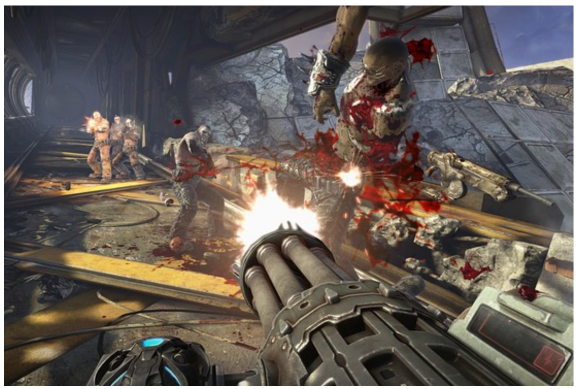

?

“how are violent video games, ADHD, and aggression related?”

Many important research questions are rooted in causality

- Sometimes the difference between '*connection*' and '*causality*' is blurred …

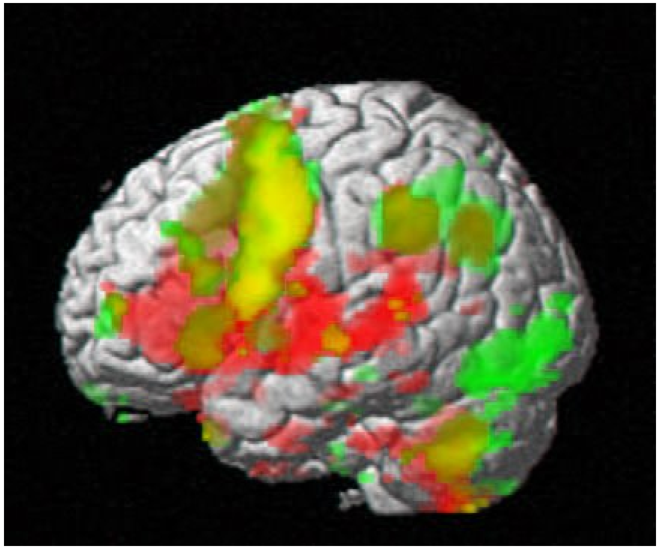

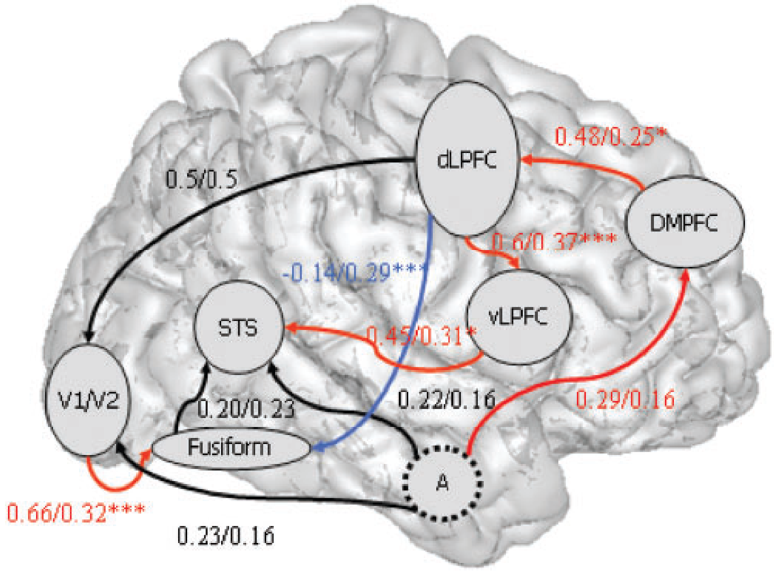

"can we infer functional brain connectivity from fMRI data?"

Many important research questions are rooted in causality

- Sometimes it is not even clear if the concept '*cause*' makes sense

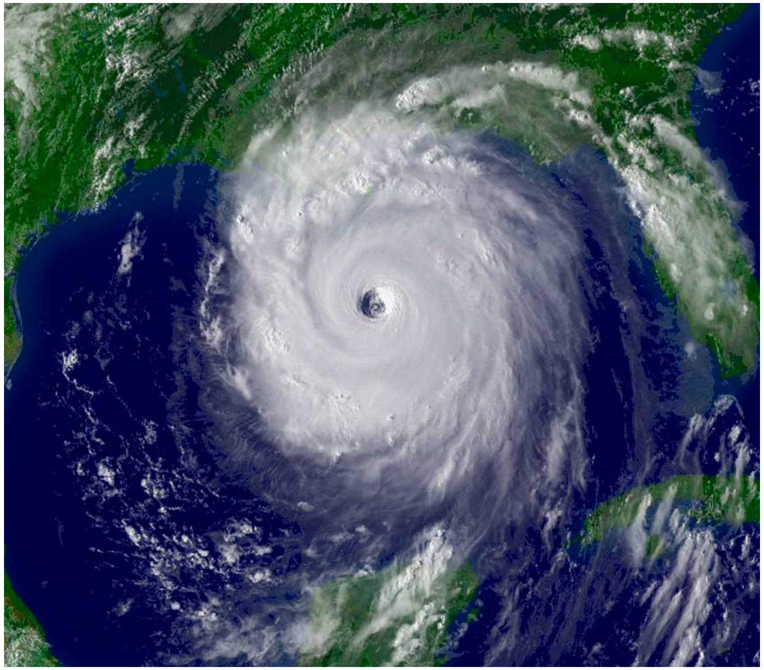

"can butterflies cause hurricanes?"

Causality: what is it?

How do we recognize causality?
(apparently so simple we don't teach this at school/university)
"*Of course I know cause and effect!*"

Causality: what is it?

How do we recognize causality?
(apparently so simple we don't teach this at school/university)
"*Of course I know cause and effect …*"

What exactly do we mean by 'cause' and 'effect'?
Intuitively obvious, yet curiously hard to define. Often involves aspects of

- things that *occur together*
- things that *follow* each other *in time*
- things that are somehow *necessary* and/or *sufficient* to lead to another
- things that *change the probability* of something else happening
- things connected by a *mechanistic chain of events*, etc. etc.

Most definitions run into trouble somewhere …

➡ Main 'cause' behind a huge amount of philosophical controversy!

Hume on causality

The subject of causality has a long history in philosophy. For example this is what Hume had to say about it:

"Thus we remember to have seen that species of object we call *flame*, and to have felt that species of sensation we call *heat*. We likewise call to mind their constant conjunction in all past instances. Without any farther ceremony, we call the one *cause* and the other *effect*, and infer the existence of the one from that of the other."

David Hume, *Treatise of Human Nature* (1739)

Russell on causality

Some philosophers even proposed to abandon the concept of causality altogether

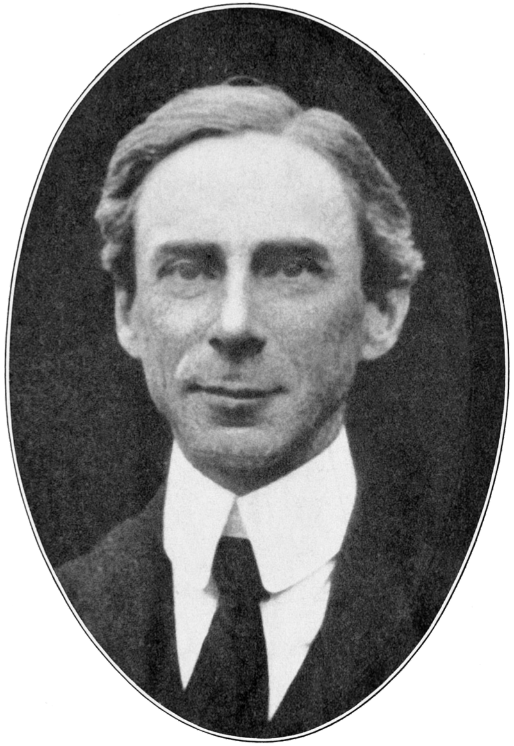

"All philosophers, of every school, imagine that causation is one of the fundamental axioms or postulates of science, yet, oddly enough, in advanced sciences such as gravitational astronomy, the word `cause' never occurs. The law of causality, I believe, like much that passes muster among philosophers, is a relic of a bygone age, surviving, like the monarchy, only because it is erroneously supposed to do no harm".

Bertrand Russell, *On the Notion of Cause* (1913)

Causality in statistics

Karl Pearson (one of the founders of modern statistics, well-known from his work on the *correlation coefficient*) writes:

"Beyond such discarded fundamentals as `matter' and `force' lies still another fetish amidst the inscrutable arcana of even modern science, namely, the category of cause and effect."

Karl Pearson, *The Grammar of Science* (1892)

Causality in statistics

Karl Pearson (one of the founders of modern statistics, well-known from his work on the *correlation coefficient*) writes:

"Beyond such discarded fundamentals as `matter' and `force' lies still another fetish amidst the inscrutable arcana of even modern science, namely, the category of cause and effect."

Karl Pearson, *The Grammar of Science* (1892)

Since then, many statisticians tried to avoid causal reasoning:

- "*Considerations of causality should be treated as they have always been in statistics: preferably not at all*." (Terry Speed, former president of the Biometric Society).
- "*It would be very healthy if more researchers abandon thinking of and using terms such as cause and effect.*" (Prominent social scientist).

Pragmatic approach

Causality = ‘Effective manipulability’

- focus on *relevant, measurable influence*
- understand *why* things happen
- predict how things change if we *intervene* (effect computation)
- not about truth, but about *validity* (given *assumptions*)
- allows principled use of *maths*, *statistics* & *logic* on data and models
- verify by *experiment* (‘Randomized Controlled Trial’)

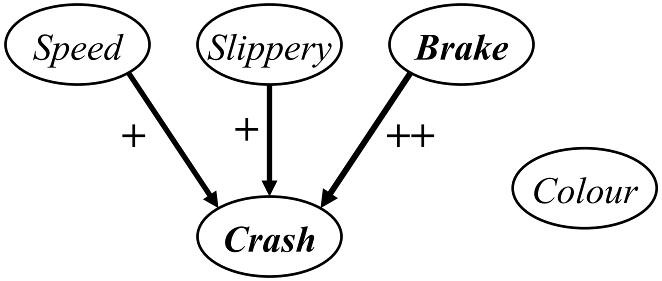

key risk factors in predicting car crashes …

Effective manipulability

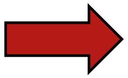

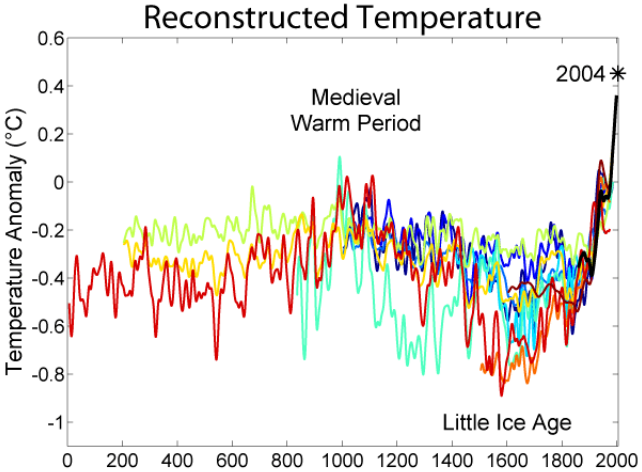

"does human activity cause climate change?"

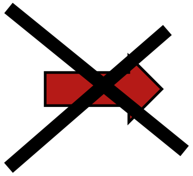

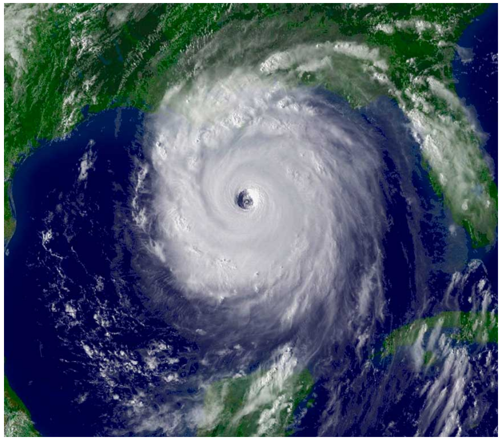

"do butterflies cause hurricanes?"

Outline

1. Introduction to causality
2. **Prediction vs. causation**
3. Causal graphs and how to read them
4. Cause-effect estimation
5. The missing link & conclusion

Example - Gene regulation

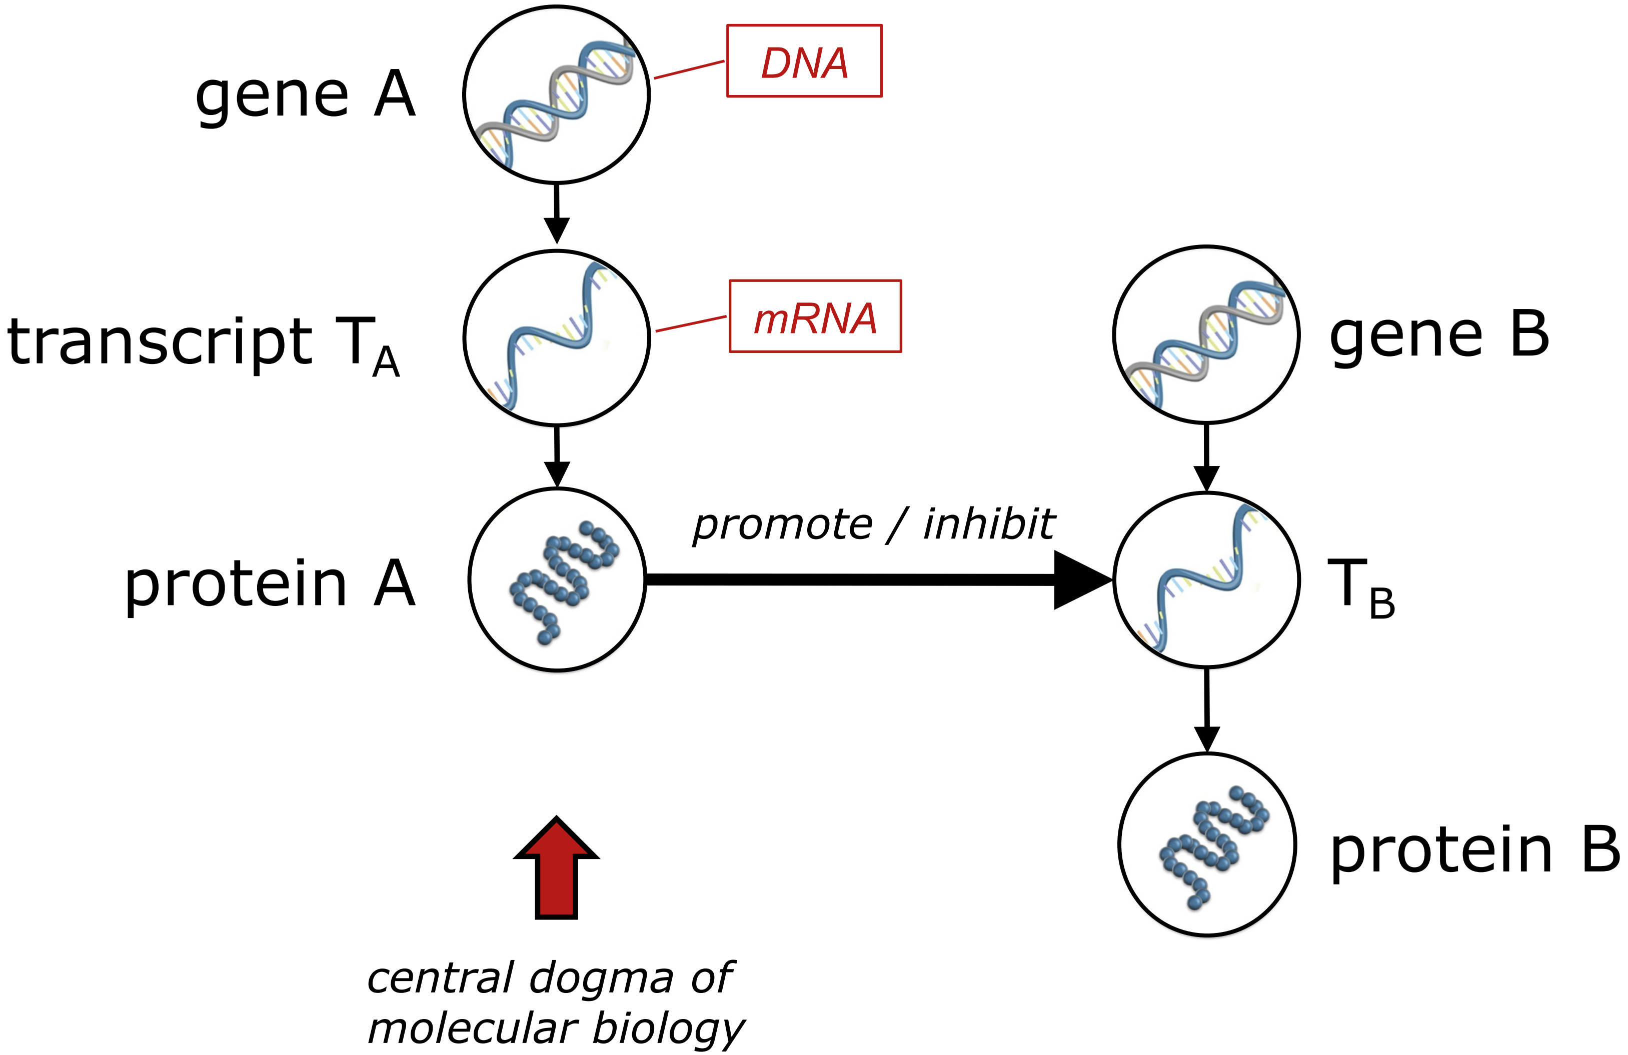

Example - Gene regulation

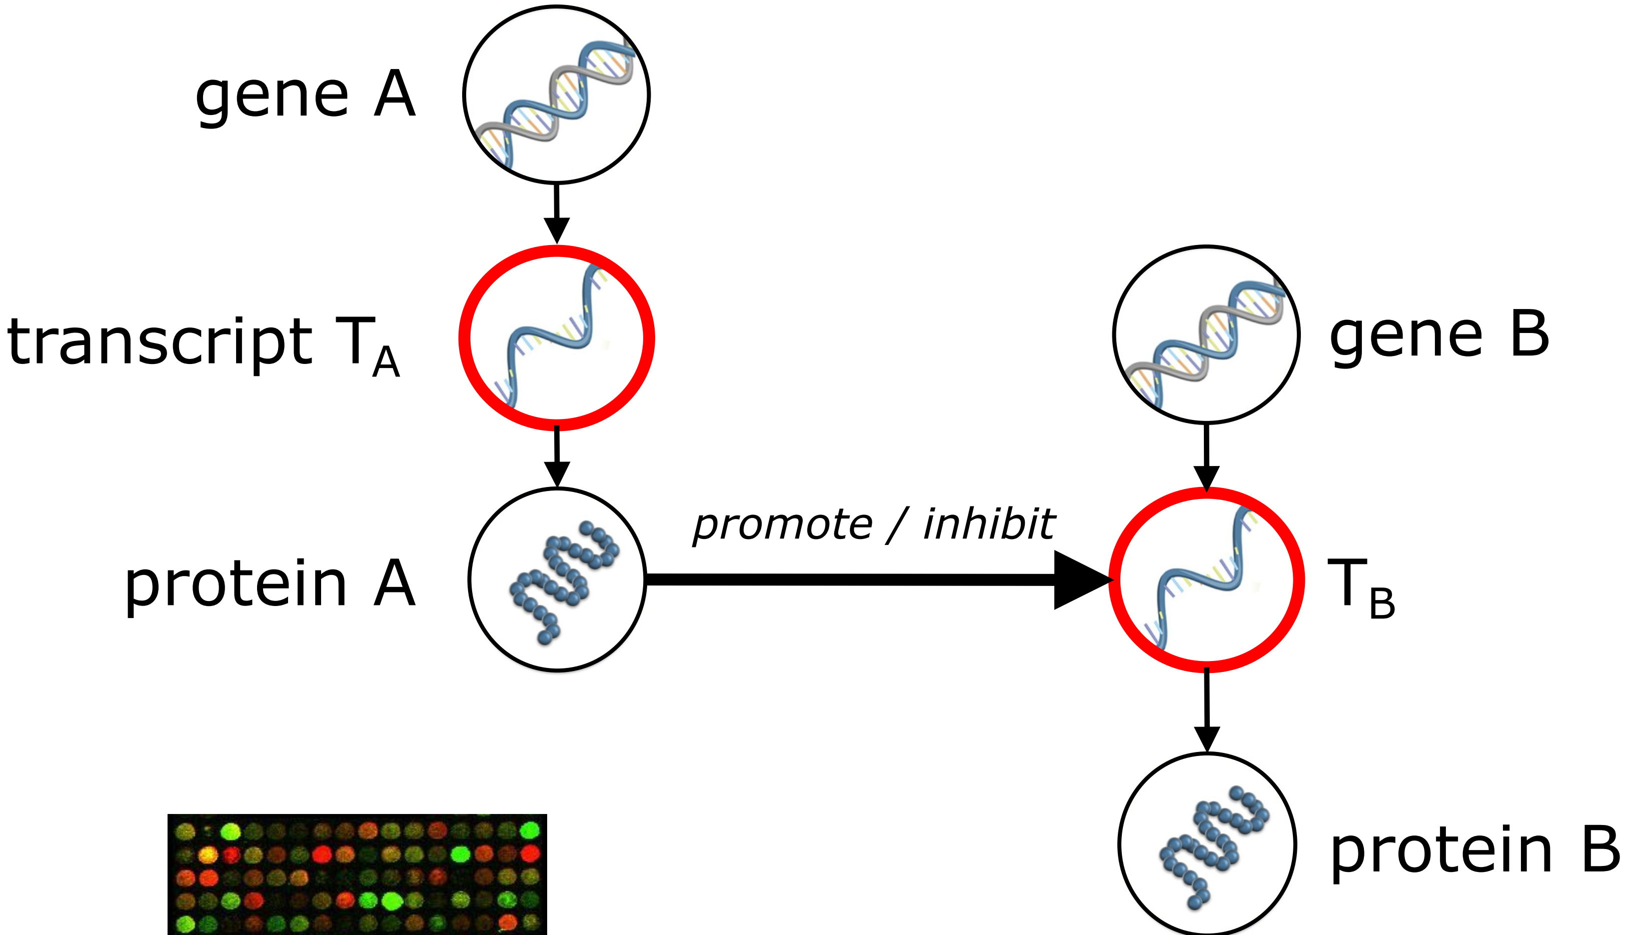

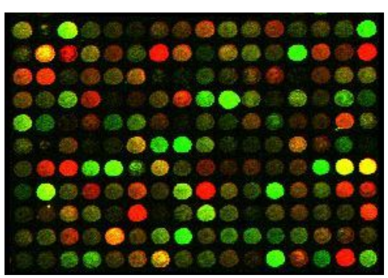

microarray to measure transcription levels

Observed gene expression levels

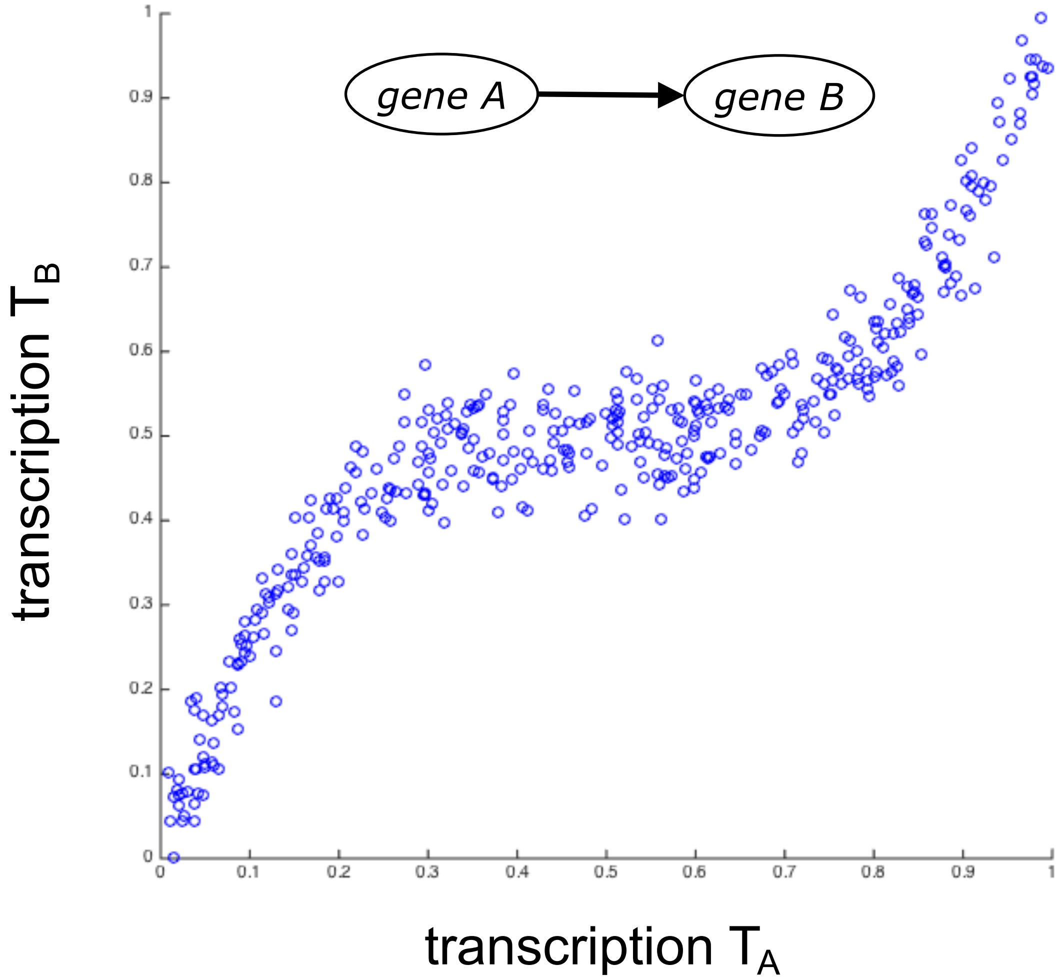

Predicting gene expression levels

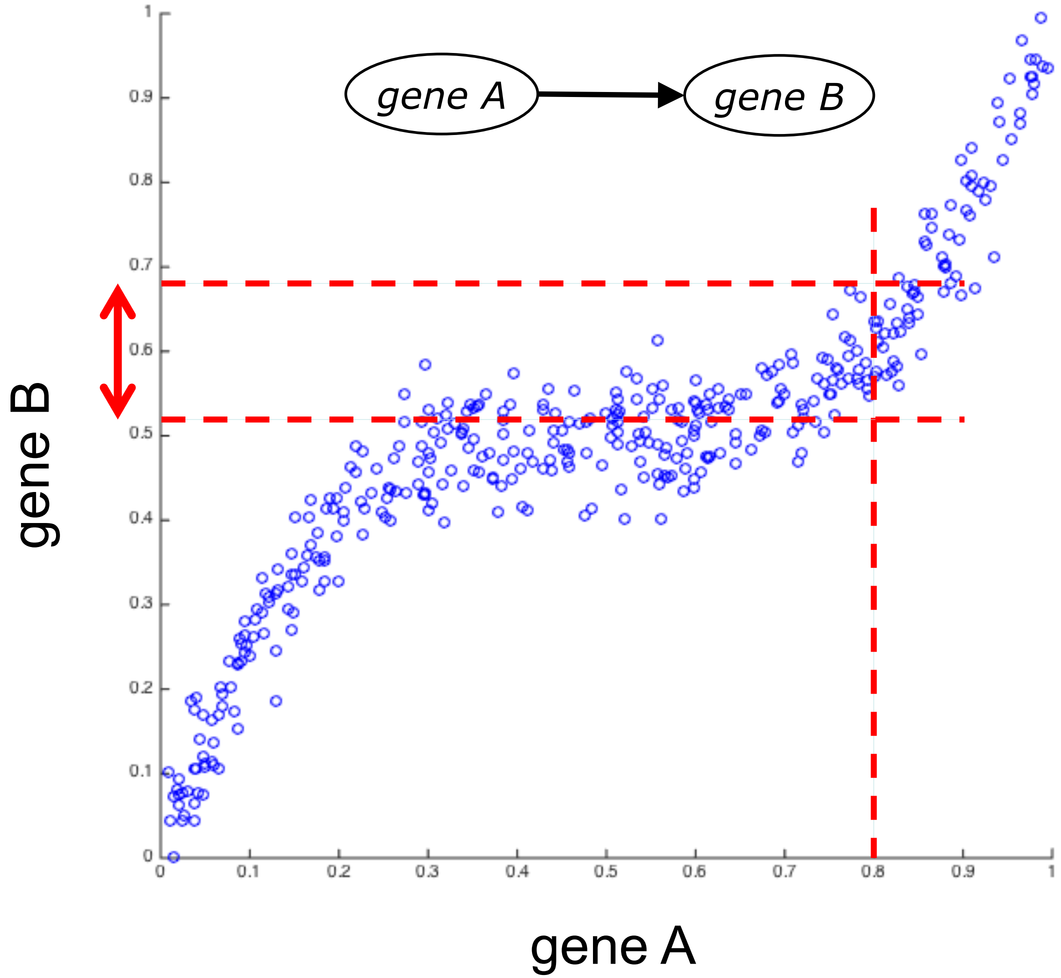

$$p\left(T_B \middle| T_A = 0.8\right)$$

Predicting gene expression levels

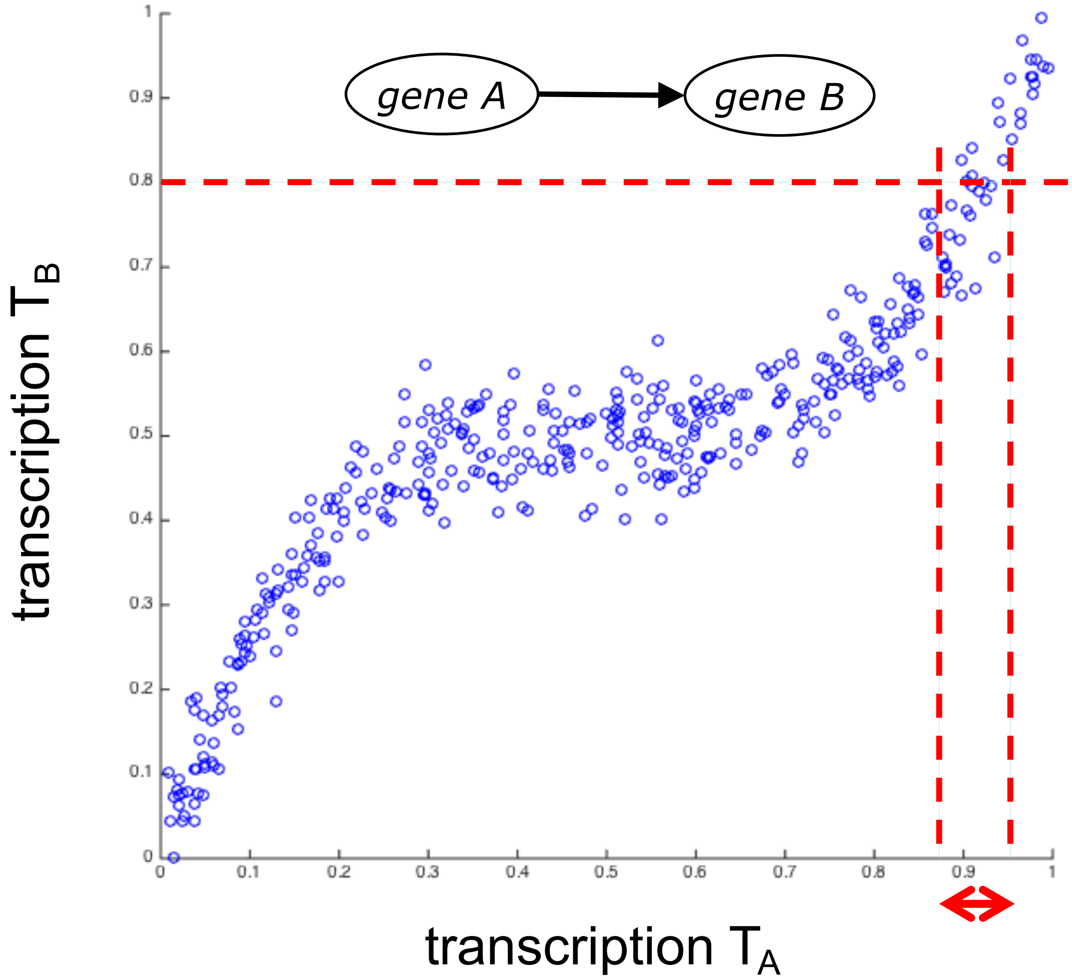

$$p\left(T_A \middle| T_B = 0.8\right)$$

Observation vs. intervention: gene knock-out experiments

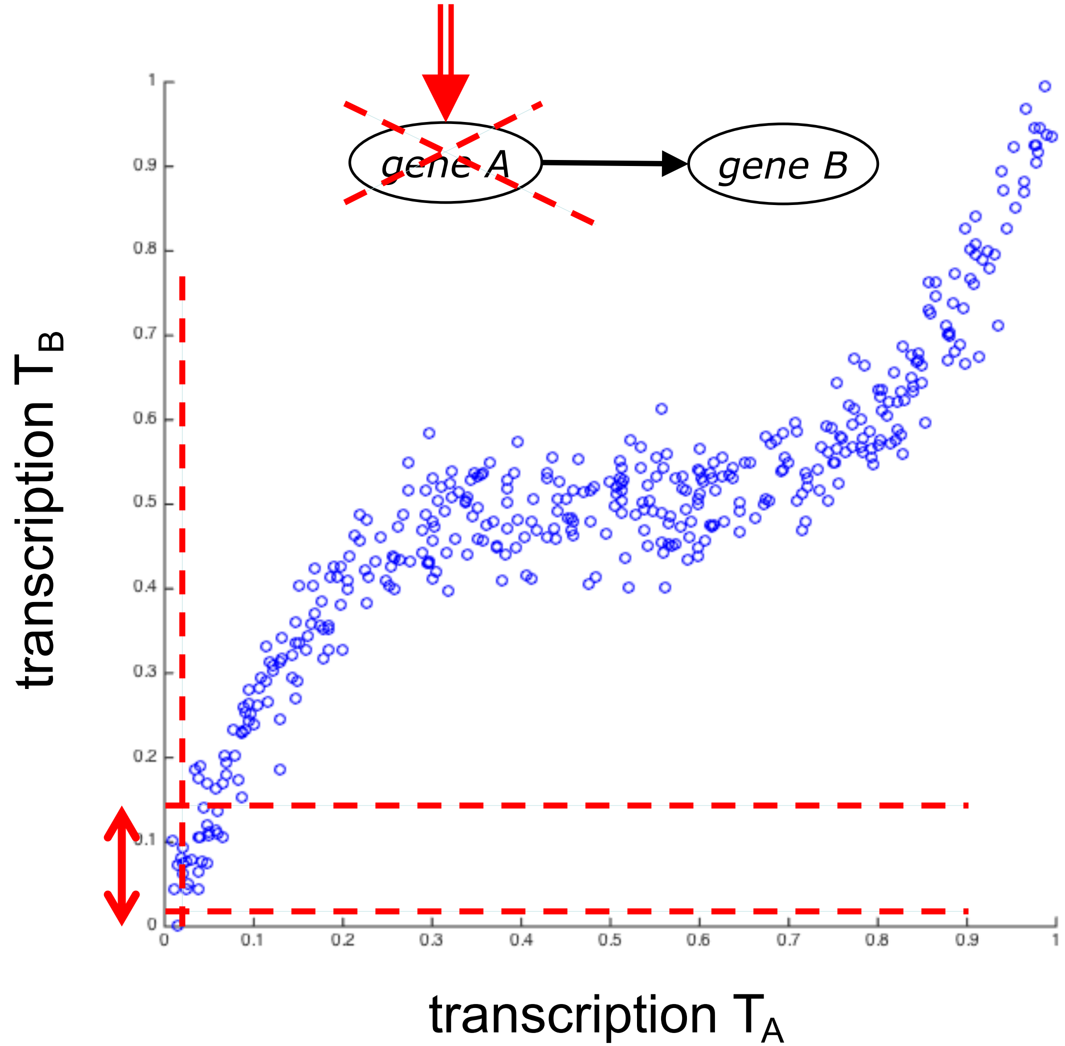

$$p\left(T_B \middle| do\left(T_A = 0.0\right)\right)$$

Observation vs. intervention: gene knock-out experiments

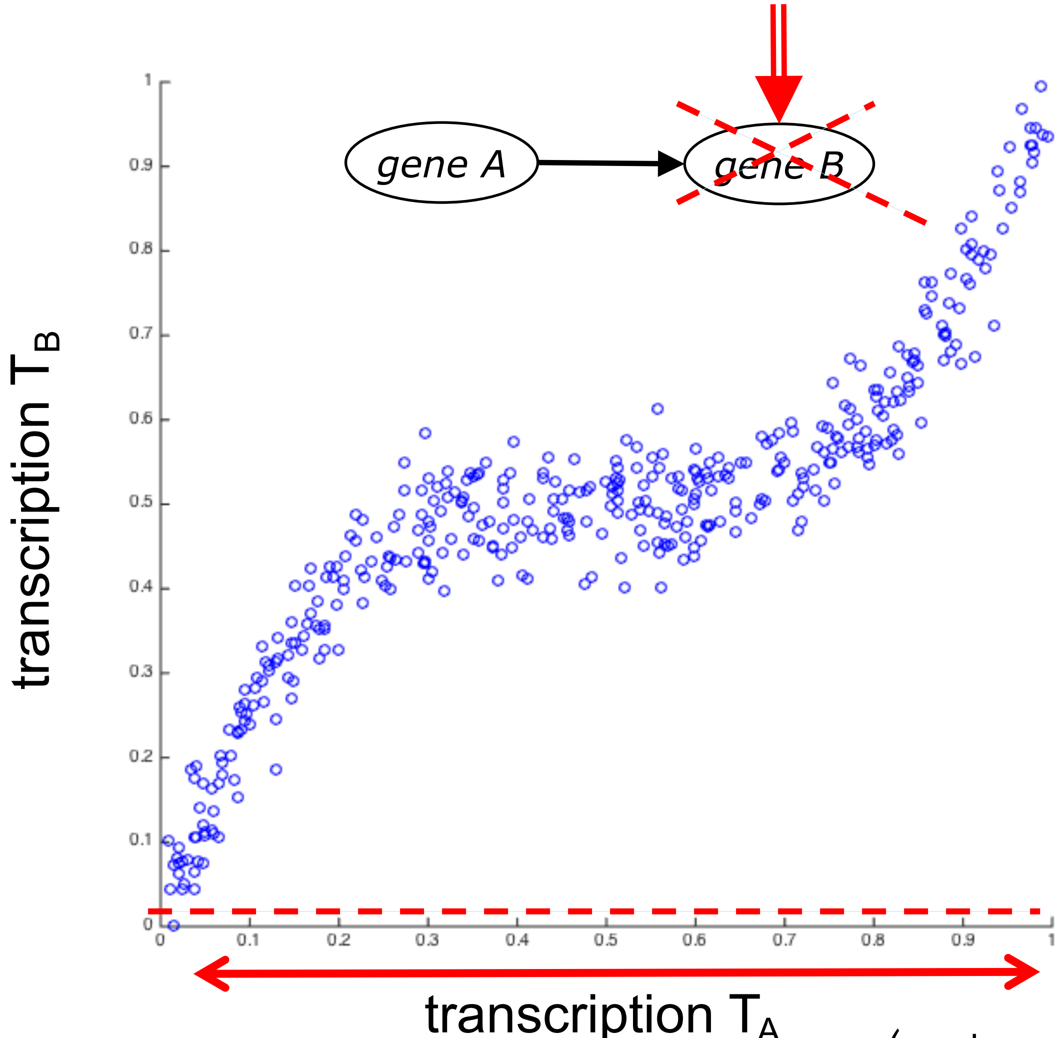

$$p\left(T_A \middle| do\left(T_B = 0.0\right)\right)$$

What if we do not know the model?

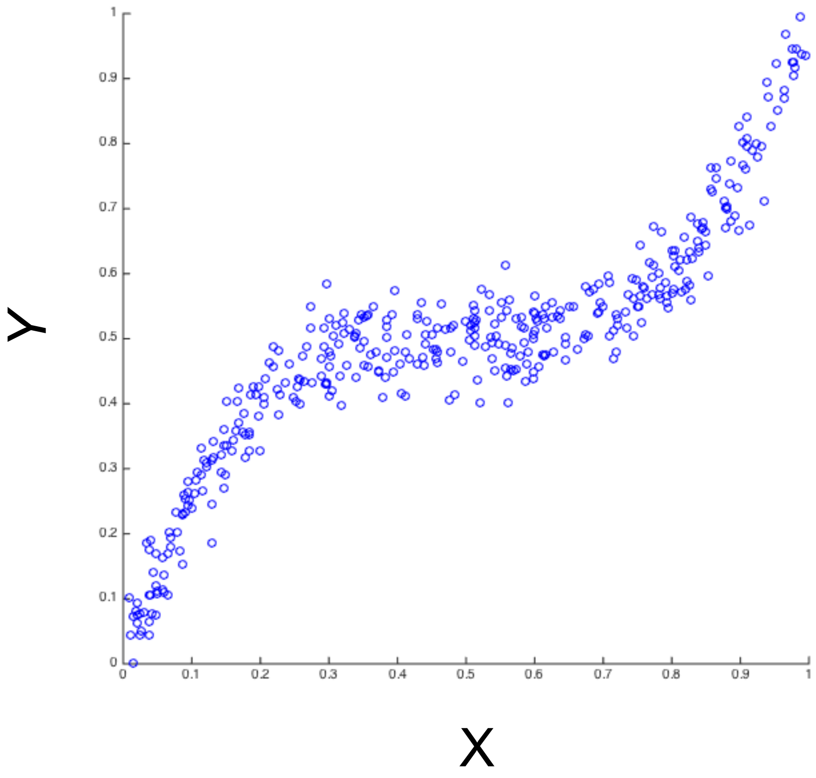

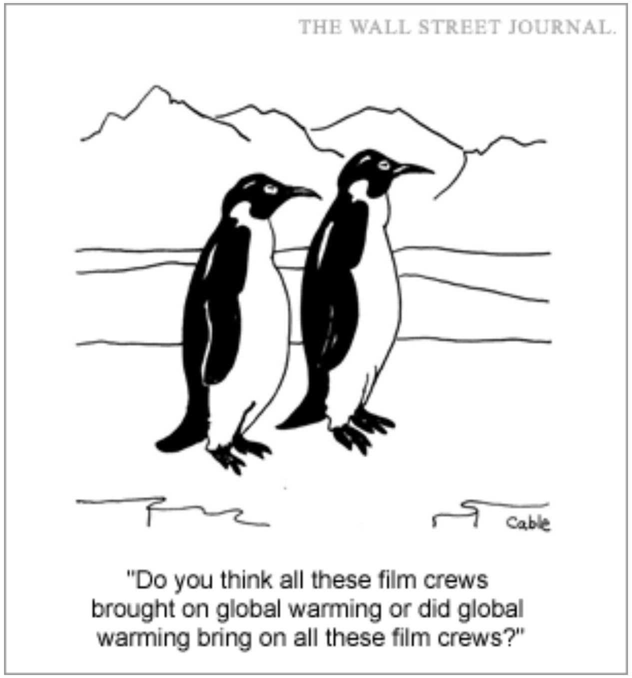

Q: Does X cause Y or does Y cause X? … or "can't tell"?

Causal direction from model simplicity

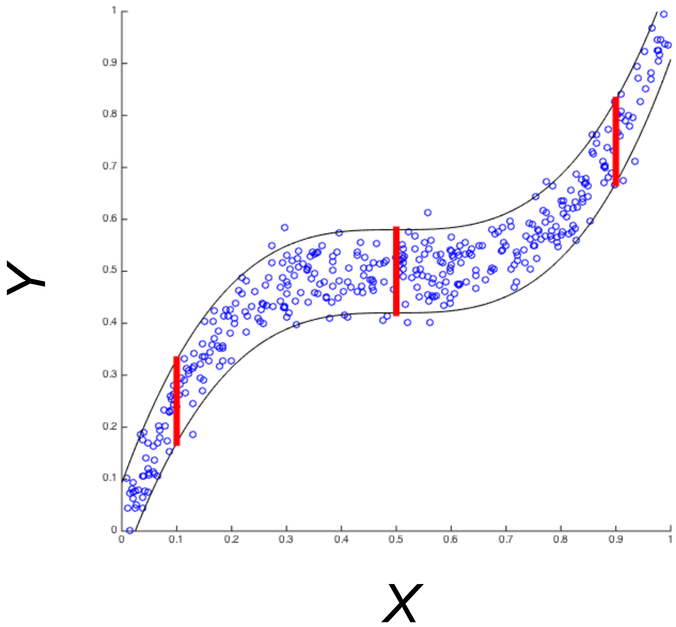

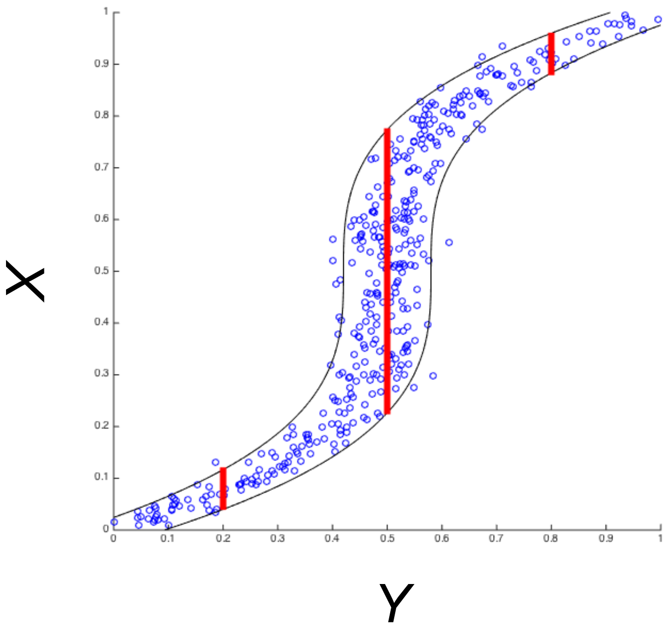

Easy to explain as

Y = *f*(*X*) + noise

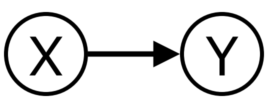

Difficult to explain as

X = *f*(*Y*) + noise

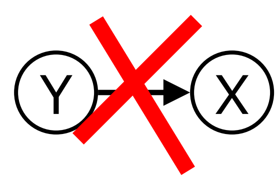

Chocolate consumption and Nobel prizes …

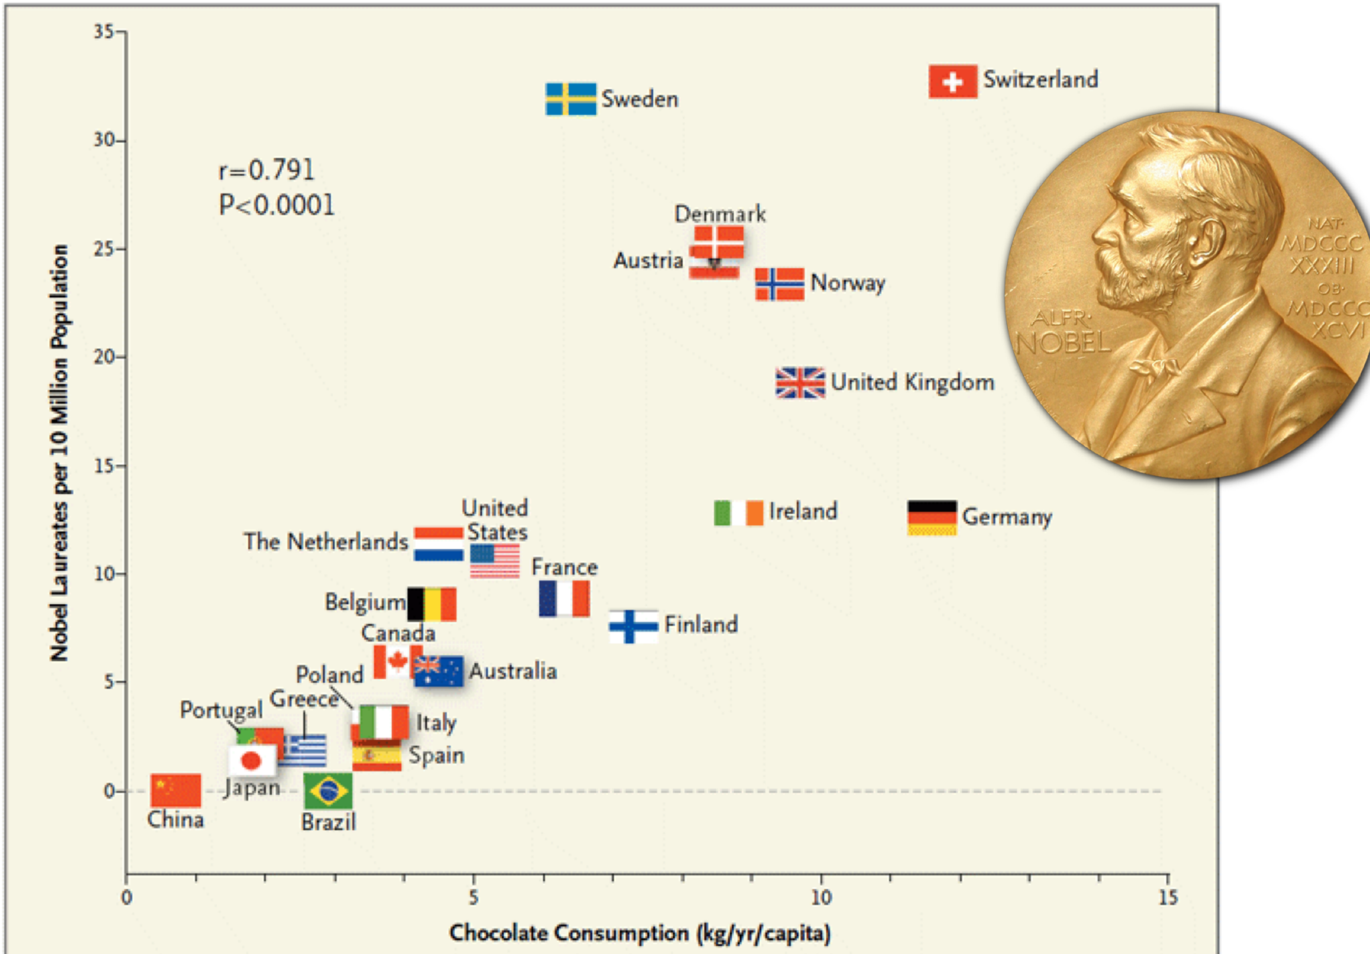

Figure 1. Correlation between Countries' Annual Per Capita Chocolate Consumption and the Number of Nobel Laureates per 10 Million Population.

Messerli, F. H., et al. "Chocolate Consumption, Cognitive Function, and Nobel Laureates." N Engl J Med 367.16 (2012): 1562-4.

Chocolate consumption and Nobel prizes …

underlying mechanism. However, since chocolate consumption has been documented to improve cognitive function, it seems most likely that in a dose-dependent way, chocolate intake provides the abundant fertile ground needed for the sprouting of Nobel laureates. Obviously, these findings are hypothesis-generating only and will have to be tested in a prospective, randomized trial.

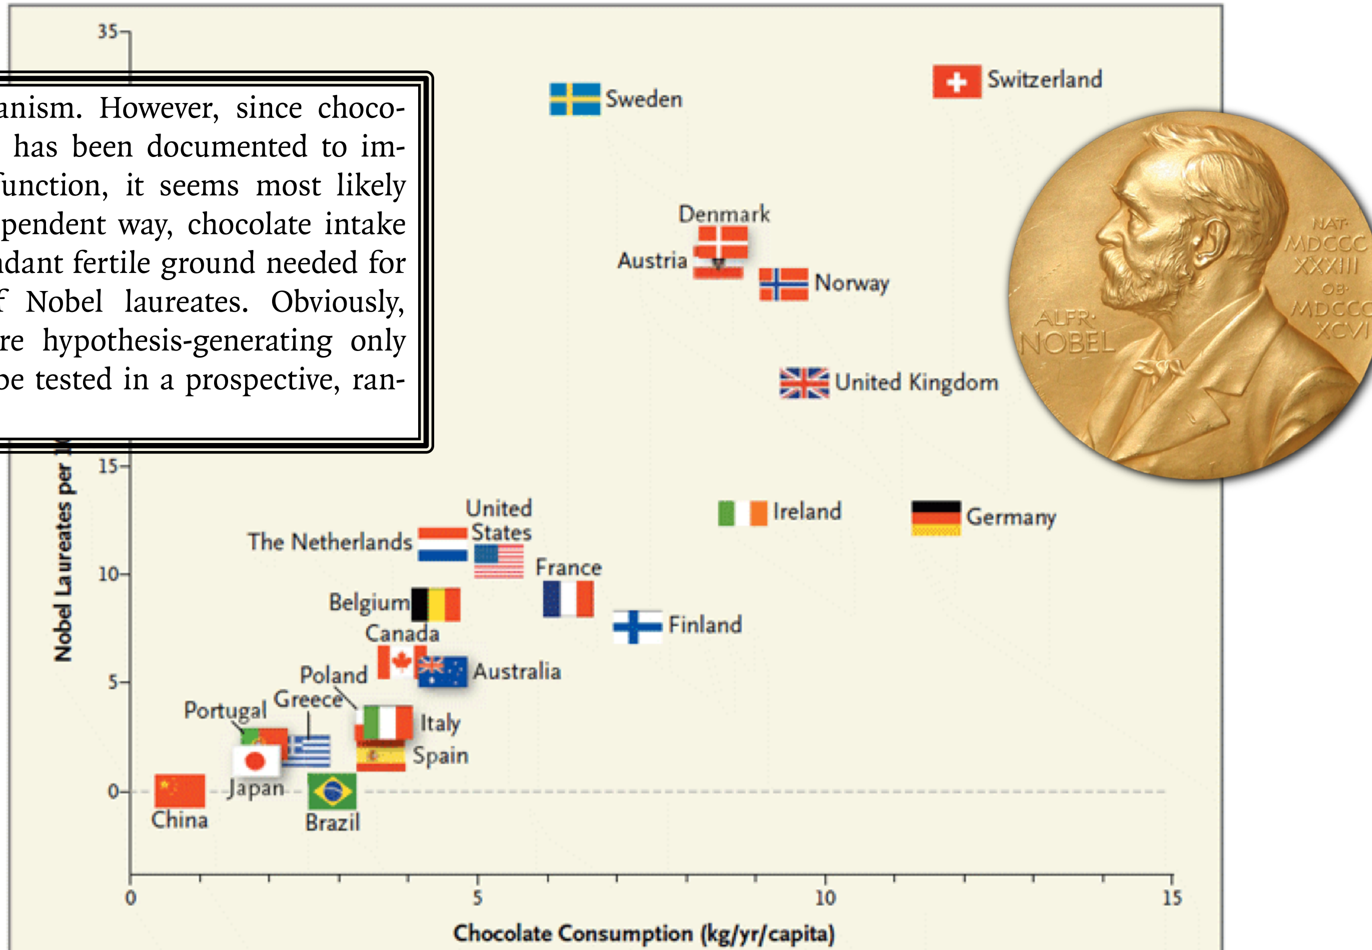

Figure 1. Correlation between Countries' Annual Per Capita Chocolate Consumption and the Number of Nobel Laureates per 10 Million Population.

NEW ENGLAND JOURNAL OF MEDICINE

Messerli, F. H., et al. "Chocolate Consumption, Cognitive Function, and Nobel Laureates." N Engl J Med 367.16 (2012): 1562-4.

Chocolate consumption and Nobel prizes ...

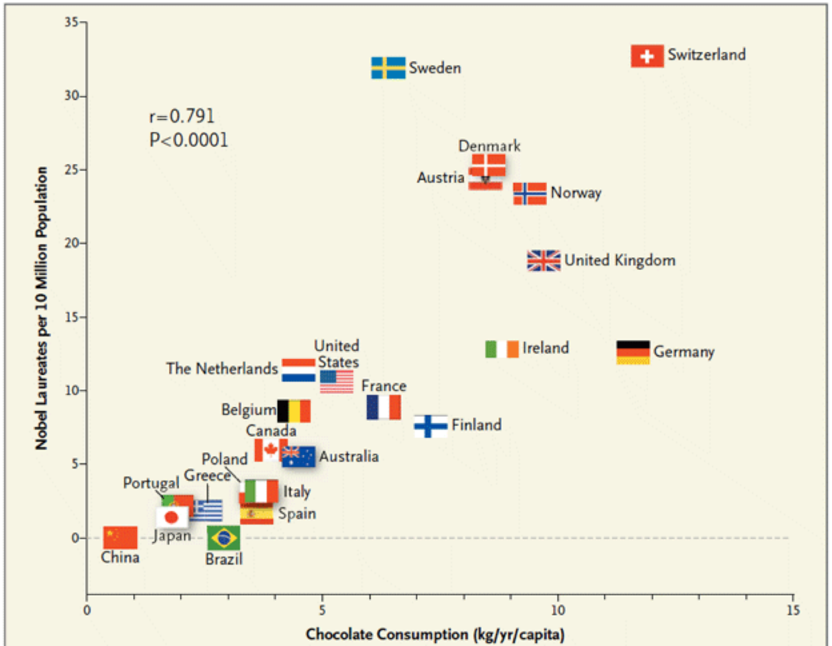

Figure 1. Correlation between Countries' Annual Per Capita Chocolate Consumption and the Number of Nobel Laureates per 10 Million Population.

NEW ENGLAND JOURNAL OF MEDICINE

Alternative explanation

- unobserved *confounder*

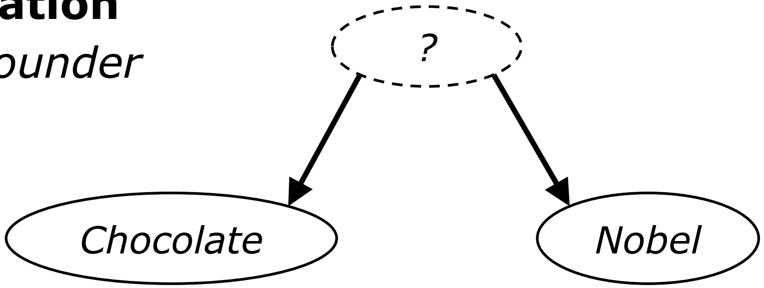

BRITISH MEDICAL JOURNAL

LONDON SATURDAY SEPTEMBER 30 1950

SMOKING AND CARCINOMA OF THE LUNG

PRELIMINARY REPORT

BY

RICHARD DOLL, M.D., M.R.C.P.

Member of the Statistical Research Unit of the Medical Research Council

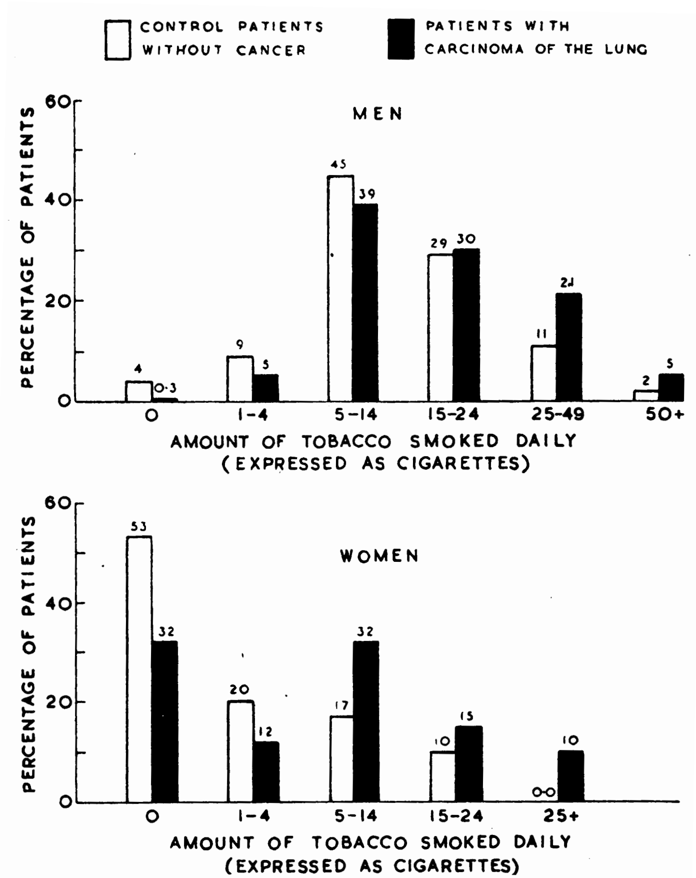

FIG. 1.—Percentage of patients smoking different amounts of tobacco daily.

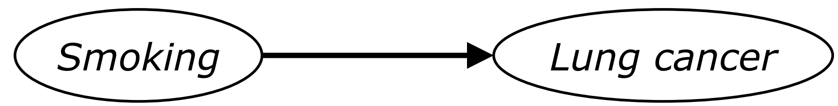

BRITISH MEDICAL JOURNAL

LONDON SATURDAY SEPTEMBER 30 1950

SMOKING AND CARCINOMA OF THE LUNG

PRELIMINARY REPORT

BY

RICHARD DOLL, M.D., M.R.C.P.

Member of the Statistical Research Unit of the Medical Research Council

Tobacco industry:

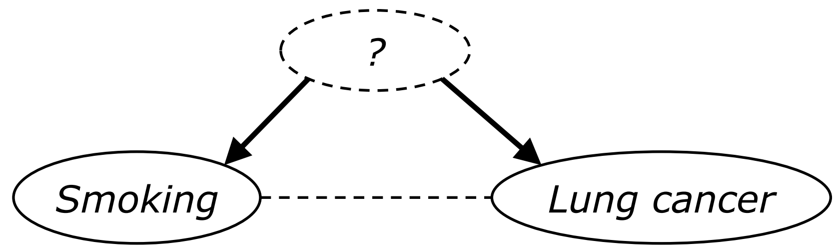

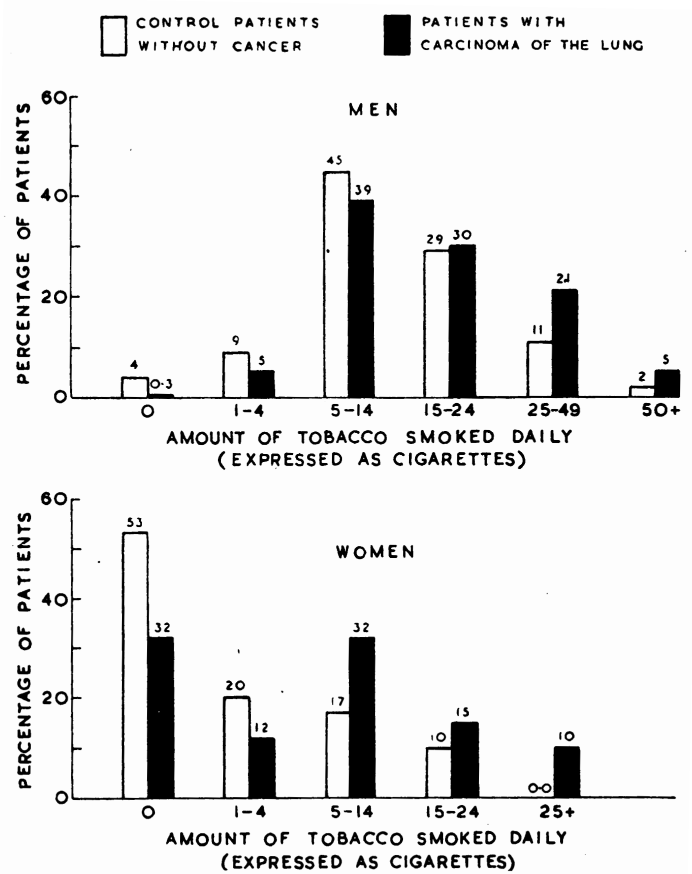

FIG. 1.—Percentage of patients smoking different amounts of tobacco daily.

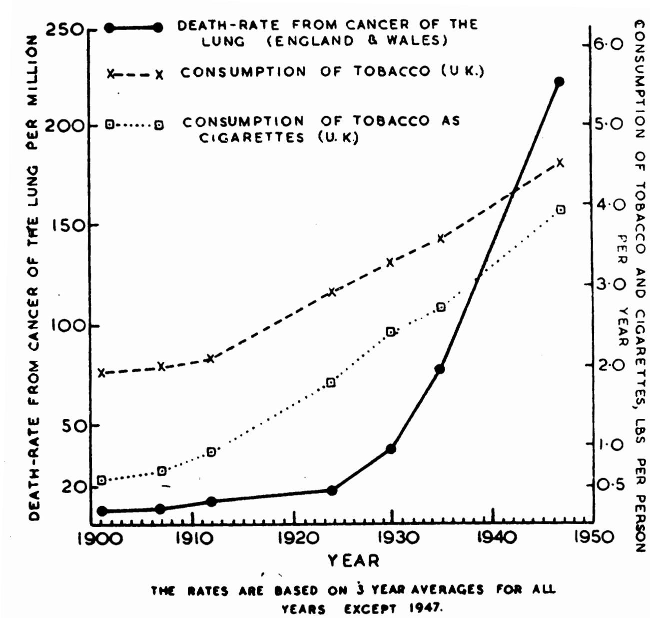

FIG. 2.—Death rate from cancer of the lung and rate of consumption of tobacco and cigarettes.

Treatment of kidney stones

	Treatment A	Treatment B
Recoveries	273/350 (78%)	289/350 (**83%**)
Total	562/700 (80%)	

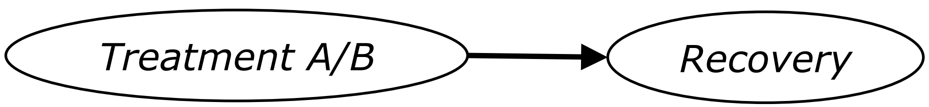

Question: What treatment would you prefer?

Charig, Clive R., et al. "Comparison of treatment of renal calculi by open surgery, percutaneous nephrolithotomy, and extracorporeal shockwave lithotripsy." Br Med J (Clin Res Ed) 292.6524 (1986): 879-882.

Treatment of kidney stones

	Treatment A	Treatment B
Small stones	81/87 (**93%**)	234/270 (87%)
Large stones	192/263 (**73%**)	55/80 (69%
Recoveries	273/350 (78%)	289/350 (**83%**)
Total	562/700 (80%)	

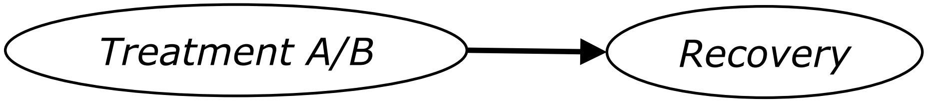

Question: What treatment would you prefer now?

Treatment of kidney stones

	Treatment A	Treatment B
Small stones	81/87 (**93%**)	234/270 (87%)
Large stones	192/263 (**73%**)	55/80 (69%
Recoveries	273/350 (78%)	289/350 (**83%**)
Total	562/700 (80%)	

"Simpson's paradox"

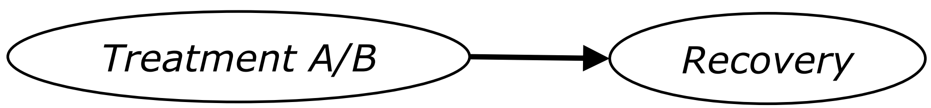

Question: What treatment would you prefer now?

Treatment of kidney stones

	Treatment A	Treatment B
Small stones	81/87 (**93%**)	234/270 (87%)
Large stones	192/263 (**73%**)	55/80 (69%
Recoveries	273/350 (78%)	289/350 (**83%**)
Total	562/700 (80%)	

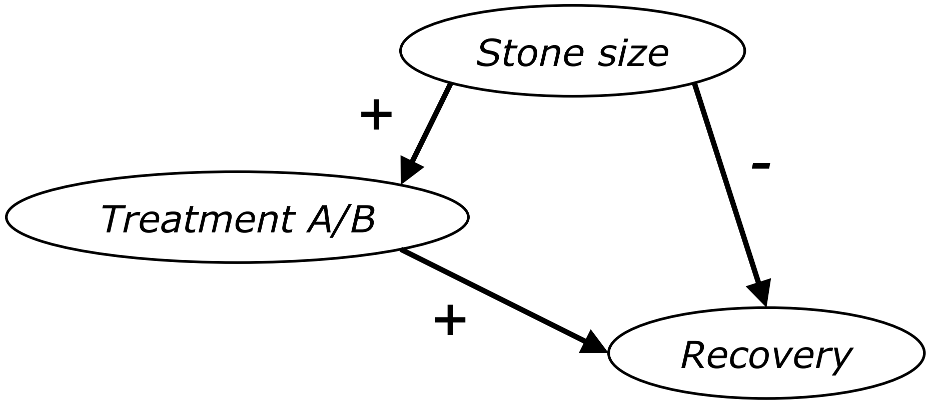

Treatment of kidney stones

	Treatment A	Treatment B
Small stones	81/87 (**93%**)	234/270 (87%)
Large stones	192/263 (**73%**)	55/80 (69%
Recoveries	273/350 (78%)	289/350 (**83%**)
Total	562/700 (80%)	

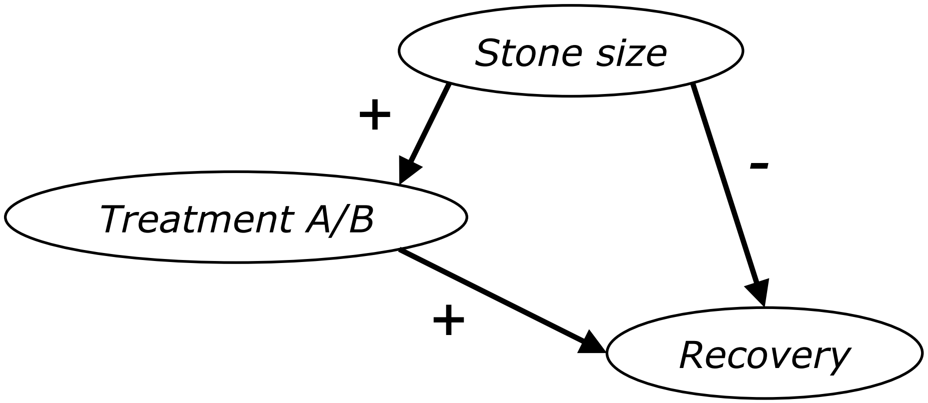

Question: How to compute the actual effect?

Computing causal effect sizes from observations

- split observed correlation in causal effect and confounding

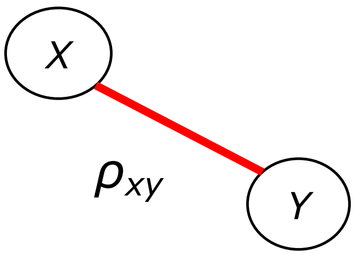

observed correlation

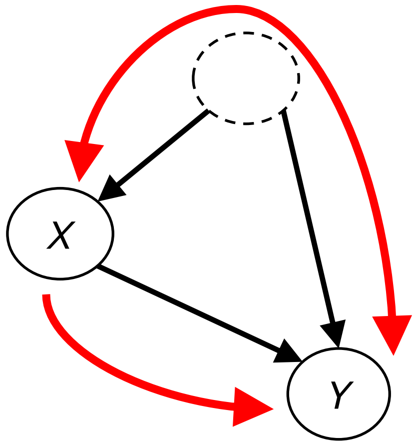

combination of (possible) causal effect and (possible) confounding

Computing causal effect sizes from observations

- split observed correlation in causal effect and confounding

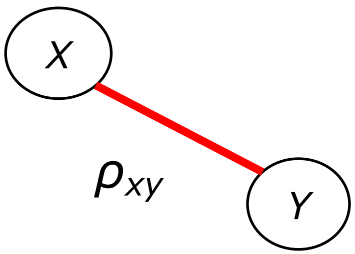

observed correlation

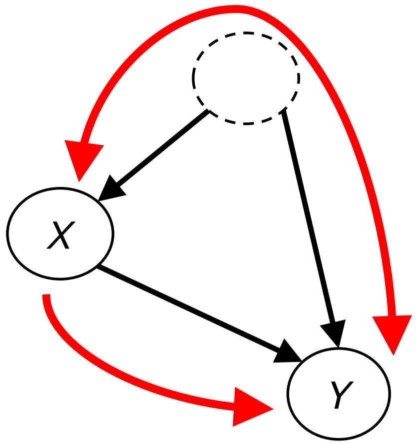

combination of (possible) causal effect and (possible) confounding

How to compute the causal effect?

- gold standard: randomized controlled trial!

otherwise

- *adjustment formula* to compensate for confounding (later this session)
- more general: *do*-calculus *[Pearl, Causality 2009]*
- not always possible!

Outline

1. Introduction to causality
2. Prediction vs. causation
3. **Causal graphs and how to read them**
4. Cause-effect estimation
5. The missing link & conclusion

Key model assumption: causal DAG

- real world consists of networks of causally interacting variables,
- structure corresponds to a *directed acyclic graph* (DAG)
- arcs represent *direct* causes between variables in the system

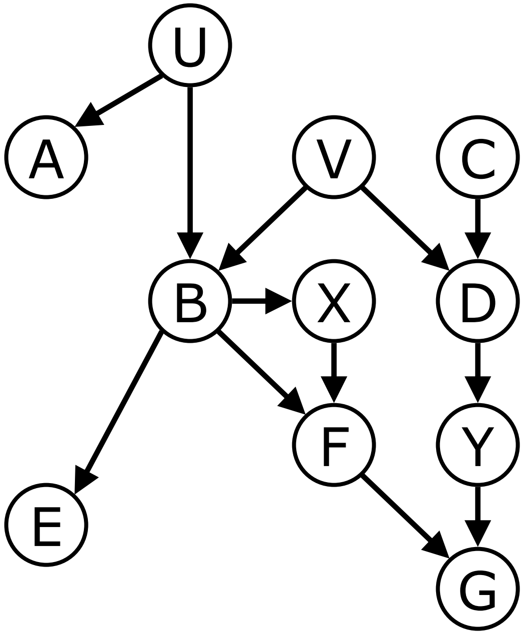

causal DAG *G*
(Directed Acyclic Graph)

Key model assumption: causal DAG

- real world consists of networks of causally interacting variables,
- structure corresponds to a *directed acyclic graph* (DAG)
- arcs represent *direct* causes between variables in the system
- subset of these variables observed in experiments

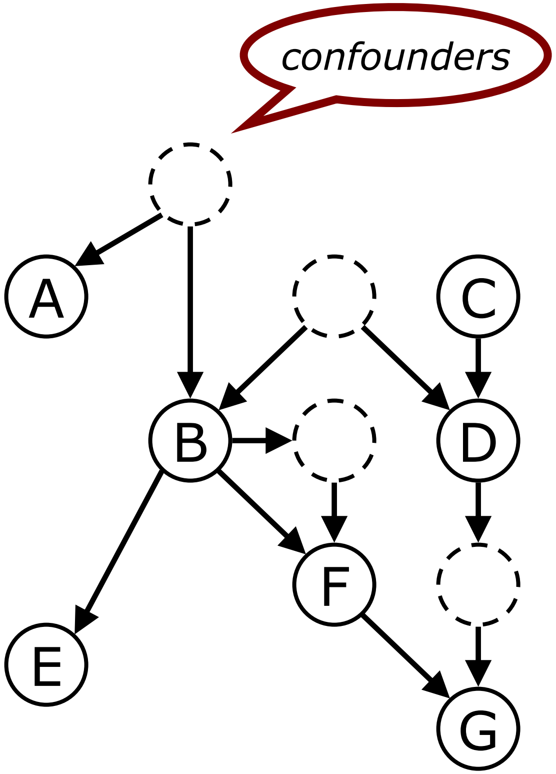

*underlying **causal DAG** G*
(Directed Acyclic Graph)

Key model assumption: causal DAG

- real world consists of networks of causally interacting variables,
- structure corresponds to a *directed acyclic graph* (DAG)
- arcs represent *direct* causes between variables in the system
- subset of these variables observed in experiments

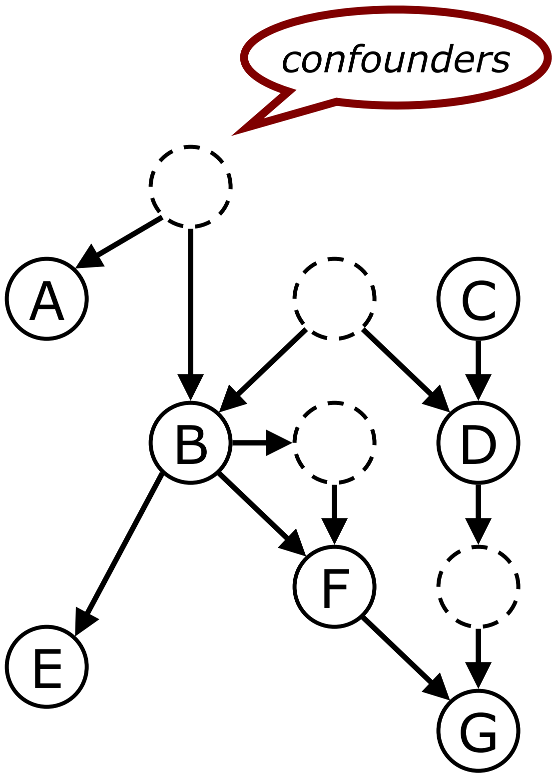

=

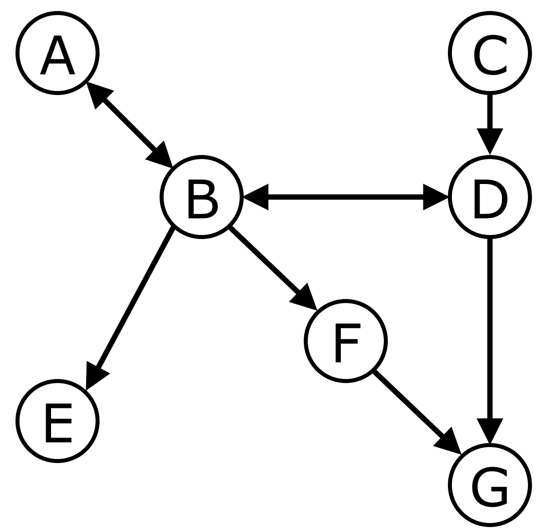

*underlying **causal DAG** G*
(Directed Acyclic Graph)

*equivalent **ADMG** representation*
(Acyclic Directed Mixed Graph)

Basic graphical model terminology

- nodes and edges

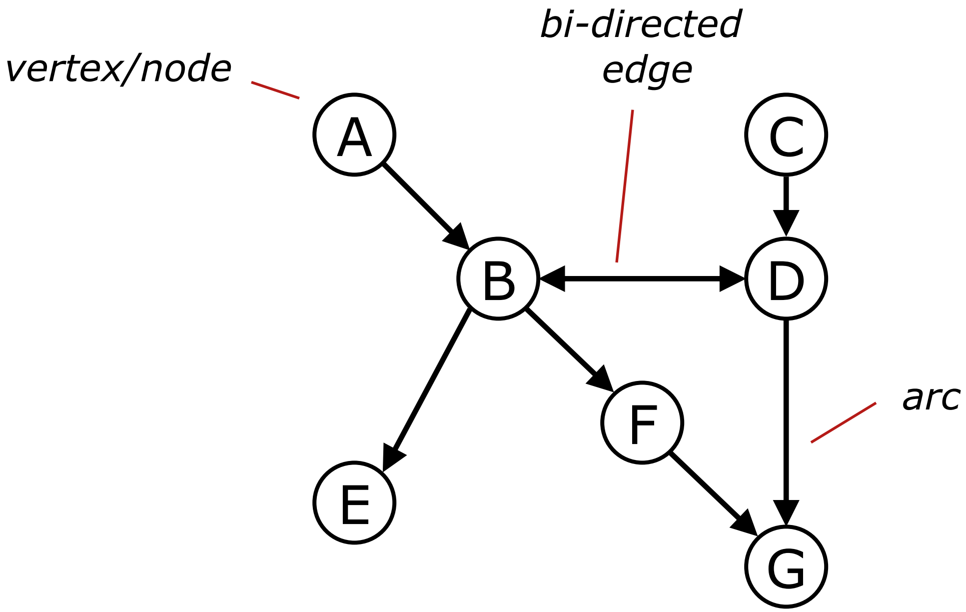

Paths

- path - sequence of (distinct) nodes $\pi = \langle X_1, X_2, .., X_k \rangle$ where each successive pair of nodes along the path is adjacent (connected by an edge) in graph *G*

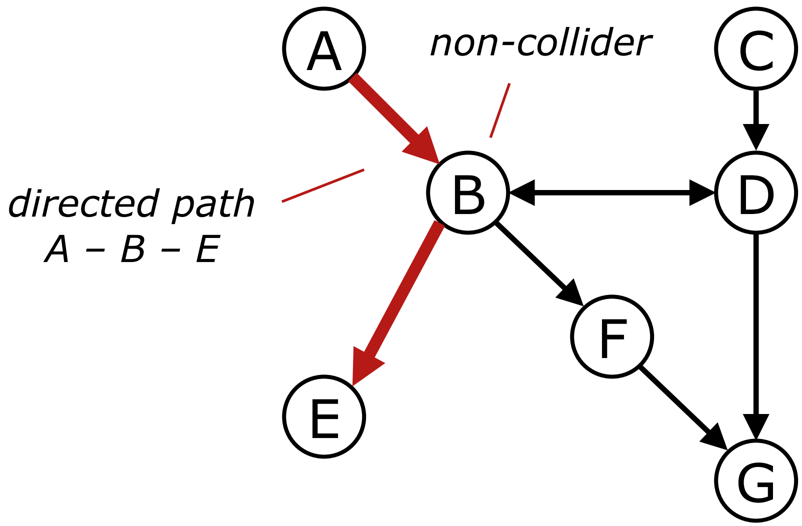

Collider and non-collider triples

- collider - triple of successive nodes ⟨X,Y,Z⟩ along a path, where the edges from X and Z have an arrowhead ('collide') at Y, e.g. $\mathbf{X} \leftrightarrow \mathbf{Y} \leftarrow \mathbf{Z}$
- non-collider - any such triple that is not a collider, e.g. $\mathbf{X} \rightarrow \mathbf{Y} \rightarrow \mathbf{Z}$, $\mathbf{X} \leftarrow \mathbf{Y} \leftarrow \mathbf{Z}$, or $\mathbf{X} \leftarrow \mathbf{Y} \rightarrow \mathbf{Z}$

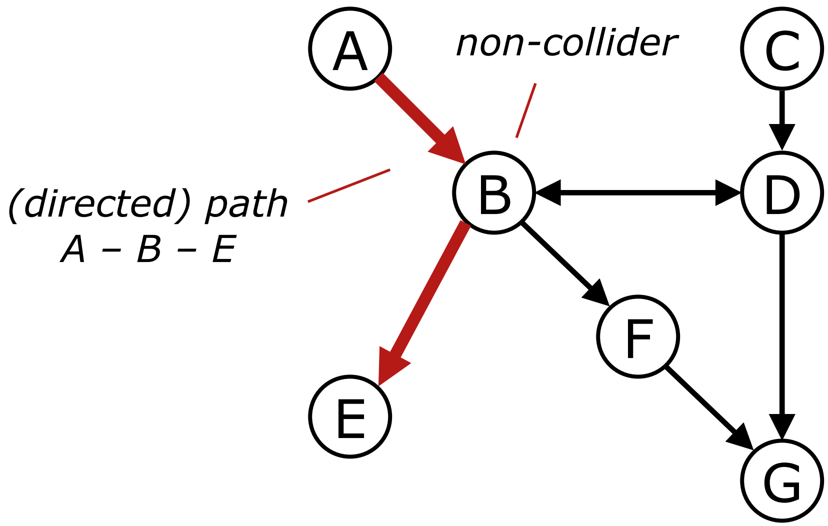

Collider and non-collider triples

- collider - triple of successive nodes ⟨X,Y,Z⟩ along a path, where the edges from X and Z have an arrowhead ('collide') at Y, e.g. $\mathbf{X} \leftrightarrow \mathbf{Y} \leftarrow \mathbf{Z}$
- non-collider - any such triple that is not a collider, e.g. $\mathbf{X} \rightarrow \mathbf{Y} \rightarrow \mathbf{Z}$, $\mathbf{X} \leftarrow \mathbf{Y} \leftarrow \mathbf{Z}$, or $\mathbf{X} \leftarrow \mathbf{Y} \rightarrow \mathbf{Z}$

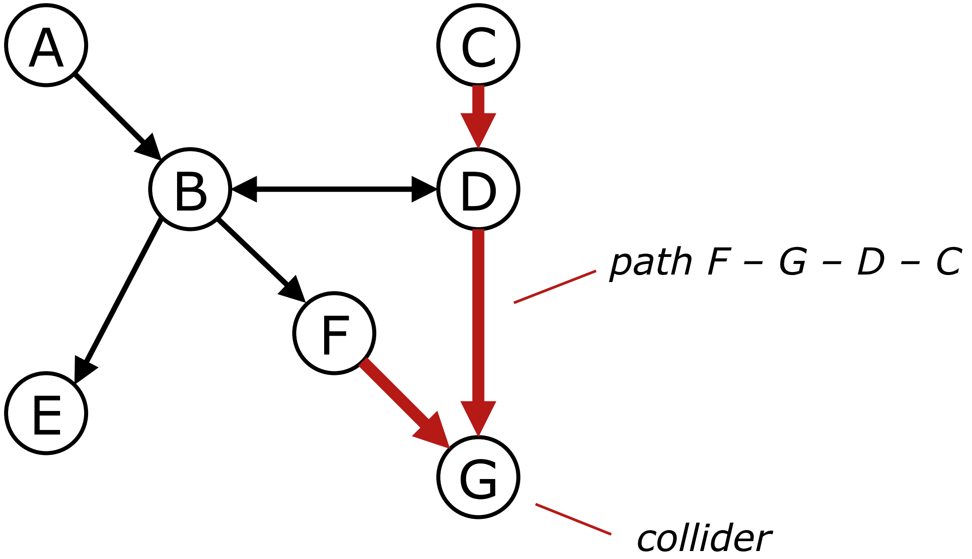

Ancestral relations

- if $\mathbf{X} \rightarrow \mathbf{Y}$ is in graph G, then X is a parent of Y, and Y is a child of X
- if $\mathbf{X} \leftrightarrow \mathbf{Y}$ is in graph G, then X is a spouse of Y (and v.v.)
- if there is a *directed path* $\mathbf{X} \rightarrow \mathbf{..} \rightarrow \mathbf{Y}$ in G, then X is ancestor of Y, and Y is a descendant of X

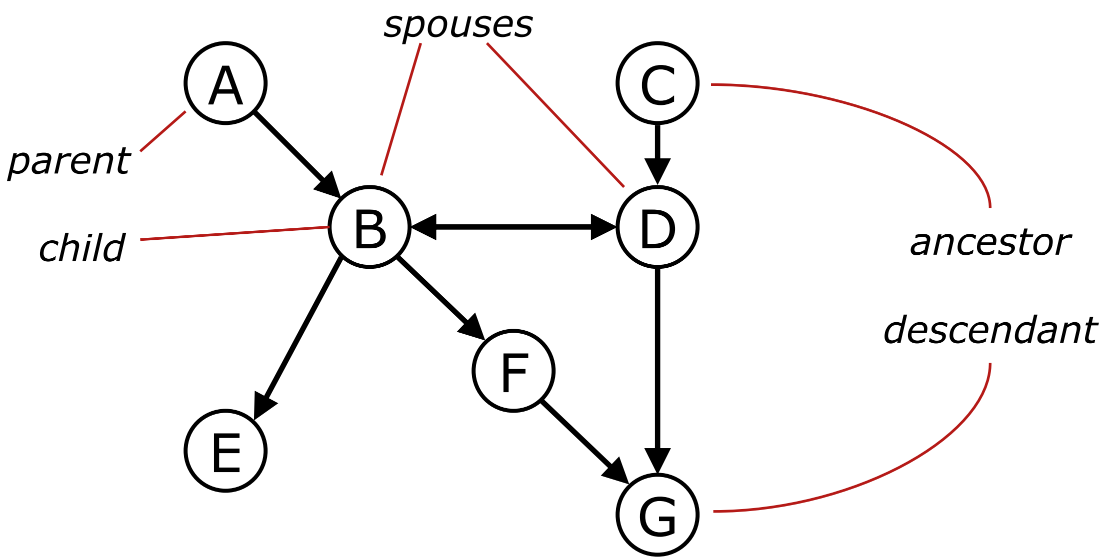

Blocked and unblocked paths

- a path $\pi = \langle X,..,Y \rangle$ is unblocked given set of nodes $\mathbf{Z}$ iff:
  - all non-colliders along π are *not* in $\mathbf{Z}$
  - all colliders along π are in $\mathbf{Z}$ or are ancestor of some $Z \in \mathbf{Z}$

  otherwise the path is blocked

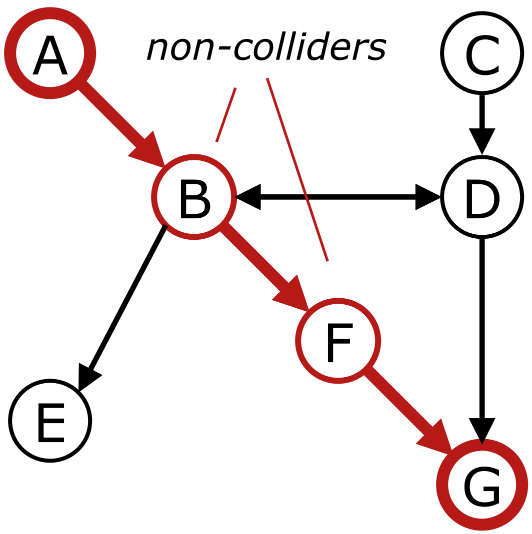

Path ⟨A,B,F,G⟩ is unblocked given the empty set …

Blocked and unblocked paths

- a path $\pi = \langle X,..,Y \rangle$ is unblocked <u>given set of nodes $\mathbf{Z}$</u> iff:
  - all non-colliders along π are *not* in $\mathbf{Z}$
  - all colliders along π are in $\mathbf{Z}$ or are ancestor of some $Z \in \mathbf{Z}$

  otherwise the path is blocked

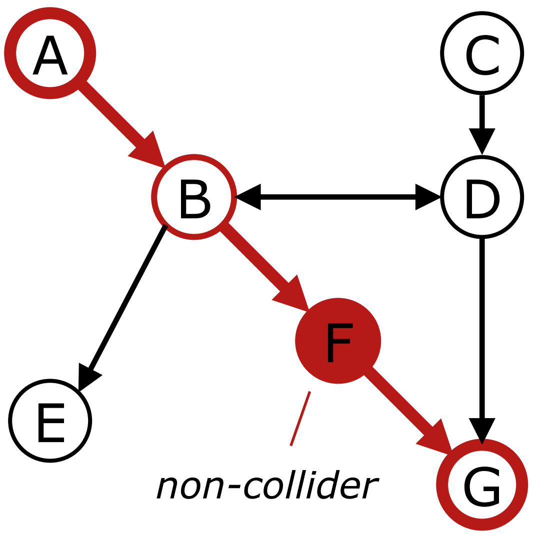

… path ⟨A,B,F,G⟩ is ***blocked*** *given F …*

Blocked and unblocked paths

- a path $\pi = \langle X,..,Y \rangle$ is unblocked <u>given set of nodes $\mathbf{Z}$</u> iff:
  - all non-colliders along π are *not* in $\mathbf{Z}$
  - all colliders along π are in $\mathbf{Z}$ or are ancestor of some $Z \in \mathbf{Z}$

  otherwise the path is blocked

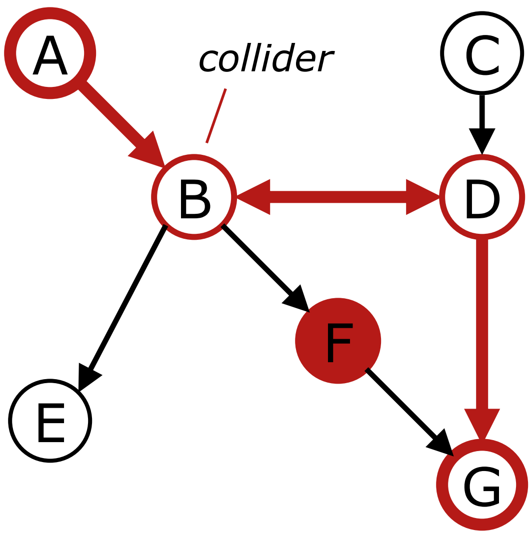

*… but path ⟨A,B,D,G⟩ becomes **unblocked** given F …*

Blocked and unblocked paths

- a path $\pi = \langle X,..,Y \rangle$ is unblocked given set of nodes $\mathbf{Z}$ iff:
  - all non-colliders along π are *not* in $\mathbf{Z}$
  - all colliders along π are in $\mathbf{Z}$ or are ancestor of some $Z \in \mathbf{Z}$

  otherwise the path is blocked

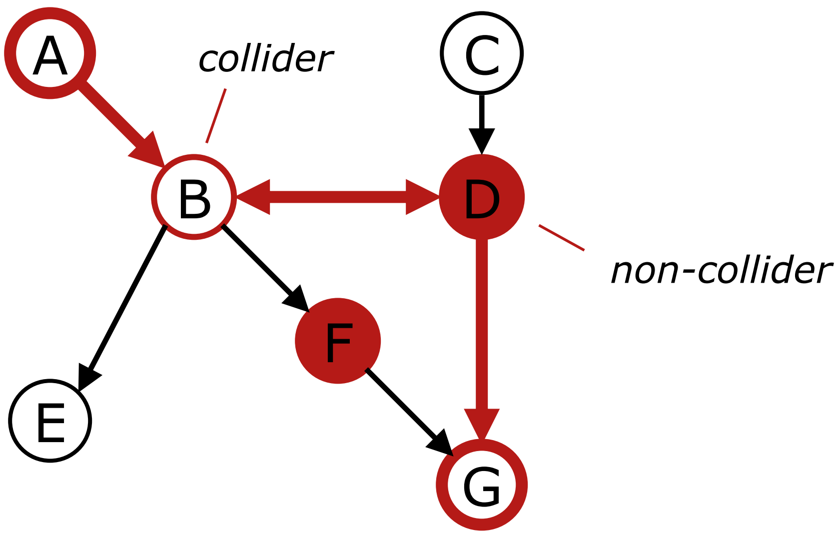

*… and path ⟨A,B,D,G⟩ is again **blocked** given {D,F}.*

d-separation

- in a graph *G*, nodes X and Y are *d*-separated given $\mathbf{Z}$, iff

  there are no unblocked paths in *G* between X and Y given $\mathbf{Z}$,

  otherwise they are *d*-connected

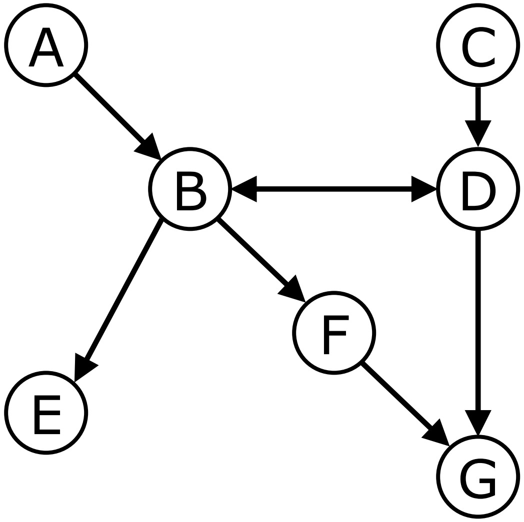

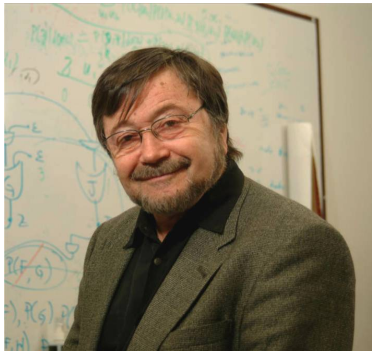

Judea Pearl
(Winner Turing Award 2012)

d-separation

- in a graph *G*, nodes X and Y are *d*-separated given $\mathbf{Z}$, iff

  there are no unblocked paths in *G* between X and Y given $\mathbf{Z}$,

  otherwise they are *d*-connected

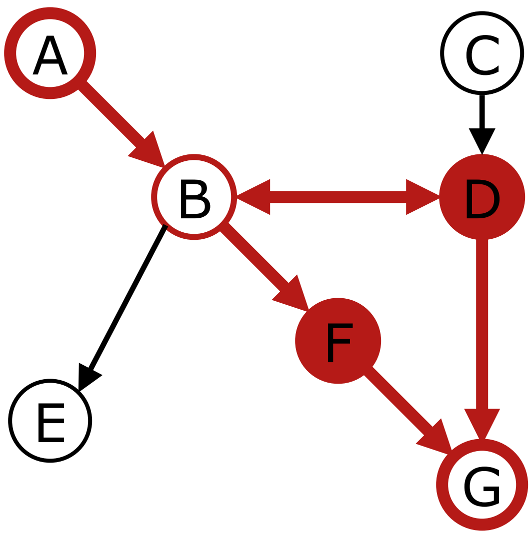

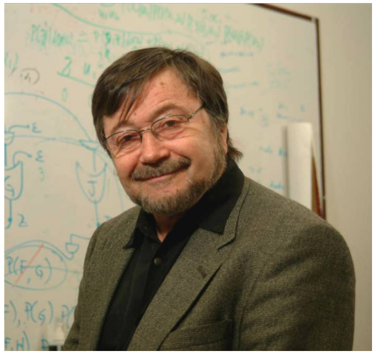

Judea Pearl
(Winner Turing Award 2012)

Nodes A and G are d-separated given {D,F},
but d-connected given {}, D, or F.

Exercise 1a – Paths and colliders

- collider - triple of successive nodes ⟨X,Y,Z⟩ along a path, where the edges from X and Z have an arrowhead ('collide') at Y, e.g. $\mathbf{X} \leftrightarrow \mathbf{Y} \leftarrow \mathbf{Z}$

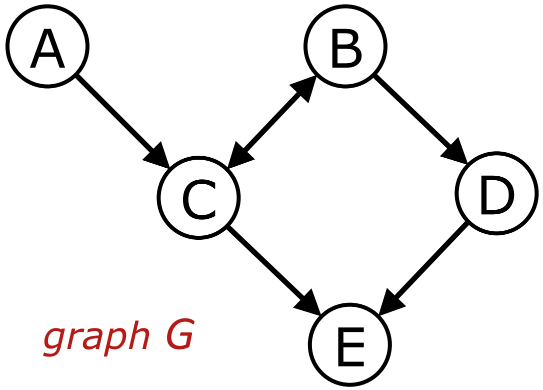

1. Is ⟨A,C,B,A⟩ a path?
2. Is ⟨A,C,E,D,B⟩ a (directed) path?
3. Is A an ancestor of D?
4. What are descendants of B?
5. Which nodes on the path ⟨B,D,E,C,A⟩ are non-colliders?
6. A v-structure is a collider between non-adjacent nodes. How many v-structures are in the graph *G*?

Exercise 1a – Paths and colliders

- collider - triple of successive nodes ⟨X,Y,Z⟩ along a path, where the edges from X and Z have an arrowhead ('collide') at Y, e.g. **X** $\leftrightarrow$ **Y** $\leftarrow$ **Z**

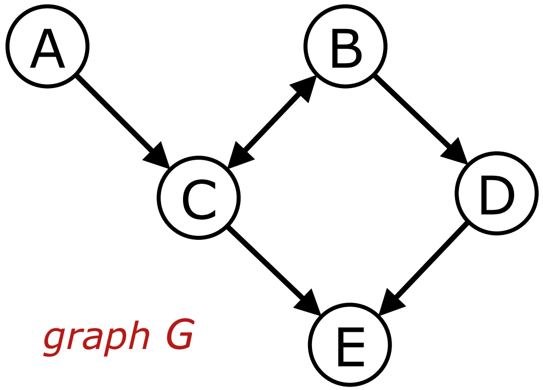

graph G

1. Is ⟨A,C,B,A⟩ a path? *No: A and B are not adjacent and A occurs twice.*
2. Is ⟨A,C,E,D,B⟩ a (directed) path? *It is a path, but not a directed path.*
3. Is A an ancestor of D? *No: there is no directed path from A to D.*
4. What are descendants of B? *Nodes {B,C,D} (B is also its own descendant!)*
5. Which nodes on the path ⟨B,D,E,C,A⟩ are non-colliders? *Nodes C and D.*
6. A v-structure is a collider between non-adjacent nodes. How many v-structures are in the graph *G*? *Two: A $\rightarrow$ C $\leftrightarrow$ B, and C $\rightarrow$ E $\leftarrow$ D.*

Exercise 1b – Blocked and unblocked paths

- a path $\pi = \langle X,..,Y \rangle$ is unblocked given set of nodes $\mathbf{Z}$ iff:
  - all non-colliders along π are *not* in $\mathbf{Z}$
  - all colliders along π are in $\mathbf{Z}$ or are ancestor of some $Z \in \mathbf{Z}$

  otherwise the path is blocked

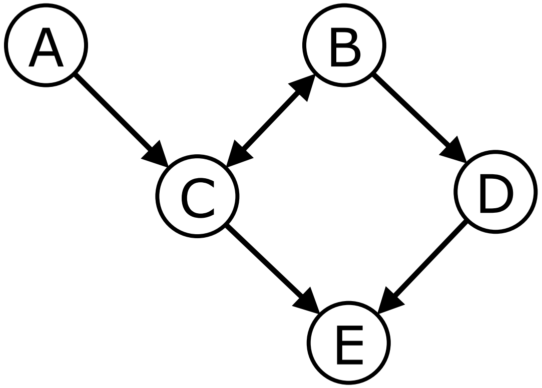

1. Is C $\leftrightarrow$ B $\rightarrow$ D blocked by (`given') B?
2. Is A $\rightarrow$ C $\leftrightarrow$ B blocked given E?
3. Is A $\rightarrow$ C $\rightarrow$ E $\leftarrow$ D blocked? (given empty set **Z** = {})
4. Is path $\langle$A,C,B,D$\rangle$ blocked by {C,E}?
5. Which set(s) of nodes (if any) unblock a path from A to B?
6. Claim: 'A path between two nodes can be blocked, iff they are non-adjacent'. True or false?

Exercise 1b – Blocked and unblocked paths

- a path $\pi = \langle X,..,Y \rangle$ is unblocked given set of nodes $\mathbf{Z}$ iff:
  - all non-colliders along π are *not* in $\mathbf{Z}$
  - all colliders along π are in $\mathbf{Z}$ or are ancestor of some $Z \in \mathbf{Z}$

  otherwise the path is blocked

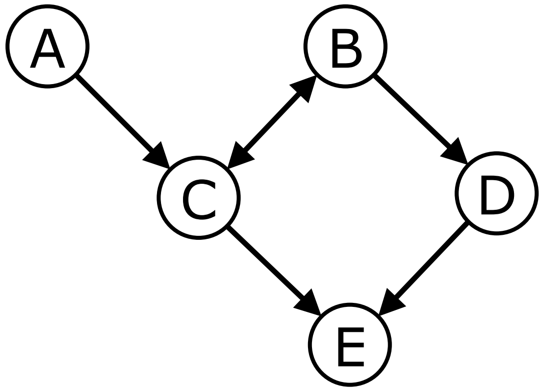

1. Is C $\leftrightarrow$ B $\rightarrow$ D blocked by (`given') B? *Yes.*
2. Is A $\rightarrow$ C $\leftrightarrow$ B blocked given E? *No.*
3. Is A $\rightarrow$ C $\rightarrow$ E $\leftarrow$ D blocked? (given empty set **Z** = {}) *Yes.*
4. Is path $\langle$A,C,B,D$\rangle$ blocked by {C,E}? *No.*
5. Which set(s) of nodes (if any) unblock a path from A to B? *Any subset of {C,D,E} containing at least one node from {C,E}.*
6. Claim: 'A path between two nodes can be blocked, iff they are non-adjacent'. True or false? *False: reverse counter example $\langle$B,D,E,C$\rangle$.*

Exercise 1c – *d*-separation

- in a graph *G*, nodes X and Y are *d*-separated given $\mathbf{Z}$, iff

  there are no unblocked paths in *G* between X and Y given $\mathbf{Z}$,

  otherwise they are *d*-connected

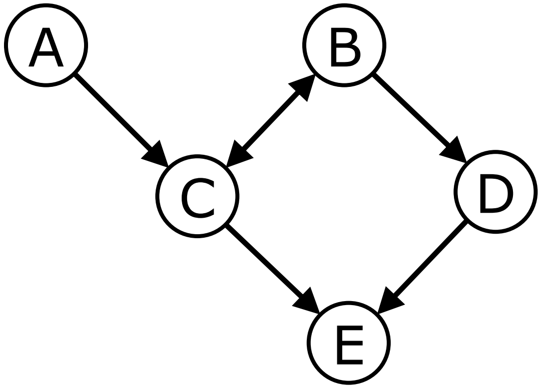

1. Are A and B *d*-separated? (given empty set {})
2. Are C and D *d*-separated by B?
3. Are A and E *d*-separated by C?
4. Are A and D *d*-separated by {B,E}?
5. Which set(s) of nodes (if any) would *d*-separate B and E?
6. True or false: ‘Two nodes can be *d*-separated, iff they are non-adjacent’?

Exercise 1c – *d*-separation

- in a graph *G*, nodes X and Y are *d*-separated given **Z**, iff

  there are no unblocked paths in *G* between X and Y given **Z**,

  otherwise they are *d*-connected

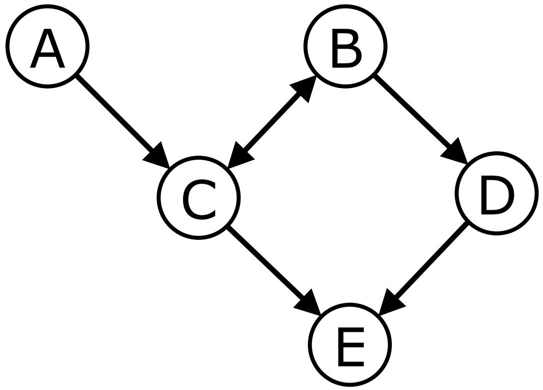

1. Are A and B *d*-separated? (given empty set {}) *Yes.*
2. Are C and D *d*-separated by B? *Yes.*
3. Are A and E *d*-separated by C? *No: path ⟨A,C,B,D,E⟩ is unblocked by C.*
4. Are A and D *d*-separated by {B,E}? *No: ⟨A,C,E,D⟩ remains unblocked.*
5. Which set(s) of nodes (if any) would *d*-separate B and E? *{C,D},{A,C,D}*
6. True or false: 'Two nodes can be *d*-separated, iff they are non-adjacent'?
   True for DAGs, but not for ADMGs in general!

Linking graphs to data

- graphical models offer an intuitive means to model causal interactions
- so far we only considered the causal structure …
- … now we need to link the graphs to data

⇒ enter the Causal Bayesian Network!

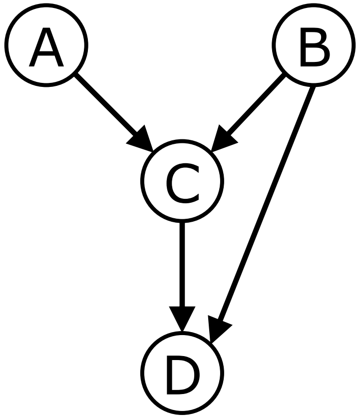

causal DAG G

Bayesian network

A **Bayesian Network** (BN) is a pair (G, p), where

- G is a directed acyclic graph over variables $\mathbf{X} = \{X_1, X_2, .. , X_K\}$
- p is a joint probability distribution over $\mathbf{X}$ that factorizes according to G

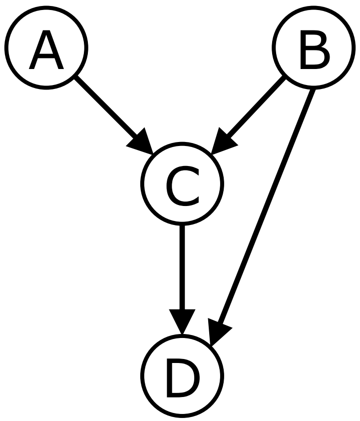

$$p(\mathbf{X}) = \prod_{k=1}^{K} p\left(X_k \middle| pa\left(X_k\right)\right)$$

parents of X_k in G

causal DAG G

factorized joint probability distribution

Bayesian network

A **Bayesian Network** (BN) is a pair (G, p), where

- G is a directed acyclic graph over variables $\mathbf{X} = \{X_1, X_2, .. ,X_K\}$
- p is a joint probability distribution over $\mathbf{X}$ that factorizes according to G

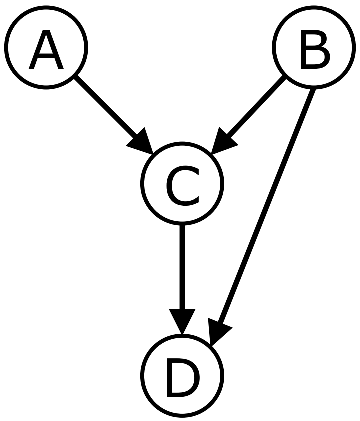

causal DAG G

$$p(\mathbf{X}) = \prod_{k=1}^{K} p\Big(X_k \,\Big|\, pa\big(X_k\big)\Big)$$

$$p(A,B,C,D) = p(A)\,p(B)\,p(C|A,B)\,p(D|B,C)$$

factorized joint probability distribution

Causal Bayesian network

A Bayesian Network (G, p) is causal if

- all and only the directed edges in G correspond to direct causal relations,
- it satisfies the **Causal Markov condition**:

> "In a causal DAG G, every node is probabilistically independent of its non-descendants given its parents (direct causes) in G."

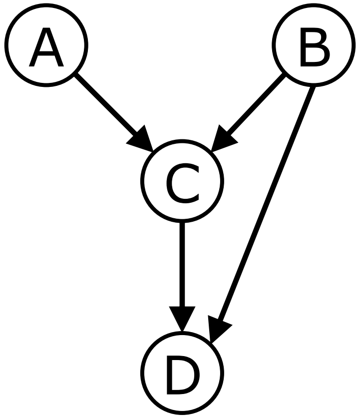

causal DAG G

$$X_i \perp\!\!\!\perp nd(X_i) \,|\, pa(X_i)$$

Causal Bayesian network

A Bayesian Network (G, p) is causal if

- all and only the directed edges in G correspond to direct causal relations,
- it satisfies the **Causal Markov condition**:

"In a causal DAG G, every node is probabilistically independent of its non-descendants given its parents (direct causes) in G."

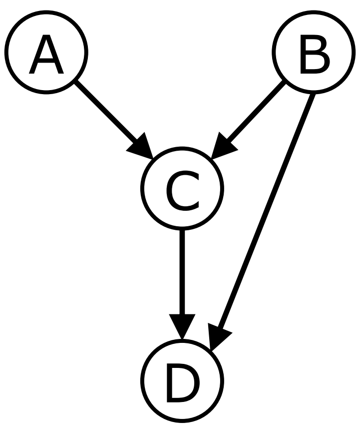

causal DAG G

$$A \perp\!\!\!\perp B \mid -$$

$$B \perp\!\!\!\perp A \mid -$$

$$D \perp\!\!\!\perp A \mid B, C$$

As a result

- *d*-separation $\Rightarrow$ probabilistic independence

Structural Causal Model

- each child-parent family in the causal DAG G corresponds to a deterministic function

$$X_i = f_i\left(pa\left(X_i\right), \varepsilon_i\right)$$

  with ε_i representing all exogenous influences (noise) on X_i
- collection is a Structural Causal/Equation Model (SCM/SEM)

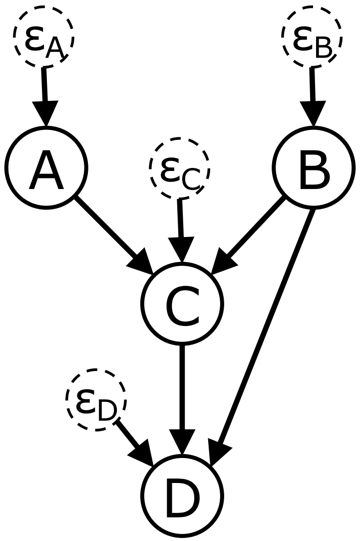

causal DAG G

$$A = f_A\left(\varepsilon_A\right)$$

$$B = f_B\left(\varepsilon_B\right)$$

$$C = f_C\left(A, B, \varepsilon_C\right)$$

$$D = f_D\left(B, C, \varepsilon_D\right)$$

corresponding Structural Causal Model

Structural Causal Model

- each child-parent family in the causal DAG G corresponds to a deterministic function

$$X_i = f_i\left(pa\left(X_i\right), \varepsilon_i\right)$$

with ε_i representing all exogenous influences (noise) on X_i

- collection is a Structural Causal/Equation Model (SCM/SEM)

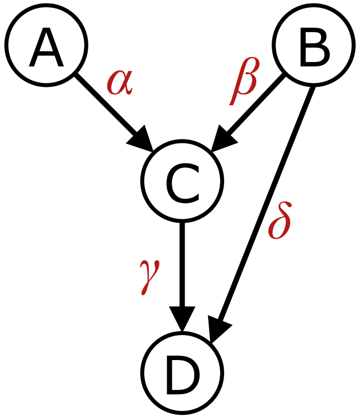

causal DAG G

$$A = \varepsilon_A \qquad \varepsilon_i \sim N\left(0, \sigma_i\right)$$

$$B = \varepsilon_B$$

$$C = \alpha A + \beta B + \varepsilon_C$$

$$D = \gamma B + \delta D + \varepsilon_D$$

Example: multivariate Gaussian model

Outline

1. Introduction to causality
2. Prediction vs. causation
3. Causal graphs and how to read them
4. **Cause-effect estimation**
5. The missing link & conclusion

Interventions

- intervening = actively changing the world

not this one …

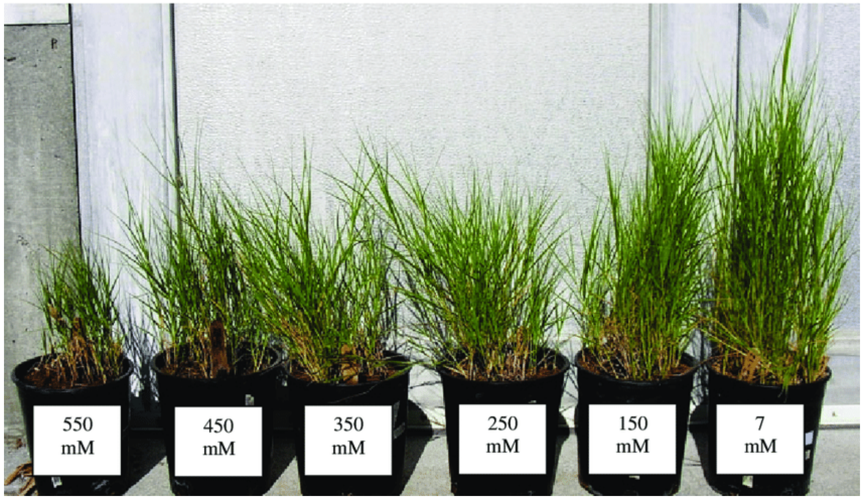

but this one

Interventions

- intervening = actively changing the world

Examples

- prescribing a treatment (or placebo) in an RCT
- gene knock-out experiment
- deciding to quit smoking
- governments changing laws / taxation levels,
- lowering home room temperatures
- adding a catalyst to a chemical reaction, etc.

Interventions

- intervening = actively changing the world

Examples

- prescribing a treatment (or placebo) in an RCT
- gene knock-out experiment
- deciding to quit smoking
- governments changing laws / taxation levels,
- lowering home room temperatures
- adding a catalyst to a chemical reaction, etc.

Common types of interventions

- hard/soft - (directly forcing a variable to a specific value vs. indirectly stimulating a variable to e.g. higher/lower values)
- surgical/fat-hand - (very precisely affecting only the target intervention variable vs. having possible unintended side-effects)
- perfect interventions = hard+surgical (Pearl's *do*-operator, see next)
- mechanism interventions (acting on the functional form of the relations)

Intervention in a Structural Causal Model

Perfect intervention in SCM

- externally force a node to a specific value: do(X_i = x_i)
- replace structural equation $f_i(..)$ with constant x_i
- corresponds to removing all incoming arcs to X_i in causal DAG G

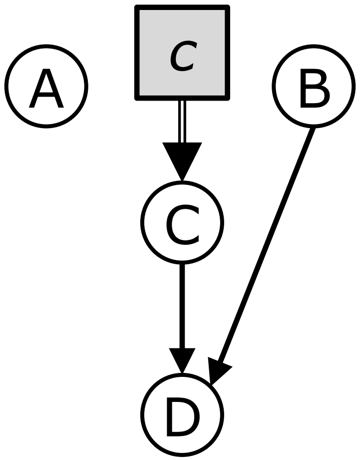

causal DAG G *with intervention on C*

$$A = f_A(\varepsilon_A)$$

$$B = f_B(\varepsilon_B)$$

$$C = c$$

$$D = f_D(B, C, \varepsilon_D)$$

intervention do(*C* = *c*)

Computing what happens after an intervention

We can understand / predict the effect of an intervention if we can rewrite the (unknown) interventional distribution in terms of the known observed distribution.

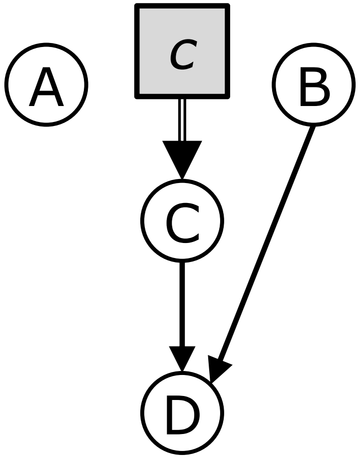

causal DAG G with intervention on C

$$p(A,B,C,D) = p(A)\,p(B)\,p(C|A,B)\,p(D|B,C)$$

original observed joint probability distribution

$$p(A,B,C,D|do(C=c)) = \ldots?$$

interventional distribution under do(C = c)

Computing the causal effect: adjustment

- The difference between the expectation under pre- and post-interventional distribution then corresponds to the *causal effect*
- Difficult to compute in general : Pearl's **do-calculus**

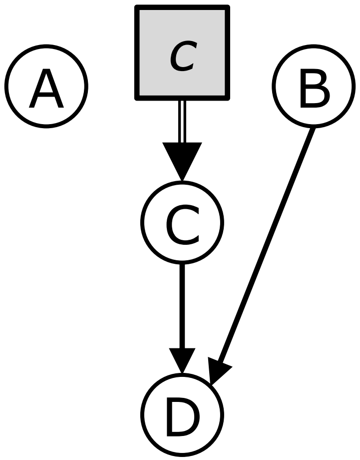

causal DAG G with intervention on C

Computing the causal effect: adjustment

- The difference between the expectation under pre- and post-interventional distribution then corresponds to the *causal effect*
- Difficult to compute in general : Pearl's **do-calculus**

- Fortunately, for a large class of problems there exists a relatively straightforward procedure: 'adjusting for the parents'

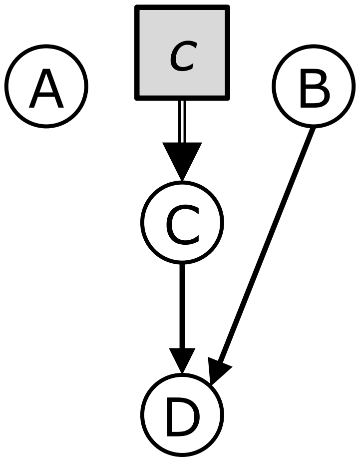

causal DAG G with intervention on C

$$p\left(Y=y \middle| do\left(X=x\right)\right) = \sum_{Pa(X)} p\left(y \middle| x, Pa(X)\right) p\left(Pa(X)\right)$$

adjustment formula for intervention on X

Back-door criterion

- we can generalize adjustment to 'admissible' sets (instead of just parents)

Theorem: A set of nodes **S** is *admissible* for adjustment to find the causal effect of X on Y, if:

- $X, Y \notin$ **S**
- no element of **S** is a *descendant* of X
- **S** blocks all back-door paths $X \leftarrow .. Y$ (all paths between X and Y that start with an incoming arc on X)

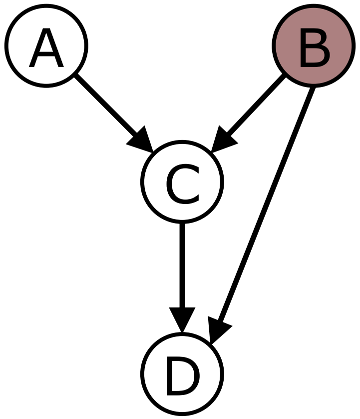

B is admissible for computing the causal effect of C (or A) on D

Back-door criterion

- we can generalize adjustment to 'admissible' sets (instead of just parents)

Theorem: A set of nodes **S** is *admissible* for adjustment to find the causal effect of X on Y, if:

- $X, Y \notin \mathbf{S}$
- no element of **S** is a *descendant* of X
- **S** blocks all back-door paths $X \leftarrow .. \, Y$ (all paths between X and Y that start with an incoming arc on X)

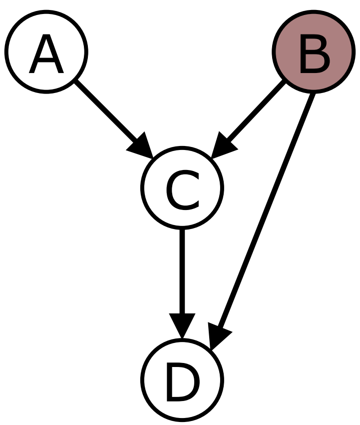

B is admissible for computing the causal effect of C (or A) on D

$$p\left(Y = y \middle| do\left(X = x\right)\right) = \sum_{\mathbf{s}} p\left(y \middle| x, \mathbf{S} = \mathbf{s}\right) p\left(\mathbf{S} = \mathbf{s}\right)$$

$$\left(\; = \int p\left(y \middle| x, \mathbf{s}\right) p\left(\mathbf{s}\right) d\mathbf{s} \; \right)$$

general adjustment formula

Average Causal Effect (ACE)

- if we can predict what happens on an intervention we can consider quantifying the causal impact of one variable on another

- the **Average Causal Effect** (ACE) quantifies the causal effect of *X* on *Y* as the difference in expectation of *Y* under different interventions on *X*

$$
\begin{aligned}
ACE(X \rightarrow Y) &= E\left[Y \middle| do(X=1)\right] - E\left[Y \middle| do(X=0)\right] \\
&= \sum_{Y} Y \cdot p\left(Y \middle| do(X=1)\right) - \sum_{Y} Y \cdot p\left(Y \middle| do(X=0)\right)
\end{aligned}
$$

ACE for causal effect of binary variable X on ordinal variable Y

Exercise 2a – Admissible sets

A set of nodes **S** is *admissible* for adjustment for the causal effect of *X* on *Y*, if:

- *X*,*Y* $\notin$ **S**
- no element of **S** is a *descendant* of *X*
- **S** blocks all back-door paths *X* $\leftarrow$ .. *Y* (all paths between *X* and *Y* that start with an incoming arc on *X*)

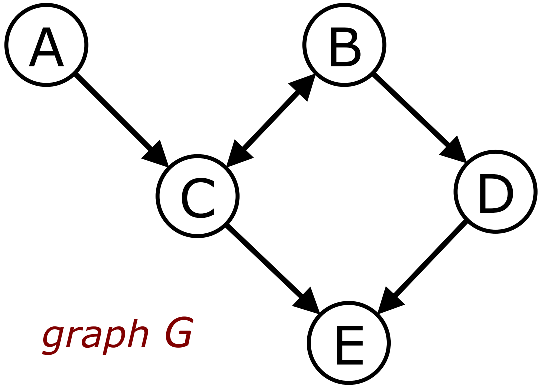

graph G

1. Is B admissible for adjustment to find the causal effect of D on E?
2. Is {} admissible for the causal effect of A on E?
3. Is B admissible for the causal effect of A on E?
4. Is {B,D} admissible for the causal effect of A on E?
5. Is C admissible for the causal effect of A on E?
6. Is {B,C} admissible for the causal effect of E on A?

Exercise 2a – Admissible sets

A set of nodes **S** is *admissible* for adjustment for the causal effect of X on Y, if:

- $X, Y \notin$ **S**
- no element of **S** is a *descendant* of X
- **S** blocks all back-door paths $X \leftarrow .. Y$ (all paths between X and Y that start with an incoming arc on X)

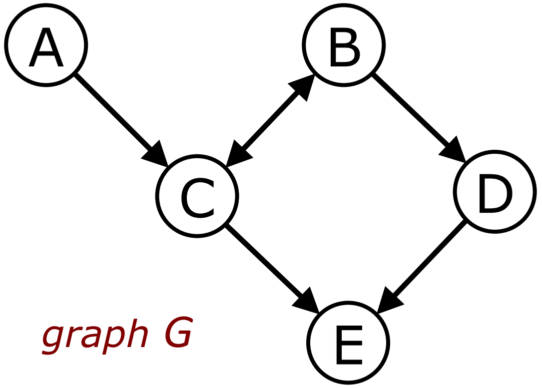

graph G

1. Is B admissible for adjustment to find the causal effect of D on E? *Yes.*
2. Is {} admissible for the causal effect of A on E? *Yes.*
3. Is B admissible for the causal effect of A on E? *Yes.*
4. Is {B,D} admissible for the causal effect of A on E? *Yes.*
5. Is C admissible for the causal effect of A on E? *No.*
6. Is {B,C} admissible for the causal effect of E on A? *Yes.*

Exercise 2b – Kidney stones revisited

	Treatment A	Treatment B
Small stones	81/87 (**93%**)	234/270 (87%)
Large stones	192/263 (**73%**)	55/80 (69%
	273/350 (78%)	289/350 (**83%**)
Total	562/700 (80%)	

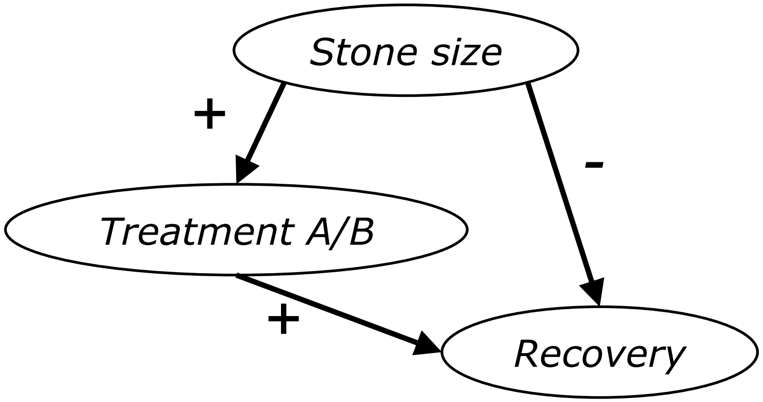

causal graph for kidney stone trial

$$p\left(Y = y \middle| do\left(X = x\right)\right)$$
$$= \sum_{Pa(X)} p\left(y \middle| x, Pa(X)\right) p\left(Pa(X)\right)$$

adjustment formula for intervention on X

1. Confirm that *Stone size* is a valid and necessary adjustment variable for the causal effect of *Treatment A/B* on *Recovery*.
2. Match the variables and values in the table above to the adjustment formula. In particular: what values need to be summed over?
3. Compute the causal effect of choosing *Treatment A* on *Recovery.*
4. Idem for the causal effect of *Treatment B*, and compare. What is the expected improvement (ACE) of choosing the optimal treatment?

Exercise 2b – Kidney stones revisited

	Treatment A	Treatment B
Small stones	81/87 (**93%**)	234/270 (87%)
Large stones	192/263 (**73%**)	55/80 (69%
	273/350 (78%)	289/350 (**83%**)
Total	562/700 (80%)	

$$p\left(Y=y \middle| do\left(X=x\right)\right)$$
$$=\sum_{Pa(X)} p\left(y \middle| x, Pa\left(X\right)\right) p\left(Pa\left(X\right)\right)$$

adjustment formula for intervention on X

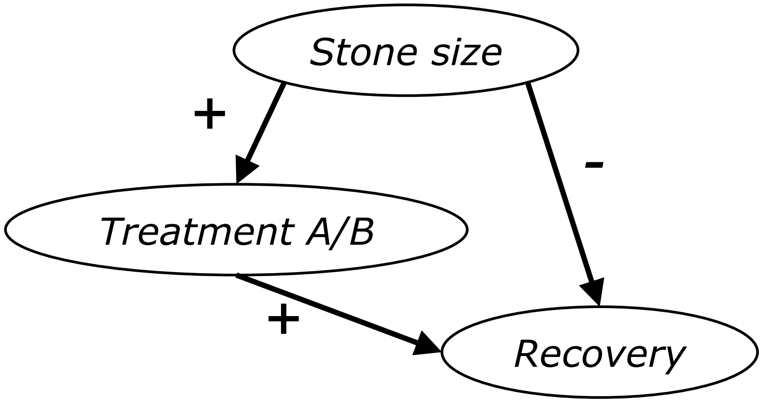

causal graph for kidney stone trial

Causal effect via adjustment

$$p\left(\text{Recovery} \middle| \hat{A}\right) = \sum_{S \in \{small, large\}} p\left(R \middle| T=A, Size=S\right) p\left(S\right) = 0.93 * 0.51 + 0.73 * 0.49 = 0.832$$

$$p\left(\text{Recovery} \middle| \hat{B}\right) = \sum_{S \in \{small, large\}} p\left(R \middle| T=B, Size=S\right) p\left(S\right) = 0.87 * 0.51 + 0.69 * 0.49 = 0.782$$

Outline

1. Introduction to causality
2. Prediction vs. causation
3. Causal graphs and how to read them
4. Cause-effect estimation
5. **The missing link & conclusion**

Another experiment: preventing eclampsia

	Treatment A	**Treatment B**
Low blood pressure	81/87 (**93%**)	234/270 (87%)
High blood pressure	192/263 (**73%**)	55/80 (69%
	273/250 (78%)	289/250 (**83%**)
Total recoveries	562/700 (80%)	

- different labels, exact same numbers … same conclusion?

Another experiment: preventing eclampsia

	Treatment A	**Treatment B**
Low blood pressure	81/87 (**93%**)	234/270 (87%)
High blood pressure	192/263 (**73%**)	55/80 (69%
	273/250 (78%)	289/250 (**83%**)
Total recoveries	562/700 (80%)	

- different labels, exact same numbers … same conclusion?

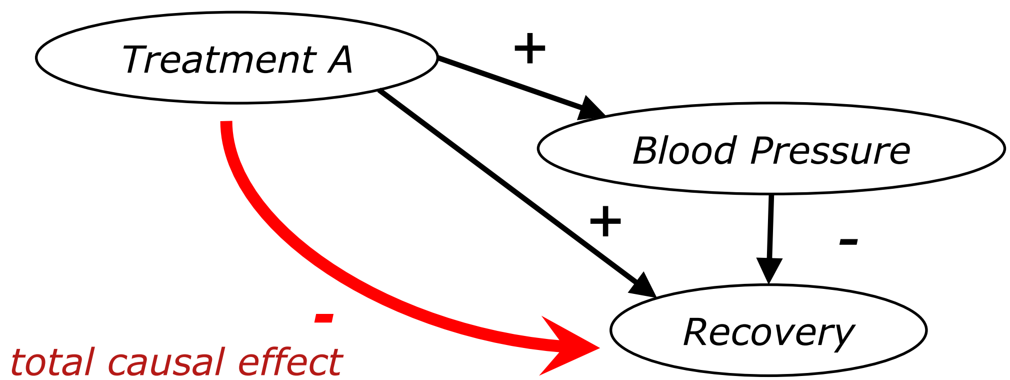

Conclusion

⇒ we need to know the true underlying causal graph to compute causal effects!

Take home messages so far …

- causality is a very useful concept
- if we want to tap into its potential we can and should use methods that treat it in a principled manner (we aim for validity, not truth)
- key feature is distinguishing between association and causation
- not always easy, but often doable

Take home messages so far …

- causality is a very useful concept
- if we want to tap into its potential we can and should use methods that treat it in a principled manner (we aim for validity, not truth)
- key feature is distinguishing between association and causation
- not always easy, but often doable

- graphical causal models offer an intuitive way to model causal structure
- we can link structure to data via structural equations / causal BNs
- if we know the causal model we can use e.g. the *back-door criterion* and *adjustment* to compute/predict post-interventional distributions
- leading to quantities of interest such as the *Average Causal Effect*

But much more to follow in the next two days!

Take home messages so far …

- causality is a very useful concept
- if we want to tap into its potential we can and should use methods that treat it in a principled manner (we aim for validity, not truth)
- key feature is distinguishing between association and causation
- not always easy, but often doable

- graphical causal models offer an intuitive way to model causal structure
- we can link structure to data via structural equations / causal BNs
- if we know the causal model we can use e.g. the *back-door criterion* and *adjustment* to compute/predict post-interventional distributions
- leading to quantities of interest such as the *Average Causal Effect*

But much more to follow in the next two days!

Thank you!