The Multi-Mind Effect

Selmer Bringsjord[1]
Konstantine Arkoudas[2], Deepa Mukherjee[3], Andrew Shilliday[4], Joshua Taylor[5], Micah Clark[6], Elizabeth Bringsjord[7]

Department of Cognitive Science [1-6]
Department of Computer Science [1,4,5]
Rensselaer Polytechnic Institute (RPI)
Troy NY 12180 USA
State University of New York [7]; SUNY Plaza
Albany NY 12246 USA

The Multi-Mind Effect

Selmer Bringsjord[1]
Konstantine Arkoudas[2], Deepa Mukherjee[3], Andrew Shilliday[4], Joshua Taylor[5], Micah Clark[6], Elizabeth Bringsjord[7]

Department of Cognitive Science [1-6]
Department of Computer Science [1,4,5]
Rensselaer Polytechnic Institute (RPI)
Troy NY 12180 USA
State University of New York [7]; SUNY Plaza
Albany NY 12246 USA

Outline

- Introduction to the Multi-Mind Effect
- Dearth of Context Independent Reasoning
- Initial Experiments
  - Experiment Design
  - Results
- Toward Computational Cognitive Modeling of the Multi-Mind Effect
- Implications of the Multi-Mind Effect
- Next steps and Developments

The Multi-Mind Effect

Extensive prior research has shown that logically untrained individuals cannot accurately solve problems that require context-independent reasoning.

The Multi-Mind Effect shows that *groups* of individuals can (without logical training) correctly solve problems that require context-independent reasoning, even though the members that form the groups cannot individually solve these problems correctly.

Dearth of Context-Independent Reasoning

Studies of human reasoning have shown that logically untrained humans systematically fail to reason in a context-independent manner, even when presented with stimuli that expressly call for this type of reasoning.

This failure is attributed to the lack of the appropriate reasoning machinery in humans.

The Stimuli

The Stimuli

The Stimuli

Assume that

(1) It is false that 'If the square is green, the circle is red'.

Given this assumption can you infer that the square is green?

The Stimuli

Assume that

(1) It is false that ‘If the square is green, the circle is red’.

Given this assumption can you infer that the square is green?

Most individuals answer ‘No’.

The Stimuli

Assume that

(1) It is false that ‘If the square is green, the circle is red’.

Given this assumption can you infer that the square is green?

Most individuals answer ‘No’.

The correct answer is ‘Yes’.

The Stimuli

The Stimuli

Assume that

‘If there is a King in the hand then there is an Ace in the hand’, or ‘If there is not a King in the hand, then there is an Ace in the hand’ but not both.

The Stimuli

Assume that

‘If there is a King in the hand then there is an Ace in the hand’, or ‘If there is not a King in the hand, then there is an Ace in the hand’ but not both.

Almost all individuals working alone answer

‘There is an Ace in the hand’

The Stimuli

Assume that

‘If there is a King in the hand then there is an Ace in the hand’, or ‘If there is not a King in the hand, then there is an Ace in the hand’ but not both.

Almost all individuals working alone answer

‘There is an Ace in the hand’

The correct answer is

‘There is not an Ace in the hand’

Mental MetaLogic and the Multi-Mind Effect

Mental MetaLogic (MML) predicts the phenomenon of *heterogeneous reasoning*, where an individual reasoner or groups of reasoners leverage different reasoning mechanisms to reach the normatively correct solution to such problems.

Such reasoners use proof-theoretic and model-theoretic mechanisms of reasoning and move between them to accurately solve the stimulus problems.

Experiment Design

Experiment Design

Stage 1

Subjects - A group of logically untrained individuals.

Materials - Problems that are deemed ‘unsolvable’.

Any individuals that can accurately solve the problems are identified and are not included in the next stage of the experiment.

Experiment Design

Stage 1

Subjects - A group of logically untrained individuals.

Materials - Problems that are deemed 'unsolvable'.

Any individuals that can accurately solve the problems are identified and are not included in the next stage of the experiment.

Stage 2

The individuals who did not get the right answer are randomly assigned to groups. The groups are then given problems that are isomorphic to the original problems.

We hypothesize that some of the groups will be able to accurately solve the isomorphic problems, i.e., the Multi-Mind effect will emerge.

Initial Experiments

Three pilot experiments were carried out to test for the Multi-Mind Effect.

Subjects - 13 undergraduate students from Rensselaer Polytechnic Institute.
One student reached the correct solution in Stage 1.
The rest were assigned randomly to one of four groups in Stage 2.

Materials - Variants of the stimuli, the Wason Selection Task and the Wise Men puzzle and their isomorphic problems.

Experimental Items

Experimental Items

The following item is a sample of the items used in the experiments. It is similar to the first stimulus problem.

What can you infer from the following premise:

"It's not the case that: if Jones is over six feet tall, the hat is too small."

Experimental Items

Experimental Items

The King Ace Problem described earlier was used in these experiments. Another example of a problem in this paradigm is given below.

If one of the following assertions is true then so is the other:
(1) There is a king in the hand if and only if there is an ace in the hand.
(2) There is a king in the hand.

Which is more likely to be in the hand, if either: the king or the ace?

Proof for the King-Ace problem

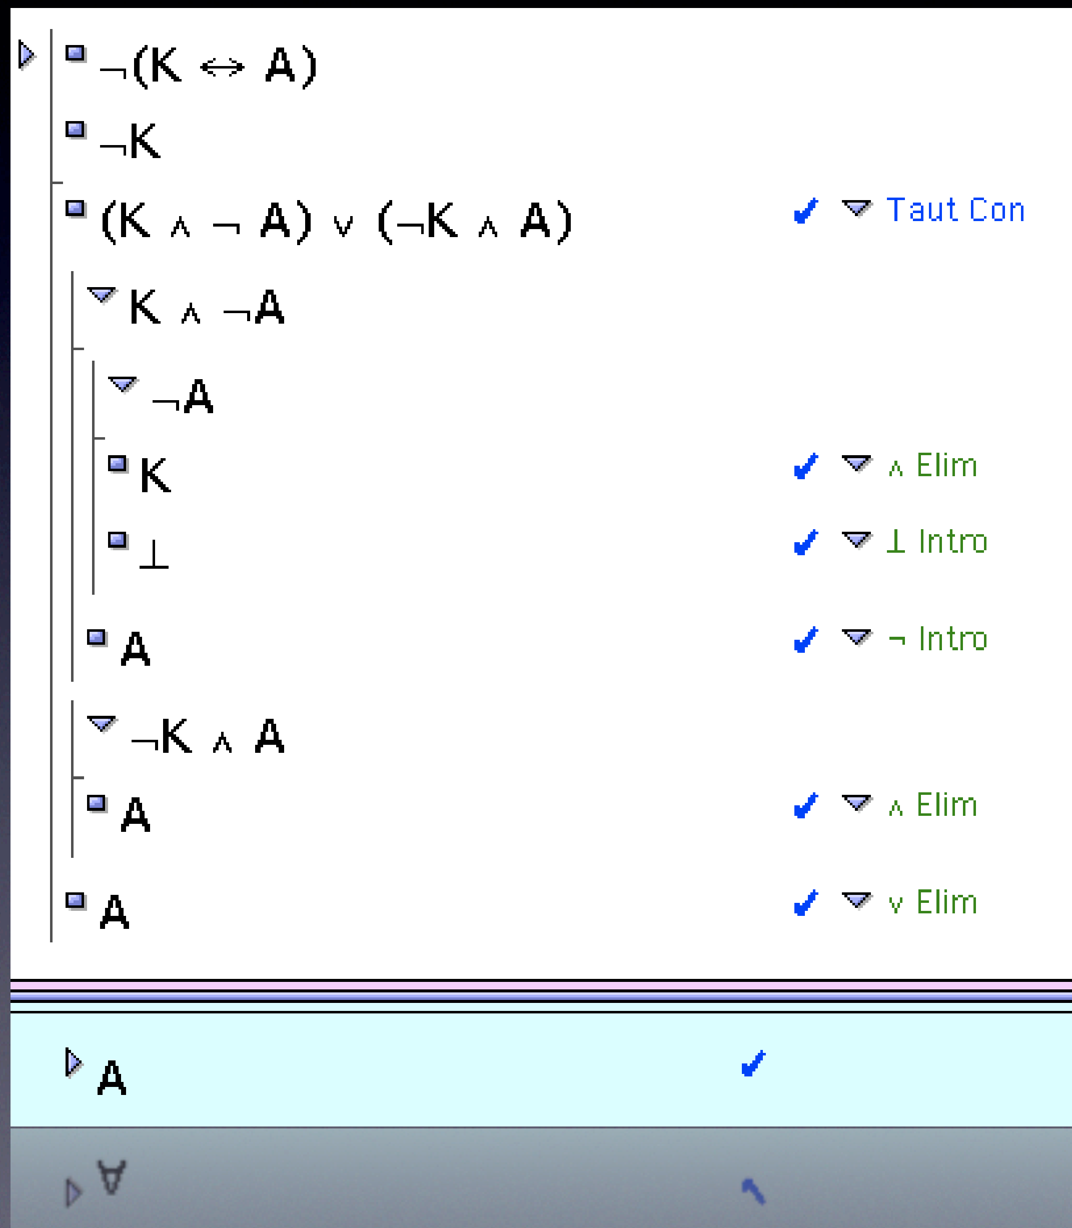

Wason Selection Task

Wason Selection Task

From a deck of cards, where each card has a capital Roman letter on one side, and a digit from 0 through 9 on the other, four cards below are dealt onto a table before you.

E T 4 7

Wason Selection Task

From a deck of cards, where each card has a capital Roman letter on one side, and a digit from 0 through 9 on the other, four cards below are dealt onto a table before you.

E   T   4   7

The following rule is given:

"If there is a vowel on one side, there is an even number on the other."

Which card or cards should be turned over in order to do your best to determine whether this rule is true?

Wason Selection Task

From a deck of cards, where each card has a capital Roman letter on one side, and a digit from 0 through 9 on the other, four cards below are dealt onto a table before you.

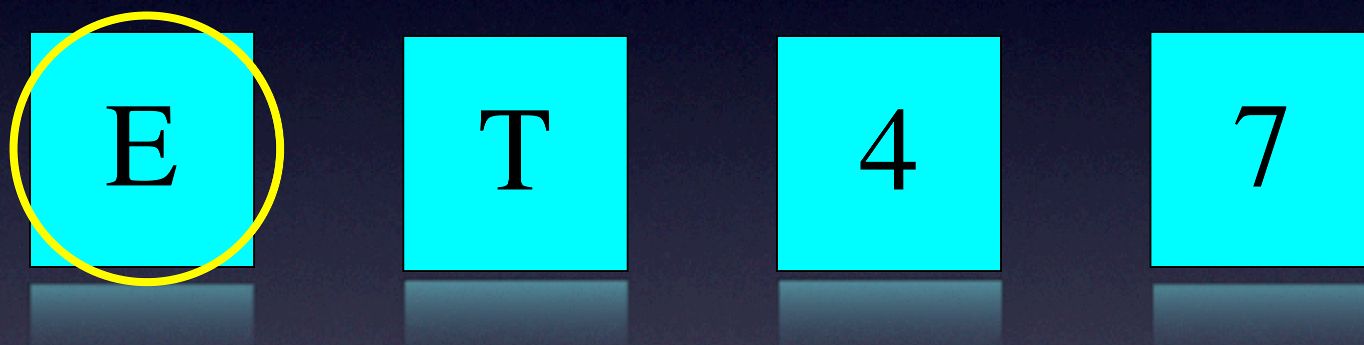

The following rule is given:

"If there is a vowel on one side, there is an even number on the other."

Which card or cards should be turned over in order to do your best to determine whether this rule is true?

Wason Selection Task

From a deck of cards, where each card has a capital Roman letter on one side, and a digit from 0 through 9 on the other, four cards below are dealt onto a table before you.

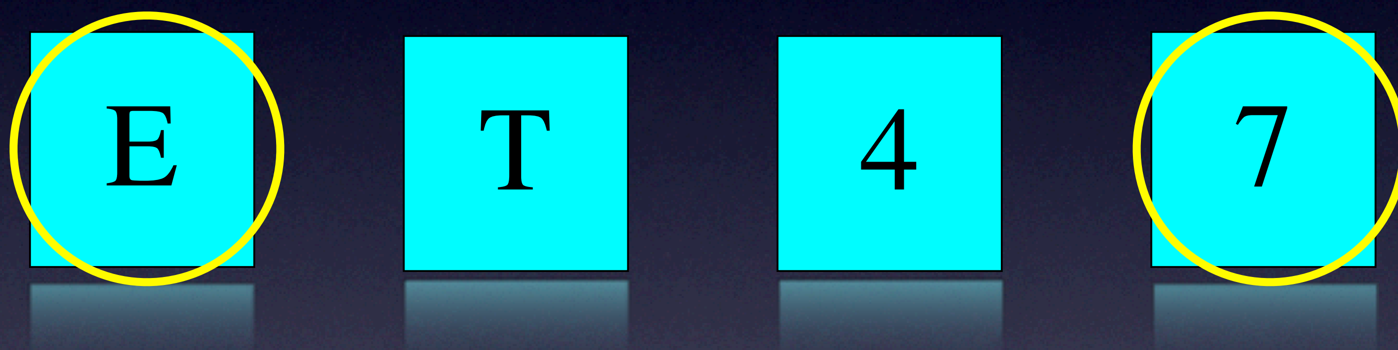

The following rule is given:

"If there is a vowel on one side, there is an even number on the other."

Which card or cards should be turned over in order to do your best to determine whether this rule is true?

Wise Men Puzzle

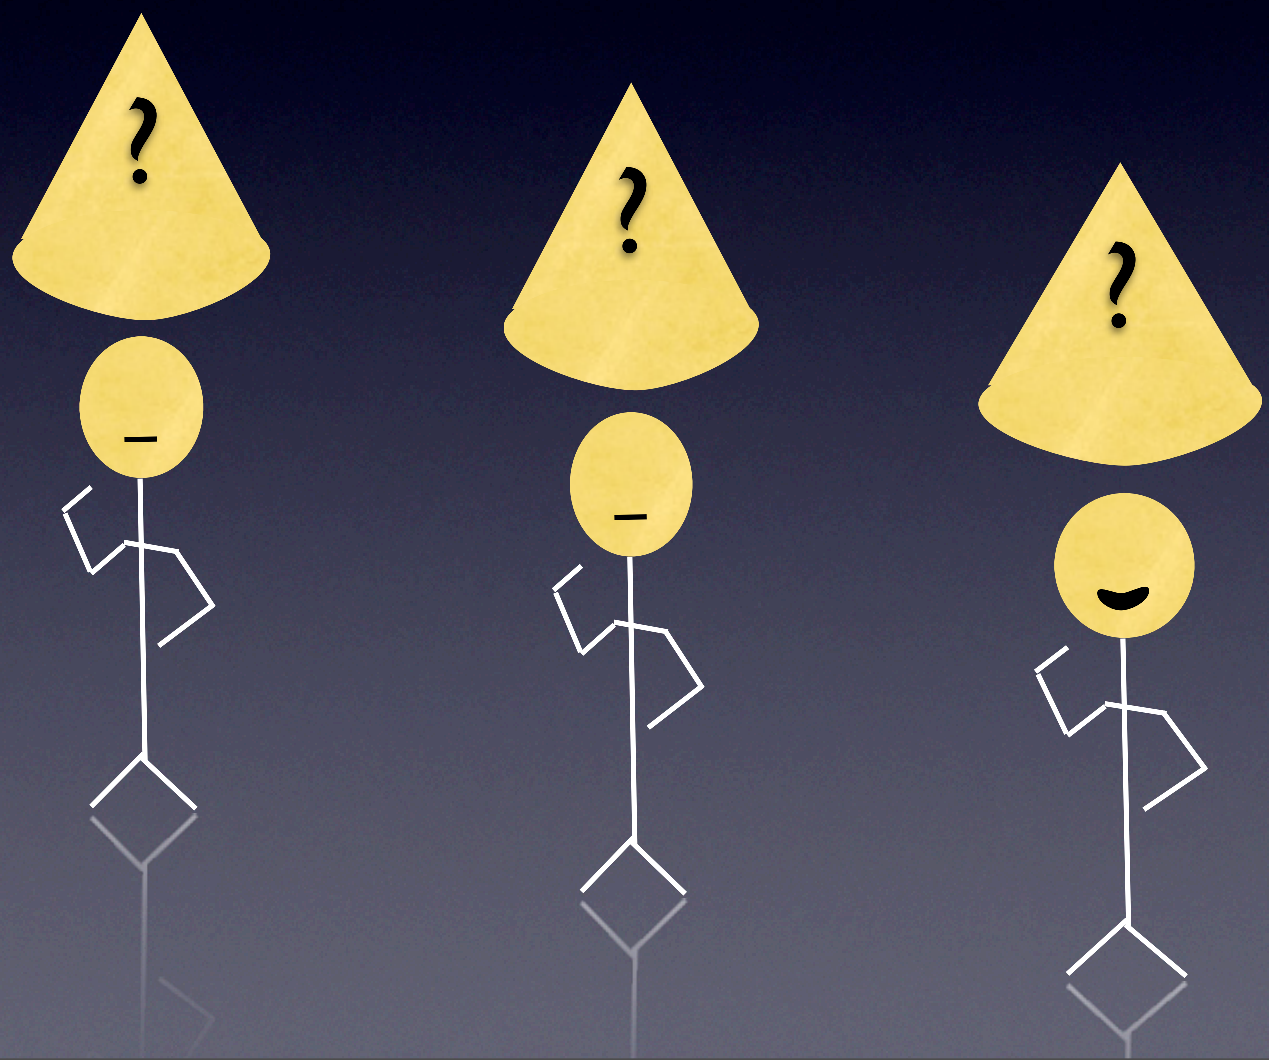

Wise Men Puzzle

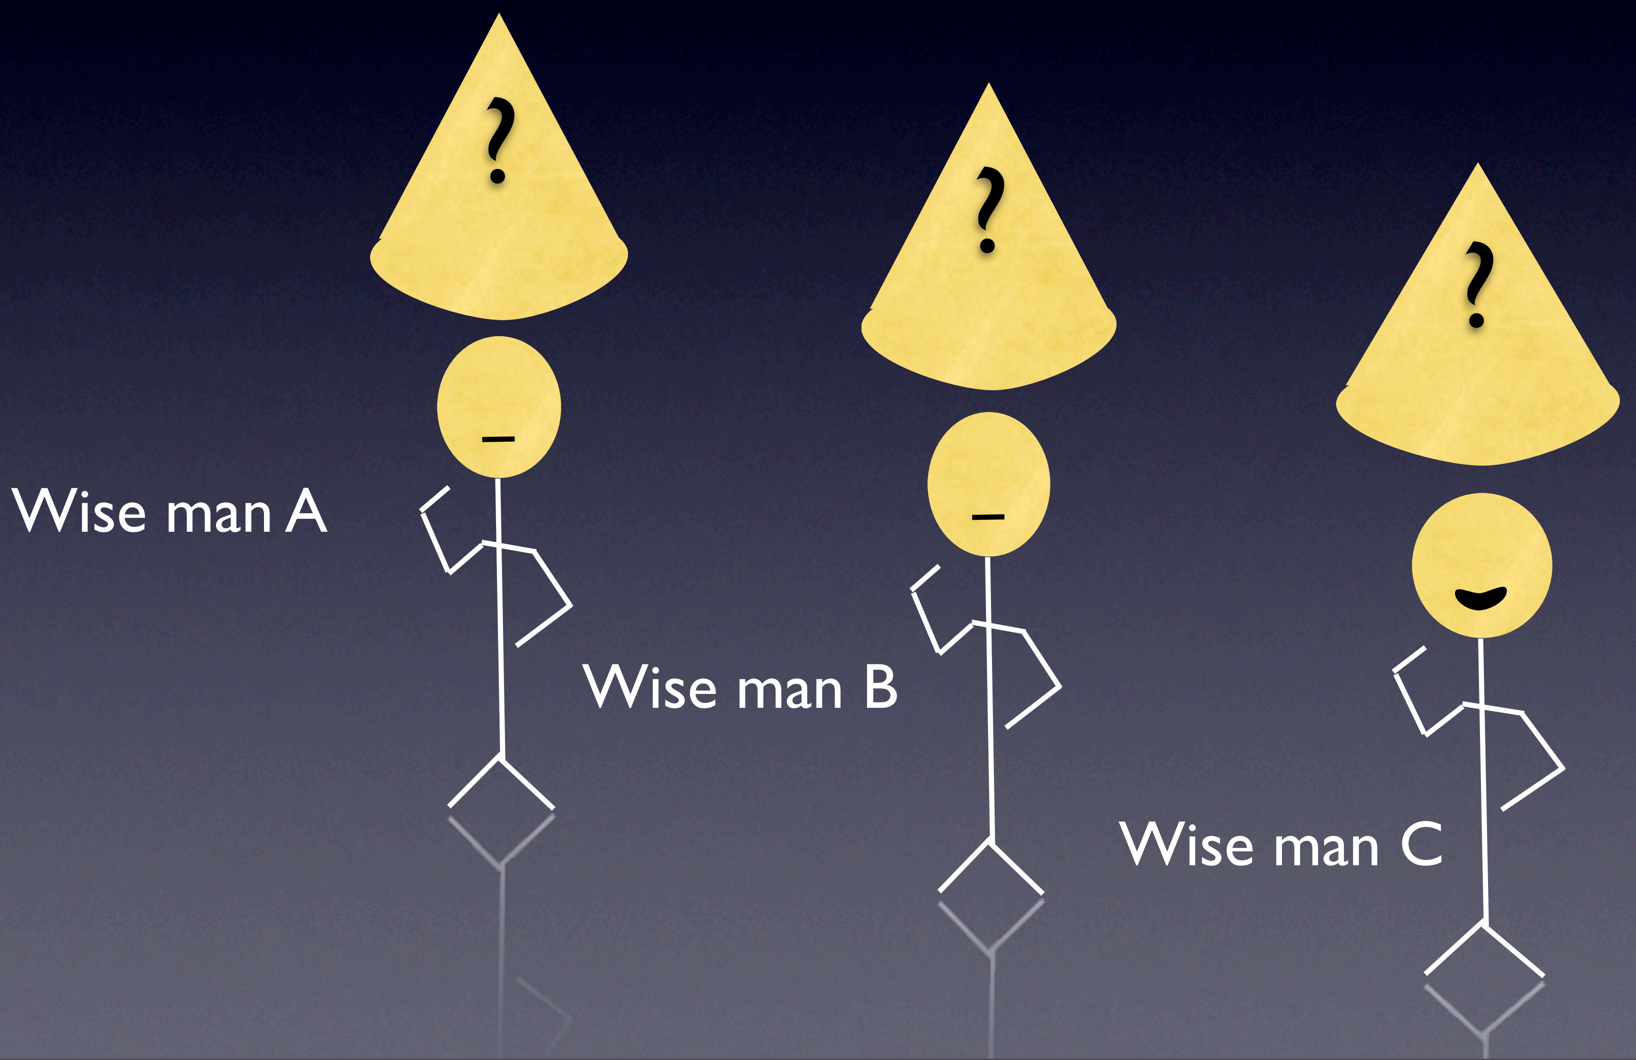

Wise Men Puzzle

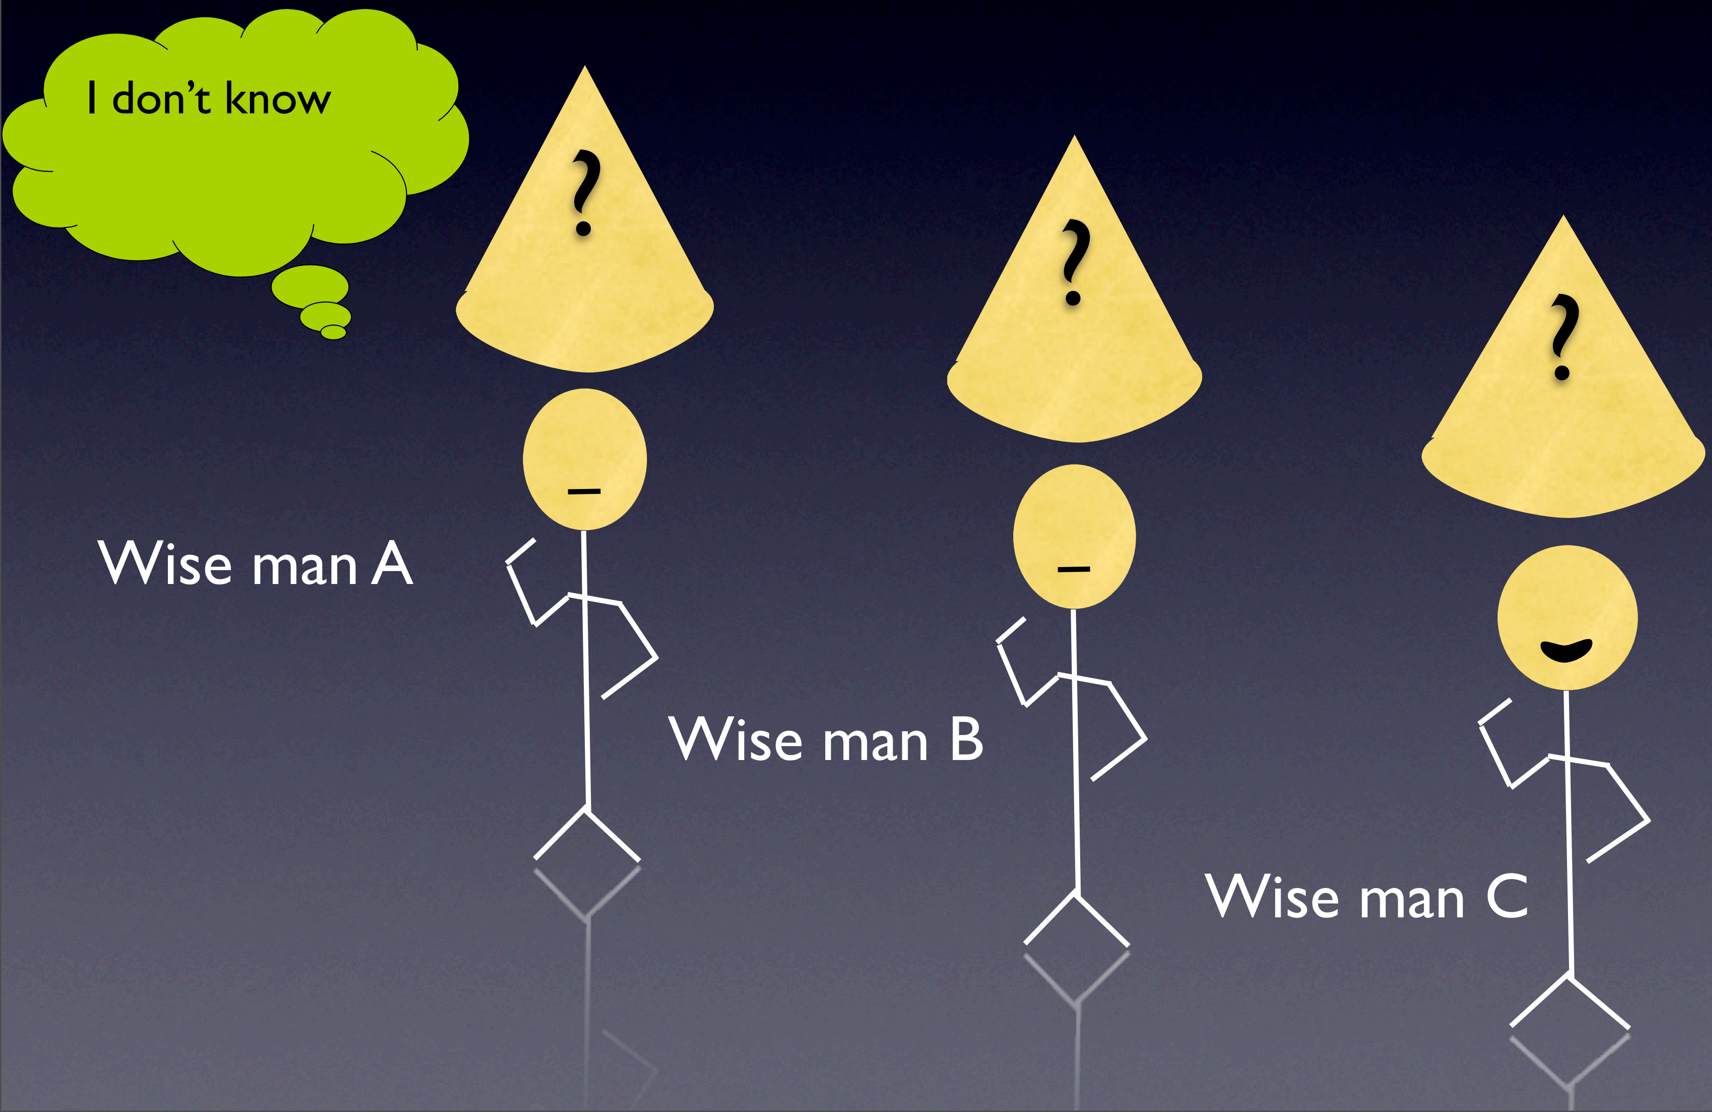

Wise Men Puzzle

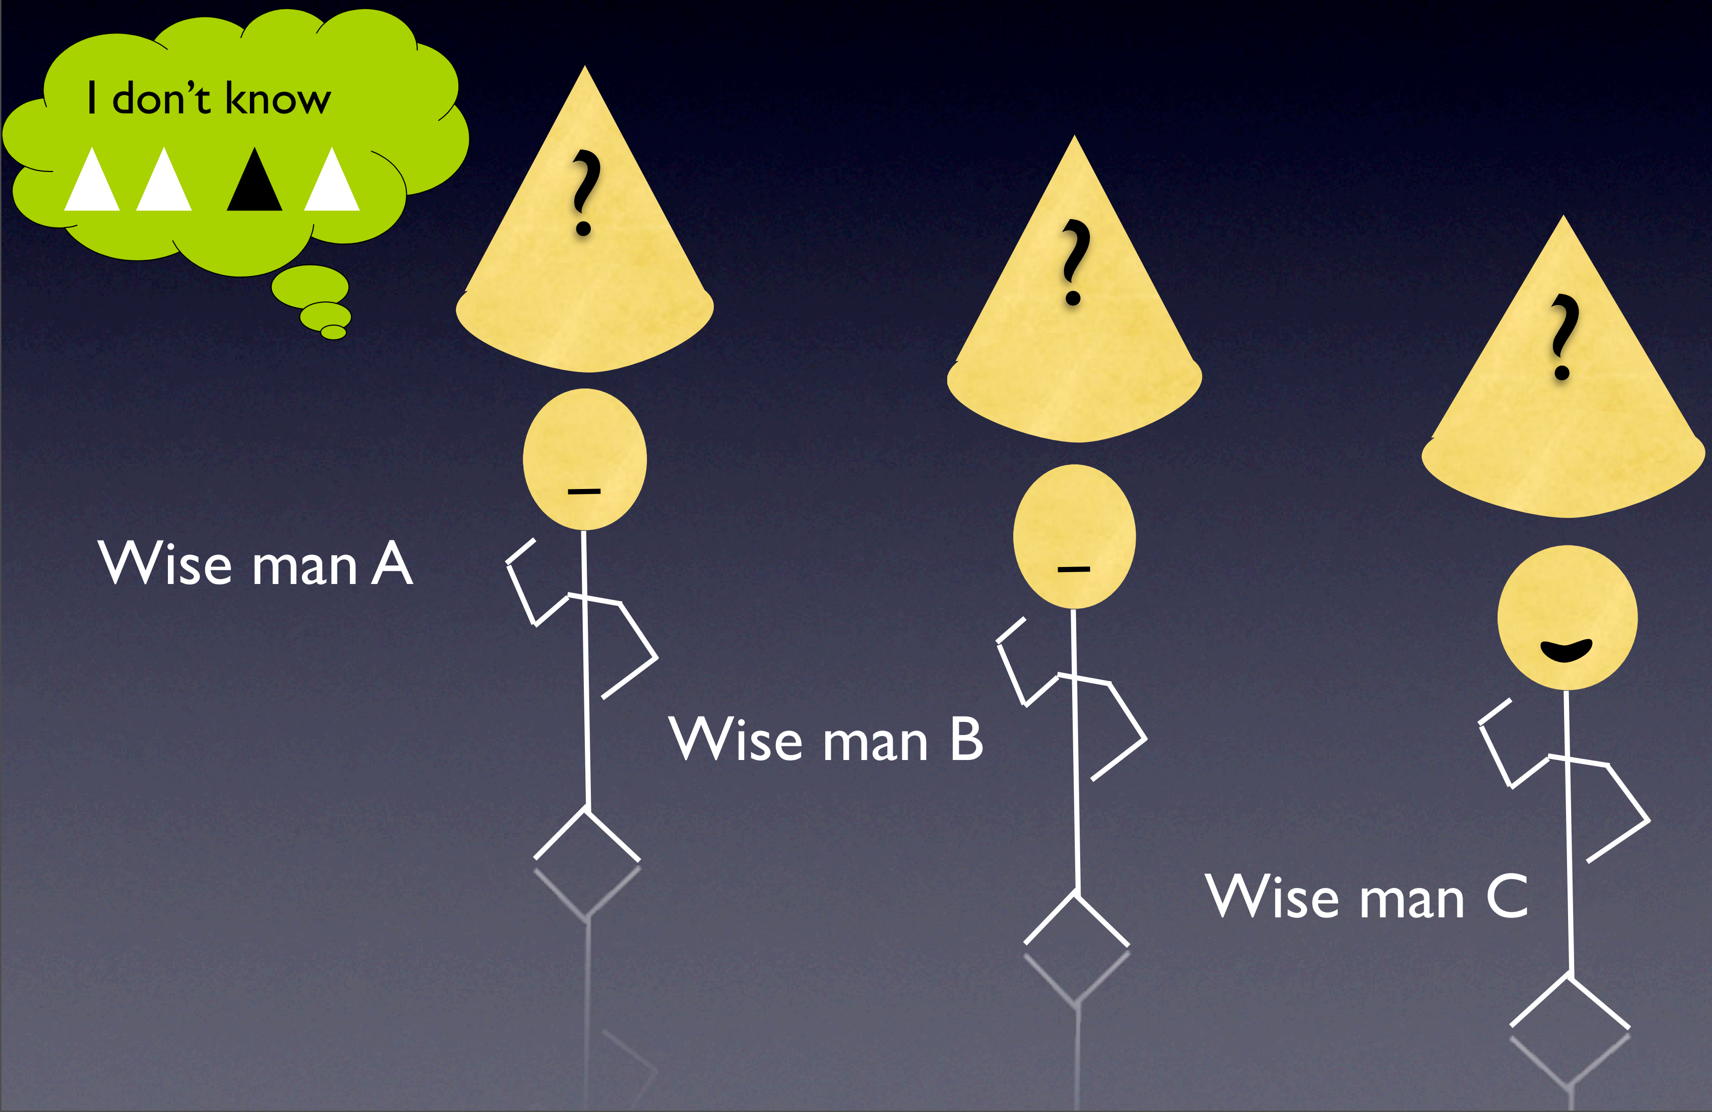

Wise Men Puzzle

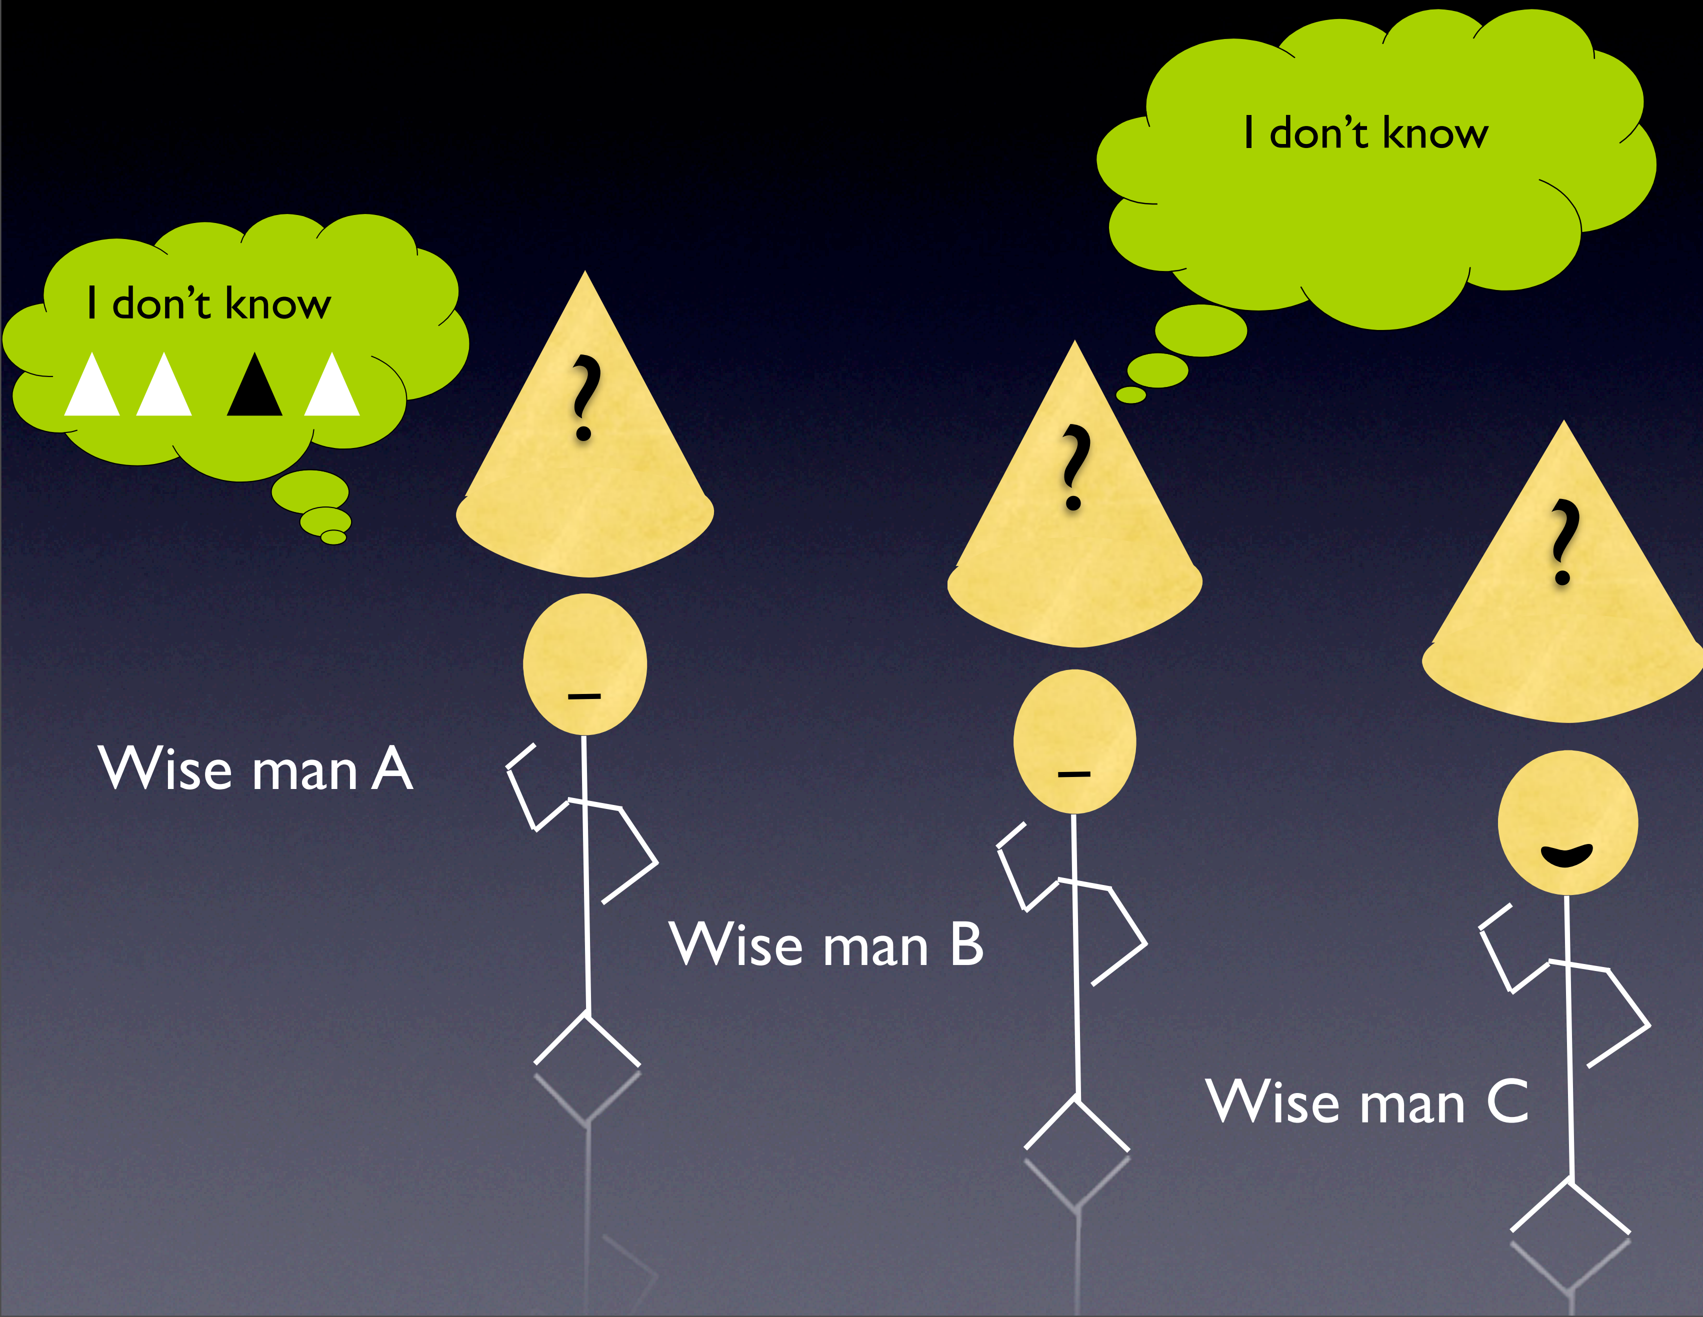

Wise Men Puzzle

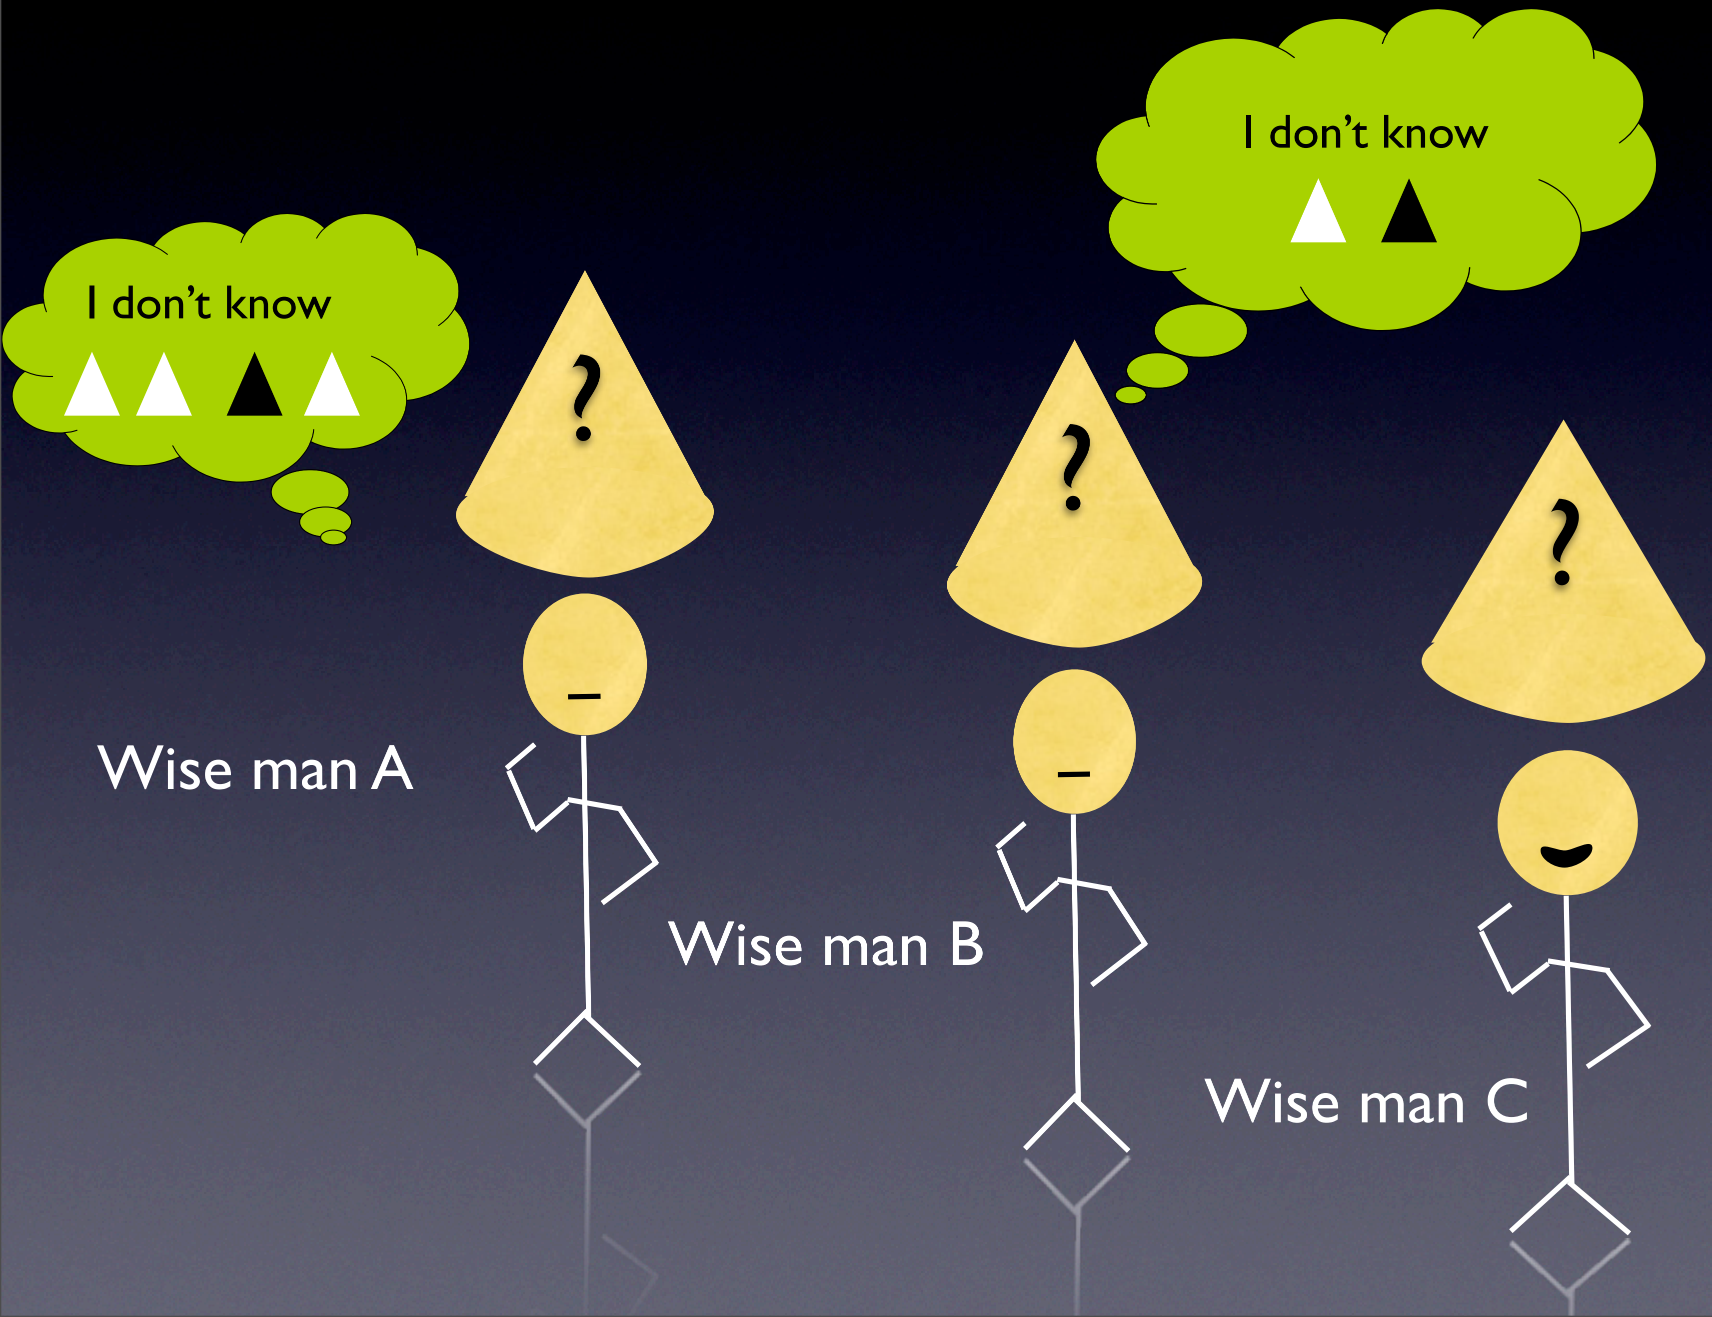

Wise Men Puzzle

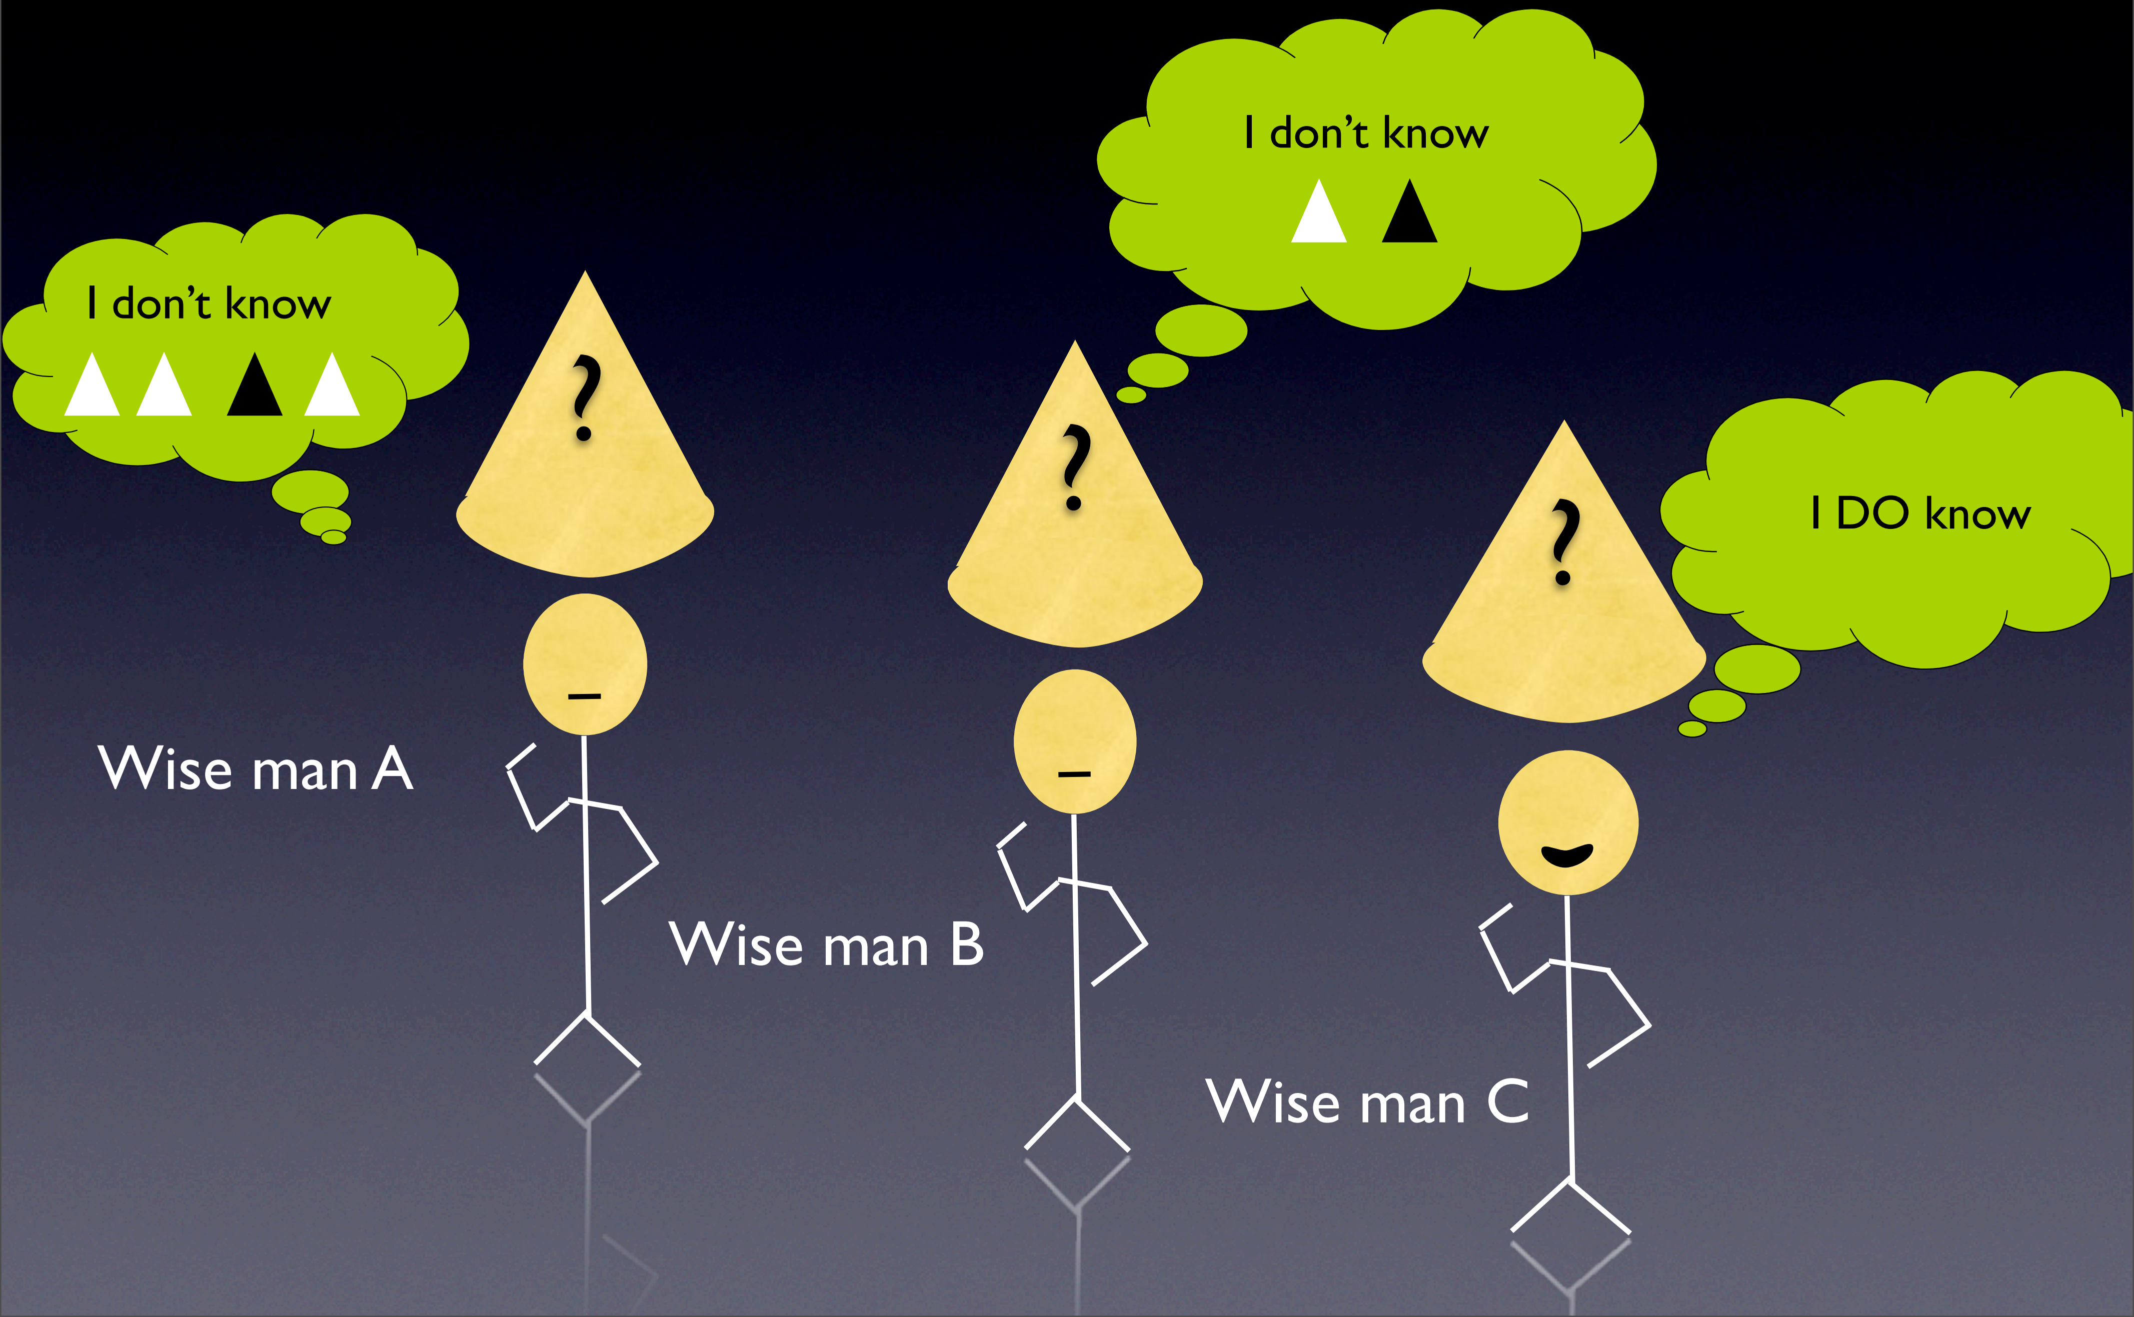

Wise Men Puzzle

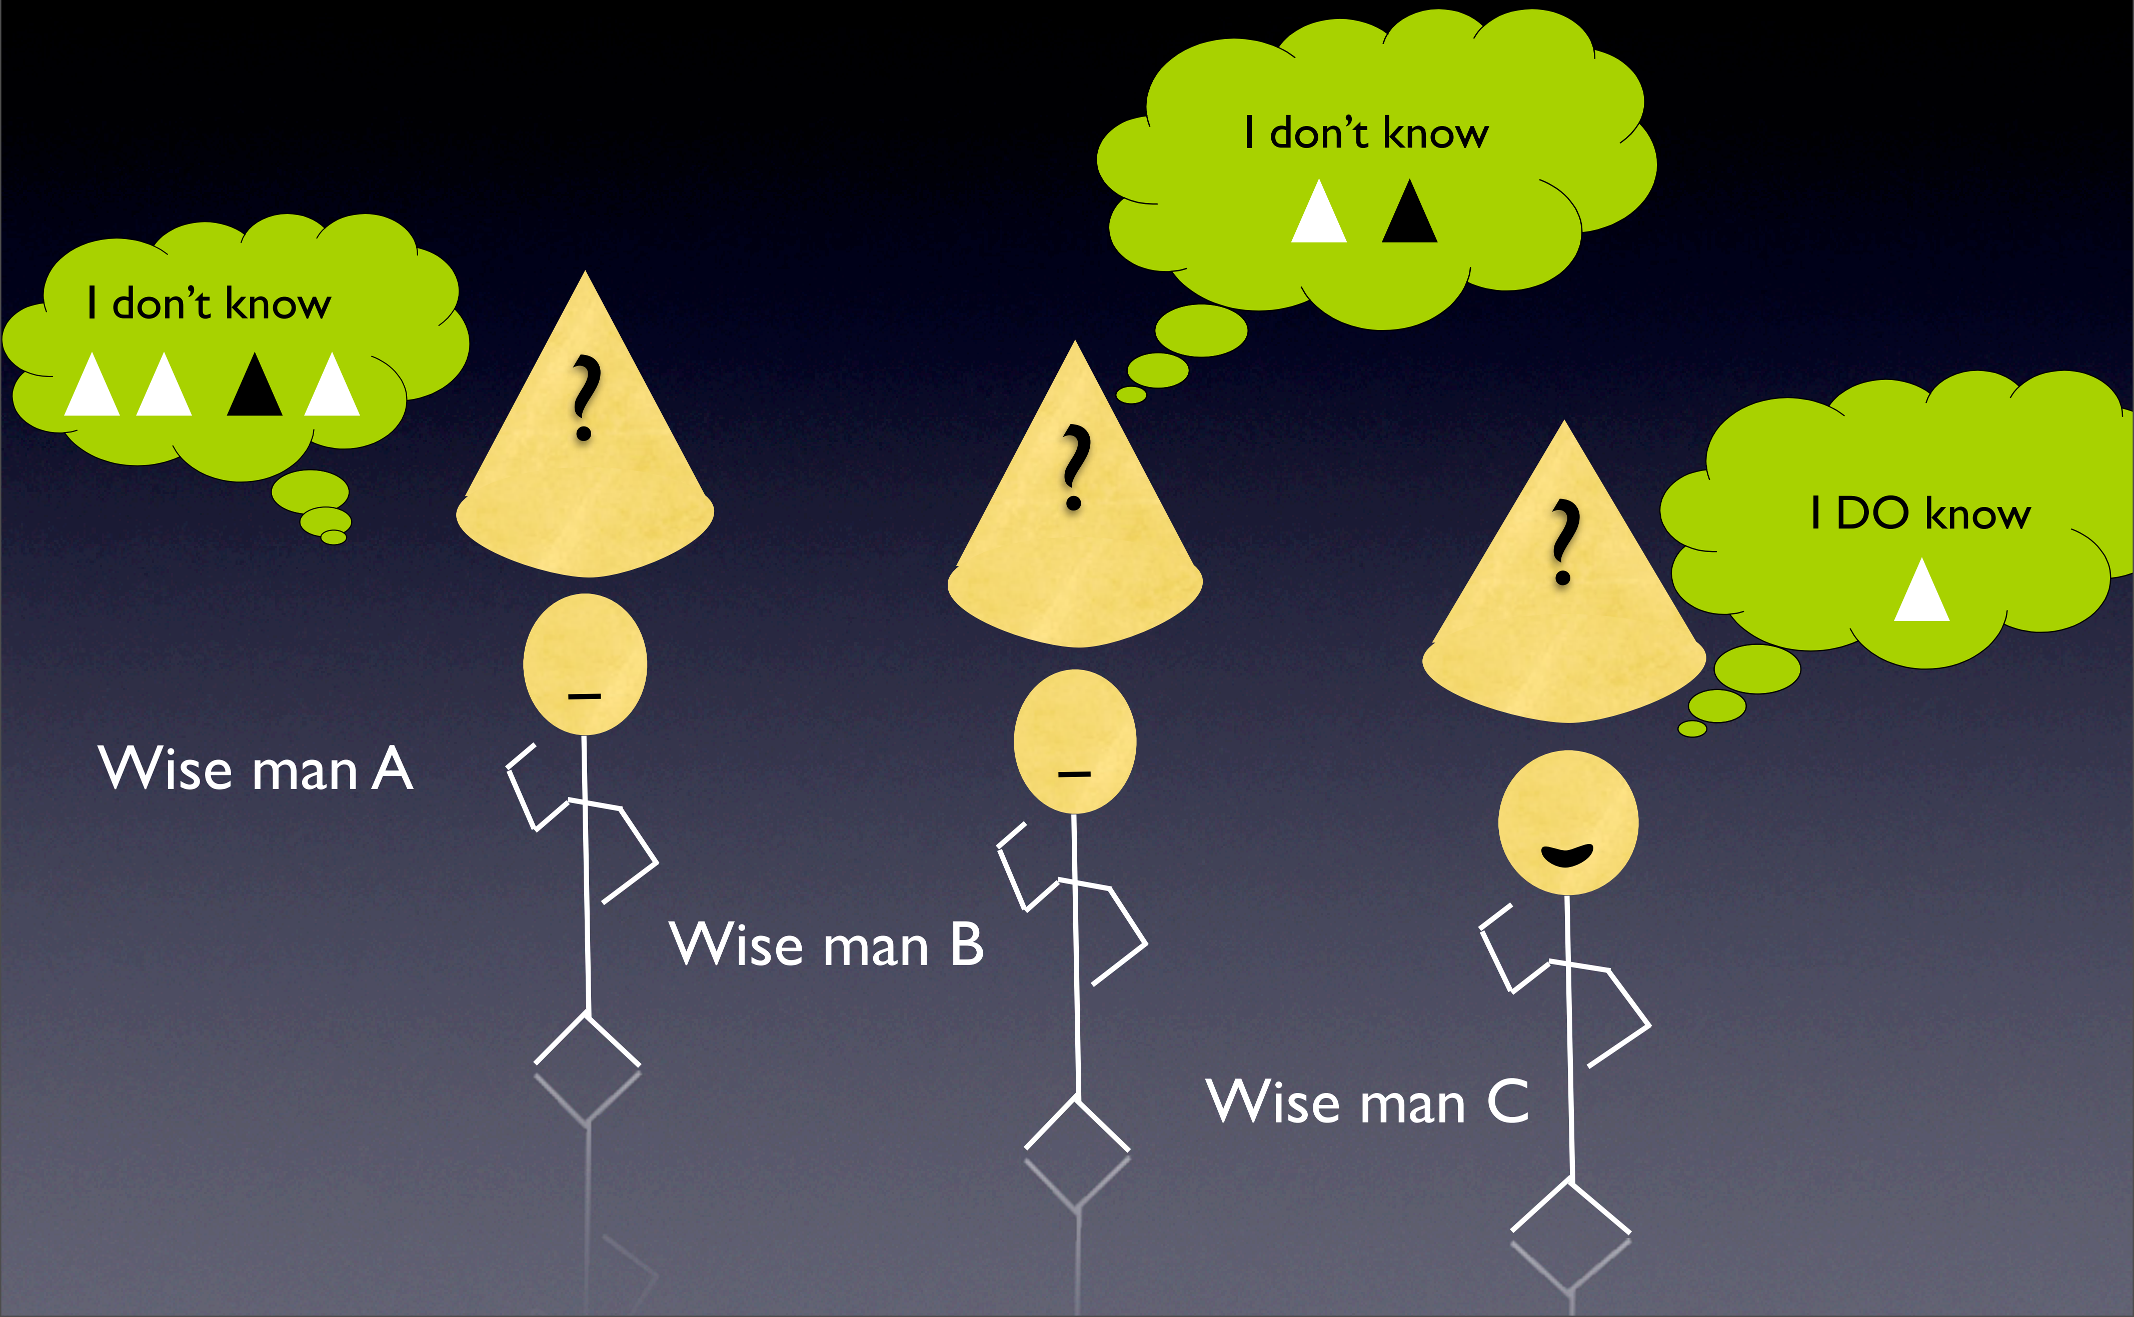

Wise Men Puzzle

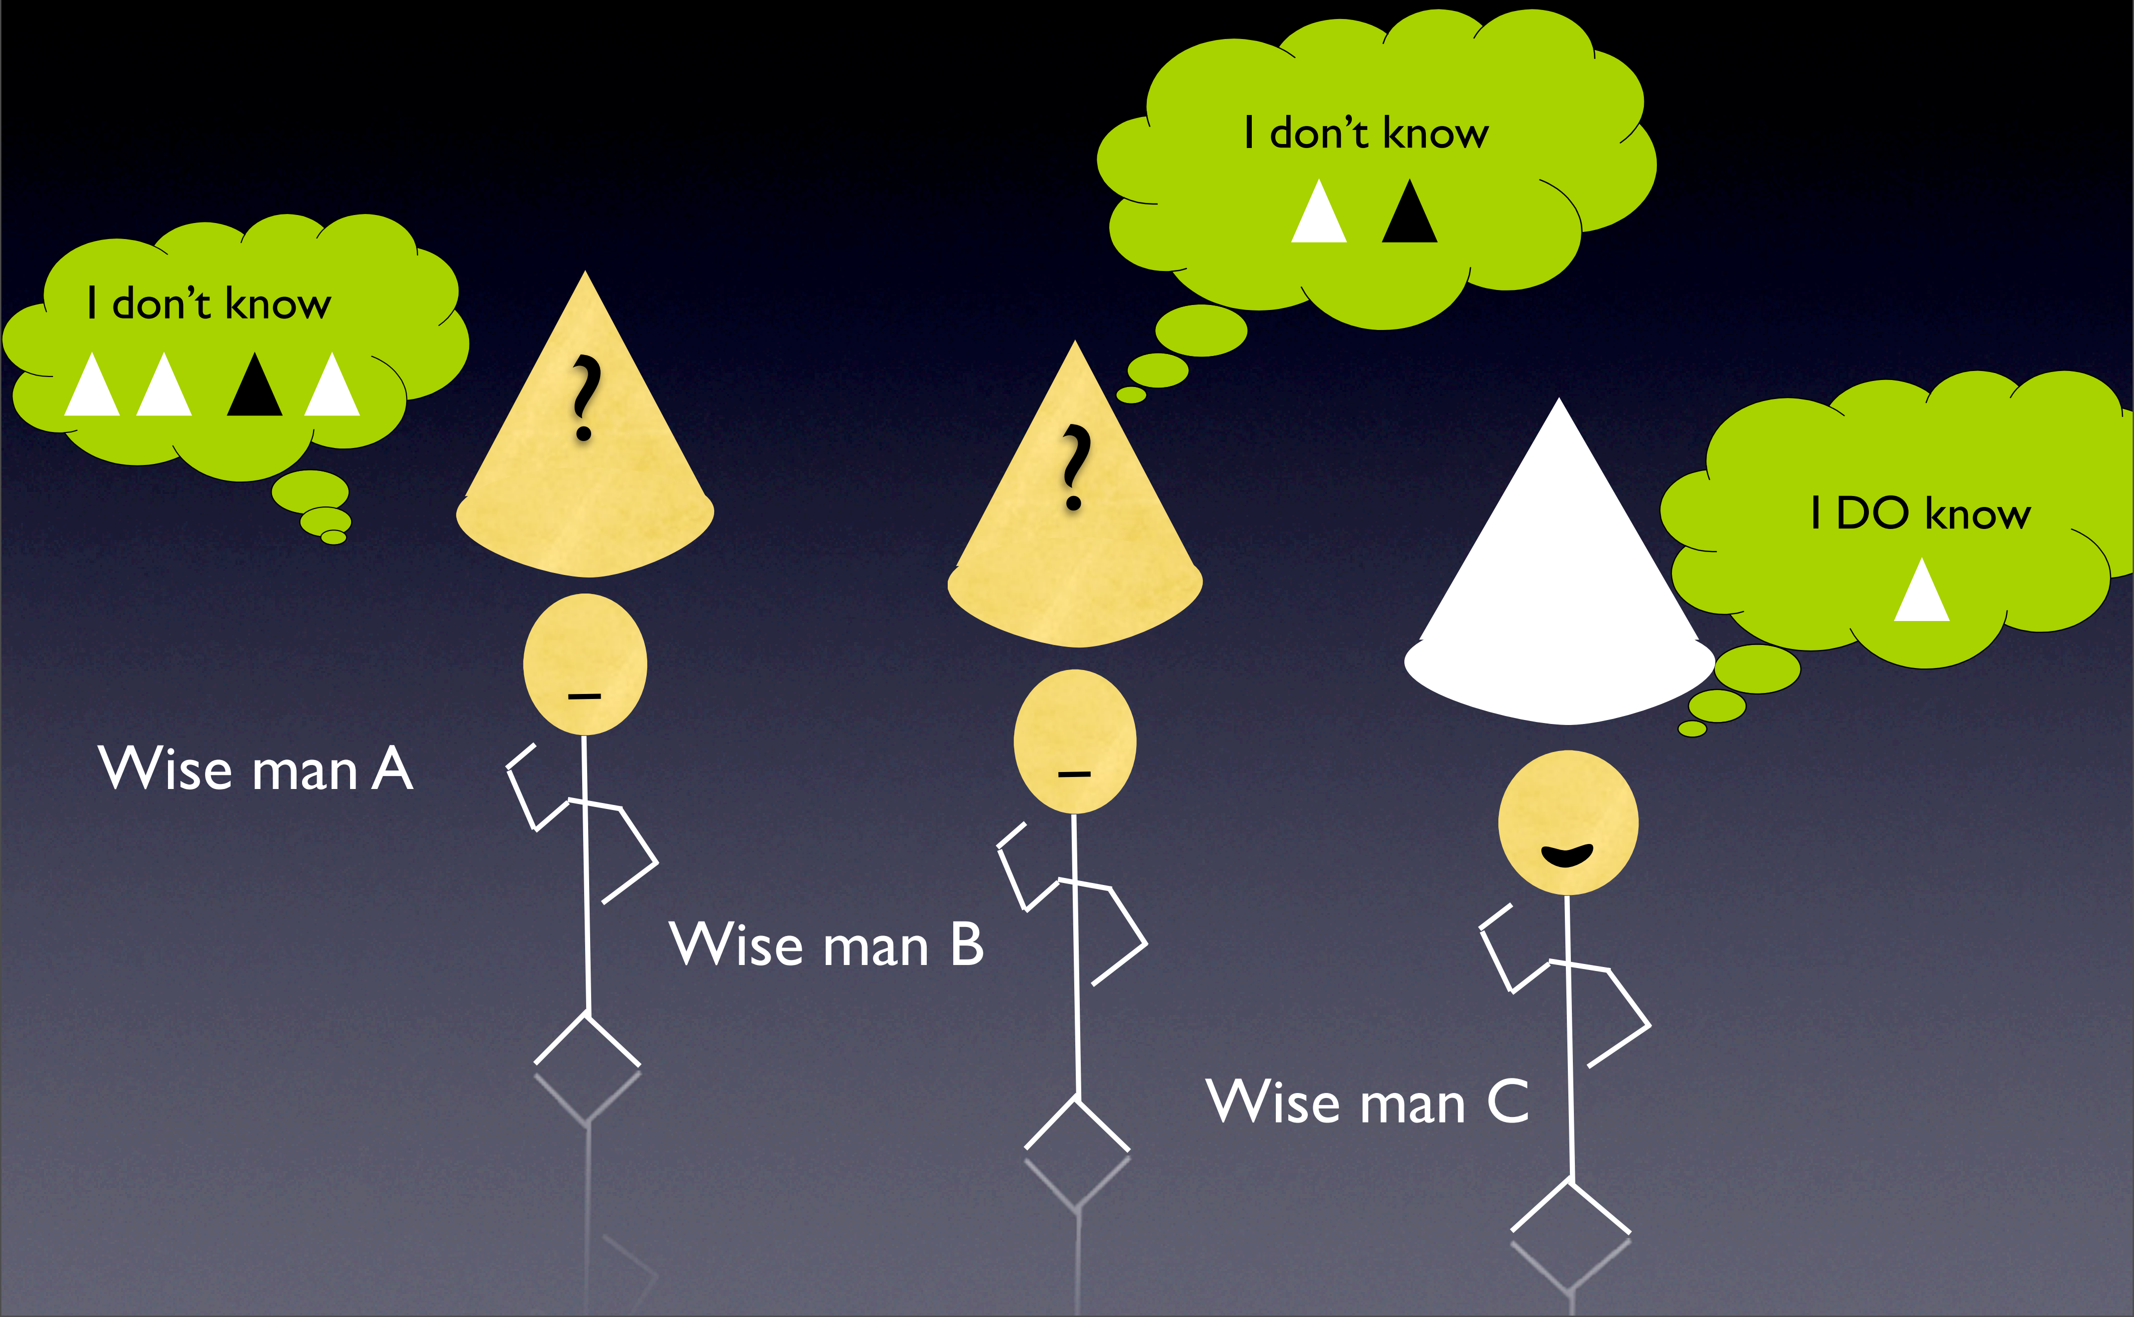

Proved-Sound Algorithm for Generating Proof-Theoretic Solution to WMP$_n$

Springer

Metareasoning for multi-agent epistemic logics

Konstantine Arkoudas and Selmer Bringsjord

RPI

{arkouk,brings}@rpi.edu

Abstract. We present an encoding of a sequent calculus for a multi-agent epistemic logic in Athena, an interactive theorem proving system for many-sorted first-order logic. We then use Athena as a metalanguage in order to reason about the multi-agent logic an as object language. This facilitates theorem proving in the multi-agent logic in several ways. First, it lets us marshal the highly efficient theorem provers for classical first-order logic that are integrated with Athena for the purpose of doing proofs in the multi-agent logic. Second, unlike model-theoretic embeddings of modal logics into classical first-order logic, our proofs are directly convertible into native epistemic logic proofs. Third, because we are able to quantify over propositions and agents, we get much of the generality and power of higher-order logic even though we are in a first-order setting. Finally, we are able to use Athena's versatile tactics for proof automation in the multi-agent logic. We illustrate by developing a tactic for solving the generalized version of the wise men problem.

1 Introduction

Multi-agent modal logics are widely used in Computer Science and AI. Multi-agent epistemic logics, in particular, have found applications in fields ranging from AI domains such as robotics, planning, and motivation analysis in natural language [13]; to negotiation and game theory in economics; to distributed systems analysis and protocol authentication in computer security [16, 31]. The reason is simple—intelligent agents must be able to reason about knowledge. It is therefore important to have efficient means for performing machine reasoning in such logics. While the validity problem for most propositional modal logics is of intractable theoretical complexity[1], several approaches have been investigated in recent years that have resulted in systems that appear to work well in practice. These approaches include tableau-based provers, SAT-based algorithms, and translations to first-order logic coupled with the use of resolution-based automated theorem provers (ATPs). Some representative systems are FaCT [24], KsatC [14], TA [25], LWB [23], and MSPASS [37].

Translation-based approaches (such as that of MSPASS) have the advantage of leveraging the tremendous implementation progress that has occurred over

[1] For instance, the validity problem for multi-agent propositional epistemic logic is PSPACE-complete [18]; adding a common knowledge operator makes the problem EXPTIME-complete [21].

$$\frac{}{\Gamma \vdash [K_\alpha(P \Rightarrow Q)] \Rightarrow [K_\alpha(P) \Rightarrow K_\alpha(Q)]}[K] \qquad \frac{}{\Gamma \vdash K_\alpha(P) \Rightarrow P}[T]$$

$$\frac{\emptyset \vdash P}{\Gamma \vdash C(P)}[C\text{-}I] \qquad \frac{}{\Gamma \vdash C(P) \Rightarrow K_\alpha(P)}[C\text{-}E]$$

$$\frac{}{\Gamma \vdash [C(P \Rightarrow Q)] \Rightarrow [C(P) \Rightarrow C(Q)]}[C_K] \qquad \frac{}{\Gamma \vdash C(P) \Rightarrow C(K_\alpha(P))}[R]$$

Fig. 2. Inference rules for the epistemic operators.

is $\Gamma \vdash P$. Intuitively, this is a judgment stating that P follows from Γ. We will write P, Γ (or Γ, P) as an abbreviation for $\Gamma \cup \{P\}$. The sequent calculus that we will use consists of a collection of inference rules for deriving judgments of the form $\Gamma \vdash P$. Figure 1 shows the inference rules that deal with the standard propositional connectives. This part is standard (e.g., it is very similar to the sequent calculus of Ebbinghaus et al. [15]). In addition, we have some rules pertaining to K_α and C, shown in Figure 2.

Rule $[K]$ is the sequent formulation of the well-known *Kripke axiom* stating that the knowledge operator distributes over conditionals. Rule $[C_K]$ is the corresponding principle for the common knowledge operator. Rule $[T]$ is the "truth axiom": an agent cannot know false propositions. Rule $[C_I]$ is an introduction rule for common knowledge: if a proposition P follows from the empty set of hypotheses, i.e., if it is a tautology, then it is commonly known. This is the common-knowledge version of the "omniscience axiom" for single-agent knowledge which says that $\Gamma \vdash K_\alpha(P)$ can be derived from $\emptyset \vdash P$. We do not need to postulate that axiom in our formulation, since it follows from $[C\text{-}I]$ and $[C\text{-}E]$. The latter says that if it is common knowledge that P then any (every) agent knows P, while $[R]$ says that if it is common knowledge that P then it is common knowledge that (any) agent α knows it. $[R]$ is a reiteration rule that allows us to capture the recursive behavior of C, which is usually expressed via the so-called "induction axiom"

$$C(P \Rightarrow E(P)) \Rightarrow [P \Rightarrow C(P)]$$

where E is the shared-knowledge operator. Since we do not need E for our purposes, we omit its formalization and "unfold" C via rule $[R]$ instead.

We state a few lemmas that will come handy later:

Lemma 1 (Cut). *If $\Gamma_1 \vdash P_1$ and $\Gamma_2, P_1 \vdash P_2$ then $\Gamma_1 \cup \Gamma_2 \vdash P_2$.*

Proof: Assume $\Gamma_1 \vdash P_1$ and $\Gamma_2, P_1 \vdash P_2$. Then, by $[\Rightarrow\text{-}I]$, we get $\Gamma_2 \vdash P_1 \Rightarrow P_2$. Further, by dilution, we have $\Gamma_1 \cup \Gamma_2 \vdash P_1 \Rightarrow P_2$ and $\Gamma_1 \cup \Gamma_2 \vdash P_1$. Hence, by $[\Rightarrow\text{-}E]$, we obtain $\Gamma_1 \cup \Gamma_2 \vdash P_2$. ❑

The proofs of the remaining lemmas are equally simple exercises:

All human-authored proofs machine-checked.

Initial Results

All the groups reached the correct solution for the problems isomorphic to the stimuli problems and the Wason Selection Task.

One group managed to correctly solve the Wise Men puzzle.

These results, though extremely preliminary, show support for the presence of the Multi-Mind Effect in multi-agent reasoning.

Computational Cognitive Modeling of the Multi-Mind Effect

Logic-based Computational Cognitive Modeling (LCCM) is the formal modeling approach that underlies top-down, declarative modeling. We use this approach to model the Multi-Mind effect.

Some of the authors have previously undertaken research designed to simulate multi-agent reasoning, where the formalisms are in line with LCCM.

http://www.cogsci.rpi.edu/slate

http://www.cogsci.rpi.edu/slate

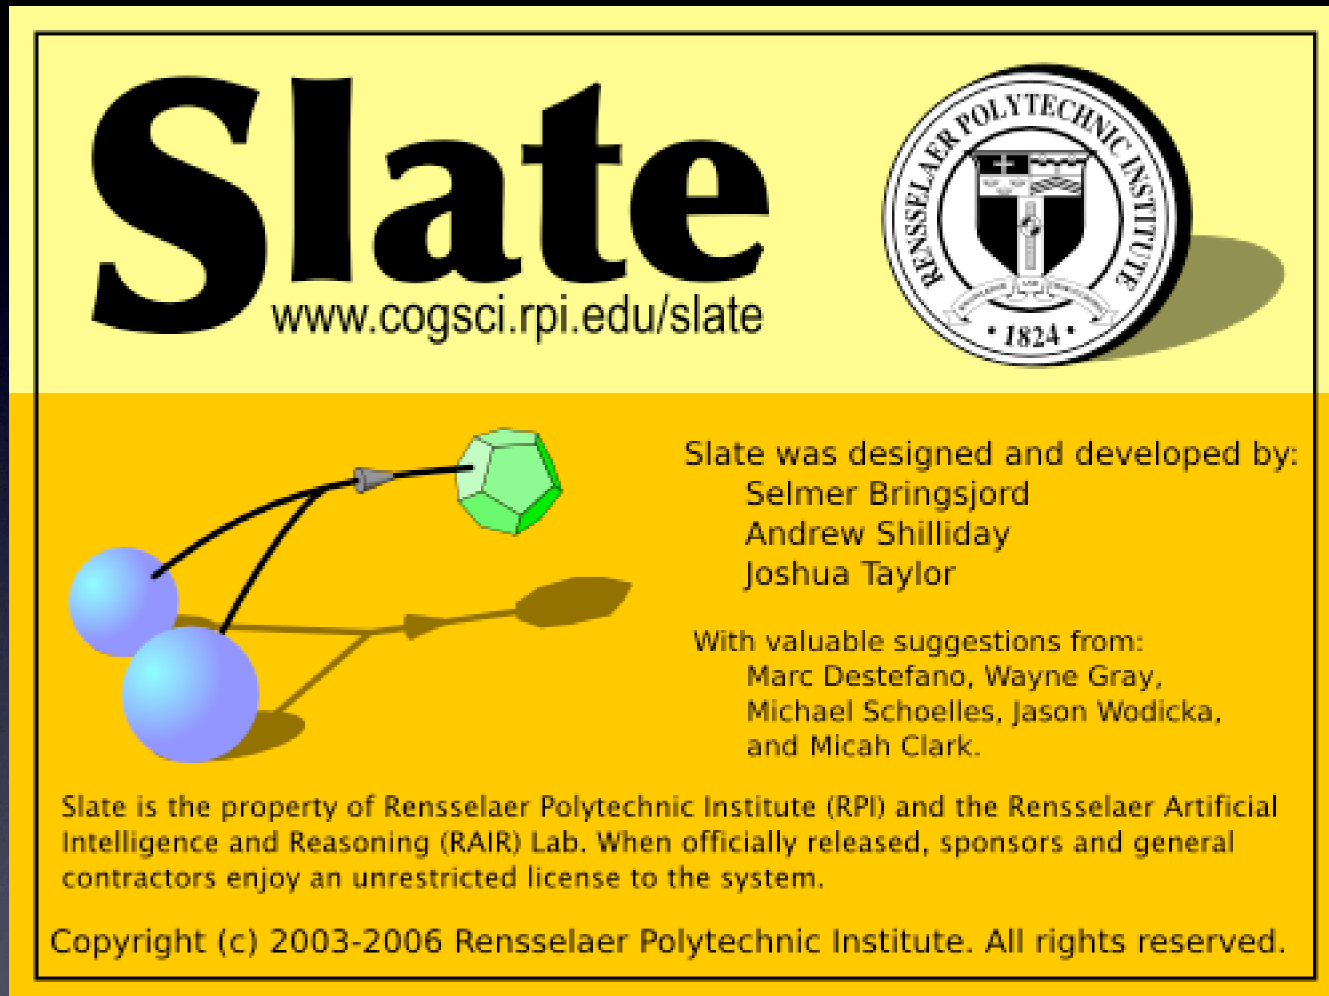

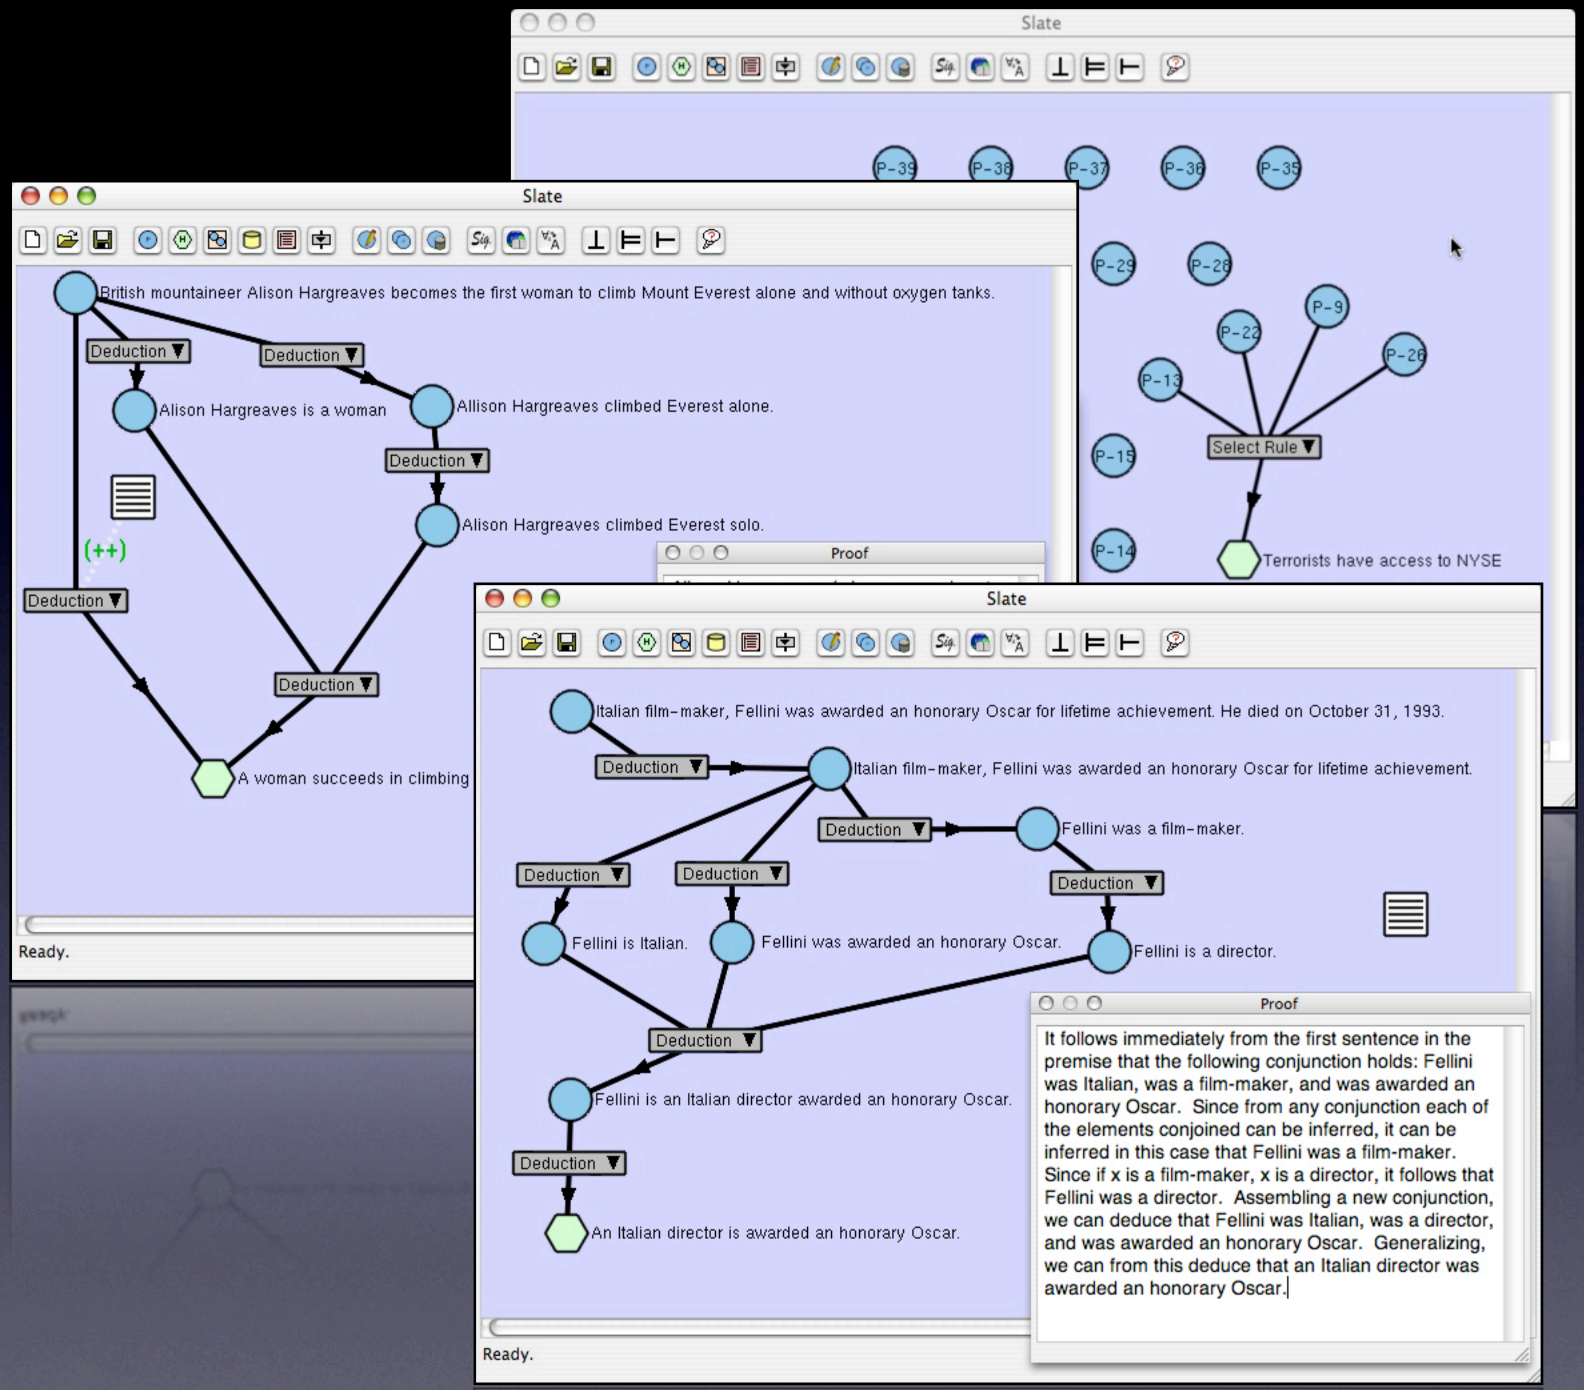

Slate
P-39
P-38
P-37
P-36
P-35
P-29
P-28
P-9
P-22
P-26
P-13
P-15
Select Rule
P-14
Terrorists have access to NYSE
Slate
British mountaineer Alison Hargreaves becomes the first woman to climb Mount Everest alone and without oxygen tanks.
Deduction
Deduction
Alison Hargreaves is a woman
Allison Hargreaves climbed Everest alone.
Deduction
Alison Hargreaves climbed Everest solo.
(++)
Deduction
Deduction
A woman succeeds in climbing
Ready.
Proof
Slate
Italian film-maker, Fellini was awarded an honorary Oscar for lifetime achievement. He died on October 31, 1993.
Deduction
Italian film-maker, Fellini was awarded an honorary Oscar for lifetime achievement.
Deduction
Fellini was a film-maker.
Deduction
Deduction
Deduction
Fellini is Italian.
Fellini was awarded an honorary Oscar.
Fellini is a director.
Deduction
Fellini is an Italian director awarded an honorary Oscar.
Deduction
An Italian director is awarded an honorary Oscar.
Ready.
Proof
It follows immediately from the first sentence in the premise that the following conjunction holds: Fellini was Italian, was a film-maker, and was awarded an honorary Oscar. Since from any conjunction each of the elements conjoined can be inferred, it can be inferred in this case that Fellini was a film-maker. Since if x is a film-maker, x is a director, it follows that Fellini was a director. Assembling a new conjunction, we can deduce that Fellini was Italian, was a director, and was awarded an honorary Oscar. Generalizing, we can from this deduce that an Italian director was awarded an honorary Oscar.

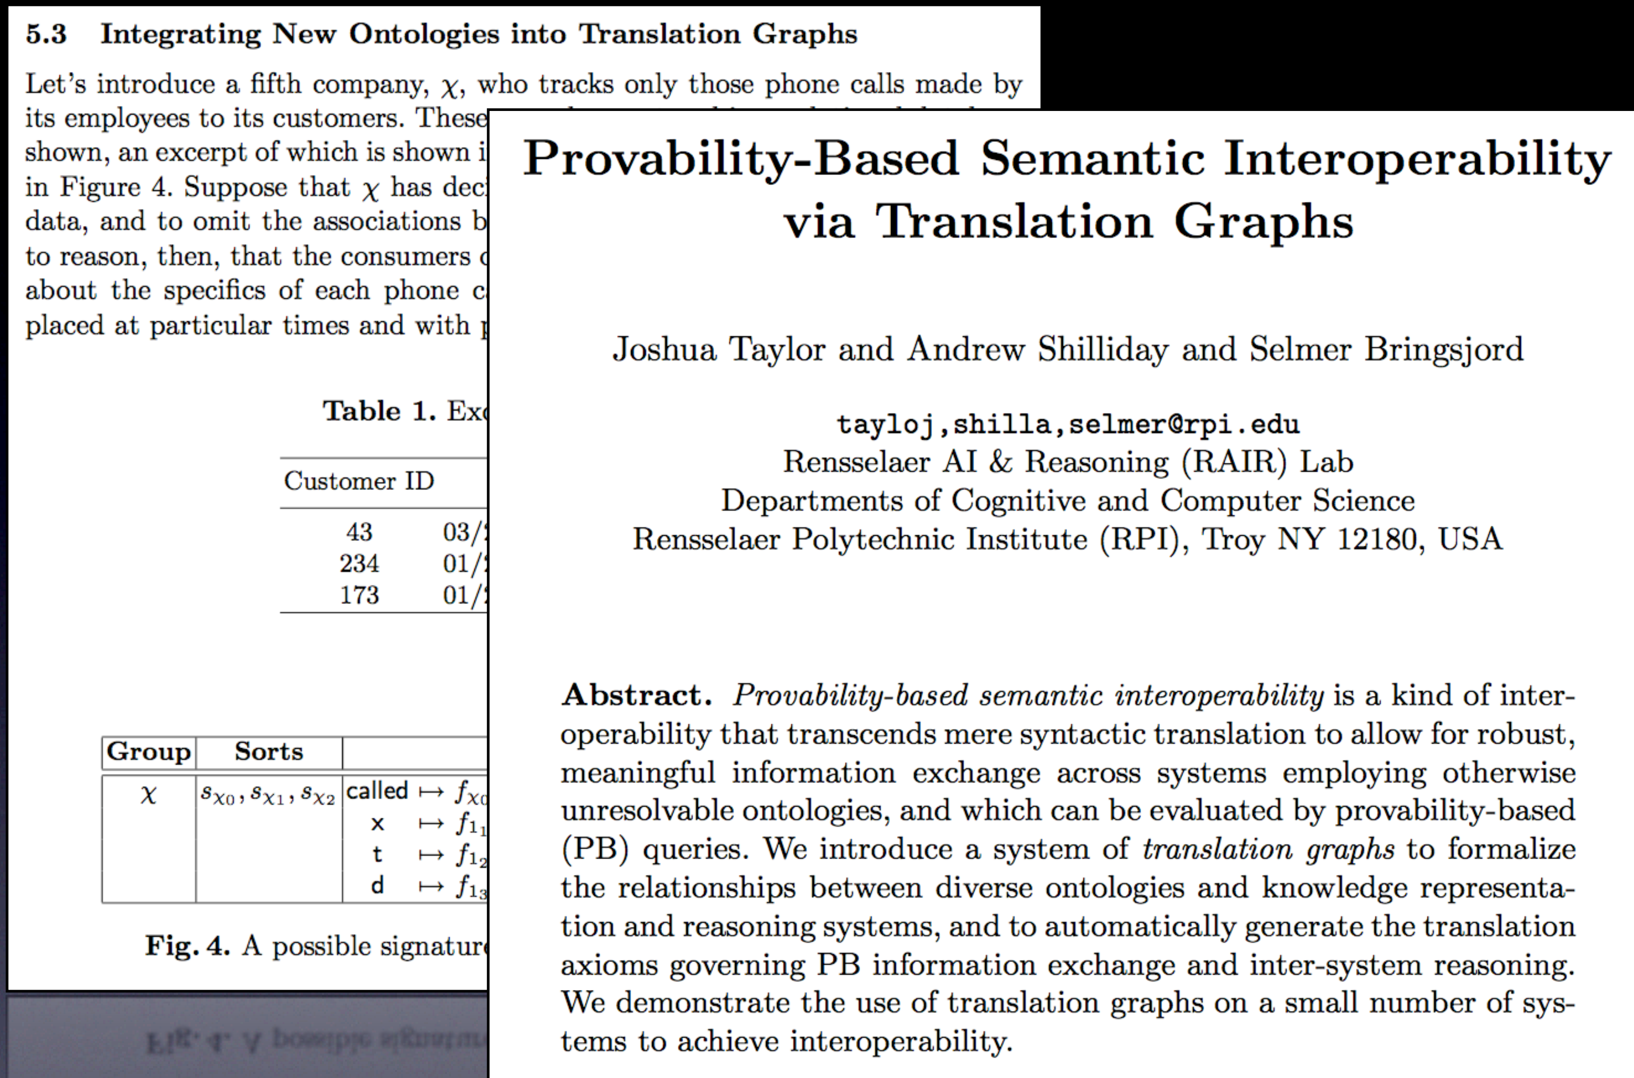

5.3 Integrating New Ontologies into Translation Graphs

Let's introduce a fifth company, χ, who tracks only those phone calls made by its employees to its customers. These

Provability-Based Semantic Interoperability via Translation Graphs

Joshua Taylor and Andrew Shilliday and Selmer Bringsjord

tayloj,shilla,selmer@rpi.edu
Rensselaer AI & Reasoning (RAIR) Lab
Departments of Cognitive and Computer Science
Rensselaer Polytechnic Institute (RPI), Troy NY 12180, USA

Abstract. *Provability-based semantic interoperability* is a kind of interoperability that transcends mere syntactic translation to allow for robust, meaningful information exchange across systems employing otherwise unresolvable ontologies, and which can be evaluated by provability-based (PB) queries. We introduce a system of *translation graphs* to formalize the relationships between diverse ontologies and knowledge representation and reasoning systems, and to automatically generate the translation axioms governing PB information exchange and inter-system reasoning. We demonstrate the use of translation graphs on a small number of systems to achieve interoperability.

Provability-Based Semantic Interoperability via Translation Graphs
for ONISW2007

Provability-Based Semantic Interoperability via Translation Graphs introduces:

Provability-Based Semantic Interoperability (PBSI), a description of interoperability at the semantic level, and why it can only be achieved using provability based techniques.

Translation Graphs, a representation agnostic tool for bridging ontologies and automatically extracting bridging axioms and translation procedures.

Mental MetaLogic Reasoning in Slate

In Slate, items in System *S* are connected with argument links to graphically depict an argument from some set of premises to a particular conclusion. Arguments can be supported by witness objects, viz. models, proofs or databases.

This mechanism can be used to simulate model-based reasoning in Slate. This process of heterogeneous reasoning is critical to the emergence of the Multi-Mind Effect.

Multi-Agent Reasoning in Slate

Slate can be used to model multi-agent reasoning analogous to the interactions between human reasoners.

Given translation graphs, the relationships between the representations used by the different agents can be explored in Slate, and a process for reconciling the representations can be constructed. A set of bridging axioms can be automatically extracted from this translation graph enabling information exchange at the semantic level.

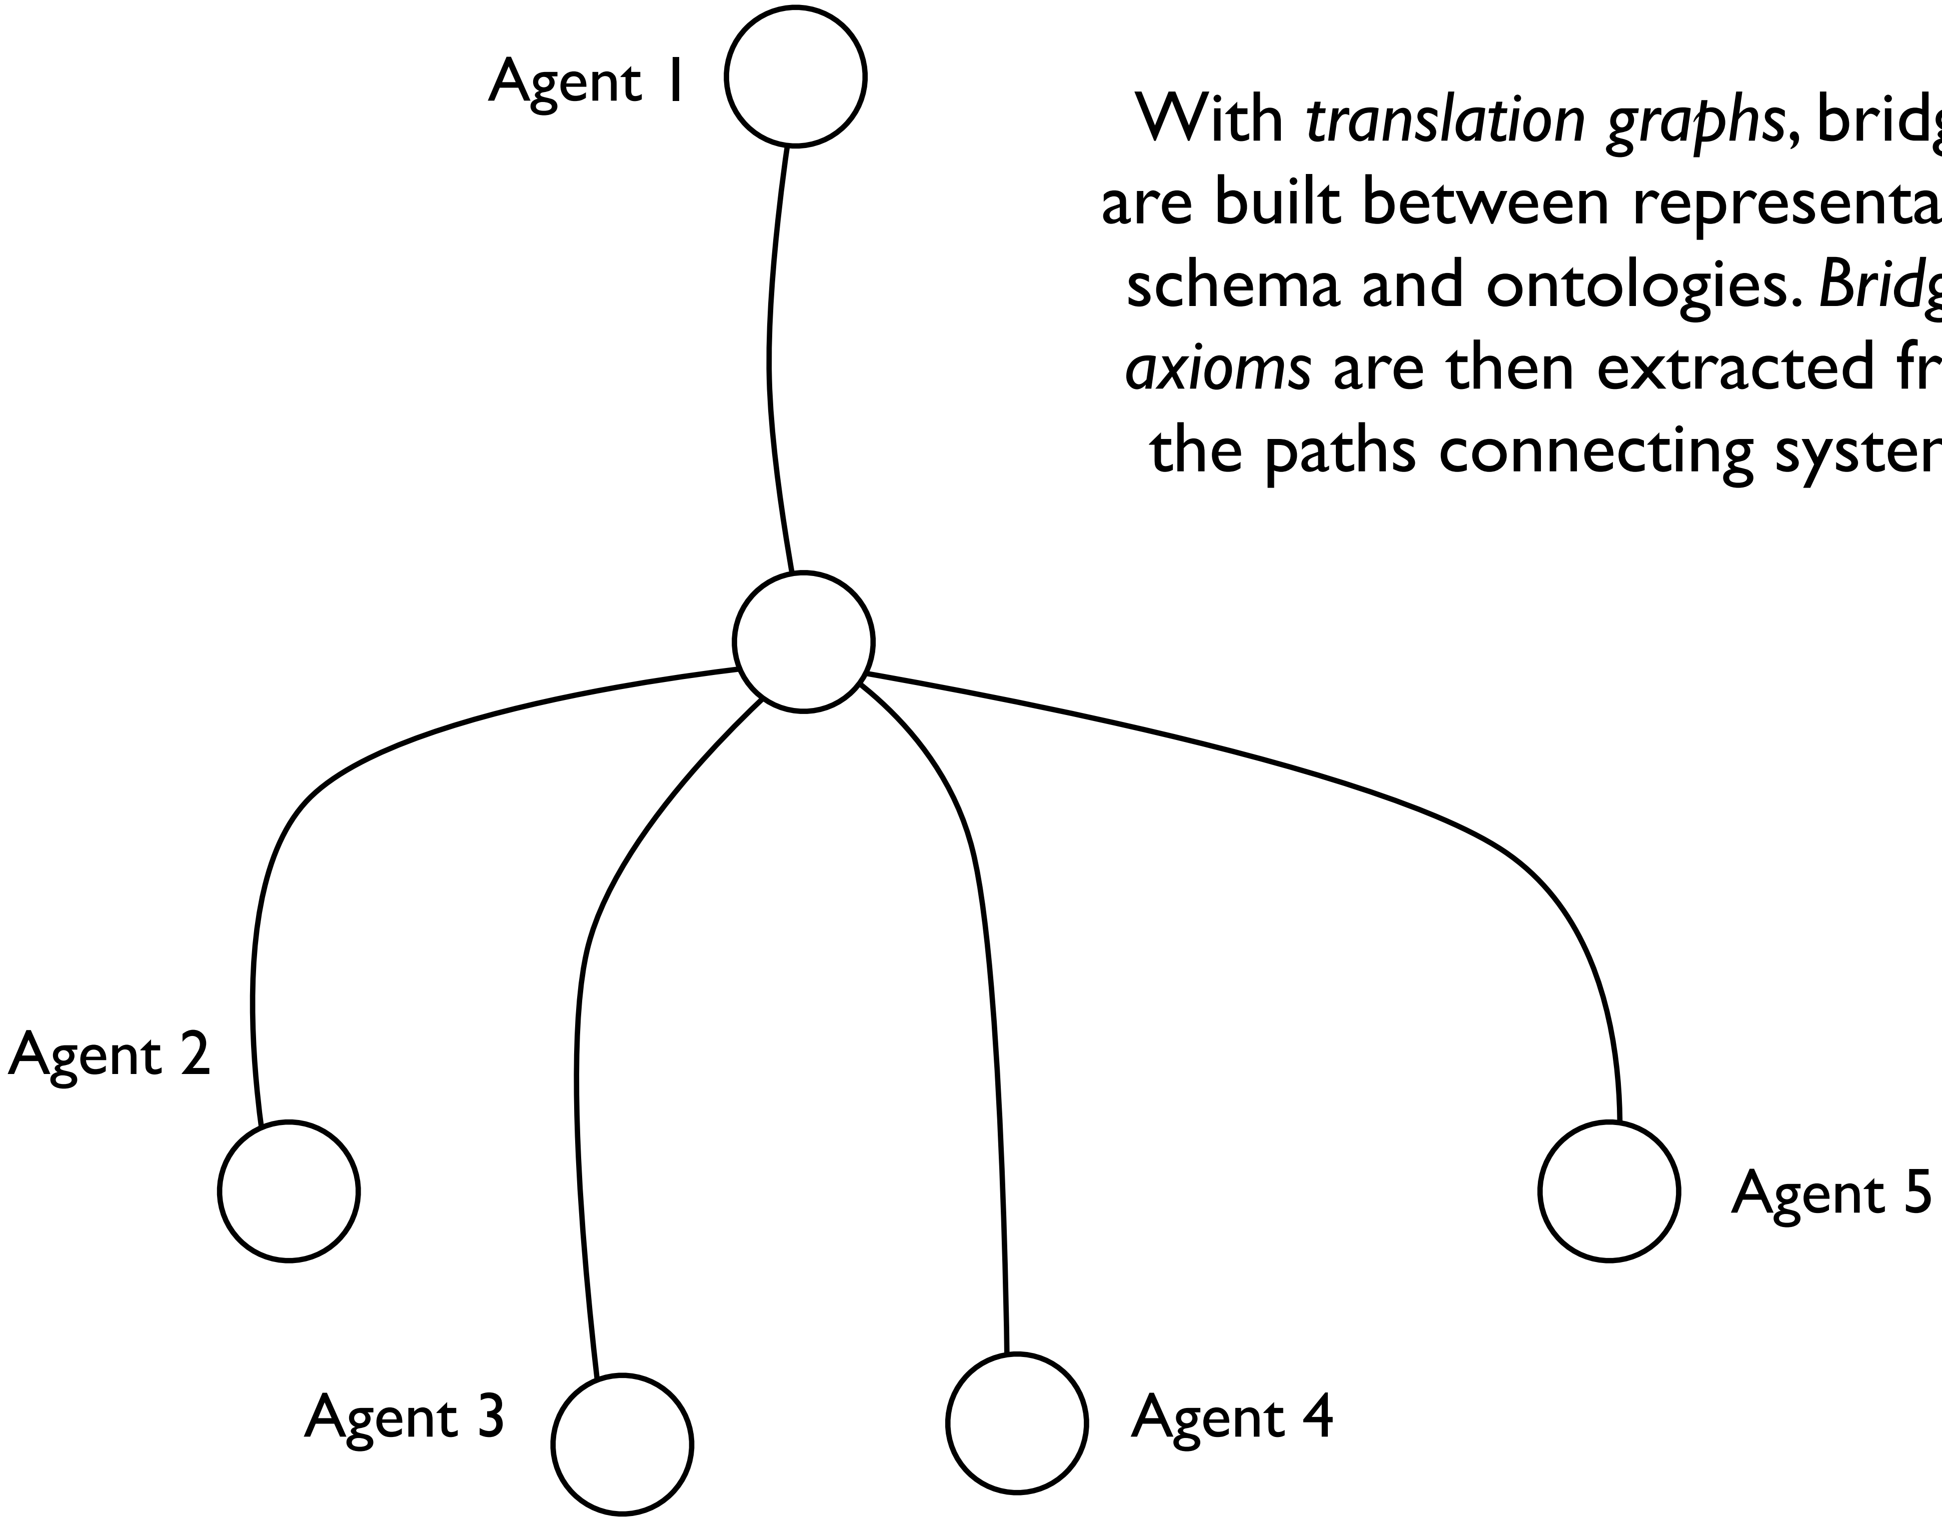

With *translation graphs*, bridges are built between representation schema and ontologies. *Bridging axioms* are then extracted from the paths connecting systems.

With *translation graphs*, bridges are built between representation schema and ontologies. *Bridging axioms* are then extracted from the paths connecting systems.

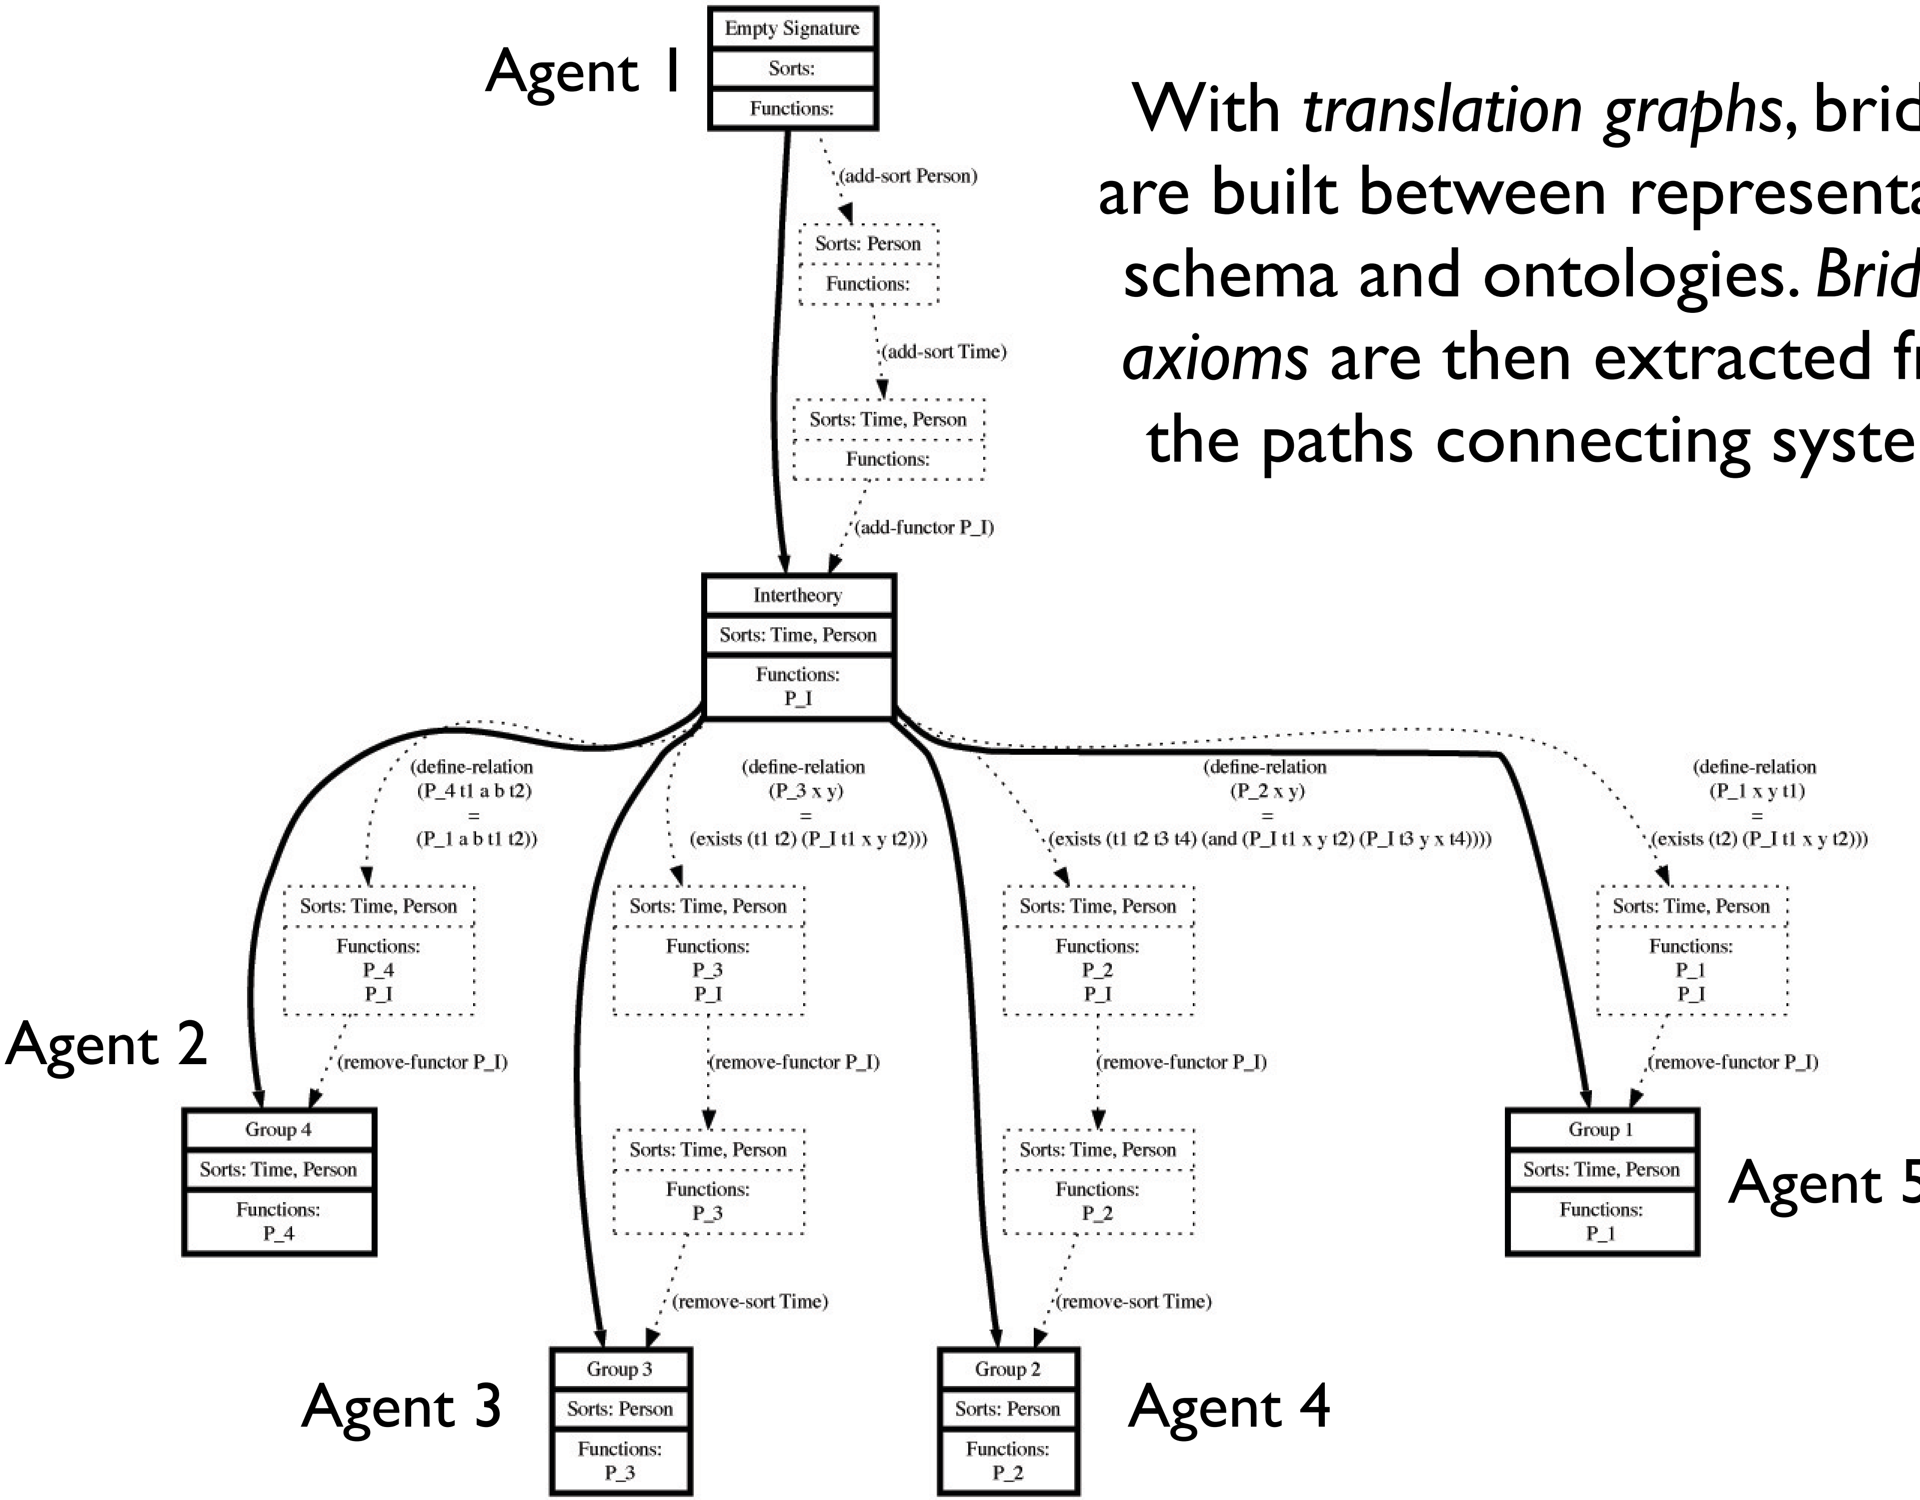

Pedagogical Implications of the Multi-Mind Effect

The Multi-Mind Effect can be very effective in creating tools that leverage multiple forms of reasoning to engage in context-independent, normatively correct reasoning. These tools can be used to improve human and machine reasoning.

It can also be of importance in decision-making, where using only one representation or one type of reasoning can lead to erroneous conclusions.

Next Steps

To study the Multi-Mind Effect in a extremely rigorous manner, through controlled experiments.

To precisely model the Multi-Mind Effect in Slate, following up on work previously done.