Developing Default Logic as a Theory of Reasons

John Horty

Netherlands Institute for Advanced Study (NIAS)
and
University of Maryland

horty@umiacs.umd.edu
www.umiacs.umd.edu/users/horty

Introduction

1. Tools for understanding deliberation/justification:

   Standard deductive logic

   Decision theory (plus probability, inductive logic)

2. But ordinairly, we seem to focus on *reasons*—in both deliberation and justification

3. Examples:

   We should eat at Obelisk tonight

   Racoons have been in the back yard again

4. This could be an allusion, or an abbreviation, or heuristic

   But it could also be right . . .

   . . . an idea common in epistemology, and especially in ethics

5. Common questions about reasons:

    Relation between reasons and motivation?

    Relation between reasons and desires?

    Relation between reasons and values?

    Objectivity of reason?

6. A different question:

    How do reasons support actions or conclusions?

    What is the mechanism of support?

7. My answer:

    Reasons are (provided by) defaults

    The *logic of defaults* tells us how reasons support conclusions

8. Talk outline:

- Prioritized default logic
  - Extensions, scenarios
  - Triggering, conflict, defeat
  - Binding defaults, proper scenarios
- Elaborating the theory
  - Variable priorities
  - Undercutting (exclusionary) defeat
- Applications and open questions
  - Exclusion and priorities
  - Exclusion by weaker defaults
  - Floating conclusions

Fixed priority default theories

1. Notation:

   Propositions: $A, B, C, \ldots, \top$

   Background language: $\wedge$, $\vee$, $\neg$, $\Rightarrow$

   Consequence: $\vdash$

   Logical closure: $Th(\mathcal{E}) = \{A : \mathcal{E} \vdash A\}$

2. Example:

   > Tweety is a bird
   > Therefore, Tweety is able to fly

   Why? There is a default that birds fly

   > Tweety is a bird
   > Tweety is a penguin
   > Therefore, Tweety is not able to fly

   Because there is a (stronger) default that penguins don't fly

3. Default rules: $X \to Y$

   Example: $B(t) \to F(t)$

   Instance of: $B(x) \to F(x)$ ("Birds fly")

4. Premise and conclusion:

   If $\delta = X \rightarrow Y$, then

$$Prem(\delta) = X$$
$$Conc(\delta) = Y$$

   If $\mathcal{D}$ set of defaults, then

$$Conc(\mathcal{D}) = \{Conc(\delta) : \delta \in \mathcal{D}\}$$

5. Priority ordering on defaults (strict, partial)

   $\delta < \delta'$ means: δ' stronger than δ

6. Priorities have different sources:

   Specificity
   Reliability
   Authority
   Our own reasoning

   For now, take priorities as fixed, leading to . . .

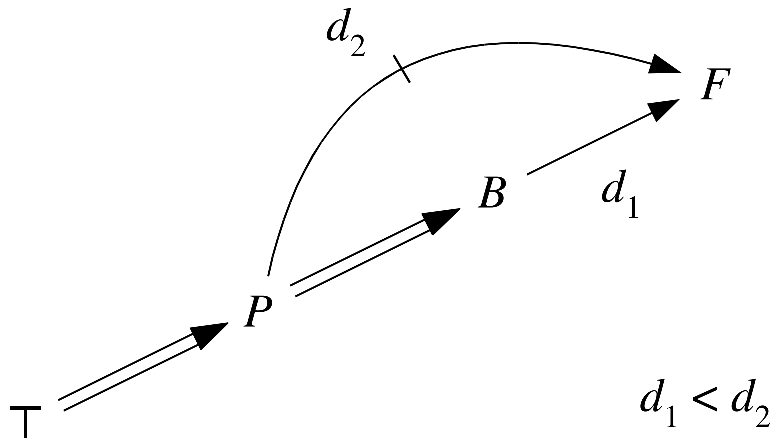

7. A *fixed priority default theory* is a tuple

$$\langle \mathcal{W}, \mathcal{D}, < \rangle$$

where $\mathcal{W}$ contains ordinary statements, $\mathcal{D}$ contains defaults, and $<$ is an ordering

8. Example (Tweety Triangle):

$$\begin{array}{lcl} \mathcal{W} & = & \{P, P \Rightarrow B\} \\ \mathcal{D} & = & \{\delta_1, \delta_2\} \\ \delta_1 & = & B \rightarrow F \\ \delta_2 & = & P \rightarrow \neg F \\ \delta_1 < \delta_2 & & \end{array}$$

(P = Penguin, B = Bird, F = Flies)

9. Main question: what can we conclude from such a theory?

10. An *extension* $\mathcal{E}$ of $\langle \mathcal{W}, \mathcal{D}, < \rangle$ is a belief set an ideal reasoner might settle on, based this information

    Usually defined directly, but we take roundabout route . . .

11. A *scenario* based on $\langle \mathcal{W}, \mathcal{D}, < \rangle$ is some subset $\mathcal{S}$ of the defaults $\mathcal{D}$

12. A *proper scenario* is the “right” subset of defaults

13. An *extension* $\mathcal{E}$ based on $\langle \mathcal{W}, \mathcal{D}, < \rangle$ is a set

$$\mathcal{E} = Th(\mathcal{W} \cup Conc(\mathcal{S}))$$

    where $\mathcal{S}$ is a proper scenario

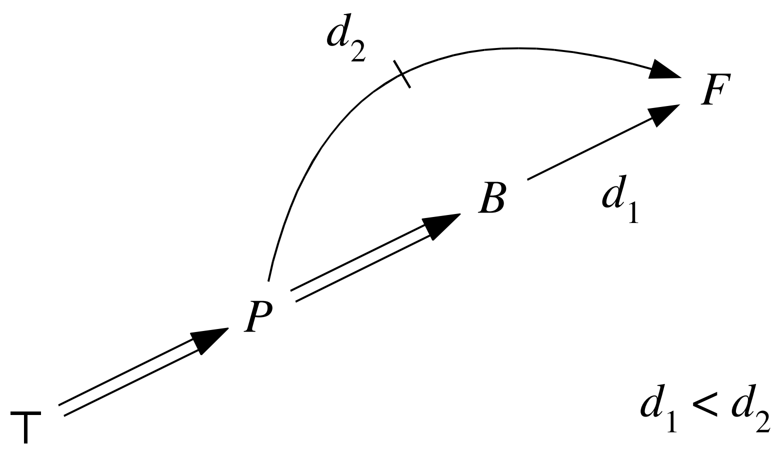

14. Returning to example: $\langle \mathcal{W}, \mathcal{D}, < \rangle$ where

$$\begin{array}{lcl} \mathcal{W} & = & \{P, P \Rightarrow B\} \\ \mathcal{D} & = & \{\delta_1, \delta_2\} \\ \delta_1 & = & B \rightarrow F \\ \delta_2 & = & P \rightarrow \neg F \\ \delta_1 < \delta_2 & & \end{array}$$

Four possible scenarios:

$$\begin{array}{l} \mathcal{S}_1 = \emptyset \\ \mathcal{S}_2 = \{\delta_1\} \\ \mathcal{S}_3 = \{\delta_2\} \\ \mathcal{S}_4 = \{\delta_1, \delta_2\} \end{array}$$

But only $\mathcal{S}_3$ proper ("right"), so extension is

$$\begin{array}{lcl} \mathcal{E}_3 & = & Th(\mathcal{W} \cup Conc(\mathcal{S}_3)) \\ & = & Th(\{P, P \supset B\} \cup \{\neg F\}) \\ & = & Th(\{P, P \supset B, \neg F\}), \end{array}$$

15. Immediate goal: specify *proper scenarios*

Binding defaults

1. If defaults provide reasons, binding defaults provide *good reasons*—forceful, or persuasive, in a context of a scenario

   Defined through preliminary concepts:

   Triggering

   Conflict

   Defeat

2. Triggered defaults:

$$\mathit{Triggered}_{\mathcal{W},\mathcal{D},<}(\mathcal{S}) = \{\delta \in \mathcal{D} : \mathcal{W} \cup \mathit{Conc}(\mathcal{S}) \vdash \mathit{Prem}(\delta)\}$$

3. Example: $\langle \mathcal{W}, \mathcal{D}, < \rangle$ with

$$\begin{array}{lcl} \mathcal{W} & = & \{B\} \\ \mathcal{D} & = & \{\delta_1, \delta_2\} \\ \delta_1 & = & B \rightarrow F \\ \delta_2 & = & P \rightarrow \neg F \\ \delta_1 < \delta_2 & & \end{array}$$

   Then

$$\mathit{Triggered}_{\mathcal{W},\mathcal{D},<}(\emptyset) \quad = \quad \{\delta_1\}$$

4. Terminology question: What are reasons?

   Answer: Reasons are antecedents of triggered defaults

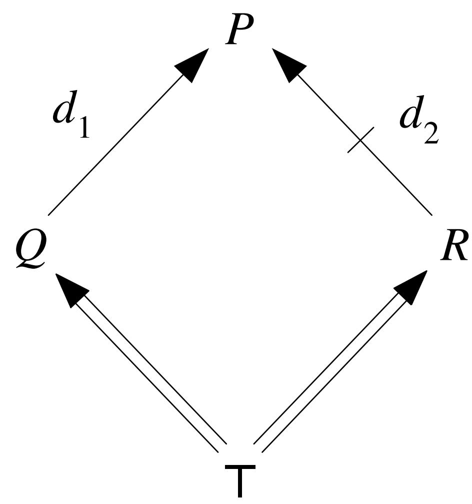

5. Conflicted defaults:

$$Conflicted_{\mathcal{W},\mathcal{D},<}(\mathcal{S}) = \{\delta \in \mathcal{D} : \mathcal{W} \cup Conc(\mathcal{S}) \vdash \neg Conc(\delta)\}$$

6. Example (Nixon Diamond):

   Take $\langle \mathcal{W}, \mathcal{D}, < \rangle$ with

$$\begin{array}{lcl} \mathcal{W} & = & \{Q, R\} \\ \mathcal{D} & = & \{\delta_1, \delta_2\} \\ \delta_1 & = & Q \to P \\ \delta_2 & = & R \to \neg P \\ < & = & \emptyset. \end{array}$$

   ($Q =$ Quaker, $R =$ Republican, $P =$ Pacifist)

   Then

$$\begin{array}{lcl} Triggered_{\mathcal{W},\mathcal{D},<}(\emptyset) & = & \{\delta_1, \delta_2\} \\ Conflicted_{\mathcal{W},\mathcal{D},<}(\emptyset) & = & \emptyset \end{array}$$

   But

$$\begin{array}{lcl} Conflicted_{\mathcal{W},\mathcal{D},<}(\{\delta_1\}) & = & \{\delta_2\} \\ Conflicted_{\mathcal{W},\mathcal{D},<}(\{\delta_2\}) & = & \{\delta_1\} \end{array}$$

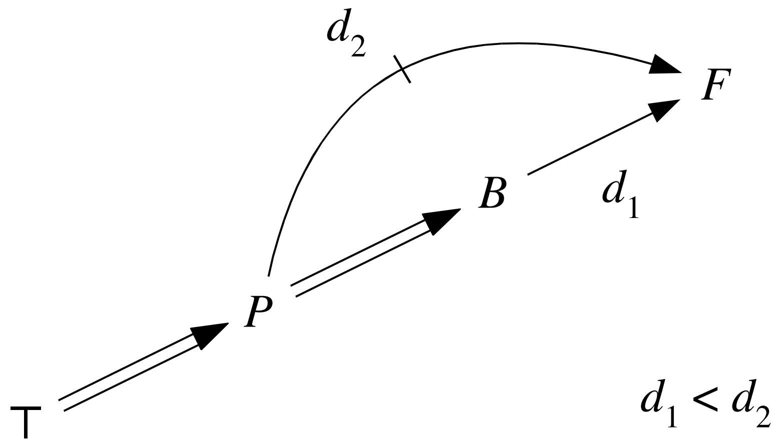

7. *Basic idea:* A default is defeated if there is a stronger reason supporting a contrary conclusion

$$Defeated_{\mathcal{W},\mathcal{D},<}(\mathcal{S}) = \{\delta \in \mathcal{D} : \exists\, \delta' \in Triggered_{\mathcal{W},\mathcal{D},<}(\mathcal{S}). \\ (1)\ \delta < \delta' \\ (2)\ Conc(\delta') \vdash \neg Conc(\delta)\}.$$

8. Example of defeat (Tweety, again):

   $\langle \mathcal{W}, \mathcal{D}, < \rangle$ where

$$\begin{array}{lcl} \mathcal{W} & = & \{P, P \Rightarrow B\} \\ \mathcal{D} & = & \{\delta_1, \delta_2\} \\ \delta_1 & = & B \rightarrow F \\ \delta_2 & = & P \rightarrow \neg F \\ \delta_1 < \delta_2 & & \end{array}$$

   Here, δ_1 is *defeated*:

$$Defeated_{\mathcal{W},\mathcal{D},<}(\emptyset) \quad = \quad \{\delta_1\}$$

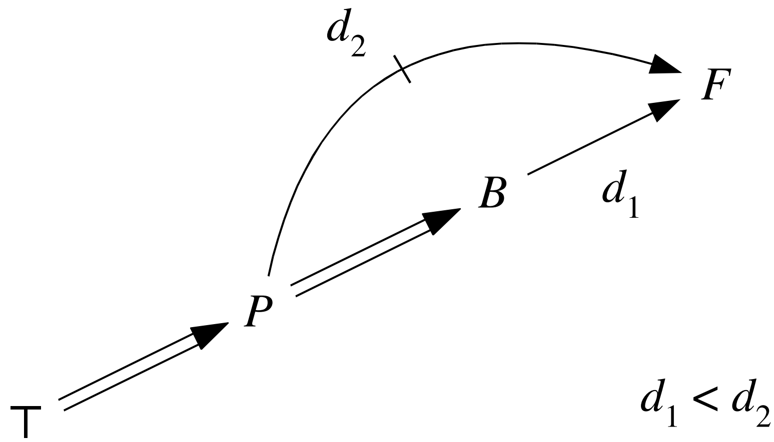

9. Finally, binding defaults:

$$\begin{aligned} \mathit{Binding}_{\mathcal{W},\mathcal{D},<}(\mathcal{S}) \quad = \quad \{\delta \in \mathcal{D} : \quad & \delta \in \mathit{Triggered}_{\mathcal{W},\mathcal{D},<}(\mathcal{S}) \\ & \delta \notin \mathit{Conflicted}_{\mathcal{W},\mathcal{D},<}(\mathcal{S}) \\ & \delta \notin \mathit{Defeated}_{\mathcal{W},\mathcal{D},<}(\mathcal{S})\} \end{aligned}$$

10. *Stable* scenarios: $\mathcal{S}$ is stable just in case

$$\mathcal{S} = \mathit{Binding}_{\mathcal{W},\mathcal{D},<}(\mathcal{S})$$

11. Example (Tweety, yet again): four scenarios

$$\begin{aligned} &\mathcal{S}_1 = \emptyset \\ &\mathcal{S}_2 = \{\delta_1\} \\ &\mathcal{S}_3 = \{\delta_2\} \\ &\mathcal{S}_4 = \{\delta_1, \delta_2\} \end{aligned}$$

Only $\mathcal{S}_3 = \{\delta_2\}$ is stable, because

$$\mathcal{S}_3 = \mathit{Binding}_{\mathcal{W},\mathcal{D},<}(\mathcal{S}_3)$$

Three complications

1. Complication #1: Can we just identify the proper scenarios with the stable scenarios?

   Almost . . . but not quite

2. Problem is "groundedness"

   Take $\langle \mathcal{W}, \mathcal{D}, < \rangle$ with

$$\begin{array}{lcl} \mathcal{W} & = & \emptyset \\ \mathcal{D} & = & \{\delta_1\} \\ \delta_1 & = & A \rightarrow A \\ < & = & \emptyset. \end{array}$$

   Then $\mathcal{S}_1 = \{\delta_1\}$ is a stable scenario, but shouldn't be proper

   The belief set generated by $\mathcal{S}_1$ is

$$Th(\mathcal{W} \cup Conc(\mathcal{S})) = Th(\{A\})$$

   but that's not right!

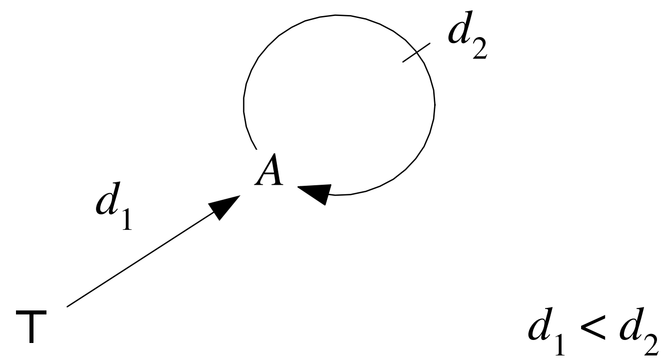

3. Complication #2: Some theories have *no* proper scenarios, and so no extensions

   Example: $\langle \mathcal{W}, \mathcal{D}, < \rangle$ with

$$
\begin{array}{lcl}
\mathcal{W} & = & \emptyset \\
\mathcal{D} & = & \{\delta_1, \delta_2\} \\
\delta_1 & = & \top \rightarrow A \\
\delta_2 & = & A \rightarrow \neg A \\
\delta_1 < \delta_2 & &
\end{array}
$$

4. Options:

   Syntactic restrictions to rule out "vicious cycles"

   Generalize definition of proper scenario, using tools from truth theory

   Live with it (benign choice if we like "skeptical" theory)

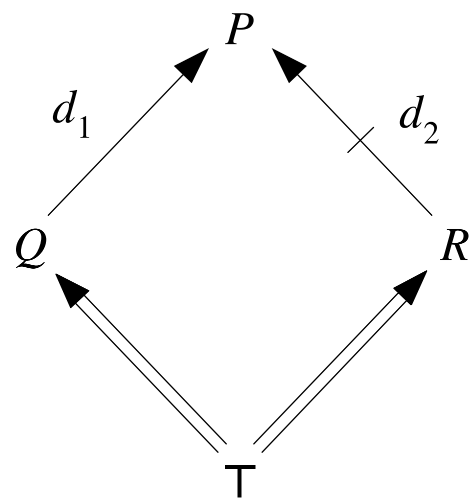

5. Complication #3: Some theories have *multiple* proper scenarios, and so multiple extensions

   Example: Nixon Diamond, again

   Take $\langle \mathcal{W}, \mathcal{D}, < \rangle$ with

$$\begin{array}{lcl} \mathcal{W} & = & \{Q, R\} \\ \mathcal{D} & = & \{\delta_1, \delta_2\} \\ \delta_1 & = & Q \rightarrow P \\ \delta_2 & = & R \rightarrow \neg P \\ < & = & \emptyset. \end{array}$$

   Then *two* proper scenarios

$$\mathcal{S}_1 = \{\delta_1\}$$
$$\mathcal{S}_2 = \{\delta_2\}$$

   and so two extensions:

$$\mathcal{E}_1 = Th(\{Q, R, P\})$$
$$\mathcal{E}_2 = Th(\{Q, R, \neg P\})$$

   So . . . what should we conclude?

6. Consider three options:

#1. Choice: pick an arbitrary proper scenario

Sensible, actually

But hard to codify as a consequence relation

#2. Brave/credulous: give some weight to any conclusion A contained in *some* extension

- Epistemic version (crazy): Endorse A whenever A is contained in some extension

  Example: P and $\neg P$ in Nixon case

- Epistemic version (not crazy): Endorse $\mathcal{B}(A)$—A is "believable"—whenever A is contained in some extension

  Example: $\mathcal{B}(P)$ and $\mathcal{B}(\neg P)$ in Nixon case

- Practical version: Endorse $\bigcirc(A)$—A is an "ought"—whenever A is contained in some extension

  Example: $\bigcirc(P)$ and $\bigcirc(\neg P)$ in Nixon case

#3. Cautious/"Skeptical": endorse A as conclusion whenever A contained in *every* extension

Defines consequence relation, and not weird: supports neither P nor $\neg P$ in Nixon case

Note: most popular option, but some problems . . .

Elaborating default logic

1. Discuss here only two things:

   Ability to reason about priorities

   Treatment of "undercutting" or "exclusionary" defeat

2. Begin with first problem

   So far, fixed priorities on default rules

   But we can reason about default priorities . . . and then use the priorities we arrive at to control our reasoning

3. Five steps:

#1. Add priority statements ($\delta_7 < \delta_9$) to object language

#2. Introduce new *variable priority* default theories

$$\langle \mathcal{W}, \mathcal{D} \rangle$$

with priority statements now belonging to $\mathcal{W}$ and $\mathcal{D}$

#3. Add strict priority axioms to $\mathcal{W}$:

$$\delta < \delta' \Rightarrow \neg(\delta' < \delta)$$
$$(\delta < \delta' \wedge \delta' < \delta'') \Rightarrow \delta < \delta''$$

#4. Lift priorities from object to meta language

$$\delta <_{\mathcal{S}} \delta' \quad \text{iff} \quad \mathcal{W} \cup Conc(\mathcal{S}) \vdash \delta < \delta'.$$

#5. Proper scenarios for new default theories:

$\mathcal{S}$ is a *proper scenario* based on $\langle \mathcal{W}, \mathcal{D} \rangle$

iff

$\mathcal{S}$ is a proper scenario based on $\langle \mathcal{W}, \mathcal{D}, <_{\mathcal{S}} \rangle$

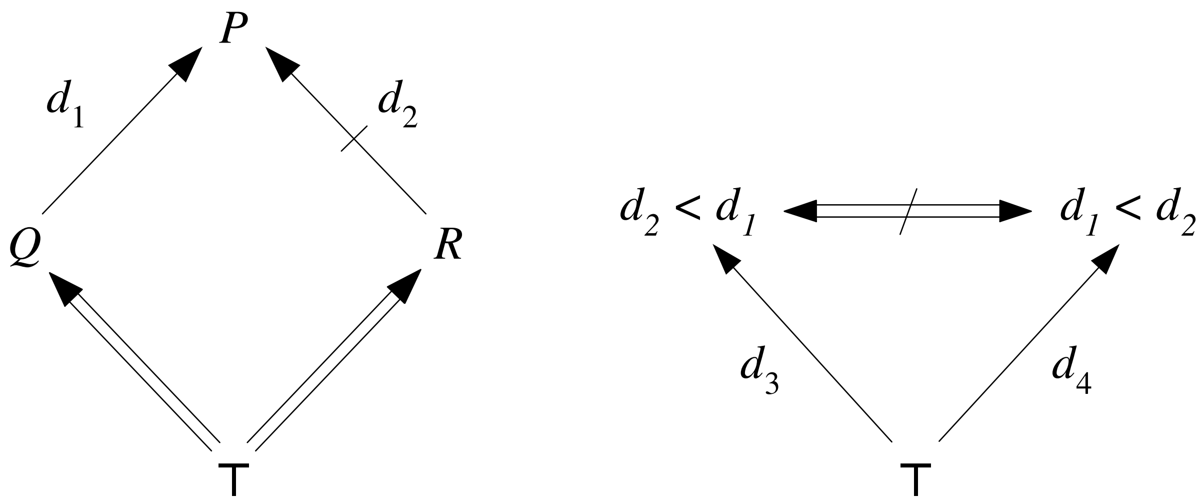

4. Example (Extended Nixon Diamond):

   Consider $\langle \mathcal{W}, \mathcal{D} \rangle$ where

   $\mathcal{W}$ contains Q, P

   $\mathcal{D}$ contains

$$\begin{aligned}
\delta_1 &= Q \rightarrow P \\
\delta_2 &= R \rightarrow \neg P \\
\delta_3 &= \top \rightarrow \ \delta_2 < \delta_1 \\
\delta_4 &= \top \rightarrow \ \delta_1 < \delta_2 \\
\delta_5 &= \top \rightarrow \ \delta_4 < \delta_3
\end{aligned}$$

   Then unique proper scenario is

$$\mathcal{S} = \{\delta_1, \delta_3, \delta_5\}$$

5. *Undercutting* defeat (epistemology), compared to rebutting defeat

   Example:

   > The object looks red
   >
   > My reliable friend says it is not red
   >
   > Drug 1 makes everything look red

6. *Exclusionary* reasons (practical reasoning)

   Example (Colin's dilemma, from Raz):

   Should son go to private school??

   > The school provides good education
   >
   > He'll meet fancy friends
   >
   > The school is expensive
   >
   > Decision would undermine public education
   >
   > Promise: only consider son's interests . . .

7. How can this be represented?

   One view (Pollock): undercutting a separate form of defeat

   My suggestion:

   > Only ordinary (rebutting) defeat
   >
   > Enhance the language slightly
   >
   > Tweak the notion of triggering

8. Four steps:

#1. New predicate Out, so that $Out(\delta)$ means that δ is undercut, or excluded

#2. Introduce new *exclusionary* default theories as theories in a language containing Out.

#3. Lift notion of exclusion from object to meta language: where $\mathcal{S}$ is scenario based on theory with $\mathcal{W}$ as hard information

$$\delta \in Excluded_{\mathcal{S}} \quad \text{iff} \quad \mathcal{W} \cup Conc(\mathcal{S}) \vdash Out(\delta).$$

#4. Only defaults that are not excluded can be triggered:

$$Triggered_{\mathcal{W},\mathcal{D},<}(\mathcal{S}) = \{\delta \notin Excluded_{\mathcal{S}} \text{ and } \mathcal{W} \cup Conc(\mathcal{S}) \vdash Prem(\delta)\}$$

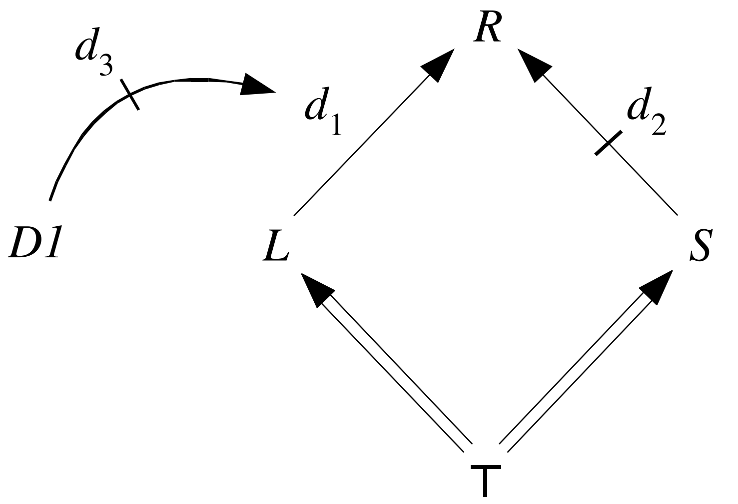

$d_1 < d_2$
$d_1 < d_3$

9. Example: For ordinary rebutting defeat, take $\langle \mathcal{W}, \mathcal{D} \rangle$ where

   $\mathcal{W}$ contains L, S, and $\delta_1 < \delta_2$, $\delta_1 < \delta_3$

   $\mathcal{D}$ contains

$$\begin{aligned} \delta_1 &= L \rightarrow R \\ \delta_2 &= S \rightarrow \neg R \\ \delta_3 &= D1 \rightarrow Out(\delta_1) \end{aligned}$$

   (L = Looks red, R = Red, S = Statement by friend, $D1$ = Drug 1)

   So proper scenario is

$$\mathcal{S} = \{\delta_2\}$$

   generating the extension

$$\mathcal{E} = Th(\mathcal{W} \cup \{\neg R\})$$

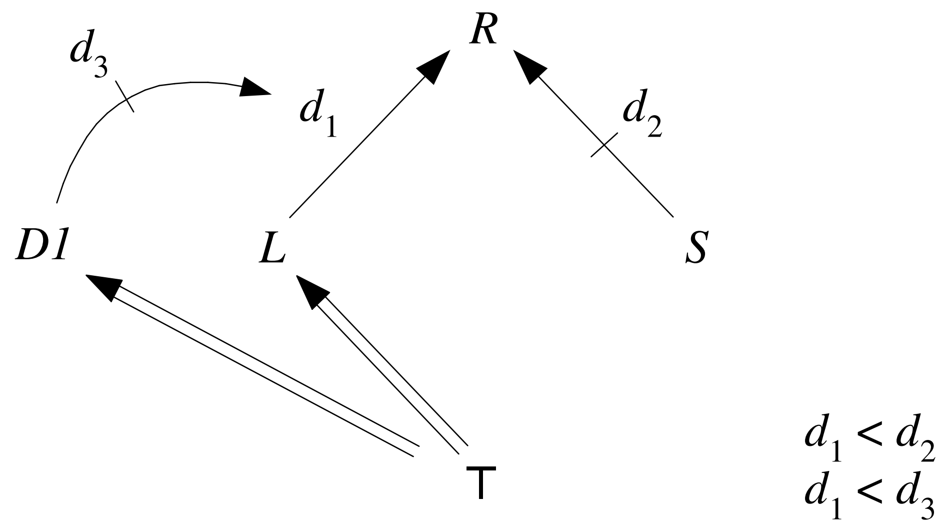

10. Example: For undercutting, or exclusionary, defeat, take $\langle \mathcal{W}, \mathcal{D} \rangle$ where

$\mathcal{D}$ contains

$$\begin{aligned} \delta_1 &= L \to R \\ \delta_2 &= S \to \neg R \\ \delta_3 &= D1 \to Out(\delta_1) \end{aligned}$$

$\mathcal{W}$ contains L, $D1$, and $\delta_1 < \delta_2$, $\delta_1 < \delta_3$

So proper scenario is

$$\mathcal{S} = \{\delta_3\}$$

generating the extension

$$\mathcal{E} = Th(\mathcal{W} \cup \{Out(\delta_1)\})$$

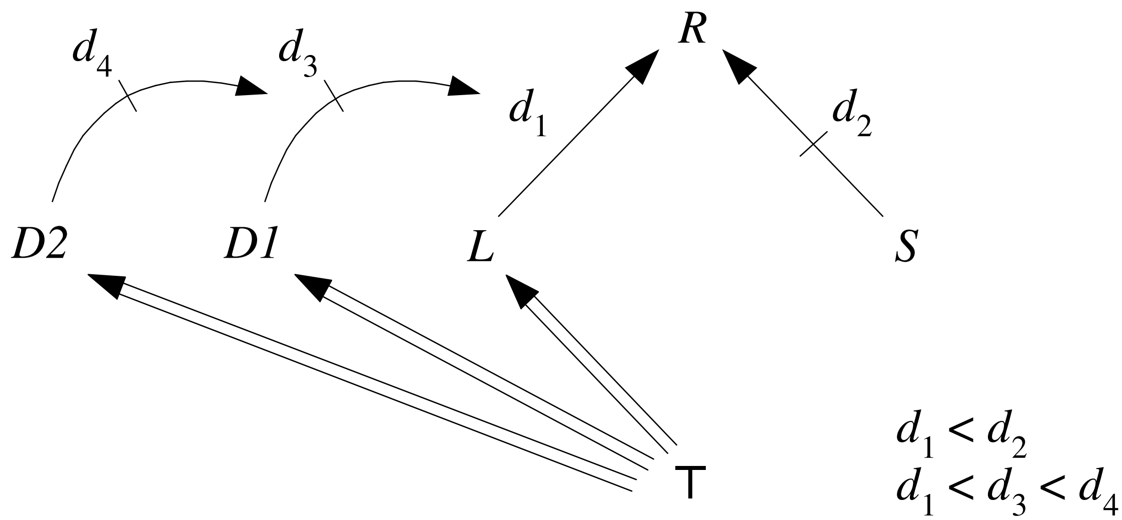

11. Example: Drug 2 is an antidote to Drug 1, so for an excluder excluder, take $\langle \mathcal{W}, \mathcal{D} \rangle$ where

    $\mathcal{W}$ contains L, $D1$, $D2$, $\delta_1 < \delta_2$, and $\delta_1 < \delta_3 < \delta_4$

    $\mathcal{D}$ contains

$$\begin{aligned} \delta_1 &= L \rightarrow R \\ \delta_2 &= S \rightarrow \neg R \\ \delta_3 &= D1 \rightarrow Out(\delta_1) \\ \delta_4 &= D2 \rightarrow Out(\delta_3) \end{aligned}$$

    So proper scenario is

$$\mathcal{S} = \{\delta_1, \delta_4\}$$

    generating the extension

$$\mathcal{E} = Th(\mathcal{W} \cup \{R, Out(\delta_3)\})$$

12. Example (Colin's dilemma, simplified):

Let $\mathcal{D}$ contain

$$\begin{aligned}\delta_1 &= E \rightarrow S \\ \delta_2 &= U \rightarrow \neg S \\ \delta_3 &= \neg Welfare(\delta_2) \rightarrow Out(\delta_2)\end{aligned}$$

(E = Provides good education, S = Send son to private school, U = Undermine support for public education)

The default δ_3 is itself an instance of

$$\neg Welfare(\delta) \rightarrow Out(\delta),$$

Let $\mathcal{W}$ contain E, U, and $\neg Welfare(\delta_2)$

Then proper scenario is

$$\mathcal{S} = \{\delta_1, \delta_3\}$$

generating the extension

$$\mathcal{E} = Th(\mathcal{W} \cup \{S, Out(\delta_2)\})$$

Exclusion and priorities

1. Can exclusion be defined in terms of priority adjustment?

   Many people think so...

   > Perry: "an exclusionary reason is simply the special case where one or more first-order reasons are treated as having zero weight"
   >
   > Dancy: "If we are happy with the idea that a reason can be attenuated . . . , why should we fight shy of supposing that it can be reduced to nothing"
   >
   > Schroeder: "undercutting" is best analyzed as an extreme case of attenuation in the strength of reasons; he refers to this thesis as the "undercutting hypothesis"
   >
   > Horty: developed a formal theory of exclusion as the assignment to a default of a priority that falls below some particular threshold

2. But this idea entials

   > *Downward closure for exclusion*: if δ is excluded and a $\delta' < \delta$, then δ' is excluded.

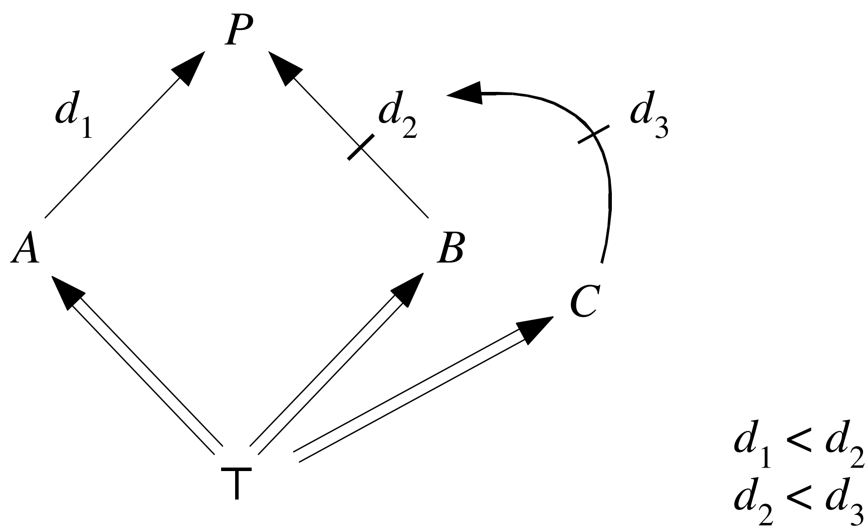

3. Example:

   Take $\langle \mathcal{W}, \mathcal{D} \rangle$ with

$$\begin{aligned} \mathcal{W} &= \{A, B, C, \delta_1 < \delta_2, \delta_2 < \delta_3\} \\ \mathcal{D} &= \{\delta_1, \delta_2, \delta_3\} \\ \delta_1 &= A \rightarrow P \\ \delta_2 &= B \rightarrow \neg P \\ \delta_3 &= C \rightarrow Out(\delta_2) \end{aligned}$$

   Then the proper scenario is

$$\mathcal{S} = \{\delta_1, \delta_3\}$$

   generating the extension

$$\mathcal{E} = Th(\mathcal{W} \cup \{P, Out(\delta_2)\})$$

   So downward closure fails, but is that right?

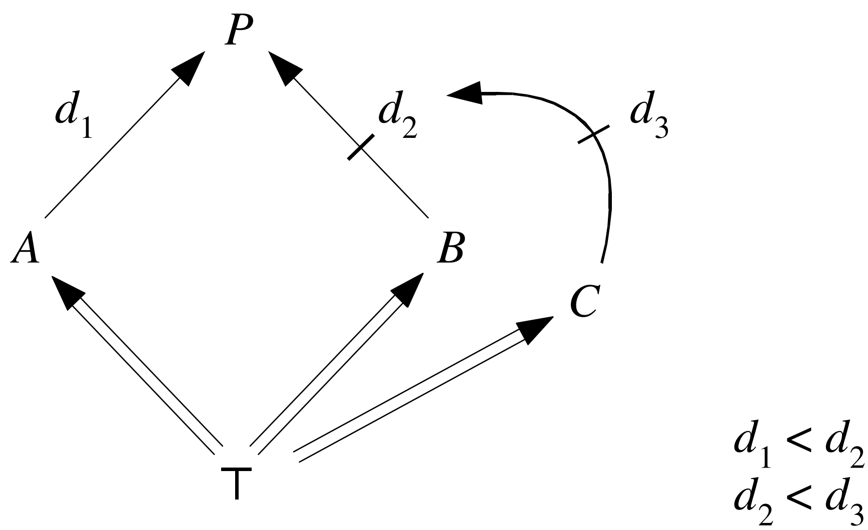

4. First interpretation (mathematicians):

   Priority ordering represents reliability of the mathematicians

   $P =$ Some conjecture

   $A =$ First mathematician's assertion that he has proved P

   $B =$ Second mathematician's assertion that she has proved $\neg P$

   $C =$ Third mathematician's assertion that second mathematician is too unreliable to be trusted

   Here downward closure seems to hold

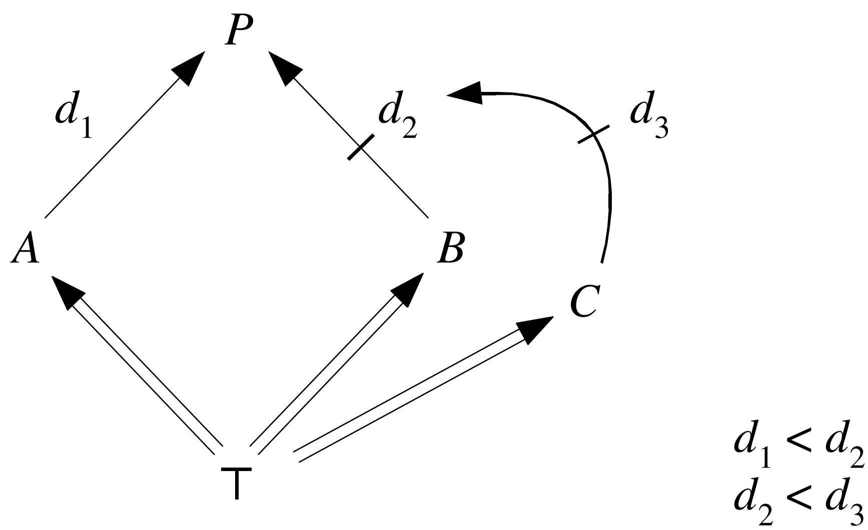

5. Second interpretation (officers):

   Captain < Major < Colonel

   $P =$ Some action to perform (or not)

   $A =$ Captain's command to perform P

   $B =$ Major's command not to perform P

   $C =$ Colonel's command to ignore Major's command

   Here downward closure seems to fail

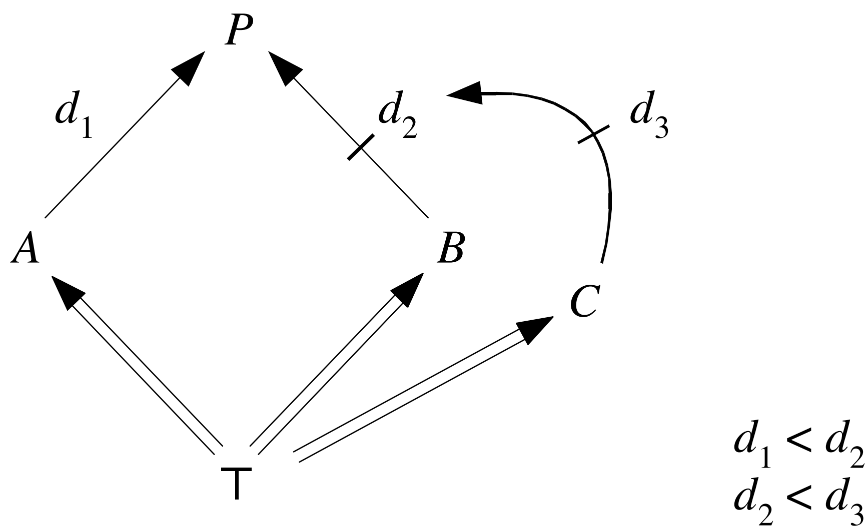

6. So if downward closure fails, what do we do when we *want* downward closure?

   Answer: Supplement hard information with

   $$(Out(\delta) \wedge \delta' < \delta) \supset Out(\delta'),$$

   This give us the proper scenario

   $$\mathcal{S} = \{\delta_3\}$$

   generating the extension

   $$\mathcal{E} = Th(\mathcal{W} \cup \{Out(\delta_2)\})$$

Exclusion by weaker defaults

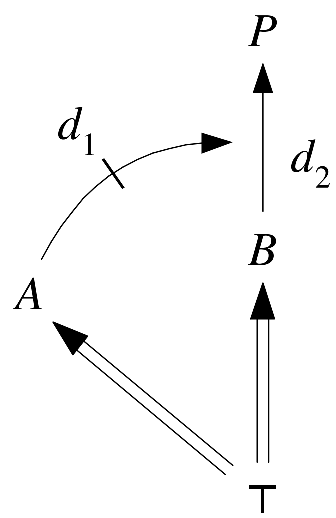

$d_1 < d_2$

1. Next question: a default cannot be *defeated* by a weaker default, but can it be *excluded* by a weaker default?

   Yes, on current account.

   Take $\langle \mathcal{W}, \mathcal{D} \rangle$ with

$$\begin{array}{lcl} \mathcal{W} & = & \{A, B, \delta_1 < \delta_2\} \\ \mathcal{D} & = & \{\delta_1, \delta_2\} \\ \delta_1 & = & A \rightarrow Out(\delta_2) \\ \delta_2 & = & B \rightarrow P \end{array}$$

   Then the proper scenario is

$$\mathcal{S} = \{\delta_1\}$$

   generating the extension

$$\mathcal{E} = Th(\mathcal{W} \cup \{Out(\delta_2)\})$$

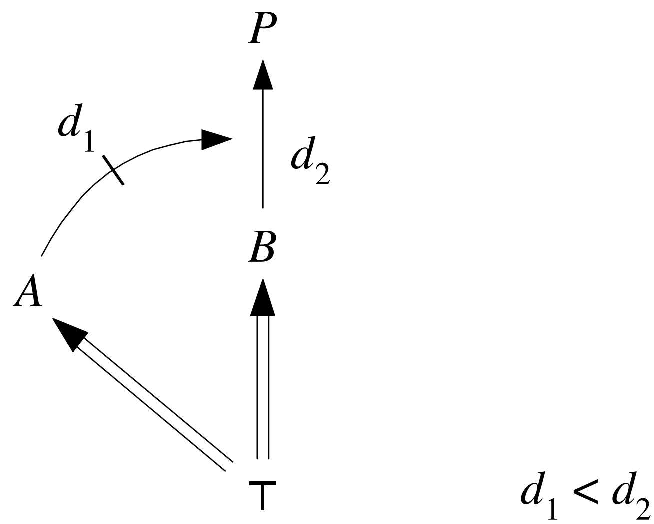

2. Pollock's answer: No

> It seems apparent that any adequate account of justification must have the consequence that if a belief is unjustified relative to a particular degree of justification, then it is unjustified relative to any highter degree of justification. (*Cognitive Carpentry*, p 104)

3. I disagree: different standards of legal evidence, jailhouse snitch

4. Question: How do we represent changing standards of evidence?

5. Another question: what do we do when we *want* to rule out exclusion by weaker defaults? (Eg, military officer interpretation)

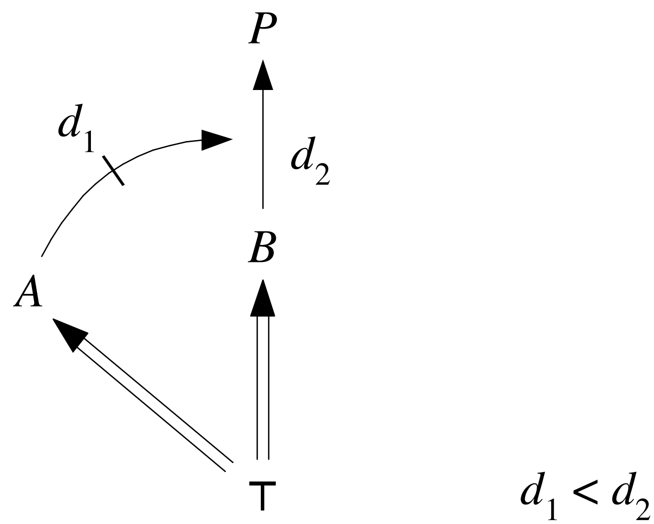

6. My (tentative) suggeston: suppose defaults are *protected from exclusion*

   Begin with $\langle \mathcal{W}, \mathcal{D} \rangle$, where

$$\begin{array}{lcl} \mathcal{W} & = & \{A, B, \delta_1 < \delta_2\} \\ \mathcal{D} & = & \{\delta_1, \delta_2\} \\ \delta_1 & = & A \rightarrow Out(\delta_2) \\ \delta_2 & = & B \rightarrow P \wedge \neg Out(\delta_2) \end{array}$$

   Then the proper scenario is

$$\mathcal{S} = \{\delta_2\}$$

   generating the extension

$$\mathcal{E} = Th(\mathcal{W} \cup \{P \wedge \neg Out(\delta_2)\})$$

Floating conclusions

1. Getting from scenario $\mathcal{S}$ to extension $\mathcal{E}$

   Direct route:

   $$\mathcal{E} = Th(\mathcal{W} \cup Conc(\mathcal{S}))$$

   Indirect route:

   Defaults are rules of inference
   Construct arguments to *support* conclusions
   Examples:

   $$\top \Rightarrow A \rightarrow B \Rightarrow \neg C$$

   $$\top \Rightarrow Q \rightarrow P$$

   $$\top \Rightarrow R \rightarrow \neg P$$

   support conclusion $\neg C$, P, $\neg P$.

   So, first form *argument extension* Φ
   Then take conclusions supported by Φ

2. Function * maps arguments into conclusions:

   $^*\alpha =$ conclusion supported by α

   $^*\Phi = \{^*\alpha : \alpha \in \Phi\}$

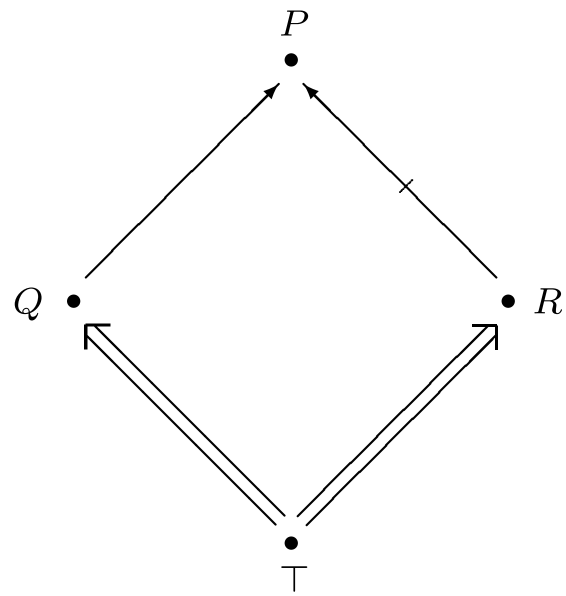

3. Consider multiple argument extensions

$$\Phi_1 = \{\top \Rightarrow Q, \top \Rightarrow R, \top \Rightarrow Q \rightarrow P\}$$

$$\Phi_2 = \{\top \Rightarrow Q, \top \Rightarrow R, \top \Rightarrow R \rightarrow \neg P\}$$

4. Skeptical—or "intersect extensions"—option now bifurcates

   Alternative #1:

   $$^*(\bigcap\{\Phi : \Phi \text{ is an argument extension of } \Gamma\})$$

   Alternative #2:

   $$\bigcap\{^*\Phi : \Phi \text{ is an argument extension of } \Gamma\}$$

5. In this case, same result:

   $$\{Q, R\}$$

   But not always, due to the phenomena of *floating conclusions*

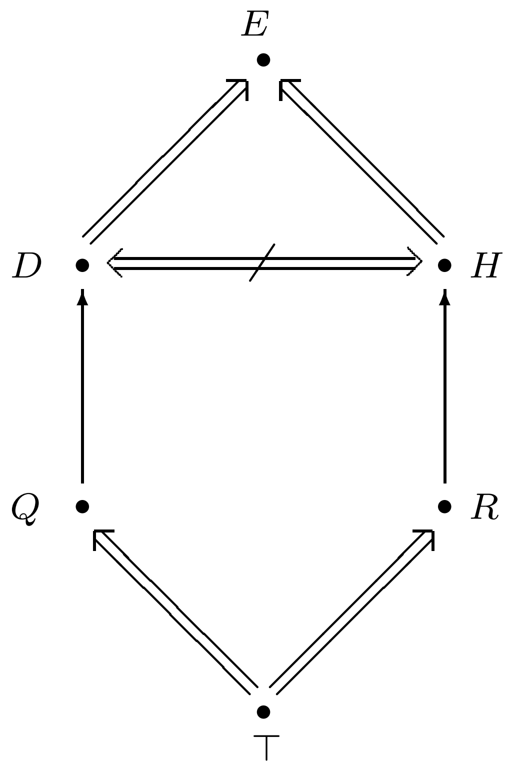

Argument Extensions:

$$\begin{aligned}\Phi_1 &= \{\ \top \Rightarrow Q,\ \ \top \Rightarrow R,\\ &\quad\ \top \Rightarrow Q \rightarrow D,\\ &\quad\ \top \Rightarrow Q \rightarrow D \not\Rightarrow H,\\ &\quad\ \top \Rightarrow Q \rightarrow D \Rightarrow E\ \}\end{aligned}$$

$$\begin{aligned}\Phi_2 &= \{\ \top \Rightarrow Q,\ \ \top \Rightarrow R,\\ &\quad\ \top \Rightarrow R \rightarrow H,\\ &\quad\ \top \Rightarrow R \rightarrow H \not\Rightarrow D,\\ &\quad\ \top \Rightarrow R \rightarrow H \Rightarrow E\ \}\end{aligned}$$

Alternative #1 yields:

$$\{Q, R\}$$

Alternative #2 yields:

$$\{Q, R, E\}$$

6. Conventional view is that floating conclusions *should* be accepted (so Alternative #2 is correct).

   Ginsberg:

   > Given that both hawks and doves are politically [extreme], Nixon certainly should be as well." (*Essentials of Artificial Intelligence, 1993*)

   Makinson and Schlechta:

   > It is an oversimplification to take a proposition A as acceptable . . . iff it is supported by some [argument] path α in the intersection of all extensions. Instead A must be taken as acceptable iff it is in the intersection of all *outputs* of extensions, where the output of an extension is the set of all propositions supported by some path within it. (*Artificial Intelligence, 1991*)

   Stein:

   > The difficulty lies in the fact that some conclusions may be true in every credulous extension, but supported by different [argument] paths in each. Any path-based theory must either accept one of these paths, and be unsound, or reject all such paths, and with them the ideally skeptical conclusion (*Resolving Ambiguity . . .*, 1991)

   Pollock:

   > (*Defeasible reasoning*, unpublished) makes it clear that desire for floating conclusions motivated 1995 semantics

7. Yacht example:

- Both (elderly) parents have $500K
- I want a yacht, requires large deposit, balance due later—otherwise, lose deposit
- Utilities determine conditional preferences:
  - If I will inherit at least half a million dollars, I should place a deposit on the yacht
  - Otherwise, I should not place a deposit

  So decision hinges on truth of

$$F \vee M$$

- Brother says: "Father will leave his money to me, but Mother is leaving her money to you"

$$BA(\neg F \wedge M)$$

- Sister says: "Mother will leave her money to me, but Father is leaving his money to you"

$$SA(F \wedge \neg M)$$

- Both brother and sister reliable, so have defaults:

$$BA(\neg F \wedge M) \rightarrow (\neg F \wedge M)$$

$$SA(F \wedge \neg M) \rightarrow (F \wedge \neg M)$$

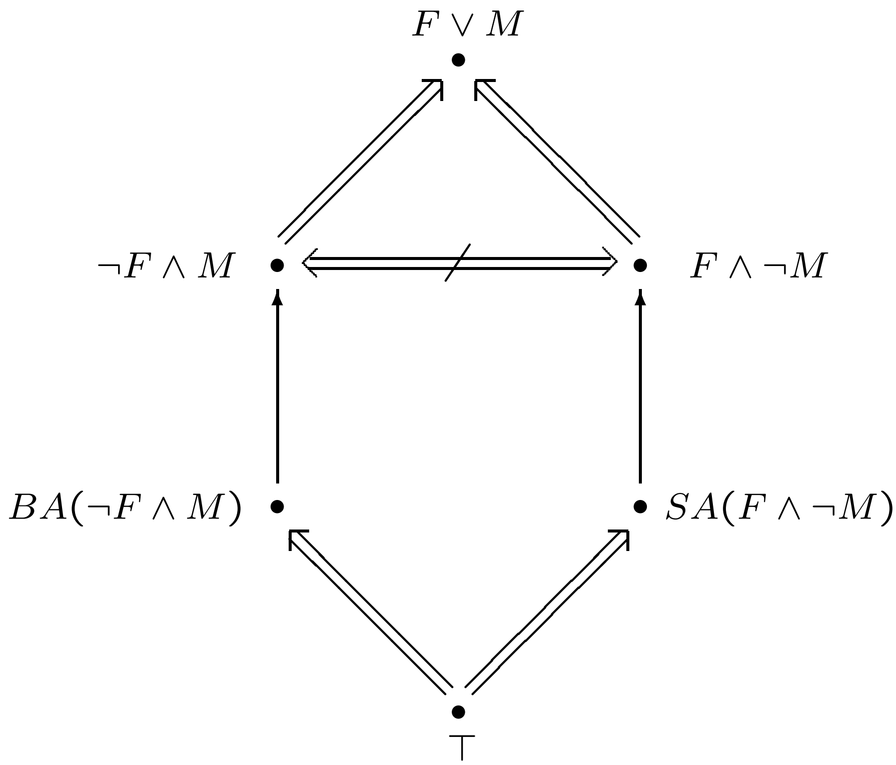

Argument Extensions:

$$\Phi_1 = \{ \top \Rightarrow BA(\neg F \wedge M), \; \top \Rightarrow SA(F \wedge \neg M), \; \top \Rightarrow BA(\neg F \wedge M) \rightarrow \neg F \wedge M, \; \top \Rightarrow BA(\neg F \wedge M) \rightarrow \neg F \wedge M \not\Rightarrow F \wedge \neg M, \; \top \Rightarrow BA(\neg F \wedge M) \rightarrow \neg F \wedge M \Rightarrow F \vee M \}$$

$$\Phi_2 = \{ \top \Rightarrow BA(\neg F \wedge M), \; \top \Rightarrow SA(F \wedge \neg M), \; \top \Rightarrow SA(F \wedge \neg M) \Rightarrow F \wedge \neg M, \; \top \Rightarrow SA(F \wedge \neg M) \Rightarrow F \wedge \neg M \not\Rightarrow \neg F \wedge M, \; \top \Rightarrow SA(F \wedge \neg M) \Rightarrow F \wedge \neg M \Rightarrow F \vee M \}$$

Alternative #1 yields:

$$\{BA(\neg F \wedge M), SA(F \wedge \neg M)\}$$

Alternative #2 yields:

$$\{BA(\neg F \wedge M), SA(F \wedge \neg M), F \vee M\}$$

8. Other examples:

- Military example:
  - You want to press ahead if enemy has retreated from defensive position
  - Spy 1 says: enemy retreating over mountains, diversionary force feigns retreat along river
  - Spy 2 says: enemy retreating along river, diversionary force feigns retreat over mountains

- Economics example:
  - Economic health, low inflation, strong growth
  - Prediction 1: strong growth will trigger high inflation, leading to recession
  - Prediction 2: inflation will continue to decline, resulting in deflation and so recession

- Ginsberg's original example:
  - Why not suppose that the extreme tendencies serve to moderate each other?

9. Why accept floating conclusions?

Maybe an analogy between

$$\begin{array}{l} B \supset A \\ C \supset A \\ B \vee C \\ \hline A \end{array}$$

and (supposing two extensions, $\mathcal{E}_1$ and $\mathcal{E}_2$)

$$\begin{array}{ll} A \in \mathcal{E}_1 & (\text{so}: \mathcal{E}_1 \supset A) \\ A \in \mathcal{E}_2 & (\text{so}: \mathcal{E}_2 \supset A) \\ \mathcal{E}_1 \vee \mathcal{E}_2 & ?????? \\ \hline A & \end{array}$$

First argument relies on premise that $A \vee B$; must skeptical reasoner suppose "$\mathcal{E}_1 \vee \mathcal{E}_2$"?

Sometimes appropriate to think

- One or another extension must be (entirely) right—we just don't know which

But other times

- Real possibility that they might all be wrong

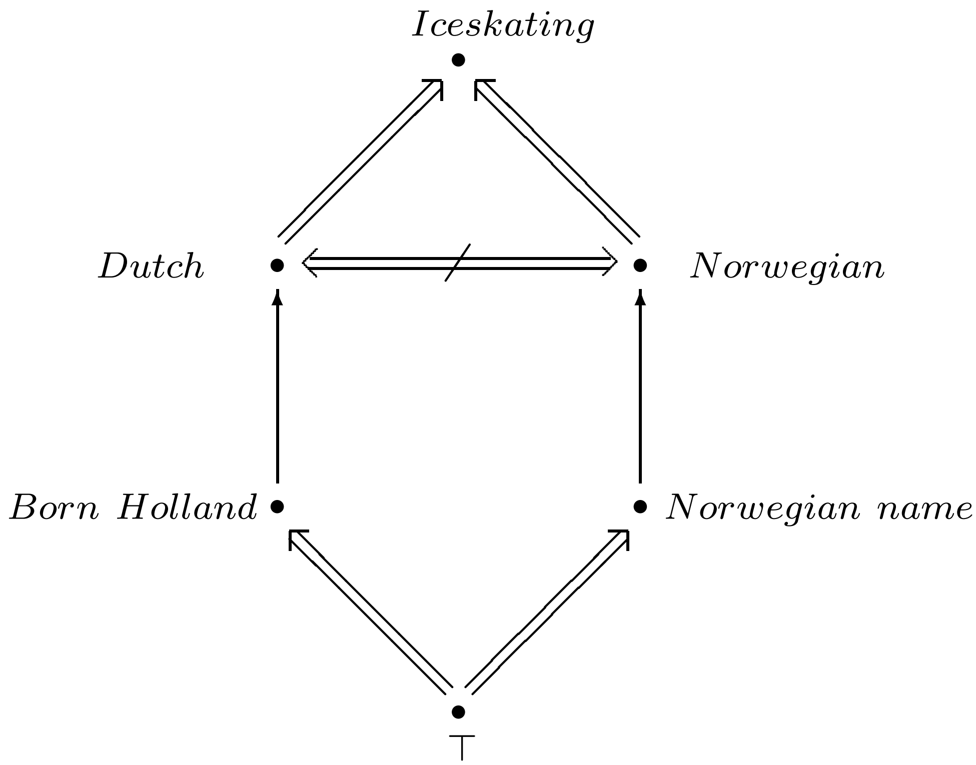

10. Prakken's example:

    Brygt Rykkje was born in Holland
    Brygt Rykkje has a Norwegian name
    Brygt Rykkje is Dutch
    Brygt Rykkje is Norwegian

    Here we *do* like the floating conclusion

11. Open question: how do we distinguish cases in which we do like floating conclusion from cases in which we don't?