Representing Causal Structures

Qualitative Representation

- We want a representation that captures many qualitative features of causality
  - Causation occurs among variables ⇒ One node per variable

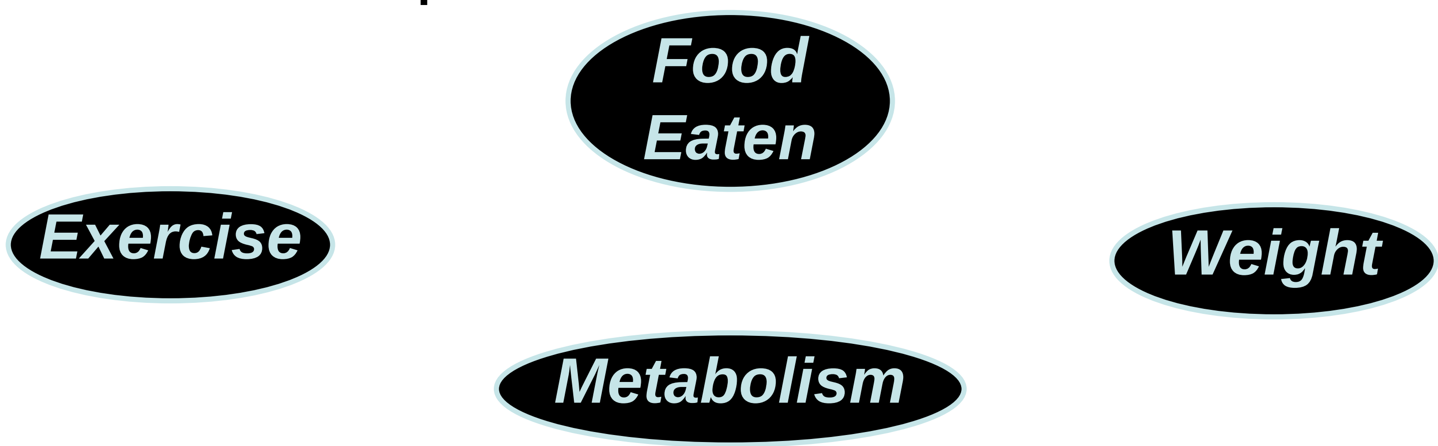

Qualitative Representation

- We want a representation that captures many qualitative features of causality
  - Asymmetry of causation ⇒ Need an asymmetric connection in the graph

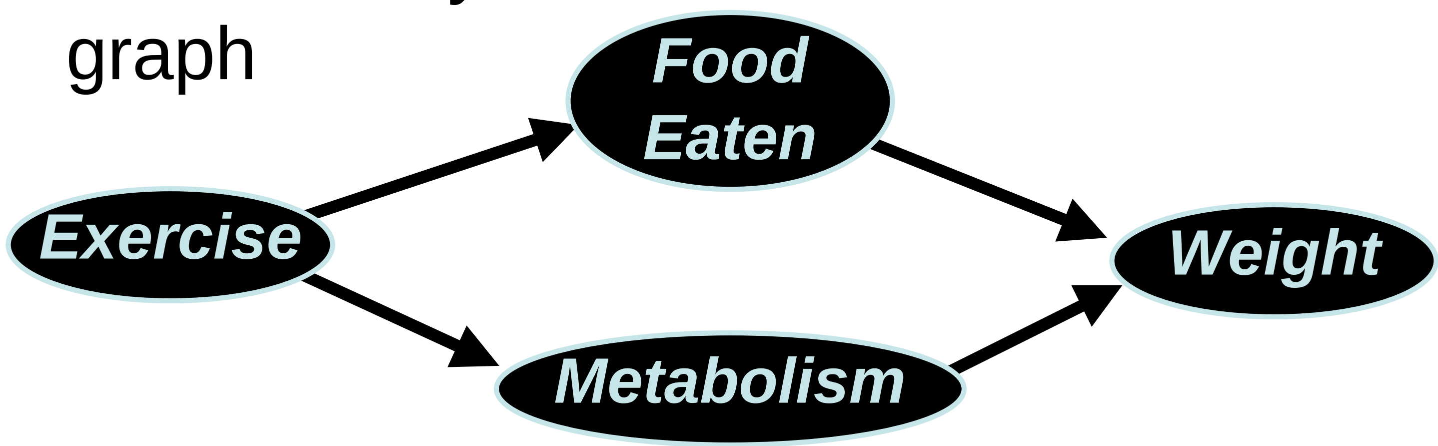

Qualitative Representation

- We want a representation that captures many qualitative features of causality
  - No (immediate) reciprocal causation ⇒ No cycles (without explicit temporal indexing)

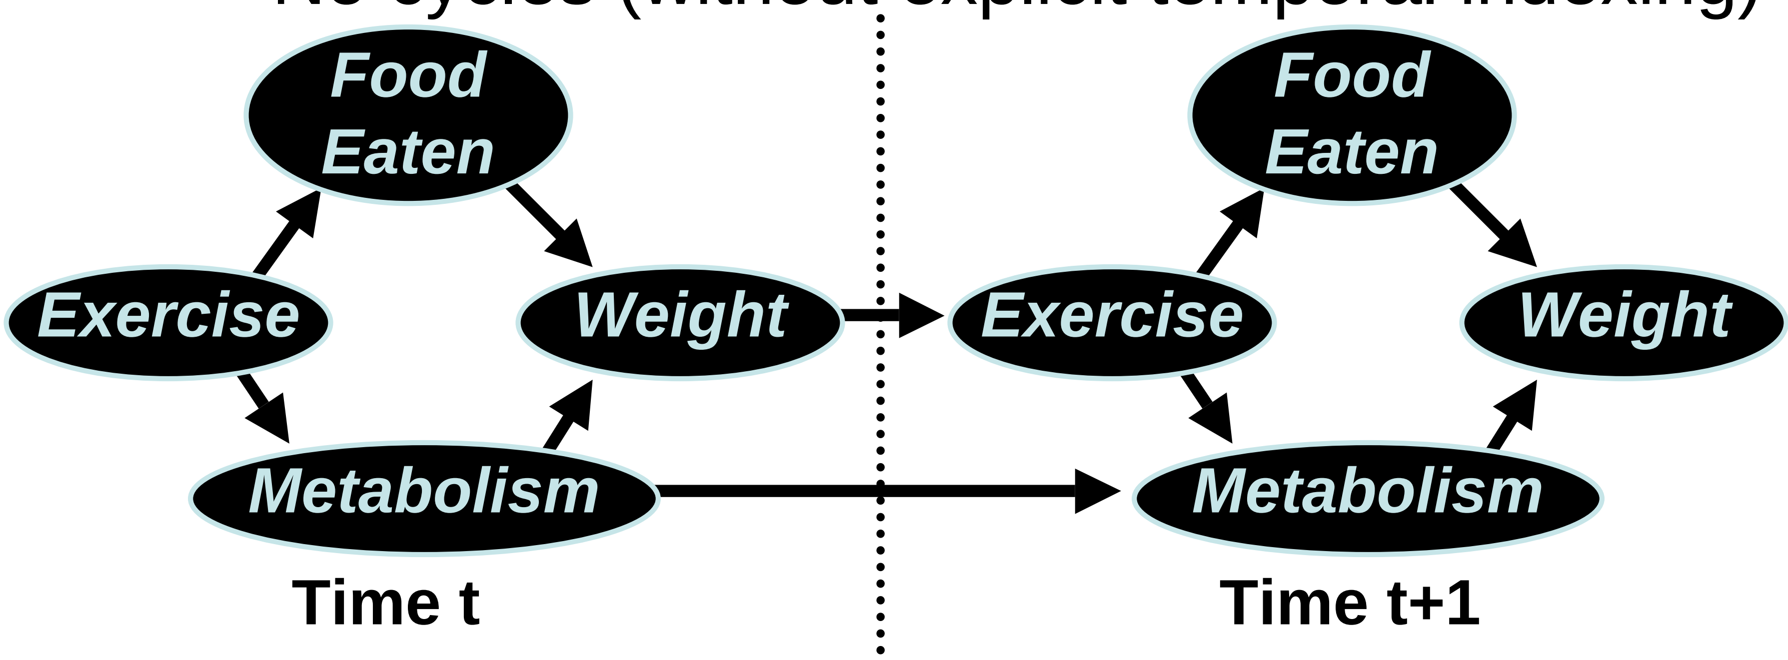

Directed Acyclic Graphs

- In general, Directed Acyclic Graphs are:
  - Defined for a set of variables
    - Some variables might be unobserved (latents)
  - One node per variables, connected by directed edges, without directed loops
    - I.e., can’t start at a node, follow arrow directions, and get back to the original node
  - Note: We use genealogical terminology when talking about relationships within a graph

Quantitative Representation

- DAGs alone can represent “A causes B”…
  but not “strength” or “form” of causation
  - Need to represent the relationships between the various variables *states*
  - Exact quantitative representation will depend on the type of variables being represented

Bayesian Networks

- All variables are discrete/categorical
- Represent quantitative causation using a *joint probability distribution*
  - I.e., a specification of the probability of any combination of variable values, such as:
    - *P*(*E*=Hi & *FE*=Lo & *M*=Hi & *W*=Hi) = 0.001;
      P(*E*=Hi & *FE*=Lo & *M*=Hi & *W*=Lo) = 0.03;
      etc.

- Note: Nothing inherently Bayesian about Bayes nets!

Structural Equation Models

- All variables are continuous/real-valued
- Represent quantitative causation using systems of linear equations
  - For example:
    $Exercise = a_1 FE + a_2 M + a_3 W + \varepsilon_{E\_noise}$
    $FE = b_1 E + b_2 M + b_3 W + \varepsilon_{FE\_noise}$
    etc.

Connecting the Pieces

- Causal Markov assumption:
  - Variables are independent of their non-effects conditional on their direct causes
    - Use the qualitative graph to constrain the quantitative relationships
  - Encodes the intuition of “screening off”
    - Given the values of the direct causes, learning the value of a non-effect doesn’t help me predict

Connecting the Pieces

- Markov assumption for Bayes nets:
  - Probability distribution can be expressed as product of: $P(X \mid \text{parents}(X))$ for all X
  - Example:

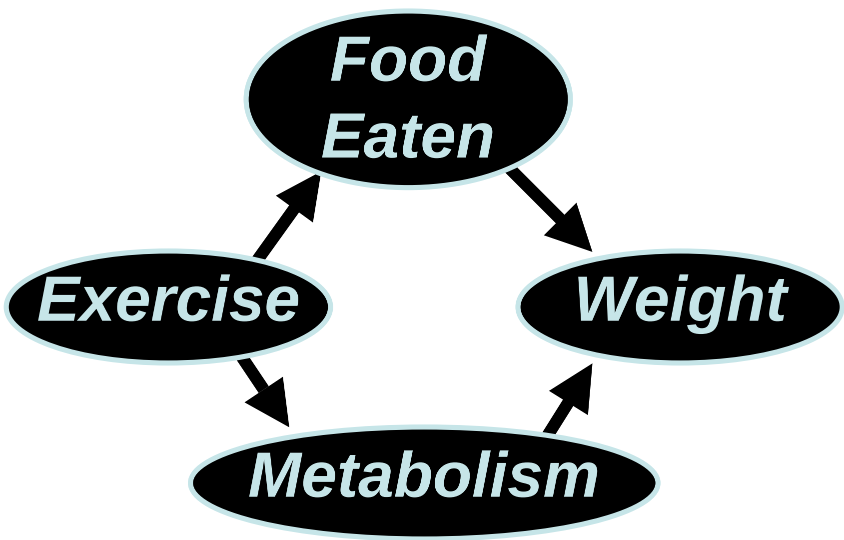

$\Rightarrow$

$$P(E, FE, W, M) = P(E) \times P(FE \mid E) \times P(M \mid E) \times P(W \mid M, FE)$$

Connecting the Pieces

- Markov assumption for SEMs:
  - Each variable's linear equation only takes its parents as arguments
  - Example:

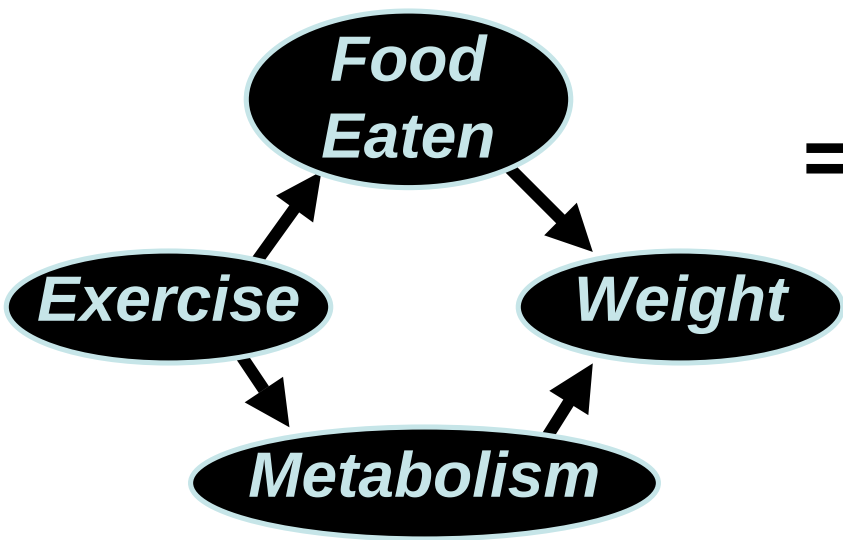

$\Rightarrow$

$$E = \varepsilon_{E\_noise}$$
$$FE = a_1 E + \varepsilon_{FE\_noise}$$
$$M = b_1 E + \varepsilon_{M\_noise}$$
$$W = c_1 FE + c_2 M + \varepsilon_{C\_noise}$$

Connecting the Pieces

- Causal Faithfulness assumption
  - The only independencies are those predicted by the Markov assumption
    - Uses the quantitative relations to constrain the qualitative graph
    - Most important implication: No exactly counter-balancing causal paths
      - *Exercise* → *Food Eaten* → *Weight* and *Exercise* → *Metabolism* → *Weight* do not exactly offset one another

Causal vs. Statistical Models

- Bayes nets and SEMs are not inherently causal models
  - Markov and Faithfulness assumptions can be expressed purely as graph-quant. constraints

- Assuming a non-causal version of the assumptions $\Rightarrow$ purely statistical model
  - I.e., a compact representation of statistical independencies among some set of variables

Three Uses

- Efficiently determine independencies
  - I.e., which variables are informationally relevant for which other ones?
- Use those independencies to rapidly update beliefs in light of evidence
- Represent (and predict the effects of) interventions on variables
  - Causal models only, of course

Determining Independence

- Given the assumptions, graph structure alone determines *all* of the statistical independencies and associations
- Graphical criterion: d-separation
  - *X* and *Y* are independent given ***S*** iff
    X and *Y* are d-separated given ***S*** iff
    X and *Y* are not d-connected given ***S***
- Intuition: *X* and *Y* are d-connected iff information can "flow" from *X* to *Y*

d-separation

- Formally:
  - *C* is a collider on a path iff $A \rightarrow C \leftarrow B$
  - A path between *A* and *B* is *activated by* ***S*** iff
    - Every non-collider on the path is not in ***S***; and
    - Every collider on the path is either in ***S***, or else one of its descendants is in ***S***
  - *X* and *Y* are d-connected by ***S*** iff there is an active path between *X* and *Y* given ***S***

d-separation

- Surprising feature being exploited here:
  - Conditioning on a common effect induces an association between independent causes
  - Motivating example:

    Gas Tank → *Car Starts* ← *Spark Plugs*

    - *Gas* and *Plugs* are independent, but if we know that the car doesn't start, then they're associated
      - In that case, learning *Gas* = Full changes the likelihood that *Plugs* = Bad
  - And similarly if *Car Starts* → *Emits Exhaust*

d-separation

- *Exercise* and *Weight* given *Metabolism*?
  - $E \rightarrow M \rightarrow W$
    - Blocked! M is an included non-collider
  - $E \rightarrow FE \rightarrow W$
    - Unblocked! FE is a non-included non-collider

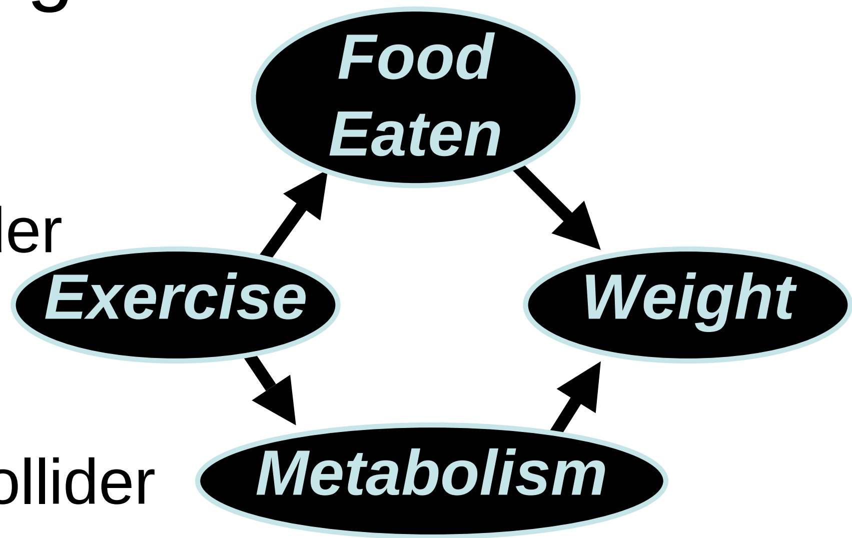

- $\Rightarrow E \not\perp\!\!\!\perp W \mid M$

d-separation

- *Metabolism* and *FE* given *Exercise*?
  - $M \rightarrow W \leftarrow FE$
    - Blocked! *W* is a non-included collider
  - $M \leftarrow E \rightarrow FE$
    - Blocked! *E* is an included non-collider

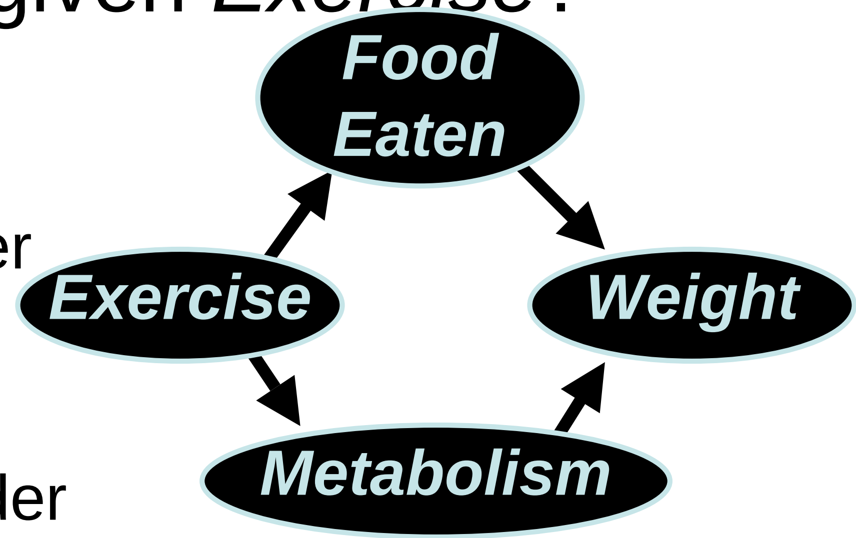

- $\Rightarrow M \perp\!\!\!\perp FE \mid E$

d-separation

- *Metabolism* and *FE* given *Weight*?
  - $M \rightarrow W \leftarrow FE$
    - Unblocked! *W* is an included collider
  - $M \leftarrow E \rightarrow FE$
    - Unblocked! *E* is a non-included non-collider

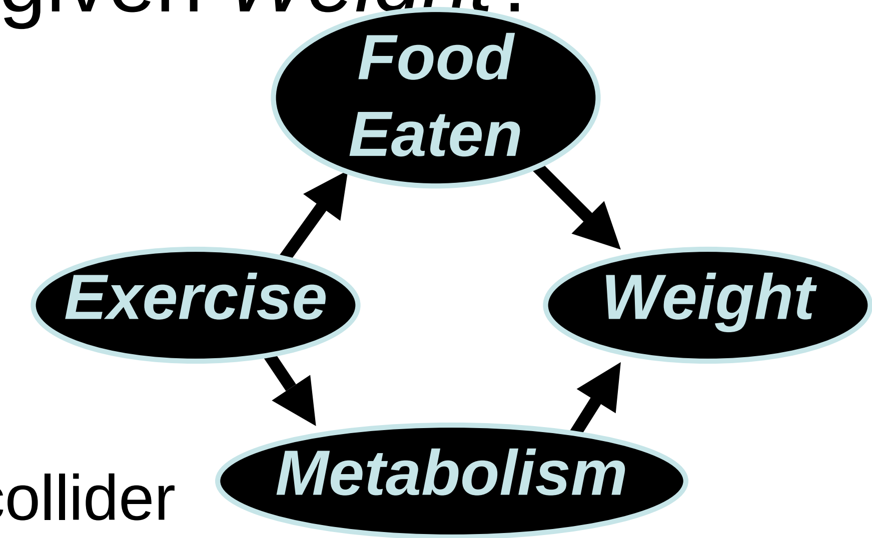

- $\Rightarrow M \not\perp\!\!\!\perp FE \mid W$

Updating Beliefs

- For both statistical and causal models, efficient computation of independencies
  $\Rightarrow$efficient prediction from observations
  - Specific instance of *belief updating*
  - "Just" compute conditional probabilities
    - Significantly easier if we have (conditional) independencies, since we can ignore variables

Updating Beliefs

- Compute: $P(M = \text{Hi} \mid E = \text{Hi}, FE = \text{Lo})$
  - $FE \perp\!\!\!\perp M \mid E \Rightarrow$
    $P(M \mid E, FE) = P(M \mid E)$
    - And $P(M \mid E)$ is a term in the Markov decomposition!

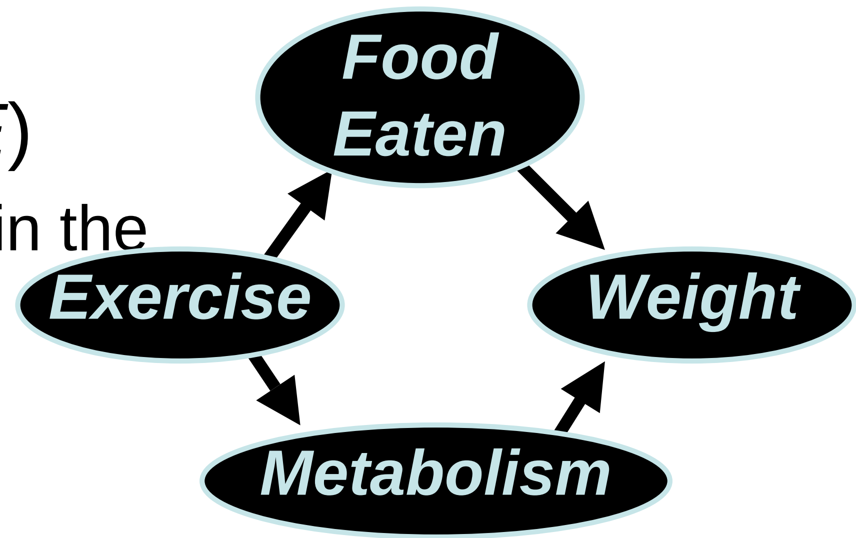

Representing Interventions

- <u>Central intuition:</u> When we intervene, we control the state of the target variable
  - And so the direct causes of the target variable no longer matter
  - But the target still has its usual effects
    - Directly applying current to the light bulb ⇒ light switch doesn’t matter, but the plant still grows

Representing Interventions

- Formal implementation:
  - Add a variable representing the intervention, and make it a direct cause of the target
  - When the intervention is “active”, remove all other edges into the target
  - Leave intact all edges directed out of the target, even when the intervention is “active”

Representing Interventions

- Example:
  - Add a manipulation variable as a “cause”

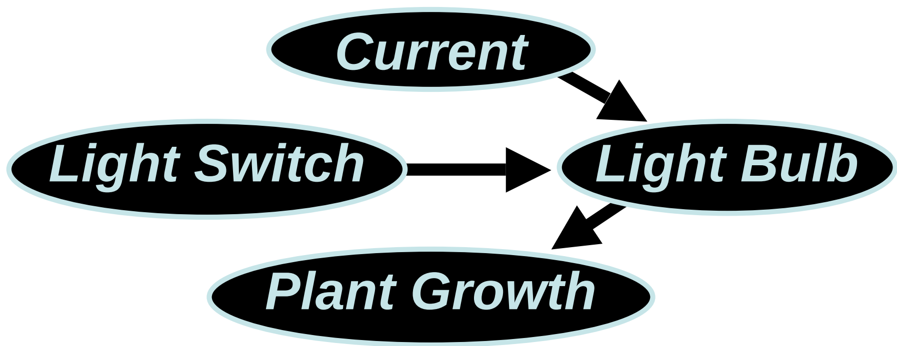

Representing Interventions

- Example:
  - Add a manipulation variable as a “cause” that does not matter when it is inactive

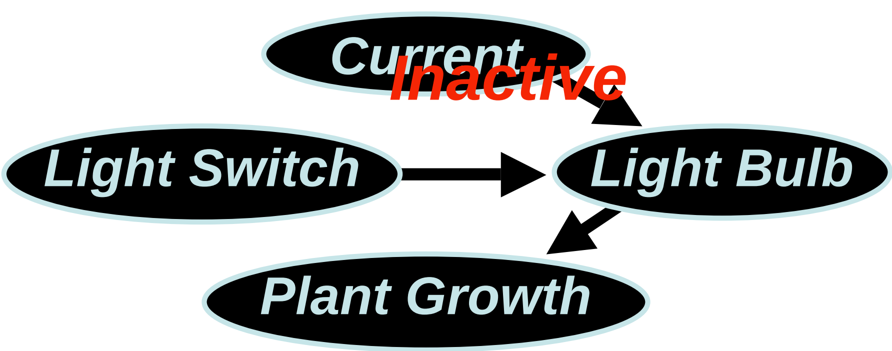

Inactive Manipulation

Representing Interventions

- Example:
  - Add a manipulation variable as a “cause” that does not matter when it is inactive
  - When it is active,

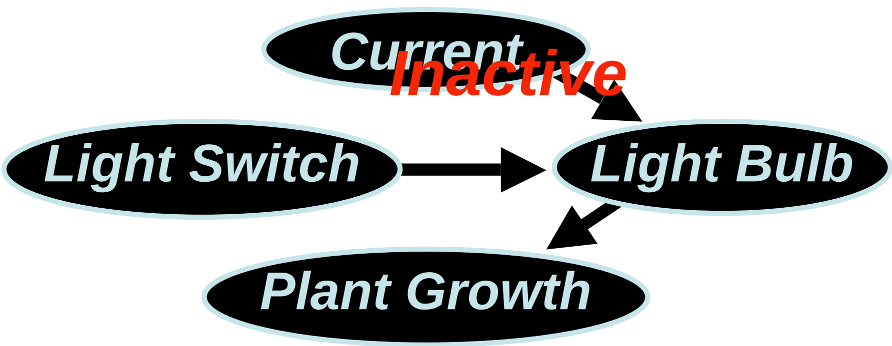

Inactive Manipulation

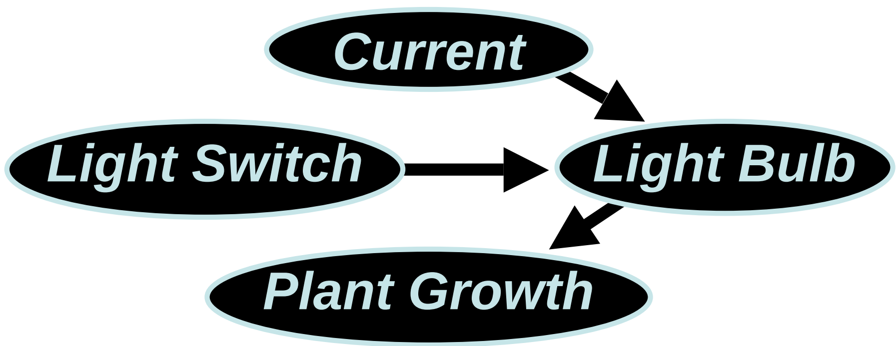

Active Manipulation

Representing Interventions

- Example:
  - Add a manipulation variable as a “cause” that does not matter when it is inactive
  - When it is active, break the incoming edges, but leave the outgoing edges

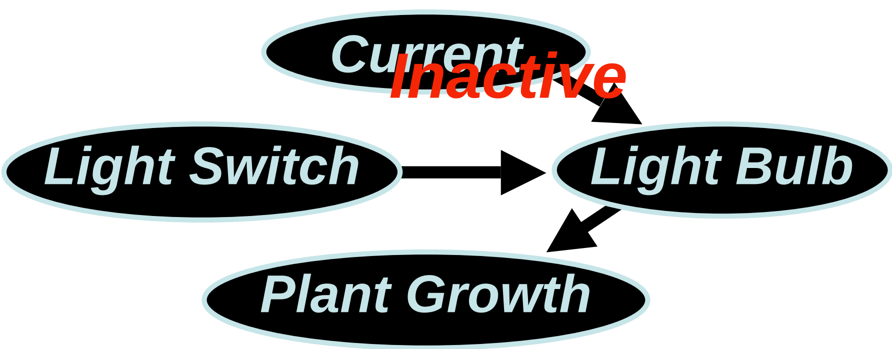

Inactive Manipulation

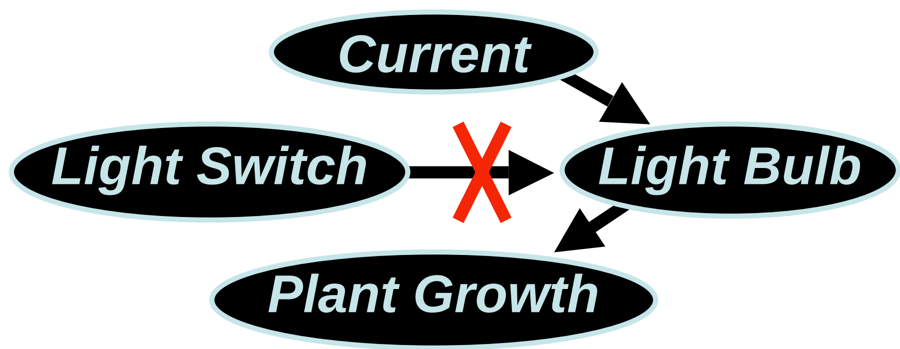

Active Manipulation

Representing Interventions

- Straightforward extension to more interesting types of interventions
  - Interventions away from current state
  - Multi-variate interventions
  - Etc.

Why Randomize?

- Standard scientific practice: randomize *Treatment* to find its *Effect*s
  - E.g., don't let people decide on their own whether to take the drug or placebo
- What is the value of randomization?
  - Randomization is an intervention
    - $\Rightarrow$ All edges into *T* will be broken, including from any common causes of *T* and *E*!
    - $\Rightarrow$ If $T \not\perp\!\!\!\perp E$, then we must have: $T \rightarrow E$

Why Randomize?

- Graphically,

Why Randomize?

- Graphically,

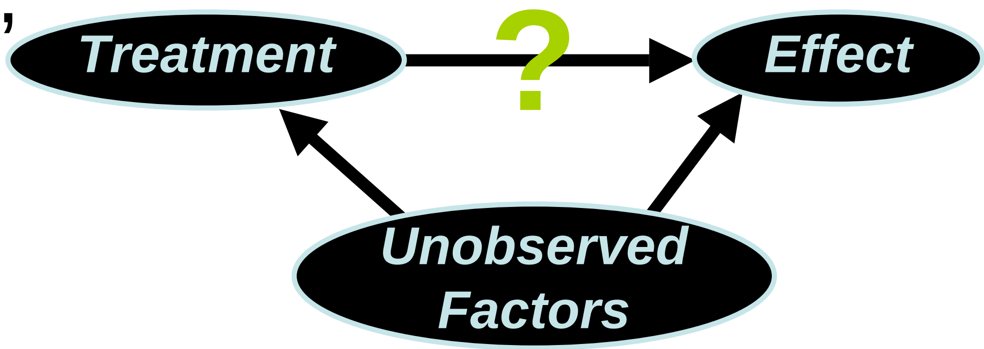

Why Randomize?

- Graphically,

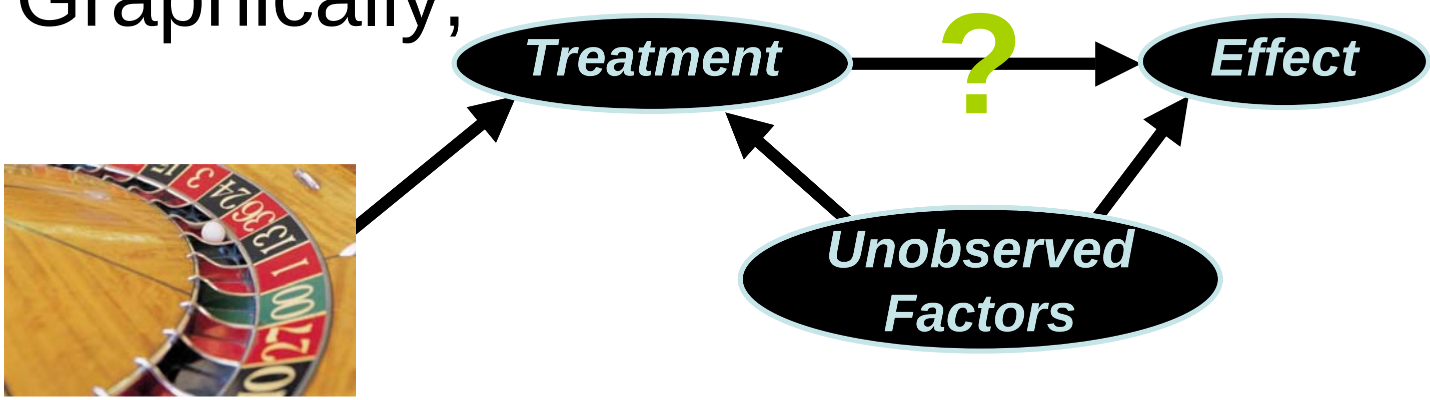

Why Randomize?

- Graphically,

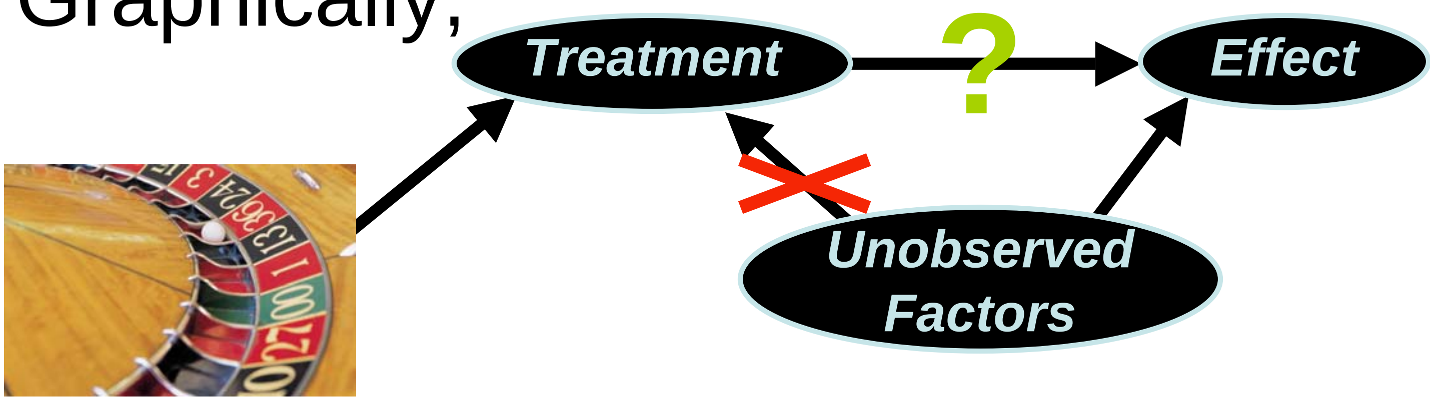

Why Randomize?

- Graphically,