Causality and the semantics of provenance

James Cheney

DCM 2010

July 9, 2010

What is provenance?

- Generally:
  - history, record of ownership, origins
- Computationally:
  - metadata needed to understand process that created some result
  - information that makes computation/data more "transparent", "trustworthy"

Why is provenance important?

- long-term record keeping
- debugging, data cleaning, error diagnosis
- scientific repeatability
  - data & provenance required by some journals
- trust, accountability, transparency
  - i.e., climategate-prevention

Goal: Semantic foundations for provenance

Goal: Semantic foundations for provenance

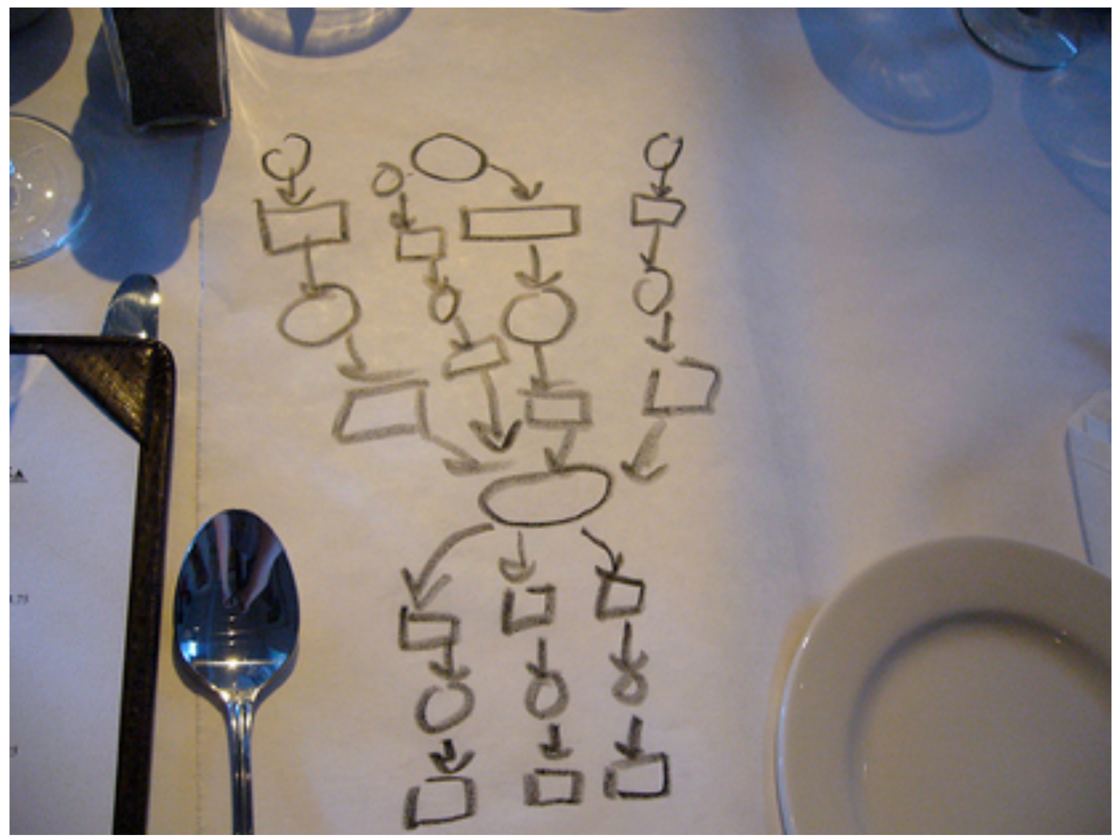

http://www.flickr.com/photos/schneertz/679692806/

Why is provenance semantics important?

- Most work on provenance is of the "follow your nose" school.
- Little attention to semantics, foundations.
- If this information is important, then its meaning should be clear.
- If its meaning is not clear, it is unlikely to have long-term value.

Causality

- Causality is frequently invoked as a motivation
- For example, Open Provenance Model (OPM) says:
    - “edges denote causal relationships linking the cause to the effect”
- This seems a bit cavalier
    - not made clear in what sense a provenance graph "describes" a computational process

Causality

- Causality has long been studied by philosophers
  - Hume, many others

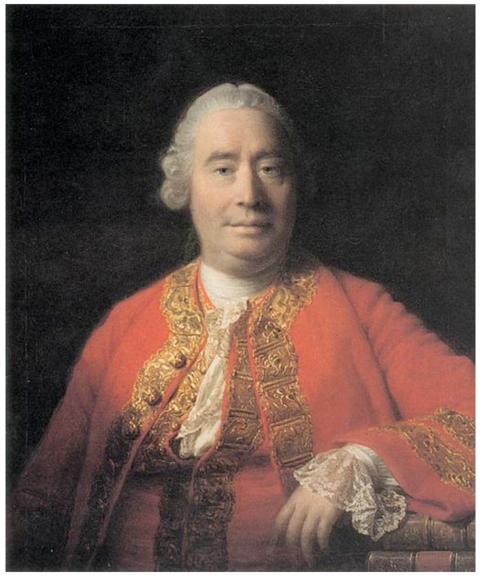

- More recently, also in AI/CS
  - Halpern, Pearl, many others

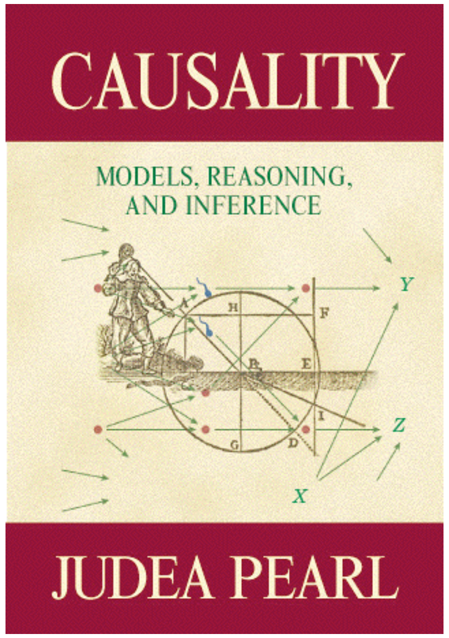

This talk

- Quick review of
  - Open Provenance Model-style graphs
  - *structural causal models*
  - Halpern-Pearl definition of "actual cause"
- Using causal models to interpret provenance graphs
  - how they match and don't

Warning

Cake

Baking a cake: The OPM way

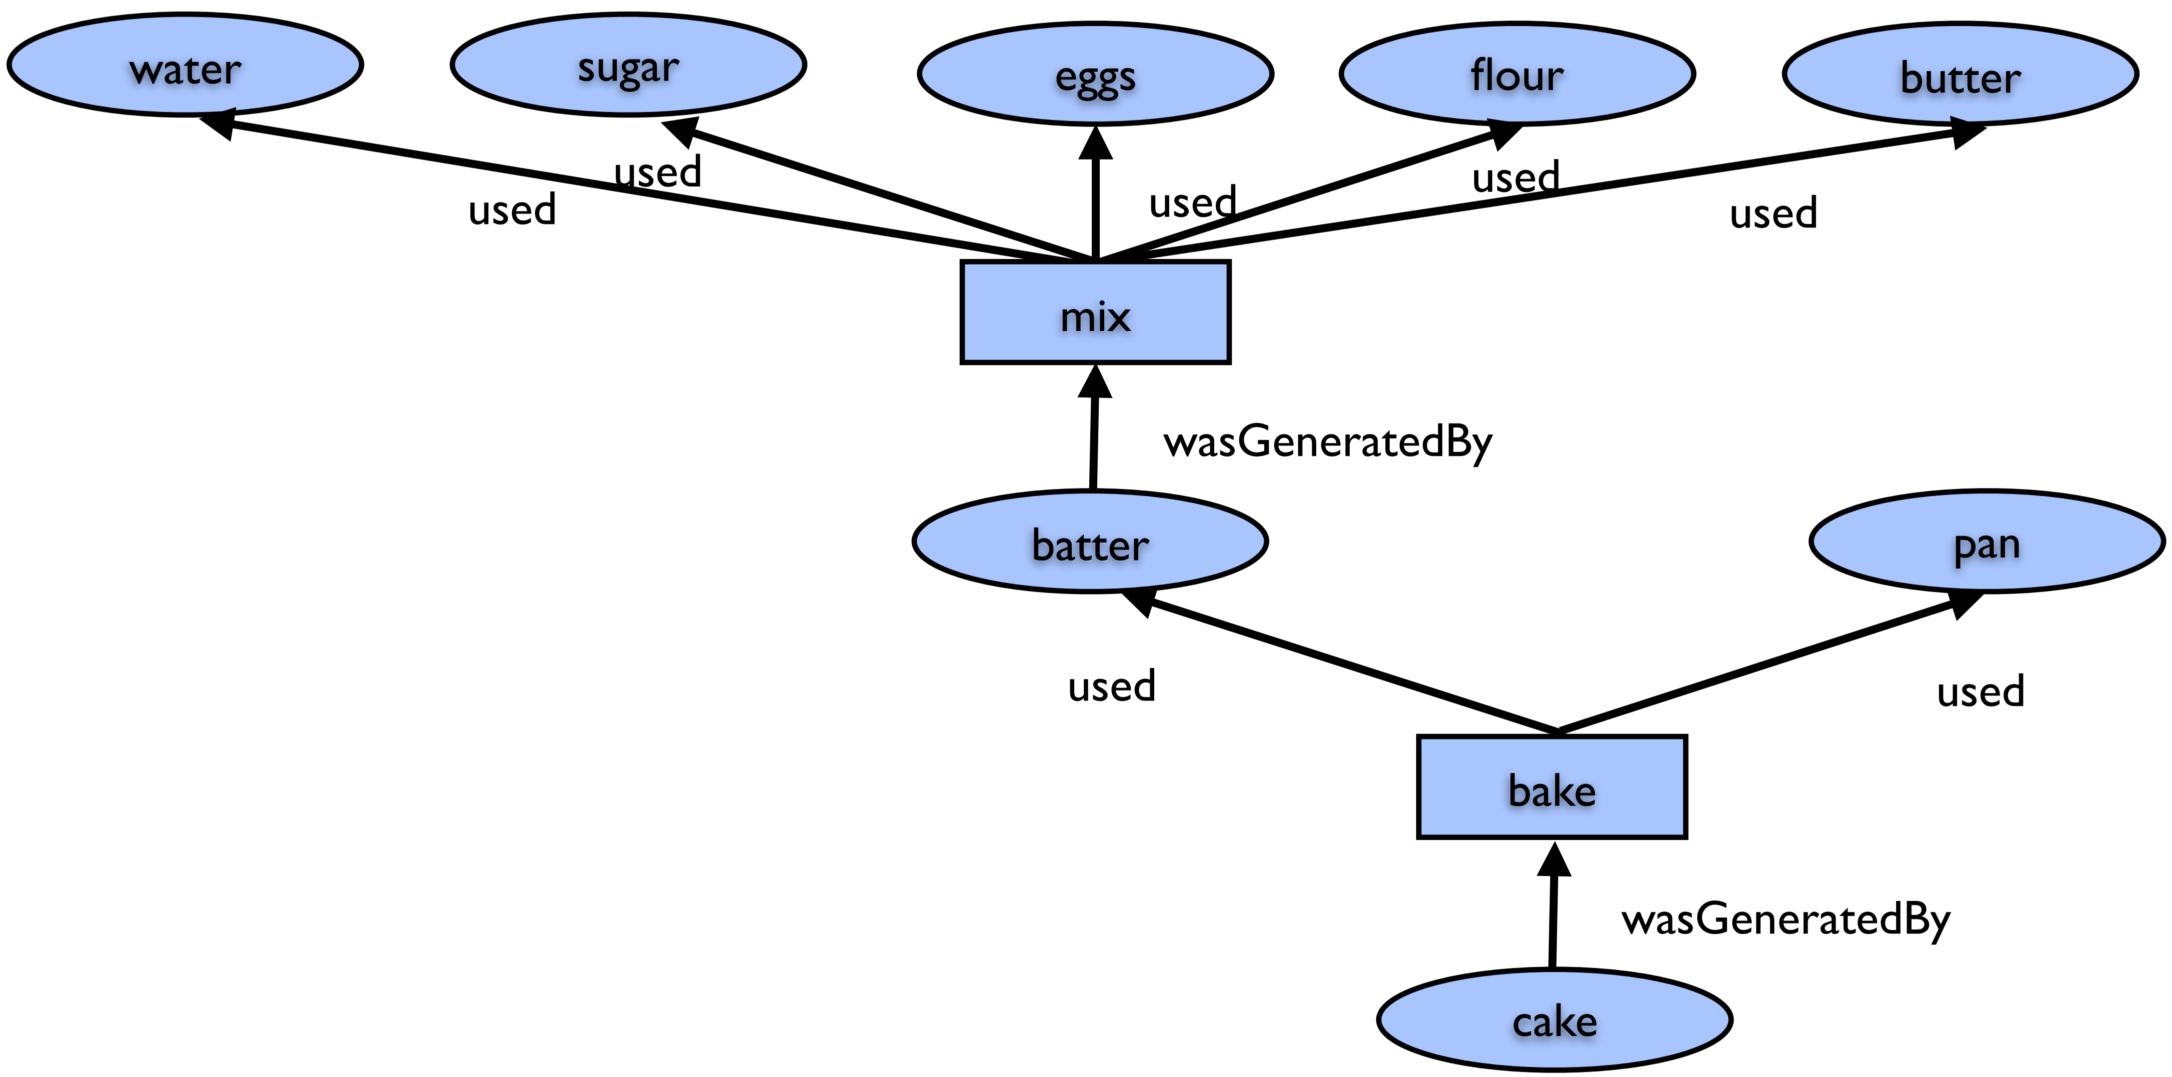

Baking a cake: The OPM way

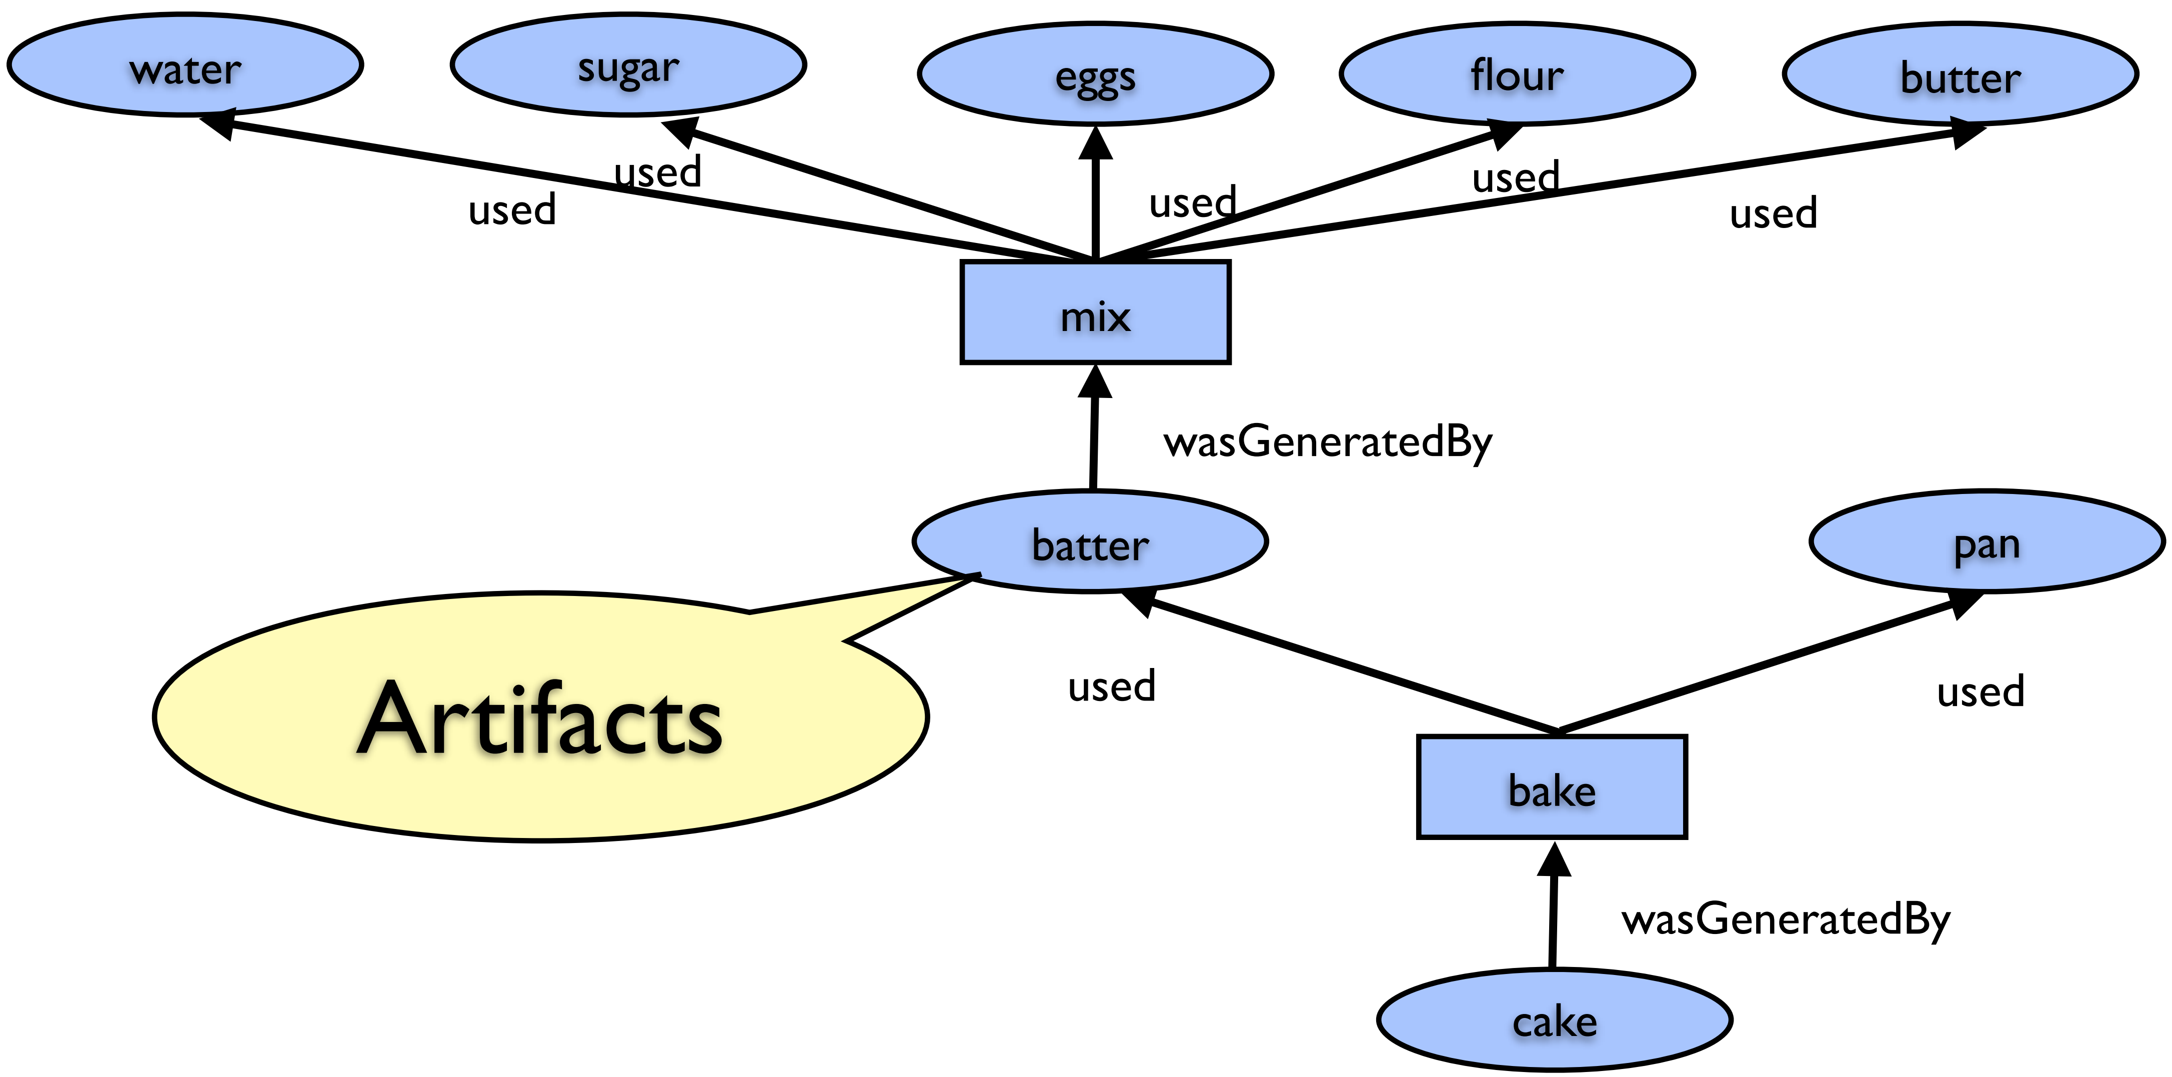

Baking a cake: The OPM way

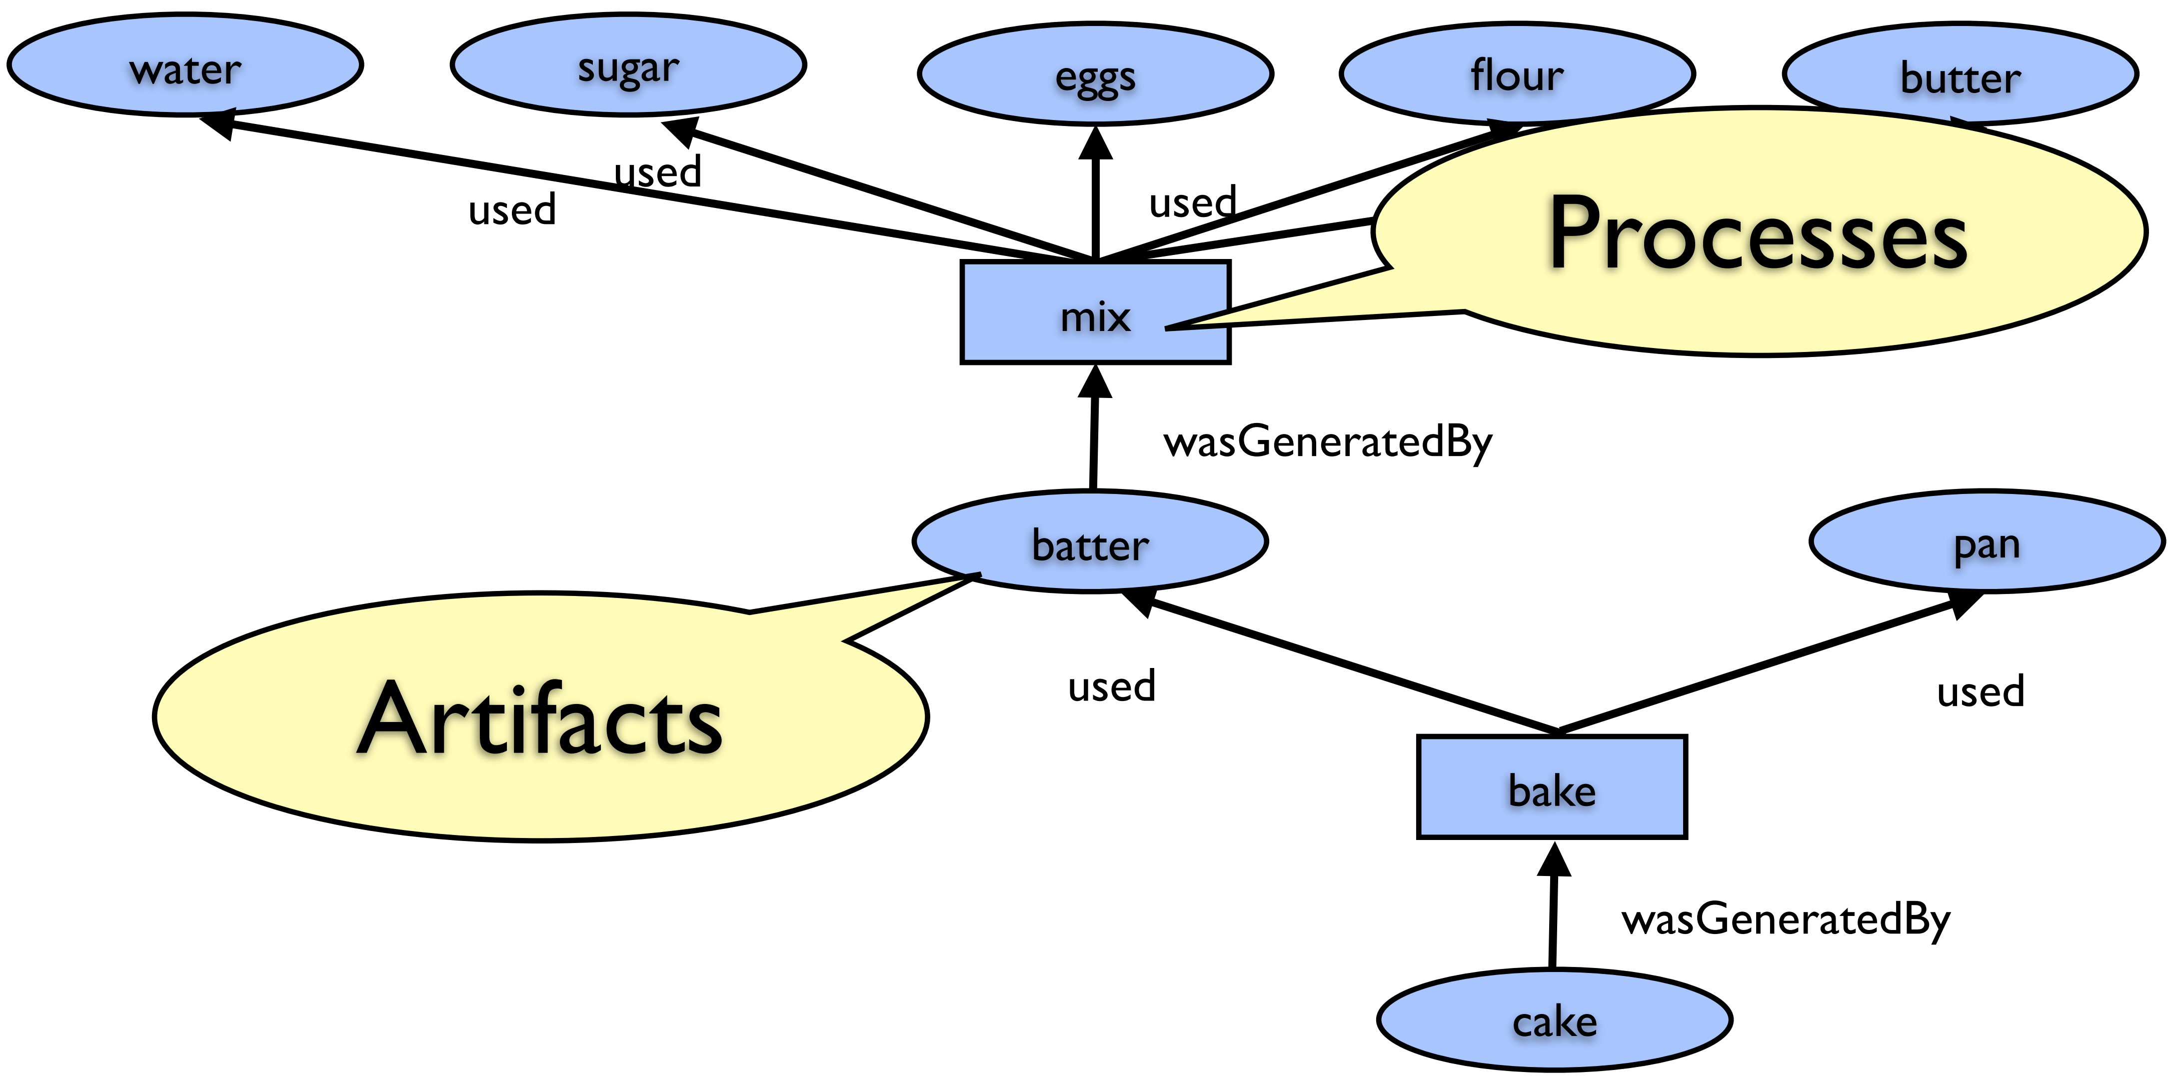

Functional interpretations

- We can interpret a provenance graph as a function in the obvious way
  - assign functions to nodes (matching arities)
- Then a "correct provenance graph" describing a function is one that has the same interpretation.
- However, this is not very satisfying...

Example

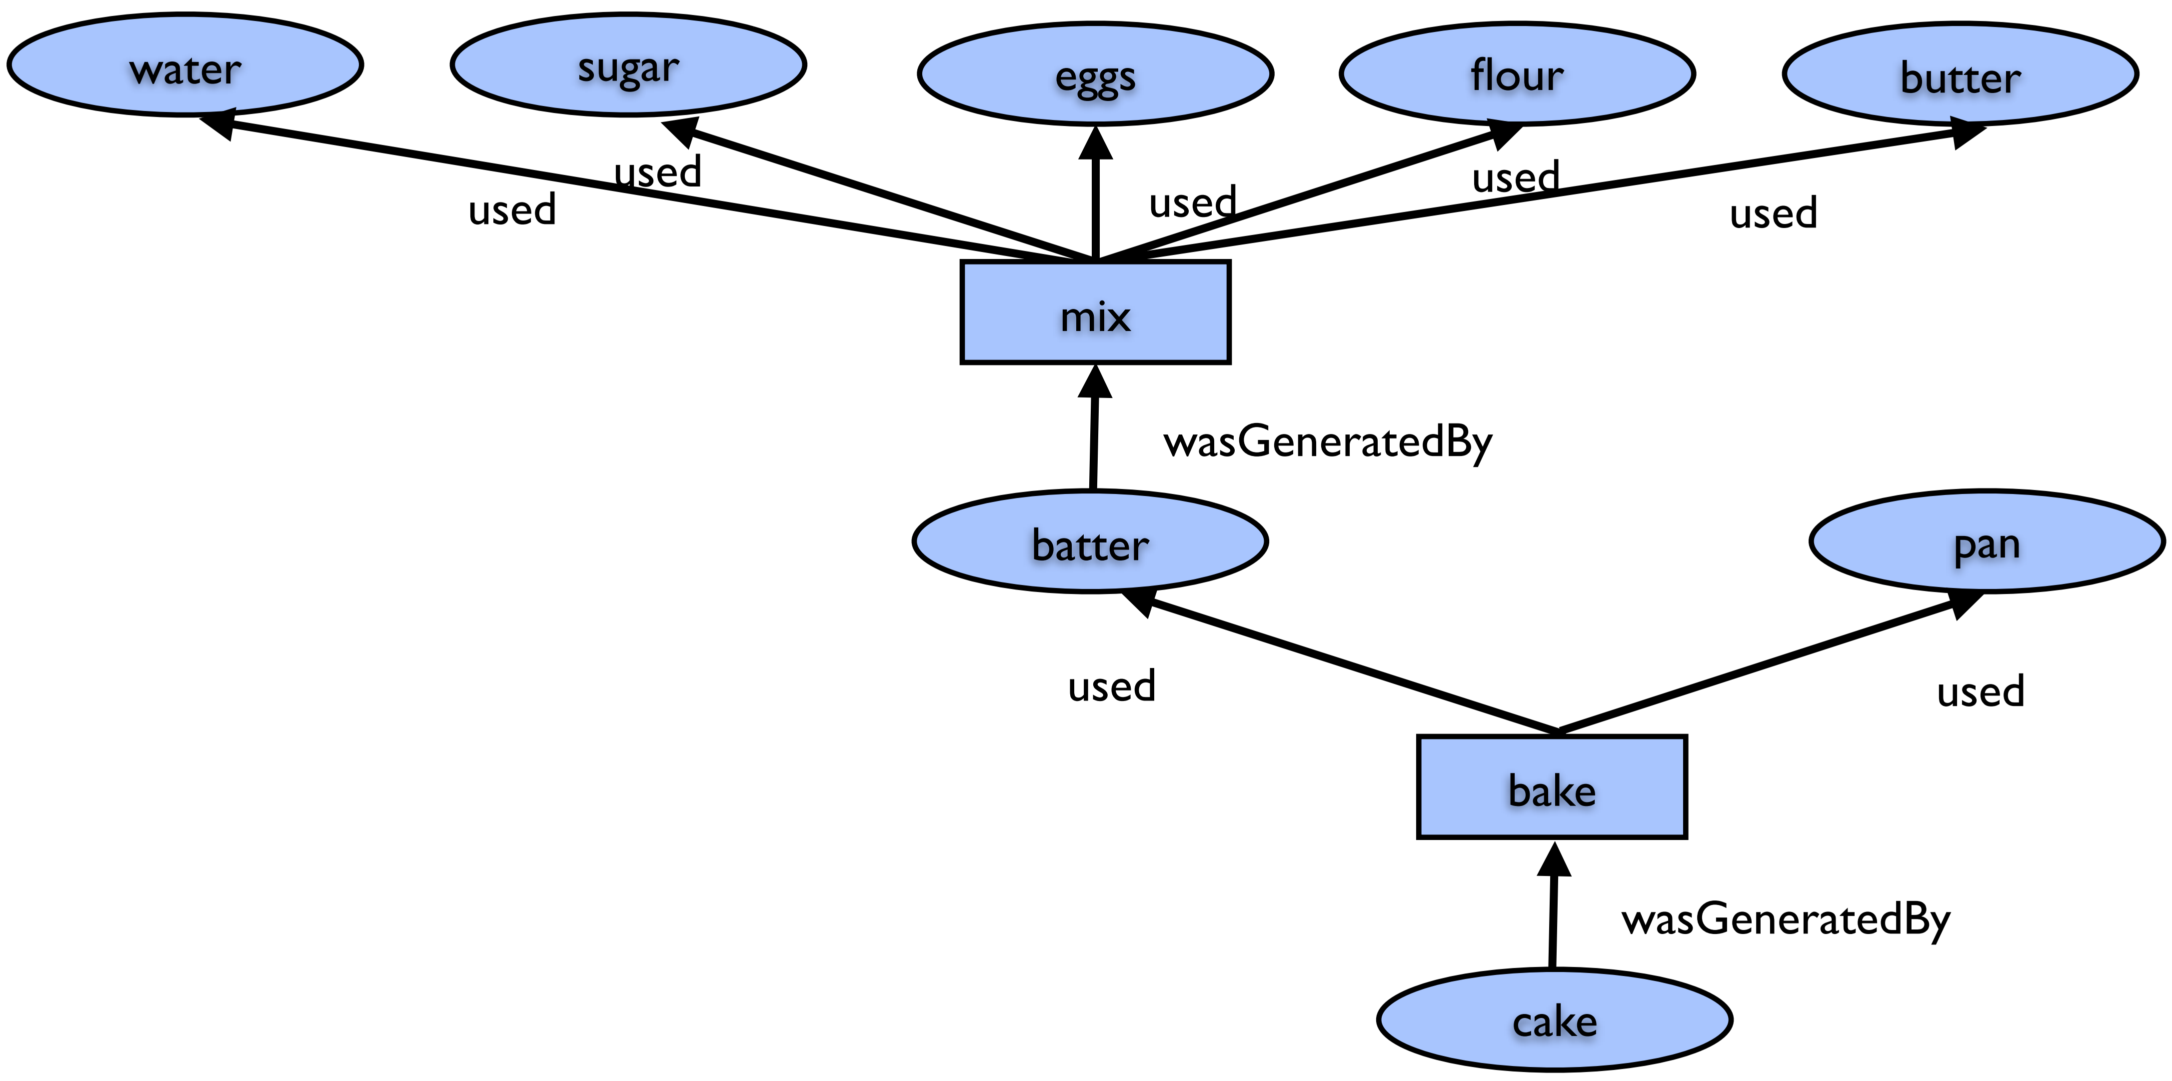
water
sugar
eggs
flour
butter
used
used
used
used
used
mix
wasGeneratedBy
batter
pan
used
used
bake
wasGeneratedBy
cake

Example

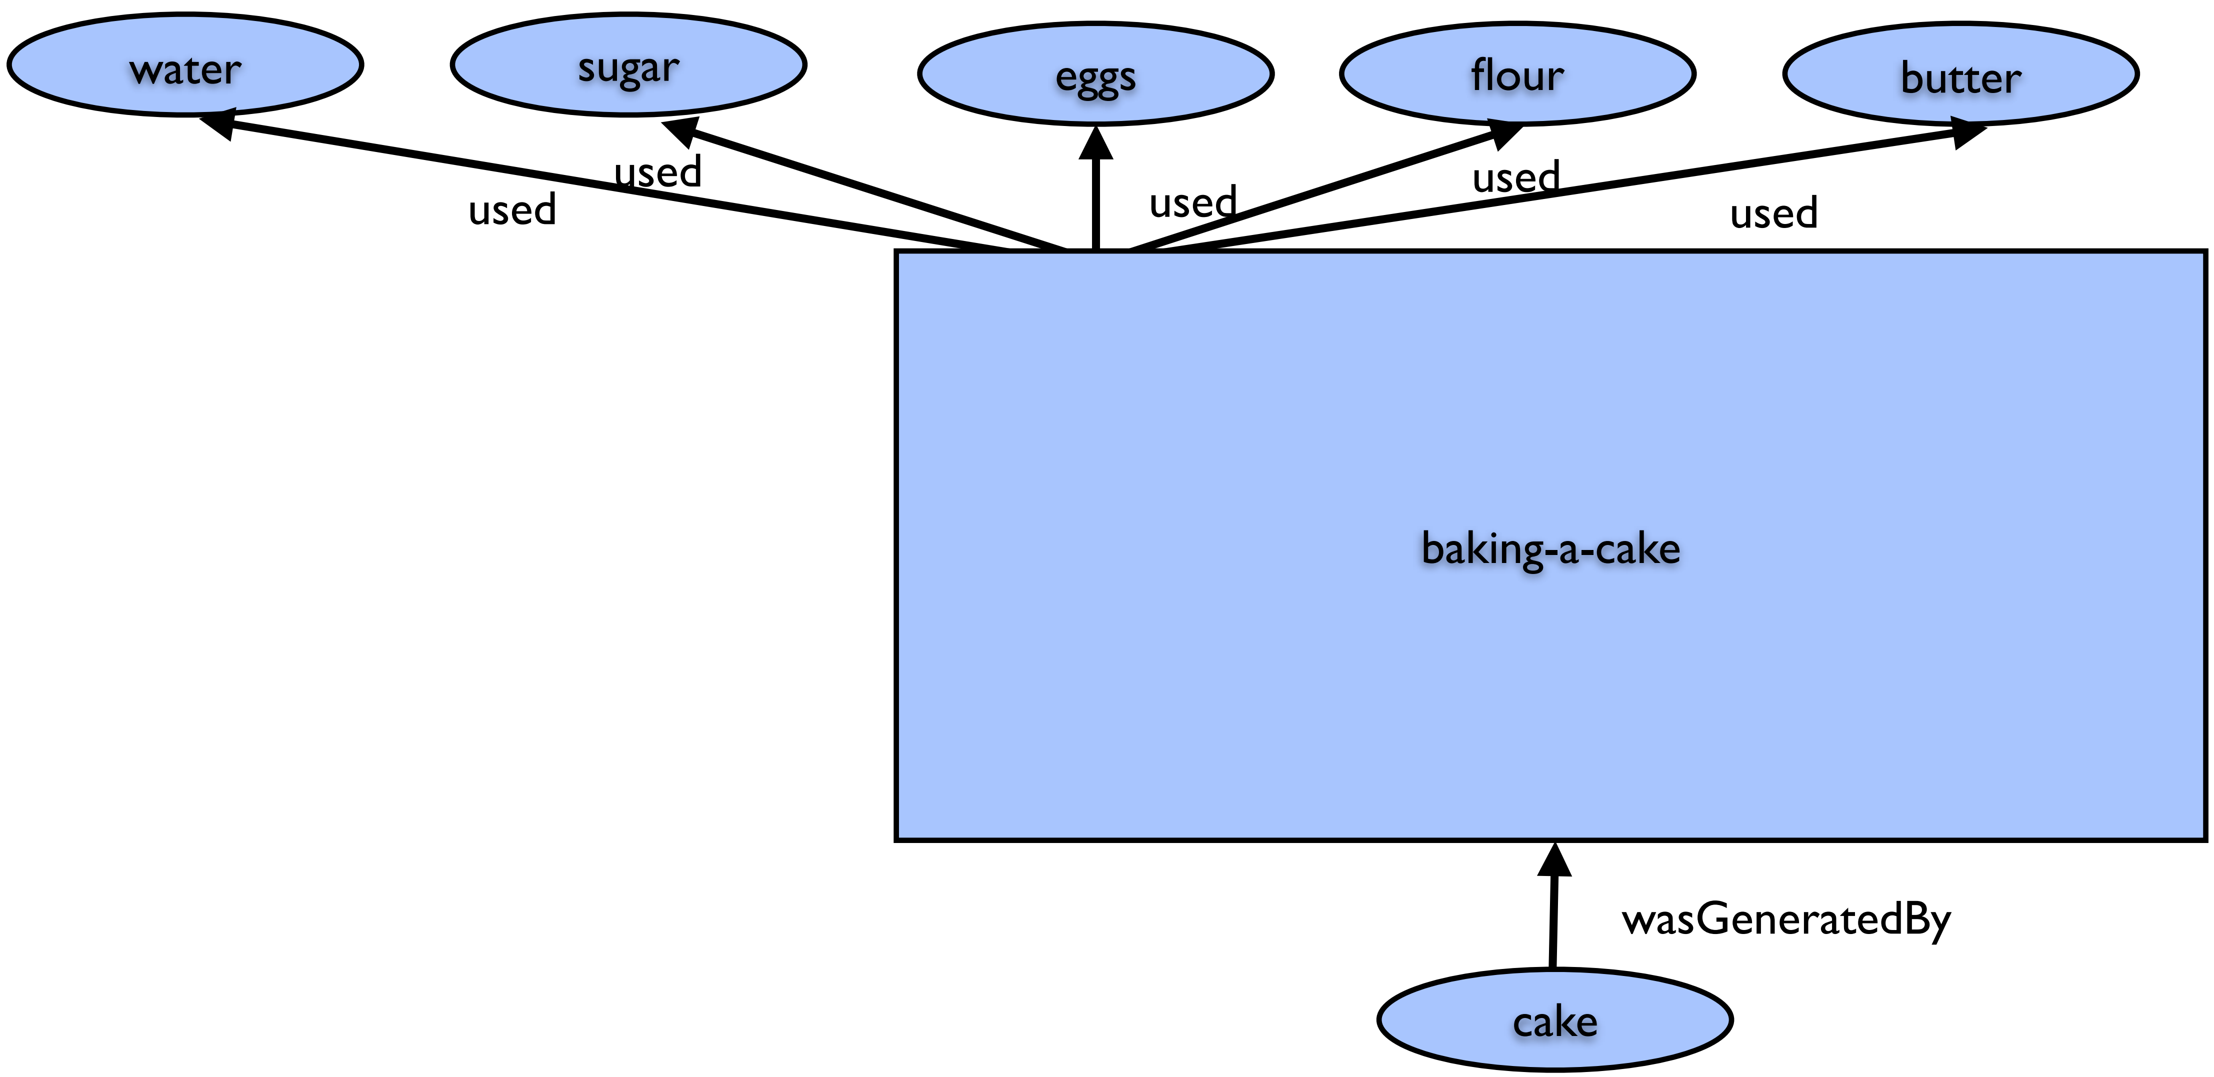

water
sugar
eggs
flour
butter
used
used
used
used
used
baking-a-cake
wasGeneratedBy
cake

Example

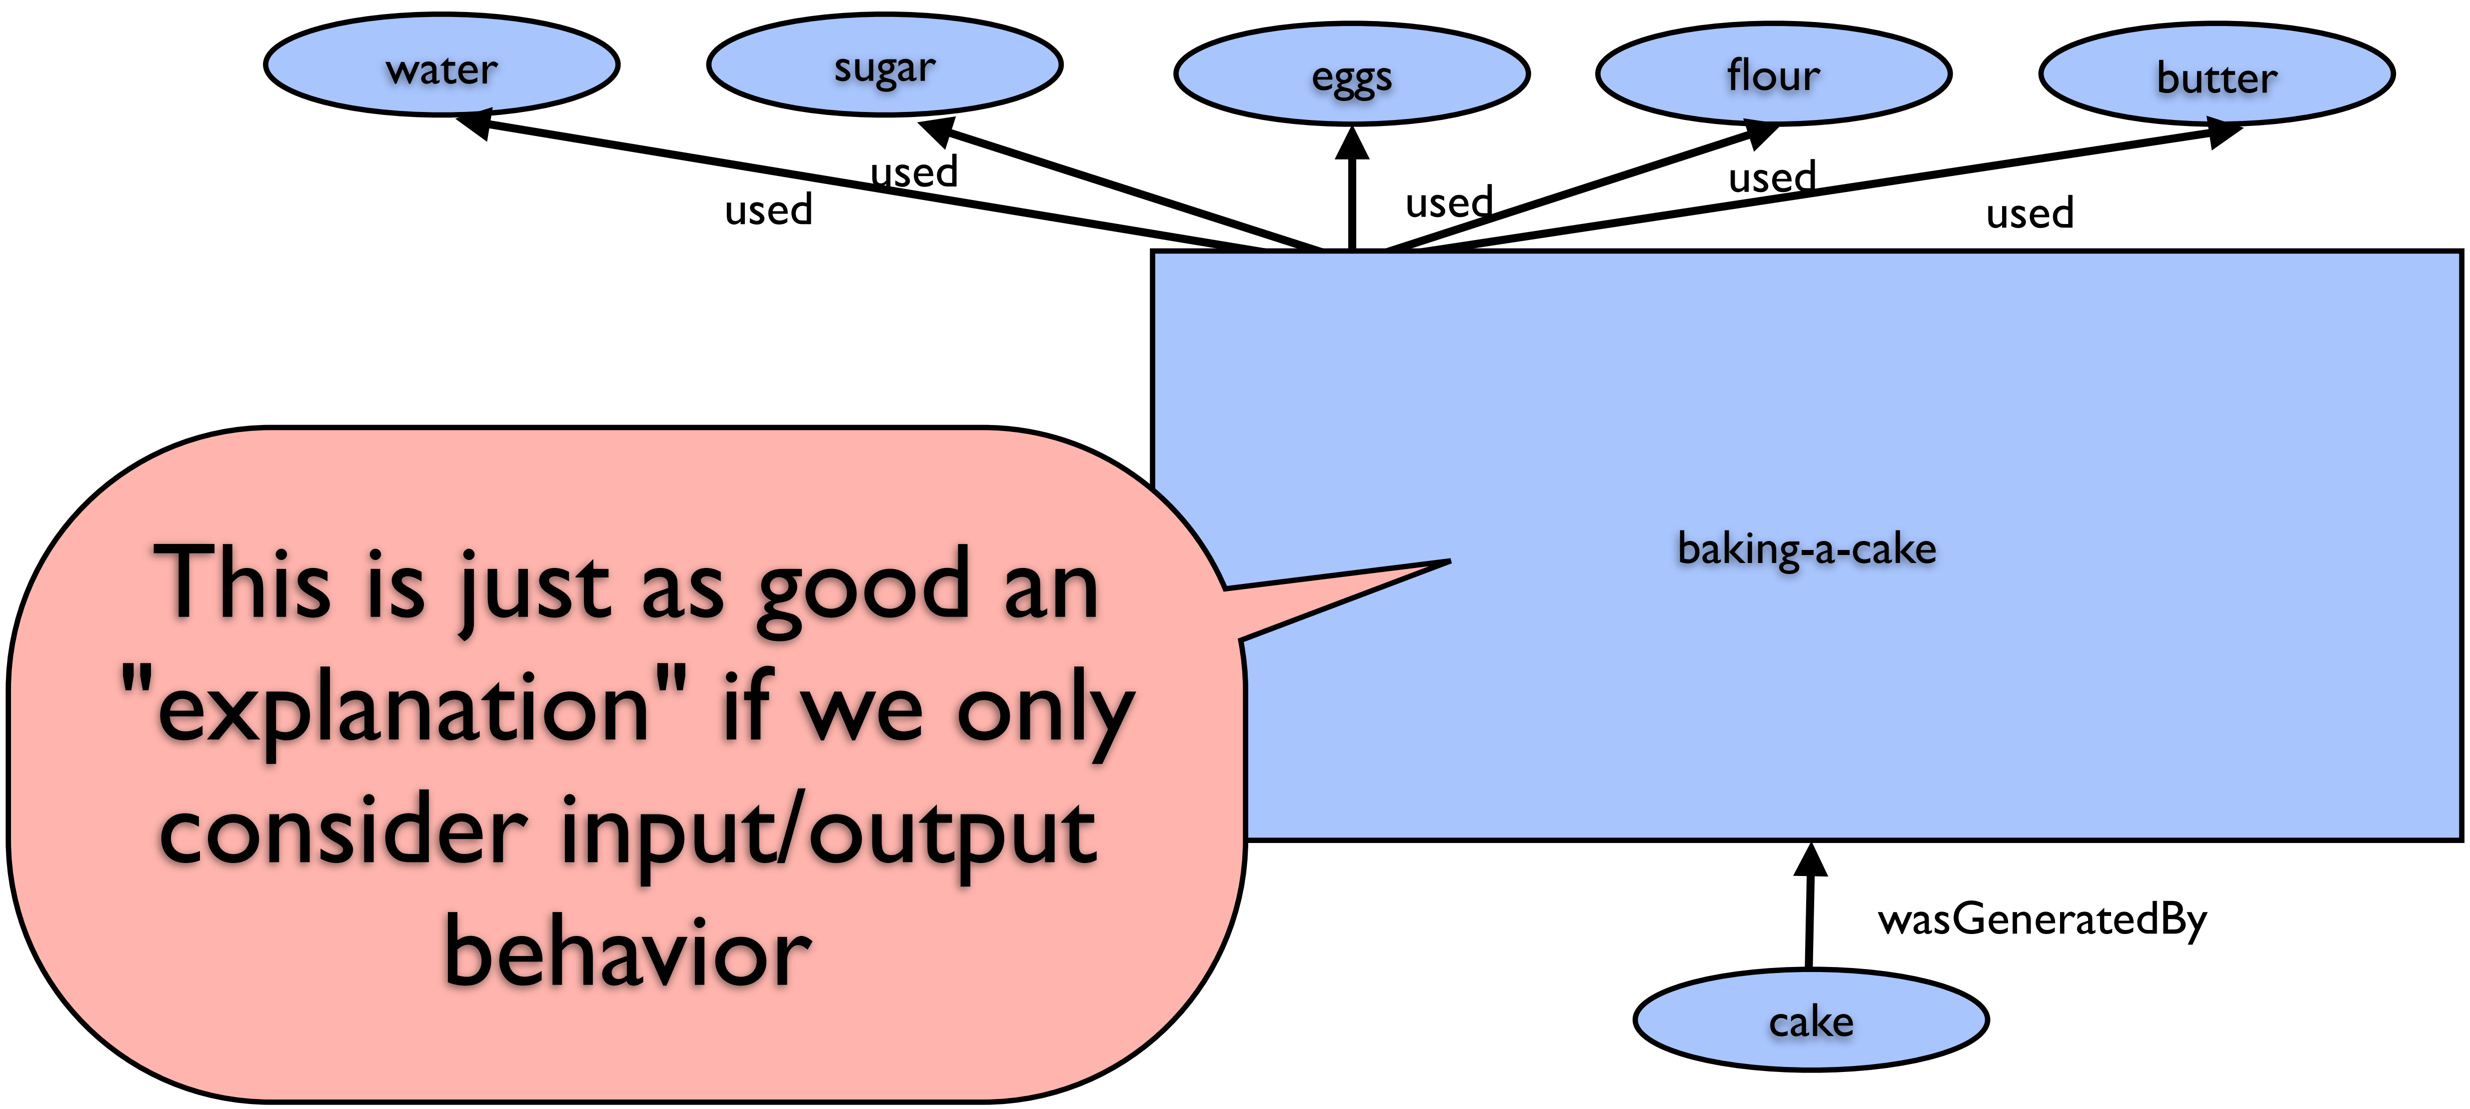
water
sugar
eggs
flour
butter
used
used
used
used
used
baking-a-cake
wasGeneratedBy
cake
This is just as good an "explanation" if we only consider input/output behavior

Baking a cake: a simple causal model

$$\begin{aligned} Mix &:= (Water \wedge Sugar \wedge Eggs \wedge Flour \wedge Butter) \oplus U_1 \\ Batter &:= Mix \oplus U_2 \\ Bake &:= (Batter \wedge Pan) \oplus U_3 \\ Cake &:= Bake \oplus U_4 \end{aligned}$$

Baking a cake: a simple causal model

$$
\begin{aligned}
Mix &:= (Water \wedge Sugar \wedge Eggs \wedge Flour \wedge Butter) \oplus U_1 \\
Batter &:= Mix \oplus U_2 \\
Bake &:= (Batter \wedge Pan) \oplus U_3 \\
Cake &:= Bake \oplus U_4
\end{aligned}
$$

Success
requires all ingredients
(conjunction)

Baking a cake: a simple causal model

$$
\begin{aligned}
Mix &:= (Water \wedge Sugar \wedge Eggs \wedge Flour \wedge Butter) \oplus U_1 \\
Batter &:= Mix \oplus U_2 \\
Bake &:= (Batter \wedge Pan) \oplus U_3 \\
Cake &:= Bake \oplus U_4
\end{aligned}
$$

U_1 ... U_4 represent unmodeled external factors

Causal models

- $M = (U,V,F)$
  - U is set of endogenous variables (nonmeasurable, external factors)
  - V is set of exogenous variables (explicitly modeled/ measurable things)
  - F is family of transfer functions F_X, one for each X in V
- In AI, often use probabilistic interpretation, here we just consider discrete behavior

Causal model, graphically

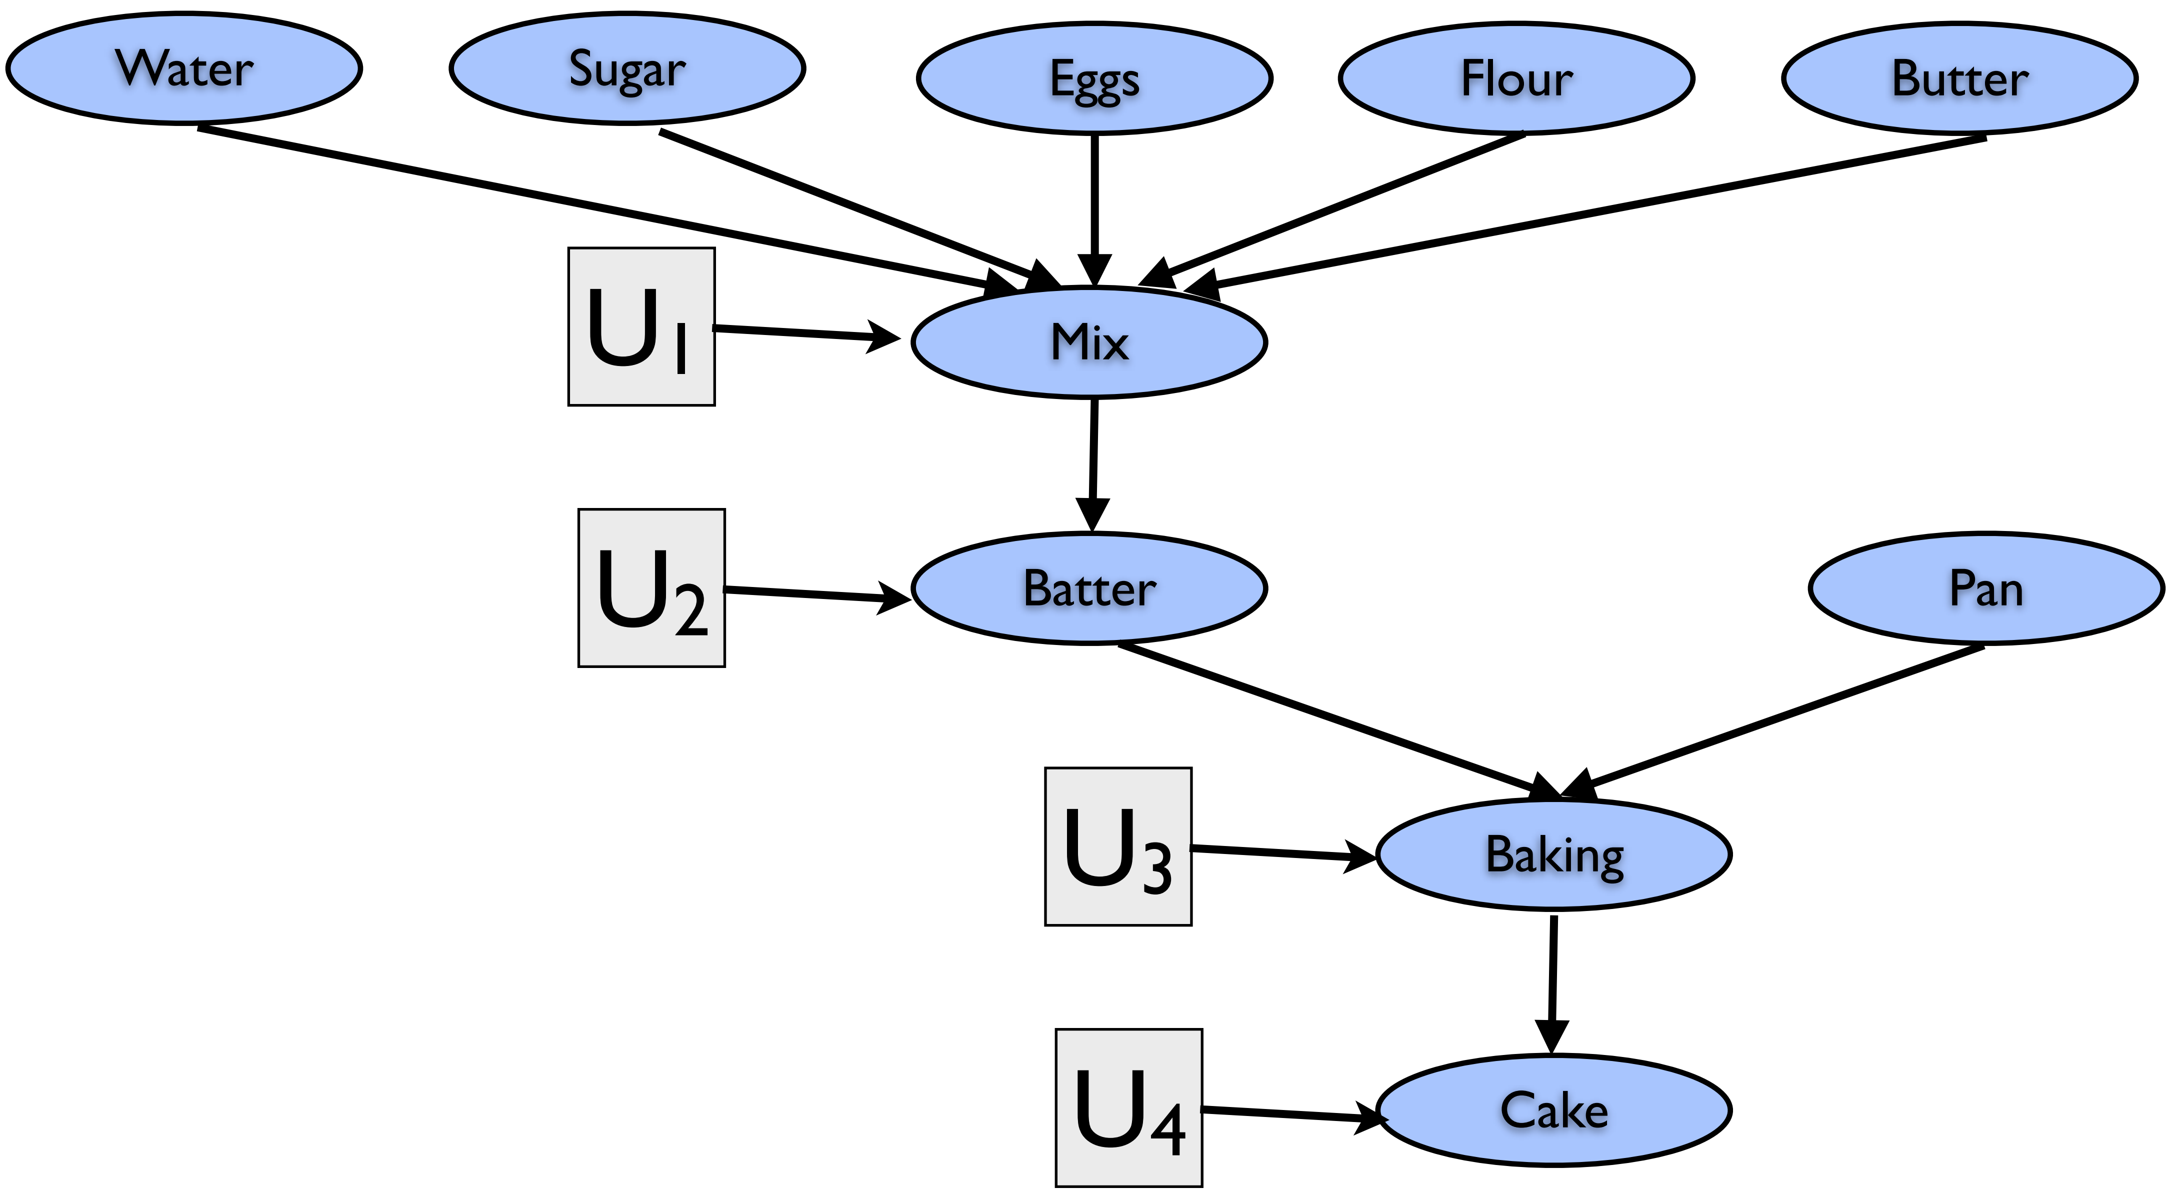

Causal situation

- (M,σ)
  - A causal model M
  - values $\sigma(X)$ for the variables
- Describes "what actually happened"
- Needed to talk about "actual causes"

Causal model, graphically

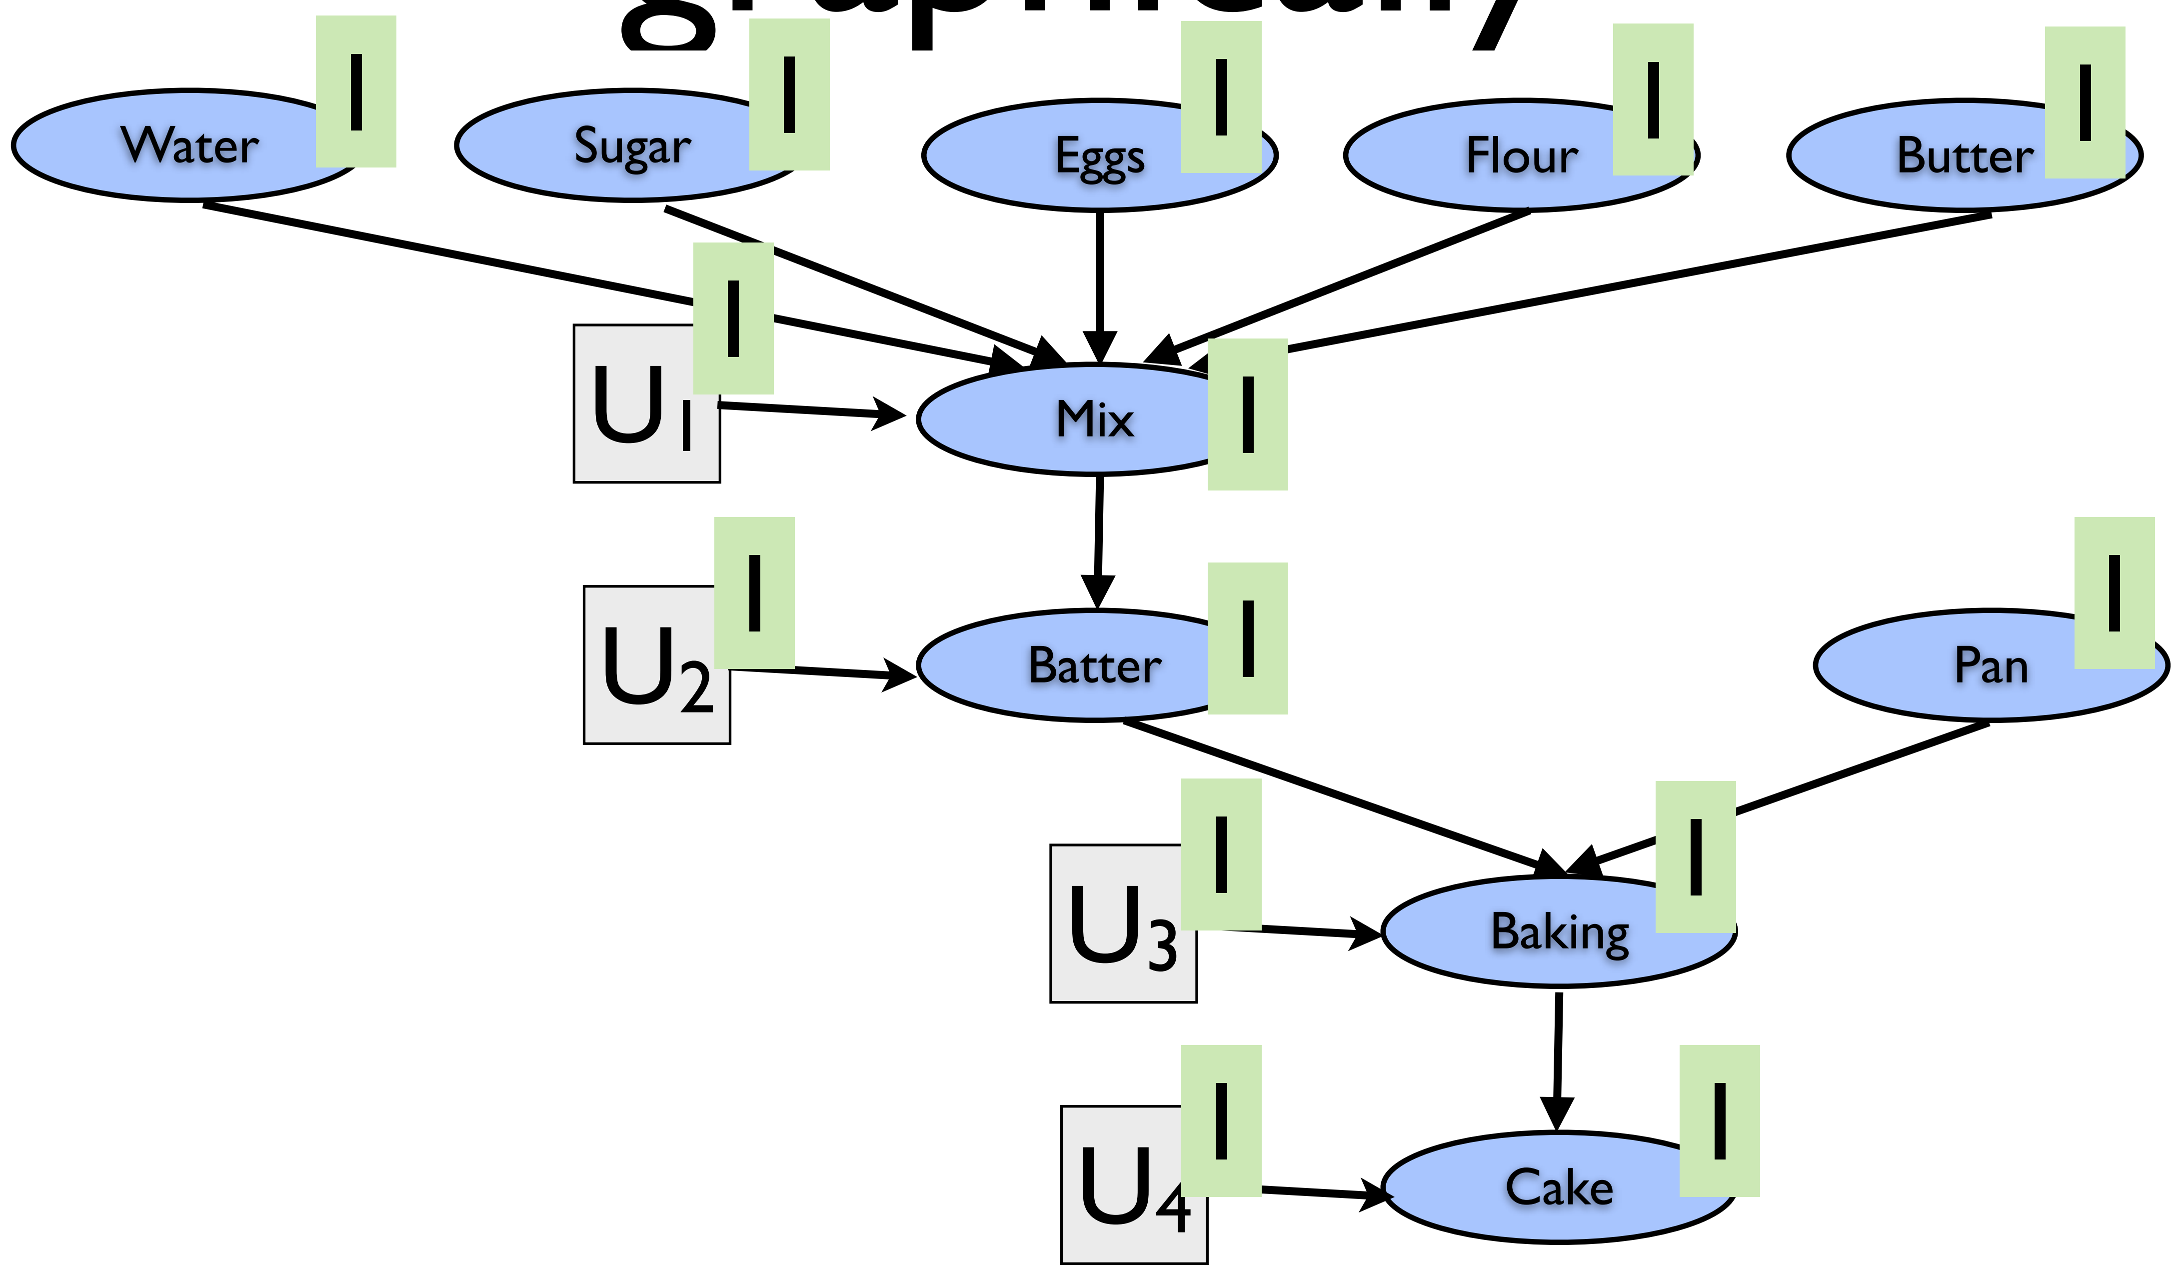

Interventions

- Causal models allow **interventions**
  - considering ramifications of hypothetical / counterfactual possibilities at any stage
- Formally, $M_{[X:=x]}$ (or just M_x) is "*M* with *X* set to *x*"
  - **re-evaluate** anything that depends on *X*
  - **disconnect** anything that feeds into *X*

Intervention, graphically

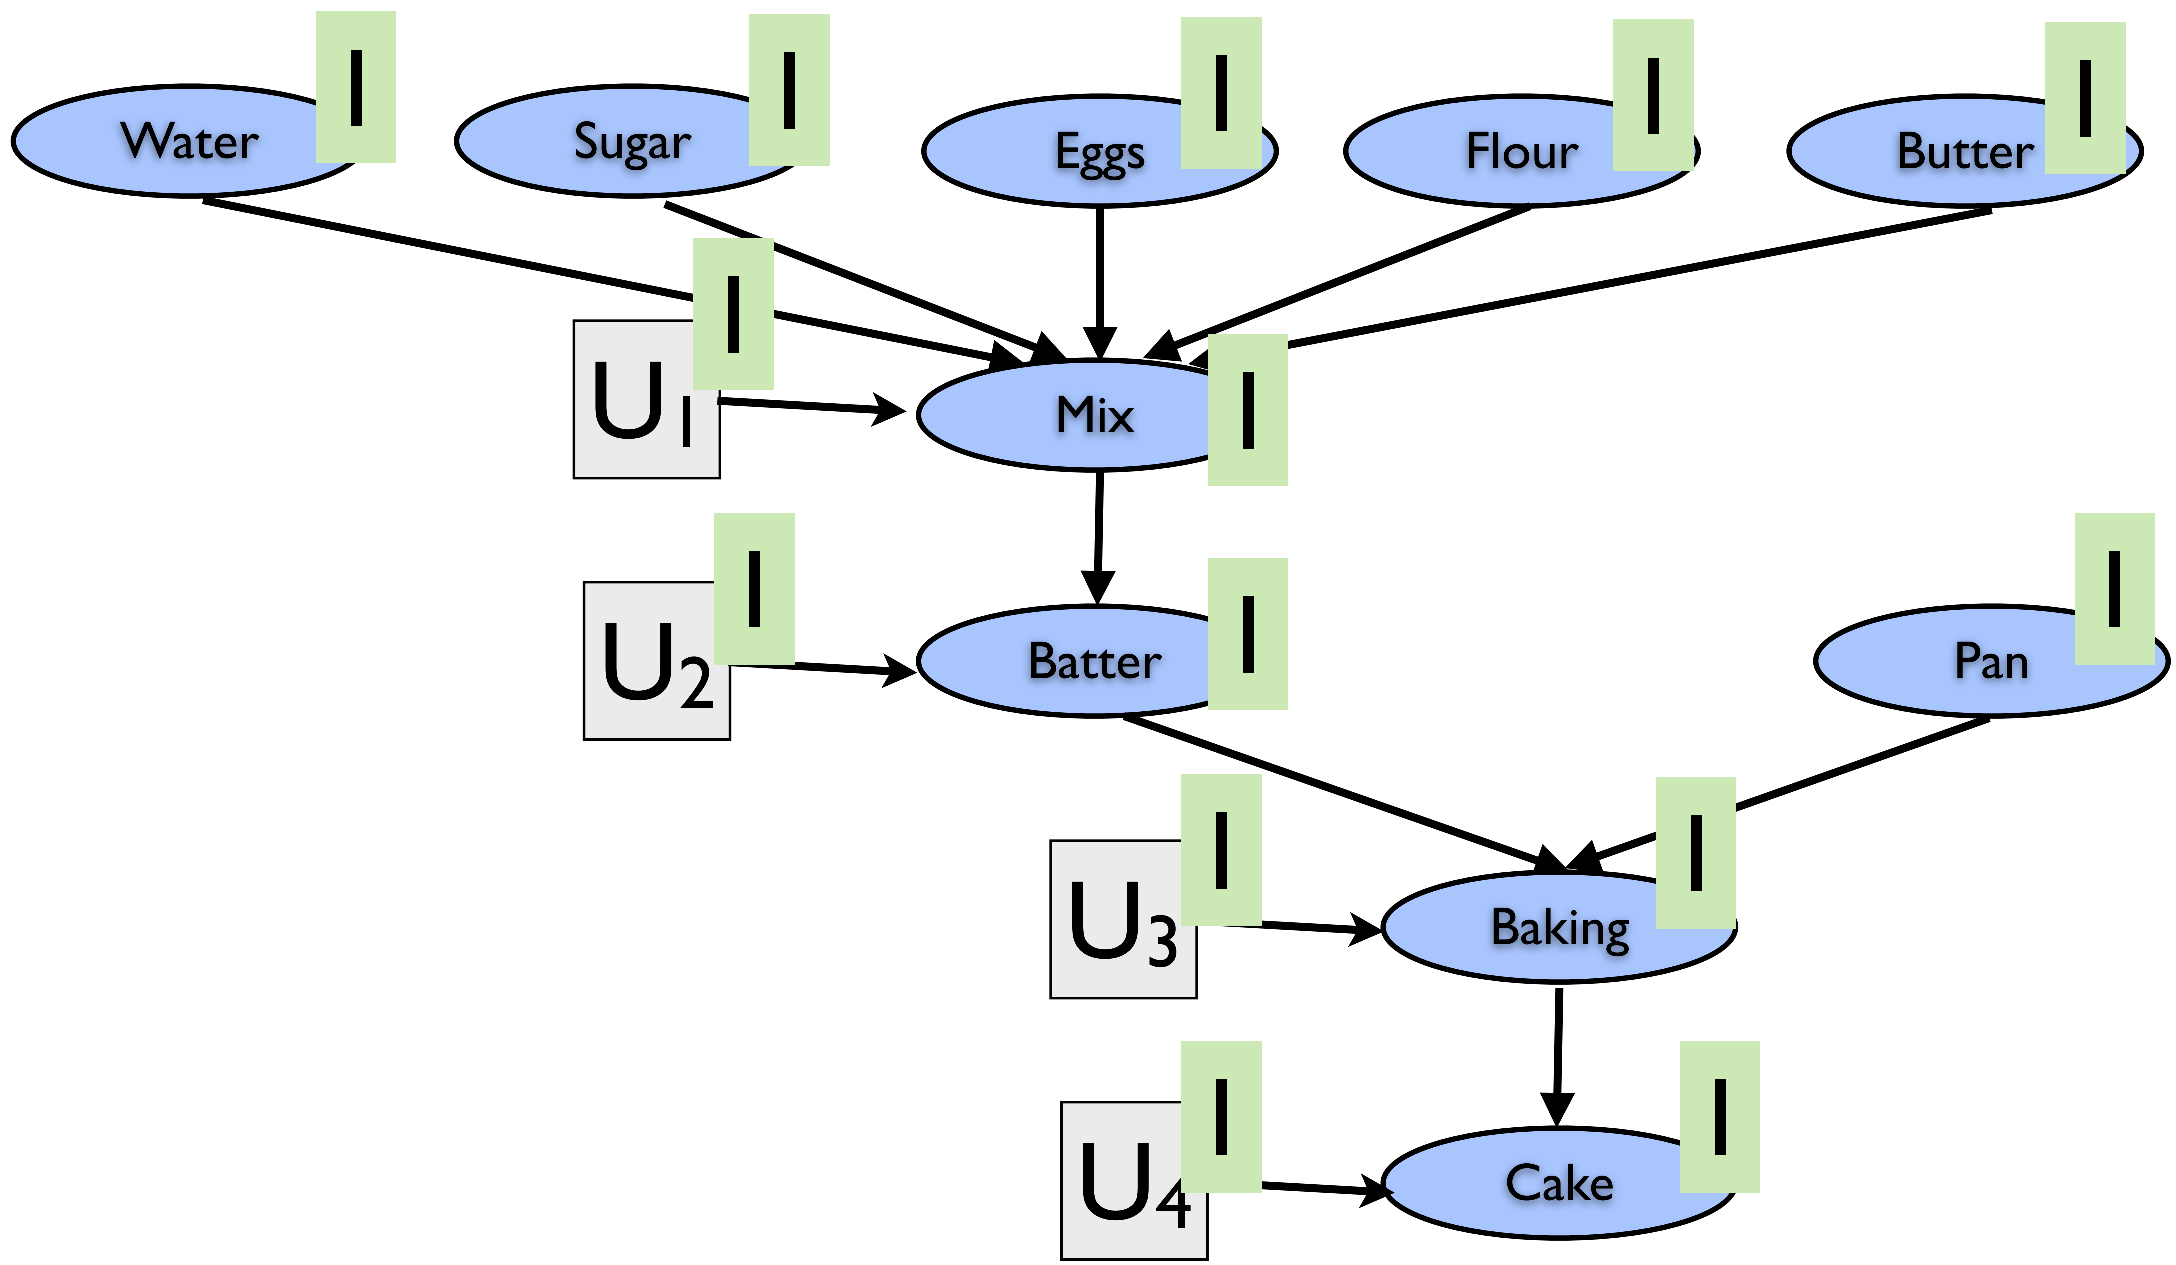

Intervention, graphically

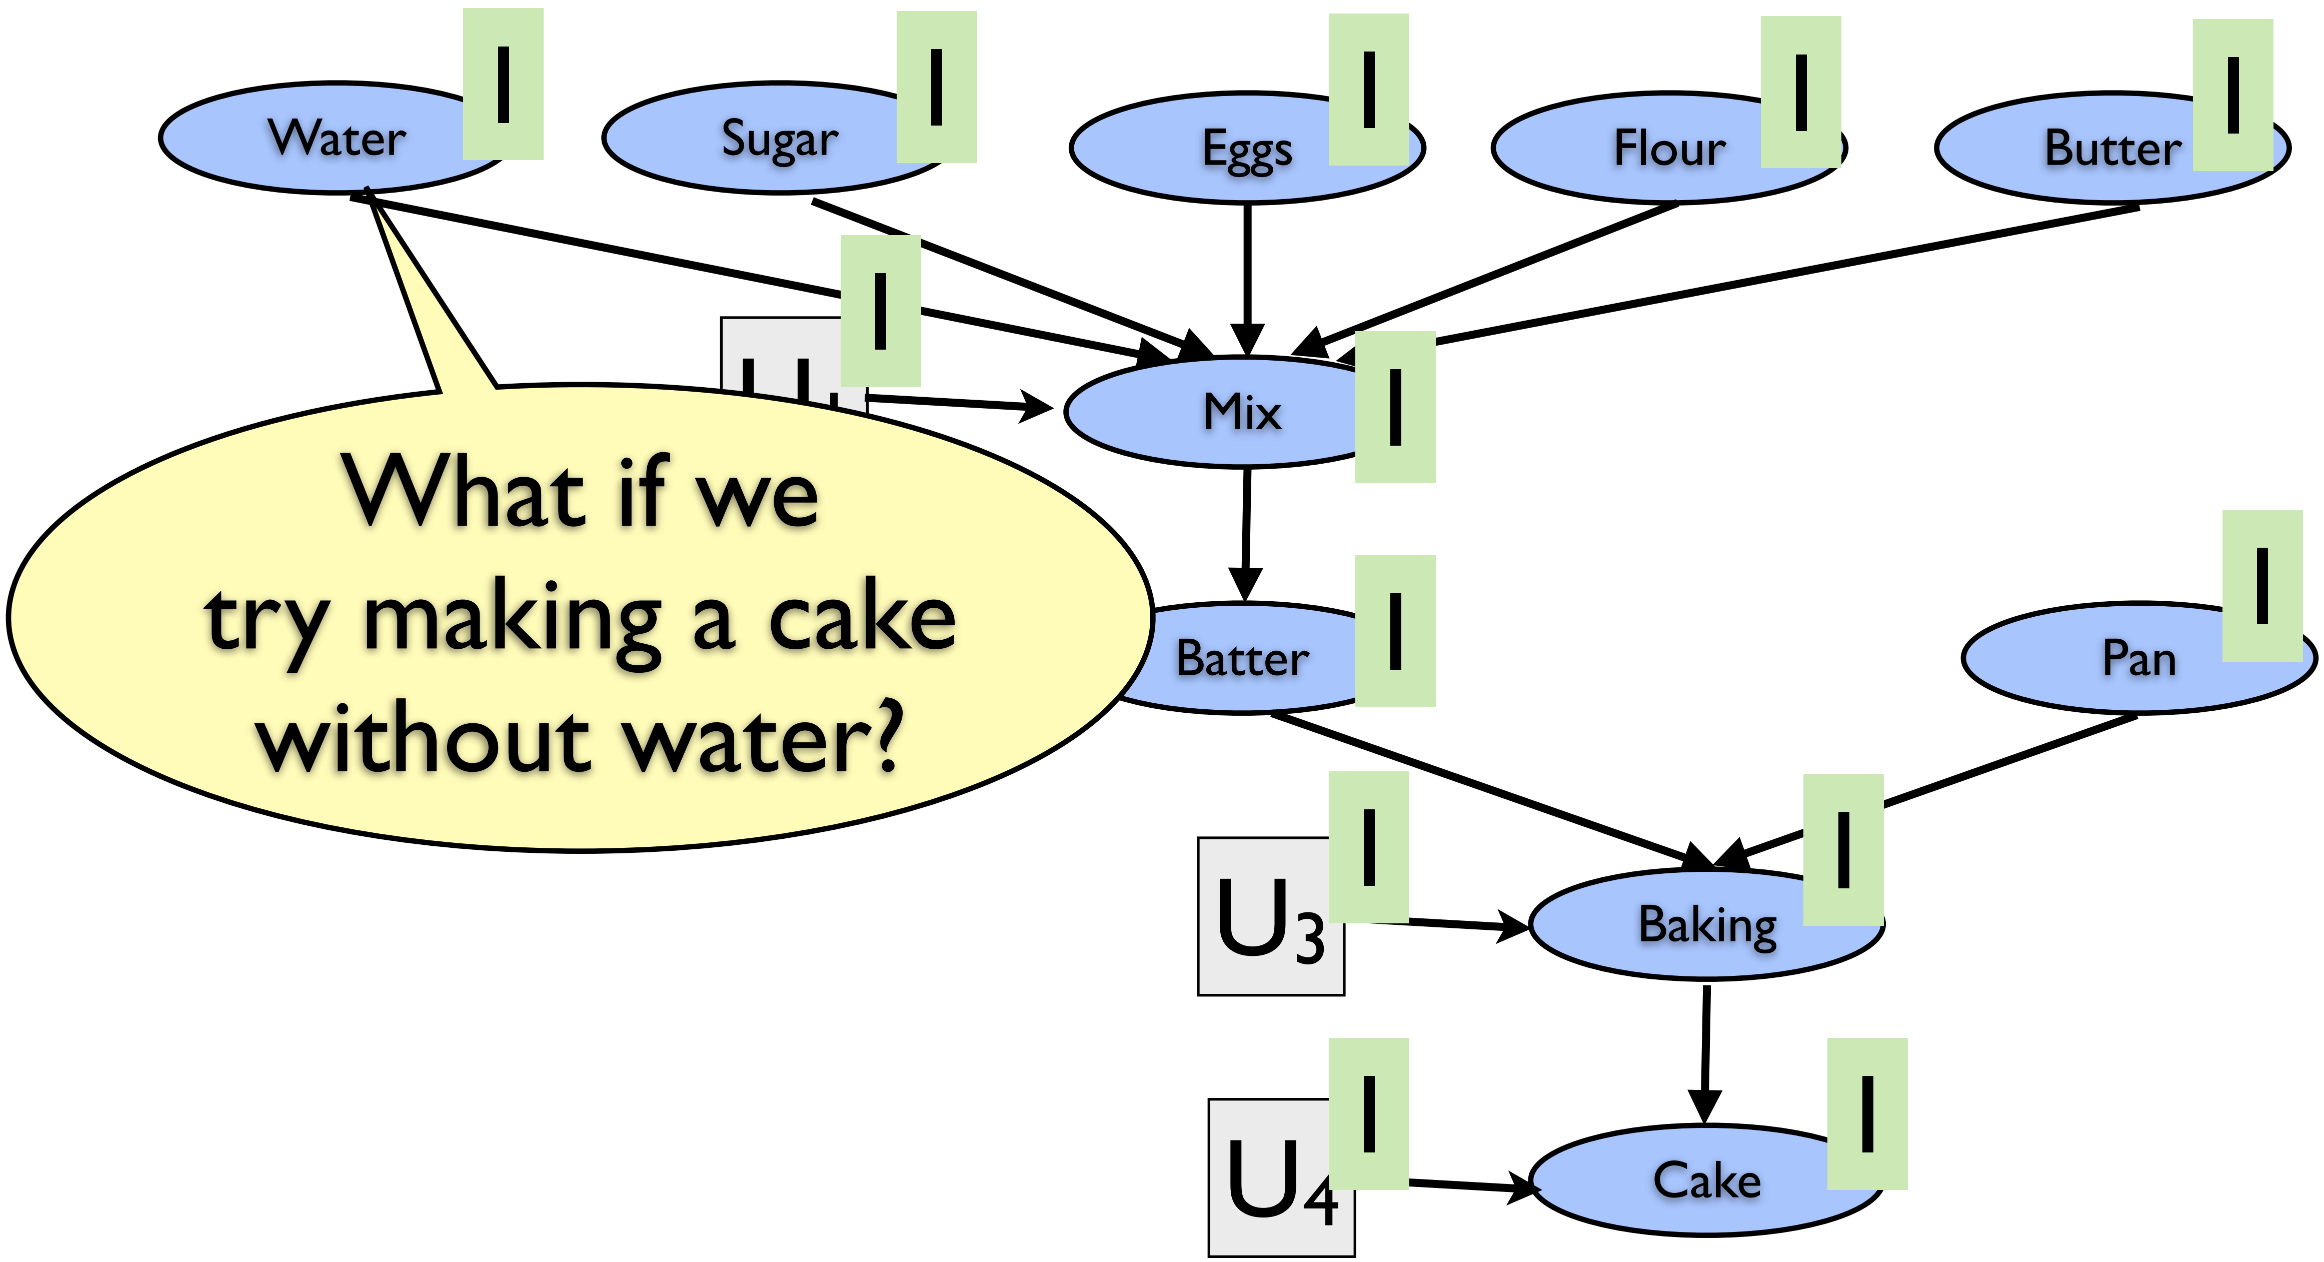

Intervention, graphically

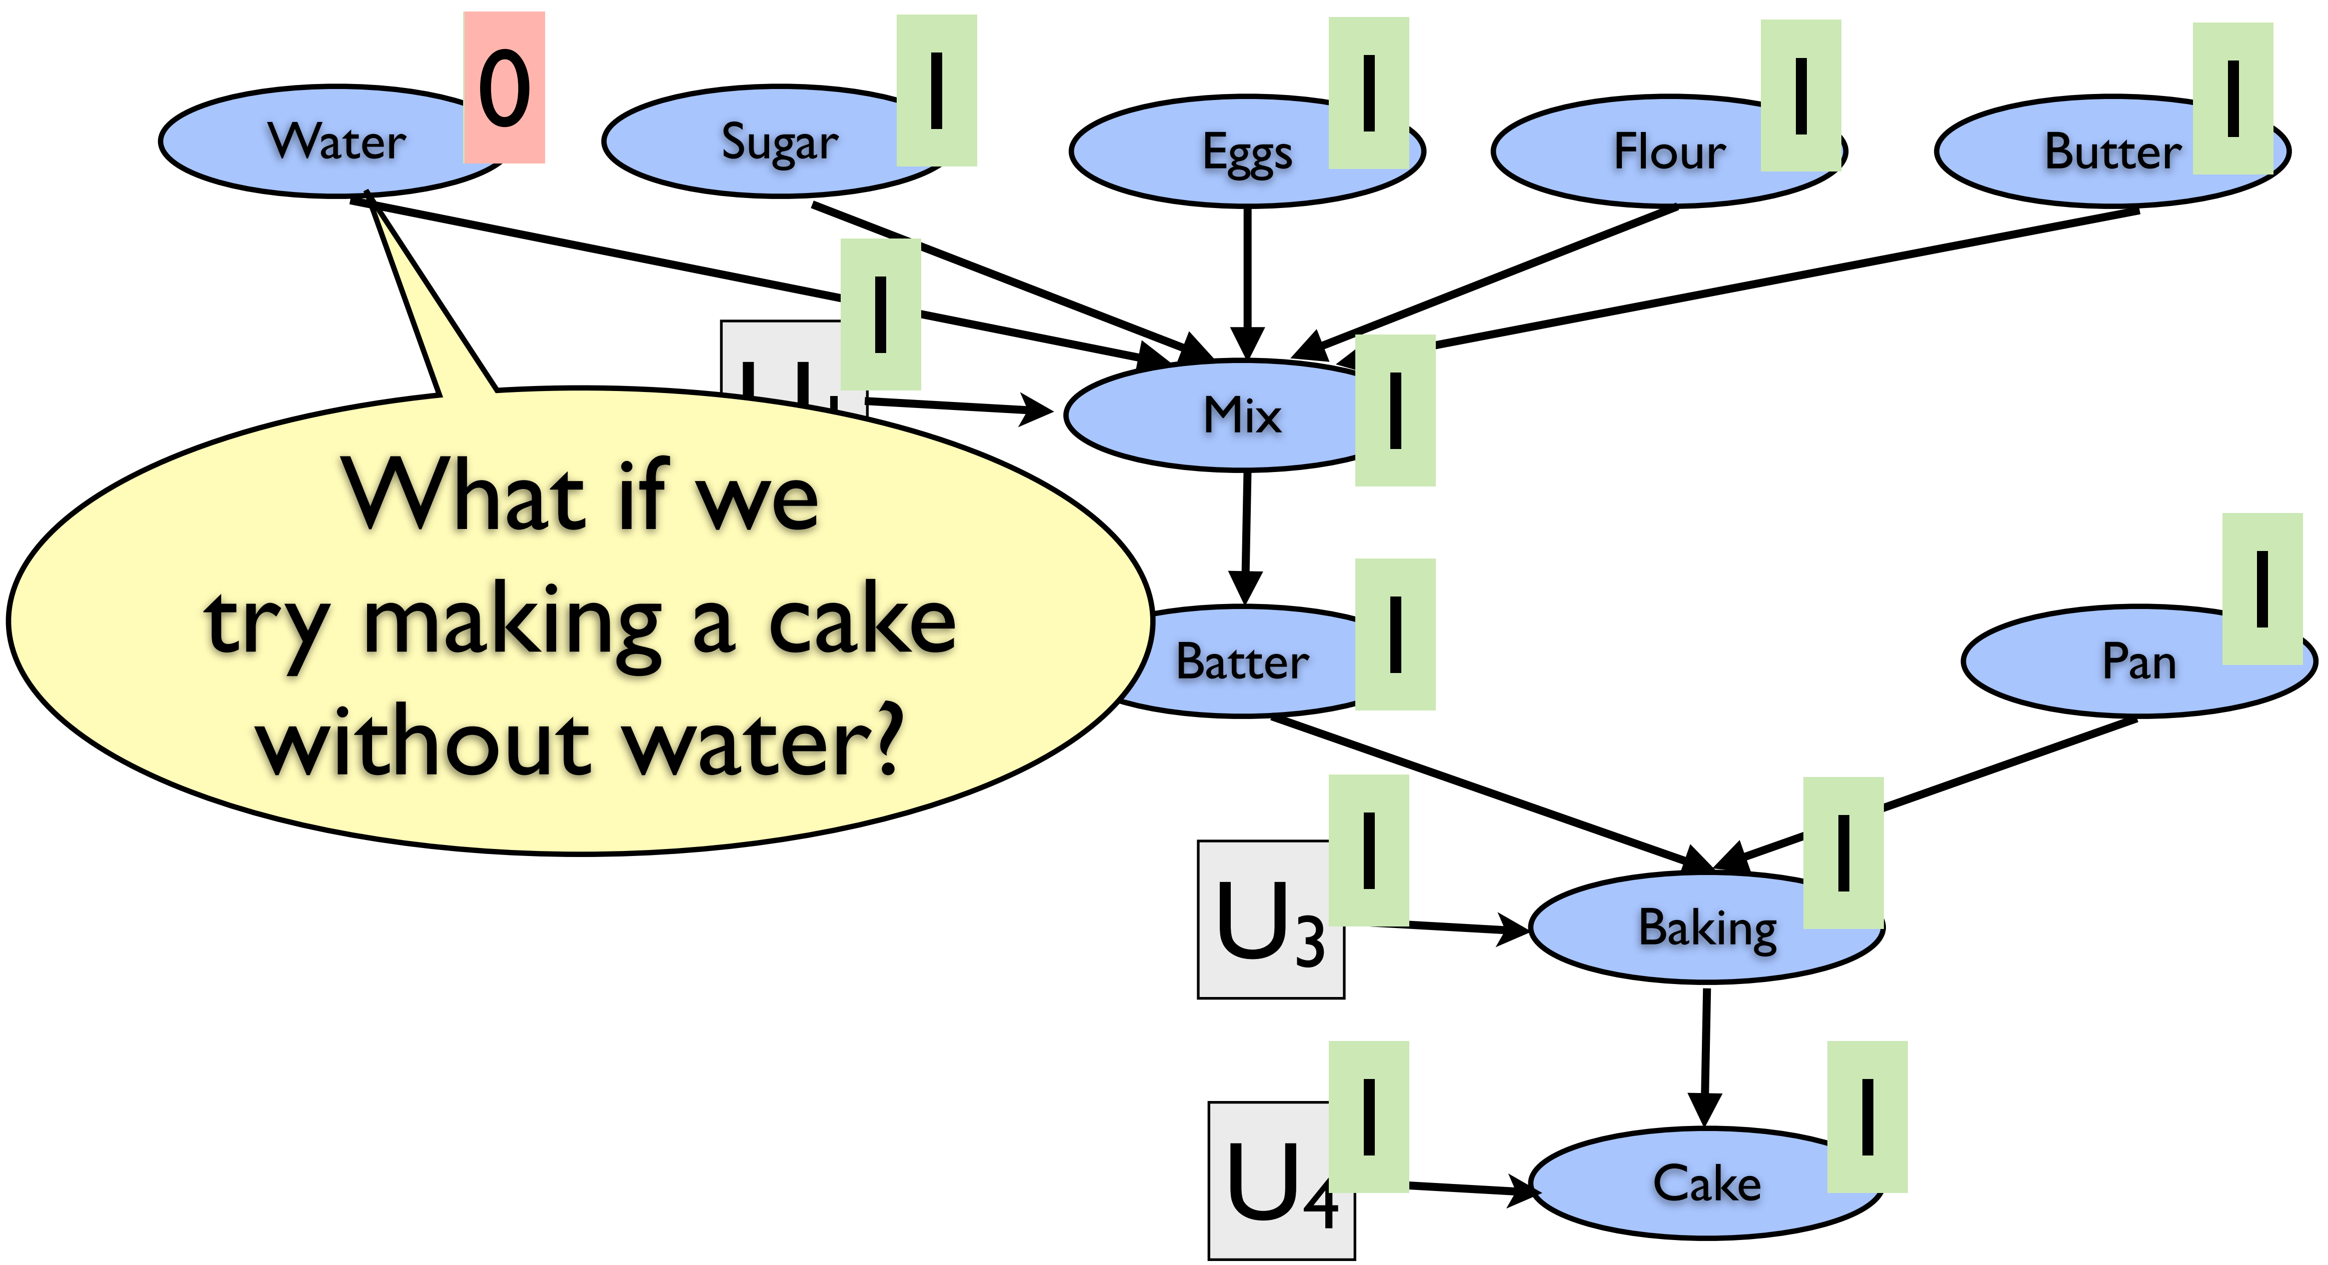

Intervention, graphically

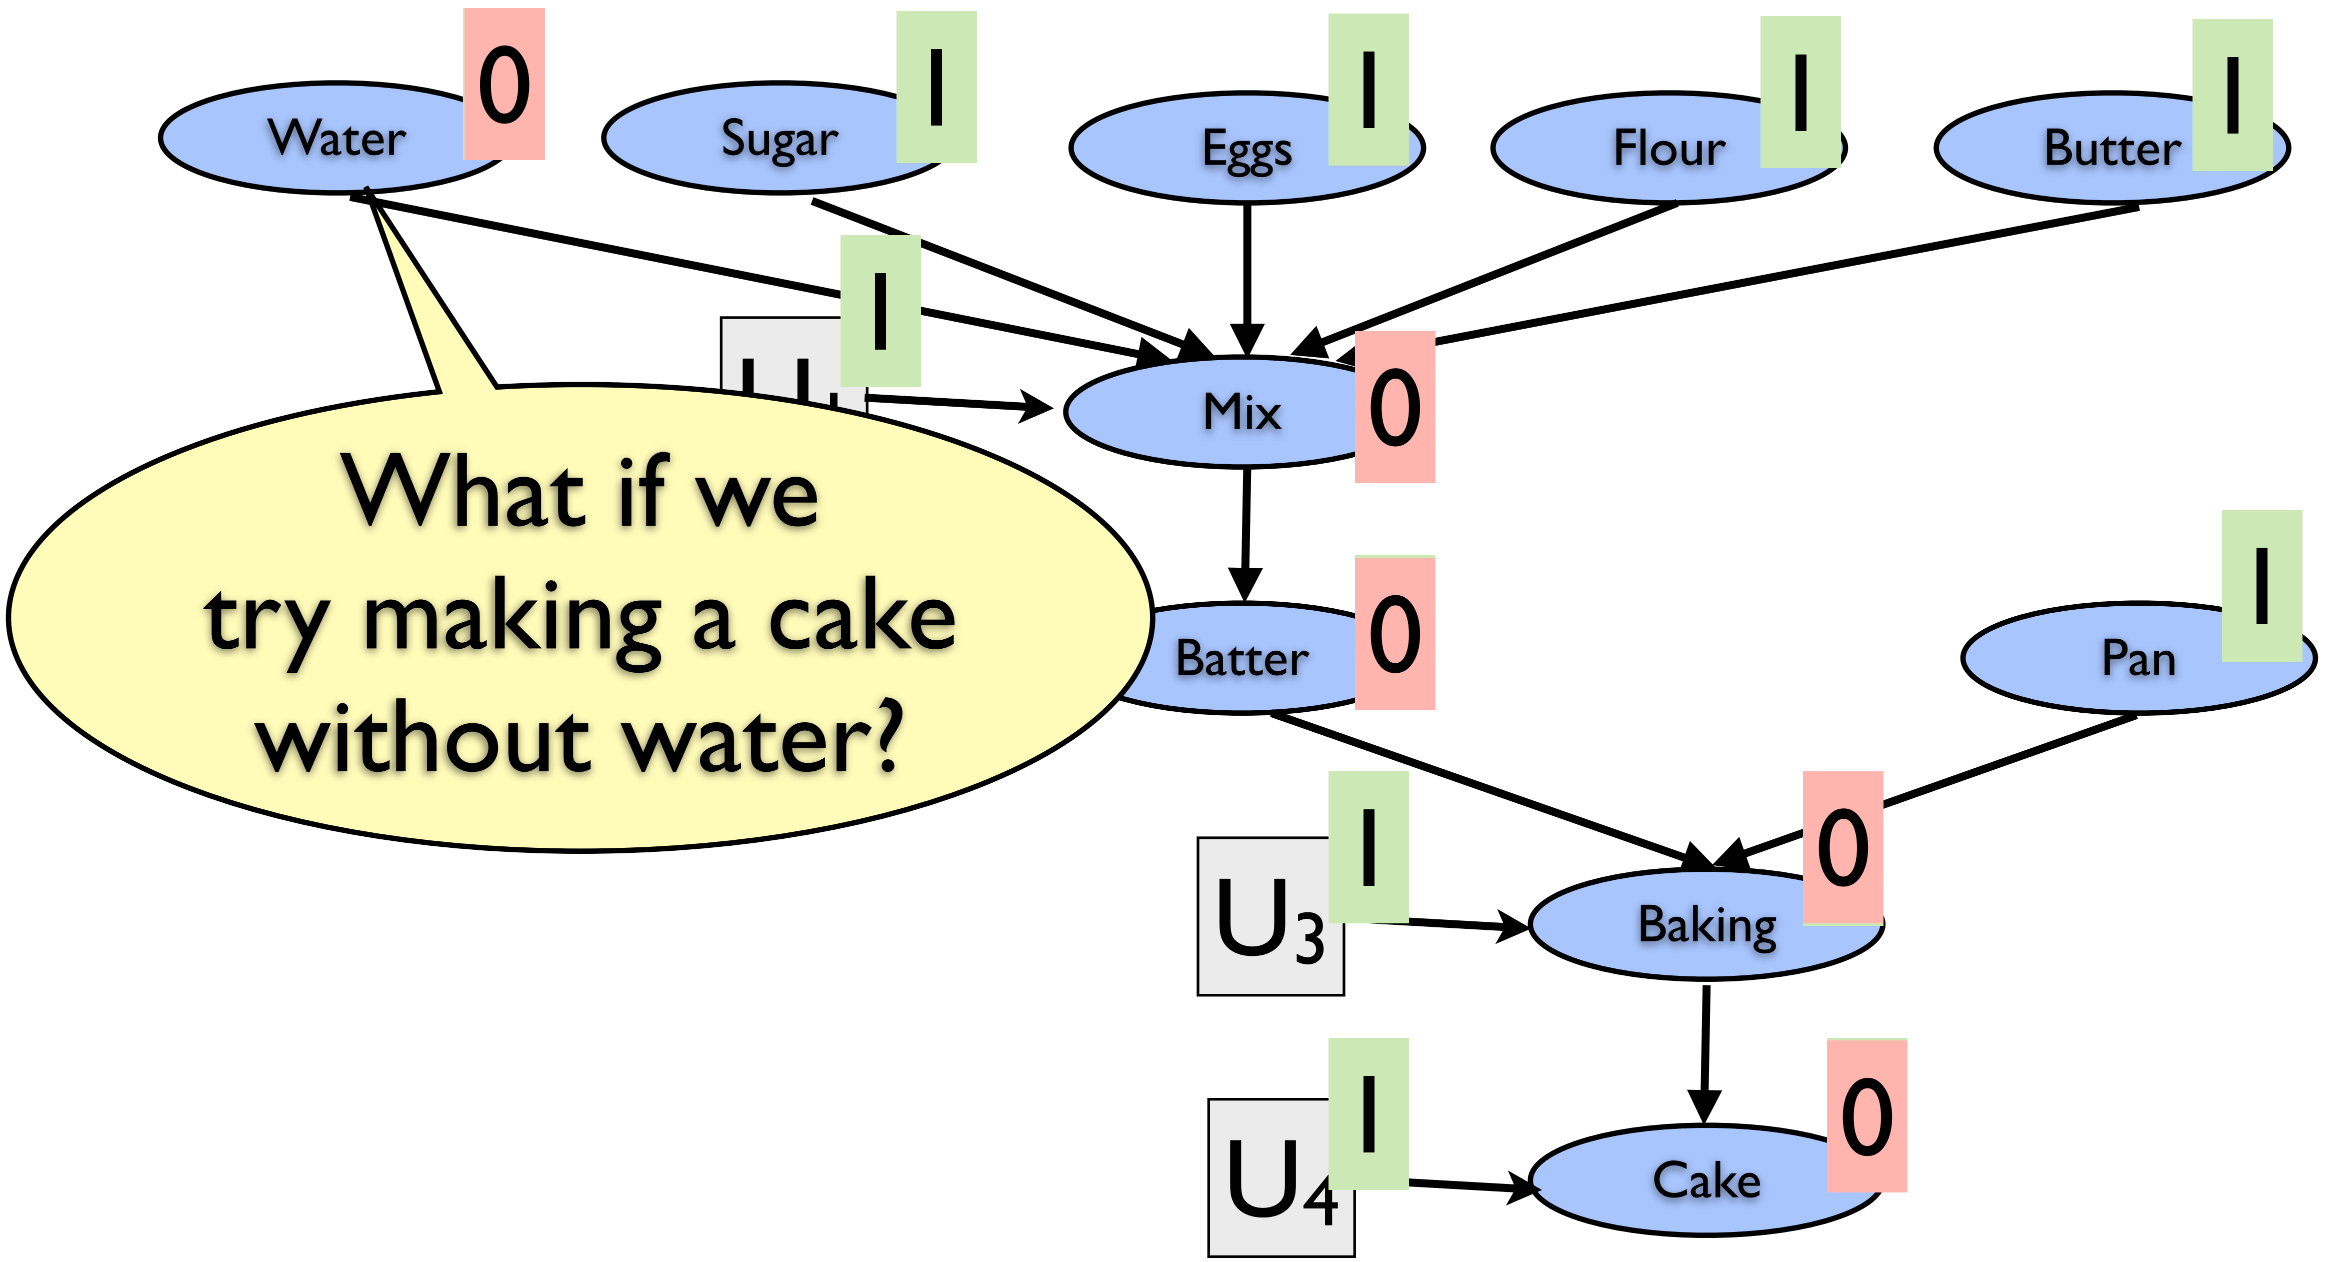

Intervention, graphically

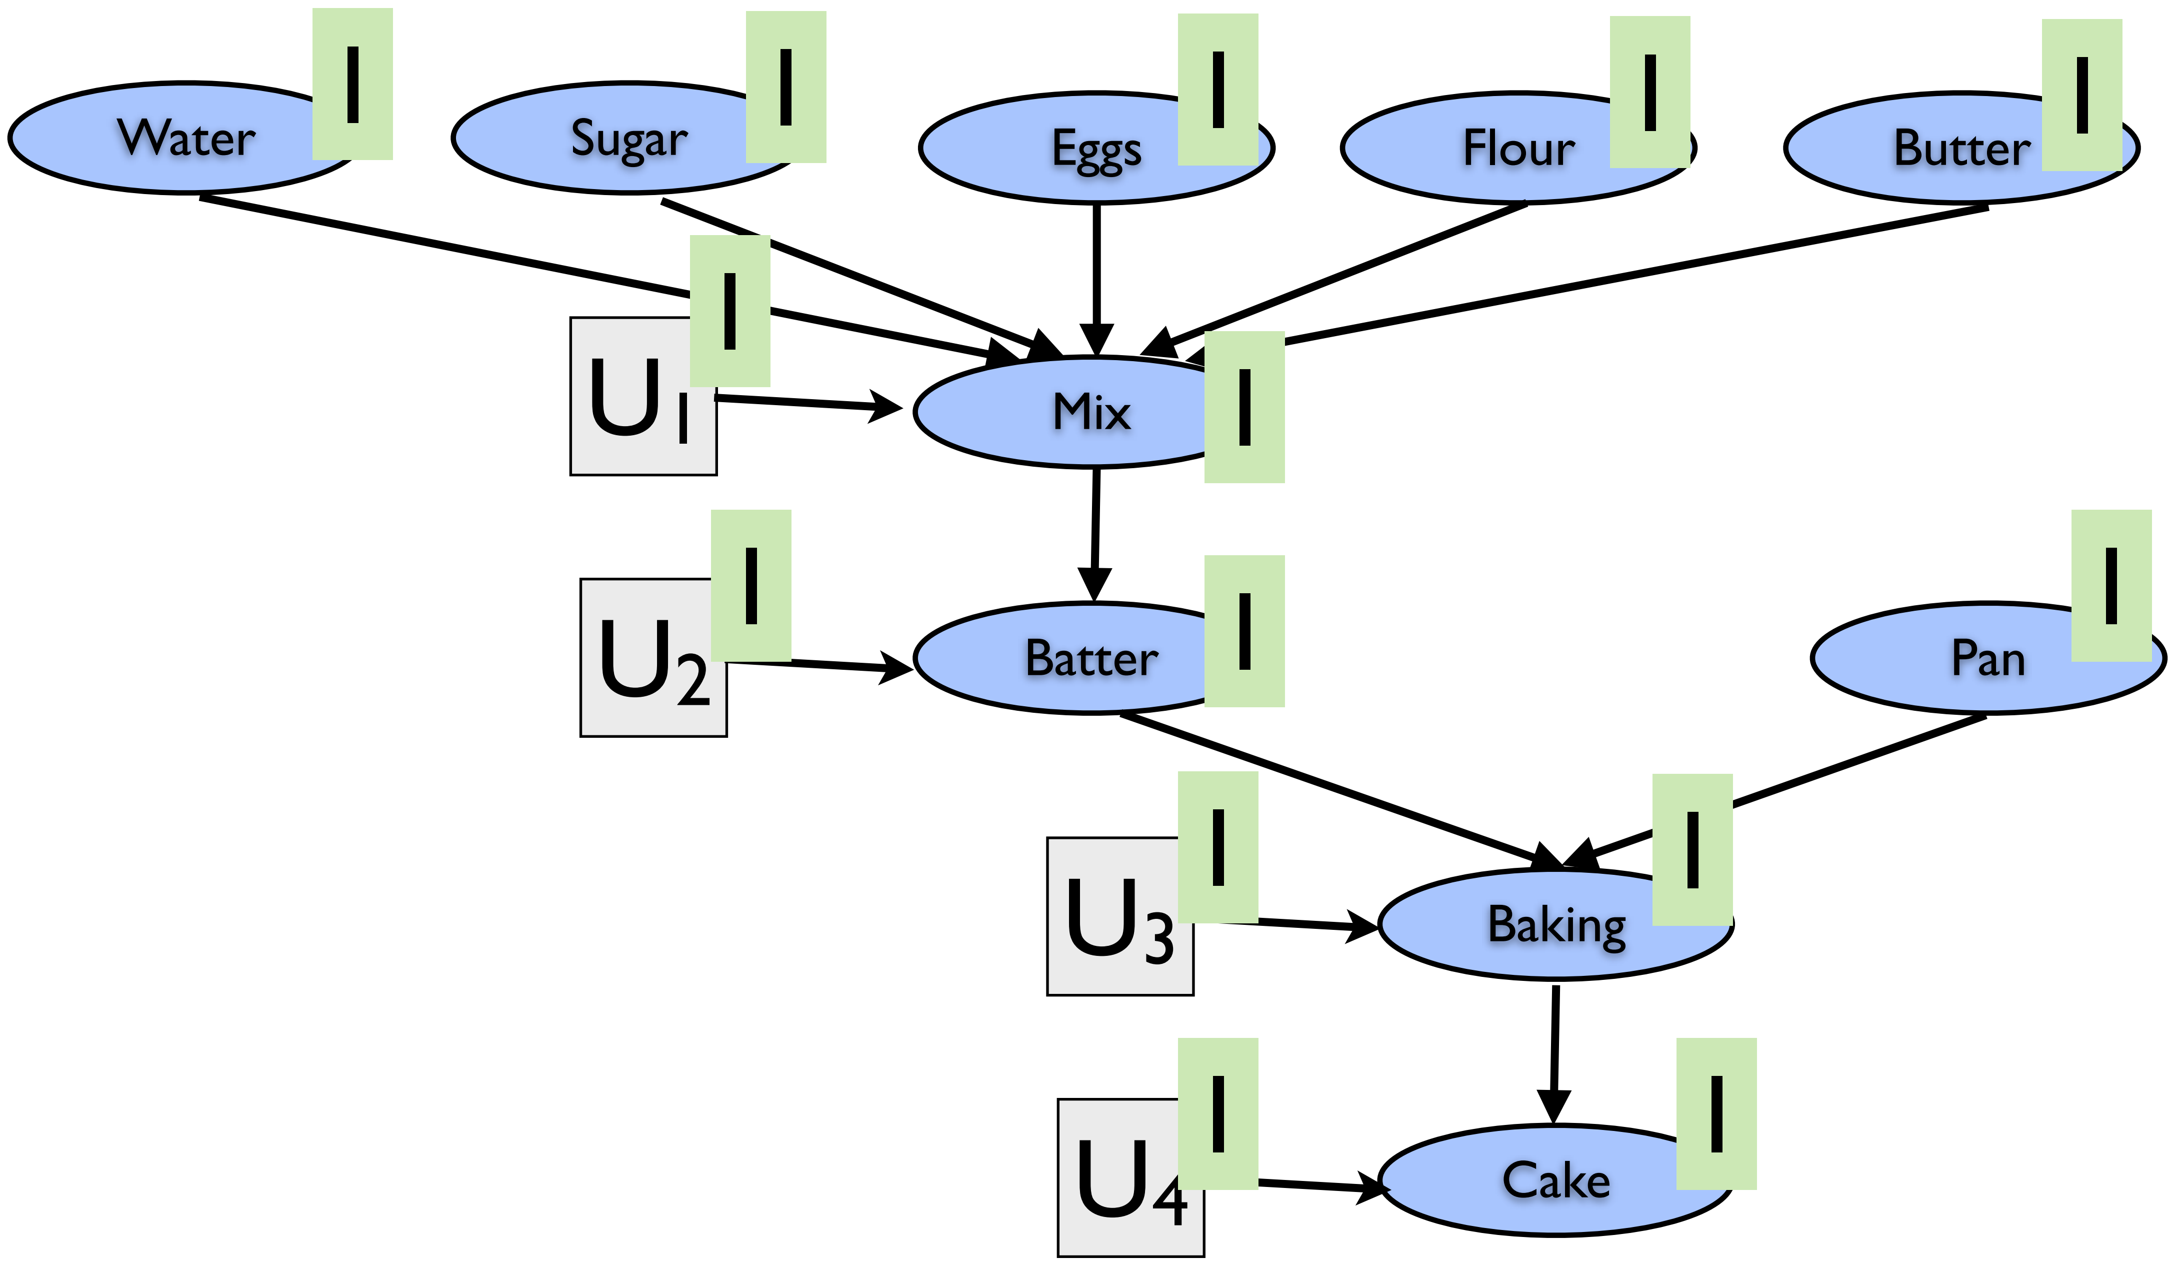

Intervention, graphically

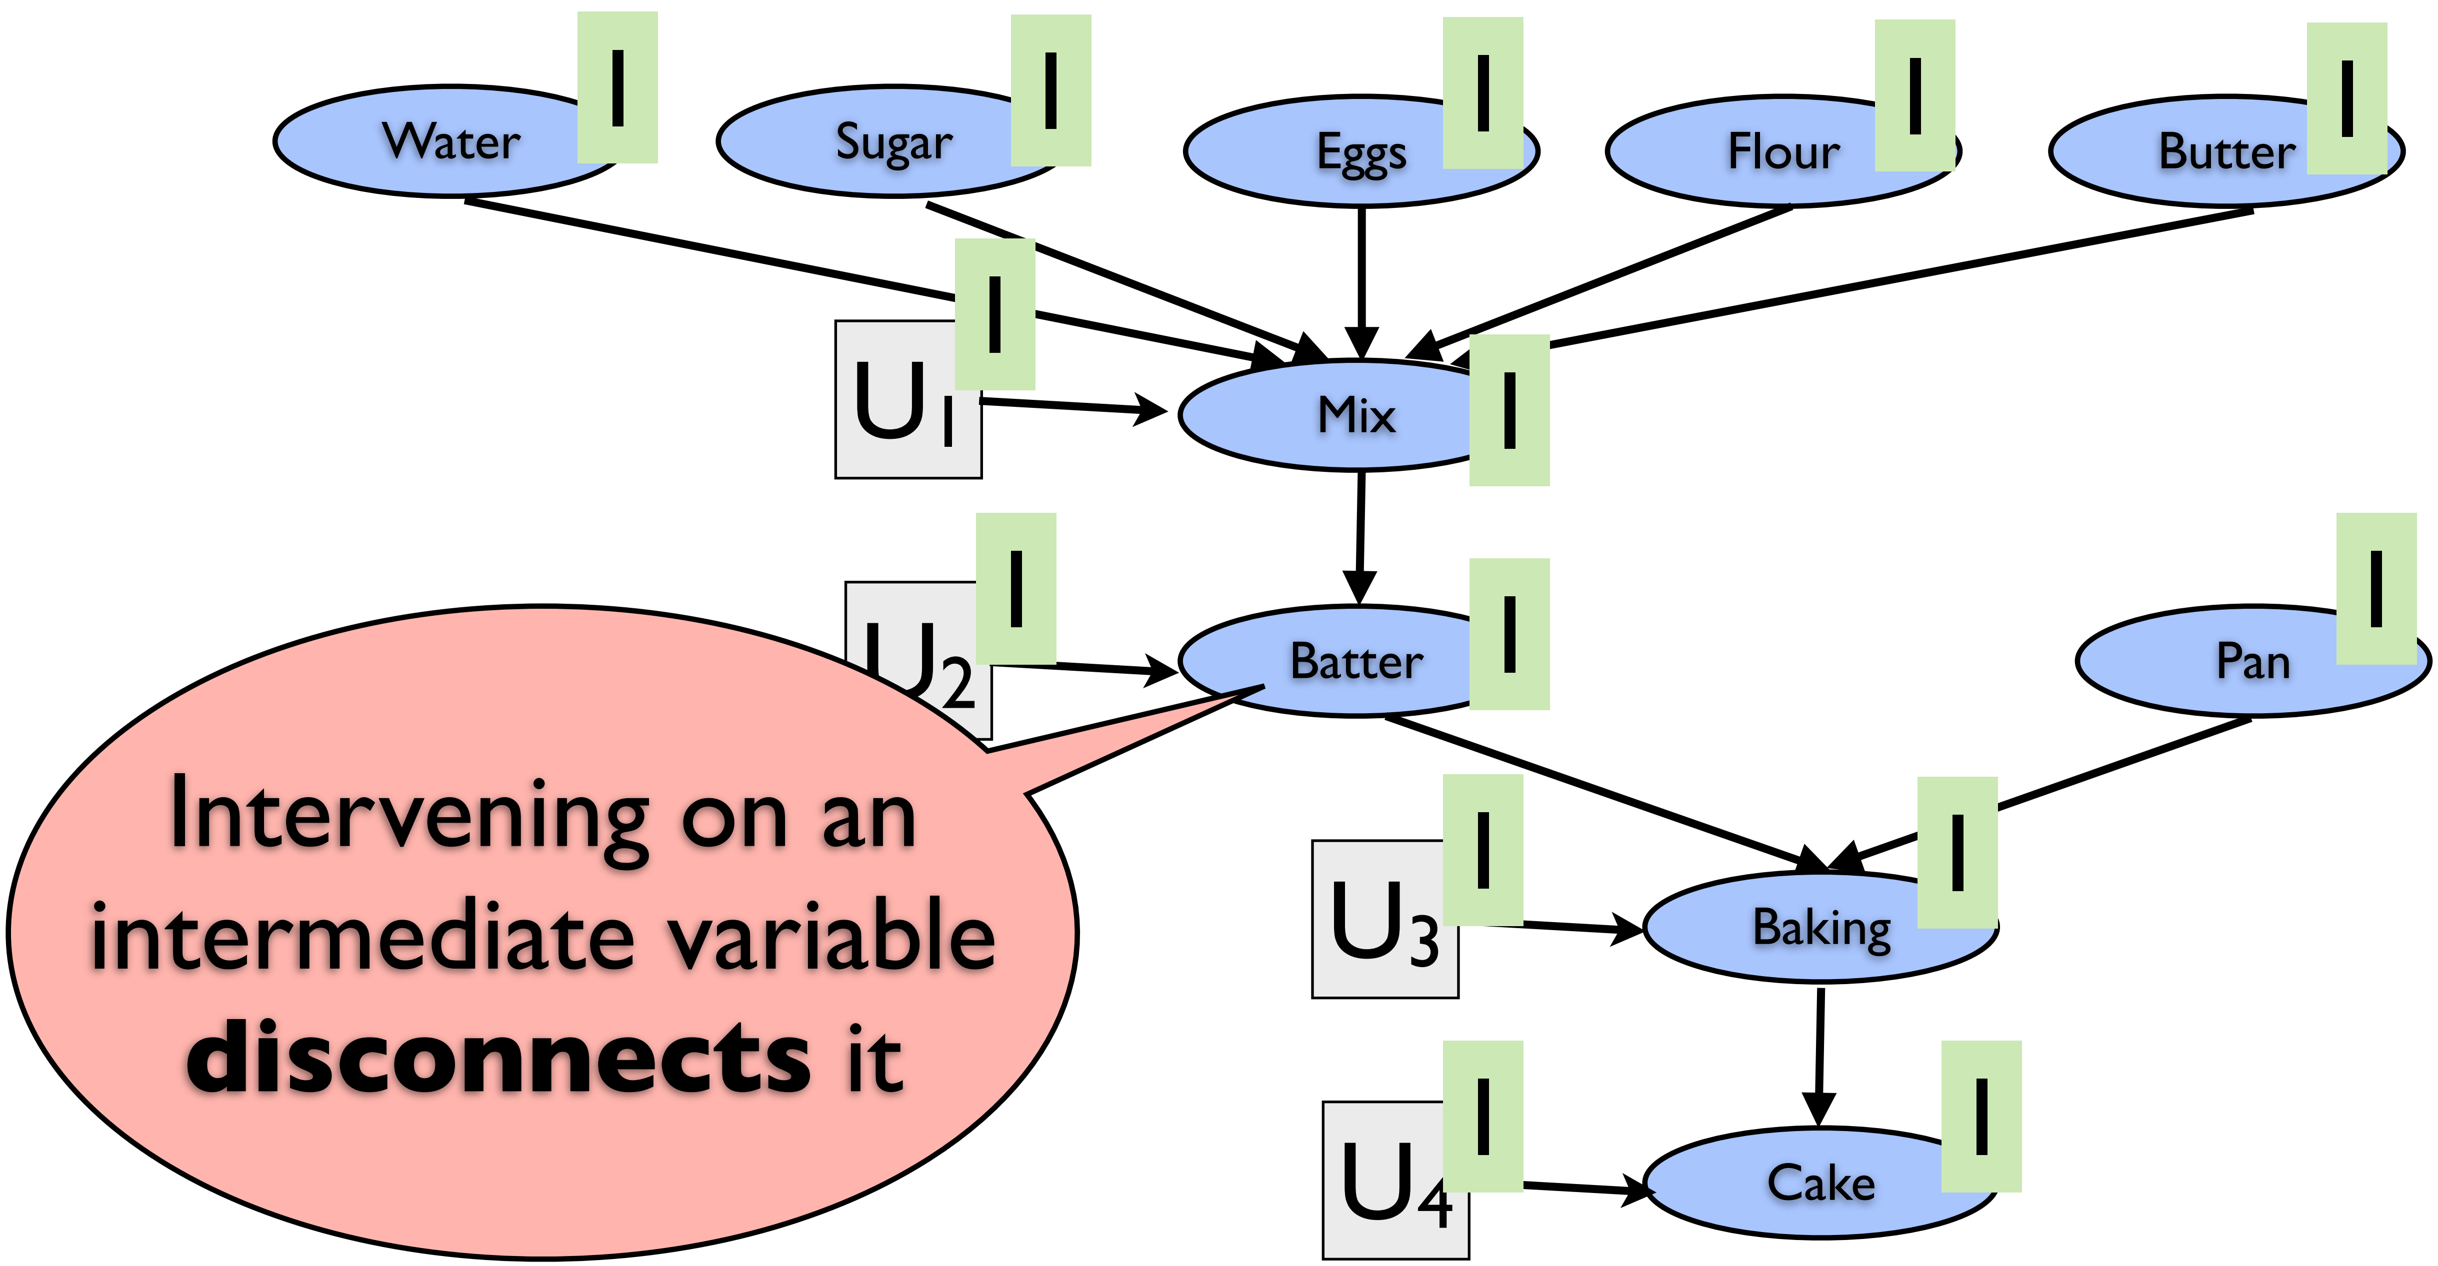

Intervention, graphically

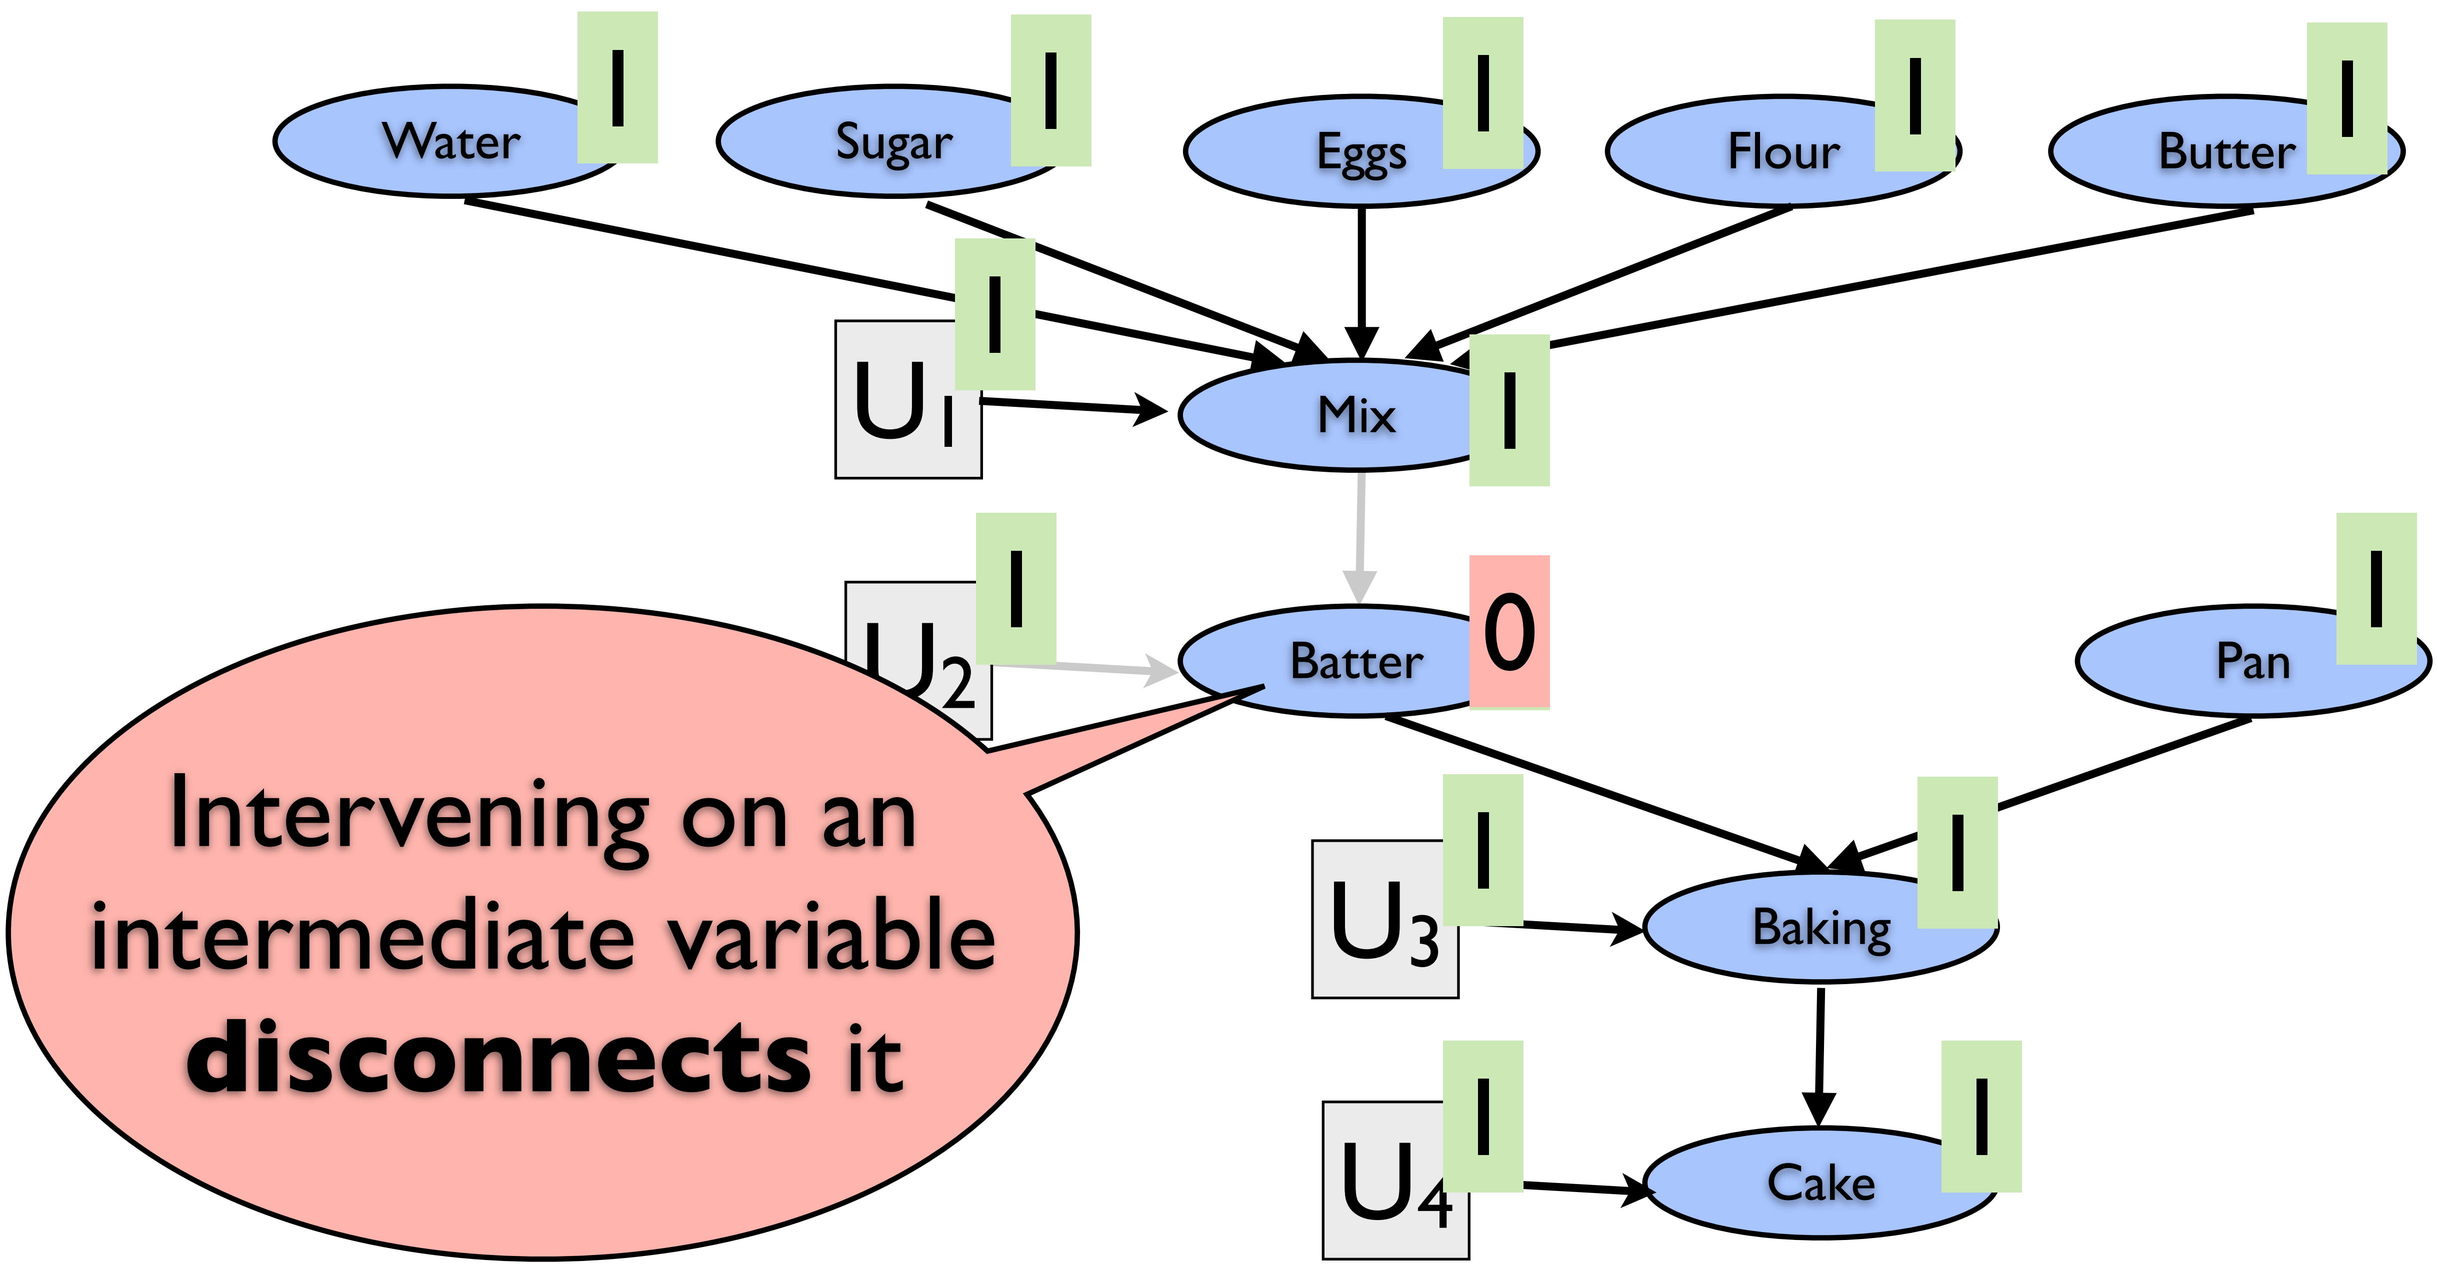

Intervention, graphically

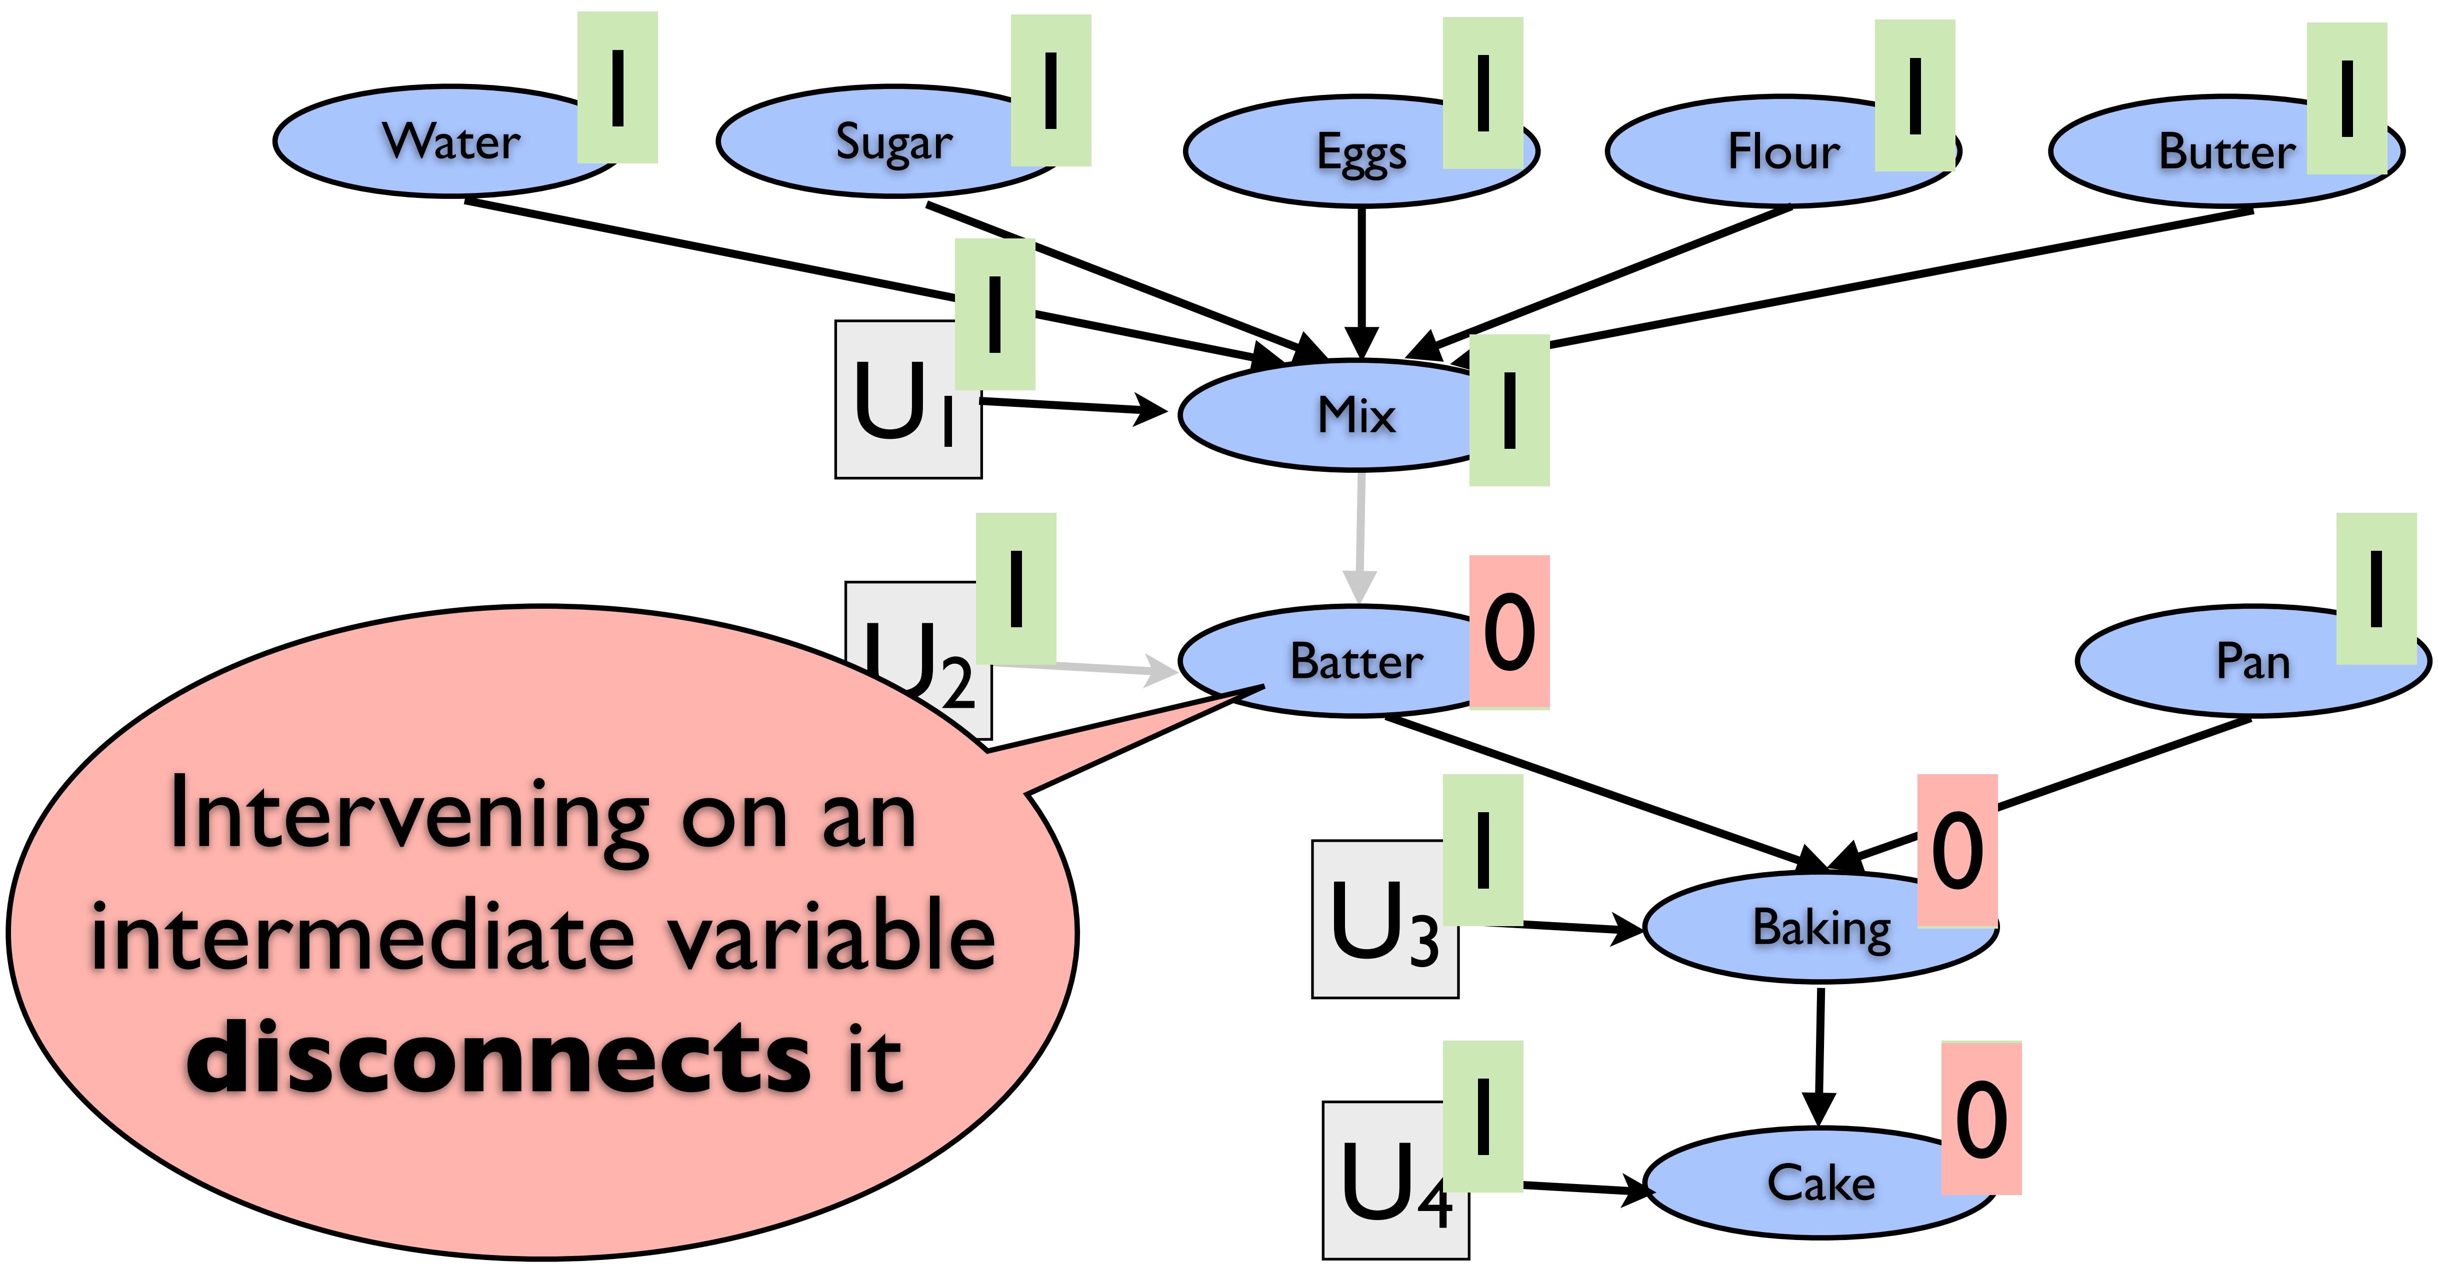

Actual causes

- Halpern-Pearl (2005) give a definition of **actual cause**

Definition 2 (Actual cause). Let (M,σ) be a causal situation. Let $\vec{X}$ be a subset of V and $Y \in V$, and suppose $\vec{x} = \sigma(\vec{X})$ and $y = \sigma(Y)$. Suppose that:

1. $\sigma(\vec{X}) = \vec{x}$ and $\sigma(Y) = y$.
2. Some set of variables $W \subseteq V - X$ and values $\vec{x}' \in D$, and $\vec{w}' \in D$ exist such that:
   (a) $Y \neq y$ holds in $M_{\vec{x}',\vec{w}'}$
   (b) $Y = y$ holds in $M_{\vec{x},\vec{w}',\vec{z}}$ for all $Z \subseteq V - (X \cup W)$, where $\vec{z}$ are the values of $\vec{Z}$ in M.

Then we say that $\vec{X} = \vec{x}$ is a *weak cause* of $Y = y$. Moreover, if no proper subset of $\vec{X} = \vec{x}$ is a weak cause, then $\vec{X} = \vec{x}$ is an *actual cause* of $Y = y$.

- Don't look directly at it! Easier via pictures

Actual causes, graphically

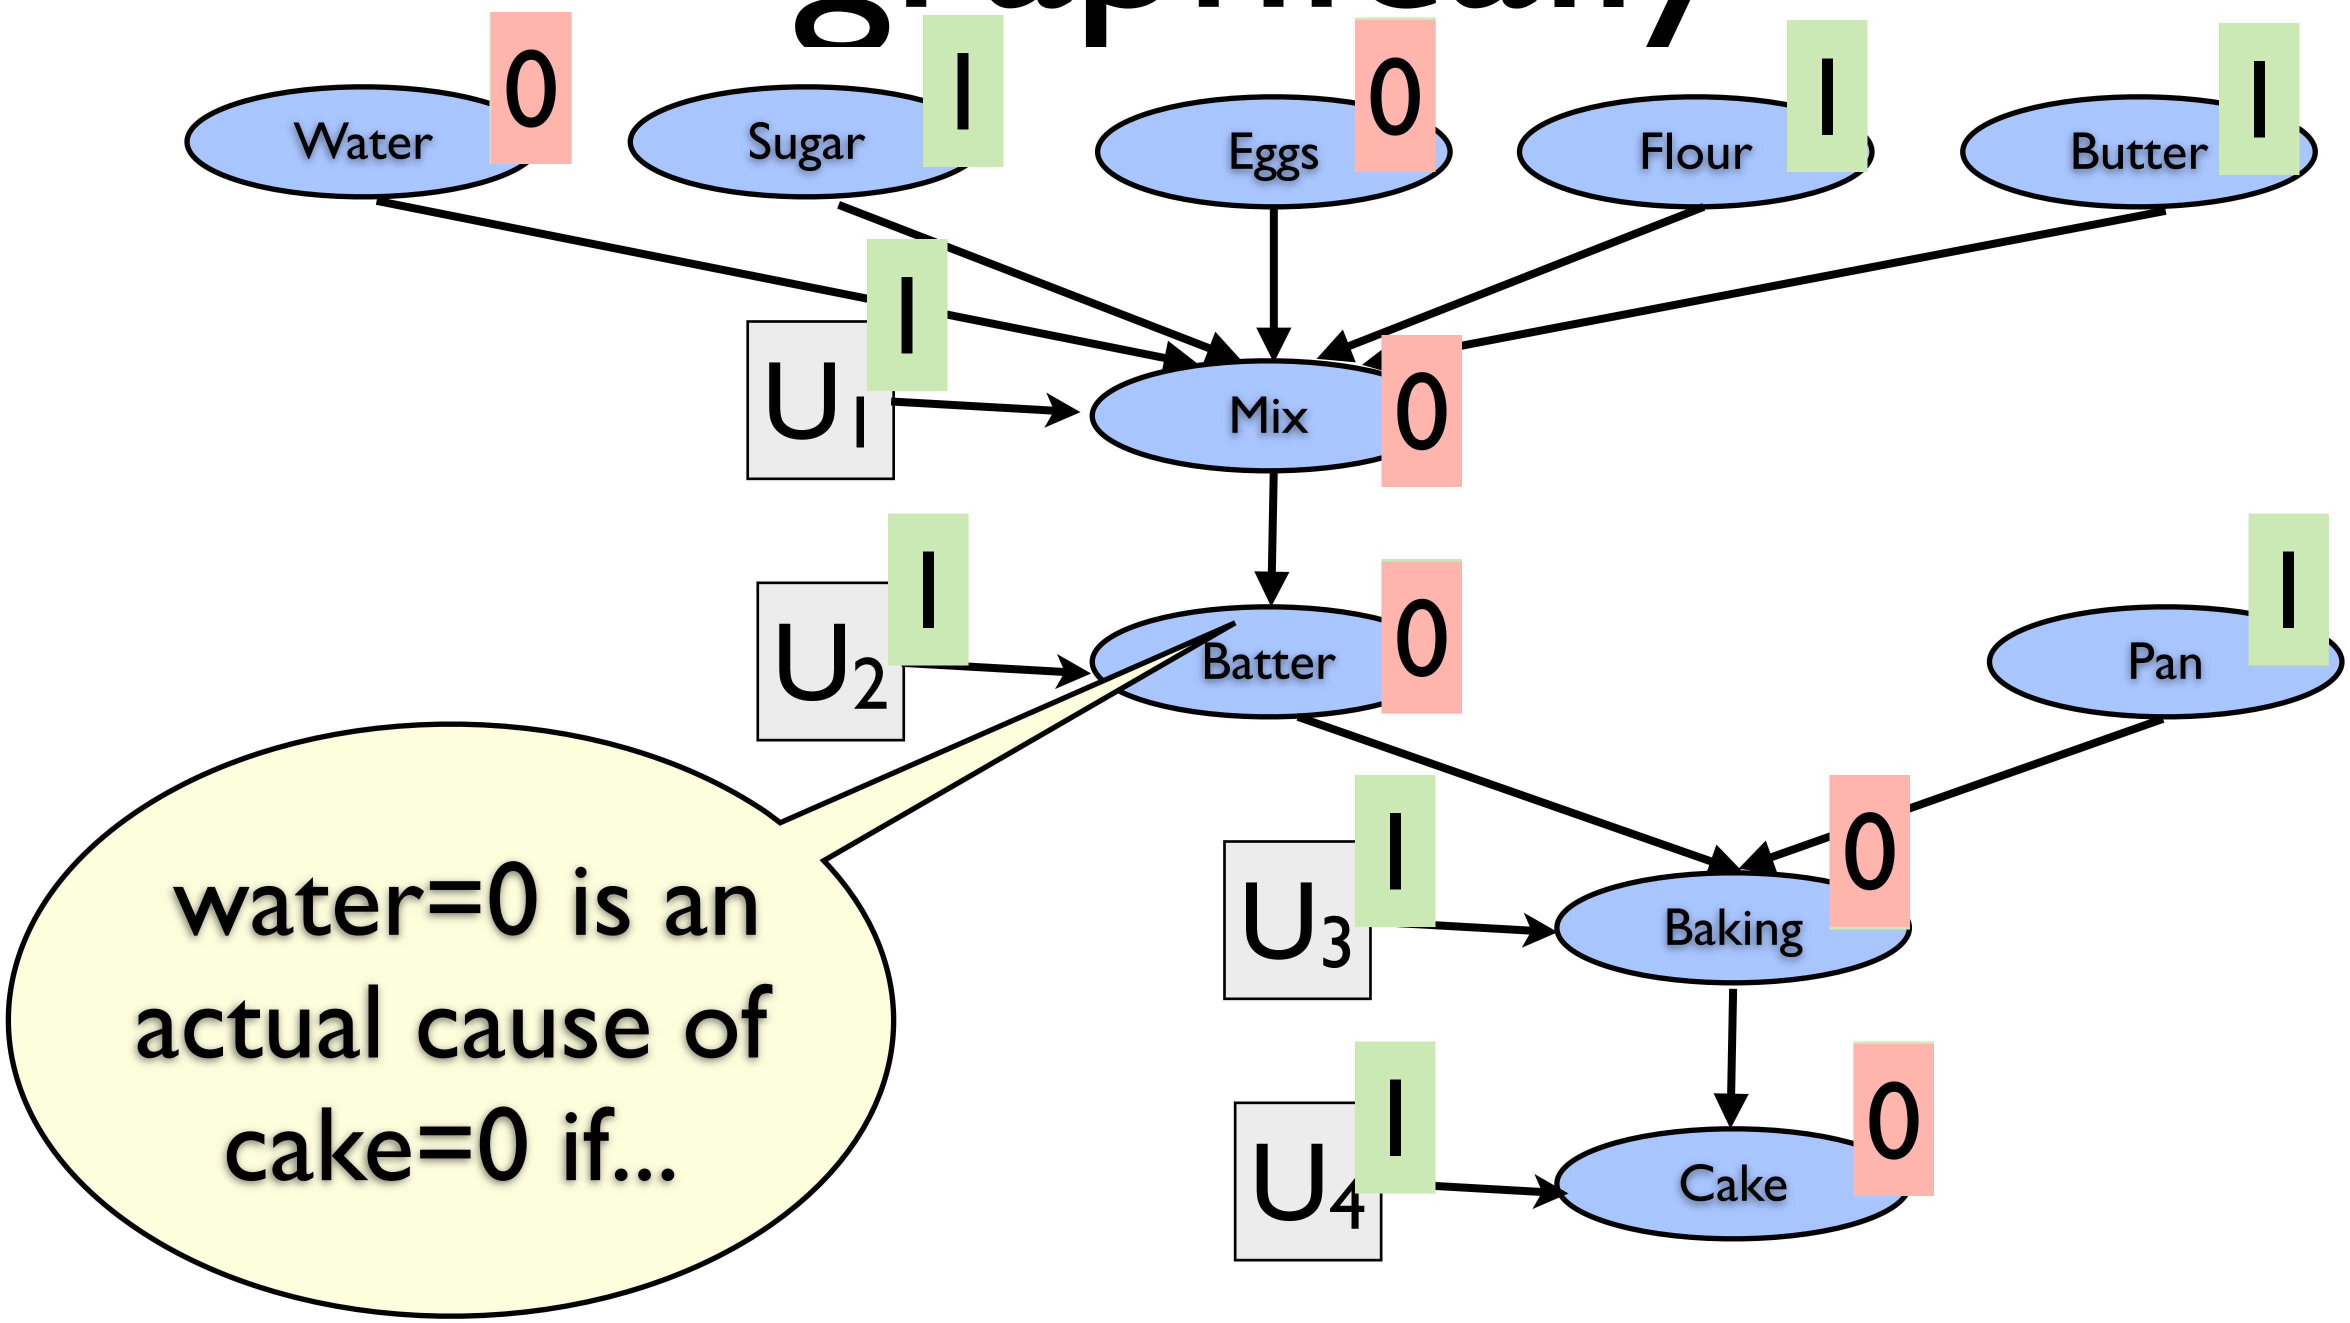

Actual causes, graphically

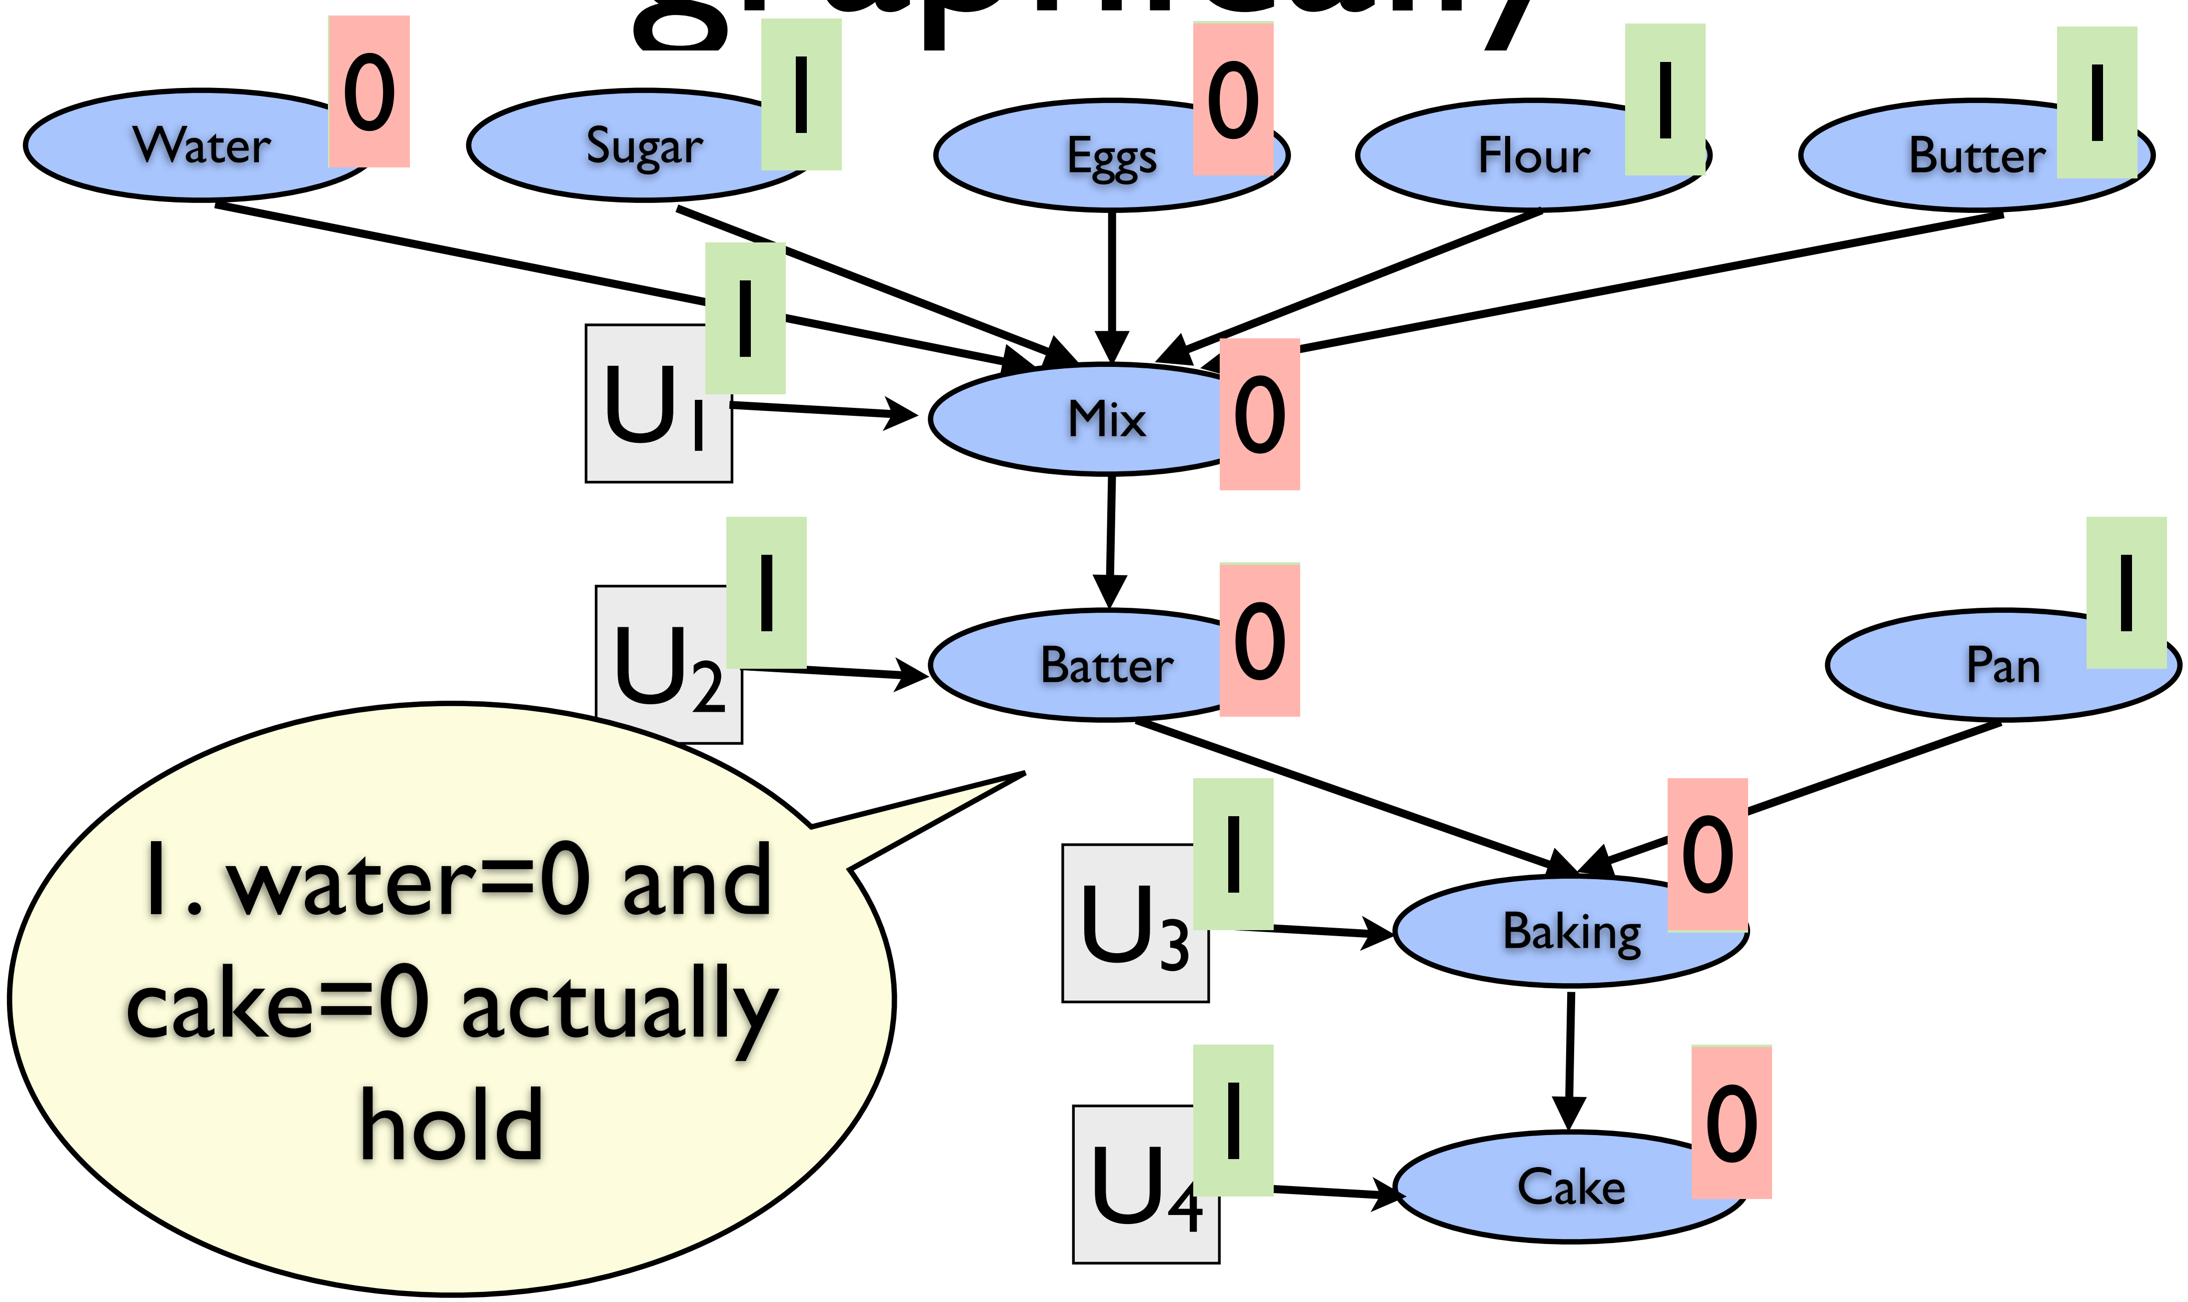

Actual causes, graphically

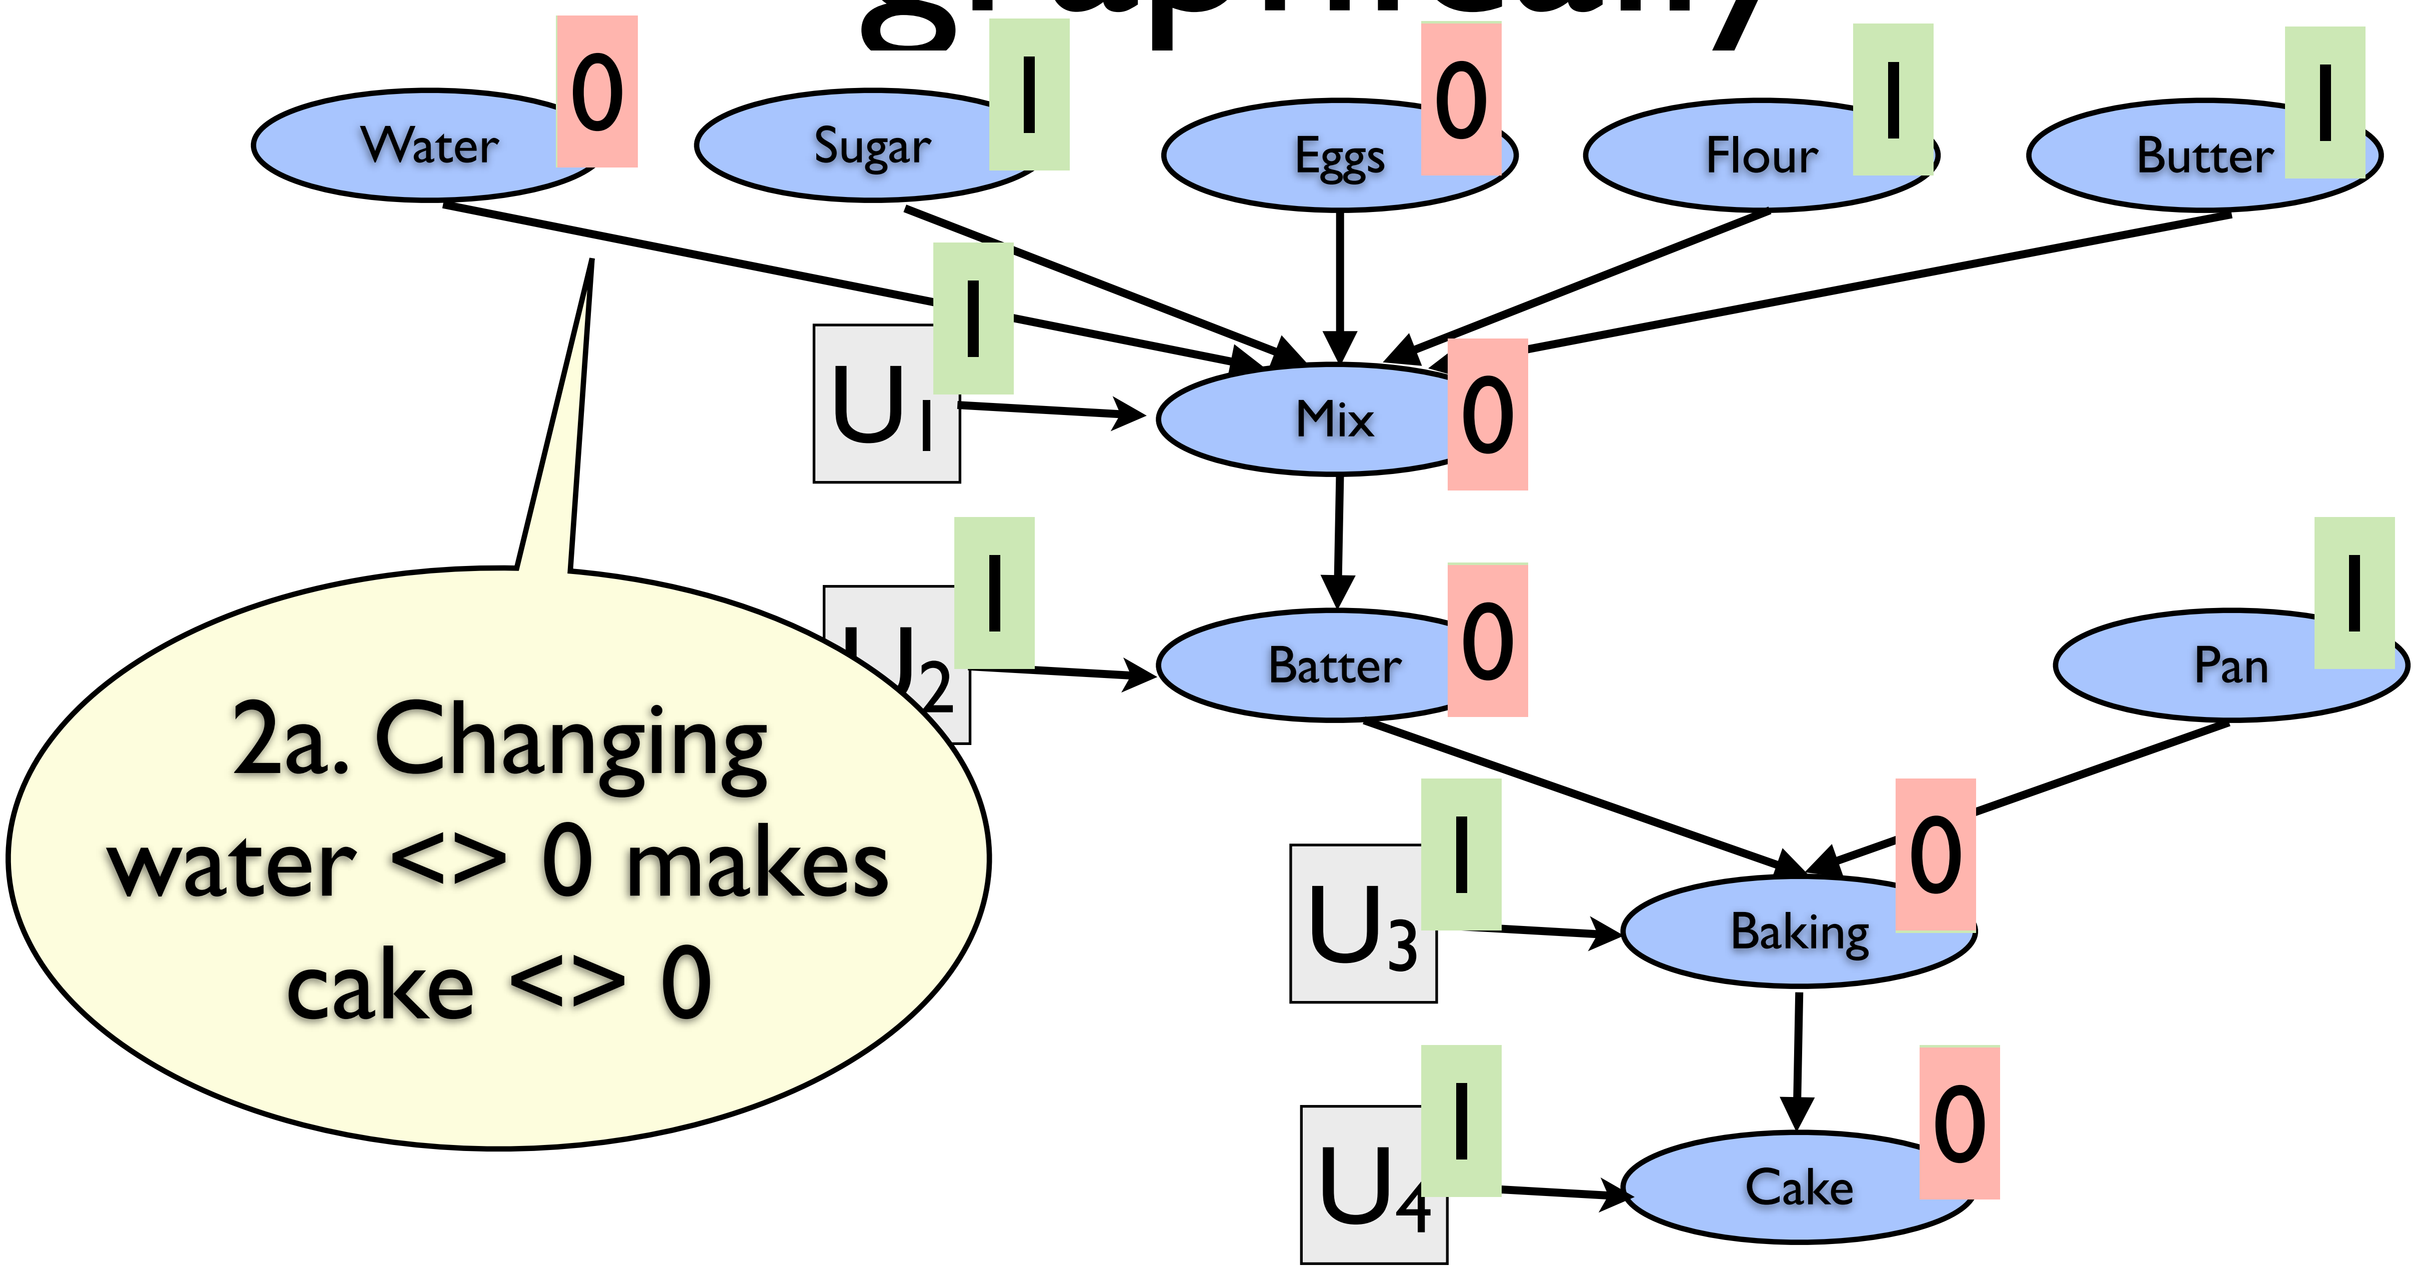

Actual causes, graphically

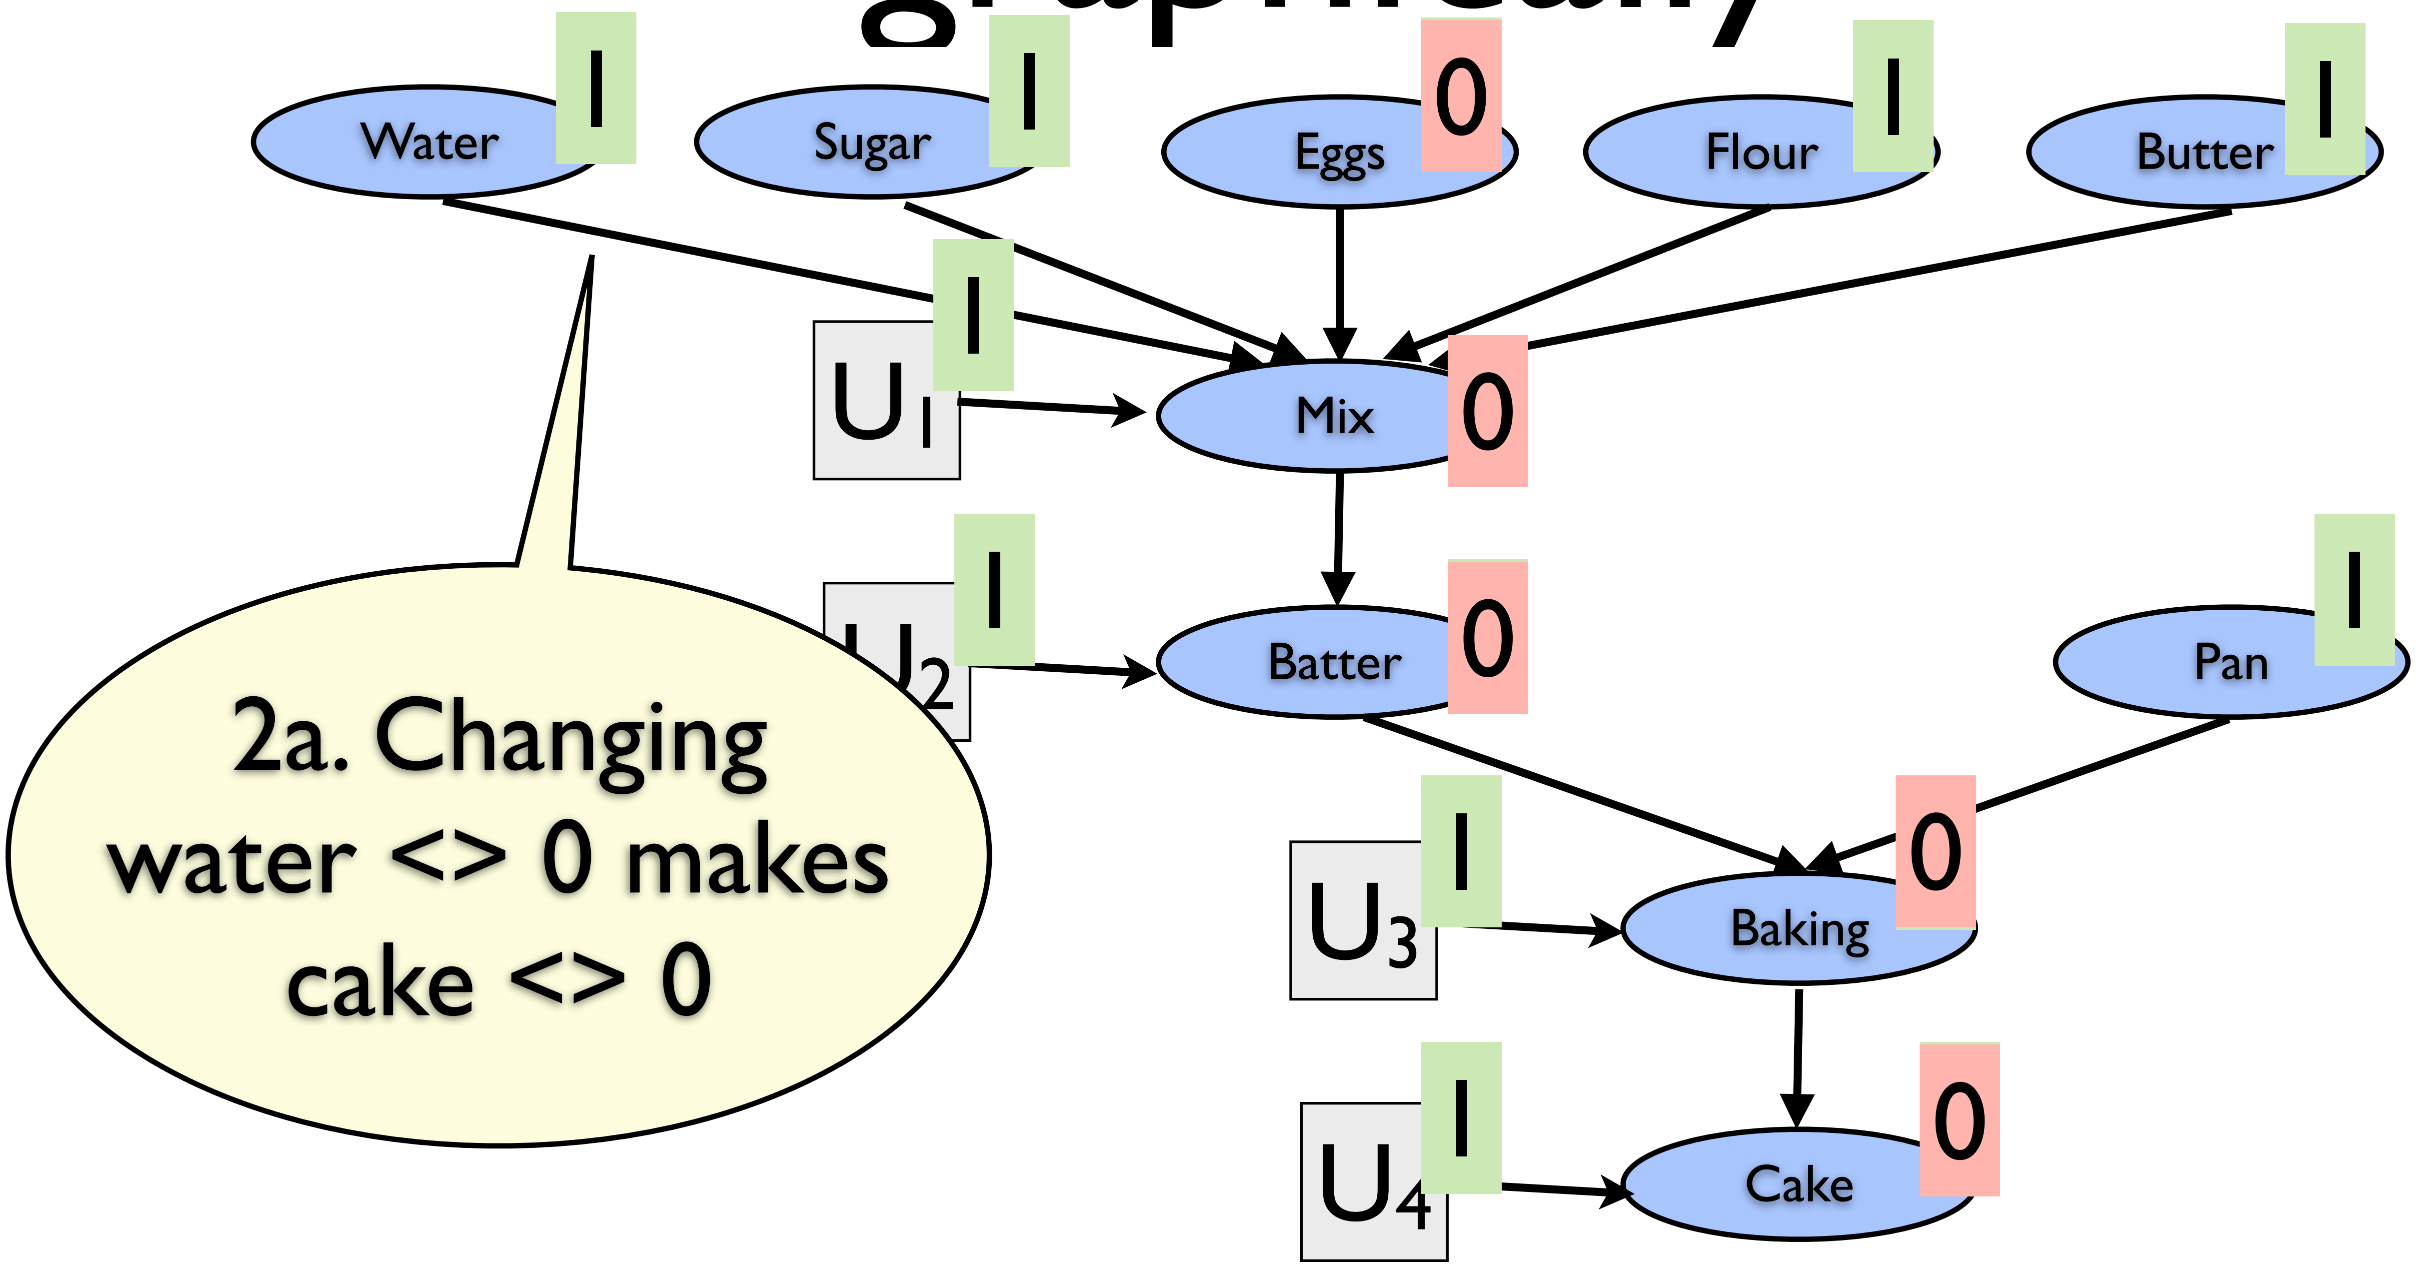

Actual causes, graphically

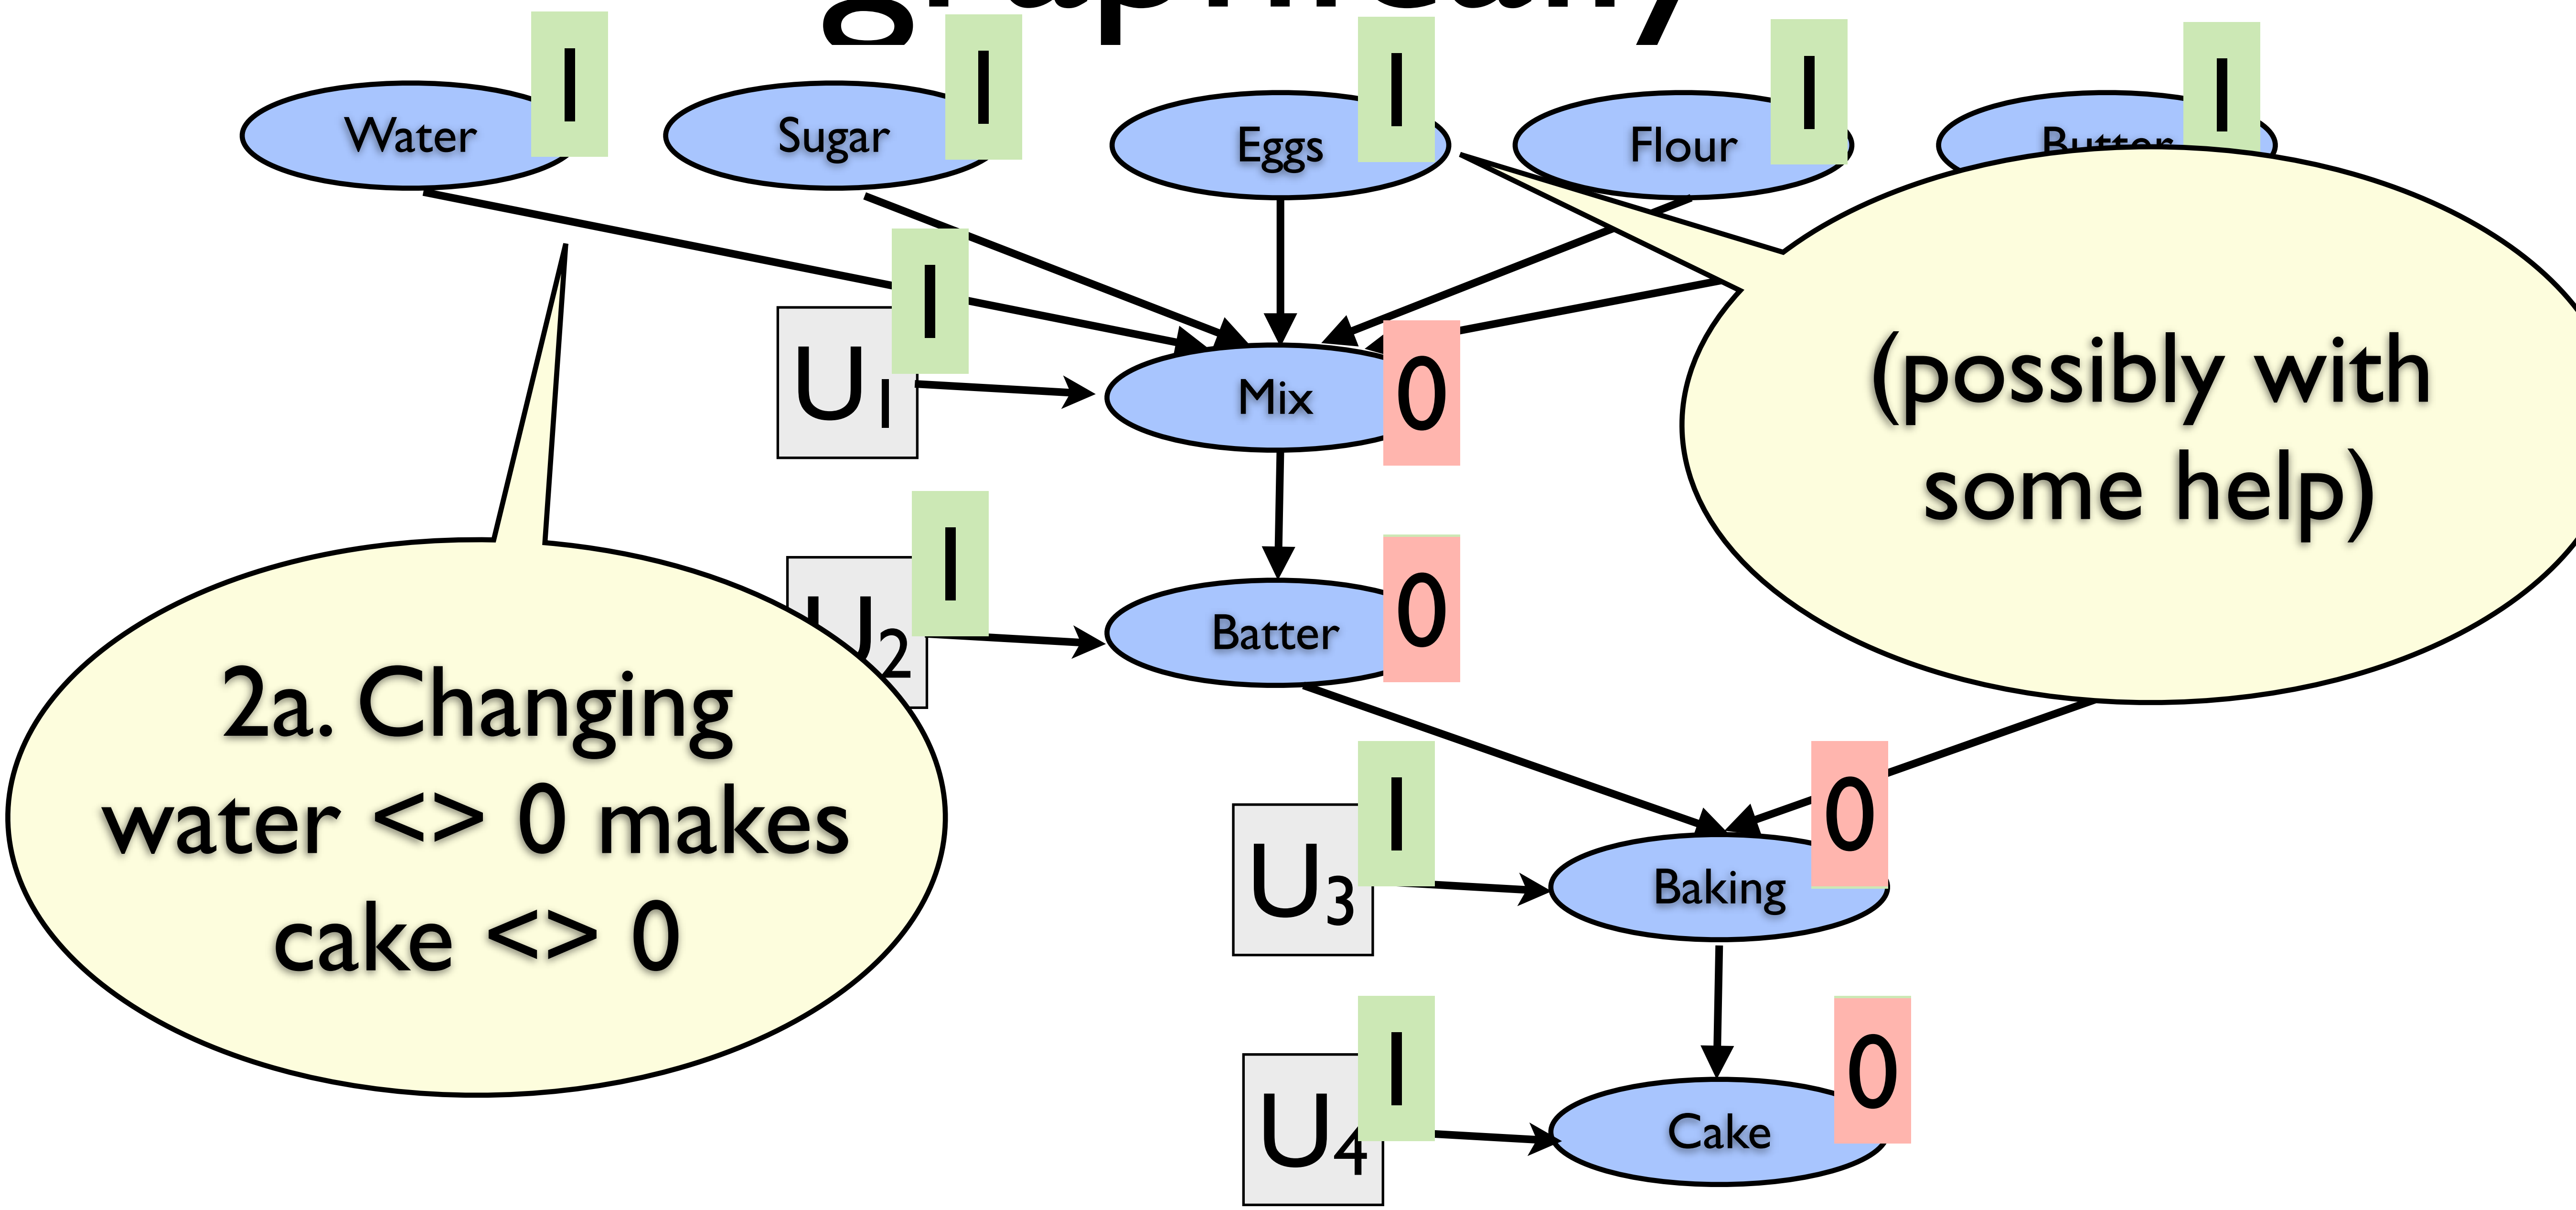

Actual causes, graphically

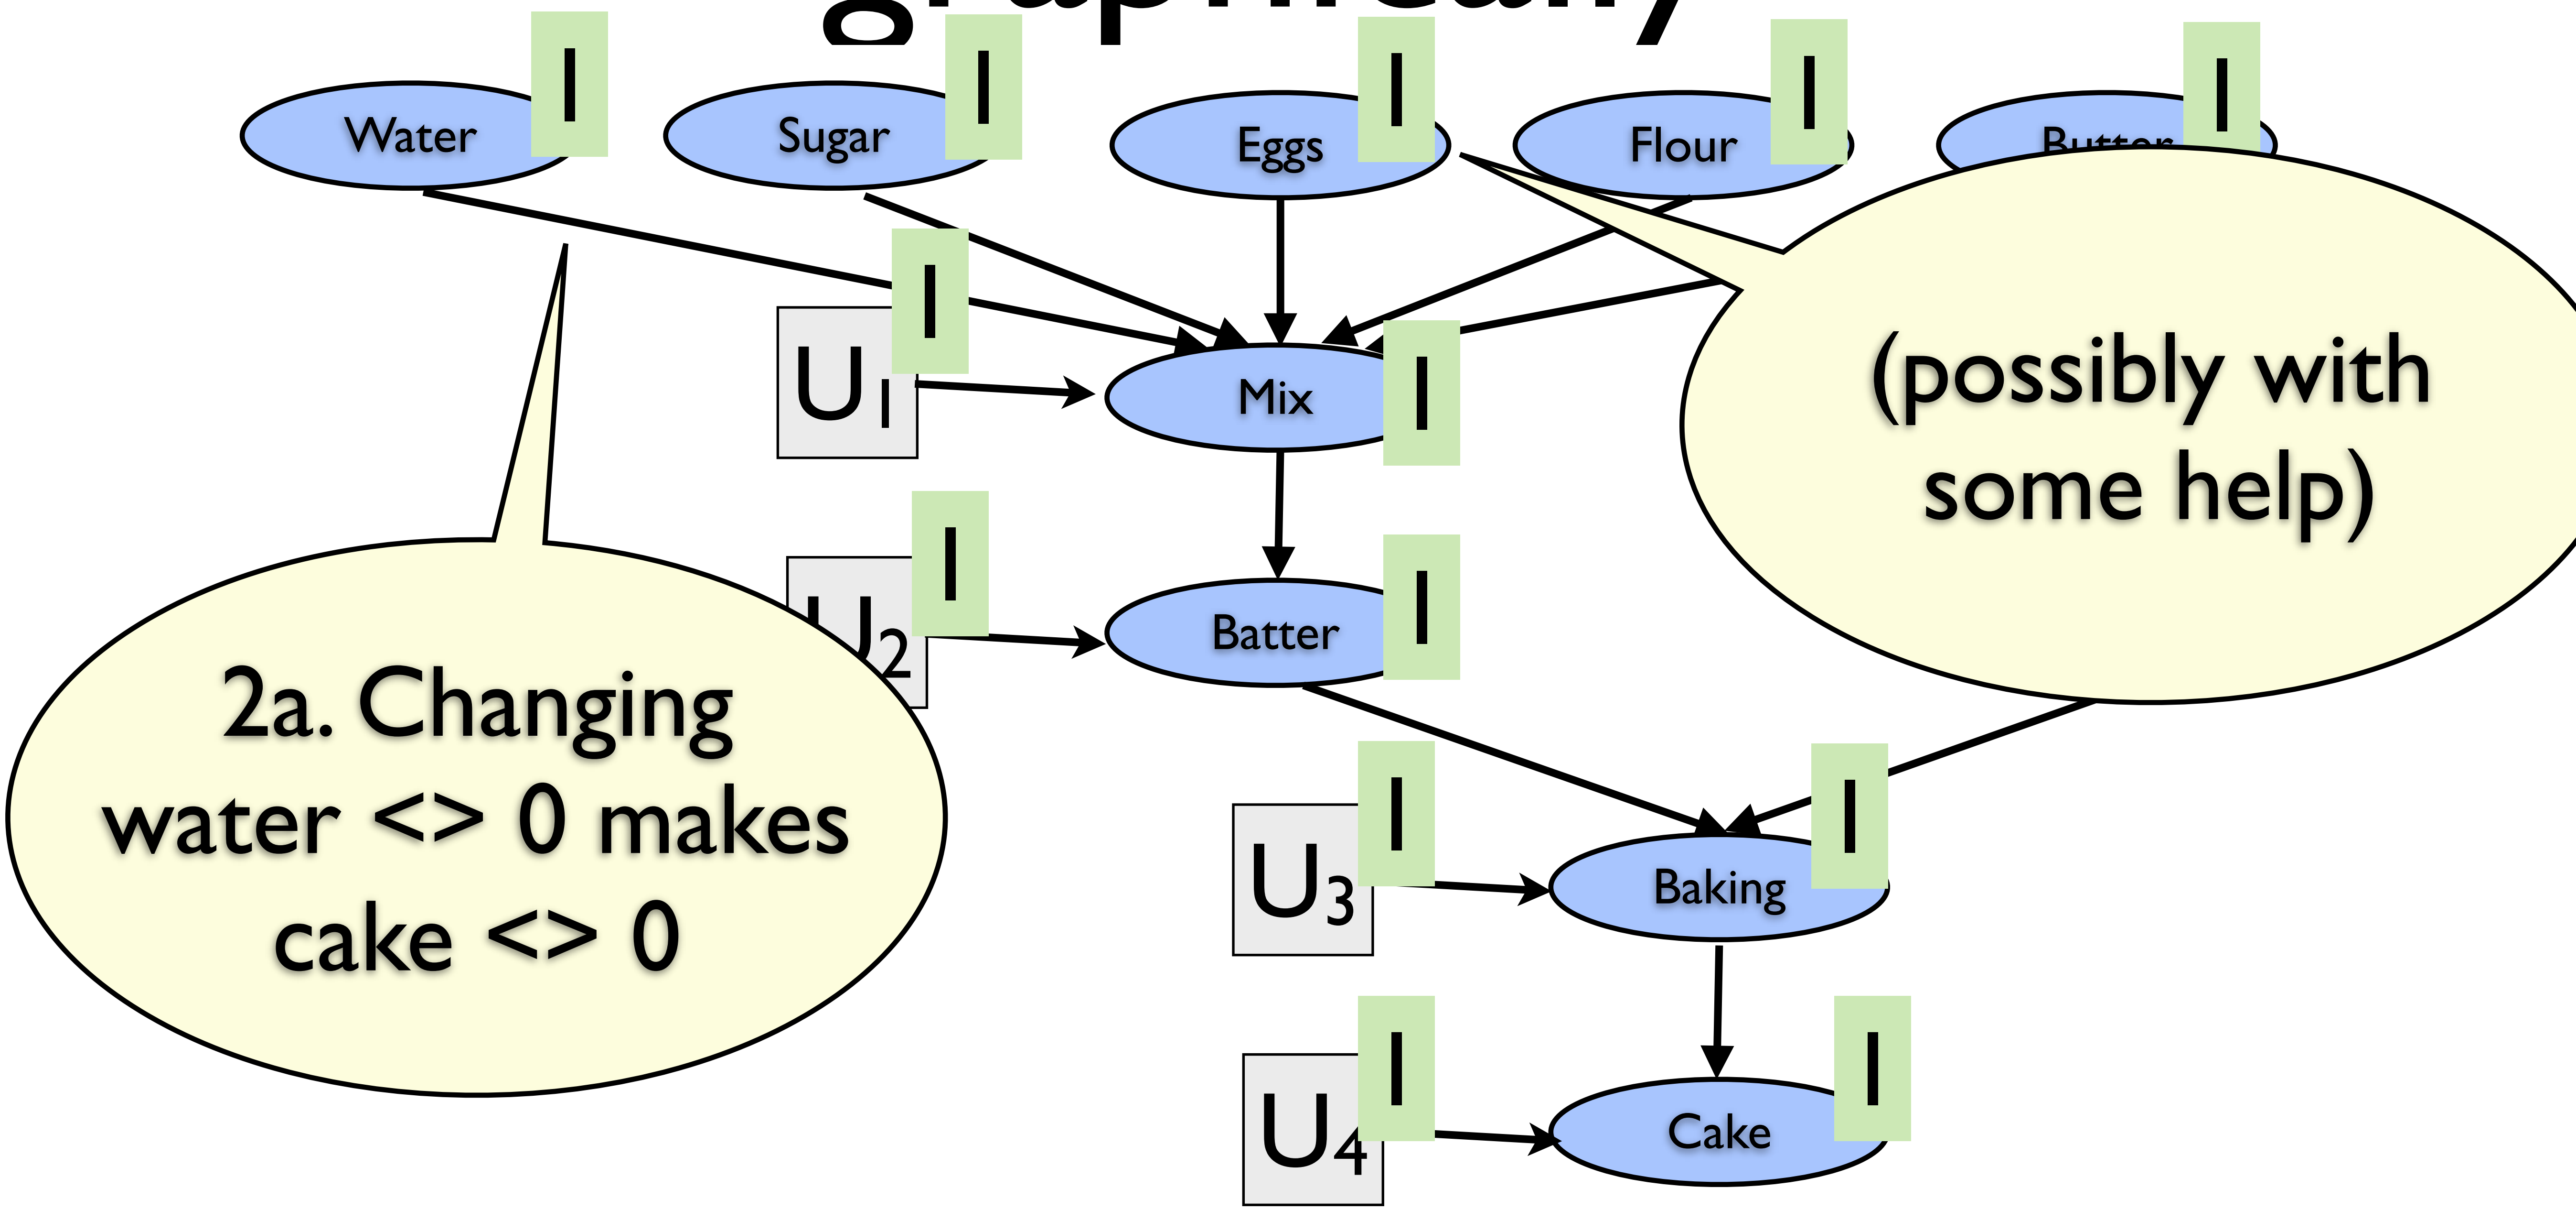

Actual causes, graphically

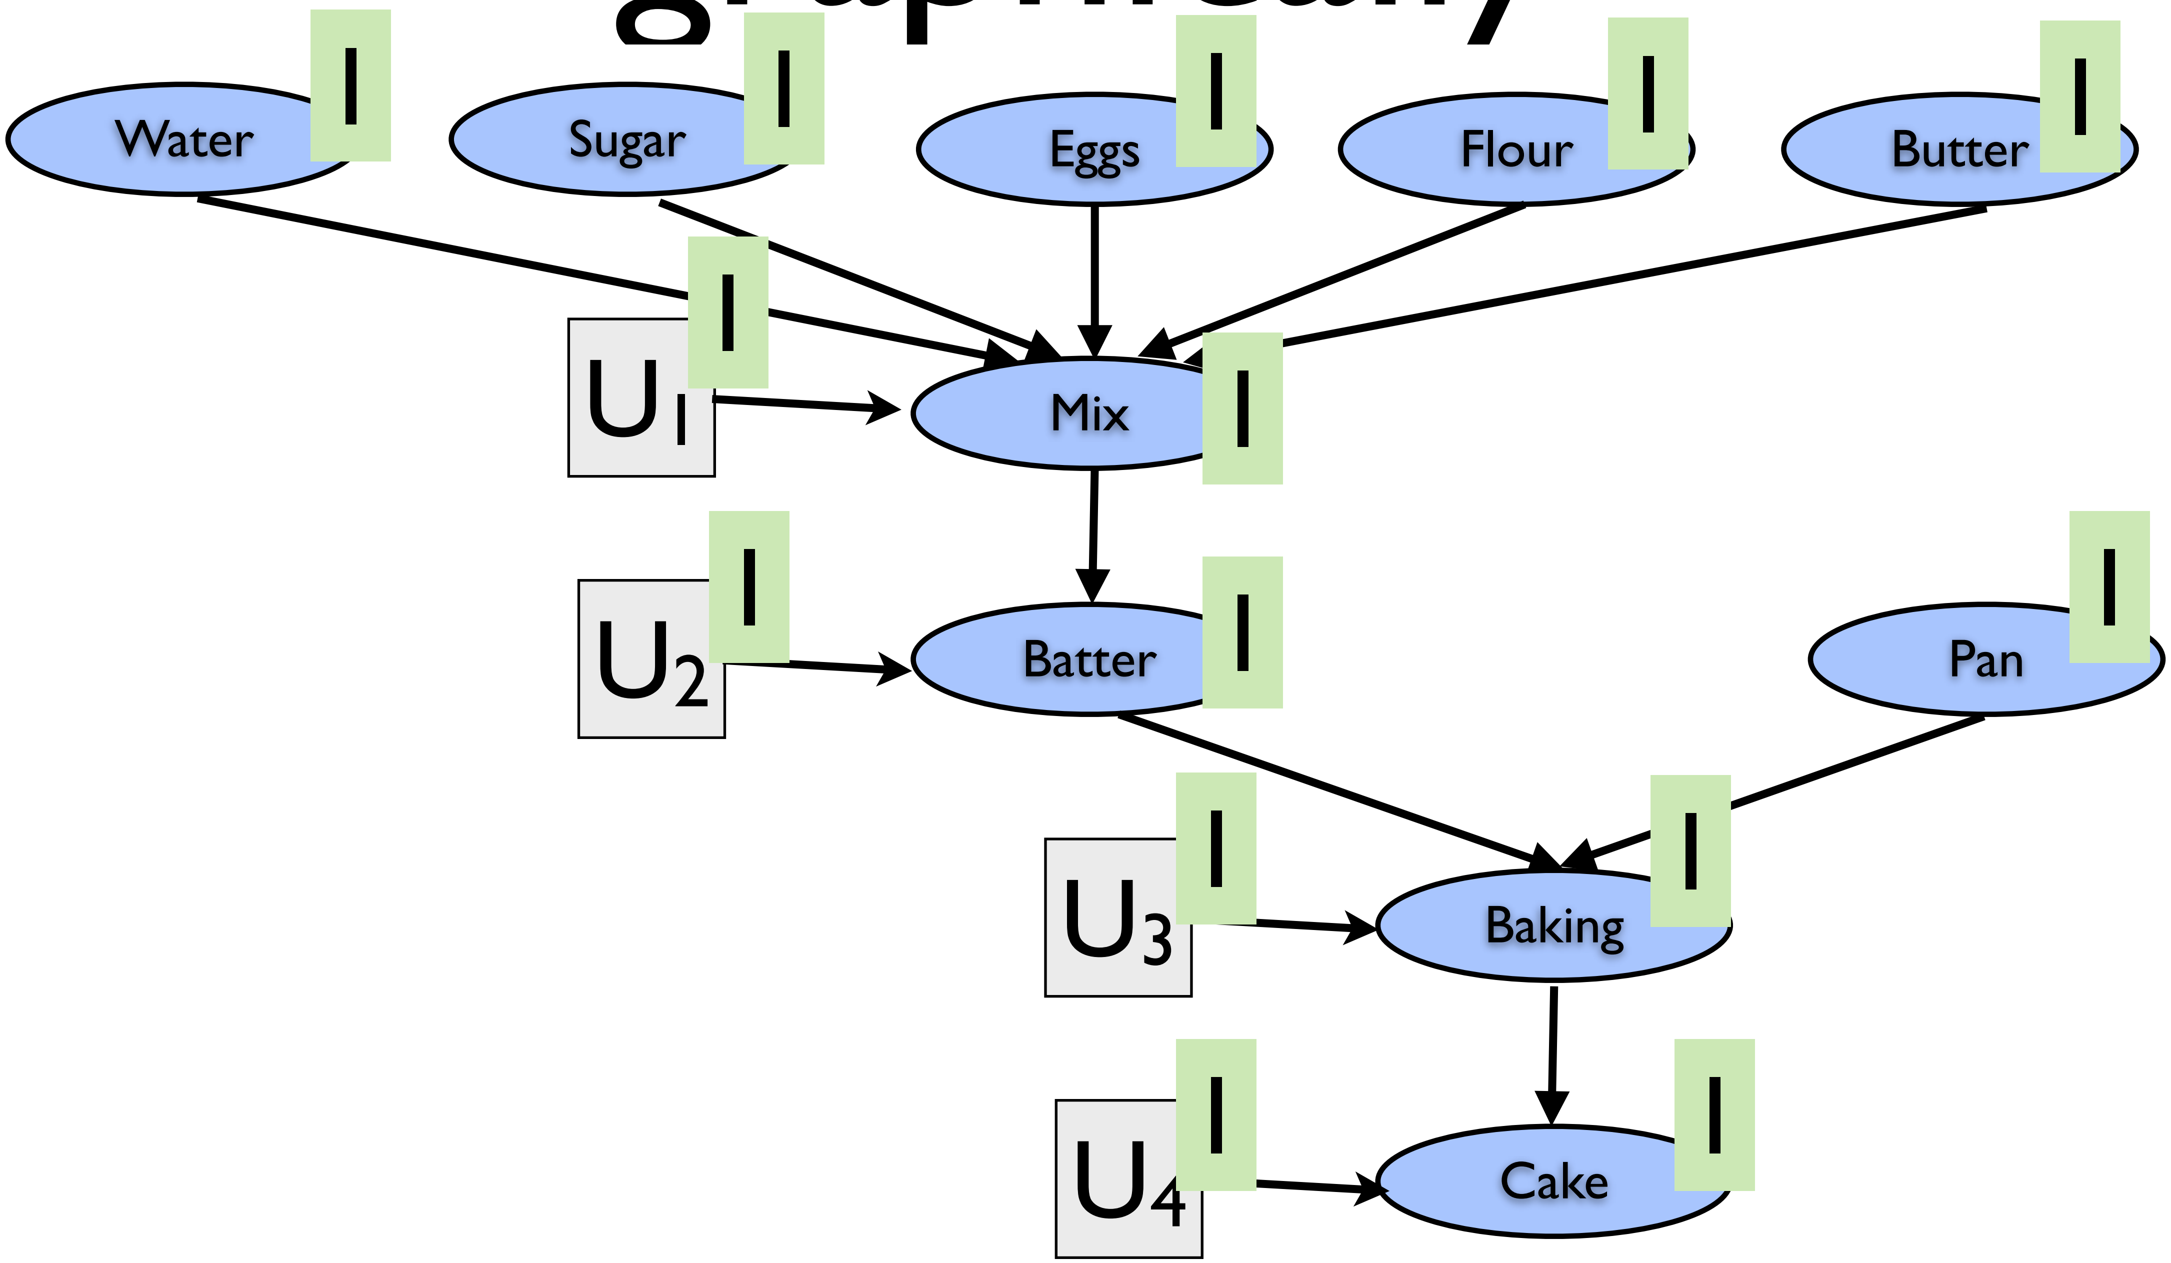

Actual causes, graphically

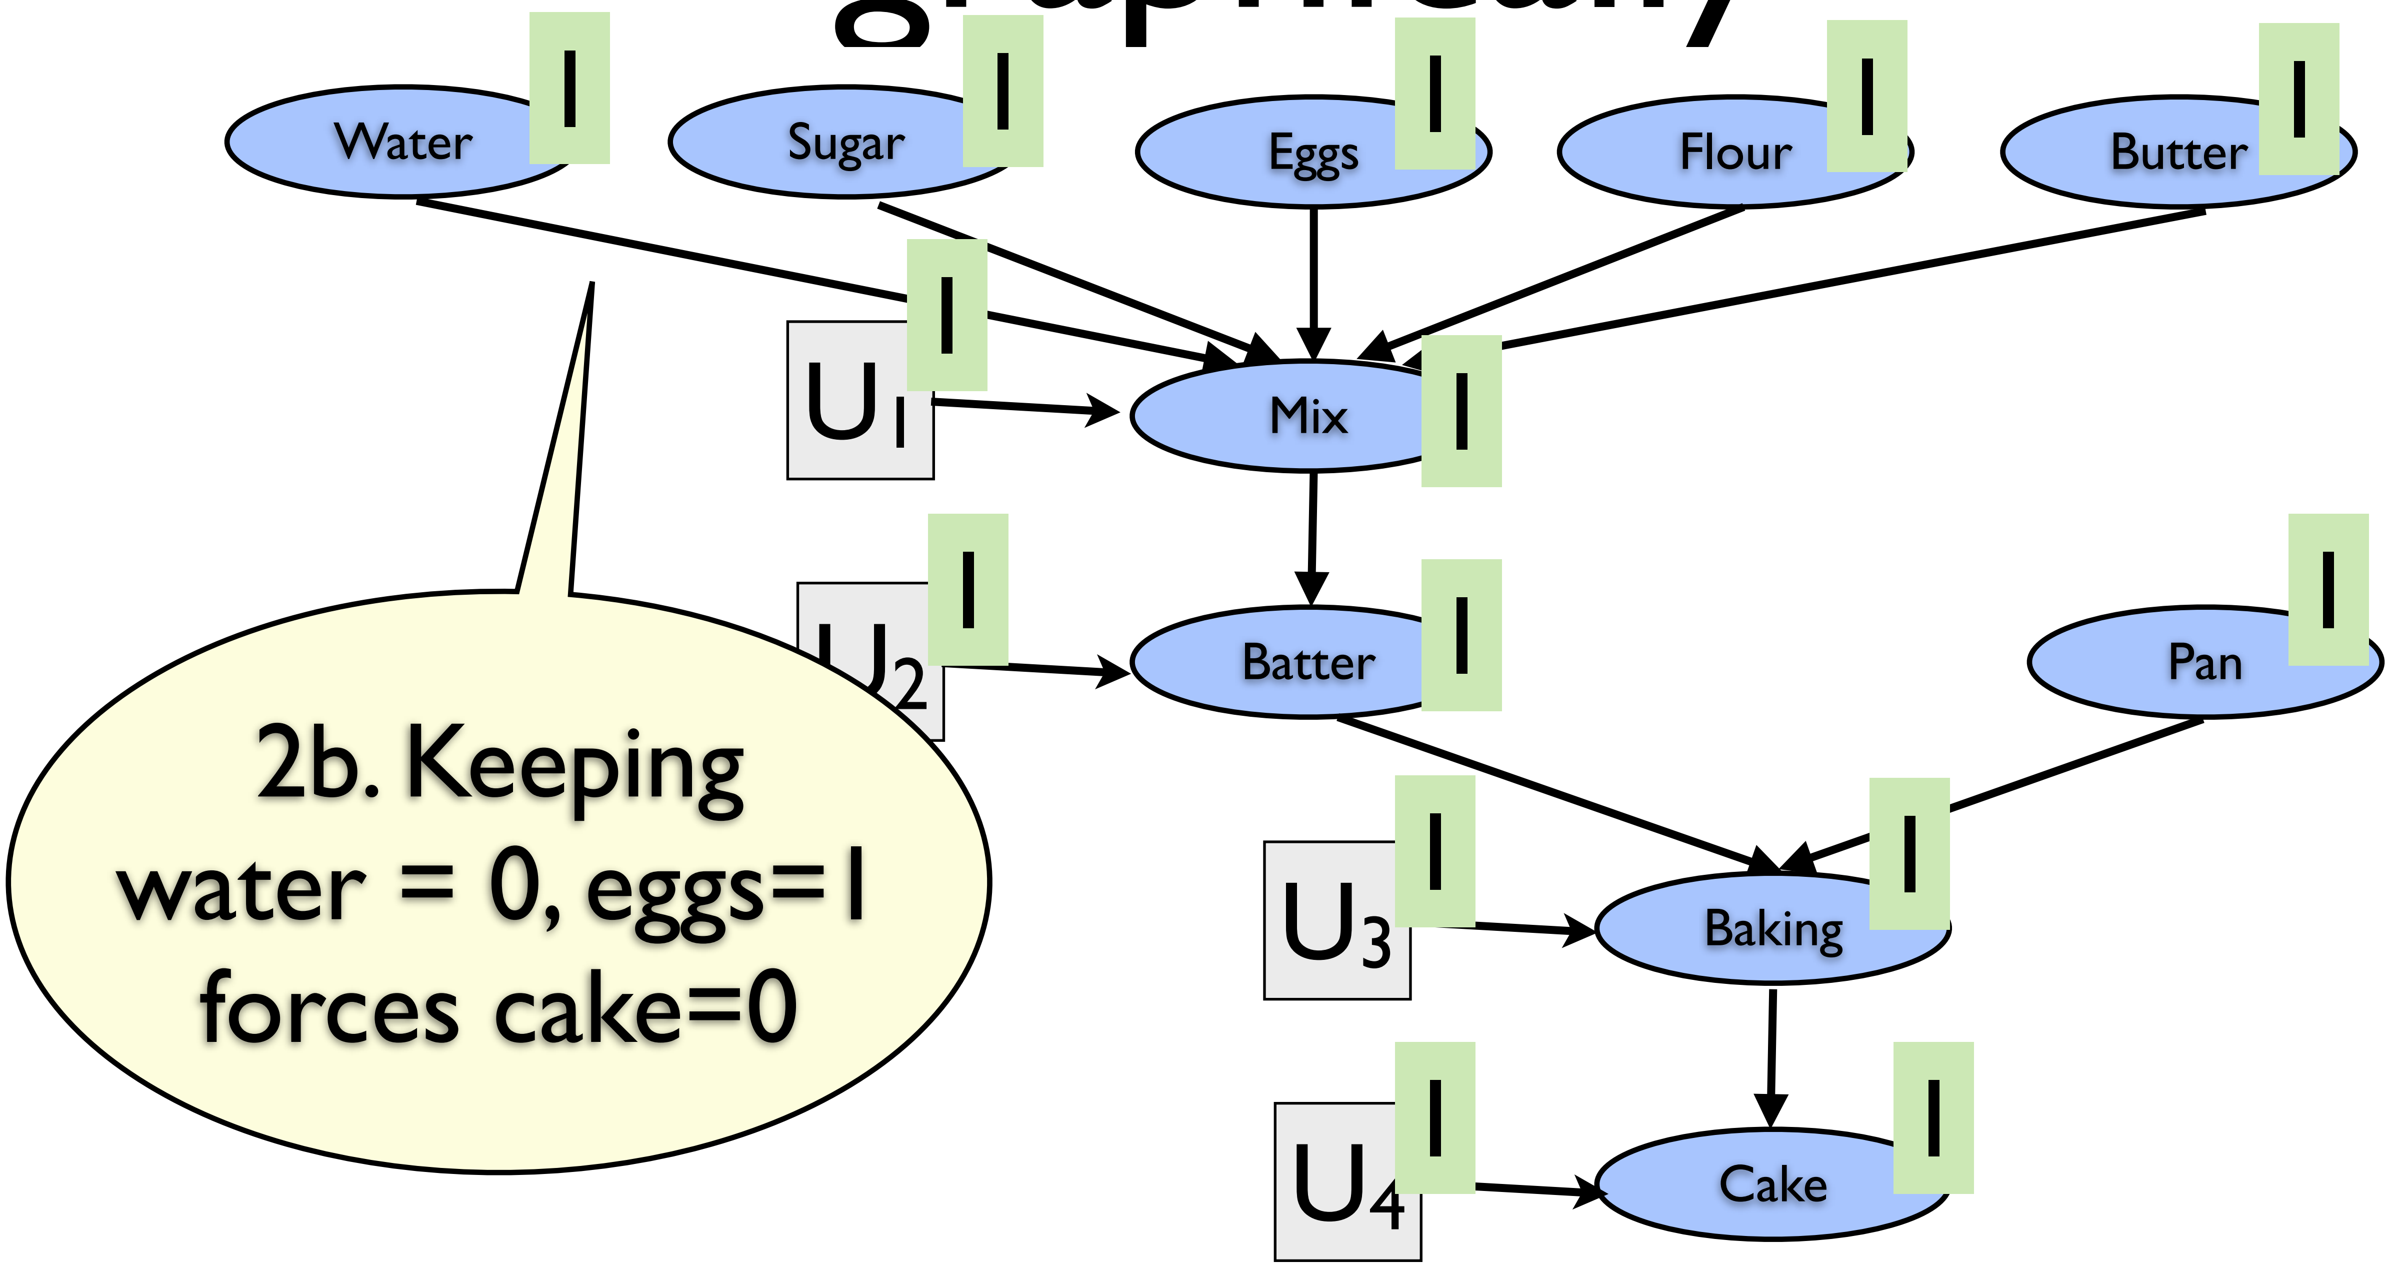

Actual causes, graphically

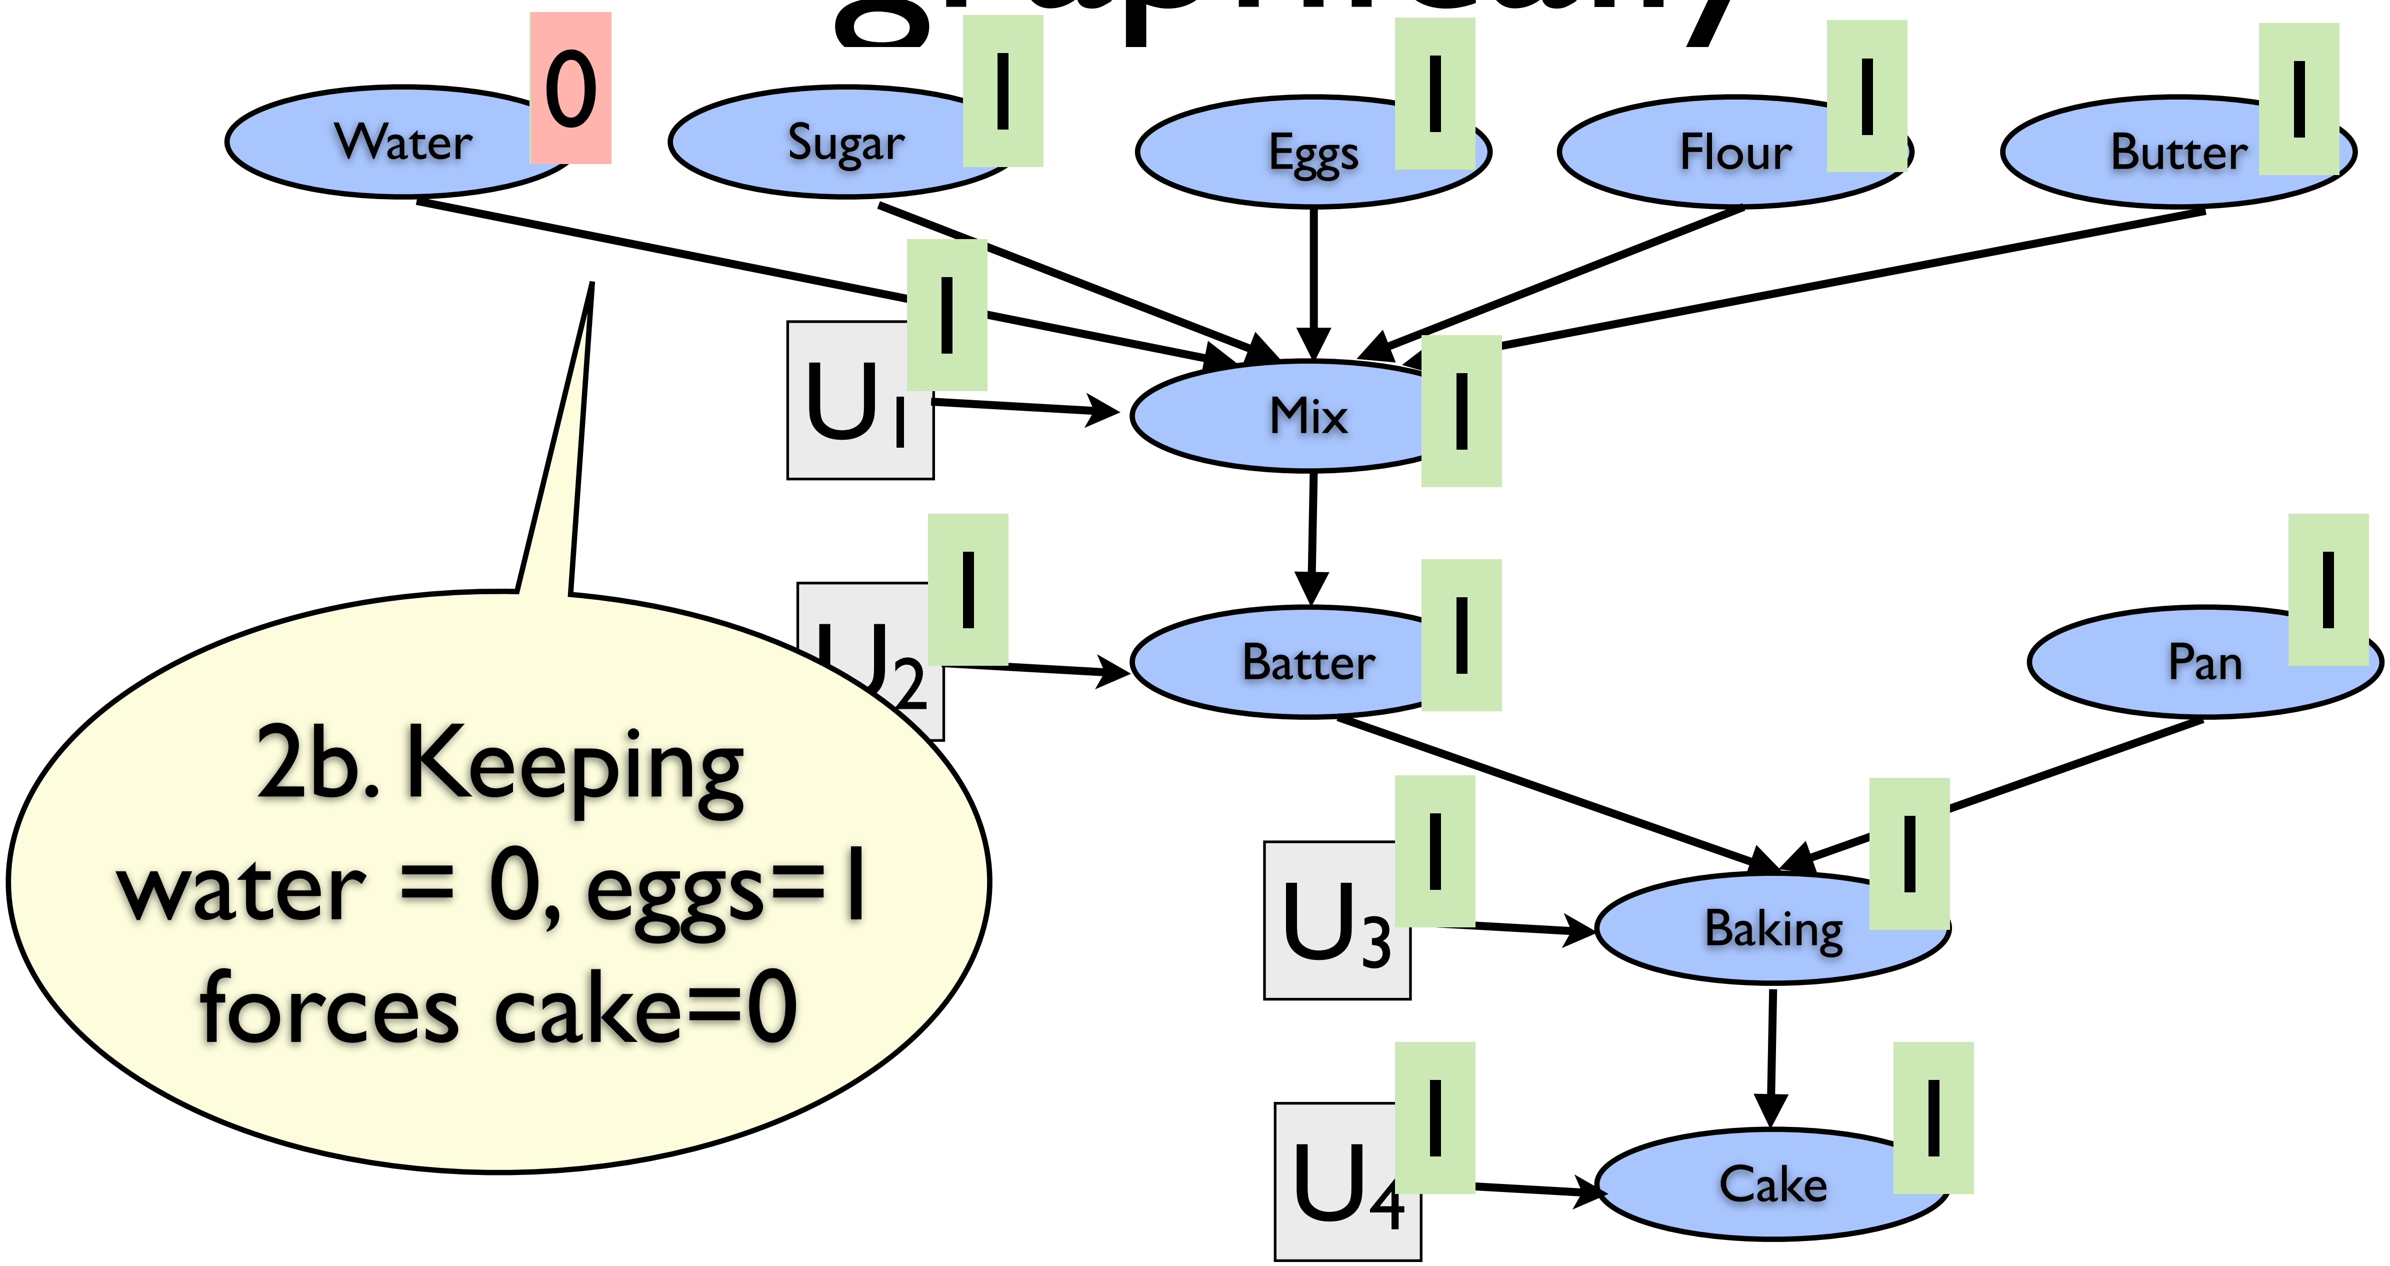

Actual causes, graphically

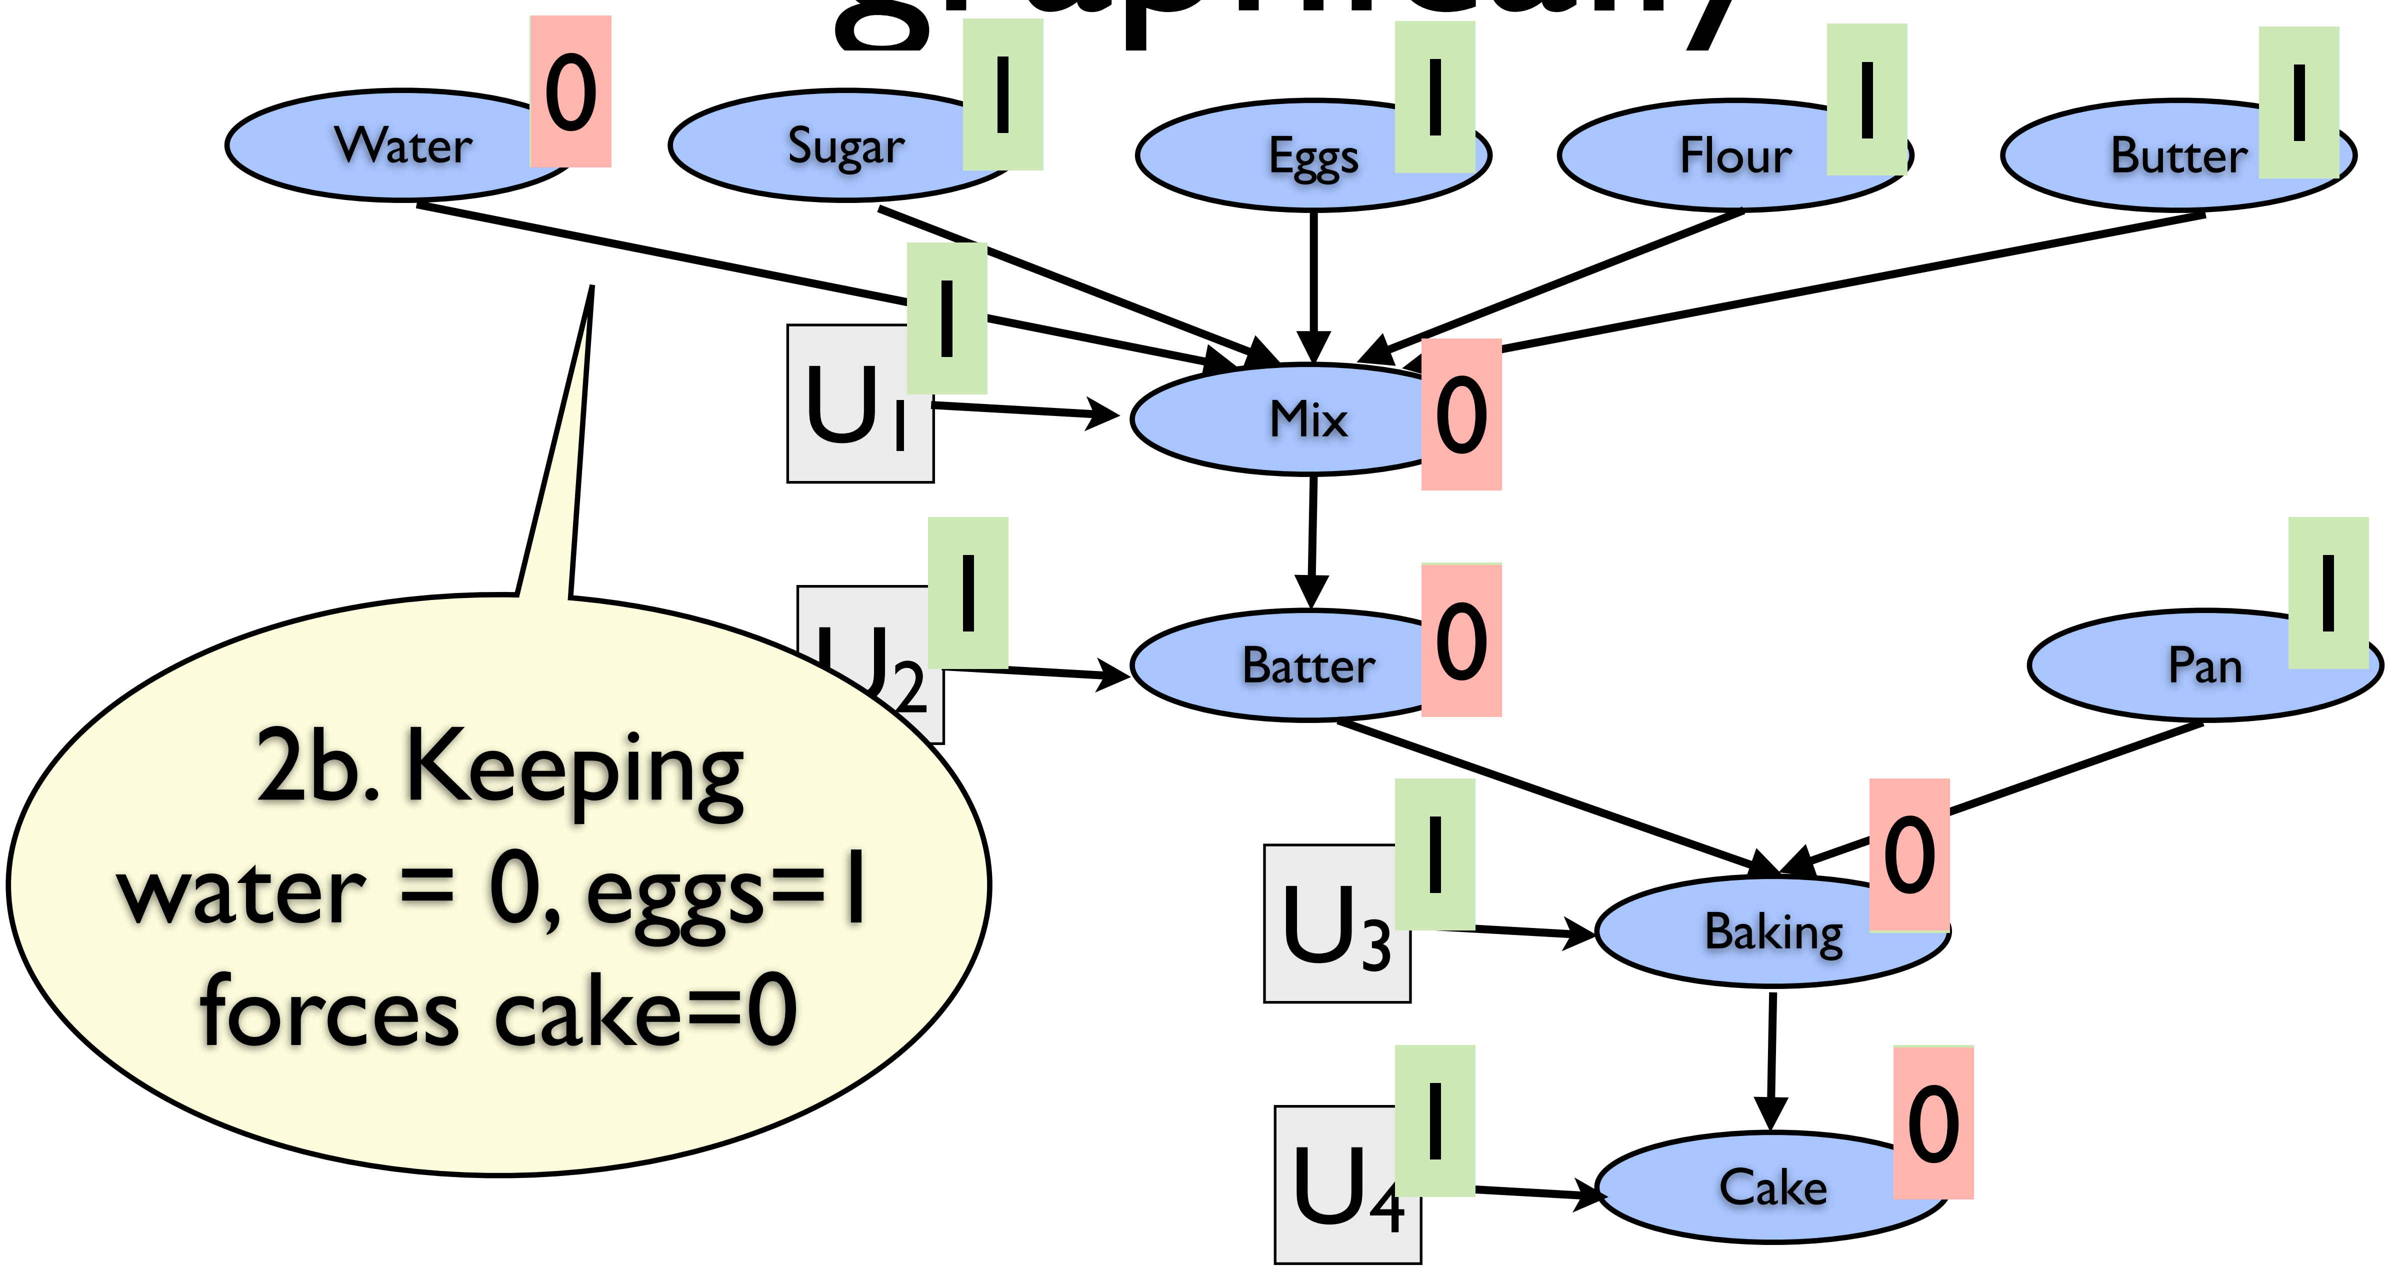

Icing on the cake

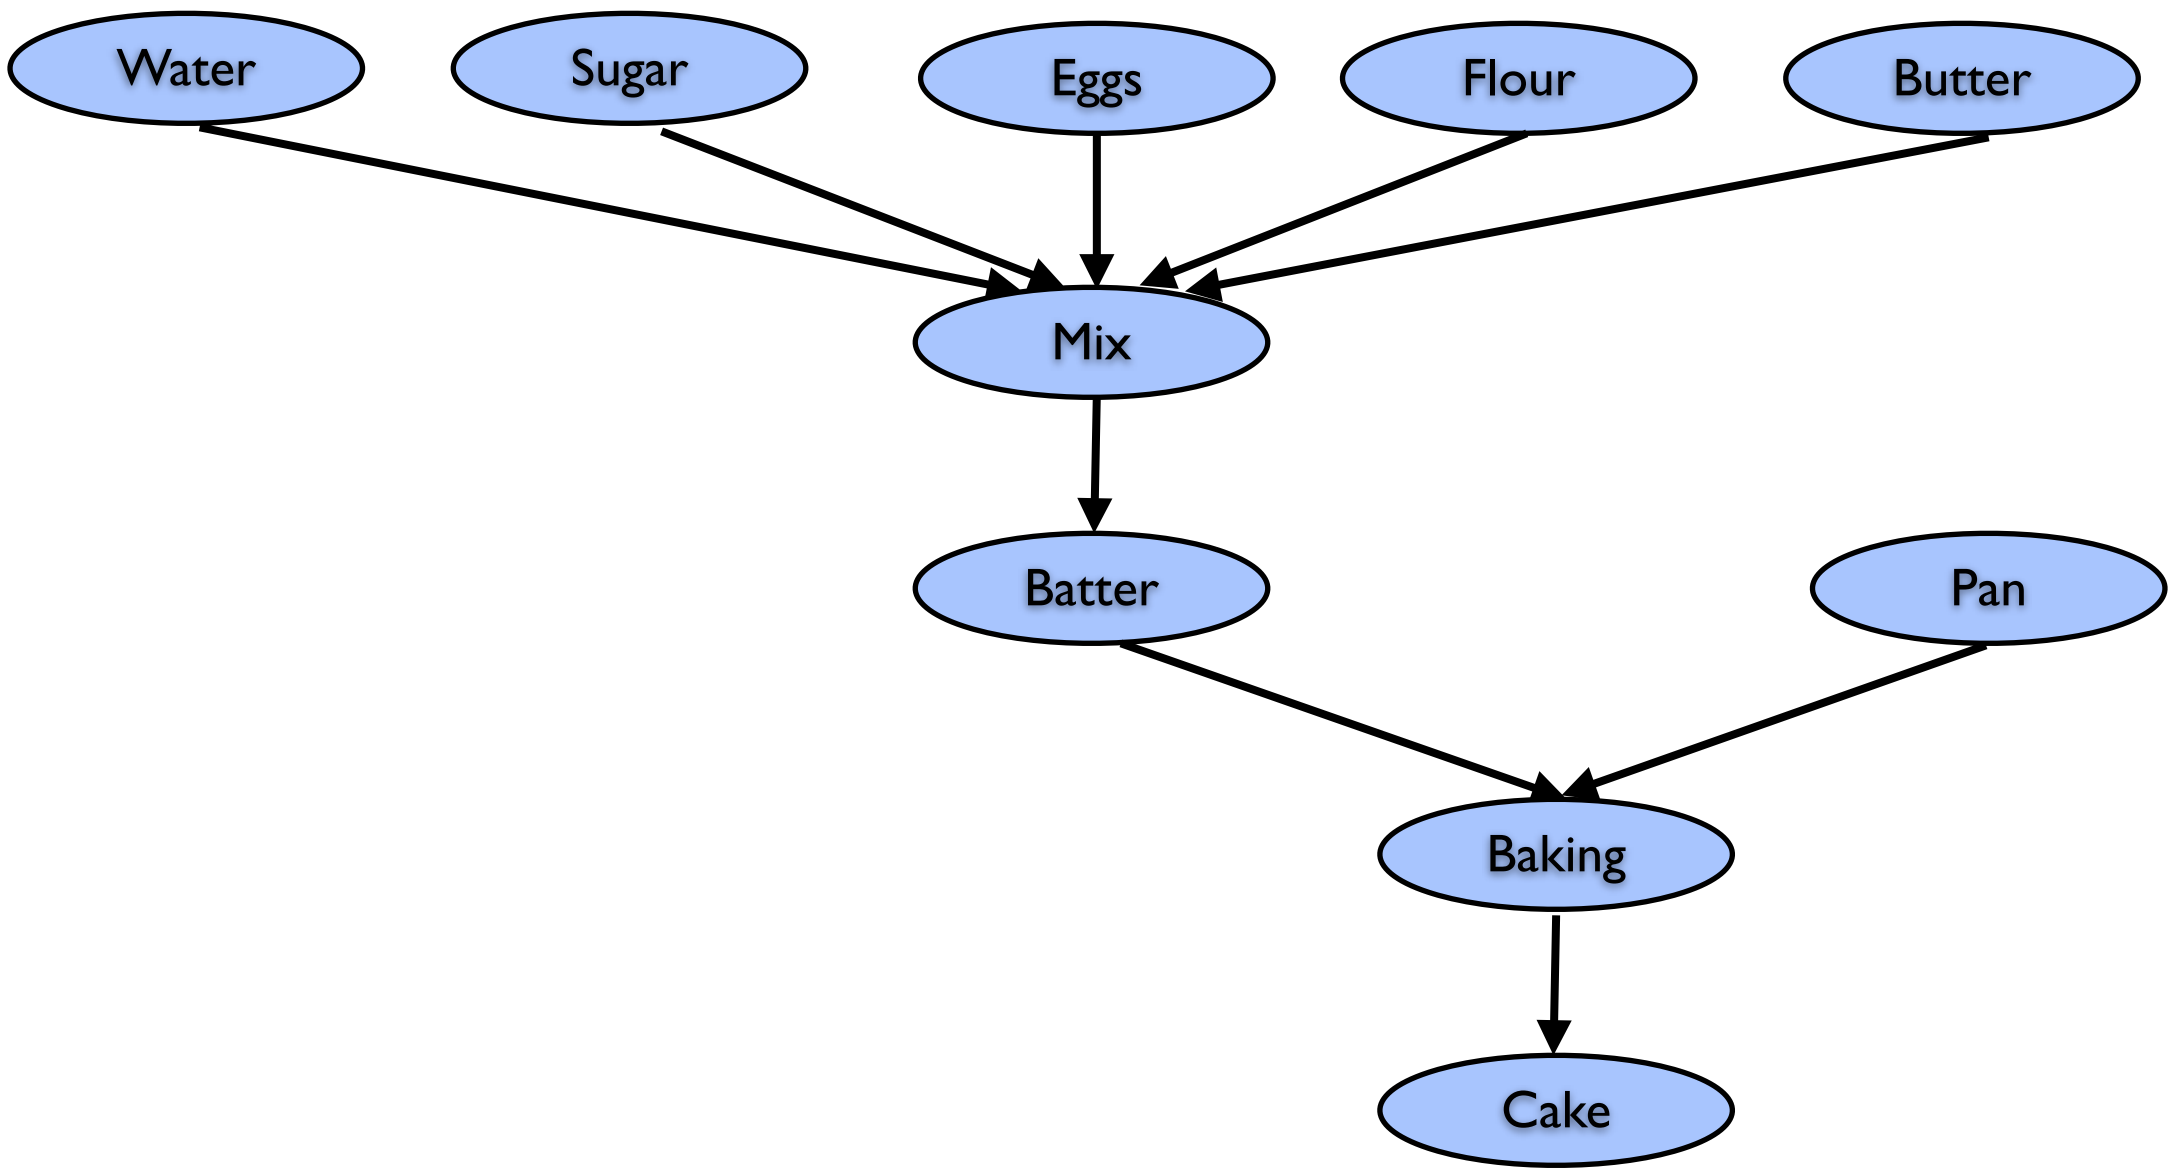

Icing on the cake

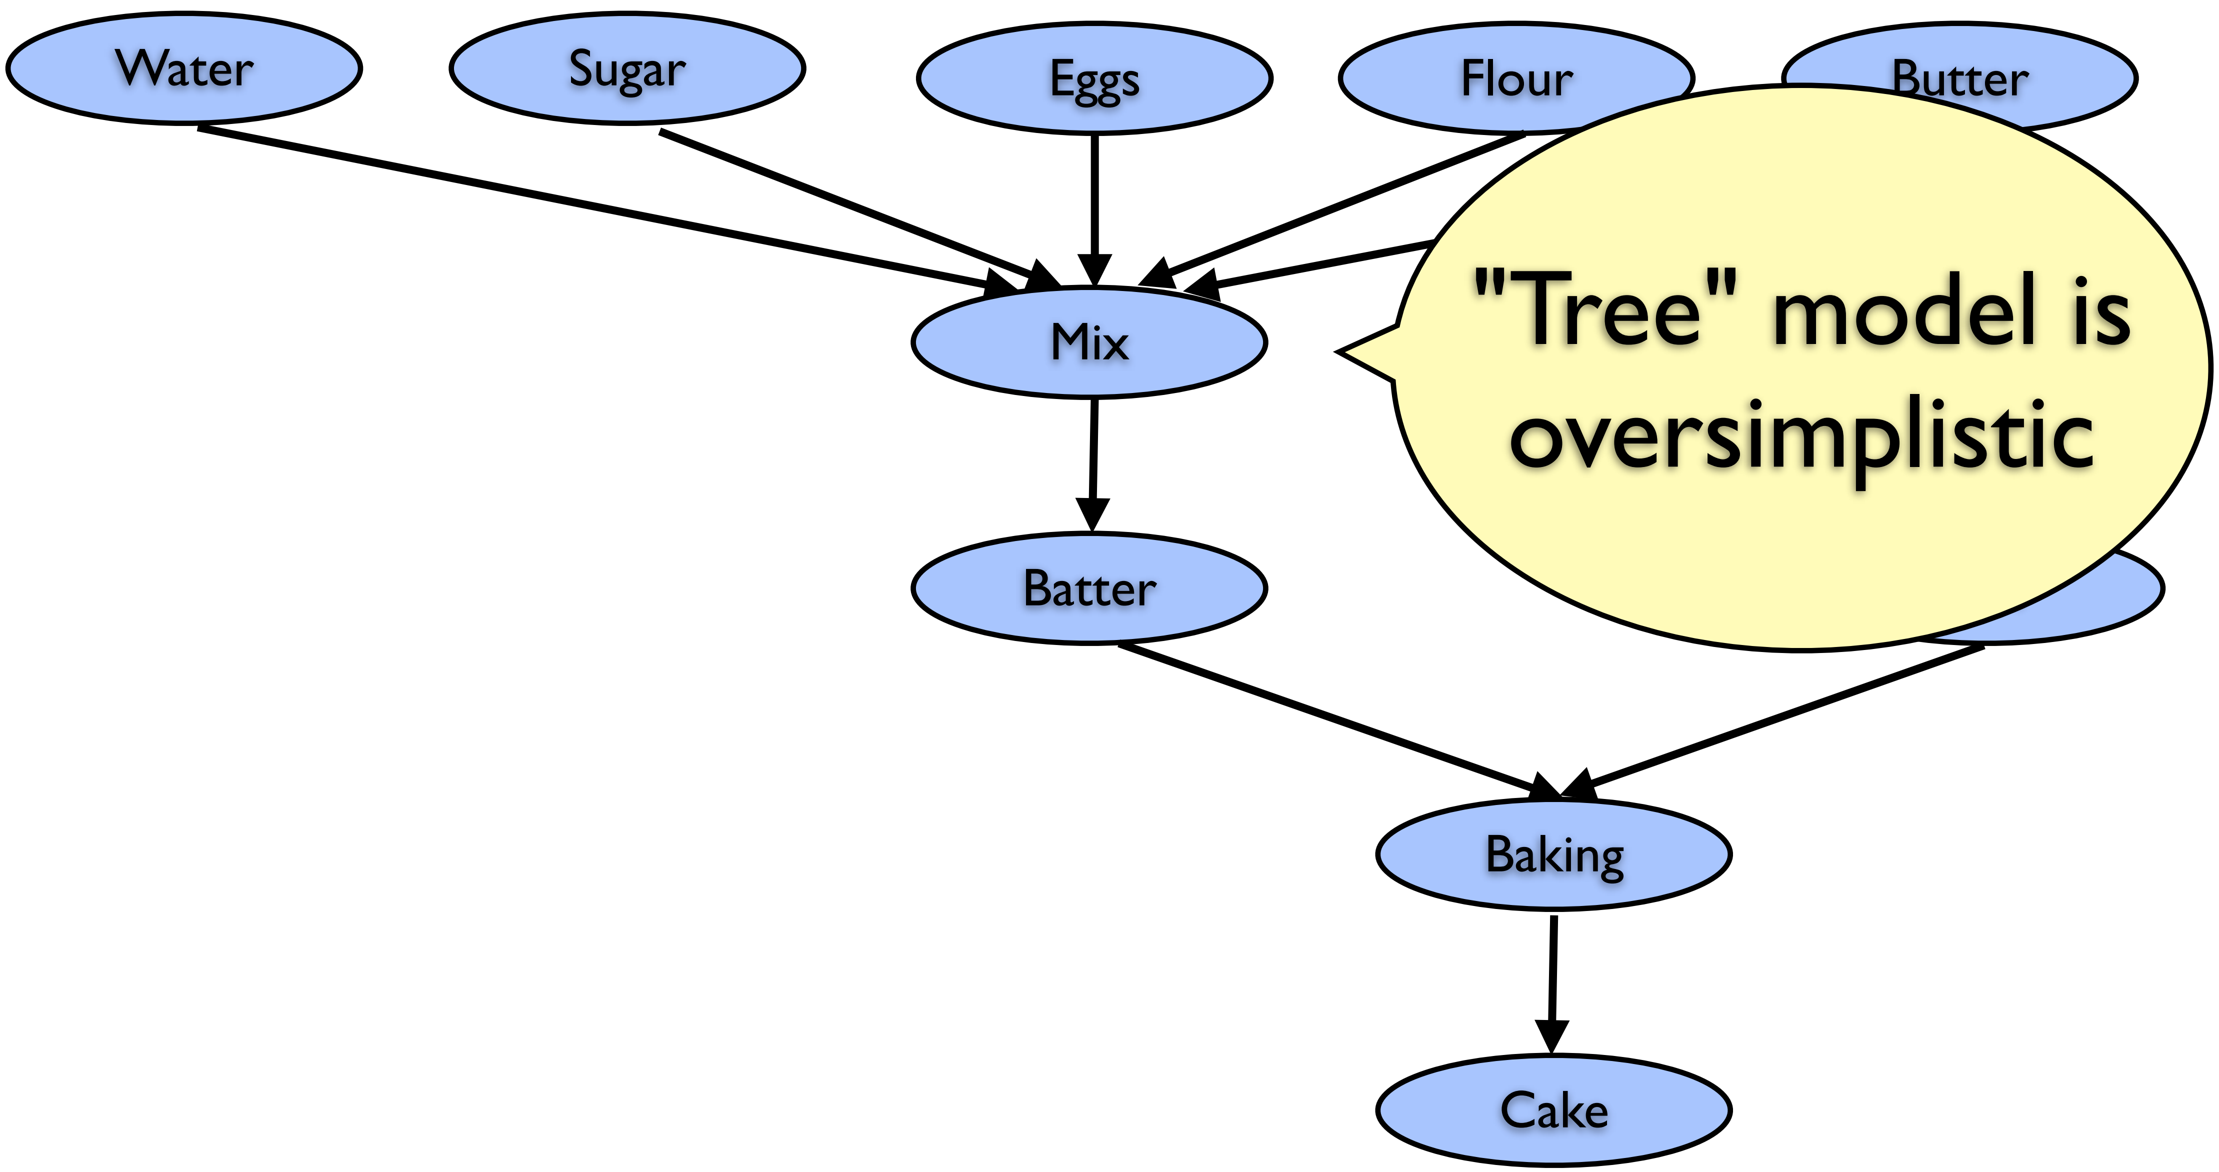

Icing on the cake

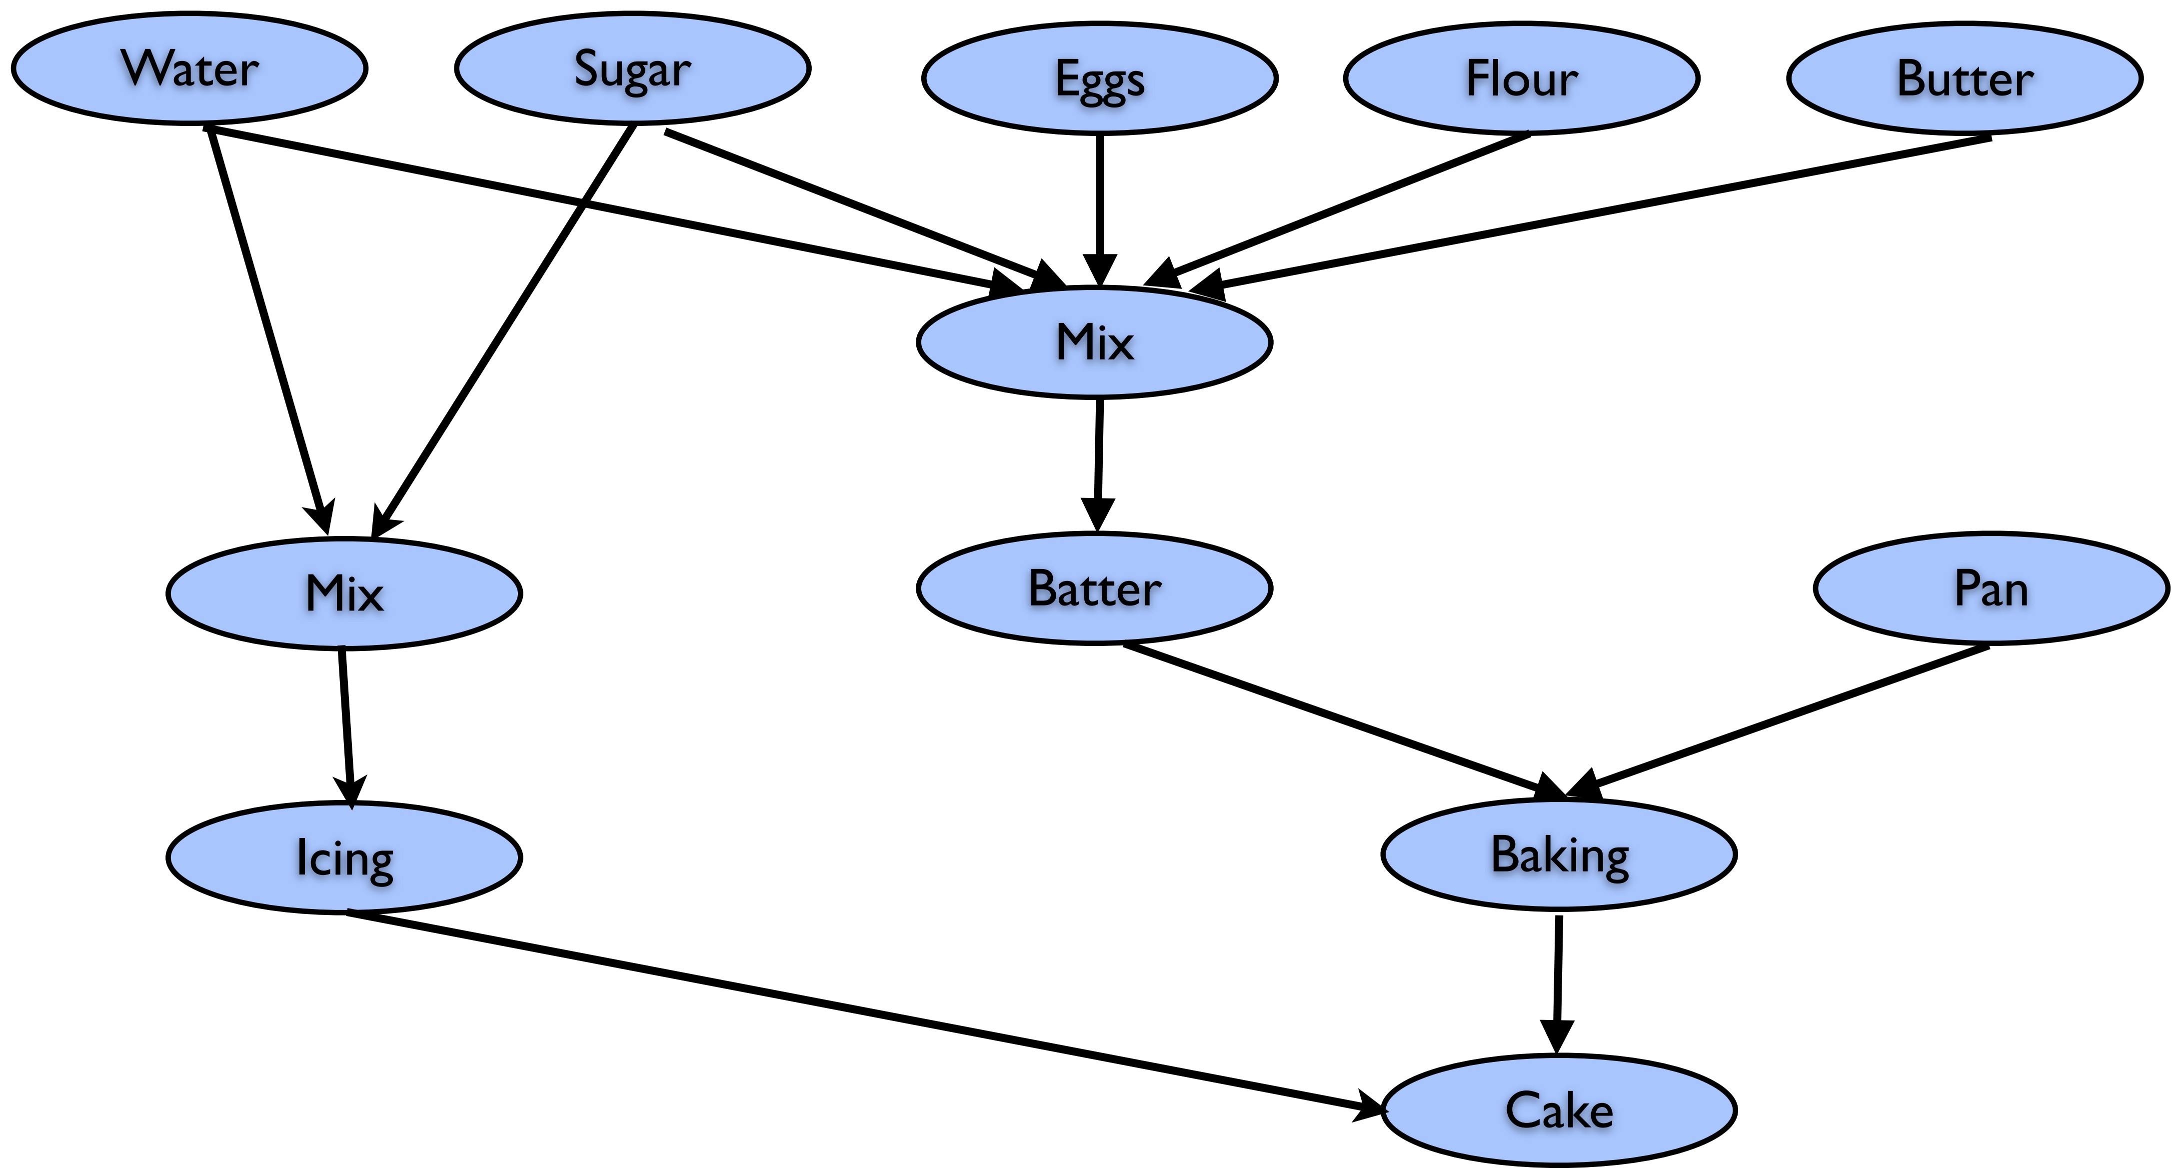

Icing on the cake

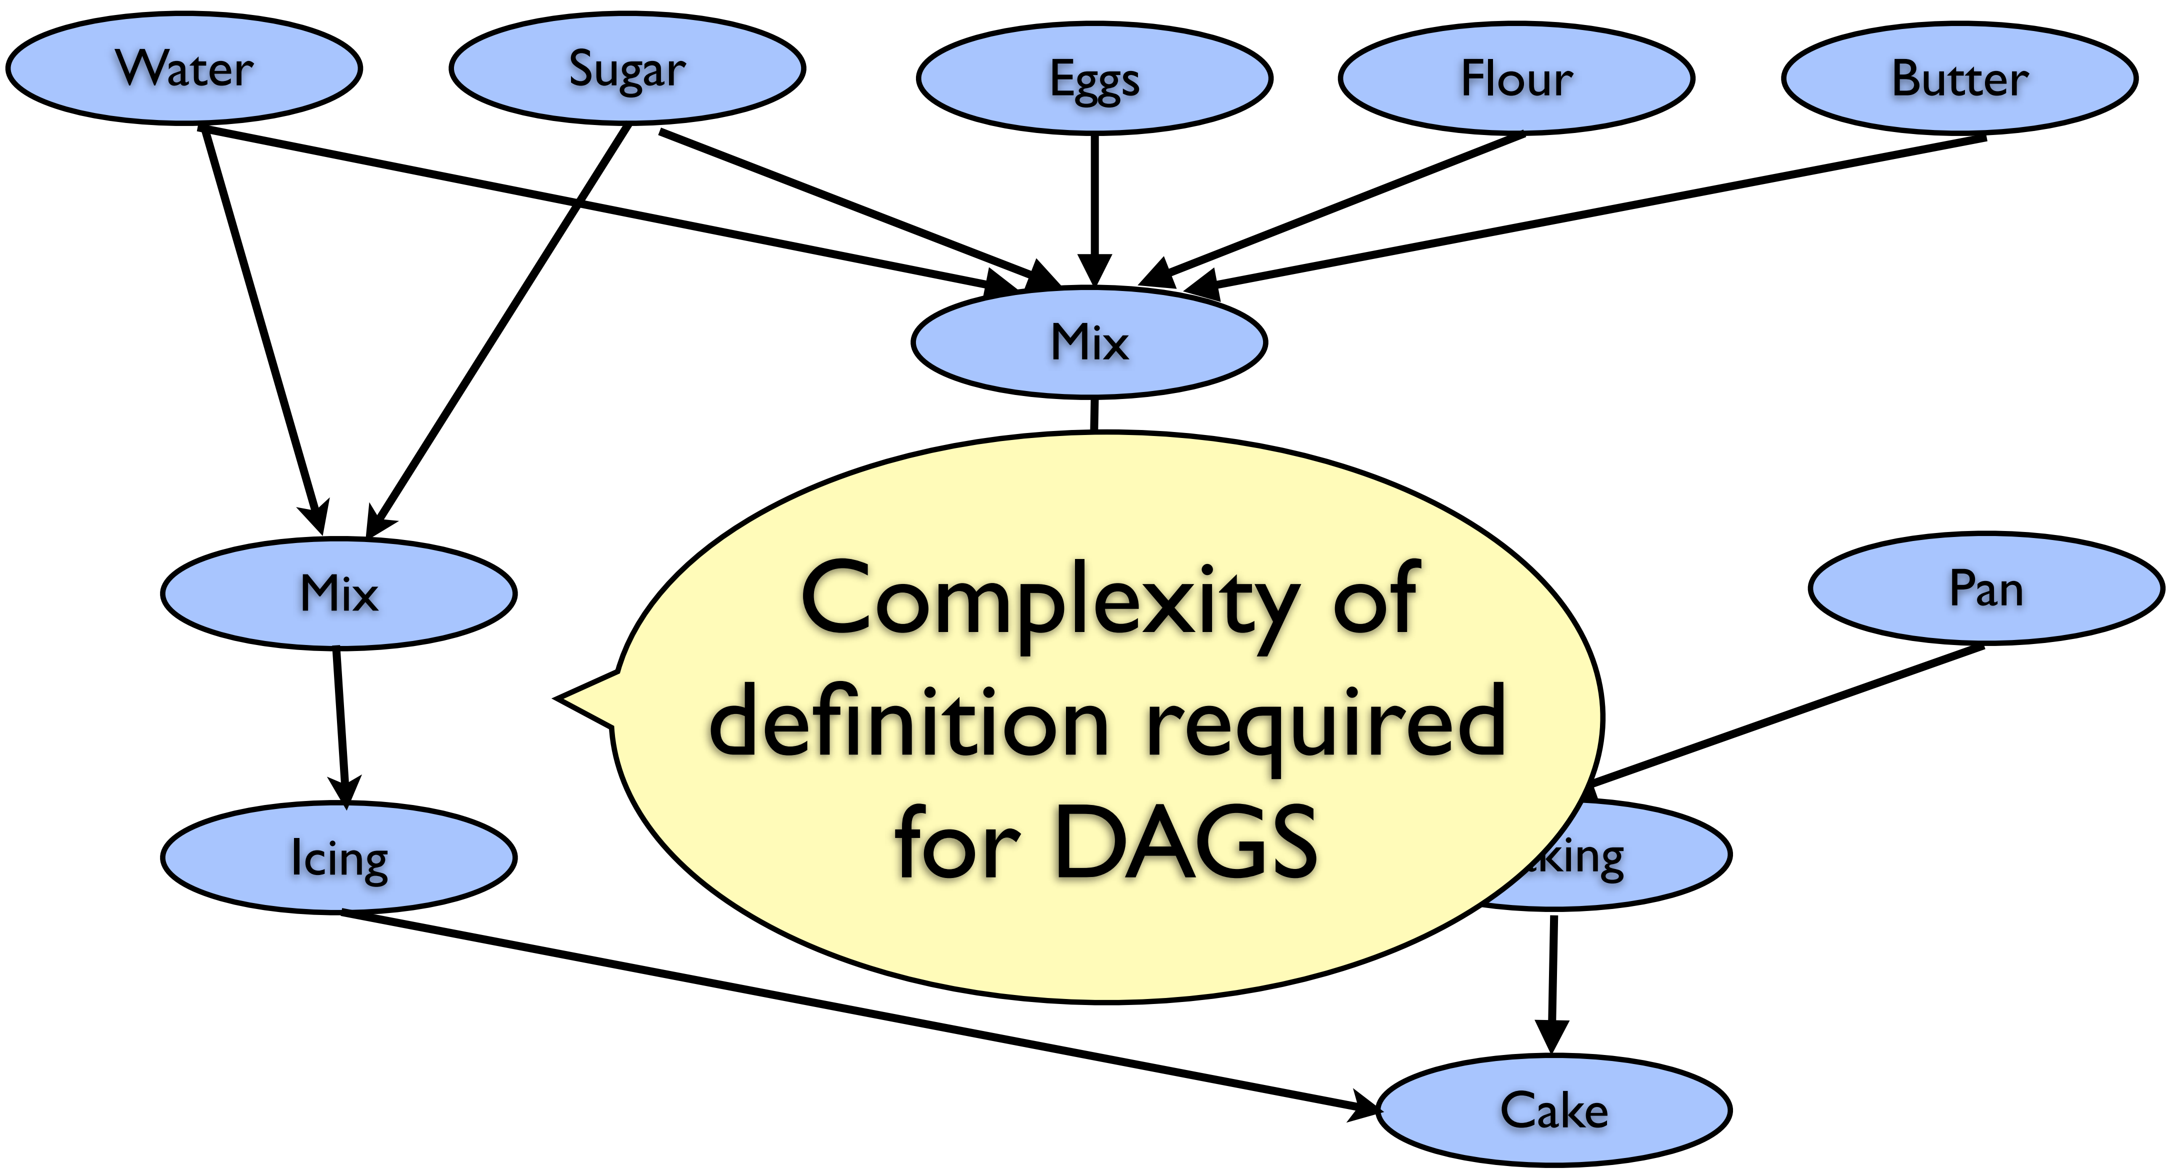

Causal interpretations of provenance graphs

- Seems appropriate to interpret provenance as causal models
- Or, more generally, **causal functions**
  - functions that support both ordinary evaluation and **intervention**
- Definition of actual cause, etc. can be formulated abstractly using causal functions

Causal interpretations vs. functional interpretations

- Causal interpretations are "richer"
  - more internal structure
  - built-in definition of "actual cause"
- Still not perfect
  - Causal models can't easily model hypotheticals that change "structure" of process
  - Ideas from Bayes nets, causal literature may help
    - but I'm way out of my depth there.

Inference over provenance graphs

- Datalog-style rules

$$
\begin{array}{rcl}
x \text{ wasDerivedFrom } y & :- & x \text{ wasGeneratedBy } p \wedge p \text{ used } y \\
p \text{ wasTriggeredBy } q & :- & p \text{ used } x \wedge x \text{ wasGeneratedBy } q \\
x \text{ wasDerivedFrom}^{+} \, y & :- & x \text{ wasDerivedFrom } y \vee (x \text{ wasDerivedFrom } z \wedge z \text{ wasDerivedFrom}^{+} \, y) \\
p \text{ wasTriggeredBy}^{+} \, q & :- & p \text{ wasTriggeredBy } q \vee (p \text{ wasTriggeredBy } r \wedge r \text{ wasTriggeredBy}^{+} \, q)
\end{array}
$$

- These talk about "syntax" (edges) of provenance graph, not "semantics"
  - (OPM does not specify **any** interpretation)

Some prior complexity results

- Eiter-Lukasiewicz (2002, 2004, 2006):
  - determining "actual cause" relationships is (at least) NP-hard
  - can be PTIME for simple classes of models
- Hence, deciding whether a "used" edge is really an actual cause is nontrivial.
- Transitive inference rules for provenance graphs complete, but not sound

Future work

- Really, this is a first step
  - provenance of "straight-line code"
- Many other possible approaches to provenance semantics
  - e.g. modeling linearity/resource sensitive situations
  - e.g. stream/concurrent programming models

Conclusions

- There is a **strong analogy** between causal models and provenance graphs
  - Provenance graphs can be interpreted fruitfully as "causal functions"
  - Edges in graph are not always "actual cause" relationships in particular situations
- Further study needed to understand how to represent provenance of richer computational models