Analysis of Competing Hypotheses using Subjective Logic (ACH-SL)

Simon Pope and Audun Jøsang,

DSTC, CRC for Enterprise Distributed Systems Technology

ICCRTS, 13-16 June 2005

ACH-SL Overview

- **Assists analysts to assess competing hypotheses** (possible outcomes, courses of action, etc.) by constructing a model of impact and relevance of evidence.
- **Designed for strategic and tactical intelligence analysis** (national security, defence, trade, etc.) and law enforcement investigations.
- **Reduces cognitive load of analysts**, allowing them to focus on the judgements they make and produce higher quality analysis.
- **Extension of Heuer's original ACH approach** to intelligence analysis.
- **Uses the *Subjective Logic* calculus** for approximate reasoning under conditions of uncertainty.
- **Implemented in DSTC's S*h*EBA technology** for structured, evidence based analysis of hypotheses.

Unassisted human reasoning is not robust enough for intelligence analysis

- Humans systematically make substantive errors in reasoning due to:
  - problems of framing;
  - resistance of mental models to change;
  - risk aversion;
  - limitations of short-term memory;
  - other cognitive and perceptual biases.
- Consequences of intelligence analysis failures can be astronomically high.
  - e.g. Pearl Harbour, Midway Island
- Tools and technology are needed to augment human reasoning and mitigate cognitive defects.

ShEBA – Basic Concept

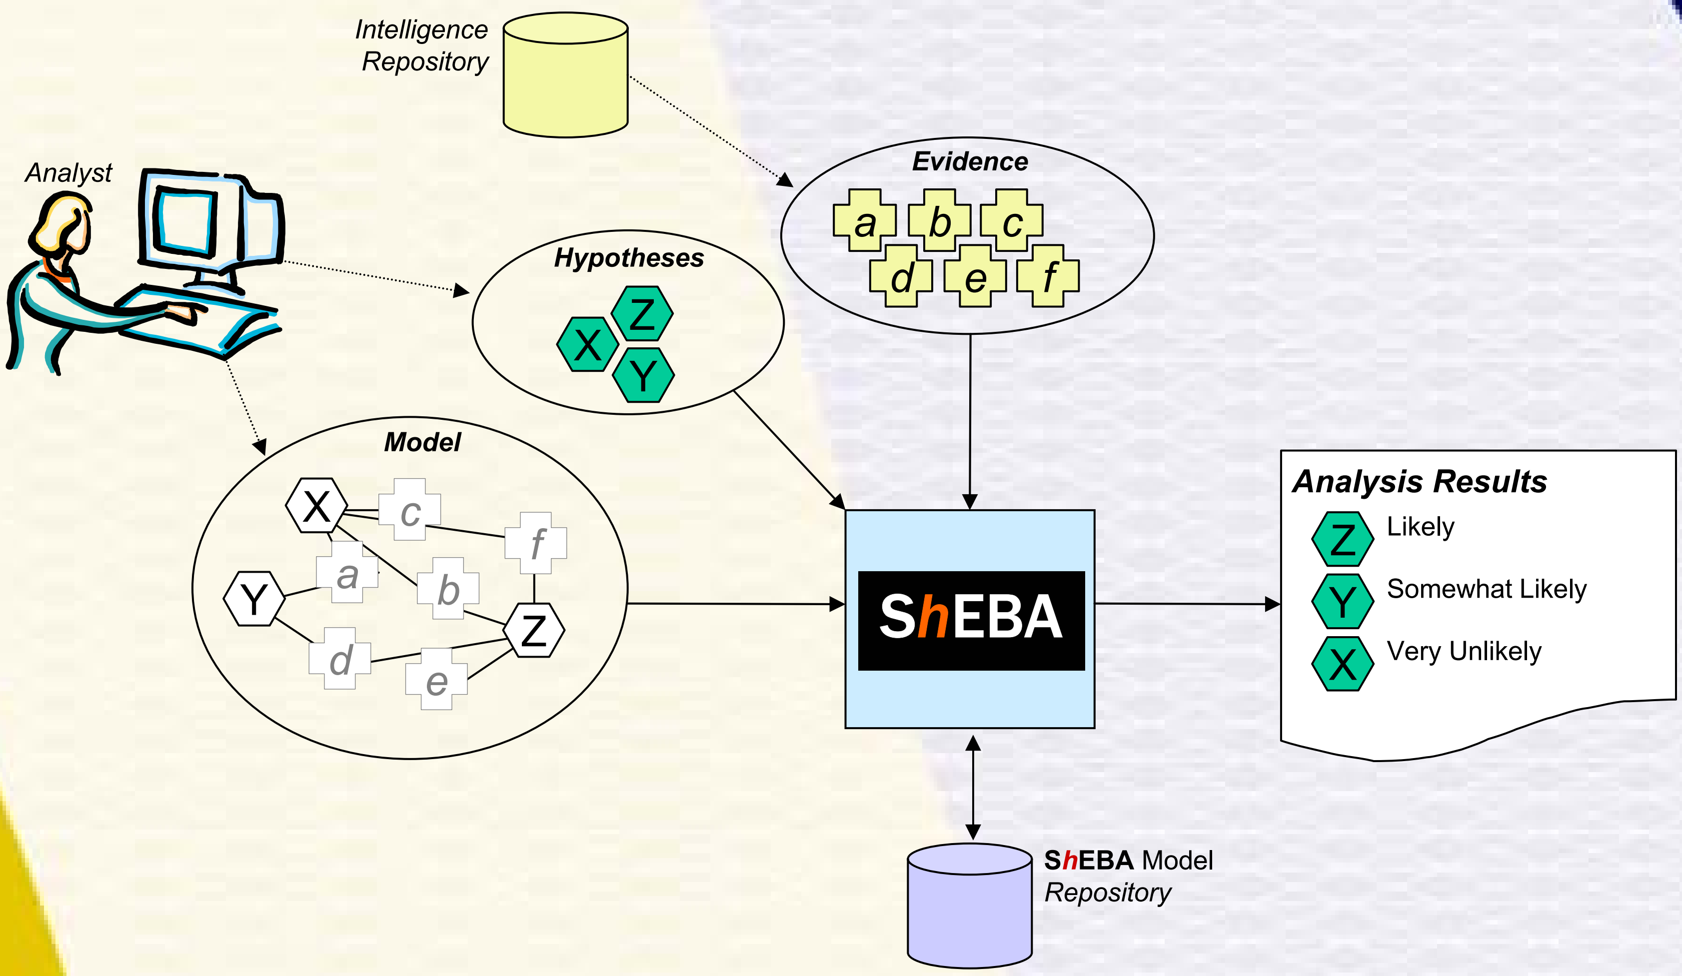

S*h*EBA – Features

- Uses *Subjective Logic* – a formal mathematical calculus that explicitly deals with uncertainty.
  - Ideal for human- and machine-interpretable representations of belief.
  - Input and results can be translated to and from ‘everday’ human terms, allowing analysis results to be more easily translated into language for presentation to policy makers and other non-analysts.
- Analytical models involving large collections of evidence and hypotheses can be created and stored separately from the *actual* evidence used for reasoning.
  - Models can be adjusted to adapt to unfolding situations
  - Automatic re-evaluation of models when reliability and value of evidence changes.
- Java-based technology
  - Platform-independent.

Belief representation

- Generalization of binary logic and probability calculus.
- *Subjective Logic* beliefs are expressed as: $B(x) = (b, d, u, a)$
  - b: belief
  - d: disbelief
  - u: uncertainty
  - a: relative atomicity

  } in range [0,1]

  Where: $b + d + u = 1$
- Expectation value: $E(x) = b+au$
- 1:1 correspondence with Bayesian representation
  - i.e. (r,s) ≡ (b,d,u)

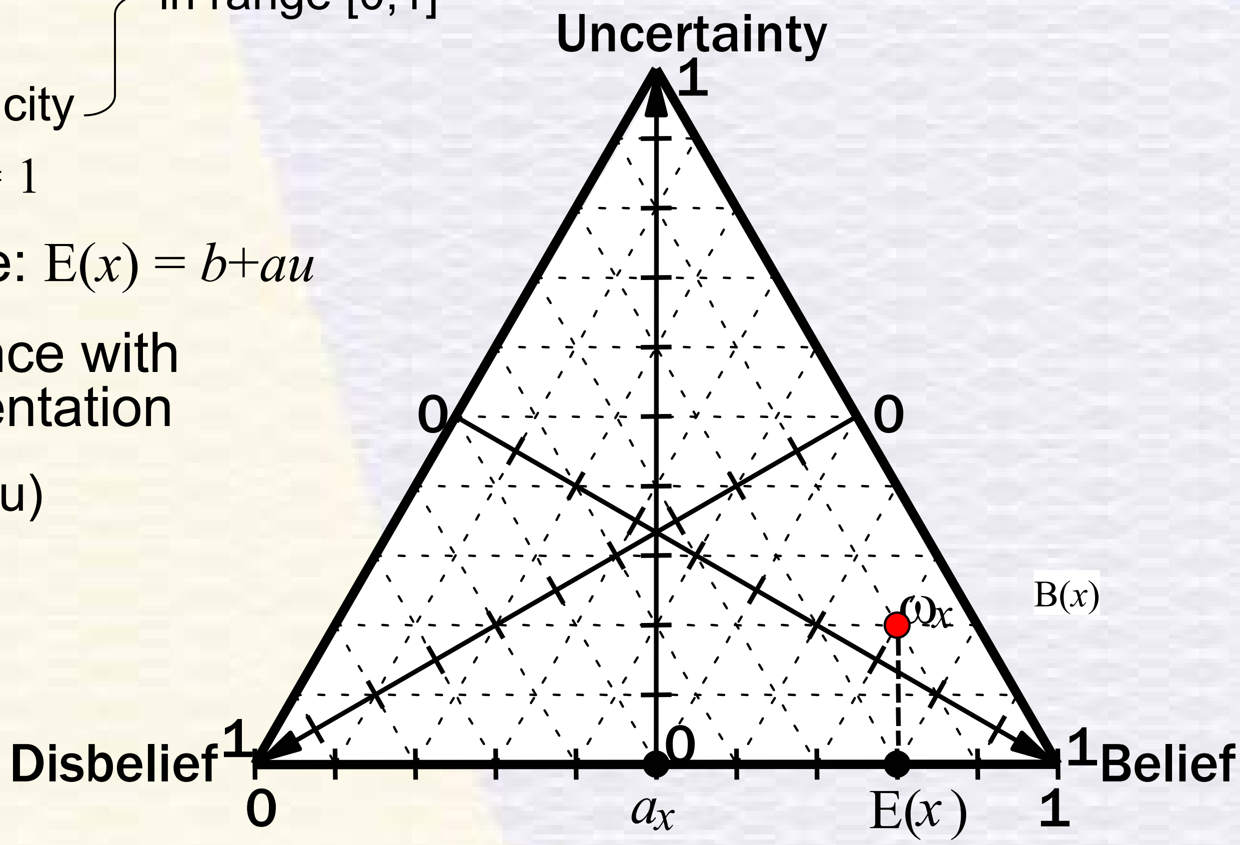

Visualization of Beliefs

Opinion Triangle representation

The three corners of the triangle represent absolute values (d=1, u=1 and b=1 clockwise from left to right). Any opinion may be expressed as a (b,d,u) tuple where b+d+u = 1.

The red **atomicity line** indicates the atomicity of the opinions .The atomicity represents the portion of the state space that is accounted for by the belief, or can represent the base rate as appropriate.

The point at which the line extending from each point meets the baseline indicates the opinion **expectation**. It is always parallel to the atomicity line.

In the example shown, the atomicity of the opinions (*x*, *y* and *z*) are all 0.5, indicating that each accounts for half of all states within their respective state spaces.

Expectation Bar representation

The length of the coloured area of the bar represents the expectation, while the dark blue area represents the portion accounted for by **belief**, and the light blue area represents the portion accounted for by the uncertainty (equal to the *uncertainty* x *atomicity*).

Bayesian Bar representation

This provides an visualization expressed as a bayesian tuple (r,s), where the length of the **green bar** is the number of positive samples, *r;* and the length of the **red bar** is the number of negative samples, *s*.

The **yellow bar** in between expresses the amount of uncertainty, and the distance of the **vertical line** is the expectation.

Fuzzy Bar representation

The length of the coloured bar indicates the expectation, while the text is the member of the fuzzy set that is closest to the actual value.

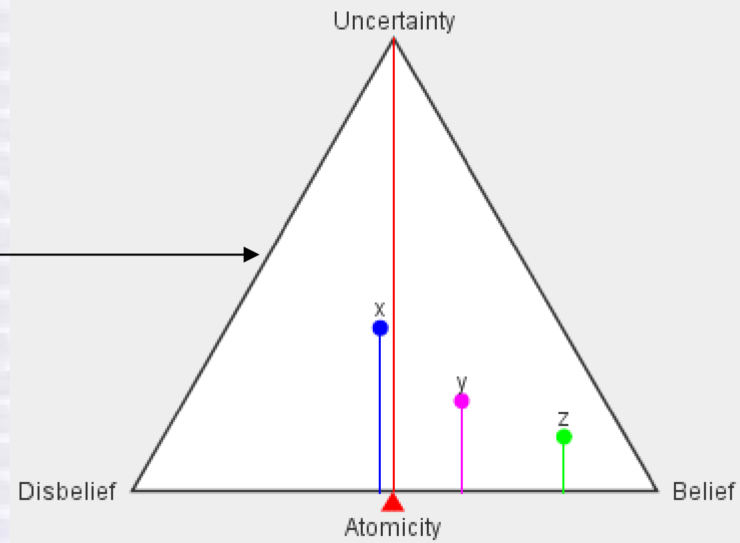

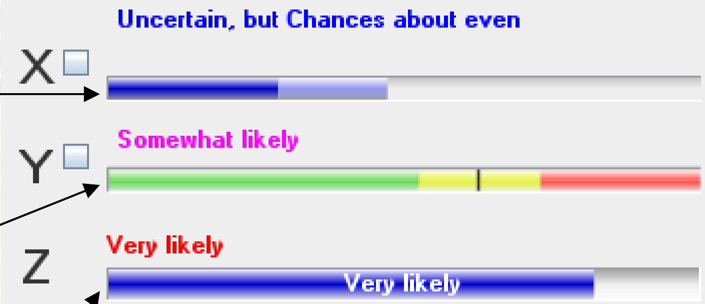

Opinion about	X	Y	Z
Belief	0.29	0.53	0.76
Disbelief	0.35	0.27	0.12
Uncertainty	0.37	0.21	0.13
Atomicity	0.50	0.50	0.50
Expectation	0.47	0.63	0.82

Step-By-Step Outline
(Abridged version)

1. Identify the hypotheses.
2. Prepare an analysis model consisting of:
   a) A set of exhaustive and exclusive hypotheses.
   b) A set of items of evidence that are of possible relevance.
3. Complete the analysis model:
   a) Decide if evidence is ***causal*** or ***derivative***.
   b) Make appropriate judgments about the hypotheses in respect of the evidence.
4. Enter the *actual* values of the evidence into the model and compute the overall likelihood for each hypothesis.
5. Analyse how sensitive the conclusion is to a few critical items of evidence. Changes in the value of evidence with high diagnosticity will alter the calculated likelihoods more than evidence with low diagnosticity.

About Evidence...

Causal evidence
directly influences the likelihood of one or more hypotheses.

***Deductive* reasoning uses likelihood of each hypothesis‡, for each piece of evidence, i.e. p(h|e) and p(h|¬e).**

Derivative evidence
is usually observed in conjunction with one or more hypotheses.

***Abductive* reasoning uses likelihood of evidence‡, for each hypothesis, i.e. p(e|h) and p(e|¬h).**

- To avoid erroneous conclusions, both *affirmative* and *negative* cases must be considered.
  - Reasoning about counterfactual conditions is absolutely necessary to avoid base rate errors, so either:

    p(h|e), p(h|¬e),*or*

    p(e|h), p(e|¬h) are necessary.

‡ *plus knowledge of the base rates of the hypotheses and evidence.*

Deductive vs. Abductive reasoning

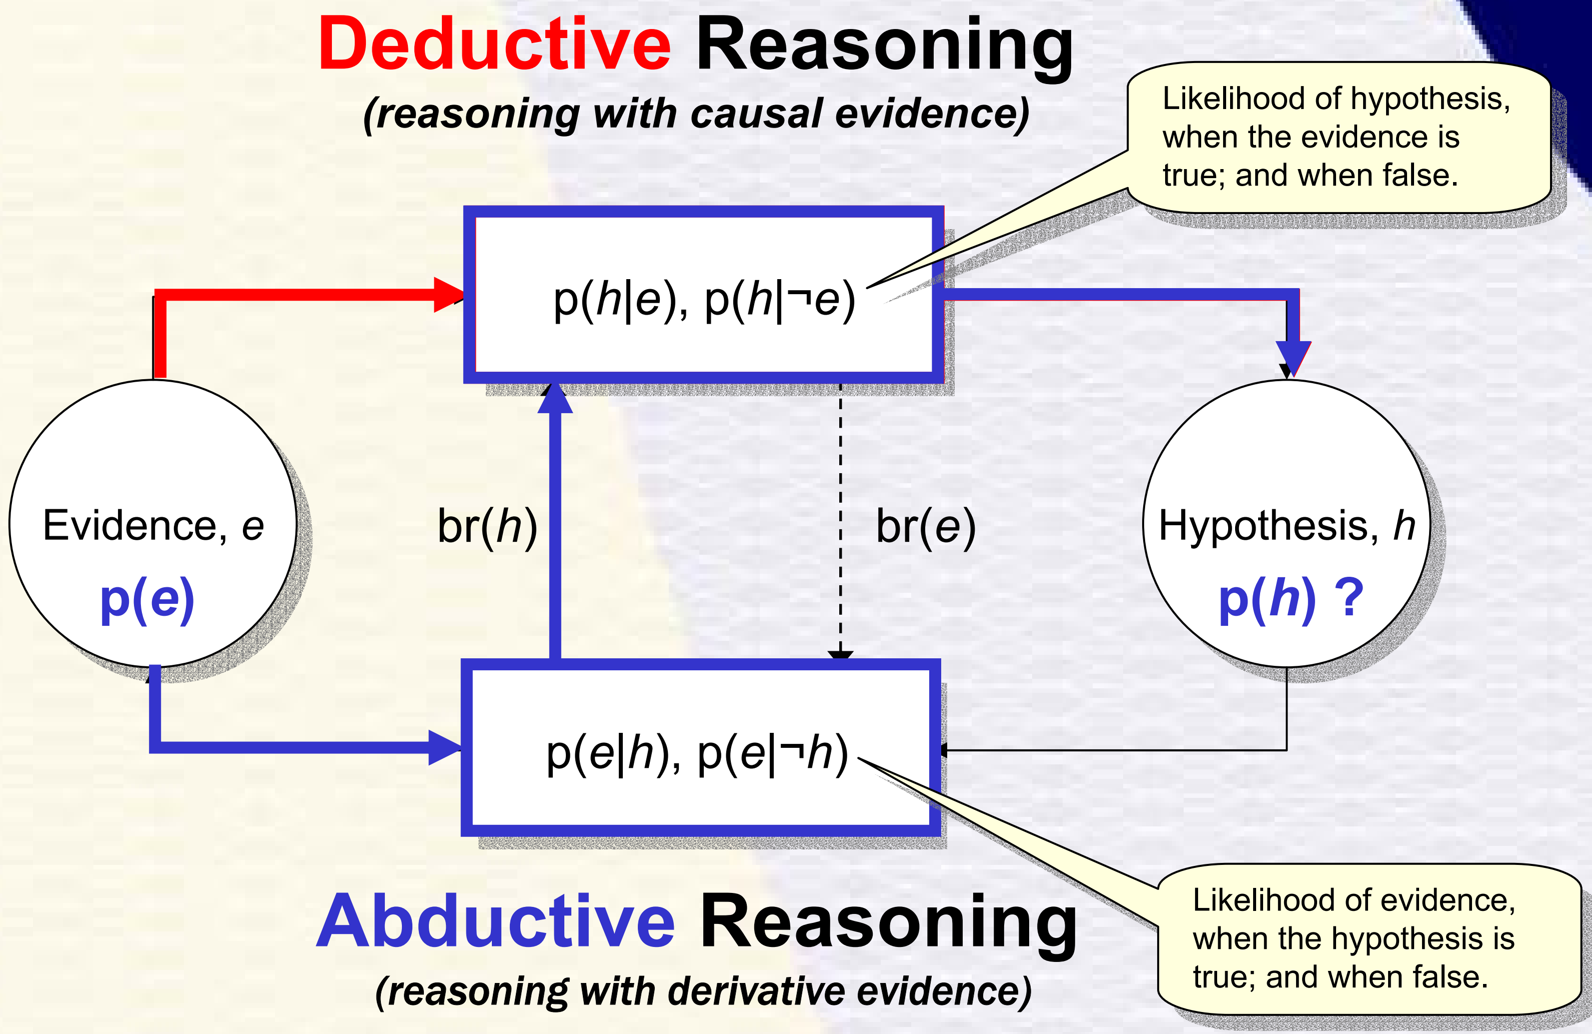

Example

- ***Background:***
  - In March 1993, an assassination attempt by Iraqi agents on George Bush (Sr.) was foiled by the United States intelligence services.
  - In April, the US considered launching a retaliatory air strike on the Iraqi Intelligence Headquarters in Baghdad.

- ***Question:***
  - “What will be Iraq’s response for the bombing of its Intelligence Headquarters?”

Example (cont'd)

Exhaustive and exclusive† *hypotheses:*

- **No retaliation –** *No military/terrorist action*
- **Minor retaliation –** *Sponsor some form of minor terrorist act without direct involvement*
- **Major Retaliation –** *Directly undertake some form of terrorist act, probably major in scope, such as an attempted bombing of a US Intelligence agency, within the United States.*

Evidence to be considered, whether:

1. There was an Iraqi-sponsored terrorist offensive during the First Gulf War.
2. The Iraqi Governing Council wants to provoke the United States into further action.
3. Saddam would personally lose face with his own people if he did not retaliate after being attacked.
4. Saddam makes a public statement to 'not retaliate if attacked'.
5. Iraqi embassies are instructed to take increased security precautions.
6. There is an increase in traffic of monitored communications between Iraqi handlers and their agents.

† ***i.e. one, and only one, of the hypotheses must be true.***

Example – Causal Evidence

- **The Iraqi Governing Council wants to provoke the United States.**
  - **Affirmative**: Given that the Iraqi Governing Council wants to provoke the U.S. into further action, what is the likelihood that there will be 'No Retaliation'.
  - **Negative**: Given that the Iraqi Governing Council does not want to provoke the U.S. into further action, what is the likelihood that there will be 'No Retaliation'.
- Positive conditional considerations – $p(h|e)$
  - The desire to provoke further reaction *significantly decreases* the likelihood of the 'No Retaliation' hypothesis (and therefore increases the likelihood of some retaliation).
- Negative conditional considerations – $p(h|\neg e)$
  - The desire to not provoke further reaction *slightly increases* the likelihood of the 'No retaliation' hypothesis (and therefore decreases the likelihood of retaliation).

Example – Derivative Evidence

- **Iraqi embassies are instructed to take increased security precautions.**
  - **Affirmative**: Given that there is a 'Major Retaliation', what is the likelihood that Iraqi embassies are instructed to take increased security precautions.
  - **Negative**: Given that there is no 'Major Retaliation', what is the likelihood that Iraqi embassies are instructed to take increased security precautions.
- Positive conditional considerations – $p(e|h)$
  - Increase in security precautions at Iraqi embassies are associated with high-risk Iraqi actions, such as the commencement of major military operations.
- Negative conditional considerations – $p(e|\neg h)$
  - Increase in security precautions at Iraqi embassies are also associated with increased internal security threats, such as Shi'ite uprisings and Kurdish guerrillas.

How to derive belief for a hypothesis (h_1)

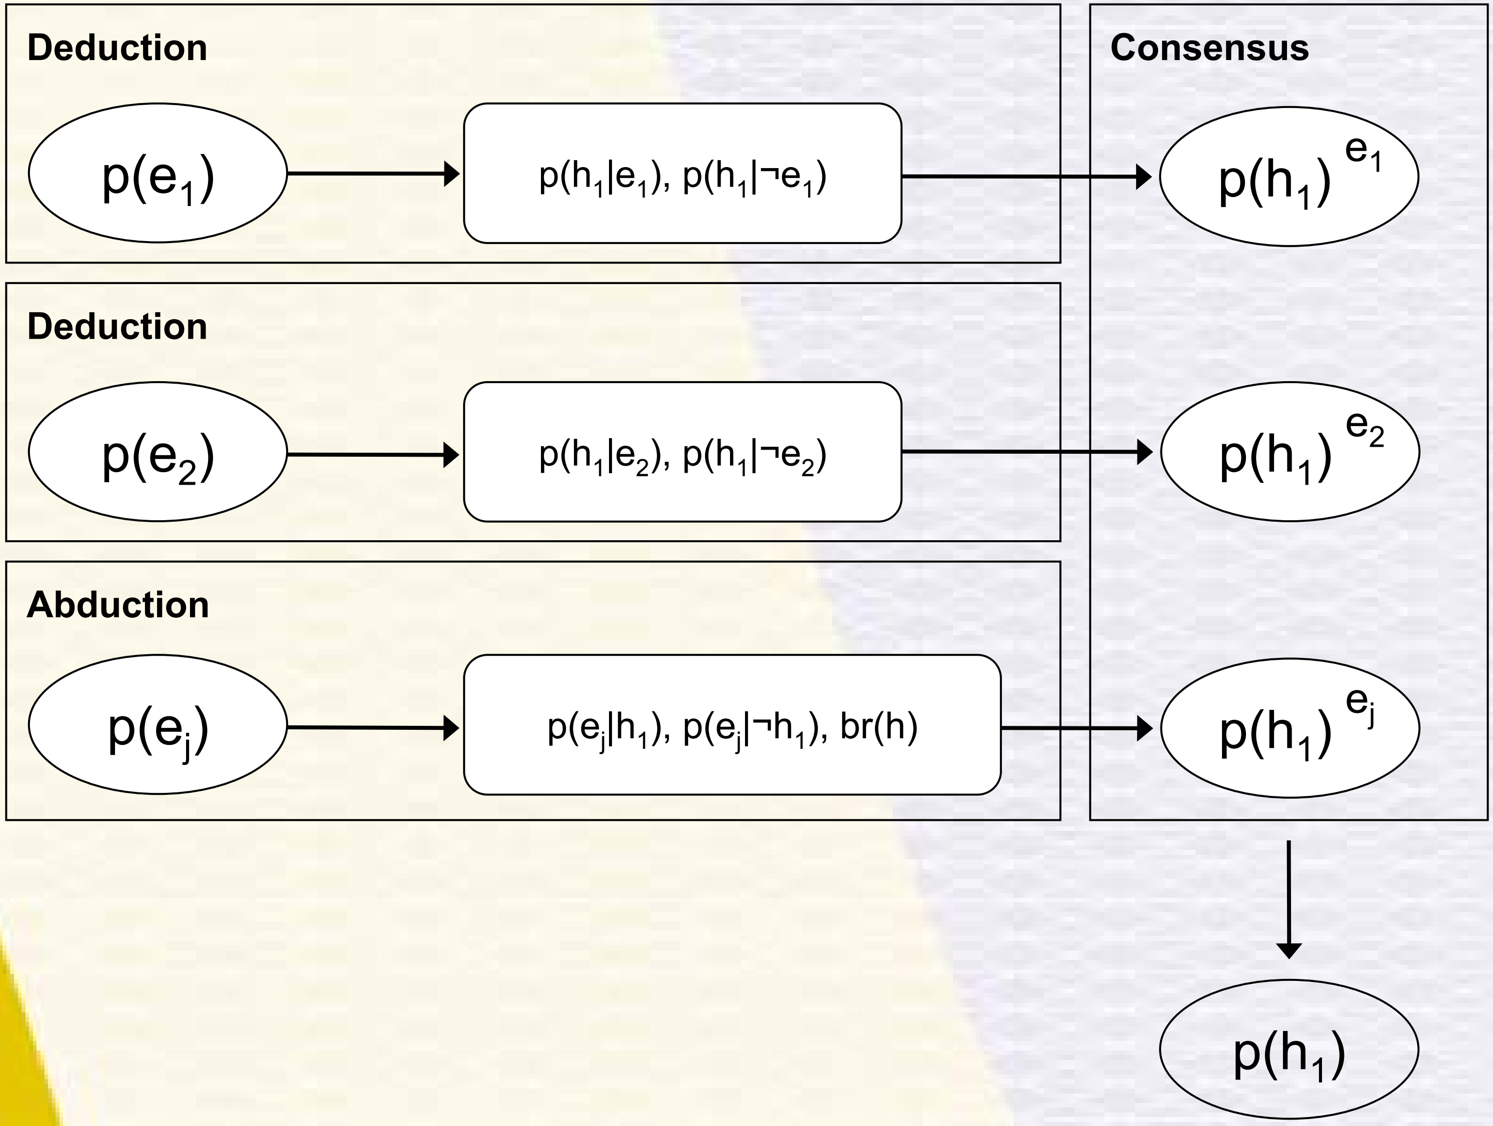

Evidence Diagnosticity

- Measures how well an item of evidence can distinguish between some or all hypotheses.
- Provides ***a priori*** indication of what evidence is most important in determining the conclusions.
- Derived from “first-order” judgements made by the analyst and accurately reflects base rates and counterfactual conditions.
  - Does not require the value of evidence to be known – diagnosticity can be derived from the model with complete ignorance of the actual value of the evidence.
  - With other approaches, *diagnosticity* is usually a “second-order” cognitive task requiring greater cognitive effort, with increased risk of errors due to cognitive limitations.
  - Avoids problems with analysts ignoring counterfactuals and base rates and the uncertainty of their impact on diagnosticity.

Conclusion

- **Facilitates higher-quality intelligence analysis**
  - Reduction in cognitive effort allows analysts to focus on the judgements they make, and the tasks for which they are experts.
- **Sound mathematical base for reasoning**
  - Built on a formal mathematical calculus, *Subjective Logic,* that explicitly deals with uncertainty.
- **Better ‘Transparent Analysis’**
  - Input and results can be translated to and from ‘everday’ human terms, allowing analysis results to be more easily translated into language for presentation to policy makers and other non-analysts.
- **Integration with existing intelligence processes**
  - Subjective Logic calculus provides ability to interface with collection, database and other analytical systems, such as evidentiary and investigative systems, to support the organisation’s wider intelligence processes.
- **Implemented in DSTC’s S*h*EBA technology**