Representing Agents' Beliefs, Preferences and Decision-Making Processes

Avi Pfeffer

Harvard University

The Goal

- A knowledge representation language for representing the way agents think in games
  - beliefs about the game structure
  - different preferences
  - strategies followed
  - strategic reasoning about other agents

Classical Game Theory

- Assumes common knowledge of rationality
- Assumes common knowledge of game structure
- But real agents have
  - different opinions about game structure
  - different beliefs about preferences
  - follow heuristics or social conventions

Bayesian Games

- Each player has set of types
- Type associated with
  - utility function
  - probability distribution over other agents' types
- Highly expressive
  - different preferences
  - different opinions about game structure

Bayesian Games are not a Natural Solution

- Bayesian games are deficient from the point of view of knowledge representation
  - The whole story of the game must be folded into the utility functions
  - May be exponentially larger than necessary
    - probability distribution over types
    - *n* decisions with 2 preferences for each decision
  - Equilibrium requires best response

Networks of Influence Diagrams

- Network in which each node represents a mental model used to make decisions
- Each model is a multi-agent influence diagram
- Agents have beliefs about which models other agents use
- Models linked by beliefs: “In model U, agent X believes that Y uses model V”
- Agents may have probability distribution over which model used

Nodes are Mental Models

- Models may differ in
  - Conditional probabilities of variables
  - Influence diagram structure
  - Utility functions
  - Decision variables in one model may become chance variables in another model
    - corresponding to agents following prescribed patterns of play
  - Beliefs about which models other agents use

Network of Influence Diagrams Example: Main Model

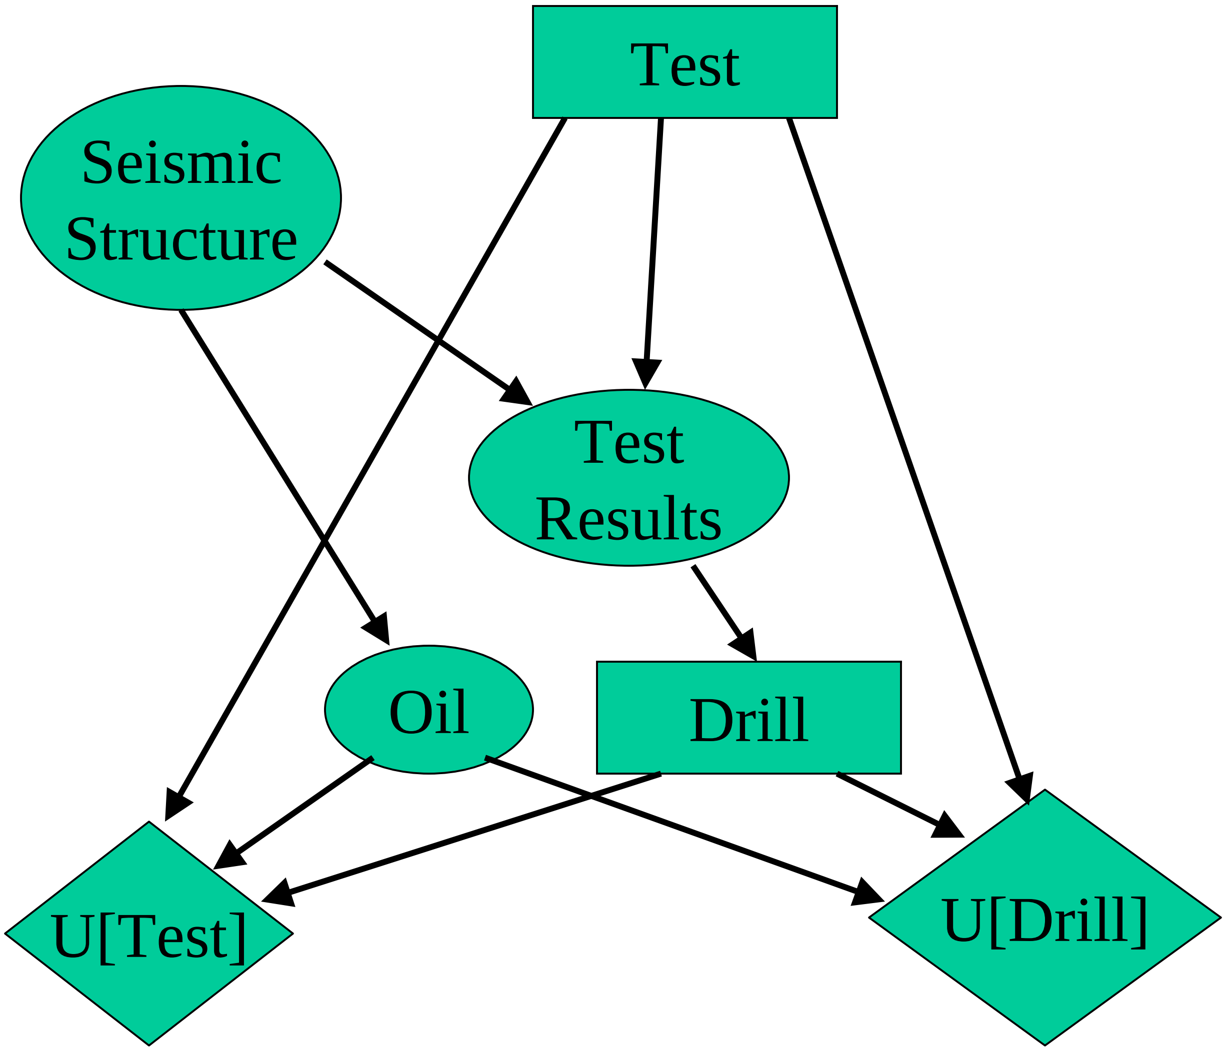

Network of Influence Diagrams Example: Model 1

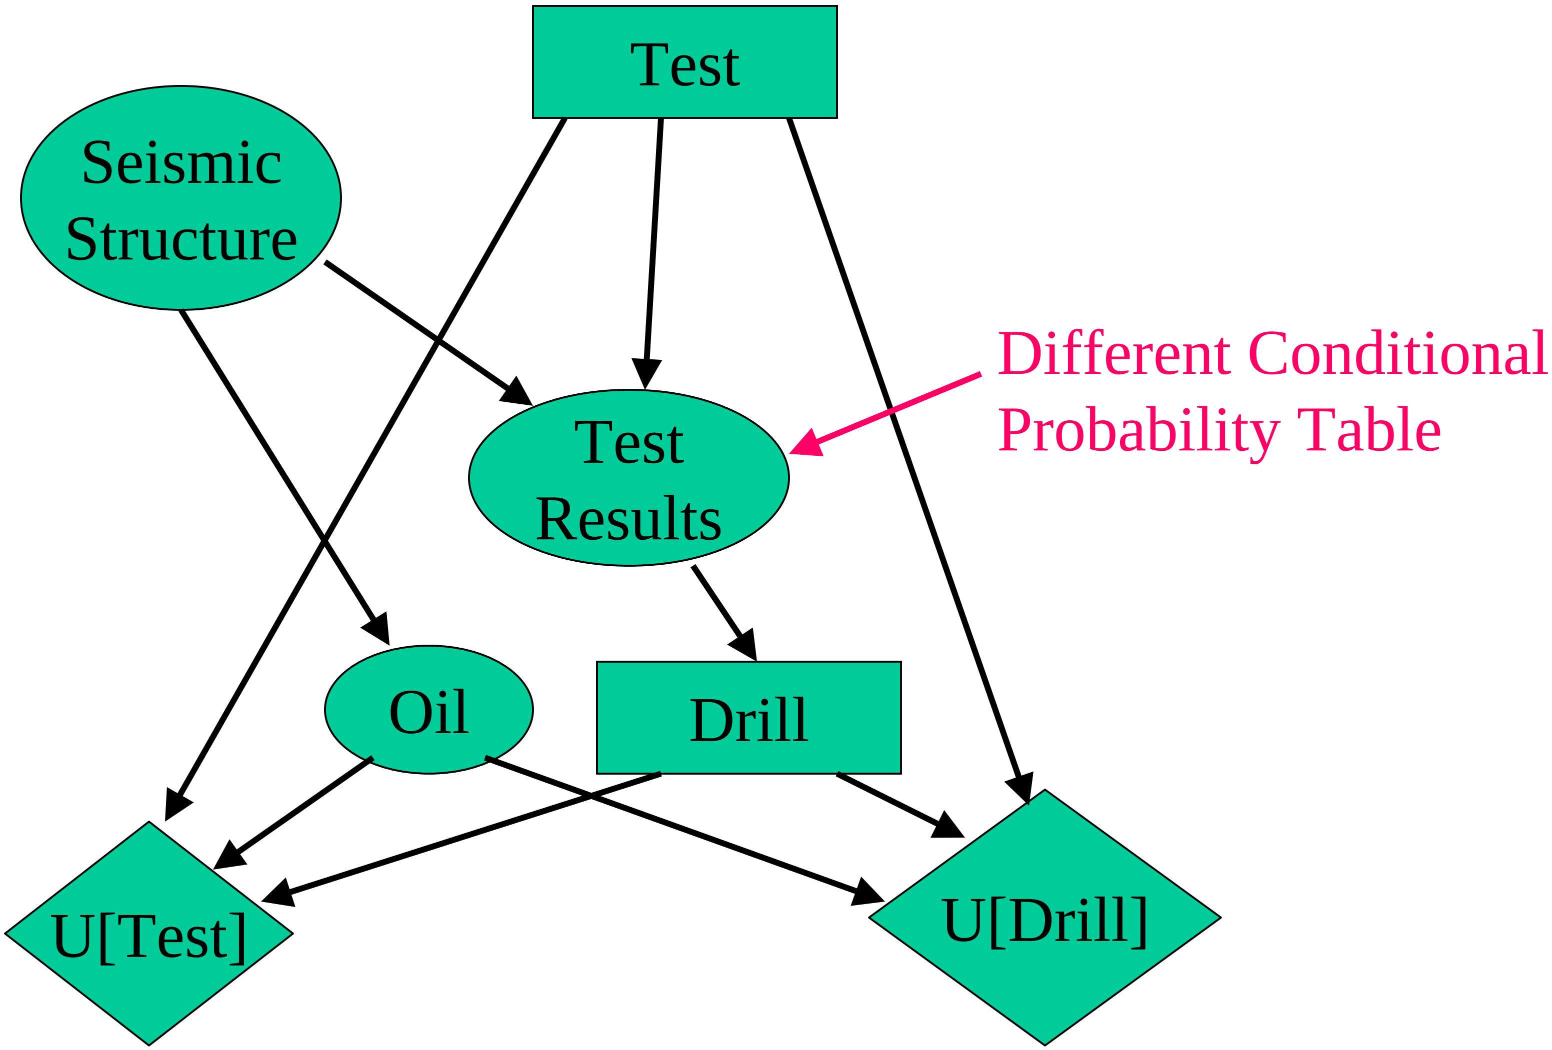

Network of Influence Diagrams Example: Model 2

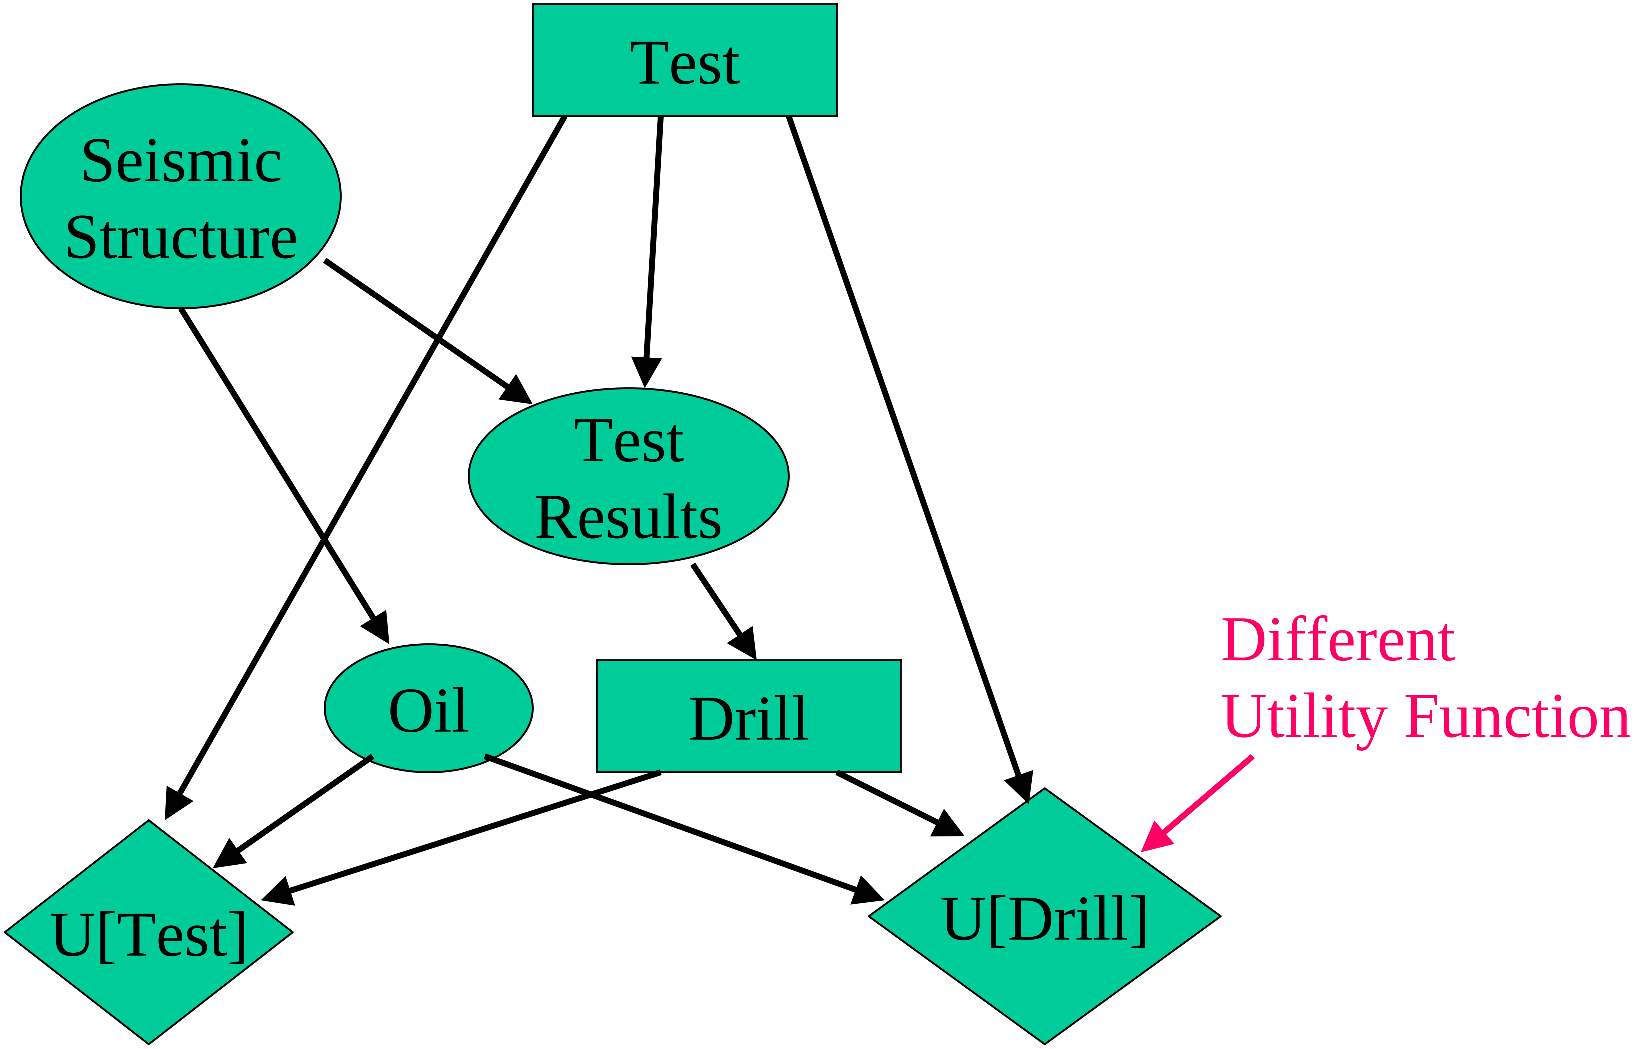

Network of Influence Diagrams Example: Model 3

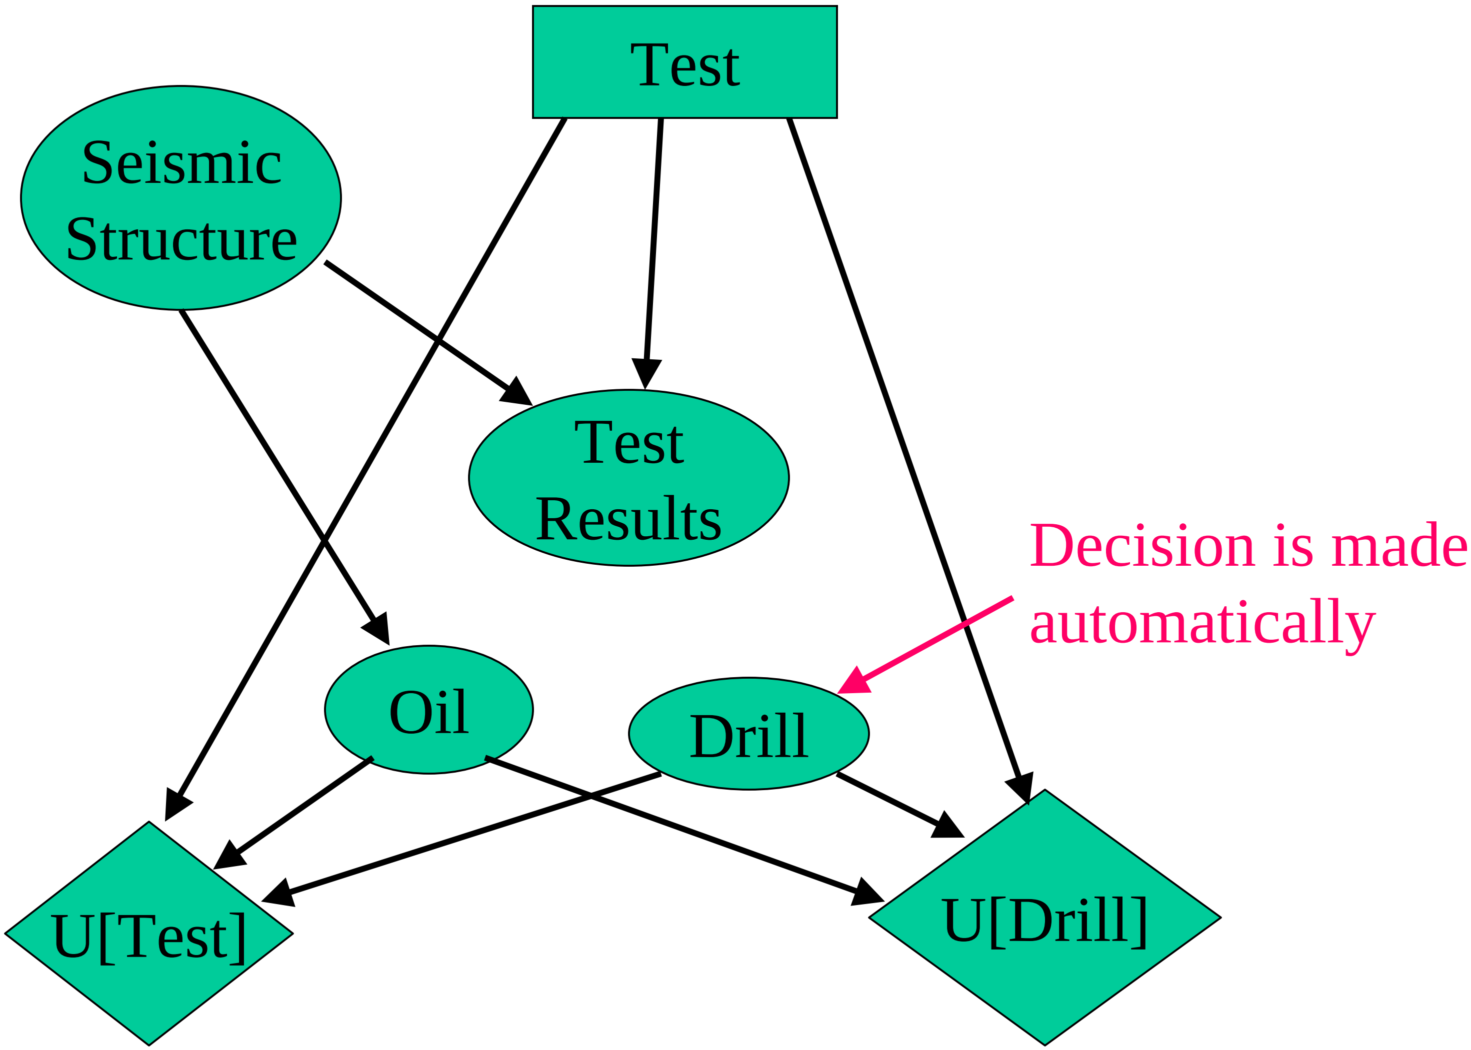

Network of Influence Diagrams Example: Model 4

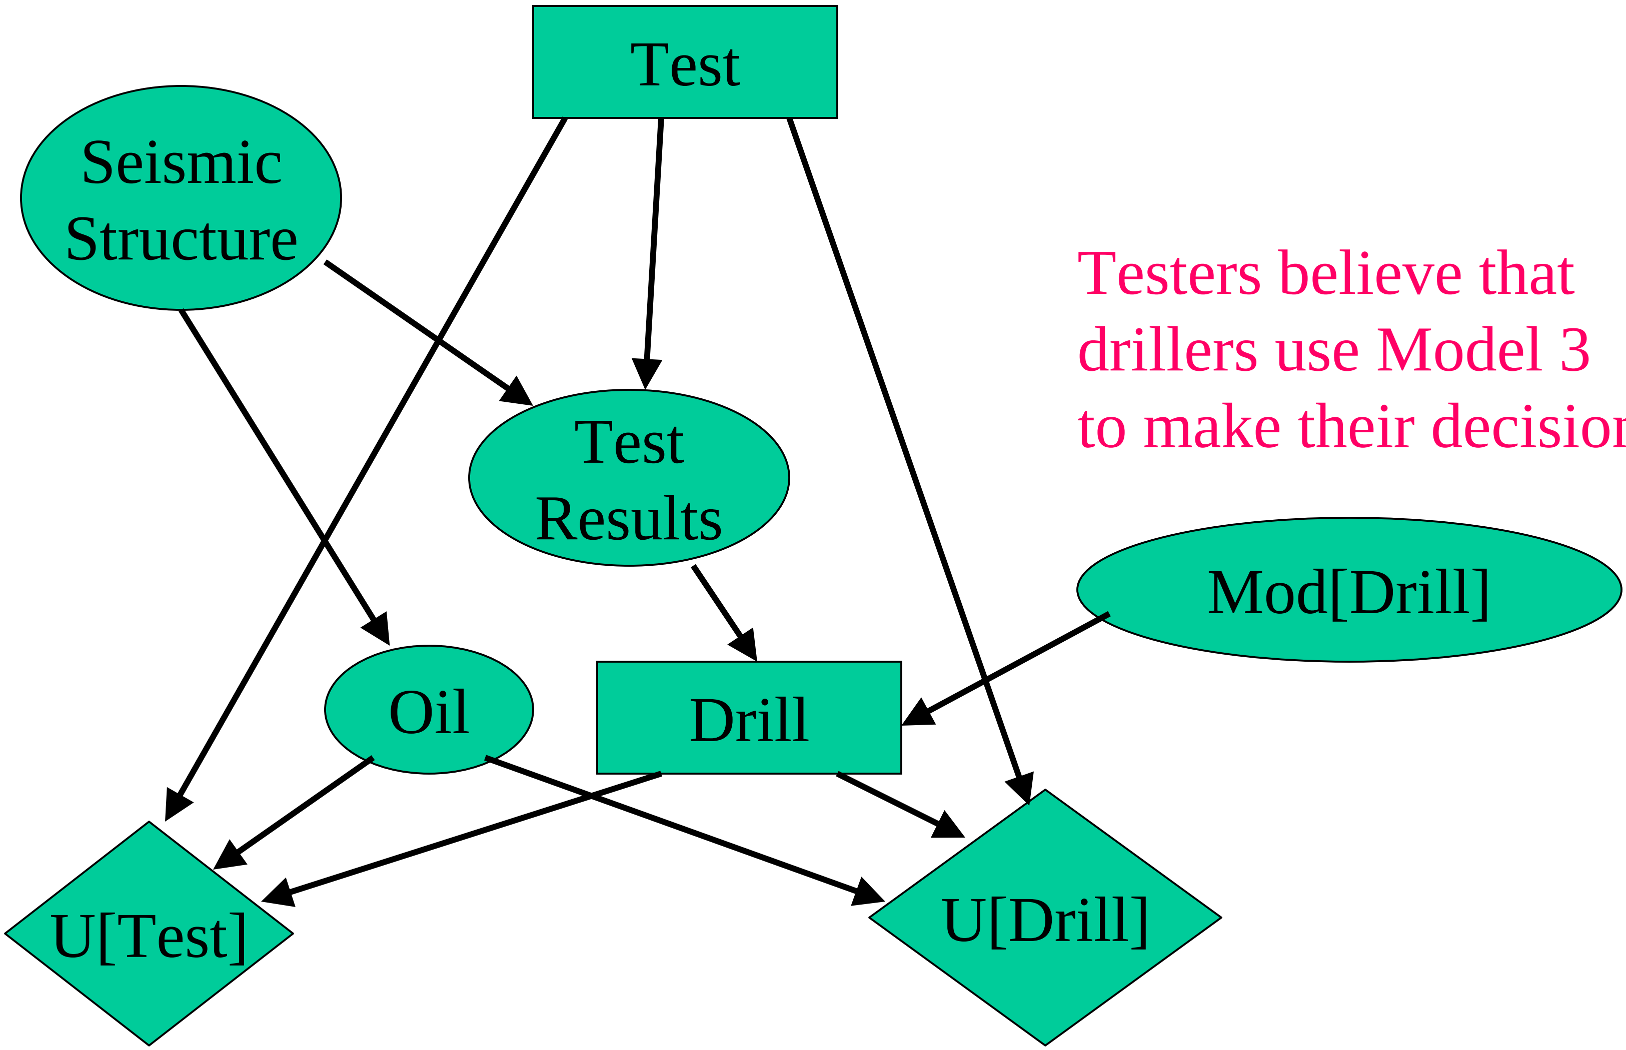

Testers believe that drillers use Model 3 to make their decision

Network of Influence Diagrams Example: Main Model Revisited

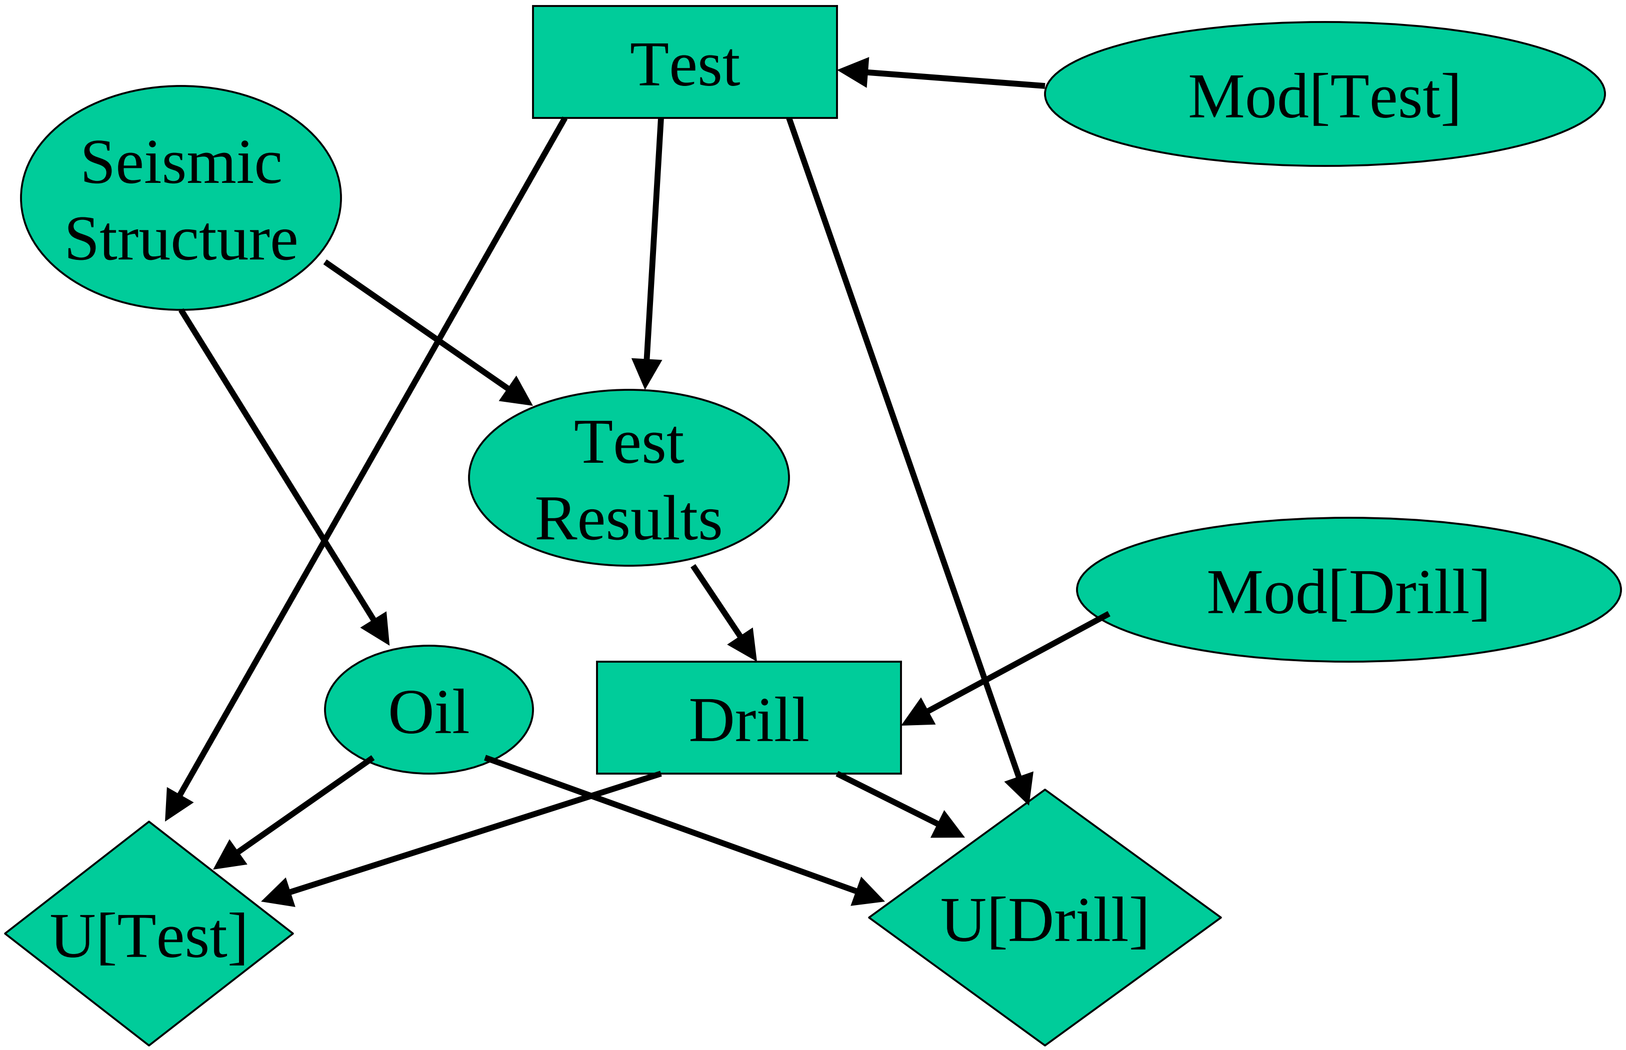

Network of Influence Diagrams Example: NID Structure

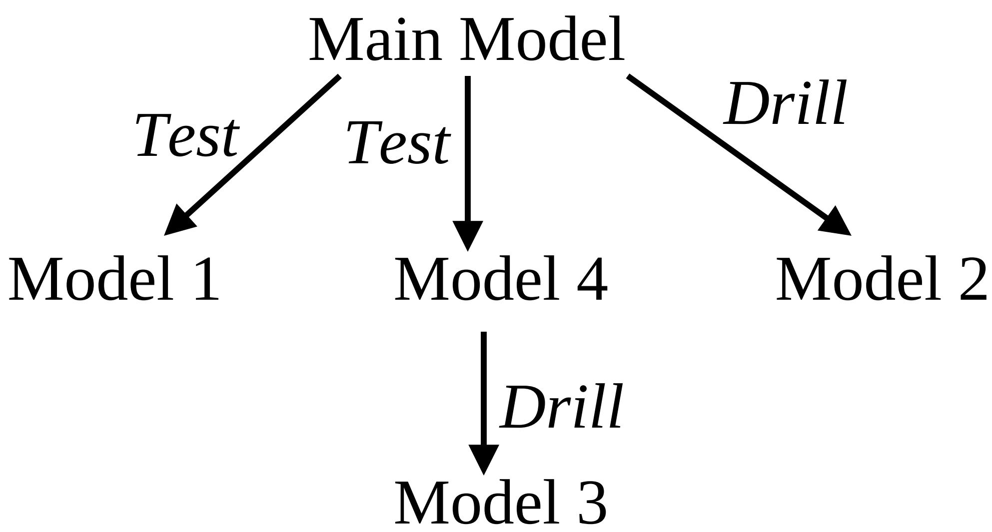

Equilibrium Conditions

- Two kinds of strategy in each mental model
  - Best response strategy
  - Actually played strategy
- These are not necessarily the same!
- Equilibrium conditions:
  - best response is best response with respect to agent's beliefs about actually played strategies
  - actually played strategies are expectations over strategies played in each of the models agent can use

Comparison to Bayesian Games

- Every Bayesian game can be represented by a NID
- Not every NID can be represented by a Bayesian game
- NIDs can be exponentially smaller

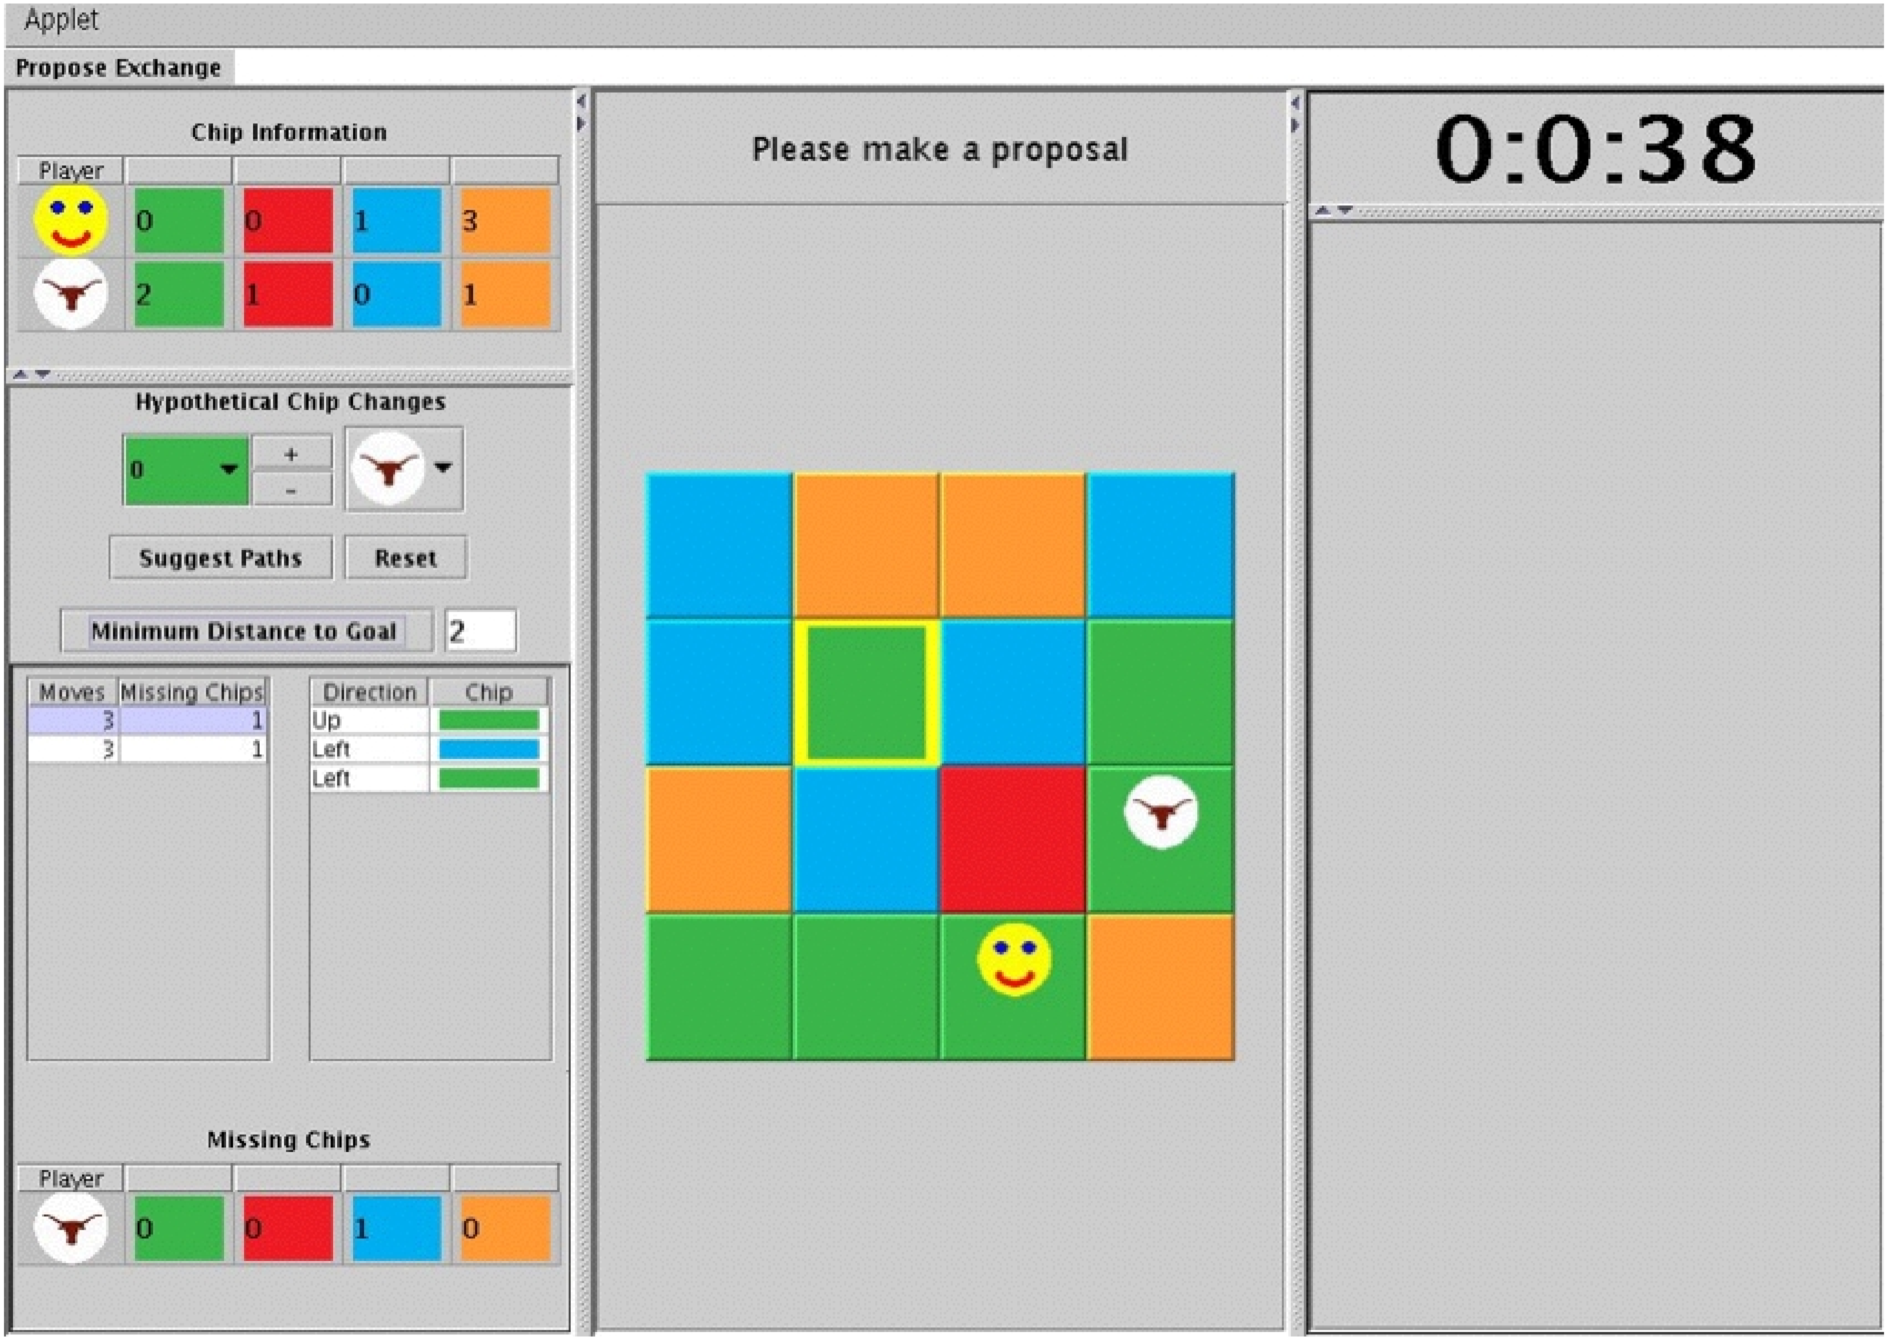
Applet
Propose Exchange
Chip Information
Player
0 0 1 3
2 1 0 1
Hypothetical Chip Changes
0
+
-
Suggest Paths
Reset
Minimum Distance to Goal
2
Moves
Missing Chips
3 1
3 1
Direction
Chip
Up
Left
Left
Missing Chips
Player
0 0 1 0
Please make a proposal
0:0:38

Our Learning Approach

- Model the type of a player by a *social utility function*
  - incorporates different *social preferences* of a player
- Learn a mixture model over types

Social Utility Function

- Individual benefit (IB) $PO_D - NN_D$
- Aggregate utility (AU) $PO_D + PO_A$
- Advantage of outcome (AO) $PO_D - PO_A$
- Advantage of trade (AT)

$$(PO_D - NN_D) - (PO_A - NN_A)$$

Utility = $w_1 IB + w_2 AU + w_3 AO + w_4 AT$

From Utilities to Actions

- Normally, a utility function determines an action
- This results in a discontinuous estimate of the action as a function of social preferences
  - does not allow for noise in play
- To get probabilistic actions, we pass the utility function through a sigmoid

$$P(accept \mid x, t) = \frac{1}{1 + e^{-u^t(x)}}$$

Implementing a Player

- Choose proposal that maximizes expected utility
- Sum over different possible types of deliberator
- Learned model provides probability of acceptance for each type

$$EU(x) = \sum_t P(t)[P(accept \mid x,t)PO_A(x) + (1 - P(accept \mid x,t))NN_A]$$

Experimental Setup

- Data gathering phase
  - 192 games played by people
- Evaluation phase
  - 21 sets of games between human responder and each of four proposers
    - our model
    - Nash equilibrium
    - Nash bargaining solution
    - human proposer

Results

	Score	Proposals accepted	Proposals declined	No offer made
Our model	2880	16	5	0
Nash equilib	2100	13	8	0
Nash bargain	2400	14	2	5
Human	2440	16	1	4