

Dilemmas and Paradoxes: Cycles in Argumentation Frameworks

T.J.M. Bench-Capon*

University of Liverpool, Department of Computer Science

October 15, 2013

Abstract

In this paper I offer an interpretation of cycles in Dung-style argumentation frameworks, in which even length cycles are treated as dilemmas and odd length cycles as paradoxes. The different properties of cycles with different parities arising from the use of preferred semantics are argued to be coherent with this interpretation.

Keywords: argumentation, preferred semantics, cycles

1 Introduction

The importance of parity is evident from many fields: the lover pulls petals from a daisy saying “she loves me, she loves me not” and whether the love is requited or not depends on whether the daisy started with an odd or even number of petals. In table tennis a rally with an odd number of shots is won by the server, while one with an even number of shots is won by the receiver. In logic if we x-or a number of propositions together, the result is true if and only if the number of true propositions so connected is odd. Most importantly for our purposes, parity is important in Argumentation Frameworks as introduced by Dung [4]¹.

A very simple structure for an argumentation framework is the *chain*, in which one argument (the *sink*) has one attacker but attacks nothing, a different argument (the *source*) has no attacker, but attacks one argument, and every other argument in the chain attacks, and is attacked by, exactly one other argument. A chain of five arguments is shown in Figure 1. Now, starting at the source, number the arguments. Those with odd numbers will be acceptable (*in* using the labelling of [7]) under grounded, preferred and stable semantics, while the arguments with even numbers will be unacceptable (*out*). If there are no cycles in an argumentation framework, things are very simple: the grounded, preferred and stable extensions coincide: these extensions are unique

*Corresponding author. Email: tbc@csc.liv.ac.uk. Postal address: Department of Computer Science, University of Liverpool, Ashton Building, Ashton Street, Liverpool L69 3BX, United Kingdom.

¹Any reader unfamiliar with the definitions of argumentation frameworks as introduced in [4], should read that paper, or another of the many excellent summaries such as that given in [1].

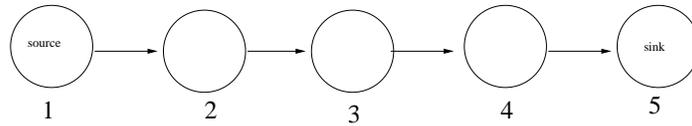


Figure 1: Argument Chain of five arguments. Odd numbered arguments are acceptable, even numbered arguments are not.

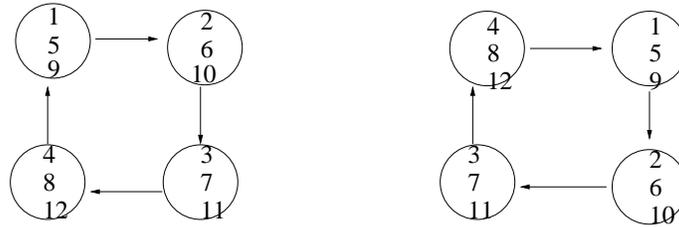


Figure 2: Even cycles: argument numbers have the same parity on every circuit

and (assuming there is at least one argument in the framework), not empty. This extension can be found using the algorithm for determining the grounded extension given in e.g. [3]. Starting with the empty set, this algorithm first includes all unattacked arguments, and then removes these arguments and those arguments they attack from the framework. The process is repeated, including the arguments not attacked in the new framework and removing those attacked by the newly included arguments. This continues until no new arguments are included. Note that each step considers each argument chain in the current framework, and includes the first and excludes the second: odd is good, even is bad; she loves me, she loves me not.

So far so straightforward. But suppose we form a cycle from a chain by adding an attack from the sink to the source. Now if we select an argument and number arguments from that argument we can continue indefinitely. Now suppose we consider an argument as a sink by initially ignoring the fact that it attacks another argument. That argument will then be the sink of an infinite set of chains, one for each time we go round the cycle, each of which has the argument it originally attacked as source. If there is an even number of arguments, we don't have any problems: the "sink" argument will have an even number in every chain and so will be *out* (see Figure 2). In an even chain the arguments partition neatly: suppose we colour alternating arguments red and blue: we can choose to accept the blue arguments or the red arguments. Wherever we start our numbering we will accept all and only the red arguments or all and only the blue arguments. Thus the grounded extension is empty (there is no source), but we have exactly two preferred extensions, one containing the red arguments and the other the blue. In other words we have a *dilemma*: we must choose which of these two subsets of the arguments in the framework to accept.

The situation with an odd length cycle is not so tidy. Now each time we go round the cycle the parity of the argument will change: in a cycle of length 5, number 1 first time round will be number 6 on the second circuit, number 11 on the third circuit

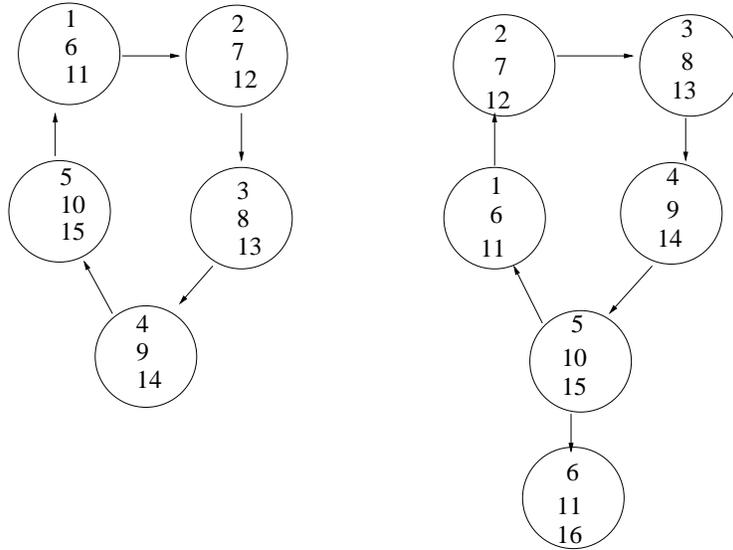


Figure 3: Odd cycles: argument parities alternate on each circuit. This ambiguity applies also to arguments attacked by a member of the cycle.

and so on, alternating between odd and even numbers (see the left hand framework of Figure 3). There is no stability here: if an argument is *in* for one circuit, it is *out* for the next. And whether it is *in* or *out* for an odd number of circuits, depends on the entirely arbitrary choice of starting argument. This means that no argument in an odd cycle can be accepted: trying to make it *in* on the first circuit will ensure that it is *out* on the second, and trying to make it *out* on the first circuit will ensure that it is *in* on the second. This makes it resemble a *paradox*, where assuming the truth of a statement enables its negation to be derived *and* assuming its negation allows its truth to be inferred. Thus the grounded and preferred extensions of an odd length cycle (of any length) are always (only) the empty set. Nothing can be proved.

So far things seem quite in accord with intuition. For some, the problem comes when we consider an argument attacked by an argument in an odd cycle. It might be argued that because the attacker is never *in*, that argument should itself be *in*. But, of course, this argument is in fact not acceptable under grounded or preferred semantics because there is no admissible set including an argument which attacks its attacker. That this is in fact may be exactly what we want can be seen by applying the reasoning we have used thus far: the argument can be seen as the sink of a chain the remainder of which is composed of the arguments in the odd length cycle (see the right hand framework of Figure 3). Suppose there are n arguments in the cycle. Now the argument outside the cycle will be the $n + 1^{th}$ argument in the chain, and so that will have an even number, making it *out*. But, if we went round the cycle twice before leaving the cycle and going to the argument outside, that argument would be numbered $2n + 1$: therefore it would have an odd number, and so should be *in*. That it was *out* on the first circuit and *in* only after two circuits is not pertinent: this depends only on the arbitrary

choice of where to start: there is no reason to consider any particular argument the source, and we can leave the cycle to go to the external argument on any lap.

Now consider an argument attacked by an argument attacked by an even length cycle. Here, once we have chosen the starting argument the parity remains the same no matter how many time we go round the cycle before going to the argument outside the cycle. The additional argument is part of the same dilemma as the arguments in the cycle: we can choose to accept it or reject it. In contrast the argument attacked by an argument in an odd cycle has no stable parity and any attempt to make it *in* will (by an extra lap of the cycle) give a reason for it to be *out*. The odd cycle poisons the whole framework, and the preferred extension will remain empty, however many arguments are added, as long as the cycle remains odd.

It is probably these consequences for arguments which are not themselves part of the cycle that lead people to think there is something unsatisfactory about the way preferred semantics treats cycles: that arguments attacked by arguments in cycles are treated differently for odd and even cycles. This dissatisfaction has motivated the development of semantics which do not have this property [2]². But dilemmas and paradoxes do arise in natural language, and so we should be able to accommodate them in argumentation frameworks, and we should expect some peculiar things to happen if we allow a paradox into our framework³.

It has to be conceded that while the above is in accordance with *my* intuitions, the intuitions of John Pollock and the others that have had different intuitions need to be respected. My intuitions have developed mainly from working with legal argument, which is a very binary domain: legal questions have to be resolved for one party or the other, and legal procedures have evolved to ensure that such decisions are possible. Where there are arguments for both sides, some person or persons are empowered to choose what is, and will be, the case. Those who have worked on argumentation in different domains may have formed different intuitions, more appropriate to those domains. Moreover, what I have presented in this section is an argument and not a proof, and so is intended to be *persuasive* rather than *coercive*. In particular it depends on being able to evaluate the arguments in cycles in a way similar to that used for the arguments in a chain. This in turn requires some argument in the cycle to be *chosen* as the source of the “chain”. Although the argument I have put forward is unaffected by which argument in the chain happens to be chosen, it can still be argued that such a choice is illegitimate, and that the arguments of a cycle should not be evaluated in this way, and that my analogy with chains should be rejected. This is all true: my concern, however, is only to defend traditional preferred semantics, and to argue that it can be useful to distinguish paradoxes and dilemmas, not to deny that alternative interpretations are possible, or that, for certain purposes, alternatives may be desirable.

The next section will offer some examples of paradoxes and dilemmas to illustrate

²John Pollock felt a similar dissatisfaction in treating arguments attacked by cycles of different parity differently [5]. He attempted to solve the problem with his “critical link” semantics.

³Because a paradox cannot be included in a description of a state of affairs which actually obtains, the situation is similar to a logical theory which includes a contradiction. It may be helpful to see the empty preferred extension of an argumentation framework with an odd cycle as akin to the so-called *Paradoxes of Material Implication*, which some also find counter-intuitive, but which are an ineluctable consequence of standard logic.

how things should go when we have a paradox as part of our framework. Again it should be borne in mind that the example it intended to support the *argument* that distinguishing paradoxes and dilemmas are useful, and that traditional semantics are adequate for this task. It is not intended to show that all other interpretations are wrong, and that there is no role at all for non-traditional semantics.

2 Liars

Perhaps the simplest paradox is the classic *Liar Paradox*. Essentially the paradox is the statement “this sentence is false”. This is clearly a paradox: the statement is true if and only if it false. The paradox has a long history: the variant known as the *Epimenides paradox* goes back to around 600 BC. Epimenides, himself a Cretan, reportedly stated that “The Cretans are always liars.” However, Epimenides’ statement that all Cretans are liars can be resolved as false, given that he knows of at least one other Cretan who does not lie, at least on some occasions. The characteristic of a paradox is that if we assume it to be true, then we can show that it is false *and* if we assume it to be false we can show that it is true. Perhaps the earliest version of the liar paradox as we now know it is attributed to Eubulides of Miletus, a Greek philosopher of the 4th century BC. Eubulides reportedly asked, "A man says that he is lying. Is what he says true or false?".

Here we will use a version that looks more like an argument⁴. Suppose that John says “Everything I, John, say is false”. Now let P be the statement that what John said was true; and $JL1$ the argument that P is true, because John asserted it (a sort of argument from authority). $JL1$ is attacked by itself, since if the claim of $JL1$ is true, the claim of $JL1$ is false. We have a simple 1-cycle, and should believe nothing. Moreover, suppose we have another argument, $JL2$ attacked by $JL1$ such as “John was born in Colchester because John says so”. Do we want to accept $JL2$ because we can’t accept its attacker? My intuition says not: it would seem odd (to me at least) to explain that we accept something on the basis that John said it simply because he himself claimed that everything he says is false, and we cannot accept that. Rather, John’s utterance should cast doubt on anything he says - $JL1$ undercuts any argument which has *because John said it* as its premise. While in general, we are inclined to accept what people say, and so this argumentation scheme can be useful, in the case of John we have a good reason to doubt what he says.

Now suppose we extend the situation: John says “what Mary is saying is false (Q)” and Mary says “what John is saying is false (P)”. Now our arguments are,

- $JM1$: Q because John says so; and
- $JM2$: P because Mary says so.

In this case $JM1$ and $JM2$ undercut one another, so that we have a two cycle. Here we can accept either argument (using preferred semantics) but not both. We have a classic dilemma: we can believe either John or Mary, but not both. Of course, we

⁴The examples in this section derive from a personal communication of Henry Prakken, reporting a personal communication to him from John Horty. I am grateful to both.

are not compelled to believe either, adopting grounded semantics. Note that this is also appropriate if we consider arguments attacked by *JM1* and *JM2*. Suppose we have

- *JM3*: John was born in Colchester because John says so; and
- *JM4*: John was born in Chelmsford because Mary says so.

JM3 is attacked by *JM2* and *JM4* is attacked by *JM1*. If we accept *JM1*, we will accept *JM3* (that is we accept the John statements), while if we choose *JM2* we have to accept *JM4* also (that is, we believe Mary rather than John). That is, we make a choice with respect to the dilemma represented by the two cycle, and this determines which partition we will accept.

Let us continue the process, as a third person arrives. Now John says “what Robin is saying is false (*R*)” and Mary says “what John is saying is false (*P*)” and Robin says “what Mary is saying is false (*Q*)”. The arguments are now:

- *JMR1*: *R* because John says so; and
- *JMR2*: *P* because Mary says so; and
- *JMR3*: *Q* because Robin says so.

and each of these arguments will undercut one of the others to form a three cycle.

Now, as with the original Liar Paradox, we have an odd cycle and so again we cannot accept anything: if we choose to try to accept what Mary says, we must reject what John says, which means we have no reason to reject Robin, and so we should reject Mary. Now of course, John is reinstated and so Robin has to go, so that Mary returns and so on *ad tedium*. We are back with a paradox: nothing can be shown. Again consider an argument attacked by one of these arguments: *JMR2* attacks *JM3*. Do we want to believe John because we cannot believe Mary? One would think not: if we believe John, we find ourselves with no reason to disbelieve Mary, and so no reason not to reject *JM3*.⁵

Extending the group again will give us four people: John says that what Ann says is false, (*S*), Mary again says that what John says is false (*P*), George says that what Mary says is false (*Q*) and Ann says that what George says is false (*U*). The arguments are now

- *JMGA1*: *S* because John says so; and
- *JMGA2*: *P* because Mary says so; and
- *JMGA3*: *Q* because George says so; and

⁵Suppose all three people assert the same proposition, say that Mary was born in Cheltenham (*M1*). Now we have arguments *JMRA* – *JMR6* each of which attacks, and is attacked by, any argument concluding that Mary was not born in Cheltenham (*NMC*). Now *NMC* will be credulously acceptable because it can defend itself against the attacks *JMRA* – *JMR6*, giving two preferred extensions, based on the even length cycles containing *NMC* and *JMRA* etc. One will contain *NMC*, and the other will be the empty set. Of course, given an even number of people *NMC* would be indefensible since it would be defeated by whichever group we chose to believe, and so appear in neither of the preferred extensions of such a framework.

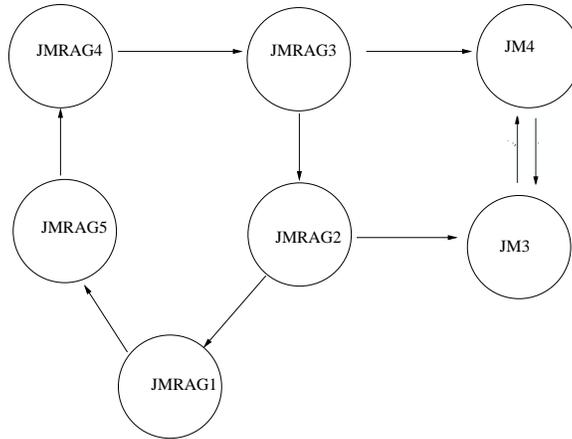


Figure 4: Cycle with five people

- *JMGAR4*: *U* because Ann says so.

Once more each of these arguments will undercut one of the others to form a four cycle. Now again we have a dilemma: We can choose to accept all the arguments of the men, or all the arguments of the women. As with the two cycle, and all cycles of even length, there are two and only two preferred extensions, and all the arguments belong to one and only of these extensions, so that a single choice suffices to determine the status of every argument. Again under grounded semantics the extension is empty, reflecting the fact there is no basis or reason for choosing one set of arguments over the other.

Finally is we reintroduce Robin, who will say that what Ann says is false, while John says that what Robin says is false, so that the arguments are:

- *JMGAR1*: *R* because John says so; and
- *JMGAR2*: *P* because Mary says so; and
- *JMGAR3*: *Q* because George says so; and
- *JMGAR4*: *U* because Ann says so; and
- *JMGAR5*: *S* because Robin says so.

This produces a five cycle (shown in Figure 4), and like all cycles of odd length, the preferred extension is empty, and we have a paradox. If we try to choose to accept *JMGAR1*, we are obliged to accept its attacker. Similarly if we try to choose its attacker, *JMGAR2*, we are obliged to accept the attacker of that argument, so that nothing can be accepted. And once more if we consider arguments attacked by arguments caught in this cycle, they turn out to be no more worthy of acceptance than the argument within the cycle that attacks the argument: We have no reason to accept what John says regarding his place of birth simply because we cannot accept Mary's claim that John is stating falsehoods, because accepting what John says on this basis

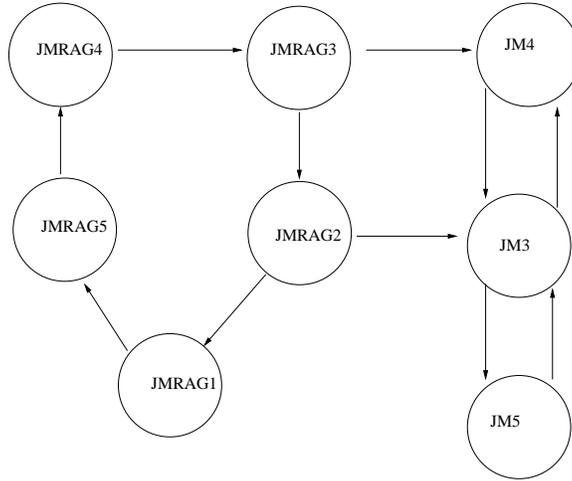


Figure 5: Attacking cycles with six people

will also oblige us to accept *JMGAR2*, and so believe Mary rather than John. Consideration of the content leads to the realisation that *JM3* and *JMGAR1* must stand and fall together because they have the same reason. But *JM3* and *JMGAR1* cannot be accepted together, because accepting *JMGAR1* means that we cannot defeat the attacker of *JM3*. In this way we can see how the paradox spreads beyond the propositions which compose it, and blights all the arguments with which it comes into contact. And giving due consideration to the content is important: the relevant intuitions concern the concrete arguments we represent in the framework, not the abstract structure.

Finally consider the situation shown in Figure 5. Here we have the five cycle of Figure 4, attacking *JM3* and *JM4*, which themselves form a two cycle since they rebut each other. We now add a new argument based on the statement of a new person, who is not involved with people in the cycle:

JM5 John was not born in Colchester, because Lesley says so

JM5 rebuts *JM3* but is unrelated to any other argument.

Now, since we have an odd cycle, every argument except *JM5* cannot be accepted because it is attacked by an argument which has no admissible defeater. *JM5* is, however, able to defend itself. It forms a two cycle - a dilemma - with *JM3*, and we can (under preferred semantics) choose it rather than *JM3*. Thus when it turns out the *JM3* cannot be chosen because the cycle has cast doubt on arguments based on the assertions of John, we simply accept *JM5* because we have no alternative, unless we want to use grounded semantics instead, in which case we can reject *JM5* also. Note that if the cycle had been an even length cycle, such as the four person cycle *JMGA1 – 4*, the situation would have been different. With an even cycle it is possible to choose to accept the utterances of the men or the women. And here, if we choose the men, we reinstate the horn of the dilemma which includes *JM3*, and so *JM5* drops

from being sceptically acceptable to being only credulously acceptable. Again some have found this counter intuitive: why should the status of an argument depend on the parity of a cycle to which it is not related? But in the concrete example, we again see that this is as it should be: the argument scheme on which all our arguments is based is undercut for people caught up in the cycle of mutual denunciation: everyone in the case of the odd cycle; the men or the women for the even cycle. Lesley, however, remains untainted, for the odd cycle and for both horns of the dilemma, and so we can still always use the scheme to support acceptance of what she asserts. Whether *JM5* is challenged depends on whether John can be seen as credible or not, and this *does* depend on the parity of the cycle.

3 Discussion

I have argued that even cycles can be seen as representing dilemmas and odd cycles can be seen as representing paradoxes, and that in many contexts we want to be able to represent and distinguish paradoxes and dilemmas. Traditional preferred semantics enable this: in the case of paradoxes we cannot believe anything, because if we try to assume anything to be true, we can, on the basis of that assumption, show it to be false and *vice versa*. Therefore the grounded and preferred extensions should be the empty set, which is the case for Dung-style argumentation frameworks. For dilemmas we can choose to accept either horn of the dilemma, or refuse to commit ourselves. Adopting preferred semantics gives us two preferred extensions to choose between, and refusal to commit is equivalent to adopting grounded semantics gives us the empty set. So all is as it should be.

For those who see a problem in the way arguments attacked by an argument in a cycle are treated differently for cycles of different parities, I have offered a structural argument in section 1, and argued on the basis of the content of the arguments in section 2. I would further contend that many examples using real arguments will also be seen to be correctly treated when the content is taken into account. Constructing cycles with plausible arguments is not a simple task: much care must be taken if symmetric attacks resulting from conflict based on rebuttal or undermining are to be avoided. It is a pitfall of logic that once we have abstracted away the content we begin to think only in terms of the abstract structure. But sometimes what seems intuitive with respect to the abstraction, no longer appears intuitive when we attempt to instantiate the structure with plausible content. The paradoxes of material implication can appear problematic until we propose examples: on the one hand the idea that if a proposition p is false, then $P \rightarrow Q$ is true for any Q does seem odd: but since P needs to be false to make the implication true, and then true to draw the conclusion that Q , making use of it in practice involves believing a contradiction, drastically limiting its application in practice. Classical logic was designed to reason with theories free from contradictions, and in such cases the behaviour of material implication, including the paradoxes, is exactly what we need. If we do wish to reason with contradictions we need special machinery, such as paraconsistent logics, in which these problems do not arise. Recognising, however, that paraconsistent (and relevance) logics may be required in some circumstances should not lead us to abandon classical logic for those cases where it is appropriate.

Similarly if we are confronted with a plausible concrete example that both represents an argument (or set of related arguments) attacked by an odd cycle, and which should intuitively, considering its content, behave differently from how it behaves under the traditional semantics then we will need an alternative to preferred semantics for that particular case. But even if such a candidate is found, serious consideration should be given to how the arguments have been represented as it may be that representing the arguments in a different way may mean that a new semantics is not essential to handle the problem. The alternative semantics are often mathematically interesting, but are not always necessary for representing naturally occurring arguments, even if there are some special cases for which they are required.

An interesting general discussion of the role that should be played by intuition in semantics of argumentation can be found in [6]. There Prakken points out that the intuitions “of logicians are hopelessly corrupted by overexposure to formalism”, while those of ordinary language speakers often “reveal a lack of understanding of the reasoning patterns” involved. Even if we choose the appropriate group, we still find that intuitions vary. An additional difficulty is that the way the problem is stated can steer people towards one conclusion rather than another ⁶. If we place reliance on intuitions we will find ourselves developing a series of patches to meet (our) intuitions. And when we “solve” one problem, almost certainly it will be possible to construct a situation for which our solution is inappropriate and a further patch must be produced. This was very much the story of calculi for non-monotonic logics in the 1990s. One result is that our semantics may become very much more complicated as we attempt to accommodate examples which rarely (if ever) occur in practice. Dung’s preferred semantics offer a simple and elegant theory, which is adequate for most purposes, and which, as I have argued in this paper, can be given a useful interpretation in some of the apparently difficult cases.

I believe that providing a new semantics should be a last resort not a first resort. Often we can use the insights provided by the defence of the traditional semantics to explain in which cases new semantics are really required and when intuitions may be misleading. I believe that in this paper I have been able to use preferred semantics to explain why it may be both sensible and useful to treat cycles of different parity differently in Dung-style argumentation frameworks, even when consideration of the formalism may initially suggest that this is counter intuitive.

Acknowledgements

Thanks are due to Henry Prakken, both for several conversations on this topic over the years, and for specific suggestions which have improved this paper. Thanks also to Michael Bench-Capon for his comments on a draft of this paper. Finally thanks also to the anonymous reviewers of the original draft of this paper, who made a number of valuable suggestions which have helped improve, clarify and impart more balance to this version.

⁶For example compare “Is it right that arguments are treated differently according to the parities of cycles of arguments to which they are unrelated?” with “Should we expect the parity of a cycle to be irrelevant to the way it behaves in an argumentation framework?”.

References

- [1] P. Baroni and M. Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artif. Intell.*, 171(10-15):675–700, 2007.
- [2] P. Baroni, M. Giacomin, and G. Guida. Scc-recursiveness: a general schema for argumentation semantics. *Artif. Intell.*, 168(1-2):162–210, 2005.
- [3] T. J. M. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *J. Log. Comput.*, 13(3):429–448, 2003.
- [4] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, 77(2):321–358, 1995.
- [5] J. L. Pollock. Defeasible reasoning with variable degrees of justification. *Artif. Intell.*, 133(1-2):233–282, 2001.
- [6] H. Prakken. Intuitions and the modelling of defeasible reasoning: some case studies. In *9th International Workshop on Non-Monotonic Reasoning*, pages 91–102, 2002.
- [7] B. Verheij. A labeling approach to the computation of credulous acceptance in argumentation. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 623–628, 2007.