Introduction to formal models of argumentation

Henry Prakken

Dundee (Scotland)

September 4th, 2014

Universiteit Utrecht

rijksuniversiteit groningen

What is argumentation?

- Giving reasons to support claims that are open to doubt
- Defending these claims against attack
- NB: Inference + dialogue

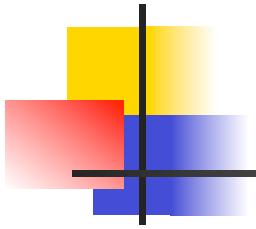

Why study argumentation?

- In linguistics:
  - Argumentation is a form of language use
- In Artificial Intelligence:
  - Our applications have humans in the loop
    - We want to model rational reasoning but with standards of rationality that are attainable by humans
    - Argumentation is natural for humans
  - Trade-off between rationality and naturalness
- In Multi-Agent Systems:
  - Argumentation is a form of communication

Today: formal models of argumentation

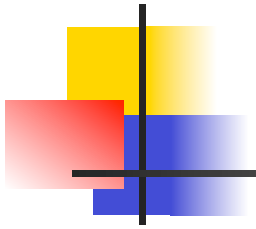

- Abstract argumentation
- Argumentation as inference
  - Frameworks for structured argumentation
    - Deductive vs. defeasible inferences
  - Argument schemes
- Argumentation as dialogue

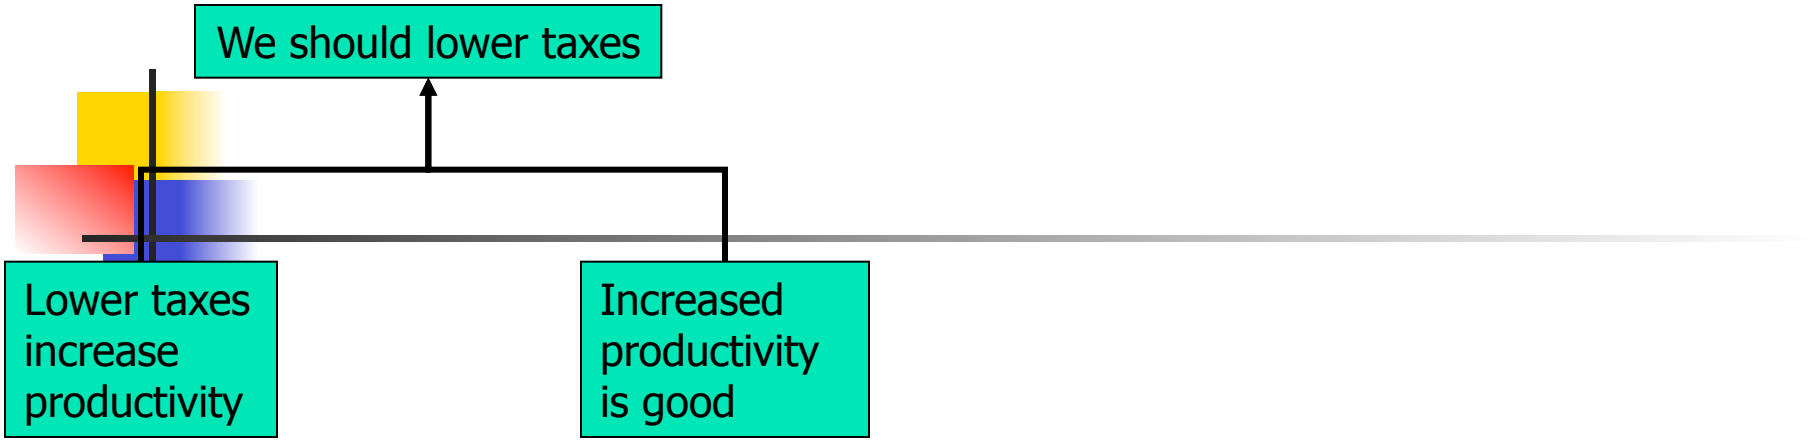
We should lower taxes
Lower taxes increase productivity
Increased productivity is good

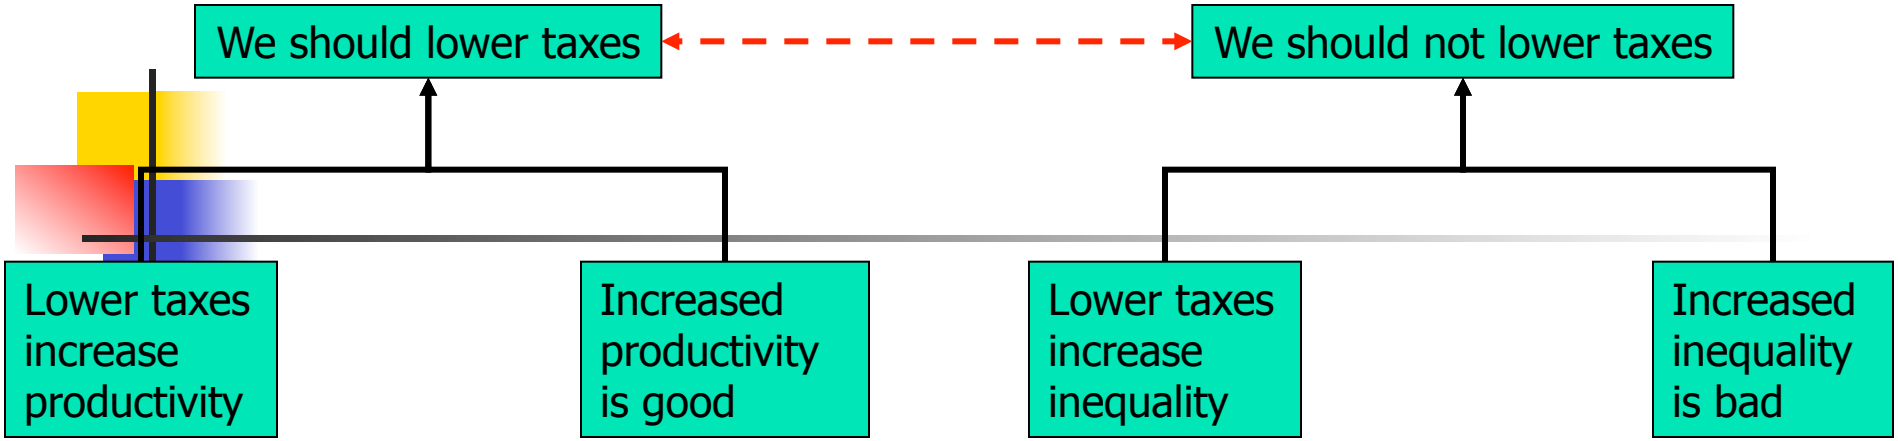

Attack on conclusion

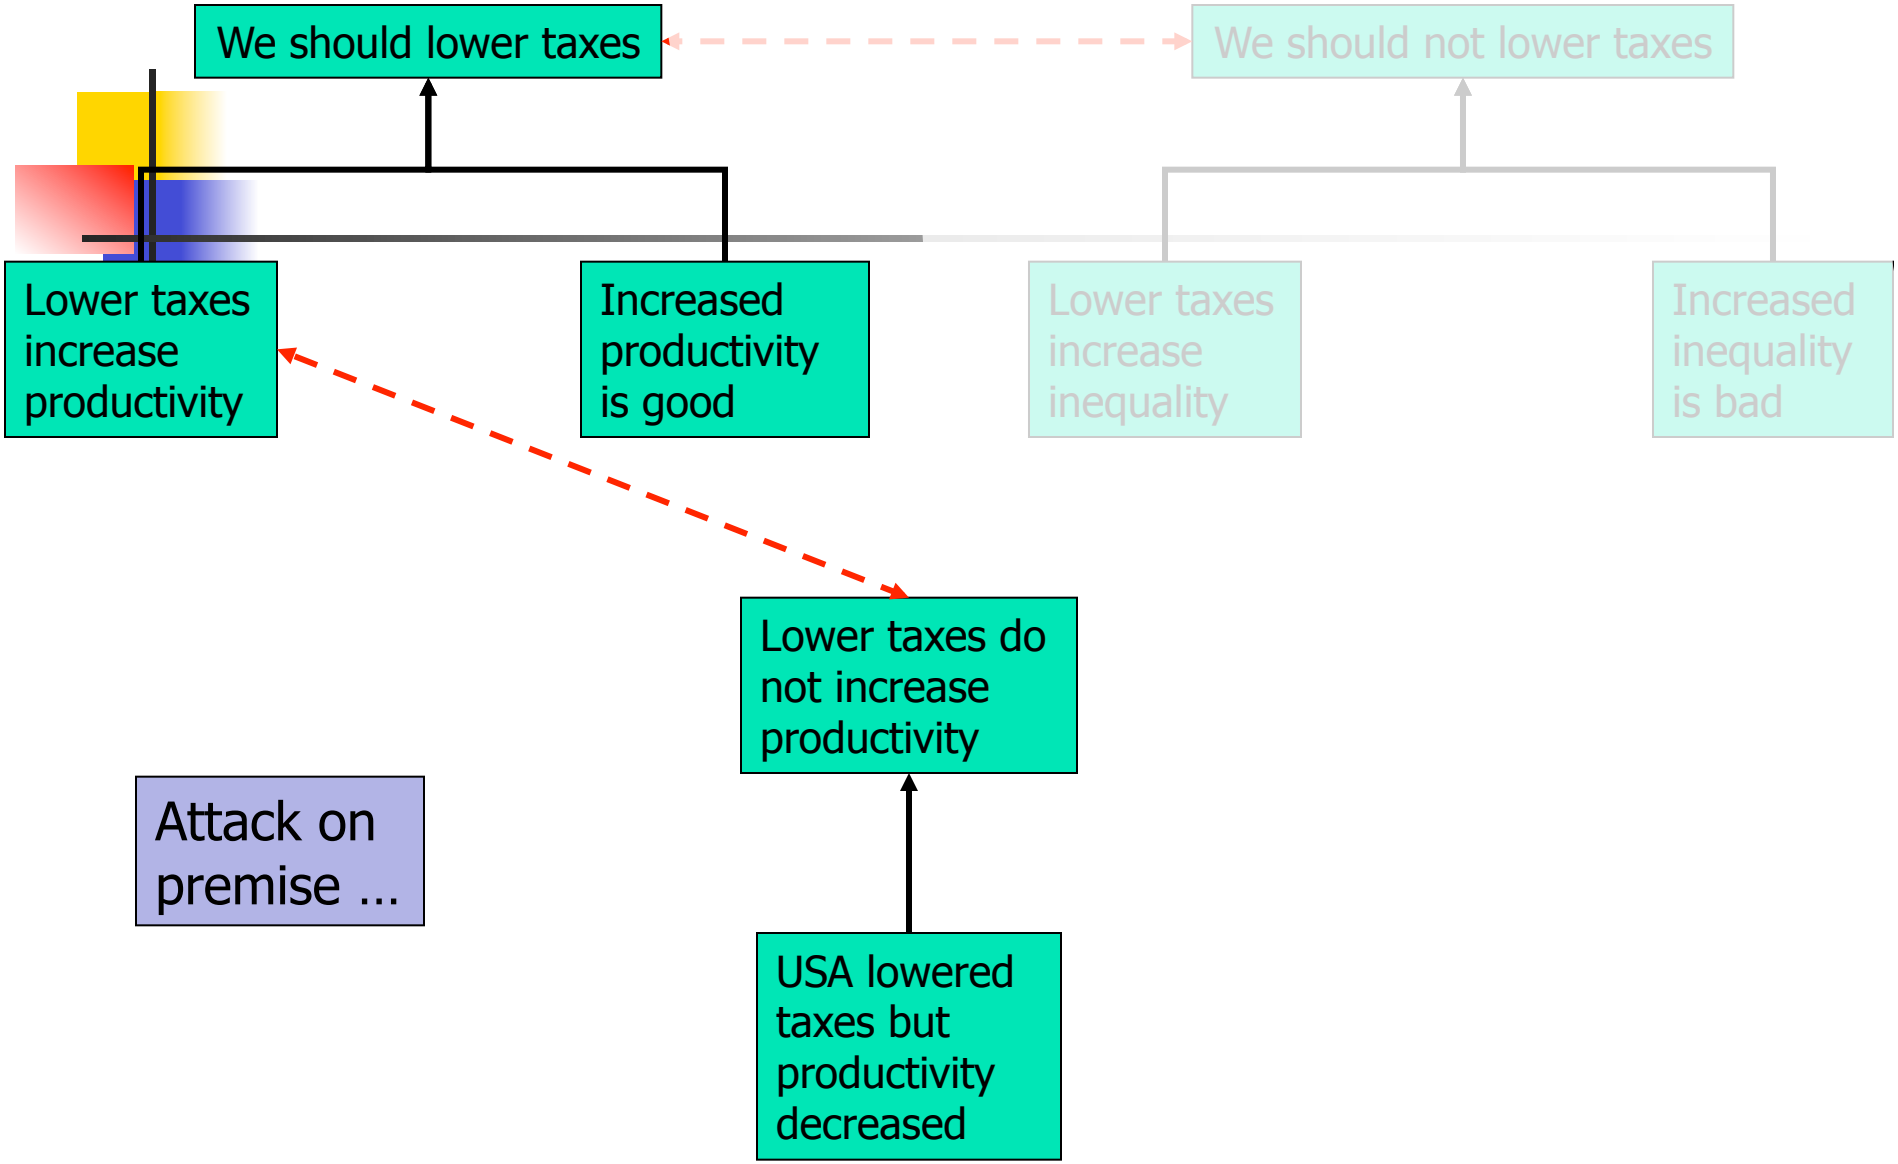
We should lower taxes
We should not lower taxes
Lower taxes increase productivity
Increased productivity is good
Lower taxes increase inequality
Increased inequality is bad
Lower taxes do not increase productivity
Attack on premise …
USA lowered taxes but productivity decreased

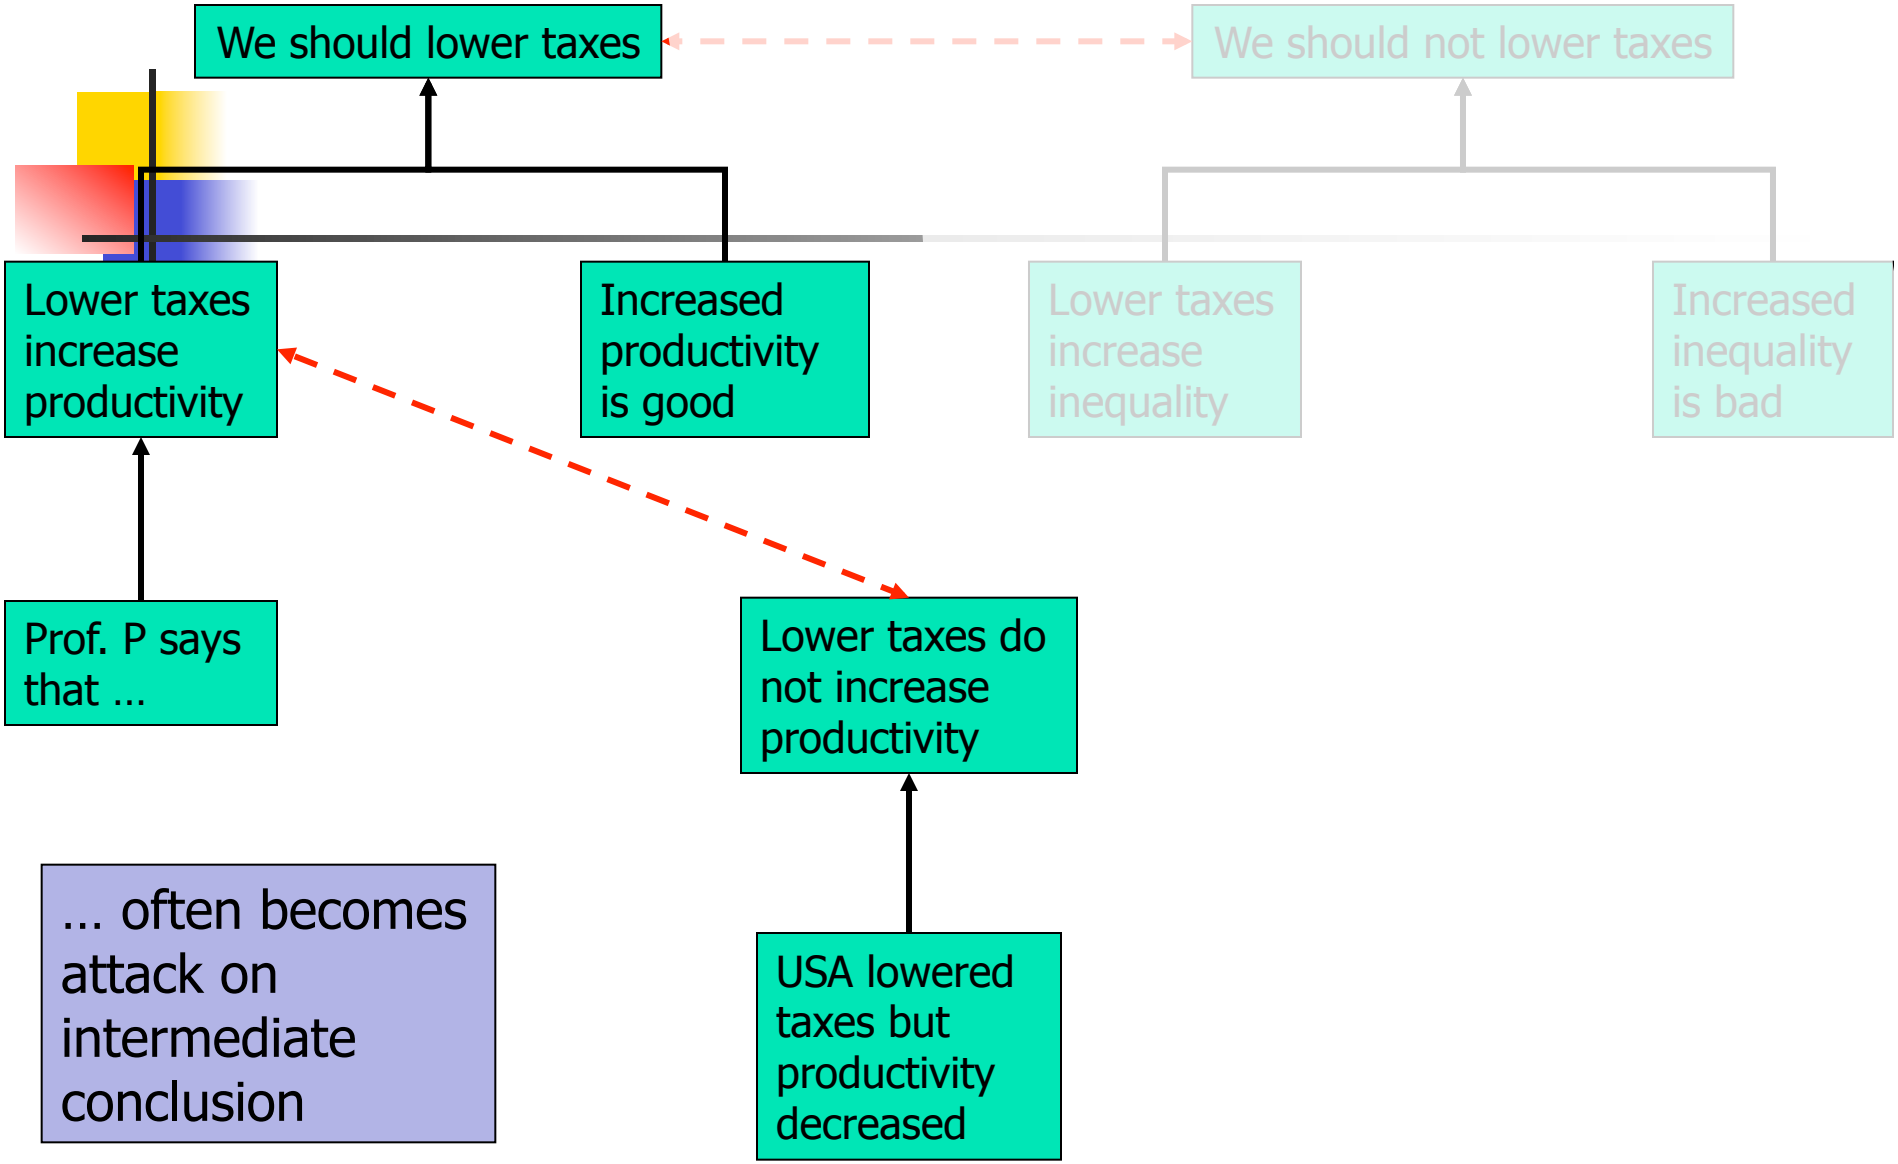
We should lower taxes
We should not lower taxes
Lower taxes increase productivity
Increased productivity is good
Lower taxes increase inequality
Increased inequality is bad
Prof. P says that …
Lower taxes do not increase productivity
… often becomes attack on intermediate conclusion
USA lowered taxes but productivity decreased

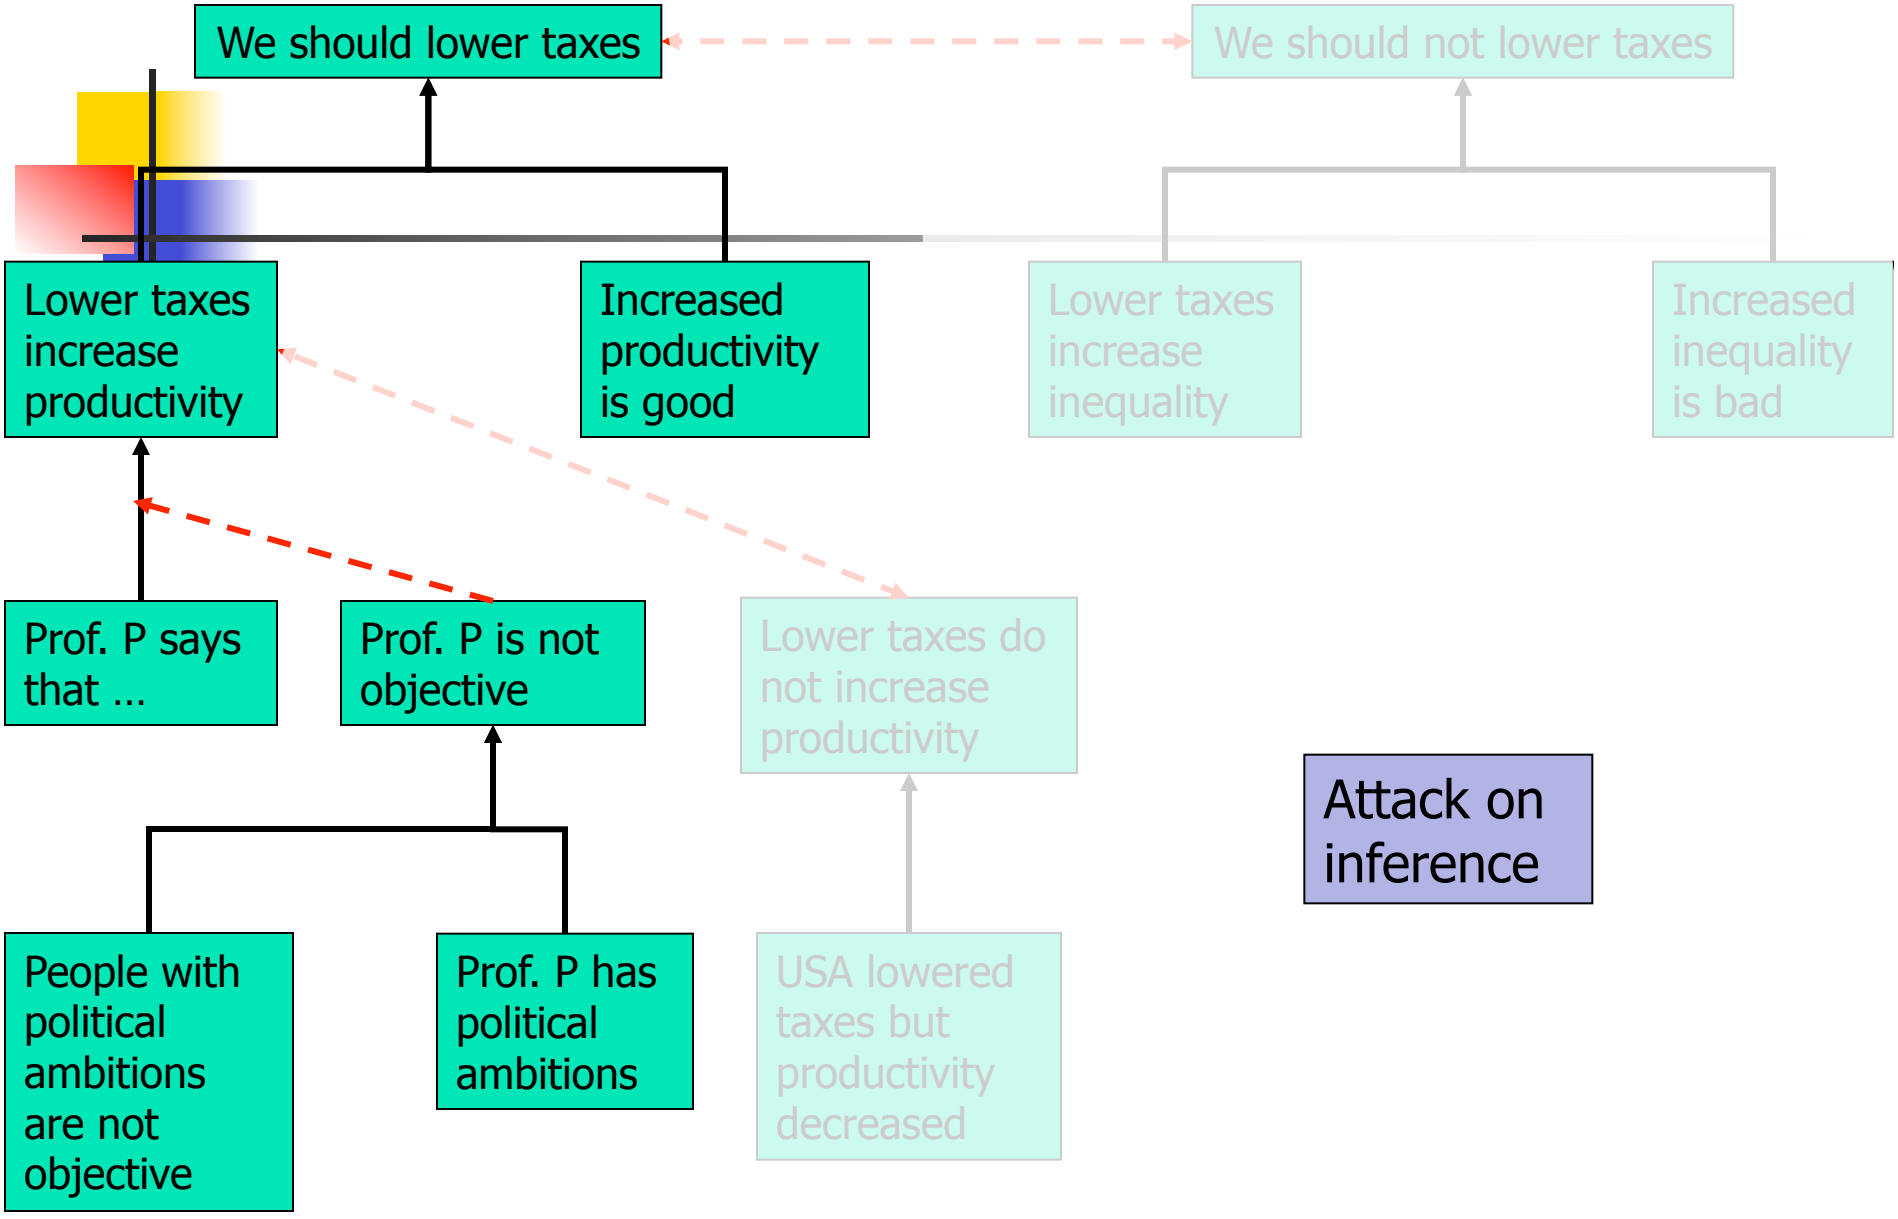
We should lower taxes
We should not lower taxes
Lower taxes increase productivity
Increased productivity is good
Lower taxes increase inequality
Increased inequality is bad
Prof. P says that …
Prof. P is not objective
Lower taxes do not increase productivity
Attack on inference
People with political ambitions are not objective
Prof. P has political ambitions
USA lowered taxes but productivity decreased

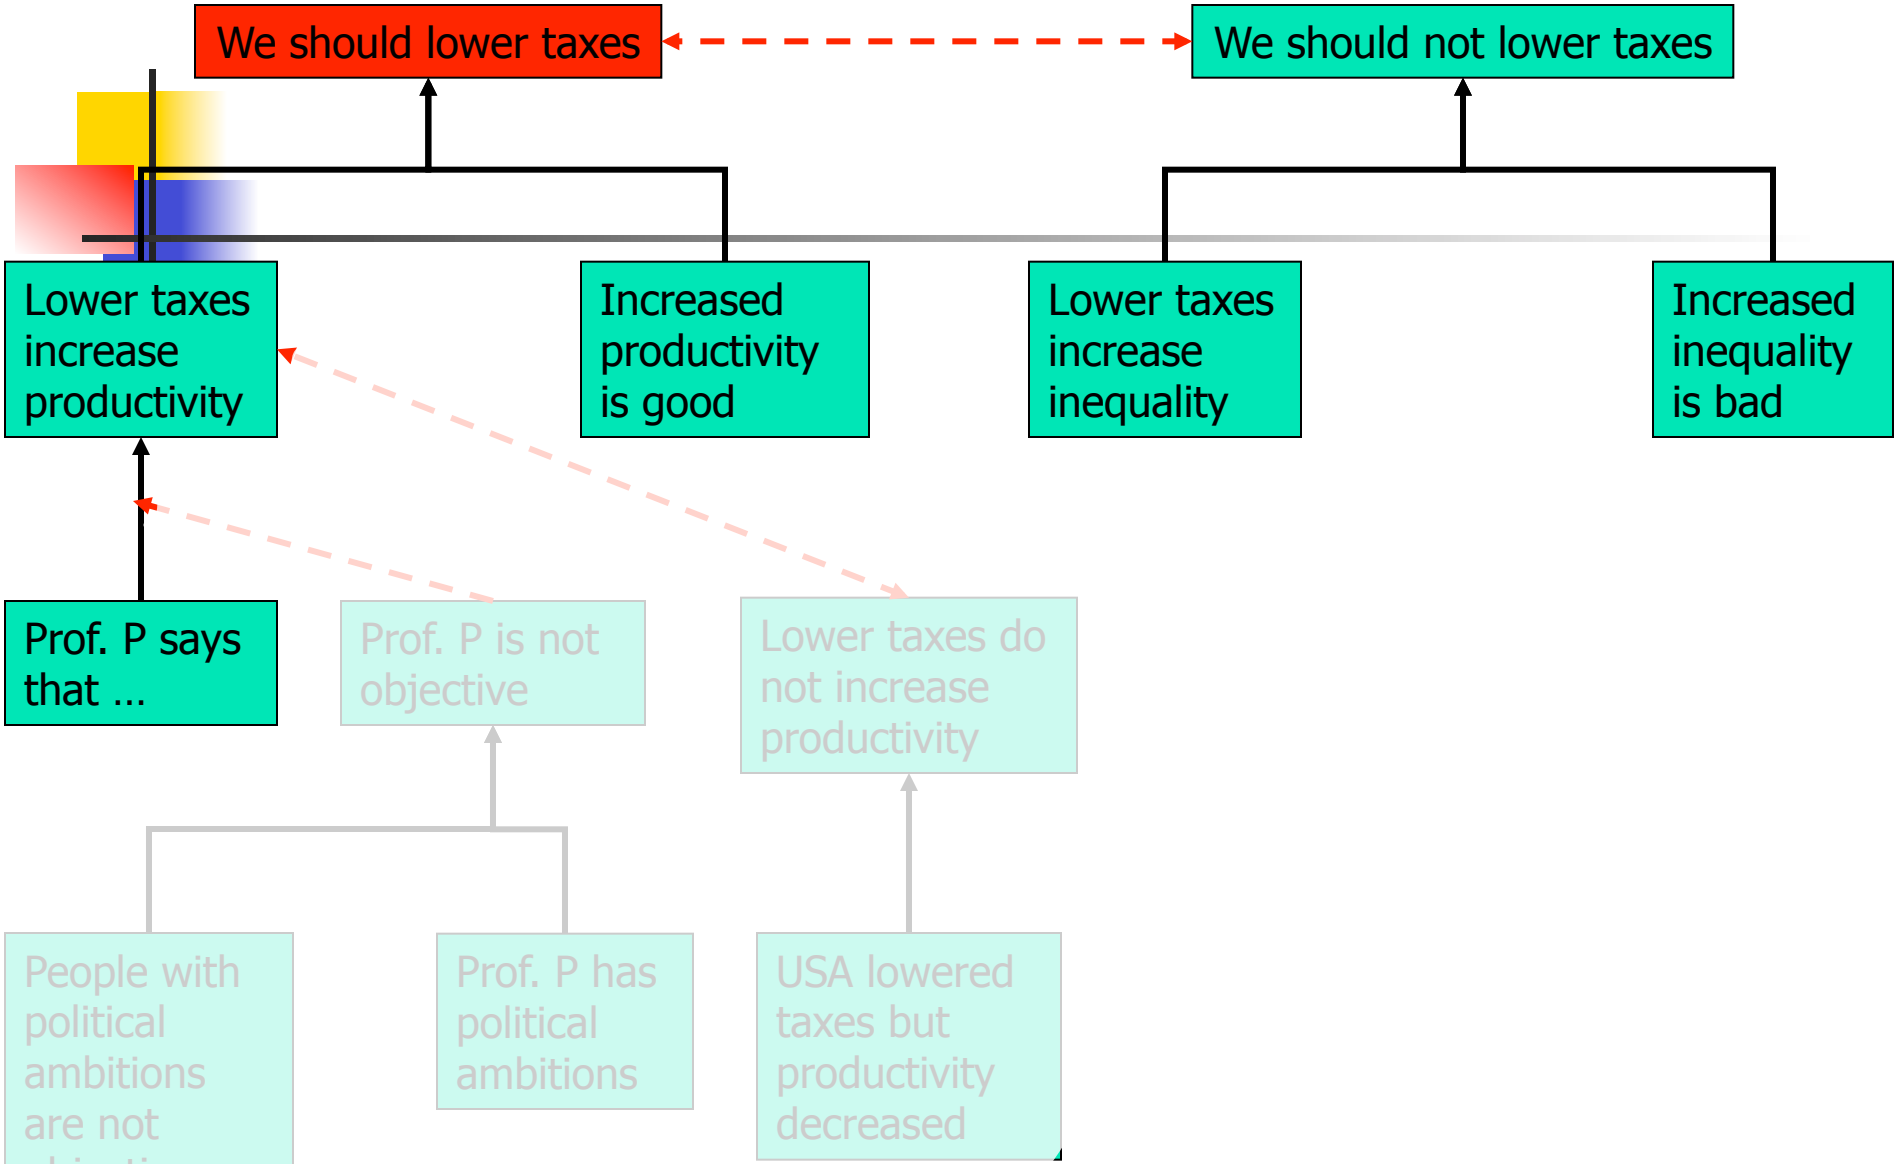
We should lower taxes
We should not lower taxes
Lower taxes increase productivity
Increased productivity is good
Lower taxes increase inequality
Increased inequality is bad
Prof. P says that …
Prof. P is not objective
Lower taxes do not increase productivity
People with political ambitions are not objective
Prof. P has political ambitions
USA lowered taxes but productivity decreased

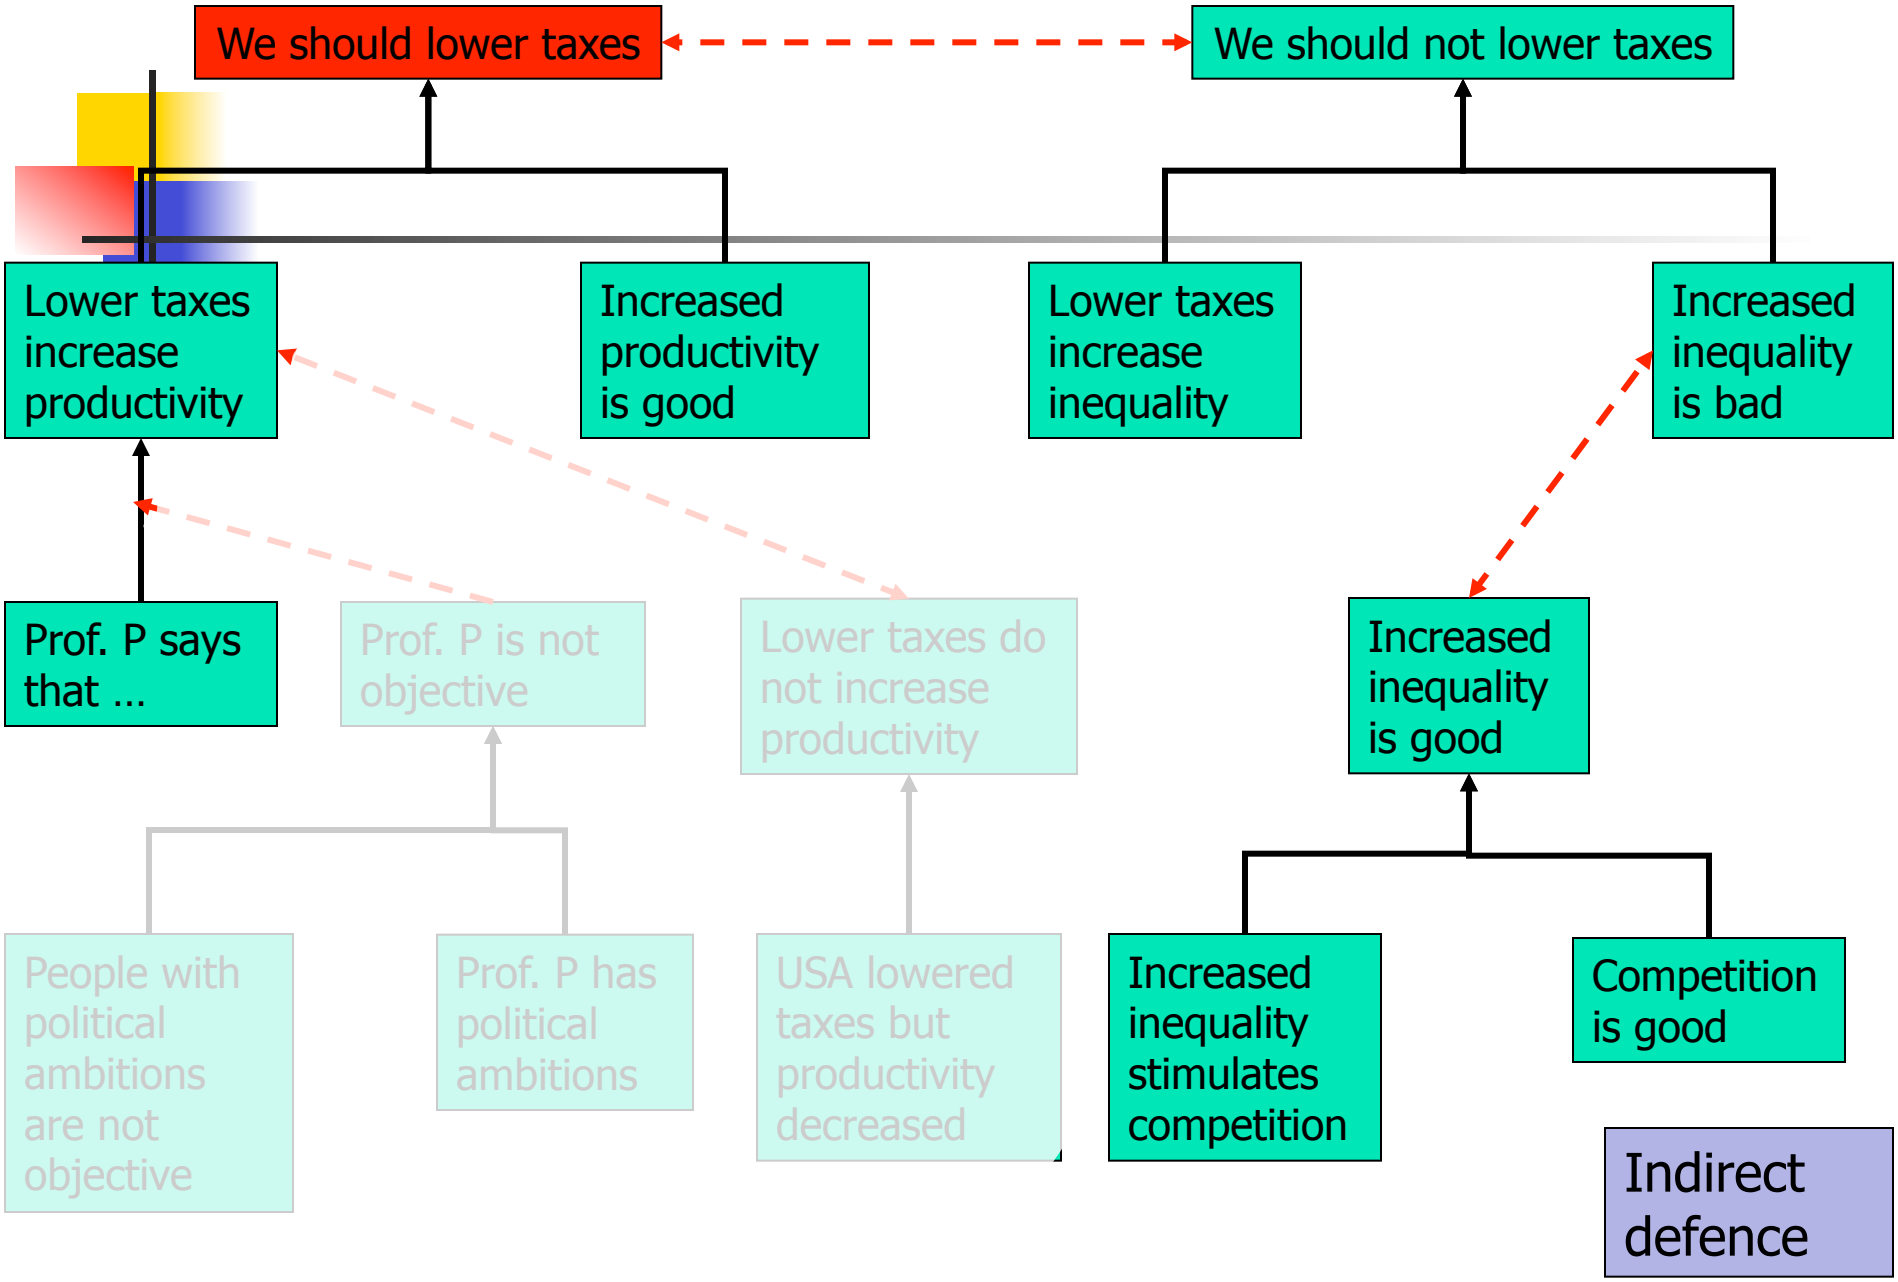
We should lower taxes
We should not lower taxes
Lower taxes increase productivity
Increased productivity is good
Lower taxes increase inequality
Increased inequality is bad
Prof. P says that …
Prof. P is not objective
Lower taxes do not increase productivity
Increased inequality is good
People with political ambitions are not objective
Prof. P has political ambitions
USA lowered taxes but productivity decreased
Increased inequality stimulates competition
Competition is good
Indirect defence

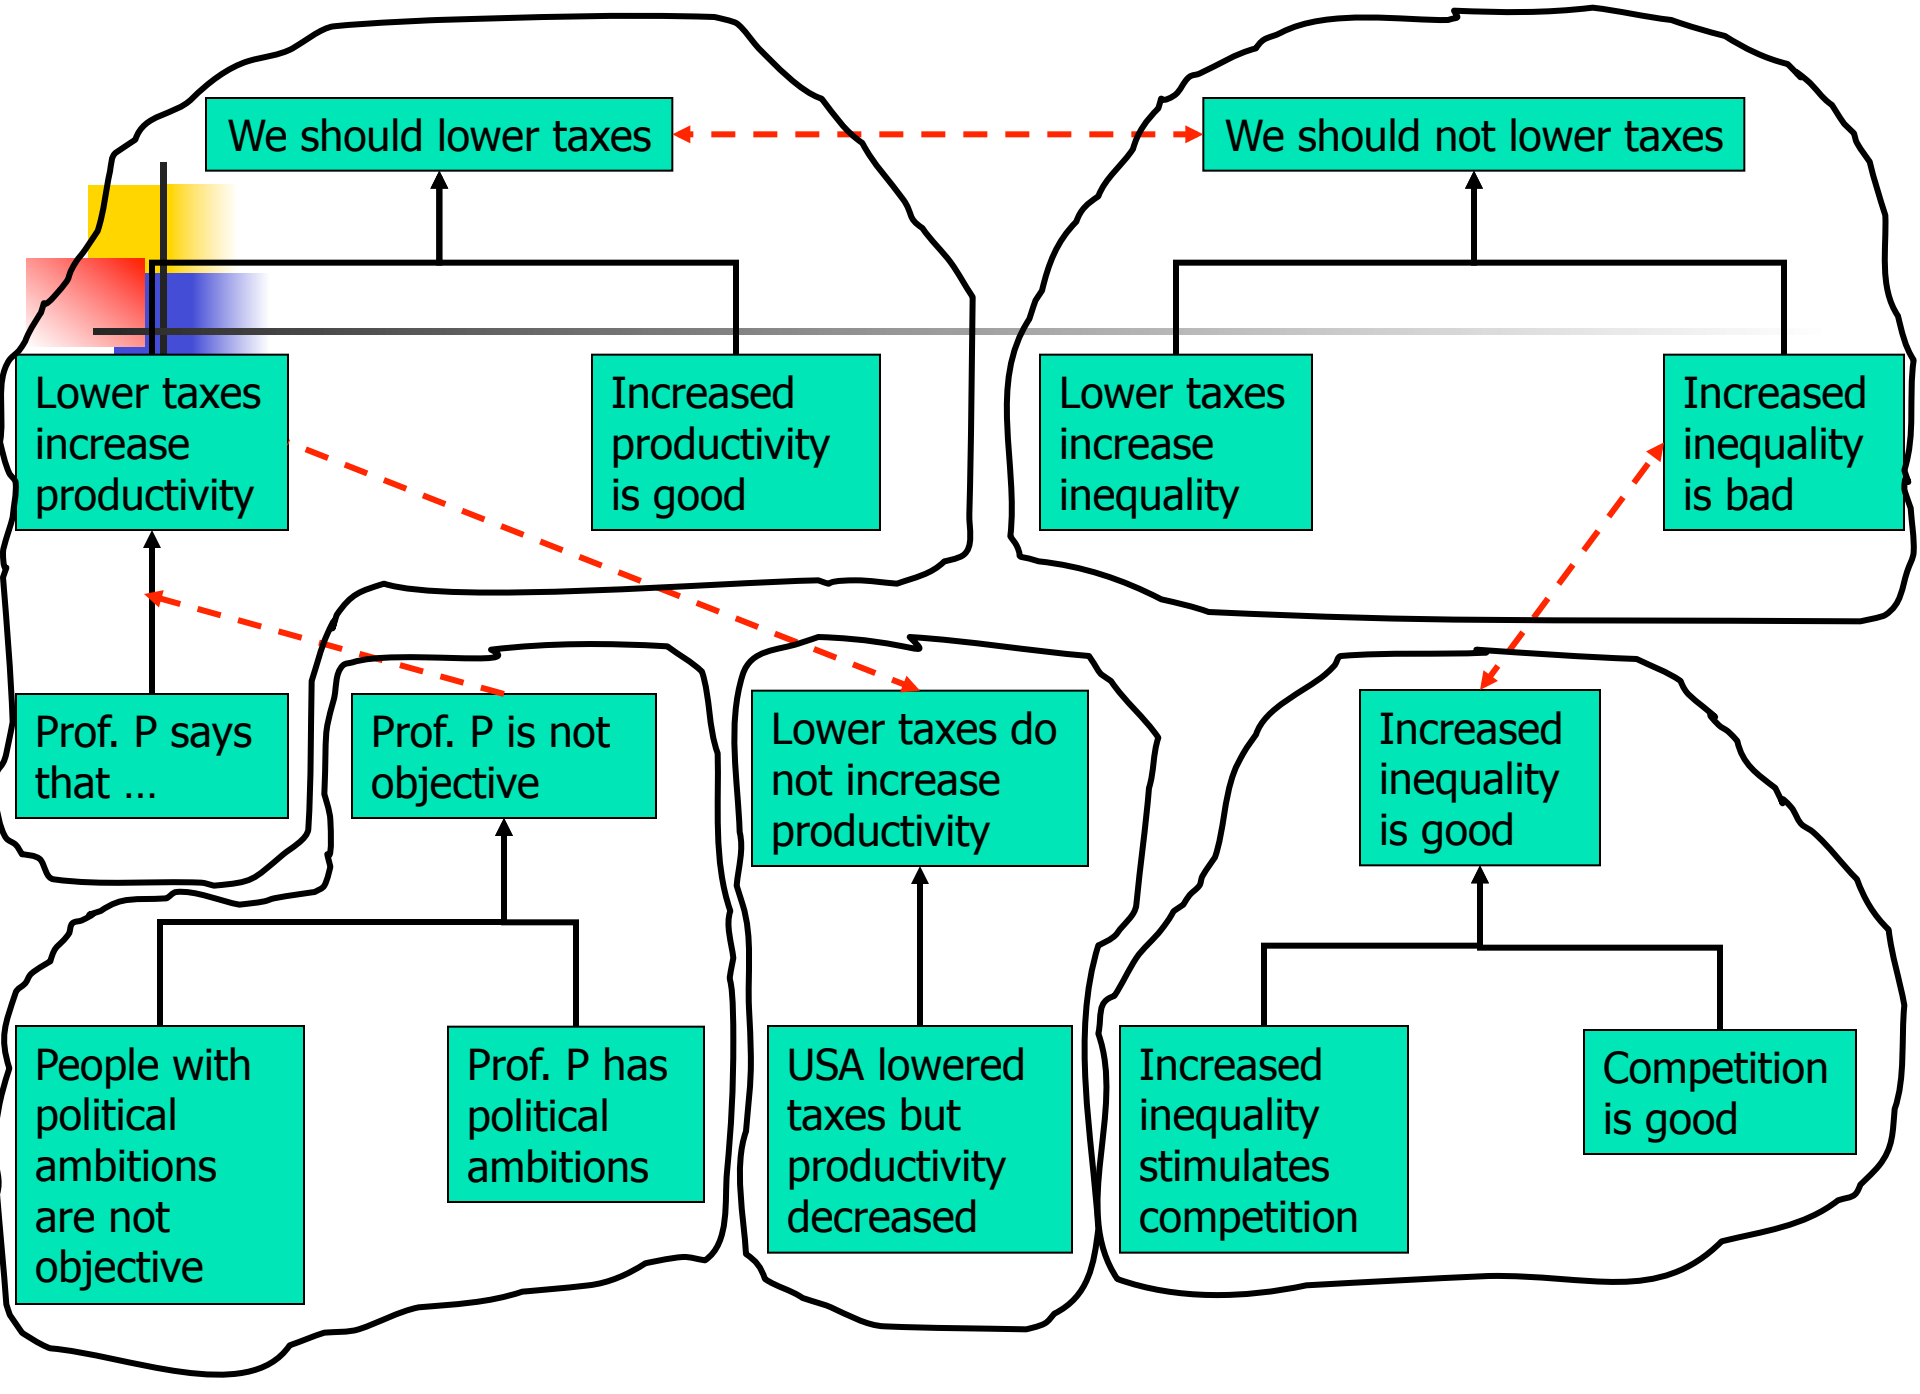
We should lower taxes
We should not lower taxes
Lower taxes increase productivity
Increased productivity is good
Lower taxes increase inequality
Increased inequality is bad
Prof. P says that …
Prof. P is not objective
Lower taxes do not increase productivity
Increased inequality is good
People with political ambitions are not objective
Prof. P has political ambitions
USA lowered taxes but productivity decreased
Increased inequality stimulates competition
Competition is good

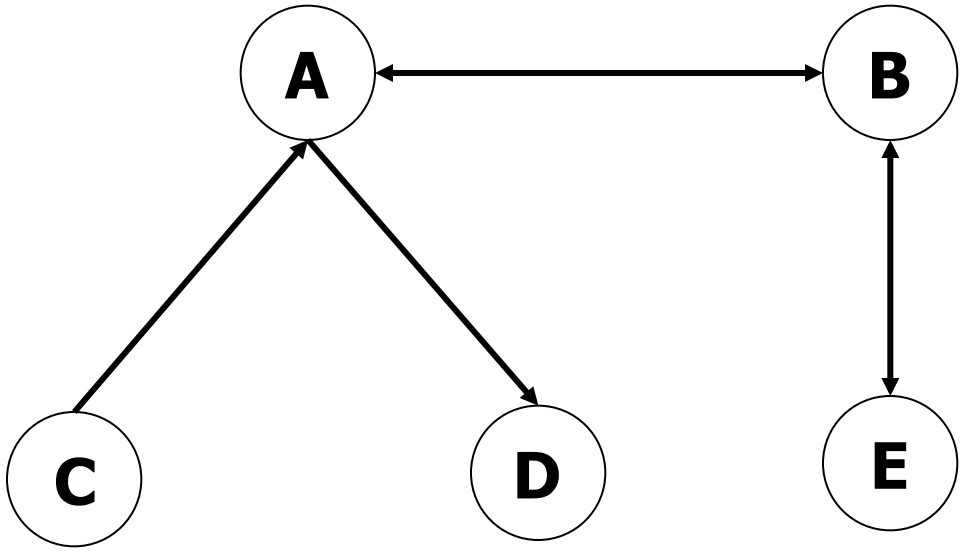

P.M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and *n*–person games. *Artificial Intelligence*, 77:321–357, 1995.

1. An argument is *In* iff all arguments that attack it are *Out*.
2. An argument is *Out* iff some argument that attacks it is *In*.

Grounded semantics *minimises In* labelling
Preferred semantics *maximises In* labelling
Stable semantics labels *all* nodes

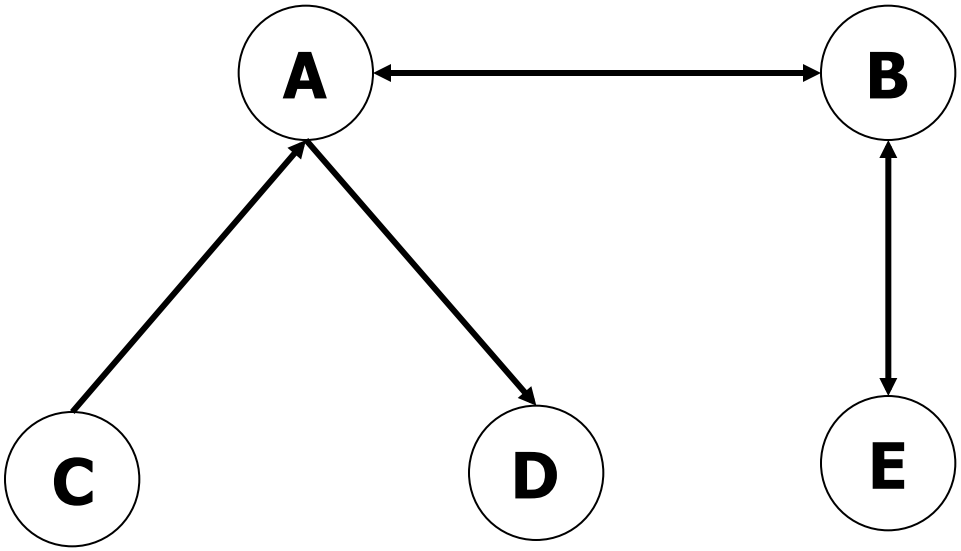

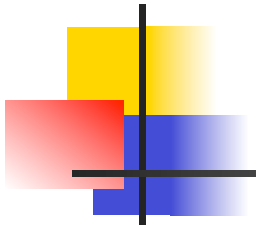

Properties

- There always exists exactly one grounded labelling
- There exists at least one preferred labelling
- Every stable labelling is preferred (but not v.v.)
- The grounded labelling is a subset of all preferred and stable labellings
- Every finite Dung graph without attack cycles has a unique labelling (which is the same in all semantics)
- ...

1. An argument is *In* iff all arguments that attack it are *Out*.
2. An argument is *Out* iff some argument that attacks it is *In*.

Stable semantics labels *all* nodes
Grounded semantics *minimises In* labelling
Preferred semantics *maximises In* labelling

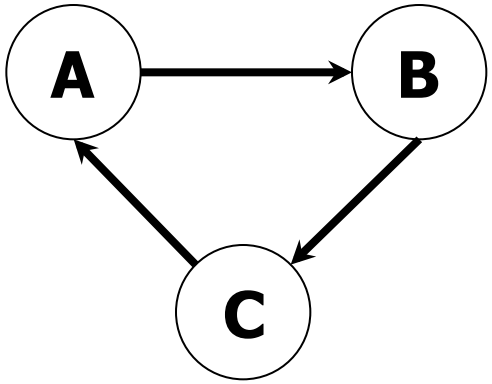

1. An argument is *In* iff all arguments that attack it are *Out*.
2. An argument is *Out* iff some argument that attacks it is *In*.

Stable semantics labels *all* nodes
Grounded semantics *minimises In* labelling
Preferred semantics *maximises In* labelling

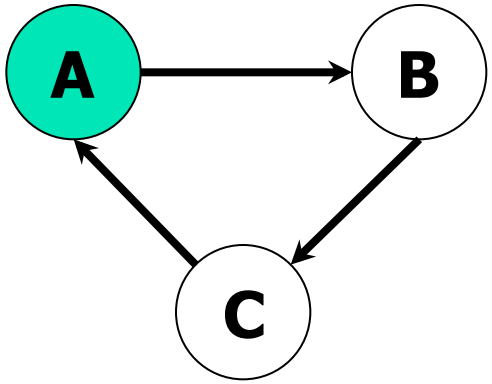

1. An argument is *In* iff all arguments that attack it are *Out*.
2. An argument is *Out* iff some argument that attacks it is *In*.

Stable semantics labels *all* nodes
Grounded semantics *minimises In* labelling
Preferred semantics *maximises In* labelling

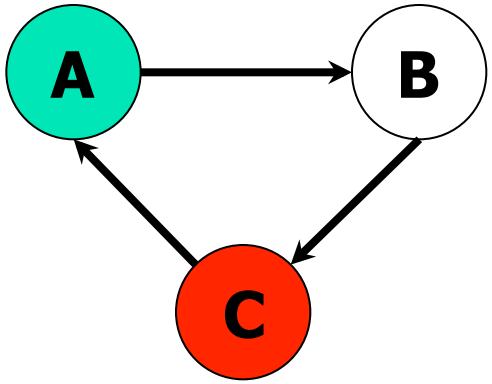

1. An argument is *In* iff all arguments that attack it are *Out*.
2. An argument is *Out* iff some argument that attacks it is *In*.

Stable semantics labels *all* nodes
Grounded semantics *minimises In* labelling
Preferred semantics *maximises In* labelling

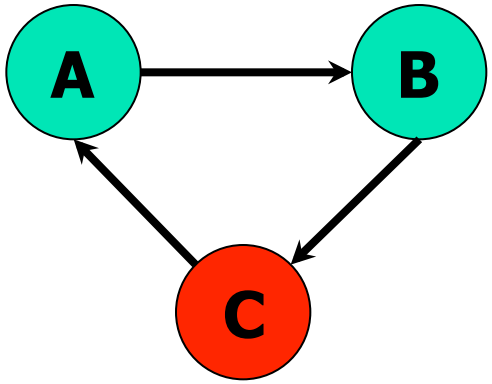

1. An argument is *In* iff all arguments that attack it are *Out*.
2. An argument is *Out* iff some argument that attacks it is *In*.

Stable semantics labels *all* nodes
Grounded semantics *minimises In* labelling
Preferred semantics *maximises In* labelling

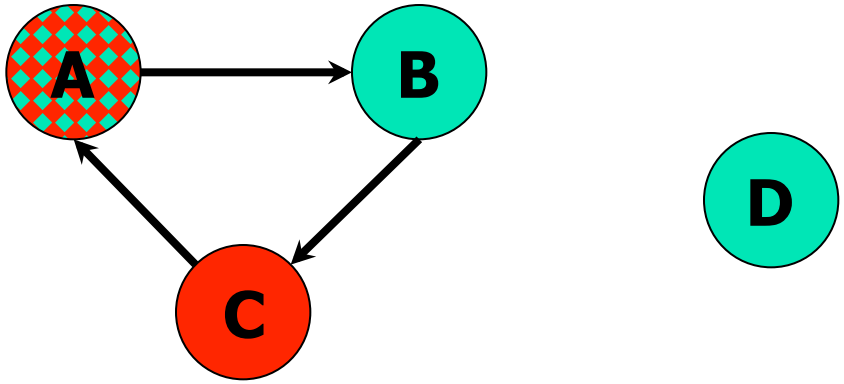

1. An argument is *In* iff all arguments that attack it are *Out*.
2. An argument is *Out* iff some argument that attacks it is *In*.

Stable semantics labels *all* nodes
Grounded semantics *minimises In* labelling
Preferred semantics *maximises In* labelling

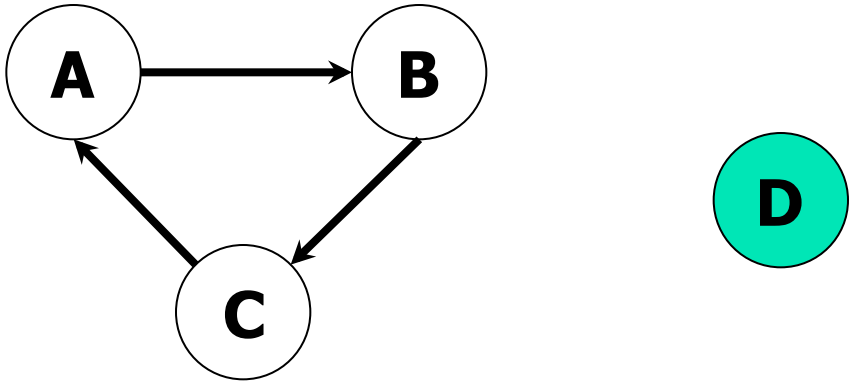

Difference between grounded and preferred labellings

1. An argument is *In* iff all arguments that attack it are *Out*.
2. An argument is *Out* iff some argument that attacks it is *In*.

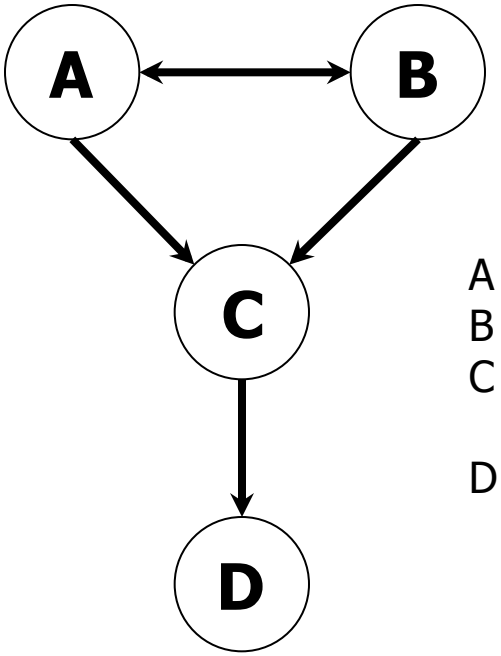

A = Merkel is German since she has a German name
B = Merkel is Belgian since she is often seen in Brussels
C = Merkel is a fan of Oranje since she wears an orange shirt (unless she is German or Belgian)
D = Merkel is not a fan of Oranje since she looks like someone who does not like football

(Generalisations are left implicit)

The grounded labelling

1. An argument is *In* iff all arguments that attack it are *Out*.
2. An argument is *Out* iff some argument that attacks it is *In*.

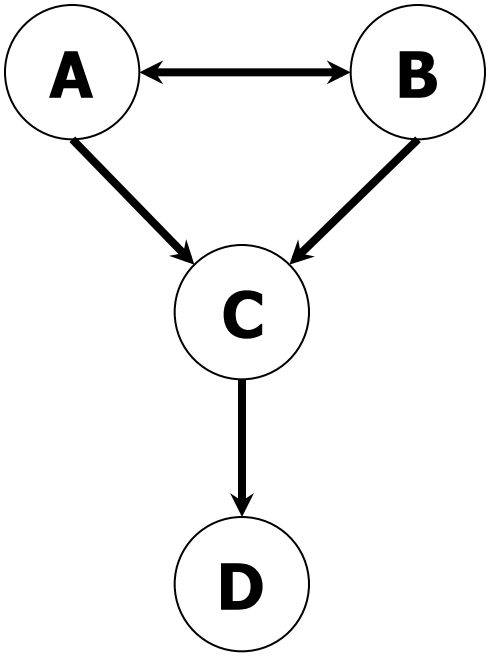

The preferred labellings

1. An argument is *In* iff all arguments that attack it are *Out*.
2. An argument is *Out* iff some argument that attacks it is *In*.

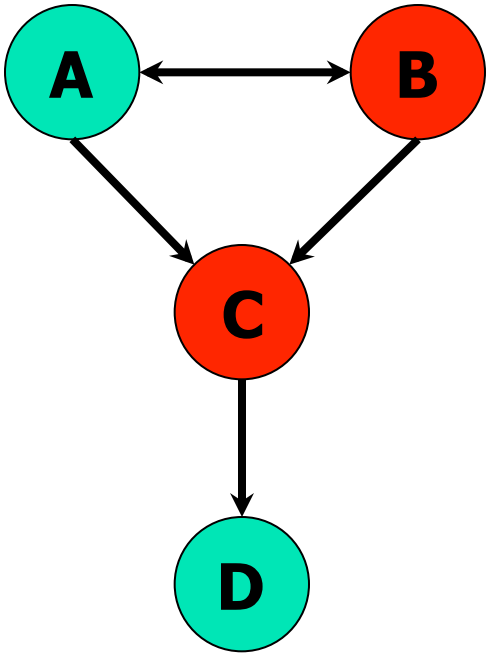

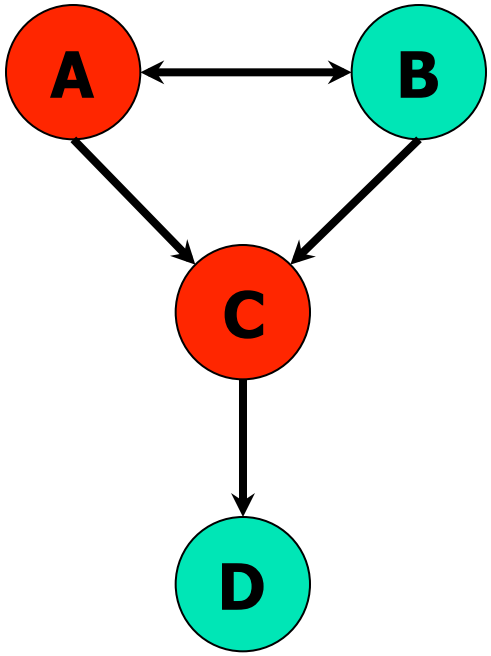

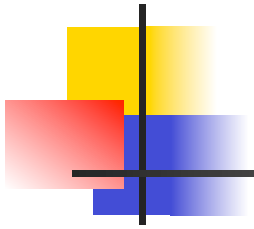

Justification status of arguments

- A is justified if A is *In* in all labellings
- A is overruled if A is *Out* in all labellings
- A is defensible otherwise

Argument status in grounded and preferred semantics

Grounded semantics:
all arguments defensible

Preferred semantics:
A and B defensible
C overruled
D justified

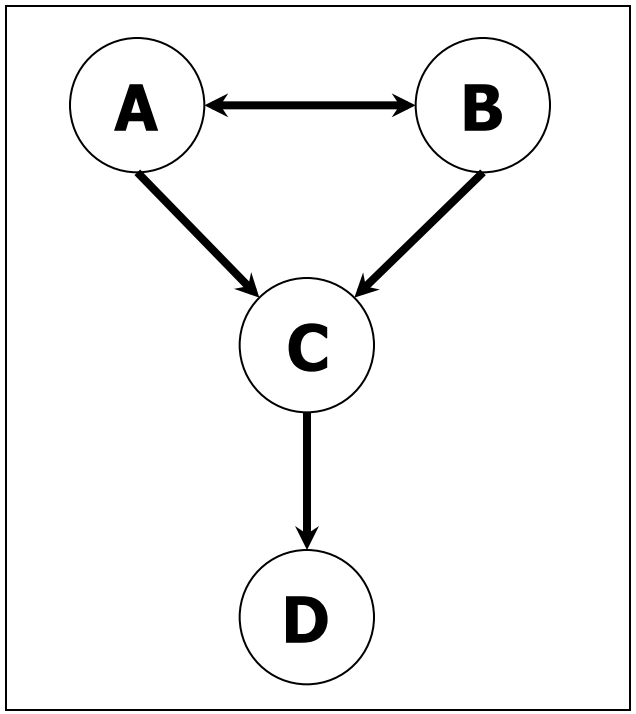

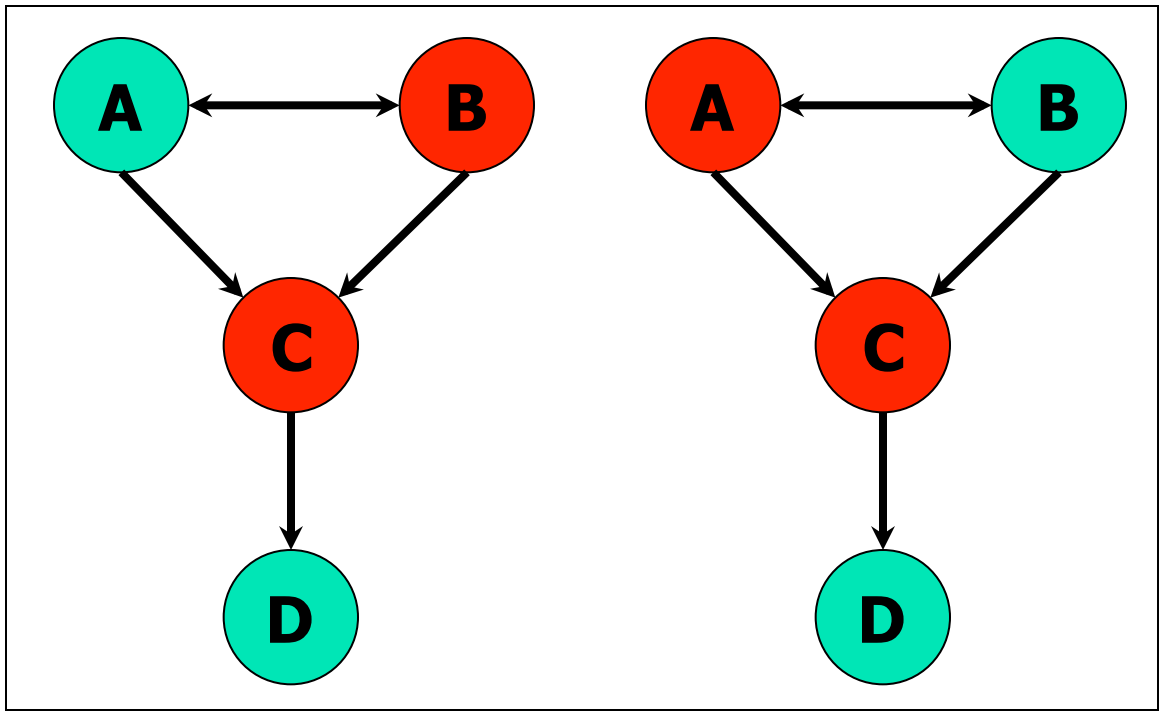

Labellings and extensions

Given an argumentation framework $AF = \langle Args, attack \rangle$:

$S \subseteq Args$ is a stable/preferred/grounded argument extension iff $S = In$ for some stable/preferred/grounded labelling

Grounded extension

- *A* is acceptable wrt *S* (or *S* defends *A*) if all attackers of *A* are attacked by *S*
  - *S* attacks *A* if an argument in *S* attacks *A*
- Let *AF* be an abstract argumentation framework
  - $F^0_{AF} = \emptyset$
  - $F^{i+1}_{AF} = \{A \in Args \mid A \text{ is acceptable wrt } F^i_{AF}\}$
  - $F^\infty_{AF} = \cup^\infty_{i=0} (F^{i+1}_{AF})$
- If no argument has an infinite number of attackers, then F^∞_{AF} is the grounded extension of AF (otherwise it is included)

S defends A if all attackers of A are attacked by a member of S

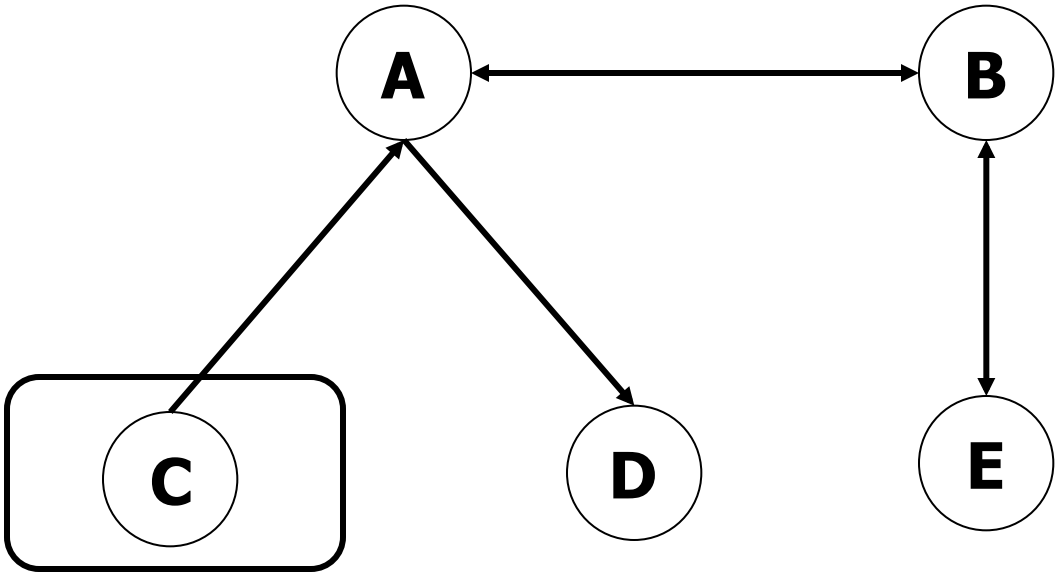

$F^1 = \{A\}$

S defends A if all attackers of A are attacked by a member of S

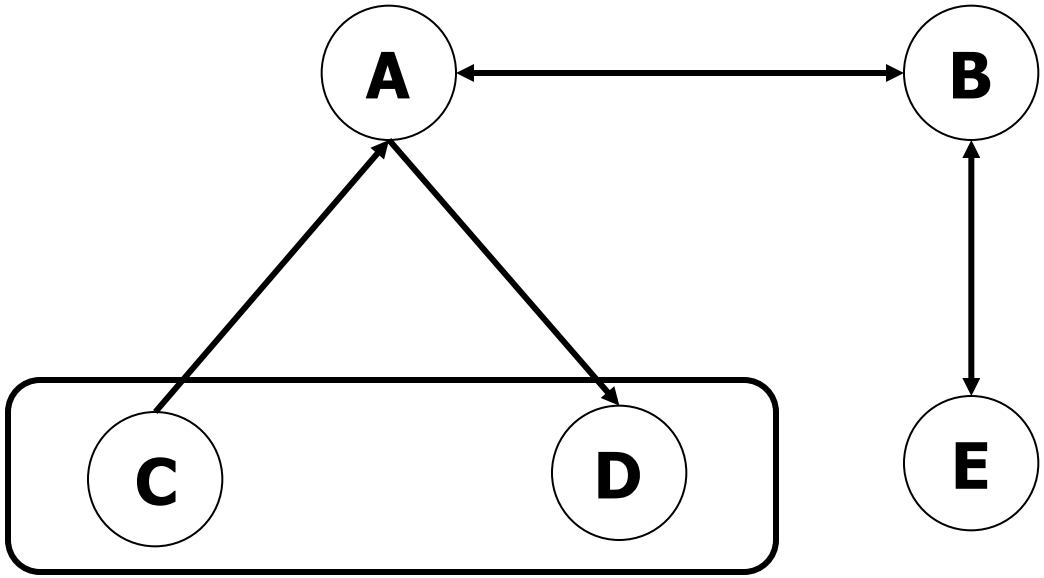

$F^1 = \{A\}$

$F^2 = \{A,D\}$

S defends A if all attackers of A are attacked by a member of S

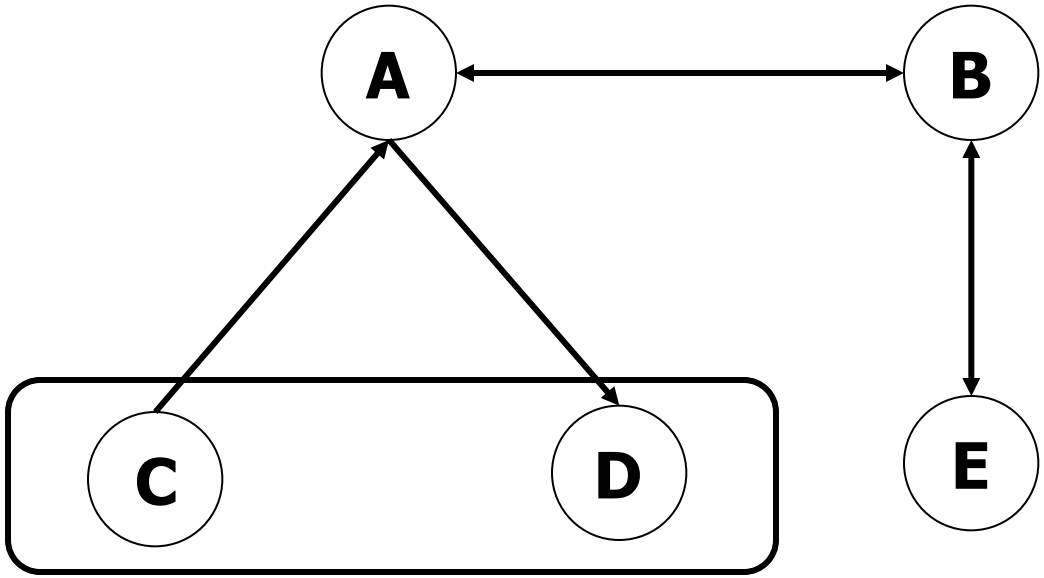

$F^1 = \{A\}$
$F^2 = \{A,D\}$
$F^3 = F^2$

S defends A if all defeaters of A are attacked by a member of S

S is admissible if it is conflict-free and defends all its members

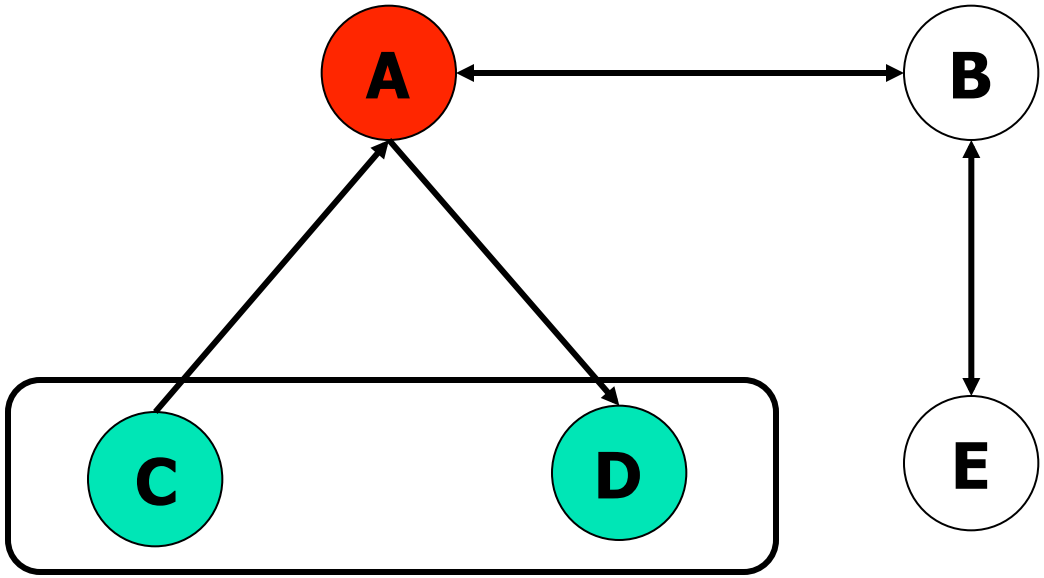

Grounded

Stable extensions

- Dung (1995):
  - *S* is conflict-free if no member of *S* attacks a member of S
  - *S* is a stable extension if it is conflict-free and attacks all arguments outside it
- Recall:
  - *S* is a stable argument extension if *S* = *In* for some stable labelling
- **Proposition**: *S* is a stable argument extension iff *S* is a stable extension

Preferred extensions

- Dung (1995):
  - *S* is conflict-free if no member of *S* attacks a member of S
  - *S* is admissible if it is conflict-free and all its members are acceptable wrt *S*
  - *S* is a preferred extension if it is $\subseteq$-maximally admissible
- Recall:
  - *S* is a preferred argument extension if *S* = *In* for some preferred labelling
- **Proposition**: *S* is a preferred argument extension iff *S* is a preferred extension

S defends A if all attackers of A are attacked by a member of S

S is admissible if it is conflict-free and defends all its members

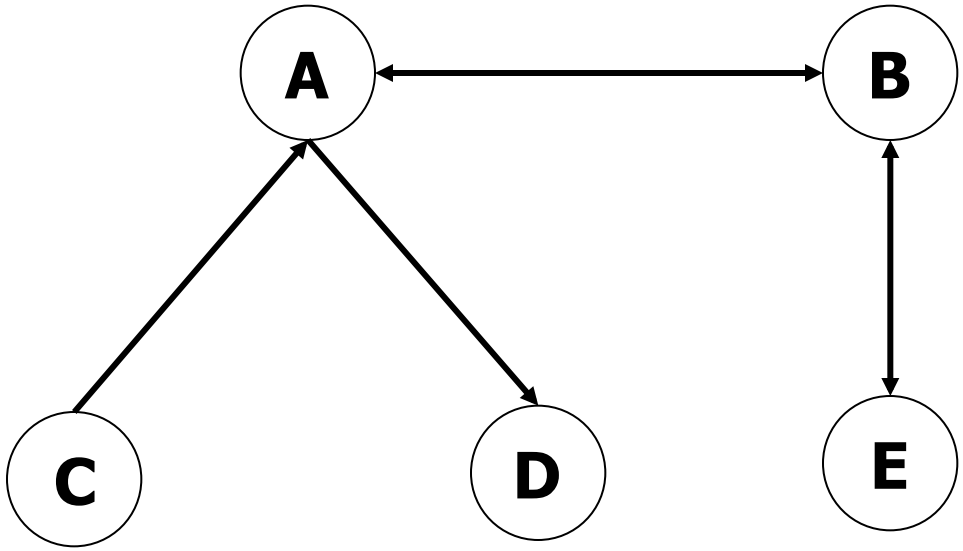

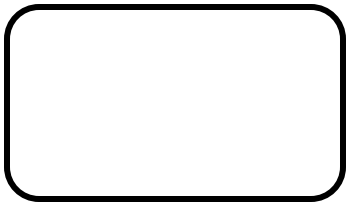

Admissible?

S defends A if all defeaters of A are attacked by a member of S

S is admissible if it is conflict-free and defends all its members

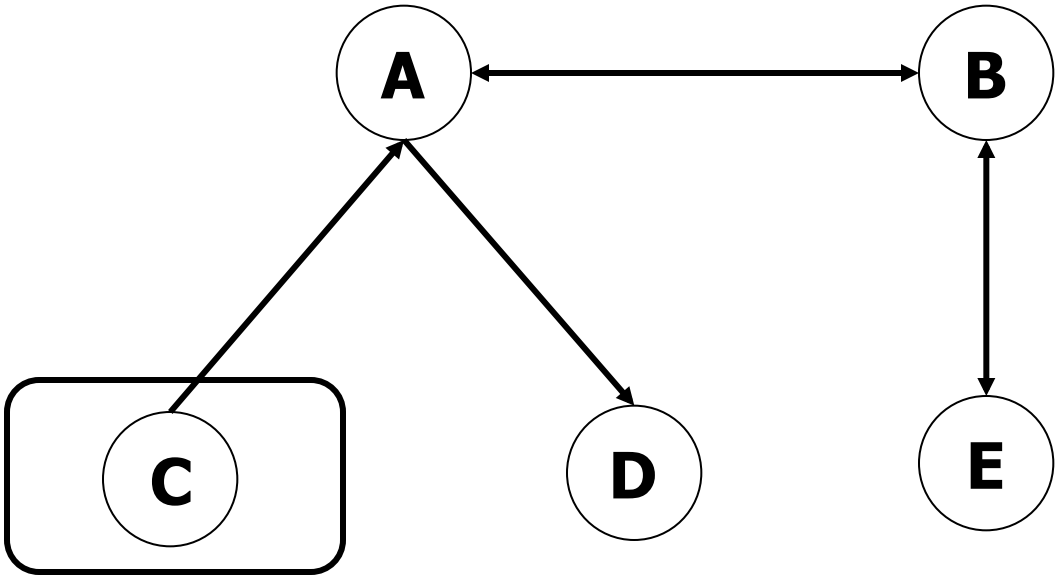

Admissible?

S defends A if all defeaters of A are attacked by a member of S

S is admissible if it is conflict-free and defends all its members

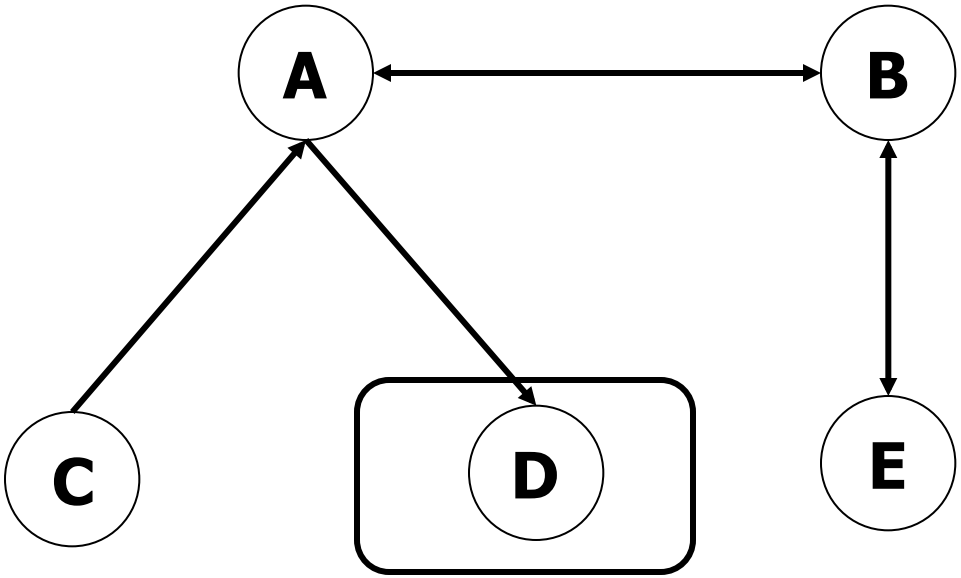

Admissible?

S defends A if all defeaters of A are attacked by a member of S

S is admissible if it is conflict-free and defends all its members

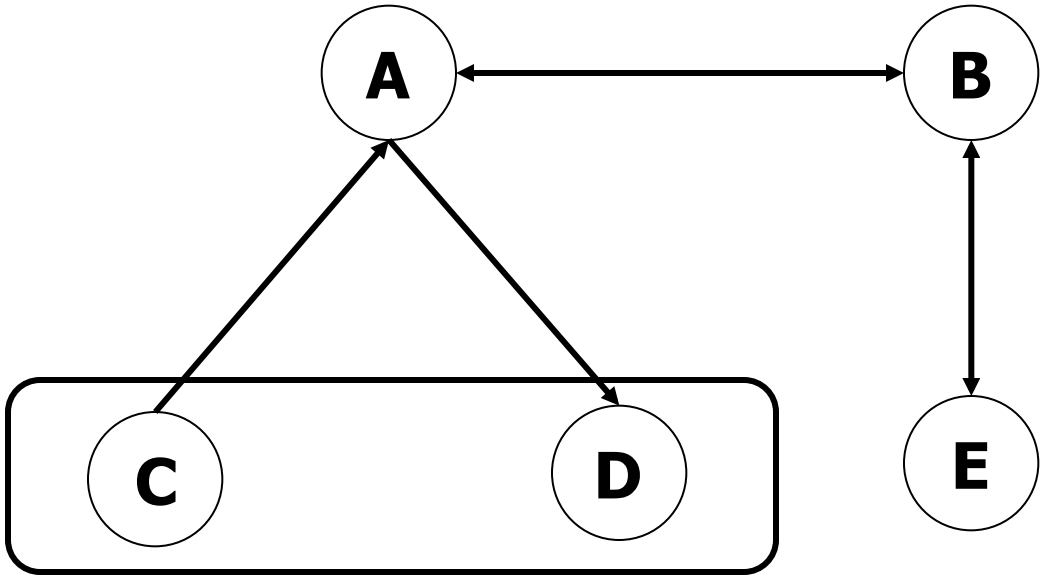

Admissible?

S defends A if all defeaters of A are attacked by a member of S

S is admissible if it is conflict-free and defends all its members

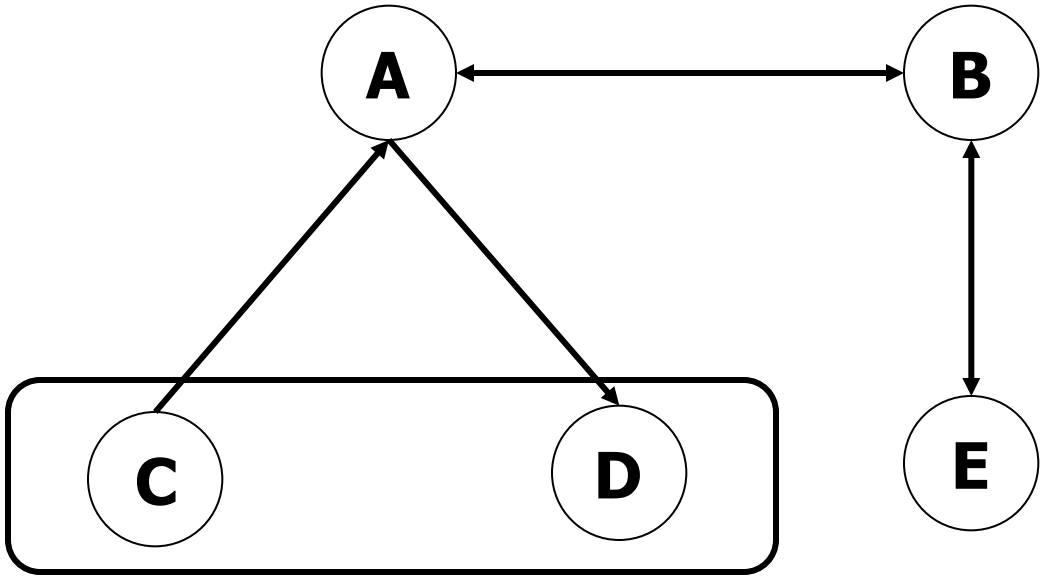

Preferred?

S is preferred if it is maximally admissible

S defends A if all defeaters of A are attacked by a member of S

S is admissible if it is conflict-free and defends all its members

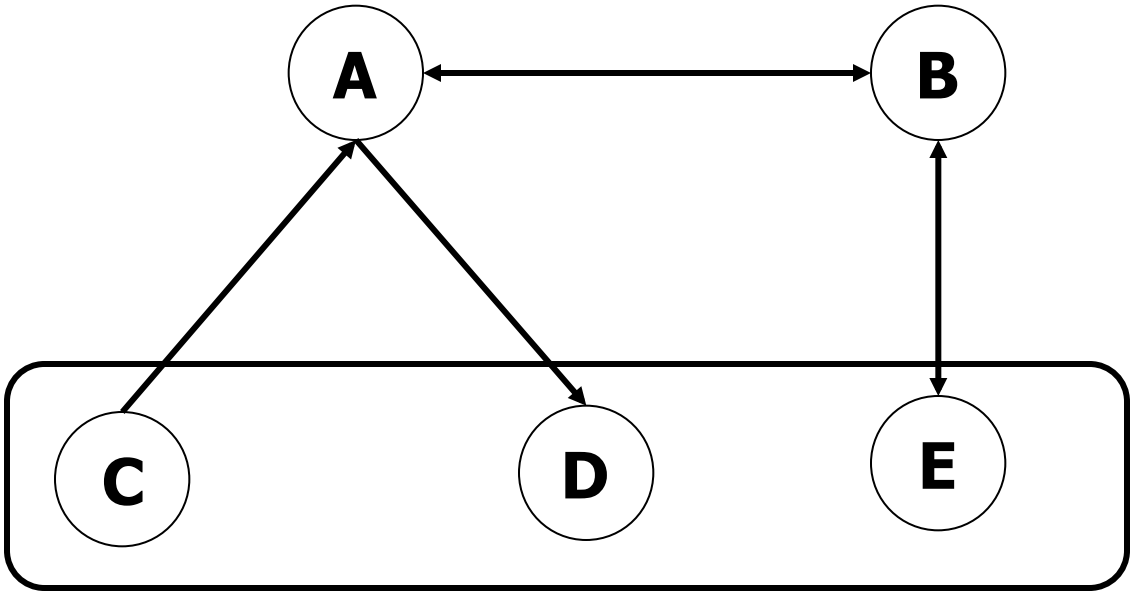

Preferred?

S is preferred if it is maximally admissible

S defends A if all defeaters of A are attacked by a member of S

S is admissible if it is conflict-free and defends all its members

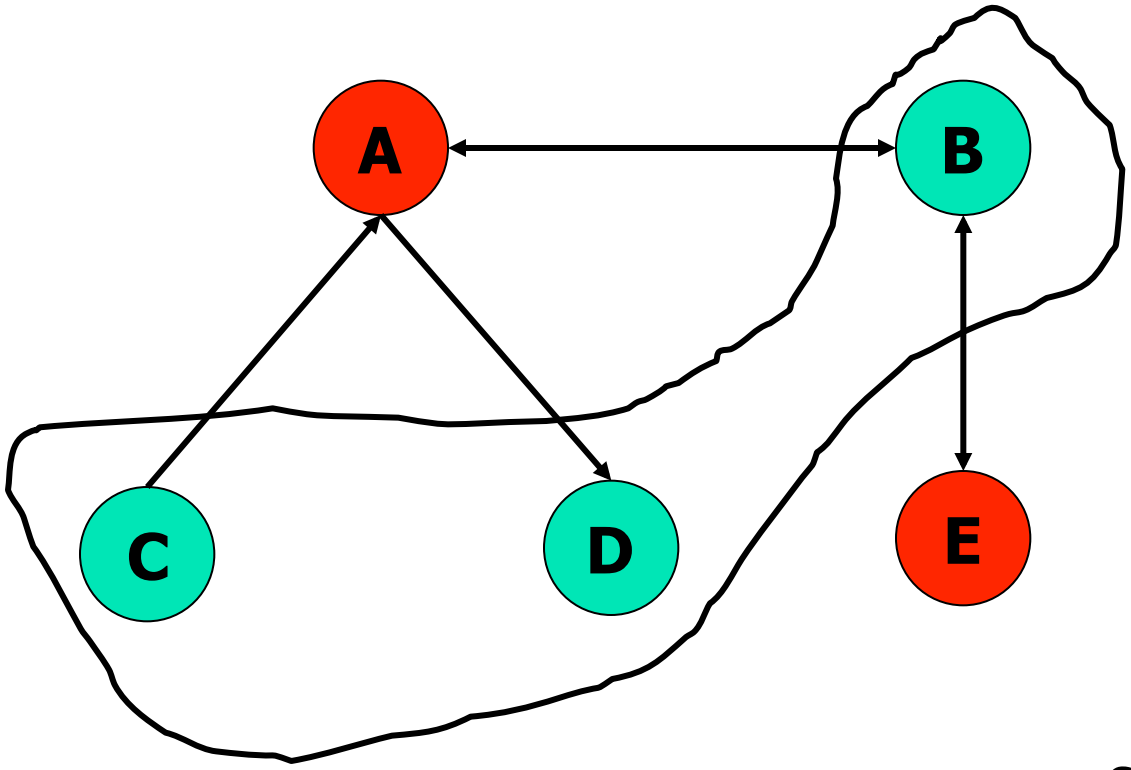

Preferred?

S is preferred if it is maximally admissible

Proof theory for abstract argumentation

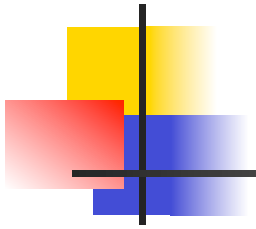

- Argument games between proponent *P* and opponent *O*:
  - Proponent starts with an argument
  - Then each party replies with a suitable attacker
  - A winning criterion
    - E.g. the other player cannot move
- Acceptability status corresponds to existence of a winning strategy.

Strategies

- A strategy for player *p* is a partial game tree:
  - Every branch is a game (sequence of allowable moves)
  - The tree only branches after moves by *p*
  - The children of *p*'s moves are all the legal moves by the other player

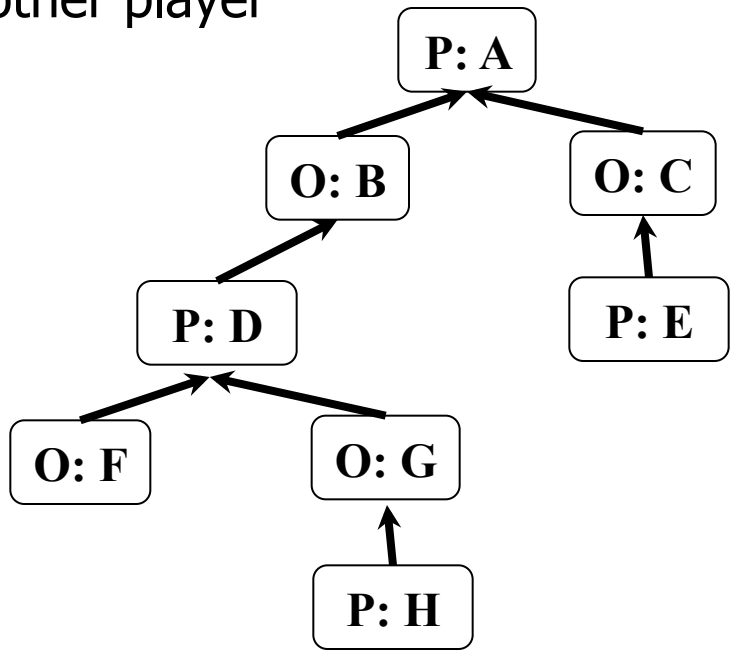

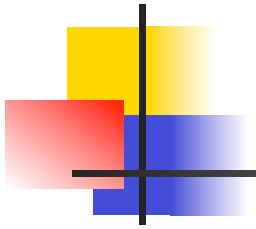

Strategies

- A strategy for player *p* is winning iff *p* wins all games in the strategy
- Let ***S*** be an argument game: *A* is ***S***-provable iff *P* has a winning strategy in an ***S***-game that begins with *A*

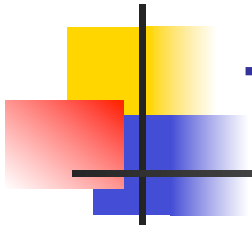

The G-game for grounded semantics:

- A sound and complete game:
  - Each move must reply to the previous move
  - Proponent cannot repeat his moves
  - Proponent moves strict attackers, opponent moves attackers
  - A player wins iff the other player cannot move
- **Proposition:** *A* is in the grounded extension iff A is G-provable

An attack graph

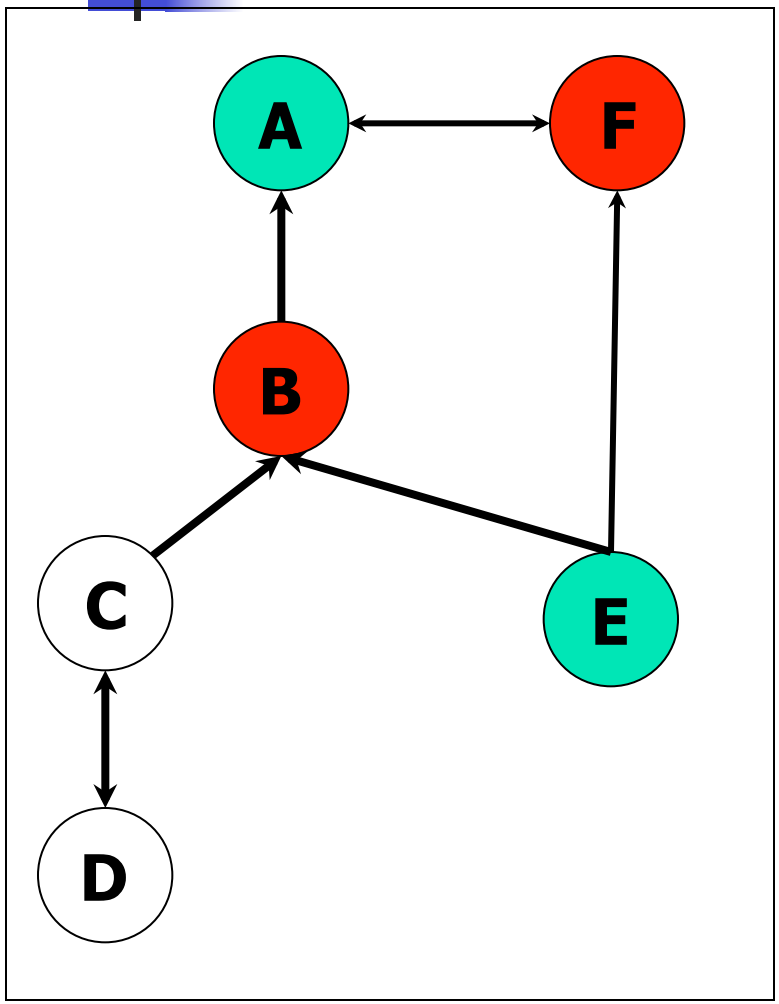

A game tree

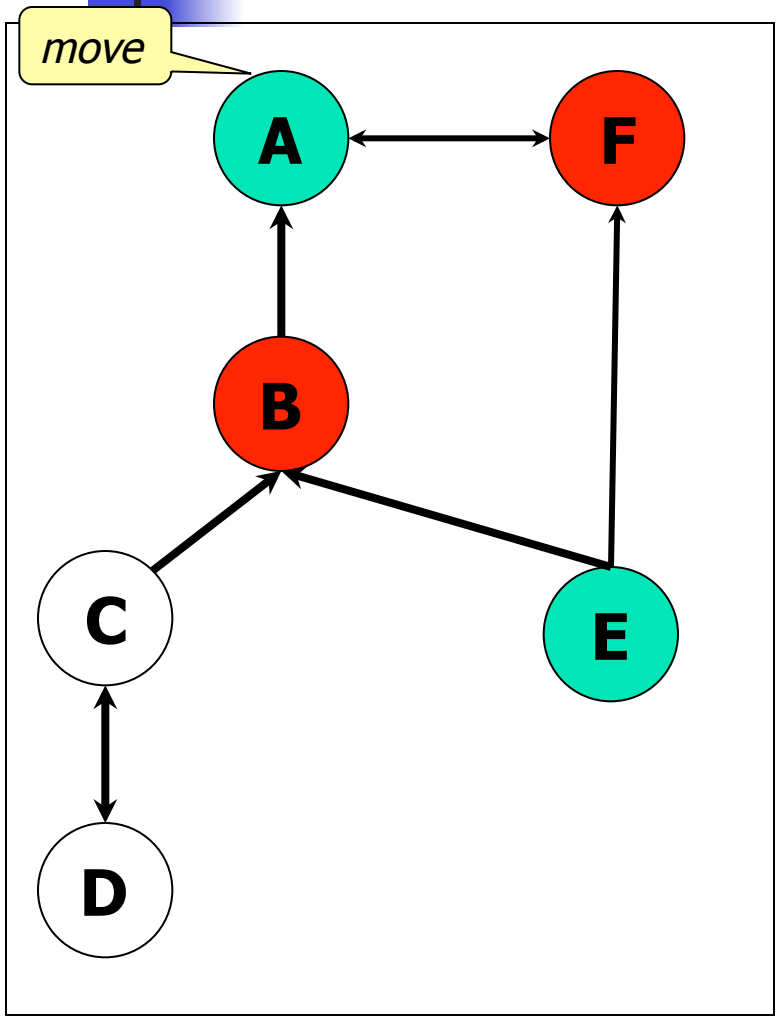

P: A

A game tree

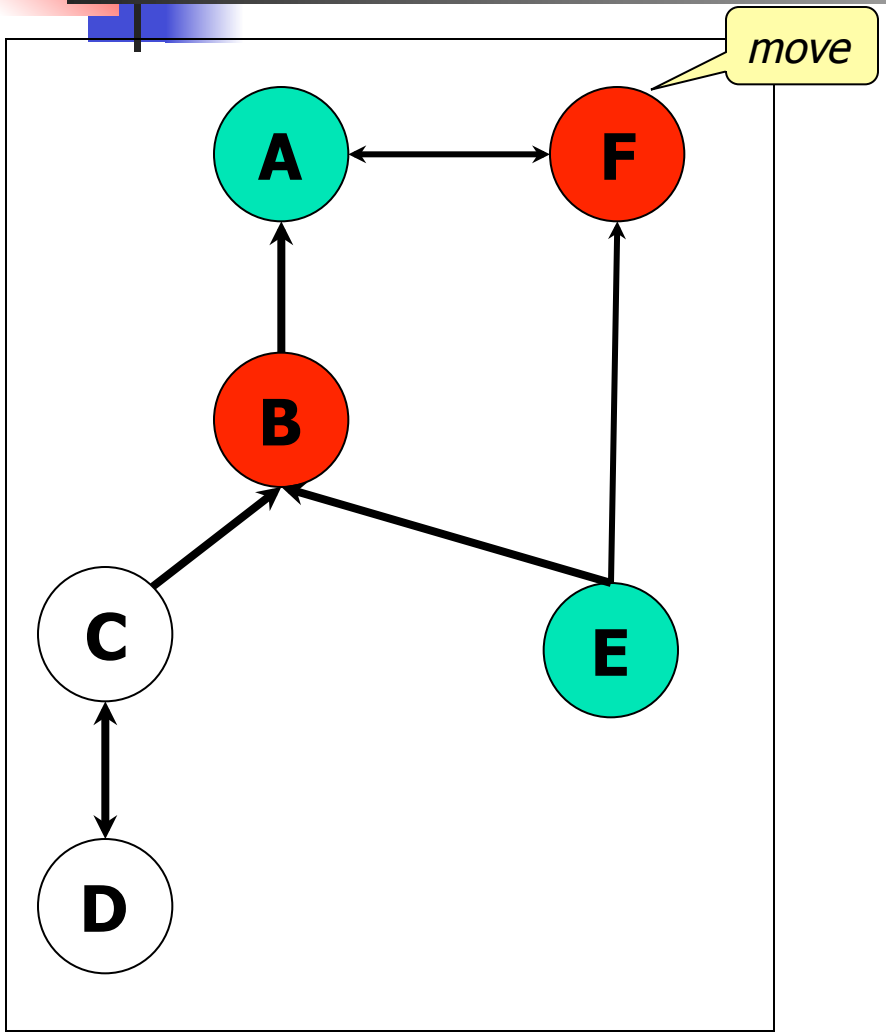

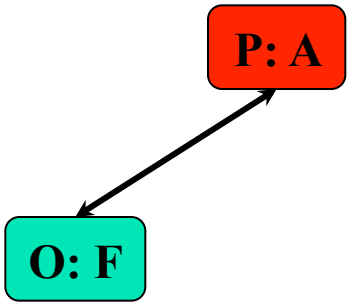

A game tree

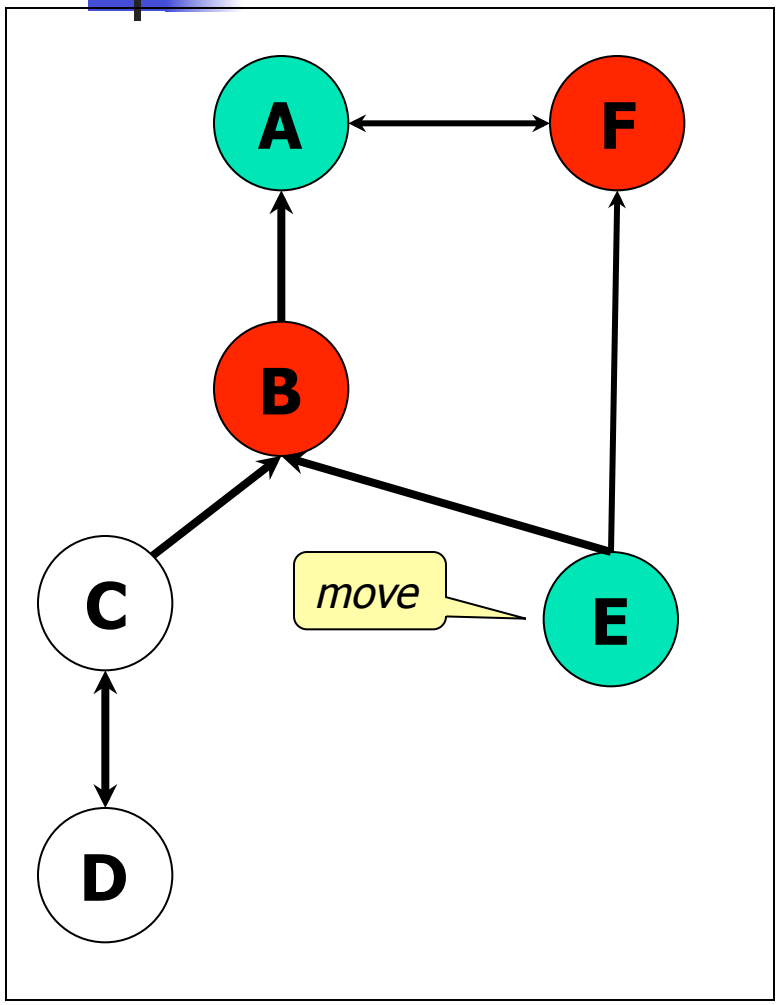

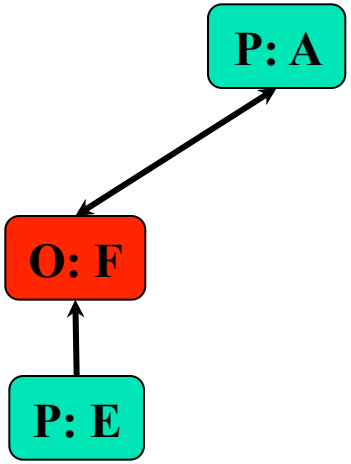

A game tree

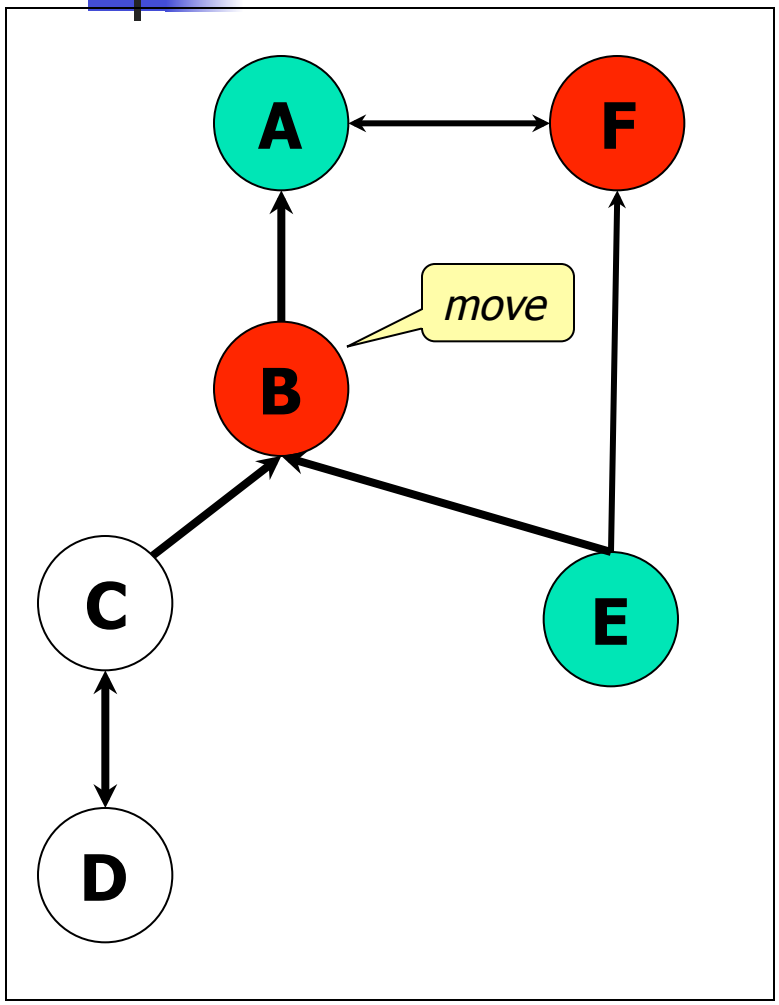

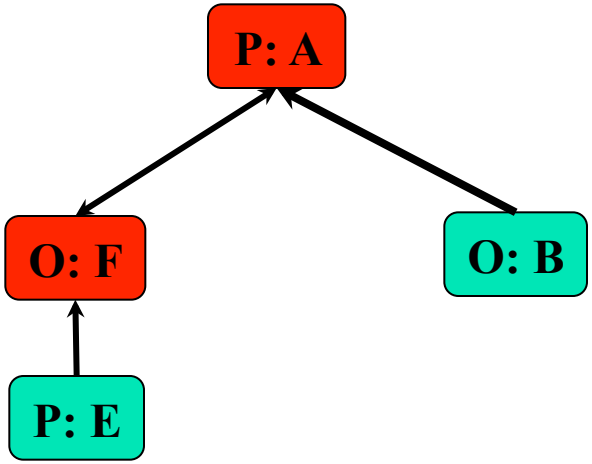

A game tree

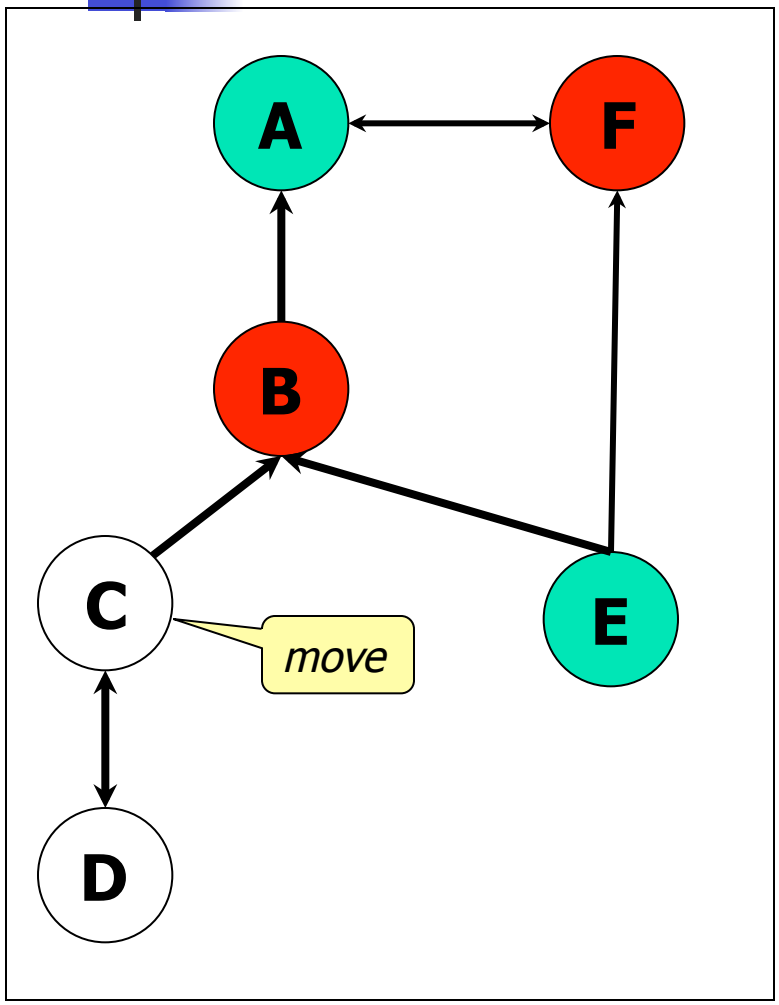

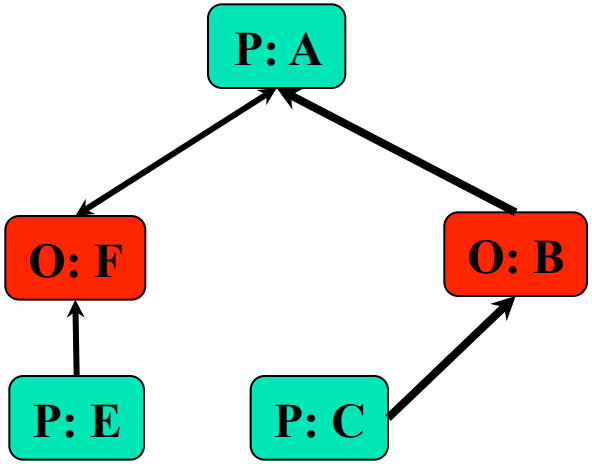

A game tree

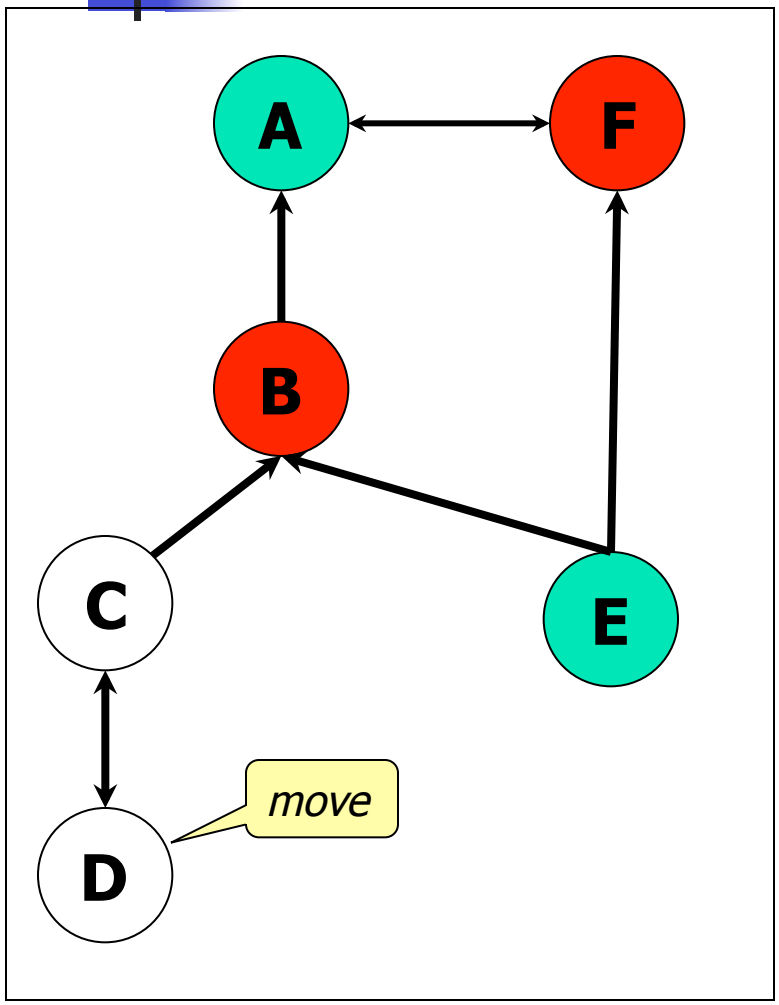

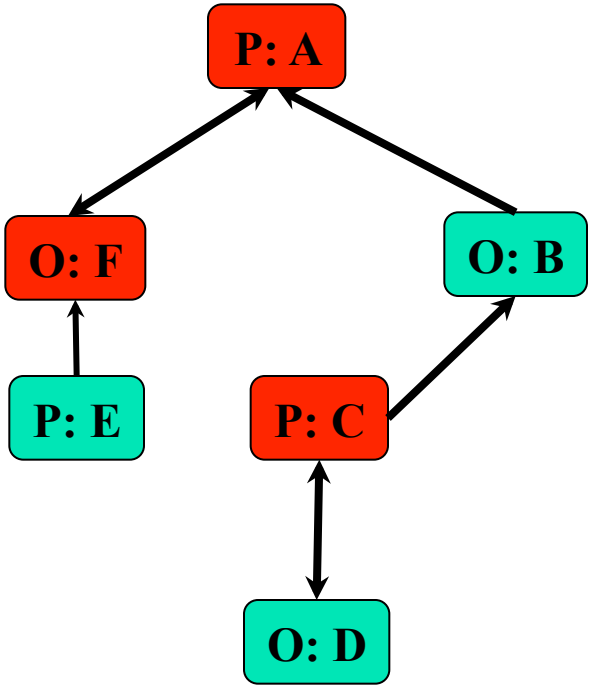

A game tree

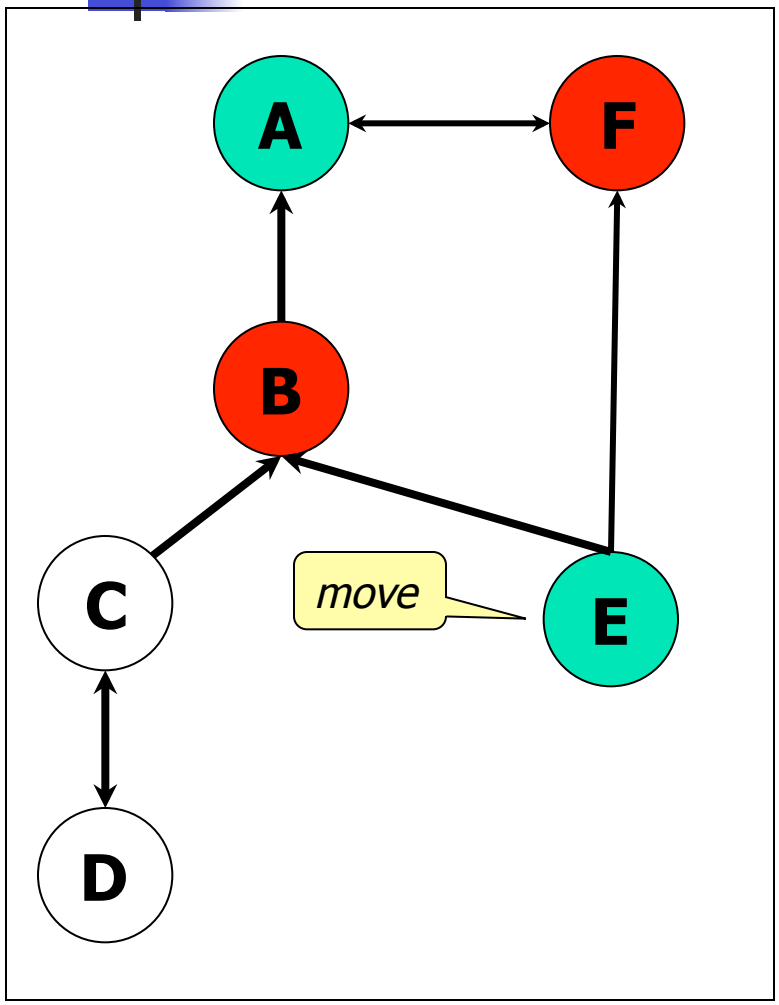

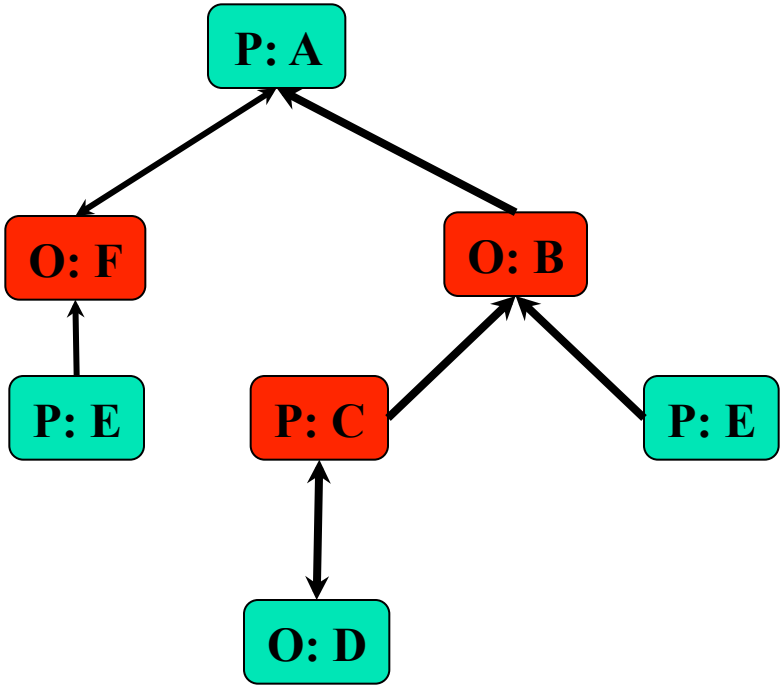

Proponent's winning strategy

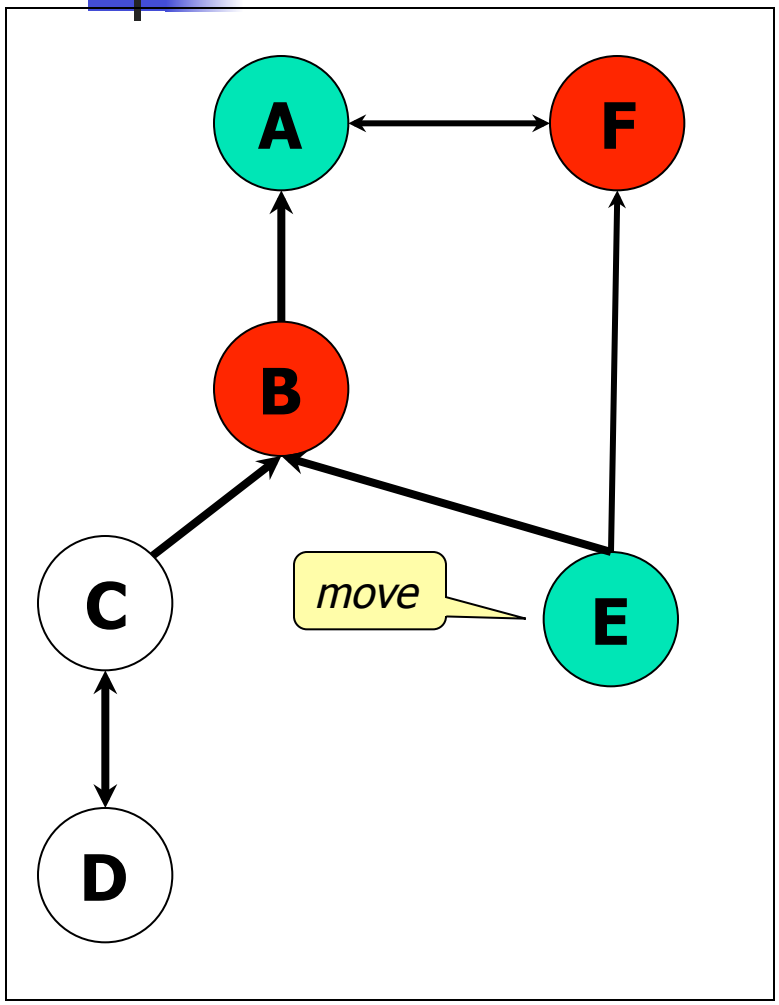

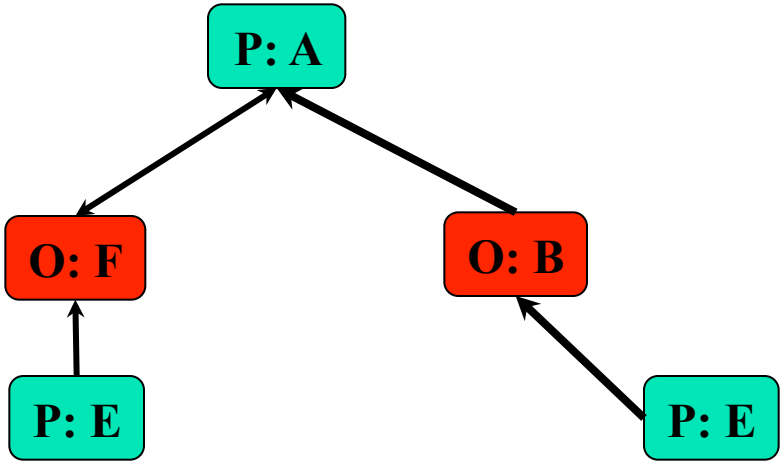

Exercise

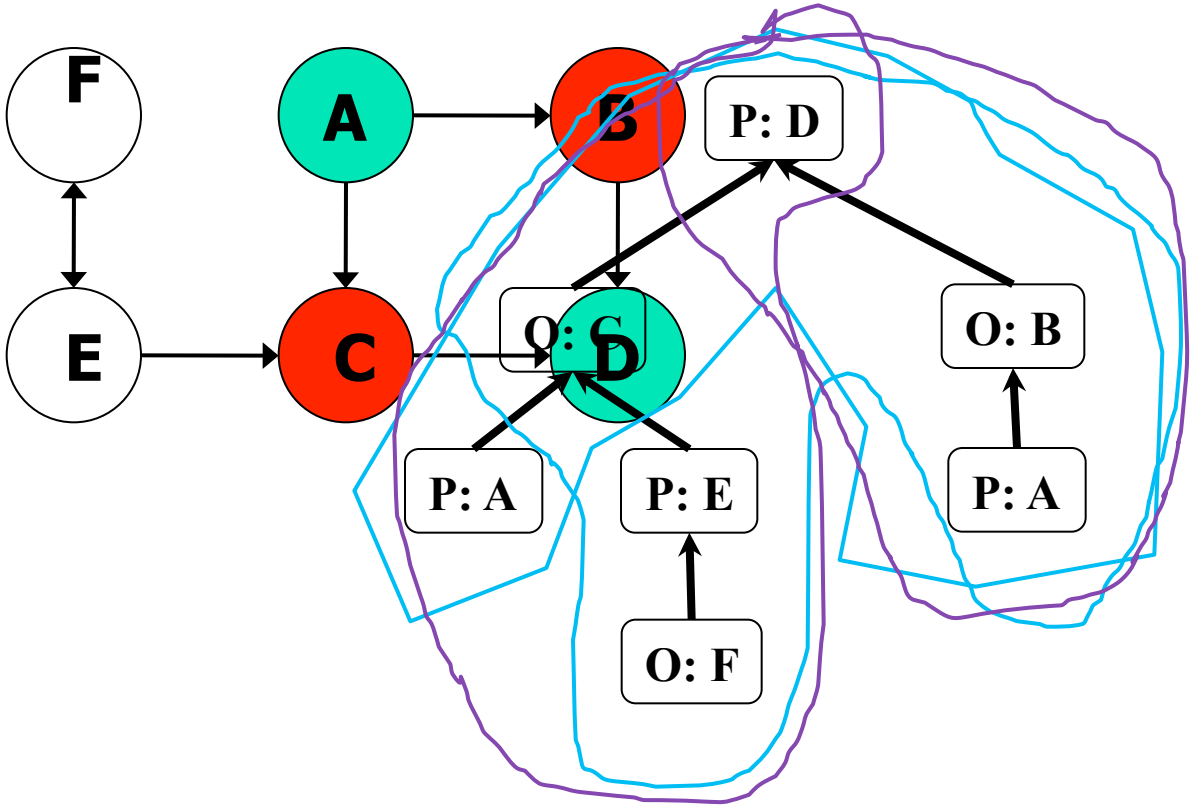

Slide made by Liz Black

Research on abstract argumentation

- New semantics
- Algorithms
  - Finding labellings (extensions)
  - Games
- Complexity
- Dynamics (adding or deleting arguments or attacks)
- Addition of new elements to AFs:
  - abstract support relations
  - preferences
- Reasons to be sceptical:
  - S. Modgil & H. Prakken, Resolutions in structured Argumentation. In *Proceedings of COMMA 2012*.
  - H. Prakken, Some reflections on two current trends in formal argumentation. In *Festschrift for Marek Sergot*, Springer 2012.
  - H. Prakken, On support relations in abstract argumentation as abstractions of inferential relations. In *Proceedings ECAI 2014*

Arguing about attack relations

- Standards for determining defeat relations are often:
  - Domain-specific
  - Defeasible and conflicting
- So determining these standards is argumentation!
- Recently Modgil (AIJ 2009) has extended Dung's abstract approach
  - Arguments can also attack attack relations

Will it rain in Calcutta?

Modgil 2009

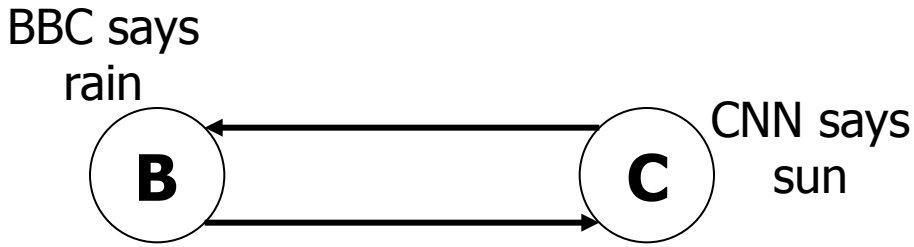

Will it rain in Calcutta?

Modgil 2009

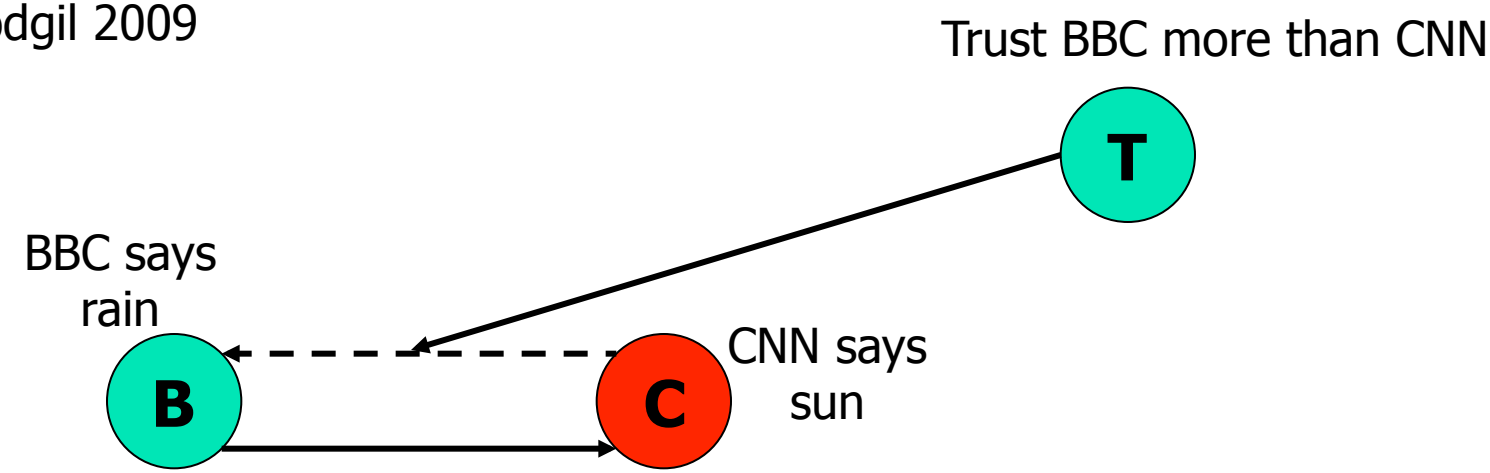

Will it rain in Calcutta?

Modgil 2009

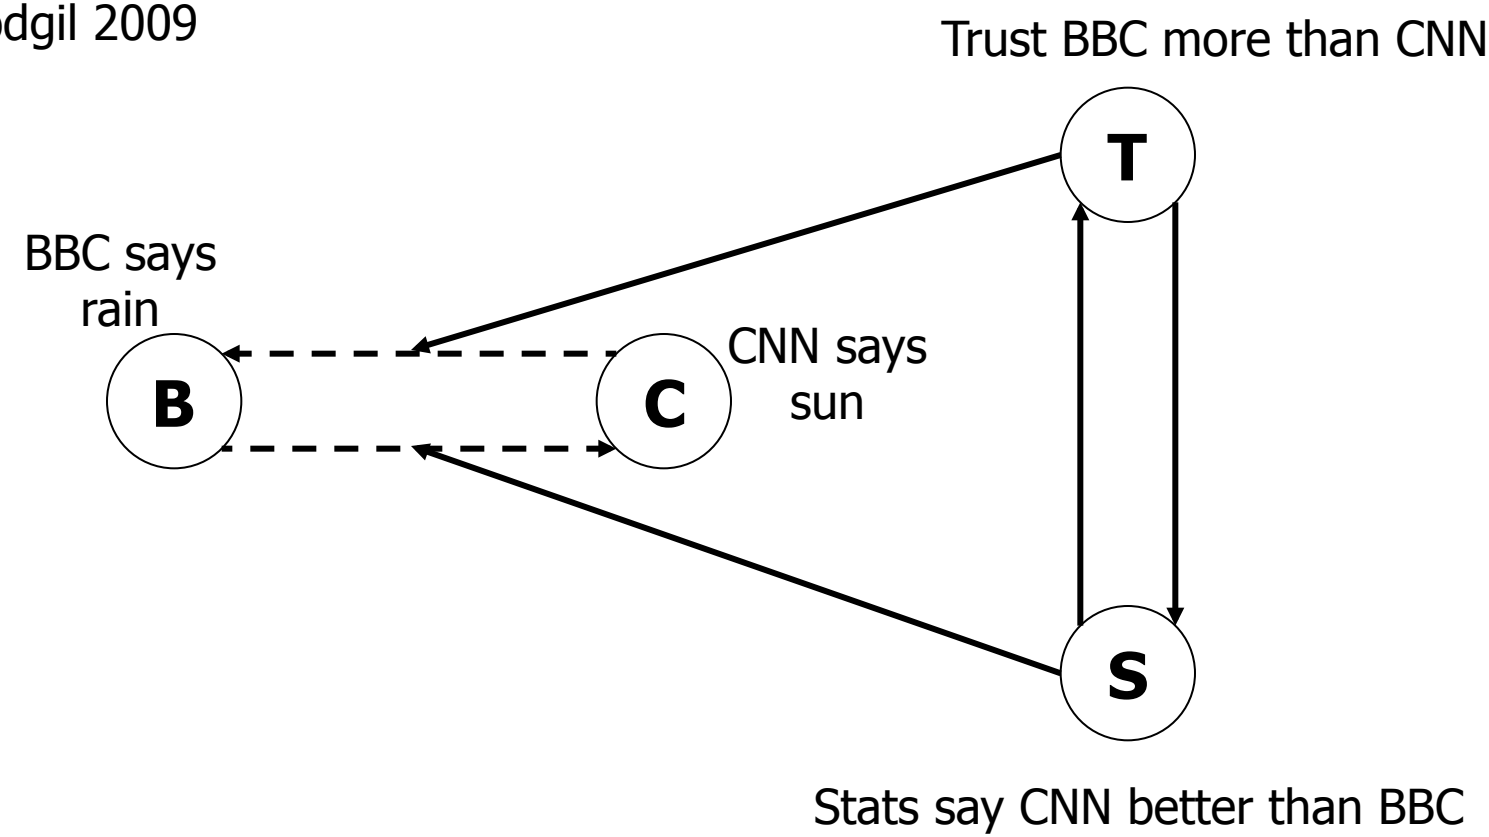

Will it rain in Calcutta?

Modgil 2009

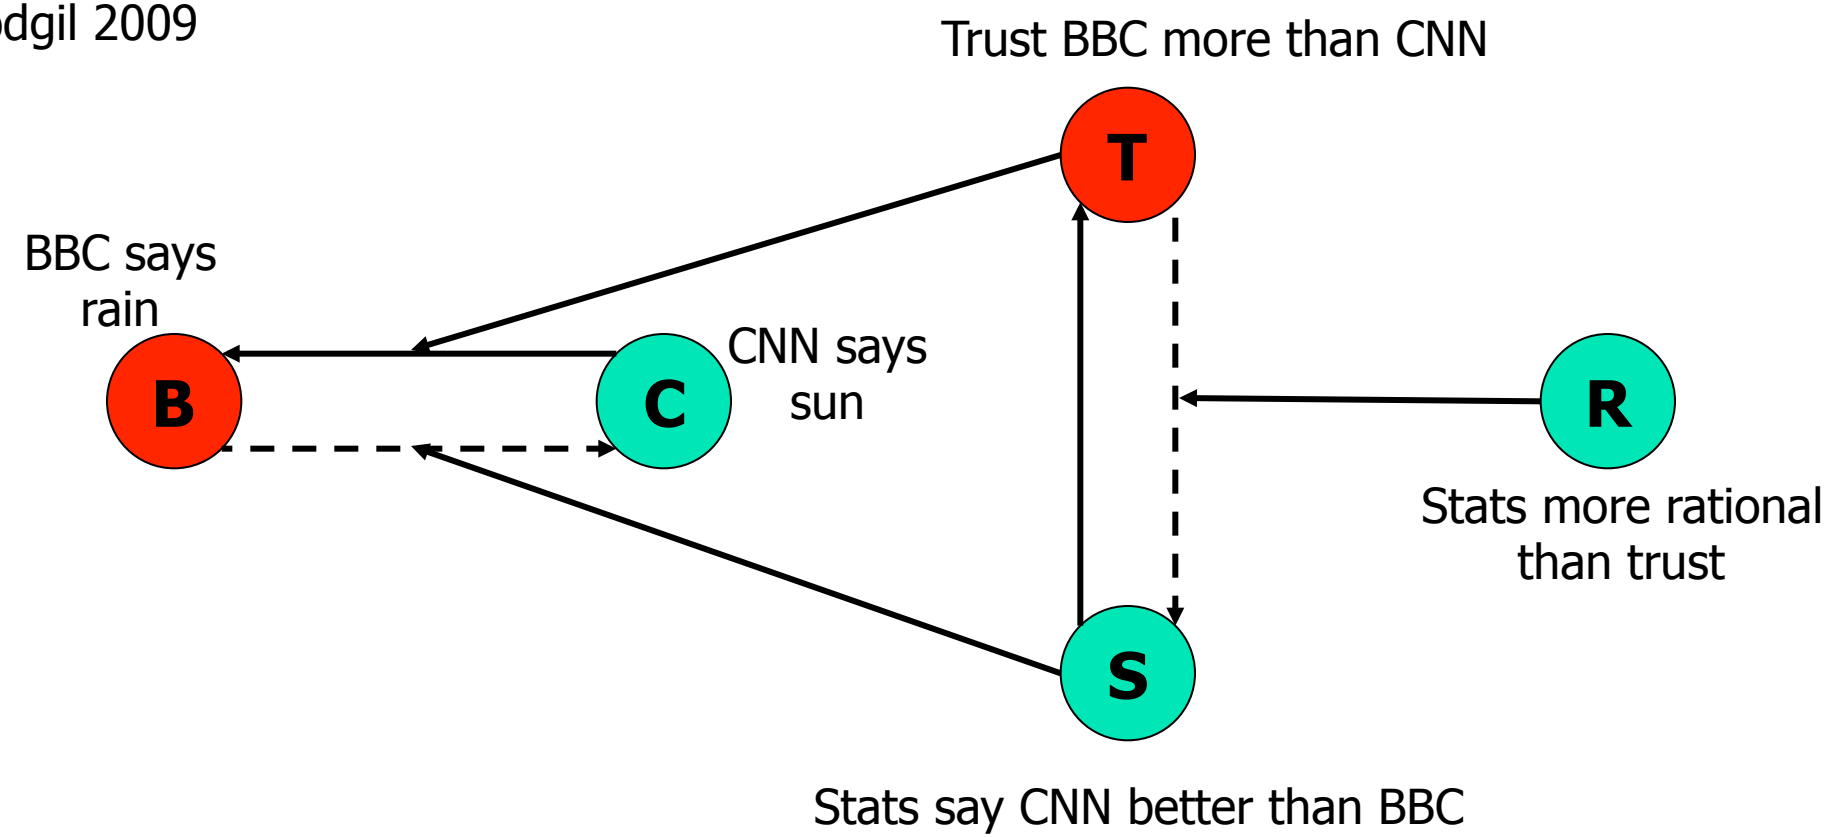

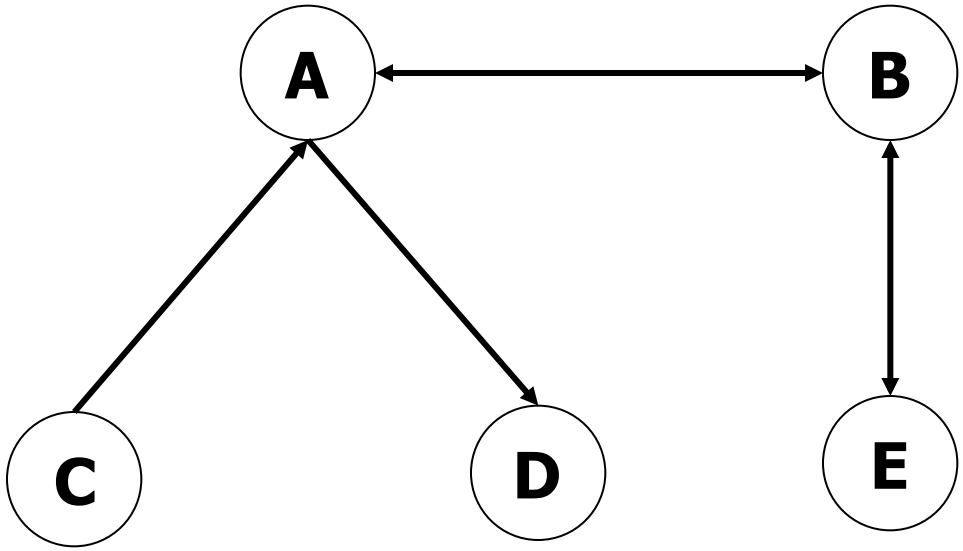
A
B
C
D
E

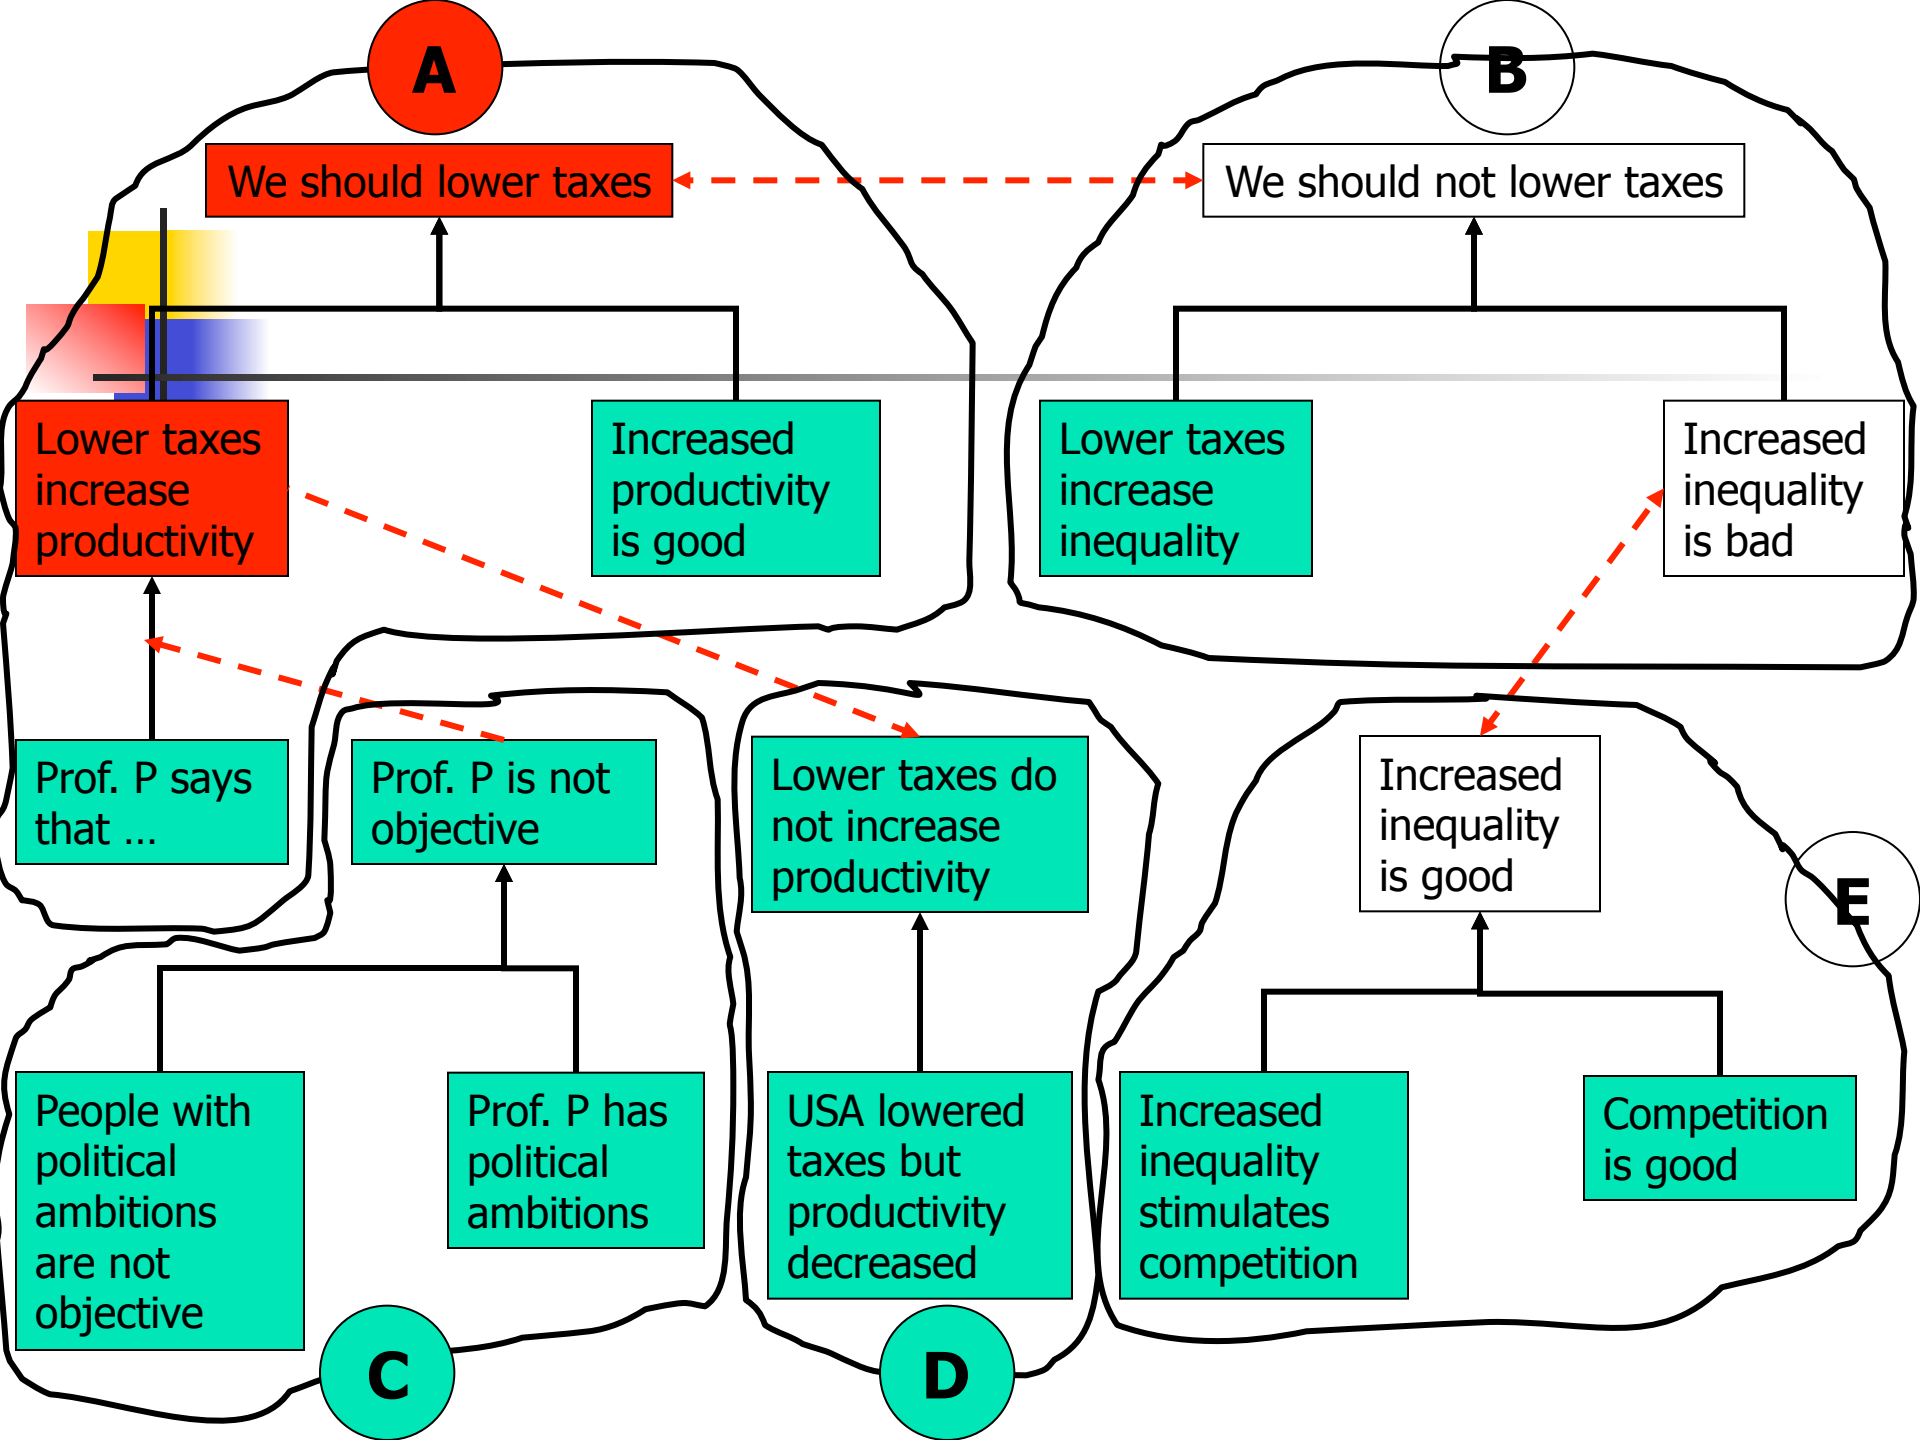
A
We should lower taxes
Lower taxes increase productivity
Increased productivity is good
B
We should not lower taxes
Lower taxes increase inequality
Increased inequality is bad
Prof. P says that …
Prof. P is not objective
People with political ambitions are not objective
Prof. P has political ambitions
C
Lower taxes do not increase productivity
USA lowered taxes but productivity decreased
D
Increased inequality is good
Increased inequality stimulates competition
Competition is good
E

The ultimate status of conclusions of arguments

- Arguments:
  - A is justified if A is *In* in all labellings
  - A is overruled if A is *Out* in all labellings
  - A is defensible otherwise
- Conclusions:
  - ϕ is justified if ϕ is the conclusion of some justified argument
  - ϕ is defensible if ϕ is not justified and ϕ is the conclusion of some defensible argument
  - ϕ is overruled if ϕ is not justified or defensible and there exists an overruled argument for ϕ
- Justification is nonmonotonic!
  - Cn over $\mathcal{L}$ is monotonic iff for all $p \in \mathcal{L}$, $S,S' \subseteq \mathcal{L}$: If $p \in Cn(S)$ and $S \subseteq S'$ then $p \in Cn(S')$

Two accounts of the fallibility of arguments

- Plausible Reasoning: all fallibility located in the premises
  - Assumption-based argumentation (Kowalski, Dung, Toni,…
  - Classical argumentation (Cayrol, Besnard, Hunter, …)
- Defeasible reasoning: all fallibility located in the inferences
  - Pollock, Loui, Vreeswijk, Prakken & Sartor, DeLP, …
- ASPIC+ combines these accounts

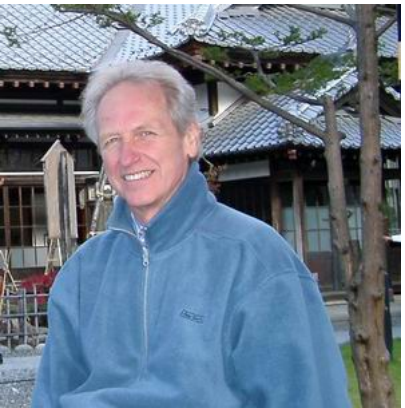

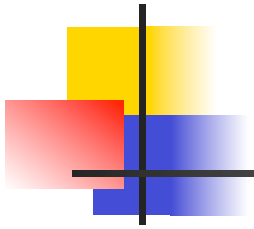

“Nonmonotonic” v. “Defeasible”

- Nonmonotonicity is a property of consequence notions
- Defeasibility is a property of inference rules
  - An inference rule is defeasible if there are situations in which its conclusion does not have to be accepted even though all its premises must be accepted.

Rationality postulates for structured argumentation

- Extensions should be closed under subarguments
- Their conclusion sets should be:
  - Consistent
  - Closed under deductive inference

M. Caminada & L. Amgoud, On the evaluation of argumentation formalisms. *Artificial Intelligence* 171 (2007): 286-310

The ‘base logic’ approach (Hunter, COMMA 2010)

- Adopt a single base logic
- Define arguments as consequence in the adopted base logic
- Then the structure of arguments is given by the base logic

Classical argumentation (Besnard, Hunter, …)

- Assume a possibly inconsistent KB in the language of classical logic
- Arguments are classical proofs from consistent (and subset-minimal) subsets of the KB
- Various notions of attack
- Possibly add preferences to determine which attacks result in defeat
  - E.g. Modgil & Prakken, AIJ-2013.
- Approach recently abstracted to Tarskian abstract logics
  - Amgoud & Besnard (2009-2013)

Classical argumentation formalised

- Given $\mathcal{L}$ a propositional logical language and |- standard-logical consequence over $\mathcal{L}$:
- An argument is a pair (S,p) such that
  - $S \subseteq \mathcal{L}$ and $p \in \mathcal{L}$
  - S |- p
  - S is consistent
  - No $S' \subset S$ is such that S' |- p
- Various notions of attack, e.g.:
- "Direct defeat": argument (S,p) attacks argument (S',p') iff p |- ¬q for some $q \in S'$
- "Direct undercut": argument (S,p) attacks argument (S',p') iff p = ¬q for some $q \in S'$
- Only these two attacks satisfy consistency, so classical argumentation is only optimal for plausible reasoning

Modelling default reasoning in classical argumentation

- Quakers are usually pacifist
- Republicans are usually not pacifist
- Nixon was a quaker and a republican

A modelling in classical logic

- Quaker ⊃ Pacifist
- Republican ⊃ → ¬Pacifist
- Facts: Quaker, Republican

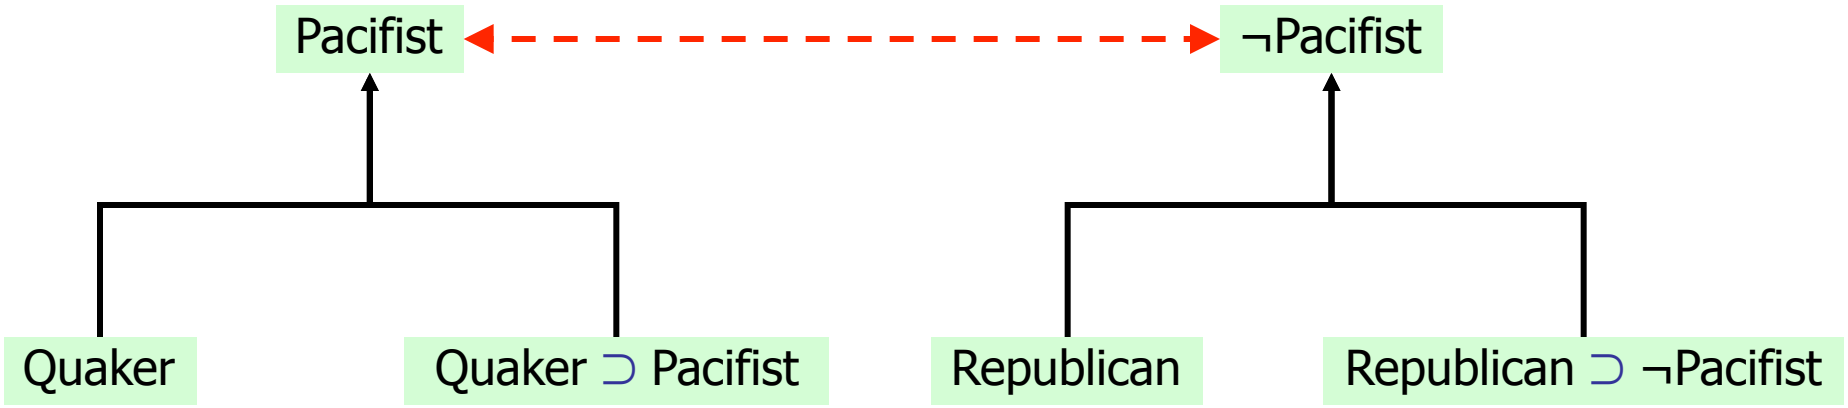

A modelling in classical logic

- Quaker ⊃ Pacifist
- Republican ⊃ → ¬Pacifist
- Facts: Quaker, Republican

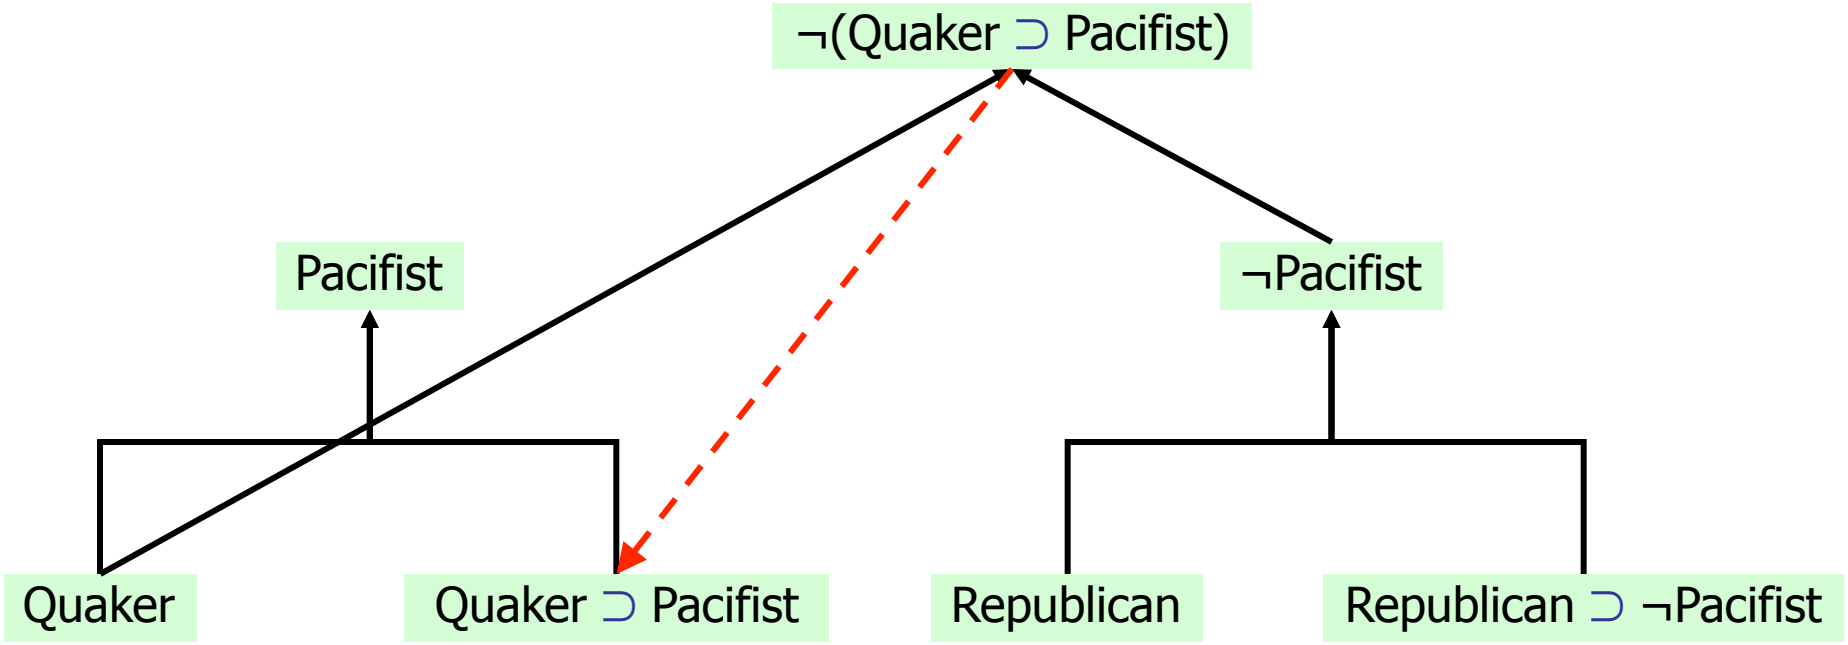

A modelling in classical logic

- Quaker & ¬Ab1 ⊃ Pacifist
- Republican & ¬Ab2 ⊃ → ¬Pacifist
- Facts: Quaker, Republican
- Assumptions: ¬Ab1, ¬Ab2 (attackable)

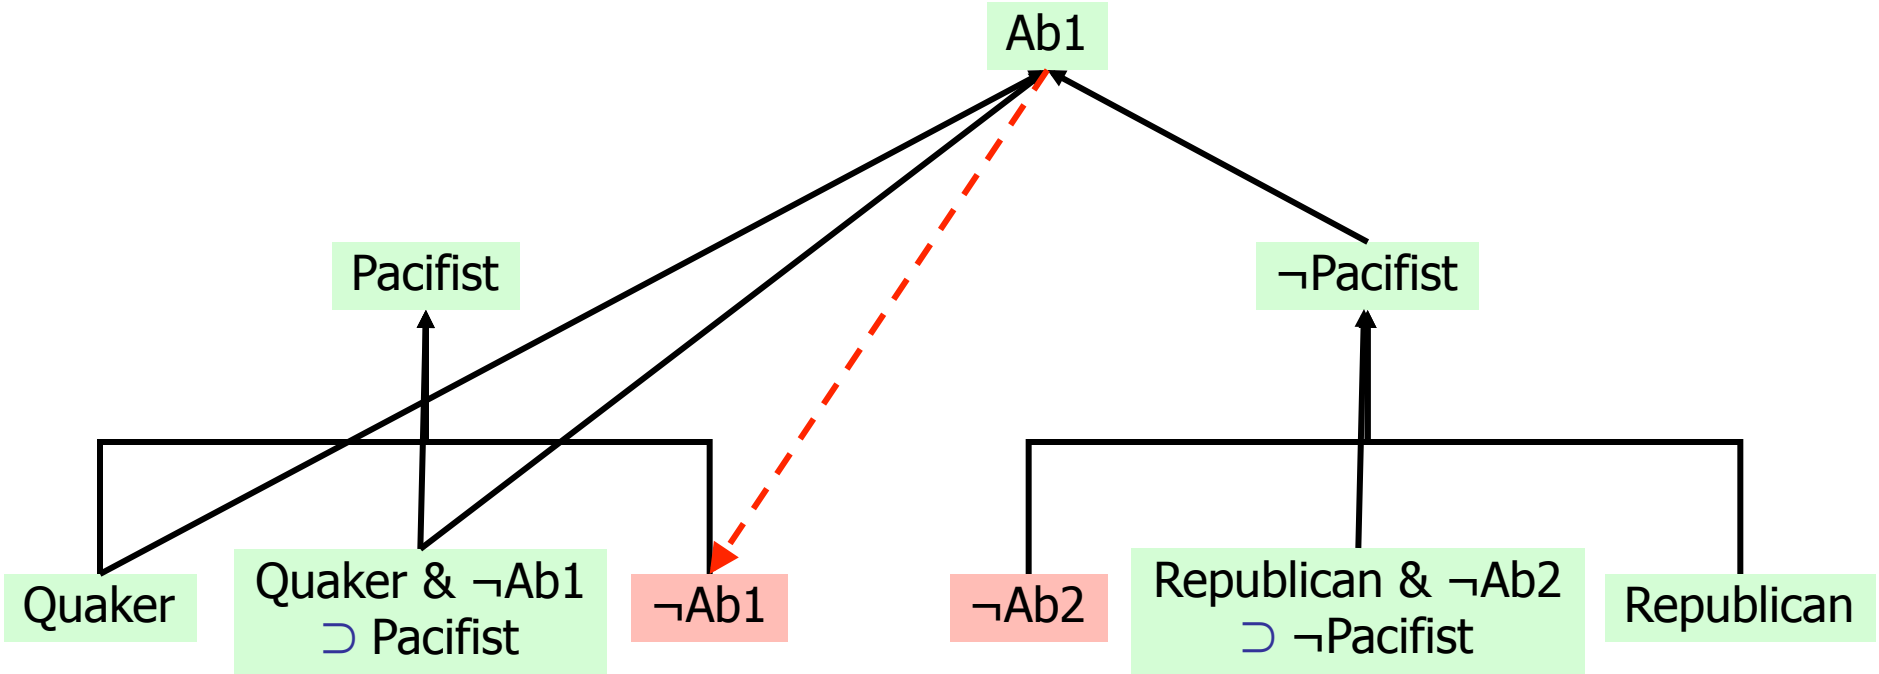

A modelling in classical logic

- Quaker & ¬Ab1 ⊃ Pacifist
- Republican & ¬Ab2 ⊃ → ¬Pacifist
- Facts: Quaker, Republican
- Assumptions: ¬Ab1, ¬Ab2 (attackable)

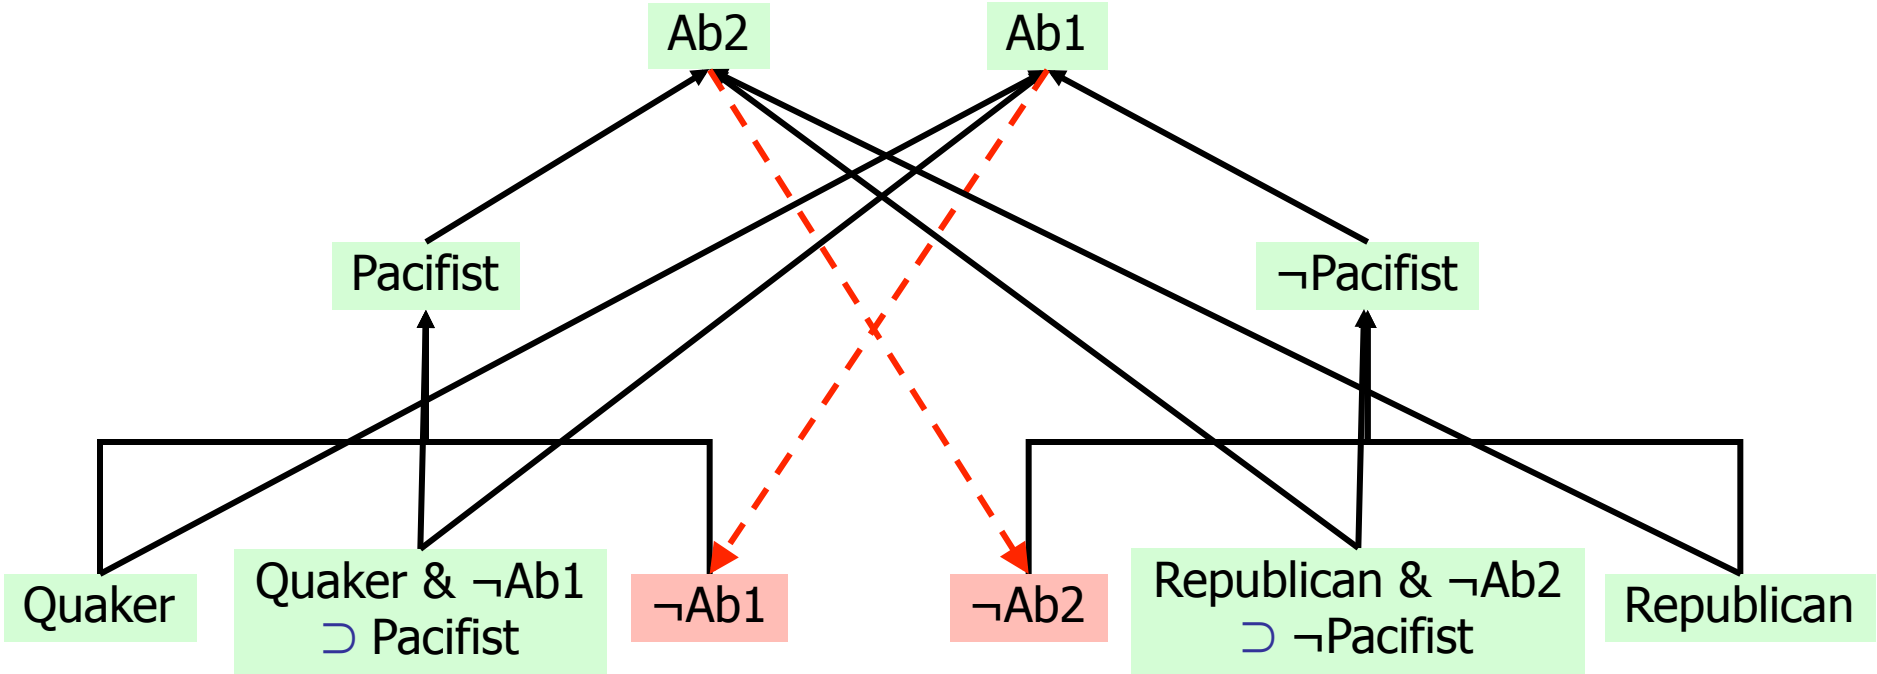

Extensions v. maximal consistent subsets

- With classical (and Tarskian) argumentation preferred and stable extensions and maximal conflict-free sets coincide with maximal consistent subsets of the knowledge base
  - Cayrol (1995)
  - Amgoud & Besnard (2013)
- If 'real' argumentation is more than identifying mcs, then deductive argumentation when combined with Dung misses something.
  - Modgil (& Prakken) 2013: with preferences they coincide with Brewka's preferred subtheories
  - But is real argumentation identifying preferred subtheories?

The ASPIC+ framework

- Arguments: Trees where
  - Nodes are statements in some logical language $\mathcal{L}$
  - Links are applications of inference rules
    - Strict rules $\rightarrow$
    - Defeasible rules $\Rightarrow$
- Constructed from a knowledge base $\mathcal{K} \subseteq \mathcal{L}$
  - Axiom (necessary) premises + ordinary (contingent) premises
- Attack:
  - On ordinary premises
  - On defeasible inferences (undercutting)
  - On conclusions of defeasible inferences (rebutting)
- Defeat: attack + argument ordering
- Argument evaluation with Dung (1995)

$\mathcal{R}_s$: $\mathbf{R_d}$:

r,s → t p,q ⇒ r

$\mathcal{K}_n$ = {p,q} K_p = {s}

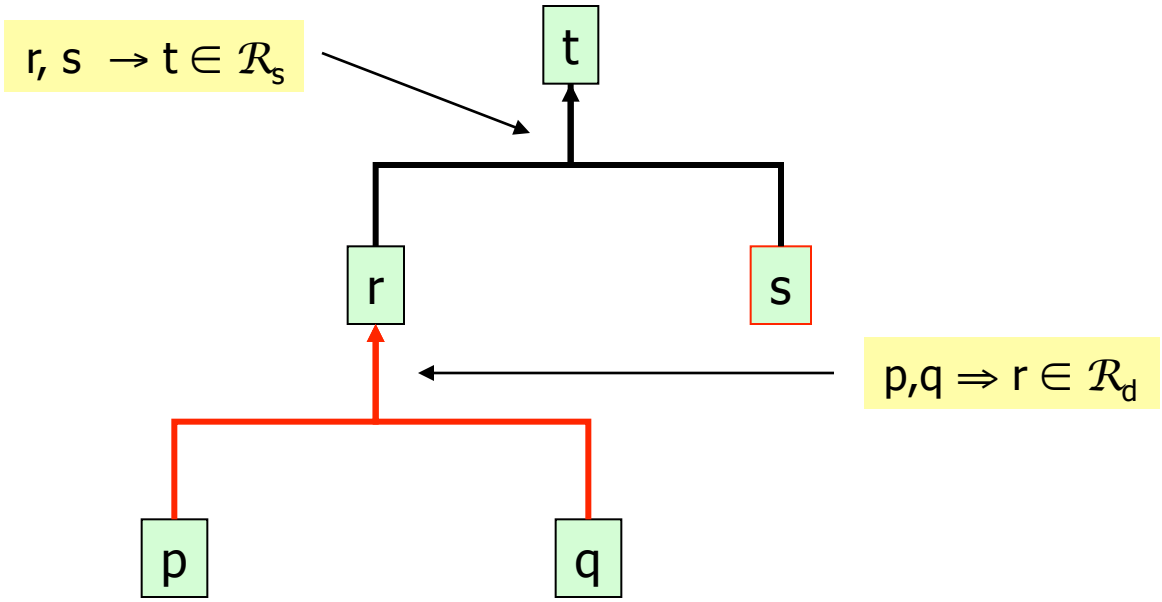

$\mathcal{R}_s$: $\mathbf{R_d}$:

r,s → t p,q ⇒ r

$\mathcal{K}_n$ = {p,q} K_p = {s}

Attack:
Undermining: on ordinary premises
Rebutting: on defeasible inferences
Undercutting: on conclusions of defeasible inferences

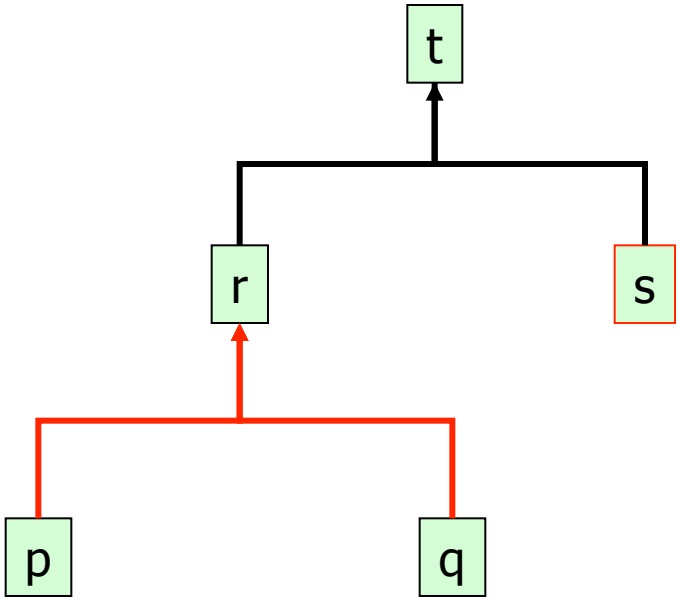

$n(\phi_1, \ldots, \phi_n \Rightarrow \phi) \in \mathcal{L}$

Attack + preferences = defeat

Consistency in ASPIC+ (with symmetric negation)

- For any $S \subseteq \mathcal{L}$
  - S is (directly) consistent iff S does not contain two formulas ϕ and $-\phi$.
  - …

Rationality postulates for ASPIC+

- Subargument closure always satisfied
- Consistency and strict closure:
  - without preferences satisfied if
    - $\mathcal{R}_s$ closed under transposition or closed under contraposition; and
    - $\mathcal{K}_n$ is indirectly consistent
  - with preferences satisfied if in addition the argument ordering is 'reasonable'
    - Versions of the weakest- and last link ordering are reasonable
- So ASPIC+ is good for both plausible and defeasible reasoning

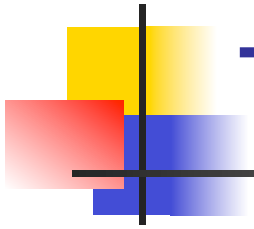

Two uses of defeasible rules

- For domain-specific information
  - Defeasible generalisations, norms, …
- For general patterns of presumptive reasoning
  - Pollock's defeasible reasons:
    - perception, memory, induction, statistical syllogism, temporal persistence
  - Argument schemes

Domain-specific vs. inference general inference rules

- d1: Bird $\Rightarrow$ Flies
- s1: Penguin $\rightarrow$ Bird
- Penguin $\in \mathcal{K}$

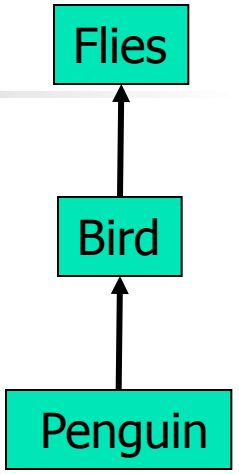

- $\mathcal{R}_d = \{\phi, \phi \dashrightarrow \psi \Rightarrow \psi\}$
- $\mathcal{R}_s$ = all valid inference rules of prop. l.
- Bird $\dashrightarrow$ Flies $\in \mathcal{K}$
- Penguin $\supset$ Bird $\in \mathcal{K}$
- Penguin $\in \mathcal{K}$

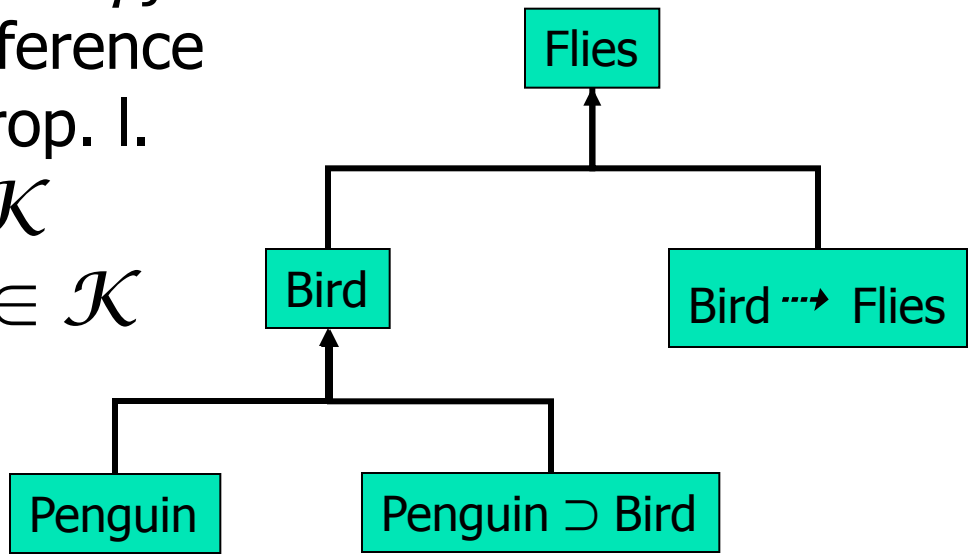

Preferred extensions do not always coincide with mcs

- r1: Quaker $\Rightarrow$ Pacifist
- r2: Republican $\Rightarrow$ ¬Pacifist
- S $\rightarrow$ p $\in \mathcal{R}_s$ iff S |- p in Prop. L and S is finite
- $\mathcal{K}$: Quaker, Republican

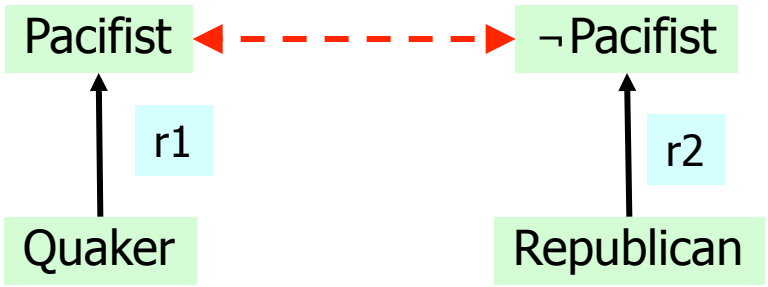

Preferred/stable extensions do not always coincide with mcs

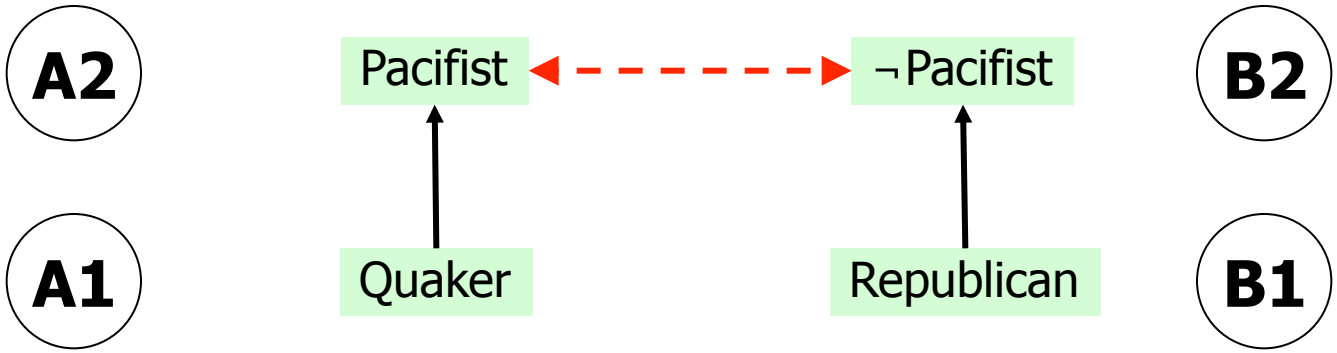

E1 = {A1,A2,B1,…}

E2 = {A1,B1,B2,…}

Preferred/stable extensions do not always coincide with mcs

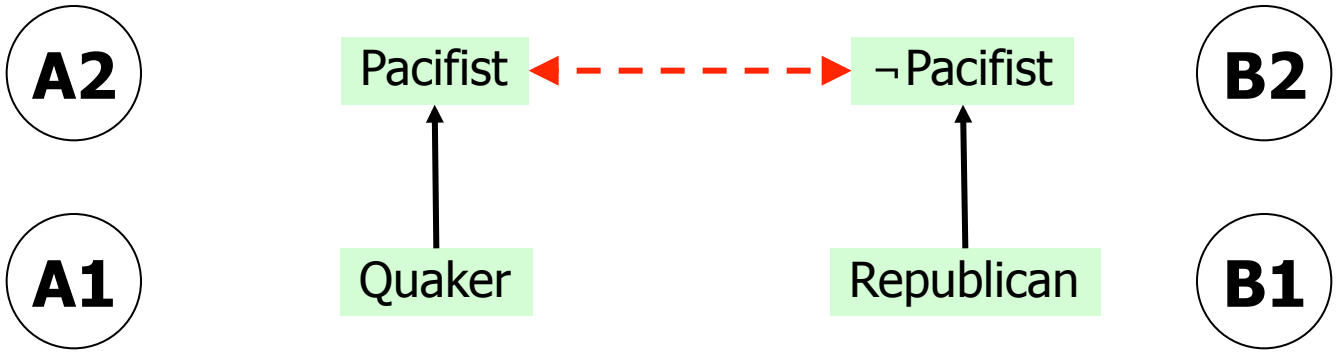

Conc(E1) = Th({Quaker, Republican, Pacifist})

Conc(E2) = Th({Quaker, Republican, ¬Pacifist})

mcs($\mathcal{K}$) = {{$\mathcal{K}$}} = {{Quaker, Republican}}

Preferred extensions do not always coincide with mcs

- $\mathcal{R}_d = \{\phi, \phi \dashrightarrow \psi \Rightarrow \psi\}$
- S $\rightarrow$ p $\in \mathcal{R}_s$ iff S |- p in Prop. L and S is finite
- $\mathcal{K}$: Quaker, Republican,
  Quaker $\dashrightarrow$ Pacifist, Republican $\dashrightarrow$ ¬Pacifist

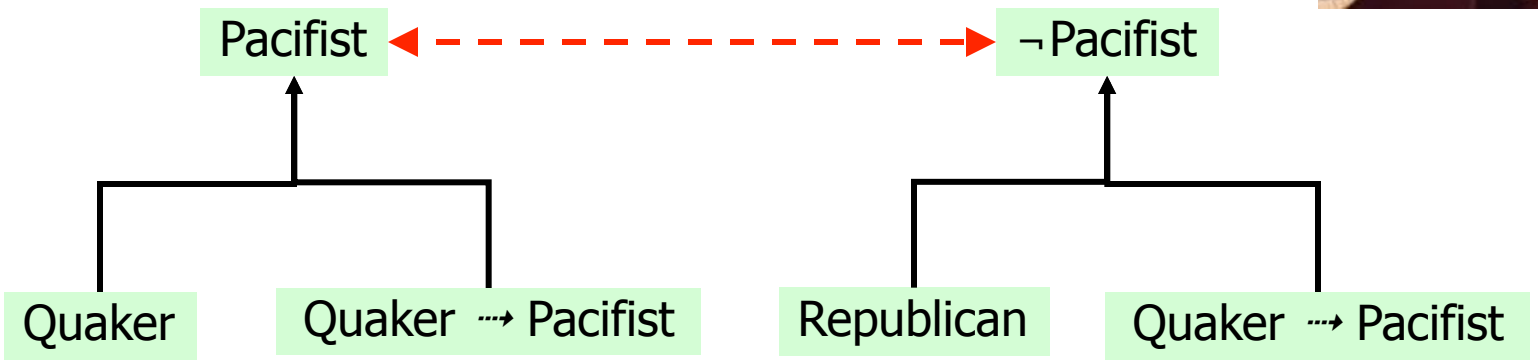

Can defeasible reasoning be reduced to plausible reasoning?

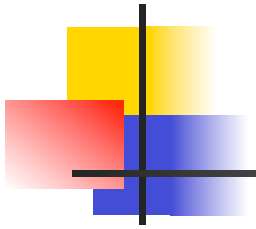

- To classical argumentation?
  - Problems with contrapositive inferences
- To assumption-based argumentation?
  - Problems with preferences
- In both cases:
  - less complex metatheory
  - but more complex representations

Default contraposition in classical argumentation

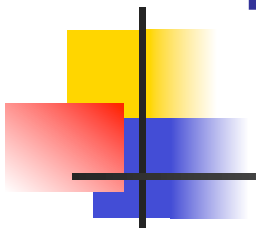

- Men are usually not rapists
  - .
- John is a rapist
- Assume when possible that things are normal

- What can we conclude about John's sex?

Default contraposition in classical argumentation

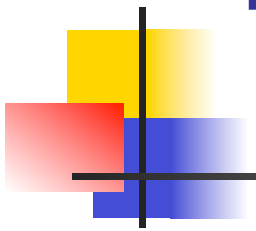

- Men are usually not rapists
  - $M \,\&\, \neg Ab \supset \neg R$
- John is a rapist (R)
- Assume when possible that things are normal
  - $\neg Ab$

Default contraposition in classical argumentation

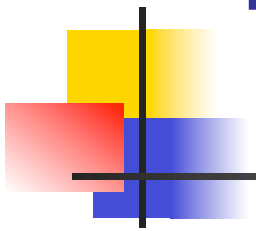

- Men are usually not rapists
  - $M \,\&\, \neg Ab \supset \neg R$
- John is a rapist (R)
- Assume when possible that things are normal
  - $\neg Ab$
- The first default implies that rapists are usually not men
  - $R \,\&\, \neg Ab \supset \neg M$
- So John is not a man

Default contraposition in classical argumentation

- Heterosexual adults are usually not married =>
  - Non-married adults are usually not heterosexual
- This type of sensor usually does not give false alarms =>
  - False alarms are usually not given by this type of sensor

Statisticians call these inferences "base rate fallacies"

Assumption-based argumentation (Dung, Mancarella & Toni 2007)

- A deductive system is a pair $(\mathcal{L}, \mathcal{R})$ where
  - $\mathcal{L}$ is a logical language
  - $\mathcal{R}$ is a set of rules $(\phi_1, \ldots, \phi_n \rightarrow \phi)$ over $\mathcal{L}$
- An assumption-based argumentation framework is a tuple $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \bar{})$ where
  - $(\mathcal{L}, \mathcal{R})$ is a deductive system
  - $\mathcal{A} \subseteq \mathcal{L}$, $\mathcal{A} \neq \emptyset$, a set of assumptions
  - No rule has an assumption as conclusion
  - ~ is a total mapping from Pow(L) into L. ~a is the contrary of a.
- An argument S |- p is a deduction of p from a set $S \subseteq \mathcal{A}$.
- Argument S |- p attacks argument S' |-p' iff p = ~q for some $q \in S'$

Reduction of ASPIC+ defeasible rules to ABA rules (Dung & Thang, JAIR 2014)

- Assumptions:
  - $\mathcal{L}$ consists of literals
  - No preferences
  - No rebuttals of undercutters

1-1 correspondence between complete extensions of ASPIC+ and ABA

$p_1, \ldots, p_n \Rightarrow q$

becomes

$d_i, p_1, \ldots, p_n, \text{not}\neg q \rightarrow q$

where:
$d_i = n(p_1, \ldots, p_n \Rightarrow q)$
d_i,not$\neg q$ are assumptions
$\phi = \sim\text{not}\phi$, $\phi = \sim\neg\phi$, $\neg\phi = \sim\phi$

From defeasible to strict rules: example

- r1: Quaker ⇒ Pacifist
- r2: Republican ⇒ ¬Pacifist

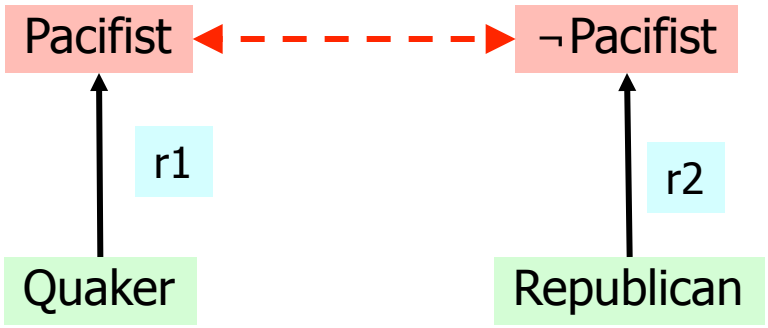

From defeasible to strict rules: example

- s1: Appl(s1), Quaker, not¬Pacifist → Pacifist
- s2: Appl(s2), Republican, notPacifist → ¬Pacifist

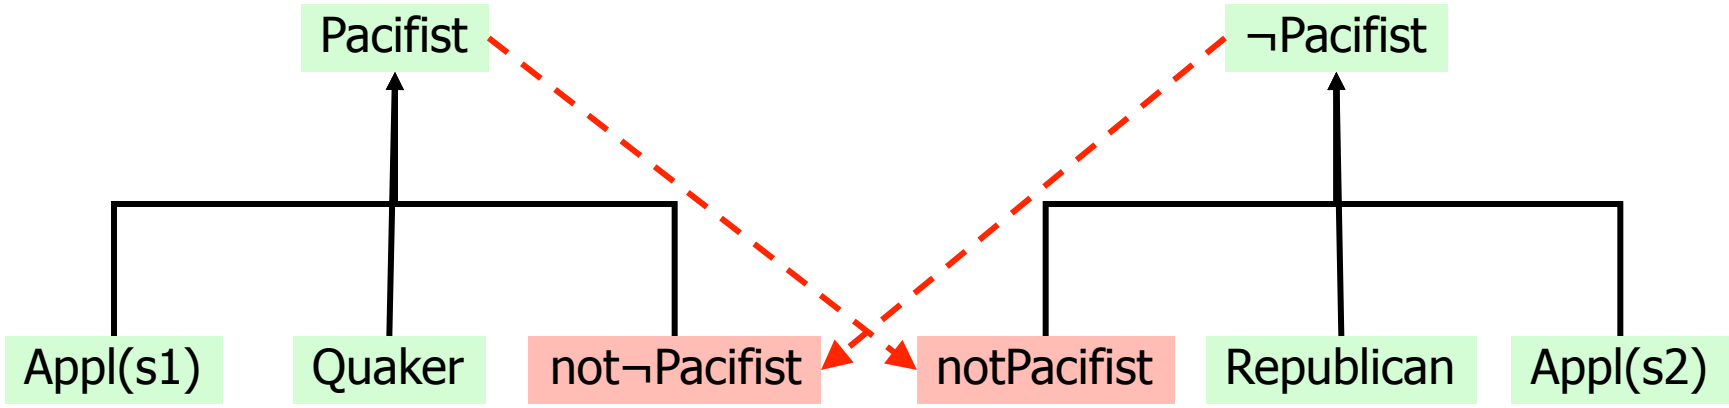

Can ASPIC+ preferences be reduced to ABA assumptions?

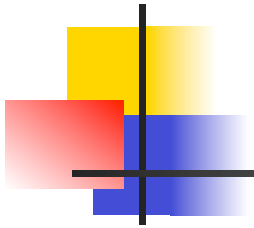

d1: Bird ⇒ Flies
d2: Penguin ⇒ ¬Flies
d1 < d2

Becomes

d1: Bird, notPenguin ⇒ Flies
d2: Penguin ⇒ ¬Flies

Only works in special cases, e.g. not with weakest-link ordering

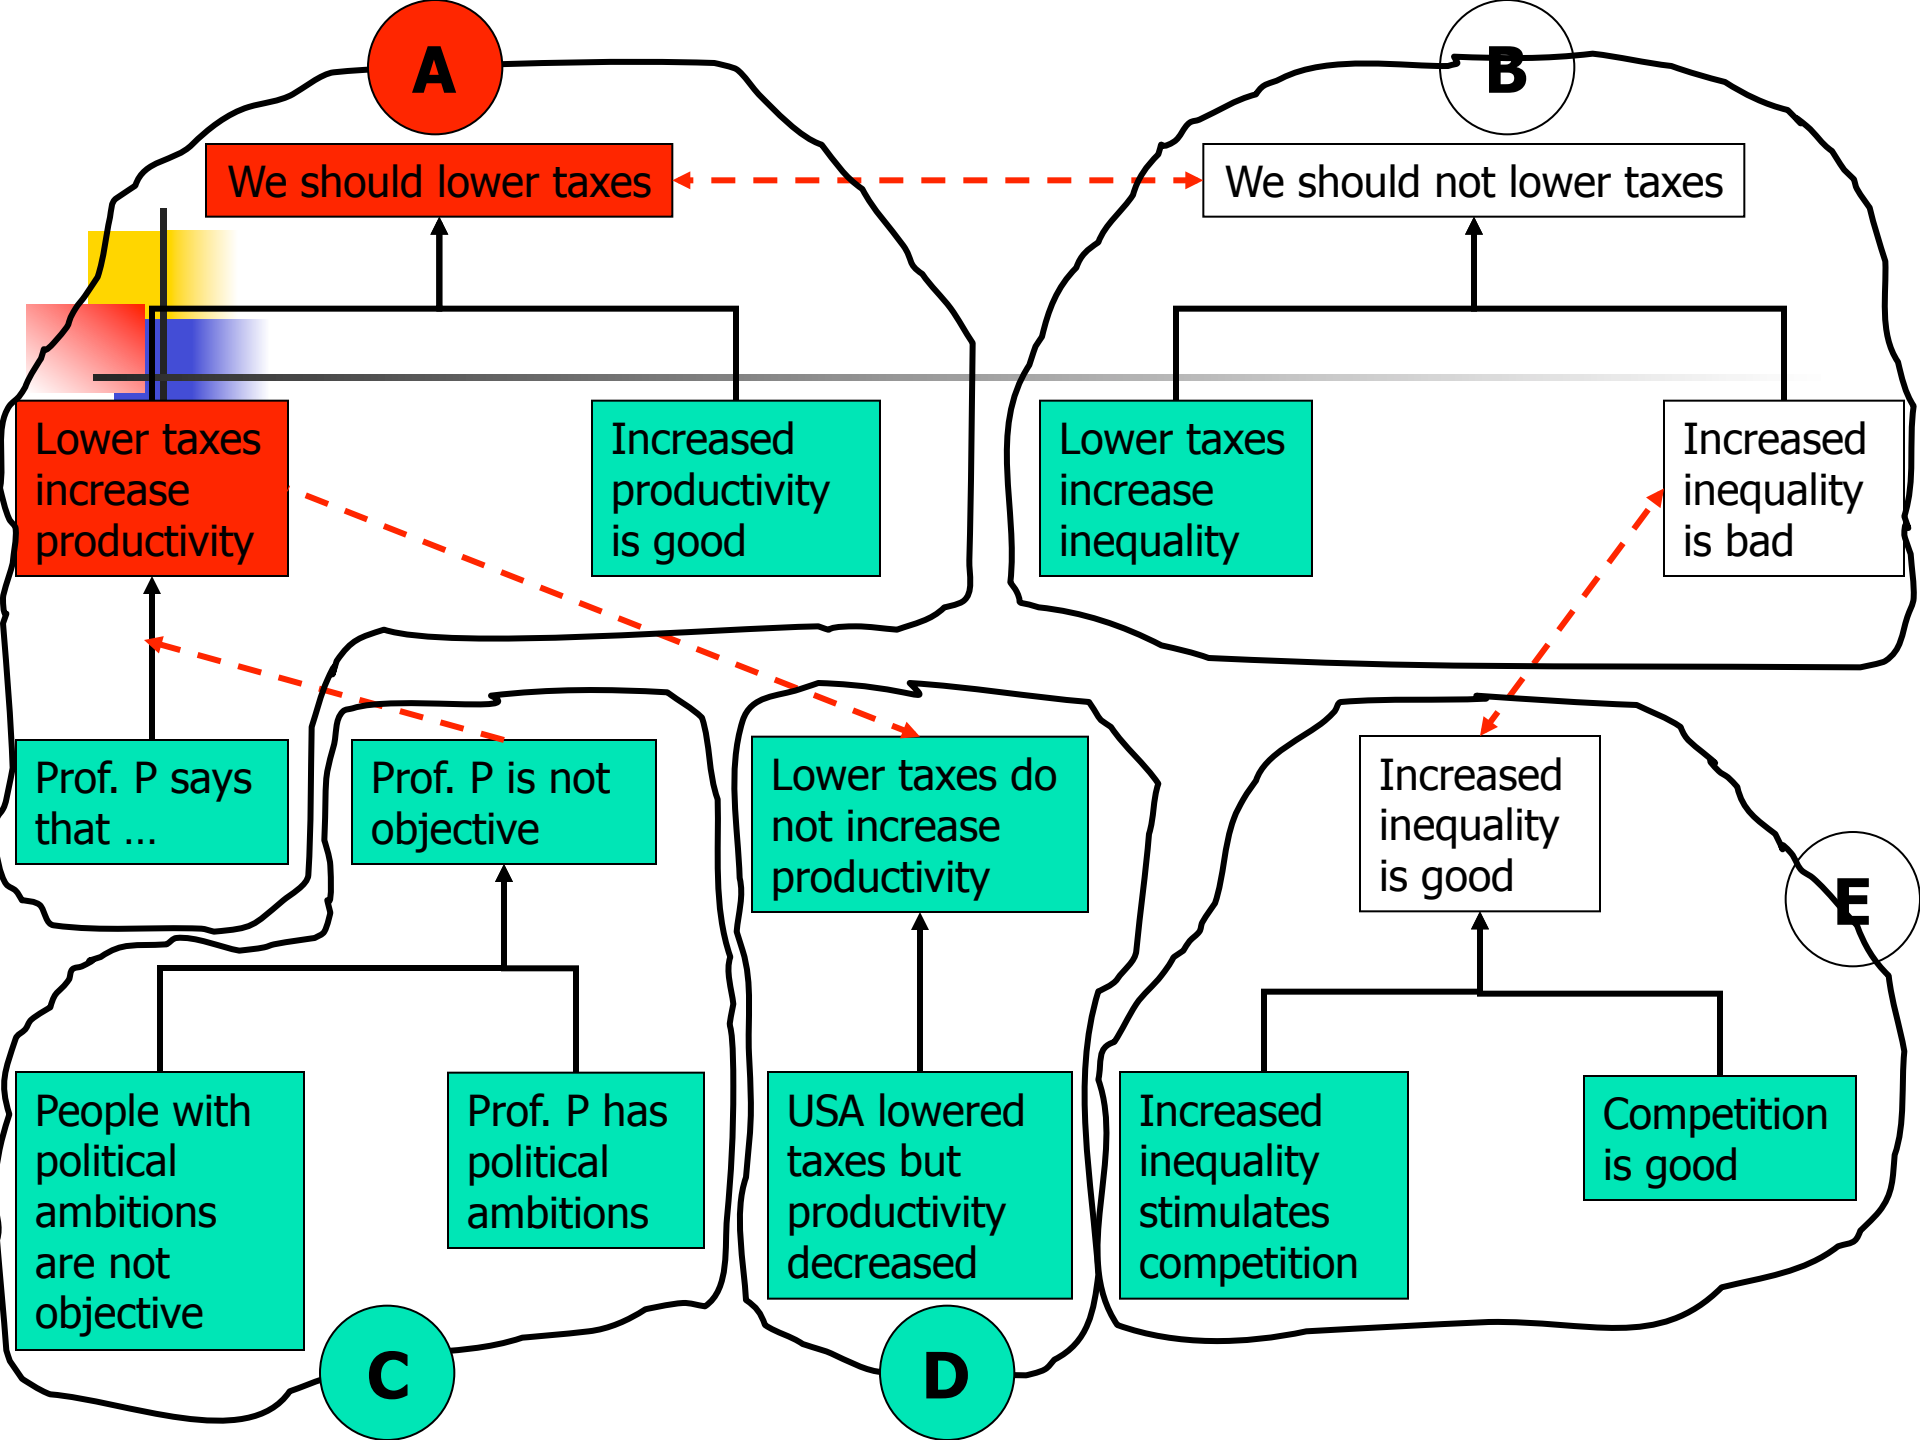
A
We should lower taxes
Lower taxes increase productivity
Increased productivity is good
B
We should not lower taxes
Lower taxes increase inequality
Increased inequality is bad
Prof. P says that …
Prof. P is not objective
People with political ambitions are not objective
Prof. P has political ambitions
C
Lower taxes do not increase productivity
USA lowered taxes but productivity decreased
D
Increased inequality is good
Increased inequality stimulates competition
Competition is good
E

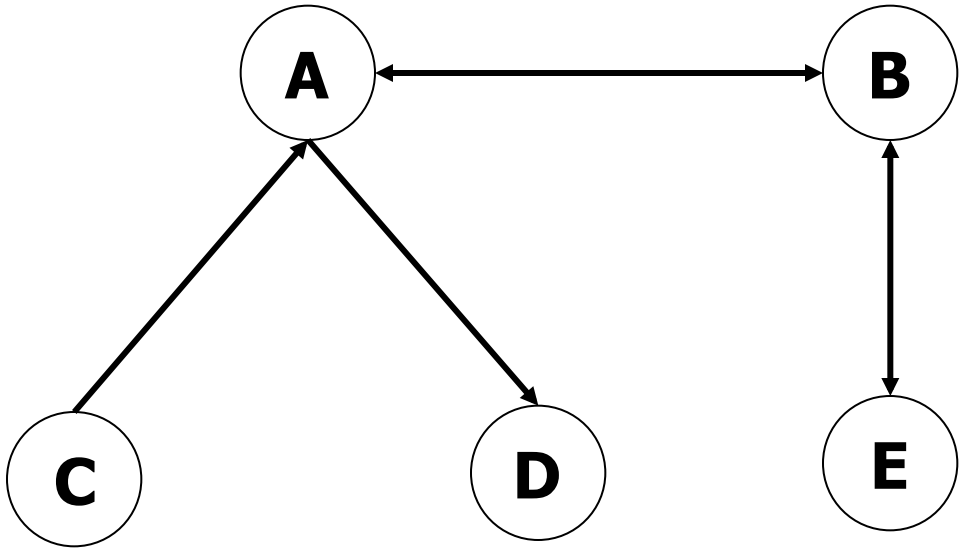
A
B
C
D
E

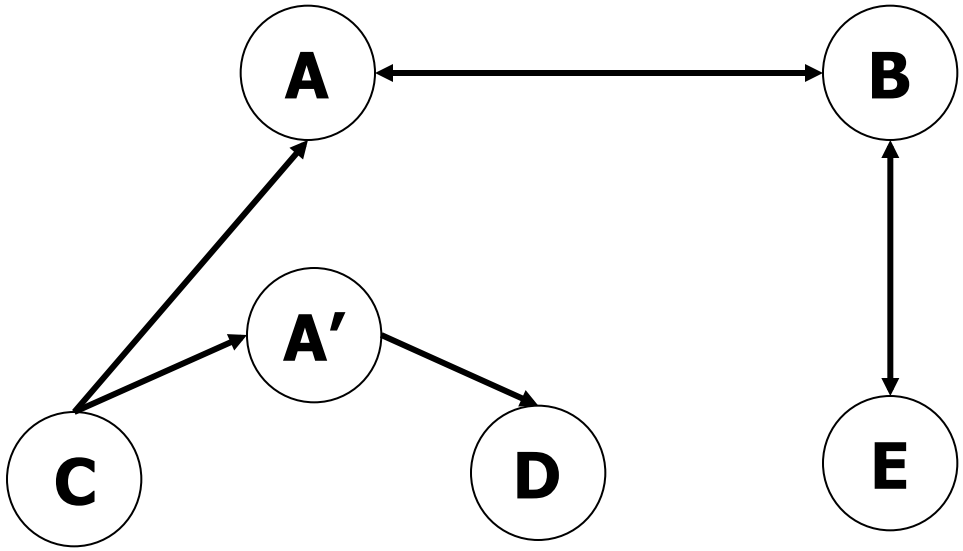
A
B
A′
C
D
E

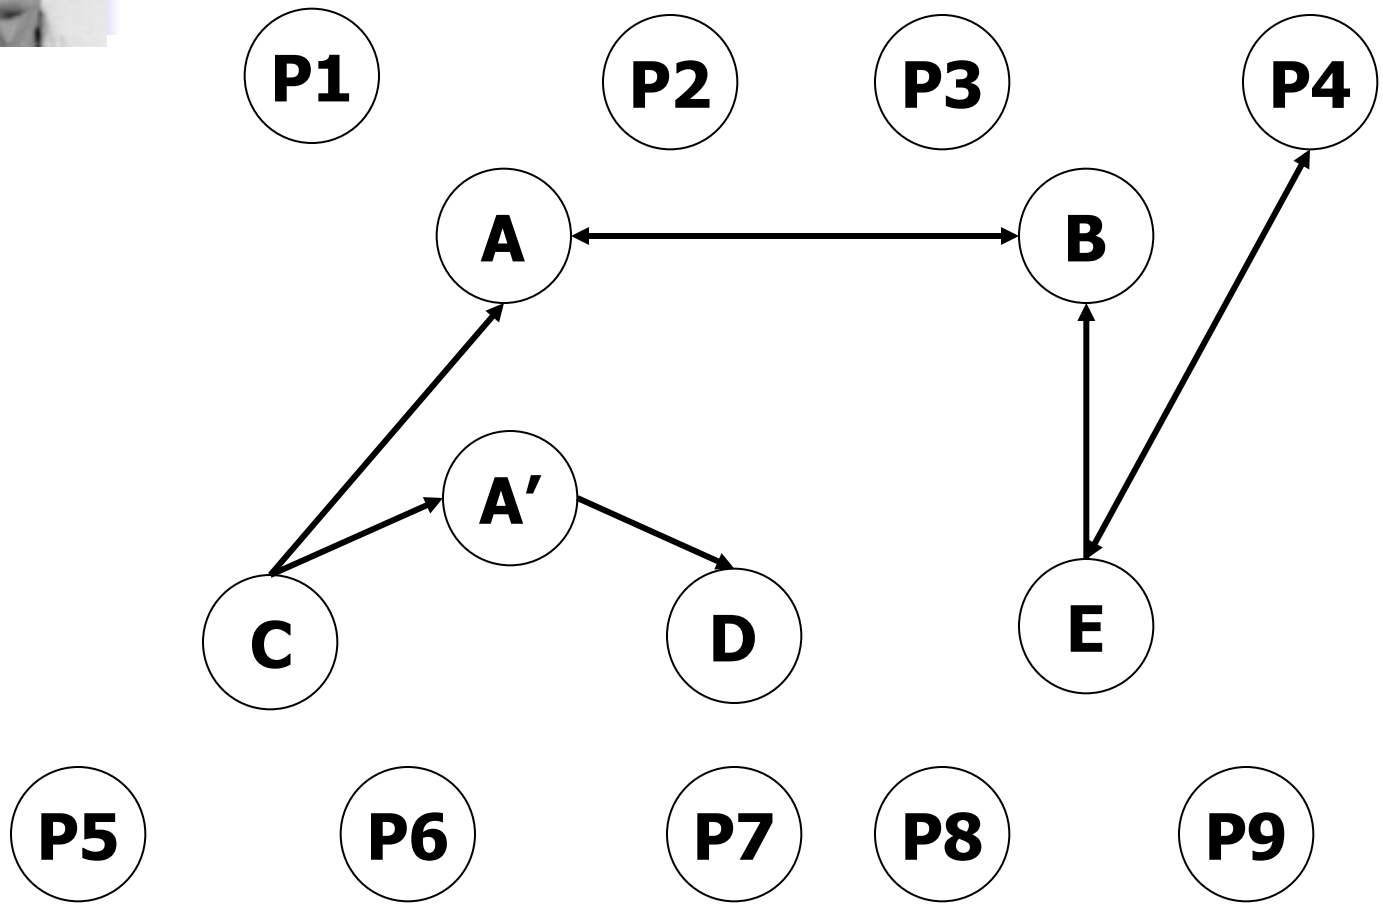
P1
P2
P3
P4
A
B
A’
C
D
E
P5
P6
P7
P8
P9

Preferences in abstract argumentation

- PAFs: extend (*args,attack*) to (*args,attack*, $\leq_a$)
  - $\leq_a$ is an ordering on args
  - A *defeats* B iff A attacks B and not A < B
  - Apply Dung's theory to (*args,defeat*)
- Implicitly assumes that:
  - All attacks are preference-dependent
  - All attacks are independent from each other
- Assumptions not satisfied in general =>
  - Properties not inherited by all instantiations
  - possibly violation of rationality postulates

R1: If you snore, you misbehave
R2: If you snore when nobody else is around, you don't misbehave
R3: If you misbehave in the library, the librarian may remove you

R1 < R2 < R3

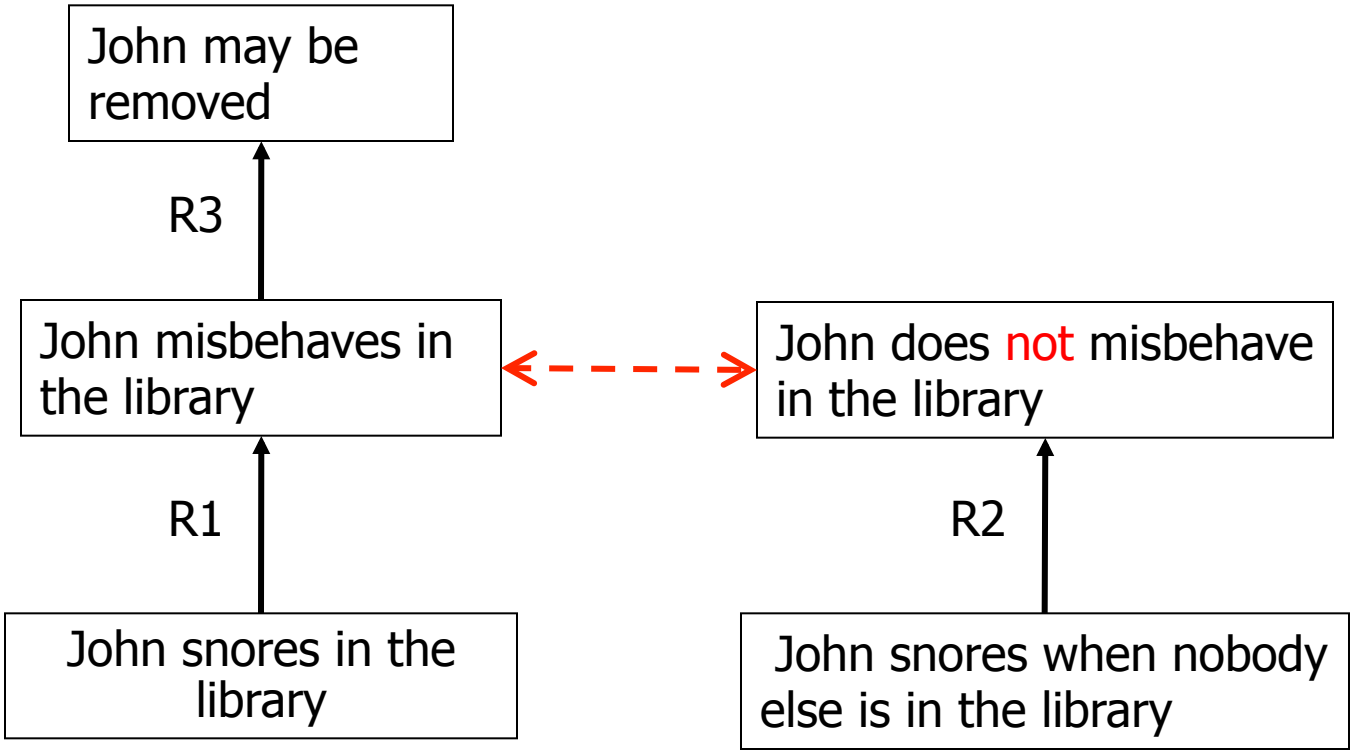

R1: If you snore, you misbehave
R2: If you snore when nobody else is around, you don't misbehave
R3: If you misbehave in the library, the librarian may remove you

R1 < R2 < R3

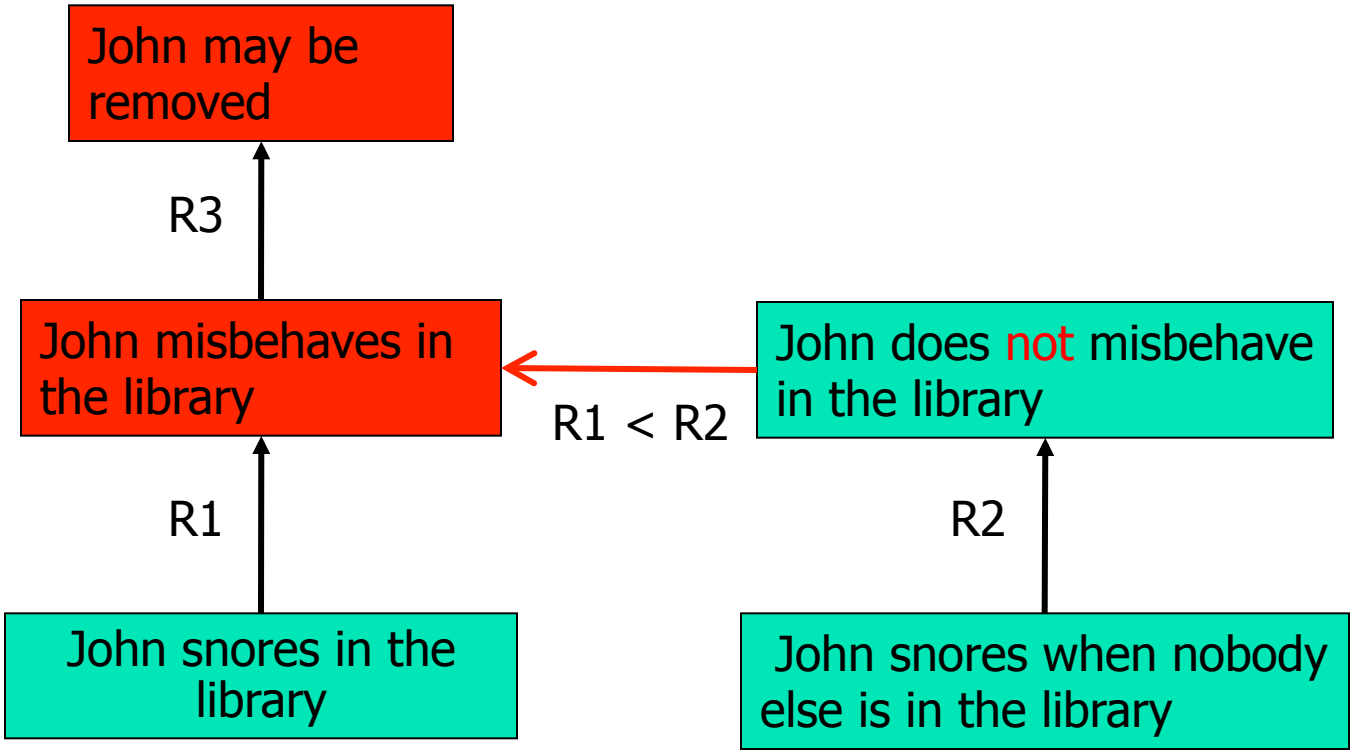

R1: If you snore, you misbehave
R2: If you snore when nobody else is around, you don't misbehave
R3: If you misbehave in the library, the librarian may remove you

R1 < R2 < R3

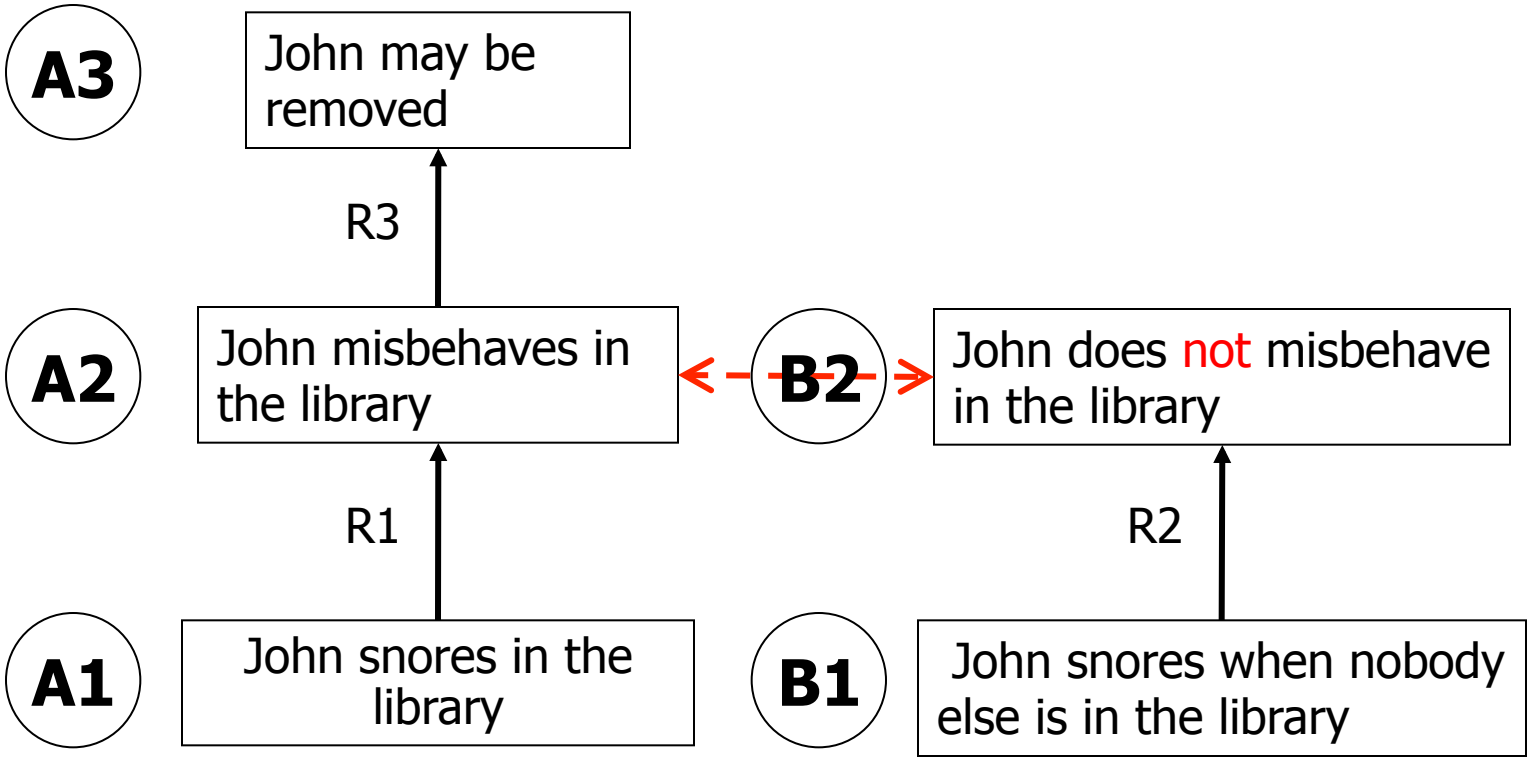

R1: If you snore, you misbehave
R2: If you snore when nobody else is around, you don't misbehave
R3: If you misbehave in the library, the librarian may remove you

R1 < R2 < R3 so A2 < B2 < A3 (with last link)

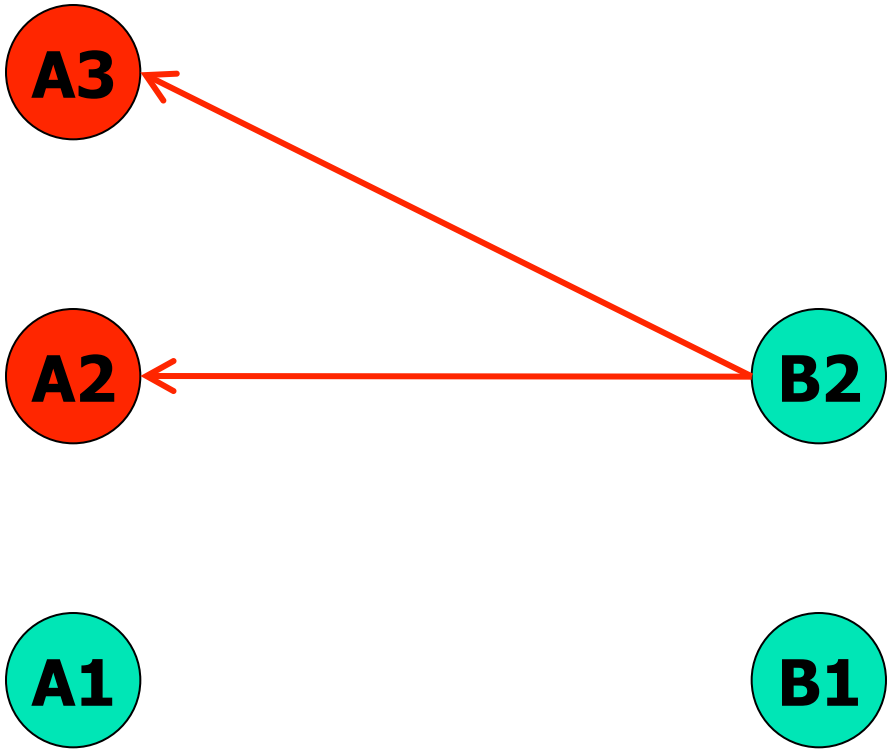

The defeat graph in ASPIC+

R1: If you snore, you misbehave
R2: If you snore when nobody else is around, you don't misbehave
R3: If you misbehave in the library, the librarian may remove you

R1 < R2 < R3 so A2 < B2 < A3 (with last link)

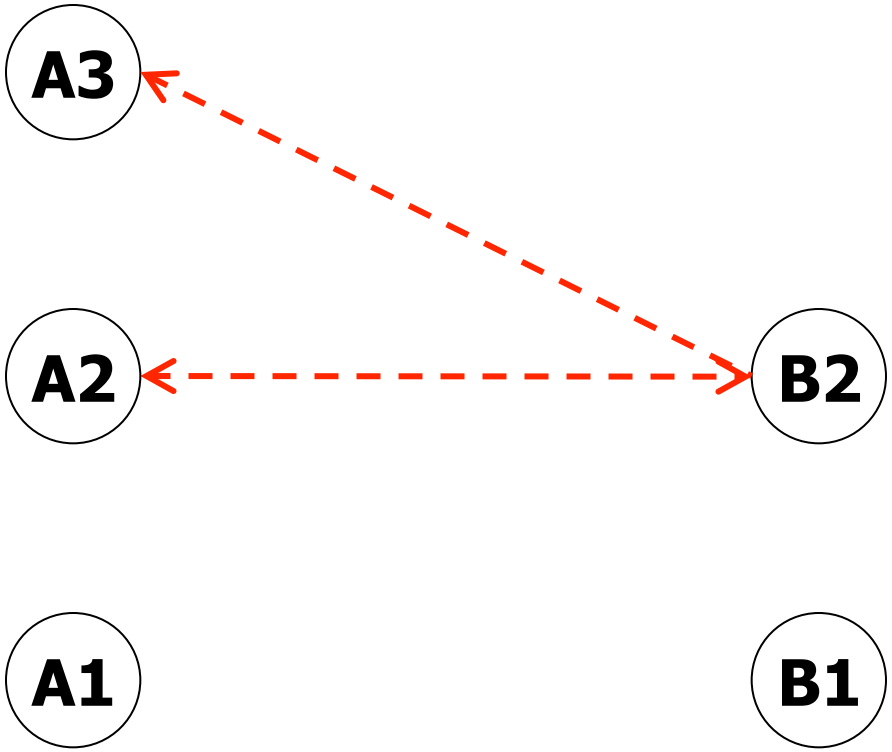

The attack graph in PAFs

R1: If you snore, you misbehave
R2: If you snore when nobody else is around, you don't misbehave
R3: If you misbehave in the library, the librarian may remove you

R1 < R2 < R3 so A2 < B2 < A3 (with last link)

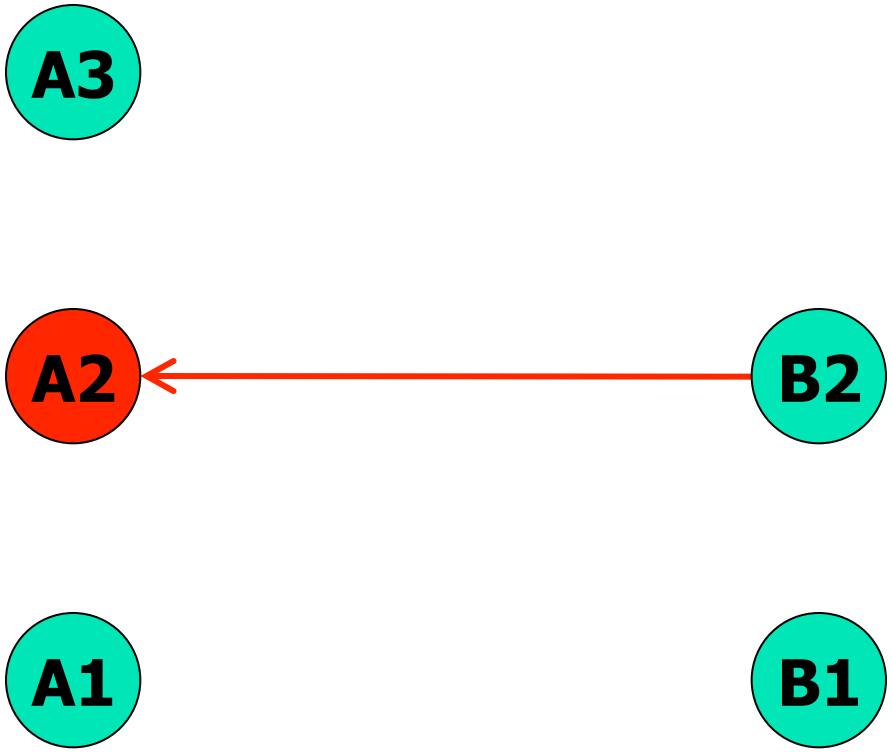

The defeat graph in PAFs

R1: If you snore, you misbehave
R2: If you snore when nobody else is around, you don't misbehave
R3: If you misbehave in the library, the librarian may remove you

R1 < R2 < R3 so A2 < B2 < A3 (with last link)

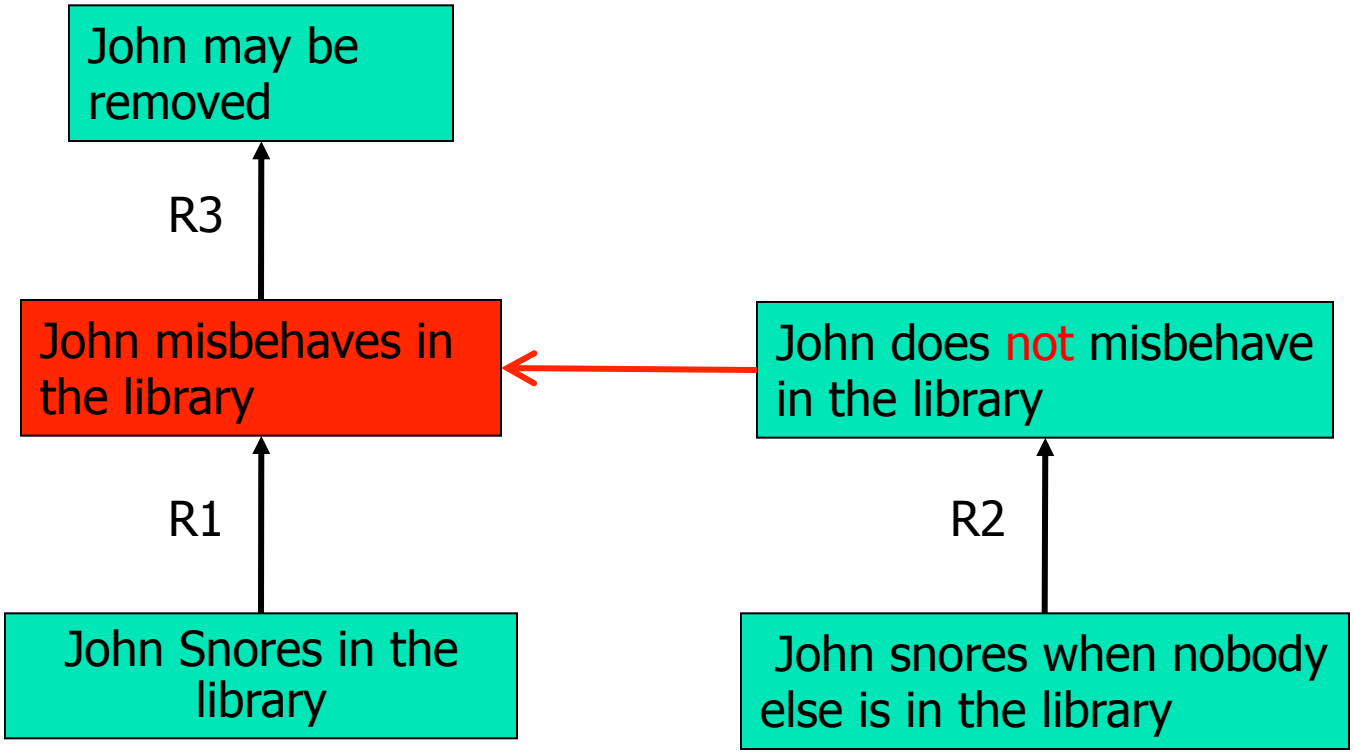

R1: If you snore, you misbehave
R2: If you snore when nobody else is around, you don't misbehave
R3: If you misbehave in the library, the librarian may remove you

R1 < R2 < R3 so A2 < B2 < A3 (with last link)

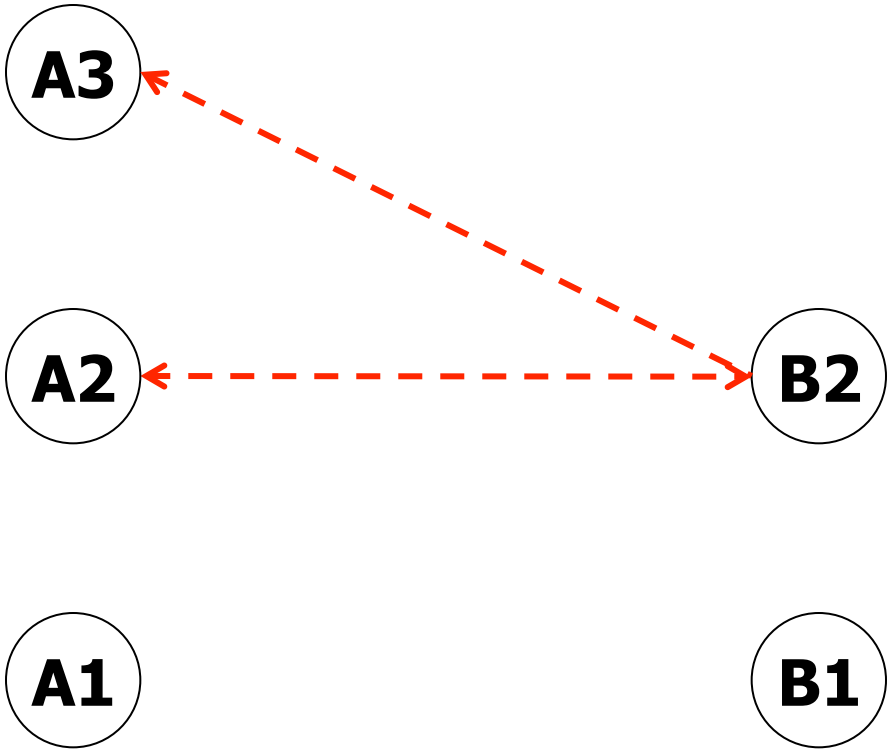

PAFs don't recognize that B2's attacks on A2 and A3 are the same

Work outside the Dung paradigm

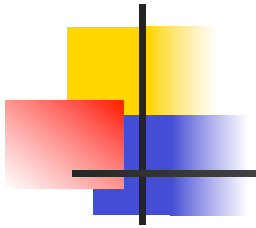

- Defeasible Logic Programming (Simari et al.)
  - Arguments roughly as in ASPIC+ but no Dung semantics
- Carneades (Gordon et al.)
  - Arguments pro and con a claim
- Abstract Dialectical Frameworks (Brewka & Woltran)
- …

Argument(ation) schemes: general form

Premise 1,

… ,

Premise n

Therefore (presumably), conclusion

- But also critical questions

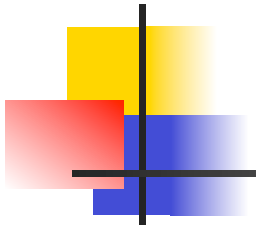

Argument schemes in ASPIC

- Argument schemes are defeasible inference rules
- Critical questions are pointers to counterarguments
  - Some point to undermining attacks
  - Some point to rebutting attacks
  - Some point to undercutting attacks

Reasoning with default generalisations

P
If P then normally/usually/typically Q
So (presumably), Q

- What experts say is usually true
- People with political ambitions are usually not objective about security
- People with names typical from country C usually have nationality C
- People who flea from a crime scene when the police arrives are normally involved in the crime
- Chinese people usually don't like coffee

- But defaults can have exceptions
- And there can be conflicting defaults

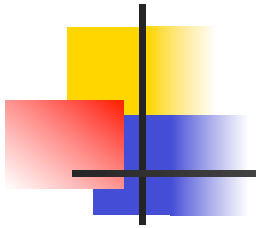

Perception

P is observed
Therefore (presumably), P

- Critical questions:
  - Are the observer's senses OK?
  - Are the circumstances such that reliable observation of P is impossible?
  - …

Inducing generalisations

Almost all observed P's were Q's
Therefore (presumably), If P then usually Q

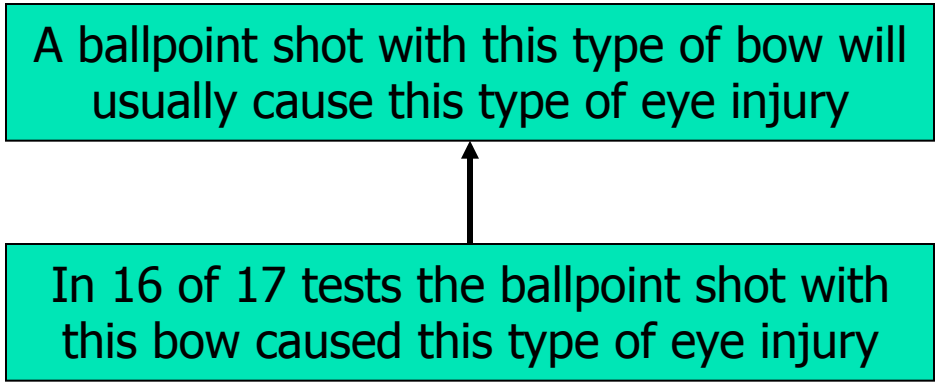

- Critical questions:
  - Is the size of the sample large enough?
  - was the sample selection biased?

Expert testimony (Walton 1996)

E is expert on D
E says that P
P is within D
Therefore (presumably), P is the case

- Critical questions:
  - Is E biased?
  - Is P consistent with what other experts say?
  - Is P consistent with known evidence?

Supporting and using generalisations

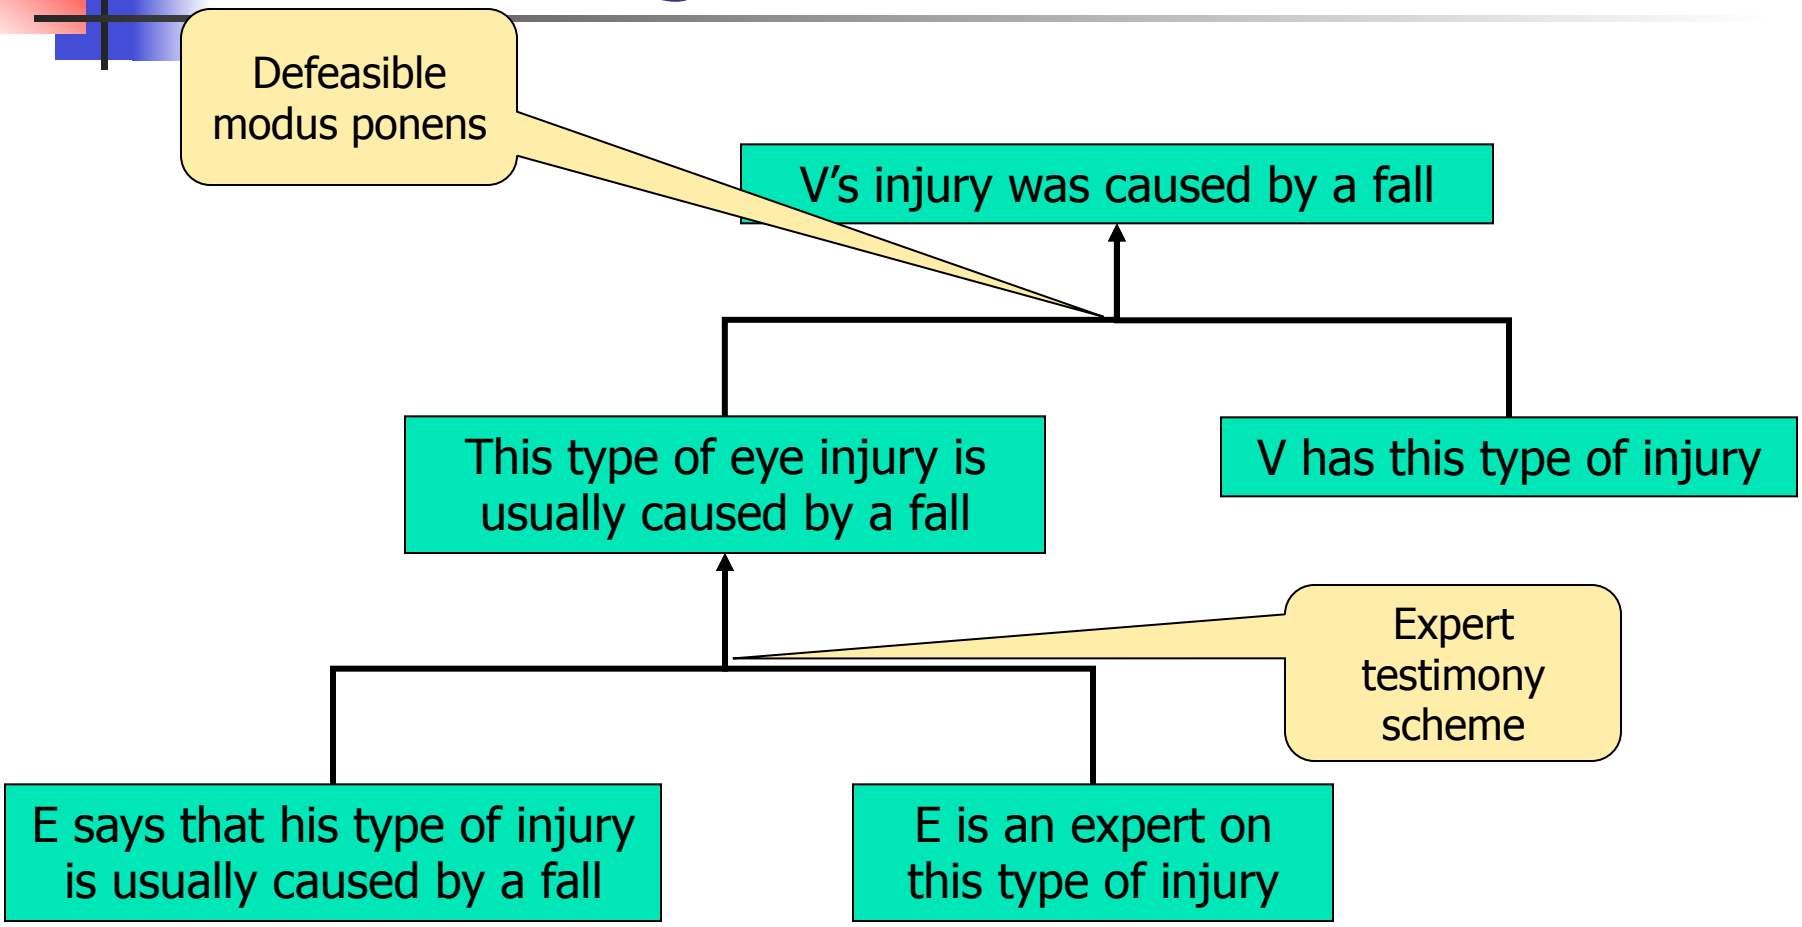

Arguments from consequences

Action A causes G,
G is good (bad)
Therefore (presumably), A should (not) be done

- Critical questions:
  - Does A also have bad (good) consequences?
  - Are there other ways to bring about G?
  - ...

Combining multiple good/bad consequences

Action A results in C1

…

Action A results in Cn

C1 is good

…

Cn is good

Therefore,

Action A is good

Action A results in C1

…

Action A results in Cn

C1 is bad

…

Cm is bad

Therefore,

Action A is bad

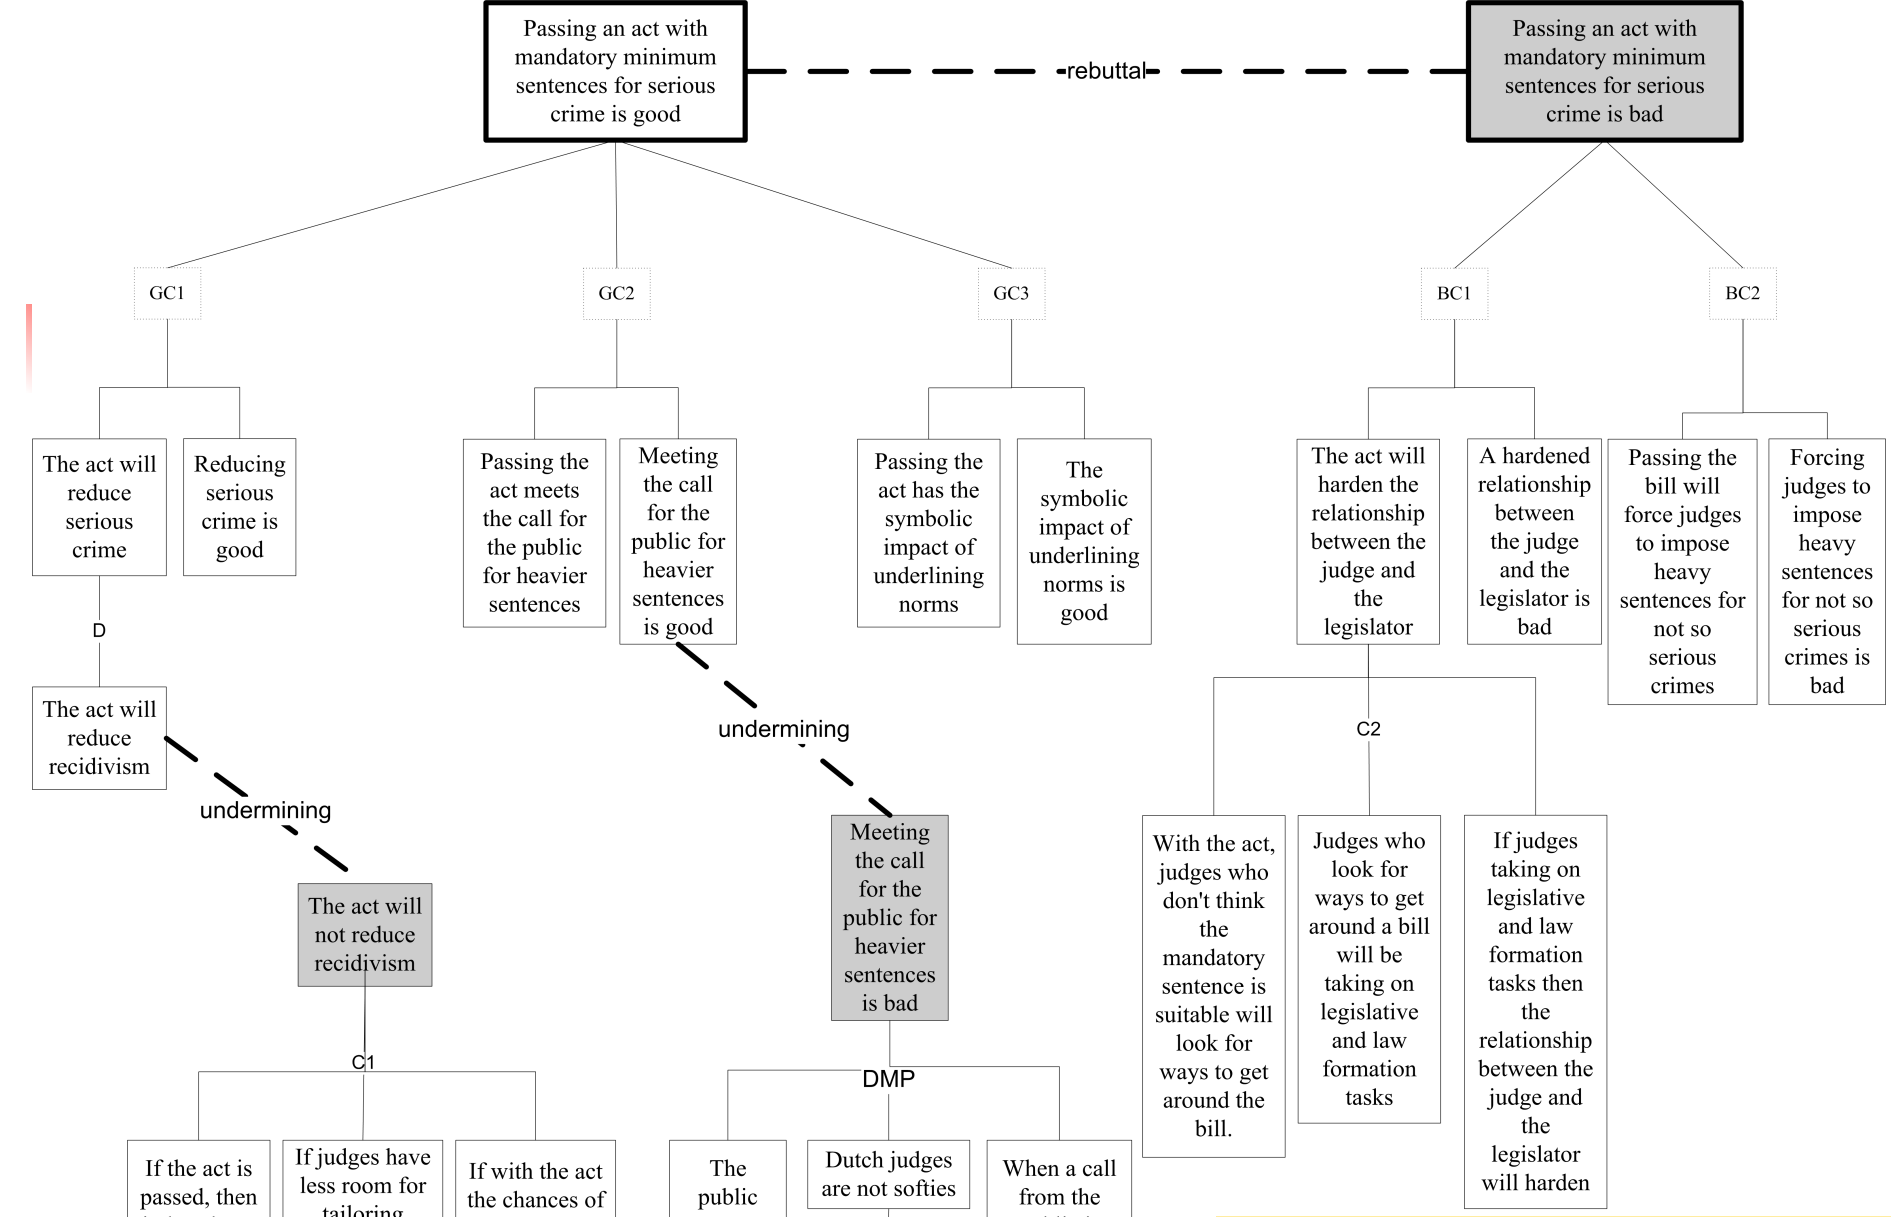

H. Prakken, Formalising a legal opinion on a legislative proposal in the ASPIC+ framework. *Proc. JURIX 2012*

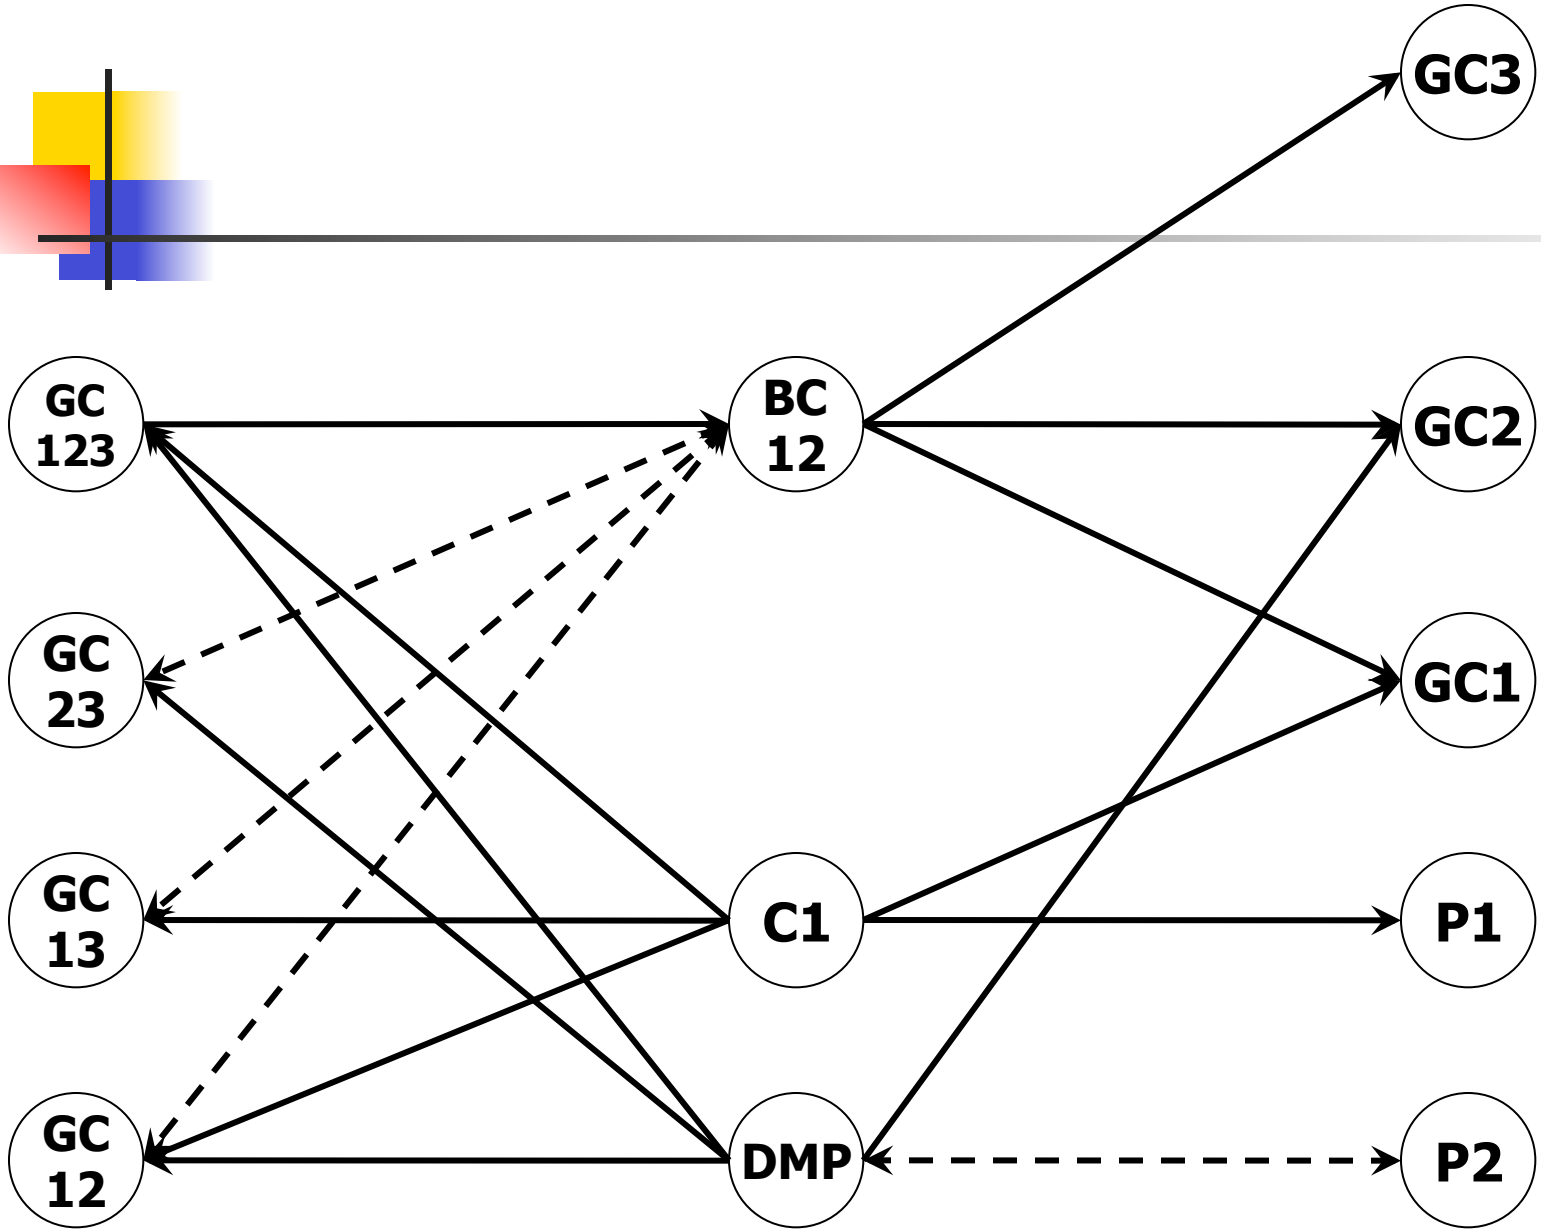
GC3
GC 123
BC 12
GC2
GC 23
GC1
GC 13
C1
P1
GC 12
DMP
P2

Preferred labelling 1

1. An argument is *In* iff all arguments that defeat it are *Out*.
2. An argument is *Out* iff some argument that defeats it is *In*.

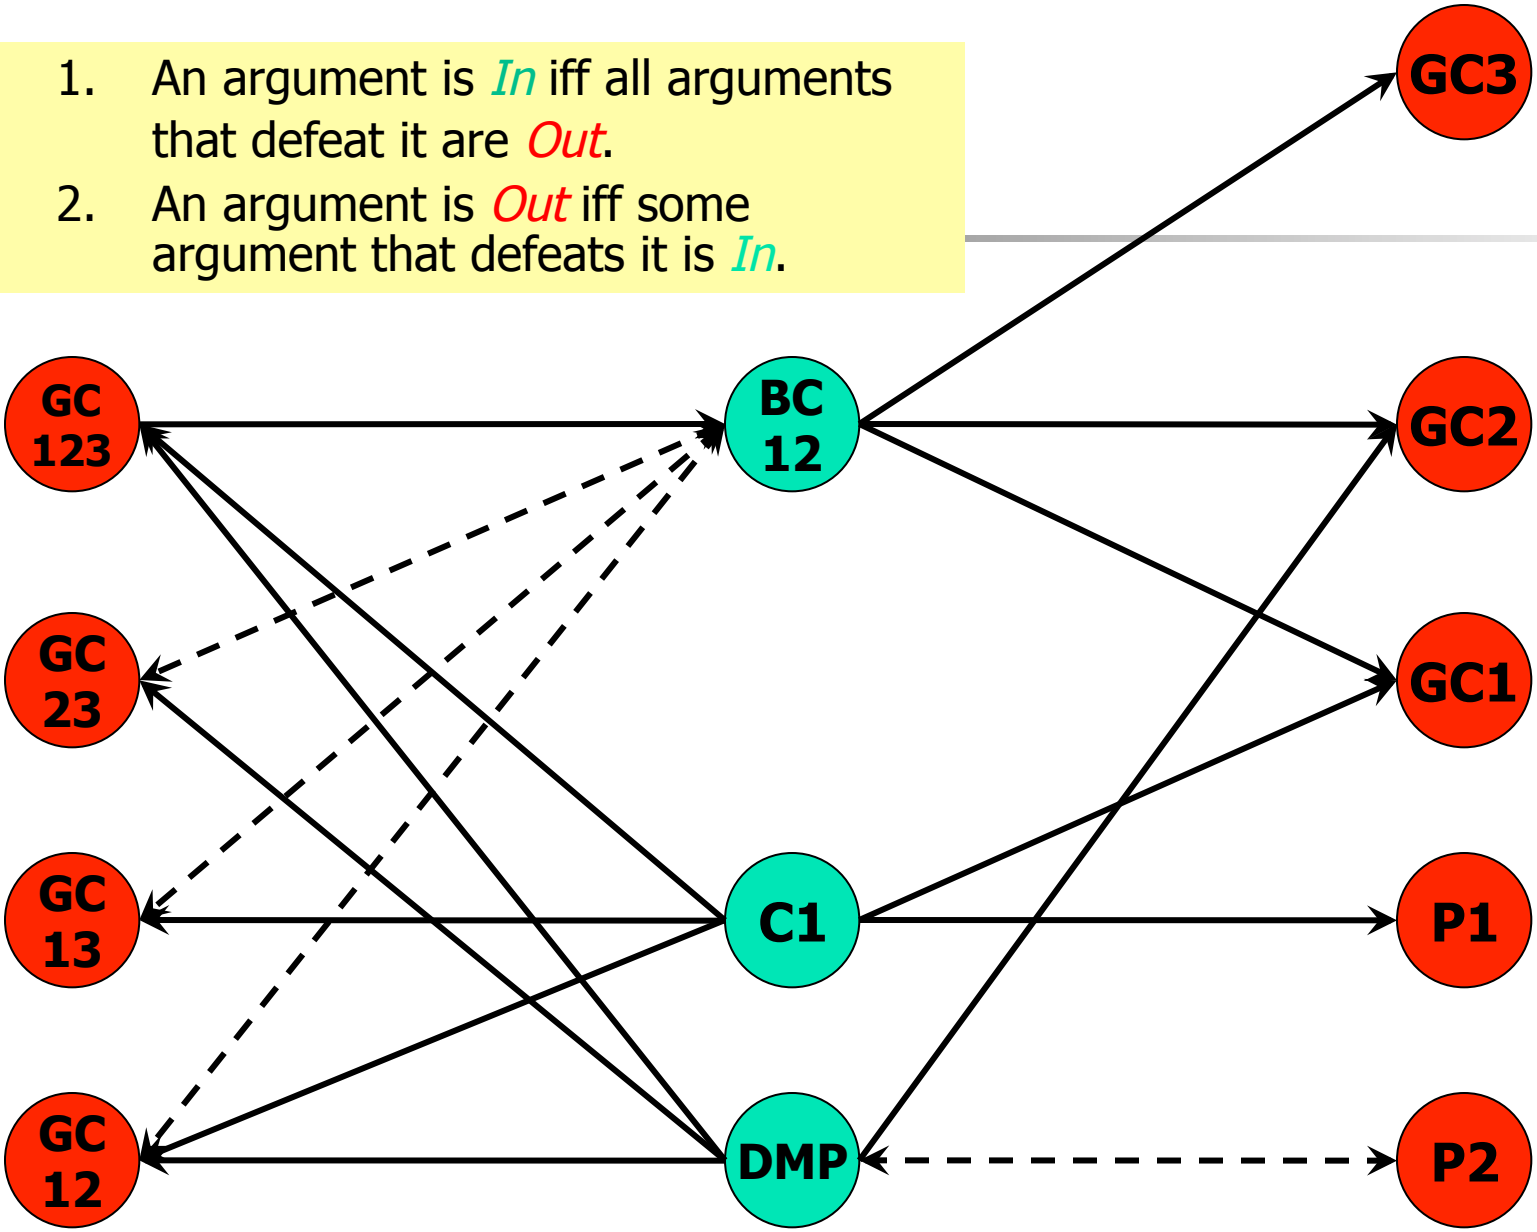

Preferred labelling 2

1. An argument is *In* iff all arguments that defeat it are *Out*.
2. An argument is *Out* iff some argument that defeats it is *In*.

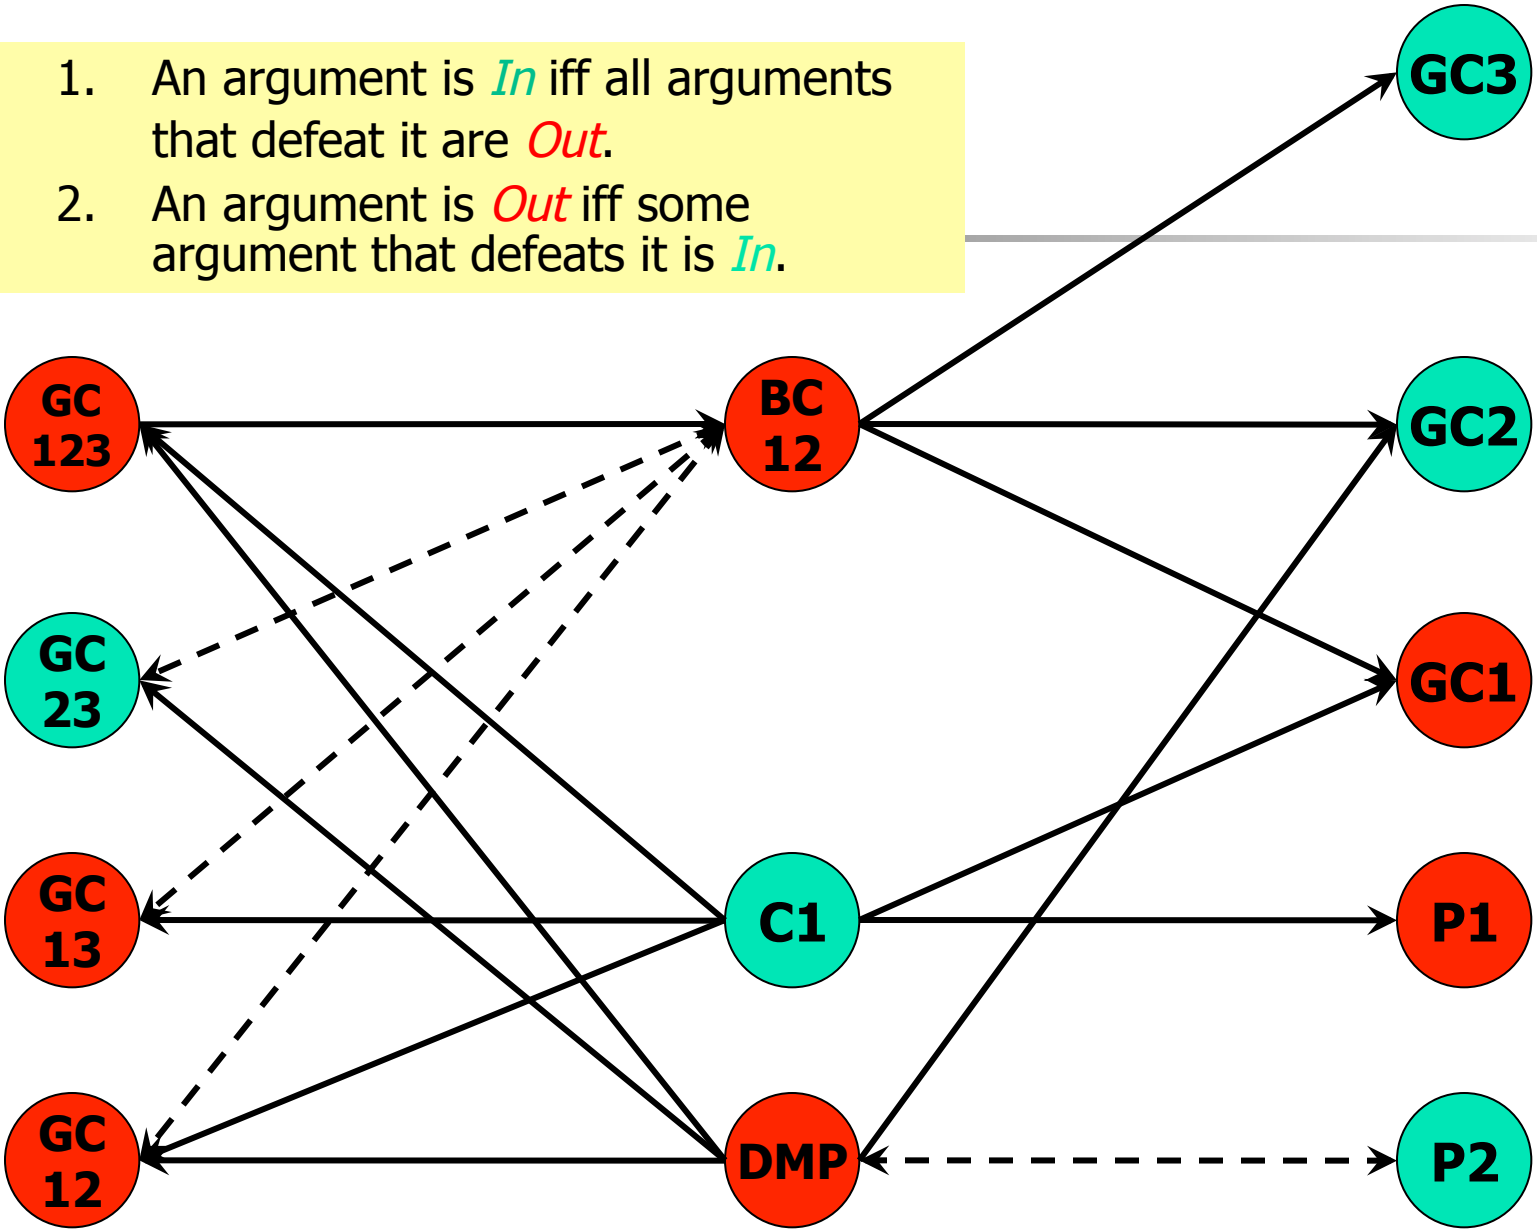

Grounded labelling

1. An argument is *In* iff all arguments that defeat it are *Out*.
2. An argument is *Out* iff some argument that defeats it is *In*.

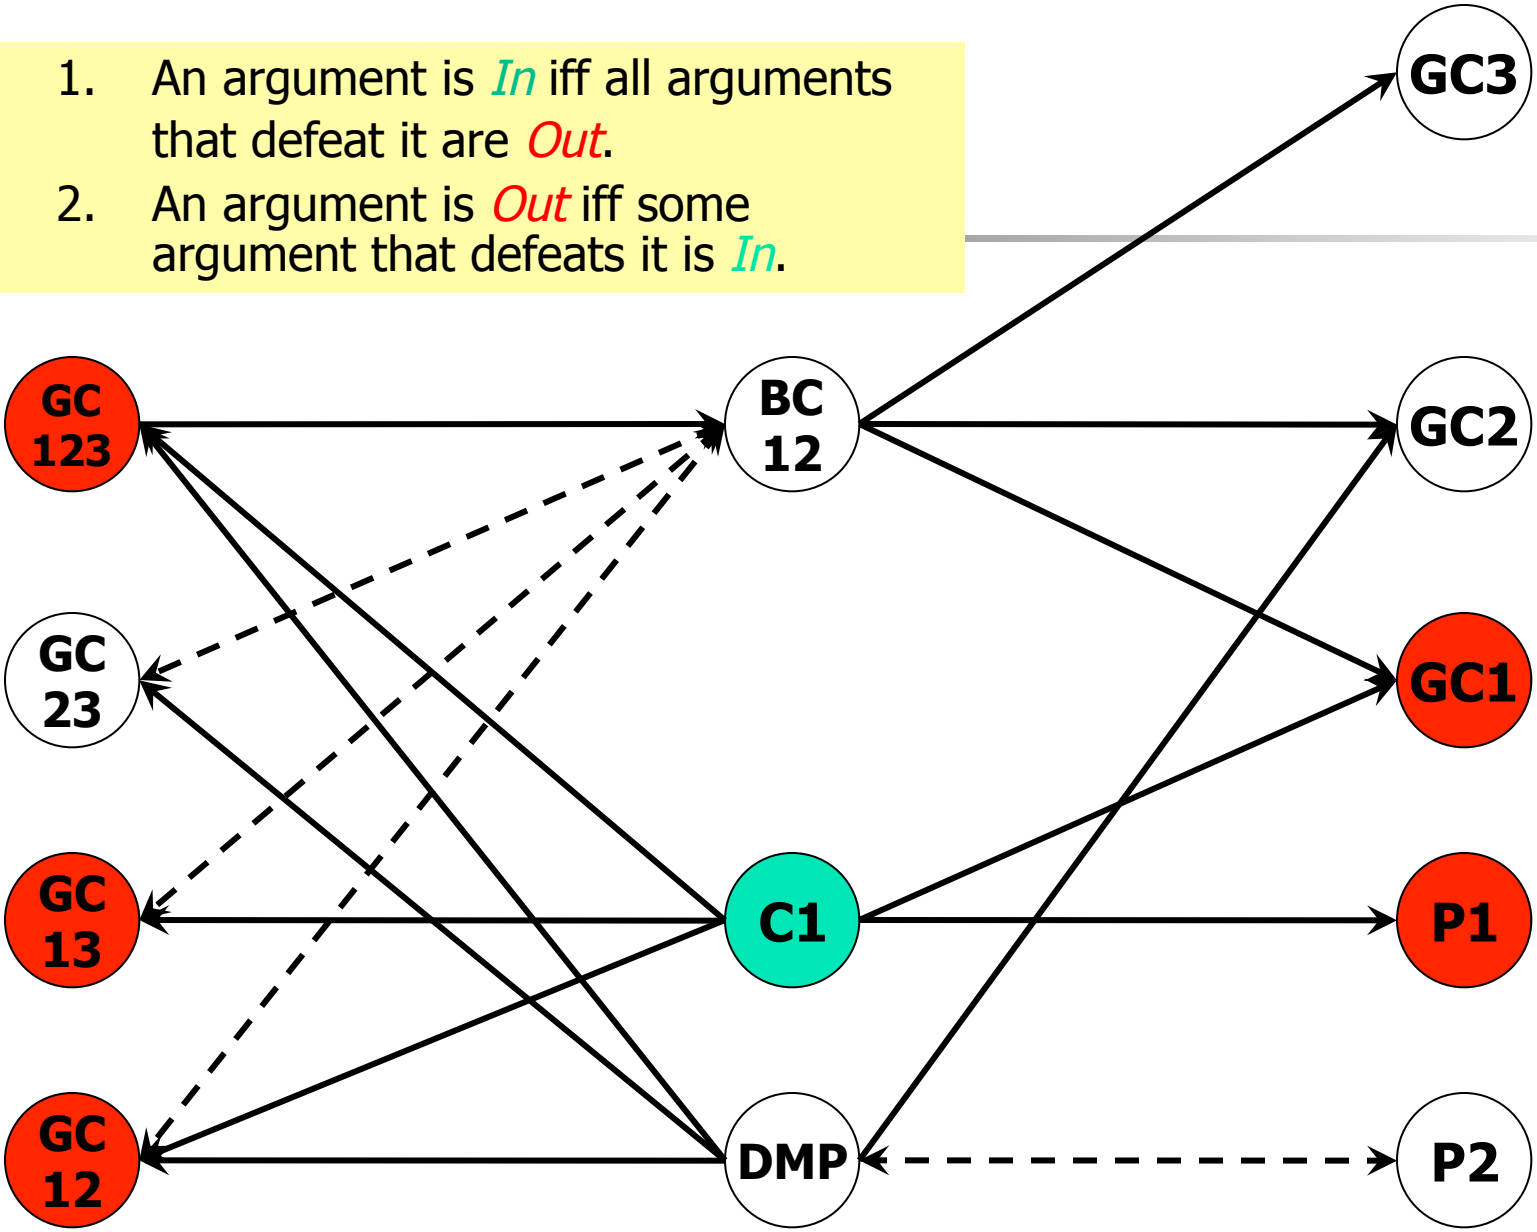

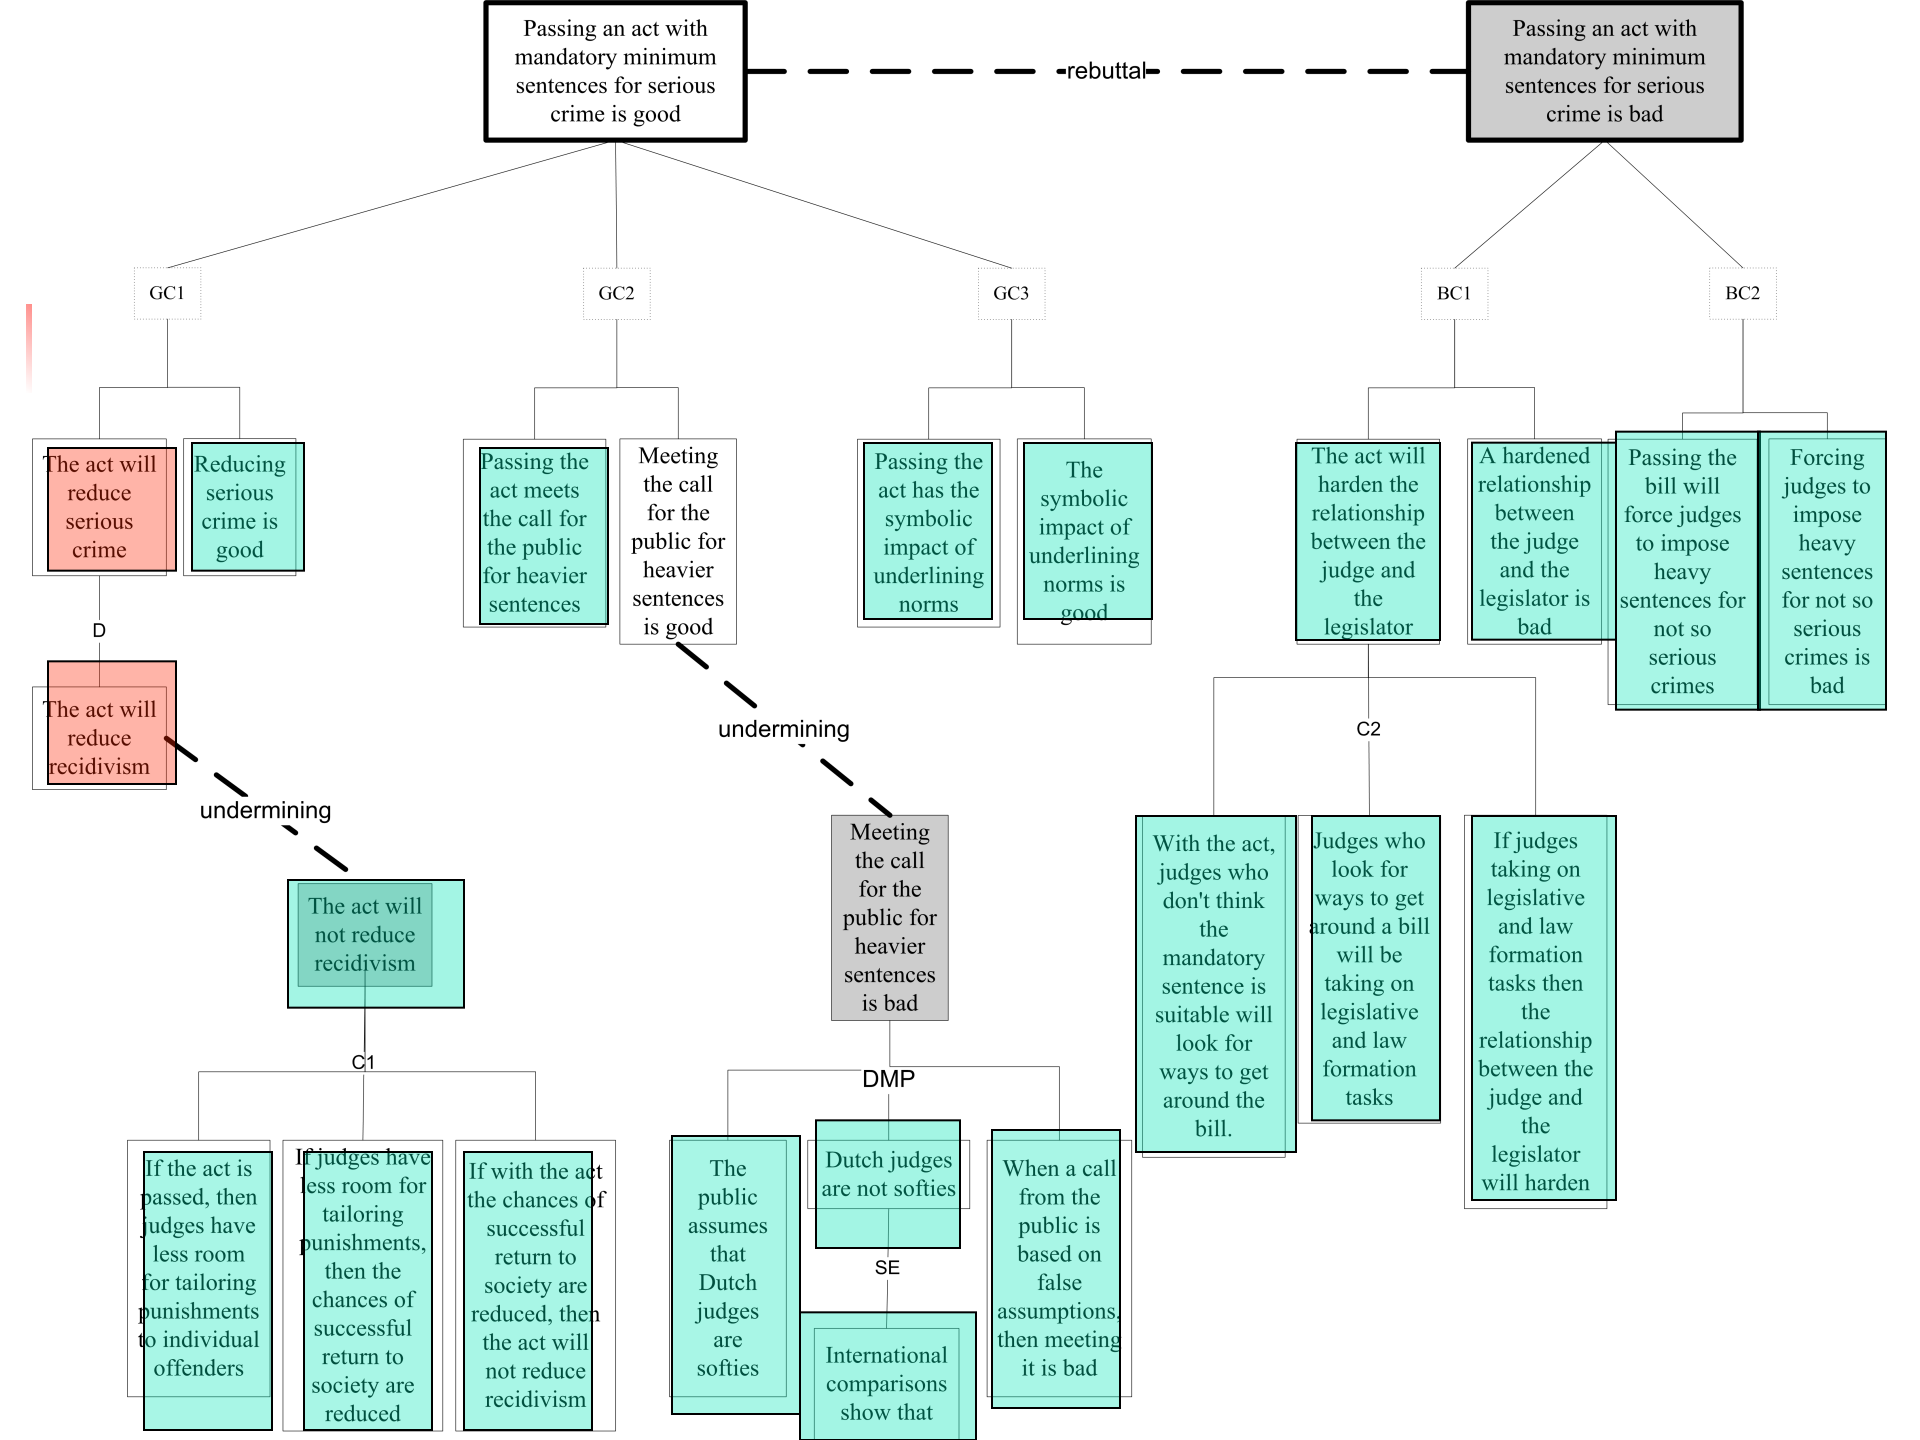
Passing an act with mandatory minimum sentences for serious crime is good
rebuttal
Passing an act with mandatory minimum sentences for serious crime is bad
GC1
GC2
GC3
BC1
BC2
The act will reduce serious crime
Reducing serious crime is good
Passing the act meets the call for the public for heavier sentences
Meeting the call for the public for heavier sentences is good
Passing the act has the symbolic impact of underlining norms
The symbolic impact of underlining norms is good
The act will harden the relationship between the judge and the legislator
A hardened relationship between the judge and the legislator is bad
Passing the bill will force judges to impose heavy sentences for not so serious crimes
Forcing judges to impose heavy sentences for not so serious crimes is bad
D
The act will reduce recidivism
undermining
undermining
C2
With the act, judges who don't think the mandatory sentence is suitable will look for ways to get around the bill.
Judges who look for ways to get around a bill will be taking on legislative and law formation tasks
If judges taking on legislative and law formation tasks then the relationship between the judge and the legislator will harden
The act will not reduce recidivism
Meeting the call for the public for heavier sentences is bad
C1
DMP
If the act is passed, then judges have less room for tailoring punishments to individual offenders
If judges have less room for tailoring punishments, then the chances of successful return to society are reduced
If with the act the chances of successful return to society are reduced, then the act will not reduce recidivism
The public assumes that Dutch judges are softies
Dutch judges are not softies
When a call from the public is based on false assumptions, then meeting it is bad
SE
International comparisons show that

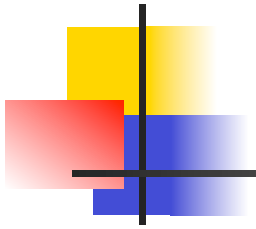

Summary

- A formal metatheory of structured argumentation is emerging
- Better understanding needed of philosophical underpinnings and practical applicability
  - Not all argumentation can be naturally reduced to plausible reasoning
  - The 'one base logic' approach is only suitable for plausible reasoning
- Important research issues:
  - Aggregation of arguments
  - Relation with probability theory

Interaction

- Argument games verify status of argument (or statement) given a single theory (knowledge base)
- But real argumentation dialogues have
  - Distributed information
  - Dynamics
  - Real players!
  - Richer communication languages

Example

P: Tell me all you know about recent trading in explosive materials (request)

P: why don't you want to tell me?

P: why aren't you allowed to tell me?

P: You may be right in general (concede) but in this case there is an exception since this is a matter of national importance

P: since we have heard about a possible terrorist attack

P: OK, I agree (offer accepted).

O: No I won't (reject)

O: since I am not allowed to tell you

O: since sharing such information could endanger an investigation

O: Why is this a matter of national importance?

O: I concede that there is an exception, so I retract that I am not allowed to tell you. I will tell you on the condition that you don't exchange the information with other police officers (offer)

Example

P: Tell me all you know about recent trading in explosive materials (request)

P: why don't you want to tell me?

P: why aren't you allowed to tell me?

P: You may be right in general (concede) but in this case there is an exception since this is a matter of national importance

P: since we have heard about a possible terrorist attack

P: OK, I agree (offer accepted).

O: No I won't (reject)

O: since I am not allowed to tell you

O: since sharing such information could endanger an investigation

O: Why is this a matter of national importance?

O: I concede that there is an exception, so I retract that I am not allowed to tell you. I will tell you on the condition that you don't exchange the information with other police officers (offer)

Example

P: Tell me all you know about recent trading in explosive materials (request)

P: why don't you want to tell me?

P: why aren't you allowed to tell me?

P: You may be right in general (concede) but in this case there is an exception since this is a matter of national importance

P: since we have heard about a possible terrorist attack

P: OK, I agree (offer accepted).

O: No I won't (reject)

O: since I am not allowed to tell you

O: since sharing such information could endanger an investigation

O: Why is this a matter of national importance?

O: I concede that there is an exception, so I retract that I am not allowed to tell you. I will tell you on the condition that you don't exchange the information with other police officers (offer)

Types of dialogues (Walton & Krabbe)

Dialogue Type	*Dialogue Goal*	*Initial situation*
Persuasion	resolution of conflict	conflict of opinion
Negotiation	making a deal	conflict of interest
Deliberation	reaching a decision	need for action
Information seeking	exchange of information	personal ignorance
Inquiry	growth of knowledge	general ignorance

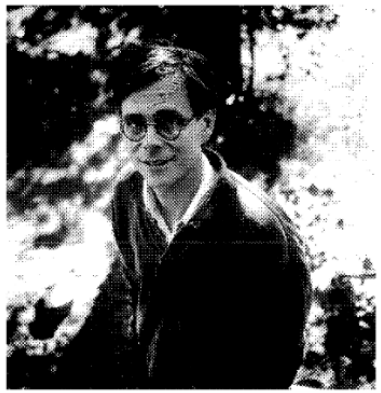

Dialogue systems (according to Carlson 1983)

- Dialogue systems define the conditions under which an utterance is appropriate
- An utterance is appropriate if it promotes the goal of the dialogue in which it is made

⇒

- Appropriateness defined not at speech act level but at dialogue level

⇒

- Dialogue game approach
  - Protocol should promote the goal of the dialogue

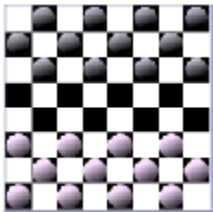

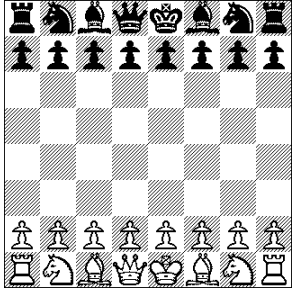

Dialogue game systems

- A communication language
  - Well-formed utterances
- Rules for when an utterance is allowed
  - Protocol
- Effect rules
- Turntaking rules
- Termination + outcome rules

Agent design:
strategies for selecting from the allowed utterances

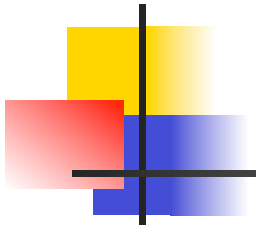

Effect rules

- Specify commitments
  - "Claim p" and "Concede p" commits to p
  - "p since Q" commits to p and Q
  - "Retract p" ends commitment to p
  - …
- Commitments used for:
  - Determining outcome
  - Enforcing 'dialogical consistency'
  - …

Public semantics for dialogue protocols

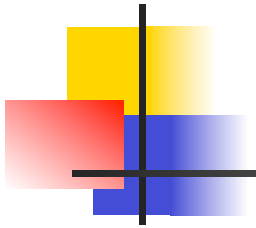

- Public semantics: can protocol compliance be externally observed?
- Commitments are a participant's publicly declared standpoints, so not the same as beliefs!
- Only commitments and dialogical behaviour should count for move legality:
  - "Claim p is allowed only if you believe p"

    vs.
  - "Claim p is allowed only if you are not committed to ¬p and have not challenged p"

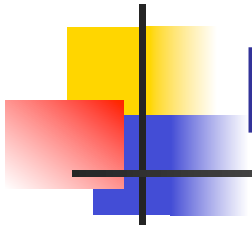

More and less strict protocols

- Single-multi move: one or more moves per turn allowed
- Single-multi-reply: one or more replies to the same move allowed
- Deterministic: no choice from legal moves
- Deterministic in communication language: no choice from speech act types
- Only reply to moves from previous turn?

Some properties that can be studied

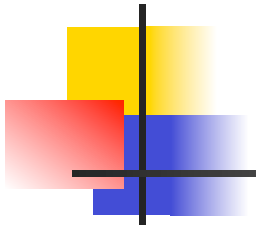

- Correspondence with players' beliefs
  - If union of beliefs implies p, can/will agreement on p result?
  - If players agree on p, does union of beliefs imply p?
  - Disregarding vs. assuming player strategies

Example 1

Knowledge bases

Paul: r

Olga: s

Inference rules

$p \Rightarrow q$

$r \Rightarrow p$

$s \Rightarrow \neg r$

P1: q since p

Paul ∪ Olga does not justify q but they could agree on q

Olga is credulous: she concedes everything for which she cannot construct a (defensible or justified) counterargument

Example 1

Knowledge bases

Paul: r

Olga: s

Inference rules

$p \Rightarrow q$

$r \Rightarrow p$

$s \Rightarrow \neg r$

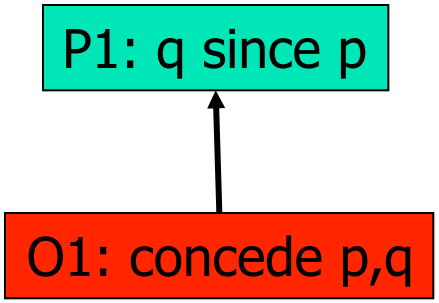

Paul $\cup$ Olga does not justify q
but they could agree on q

Example 1

Knowledge bases

Paul: r

Olga: s

Inference rules

$p \Rightarrow q$

$r \Rightarrow p$

$s \Rightarrow \neg r$

P1: q since p

Paul ∪ Olga does not justify q but they could agree on q

Olga is sceptical: she challenges everything for which she cannot construct a (defensible or justified) argument

Example 1

Knowledge bases

Paul: r

Olga: s

Inference rules

$p \Rightarrow q$

$r \Rightarrow p$

$s \Rightarrow \neg r$

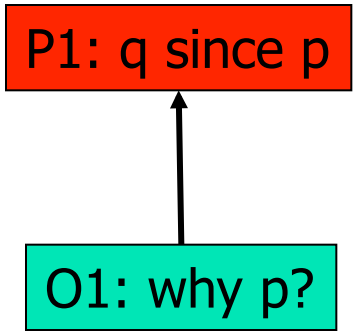

Paul $\cup$ Olga does not justify q
but they could agree on q

Example 1

Knowledge bases

Paul: r

Olga: s

Inference rules

$p \Rightarrow q$

$r \Rightarrow p$

$s \Rightarrow \neg r$

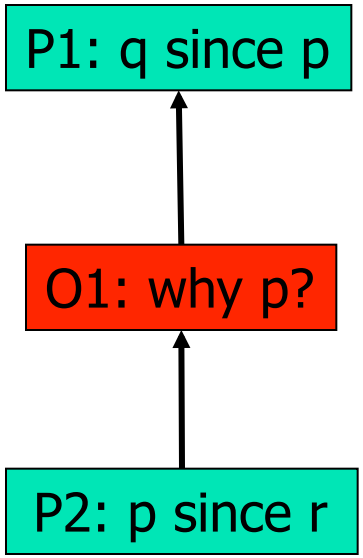

Paul $\cup$ Olga does not justify q
but they could agree on q

Example 1

Knowledge bases

Paul: r

Olga: s

Inference rules

$p \Rightarrow q$

$r \Rightarrow p$

$s \Rightarrow \neg r$

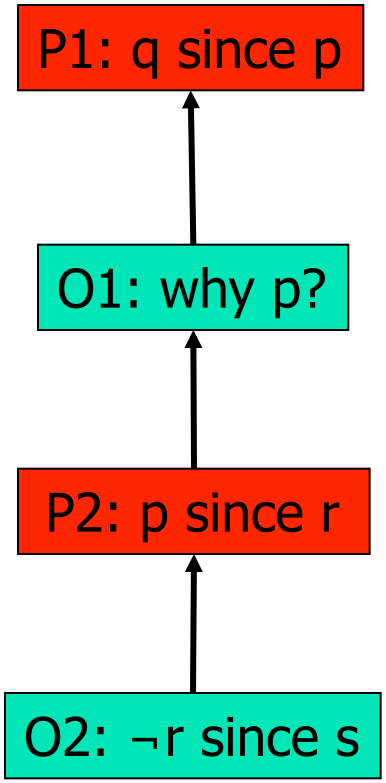

Paul ∪ Olga does not justify q
but they could agree on q

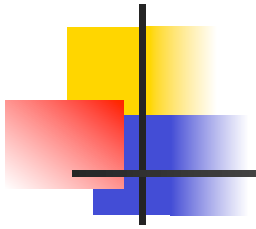

Example 2

Knowledge bases

Paul:
p
q

Olga:
p
$q \supset \neg p$

Inference rules

Modus ponens
…

P1: claim p

Paul $\cup$ Olga does not justify p but they will agree on p if players are conservative, that is, if they stick to their beliefs if possible

Example 2

Knowledge bases

Paul:
p
q

Olga:
p
$q \supset \neg p$

Inference rules

Modus ponens
…

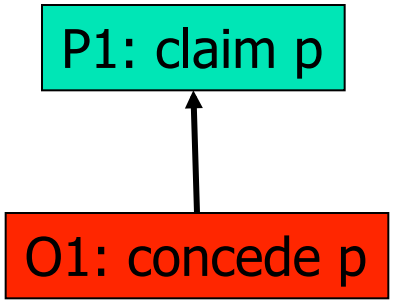

Paul ∪ Olga does not justify p but they will agree on p if players are conservative, that is, if they stick to their beliefs if possible

Example 2

Knowledge bases

Paul:
p
q

Olga:
p
$q \supset \neg p$

Inference rules

Modus ponens
…

P1: claim p

O1: what about q?

Possible solution (for open-minded agents, who are prepared to critically test their beliefs):

Example 2

Knowledge bases

Paul:
p
q

Olga:
p
$q \supset \neg p$

Inference rules

Modus ponens
…

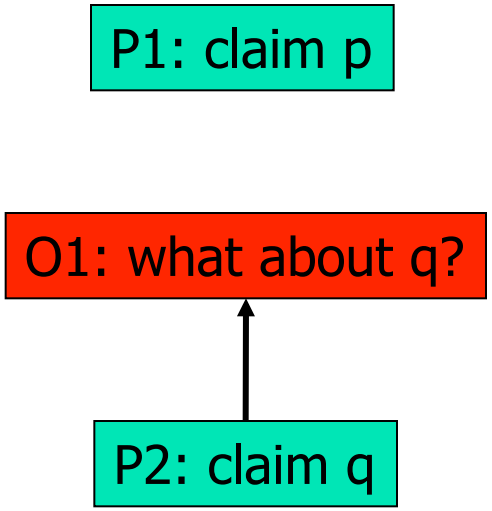

Possible solution (for open-minded agents, who are prepared to critically test their beliefs):

Example 2

Knowledge bases

Paul:
p
q

Olga:
p
q ⊃ ¬p

Inference rules

Modus ponens
…

Problem: how to ensure relevance?

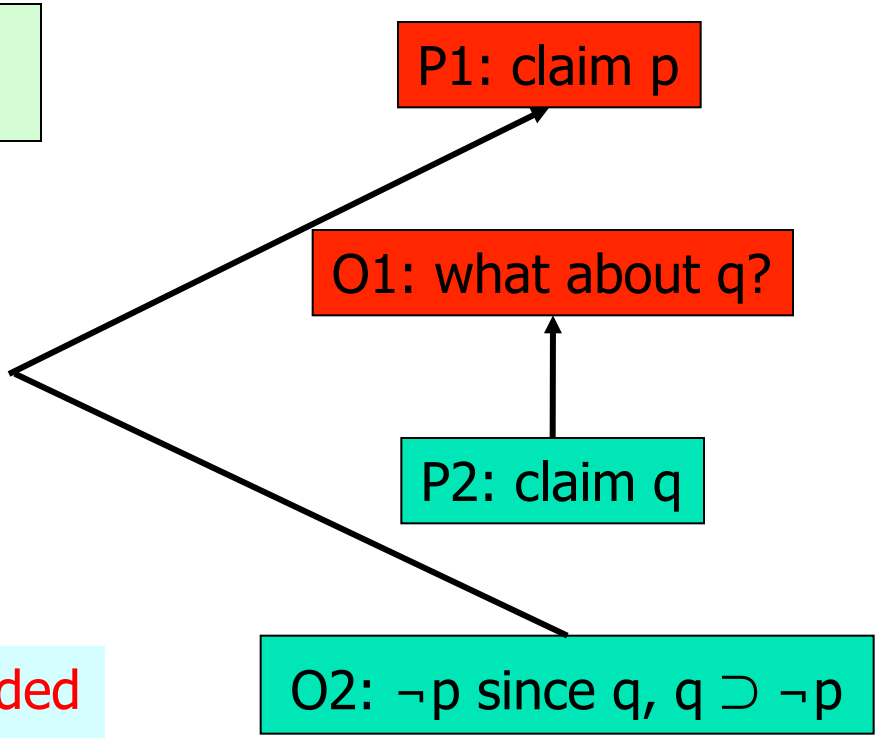

Possible solution (for open-minded agents, who are prepared to critically test their beliefs):

Automated Support of Regulated Data Exchange. A Multi-Agent Systems Approach

PhD Thesis Pieter Dijkstra (2012)
Faculty of Law
University of Groningen

rijksuniversiteit groningen

The communication language

Speech act	***Attack***	***Surrender***
request(φ)	offer (φ′), reject(φ)	-
offer(φ)	offer(φ′) (φ ≠ φ′), reject(φ)	accept(φ)
reject(φ)	offer(φ′) (φ ≠ φ′), why-reject (φ)	-
accept(φ)	-	-
why-reject(φ)	claim (φ′)	-
claim(φ)	why(φ)	concede(φ)
why(φ)	φ since S (an argument)	retract(φ)
φ since S	why(φ) (φ ∈ S) φ′ since S′ (a defeater)	concede(φ) concede φ′ (φ′ ∈ S)
concede(φ)	-	-
retract(φ)	-	-
deny(φ)	-	-

The protocol

- Start with a request
- Repy to an earlier move of the other agent
- Pick your replies from the table
- Finish persuasion before resuming negotiation
- Turntaking:
  - In nego: after each move
  - In pers: various rules possible
- Termination:
  - In nego: if offer is accepted or someone withdraws
  - In pers: if main claim is retracted or conceded

Example dialogue formalised

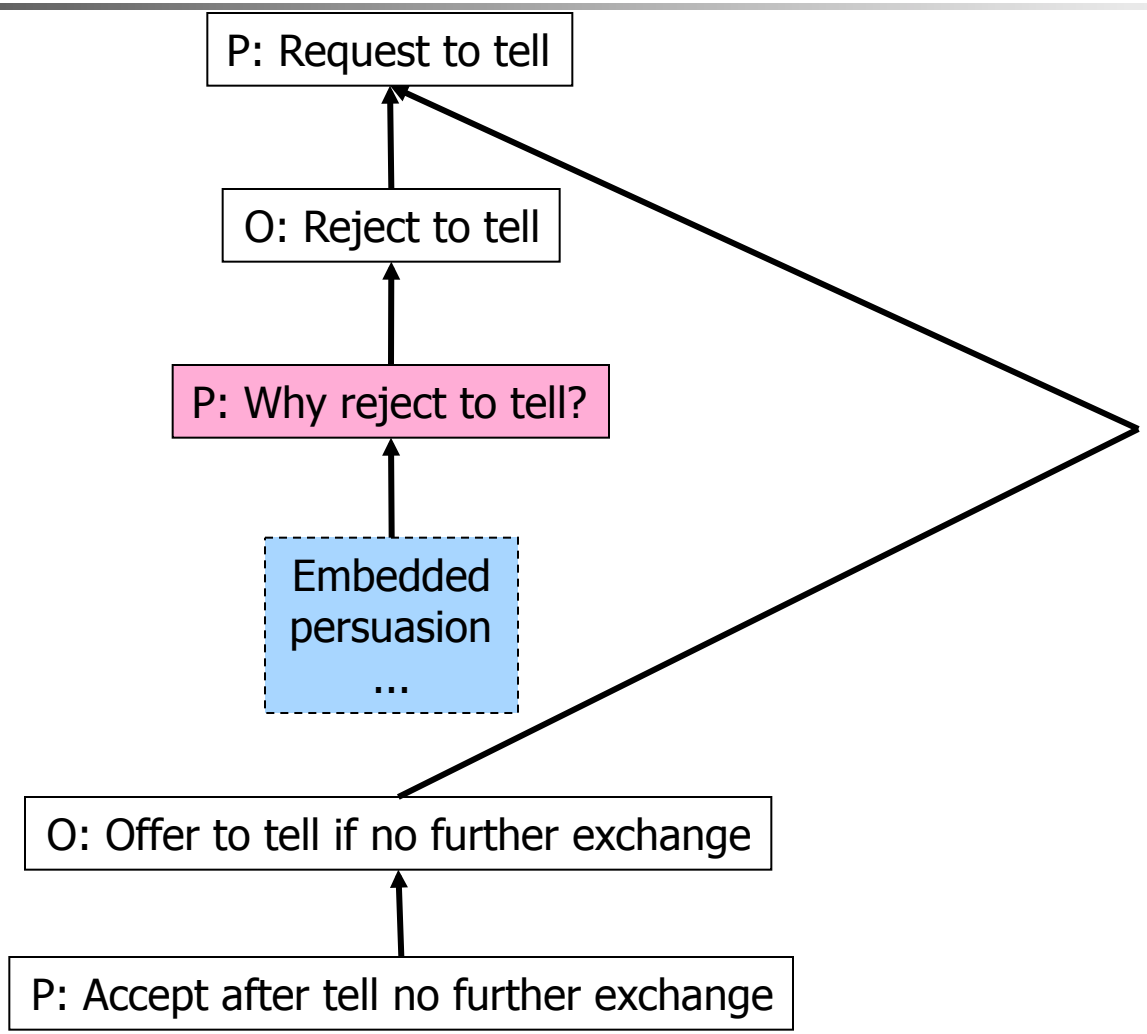

Persuasion part formalised

O: ***Claim*** Not allowed to tell

Persuasion part formalised

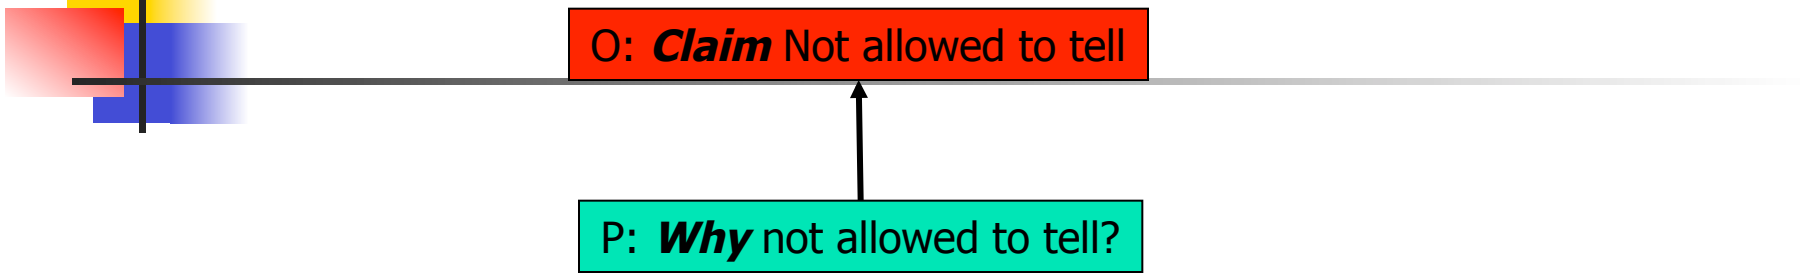

Persuasion part formalised

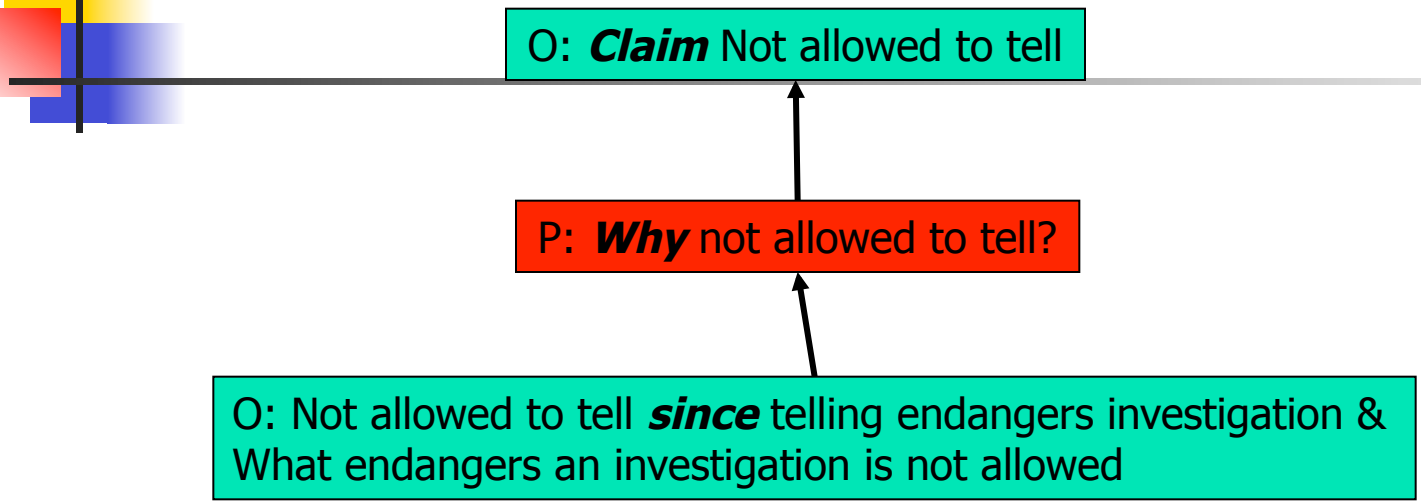

Persuasion part formalised

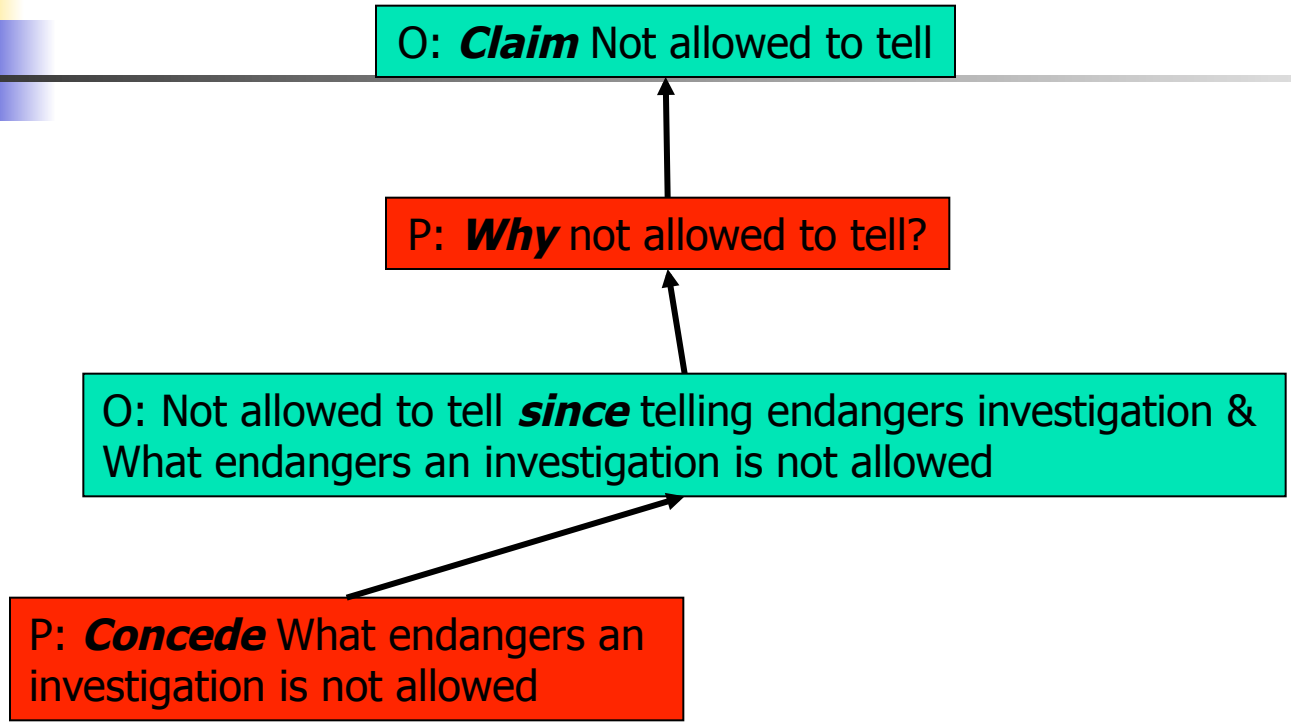

Persuasion part formalised

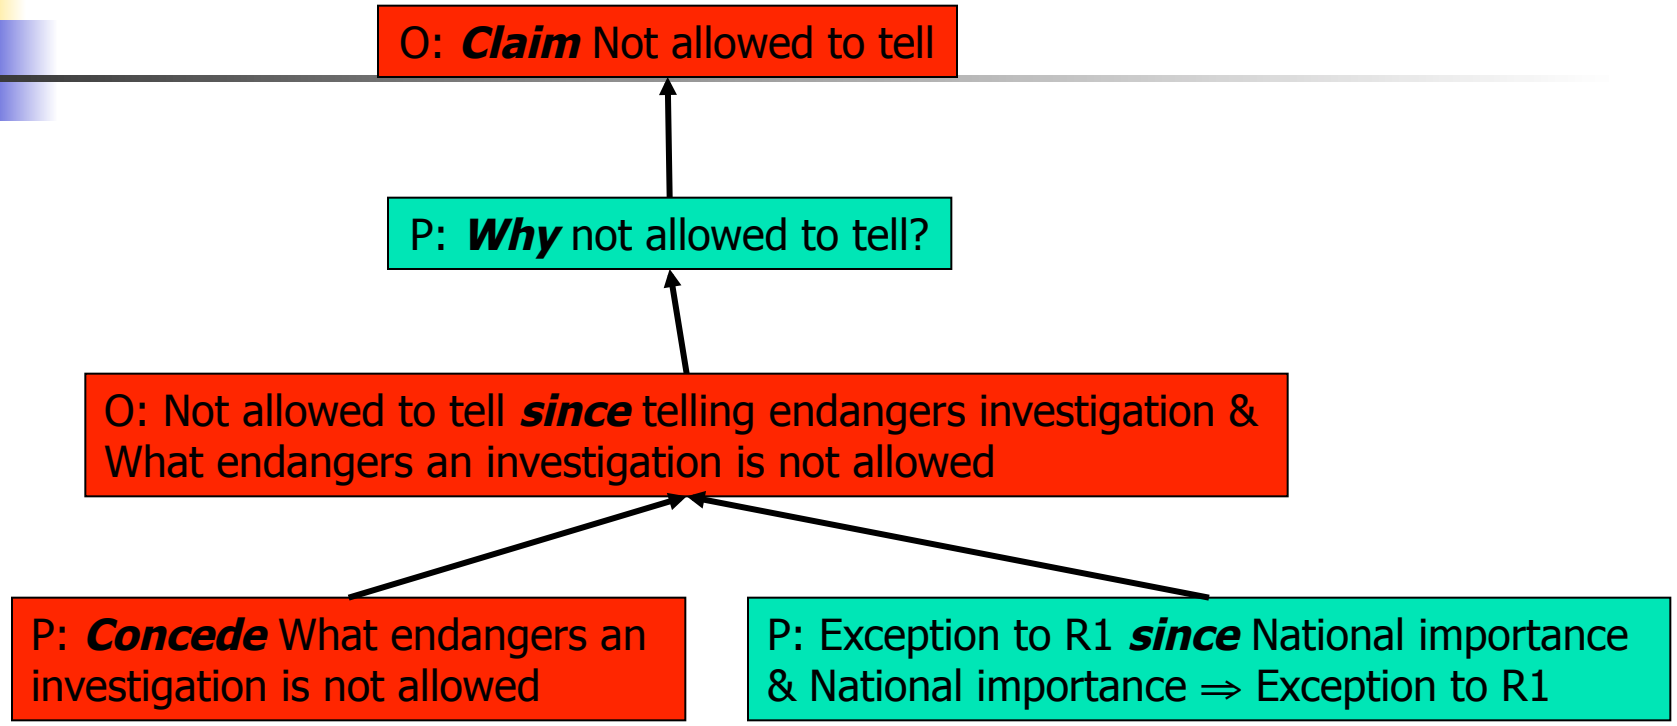

Persuasion part formalised

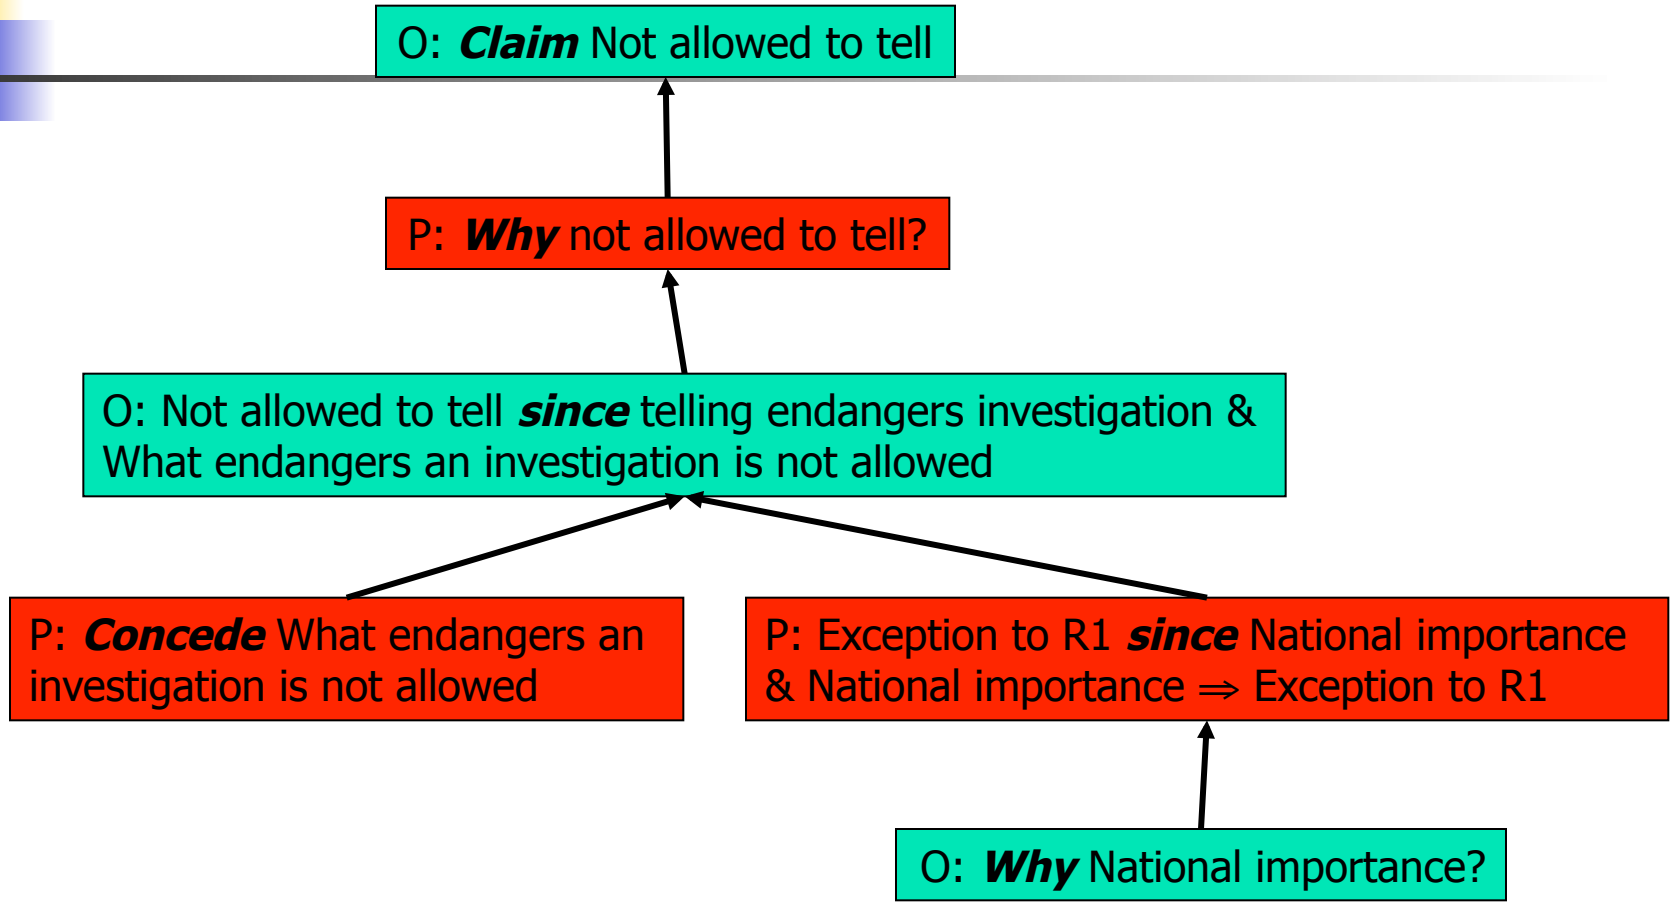

Persuasion part formalised

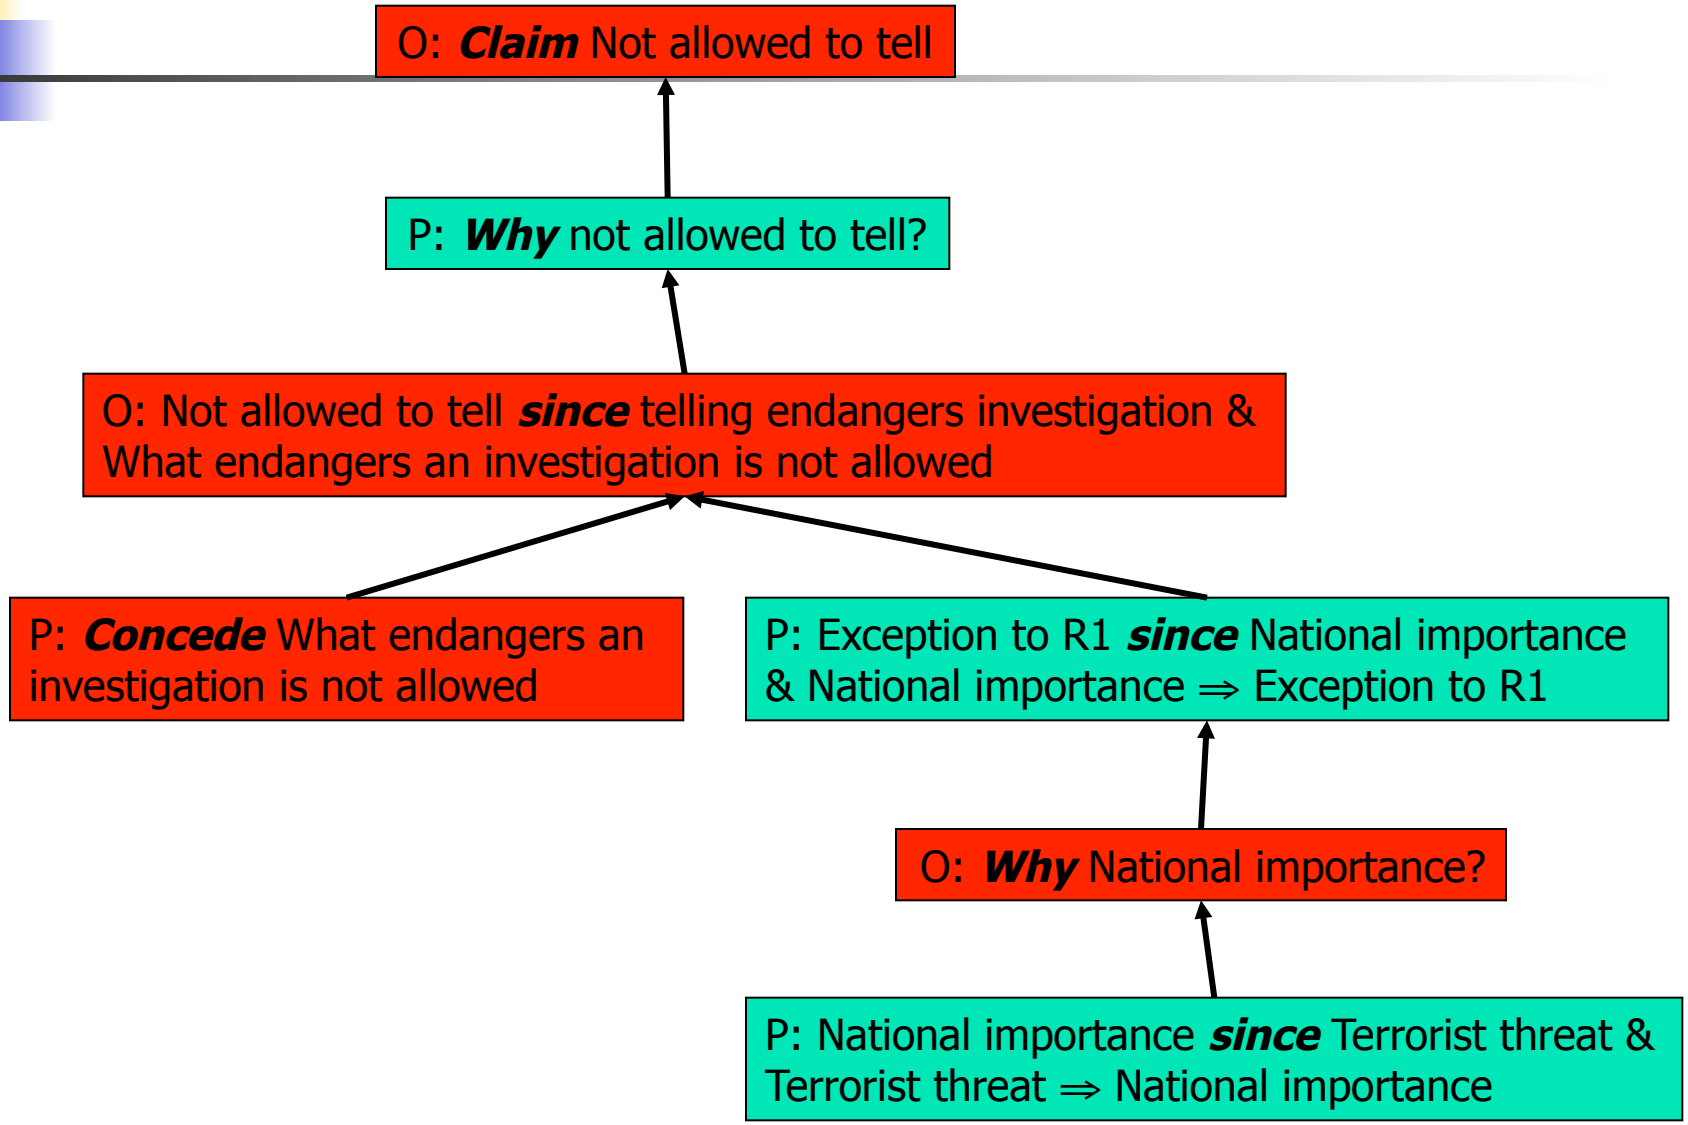

Persuasion part formalised

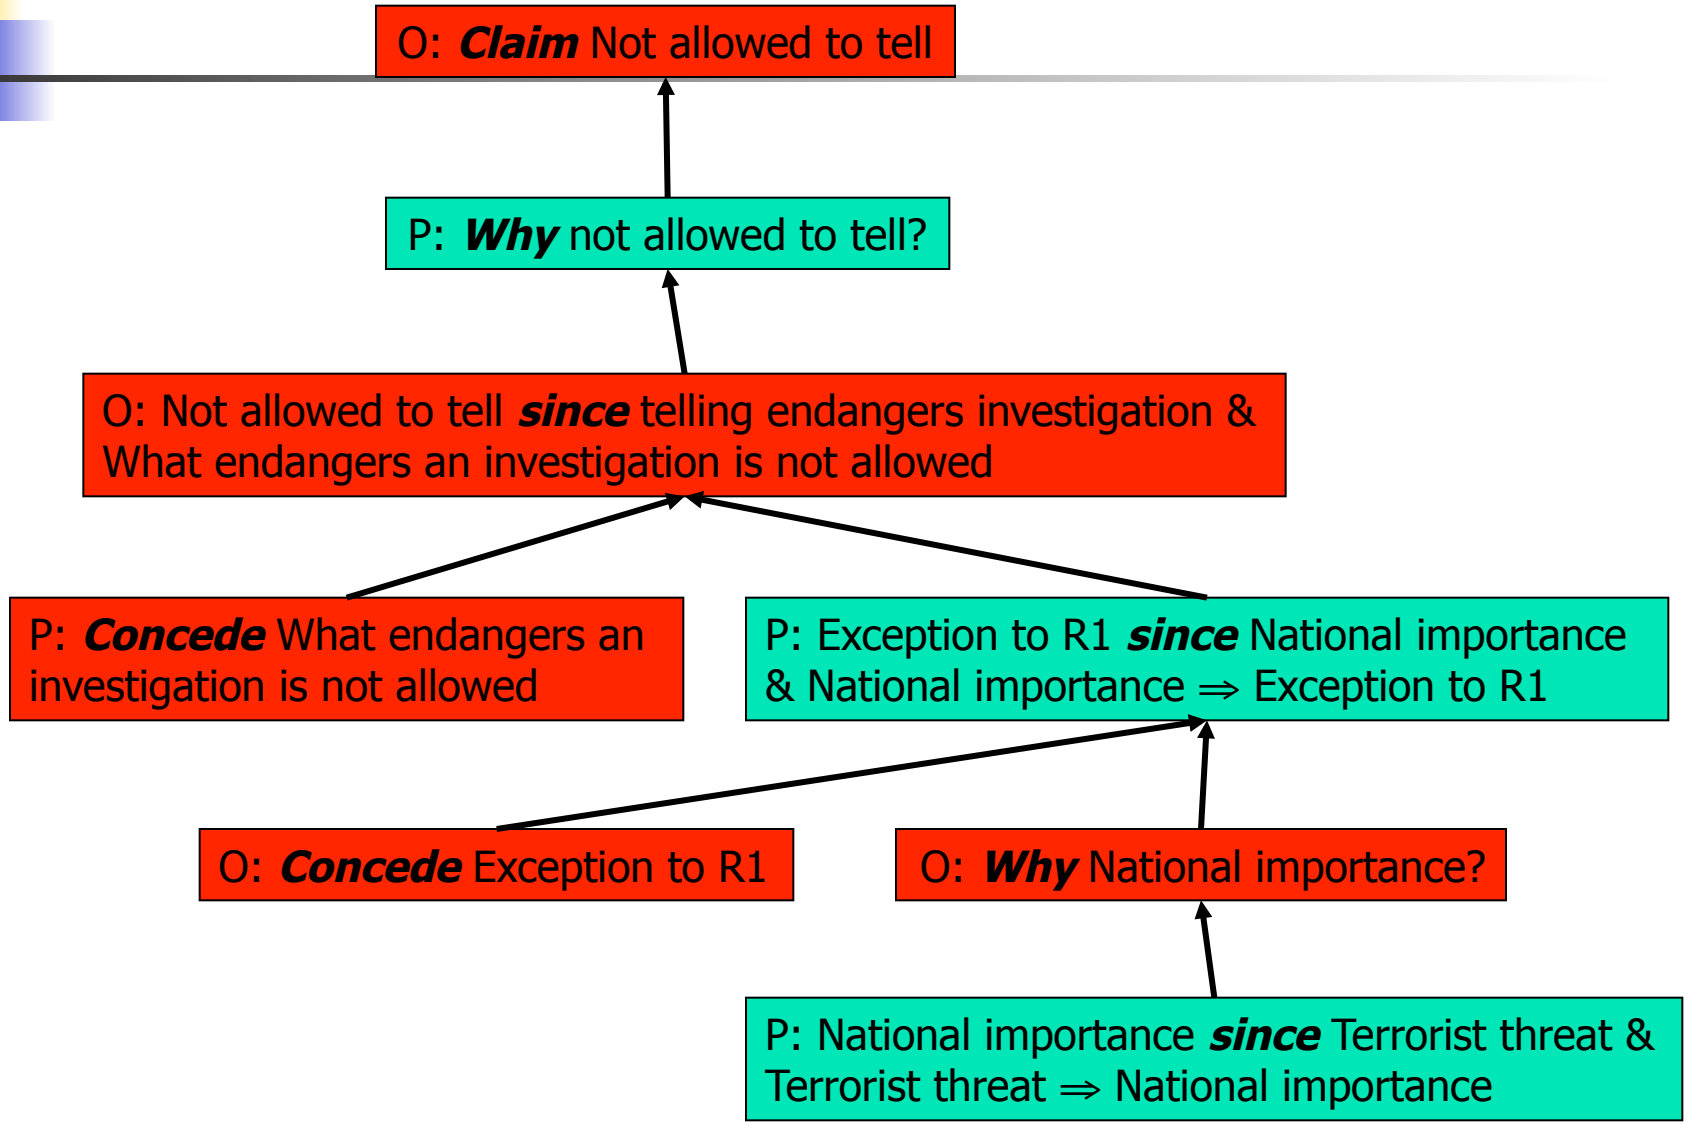

Persuasion part formalised

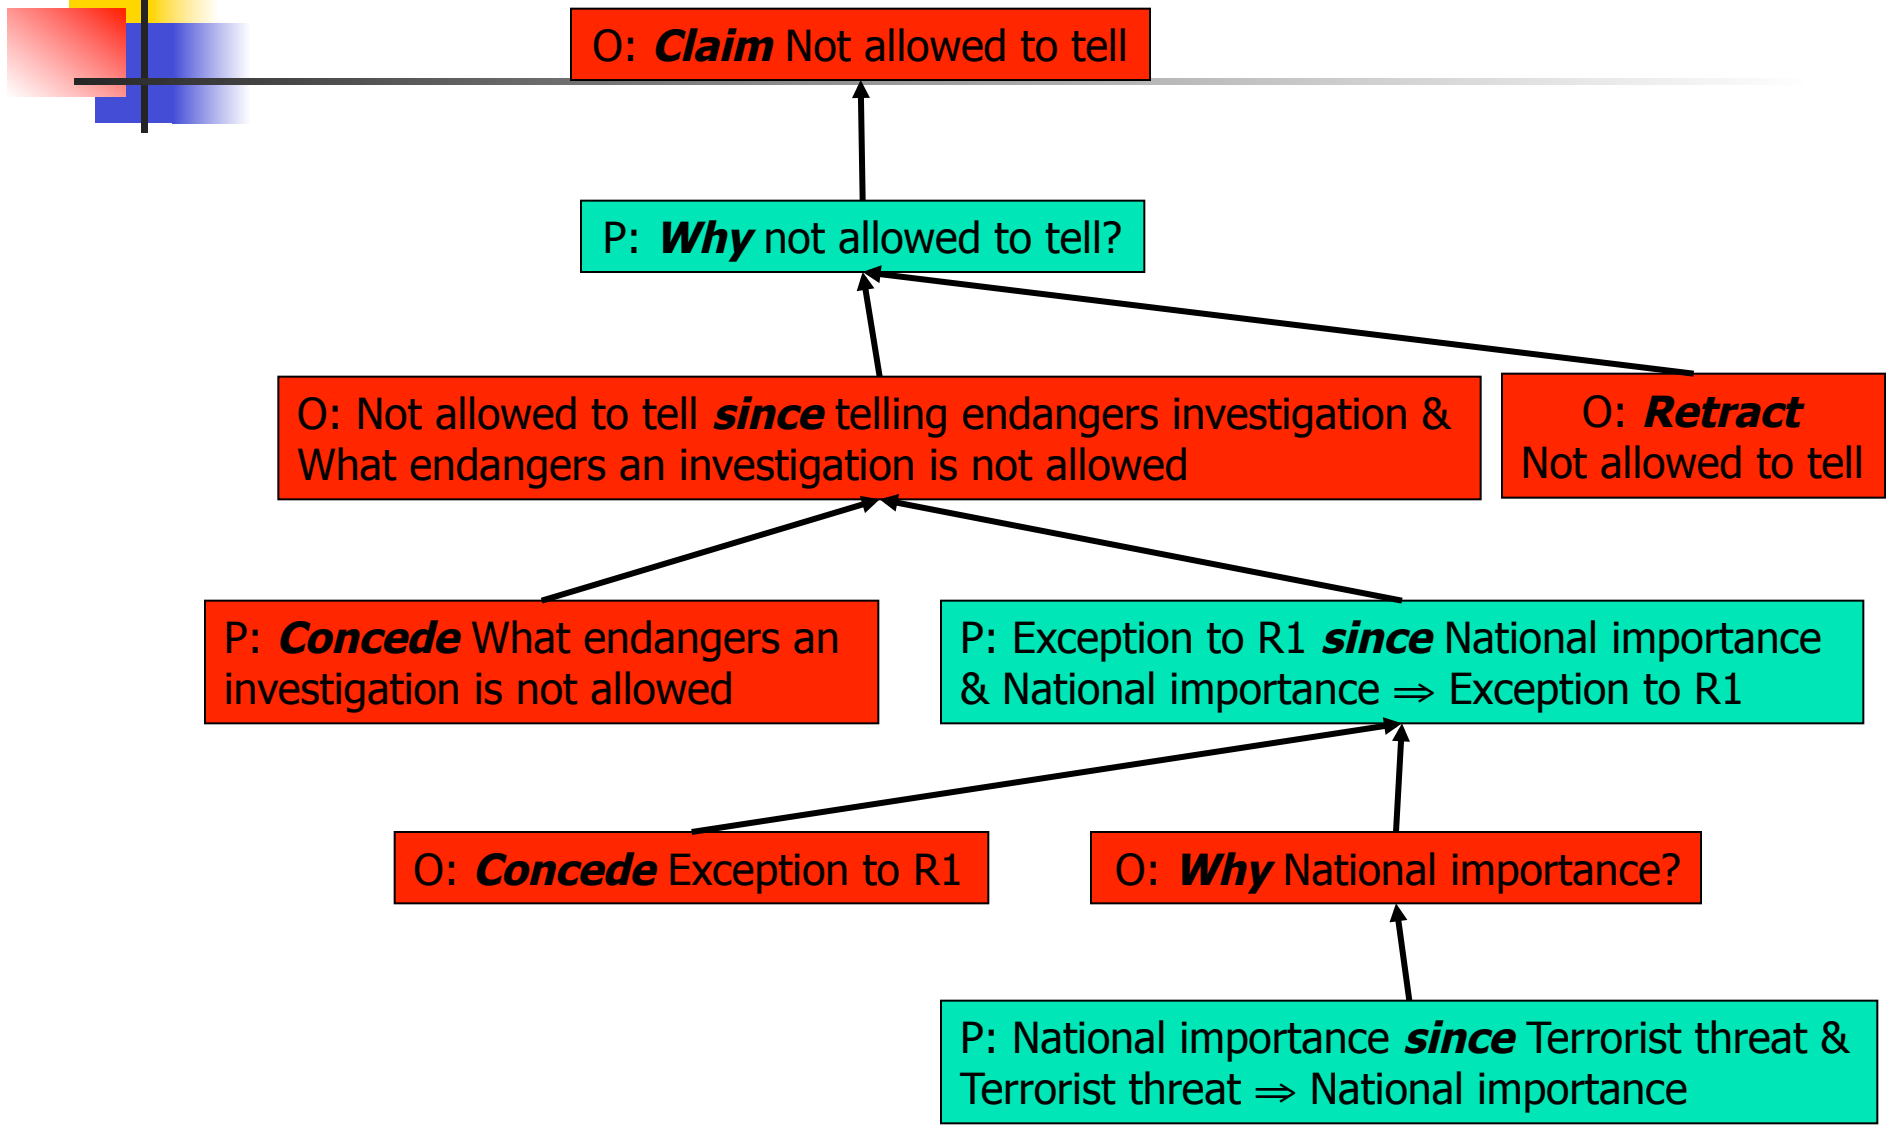

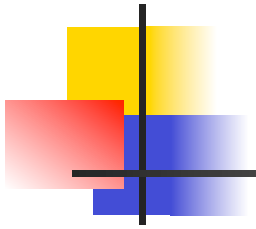

Conclusion

- Argumentation has two sides:
  - Inference
    - semantics
    - strict vs defeasible inferences
    - preferences
  - Dialogue
    - language + protocol
    - agent design
- Both sides can be formally and computationally modelled
  - But not in the same way
  - Metatheory of inference much more advanced than of dialogue

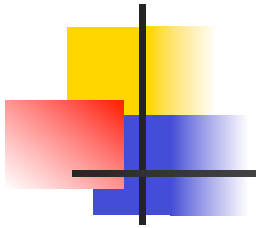

Reading (1)

- Collections
  - T.J.M. Bench-Capon & P.E. Dunne (eds.), *Artificial Intelligence* 171 (2007), Special issue on Argumentation in Artificial Intelligence
  - I. Rahwan & G.R. Simari (eds.), *Argumentation in Artificial Intelligence*. Berlin: Springer 2009.
  - A. Hunter (ed.), *Argument and Computation* 5 (2014), special issue on Tutorials on Structured Argumentation
- Abstract argumentation
  - P.M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77 (1995): 321-357
  - P. Baroni, M.W.A. Caminada & M. Giacomin. An introduction to argumentation semantics. *The Knowledge Engineering Review* 26: 365-410 (2011)

Reading (2)

- Classical and Tarskian argumentation
  - Ph. Besnard & A. Hunter, *Elements of Argumentation*. Cambridge, MA: MIT Press, 2008.
  - N Gorogiannis & A Hunter (2011) Instantiating abstract argumentation with classical logic arguments: postulates and properties, *Artificial Intelligence* 175: 1479-1497.
  - L. Amgoud & Ph. Besnard, Logical limits of abstract argumentation frameworks. *Journal of Applied Non-Classical Logics* 23(2013): 229-267.
- ASPIC+
  - H. Prakken, An abstract framework for argumentation with structured arguments. *Argument and Computation* 1 (2010): 93-124.
  - S. Modgil & H. Prakken, A general account of argumentation with preferences. *Artificial Intelligence* 195 (2013): 361-397

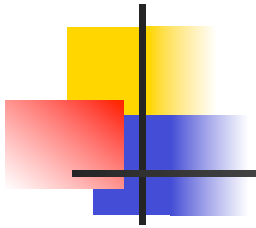

Reading (3)

- Assumption-based argumentation
  - A. Bondarenko, P.M. Dung, R.A. Kowalski & F. Toni, An abstract, argumentation-theoretic approach to default reasoning, *Artificial Intelligence* 93 (1997): 63-101.
  - P.M. Dung, P. Mancarella & F. Toni, Computing ideal sceptical argumentation, *Artificial Intelligence* 171 (2007): 642-674.
- Dialogue
  - S. Parsons, M. Wooldridge & L. Amgoud, Properties and complexity of some formal inter-agent dialogues. *Journal of Logic and Computation* 13 (2003): 347-376.
  - H. Prakken, Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation* 15 (2005): 1009-1040.
  - H. Prakken, Formal systems for persuasion dialogue. *The Knowledge Engineering Review* 21 (2006): 163-188.