Jun 1st, 9:00 AM - 5:00 PM

Arguments as Belief Structures: Towards a Toulmin Layout of Doxastic Dynamics?

Fabio Paglieri
University of Siena

Cristiano Castelfranchi
National Research Council (CNR)

Arguments as Belief Structures: Towards a Toulmin Layout of Doxastic Dynamics?

FABIO PAGLIERI[1] AND CRISTIANO CASTELFRANCHI[2]

[1]*Department of Communication Sciences*
University of Siena
Piazza S. Francesco 8
00137 Roma, Italy
paglieri@media.unisi.it

[2]*Institute of Cognitive Sciences and Technologies (ISTC)*
National Research Council (CNR)
Viale Marx 15
53100 Siena, Italy
cristiano.castelfranchi@istc.cnr.it

ABSTRACT: Argumentation is a dialogical attempt to bring about a desired change in the beliefs of another agent – that is, to trigger a specific belief revision process in the mind of such agent. However, so far formal models of belief revision widely neglected any systematic comparison with argumentation theories, to the point that even the simplest argumentation structures cannot be captured within such models. In this essay, we endeavour to bring together argumentation and belief revision in the same formal framework, and to highlight the important role played by Toulmin's layout of argument in fostering such integration.

KEY WORDS: argumentation, belief revision, Toulmin's layout of argument, data structures

1. INTRODUCTION

This research aims to bring together two separate threads in the formal study of doxastic change: *belief revision* and *argumentation theories*. Belief revision describes the way in which an agent is supposed *to change his own mind* in the face of new information, while argumentation deals with persuasive strategies employed *to change the mind of other agents*, by providing reasons for such change: hence belief revision and argumentation are two sides (*cognitive* and *social*) of the same epistemic coin. Argumentation theories remain incomplete, if they cannot be grounded in belief revision models: they describe dialogical patterns and their effects, but cannot explain *how such effects are produced*. Without an underlying model of belief dynamics, argumentation theories are forced to remain 'out of the black box.'

In spite of the vast literature on belief change in epistemology (Rescher, 1976; Harman, 1986; Levi, 1991; Pollock and Gillies, 2000), logic (Alchurrón et al., 1985; Gärdenfors, 1988; Rott, 2001) and Artificial Intelligence (Doyle, 1992; Boutilier, 1998; Dragoni and Giorgini, 2003), so far the formal treatment of belief revision avoided systematic comparison with argument structures. Conversely, current argumentation theories (Toulmin, 1958[2003]; Perelman and Olbrechts-Tyteca, 1969; van Eemeren and Grootendorst, 1992; Walton, 1996) have not yet been integrated with belief dynamics of the arguing agents.

Our basic claim is that such *lack of integration* poses severe limitations to our understanding of both doxastic change and argumentation, so that *more comprehensive models should instead be devised* (Pollock and Gillies, 2000; Falappa et al., 2002; Paglieri and Castelfranchi, 2004). As preliminary step in this direction, here current belief revision formalisms are tested against Toulmin's layout of argument (Toulmin, 1958[2003]), to check whether or not they are able to express the basic features of argumentation. We submit two different belief revision models to the 'Toulmin's test': the classical set-theoretical approach of the *AGM theory* (Alchurrón et al., 1985; Gärdenfors, 1988), and *Data-oriented Belief Revision* (DBR from now on; Castelfranchi, 1996; Paglieri, 2004), a cognitive theory of doxastic processing based on critical assessment of available information (cf. 3). While the AGM approach fails to capture even the most basic features of argumentation (cf. 2), the DBR framework provides an effective mapping for several facets of Toulmin's model (cf. 4). So Toulmin's layout is first showed to act as a *litmus test* for assessing expressive power of belief revision formalisms: failure in integrating such layout in the formal description of belief revision reveals lack of structural analysis, which in turn prevents more comprehensive accounts of doxastic processing. Finally, Toulmin's model itself is argued to improve through comparison with the DBR framework (cf. 5).

2. THE AGM MODEL: BELIEF REVISION WITHOUT ARGUMENTATION?

The AGM paradigm (Alchourrón et al., 1985; Gärdenfors, 1988) has been the most influential model of belief revision so far, serving as a frame of reference for both refinements and criticisms of the original proposal. Roughly summarizing (for details, see Gärdenfors, 1988; Rott, 2001), this model is conceived as an idealistic theory of rational belief change: belief states are characterized as *sets of propositions* (infinite and deductively closed), three basic types of change are described (*expansion*, *contraction*, *revision*), and rationality is expressed by a set of postulates binding these operators. To decide between different outcomes of the revision process (i.e. different sets of propositions consistent with the rationality postulates), an ordering criterion is introduced in the original belief state, ranking propositions according to their relative importance (*epistemic entrenchment*): hence, different epistemic entrenchments would result in different outcomes of the same revision process. Basic aim of this framework is to explain how an agent might change her belief set, while (1) effectively avoiding contradictions and (2) minimizing the loss of information (Harman, 1986).

However, what remains out of the picture in the AGM approach is an account of the *reasons behind epistemic entrenchments*: why and how should a given belief become more important (i.e. entrenched) than another one – besides, why and how should a proposition come to be believed at all? The AGM model, and in general all *coherence theories* of belief revision, as opposed to *foundations theories* (Doyle, 1992), do not take in account 'reasons to believe' – and this is why they cannot capture argumentation structures. In fact, AGM belief revision fails to integrate with argumentation theories because (1) it does not make any predictions or assumptions about how and why some propositions come to be believed, rather than others, and how and why some of them happen to be held more dearly than others by the agent; (2) there is a deliberate lack of structural properties in the AGM characterization of doxastic states. Argumentation theories capture the ways in which a desired change in the audience beliefs is brought about by the arguer: therefore, without an explicit theory of *the reasons to believe something*, the main point of argumentation is lost. Moreover, structural properties of doxastic states play a crucial role in accounting for argumentative strategies (cf. 4 and 5). AGM-style belief revision simply lacks the necessary internal structure to describe such strategies, since it relies on an over-simplified conceptual model of belief dynamics.

If we now turn to consider Toulmin's layout of argument, it is evident there is no way of capturing such structure within the AGM framework. The functional relations between different parts of an argument (its 'physiology', according to Toulmin's metaphor; 1985, 87) cannot be expressed by an ordering criterion such as epistemic entrenchment, nor by simple inferential rules – in fact, it was mainly to contrast the logical-mathematical analysis of arguments in terms of syllogisms that Toulmin conceived his alternative proposal (1985, pp. 88-89, pp. 100-105). Hence, an in-depth analysis of supporting and conflicting relations among different information is required, to cope with the Toulmin's layout (cf. 4).

3. DATA-ORIENTED BELIEF REVISION (DBR)

The following sections outline very shortly a model of belief dynamics alternative to the AGM approach: *Data-oriented Belief Revision (DBR)*. This introduction to DBR is meant to provide only the basics needed to frame argumentation structures within the model: a more comprehensive overview of DBR can be found in Castelfranchi (1996) and Paglieri (2004).

3.1. Data and beliefs: properties and interactions

Two basic informational categories, *data* and *beliefs* (Rescher, 1976; Castelfranchi, 1996; Paglieri, 2004), are put forward in this model, to account for the distinction between pieces of information that are simply *gathered and stored* by the agent (data), and pieces of information that the agent considers (possibly up to a certain degree) *truthful representations of states of the world* (beliefs) – hence taking them as a reliable bases for action, decision, and specific reasoning tasks, e.g. prediction and explanation. Clearly, the latter are a subset of the former: the agent might well be aware of a datum that he does not believe (i.e. he does not consider reliable enough); on the other hand, the agent should not be forced to forget (i.e. to lose as a datum) a piece of information which he temporarily rejects as a belief (Castelfranchi, 1996).

The distinction between data and beliefs yields a number of consequences for the formal study of doxastic dynamics: to start with, it leads to conceive *belief change as a two-step process*. Whenever a new piece of evidence is acquired through perception or communication, it affects *directly* the agent's data structure, and only *indirectly* his belief set. In other words, the effects (if any) of the new datum on the agent's beliefs depend (1) on its effects on the other data, and (2) on the process of belief selection applied by the agent over such data (cf. 3.2). More generally, data and beliefs define two basic cognitive layers of doxastic processing, as summarized in Figure 1.

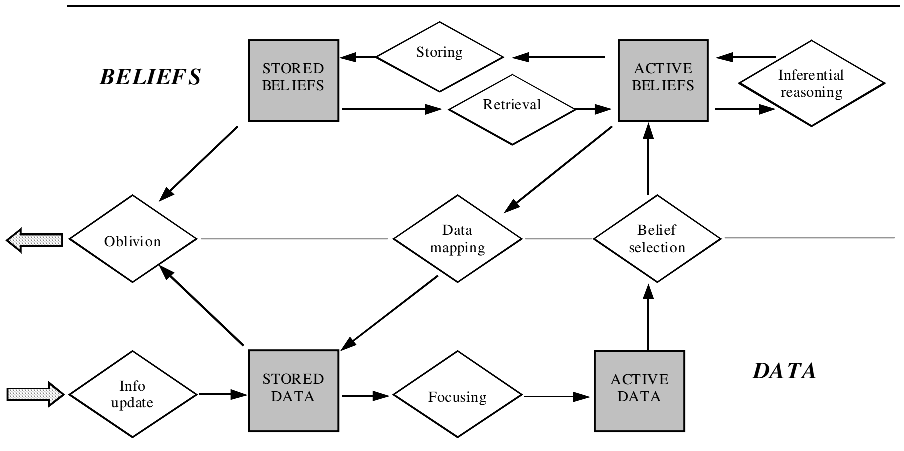

Fig. 1. Doxastic processing in Data-oriented Belief Revision

In this model data are selected (or rejected) as beliefs on the basis of their properties, i.e. the *cognitive reasons to believe* them (Castelfranchi, 1996; Paglieri, 2004), such as:

I. *Relevance*: a measure of the pragmatic utility of the datum, i.e. the number and values of the (pursued) goals that depends on that datum;
II. *Credibility*: a measure of the number and values of all supporting data, contrasted with all conflicting data, down to external and internal sources;
III. *Importance*: a measure of the epistemic connectivity of the datum, i.e. the number and values of the data that should be revised, if the agent revises that single one;
IV. *Likeability*: a measure of the motivational appeal of the datum, i.e. the number and values of the (pursued) goals that are directly fulfilled by that datum.

The assessment of such properties in DBR is detailed in Paglieri (2004). As for their function in belief dynamics, credibility, importance and likeability determine the outcomes of belief selection (cf. 3.2); relevance instead is crucial in pre-selecting the sub-set of active data (focusing), i.e. what data in the agent's memory are useful/appropriate for the current task, and should therefore be taken in consideration as candidate beliefs. While relevance and likeability depend on comparison between data and goals, credibility and importance *depend on structural relations between data* (Castelfranchi, 1996). DBR data bases are structured domains, best conceived as *networks*: data are represented as *nodes*, interconnected through functional relations (cf. 3.3), i.e. *links* in the network. In contrast, beliefs are characterized by a given *strength*, which reflects their implicit ordering after the selection process (cf. 3.2): hence beliefs are organized in *ordered sets*, rather than networks (Gärdenfors, 1988).

3.2. Belief selection in DBR

Once the informational values of the active data are assessed (Paglieri, 2004), a selection over such data is performed, to determine the subset of reliable information (i.e. beliefs) and their degree of strength. When new relevant information is gathered by the agent, modifying his data network and the subset of active data, the belief selection takes place anew, possibly (but not necessarily) changing the agent's belief set.

This process of belief selection regulates the interaction from data to beliefs, determining (1) what data are to be believed, given the current informational state, and (2) which degree of strength is to be assigned to each of them. The outcome of belief selection is determined by the informational values of the candidate data (credibility, importance, likeability) and by the nature of the agent's selection process.

In DBR the agent's belief selection is represented by a mathematical system, including a *condition C*, a *threshold k*, and a *function F*. Condition C and threshold k together express the minimal informational requirements for a datum to be selected as belief. The function F assigns a value of strength to the accepted beliefs. Both C and F are mathematical functions with credibility and/or importance and/or likeability as their arguments. Given a datum ϕ, c^{ϕ}, i^{ϕ}, l^{ϕ} are, respectively, its credibility, importance, and likeability. Let $\boldsymbol{B}$ represents the set of the agent's beliefs, and $B^s\phi$ represents the belief ϕ with strength s. The general form of the selection process is:

if $C(c^{\phi}, i^{\phi}, l^{\phi}) \leq k$ then $B^s\phi \notin \boldsymbol{B}$
if $C(c^{\phi}, i^{\phi}, l^{\phi}) > k$ then $B^s\phi \in \boldsymbol{B}$ with $s^{\phi} = F(c^{\phi}, i^{\phi}, r^{\phi})$

The setting of C, F and k is an individual parameter, which might vary in different agents. Paglieri (2004) provides examples of individual variation in belief selection, and a specific case on argumentation is discussed in Paglieri and Castelfranchi (2004).

3.3. Information update and data structure

Belief revision is often triggered by *information update* either on a fact or on a source: the agent receives a new piece of information, rearranges his data structure accordingly, and possibly changes his belief set. Information update specifies the way in which new evidences are integrated in the agent's data structure. *Data structures* are conceived as networks of nodes (data), linked together by characteristic relations. For the purposes of the present discussion, we define three different types of data relations: support, contrast, and union.

I. *Support*: ϕ supports ψ ($\phi \Rightarrow \psi$) iff $c^{\psi} \propto c^{\phi}$, the credibility of ψ is directly proportional to the credibility of ϕ.
II. *Contrast*: ϕ contrasts ψ ($\phi \perp \psi$) iff $c^{\psi} \propto 1/c^{\phi}$, the credibility of ψ is conversely proportional to the credibility of ϕ.
III. *Union*: ϕ and ψ are united ($\phi \& \psi$) iff c^{ψ} and c^{ϕ} jointly (not separately) determine the credibility of another datum γ.

As we shall discuss in 4 and 5, data relations and their dynamics are the key to represent argumentation in DBR. Moreover, they are also needed to account for information update, since a new input generates not only a datum concerning its *content*, but also data concerning *source attribution* and *source reliability*, and the *structural relations* among them. More precisely, information update brings together:

I. a datum concerning the content (object datum, *O-datum*);
II. a datum identifying the information source (*S-datum*);
III. a datum concerning the reliability of the source (*R-datum*).

These data are closely related, since the credibility of the new information depends on the jointed credibility of the other two data: i.e. the union of the S-datum and the R-datum supports the O-datum (Fig. 2). Once an agent has been told by x that ϕ holds, his confidence in ϕ will depend on the reliability he assigns to x, provided he is sure enough that the source of ϕ was indeed x. The environmental input is characterized by a content ϕ (e.g. its propositional meaning), a source x (e.g. another agent), and a noise n (affecting both source identification and content understanding).

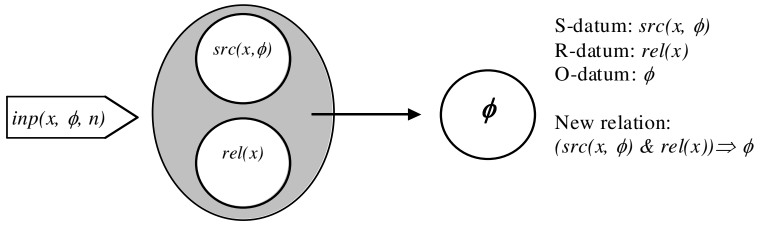

Fig. 2. Information update: integrating new external data

4. TOULMIN'S LAYOUT OF ARGUMENT IN DBR

Toulmin's layout distinguishes and analyzes six features of an argument: data, claim, warrant, backing, qualifier, rebuttal. *Data* are the facts (e.g. 'John loved his wife') which support the arguer's *claim* (e.g. 'John did not murder her'), while the *warrant* ensures the connection between data and claim (e.g. 'people do not murder the ones they love'), on the basis of some *backing* (e.g. statistical evidences); the *qualifier* specifies to what extend the warrant applies (e.g. usually), and the *rebuttal* describes exceptions to the warrant (e.g. 'John aimed to benefit from her insurance'). This schema is liable of immediate implementation in DBR, since it defines a specific data structure (Fig. 3). The *union* of data and warrant *supports* the claim, and the warrant is in turn *supported* by its backing and *contrasted* by the rebuttal, i.e. supports to the rebuttal make the warrant less reliable. The qualifier is represented by the *degree of credibility* assigned to the claim by this data structure.

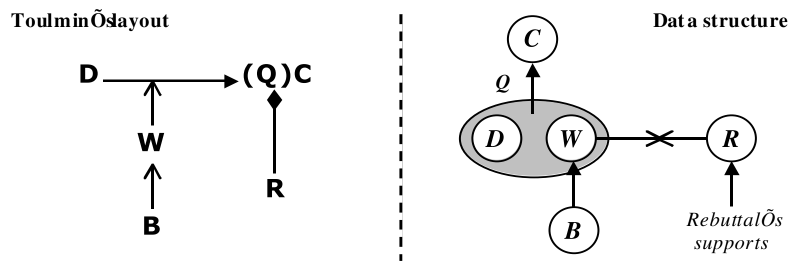

Fig. 3. Toulmin's layout in data structure

4.1. Assessing defeasible reasoning in DBR

Argumentation is often modelled in the framework of *defeasible reasoning* (Pollock and Gillies, 2000; Falappa et al., 2002), distinguishing between two kinds of defeaters (i.e. possible counterarguments against a reason-schema): *rebutting* vs. *undercutting defeaters*. Applying to defeasible reasoning the terminology proposed by Toulmin (1958[2003]), a rebutting defeater is any reason which directly denies the claim of the argument, while an undercutting defeater undermines the validity of the relevant warrant. Actually, with reference to Toulmin's, the expression 'rebutting defeaters' is rather misleading, since «conditions of rebuttal (R) indicat[es] circumstances in which the general authority of the warrant would have to be set aside» (Toulmin, 1958[2003], p. 94). Hence rebuttals are in fact undercutting defeaters. To avoid confusion, we propose the expression *direct defeaters*, instead of rebutting defeaters, as a less ambiguous term for defeaters which directly affect the claim.

In DBR, different defeaters are effectively shown to target different nodes in the data network (Fig. 4): direct defeaters are data which contrast the claim-node (in the previous example, 'John was seen shooting his wife'), while undercutting defeaters are data contrasting the warrant-node (e.g. 'jealousy can make you kill the ones you love most'). Moreover, a third category of defeaters can be expressed: *premise defeaters*, i.e. reasons which contrast the data-node (e.g. 'John did not love his wife'). Undercutting and premise defeaters have similar function but different targets: the former attack the connection between data and claim, while the latter question the statement of fact supporting the conclusion.

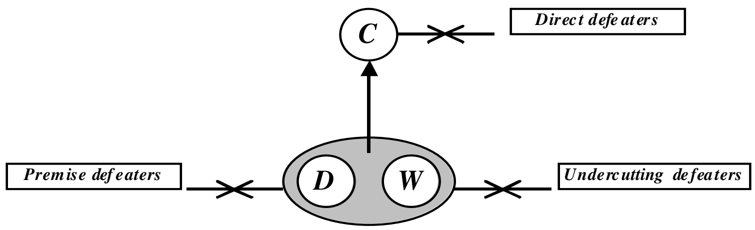

Fig. 4. Defeasible reasoning in data structure

4.2. Recursion in argumentation and belief change

Toulmin was well aware of possible *recursion* in argumentation. He explicitly mentioned this issue with reference to, respectively, data and warrants:

> [I]f the claim is challenged, it is up to us to appeal to these facts [i.e. data], and present them as the foundation upon which our claim is based. Of course we may not get the challenger even to agree about the correctness of these facts, and in that case we have to clear his objection out of the way by a preliminary argument: only when this prior issue or 'lemma', as geometers would call it, has been dealt with, are we in a position to return to the original argument. (1958[2003], p. 90)

> Indeed, if we demanded the credentials of all warrants at sight and never let one pass unchallenged, argument could scarcely begin. Jones puts forward an argument invoking warrant W_1, and Smith challenges that warrant; Jones is obliged, as a lemma, to produce another argument in the hope of establishing the acceptability of the first warrant, but in the course of this lemma employs a second warrant W_2; Smith challenges the credentials of this second warrant in turn; and so the game goes on. (1958[2003], p. 98)

Therefore, Toulmin's layout of argument is not a finite and self-sustaining structure – and the same applies to any other structural account of argumentation. Recursive processes are always possible: all features supporting or challenging a claim can be supported or challenged in turn, hence becoming the (disputed) claim of a specific sub-argument. This holds not only for data and warrants, but also for backings, rebuttals, and even qualifiers. All these recursive processes are liable of representation in DBR.

Consider again the example introduced before, and imagine I am determined to prove John guilty of the murder of his wife, defeating the argument on his innocence defended by my opponent, John's lawyer. I might decide to challenge the backing of such argument, i.e. statistical evidences, by questioning their accuracy, since they refer only to passional crimes and not to premeditate homicides, henceforth being biased and *ad hoc*; to which the opponent might reply that John stands accused to have murdered his wife out of jealousy, hence statistics on passional crimes are absolutely to the point; and so on. In this case, the backing becomes the claim of a complex sub-argument aimed to assess its validity (Fig. 5).

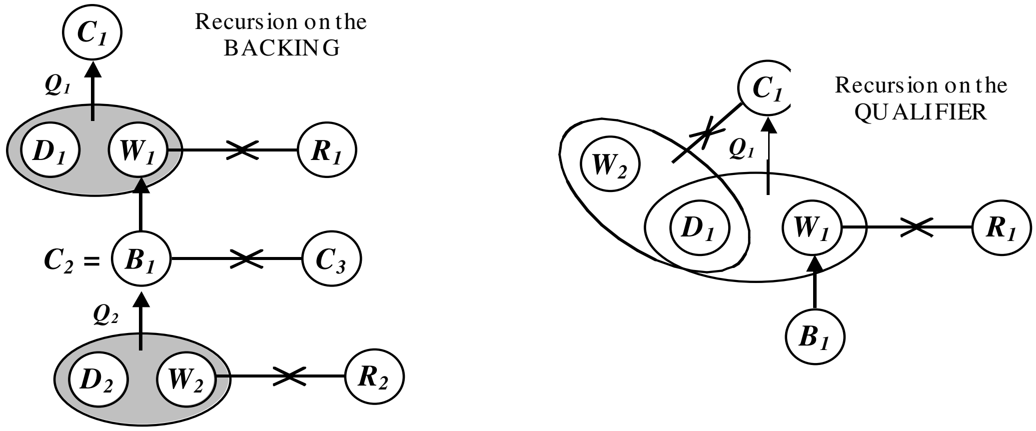

Fig. 5. Two cases of recursion in argument structures

Alternatively, I might attack the qualifier of the original argument (usually), by pointing at specific circumstances, such as jealousy, in which the very fact that John did love his wife might had been a motive to kill her – hence concluding that loving one's own wife does not 'usually' entail not being ready to kill her. For the structural analysis of argumentation, this case is less obvious than the previous one: since the qualifier is not an independent node in the structure, but rather a specification of the triadic relation between data, warrant and claim, it cannot be attacked by opposing to it some counter-argument. To challenge the qualifier, I must show that the *same data* presented by my opponent, considered in the light of a *different warrant*, do not support the conclusion he wanted to draw, or support it to a minor extend, or even contrast it (as in our example). In this way, I indirectly argue that the original qualifier overestimated the bearing of data and warrant on the claim (Fig. 5).

A similar analysis applies also to recursion on data, warrants and rebuttals. The explanation of such a precise and 'natural' mapping between recursive arguments and data structures is straightforward: arguments are dialogical devices to provide reasons to believe a given claim – but rational belief is not born in isolation, i.e. it always requires a network of mutual support to be formed and maintained. Isolated beliefs are mere matters of faith – and even in this case, the human mind struggles to provide some rational foundation to its conviction, and it is not by chance that the most famous philosophical argument was devoted to prove the existence of God. Structural integration of available information is a basic feature of our mind, which constraints both belief dynamics and argumentative persuasion.

A final interesting point on recursion is that it has to *stop* somewhere – and, indeed, it does. As remarked by Toulmin, without some arbitrary limitation to recursive arguments, argumentation itself would be a pointless (and endless) occupation.

> Some warrants must be accepted provisionally without further challenge, if argument is to open to us in the field in question: we should not even know what sort of data were of the slightest relevance to a conclusion, if we had not at least a provisional idea of the warrants acceptable in the situation confronting us. The existence of considerations such as would establish the acceptability of the most reliable warrants is something we are entitled to take for granted. (1958[2003], pp. 98-99)

Exactly the same might be said of belief: we would never believe anything, nor we would decide anything, without first taking some background knowledge for granted (Stalnaker, 1984; Bratman, 1992). In order to reason at all, we must reason finitely, and much of whatever goes on in our mind, it goes on unchallenged and unchecked – although usually it *could* be challenged and checked, should the need arise. So the fact that argumentation conforms to the limitations of our recursive reasoning is neither surprising nor inspiring. But what *is* inspiring are rather the specific ways in which these very limitations can actively be exploited to achieve persuasion. While the mere fact that we need to start somewhere, both in believing and arguing, is trivial, *where we decide to start* is not trivial at all.

5. BEYOND TOULMIN'S MODEL

The utility of DBR in understanding argumentation is not limited to represent some features of Toulmin's layout. In a previous work (Paglieri and Castelfranchi, 2005) we discussed how representing argumentation structures in DBR sheds light on some problems that were left unsolved, vague, or simply unnoticed in Toulmin's account of arguments. Here we focus only on a specific feature: i.e., *implicitness vs. explicitness* of arguments and beliefs. We already argued (cf. 4.2) that any argument, as any belief, is developed against a background of claims that are taken for granted (Burke, 1985; Gough and Tindale, 1985), i.e. which are not in themselves to be disputed and settled – although they might be. However, how each opponent choose and manipulate such background is an important part of argumentation, which relates to the process of *focusing* in DBR (cf. 3.1). Arguing is not only a matter of supporting or challenging specific claims, but also a struggle to *conveniently direct the attention of our opponent*, to make him focus on the strengths of our arguments, while letting pass unnoticed their weaknesses. We often hide some of our premises exactly because we would have no sound reasons to support them, should they be challenged explicitly. Consider again the example in 4.2: when I, trying to convict John for murdering his wife, argue that the very fact he loved her might be a motive for the crime, since jealousy can make us kill the ones we love most, I am grounding this argument on the implicit premise that John was indeed jealous of his wife – which is not the same as to say he loved her. In fact, loving her was just a possible (not sufficient) reason to be jealous of her, and much more would be needed to effectively support my suggestion of a passional crime: e.g., John should be proved to be inclined to violent passion, and ready to act on it; and I should also show that there were indeed overt hints in John's wife behaviour as to ignite his alleged jealousy. By hiding the premise 'John was jealous of his wife', I am just trying to dispense with all this complex (and slippery) sub-argument. If the opponent let go this step unchallenged, my strategy is successful, since I will have forced him to cope with a counter-argument that he might have easily disposed of.

Neither we want to provide our opponent with fuel for his own arguments, i.e. hints and suggestions that might help him supporting his claim or challenging ours. So we try to avoid giving away information that might be harmful for the case we want to make, and we actively attempt to distract the opponent's attention from information he might already possess, but whose salience he has not yet realized. Imagine our friend John has an elderly mother, of whom he is incredibly fond, who in turn cherished John's wife as the apple of her eye. My opponent, John's lawyer, is perfectly aware of this issue, but he has not yet raised it, nor he seems likely to do so in the future: he might possibly underestimate the impact of such emotional argument on the jury (e.g., 'John would have never risked to heartbreak his dear mother by killing her beloved daughter-in-law!'), or he is simply missing the point. In any case, I shall pay great attention to avoid mentioning John's mother in my own arguments, lest my opponent is reminded of such potential counter-argument and decide to use it.

Toulmin mentioned only in passing the parts of an argument that are left untold, i.e. *implicit*, and claimed that implicitness should be considered a distinctive feature of warrant, in contrast with data (1958[2003], p. 92). Although it might be true that warrants are more frequently taken for granted than data, this is not a necessary rule: the same argument can easily, and sometimes effectively, be presented by stating explicitly the warrant rather than the data, as we have just seen ('It is because jealousy can make you kill your most beloved ones, that I suspect John to be the murderer of his wife'). Here two different issues influence each other: (1) whether a given information is *assessed* (i.e. claimed to be or not to be the case) or *not assessed*; and (2) the fact that this happens *explicitly* or *implicitly*, i.e. within or outside the scope of the agent attention. Integrating these independent distinctions, we obtain a taxonomy of four *degrees of focusing*: whenever an information is explicitly assessed, it is *argued for*, i.e. submitted to full scrutiny, focusing both the claim and its supporting structure; in contrast, a claim is *presupposed* whenever some commitment on its value is assumed but left implicit, i.e. the claim itself is focused by the arguers, but not the structural relations that might verify or falsify it (cf. also issues vs. assumptions in Verheij, 2001). In parallel, some claims are not assessed but just left aside: they are either available in principle to the arguer, but not explicitly focused (i.e., *passed under silence*), or they are utterly *unknown*.

Such taxonomy yields a topological representation of focusing, in relation to the data available to the agent (Fig. 6). A claim which is argued for is directly under the spotlight of the arguers' attention, while a presupposed fact is (literally) a borderline case: the claim itself is under consideration, but not the doxastic structure behind it, i.e. its supports and contrasts. A fact passed under silence stands in a shadow zone, where it remains available within the mind of the arguers, but it is not being currently focused; finally, an unknown fact is simply represented as falling out of the agent data structure.

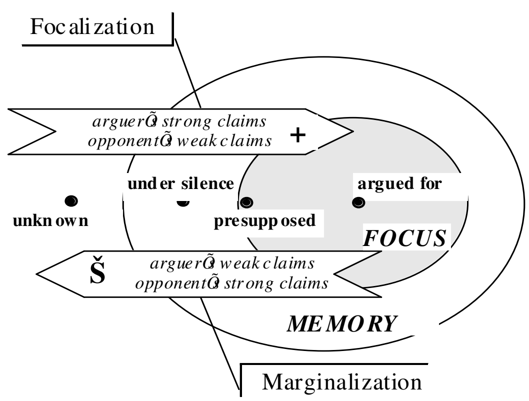

Fig. 6. Degree of focusing: topological and dynamical model

This results in an elementary prescriptive model of *focusing dynamics* in argumentation (represented by the arrows in Fig. 6), depending on the arguer's perceived soundness of her claims. Whenever she considers one of her claims to be well-founded, she will try to focus it in the debate; conversely, she will try to marginalize her weakest claims, and to keep out of the field of discussion facts that might endanger her own case. As for the claims of her opponent, this strategy is reversed: she will try to focus on the weakest spots in the opponent's arguments, while putting aside the strongest claims. This is achieved by shifting the focus of the opponent's attention from a certain set of issues to a different one, e.g. by manipulating the surface form of our arguments, i.e. how they are phrased.

6. CONCLUSIONS AND FUTURE WORKS

The main implication of this preliminary proposal is to initiate a systematic effort of integrating research areas necessarily connected with each other, i.e. argumentation and belief revision, which only rarely have been modelled in the same framework (Pollock and Gillies, 2000; Falappa et al., 2002). To this purpose, DBR makes a first step towards richer models of the complex socio-cognitive dynamics involved in belief change, in contrast with the dominant idealistic approach (Castelfranchi, 1996; Pollock and Gillies, 2000; Paglieri, 2004).

With reference to Toulmin's layout of argument, our analysis emphasizes its relevance for understanding not only argumentation, but belief dynamics as well. Toulmin's theory turns out to be crucial both for a critical assessment of the existing formalisms of doxastic change (cf. 2), and as an inspiring frame of reference for new models (cf. 4 and 5). However, while Toulmin's layout was ideal to preliminarily check the convergence between belief change and argumentation, more sophisticated theories of argumentation will have to be compared with and within the DBR framework, such as pragma-dialectics (van Eemeren and Grootendorst, 1992) and argumentation schemes (Walton, 1996) – with special emphasis on computational models of argumentation (Falappa et al., 2002; Reed and Norman, 2004).

Finally, the DBR model itself requires refinements and fine-tuning: in current and future works, we aim to extend the computational treatment of data properties to motivational and emotional features, i.e. relevance and likeability (Castelfranchi, 1996; Paglieri, 2004), to explore the interaction between Truth-Maintenance Systems (Doyle, 1992), DBR and belief structures, and to move towards implementation in agent-based cognitive and social simulation (Paglieri and Castelfranchi, 2004). As a starting point, we plan to use argumentation tasks as testing ground for belief revision algorithms, and vice versa – building on the general results discussed here.

REFERENCES

Alchourrón, Carlos, Peter Gärdenfors and David Makinson: 1985, 'On the logic of theory change: Partial meet contraction and revision functions', *Journal of Symbolic Logic* **50**, 510-530.

Boutilier, Craig: 1998, 'A unified model of qualitative belief change: a dynamical systems perspective', *Artificial Intelligence* **98**, 281-316.

Bratman, Michael: 1992, 'Practical reasoning and acceptance in a context', *Mind* **101**, 1-15.

Burke, Michael: 1985, 'Unstated premises', *Informal Logic* **7**, 107-118.

Castelfranchi, C.: 1996, 'Reasons: Belief support and goal dynamics', *Mathware & Soft Computing* **3**, 233-247.

Doyle, Jon: 1992, 'Reason maintenance and belief revision: Foundations vs. coherence theories', in Peter Gärdenfors (ed.), *Belief revision*, Cambridge University Press, Cambridge MA, pp. 29-51.

Dragoni, Aldo F. and Paolo Giorgini: 2003, 'Distributed belief revision', *Autonomous Agents and Multi-Agent Systems* **6**, 115-143.

van Eemeren, Frans and Rob Grootendorst: 1992, *Argumentation, communication and fallacies*, LEA, Hillsdale.

Falappa, Marcelo A., Gabriele Kern-Isberner and Guillermo Simari: 2002, 'Explanations, belief revision and defeasible reasoning', *Artificial Intelligence* **141**, 1-28.

Gärdenfors, Peter: 1988, *Knowledge in flux: Modelling the dynamics of epistemic states*, MIT Press, Cambridge.

Gough, James and Christopher Tindale: 1985, 'Hidden or missing premises', *Informal Logic* **7**, 99-106.

Harman, Gilbert: 1986, *Changes in view: principles of reasoning*, The MIT Press, Cambridge MA.

Levi, Isaac: 1991, *The fixation of belief and its undoing*, Cambridge University Press, Cambridge MA.

Paglieri, Fabio: 2004, 'Data-oriented belief revision: Towards a unified theory of epistemic processing', in Eva Onaindia and Stephen Staab (eds.), *Proceedings of STAIRS 2004*, IOS Press, Amsterdam, pp. 179-190.

Paglieri, Fabio and Cristiano Castelfranchi: 2004, ‘Revising beliefs through arguments: Bridging the gap between argumentation and belief revision in MAS’ in Iyad Rahwan, Pavlos Moraïtis and Cris Reed (eds.): *Argumentation in Multi-Agent Systems*, Springer, Berlin.

Perelman, Chaim and Lucie Olbrechts-Tyteca: 1969, *The new rhetoric*, Univ. of Notre Dame Press, Notre Dame.

Pollock, John L. and Anthony S. Gillies: 2000, ‘Belief revision and epistemology’, *Synthese* **122**, 69-92.

Reed, Cris and Timothy J. Norman (eds.): 2004, *Argumentation machines*, Kluwer, Dordrecht.

Rescher, Nicholas: 1976, *Plausible reasoning*, Van Gorcum, Amsterdam.

Rott, Hans: 2001, *Change, choice and inference: A study of belief revision and nonmonotonic reasoning*, Oxford University Press, Oxford.

Stalnaker, Robert C.: 1984, *Inquiry*, The MIT Press, Cambridge MA.

Toulmin, Stephen: 1958/2003, *The uses of argument*, Cambridge University Press, Cambridge MA.

Verheij, Bart: 2001, ‘Evaluating argument based on Toulmin’s scheme’ in *Proceedings of OSSA 2001: Argumentation and its applications*.

Walton, Douglas N.: 1996, *Argumentation schemes for presumptive reasoning*, LEA, Mahwah NJ.