An introduction to abstract argumentation

Pietro Baroni
DII - Dip. di Ingegneria dell’Informazione
University of Brescia (Italy)

Roadmap

- What's in the words?
- From the Big Bang to now
- Dung's framework
- A "conflict calculus": argumentation semantics
- "Basic" semantics properties
- A catalogue of semantics
- Properties of semantics (so far)
- Beyond Dung's framework
- (Advanced topics ?)
- Conclusions

Possible advanced topics

- Taking topology seriously
- Differences vs commonalities: semantics agreement
- Comparing argumentation frameworks
- Skepticism and related properties
- A richer notion of justification status
- Skepticism-related criteria for semantics
- Satisfying all criteria
- Computational issues

Roadmap

- What’s in the words ?

Abstract (for good)

- General, independent of specific details,able to capture a variety of situations
(MW: disassociated from any specific instance, theoretical)
- In computer science, abstraction is the process by which data and programs are defined with a representation similar in form to its meaning (semantics), while hiding away the implementation details. Abstraction tries to reduce and factor out details so that the programmer can focus on a few concepts at a time. A system can have several abstraction layers whereby different meanings and amounts of detail are exposed to the programmer (Wikipedia).

Abstract (for bad)

- Poorly related (or totally unrelated) with real problems
(MW: insufficiently factual, difficult to understand, abstruse)

Argumentation (for what?)

- Argumentation is a multi-faceted word, with a variety of informal/intuitive and also formal meanings
- The abstraction process detaches the word from some/most/all of its meanings and properties, keeping only those required by the desired abstraction level (and possibly adding other ones)

- Abstract arguments are not arguments
(or need not to be arguments)

Abstract argumentation in context

- Abstract argumentation is just one of the formal components that are needed to define a computational model of actual argumentation processes
- It is not sufficient for a comprehensive modeling
- … but can stand on its feet and be studied autonomously by its nature
- It is not even necessary for a comprehensive modeling (though its role should then be covered by some other component)

A generic model of argumentation processes

0. Looking for argument in natural sources
1. Constructing arguments (e.g. from a set of logical formulae, from a rule base, …)
2. Identifying important relations between arguments
3. Possibly exchanging arguments
4. Evaluating the status of the arguments
5. Evaluating the status of the conclusions of the arguments

The order of these activities is not rigid and there can be loops or jumps and skips

A generic model of argumentation processes

Argument mining

0. Looking for argument in natural sources

Structured argumentation

1. Constructing arguments
2. Identifying important relations between arguments

Dialogues

3. Possibly exchanging arguments

Abstract argumentation

4. Evaluating the status of the arguments

Structured argumentation

5. Evaluating the status of the conclusions of the arguments

Another view of the process

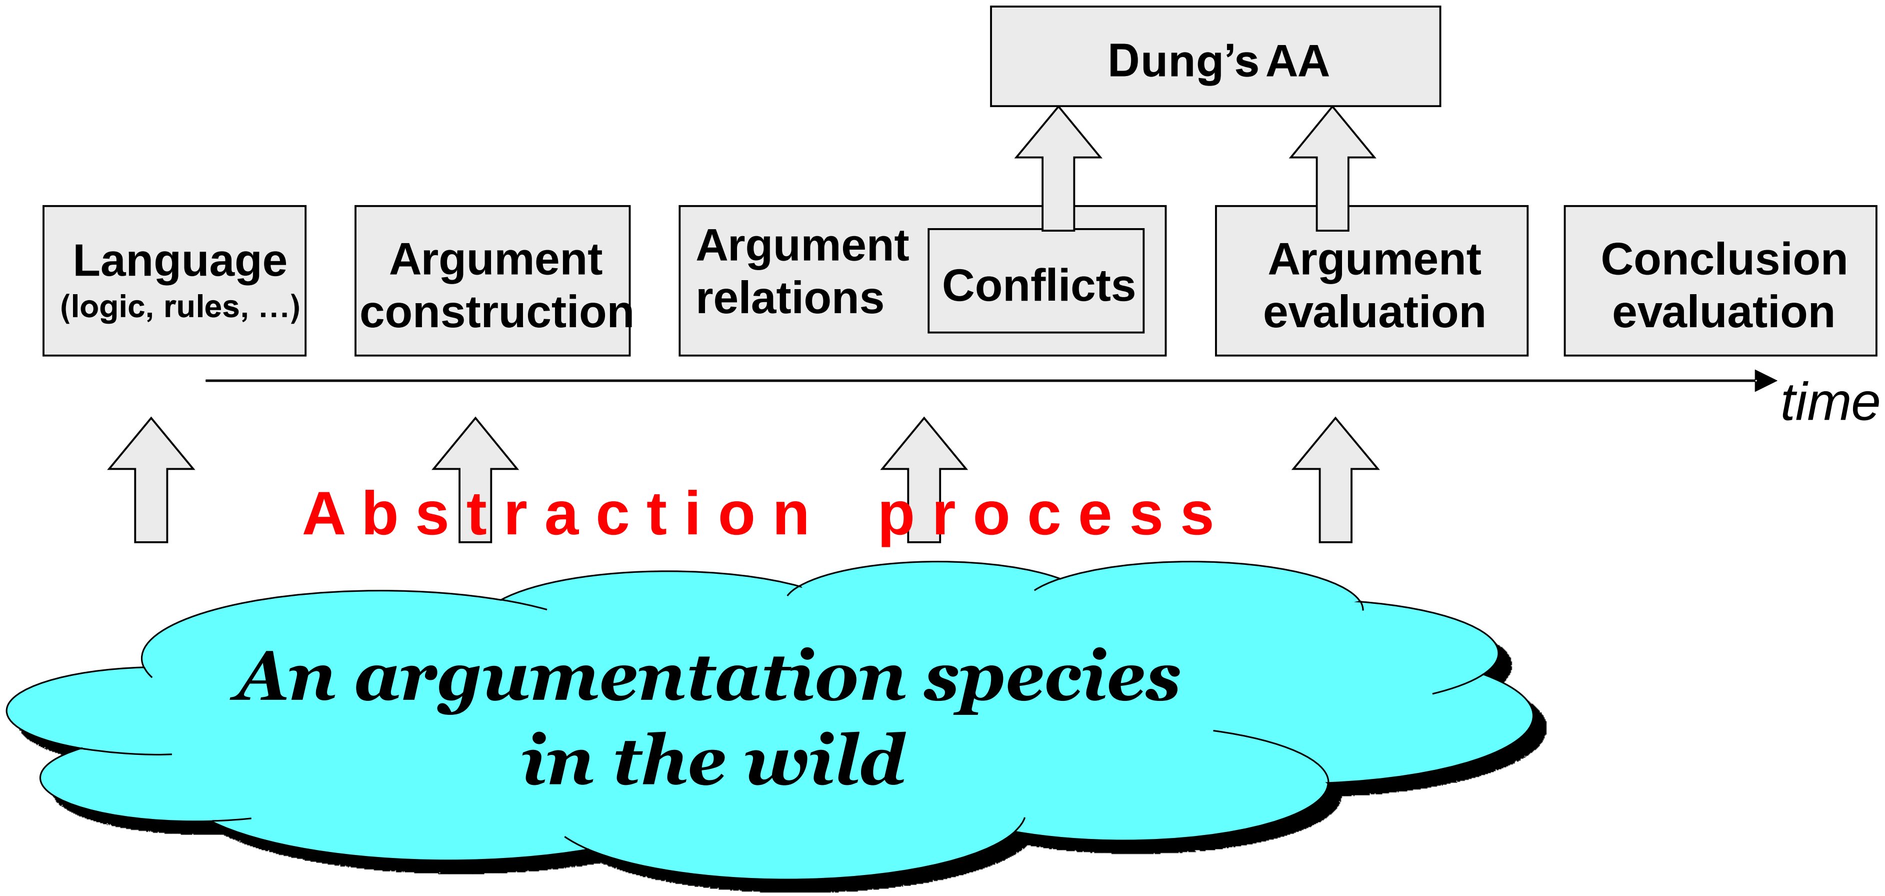

Roadmap

- What’s in the words?
- From the Big Bang to now

The Big Bang
About 25 years ago

- Dung's 1995 paper on Artificial Intelligence Journal
- Theory of Abstract Argumentation (AA) Frameworks
- Basic notions related to Abstract Argumentation Semantics (conflict freeness, admissibility, 4 "traditional" semantics: complete, grounded, stable, preferred)
- Ability to capture a variety of other (less abstract) formal contexts:
  - » default logic
  - » logic programming with negation as failure
  - » defeasible reasoning
  - » N-person games
  - » stable-marriage problem

The explosion

- Not an easy to read paper, but a very general (and basically simple) formalism with ideas and results regarded as extremely powerful
- Huge impact on the literature (1818 citations WoS, 2657 Scopus, 4450 Google Scholar)
- Probably the most cited paper in the computational argumentation literature
- Originated entire new research lines: many followers (and some criticisms)

The AA universe

Abstract argumentation framework

"Traditional" semantics notions

Extensions/variations of the definition of the basic framework (and relevant semantics)

"New" semantics
Formalisms and properties for old and new semantics

A note on notation

- There is no universally adopted notation even for the most basic concepts
- The slides include excerpts taken directly from the original papers and reflect all these differences
- Should not be too mysterious …
- … some notes along the way
- for any doubt, please ask

Roadmap

- What's in the words?
- From the Big Bang to now
- Dung's framework

Dung's framework is (almost) nothing

Definition 2. An *argumentation framework* is a pair

$$AF = \langle AR, attacks \rangle$$

where AR is a set of arguments, and *attacks* is a binary relation on AR, i.e. $attacks \subseteq AR \times AR$.

- A directed graph (called *defeat graph*) where:
  - » arcs are interpreted as attacks
  - » nodes are called arguments "by chance" (let say historical reasons)

Here, an argument is an abstract entity whose role is solely determined by its relations to other arguments. No special attention is paid to the internal structure of the arguments.

Dung's framework is (almost) nothing

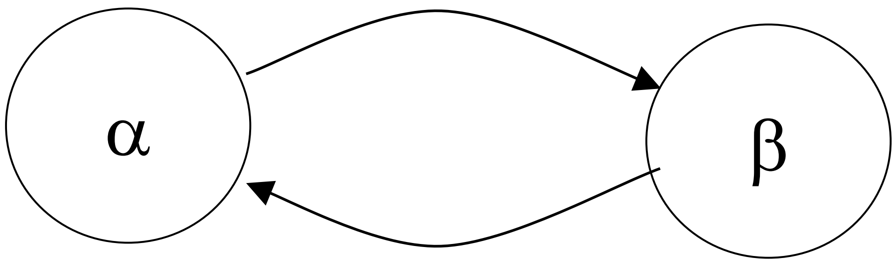

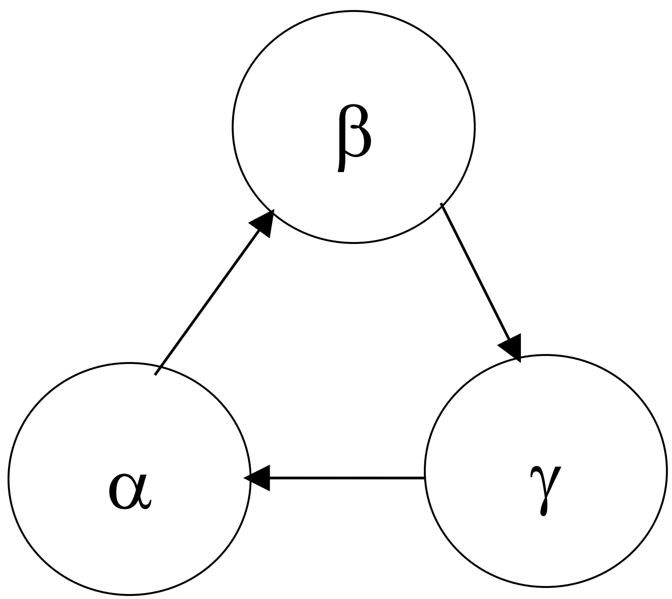

Dung's framework is (almost) everything

- Arguments are simply "conflictables"

- Conflicts are everywhere
- Conflict management is a fundamental need with potential spectacular/miserable failures both in real life and in formal contexts (e.g. in classical logic)
- A general abstract framework centered on conflicts has a wide range of potential applications

Dung's framework is (almost) everything

- The pervasiveness of Dung's framework and semantics is witnessed by the correspondences drawn in the original paper with a variety of other formal contexts
- Many extensions and variations of Dung's framework allow a translation procedure back to the original framework to exploit its basic features

A conflict calculus: abstract argumentation semantics

- A way to identify sets of arguments “surviving the conflict together” given the conflict relation only
- In general, several choices of sets of “surviving arguments” are possible
- Two main styles for semantics definition: extension-based and labelling-based
- These points will be discussed extensively after some “user instructions”

AA user manual

1. Identify an interesting application domain where conflict management plays a key role
2. Define a suitable formalization of problem instances in the selected domain
3. Define the notions of argument and attack in your formalization, i.e. a translation/abstraction method from problem instances to argumentation frameworks
4. Play at your will with "conflict calculus" at abstract level
5. Map back the results of "conflict calculus" (extensions or labellings) into entities at the problem level
6. Are they meaningful? Do they provide useful/original perspectives? Did you avoid to reinvent the wheel?

An example: the stable marriage problem

- Steps 1 and 2

Given two sets M and W of n men and n women respectively. The stable marriage problem (SMP) is the problem of finding a way to arrange the marriage for the men and women in M and W, where it is assumed that all the men and women in M and W have expressed mutual preference (each man must say how he feels about each woman and vice versa).[9] The marriages have to be stable in the sense that, if for example A is married to B, then all those whom A prefers to B must be married to someone whom they prefer to A. Formally, a solution to the SMP is a one–one correspondence $S: M \rightarrow W$ such that there exists no pair $(m, w) \in M \times W$ such that m prefers w to $S(m)$ and w prefers m to $S^{-1}(w)$.

An example: the stable marriage problem

- Step 3

$$AR = M \times W\ ,$$

$$attacks \subseteq AR \times AR:$$

$$(C, D) \text{ attacks } (A, B) \quad \text{iff} \quad \begin{array}{l} (1)\ A = C \text{ and } A \text{ prefers } D \text{ to } B\text{, or} \\ (2)\ D = B \text{ and } B \text{ prefers } C \text{ to } A\ . \end{array}$$

An example: the stable marriage problem

- Steps 4 and 5

Theorem 39. *A set $S \subseteq AR$ constitutes a solution to the SMP iff S is a stable extension of the corresponding argumentation framework.*

An example: the stable marriage problem

- Step 6

To demonstrate once more that there are practically relevant argumentation systems which have no stable semantics, in the following we introduce the Stable Marriage Problem with Gays (SMPG) which is a modification of the SMP in which individuals of the same sex can be married to each other. The condition for the stability of a marriage is defined as in the SMP. The problem now is finding a way to arrange the marriage for a maximal number of persons. In contrast to the SMP, the SMPG corresponds to the problem of finding a preferred extension in an argumentation framework $AF = (AR, attacks)$ with $AR = P \times P$ where P is the set of persons involved and *attacks* is defined as in the SMP. The following example shows that in general, the argumentation framework corresponding to an SMPG has no stable semantics.

An example: the stable marriage problem

- Step 6

Let $P = \{m, w, p_1, p_2, p_3\}$ where m is a man, w is a woman. For short we say that x loves y if x prefers y to all others. Suppose that m and w are in love with each other. Further suppose that there is a love triangle between p_1, p_2 and p_3 as follows: p_1 loves p_2, p_2 loves p_3 and p_3 loves p_1. So it is not difficult to see that there is no way to arrange a stable marriage for any among p_1, p_2 and p_3. The only stable marriage is between m and w. Indeed, the corresponding argumentation framework has exactly one preferred extension containing only the pair (m, w).

- The notion of preferred extension provides a natural solution to a non traditional version of the marriage problem

Roadmap

- What's in the words?
- From the Big Bang to now
- Dung's framework
- A "conflict calculus": argumentation semantics

Abstract argumentation semantics

- A way to identify sets of arguments “surviving the conflict together” given the conflict relation only
- Two main styles for semantics definition: extension-based and labelling-based
- In general, several choices of sets of “surviving arguments” are possible (multiple-status semantics) but some semantics prescribe exactly one extension/labelling (single status semantics)

Extension-based semantics

- A set of extensions is identified
- Each extension is a set of arguments which can "survive together" or are "collectively acceptable" i.e. represent a reasonable viewpoint
- The justification status of each argument can be defined on the basis of its extension membership
  - » skeptical justification = membership in all extensions
  - » credulous justification = membership in one extension

Sets of extensions

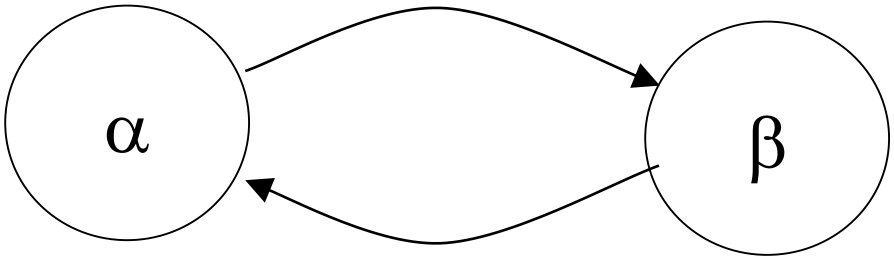

$E_1 = \{\{\alpha\},\{\beta\}\}$
$E_2 = \{\varnothing\}$

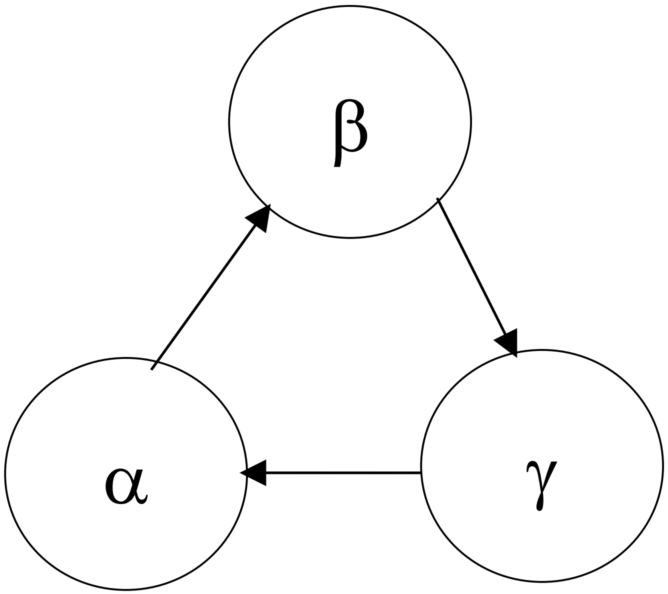

$E_1 = \{\{\alpha\},\{\beta\},\{\gamma\}\}$
$E_2 = \{\varnothing\}$

Labelling-based semantics

- A set of labels is defined (e.g. IN, OUT, UNDECIDED) and criteria for assigning labels to arguments are given
- Several alternative labellings are possible
- The justification status of each argument can be defined on the basis of its labels

Labelling-based semantics

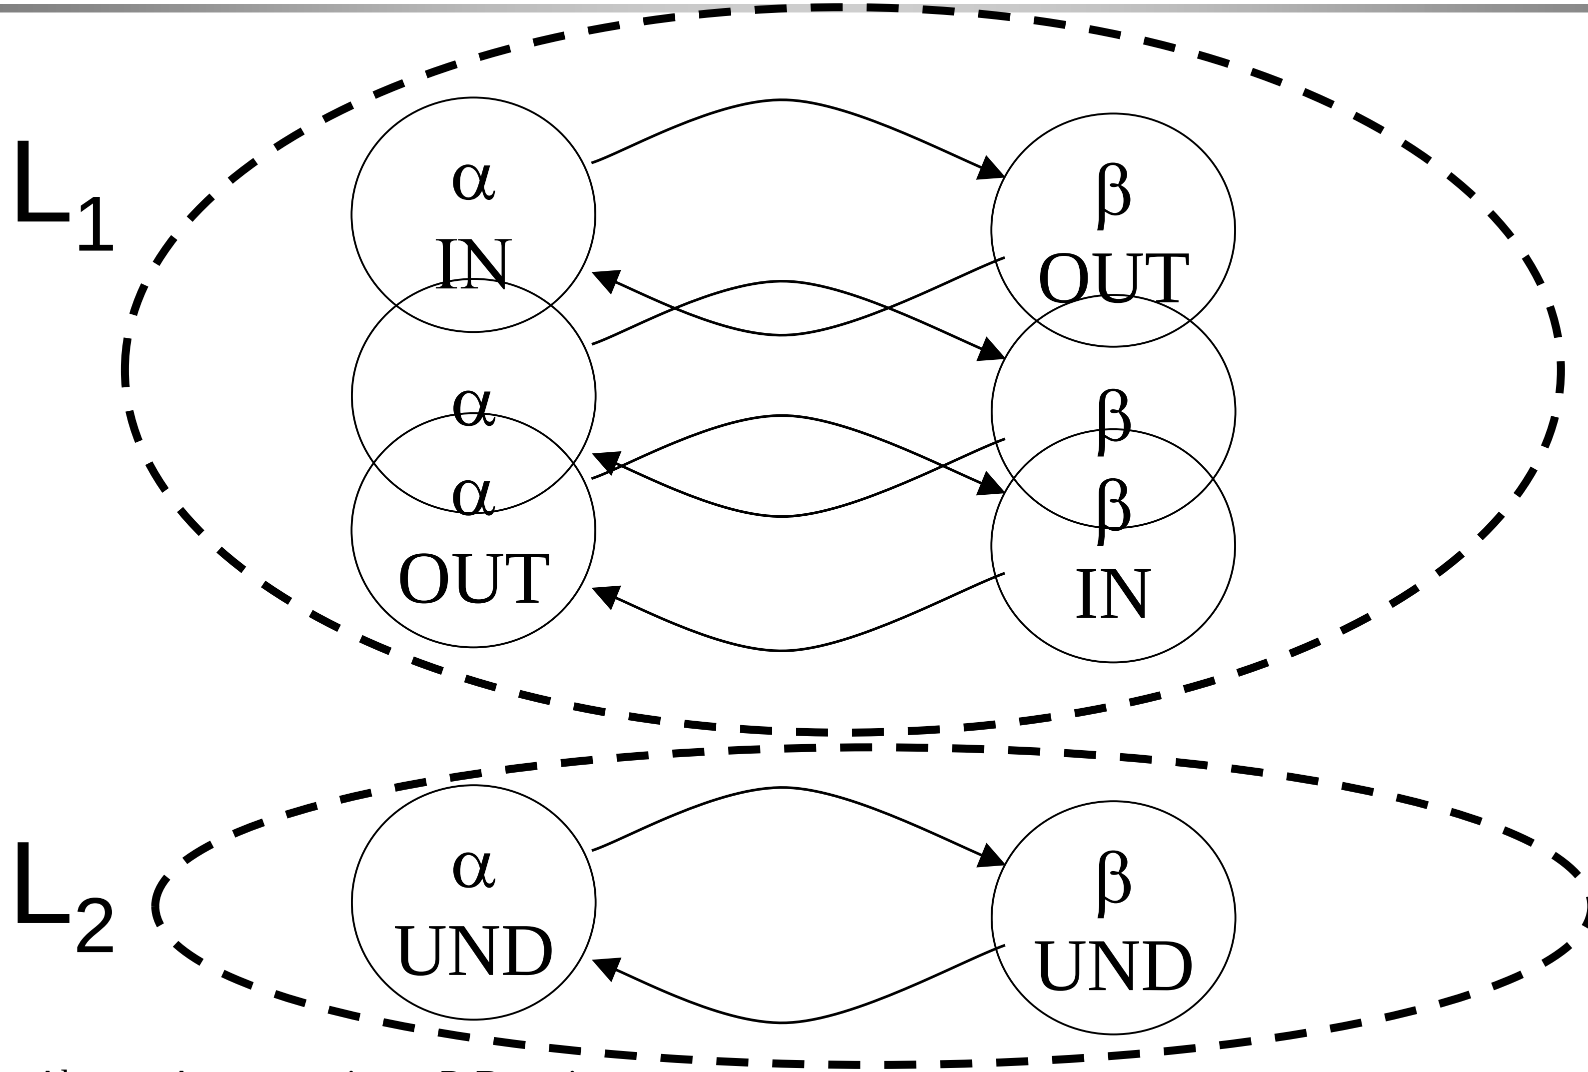

Labelling-based semantics

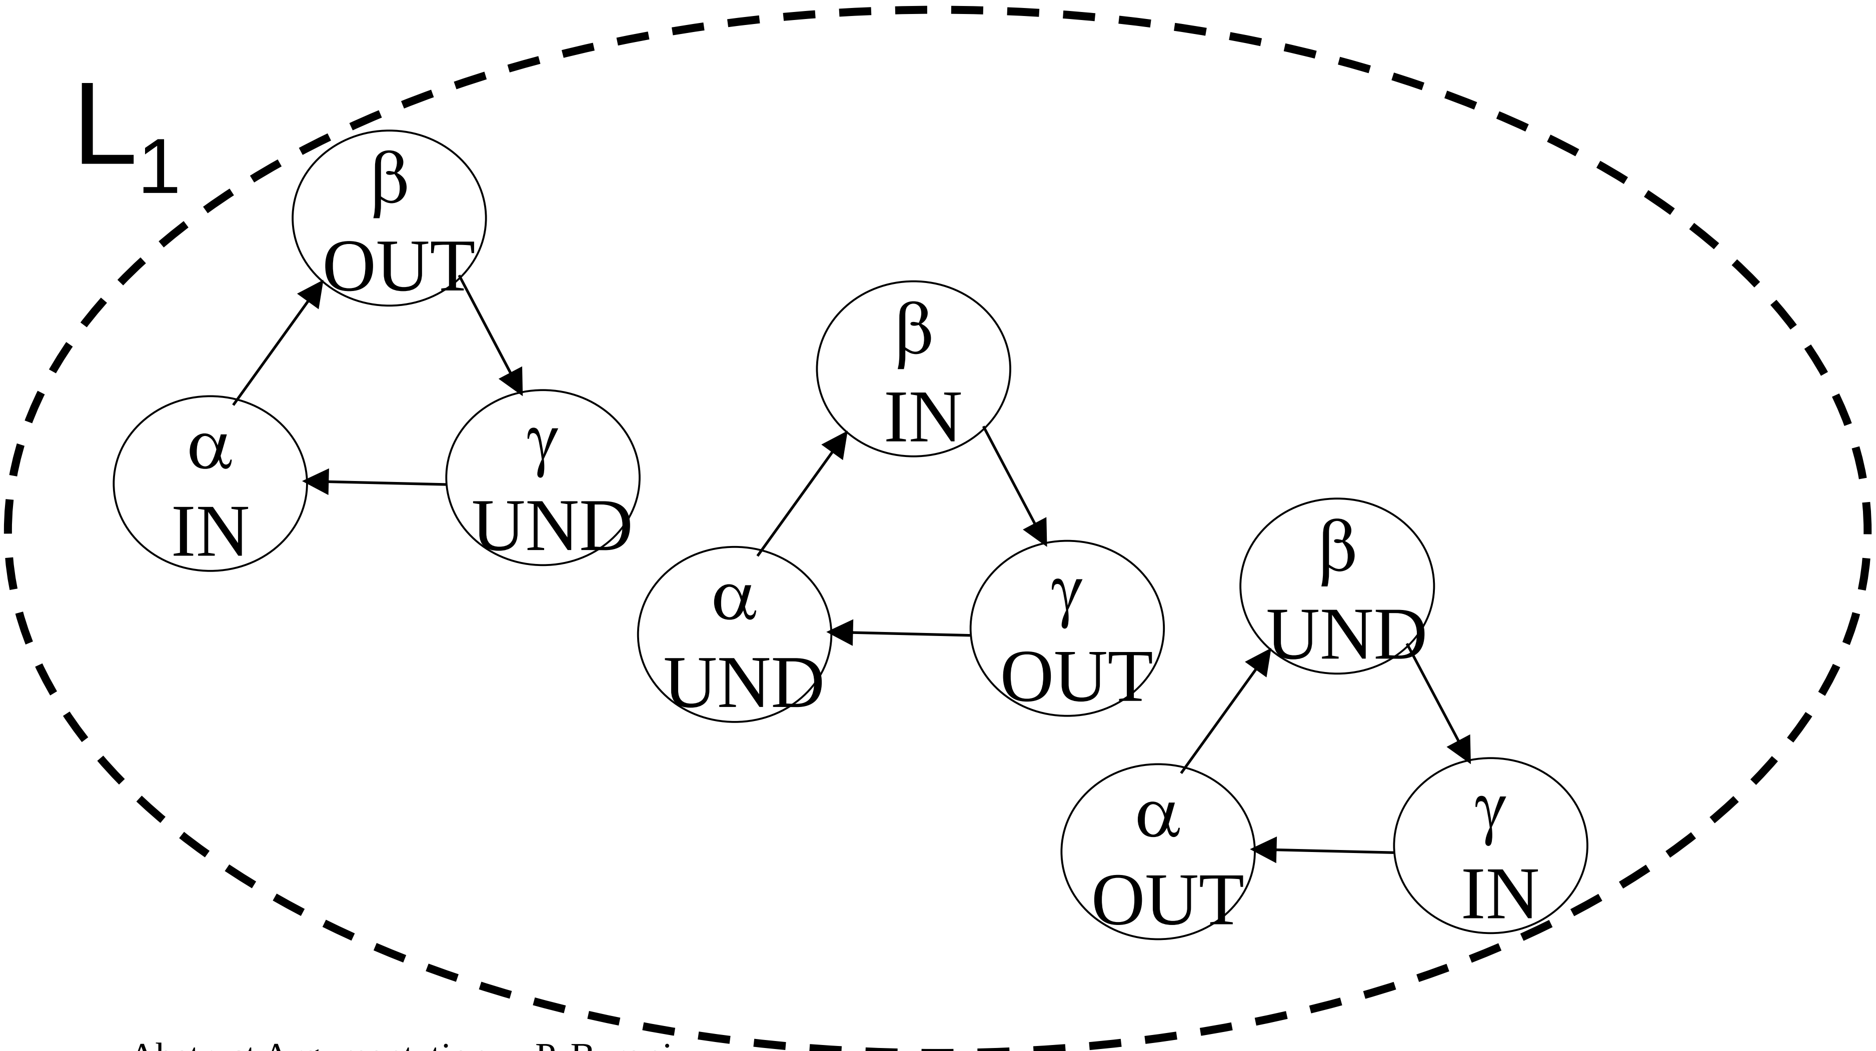

Labelling-based semantics

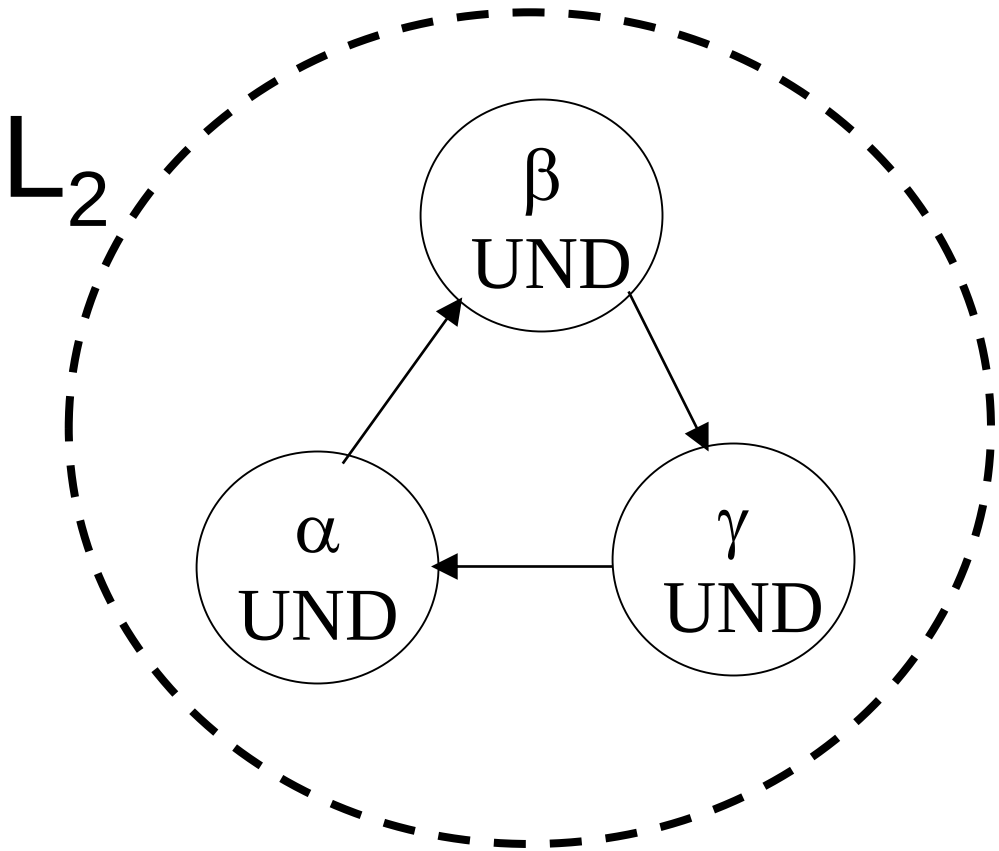

Labellings vs. extensions

- Labellings based on {IN, OUT, UNDEC} and extensions can be put in direct correspondence

- Given a labelling L, LabToExt(L) = in(L)

- Given an extension E, a labelling L=ExtToLab(E) can be defined as follows:
  in(L)=E
  out(L)=attacked(E)
  undec(L)=all other arguments

Labellings vs. extensions

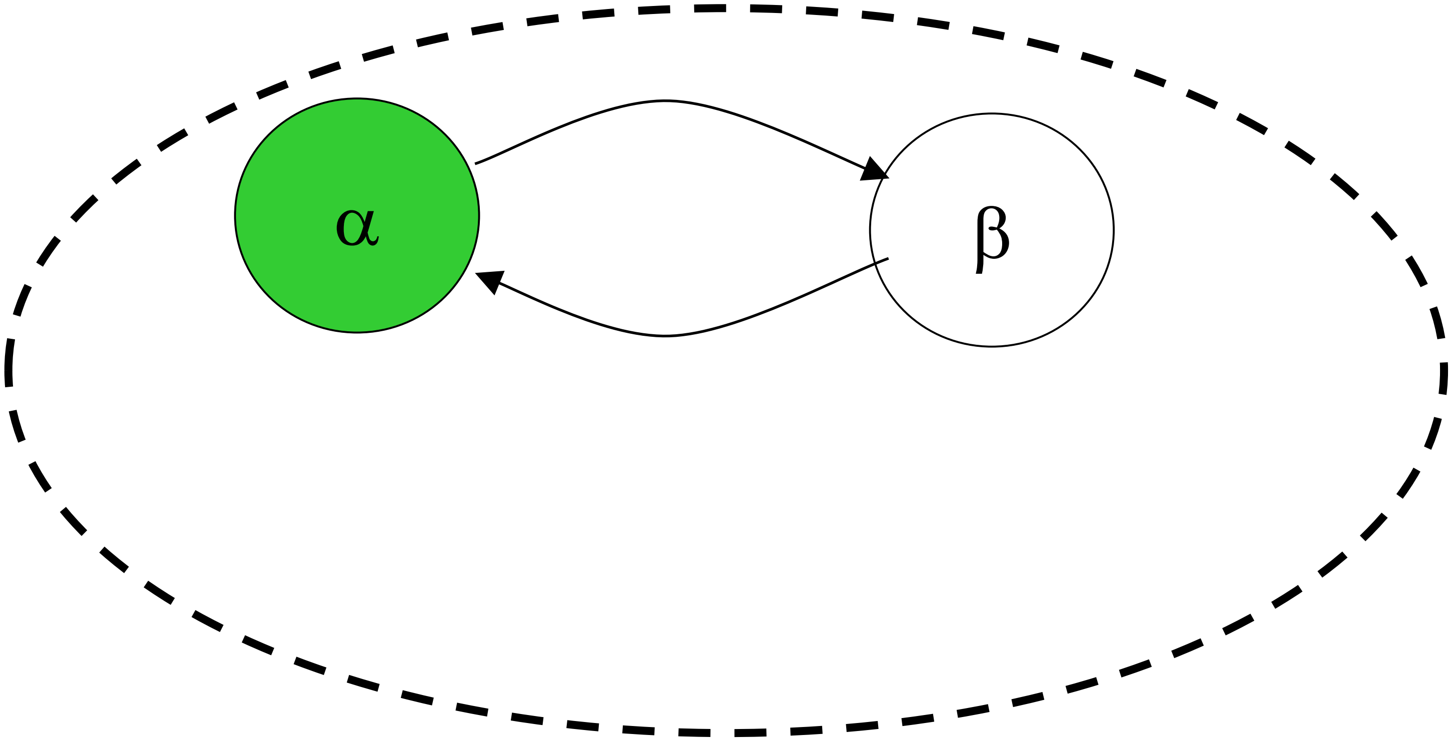

Labellings vs. extensions

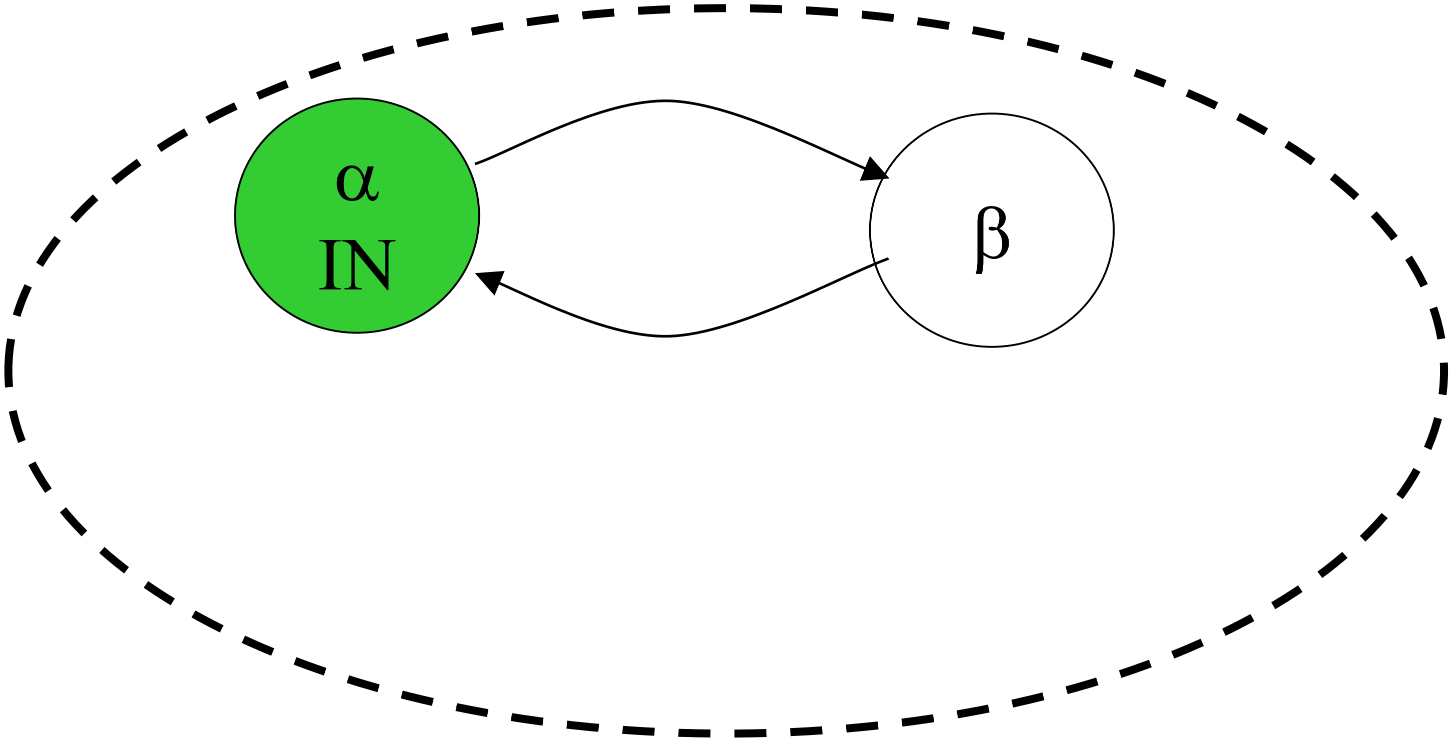

Labellings vs. extensions

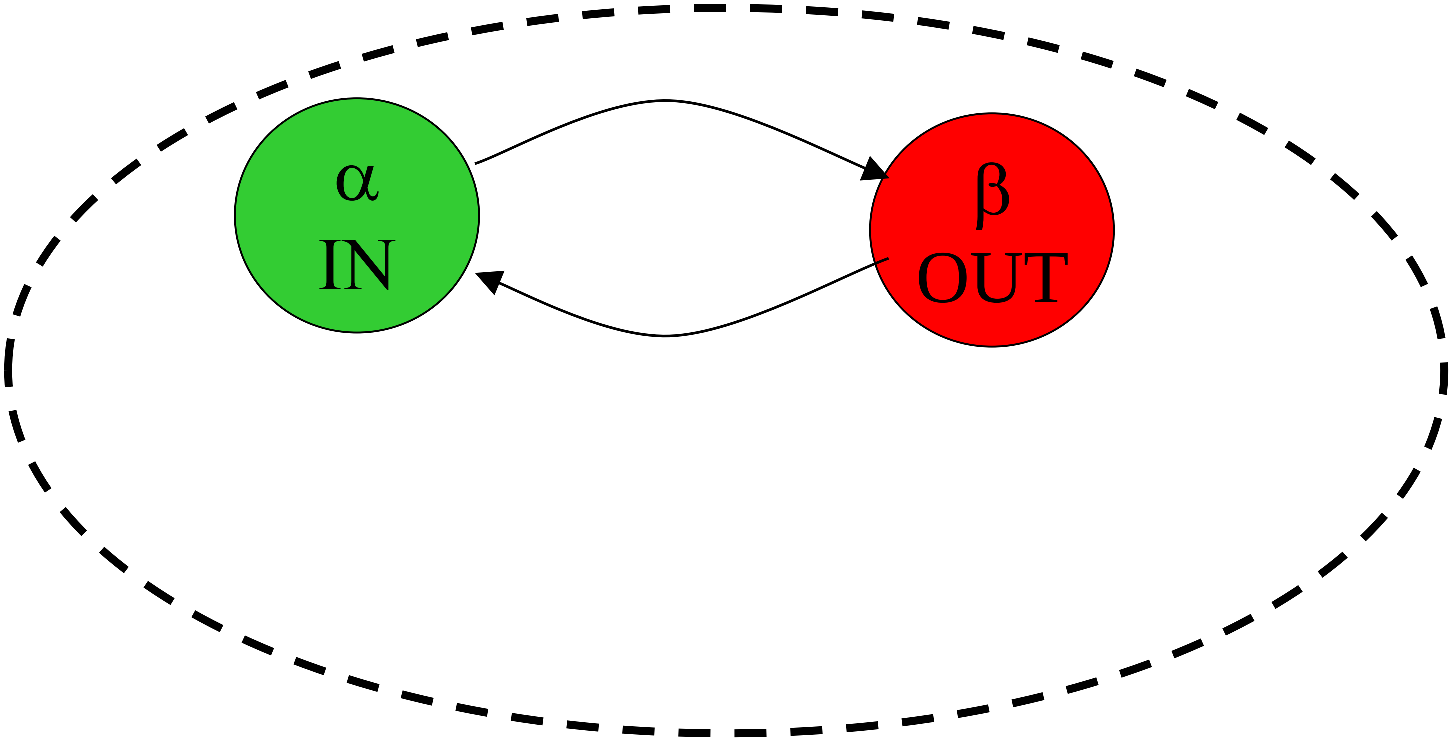

Labellings vs. extensions

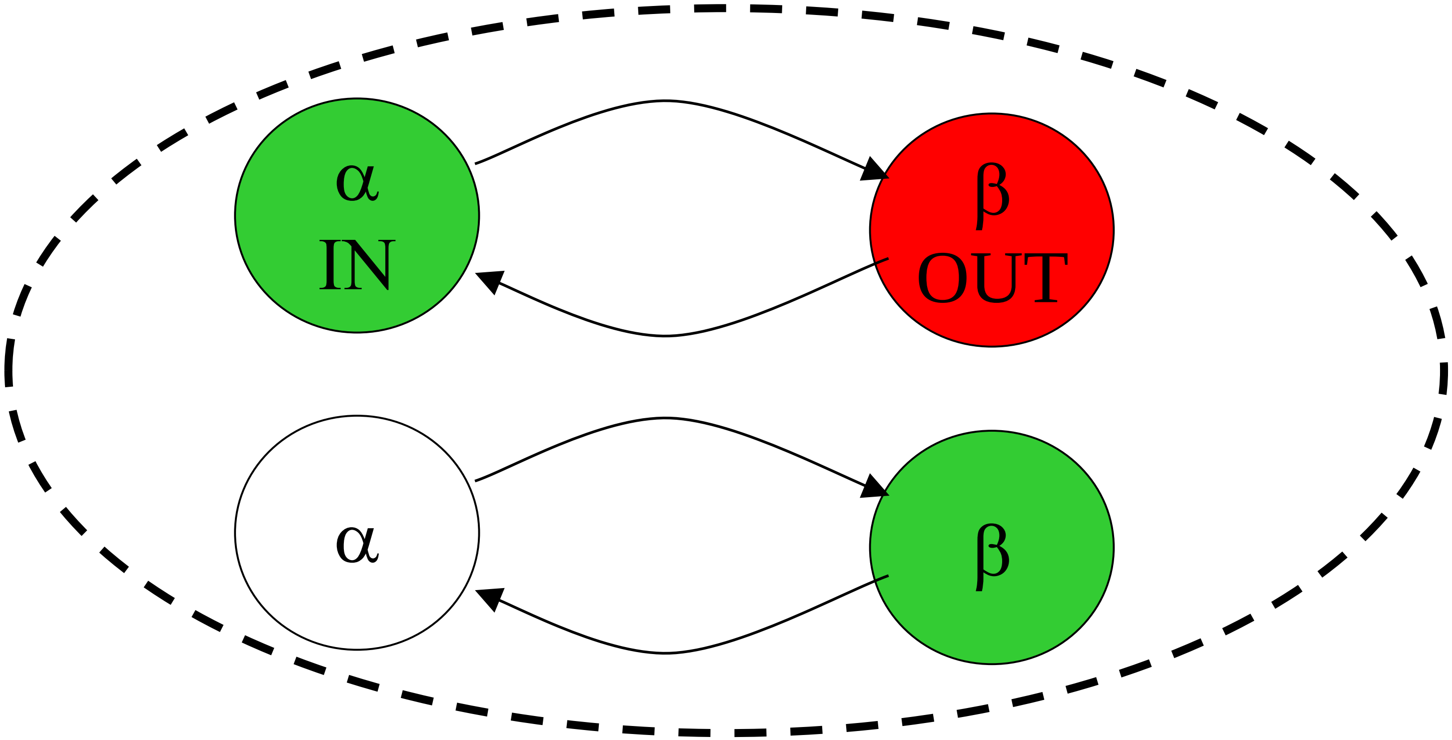

Labellings vs. extensions

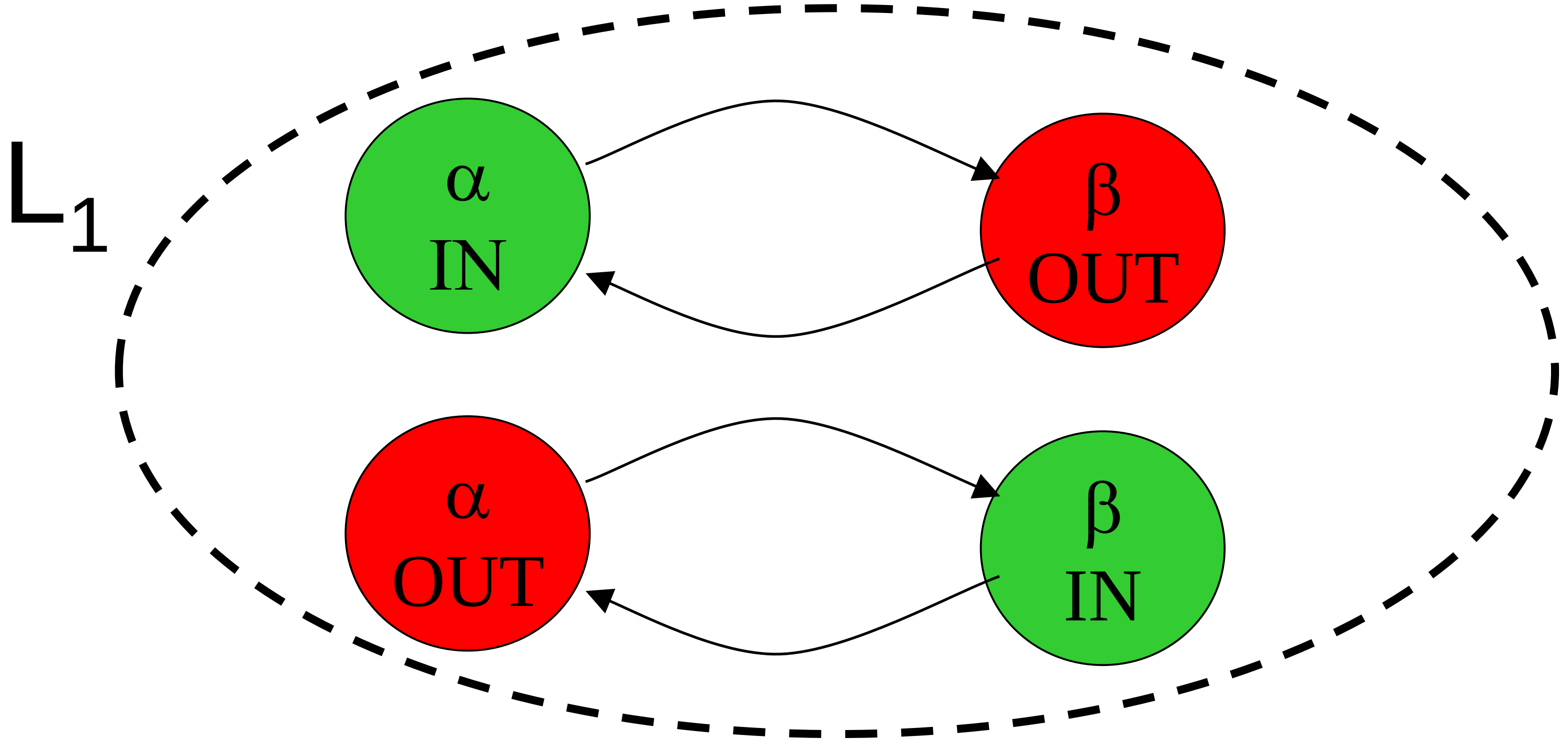

Labellings vs. extensions

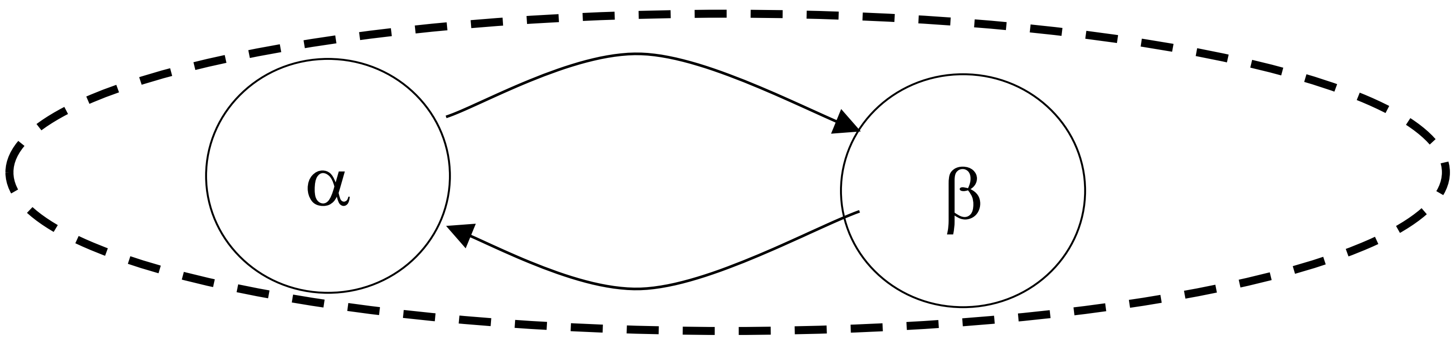

Labellings vs. extensions

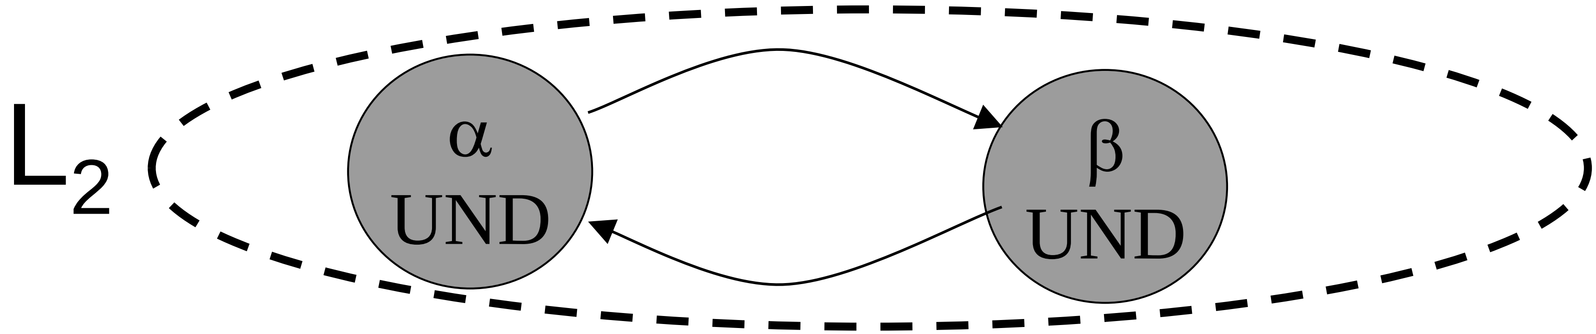

Labellings vs. extensions

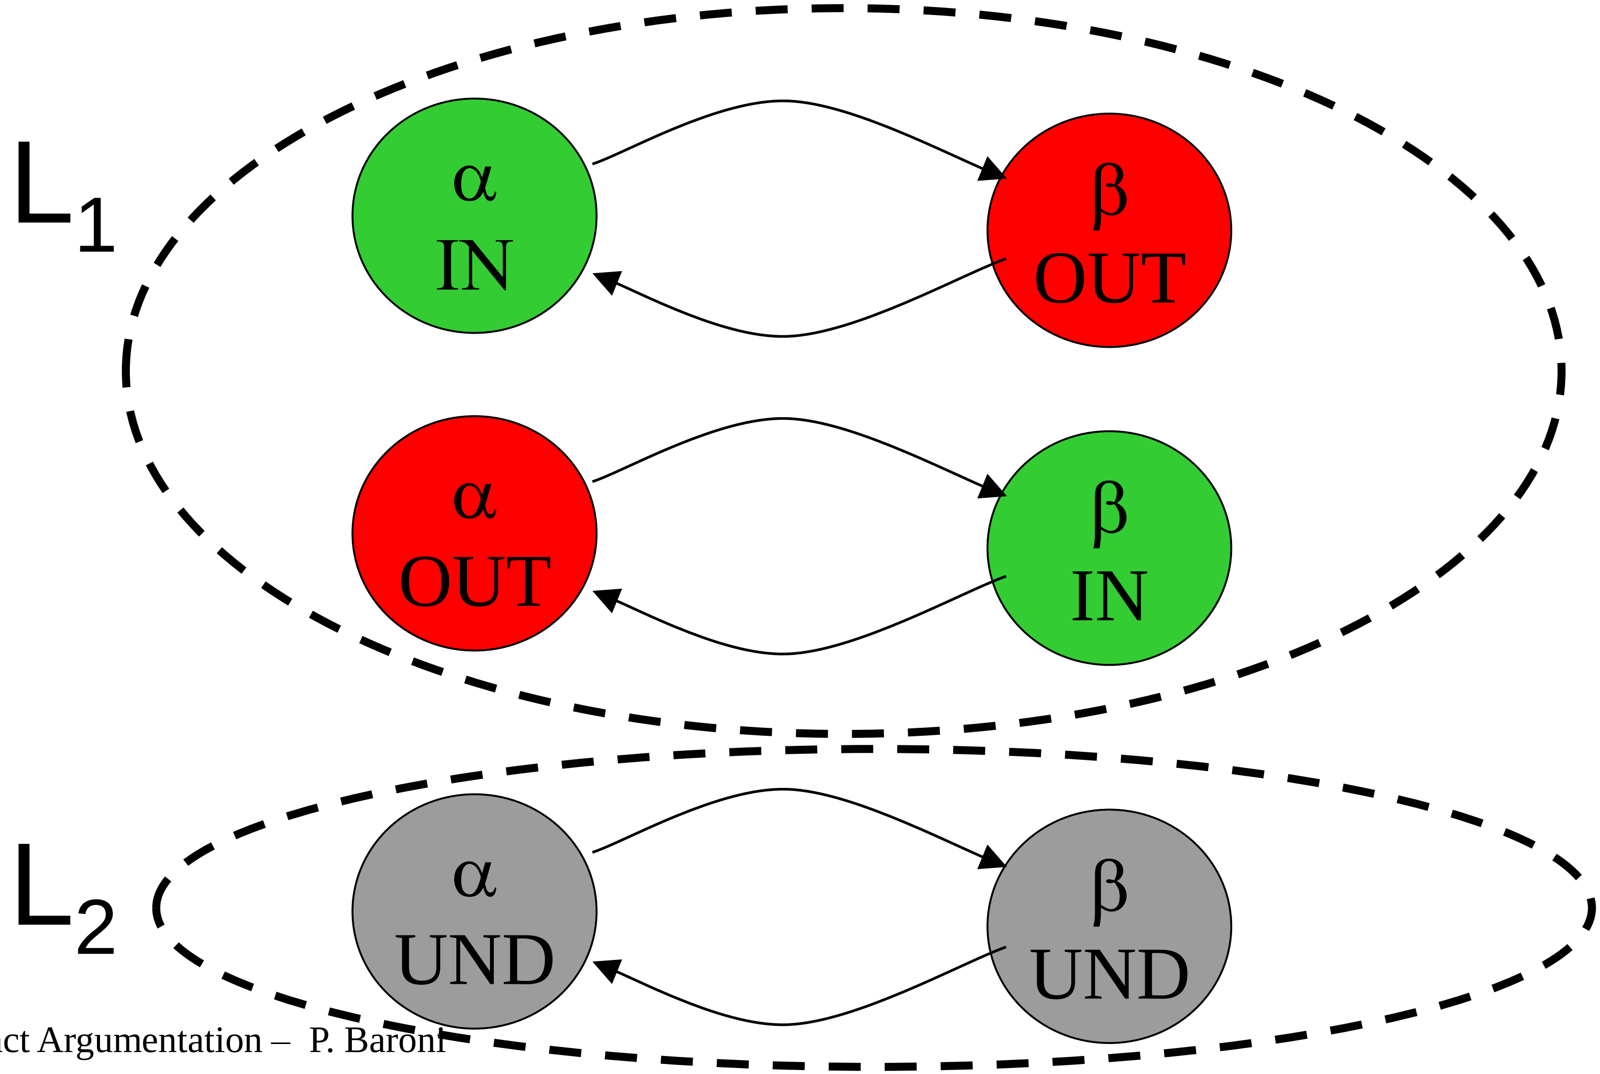

Labellings vs. extensions

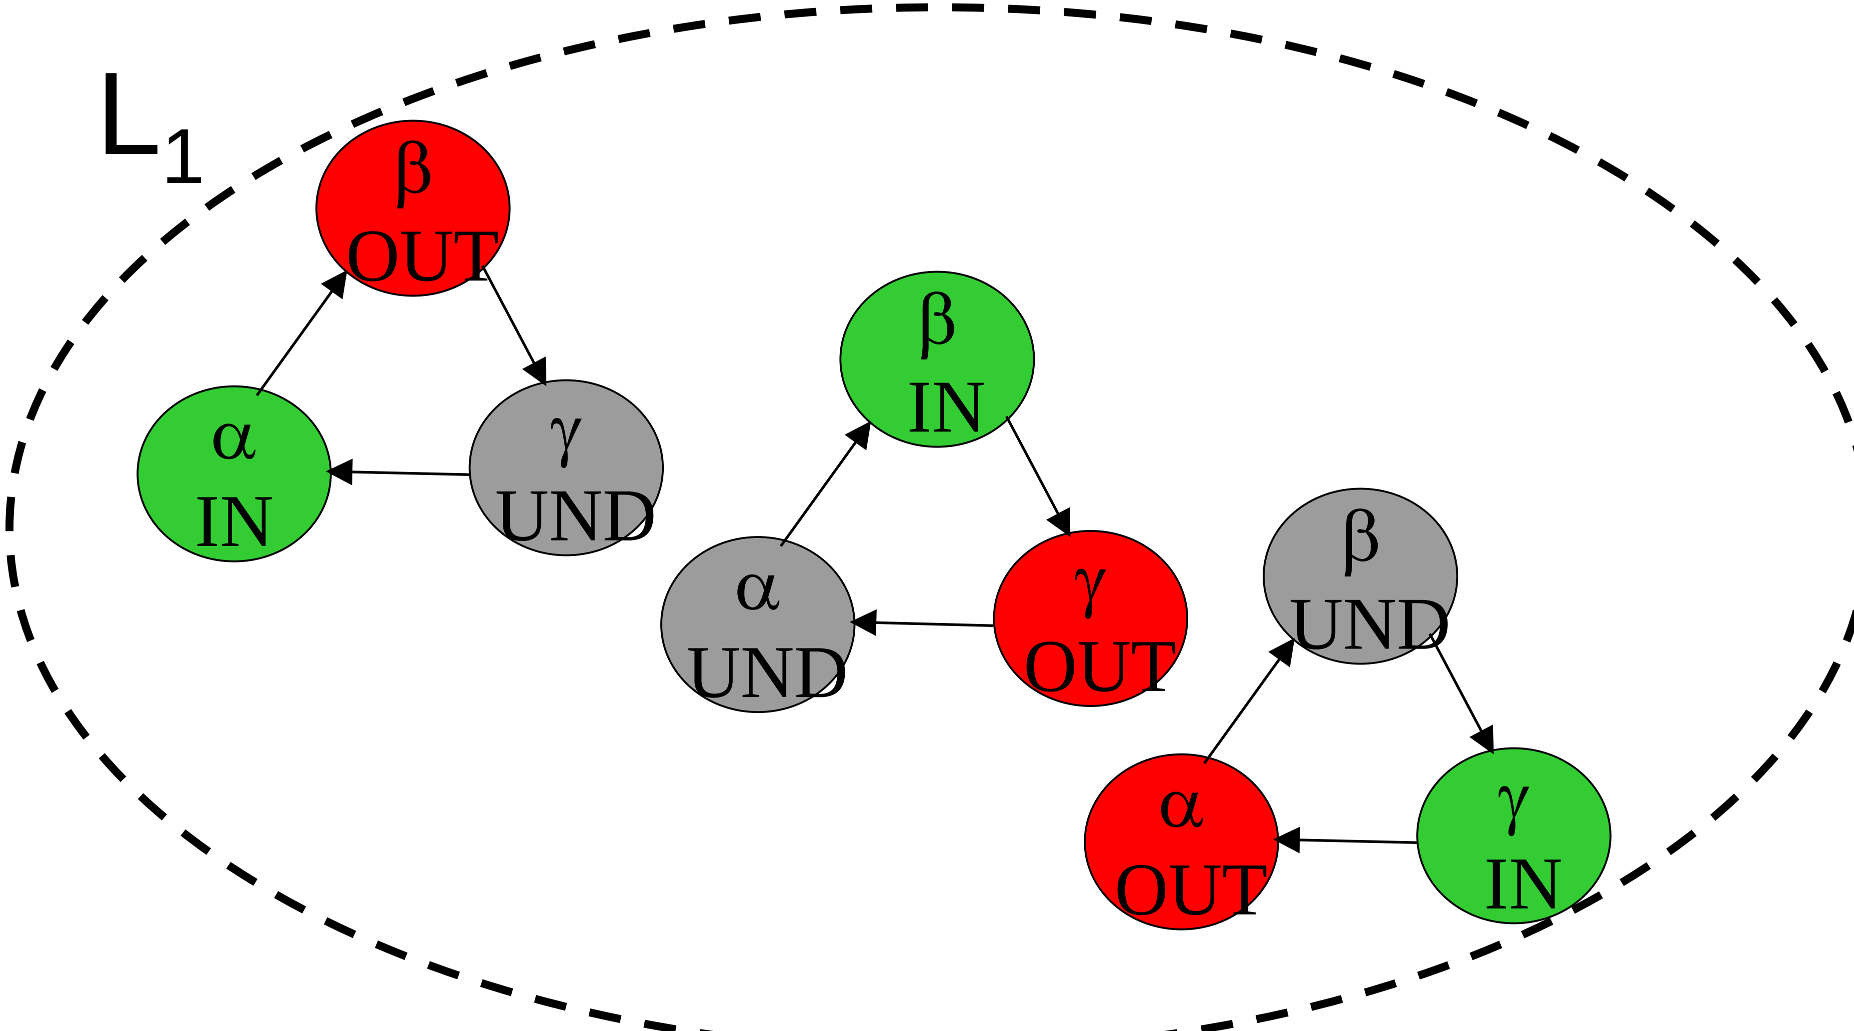

Labellings vs. extensions

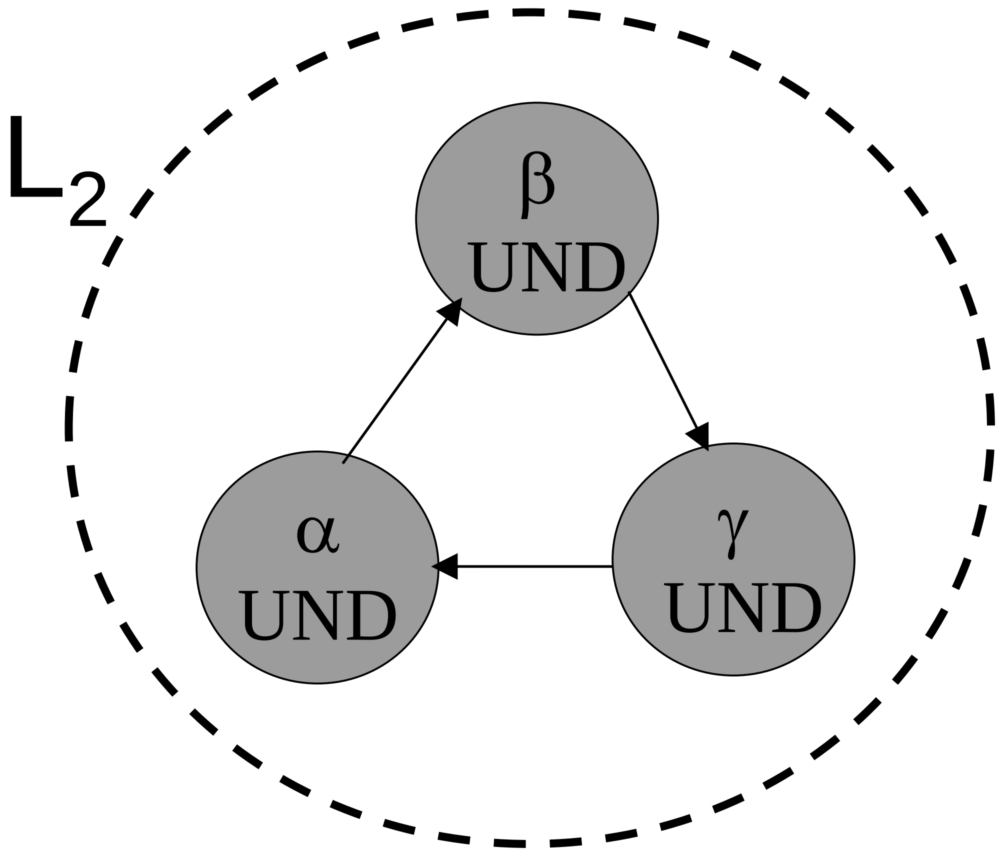

Defining argumentation semantics

- Many different proposals in the literature, corresponding to different intuitions
- There are cases where "every" literature semantics gives a different outcome
- How (and in how many ways) would you colour this graph?

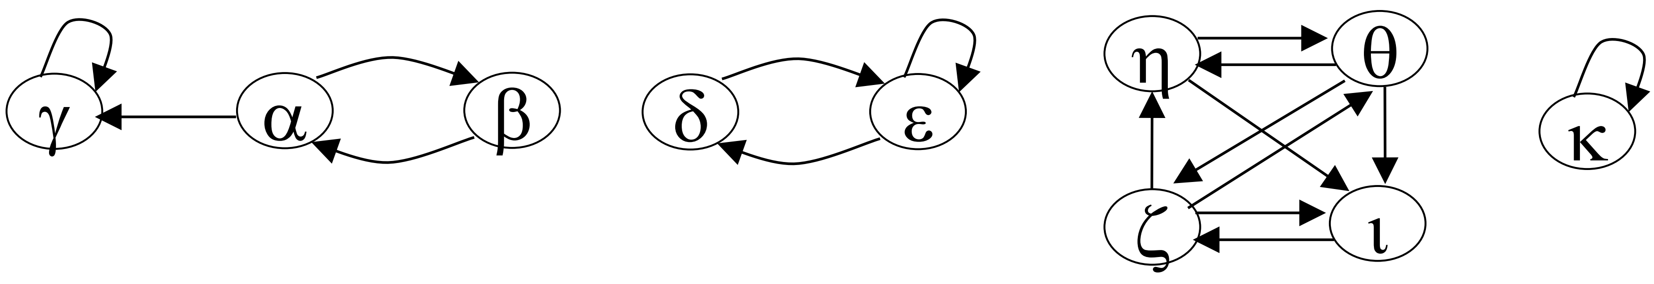

Analyzing argumentation semantics

- A catalogue of literature semantics is not immediately useful since different semantics definitions don’t follow a common pattern and are not directly comparable
- To start, examining a set of general properties driving the definition (or marking the differences) of argumentation semantics is more useful
- Focus on extension-based definition
- We will look at the semantics catalogue after that

Roadmap

- What's in the words?
- From the Big Bang to now
- Dung's framework
- A "conflict calculus": argumentation semantics
- "Basic" semantics properties

Conflict freeness

- The conflict free principle states that an attacking and an attacked argument can not stay together

Definition 8. Given an argumentation framework $\mathrm{AF} = \langle \mathcal{A}, \rightarrow \rangle$, a set $S \subseteq \mathcal{A}$ is *conflict-free*, denoted as $cf(S)$, iff $\nexists \alpha, \beta \in S$ such that $\alpha \rightarrow \beta$. A semantics $\mathcal{S}$ satisfies the $\mathcal{CF}$ *principle* if and only if $\forall\,\mathrm{AF}, \forall E \in \mathcal{E}_{\mathcal{S}}(\mathrm{AF})$ E is conflict-free.

Notational remark

$\mathcal{E}_{\mathcal{S}}(\mathrm{AF})$ usually denotes the set of extensions prescribed by semantics $\mathcal{S}$ for the argumentation framework AF

Conflict freeness

- The conflict free principle states that an attacking and an attacked argument can not stay together

Definition 8. Given an argumentation framework $\mathrm{AF} = \langle \mathcal{A}, \rightarrow \rangle$, a set $S \subseteq \mathcal{A}$ is *conflict-free*, denoted as $cf(S)$, iff $\nexists \alpha, \beta \in S$ such that $\alpha \rightarrow \beta$. A semantics $\mathcal{S}$ satisfies the $\mathcal{CF}$ *principle* if and only if $\forall\, \mathrm{AF}, \forall E \in \mathcal{E}_{\mathcal{S}}(\mathrm{AF})$ E is conflict-free.

- Very basic idea, followed by all semantics
- Minimal use of the attack relation (in fact the attack direction does not count)
- The empty set is conflict free by definition

Conflict freeness

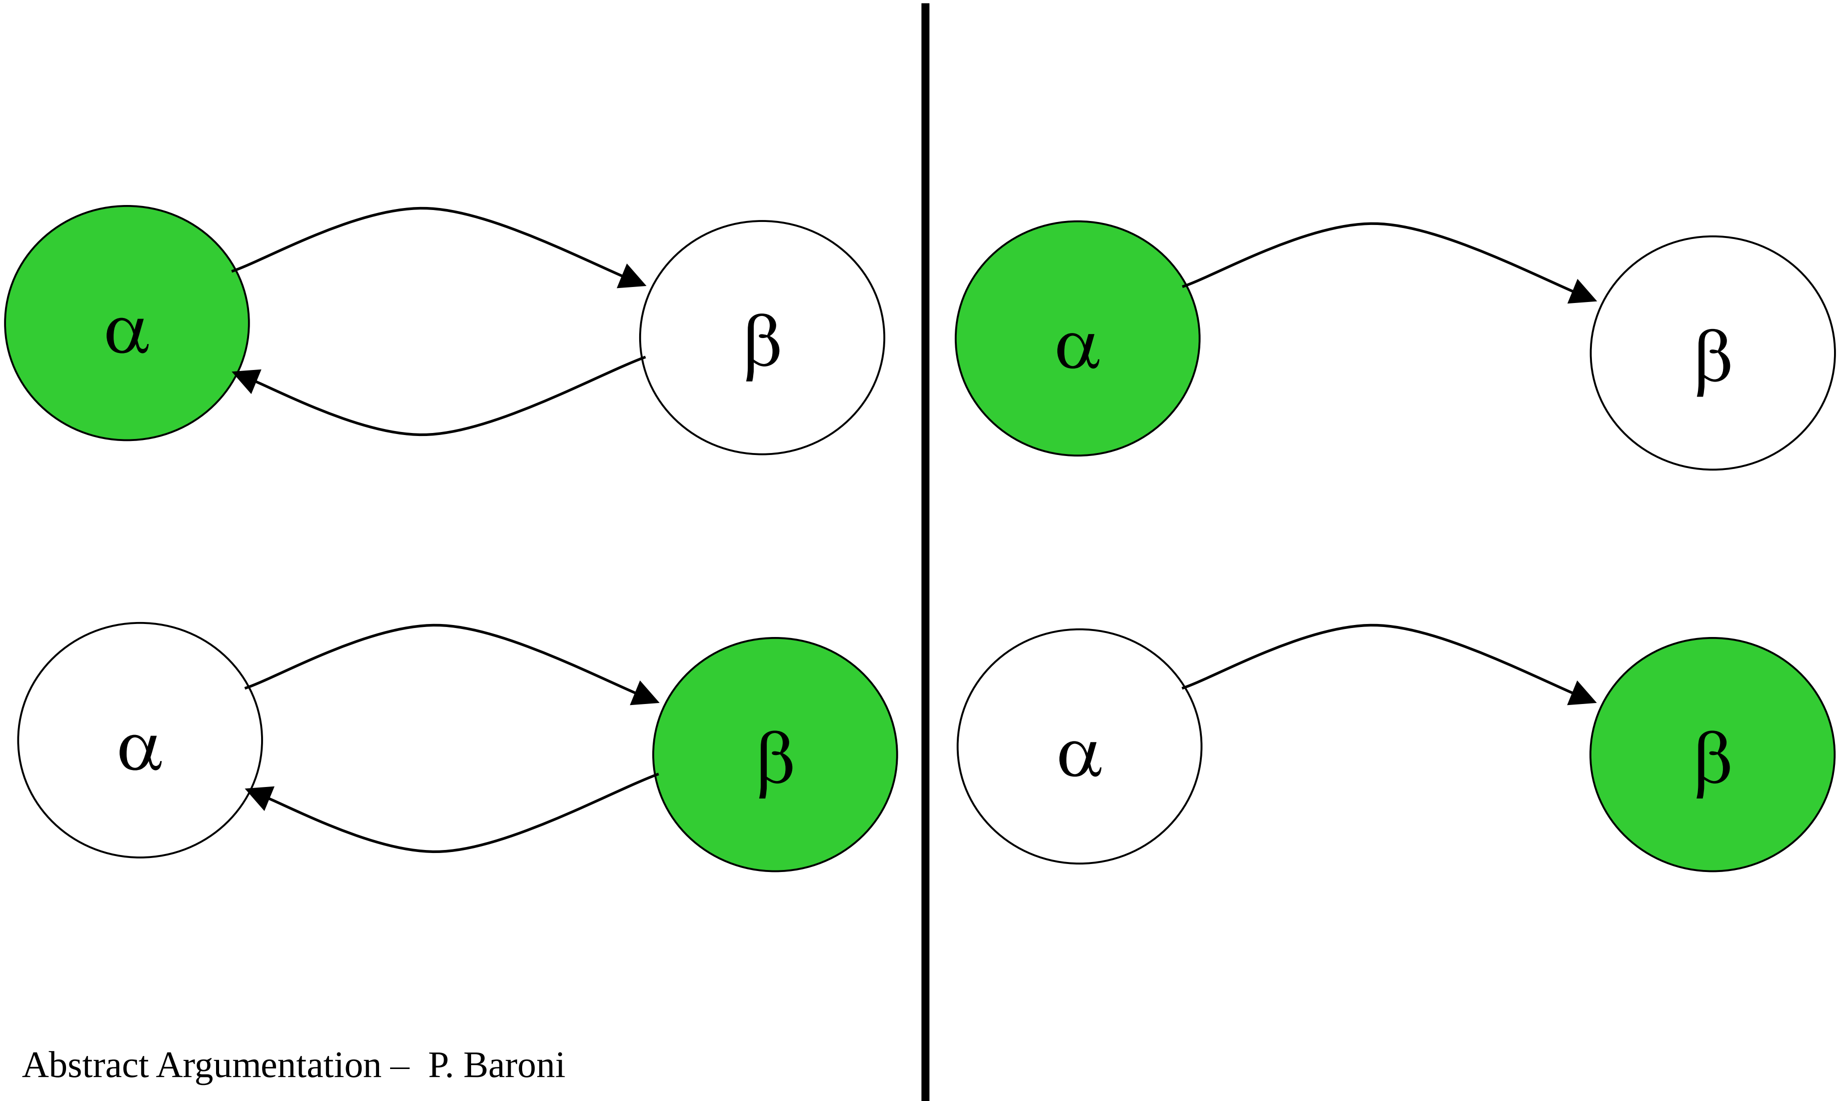

Conflict freeness

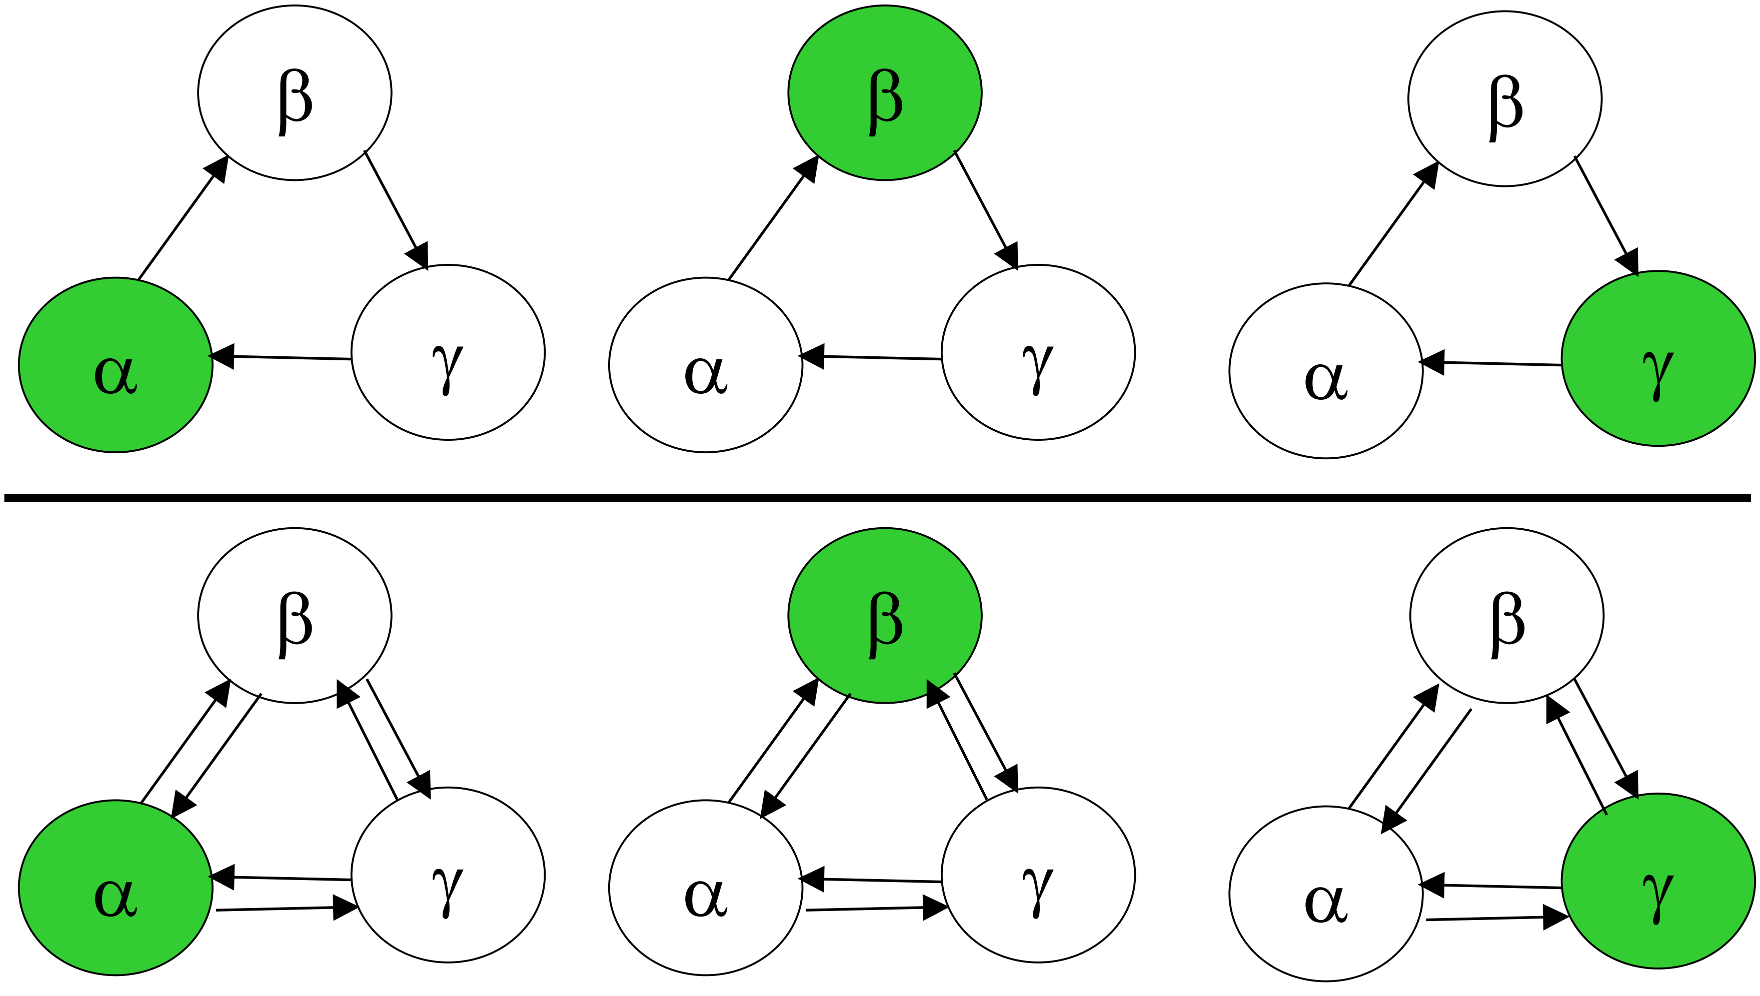

I-maximality

- A given criterion can select sets which are in a relation of proper inclusion (this holds for conflict freeness and for more articulated notions)
- One may consider as a constraint that no extension is a proper subset of another one

Definition 9. A set of extensions $\mathcal{E}$ is I-maximal iff $\forall E_1, E_2 \in \mathcal{E}$, if $E_1 \subseteq E_2$ then $E_1 = E_2$. A semantics $\mathcal{S}$ satisfies the I-*maximality criterion* if and only if $\forall$ AF $\mathcal{E}_{\mathcal{S}}(\mathrm{AF})$ is I-maximal.

- Important for some properties (e.g. to avoid triviality in the notion of skeptical justification) but not fundamental and poorly related with other principles

Defense related principles: admissibility

- To survive the conflict you should be able to defend yourself, namely to reply to every attack with a counterattack
- A conflict free set is admissible if it defends itself

Definition 11. Given an argumentation framework $\mathrm{AF} = \langle \mathcal{A}, \rightarrow \rangle$, a set $E \subseteq \mathcal{A}$ is *admissible* if and only if E is conflict-free and $\forall \beta \in \mathcal{A}$: $\beta \rightarrow E, E \rightarrow \beta$. The set made up of all the admissible sets of AF will be denoted as $\mathcal{AS}(\mathrm{AF})$.

Definition 12. A semantics $\mathcal{S}$ satisfies the *admissibility criterion* if $\forall\, \mathrm{AF} \in \mathcal{D}_{\mathcal{S}}, \forall E \in \mathcal{E}_{\mathcal{S}}(\mathrm{AF})\ E \in \mathcal{AS}(\mathrm{AF})$, namely:

$$\alpha \in E \Rightarrow \forall \beta \in \mathrm{par}_{\mathrm{AF}}(\alpha)\ E \rightarrow \beta \qquad (1)$$

- The attack direction counts
- The empty set is admissible by definition

Admissibility

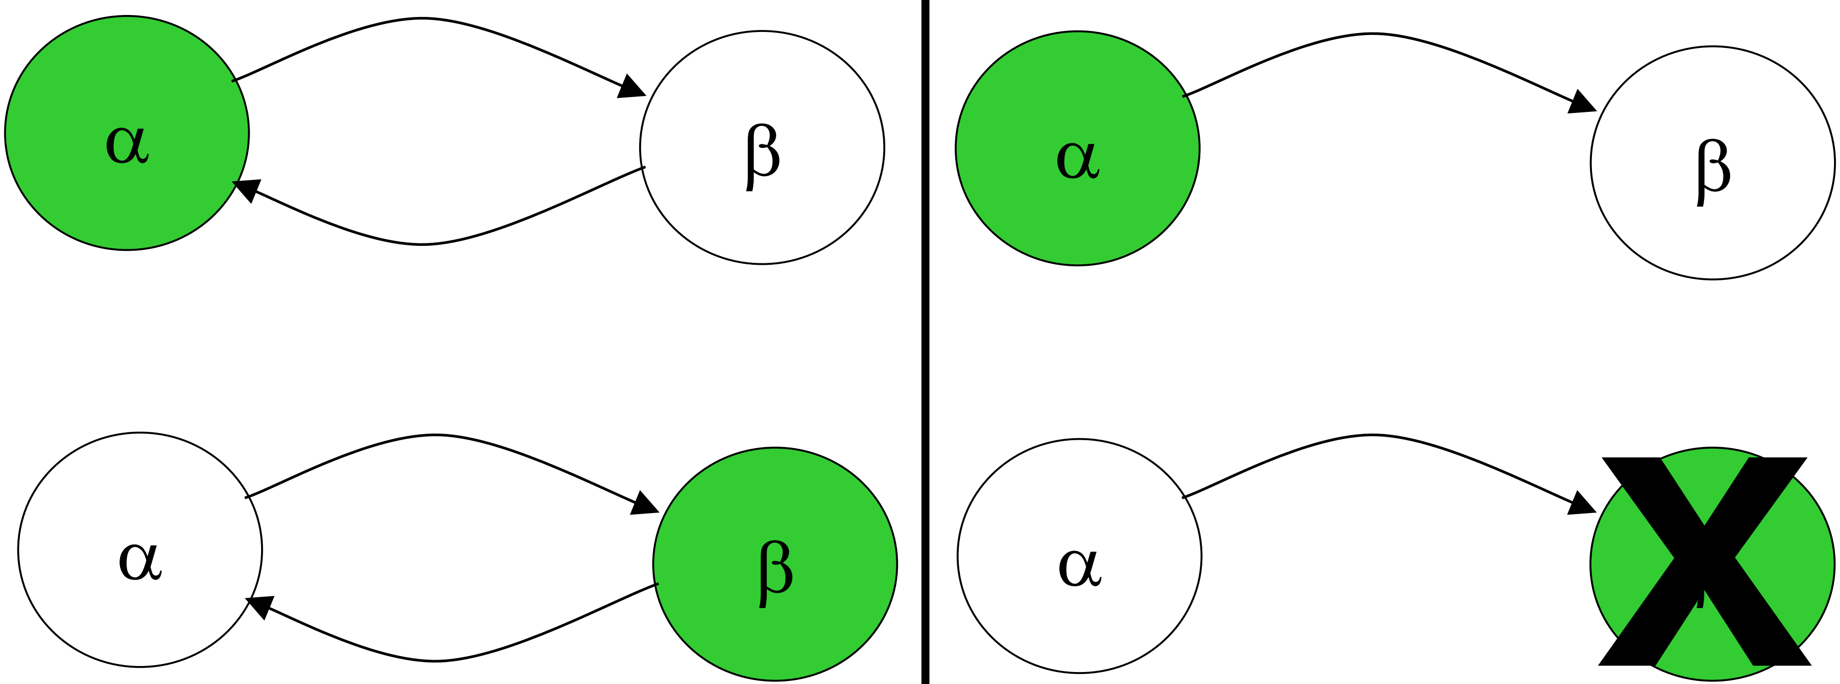

Admissibility

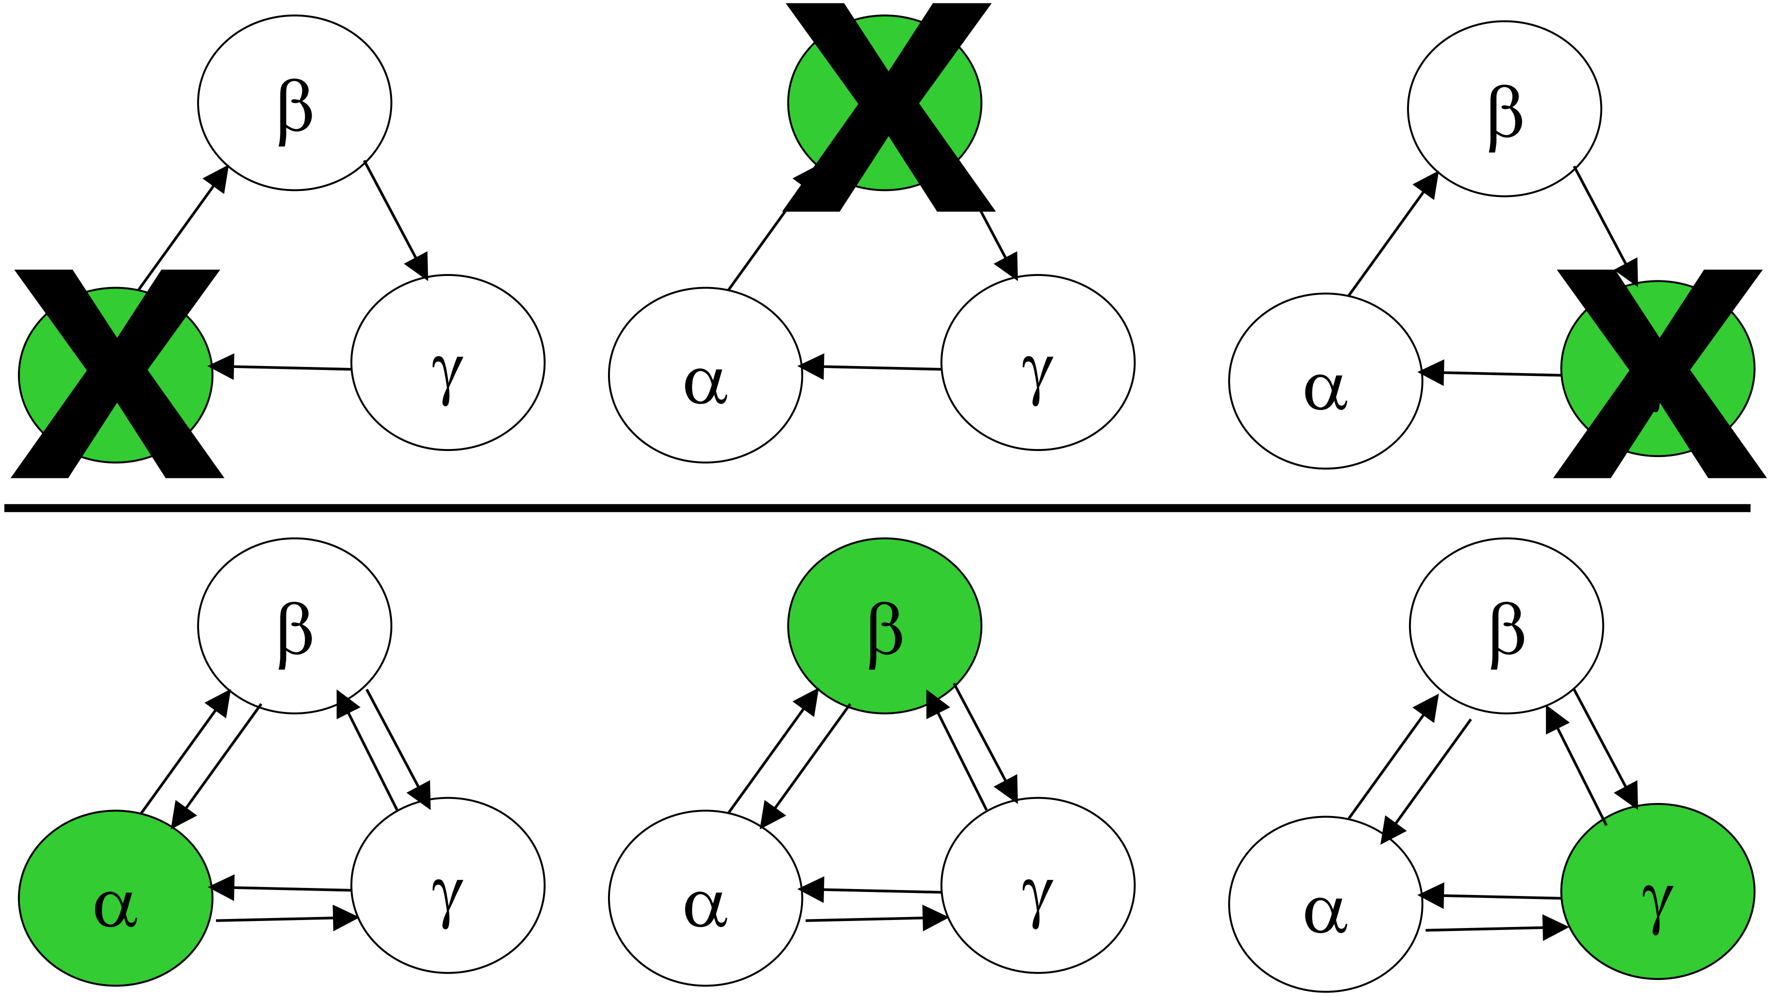

Admissibility from a labelling perspective

- You should have good reasons to assign IN or OUT labels, but you are free to be undecided
- All arguments UNDECIDED is an "admissible" labelling, as the empty set is always an admissible set

DEFINITION 9 *Let $\mathcal{L}ab$ be a labelling of argumentation framework (Ar, att):*

- *An* `in`*-labelled argument is said to be* legally `in` *iff all its attackers are labelled* `out`.
- *An* `out`*-labelled argument is said to be* legally `out` *iff it has at least one attacker that is labelled* `in`.

DEFINITION 10 *Let $AF = (Ar, att)$ be an argumentation framework. An* admissible labelling *is a labelling $\mathcal{L}ab$ where each* `in`*-labelled argument is legally* `in`*- and each* `out`*-labelled argument is legally* `out`.

Defense related principles: strong admissibility/defense

- Admissibility includes self-defense
- A stronger notion of defense requires that defense comes from other arguments which are in turn strongly defended by other arguments

Definition 13. Given an argumentation framework $\mathrm{AF} = \langle \mathcal{A}, \rightarrow \rangle$, $\alpha \in \mathcal{A}$ and $S \subseteq \mathcal{A}$, we say that α is strongly defended by S (denoted as $sd(\alpha, S)$) iff $\forall \beta \in \mathrm{par}_{\mathrm{AF}}(\alpha)\ \exists \gamma \in S \setminus \{\alpha\}$: $\gamma \rightarrow \beta$ and $sd(\gamma, S \setminus \{\alpha\})$.

Definition 14. A semantics $\mathcal{S}$ satisfies the *strong admissibility criterion* if $\forall\,\mathrm{AF} \in \mathcal{D}_{\mathcal{S}}$, $\forall E \in \mathcal{E}_{\mathcal{S}}(\mathrm{AF})$ it holds that:

$$\alpha \in E \Rightarrow sd(\alpha, E) \qquad (2)$$

- Rather strong requirement: defense chains rooted in unattacked arguments
- The empty set satisfies strong admissibility

Strong defense

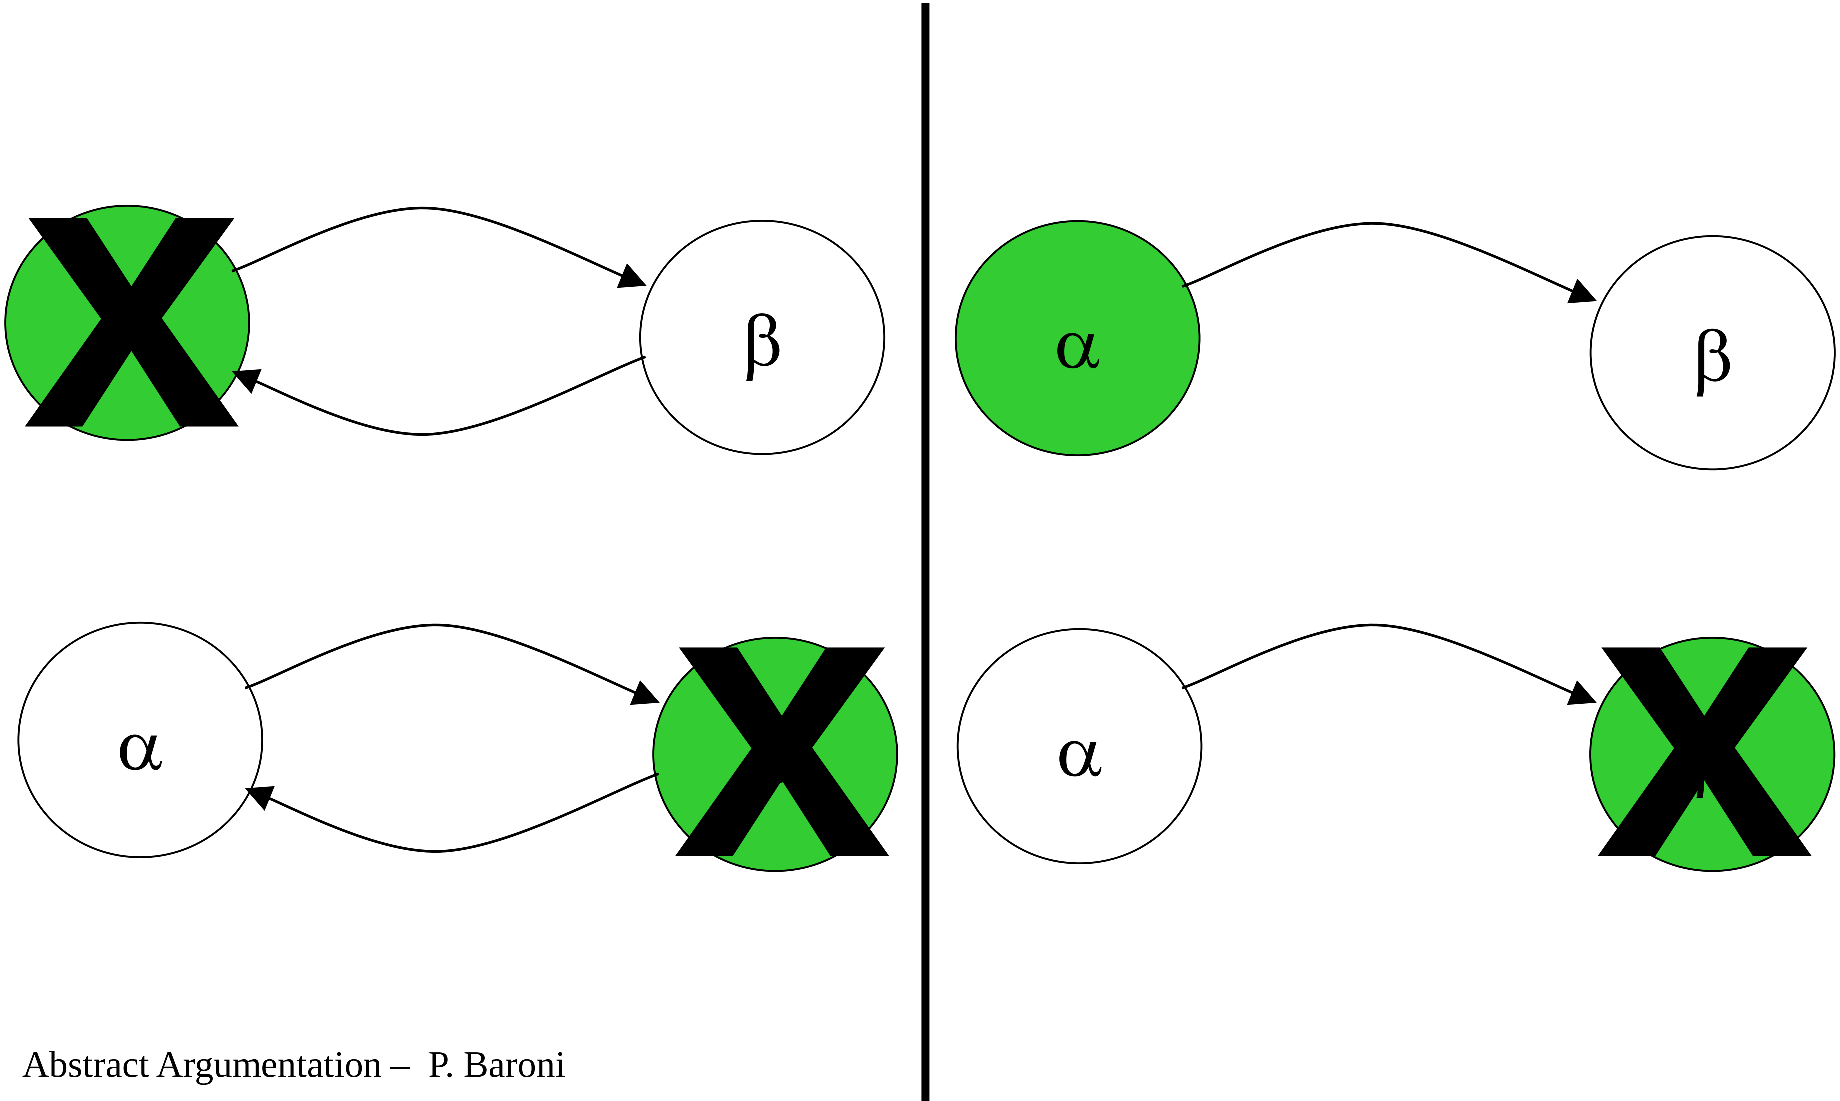

Strong defense

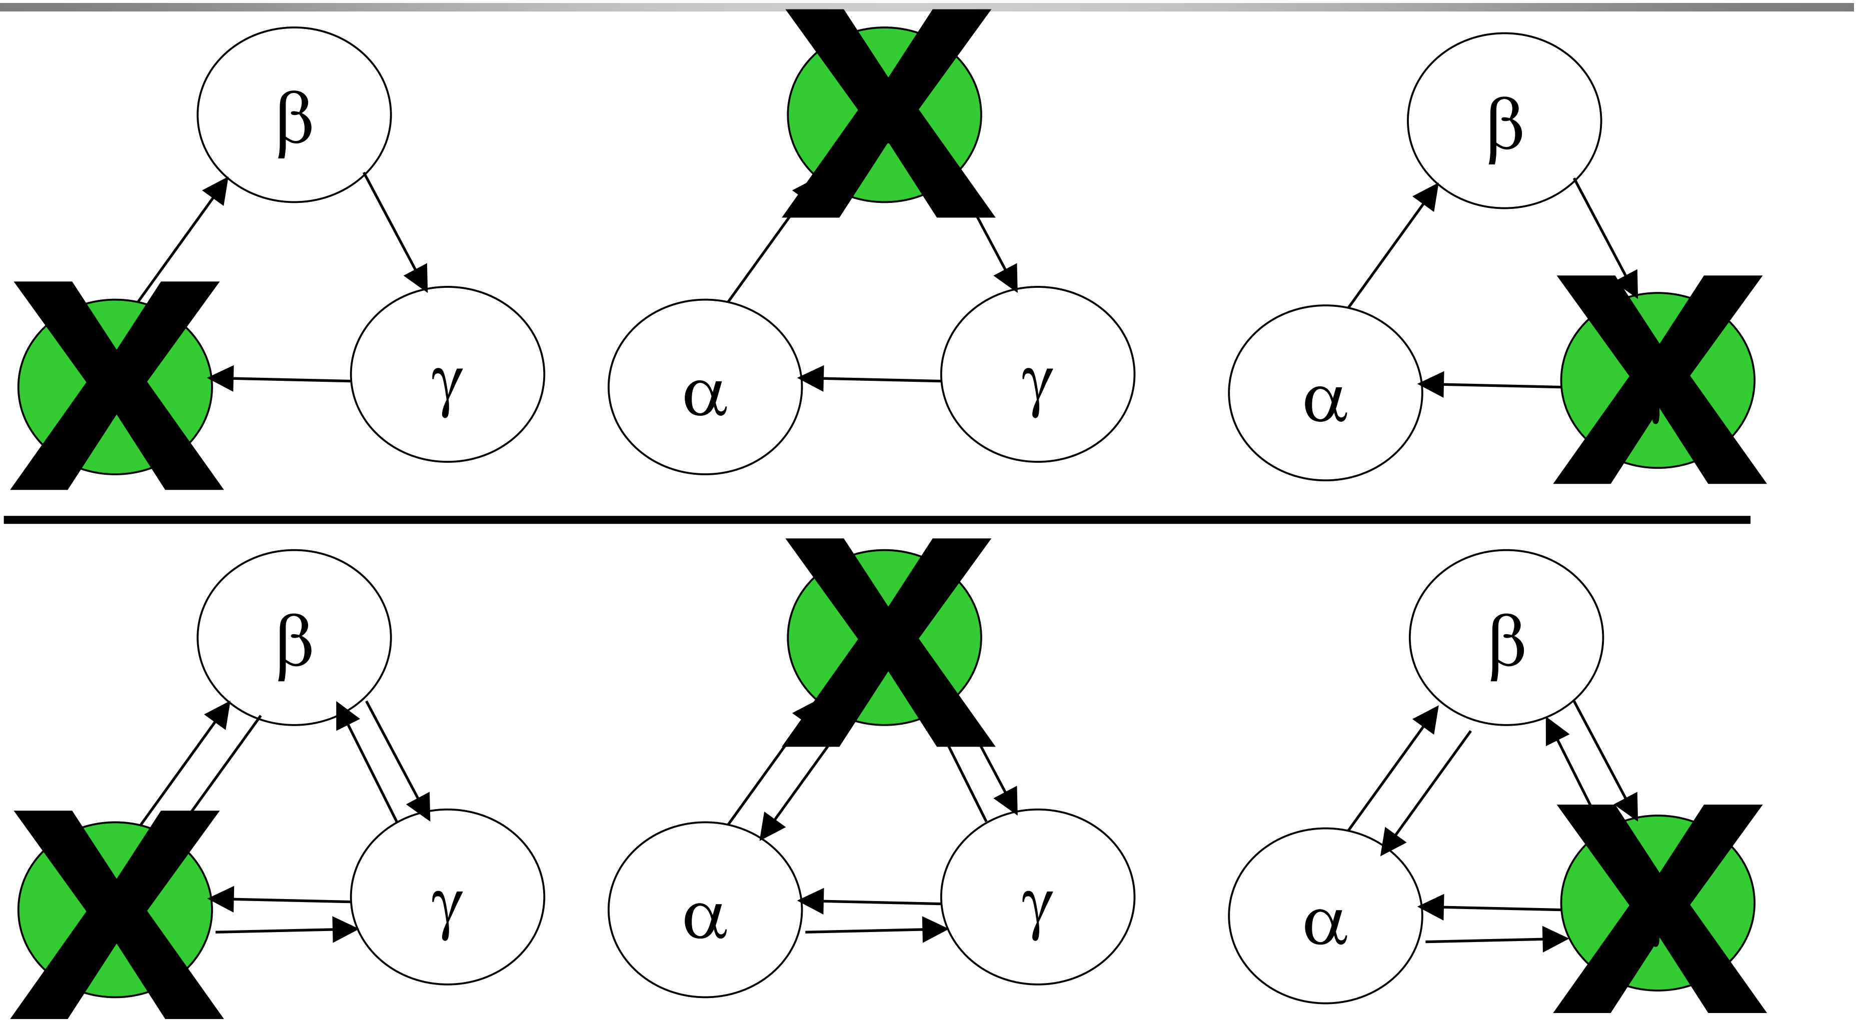

Strong defense

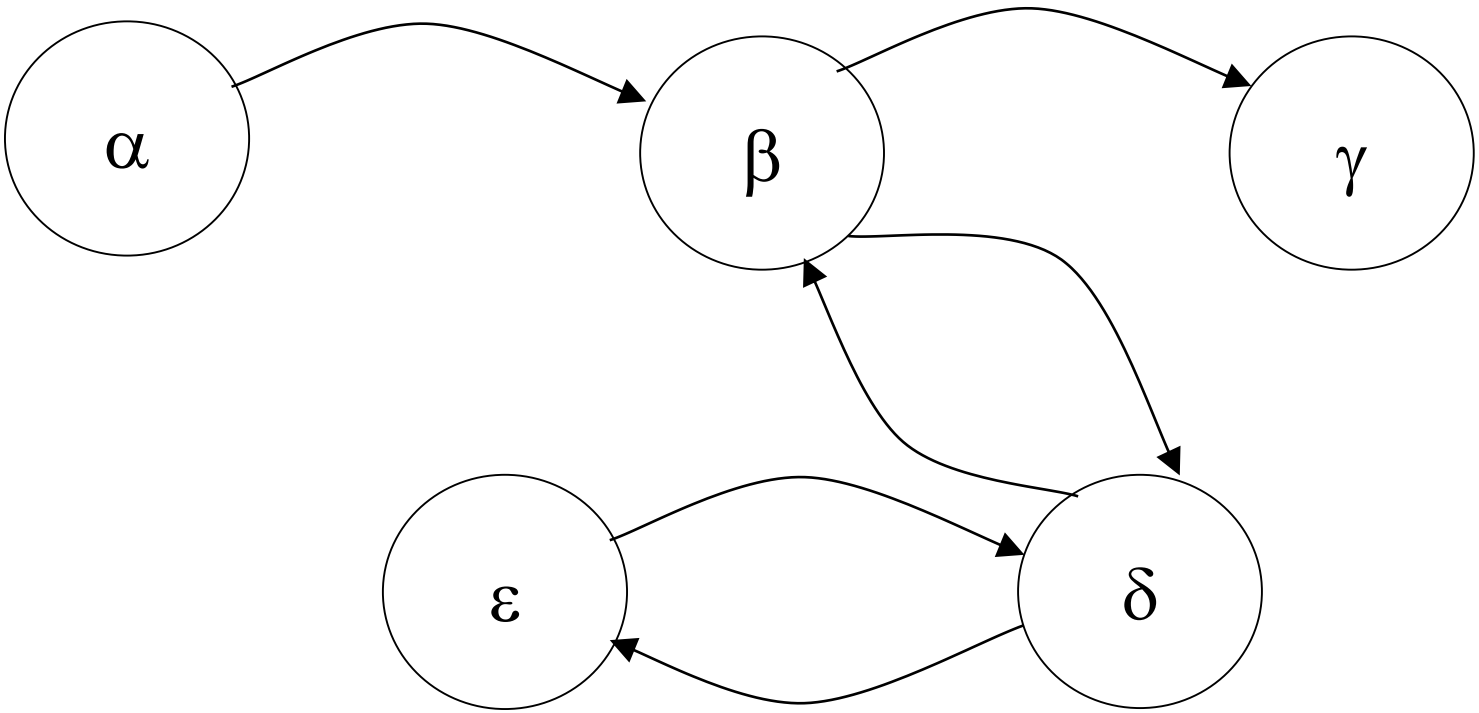

Strong defense

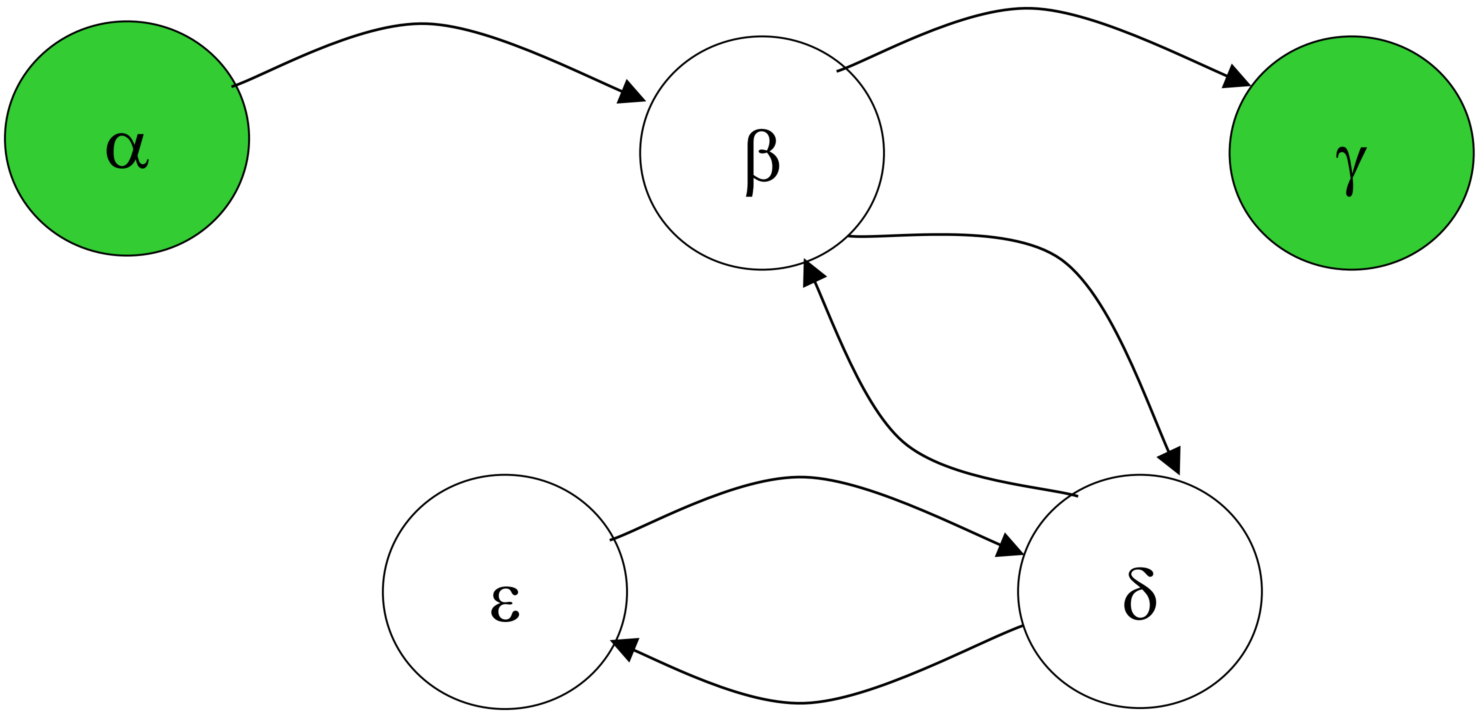

Strong defense

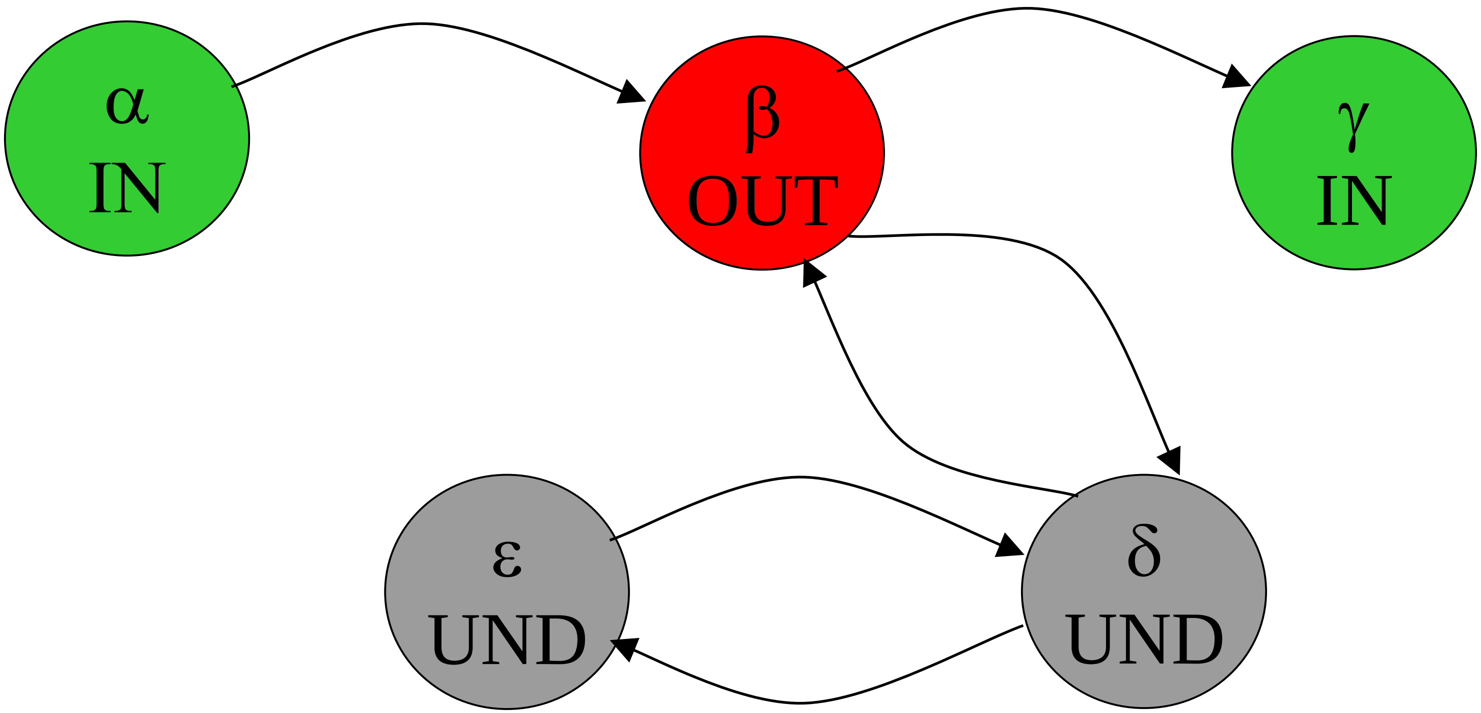

Defense related principles: reinstatement

- Reinstatement concerns effectiveness (or altruism) of defense: if you defend some argument you should take it on board (include it in the extension)

Definition 15. A semantics $\mathcal{S}$ satisfies the *reinstatement criterion* if $\forall\, \mathrm{AF} \in \mathcal{D}_{\mathcal{S}}$, $\forall E \in \mathcal{E}_{\mathcal{S}}(\mathrm{AF})$ it holds that:

$$\big(\forall \beta \in \mathrm{par}_{\mathrm{AF}}(\alpha)\ E \rightarrow \beta\big) \Rightarrow \alpha \in E$$

- Completeness requirement: can not leave out your protected ones
- The empty set satisfies reinstatement only if there are no unattacked arguments

Reinstatement

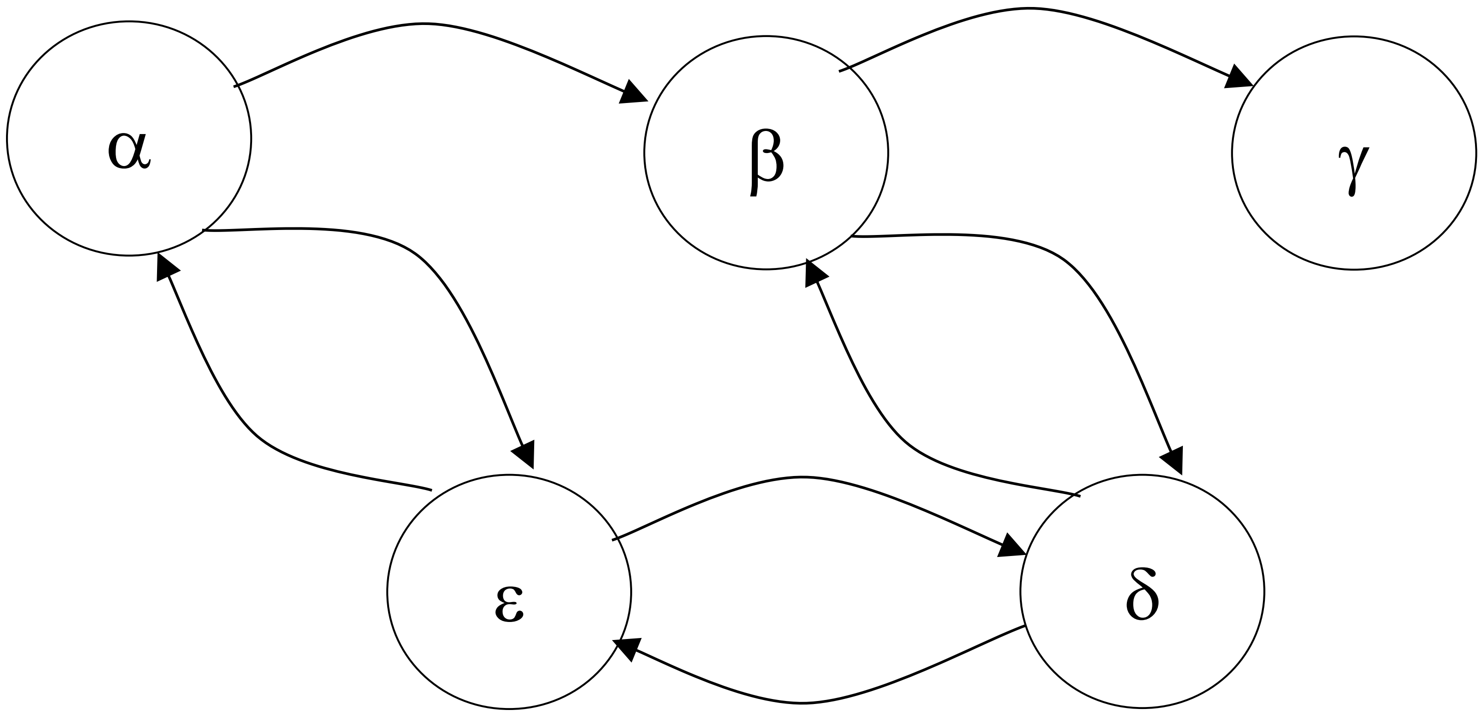

Reinstatement

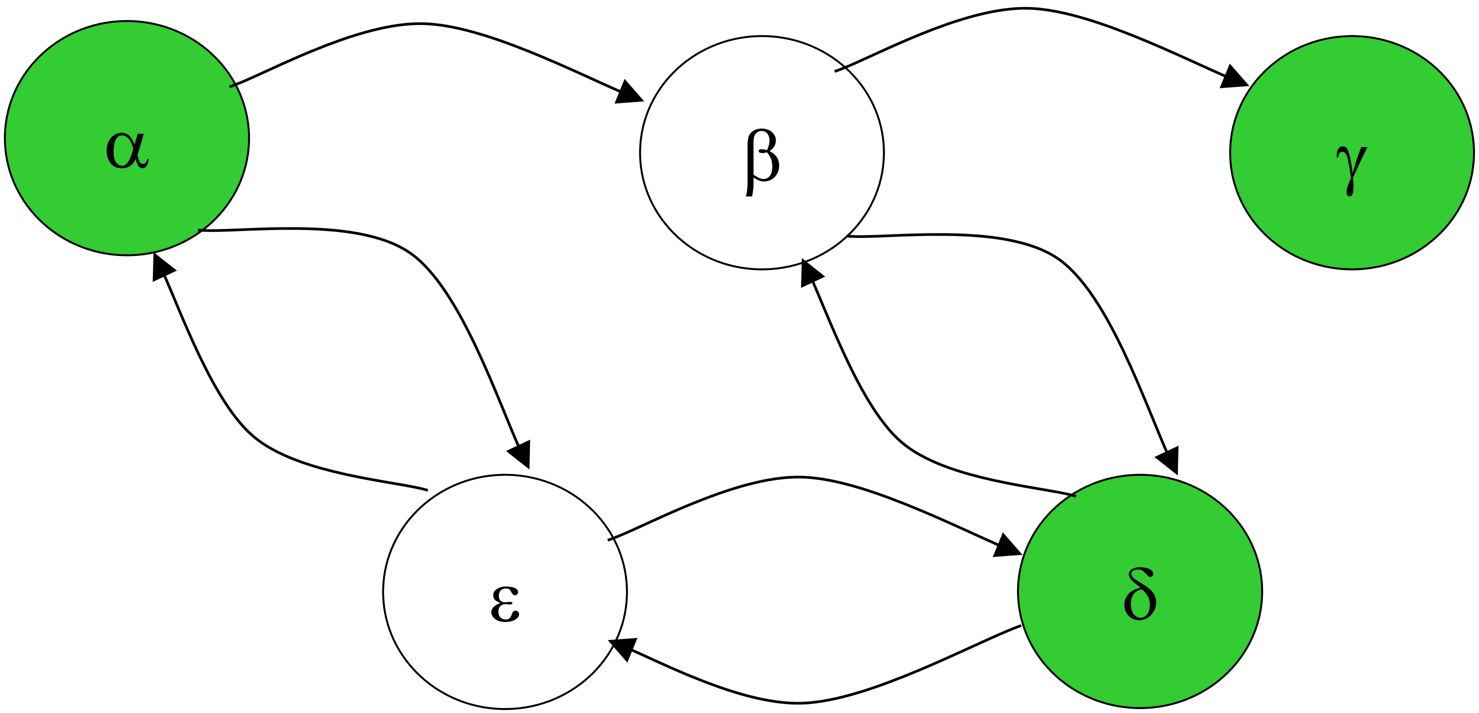

Reinstatement

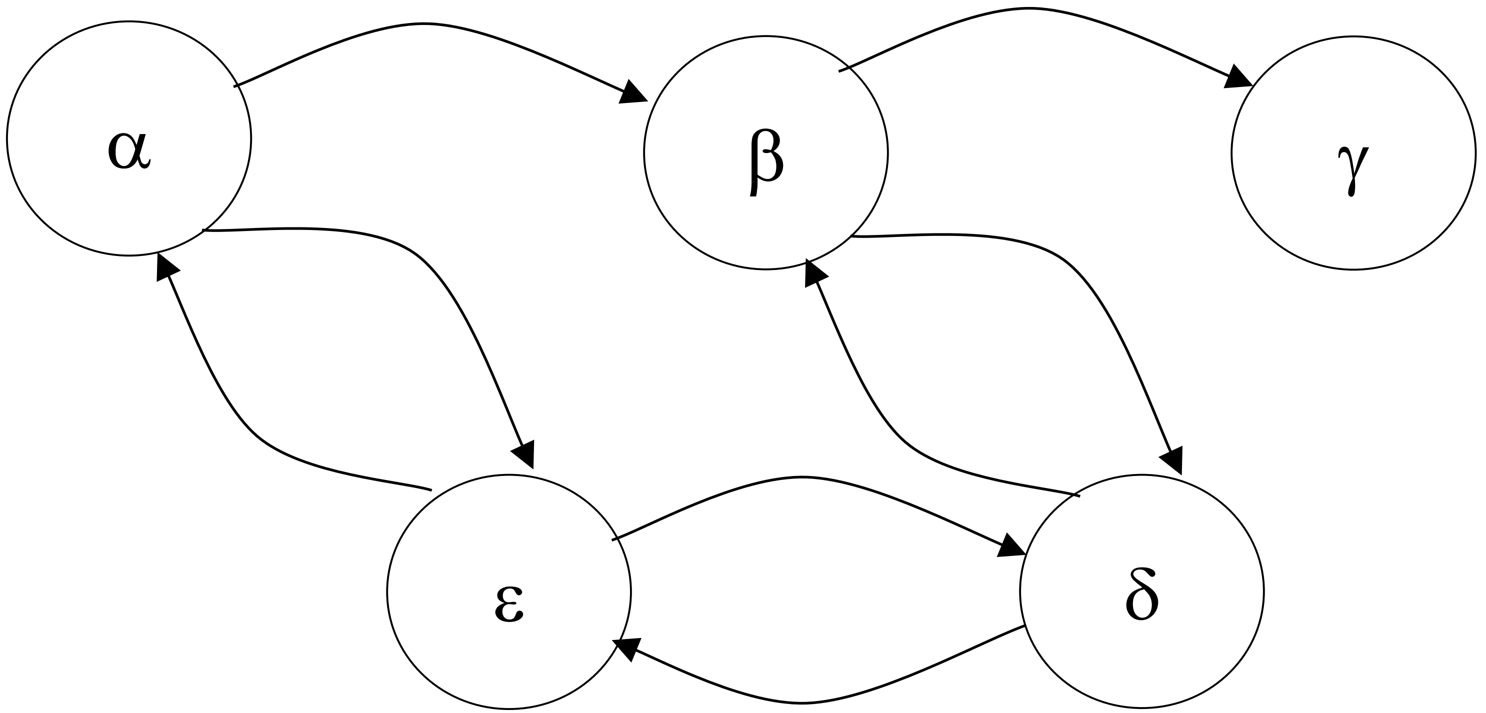

Reinstatement

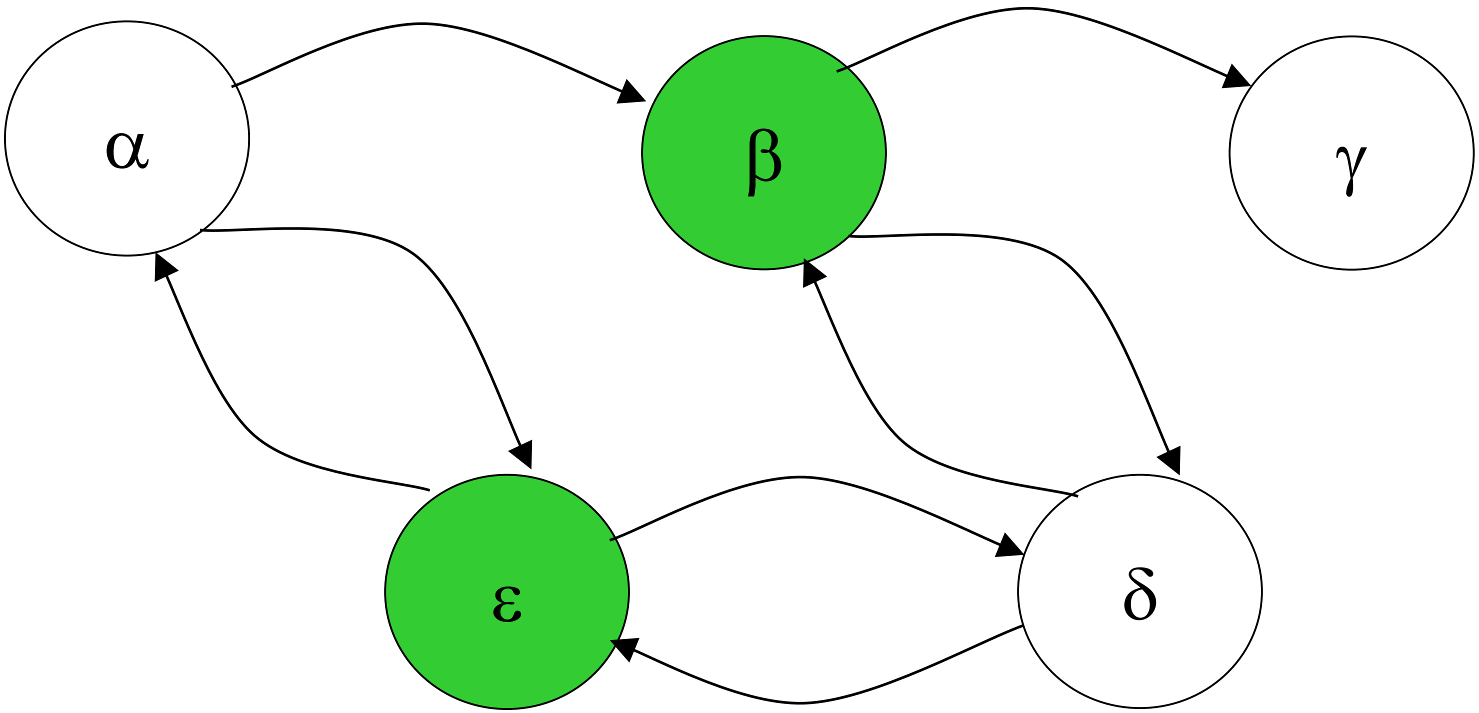

Defense related principles: weak reinstatement

- A weaker notion of reinstatement is obtained considering strong defense instead of defense

Definition 16. A semantics $\mathcal{S}$ satisfies the *weak reinstatement criterion* if $\forall$ AF $\in \mathcal{D}_{\mathcal{S}}$, $\forall E \in \mathcal{E}_{\mathcal{S}}(\mathrm{AF})$ it holds that:

$$sd(\alpha, E) \Rightarrow \alpha \in E \qquad (4)$$

Defense related principles: CF reinstatement

- Reinstatement does not require explicitly conflict freeness with the defended argument. Adding this gives rise to CF reinstatement

Definition 17. A semantics $\mathcal{S}$ satisfies the $\mathcal{CF}$*-reinstatement criterion* if $\forall\, \mathrm{AF} \in \mathcal{D}_{\mathcal{S}}$, $\forall E \in \mathcal{E}_{\mathcal{S}}(\mathrm{AF})$ it holds that:

$$\big(\big(\forall \beta \in \mathrm{par}_{\mathrm{AF}}(\alpha)\ E \to \beta\big) \wedge cf\big(E \cup \{\alpha\}\big)\big) \Rightarrow \alpha \in E \qquad (5)$$

Directionality

- Basic idea: arguments affect each other following the direction of attacks.
- An argument (or set of arguments) is affected only by its ancestors in the attack relation

Definition 18. Given an argumentation framework $\mathrm{AF} = \langle \mathcal{A}, \rightarrow \rangle$, a set $U \subseteq \mathcal{A}$ is *unattacked* if and only if $\nexists \alpha \in (\mathcal{A} \setminus U)\colon \alpha \rightarrow U$. The set of unattacked sets of AF is denoted as $\mathcal{US}(\mathrm{AF})$.

Definition 19. A semantics $\mathcal{S}$ satisfies the directionality criterion if and only if $\forall\, \mathrm{AF} = \langle \mathcal{A}, \rightarrow \rangle$, $\forall U \in \mathcal{US}(\mathrm{AF})$, $\mathcal{AE}_{\mathcal{S}}(\mathrm{AF}, U) = \mathcal{E}_{\mathcal{S}}(\mathrm{AF}{\downarrow_U})$, where $\mathcal{AE}_{\mathcal{S}}(\mathrm{AF}, U) \triangleq \{(E \cap U) \mid E \in \mathcal{E}_{\mathcal{S}}(\mathrm{AF})\} \subseteq 2^U$.

Directionality

- In words the restrictions of the extensions to an unattacked part of the graph are the same whatever is the remaining part of the graph (if any)
- Extensions (or labellings) can be constructed locally in suitable parts of the graph*

*Directionality implies that you may compute the restrictions of extensions to unattacked sets without considering the rest of the graph, it does not however always imply that there is an easy way to proceed incrementally to compute the extensions for the rest of the graph

Directionality

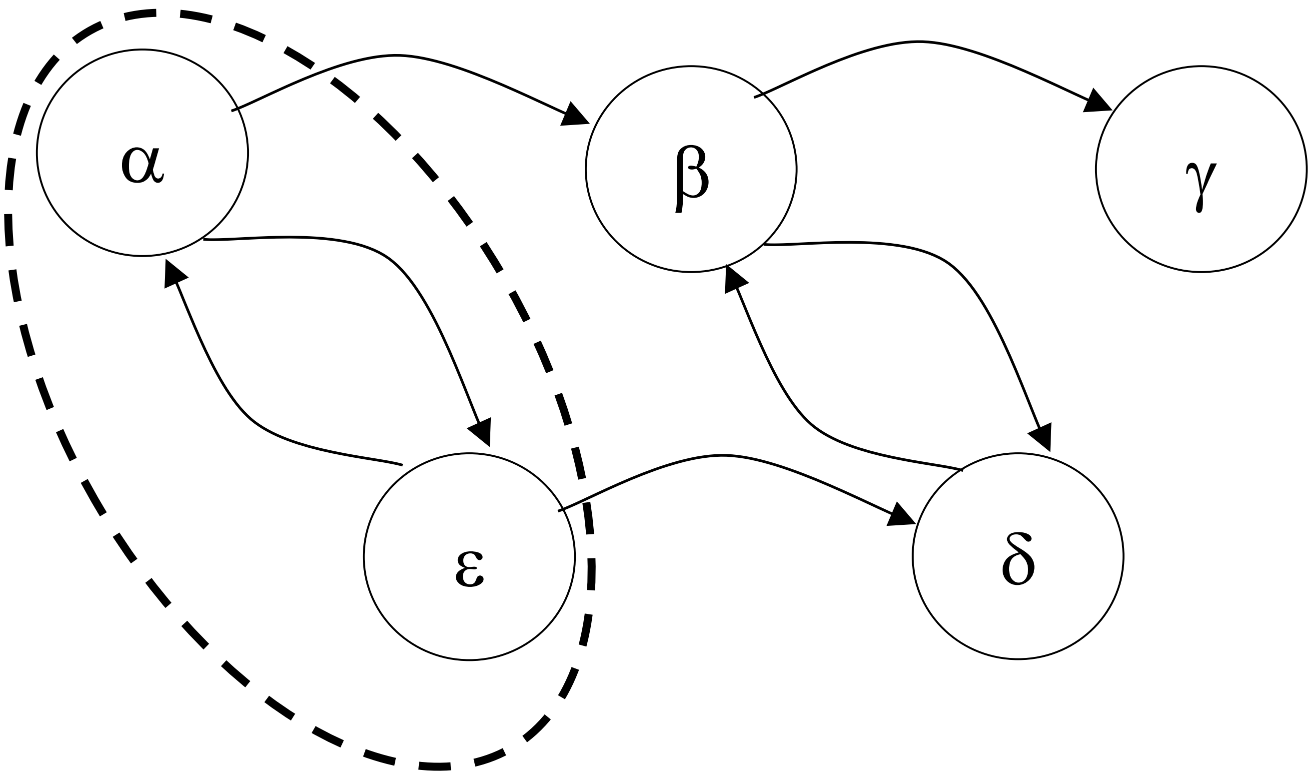

Directionality

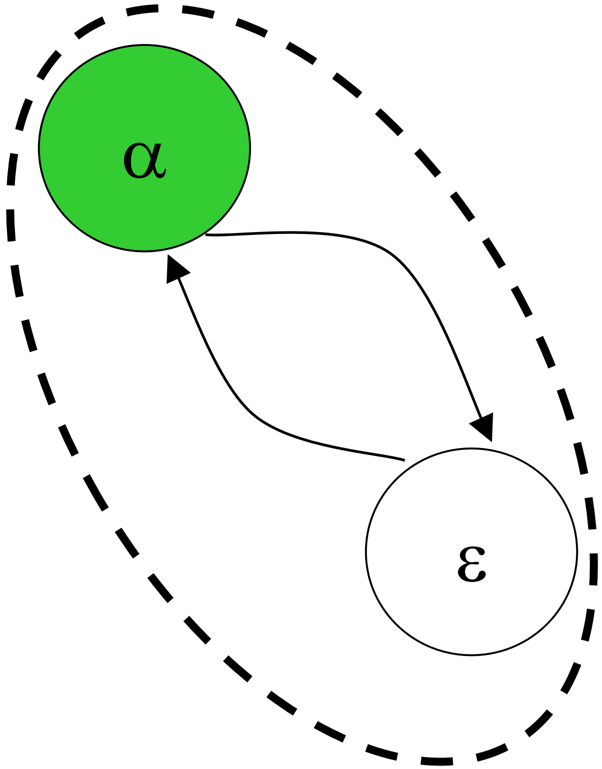

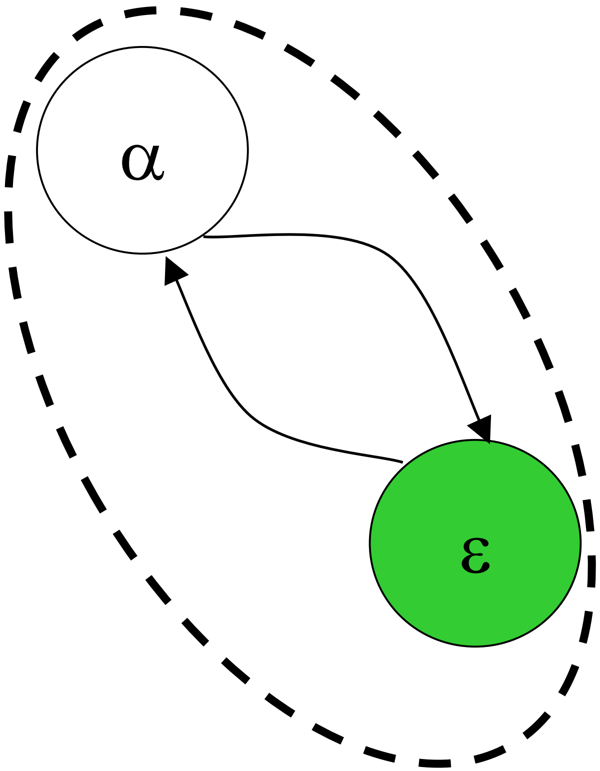

Directionality

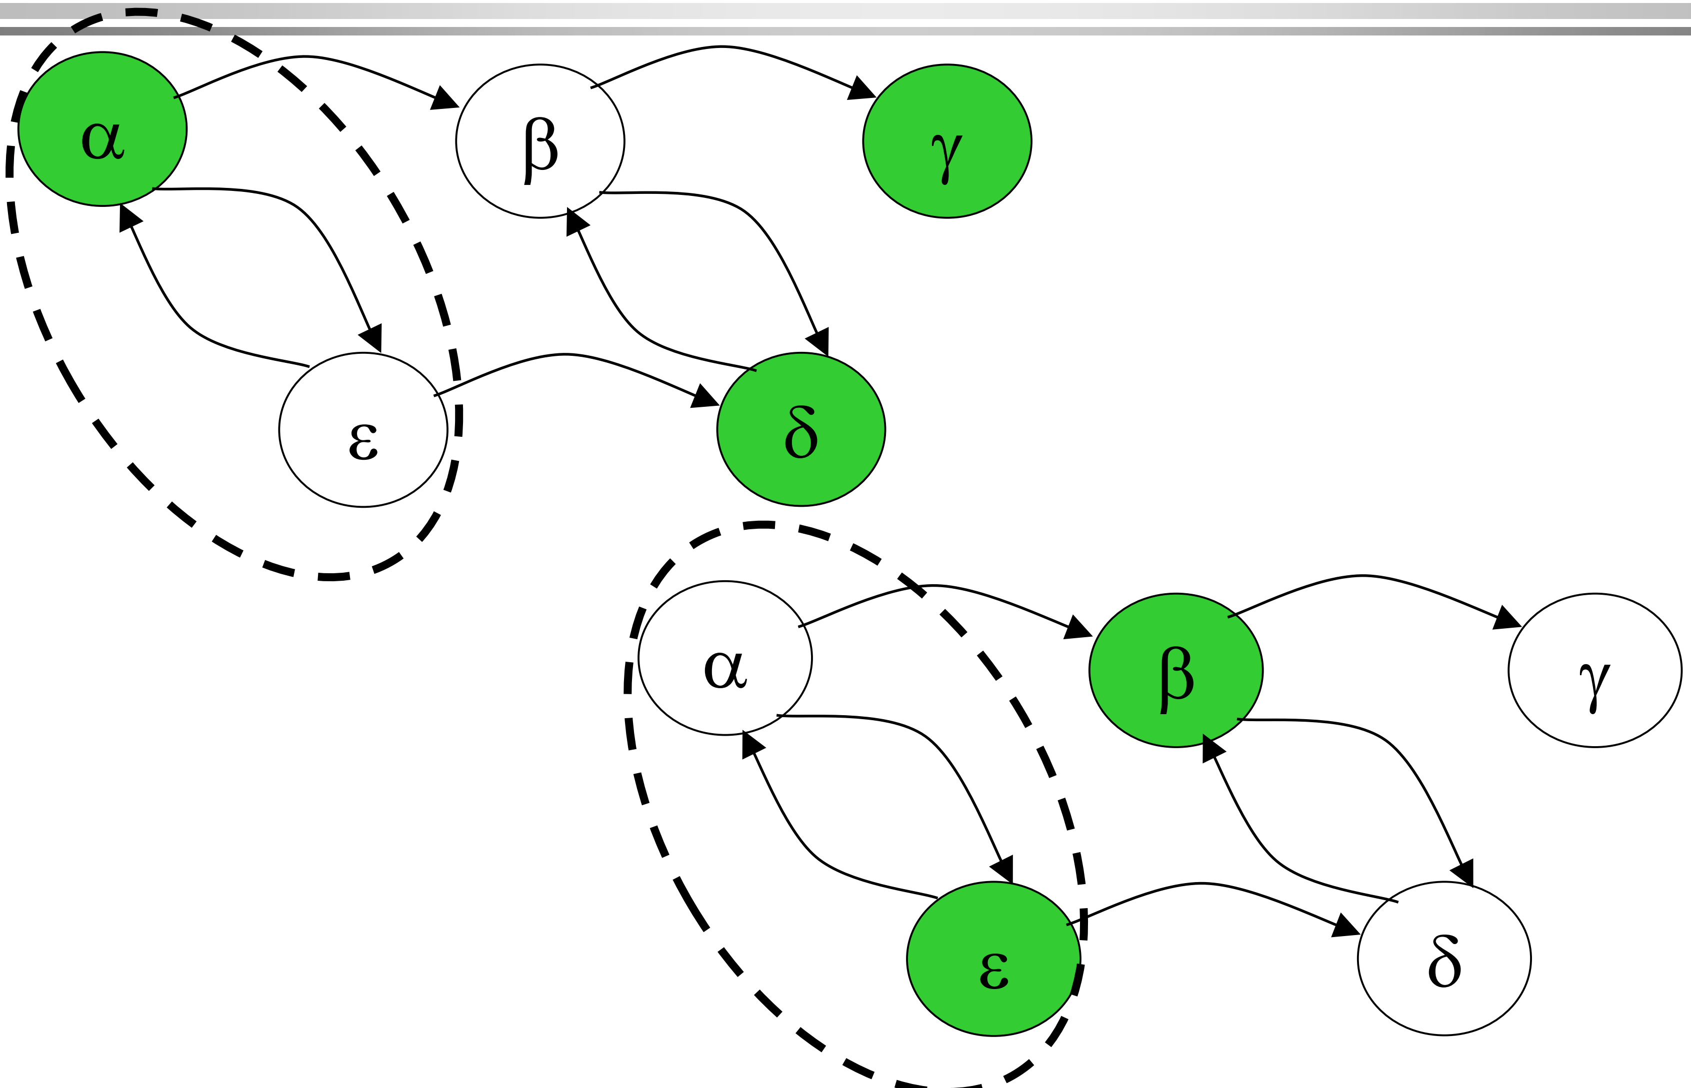

Directionality

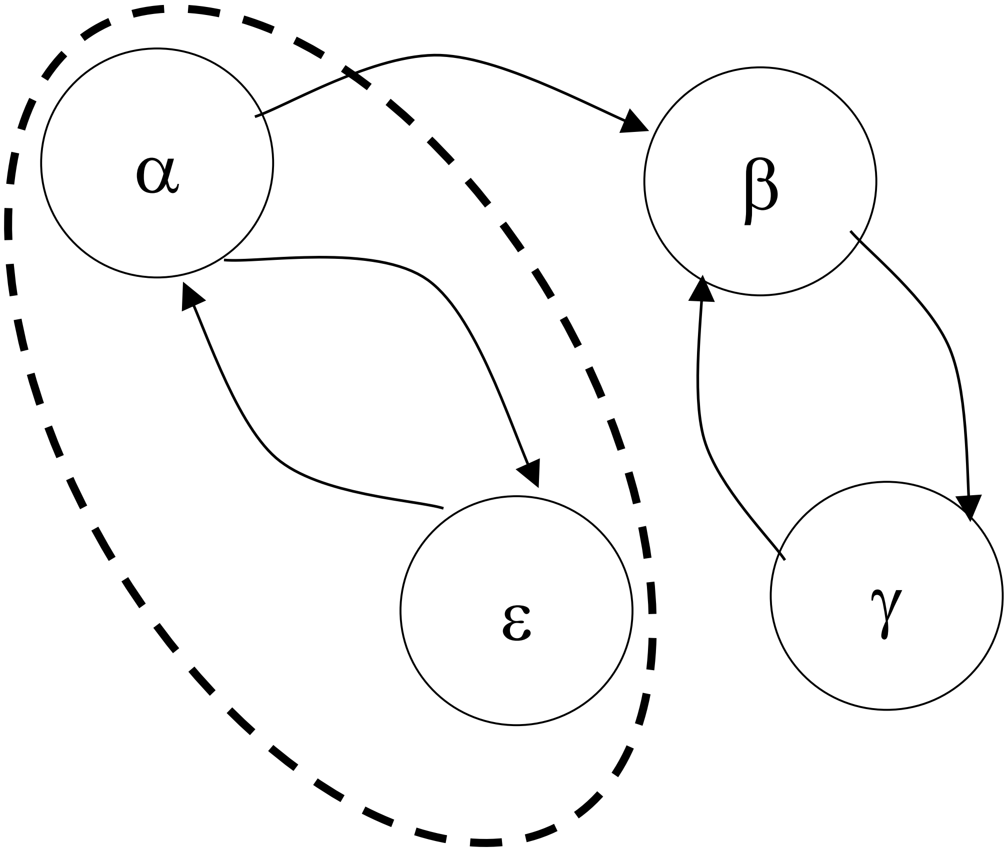

Directionality

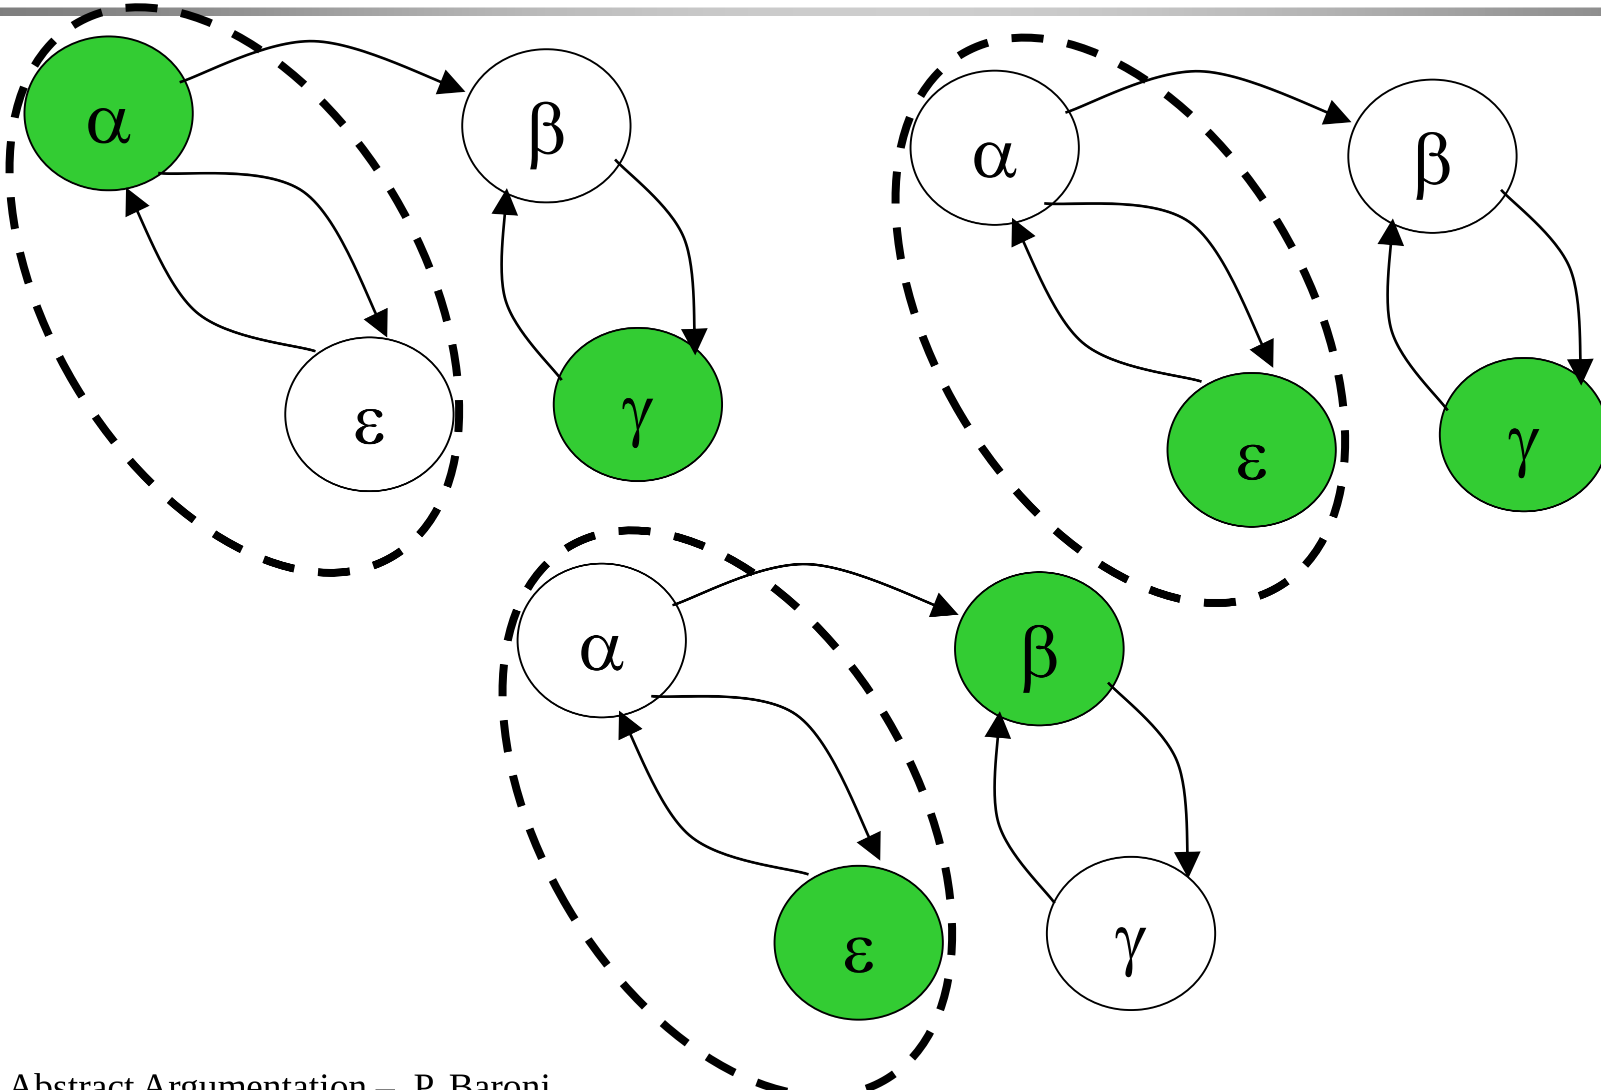

A hidden principle: language independence

- Extensions only depend on the attack relation, i.e. on the graph topology not on argument "names" or on other underlying properties
- Isomorphic frameworks have the same (modulo the isomorphism) extensions

Definition 6. Two argumentation frameworks $\mathrm{AF}_1 = \langle \mathcal{A}_1, \rightarrow_1 \rangle$ and $\mathrm{AF}_2 = \langle \mathcal{A}_2, \rightarrow_2 \rangle$ are isomorphic if and only if there is a bijective mapping $m : \mathcal{A}_1 \rightarrow \mathcal{A}_2$, such that $(\alpha, \beta) \in \rightarrow_1$ if and only if $(m(\alpha), m(\beta)) \in \rightarrow_2$. This is denoted as $\mathrm{AF}_1 \stackrel{\circ}{=}_m \mathrm{AF}_2$.

Definition 7. A semantics $\mathcal{S}$ satisfies the language independence principle if and only if $\forall\, \mathrm{AF}_1 = \langle \mathcal{A}_1, \rightarrow_1 \rangle, \forall\, \mathrm{AF}_2 = \langle \mathcal{A}_2, \rightarrow_2 \rangle$: $\mathrm{AF}_1 \stackrel{\circ}{=}_m \mathrm{AF}_2$, $\mathcal{E}_{\mathcal{S}}(\mathrm{AF}_2) = \{M(E) \mid E \in \mathcal{E}_{\mathcal{S}}(\mathrm{AF}_1)\}$, where $M(E) = \{\beta \in \mathcal{A}_2 \mid \exists \alpha \in E, \beta = m(\alpha)\}$.

Implicit equal treatment

- The language independence principle implies “equal opportunity” for arguments, i.e. that “topologically equivalent” arguments are treated in the same way
- No semantics can prescribe an “asymmetric” set of extensions for a “symmetric” framework

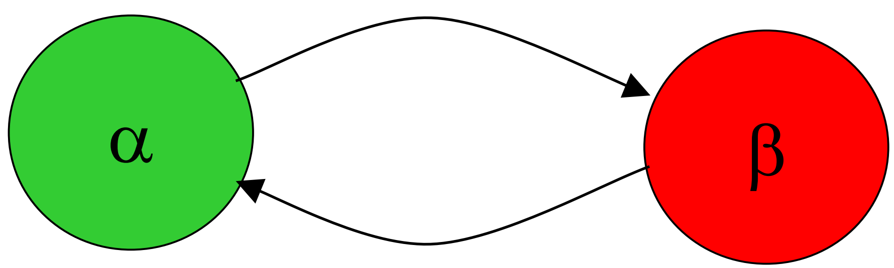

Implicit equal treatment

- If there are reasons to prefer α to β, they should be reflected in the topology

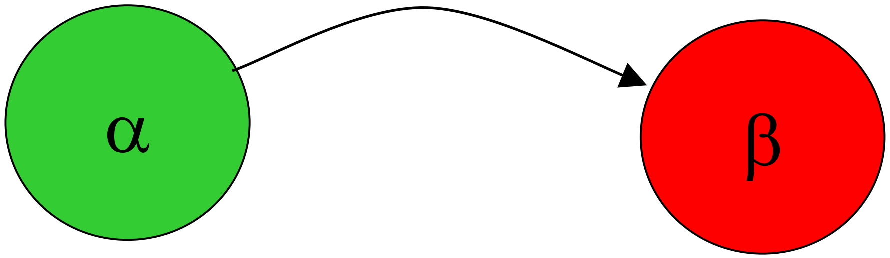

Roadmap

- What's in the words?
- From the Big Bang to now
- Dung's framework
- A "conflict calculus": argumentation semantics
- "Basic" semantics properties
- A catalogue of semantics

Conflict-freeness + maximality= naive semantics

- The so-called naive semantics prescribes as extensions the maximal conflict free sets of a framework
- As already remarked, it ignores the attack direction an so does not reflect the whole amount of available information
- It blatantly violates several basic principles like admissibility and reinstatement (even for unattacked arguments)
- In a sense, it is not really a semantics and is not considered in Dung’s paper

Naive semantics

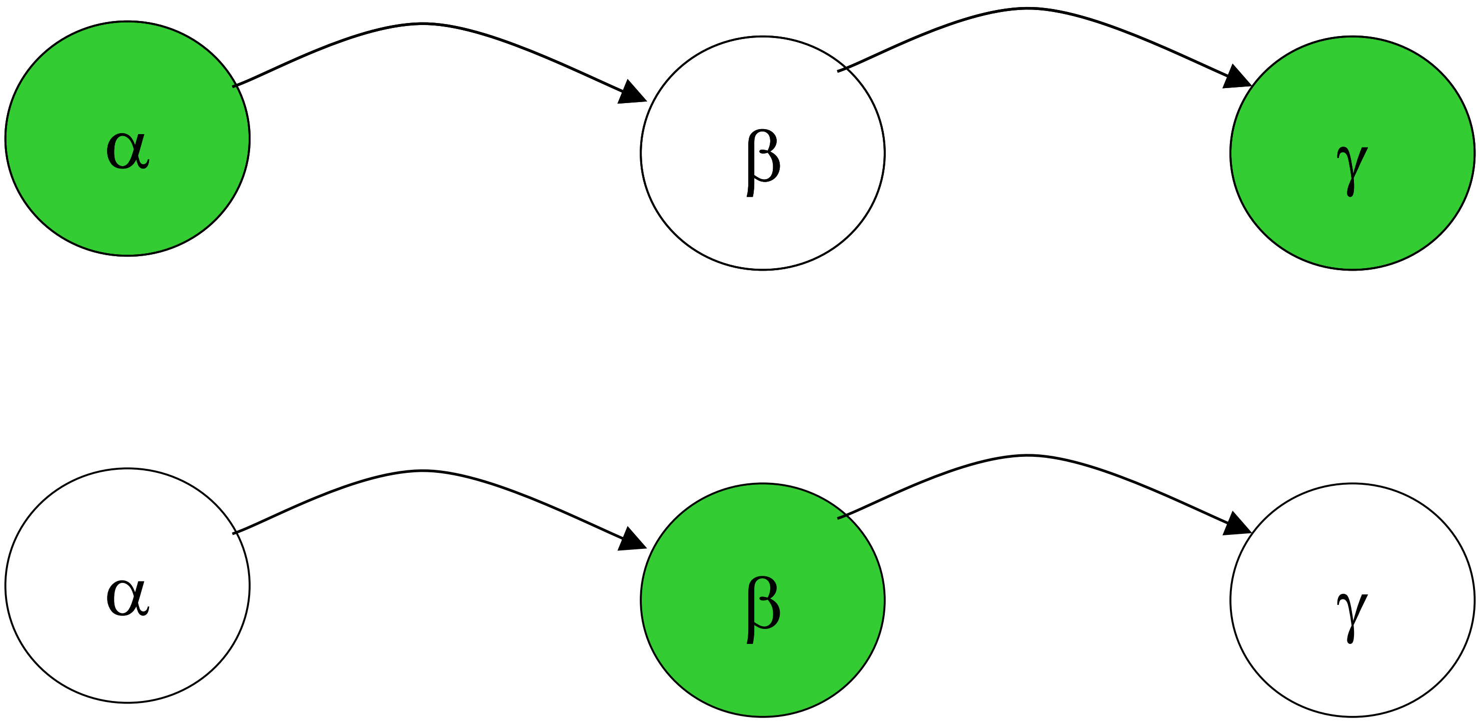

Can naive semantics work?

- Naive semantics makes sense when the direction of the attacks does not count: symmetric frameworks
- In any symmetric framework maximal conflict free-sets are reasonable extensions for any multiple-status semantics
- If the underlying “instantiated” formalism gives rise only to symmetric frameworks more sophisticated semantics notions do not make any difference
- Limited utility of semantics notions or limited expressiveness of the underlying formalism?

A bit less naive: stage semantics

- Stage semantics prescribes as extensions the conflict free sets with maximal range, i.e. those maximizing the union of the extension with the arguments it attacks (attack direction "partially counts")

DEFINITION 32 *Let $AF = (Ar, att)$ be an argumentation framework. A* stage extension *of AF is a conflict-free set $Args \subseteq Ar$ where $Args \cup Args^+$ is maximal (w.r.t. set inclusion) among all conflict-free sets.*

- Some "extremely naive" behavior of naive semantics are avoided, but there are still cases where unattacked arguments are excluded from some extensions

A bit less naive: stage semantics

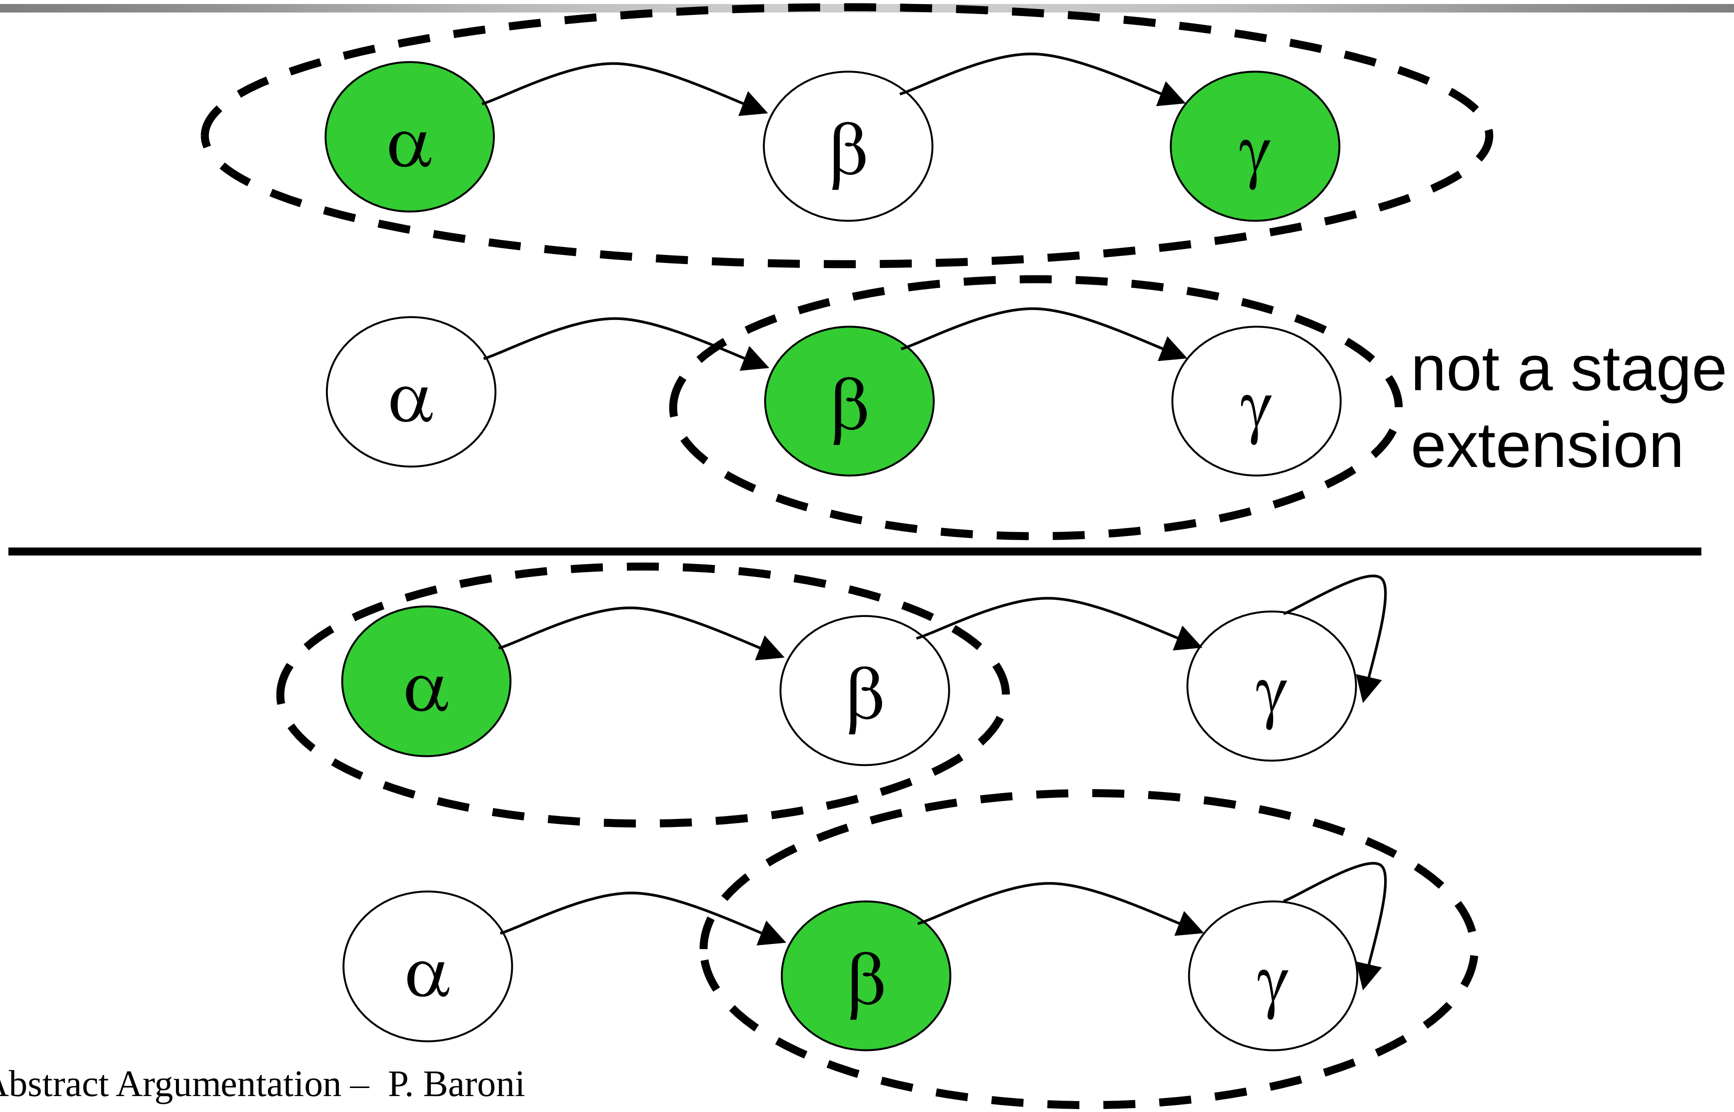

Defense based semantics

- Dung's paper introduces several notions and semantics all based on the notion of defense (admissibility principle) + conflict-freeness
- The direction of attacks counts!

Definition 6.

(1) An argument $A \in AR$ is *acceptable* with respect to a set S of arguments iff for each argument $B \in AR$: if B attacks A then B is attacked by S.[4]

(2) A conflict-free set of arguments S is *admissible* iff each argument in S is acceptable with respect to S.

Complete semantics

- Adds to admissibility the reinstatement principle: if you defend an argument you have to include it

Definition 23. An admissible set S of arguments is called a *complete extension* iff each argument, which is acceptable with respect to S, belongs to S.

- The empty set is complete only if there are no unattacked arguments
- Completeness does not imply I-maximality: a complete extension can be a proper subset of another one (obtained adding some self-defending set of arguments)

Admissibility and completeness

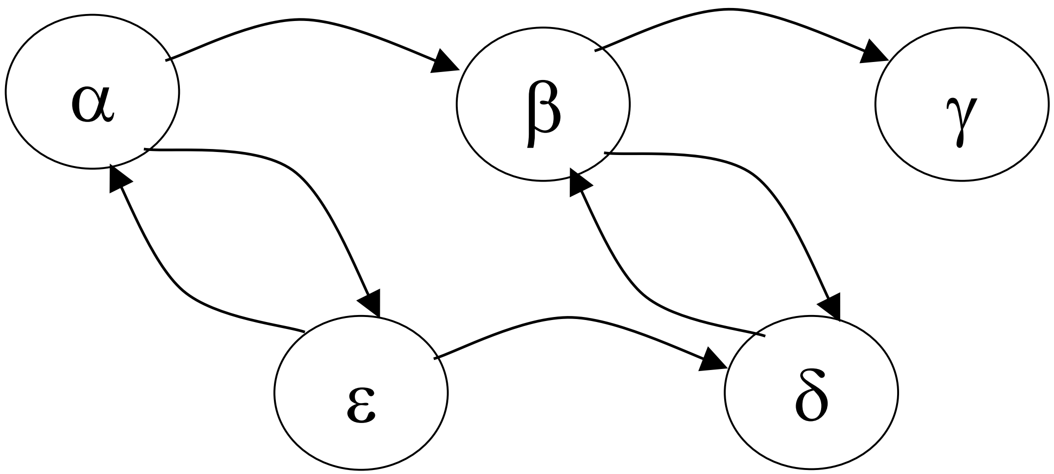

Admissible sets: { {}, {α}, {ε}, {α,γ}, {α,δ}, {ε,β}, {α,γ,δ} }

Complete extensions: { {}, {ε,β}, {α,γ,δ} }

Admissibility and completeness

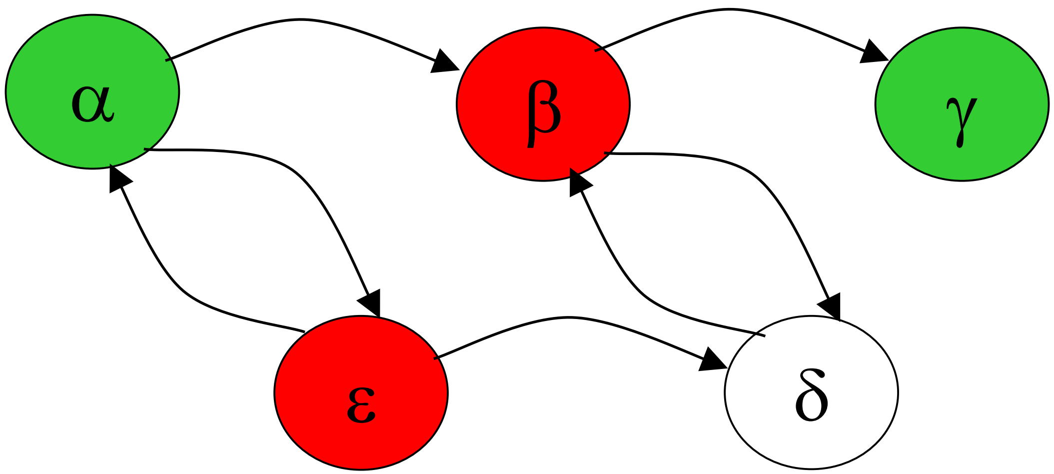

{α,γ} is admissible but not complete

Admissibility and completeness

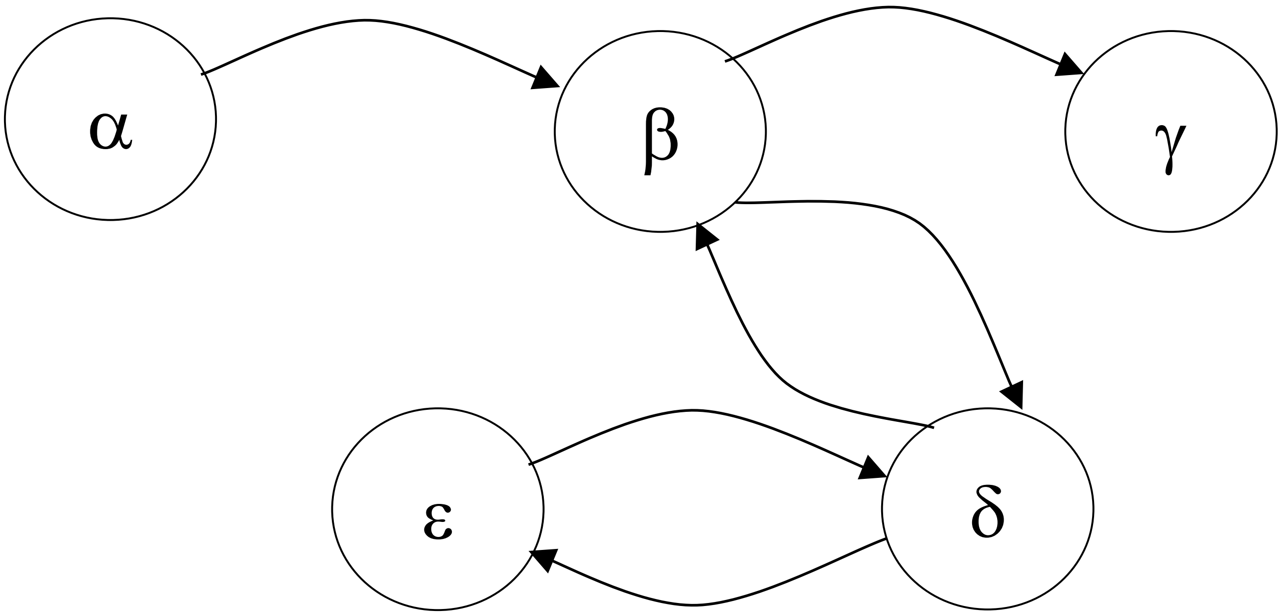

Admissible sets:

{ {}, {α}, {δ}, {ε}, {α,γ}, {α,δ}, {α,ε}, {γ,δ}, {α,γ,δ}, {α,γ,ε} }

Complete extensions: {{α,γ}, {α,γ,δ}, {α,γ,ε} }

Admissibility and completeness

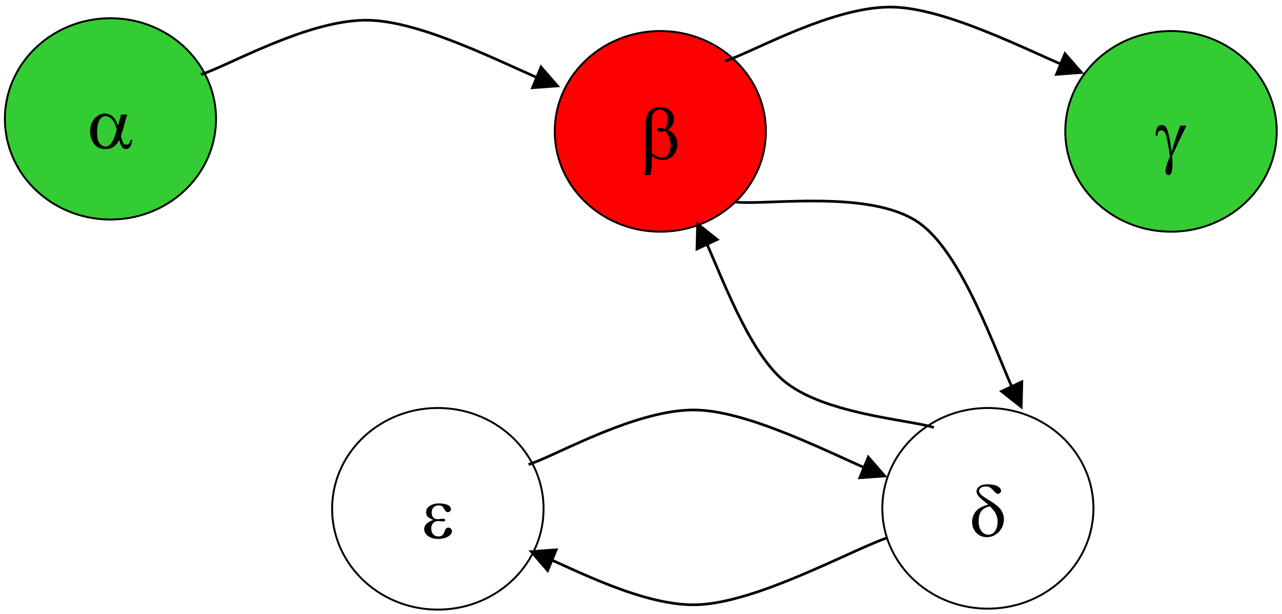

$\{\alpha,\gamma\}$ is admissible and complete

Admissibility and completeness

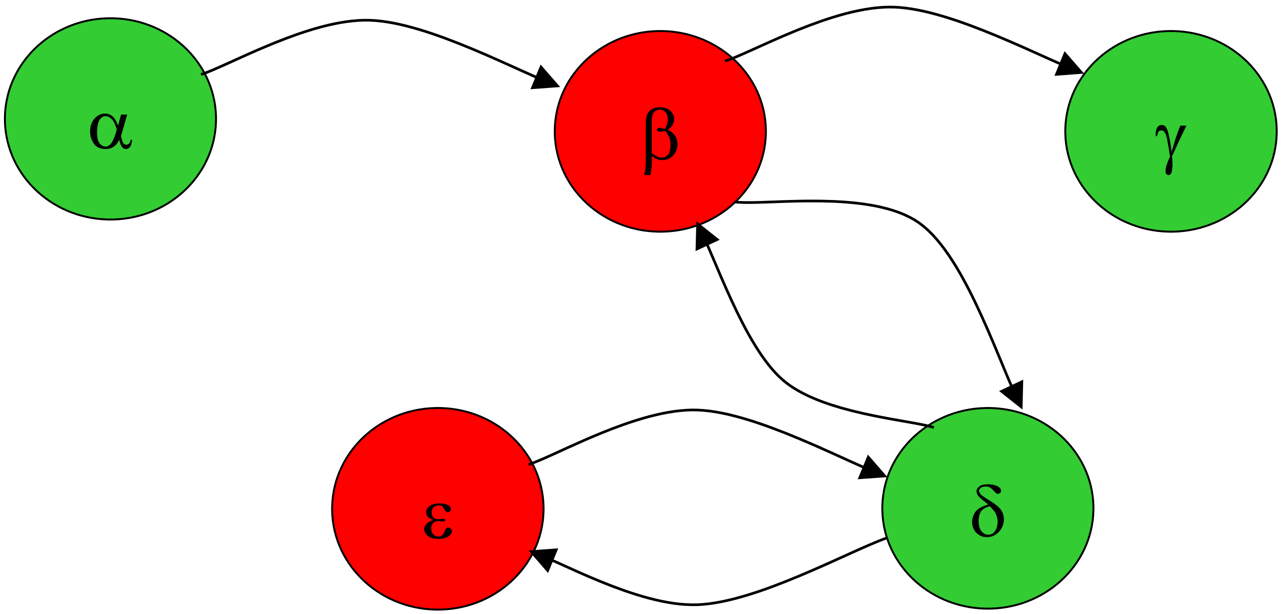

{α,γ,δ} is also admissible and complete

Grounded semantics: accepting only the unquestionable

- Given that complete extensions satisfy some basic properties, if one wants to be cautious s/he has to choose the smallest (wrt inclusion) complete extension
- Actually, it can be proved that there is a unique smallest complete extension, called grounded extension
- So grounded semantics belongs to the unique-status approach

Grounded semantics: unquestionable=strong defense

- Dung's definition of grounded semantics takes another route
- First, the characteristic function of a framework associates with a set S the arguments it defends

Definition 16. The *characteristic function*, denoted by F_{AF}, of an argumentation framework $AF = \langle AR, attacks \rangle$ is defined as follows:

$$F_{AF} : 2^{AR} \rightarrow 2^{AR},$$

$$F_{AF}(S) = \{A \mid A \text{ is acceptable with respect to } S\}.$$

Grounded semantics: unquestionable=strong defense

- The characteristic function is well-behaved

Lemma 19. F_{AF} *is monotonic (with respect to set inclusion).*

- This entails that it has a least fixed point, which, by definition, is called grounded extension

Definition 20. The *grounded extension* of an argumentation framework AF, denoted by GE_{AF}, is the least fixed point of F_{AF}.

- It can be proved that the least fixed point of the characteristic function coincides with the least complete extension

Grounded semantics: unquestionable=strong defense

- This has a rather intuitive counterpart: start computing the arguments defended by the empty set, repeat the computation (i.e. find the arguments defended by them), and so on, until you reach a fixed point
- For finitary frameworks the grounded extension can be obtained as:
  $F_{AF}(F_{AF}(\ldots F_{AF}(\emptyset)\ldots)$

PROPOSITION 7 *The grounded extension of any finitary argumentation framework is equal to* $\cup_{i=1,\ldots,\infty} F^i(\emptyset)$, *where* $F^1(\emptyset) = F(\emptyset)$ *and for* $i > 1$ $F^i(\emptyset) = F(F^{i-1}(\emptyset))$.

Computing the grounded extension

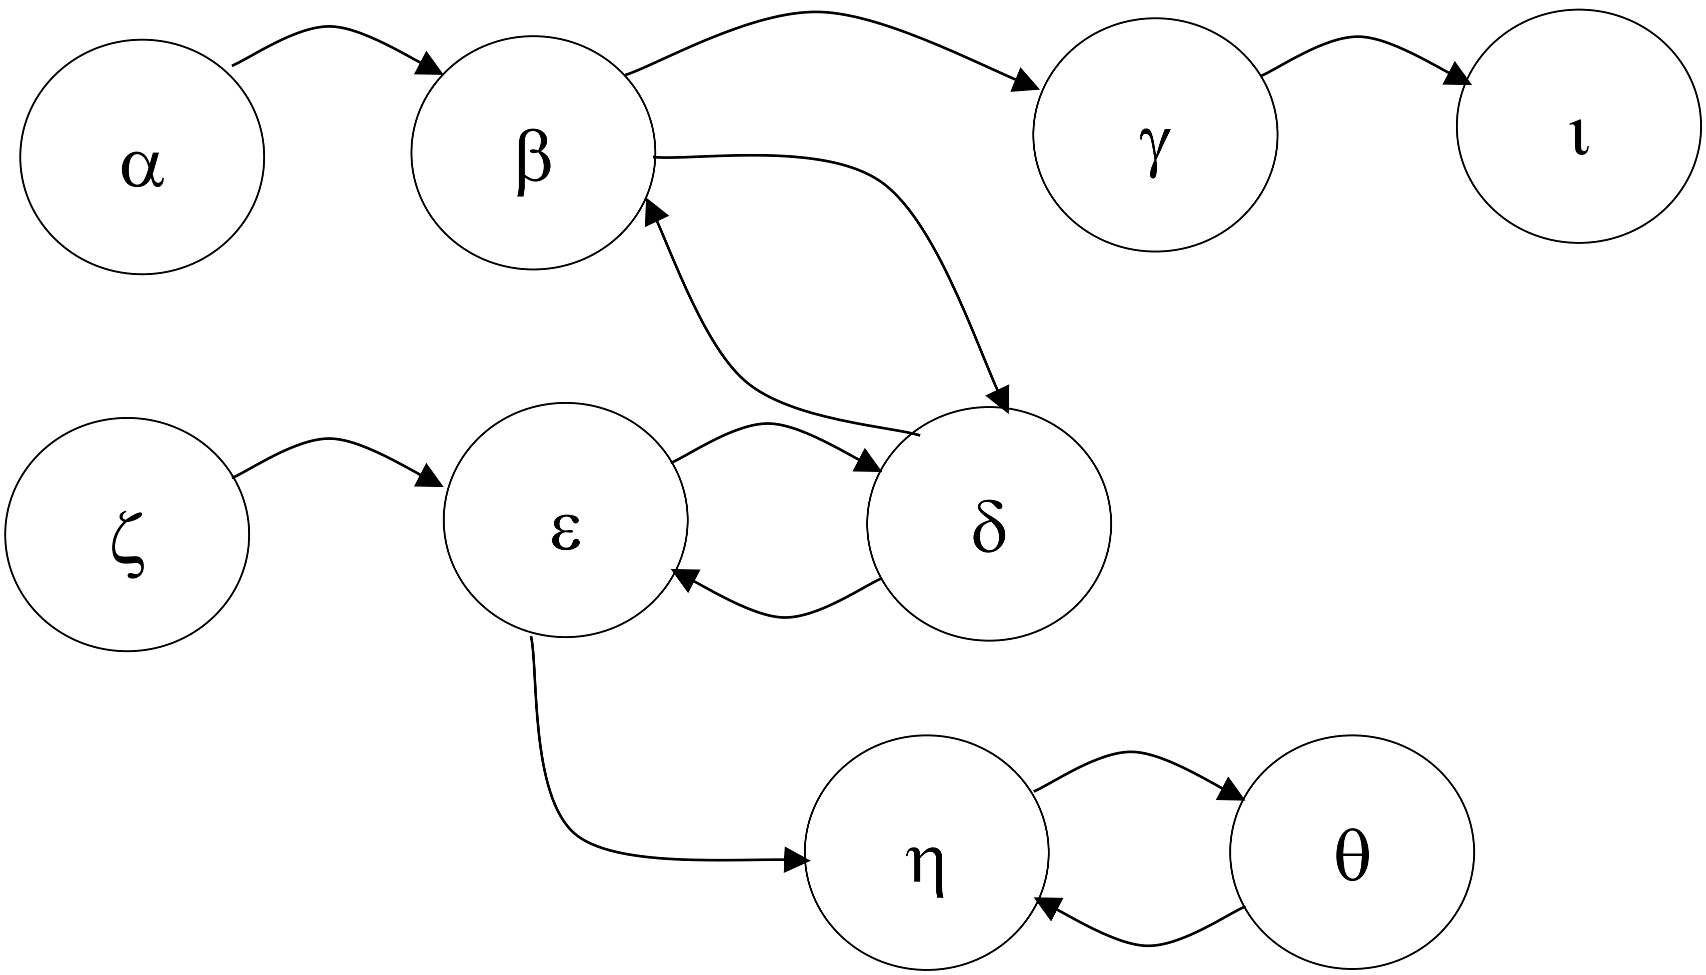

Computing the grounded extension

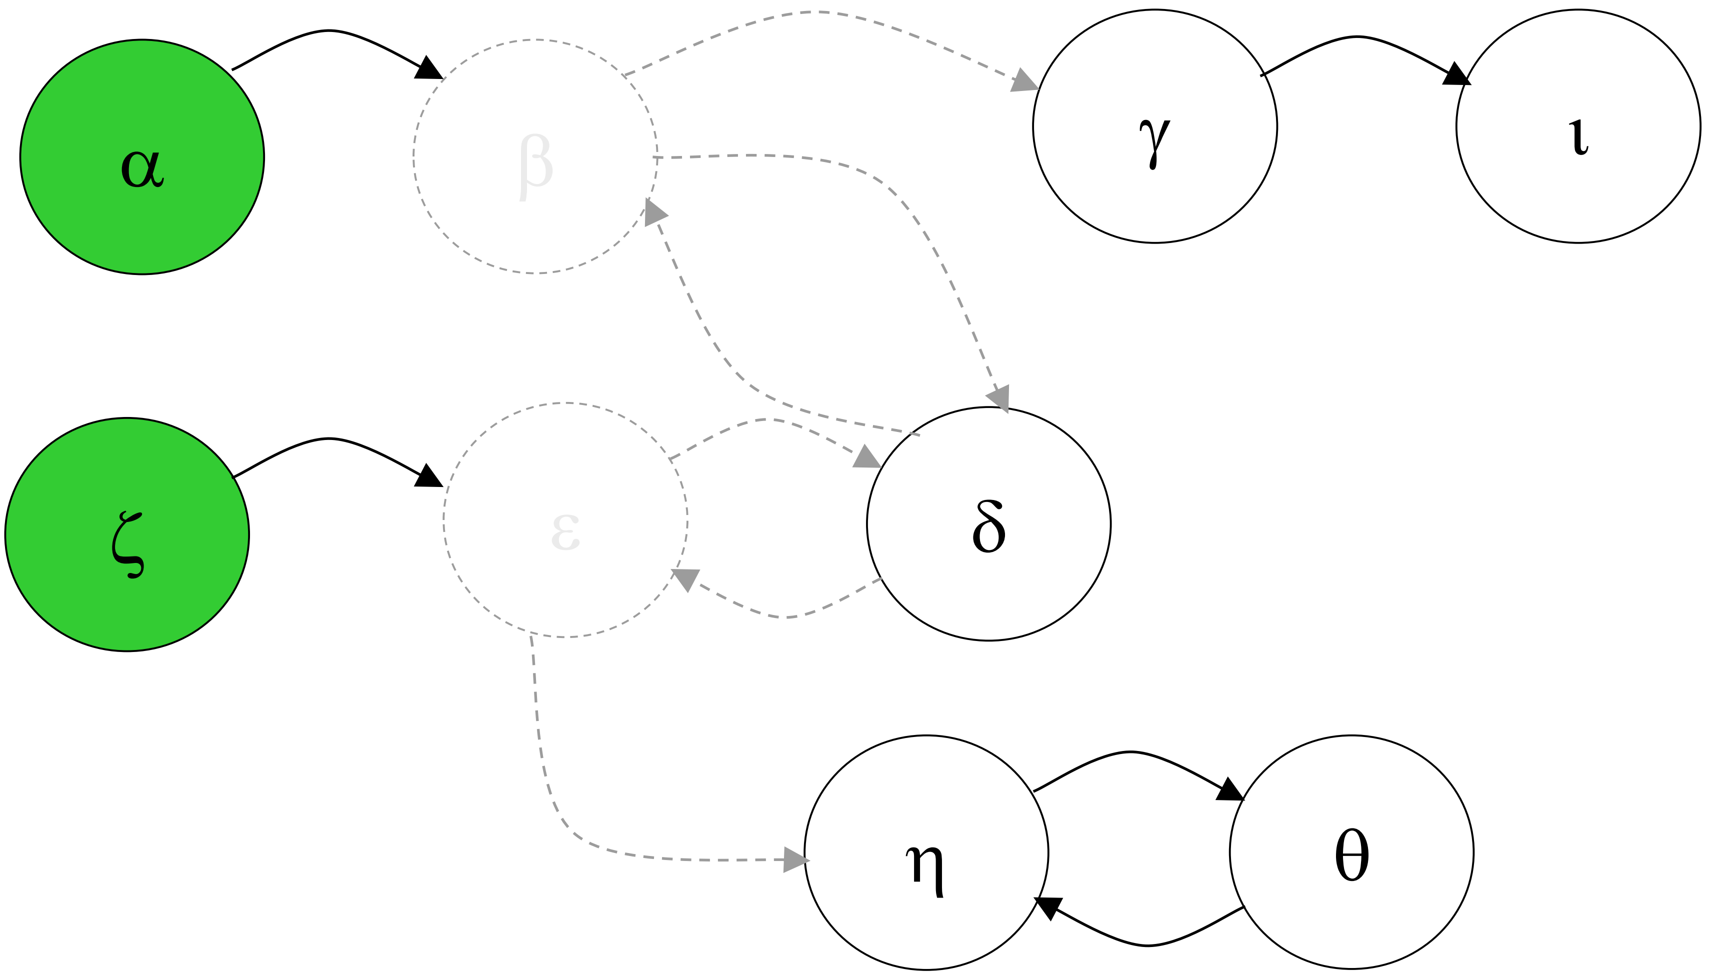

Computing the grounded extension

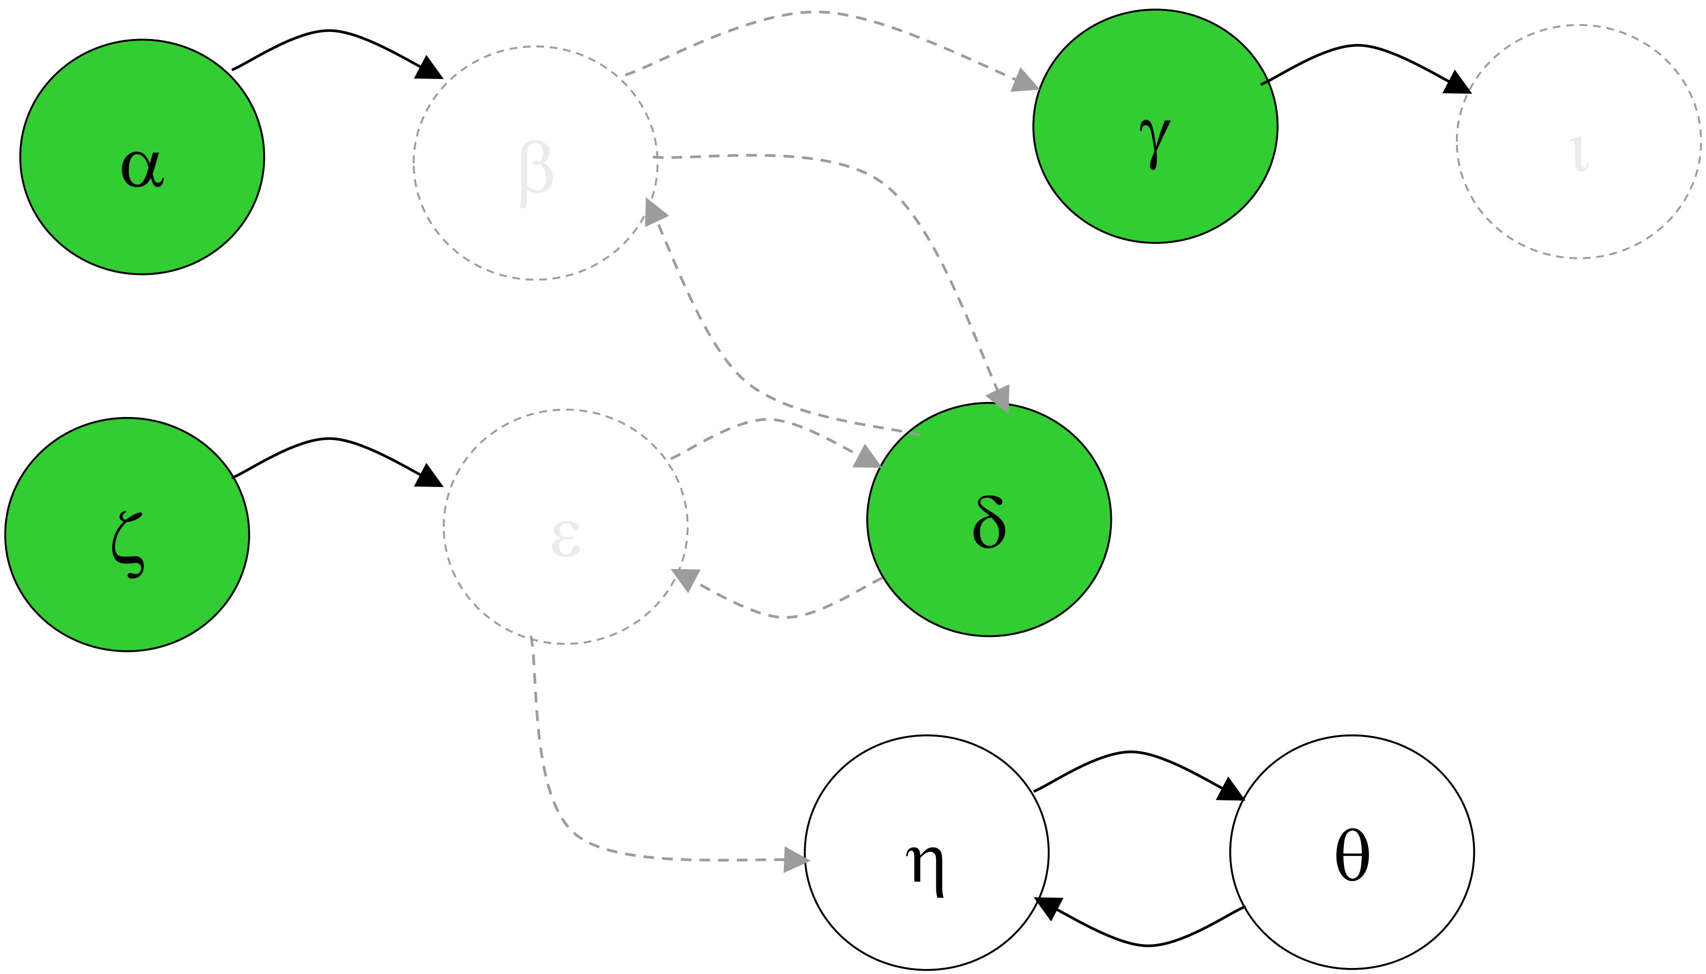

Computing the grounded extension

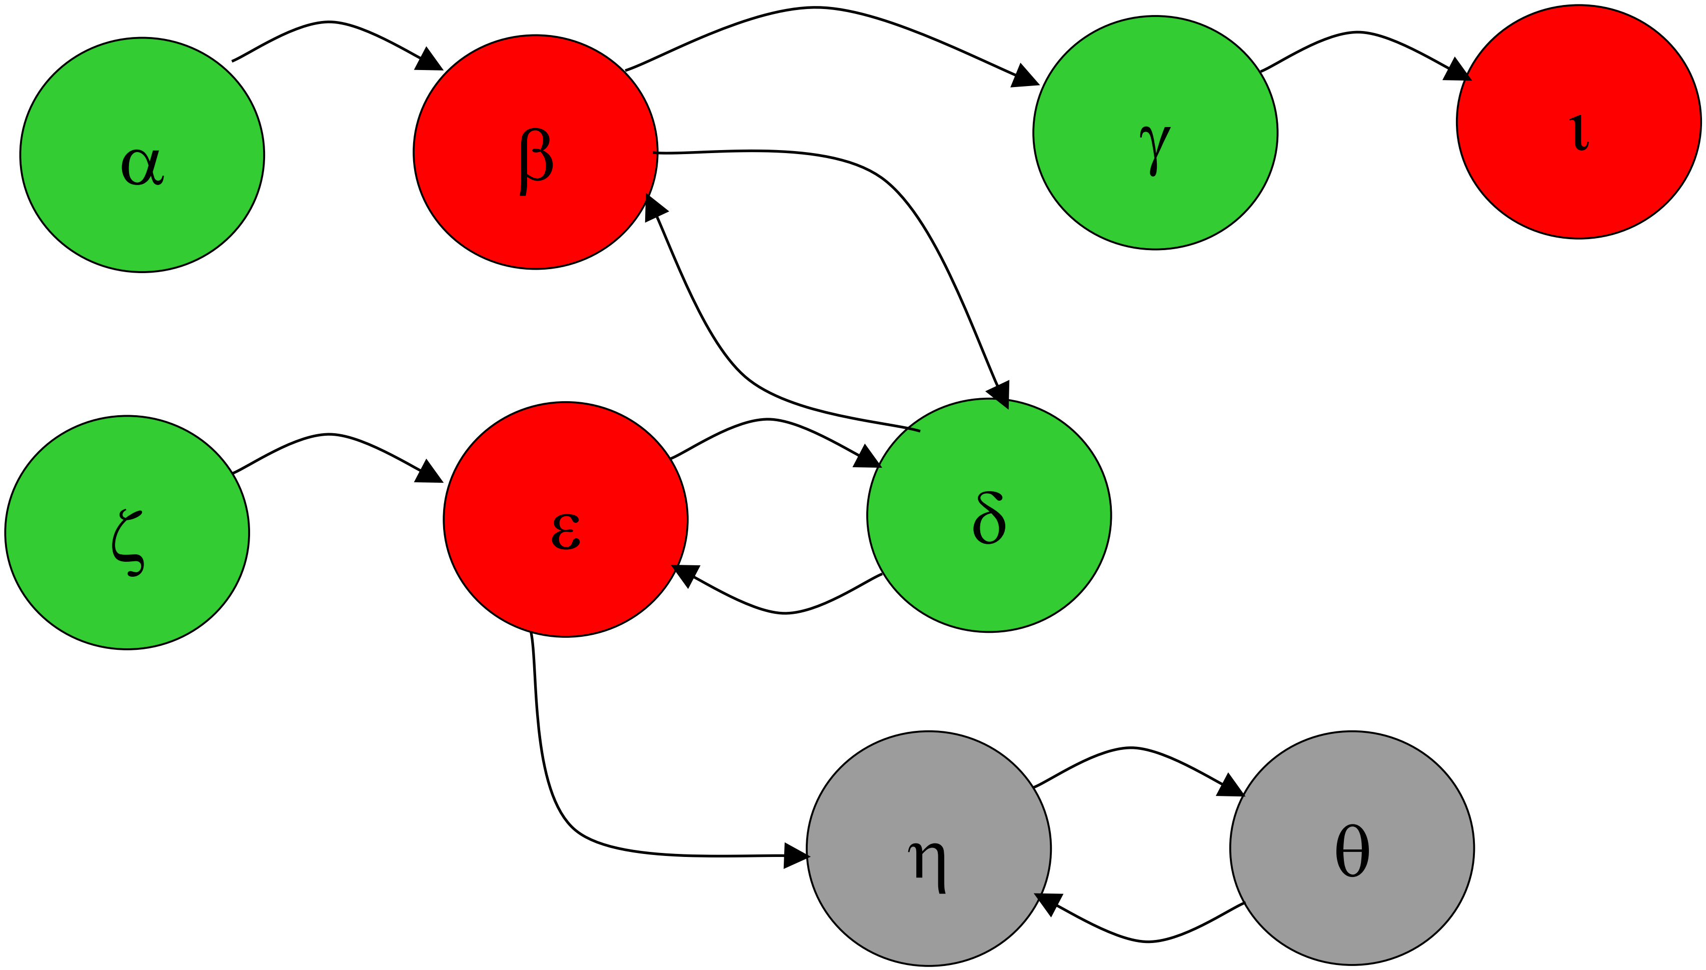

Horror vacui: stable semantics

- Stable semantics prescribes that any argument is either in the extension or attacked by the extension
- In the labelling version: no argument is undecided

Definition 13. A conflict-free set of arguments S is called a *stable extension* iff S attacks each argument which does not belong to S.

- In general there are several stable extensions
- Any stable extension is admissible and complete (hence it includes the grounded extension)

A stable extension

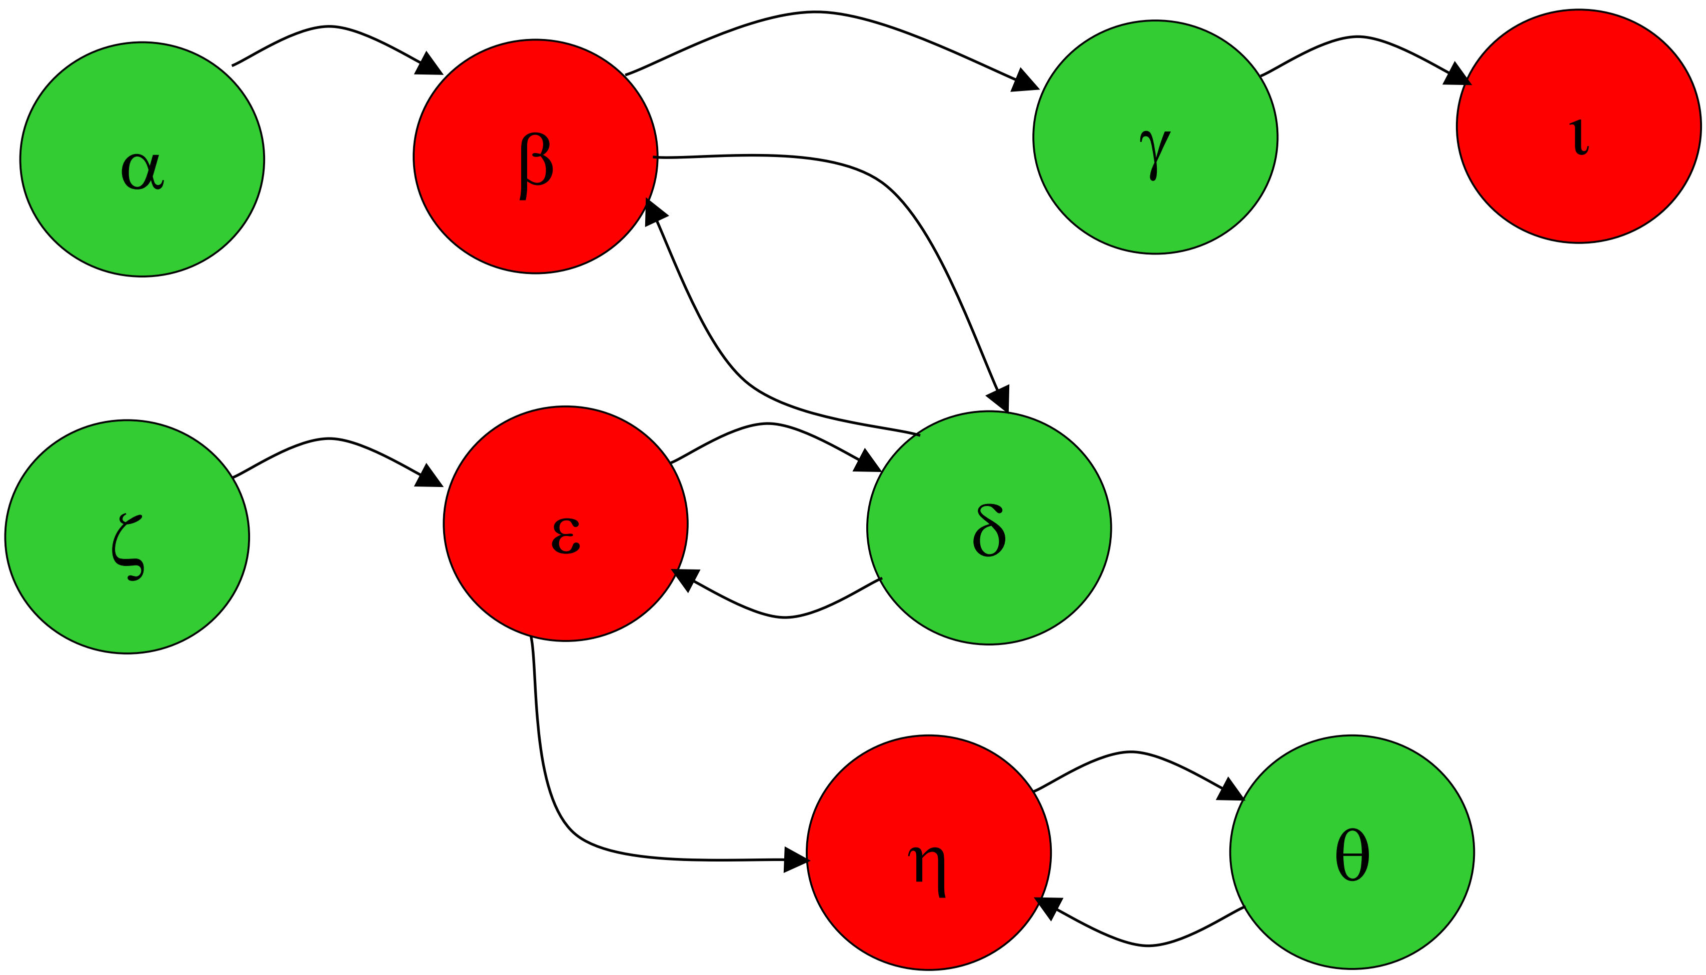

Another stable extension

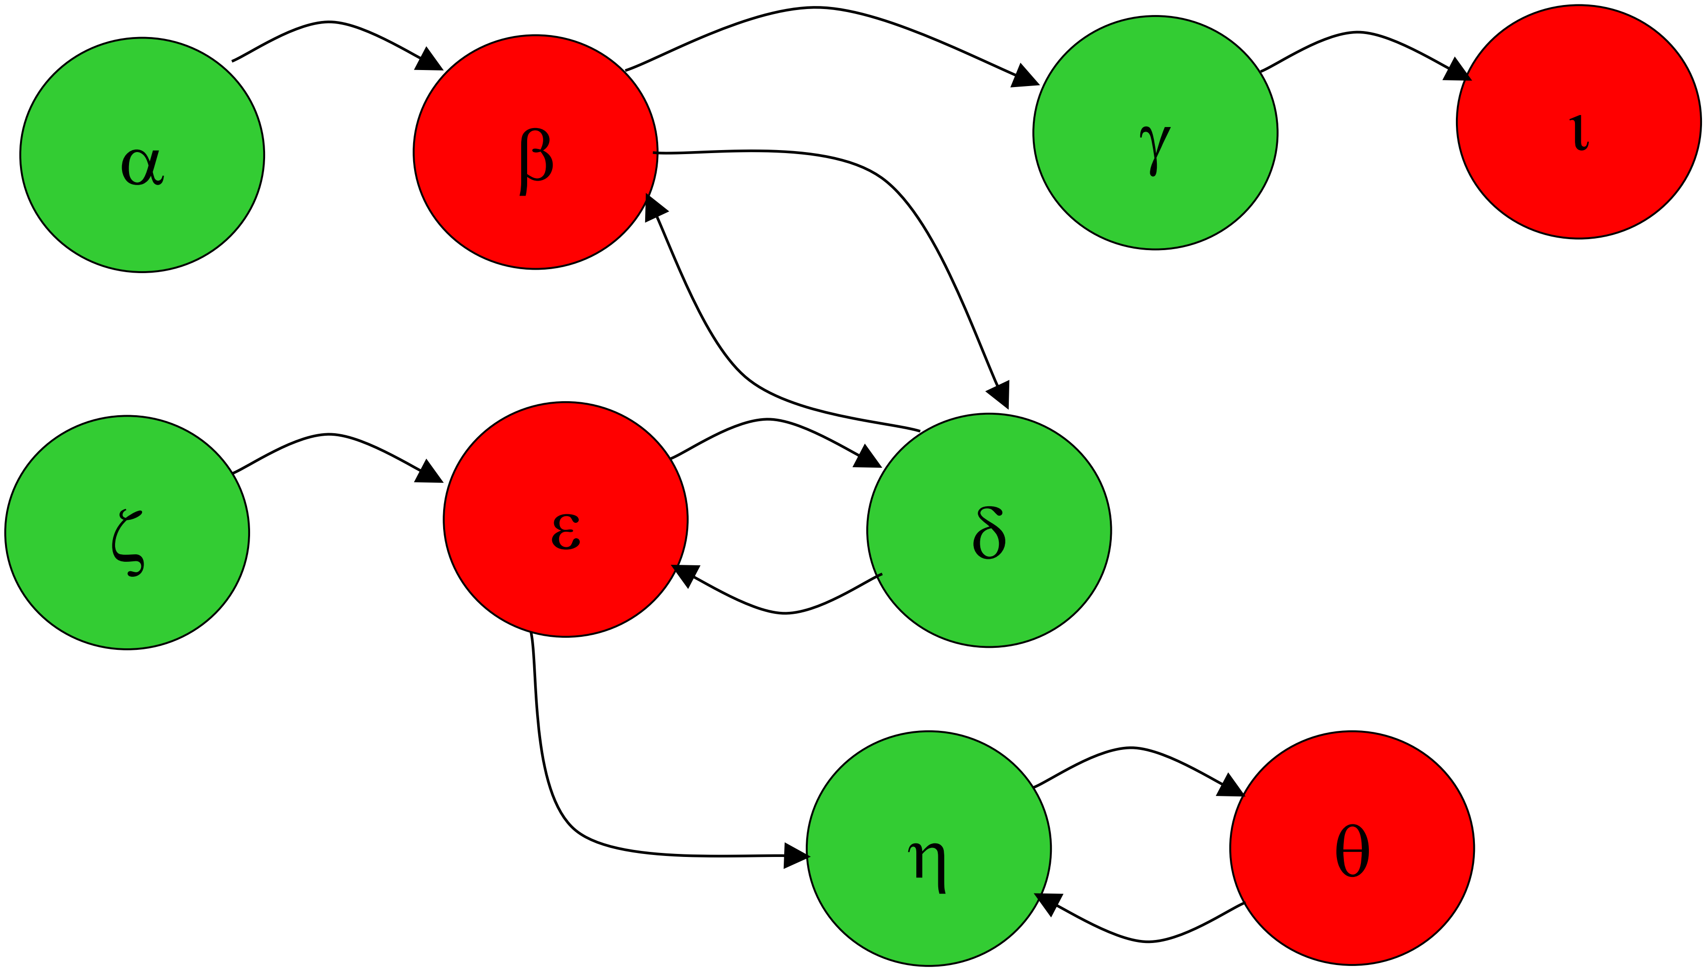

Being bold is not always possible

- The very strong requirement posed by stable semantics in some cases is not satisfiable
- The empty set is not the "most basic" extension (unless the framework is empty too)
- There are frameworks where no stable extensions exist
- The absence of odd-length cycles is a sufficient condition for existence of stable extensions

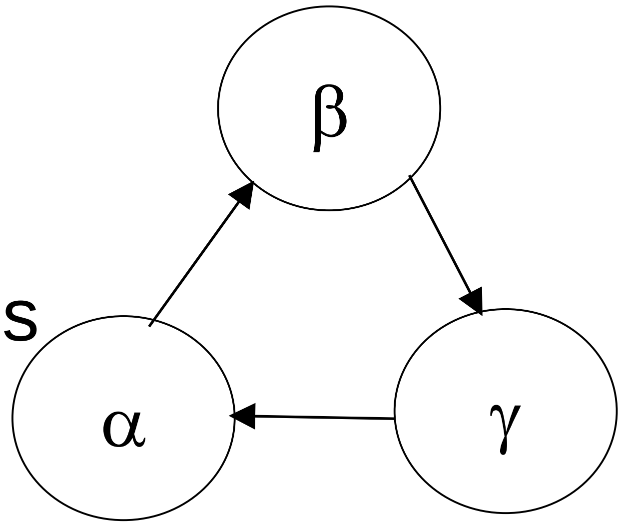

Being bold is not always possible

- A limited (possibly isolated) part of the framework may prevent the existence of extensions for the whole framework
- Stable semantics is not directional

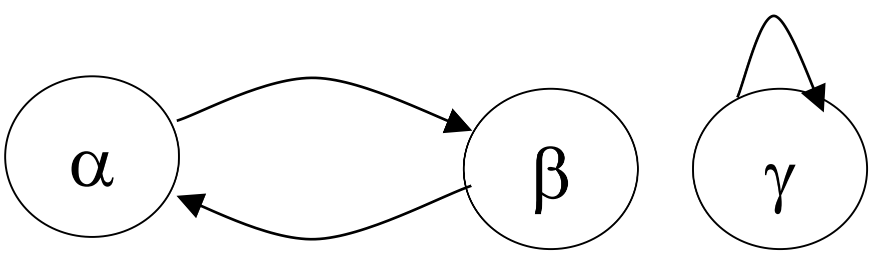

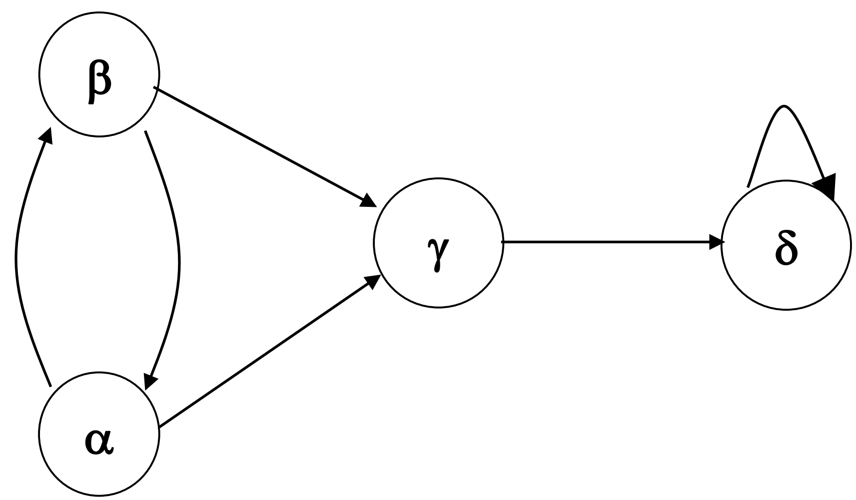

Accept as much as you can defend: preferred semantics

- Admissibility + maximality = preferred semantics

Definition 7. A *preferred extension* of an argumentation framework AF is a maximal (with respect to set inclusion) admissible set of AF.

- From a labelling perspective this corresponds to maximizing IN without banning UNDEC
- Preferred extensions always exist (in general many)
- Any preferred extension is complete (hence it includes the grounded extension)
- Stable extensions are preferred, not viceversa

Preferred vs. stable semantics

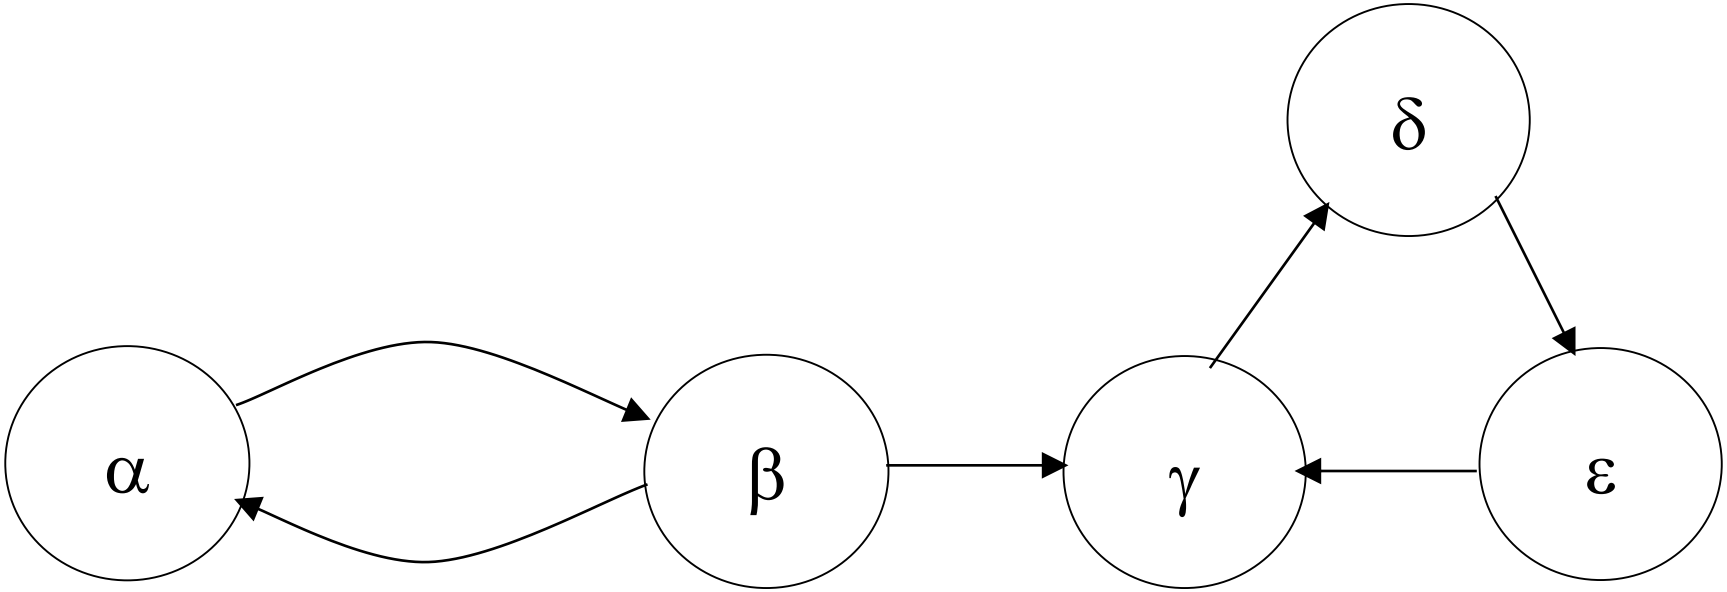

Preferred vs. stable semantics

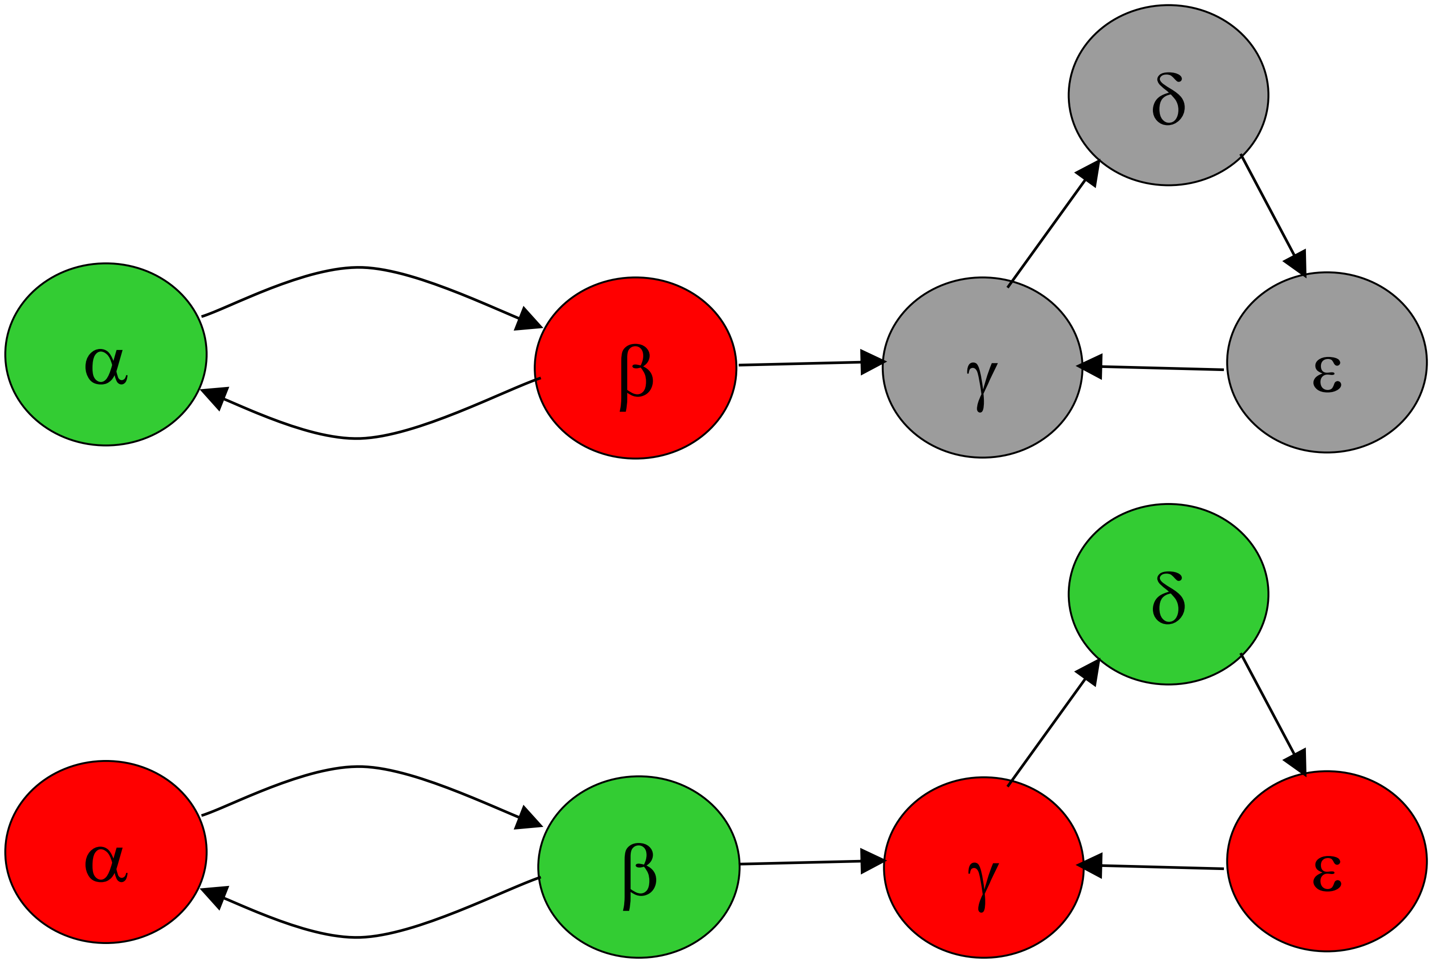

Minimizing indecision: semi-stable semantics

- Maximizing acceptance does not mean minimizing indecision: to this purpose you should maximize the union of IN and OUT
- Stable semantics achieves this implicitly by banning UNDEC, semi-stable semantics does this explicitly without banning UNDEC

DEFINITION 27 *Let $AF = (Ar, att)$ be an argumentation framework. A* semi-stable extension *of AF is a complete extension $\mathcal{A}rgs$ where $\mathcal{A}rgs \cup \mathcal{A}rgs^{+}$ is maximal (w.r.t. set inclusion) among all complete extensions.*

Minimizing indecision: semi-stable semantics

- Semi-stable extensions always exist (in general many)
- When stable extensions exist, semi-stable extensions coincide with stable extensions
- Semi-stable extensions are preferred, not viceversa
- Semi-stable semantics is not as “fragile” as stable semantics, but it is non directional too

Single status from multiple status

- A natural way to obtain a single status semantics from a multiple status semantics consists in using the intersection Int of the extensions
- For the four Dung's semantics Int is a (possibly strict) superset of the grounded extension
- Int is conflict-free but not admissible in general
- Adding the requirement of admissibility to the one of inclusion in all extensions gives rise to a generic “scheme” for semantics definition which has been explicitly considered in two cases

Ideal and eager semantics

- The ideal extension is the maximal (wrt inclusion) admissible set included in all preferred extensions
- The eager extension is the maximal (wrt inclusion) admissible set included in all semi-stable extensions
- Both are supersets of the grounded extension
- The eager extension is a superset of the ideal extension

A less cautious single status

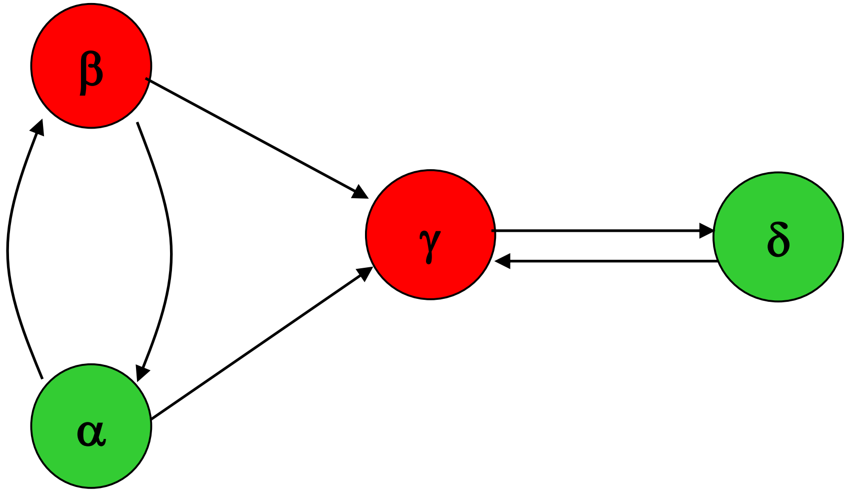

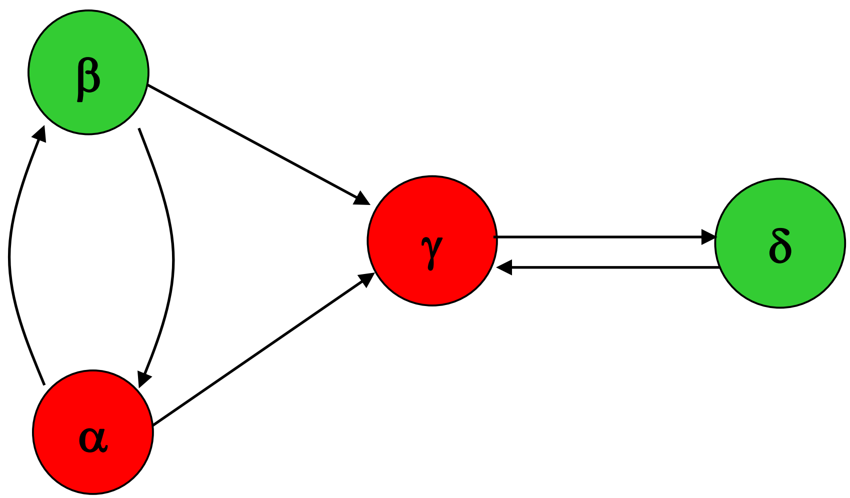

The ideal/eager extension is {δ}, while the grounded extension is empty

A less cautious single status

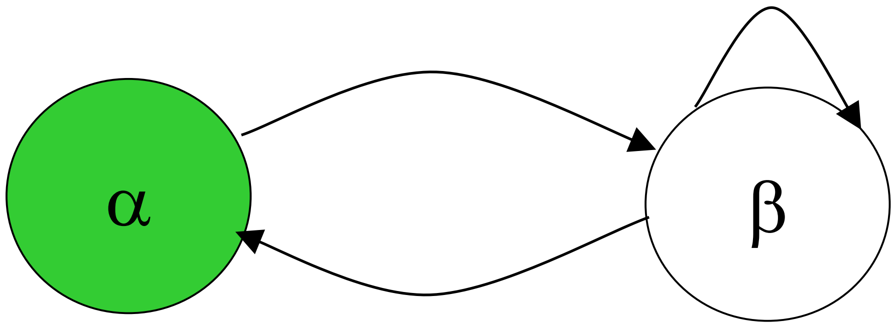

The ideal/eager extension is {α}, while the grounded extension is empty

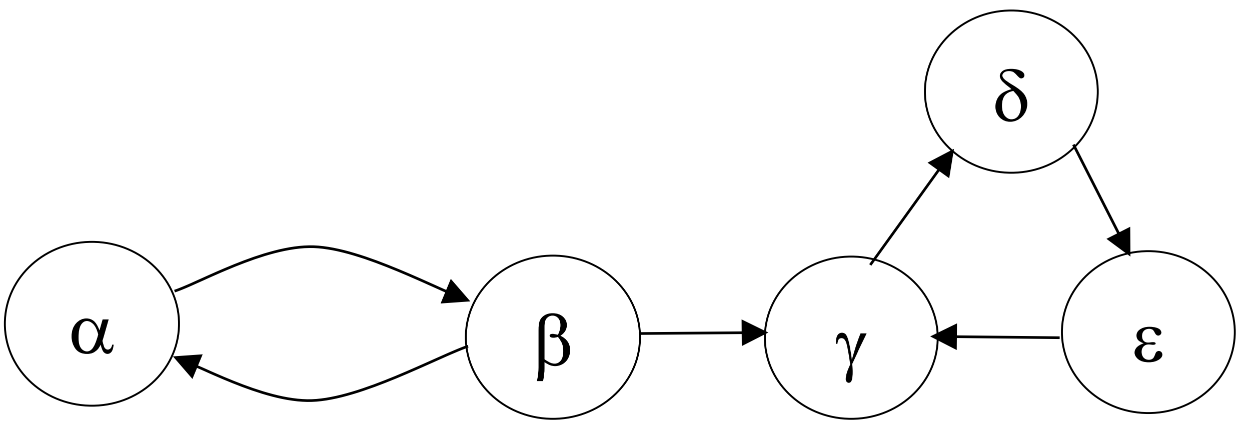

The eager extension is {β,δ}, the ideal extension is empty

Roadmap

- What's in the words?
- From the Big Bang to now
- Dung's framework
- A "conflict calculus": argumentation semantics
- "Basic" semantics properties
- A catalogue of semantics
- Properties of semantics (so far)

Universal and almost universal properties

- Conflict freeness and language independence are satisfied by all semantics
- I-maximality is satisfied by single-status semantics and is directly implied by several multiple-status definitions, only complete semantics is not I-maximal

"Conflict-free only" semantics

- Naive and stage extensions can exclude unattacked arguments and include undefended arguments hence they can not satisfy admissibility, strong defense, reinstatement, weak reinstatement, and directionality
- They only satisfy CF-reinstatement

“Complete-based” semantics

- Grounded, stable, preferred, semi-stable, ideal and eager extensions are complete extensions
- This implies that they satisfy admissibility and reinstatement (including weaker forms)
- Strong defense is only satisfied by grounded semantics

Directionality

- Disproving directionality is immediate for stable, semi-stable and eager semantics
- Proving directionality is less immediate (but possible) for the other complete-based semantics

A synopsis

	NA	STA	CO	GR	STB	PR	SST	ID	EAG
CF-principle	Y	Y	Y	Y	Y	Y	Y	Y	Y
Language indip.	Y	Y	Y	Y	Y	Y	Y	Y	Y
I-maximality	Y	Y	N	Y	Y	Y	Y	Y	Y
Admissibility	N	N	Y	Y	Y	Y	Y	Y	Y
Strong defense	N	N	N	Y	N	N	N	N	N
Reinstatement	N	N	Y	Y	Y	Y	Y	Y	Y
Weak reinstatement	N	N	Y	Y	Y	Y	Y	Y	Y
CF-reinstatement	Y	Y	Y	Y	Y	Y	Y	Y	Y
Directionality	N	N	Y	Y	N	Y	N	Y	N

Roadmap

- What's in the words?
- From the Big Bang to now
- Dung's framework
- A "conflict calculus": argumentation semantics
- "Basic" semantics properties
- A catalogue of semantics
- Properties of semantics (so far)
- Beyond Dung's framework

An abstraction of abstract argumentation

- Basic idea: a formalism where an assessment is produced for a set of entities based on their relations only
- In Dung's abstract argumentation:
  - » only one relation is considered: attack
  - » the assessment is symbolic (essentially 3-valued)
  - » the assessment is based on some assumptions (e.g. all arguments have the same initial status)
- The same basic idea can be applied going beyond some of these specific choices/assumptions

Extended frameworks

- Several extensions/variations of Dung’s framework have been considered in the literature
- They can be classified according to distinct (but not necessarily disjoint) lines:
  - » representation of additional notions
  - » finer evaluations
  - » more articulated notions of attack
  - » flexible relations
  - » …

Additional notions: support

- Basic intuition: arguments are not just “conflictables” they may also “help” or “support” each other
- A relation of support should parallel the one of attack
- Is support as “abstractable” as conflict?
- (It seems that) we can define general conflict management mechanisms independently of the underlying meaning of conflict
- Is this analogously possible independently of the underlying meaning of support?

Different meanings of support

- Support as defense: A supports B if A defends B from an attack
- Support as derivation (deductive support): A supports B if acceptance of A implies the acceptance of B
- Support as necessity: A supports B if acceptance of A is necessary for acceptance of B
- Evidential support: distinguishes "prima-facie" arguments (primitive evidences) from arguments requiring an underlying evidence as a basis

Bipolar argumentation frameworks

Formal definition of an abstract bipolar argumentation framework An *abstract bipolar argumentation framework* $<\mathcal{A}, \mathcal{R}_{def}, \mathcal{R}_{sup}>$ consists of a set $\mathcal{A}$ of arguments, a binary relation $\mathcal{R}_{def}$ on $\mathcal{A}$ called a *defeat relation* and another binary relation $\mathcal{R}_{sup}$ on $\mathcal{A}$ called a *support relation*: consider A_i and $A_j \in \mathcal{A}$, $A_i \mathcal{R}_{def} A_j$ (resp. $A_i \mathcal{R}_{sup} A_j$) means that A_i defeats A_j (resp. A_i supports A_j).

- It is required that the two relations have no intersection (an argument can not attack and support another one at the same time)
- A path of supports gives rise to indirect support
- A path of supports behind an attack gives rise to indirect (or supported) attack
- A single attack behind a path of supports gives rise to a secondary attack

Bipolar argumentation frameworks

- A lot of variants and discussions in the literature
- Possible mixing between the notions of argument and of conclusion of an argument in the informal intuitions about support
- However, bipolarity is quite natural and intuitive for humans
- As to argument assessment, traditional semantics definitions are modified to take into account both attacks and supports

Additional notions: preference-based AFs

- Basic idea: the attack relation may fail to capture the fact that arguments may have different “quality”

Dung’s definition of acceptability disregards the quality of the arguments. However, the force of an argument can be often estimated by considering the beliefs used to build this argument. For example according to the preferences which can exist between the beliefs, an argument can be more or less strong than another argument. Let’s consider the following example.

Example 2.3. Let $\langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework with $\mathcal{A} = \{A, B, C\}$ and $\mathcal{R} = \{(B, A), (C, B)\}$.

According to Dung, the set of acceptable arguments is $\{A, C\}$. Yet, if we know that the argument B is preferred to A and C according to a preference relation *Pref* between arguments, then the argument C defeats B but B defends itself against C in some sense. Hence, the argument B might be considered as acceptable and A might be rejected.

Additional notions: preference-based AFs

- The "quality" of arguments can be represent through a preference relation

Definition 3.2. A preference-based *argumentation framework (PAF)* is a triplet $\langle \mathcal{A}, \mathcal{R}, \mathit{Pref} \rangle$ where $\mathcal{A}$ is a set of arguments, $\mathcal{R}$ is a binary relation representing a defeat relationship between arguments, $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$, and Pref is a (partial or complete) preordering on $\mathcal{A} \times \mathcal{A}$.

- Attacks are "effective" only if the attacked argument is not preferred to the attacking one

Definition 3.3. Let A, B be two arguments of $\mathcal{A}$. B *attacks* A iff $B \;\mathcal{R}\; A$ and $\mathit{not}(A \gg_{\mathit{Pref}} B)$.

- A grounded-semantics-like notion of acceptability is defined accordingly

Additional notions: value-based AFs

- A more articulated way to assess the “extra-conflict” merits of arguments
- Arguments are associated with values
- There is no unique ordering of values, since they may count differently in different contexts
- A specific value ordering is called an audience
- For instance, in a political debate different audiences may assess differently arguments promoting public welfare wrt arguments promoting financial stability

Additional notions: value-based AFs

Definition 16. $\langle Args, \mathcal{R}, V, val, P\rangle$ is a value-based argumentation framework (*VAF*), where *val* is a function from *Args* to a non-empty set of values V, and P is a set $\{a_1, \ldots, a_n\}$, where each a_i names a total ordering (audience) $>_{a_i}$ on $V \times V$.

An *audience specific VAF* (*aVAF*) is a 5-tuple $\langle Args, \mathcal{R}, V, val, a\rangle$ where $a \in P$.

- Defeat (and hence any semantics evaluation) becomes an audience-specific notion

Definition 17. Let $\langle Args, \mathcal{R}, V, val, a\rangle$ be an *aVAF*. Then $A \in Args$ defeats$_a$ $B \in Args$ iff $(A, B) \in \mathcal{R}$ and it is not the case that $val(B) >_a val(A)$.

We say that $(A, B) \in \mathcal{R}_a$ iff A defeats$_a$ B.

- In fact, an audience induces a specific argumentation framework

Acceptance in value-based AFs

- Objective acceptance: a stronger skeptical justification
- An argument is objectively acceptable if it is skeptically justified in any audience-specific AF (i.e. for any audience)
- Subjective acceptance: a weaker credulous justification
- An argument is subjectively acceptable if it is credulously justified in an audience-specific AF (i.e. for at least one audience)

Finer evaluation of arguments

- Dung's framework encompasses a qualitative assessment of argument justification status (essentially three-valued)
- The use of numbers or rankings to provide a finer evaluation of arguments has been considered in several contexts and with various flavors

Argument “strength”

- According to Pollock, the strength of an argument is “the degree of justification it would confer on its conclusion” in absence of defeaters
- Here strength is seen as an a priori property of the argument, determined at the moment of its construction
- Weakest link principle: an argument can not be stronger than its subarguments
- Initial argument strengths are used to determine a final numerical evaluation of arguments, on the basis of the attack (and possibly support) relations

Argument "strength"

- The idea of argument strength has been considered in several variants (often at a non abstract level) in the literature
- No standard reference model
- Many open questions (with non univocal answer): suitability of a probabilistic-style treatment, combination operators, the role of accrual ...

Gradual (numerical) argumentation

- Arguments may have an initial "score"
- Based on the initial score and on the relations with other arguments (attacks and/or supports) a final score is derived for each argument
- A wide variety of proposals is available in the literature
- Also the range of properties and principles considered is wider than for traditional semantics
- Some attempts to provide some synthetic view have also been undertaken

Ranking-based semantics

- The produced evaluation consists of a ranking, namely a (possibly partial) order relation on arguments
- Finer than traditional 3-valued assessments but without requiring the use of numbers
- Numerical evaluations induce a ranking, but it is possible to produce rankings without associating numbers to arguments

Mind the meaning

- Many ways of having numbers "attached" to argumentation frameworks have been considered
- The meaning of these numbers may vary a lot: probabilities (of different nature), fuzzy sets, something else
- Some confusion is possible and attention should be paid (as always) to the exact intended meaning of the used numbers
- Not every appearance of number associated with arguments is a form of gradual semantics or strength evaluation

Equational approach

- In the equational approach to argumentation networks, an argumentation framework induces a system of equations which are parametric wrt a function f

Let (S, R_A) be a Dung network. So $R_A \subseteq S^2$ is the attack relation. We are looking for a function $\mathbf{f} : S \mapsto [0, 1]$ assigning to each $a \in S$ a value of $0 \leq \mathbf{f}(a) \leq 1$ such that the following holds.

(1) $(S, R_A, \mathbf{f})$ satisfies the following equations for some family of functions $\{\mathbf{h}_a\}, a \in S$:

(a) If a is not attacked (i.e. $\neg\exists x(xR_Aa)$) then $\mathbf{f}(a) = 1$.

(b) If $x_1, \ldots, x_n$ are all the attackers of a (i.e. $\bigwedge_{i=1}^{n} x_iR_Aa \wedge \forall y(yR_Aa \rightarrow \bigvee_{i=1}^{n} y = x_i))$ then we have that $\mathbf{f}(a) = \mathbf{h}_a(\mathbf{f}(x_1), \ldots, \mathbf{f}(x_n))$.

Equational approach

- Several choices for f can be considered

(1) $Eq_{\text{inverse}}(\mathbf{f})$ $\qquad \mathbf{f}(a) = \prod_i (1 - \mathbf{f}(x_i)).$

(2) $Eq_{\text{geometrical}}(\mathbf{f})$ $\qquad \mathbf{f}(a) = \dfrac{[\prod_i (1 - \mathbf{f}(x_i))]}{[\prod_i (1 - \mathbf{f}(x_i)) + \prod_i \mathbf{f}(x_i)]}.$

- In case of cycles the solutions are fixed points of the function f
- No notion of initial (nor final) argument strength involved, rather a form of “numerical semantics” reflecting the structure of the framework

Equational approach: numbers but not for strength

The numbers we get for the nodes as the result of the equation have no qualitative meaning beyond the fact that they indicate which points are in, out, or undecided. They do not indicate strength of argument and therefore cannot be used to indicate support.

To make this point crystal clear, consider again Figure 9. We got solutions for this figure, namely:

$$\mathbf{f}(a) = 1 - \frac{\sqrt{2}}{2},$$

$$\mathbf{f}(b) = \sqrt{2} - 1,$$

$$\mathbf{f}(c) = 2 - \sqrt{(2)}.$$

These numbers say that the nodes a, b, c are all undecided and may indicate the types of loops we have in the original graph (though we have not developed such a theory yet, this is postponed to a subsequent paper. See, however, Gabbay (2012a,b), where we discuss the methodology of loops). They are not indicating strength of arguments!

Weighted argumentation frameworks

- In weighted argumentation frameworks weights are on attacks rather than on arguments

Definition 4. A *weighted argument system* is a triple $W = \langle \mathcal{X}, \mathcal{A}, w \rangle$ where $\langle \mathcal{X}, \mathcal{A} \rangle$ is a Dung-style abstract argument system, and $w : \mathcal{A} \rightarrow \mathbb{R}_{>}$ is a function assigning real valued weights[3] to attacks.

- The weights represent the "strength of the attack" or the amount of inconsistency it carries
- An "inconsistency budget" β defines the amount of inconsistency one is prepared to tolerate

Weighted argumentation frameworks

- The idea is that one can "ignore" any set of attacks whose total weight is not greater than the inconsistency budget β and define β–compatible extensions accordingly

Definition 5. Let $\langle \mathcal{X}, \mathcal{A}, w \rangle$ be a weighted argument system. Given $R \subseteq \mathcal{A}$,

$$wt(R, w) = \sum_{\langle \alpha_1, \alpha_2 \rangle \in R} w\big(\langle \alpha_1, \alpha_2 \rangle\big)$$

The function $sub(\cdots)$, which takes an attack relation $\mathcal{A}$, weight function $w : \mathcal{A} \to \mathbb{R}_{>}$, and inconsistency budget $\beta \in \mathbb{R}_{\geqslant}$, returns the set of subsets R of $\mathcal{A}$ whose total weight does not exceed β, i.e.,

$$sub(\mathcal{A}, w, \beta) = \big\{R\colon\ R \subseteq \mathcal{A}\ \&\ wt(R, w) \leqslant \beta\big\}$$

We now use inconsistency budgets to introduce weighted variants of the semantics introduced for abstract argument systems, above.

Definition 6. Given a weighted argument system $\langle \mathcal{X}, \mathcal{A}, w \rangle$, let $\sigma : 2^{\mathcal{X}} \to \{\top, \bot\}$. For $\beta \in \mathbb{R}_{\geqslant}$, the subset $\mathcal{E}^{\mathrm{WT}}_{\sigma}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$ of $2^{\mathcal{X}}$ is given as

$$\mathcal{E}^{\mathrm{WT}}_{\sigma}\big(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta\big) = \big\{S \subseteq \mathcal{X}\colon\ \exists R \in sub(\mathcal{A}, w, \beta)\ \&\ S \in \mathcal{E}_{\sigma}\big(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle\big)\big\}$$

Weighted argumentation frameworks

- Weighted argumentation frameworks can be seen as a generalization of the approaches where attacks can be suppressed (or resolved) for various reasons like:
  - » Preference-based AFs
  - » Value-Based AFs
  - » AFs with attacks to attacks
  - » Resolution-based semantics

Collective attacks

- The binary relation of attack can not express situations of non-binary incompatibility
- Consider a situation where you can choose arbitrarily two out of three items (e.g. three people A,B,C, wanting to have a ride on a tandem)
- Clearly it is impossible that the three facts (A,on), (B,on) and (C,on) hold together but they are not pairwise incompatible
- If you need to capture this kind of situations you need collective attacks

Vreeswijk's abstract argumentation systems

- Vreeswijk's abstract argumentation systems are not as abstract as Dung's argumentation framework
- They include argument structure and inference rules
- In this context a set of arguments (rather than a single argument) may defeat an argument

Definition 4.2 (*Defeater*). Let P be a base set, and let σ be an argument. A set of arguments Σ is a *defeater* of σ if it is incompatible with this argument and not undermined by it; in this case σ is *defeated* by Σ, and Σ *defeats* σ. Σ is a *minimal defeater* of σ if all its proper subsets do not defeat σ.

Nielsen and Parsons' sets of attacking arguments

- Dung's framework is extended to encompass attacks from sets of arguments

Definition 1 (Argumentation System*). *An* argumentation system *is a pair* $(\mathbf{A}, \triangleright)$*, where* $\mathbf{A}$ *is a set of arguments, and* $\triangleright \subseteq (\mathcal{P}(\mathbf{A}) \setminus \{\emptyset\}) \times \mathbf{A}$ *is an* attack relation.

- The basic semantics notions of Dung's framework are extended to the case of attacking sets in a nice and rather direct way
- Further developments in this direction are being investigated in the literature to further explore the potential of Dung's framework

Attacks to attacks

- May attack be attacked in turn?
- Basic idea: if there is a reason for an argument being attacked by another one, this reason may be defeasible
- Attacks as part of the reasoning process rather than a sort of syntactic-automatic-non revisable notion
- Meta-argumentation: reasoning about reasoning, attacks at a lower level are “special arguments” at a meta-level

Modgil's EAF

- Focused on reasoning about preferences
- Given two conflicting arguments A and B let say you have a reason to prefer A wrt B. This can be expressed as an argument C attacking (suppressing) an attack from B to A.
- The reason to prefer A can be defeasible and so on …

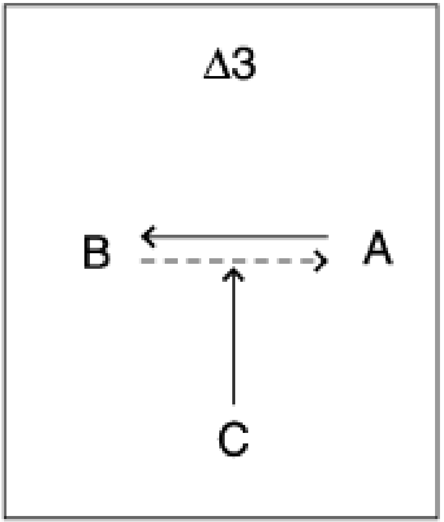

Modgil's EAF

- Arguments can attack attacks between arguments

Definition 4. An *Extended Argumentation Framework* (*EAF*) is a tuple $(Args, \mathcal{R}, \mathcal{D})$

- $\mathcal{R} \subseteq Args \times Args$,
- $\mathcal{D} \subseteq Args \times \mathcal{R}$,
- If $(X, (Y, Z)), (X', (Z, Y)) \in \mathcal{D}$ then $(X, X'), (X', X) \in \mathcal{R}$.

- Two attack levels are explicitly distinguished
- Only one level of "attack recursion" is allowed
- The underlying interpretation in terms of preferences motivates the fact that attacks to attacks can not be attacked in turn and the final condition (opposite preferences should attack each other)

Unlimited attack recursion: AFRA

- AFRA: arguments can attack any attack

Definition 3 ($AFRA$)**.** An Argumentation Framework with Recursive Attacks ($AFRA$) is a pair $\langle \mathcal{A}, \mathcal{R} \rangle$ where:

- $\mathcal{A}$ is a set of arguments;
- $\mathcal{R}$ is a set of attacks, namely pairs $(A, \mathscr{X})$ s.t. $A \in \mathcal{A}$ and ($\mathscr{X} \in \mathcal{R}$ or $\mathscr{X} \in \mathcal{A}$).

- The definition explicitly encompasses two sorts of entities (like Dung's one)
- Unlimited attack recursion levels are allowed
- No domain dependencies

Semantics notions with attack to attacks

- Traditional semantics notions have been recasted in frameworks with attacks to attacks
- Basic intuition 1: attacks which are defeated do not count as conflicts anymore
- Basic intuition 2: you can translate (flatten) an extended framework into a traditional one and then do semantics evaluation at the flattened level
- Formalising these intuitions is not immediate: it has been done differently in the context of the cited approaches

Flatten them all …

- The variants presented above provide an explicit representation of some further notions
- This does not mean that they are “more expressive” than the original framework: “flattening” procedures have been investigated to translate extended frameworks back to the original one
- Typically flattening means to introduce additional (meta-)arguments and attacks corresponding to the additional notions of the extended framework
- Correspondences can be drawn between semantics notions at the extended and flattened level

Flatten them all …

- Several perspectives:
  - » interpretation: the flattened representation may provide hints of various kinds on the extended one
  - » reuse: existing results are applicable effortless to the flattened representation
  - » correspondence/verification: a meaningful mapping should exist between semantics notions in the two levels

… with some concern

- A suitable flattening procedure shows that an extended framework is not more expressive than the traditional AF
- One may argue that extended frameworks arise from modeling carelessness (or modeling indolence) while the “right” modeling is in the flattened version
- Synthetic modeling is appropriate for knowledge and reasoning representation: a sort of high-level language to be compiled into AF

Examples of flattening: recursive attacks

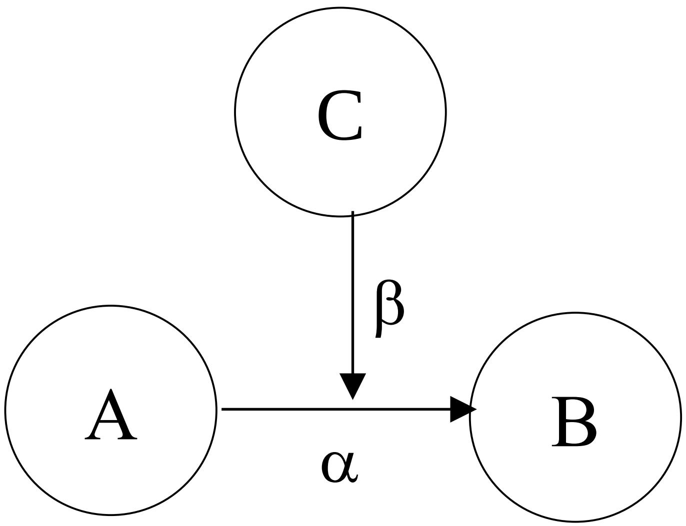

EAF

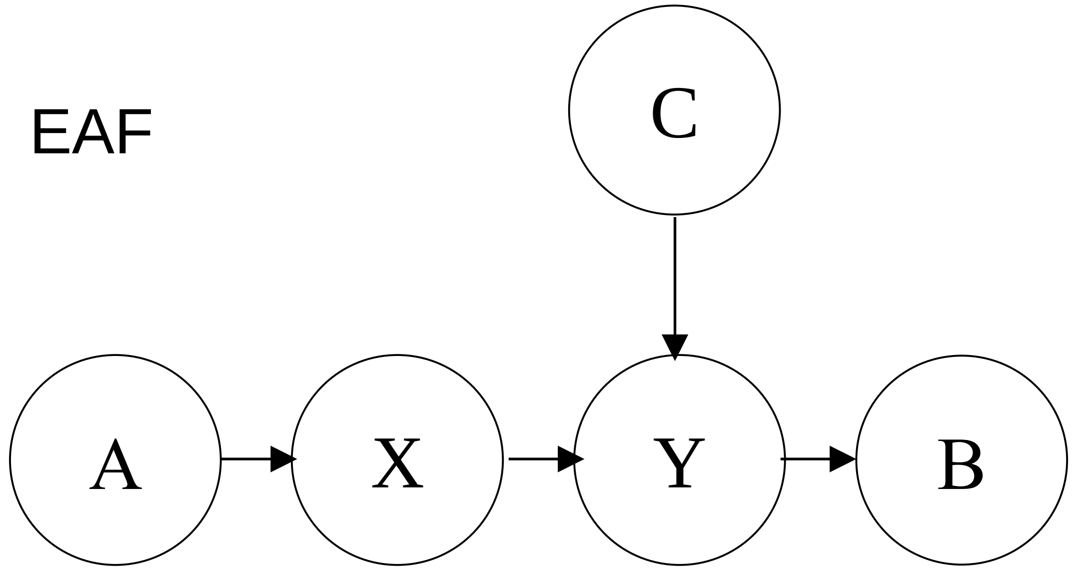

AFRA

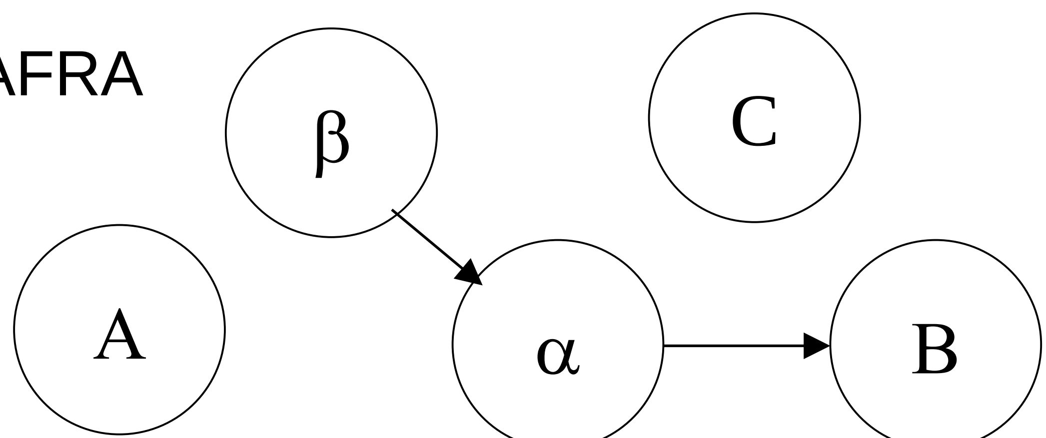

Flexible relations: abstract dialectical frameworks

Definition 5. An *abstract dialectical framework* is a tuple $D = (S, L, C)$ where

- S is a set of statements,
- $L \subseteq S \times S$ is a set of links,
- $C = \{C_s\}_{s \in S}$ is a set of total functions $C_s : 2^{par(s)} \to \{in, out\}$, one for each statement s. C_s is called acceptance condition of s.

- Even the nature of the relation between "arguments" is not specified: links of different nature (attack, support, others? …) all belong to the relation L
- All the meaning is embedded into the acceptance conditions (one for each node: heterogeneous situations may occur)

A non-Dung semantics: “unanimity of attacks”

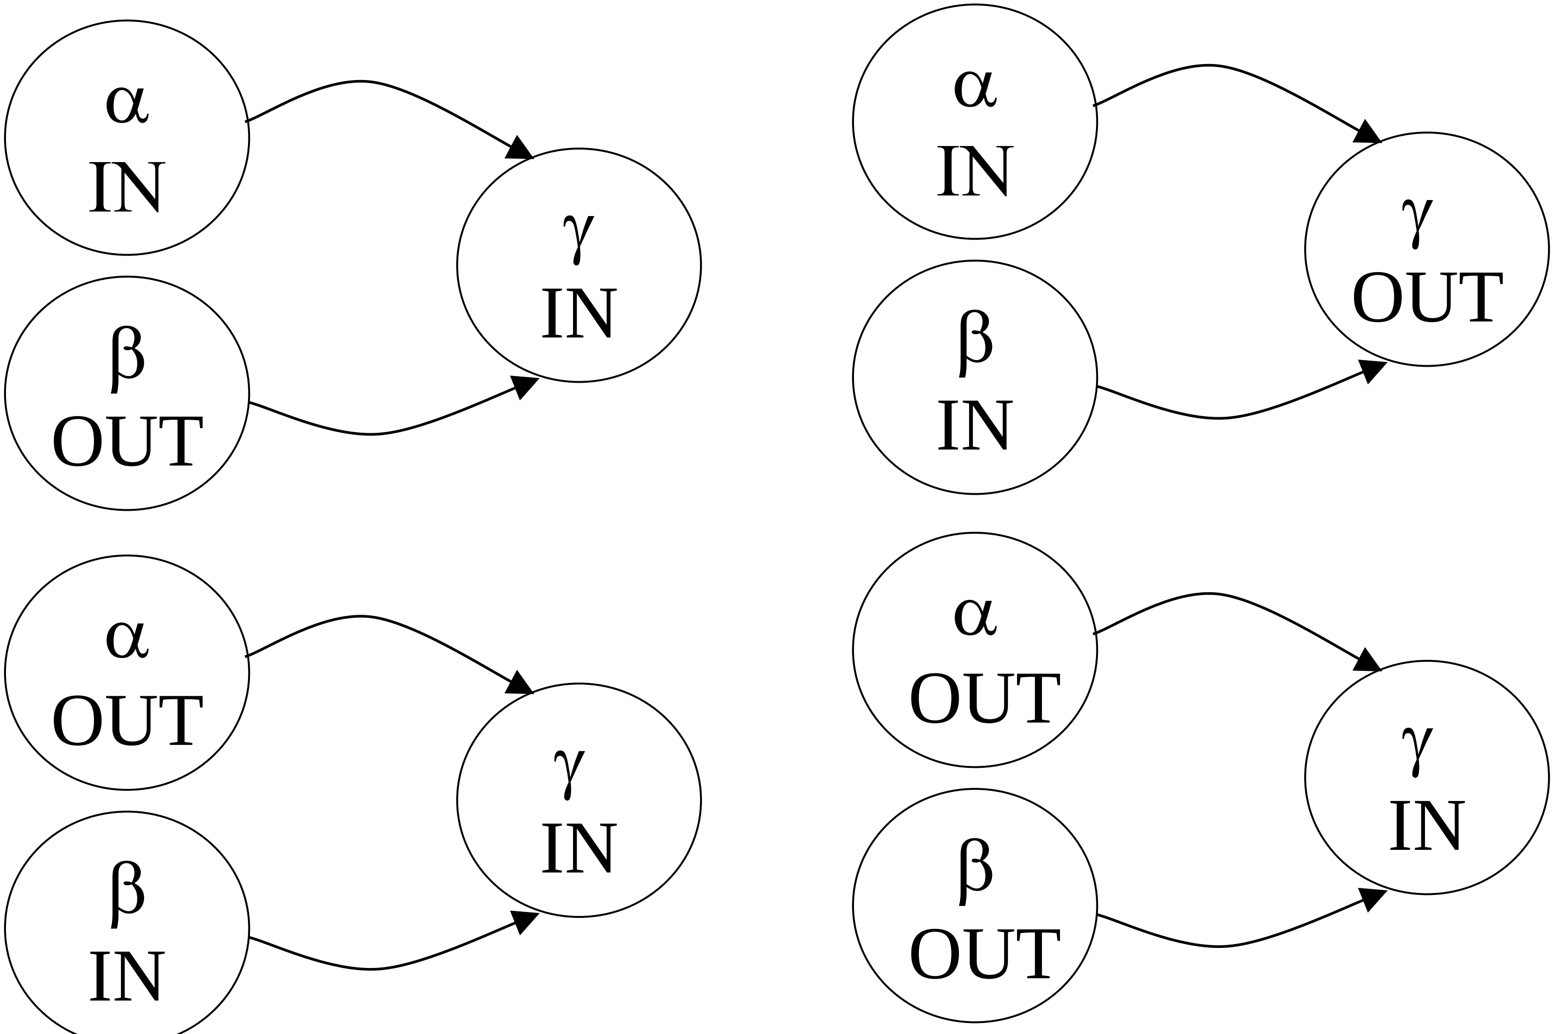

… that can be expresses in Dung's AF through additional arguments

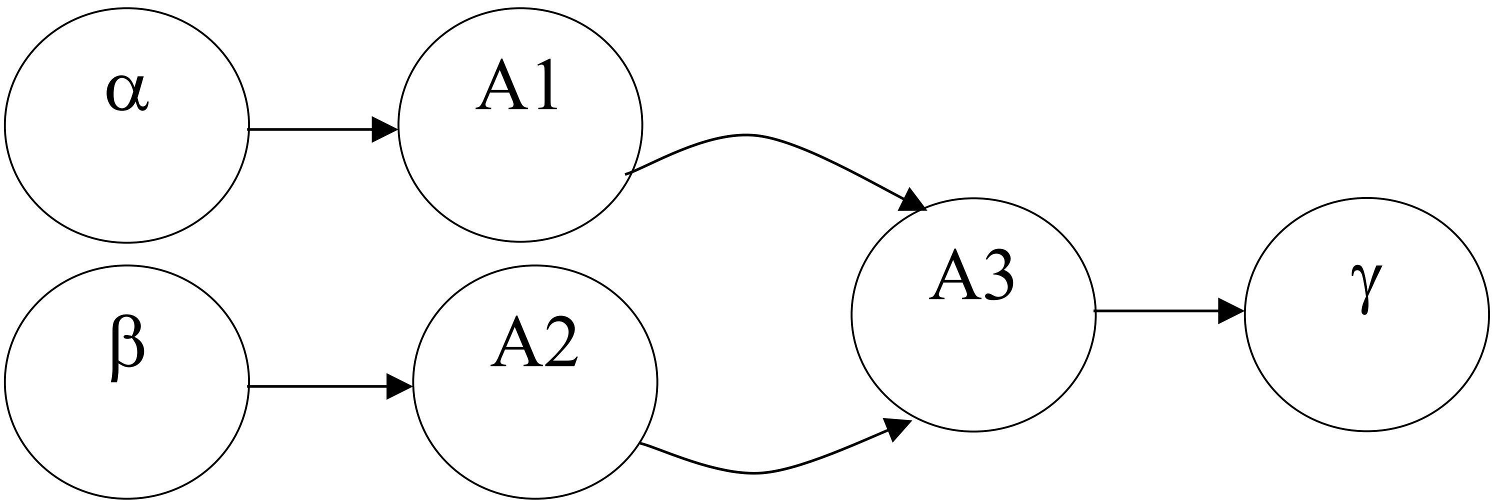

A larger universe?

- ADFs represent an alternative perspective where the only embedded principle seems the one of directionality (rather than conflict-free)
- Large variety of “semantics”, actually of acceptance functions, even inside the same framework
- An alternative universe for lovers of abstract argumentation semantics

Abstract but not too much …

- Dung's abstract argumentation forgets every detail of argument structure and every relation but attack
- In some cases it may be interesting to define an intermediate abstraction level where some argument features (e.g. the conclusion) and/or some further relations (e.g. the subargument relation) are considered
- Here "semi-abstract" arguments are arguments
- Useful to study general properties of structured argumentation in a formalism independent way

Claim-augmented frameworks

Definition 3. A *claim-augmented argumentation framework* (CAF for short) is a triple $(A, R, claim)$ where (A, R) is an AF and $claim\colon A \to \mathcal{C}$ assigns a claim to each argument of A; $\mathcal{C}$ is the set of possible claims.

A CAF $(A, R, claim)$ is called *well-formed* if, for any $a, b \in A$ with $claim(a) = claim(b)$, $\{c \mid (a, c) \in R\} = \{c \mid (b, c) \in R\}$, i.e. arguments with the same claim attack the same arguments.

- A framework where one "remembers" the argument claims (or conclusions)
- The claim is typically the most interesting part of an argument for the final user
- One may be more interested in the evaluation of claims than of arguments
- The complexity of evaluating claims has been investigated at a general level in this context

LAF-ensembles

Definition 7. Given a language $\mathcal{L}$ and a disjoint set *IDS* of *argument identifiers* ($IDS \cap \mathcal{L} = \emptyset$), a $\mathcal{L}$-argument α is a tuple $\langle \mathsf{AI}(\alpha), Conc(\alpha), \mathsf{AE}(\alpha) \rangle$ where $\mathsf{AI}(\alpha) \in IDS$ is the *argument identifier*, $Conc(\alpha) \in \mathcal{L}$ is the *argument conclusion* and $\mathsf{AE}(\alpha) \subseteq \mathcal{L}$ is the set of *attackable elements* of α. For a set of $\mathcal{L}$-arguments X we define, with a little abuse of notation, $Conc(X) \triangleq \{Conc(\alpha) \mid \alpha \in X\}$ and $\mathsf{AE}(X) \triangleq \bigcup_{\alpha \in X} \mathsf{AE}(\alpha)$.

- An argument has a conclusion and some attackable elements
- Several formalism-independent notions bridging the structured level and argument evaluation at the abstract level are defined
- General properties are introduced to investigate the satisfaction of some desiderata of structured formalisms in a formalism-independent way

LAF ensembles

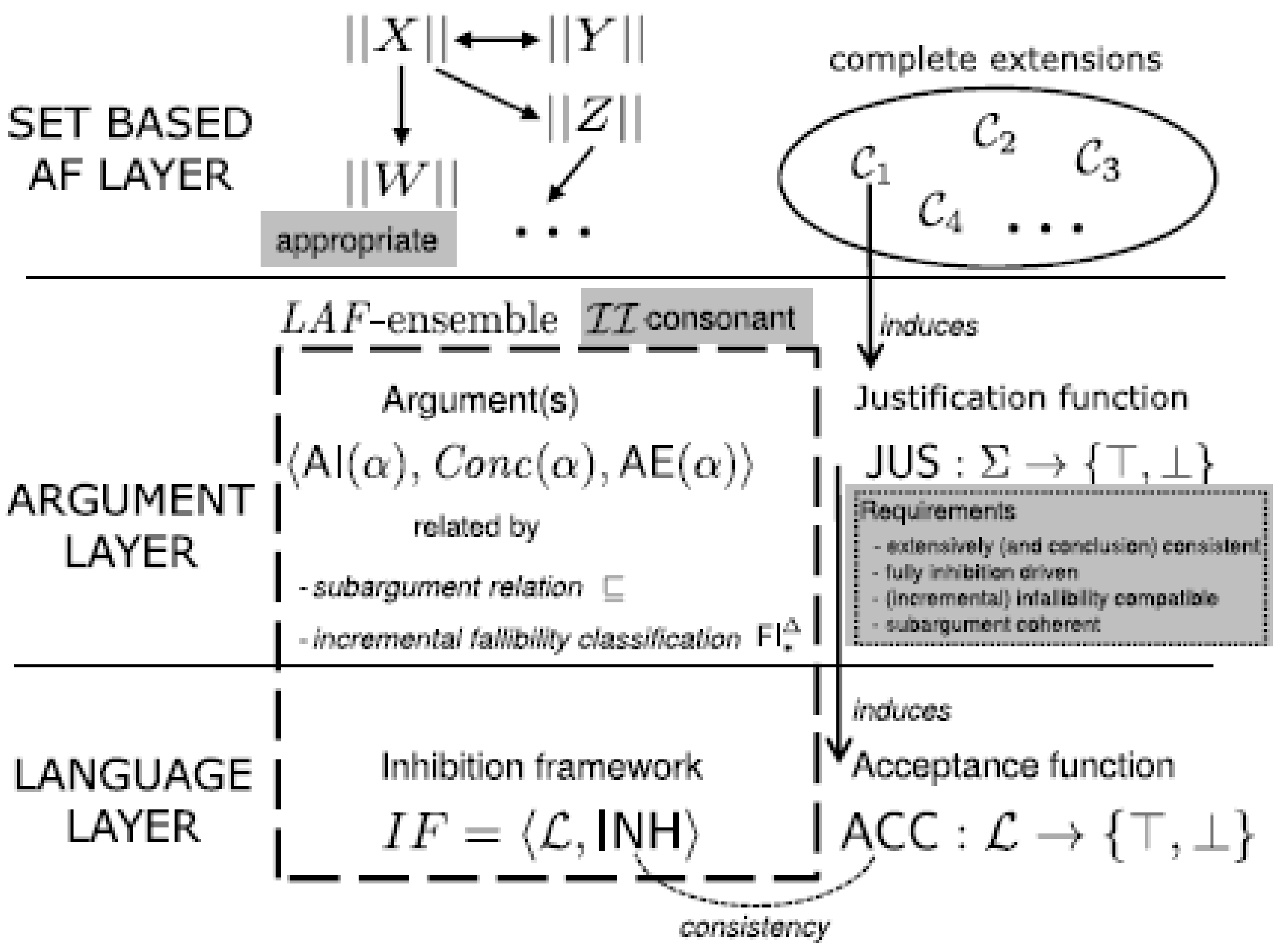

LAF ensembles

- Identification and solution of some problems in ASPIC+
- First connection between Vreeswijk's and Dung's approaches
- Analysis and comparison of further structured formalisms as future work

Conclusions

- A rich and very active research area based on "almost nothing"
- Conflict (rather than argument) is the key notion ensuring wide scope both in theory and in practice
- Large corpus of both basic and specific/advanced results with many "research avenues" and "rethinking opportunities" still open
- Many (at least in principle) reusable and (hopefully) stimulating concepts for other research fields where conflict management plays a key role

Thank you for your patience!

Any argument?

Some references

The big bang

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.

A recent special issue

Special issue: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games: 25 years later. Argument & Computation, 11(1-2), 2020

Argumentation in AI

K. Atkinson, P. Baroni, M. Giacomin, A. Hunter, H. Prakken, C. Reed, G. Simari, M. Thimm, S. Villata, Toward Artificial Argumentation, AI Magazine, 38(3), 2017, 25-36

G. R. Simari, I. Rahwan (Eds.), Argumentation in Artificial Intelligence, Springer Science+Business Media, 2009

T. J. M. Bench-Capon, P. E. Dunne: Argumentation in artificial intelligence. Artificial Intelligence 171(10-15): 619-641 (2007)

A comprehensive handbook

P. Baroni, D. Gabbay, M. Giacomin, L. van der Torre (Eds.), Handbook of Formal Argumentation, College Publications, London, 2018

Some references

Introductory material on abstract argumentation semantics

P. Baroni, M. Caminada, M. Giacomin, An introduction to argumentation semantics, Knowledge Engineering Review, 26(4), 2011, 365-410

P. Baroni, M. Giacomin, Semantics of abstract argument systems, In G. R. Simari, Iyad Rahwan (Eds.), Argumentation in Artificial Intelligence, Springer Science+Business Media, 2009, 25-44

Semantics evaluation criteria

P. Baroni, M. Giacomin, On principle-based evaluation of extension-based argumentation semantics, Artificial Intelligence, 171(10/15), 2007, 675-700

Non-traditional semantics

B. Verheij, Two approaches to dialectical argumentation: admissible sets and argumentation stages. In Proceedings of the 8th Dutch Conference on Artificial Intelligence (NAIC'96), Meyer, J.-J. & van der Gaag, L. (eds). Utrecht University, 357–368. 1996.

M. W. A. Caminada, Semi-stable semantics. In Proceedings of the 1st International Conf. on Computational Models of Argument (COMMA 2006), 121–130.

P. M. Dung, P. Mancarella, F. Toni, Computing ideal sceptical argumentation. Artificial Intelligence 171(10–15), 642–674, 2007.

M. W. A.Caminada, Comparing two unique extension semantics for formal argumentation: ideal and eager. In Proceedings of the 19th Belgian-Dutch Conf. on Artificial Intelligence (BNAIC 2007), 81–87

Some references

Beyond Dung's framework

Bipolar argumentation frameworks

L. Amgoud, C. Cayrol, M.-C. Lagasquie-Schiex, P. Livet: On bipolarity in argumentation frameworks. Int. J. Intell. Syst. 23(10): 1062-1093 (2008)

C. Cayrol, M.-C. Lagasquie-Schiex, Bipolarity in argumentation graphs: Towards a better understanding. Int. J. Approx. Reason. 54(7): 876-899 (2013)

Preference-based argumentation frameworks

L. Amgoud, C. Cayrol: A Reasoning Model Based on the Production of Acceptable Arguments. Ann. Math. Artif. Intell. 34(1-3): 197-215 (2002)

Value-based argumentation frameworks

T. J. M. Bench-Capon: Persuasion in Practical Argument Using Value-based Argumentation Frameworks. J. Log. Comput. 13(3): 429-448 (2003)

Some references

Quantitative evaluations

J. L. Pollock: Defeasible reasoning with variable degrees of justification. Artif. Intell. 133(1-2): 233-282 (2001)

D.M. Gabbay, Equational approach to argumentation networks. Argument and Computation. 3(2–3), 87–142

P. E. Dunne, A. Hunter, P. McBurney, S. Parsons, M. Wooldridge: Weighted argument systems: Basic definitions, algorithms, and complexity results. Artif. Intell. 175(2): 457-486 (2011)

L. Amgoud, J. Ben-Naim, D. Doder, S. Vesic: Ranking Arguments With Compensation-Based Semantics. KR 2016: 12-21

P. Baroni, A. Rago, F. Toni, From fine-grained properties to broad principles for gradual argumentation: A principled spectrum, International Journal of Approximate Reasoning, 105, 252–286 (2019)

Systems with collective attacks

G. Vreeswijk: Abstract Argumentation Systems. Artif. Intell. 90(1-2): 225-279 (1997)

P. Baroni, M. Giacomin, G. Guida: Extending abstract argumentation systems theory. Artif. Intell. 120(2): 251-270 (2000)

S. H. Nielsen, S. Parsons: Computing Preferred Extensions for Argumentation Systems with Sets of Attacking Arguments. Proc. 1st Int. Conf. on Computational Models of Argument (COMMA 2006) 97-108

Some references

Attacks to attacks

S. Modgil, Reasoning about preferences in argumentation frameworks. Artif. Intell. 173(9-10): 901-934 (2009)

P. Baroni, F. Cerutti, M. Giacomin, Giovanni Guida: AFRA: Argumentation framework with recursive attacks. Int. J. Approx. Reasoning 52(1): 19-37 (2011)

H. Barringer, D. M. Gabbay, J. V. Woods, Temporal dynamics of support and attack networks: From argumentation to zoology, in: D. Hutter, W. Stephan (Eds.), Mechanizing Mathematical Reasoning, Springer, 2005, 59–98.

D. Gabbay, Semantics for higher level attacks in extended argumentation frames Part 1: Overview, Studia Logica 93 (2/3) (2009) 357–381.

Abstract dialectical frameworks

G. Brewka, S. Woltran, Abstract dialectical frameworks, Proc. of the 12th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2010), 102-111

G. Brewka, H. Strass, S. Ellmauthaler, J. P. Wallner, S. Woltran, Abstract Dialectical Frameworks Revisited, Proc. of the 23rd Int. Joint Conf. on Artificial Intelligence (IJCAI 2013), 803–809

Semi-abstract (or semi-structured) frameworks

W. Dvorák, S. Woltran, Complexity of abstract argumentation under a claim-centric view. Artif. Intell. 285 (2020)

P. Baroni, M. Giacomin, B. Liao, A general semi-structured formalism for computational argumentation: Definition, properties, and examples of application. Artif. Intell. 257: 158-207 (2018)