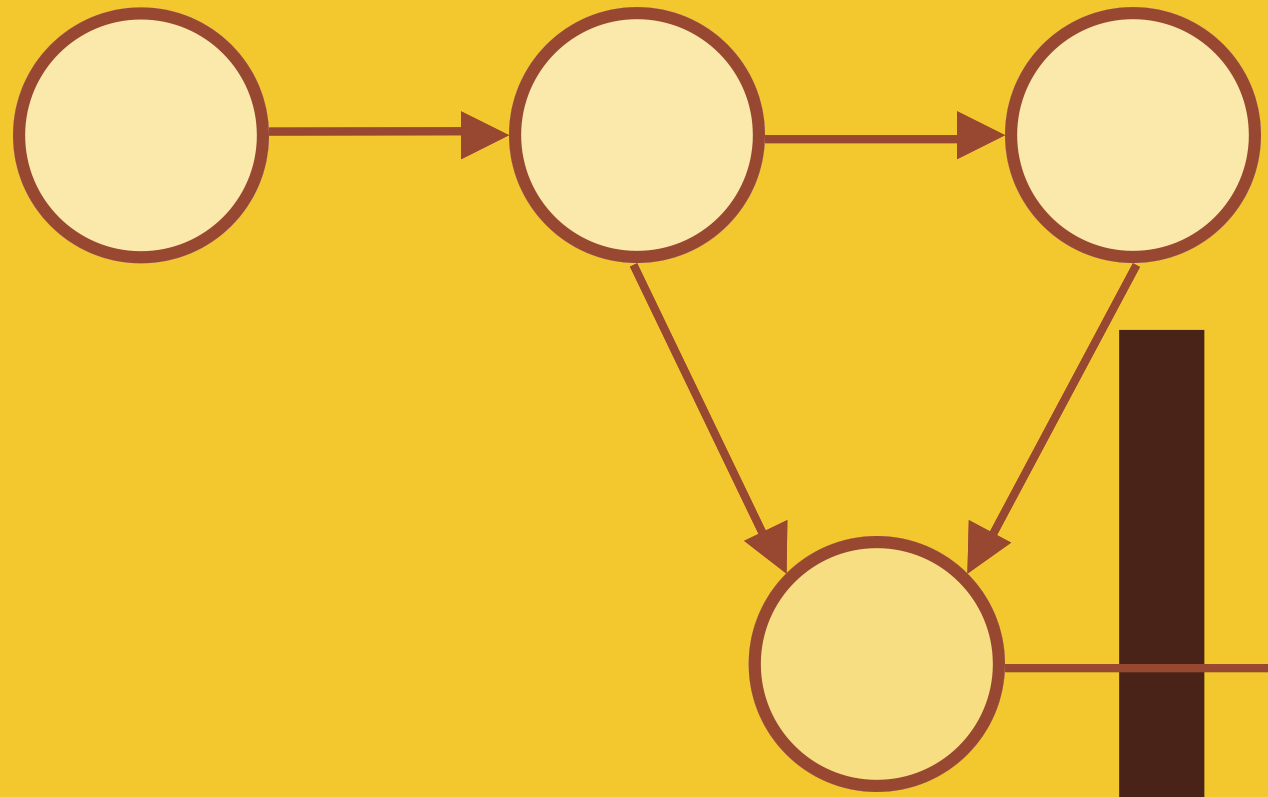

AAAI-16 Tutorial

CP-nets

- Thomas E. Allen
- Judy Goldsmith
- Francesca Rossi

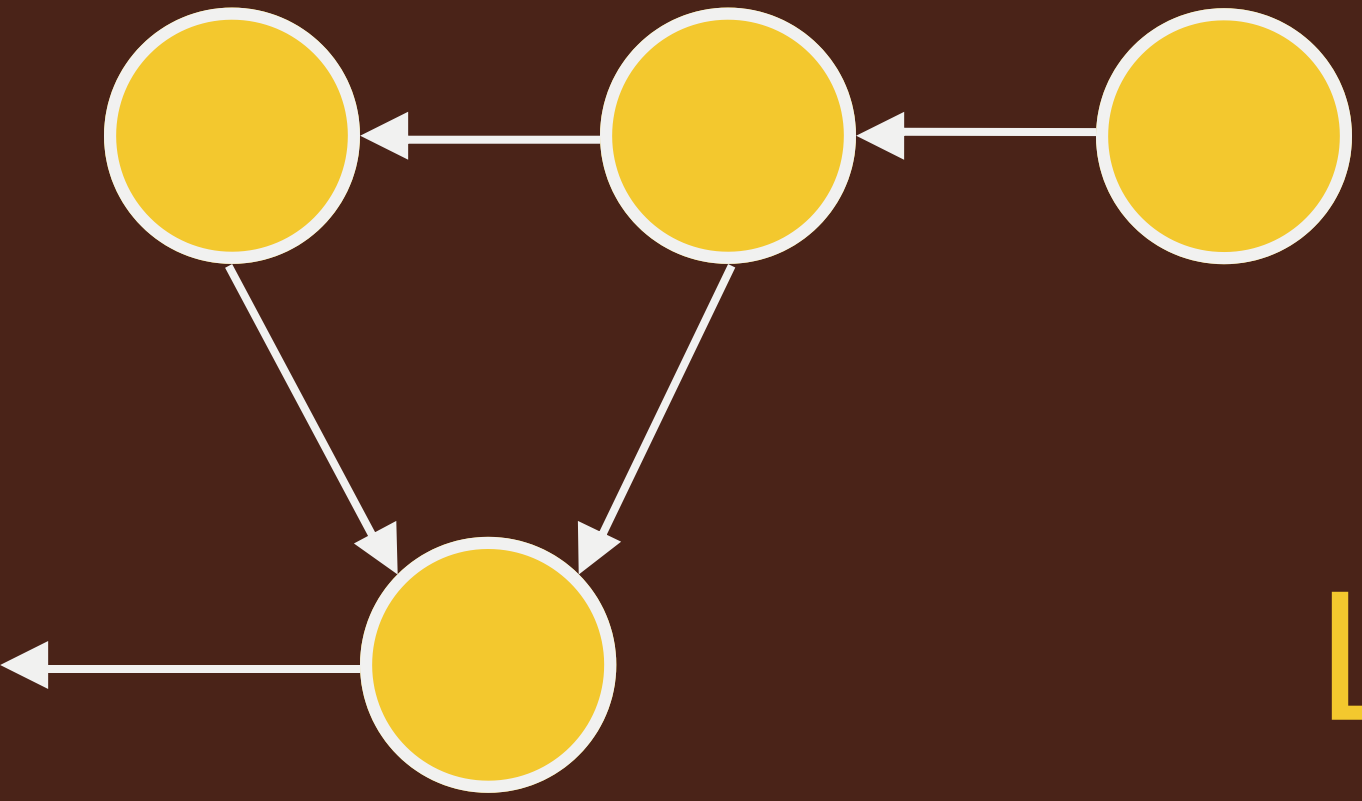

LEARNING

Thomas E. Allen

What do we mean by *learning* a CP-net?

- U User (individual subject)
- O Outcomes (set of objects)
- $\succ_U$ Preferences of the user over outcomes
- $\mathcal{N}$ Formalism (set of CP-nets as models for $\succ_U$)

Objective: *Obtain the best (or a satisfactory) model $N \in \mathcal{N}$ of the user's preferences $\succ_U$.*

Example:

Joseph prefers chocolate

- User: Joseph
- Outcomes: Desserts
  - *One binary feature: chocolate*
- Preference: chocolate $\succ$ $\overline{\text{chocolate}}$
- Available models:

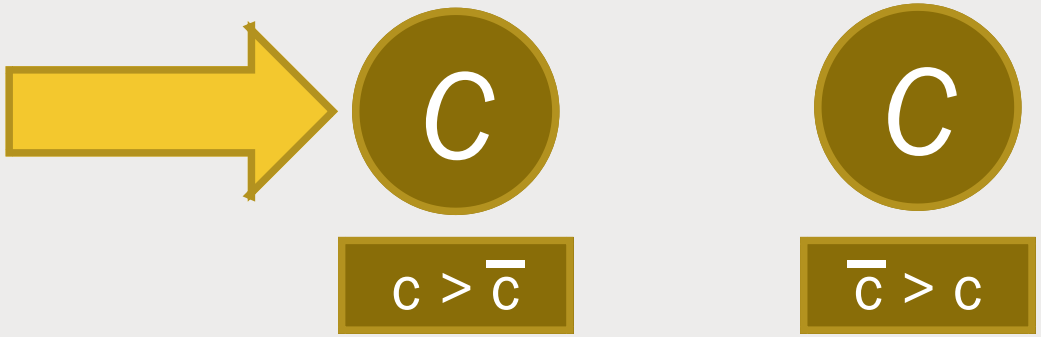

- Learning: Find best model

What if we introduce a second feature?

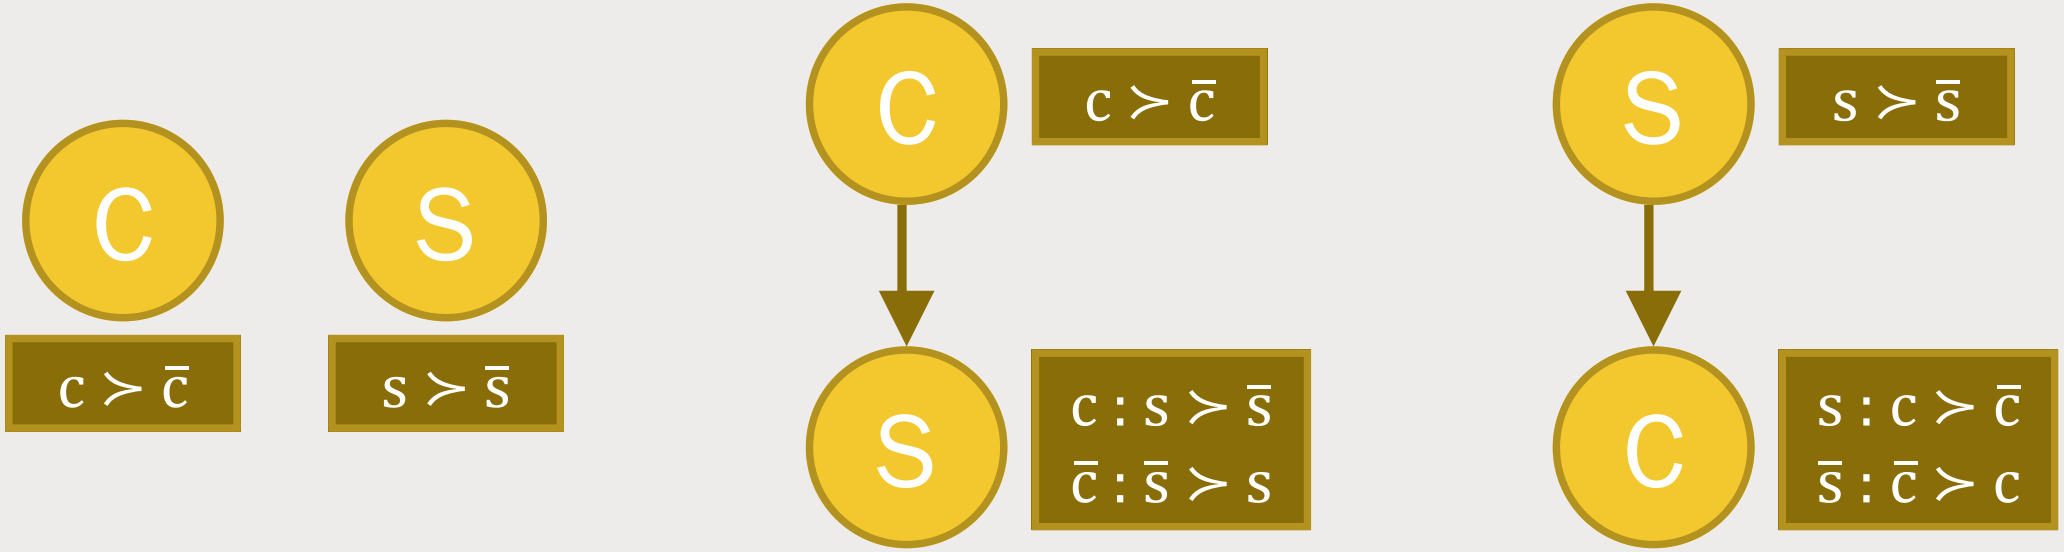

3 dependency graphs

(Labeled directed acyclic graphs)

What if we introduce a second feature?

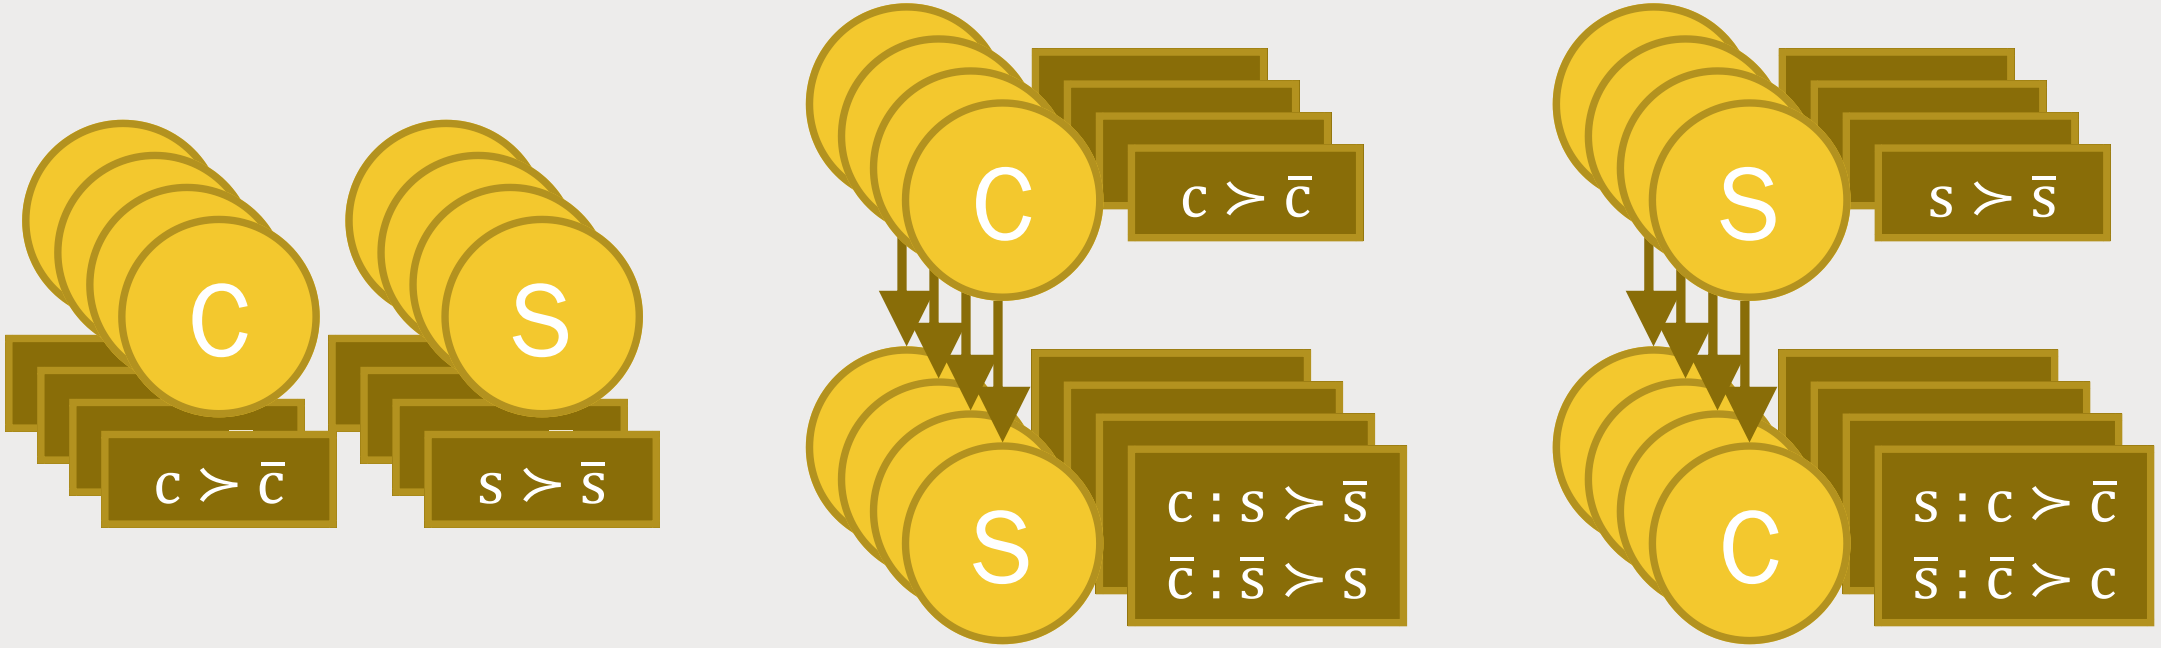

12 CP-nets

(Directed acyclic graphs labeled with non-degenerate CPTs)

Problem: There are many models

n binary features	$\lvert \mathcal{N} \rvert$ number of possible CP-nets
1	2
2	12
3	488
4	481776
5	157549032992 (1.6×10^{11})
6	4059976627283664056256 (4.1×10^{21})

How can we obtain such a model $N \in \mathcal{N}$?

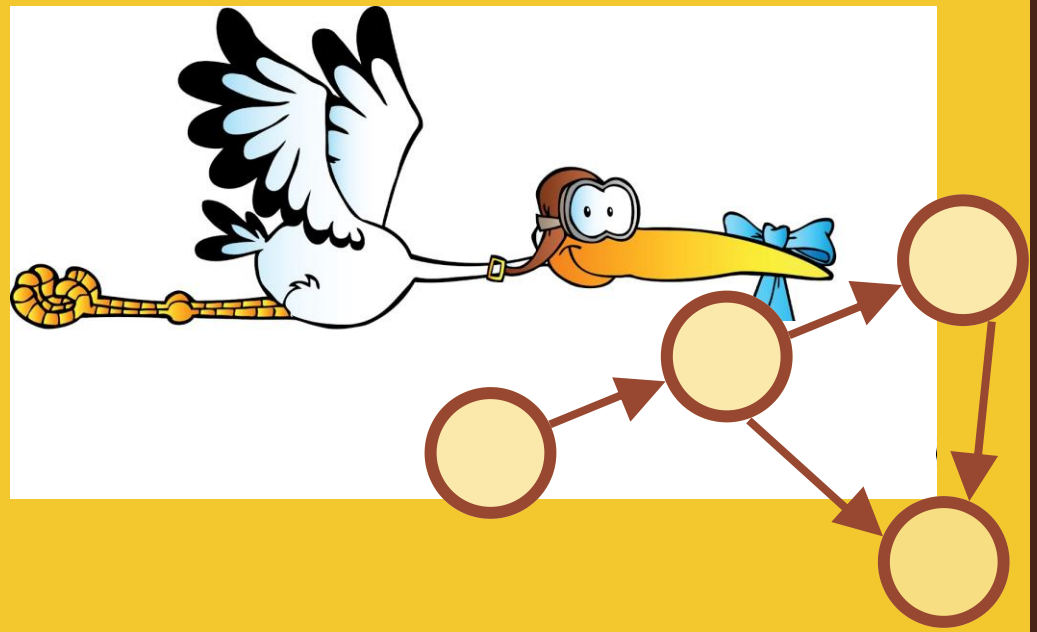

- Directly through introspection
- Elicit through queries to the user (active learning)
- Learn passively from observational data (passive learning)

Why can't we just ask people what they want?

- The original idea of eliciting CP-nets (Boutilier et al., 2004)
- It is often claimed that CP-nets are "easy to elicit"

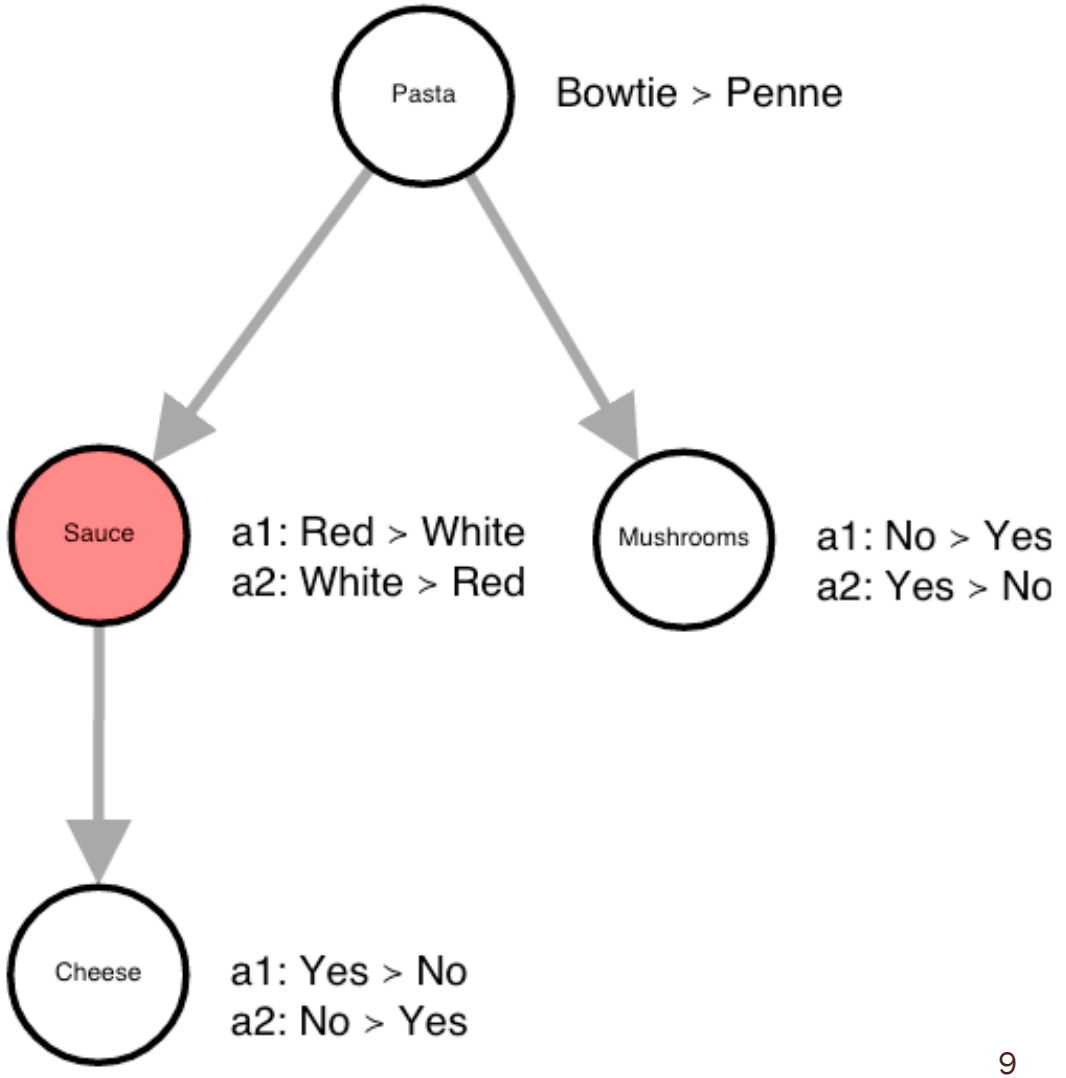

Alex Morgan

Anthony Davis

Ma Long

Gracie Gold

Draw a CP-net that expresses what you prefer

"We would like you to draw a map of how you made your decisions during the experiment. Please fill in the template diagram and tables the same way we explained for sandwiches."

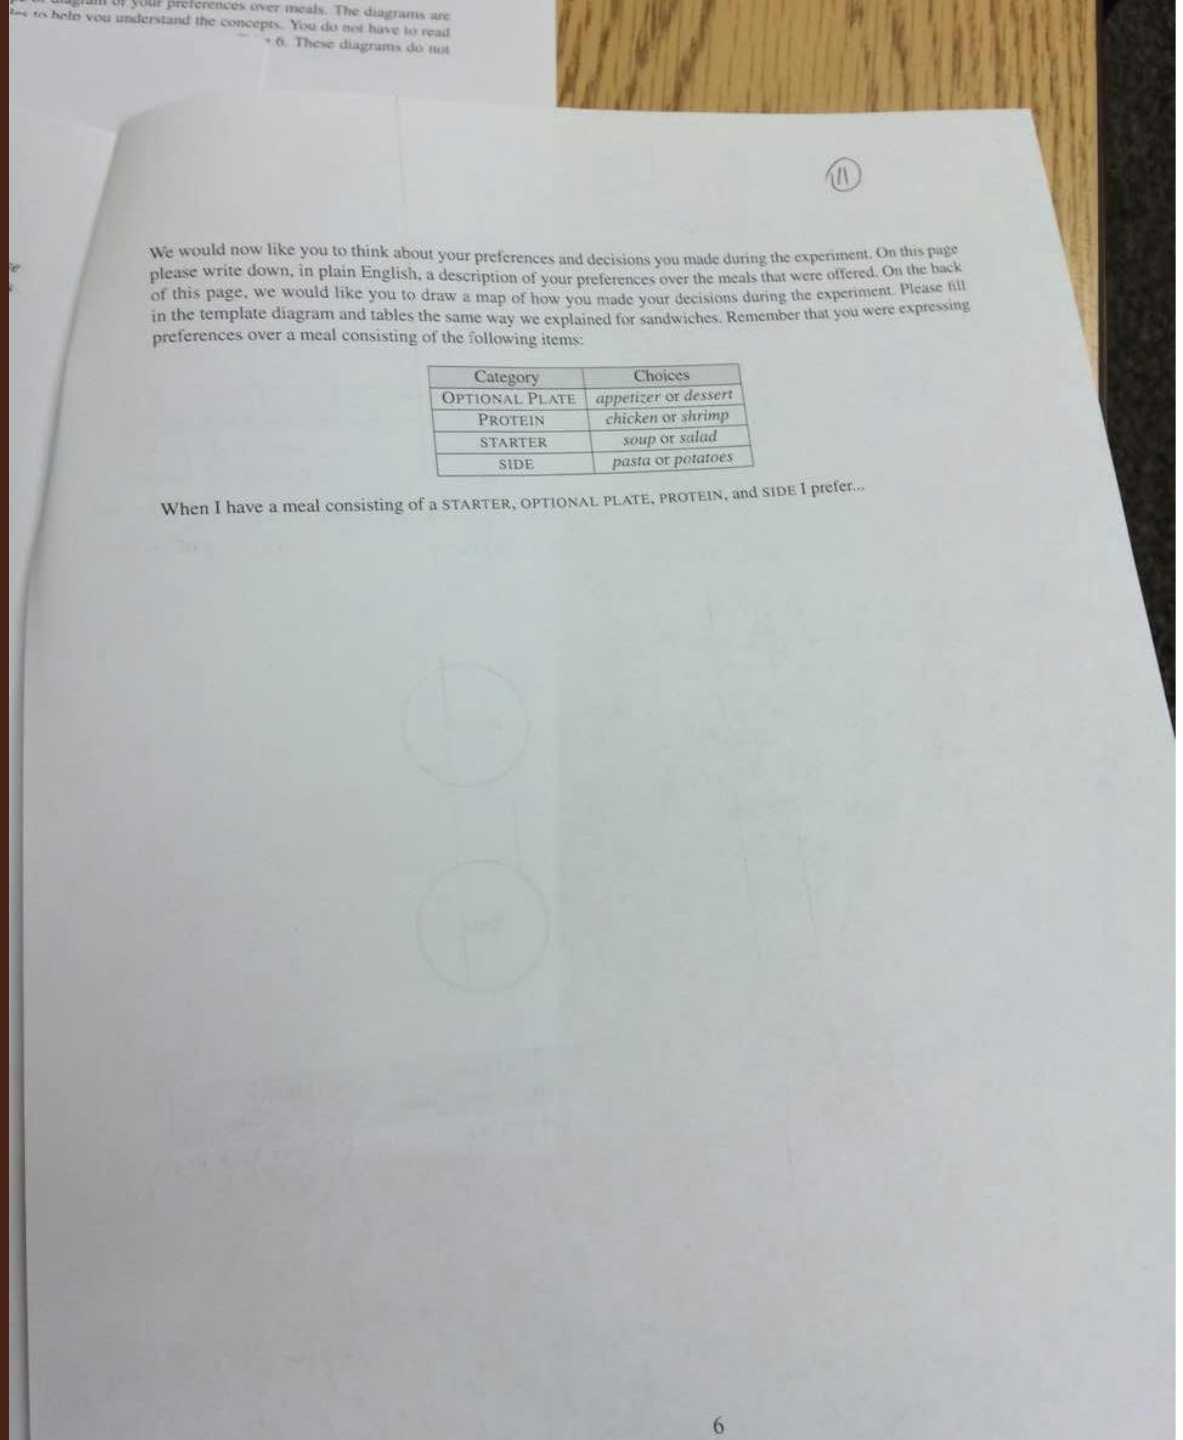

Draw a CP-net that expresses what you prefer

"We would like you to draw a map of how you made your decisions during the experiment. Please fill in the template diagram and tables the same way we explained for sandwiches."

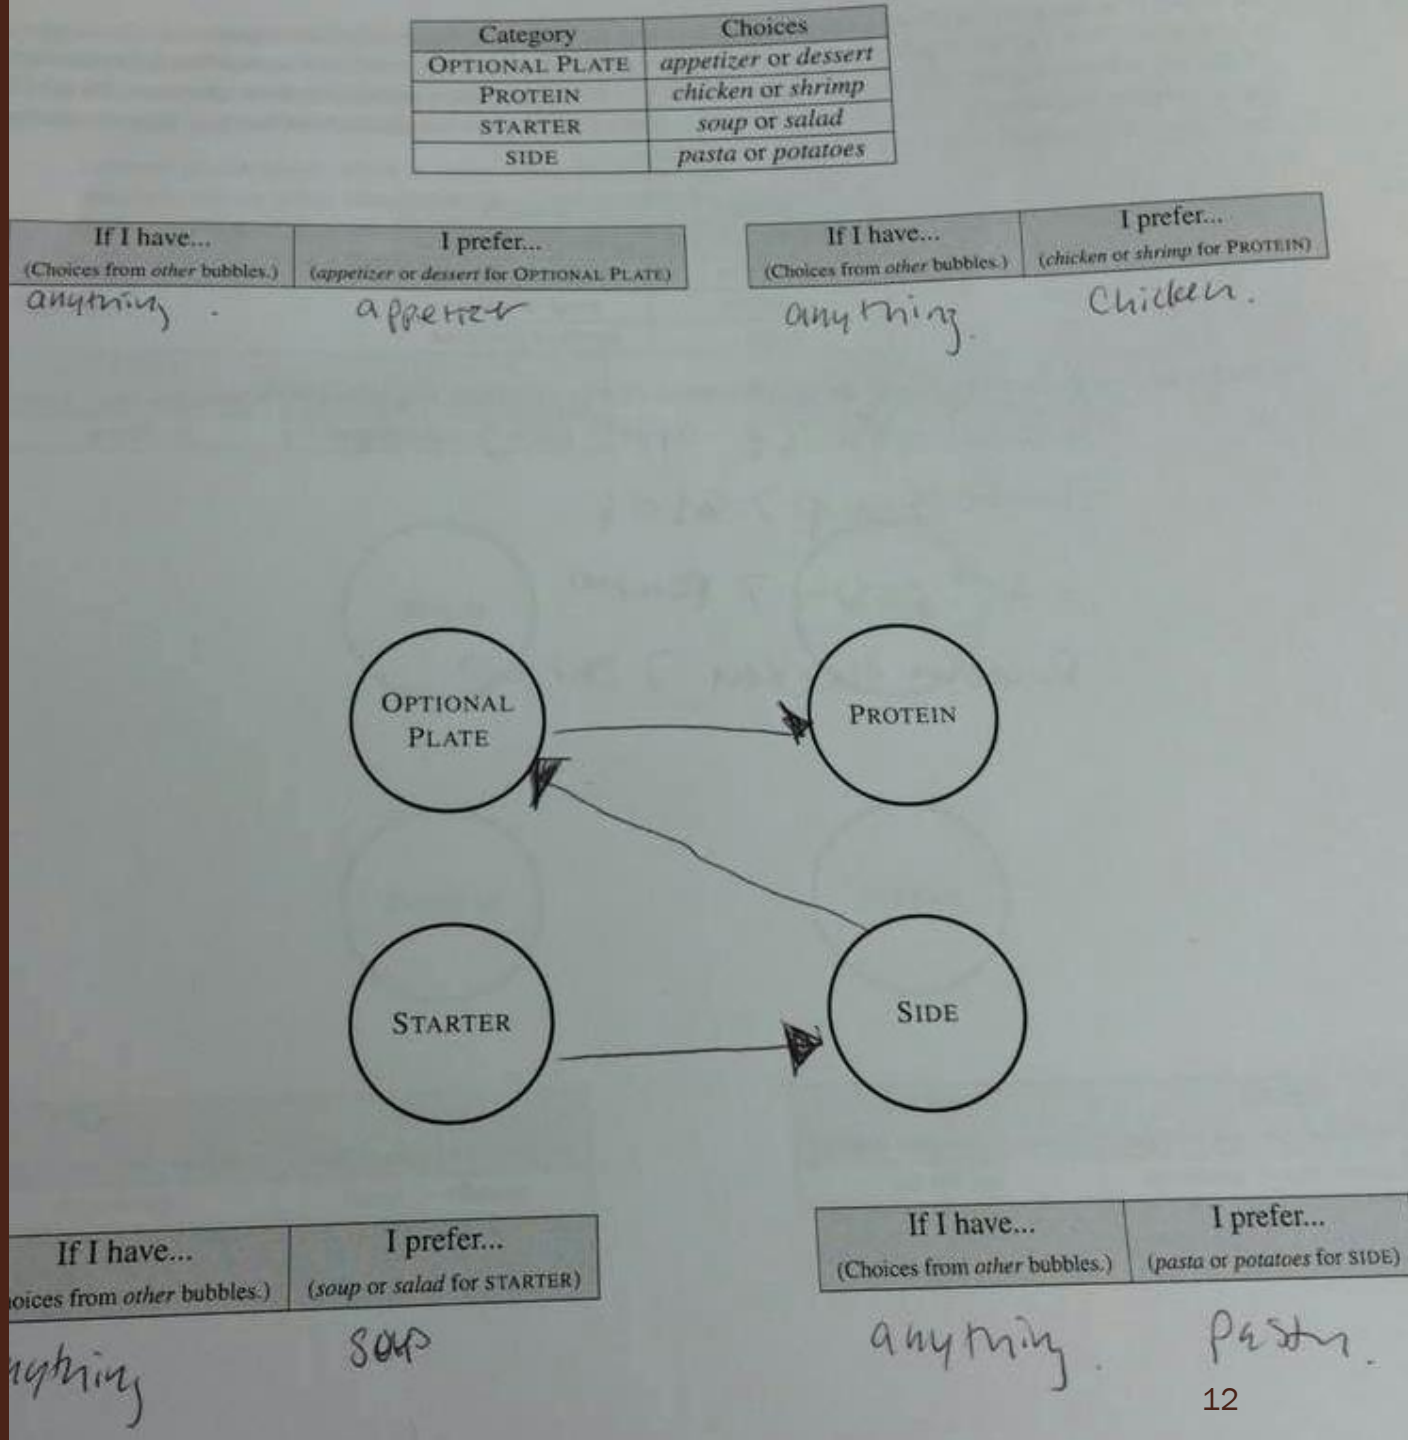

Draw a CP-net that expresses what you prefer

"We would like you to draw a map of how you made your decisions during the experiment. Please fill in the template diagram and tables the same way we explained for sandwiches."

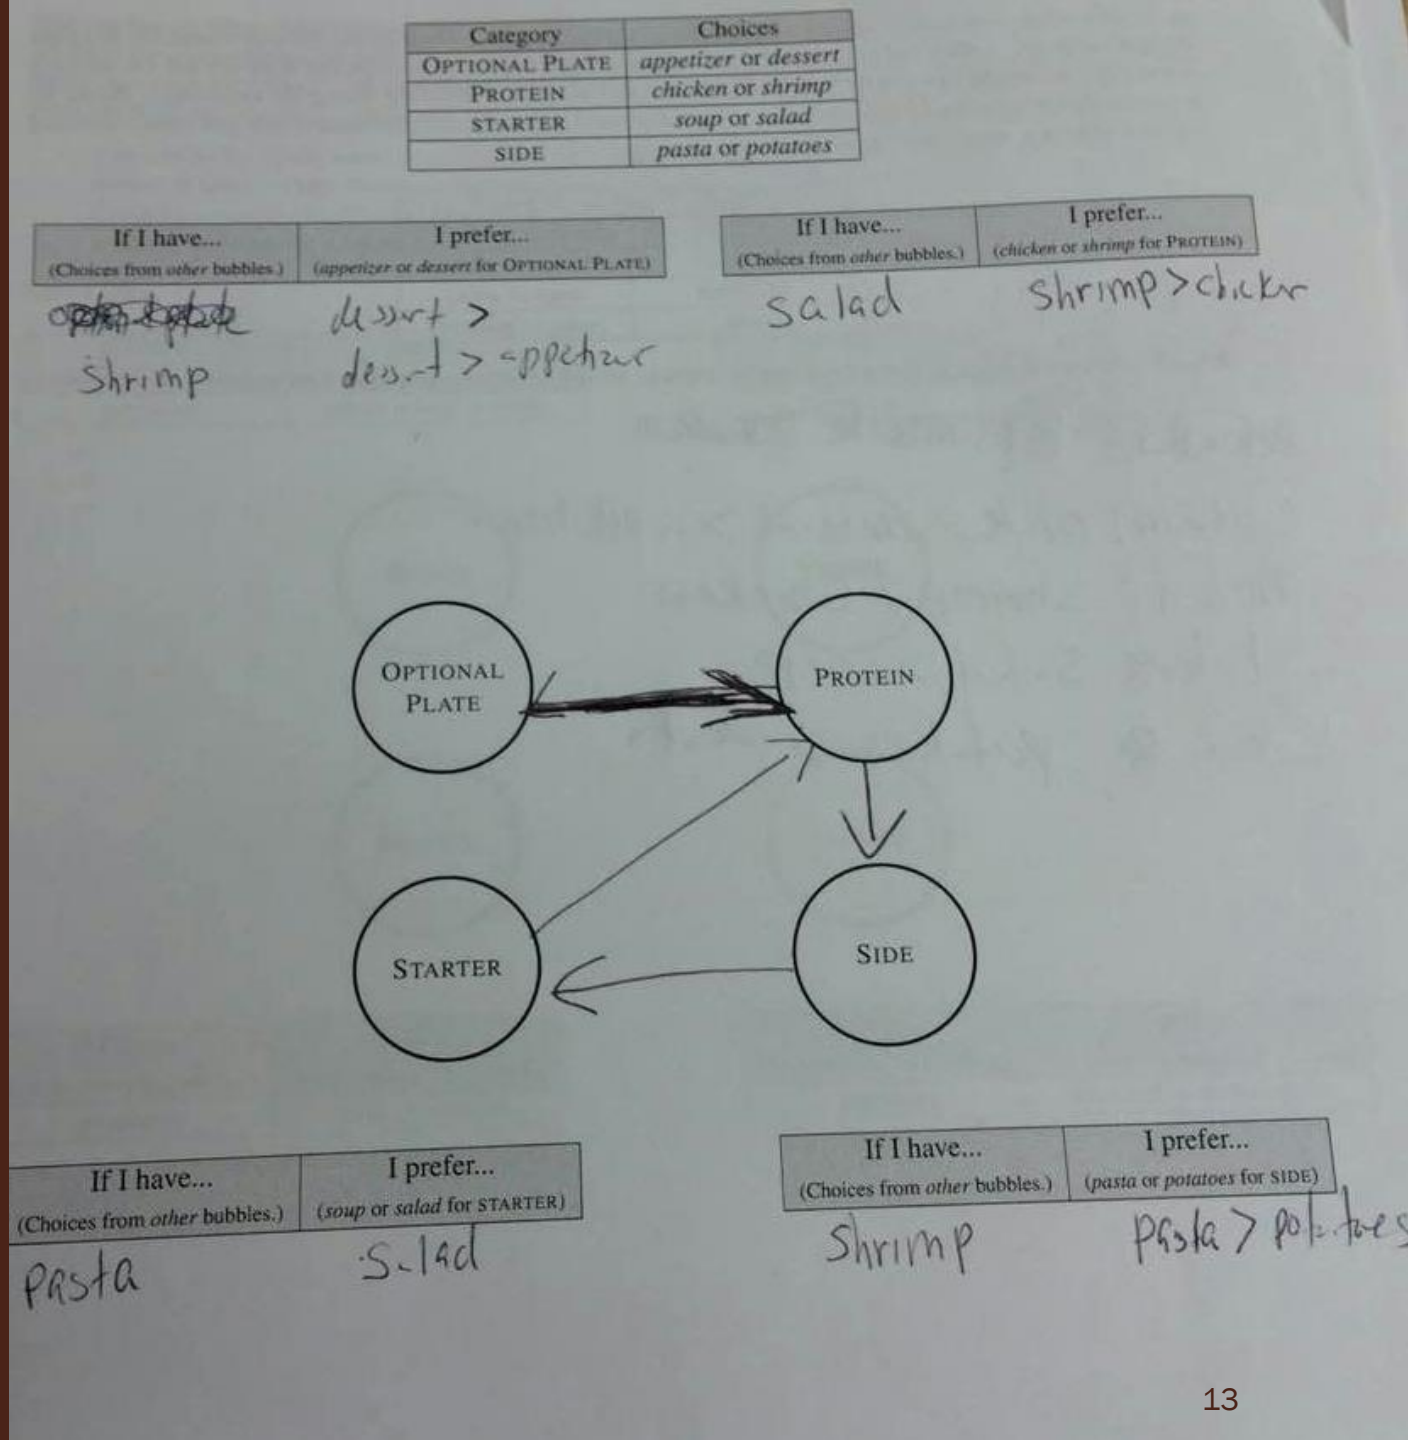

Draw a CP-net that expresses what you prefer

"We would like you to draw a map of how you made your decisions during the experiment. Please fill in the template diagram and tables the same way we explained for sandwiches."

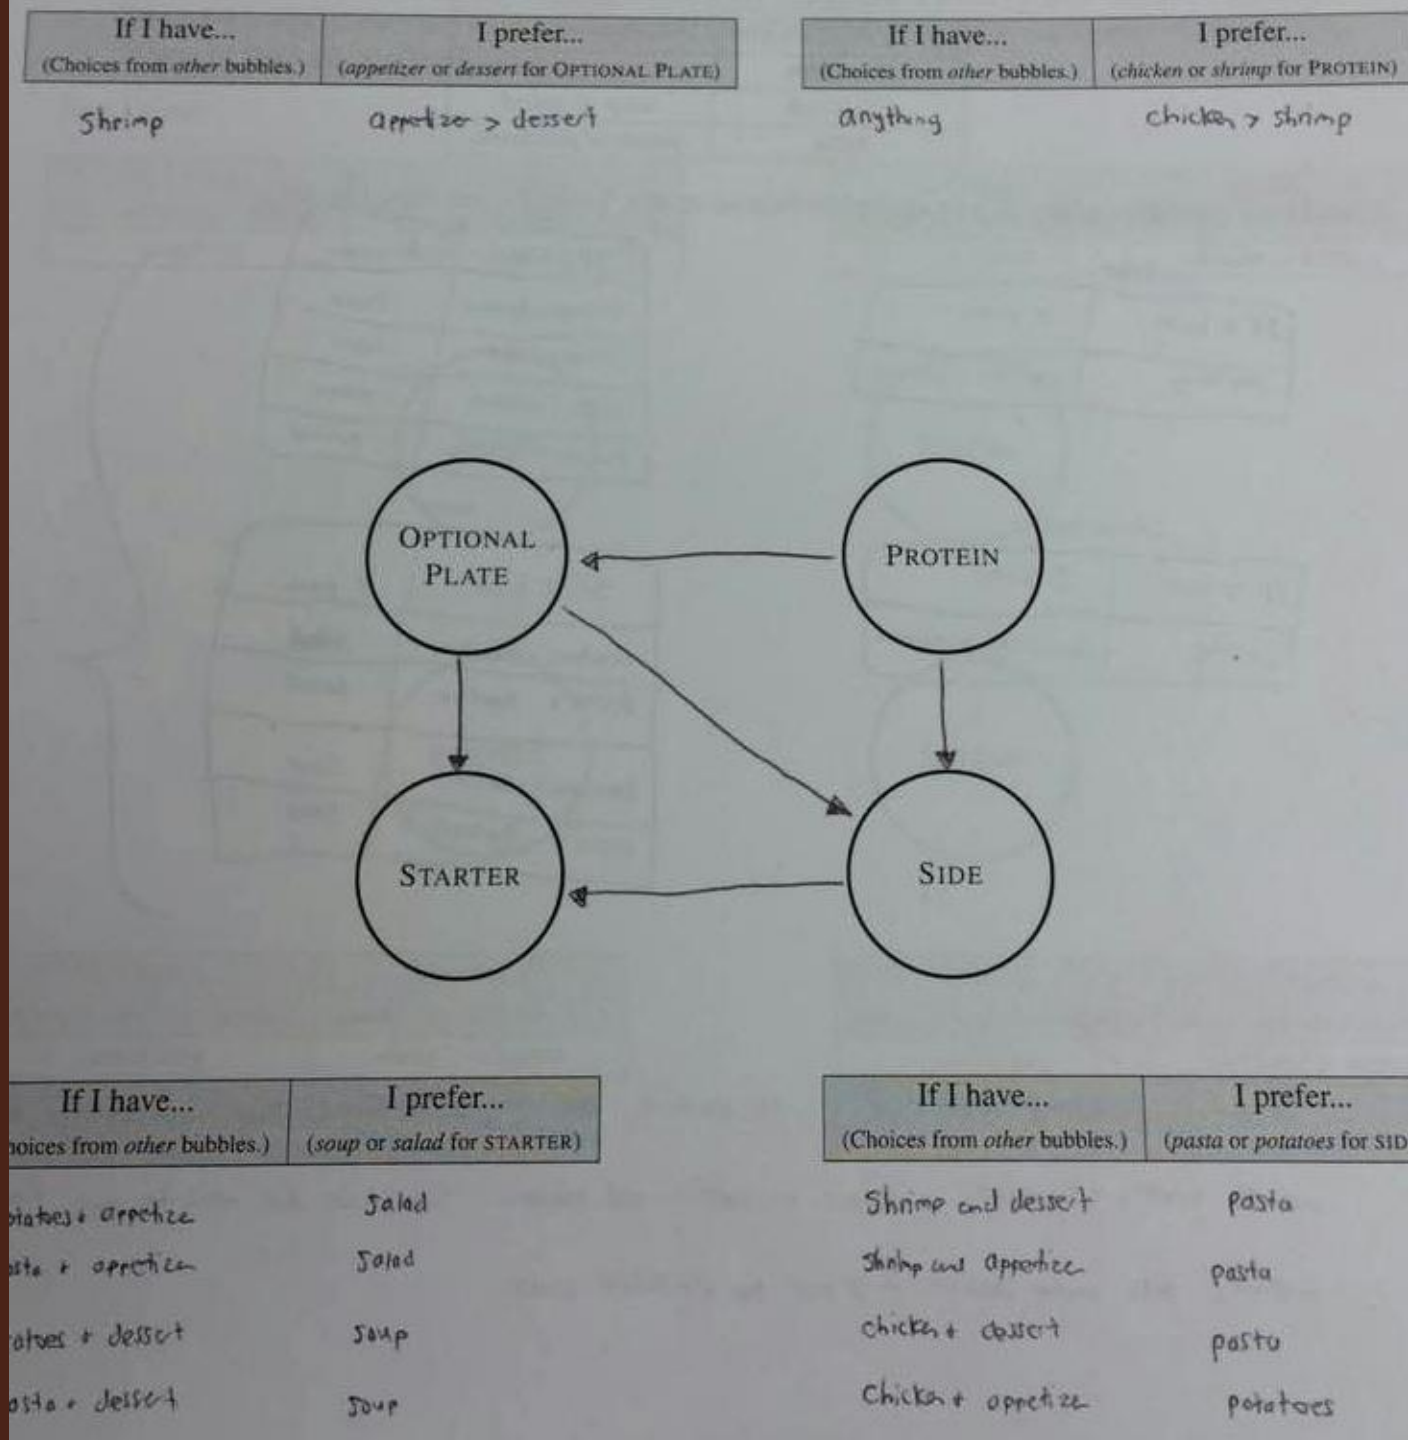

How can we obtain such a model $N \in \mathcal{N}$?

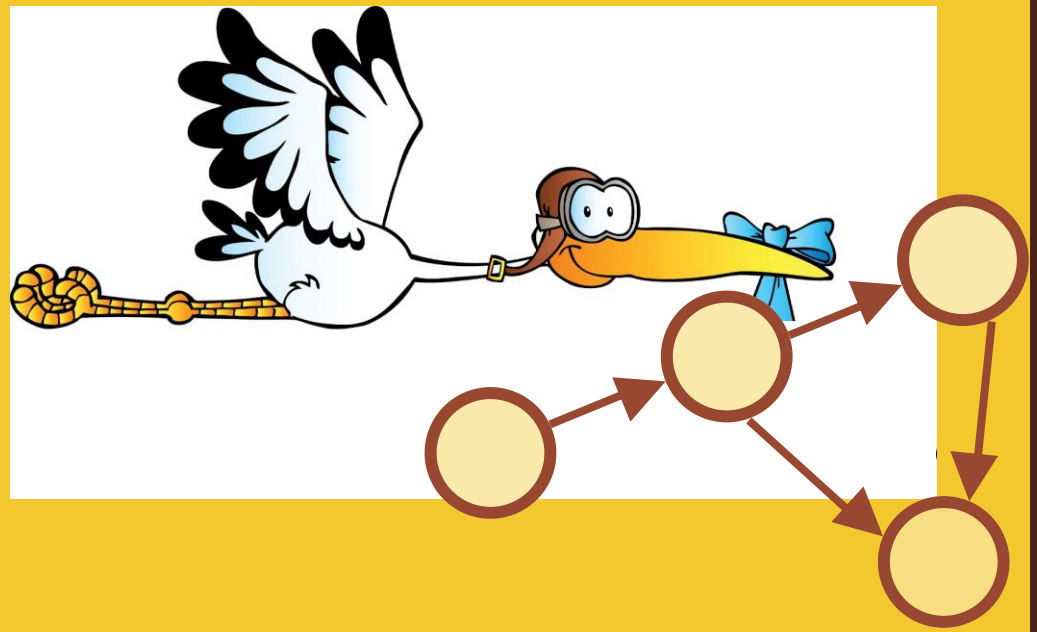

- Directly through introspection
- Elicit through queries to the user (active learning)
- Learn passively from observational data (passive learning)

Data from which to learn a CP-net

- Optimal outcomes
  - *Of all outcomes, which is best?*
- Outcome comparisons
  - *Arbitrary pairs of outcomes*
  - *Swap queries (differ in just one feature)*
- Other (e.g., introspection)

Distance of Preference Orders

- Recall that our that our objective is to obtain a best model $N \in \mathcal{N}$ of the user's preferences $\succ_U$.
- This implies some measure of **distance** between $\succ_U$ and $\succ_N$ the order the entailed by the CP-net *N*
- Problem: Find argument *N* that minimizes the distance between $\succ_U$ and $\succ_N$.

Question: How would you suggest defining this distance?

Distance of Preference Orders

- $\{ 0, \infty \}$
  - *Assumes preferences are consistent with a CP-net*
  - *Success if we recover that CP-net, else fail*
- How many outcomes are out of order? (Mallows model)
  - *But CP-nets induce a partial order*
- Compare complete transitive relations
- Compare induced preference graphs
- Edit distance for CPTs and graph

Distance of Preference Orders

Joe's true preference

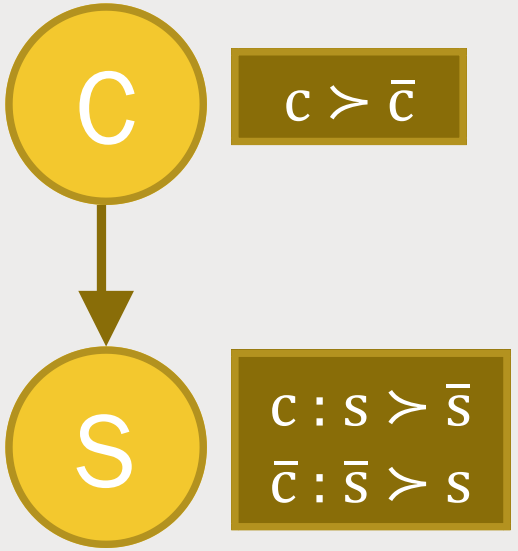

CP-net Model

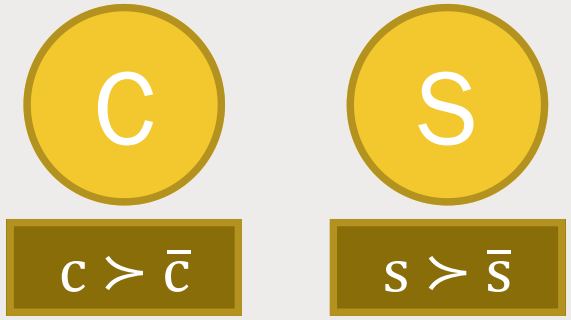

$\succ_U$	cs	$c\bar{s}$	$\bar{c}s$	$\bar{c}\bar{s}$
cs	0	1	1	1
$c\bar{s}$	0	0	1	1
$\bar{c}s$	0	0	0	0
$\bar{c}\bar{s}$	0	0	1	0

$\succ_N$	cs	$c\bar{s}$	$\bar{c}s$	$\bar{c}\bar{s}$
cs	0	1	1	0
$c\bar{s}$	0	0	0	1
$\bar{c}s$	0	0	0	1
$\bar{c}\bar{s}$	0	0	0	0

Distance of Preference Orders

$\succ_U$	cs	$c\bar{s}$	$\bar{c}s$	$\bar{c}\bar{s}$
cs	0	1	1	**1**
$c\bar{s}$	0	0	**1**	1
$\bar{c}s$	0	0	0	**0**
$\bar{c}\bar{s}$	0	0	**1**	0

$\oplus$

$\succ_N$	cs	$c\bar{s}$	$\bar{c}s$	$\bar{c}\bar{s}$
cs	0	1	1	**0**
$c\bar{s}$	0	0	**0**	1
$\bar{c}s$	0	0	0	**1**
$\bar{c}\bar{s}$	0	0	**0**	0

$=$

0	0	0	1
0	0	1	1
0	0	0	0
0	0	1	0

$$\frac{4}{16} = 0.25$$

Distance of Preference Orders

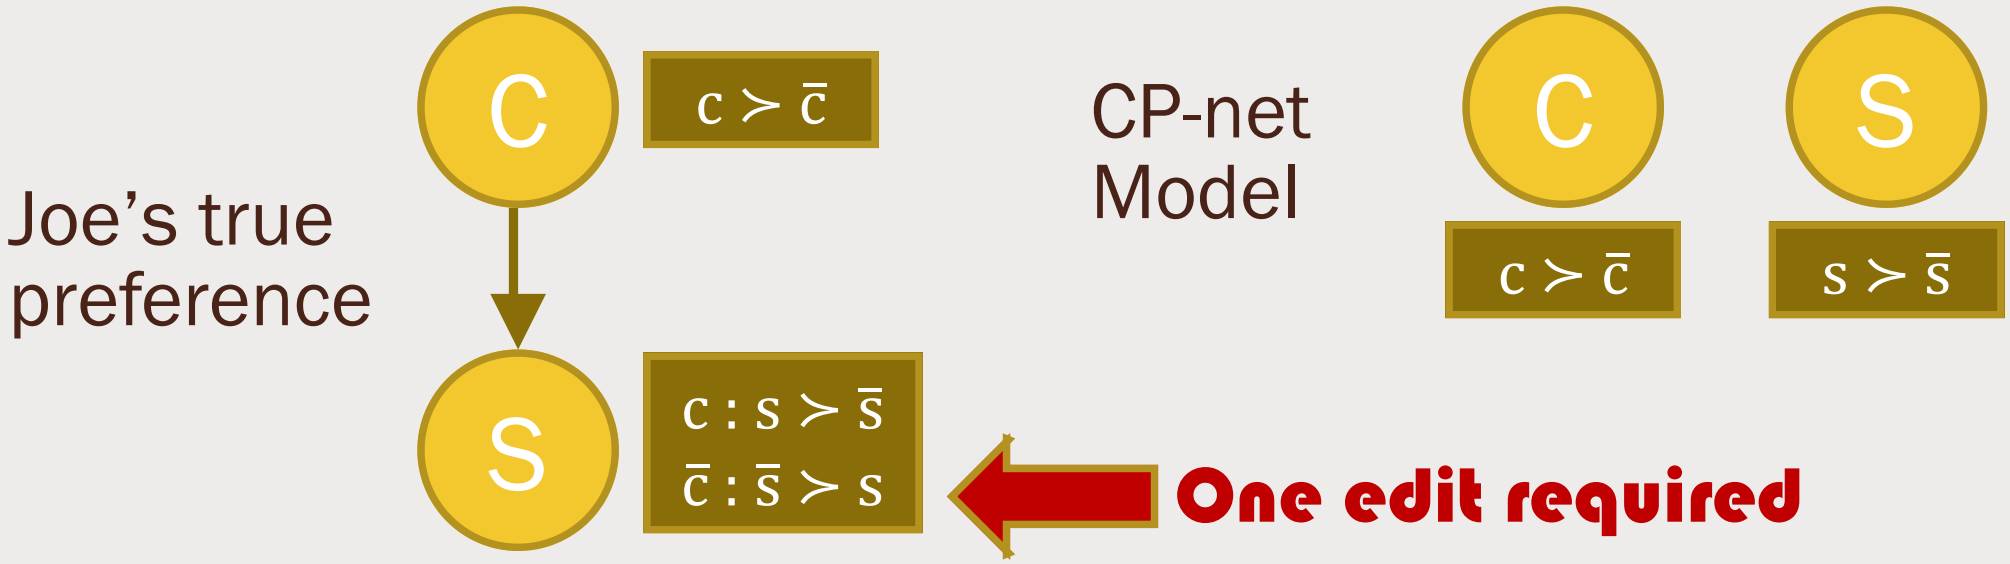

Learning CP-nets: a Brief History

- Athienitou and Dimopoulos (M-PREF 2007)
  - *Idea: Learn CP-net that* **entails** *all comparisons*
- Lang and Mengin (NMR 2008, IJCAI 2009)
  - **Consistency**, *not entailment (3 notions)*
  - ∃ ***separable*** *CP-net that* **entails** *all examples?* **(P-time)**
  - ∃ *separable CP-net that is* **weakly consistent***?* **(NP-C)**

Learning CP-nets: a Brief History

- Koriche and Zanuttini (IJCAI 2009, AIJ 2010)
  - **Eliciting** *a CP-net (active)*
  - **Attribute efficient algorithm**: *polynomially many queries*
  - *Adaptively generated* **swap** *queries*
  - *Logarithmically many queries if binary, tree-shaped*
- Dimopoulos, Michael, and Athienitou (IJCAI 2009)
  - *Learning binary acyclic CP-nets from data (passive)*
  - *Consistency not entailment;* **transparent entailment**
  - *PAC learner; 2-SAT reduction for obtaining CPTs*

Learning CP-nets: a Brief History

- Guerin, Allen, and Goldsmith (ADT 2013)
  - *Active learning*
  - *Arbitrary queries (not limited to swaps)*
  - *Attribute comparison queries*
- Allen (Allerton 2013)
  - *Extends Dimopoulos et al. (2009) to multi-valued CP-nets with indifference*
  - *Finding CPTs requires 3-SAT reduction rather than 2-SAT*

Learning CP-nets: a Brief History

Liu, et al. (TKDE 2012, KBS 2012)

- Passive learning in noisy environments: $o \succ o' \succ o'' \succ o$
- Distance: #edges that differ in **preference graphs**
- Nonlinear 0-1 program solved via branch-and-bound
  - *Preference graph*
- Chi-squared hypothesis testing

More Work Is Needed!

- Presently we have no CP-net databases
- Factored preference datasets (e.g., SUSHI) are also scarce
- Current proposals have limitations
  - *Introspection*
  - *Transparent entailment*
  - *Binary domains*
  - *Separable or tree-shaped CP-net graphs*
  - *Intractable*

Questions?

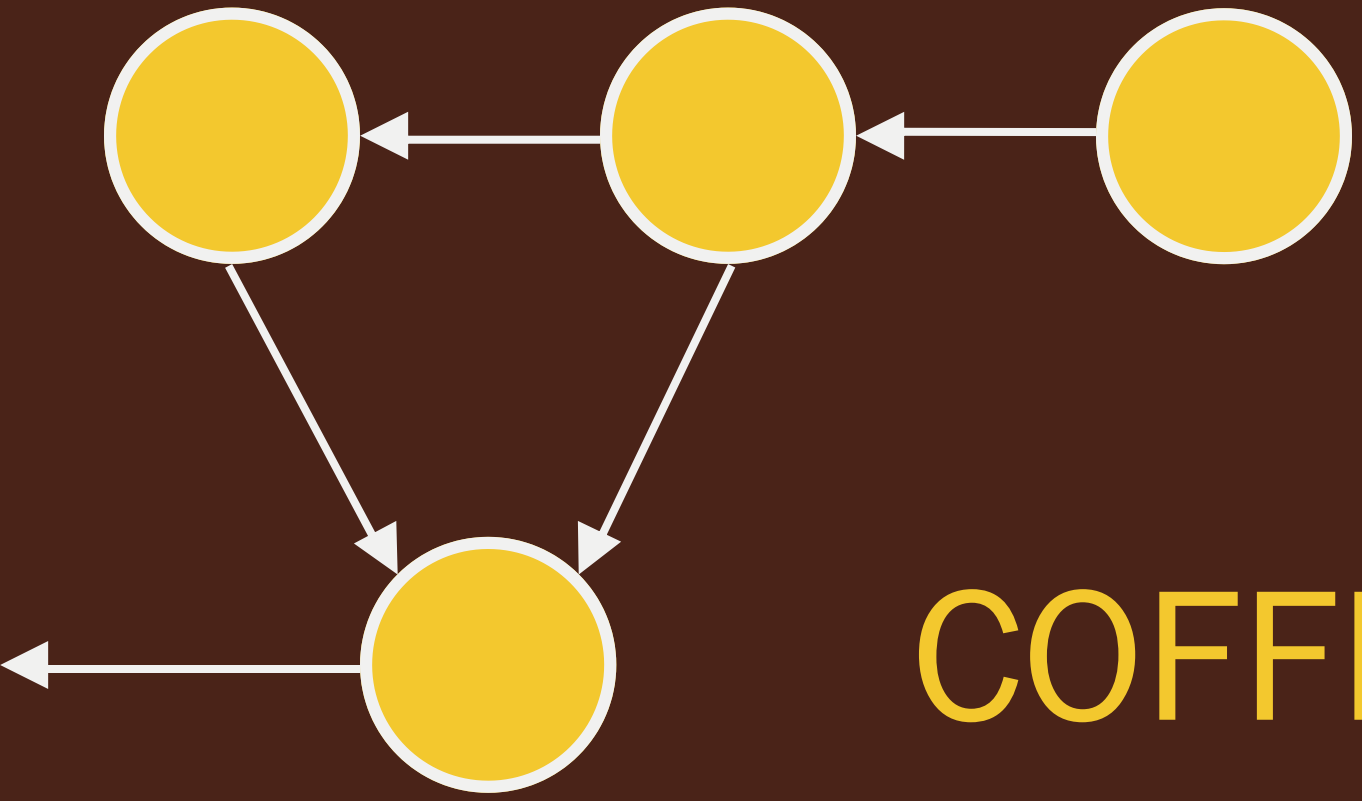

COFFEE BREAK