Breaking the Chains. Following-from and Transitivity[*]

November 5, 2013

1 Introduction and Overview

I start with some rather dogmatic statements, simply in order to fix a specific enough framework against which to investigate the topic of this paper. The reader who does not share some or all of the doctrines thereby expressed is invited to modify the rest of the discussion in this paper in accordance to her favourite views on logical consequence (I'd myself sympathize in some cases with such a reader).

Sometimes, some things logically *follow from* some things. The former are a logical *consequence* of the latter and, conversely, the latter logically *entail* the former. Something

[*]Earlier versions of the material in this paper have been presented in 2008 at the 3rd International Conference on Philosophy (ATINER, in Athens) and at the Arché Philosophy of Logic Seminar (University of St Andrews); in 2009, at the Formal Philosophy Research Seminar (Catholic University of Leuven), at a research seminar at the University of Aarhus and at the 7th GAP Conference on *Reflections and Projections—Challenges to Philosophy* (University of Bremen); in 2010, at the Workshop on *Vagueness and Self-Reference* (New University of Lisbon/University of Évora, in Lisbon); in 2011, at the 4th Workshop on *Vagueness and Physics, Metaphysics, and Metametaphysics* (University of Barcelona). I would like to thank all these audiences for very stimulating comments and discussions. Special thanks go to Cian Dorr, Hartry Field, Patrick Greenough, Richard Heck, Dan López de Sa, the late Bob Meyer, Josh Parsons, Graham Priest, Steve Read, Sven Rosenkranz, Marcus Rossberg, Ricardo Santos, Kevin Scharp, Stephen Schiffer, Moritz Schulz, Johanna Seibt, Daniele Sgaravatti, Stewart Shapiro, Martin Smith, Robbie Williams, Crispin Wright and an anonymous referee. I am also grateful to the editors Colin Caret and Ole Hjortland for inviting me to contribute to this volume and for their support and patience throughout the process. At different stages during the writing of the paper, I have benefitted from a RIP Jacobsen Fellowship, an AHRC Postdoctoral Research Fellowship, a UNAM Postdoctoral Research Fellowship and the FP7 Marie Curie Intra-European Research Fellowship 301493, as well as from partial funds from the project FFI2008-06153 of the Spanish Ministry of Science and Innovation on *Vagueness and Physics, Metaphysics, and Metametaphysics*, from the project FFI2011-25626 of the Spanish Ministry of Science and Innovation on *Reference, Self-Reference and Empirical Data*, from the project CONSOLIDER-INGENIO 2010 CSD2009-00056 of the Spanish Ministry of Science and Innovation on *Philosophy of Perspectival Thoughts and Facts* (PERSP) and from the FP7 Marie Curie Initial Training Network 238128 on *Perspectival Thoughts and Facts* (PETAF).

logically[1] following "from nothing" is a *logical truth* (see Zardini [2013d] for some reason to doubt the general correctness of this vulgate about logical truth). I will assume that the things that can stand in such relations are *sentences* (see Zardini [2013m] for a defence of this assumption).

While there may be ways of following-from other than the logical one, even restricting to the logical way of following-from this may be determined by different features of the sentences in question. Assuming a *semantic* individuation of sentences ("two" sentences are the same only if they mean exactly the same things), the relevant features may be:

- The *semantic value* of some expressions or others (together with syntactic structure and the identities of all occurring expressions), as in 'John is unmarried' 's following from 'John is a bachelor';
- The *semantic value* of some expressions belonging to a *privileged class* of "*logical constants*" (again together with syntactic structure and the identities of all occurring expressions), as in 'Snow is white' 's following from 'Snow is white and grass is green';
- The *semantic structure* of the sentences, determined by the semantic categories of the expressions occurring in them and their modes of composition (again together with syntactic structure and the identities of all occurring expressions), as in 'New York is a city' 's following from 'New York is a great city' (see Evans [1978]; Sainsbury [2001], pp. 359–64 for some discussion);
- The sheer *identities* of the sentences, as in 'Snow is white' 's following from 'Snow is white' (see Varzi [2002], pp. 213–4; Moruzzi and Zardini [2007], pp. 180–2 for critical discussion).

The last three cases are usually considered to be cases of "*formal* consequence". In the following, we focus attention on them, even though, on my view and as I have just set up things, the interesting divide between logical following-from and the rest does not coincide with the divide between formal logical following-from and the rest, the latter simply arising from a division of the theoretically basic notion of logical following-from into subnotions which are individuated by the specific features of the sentences which, in each particular case, make it so that the relation of following-from obtains (see García-Carpintero [1993] for a similar view). Unfortunately, I won't have anything more to say here about the analysis of either the basic notion of logical following-from or the derivative notion of formal logical following-from (see Moruzzi and Zardini [2007], pp. 161–74 for a critical survey of the main approaches to the analysis of logical following-from and formal logical following-from).

In consequence, sentences are often "*put together*". Their mode of being put together is signalled in English by 'and' in the locution ' 'Q_0' *and* 'Q_1' *and* 'Q_2'... follow from 'P_0'

[1] For readability's sake, I will henceforth mostly drop in such locutions qualifications like 'logically' and its like.

and 'P_1' *and* 'P_2'...'[2] and cannot be assumed to be less structured than the mode of being put together enjoyed by the coordinates of a *sequence*. This is so because many logics (i.e. those that are *non-contractive*) give divergent answers to the questions:

(i) Whether ψ follows from φ;[3]

(ii) Whether ψ follows from φ and φ

(I myself have developed and defended one such logic for solving the *semantic paradoxes*, see Zardini [2011]; [2013b]; [2013e]; [2013f]; [2013g]). But φ and φ form the same plurality, set, fusion, aggregate, compound etc. as φ. Some logics (i.e. those that are *non-commutative*) even give divergent answers to the questions:

(i′) Whether χ follows from φ and ψ;

(ii′) Whether χ follows from ψ and φ

(I myself have developed and defended one such logic for representing *cross-contextual consequence*, see Zardini [2013j]). But φ and ψ form the same multi-set as ψ and φ (see Restall [2000] and Paoli [2002] for useful overviews of non-contractive and non-commutative logics). We will thus adopt the convention of representing with sequences such a fine-grained mode of putting together sentences (and will in turn adopt the convention of representing sequences with functions whose domain is some suitable initial segment of the ordinals and whose range is a subset of the formulae of the relevant language).

An *argument* is a structure representing some or no sentences (the *conclusions* of the argument) as following from some or no sentences (the *premises* of the argument). In English, an argument is usually expressed with a discourse of the form 'P_0; (*and*) P_1; (*and*) P_2... Therefore, Q_0; (*or*) Q_1; (*or*) Q_2...'. An argument is *valid* iff the conclusions in effect follow from the premises. An *inference* is an act of drawing outputs (mostly, conclusions or whole arguments) from inputs (mostly, premises or whole arguments). (I should stress for future reference that I am not equating inferring with drawing conclusions conforming to a certain collection of syntactic rules.) A *derivation* is an abstract codification of an inference conforming to a certain collection of syntactic rules.

What does it mean to draw (accept, reject, doubt etc.) conclusions when they are not *one* conclusion, but *many* (as it might be the case in our general multiple-conclusion framework)? As a first approximation, we can say that, while premises have to be treated (accepted, rejected, doubted etc.) "*conjunctively*", conclusions have to be treated (accepted, rejected, doubted etc.) "*disjunctively*". Focussing on acceptance, it is important

[2] Most of the time, I will be assuming a *multiple-conclusion* framework: just like no/one/many premises can occupy the first place of the relation of entailment, so can no/one/many conclusions occupy the second place of that relation (see the informal explanation to follow in this section of what it means to reason in a multiple-conclusion framework).

[3] Throughout, 'φ', 'ψ' and 'χ' (possibly with numerical subscripts) are used as metalinguistic variables ranging over the set of formulae of the relevant language.

to note that accepting "disjunctively" a sequence should not be interpreted as accepting every (or even some) of its coordinates (i.e. sentences in its range). For example, in most multiple-conclusion logics, 'There are 1,963 houses in St Andrews', 'It is not the case that there are 1,963 houses in St Andrews' follows from 'Either there are 1,963 houses in St Andrews or it is not the case that there are 1,963 houses in St Andrews', and in most of these logics one may well rationally accept the latter while having no idea of how many houses there are in St Andrews. Of course, if one is in such a situation, it is not only false that one is committed to accepting *both* 'There are 1,963 houses in St Andrews' *and* 'It is not the case that there are 1,963 houses in St Andrews'; it is also false that *either* one is committed to accepting 'There are 1,963 houses in St Andrews' *or* one is committed to accepting 'It is not the case that there are 1,963 houses in St Andrews'. One is only committed as it were "disjunctively" to accepting 'There are 1,963 houses in St Andrews', 'It is not the case that there are 1,963 houses in St Andrews', which can very roughly be characterized as a commitment to accepting *that* either 'There are 1,963 houses in St Andrews' is true or 'It is not the case that there are 1,963 houses in St Andrews' is true.[4] Our multiple-conclusion framework also allows for the cases of *no* premises and *no* conclusions. For reasons we don't need to go into here, and focussing on acceptance and rejection, we can assume that one always accepts and never rejects no premises and that one always rejects and never accepts no conclusions (and, in the terminology to be introduced in section 4.3, that one does so for non-inferential reasons).

A major part of the philosophical investigation of consequence consists in an attempt at elucidating its *nature*—what consequence consists in. Yet, consequence is also a relation, and as such one can sensibly ask what its *formal properties* are.[5] Arguably, Tarski's most notorious contribution to the philosophical investigation of consequence is constituted by his theory of what consequence consists in: truth preservation in every model (see Tarski [1936]). An at least equally important contribution to such investigation is however represented by his earlier studies concerning an abstract theory of consequence relations, aimed at determining the formal behaviour of any such relation. In Tarski [1930], pp. 64–65, he mentions four properties a consequence relation worthy of this name must have:

[4]Another reason for not equating accepting (rejecting, doubting etc.) a sequence with accepting (rejecting, doubting etc.) all (or some) of its coordinates emerges already in the *conjunctive* case (where the idea is presumably that of accepting (rejecting, doubting etc.) *all* of its coordinates) if we pay heed to non-contractive and non-commutative logics. For accepting (rejecting, doubting etc.) φ, φ would then be the same as accepting (rejecting, doubting etc.) φ, and accepting (rejecting, doubting etc.) φ, ψ would be the same as accepting (rejecting, doubting etc.) ψ, φ. Notice that some philosophical applications of such logics may in fact not draw these distinctions for *every* attitude: for example, on my non-contractive view of the semantic paradoxes, accepting φ, φ is the same as accepting φ—it is for attitudes like rejecting and supposing that bearing that attitude towards φ, φ is not the same as bearing it towards φ.

[5]Compare with resemblance: a major task for resemblance theories is to determine what resemblance between two individuals consists in (sharing of universals, matching in tropes, primitive similarity etc.); yet, the study of the formal properties of resemblance (seriality, reflexivity, symmetry etc.) can fruitfully be pursued even in the absence of an answer to the question about its ultimate nature.

reflexivity, monotonicity, transitivity and compactness.[6,7]

I think all of these properties are at least questionable (see Moruzzi and Zardini [2007], pp. 180–187). But here I want to focus on *transitivity*, trying to make sense of a position according to which consequence is not transitive. Although there are many different transitivity properties of various strength (see Zardini [2013d]), for our purposes we can simply work with the following principle:[8]

(T) If, for every $\varphi \in \text{ran}(\Theta)$, $\Gamma \vdash \Delta, \varphi$ and $\Lambda, \Theta \vdash \Xi$, then $\Lambda, \Gamma \vdash \Delta, \Xi$.[9,10]

Having proposed elsewhere a solution to the sorites paradox which consists in placing some principled restrictions on transitivity (see the family of *tolerant logics* developed in Zardini [2008a]; [2008b], pp. 93–173; [2009]; [2013a]; [2013c]; [2013d]; [2013h]; [2013l]), I do have a heavy investment in the issue. However, here I will not argue directly for any position according to which consequence is not transitive. Rather, I will simply try to make it adequately intelligible and assess what impact its correctness would have on our understanding of consequence. The discussion will typically be conducted at a high level of generality, concerning non-transitive logics as such; however, I cannot deny that, heuristically, I have been led to the conclusions defended here by taking tolerant logics as my main paradigm, and this will no doubt in some cases reflect in the fact that the assumptions made and arguments developed will fit tolerant logics better than other non-transitive logics.[11]

The rest of the paper is organized as follows. To fix ideas, section 2 puts on the table a range of philosophically interesting non-transitive consequence relations, introducing briefly their rationale. Section 3 discusses and disposes of two very influential objections of principle to the use of non-transitive consequence relations. Section 4 delves into some fine details of the metaphysical and normative structures generated by non-transitivity. Section 5 concludes by placing the foregoing investigations within the wider context of the relationships between consequence and rationality.

[6]Important as they may be, it is worth stressing that they would still grossly underdetermine the nature and even the extension of consequence. Consider e.g. that the relation which holds between Γ and Δ iff Γ entails Δ and is non-empty will satisfy all the properties mentioned in the text if consequence does.

[7]Contraction and commutativity are not discussed as they are simply built into Tarski's framework, since he puts premises together in sets.

[8]Throughout, 'Γ', 'Δ', 'Θ', 'Λ', 'Ξ', 'Π', 'Σ' (possibly with subscripts) are used as metalinguistic variables ranging over the set of sequences of formulae of the relevant language; Γ, Δ is the obvious composition of Γ with Δ; $\langle . \rangle$ is the empty sequence. If no ambiguity threatens, I will identify $\langle \varphi \rangle$ with φ.

[9]Throughout, $\text{ran}(\Gamma)$ is the range of Γ; '$\vdash$' (possibly with subscripts) denotes the relevant consequence relation.

[10]In some respects, (T) does not embody a *pure idea* of transitivity, since it also implies versions of monotonicity and contraction. More guarded (and convoluted) transitivity principles which do not have suchlike implications are available. However, I will mostly stick to (T) for its relative simplicity, which will help me to focus on the issues I want to focus on here.

[11]And also fit some tolerant logics (for example, the tolerant logics in favour of whose adoption I myself have argued for in the works just referenced in the text, as well as many other relatively weak tolerant logics) better than others (for example, the very strong tolerant logics that have subsequently been studied in more detail by Cobreros et al. [2012]).

2 Non-Transitive Consequence Relations

2.1 Relevance

In the following, it will be helpful to have in mind concrete examples of non-transitive consequence relations—the different informal pictures underlying the various examples will help us to see (some of) the different ways in which one can so understand consequence as to make sense of failures of transitivity.[12] Of course, again, the aim here is not to argue for the adoption of any such logic, but only to shed light on their rational motivation.

Consider the two following arguments:

CONTRADICTION Intuitively, 'Graham Priest is wrong' does not follow from 'The Strengthened-Liar sentence is true and the Strengthened-Liar sentence is not true'. Graham Priest's error should not be entailed by the correctness of one of his most famous doctrines. Yet:

(i) Both 'The Strengthened-Liar sentence is true' and 'The Strengthened-Liar sentence is not true' do seem to follow from 'The Strengthened-Liar sentence is true and the Strengthened-Liar sentence is not true' (by simplification);

(ii) 'The Strengthened-Liar sentence is true or Graham Priest is wrong' does seem to follow from 'The Strengthened-Liar sentence is true' (by addition);

(iii) 'Graham Priest is wrong' does seem to follow from 'The Strengthened-Liar sentence is true or Graham Priest is wrong' and 'The Strengthened-Liar sentence is not true' (by disjunctive syllogism).

Two applications of (T) would then yield that 'Graham Priest is wrong' does after all follow from 'The Strengthened-Liar sentence is true and the Strengthened-Liar sentence is not true'.

LOGICAL TRUTH Intuitively, 'Casey is a male and Casey is a sibling iff it is not the case that [Casey is not a male or Casey is not a sibling]'[13] does not follow from 'Casey is a brother iff [Casey is a male and Casey is a sibling]'. No De Morgan Law should be entailed by an analysis of 'brother'. Yet:

(i′) 'Casey is a brother iff it is not the case that [Casey is not a male or Casey is not a sibling]' does seem to follow from 'Casey is a brother iff [Casey is a male and Casey is a sibling]' and 'Casey is a male and Casey is a sibling iff it is not the case that [Casey is not a male or Casey is not a sibling]' (by transitivity of the biconditional), and so from the former only (by a version of suppression of logical truths);

[12] The reader should keep in mind that, throughout, what is meant by 'failures of transitivity' and its like is simply failures of transitivity *for certain argument forms* and its like. The most interesting non-transitive logics retain transitivity for many argument forms.

[13] Throughout, I use square brackets to disambiguate English constituent structure.

(ii′) ‘Casey is a brother iff [Casey is a male and Casey is a sibling]’ does seem to follow from itself (by reflexivity);

(iii′) ‘Casey is a male and Casey is a sibling iff it is not the case that [Casey is not a male or Casey is not a sibling]’ does seem to follow from ‘Casey is a brother iff [Casey is a male and Casey is a sibling]’ and ‘Casey is a brother iff it is not the case that [Casey is not a male or Casey is not a sibling]’ (by transitivity of the biconditional).

One application of (T) would then yield that ‘Casey is a male and Casey is a sibling iff it is not the case that [Casey is not a male or Casey is not a sibling]’ does after all follow from ‘Casey is a brother iff [Casey is a male and Casey is a sibling]’.

Some (Bolzano [1837]—at least according to George [1983]; [1986]—and then Lewy [1958], pp. 123–132; Geach [1958]; Smiley [1959], pp. 238–243; Walton [1979]; Epstein [1979]; Tennant [1987], pp. 185–200, 253–265)[14] have taken taken these or similar intuitive judgements at face value and concluded that (T) does not unrestrictedly hold (see Lewy [1976], pp. 126–131; Routley et al. [1982], pp. 74–78 for critical discussion of this approach).[15] One possible way of elaborating the rationale for such judgements would be as follows (see von Wright [1957], pp. 175, 177). Apart from no-premise and no-conclusion arguments, consequence should hold in virtue of some *genuine relation* between the *contents* of the premises and the *contents* of the conclusions—its holding should never be determined merely by the independent logical status of some premises (as logical falsities) or conclusions (as logical truths). Thus, apart from no-premise and no-conclusion arguments, consequence should never discriminate between logical truths and falsities on the one hand and logical contingencies on the other hand.

This intuitive constraint can then be made precise as the “filter condition” that an argument is valid iff it can be obtained by substitution of sentences for atomic sentences from a classically valid argument none of whose premises or conclusions are logical truths or falsities (see Smiley [1959], p. 240).[16] The filter condition yields the desired results: it can easily be checked that each subargument of **CONTRADICTION** and **LOGICAL TRUTH** satisfies the filter condition, even though the overall arguments do not.

In particular, as for **CONTRADICTION**, notice that the two subarguments in (i) are valid, as they can be obtained from the valid argument ‘Snow is white and grass is green. Therefore, snow is white (grass is green)’ by substitution of ‘The Strengthened-Liar sentence is true’ for ‘Snow is white’ and of ‘The Strengthened-Liar sentence is not true’ for

[14] Sylvan [2000], pp. 47–49, 98–99 intriguingly mentions some possible medieval and early-modern sources, but, to the best of my knowledge, a satisfactory investigation into the pre-Bolzanian history of this logical tradition has yet to be undertaken.

[15] While, as noted in fn 10, (T) in general implies versions of monotonicity, and while there are a variety of logics in which considerations of relevance lead to restrictions of monotonicity, it is worth noting that the particular applications of (T) involved in **CONTRADICTION** and **LOGICAL TRUTH** are unexceptionable in many such logics.

[16] The filter condition is so-called as it acts as a filter on the classically valid arguments, selecting only the “good” ones (so that validity is in effect defined as classical validity plus something else).

‘Grass is green’; the other subarguments satisfy the filter condition already in their present form. As for **LOGICAL TRUTH**, notice that the first subargument in (i′) is valid, as it can be obtained from the valid argument ‘Snow is white iff grass is green; grass is green iff water is blue. Therefore, snow is white iff water is blue’ by substitution of ‘Casey is a brother’ for ‘Snow is white’, of ‘Casey is a male and Casey is a sibling’ for ‘Grass is green’ and of ‘It is not the case that [Casey is not a male or Casey is not a sibling]’ for ‘Water is blue’; the subargument in (iii′) is valid, as it can be obtained from the valid argument ‘Snow is white iff grass is green; snow is white iff water is blue. Therefore, grass is green iff water is blue’ by the same substitutions; the other subarguments satisfy the filter condition already in their present form. Given what counts as *meaning connection* in this framework, addition of (T) would lead to a gross overgeneration of meaning connections between sentences—the genuine intensional dependencies between the premises and conclusions of each subargument would overgenerate into the bogus intensional dependencies between the premises and conclusions of the overall arguments.

It is often claimed that the imposition of this and similar filter conditions on the consequence relation amounts to changing the subject matter of logic, and does not really engage with the traditional view according to which “*truth preservation*” is what consequence is all about. Notice that such a claim cannot be addressed in the particular case of **CONTRADICTION** by holding that some contradictions are true whilst not everything is true, so that the argument would fail to be truth preserving in the straightforward sense of having true premises and false conclusions. For the truth of only some contradictions would presumably invalidate subargument (iii) (*qua* not truth preserving in the straightforward sense) and therefore prevent a possible failure of transitivity (see Priest [2006], pp. 110–122).

The claim is however highly dubious on other grounds. For it is plausible to assume that there is a notion of *implication*[17] (expressible in the language by $\rightarrow$) such that, if Δ, ψ follows from Γ, φ, then $\Delta, \varphi \rightarrow \psi$ follows from Γ. At least on some readings, ‘If φ, then ψ’ does presumably express such a notion in English. But, on such a reading, ‘If the Strengthened-Liar sentence is true and the Strengthened-Liar sentence is not true, then Graham Priest is wrong’ seems to be false (indeed false-only) if anything is. Under the present assumption, it would, however, be a logical truth if the overall argument of **CONTRADICTION** were valid. Even though, we may assume, truth preserving *in the straightforward sense* of not having true premises and false conclusions, such an argument would thus not be truth preserving *in the only slightly less straightforward sense* of being such that its validity would imply the validity of arguments which are not truth preserving in the straightforward sense. Under the present assumption about the behaviour of implication, there is thus a perfectly good sense in which, for someone attracted by the framework just sketched, licensing the validity of the overall argument of **CONTRADICTION**, (T) would indeed lead to failures of truth preservation—what consequence is supposed to be all about (see section 3.2 for more on transitivity and truth preservation and Read [1981]; [2003] for a different reply on the issue of relevance and truth preservation).

[17] Henceforth, by ‘implication’ and its like I mean any operation expressed by some conditional or other.

2.2 Tolerance

Consider the three following premises:

(1) A man with 0 hairs is bald;

(2) A man with 1,000,000 hairs is not bald;

(3) If a man with i hairs is bald, so is a man with $i+1$ hairs.

All these premises are intuitively true, and, presumably *a fortiori*, consistent. However, from (3) we have that, if a man with 0 hairs is bald, so is a man with 1 hair, which, together with (1), yields that a man with 1 hair is bald. Yet, from (3) we also have that, if a man with 1 hair is bald, so is a man with 2 hairs, which, together with the previous lemma that a man with 1 hair is bald, yields that a man with 2 hairs is bald. With another 999,997 structurally identical arguments, we reach the conclusion that a man with 999,999 hairs is bald. From (3) we also have that, if a man with 999,999 hairs is bald, so is a man with 1,000,000 hairs, which, together with the previous lemma that a man with 999,999 hairs is bald, yields that a man with 1,000,000 hairs is bald. 999,999 applications of (T) (and of monotonicity) would then yield that the contradictory of (2) follows simply from (1) and (3).

In my works on vagueness referenced in section 1, I have taken these intuitive judgements at face value and concluded that (T) does not unrestrictedly hold (earlier, Weir [1998], pp. 792–794; Béziau [2006] had also briefly entertained this possibility).[18] One

[18]If I believe that (T) does not unrestrictedly hold for vague discourses, am I not barred from accepting any of its applications involving vague expressions? That would be extremely problematic, since most of our reasoning involves vague expressions and such reasoning often seems to involve applications of (T) (in fact, the arguments in this paper are no exception!). Before addressing this issue, it is important to realize that it does not constitute an *idiosyncratic* problem for my view on vagueness, but that it represents a *general* problem common to many other cases of logical deviance. The common problem can abstractly be described as the one of recovering what are, even from the point of view of a deviant logician, intuitively acceptable applications of a logical principle whose unrestricted validity she has rejected precisely for the discourses to which those applications belong. Thus, to take a different instance of the same problem, a deviant logician who, because of the semantic paradoxes, has rejected the unrestricted validity of the law of excluded middle for discourses about truth will still want to accept an instance of that law such as 'Either 'Socrates was in the agora on his 23$^{\text{rd}}$ birthday' is true or it is not'. (Indeed, there is arguably a *parallel* problem for transitivists who reject tolerance principles like (3), since they need to recover intuitively acceptable applications of (3) and its like, as when, for example, after noticing that one has lost 1 hair one infers, from the fact that one was not bald before, that one is still not bald.) The problem has at least two aspects: at the *philosophical* level, how to motivate the distinction between acceptable and unacceptable applications of a not unrestrictedly valid principle; at the *logical* level, how to recover the acceptable applications. In the case of tolerance-driven non-transitivism, both aspects would deserve a much more extended treatment than can be given in this paper, but let me briefly sketch my own take on them (focussing on those applications of (T) that are indeed not formally valid—as I have already mentioned in fn 12, the most interesting non-transitive logics retain (T) for many argument forms). As for the first aspect, the only fault the tolerance-driven non-transitivist finds with (T) is that, by sorites reasoning, it allows one to use intuitively correct tolerance principles like (3) to go down the slippery

possible way of elaborating the rationale for these judgements would be as follows. A major point of a vague predicate is to draw a difference in application between some cases which are far apart enough on a dimension of comparison relevant for the application of the predicate. The predicate should *discriminate* between some such cases. Hence, (1) and (2) must be enforced. Still, another major point of a vague predicate is not to draw any difference in application (from a true application to anything falling short of that) between any two cases which are close enough on a dimension of comparison relevant for the application of the predicate. The predicate should *not discriminate* between any two such cases. Hence, (3) must be enforced. Moreover, instances of (3) should allow *modus ponens*: what substance is there to the idea that there is no sharp boundary between i and $i + 1$ in matters of baldness if, given the premise that a man with i hairs is bald, it does not follow that a man with $i + 1$ hairs is bald? Given what counts as *indifference connection* in this framework, addition of (T) would lead to a gross overgeneration of indifference connections between sentences—the correct mandate of not drawing any distinction in matters of baldness between people that differ of 1 hair would overgenerate into the incorrect mandate of not drawing any distinction in matters of baldness between people that differ of 1,000,000 hairs.

2.3 Evidence

Finally, I would also like to put on the table a case of a non-deductive, *defeasible* consequence relation—that is, roughly, a relation which is supposed to hold between premises and conclusions iff the truth of some of the latter is *reasonable* by the lights of the state of information represented by the former, even if not *guaranteed* by their truth (as is supposed to be the case for a *deductive* consequence relation). For example, the inference of 'Al is a native speaker of Italian' from 'Al was born in Little Italy' is eminently reasonable (all else being equal), as well as the inference of 'Al was born in Italy' from 'Al is a native speaker of Italian'. However, the inference of 'Al was born in Italy' from 'Al was born in Little Italy' is eminently unreasonable. Under this intuitive understanding of defeasible consequence:

(I) 'Al is a native speaker of Italian' is a consequence of 'Al was born in Little Italy';

slopes associated with vague predicates. She will thus regard an application of (T) as unacceptable only if it is in this sense "*soritical*" (while intuitive enough for the purposes of this fn, in a fuller treatment the notion certainly could and should be made more precise). On this view, many applications of (T) remain acceptable even if they involve vague expressions, since such applications will not involve tolerance principles at all or, even if they do, will still be non-soritical. As for the second aspect, the tolerance-driven non-transitivist could simply take acceptable applications of (T) to be, although not *formally* valid, *materially* valid in virtue of the fact that the specific occurrences of non-logical expressions in them determines that they are non-soritical. Alternatively, and focussing on the tolerant logics advocated in my works on vagueness referenced in section 1, the tolerance-driven non-transitivist could add *further principles* to her theory which, together with the initial premises of a target acceptable application of (T), formally entail (in those tolerant logics) the final conclusions of that application (see Zardini [2013d] for some details). Thanks to an anonymous referee for pressing me on these issues.

(II) 'Al was born in Italy' is a consequence of 'Al is a native speaker of Italian';

(III) 'Al was born in Italy' is not a consequence of 'Al was born in Little Italy',

whereas, given (I) and (II), (T) would rule out (III).

One possible way of elaborating the rationale for these judgements would be as follows. Whereas on all *probability functions* reasonable in the light of how things actually are the conditional probabilities of 'Al is a native speaker of Italian' and 'Al was born in Italy' on 'Al was born in Little Italy' and 'Al is a native speaker of Italian' respectively are both very high, the conditional probability of 'Al was born in Italy' on 'Al was born in Little Italy' is very low (if not $= 0$!).

The idea can be made more precise in different specific ways. Here is a fairly general recipe. Let a model $\mathfrak{M}$ of a language $\mathscr{L}$ be a probability function on the sentences of $\mathscr{L}$. Let the conditional probability functions be *totally* defined—assume to that effect a suitable probability calculus (for well-known examples, see Popper [1955]; Rényi [1955]). The probability function corresponding to $\mathfrak{M}$ will thus be in effect determined by the set of conditional probability functions—the unconditional probability in $\mathfrak{M}$ of φ being the conditional probability in $\mathfrak{M}$ of φ on ψ, for some logical truth ψ (assuming a suitable logic in the characterization of the probability calculus). Define then Δ to be a δ*-consequence* of Γ in $\mathfrak{M}$ iff the conditional probability in $\mathfrak{M}$ of the disjunction of all the coordinates of Δ on the conjunction of all the coordinates of Γ is $\geq \delta$. Finally, define Δ to be a δ-consequence of Γ iff, for every model $\mathfrak{M}$, Δ is a δ-consequence of Γ in $\mathfrak{M}$. Under natural assumptions, it's easy to check that, as defined, for every $\delta > 0$, δ-consequence is just classical consequence (for $\delta = 0$, δ-consequence is of course trivial). The desired supra-classical strength comes by restricting the range of admissible models to a set of (contextually determined) "reasonable" probability functions.[19]

No doubt this recipe still leaves a lot of leeway in the choice of the probability calculus and of the appropriate restrictions on models. However, it seems plausible that, whichever specific implementation is eventually chosen, the consequence relation so obtained will have a decent claim to codify at least in part the (contextually determined) canon of defeasible reasoning. If so, such a canon will not satisfy (T): given what counts as *probabilification connection* in this framework, addition of (T) would lead to a gross overgeneration of probabilification connections between sentences—the reasonable rules

[19]Literature on defeasible consequence usually discusses the *cumulative-transitivity* principle:

(CT) If, for every $\varphi \in \text{ran}(\Theta)$, $\Gamma \vdash \Delta, \varphi$ and $\Gamma, \Theta \vdash \Delta, \Xi$, then $\Gamma \vdash \Delta, \Xi$

instead of (T) (for the reason I have mentioned in fn 10). Interestingly, although they constitute a counterexample to (T), (I)–(III) do not constitute a counterexample to (CT), for 'Al was born in Italy' is not a consequence of 'Al is a native speaker of Italian' *and* 'Al was born in Little Italy'. Indeed, while they fail to satisfy (T), many defeasible consequence relations do satisfy (CT). However, it is easy to see that the kind of defeasible probabilistic consequence relation defined in the text will fail to satisfy (CT) as well as (T) in many natural cases (see fn 47 for a concrete example). Thanks to Cian Dorr, Sven Rosenkranz and Martin Smith for discussions of this issue.

of thumb that people born in Little Italy (typically) speak Italian and that Italian speakers were (typically) born in Italy would overgenerate into the unreasonable rule of thumb that people born in Little Italy were (typically) born in Italy.

It might legitimately be wondered what the point is of introducing a *defeasible* consequence relation in the course of an attempt at understanding the idea that *deductive* consequence is non-transitive. Yet, it will turn out that one of the gateways to this understanding is constituted by an appreciation of the *normative import of consequence on rational attitudes*. The non-transitivist can be made sense of as interpreting such import in a particular way. However, that import is not of course a privilege of deductive consequence, and so study of defeasible consequence relations with certain properties (such as the failure of transitivity) may well shed light on some features of deductive consequence relations with the same properties. Indeed, especially in section 4.3, I will argue for the claim that the normative force of non-transitive deductive consequence shares crucial features with the normative force of defeasible consequence (in particular, of defeasible probabilistic consequence).[20,21]

3 Two Objections to the Very Idea of Non-Transitive Consequence

3.1 Consequence, Inference and Derivation

I can see (and have encountered, in print and conversation) two main objections concerning the very idea of non-transitive consequence, objections which, if correct, would seem to doom from the start any interesting use of a non-transitive consequence relation. Their rebuttal will help to dispel some misunderstandings of what non-transitivity of consequence amounts to. A more positive characterization will be offered in section 4.

There should be an uncontroversial sense in which logic is oftentimes substantially

[20]In fact, there seem to be some even deeper connections between probabilistic structure and at least some kinds of non-transitive deductive consequence relations. For example, the tolerant logics advocated in my works on vagueness referenced in section 1 fail to satisfy some principles (such as implication in the premises, reasoning by cases, conjunction in the conclusions etc.) which are also invalid from a probabilistic point of view, even though they are valid on many non-probabilistic codifications of defeasible consequence, such as mainstream non-monotonic logics (see Makinson [2005] for a useful overview of these). I reserve to future work the exploration and discussion of these more specific connections.

[21]Defeasible consequence relations have many structural similarities with *conditionals*. Conditionals arguably fail to satisfy the relevant analogues of (T) and (what is seldom noticed) of (CT) (not to speak of monotonicity), for reasons very similar to those for which certain defeasible consequence relations fail to satisfy (T) and (CT). That is generally a pertinent observation given the tight connections between consequence and conditionals (indeed, it is attractive to see consequence as a special, limit case of implication, see section 3.2). The observation is even more pertinent given that the problems and views I will discuss in sections 3 and 4 find their analogues in the case of conditionals. For lack of space, in this paper I will however have to focus on consequence relations and leave for another occasion the application of the framework developed in sections 3 and 4 to the case of non-transitivity of conditionals. Thanks to Sven Rosenkranz for impressing upon me the importance of this connection.

informative. There should be an uncontroversial sense in which reasoning oftentimes leads to the *discovery* of new truths. For example, there should be an uncontroversial sense in which the derivation from a standard arithmetical axiom system of the conclusion that there are infinitely many prime numbers is rightly regarded as a substantial discovery about natural numbers, accomplished by purely logical means (under the assumption of the truth of the axioms), no matter what one's views are about the ultimate aptness of logical vocabulary to represent features of the world. However, as Quine among others has stressed (see e. g. Quine [1986], pp. 80–94), the *elementary steps* of logic are, in a sense, *obvious*.[22] In what sense can then logic still be substantially informative?

Well, as the traditional thought about this puzzle has always been at pains to stress, a series of completely obvious elementary steps may well lead to a completely unobvious conclusion. In other words, the relation x-is-an-obvious-consequence-of-y is non-transitive. However, it is important to note that this observation by itself does not yet explain away the puzzle: a series of one-foot steps may well lead to cover a considerable distance, but no one has ever supposed this to show that one can take a much longer step than the length of one's legs would give reason to suppose. Why is then the fact that φ_i is reachable in i elementary steps from φ_0 supposed to show that there is some interesting connection between φ_i and φ_0? The crucial, implicit, auxiliary assumption must be that consequence, as opposed to obvious consequence, is indeed transitive, so that the path to an unobvious final conclusion can be obliterated, in the sense that the conclusion can be seen as already following from the initial premises. As Timothy Smiley so nicely put it, "the whole point of logic as an instrument, and the way in which it brings us new knowledge, lies in the contrast between the transitivity of 'entails' and the non-transitivity of 'obviously entails'" (Smiley [1959], p. 242). Thus, there seems to be little space for consequence not to be transitive, if logic is to preserve its function as a means of discovering new truths.

I accept that logic is to preserve such a function, but, as it stands, I contest the very coherence of the previous objection. Suppose that ψ obviously follows from φ and that χ obviously follows from ψ but not so obviously from φ. If consequence is indeed transitive, then, *at least as far as consequence is concerned*, the step from φ to χ is no more mediated, and thus no less "elementary", than those from φ to ψ and from ψ to χ: χ just follows straight from φ as it follows straight from ψ (and as ψ follows straight from φ)—the idea that χ bears the more complex relation to φ of being reachable in two "elementary" steps from it but not in one is simply a non-transitivist illusion. If then the step from φ to χ is indeed unobvious, it is just not true, in a transitivist framework, that every "elementary" logical step is obvious. Thus, the solution offered to the puzzle (transitivity) simply denies one of the elements from which the puzzle arose. To repeat, in a transitive logic the step from φ to χ is "elementary" in the sense that, *at least as far as consequence is concerned*, it is no more mediated than those from φ to ψ and

[22]Quine uses 'obvious' in such a way that an apparently deviant logician would be best translated non-homophonically (see also Quine [1960], pp. 57–61). I find very dubious that even the elementary steps of logic are "obvious" in this sense and do not mean anything that strong—only that, once one has adopted a reasonable logic, its elementary steps look most straightforward and least informative, unlike the step from a standard arithmetical axiom system to Euclid's theorem. Thanks to Graham Priest and Stewart Shapiro for stimulating discussions of Quine's views on these matters.

from ψ to χ. Indeed, in this sense, in a transitive logic *every* valid step—no matter how remotely connected the premises and the conclusions may in effect be—is an "elementary" step. The root of the trouble is then that the transitivity solution simply obliterates the *structure* presupposed by the puzzle, namely the distinction between elementary logical steps and non-elementary ones, *as long as that structure is supposed to be generated by the relation of consequence* (a supposition without which, as we are about to see, the objection against the non-transitivist evaporates).

This strongly suggests that the elementary/non-elementary distinction has been mislocated by the objection. It is not a distinction to be drawn at the level of *consequence*—rather, it is a distinction to be drawn at the level of *inference*. For our purposes, the notion of an elementary inference can be left at an intuitive level: it is understood in such a way that the inference of φ from 'φ and ψ' is elementary, whereas that of '[φ or χ] and [ψ or υ]' from 'φ and ψ' is not. This is just as good, since the puzzle was an *epistemic* one and the drawing of an inference is one of the ways in which the validity of an argument can be *recognized* by a subject. The transitivist herself has thus to acknowledge that the puzzle should properly be stated as the puzzle of how any inference can be substantially informative given that every elementary inference is least informative. Her own solution to the puzzle may then be revised as follows. x-is-elementarily-inferrable-from-y is non-transitive, yet it entails x-is-inferrable-from-y. Furthermore, by the soundness of inferrability (with respect to consequence), x-is-inferrable-from-y entails x-is-a-consequence-of-y. The transitivity of the latter in turn ensures that an unobvious final conclusion already follows from the initial premises.

The revised argument is coherent and might seem to offer a satisfactory explanation of the revised puzzle from a transitivist perspective. However, if x-is-inferrable-from-y is itself transitive, the appeal to the transitivity of consequence is now exposed as a superfluous detour. For a simpler argument can run as follows. x-is-elementarily-inferrable-from-y is non-transitive, yet it entails x-is-inferrable-from-y. The transitivity of the latter in turn ensures that an unobvious final conclusion can already be inferred from the initial premises, and so, by the soundness of inferrability (with respect to consequence), that the unobvious final conclusion follows from the initial premises.

Now, it is true that, for the little that has been said so far, x-is-inferrable-from-y need not be transitive from a non-transitivist point of view. Yet, any theoretical employment of logic by a being whose cognitive architecture resembles that of humans is likely to require a *systematization* of pre-theoretical judgements of validity in terms of a small set of syntactic rules which are at least sound (and possibly complete) with respect to consequence. It is simply a fact that such a systematization provides one of the most effective methods a human (and anyone with a similar cognitive architecture) can employ in order to explore what follows from what (just as something along the lines of the standard rules for addition, subtraction, multiplication and division provides one of the most effective methods a human (and anyone with a similar cognitive architecture) can employ in order to explore what gives what).

To stress, I don't think that there is an immediate connection between inference and syntactic rules: I think it's clear on reflection that one can draw an inference simply in virtue of one's appreciation of the logical concepts involved in it, even if one is not in possession of a set of syntactic rules which would allow the relevant derivation. Yet, as I have just noted, given the way human cognition works, there is every theoretical reason for transitivists and non-transitivists alike to accept a systematization of our pre-theoretical judgements of validity in terms of syntactic rules. And while there is no conceptual bar to the derivational system thus generated being itself non-transitive, it is of course a desirable epistemic feature also from the non-transitivist perspective that the derivational system used in the study of a non-transitive consequence relation be itself transitive, in the sense that *any new application of a syntactic rule preserves the property of being a correct derivation*. This feature is desirable exactly because non-transitive consequence can be just as unobvious as transitive consequence, whence the need arises for derivational techniques offering the epistemic gain flowing from the asymmetry between the non-transitivity of x-is-obviously-derivable-from-y (x-is-elementarily-derivable-from-y) and the transitivity of x-is-derivable-from-y (for an example of non-transitive logics with transitive sound and complete derivational systems see Zardini [2013l]).[23] No further asymmetry between the non-transitivity of x-is-an-obvious-consequence-of-y and the alleged transitivity of x-is-a-consequence-of-y is required.

3.2 Consequence and Truth Preservation

We have already encountered in section 2.1 the (rather vague) claim that consequence is all about truth preservation, in the sense that whether the consequence relation holds between certain premises and conclusions is *wholly* determined by whether the conclusions preserve the truth of the premises. On the basis of an intuitive understanding of such notion of truth preservation, the following objection can be mounted.

The non-transitivist, we may assume, claims, for some φ, ψ and χ, that ψ follows from φ and χ from ψ, but χ does not follow from φ (i.e. that (T) fails in some single-premise and single-conclusion case). So, since truth preservation is necessary for consequence, the non-transitivist should concede that ψ preserves the truth of φ and that χ preserves the truth of ψ. Yet, on the face of it, truth preservation is a transitive relation: if ψ preserves the truth of φ and χ preserves the truth of ψ, then it would seem that χ also preserves the truth of φ. For suppose that ψ preserves the truth of φ and that χ preserves the truth of ψ, and suppose that φ is true. Then, since ψ preserves the truth of φ, ψ should be true as well. But χ preserves the truth of ψ, and so, since ψ is true, χ should be true as well.

[23] To get a concrete sense of how a transitive derivational system can be sound and complete with respect to a non-transitive consequence relation, just take your favourite transitive derivational system presented in sequent-calculus style which does not admit of cut elimination. Throw the cut rule out of the derivational system (if it was included in it) and take your consequence relation to be simply the set of sequents provable in the resulting derivational system. The resulting derivational system itself is obviously sound and complete with respect to that consequence relation, although the former will still be transitive (in the specific sense just explained in the text) while the latter will be non-transitive (in the sense of failing to satisfy (T)). Similar examples could be provided for instance for derivational systems presented in natural-deduction style by suitably complicating with provisos the derivational system's syntactic rules. Thanks to Marcus Rossberg and Moritz Schulz for discussions on this issue.

Thus, under the supposition that φ is true (and that ψ preserves the truth of φ and that χ preserves the truth of ψ), we can infer that χ is true. Discharging that supposition, we can then infer (still under the suppositions that ψ preserves the truth of φ and that χ preserves the truth of ψ) that, if φ is true, so is χ, which might seem sufficient for χ's preserving φ's truth. Having thus argued that truth preservation is transitive, the objection is completed by noting that truth preservation suffices for consequence, so that, since χ preserves the truth of φ, χ follows from φ, contrary to what the non-transitivist claims.

A proper assessment of the objection requires first an adequate explication of the underlying notion of truth preservation. I know of no better way of spelling out this notion than in terms of a *conditional* statement: certain conclusions preserve the truth of certain premises iff, *if* every premise is true, *then* some conclusion is true (where, for the purposes of the discussion to follow, we can afford to remain rather neutral as to the exact behaviour of the conditional). Henceforth, I will officially adopt this explication of the notion (save for briefly considering a possible alternative understanding at the end of this section). Of course, so stated in terms of an unadorned indicative conditional, truth preservation alone is not *sufficient* for consequence: if 'A match is struck' is true, so is 'A match will light', but the latter does not in any sense logically follow from the former. Many authors would be sceptical that any sufficient non-circular strengthening of the intensional force of the conditional is available, but I remain optimist (see Zardini [2013f]; [2013g]). In the interest of generality, I will however assume that there is no interesting sense of 'truth preservation' in which truth preservation is sufficient for consequence, and show how the objection from truth preservation against the non-transitivist can be modified so as to need only the assumption that truth preservation is *necessary* for consequence (in any event, my reply to that modified objection will apply with equal force to the original one).

Clearly, if truth preservation is not sufficient for consequence, the original objection against the non-transitivist presented at the beginning of this section fails. In particular, the objection breaks down at the very last step 'since χ preserves the truth of φ, χ follows from φ', which assumes the sufficiency of truth preservation for consequence. Yet, a more sophisticated version of the objection could still be made to run against some applications that the non-transitivist envisions for her logics (for example, to the sorites paradox). To see this, let us classify three degrees of "*non-transitive involvement*":

- An application of non-transitivity which only requires a verdict of invalidity concerning a transitivistically valid target argument is *weak*;
- An application which in addition requires the rejection of a commitment to accepting the conclusions of a transitivistically valid target argument all of whose premises are accepted[24] is *intermediate*;
- An application which in addition requires the falsity of all the conclusions of a

[24] For simplicity's sake, I will henceforth take sentences as the objects of acceptance and rejection. The whole discussion may be recast, more clumsily, in terms of acceptance of propositions.

transitivistically valid target argument all of whose premises are true is *strong*.[25]

Focus then on strong applications of non-transitivity (an analogous discussion could be run for intermediate applications).[26] For some of these, the non-transitivist wishes to maintain that φ is true but χ false, even though ψ follows from φ and χ from ψ. However, without assuming the sufficiency of truth preservation for consequence, the objection from truth preservation presented at the beginning of this section can still be used to reach the conclusion that truth preservation is transitive, which, together with the necessity of truth preservation for consequence, yields that χ preserves the truth of φ—i.e., if φ is true, so is χ.[27] And that conditional sits badly with the other commitments (truth of φ, falsity of χ) the non-transitivist would wish to undertake (exactly how badly it sits will depend of course on the details of the logic—for example, these claims are jointly inconsistent in the tolerant logics advocated in my works on vagueness referenced in section 1). Now, because of various considerations relating to context dependence and to indeterminacy which would lead us too far afield to rehearse here, I would be wary of principles stating that certain versions of truth preservation are necessary for validity (see Zardini [2012]; [2013h]). Still, none of the considerations underlying that caution target the necessity of truth preservation for the kinds of arguments that are typically considered in strong applications of non-transitivity, so that I am willing to concede that, for those cases, truth preservation is indeed necessary for consequence.[28]

[25]Interestingly, our three examples of applications of non-transitivity in section 2 suffice to cover all the three degrees of non-transitive involvement: the application presented in section 2.1 is weak, the one presented in section 2.3 is intermediate and the one presented in section 2.2 is strong.

[26]Henceforth, I will mainly talk about non-transitive logics tailored to *intermediate* and *strong* applications of non-transitivity, as these are arguably the philosophically most problematic and interesting cases to be made sense of. This qualification must be understood as implicit in the following. However, although I cannot give a detailed treatment of the issue here, I hasten to add that this is very likely not a real limitation. For, in the case of *suppositional* acceptance if not in that of *flat-out* acceptance, weak applications will very likely have to replicate metaphysical and normative structures identical to those which I will argue are generated by intermediate and strong applications.

[27]In fact, at least one of the most prominent strong applications of non-transitivity, i.e. the one to the sorites paradox, presents some recalcitrance against being fit into the mould used in the text. For that is an application with *multiple-premise* arguments, so that transitivity of truth preservation is not applicable to such chains of arguments in the direct way exploited by the modified objection in the text. One way to recover an application with multiple-premise arguments would be to modify the relevant arguments, so as to bring all the premises $\varphi_0, \varphi_1, \varphi_2 \ldots \varphi_i$ used at some point or other in the chain up front, collect them together in a single, long "conjunction" (at least assuming that they are finite) and carry them over from conclusion to conclusion adding to the relevant intermediate conclusion the "conjuncts" $\varphi_0, \varphi_1, \varphi_2 \ldots \varphi_i$ (scare quotes being used here since, for example in tolerant logics, standard conjunction has not the right properties to do the job and a new operation would have to be introduced instead). An alternative way to recover an application with multiple-premise arguments would be to use a generalized version of transitivity of truth preservation, which, given the relevant chain of arguments $\mathfrak{a}_0, \mathfrak{a}_1, \mathfrak{a}_2 \ldots \mathfrak{a}_i$ and under the assumption that all the side premises in those arguments are true, says that, if each of the arguments is truth preserving, so is the argument whose premises are those of $\mathfrak{a}_0$ and whose conclusion is that of $\mathfrak{a}_i$.

[28]Some authors have recently started to talk as though the semantic paradoxes gave a reason to reject that truth preservation is necessary for consequence even in the cases that I myself would regard as unproblematic (see e.g. Field [2006]). The situation is however more neutrally described in terms of

Before giving my reply to the modified objection from truth preservation, I would like to undertake a brief digression on truth preservation which I hope will prove instructive for understanding some features of non-transitive logics. One might think that any worry about the connection between consequence and truth preservation is simply misplaced because of the fact that consequence is trivially guaranteed to be equivalent with the preservation of at least *some kind* of truth, even if not truth *simpliciter*. For consequence is often defined as preservation, for every model $\mathfrak{M}$, of truth in $\mathfrak{M}$. Consequence would then be trivially guaranteed to be equivalent with the preservation of truth in every model.

Setting aside the important question about what kind of truth "truth in a model" really is, what I want to remark on here is that already the identification of consequence with the *preservation* of something or other (in some class of structures or other) is unwarranted in our dialectic, as one can define well-behaved consequence relations without appealing to anything recognizable as preservation of something or other (in some class of structures or other).[29] Tolerant logics are an example of such a consequence relation developed in a non-transitive framework; an example driven by a completely different kind of consideration (not affecting transitivity) can be found in Martin and Meyer [1982]. The general structural point in these logics is that *the collection of designated values relevant for the premises is not identical with the collection of designated values relevant for the conclusions*: in tolerant logics, the former is in effect a proper *subcollection* of the latter, whereas the logic **S** of Martin and Meyer [1982] can be seen as a form of the dual position, in which the former is a proper *supercollection* of the latter (thereby leading to the failure not of transitivity, but of another property considered by Tarski [1930] essential for a consequence relation—namely, *reflexivity*).[30] Thus, in both tolerant logics and **S**, consequence is not equivalent with anything recognizable as preservation of something or other (in some class of structures or other)—it is rather equivalent with the *connection* of a certain collection of designated values with a certain *other* collection of designated values (in every model of the relevant class).

Let me be clear. It is not the case that a representation in terms of non-identity of the two collections of designated values *suffices* to ensure the non-transitivity of a logic (see Smith [2004] for an example of a transitive logic generated by such a representation). Nor is it the case that a representation in terms of non-identity of the two collections is *essential* for every non-transitive logic (see Smiley [1959] for an example of a non-transitive logic

these authors' theories of truth being inconsistent with the principle that truth preservation is necessary for consequence, which would seem (at least to me) to be better taken as a reason for rejecting these authors' theories rather than as a reason for rejecting the principle (at least in the cases that I myself would regard as unproblematic), especially given that there are other theories of truth on the market which are perfectly consistent with the principle and even entail it (see e.g. the theories developed in my works on the semantic paradoxes referenced in section 1).

[29]Even granting an identification of consequence with truth preservation in some class of structures or other, it is well-known that the standard set-theoretic notion of *model*, which replaces the generic notion of *structure*, has serious drawbacks in the analysis of the consequence relation of expressively rich languages. I won't go here into this further aspect of the complex relation between consequence and preservation of truth in a structure (see McGee [1992] for a good introduction to some of these issues).

[30]Thanks to the late Bob Meyer for a very helpful discussion of **S** and to Graham Priest for pointing out to me the duality connection.

where no such natural representation seems to be forthcoming). Yet such a representation, where available, is a fruitful point of entry to at least one of the key thoughts behind the logic (and as such will be deployed in section 4.4). The representation also connects up neatly with usual representations of transitive logics (preservation of designated value in some class of structures or other), showing how non-transitivity arises from a very natural and straightforward generalization of the usual model-theoretic representation of a (transitive) consequence relation.

Moving on now to my reply to the modified objection from truth preservation, I reject that the non-transitivist is committed to χ's being true if φ is. For the most natural way in which this conclusion (and so the transitivity of truth preservation) might be thought to follow from the premises that ψ preserves the truth of φ and that χ preserves the truth of ψ requires the assumption that the consequence relation of the *metalanguage* (the language in which we talk about the truth of φ, ψ and χ) is transitive (as against the non-transitivity of the consequence relation of the *object language* in which we talk about whatever φ, ψ and χ talk about). More explicitly: the validity of the metalinguistic argument 'If 'φ' is true, then 'ψ' is true; if 'ψ' is true, then 'χ' is true. Therefore, if 'φ' is true, then 'χ' is true' (and so, under the current interpretation of truth preservation, the transitivity of truth preservation) boils down to the validity of the argument form 'If P_0, then P_1; if P_1, then P_2. Therefore, if P_0, then P_2', which is however invalid in many non-transitive logics (it is for example invalid in many tolerant logics).[31]

Moreover, virtually no deviance from classical logic worthy of this name should grant the assumption that the consequence relation of the metalanguage $\mathscr{M}$ talking about the truth and falsity of the sentences of the object language $\mathscr{O}$ over which a deviation from classical logic is envisaged should itself be classical. For, given any decent theory of truth, this will be sufficient to reintroduce classical logic in $\mathscr{O}$ itself. Consider for example a deviant intuitionist logician. Were she to accept 'Either 'φ' is true or 'φ' is not true' in $\mathscr{M}$ (for φ belonging to $\mathscr{O}$), the most natural theory of truth (namely one such that φ and ''φ' is true' are fully equivalent,[32] so that in particular ''φ' is not true' entails 'It is not the case that φ') would commit her to 'Either φ or it is not the case that φ'.

The point need not exploit the full equivalence between φ and ''φ' is true' which is induced by the *enquotation/disquotation* schema:

(ED) P iff 'P' is F.

We can produce similar results only e. g. with the right-to-left direction of (ED). If a metalinguistic *necessity* predicate were to behave classically, the intuitionist would have

[31] I do not wish to suggest that the reply I am proposing is the only one available to the non-transitivist (although it is the kind of reply that, subject to the important refinement mentioned in the last point made in fn 36, I do endorse): I will present in section 4.1 a coherent kind of non-transitive position which would rather reject that truth preservation is necessary for consequence (in particular, while that position would accept that χ follows from ψ, it would reject that, if ψ is true, so is χ).

[32] Here and in what follows, someone (not me) might want to make a proviso for the *semantic paradoxes*. Even with that proviso, the argument in the text shows that all of $\mathscr{O}$'s *grounded* sentences behave classically—hardly a pleasing result for the intuitionist.

to accept in $\mathscr{M}$ 'Either 'φ' is necessary or 'φ' is not necessary' (for φ belonging to $\mathscr{O}$). Given that ''φ' is not necessary' is classically equivalent with ''It is not the case that φ' is possible', she would have to accept 'Either 'φ' is necessary or 'It is not the case that φ' is possible', whence, by substituting 'Either φ or it is not the case that φ' for 'φ', she would have to accept 'Either 'Either φ or it is not the case that φ' is necessary or 'It is not the case that either φ or it is not the case that φ' is possible', which entails in (any suitable modal extension of) intuitionist logic ''Either φ or it is not the case that φ' is necessary'. By right-to-left (ED), she would then have to accept 'Either φ or it is not the case that φ'.[33] Indeed, up to the very last step, the previous argument would go through also for necessity-like metalinguistic predicates[34] which fail to satisfy either direction of (ED), such as 'in principle acceptable'. Thus, if a metalinguistic *in-principle-acceptability* predicate were to behave classically, the intuitionist would have to accept in $\mathscr{M}$ ''Either φ or it is not the case that φ' is in principle acceptable' (for φ belonging to $\mathscr{O}$). (Notice that, although failing to satisfy either direction of (ED), 'in principle acceptable' and what is naturally assumed to be its classical dual 'in principle non-rejectable' still connect with the object language in a broad "(ED)-style" fashion in the sense that, roughly, if one knows that φ is in principle acceptable, one should accept φ, and, if φ is inconsistent, φ is not in principle non-rejectable.)

The general observation implicit in the last paragraph is that a deviance from classical logic for a language $\mathscr{L}_0$ should equally apply to a language $\mathscr{L}_1$ as soon as some of the notions expressible in $\mathscr{L}_1$ are best thought of as exhibiting the same problematic properties which motivate a deviance from classical logic for $\mathscr{L}_0$, and that, in some cases, this may be so exactly because of some systematic connections that those notions bear to notions expressible in $\mathscr{L}_0$—for example, connections established by some broadly (ED)-style principle when $\mathscr{L}_1$ is the metalanguage of $\mathscr{L}_0$. Importantly, this general observation fatally affects also revisions of the objection from truth preservation which replace the notion of truth preservation with that (structurally identical) of closure of knowledge or of other epistemic properties under logical consequence (for example, 'known' still satisfies the right-to-left direction of (ED)).[35]

[33]The argument in the text should make it clear what exactly I mean when I say that a metalinguistic predicate Φ "behaves classically". What exactly I mean is that every instance of an argument form valid in classical first-order logic in which only Φ occurs as non-logical constant (treating quotation names as logical constants) is valid and, in addition, that, as is usual in the relevant extensions of classical first-order logic, Φ has a dual Ψ satisfying the schema ''P' is not Φ iff 'It is not the case that P' is Ψ' (in the argument in the text, 'possible' is naturally assumed to be such classical dual for 'necessary'). In particular, notice that, since 'a is F' is not an argument form valid in classical first-order logic, we cannot straightforwardly assume ''Either φ or it is not the case that φ' is necessary' (maybe on the grounds that 'Either φ or it is not the case that φ' is classically valid), whose most fine-grained form is precisely 'a is F'. Only the object-language logic—which in our example is intuitionist—is allowed, as it were, to "look inside" quotation names, whence it is crucial that, in the argument in the text, ''It is not the case that either φ or it is not the case that φ' is not possible' is valid in (any suitable modal extension of) intuitionist logic.

[34]A metalinguistic predicate Φ is *necessity-like* iff Φ is closed under *adjunction* (''φ' is Φ' and ''ψ' is Φ' entail ''φ and ψ' is Φ') and φ entails ''φ' is Φ' if φ is the *trivial truth* (that is, the truth entailed by anything).

[35]Thanks to Stephen Schiffer for interesting discussions on this last kind of objection. Within a tolerant

It is worth seeing in closing how the objection from truth preservation fares if the current (and standard) explication of the notion of truth preservation in terms of the notion of implication is rejected in favour of a *primitive* relation of truth preservation. While under the former explication there was a logical guarantee that the relation is transitive (assuming a transitive logic!), now that guarantee is lost and the alleged transitivity must be postulated as a specific law governing the relation. Pending further argument, there do not seem to be clear reasons for the non-transitivist to accept this postulation. But what is most important to note is that even the acceptance of a transitivity postulate for the relation would not by itself wreck havoc at least for a strong application of non-transitivity like the one to the sorites paradox, for what is really at issue there is not so much simple transitivity, but the stronger assumption that a finite chain of elements connected by a relation R is such that its first element bears R to its last element. This is in effect the property of relations that is sometimes called '*chain transitivity*' (see Parikh [1983]), which, in a non-transitive logic, is usually *stronger* than the property of transitivity (see Zardini [2013d] for details; of course, under minimal assumptions, the two properties are equivalent in a transitive logic). Even though there may ultimately be no need for the non-transitivist to pursue this latter strategy in the case of a primitive truth-preservation relation, I think there is good reason for her to embrace it in other cases in which independent grounds support the transitivity of a particular relation, as in the case of the identity relation (see again Zardini [2013d] for details).[36]

---

logic, the point made in the text can actually be developed into a surprising defence of certain closure principles for knowledge against some very influential objections (see Zardini [2013i]).

[36]Going back to the standard explication of the notion of truth preservation in terms of the notion of implication, I have focussed on the fact that the argument 'If 'φ' is true, then 'ψ' is true; if 'ψ' is true, then 'χ' is true. Therefore, if 'φ' is true, then 'χ' is true' is invalid in many non-transitive logics. That argument was in turn supported in the objection from truth preservation by the argument ' 'φ' is true; if 'φ' is true, then 'ψ' is true; if 'ψ' is true, then 'χ' is true. Therefore, 'χ' is true', which is equally invalid in many non-transitive logics (while the step from the validity of the latter argument to the validity of the former argument—an application of the *deduction theorem*—is valid in many non-transitive logics and can be taken for granted for our purposes). One might object that, never mind what many non-transitive logics say about it, especially the latter argument (a "metalinguistic double-*modus-ponens* argument") seems *intuitively very compelling*, so that its rejection is a cost of the reply I have given on behalf of the non-transitivist. In response to this worry, I would like to make four points. First, as a preliminary clarification of the limited aims of the reply I have given on behalf of the non-transitivist, I should stress that the purpose of the reply is not to argue that there is absolutely nothing to the objection from truth preservation (or to any other objection against the non-transitivist), but only to show that the non-transitivist can take a *coherent* and (at least from her point of view) *very natural* position in reply to that objection *that is compatible with the principle that truth preservation is necessary for consequence*. To be sure, adoption of that position comes with its own costs—just as adoption of any other position in the puzzling areas for which non-transitive logics have usually been proposed does—and rejection of the metalinguistic double-*modus-ponens* argument may well be among those, but, contrary to the spirit with which the objection from truth preservation has often been raised to me, the non-transitivist is not forced to a rather improbable rejection of the principle that truth preservation is necessary for consequence. Second, it is actually unclear (at least to me) that the metalinguistic double-*modus-ponens* argument is intuitively so compelling that its rejection should be regarded as a cost of the reply I have given on behalf of the non-transitivist. What seems clear (at least to me) is that that argument is naturally analyzed as involving two applications of *modus ponens* and a transitivity step, and that both applications of *modus ponens* are indeed intuitively valid; but it is unclear (at least to me) whether chaining those applications

4 Making Sense of Non-Transitivity

4.1 Non-Logical/Logical Dualism and Non-Transitivism

We have only started to scratch the surface of the philosophical underpinnings of a non-transitive logic. The rebuttal of the two previous objections has helped to dispel some misunderstandings of what non-transitivity of consequence amounts to, but not much has yet been offered by way of a positive characterization of a conception of consequence as non-transitive.

I think more progress on this issue can be made by asking why consequence is usually assumed to be transitive. Consider the following natural picture (taking now sets to be the terms of the consequence relation and restricting our attention to single-conclusion arguments). The laws of logic can be seen as an *operation*[37] log which, given a set of facts

together and thus *obliterating the role of the intermediate conclusion* ('ψ' is true') *as a premise* in yielding the final conclusion ('χ' is true') has much intuitive compellingness. Third, it would seem that, if, contrary to the doubt I have just raised, the *metalinguistic* double-*modus-ponens* argument is regarded as intuitively compelling, the simpler *object-language* double-*modus-ponens* argument 'φ; if φ, then ψ; if ψ, then χ. Therefore, χ' should be regarded as no less intuitively compelling. But that argument usually has to be rejected by strong applications of non-transitivity, so that the worry raised would ultimately have nothing to do *in particular* with metalinguistic double-*modus-ponens* arguments, and so nothing to do *in particular* with the metalinguistic principle that consequence requires truth preservation. That is of course not to say that the worry would not be real; only that, contrary to the objection I have discussed in this section, there would be no specific problem of compatibility between non-transitivism and the principle that truth preservation is necessary for consequence. Fourth, I should mention that, at least focussing on the strong application of non-transitivity made in my works on vagueness referenced in section 1, insofar as one finds double-*modus-ponens* arguments intuitively compelling there are on the one hand well-behaved tolerant logics (not those advocated in said works) in which they are valid, and there are on the other hand prospects for upholding their validity also in other tolerant logics by switching to *non-classical* (in fact, *tolerant*) *metatheories in which it is vague what the values relevant for consequence are*. On either scheme, the result would be the analogue for implication of the move discussed in the text for relations: again, what is really at issue in the sorites paradox is not so much a simple-transitivity (i.e. double-*modus-ponens*) argument like 'A man with 0 hairs is bald; if a man with 0 hairs is bald, so is a man with 1 hair; if a man with 1 hair is bald, so is a man with 2 hairs. Therefore, a man with 2 hairs is bald', but a chain-transitivity (i.e. *iple-modus-ponens*) argument like 'A man with 0 hairs is bald; if a man with 0 hairs is bald, so is a man with 1 hair; if a man with 1 hair is bald, so is a man with 2 hairs; if a man with 2 hairs is bald, so is a man with 3 hairs...; if a man with 999,999 hairs is bald, so is a man with 1,000,000 hairs. Therefore, a man with 1,000,000 hairs is bald'. Obviously, on either scheme, truth preservation, even if explicated in terms of implication, is now harmlessly transitive (but not chain transitive). Thus, interestingly, if a non-transitivist pursuing either scheme were to identify consequence with truth preservation, consequence would itself be governed by a non-classical (indeed, non-transitive) logic and would after all be transitive (but not chain transitive—which, again, is what is really at issue in the strong application of non-transitivity to the sorites paradox). Thanks to an anonymous referee for pushing the worry discussed in this fn.

[37] The study of consequence as an *operation* rather than *relation* goes back at least as far as Tarski [1930] (see Wójcicki [1988] for a comprehensive study within this approach). Under the simplifying assumptions made at the beginning of this section, the properties of reflexivity, monotonicity and transitivity of a consequence relation correspond to the following properties of a consequence operation cons:

(i) $X \subseteq \text{cons}(X)$ (*increment*);

S,[38] applies to it yielding with logical necessity another such set $\log(S) = T$ (possibly identical with S)—just like, at least in the deterministic case, the laws of nature can be seen as an operation which, given the state of a system at a certain time, applies to it and yields with natural necessity the states of the system at subsequent times. Importantly, as in the case of truth preservation, the sense in which, given a set of facts S, the laws of logic apply to it yielding with logical necessity T is again best explicated in terms of a *conditional* statement to the effect that, if all the members of S hold, so do with logical necessity all the members of T (where the necessity operator has wide scope over the conditional).

Given this natural picture, the thought in favour of transitivity goes as follows. Presumably, the laws of logic enjoy *universal applicability*: they apply to any set of facts S *whatsoever* yielding with logical necessity $\log(S)$. Given the set of facts S, the laws of logic apply to it yielding with logical necessity T, but, by their universal applicability, given the set of facts $T = \log(S)$ they should also apply to it yielding with logical necessity another set $\log(\log(S)) = U$ (possibly identical with T).[39] Recalling the explication given in the last paragraph, we should then have that, if all the members of S hold, so do with logical necessity all the members of T and that, if all the members of T hold, so do with logical necessity all the members of U, from which it may seem to follow that, if all the members of S hold, so do with logical necessity all the members of U—that is, that the laws of logic, given the set of facts S, apply to it already yielding with logical necessity U. Supposing

(ii) If $X \subseteq Y$, $\mathrm{cons}(X) \subseteq \mathrm{cons}(Y)$ (*monotonicity*);

(iii) $\mathrm{cons}(\mathrm{cons}(X) \cup Y) \subseteq \mathrm{cons}(X \cup Y)$ (*union-adjoint subidempotency*)

when cons is an operation from sets of sentences to sets of sentences. An operation satisfying conditions (i)–(iii) is a *Tarski closure operation*. A generalization of closure operations for modelling multiple-conclusion consequence relations are *Scott closure operations* (see Scott [1974]). Further generalizations dealing with collections more fine-grained than sets are possible (see Avron [1991] for a start). Closure operations have first been identified by Kuratowski [1922], who also included the conditions:

(iv) $\mathrm{cons}(X \cup Y) \subseteq \mathrm{cons}(X) \cup \mathrm{cons}(Y)$ (*preservation of binary unions*);

(v) $\mathrm{cons}(\varnothing) \subseteq \varnothing$ (*preservation of nullary unions*),

which, while making sense for *topological* closure operations, seem rather out of place for *consequence* closure operations, since they in effect amount, respectively, to obliterating the extra logical power given by the *combination of premises* and to forcing that no sentences are *logical truths*. It may be worth mentioning that operations satisfying (i), (ii), (iv) and (v) but not necessarily (iii) (known as '*preclosure operations*' or '*Čech closure operations*') form a well-behaved, well-understood and interesting topological kind (see Čech [1966]).

[38] As usual, I assume that there are such abstract entities as *states-of-affairs* which are described by sentences and which may hold or fail to hold, and I identify *facts* with states-of-affairs that hold. Throughout, 'S', 'T', 'U' and 'V' (possibly with numerical subscripts) are used as variables ranging over the set of sets whose members are states-of-affairs (which may or may not be facts).

[39] Throughout, U is simply understood as $\{s : s$ is a state-of-affairs and, if all the members of T hold, so does with logical necessity $s\}$ (T is understood analogously). Thus, assuming that states-of-affairs always exist, U always exists. But, for every V_0 and V_1, $V_1 = \log(V_0)$ only if all the members of $V_0 \cup V_1$ hold (since log is an operation on sets of *facts*). Thus, in particular, $U = \log(\log(S))$ only if all the members of $\log(S) \cup U$ hold (a condition which is extremely plausibly satisfied under the supposition that all the members of $\log(S)$ hold).

for the rest of this paragraph that S is indeed a set of facts, this would extremely plausibly imply that the result ($U = \log(\log(S))$) of the operation of the laws of logic on the result ($T = \log(S)$) of the operation of the laws of logic on S exists and is included in the result ($\log(S)$) of the operation of the laws of logic on S—log would be subidempotent (which would thus validate a form of transitivity for the correlative relation of consequence).[40] These implications would be multiply repugnant for many applications of non-transitivity. Let us assume again, as at the beginning of section 3.2, that the non-transitivist claims, for some φ, ψ and χ, that ψ follows from φ and χ from ψ, but χ does not follow from φ, and let us also assume that φ describes a fact and S is the singleton of that fact (and that monotonicity holds). Then, one rather crude way to bring out the repugnancy is to notice that, since $\log(\log(S))$ would exist and χ belongs to it, the state-of-affairs described by χ would hold, contrary to what many applications of non-transitivity require (among which all intermediate and strong applications). A more subtle way to bring out the repugnancy is to notice that, since, by subidempotency, $U \subseteq T$, χ, which describes a state-of-affairs belonging to U, would also describe a state-of-affairs belonging to T, contrary to what many applications of non-transitivity require.

However, as in the case of the objection from truth preservation, the conclusion that the laws of logic, given the set of facts S, apply to it yielding with logical necessity U (that is, that, if all the members of S hold, so do with logical necessity all the members of U), in its characteristic obliteration of the logical role played by the implicit intermediate conclusion (that is, that all the members of T hold), implicitly relies on transitivity. Given that, in the sense explained in section 3.2, facts-talk is no less "logically penetrable" than truth-talk, the considerations developed in that section apply here: by the non-transitivist's lights, the repugnant conclusion that the laws of logic, given the set of facts S, apply to it already yielding with logical necessity U does not follow from the assumptions characterizing the natural picture, even when taken together with the further assumption that the laws of logic enjoy universal applicability. Thus, the non-transitivist should not be seen as committed to rejecting any of those assumptions.

The point about the compatibility of non-transitivism with the universal applicability of the laws of logic is crucial. Keeping fixed the natural picture of the laws of logic as an operation on facts, it is very tempting to try to make sense of the non-transitivist's position as precisely rejecting the universal applicability of the laws of logic. More specifically, it is very tempting to try to make sense of the non-transitivist's position as relying on a distinction between *two different kinds of facts*, the non-logical and the logical. *Non-logical facts* are those provided, as it were, by the world itself, such as the fact that snow is white; the fact that, if snow is white, it reflects light; the fact that every piece of snow is white (I cannot but emphasize, as the last two examples make clear, that non-logical facts need not in any sense be "atomic"). *Logical facts* are those that hold in virtue of the application of the laws of logic to the non-logical facts, such as the fact that either

[40] An operation op is *subidempotent* on a set S and ordering $\leq$ iff, for every $x \in S$, $\mathrm{op}(\mathrm{op}(x)) \leq \mathrm{op}(x)$. Under the simplifying assumptions made at the beginning of this section, in our case $\leq$ is simply subset inclusion. It's easy to check that, with conditions (i) and (ii) of fn 37 in place, subidempotency on subset inclusion implies union-adjoint subidempotency.

snow is white or grass is blue (holding in virtue of the application of the laws of logic to the fact that snow is white); the fact that snow is white and grass is green (holding in virtue of the application of the laws of logic to the fact that snow is white and the fact that grass is green); the fact that something is white (holding in virtue of the application of the laws of logic to the fact that snow is white). Of course, much more would have to be said about how to draw exactly the non-logical/logical distinction, but I take it that we have an intuitive grasp of it (as witnessed by our intuitive judgements in the foregoing cases) which will be sufficient for our purposes.

The non-transitivist would then be seen as advocating what we may call '*non-logical/logical dualism*', consisting in *rejecting the universal applicability of the laws of logic to logical facts*. Thus, in our original case, since T contains the fact described by ψ and, we may assume, such fact is logical, it is open to the dualist non-transitivist to reject that the laws of logic apply to T, and thus to reject that they yield U (in particular, that they force the state-of-affairs described by χ to hold). As should be clear, I am understanding the dualist non-transitivist as someone who not only rejects the *conditional* that, if all the members of T hold, so do all the members of U, but also refuses to *infer* the conclusion that all the members of U hold from the premise that all the members of T hold (we will see in sections 4.2 and 4.3 that, as opposed to her rejection of the conditional, her rejection of the inference is what she shares with a non-dualist non-transitivist). Still, I want to understand the dualist non-transitivist's rejection of the inference as *grounded* in her rejection of the conditional, and more specifically as grounded in her idea that the antecedent of the conditional ('All the members of T hold') might hold while its consequent ('All the members of U hold') fails to hold. I also want to understand the dualist non-transitivist as taking an analogous position in the case of truth preservation.[41] Thus, in our original case, she would maintain that ψ might be true while χ is not, and so would reject that, if ψ is true, so is χ, contrary to the position I recommended to the non-transitivist in section 3.2 (see fn 31).

It is interesting to observe that, under a certain extremely plausible assumption, non-logical/logical dualism *by itself* requires restrictions on the transitivity of consequence. For, whenever a restriction of the applicability of the laws of logic to some non-empty set[42] of logically contingent facts described by the premises $\varphi_0, \varphi_1, \varphi_2 \ldots$ and logically necessary facts described by the premises $\psi_0, \psi_1, \psi_2 \ldots$ is envisaged (so that the inference to the conclusion χ, which follows from them, is rejected), for every k φ_k must itself follow from some non-empty set X_k whose members describe non-logical facts. This is so because, even if not every logically contingent fact is non-logical (consider for example the fact that something is white), it is extremely plausible to assume that every logically contingent fact is ultimately grounded in certain non-logical facts in such a way as to be

[41]As should be clear, I see a structural identity between what a non-transitivist can or should say about *truth* and what she can or should say about *facts*. All the points made about non-transitivism and truth are meant to apply just as well to non-transitivism and facts and *vice versa*, even though sometimes the point is more easily made about non-transitivism and truth and some other times about non-transitivism and facts.

[42]If the set is empty, then it is vacuously the case that all its members are non-logical facts, and so there is no dualist bar to the application of the laws of logic to it.

yielded from them by an application of the laws of logic (here I won't try though to justify this assumption). If transitivity were then to hold unrestrictedly, χ would already follow from $X_0 \cup X_1 \cup X_2 \ldots$. Since however the members of this set all describe non-logical facts, there would be no dualist bar to the application of the laws of logic to it, and so the state-of-affairs described by χ would be forced to hold as well.

Interesting as non-logical/logical dualism may be, as I have already explained four paragraphs back a non-transitivist is not committed to it (and can indeed accept its negation). I should also remark, however, on an important point of *agreement* between the dualist and the non-dualist non-transitivist. For reflect that, in a non-transitive framework, the dualist non-transitivist can be seen as doing one thing by means of a quite different one. That is, to go back to our original case, she can be seen as rejecting a commitment to accepting a consequence (χ) of what she is committed to accepting (ψ) by rejecting that, if ψ is true, so is χ (I take it to be extremely plausible that the dualist non-transitivist is committed to accepting ψ, since it follows from φ which, we may assume, describes a non-logical fact).[43] Even if a non-transitivist can disagree with the latter, she cannot but agree with the former, since, as I will explain in sections 4.2 and 4.3, the rejection of being committed to accepting a consequence of what she is committed to accepting is arguably the *crux* of her disagreement with the transitivist, one of the places at which their different *metalinguistic judgements* about validity are finally reflected in a clash of *object-language attitudes* (in this case, the transitivist's acceptance of χ against the non-transitivist's non-acceptance of χ). A non-transitivist need not disagree with the transitivist as to whether, if ψ is true, so is χ, but she has to disagree with the transitivist's willingness to undertake a commitment to χ's being true on the sole basis of her logical commitment to ψ's being true.

It should by now be clear that, in opposition to the dualist non-transitivist, a non-dualist non-transitivist can legitimately insist on not explaining her position in terms of some *deviant conception of truth*—in particular, in terms of a denial of the principle that truth preservation is necessary for consequence (claiming that, although χ follows from ψ, it is not the case that, if ψ is true, so is χ). That is helpfully compared with the fact that someone who rejects the *law of excluded middle* need not explain her position in terms of a deviant (e. g. *gappist*) conception of the relation between the truth of a sentence and the truth of its negation (e. g. such that it might be the case that neither a sentence nor its negation are true; see Field [2008] for a non-gappist rejection of the law of excluded middle), or with the fact that someone who rejects the rule of *disjunctive syllogism* need not explain her position in terms of a deviant (e.g. *dialetheist*) conception of the relation between the truth of a sentence and the truth of its negation (e.g. such that it might be the case that both a sentence and its negation are true; see Read [1988] for a non-dialetheist rejection of the rule of disjunctive syllogism). As in all those other cases, the heart of the logical deviance being proposed by the non-dualist non-transitivist is a

[43] One might wonder how plausible this rejection on the part of the non-logical/logical dualist is. Given that χ follows from ψ, the conditional statement that, if ψ is true, so is χ follows by the *deduction theorem* (and what I take to be utterly uncontroversial instances of (ED) for truth). It would thus seem that the non-logical/logical dualist is committed to an implausible rejection of the deduction theorem.

certain conception of what counts as a *correct pattern of reasoning* rather than a certain deviant conception of what truth is.

4.2 Logical Nihilism and Non-Transitivism

Still, regarded now as a proposal as to what counts as a correct pattern of reasoning, non-transitivism may look perilously close to a *logical nihilism* which rejects the universal validity of all the argument forms that trigger failures of (T) when satisfying the second conjunct of its antecedent.[44] For take without loss of generality single-conclusion arguments $\mathfrak{a}_0, \mathfrak{a}_1, \mathfrak{a}_2 \ldots \mathfrak{a}_{i-1}$ of forms $\mathbb{F}_0$, $\mathbb{F}_1$, $\mathbb{F}_2 \ldots \mathbb{F}_{i-1}$ and a single-conclusion argument $\mathfrak{a}_i$ of form $\mathbb{F}_i$, and suppose that $\mathfrak{a}_0, \mathfrak{a}_1, \mathfrak{a}_2 \ldots \mathfrak{a}_{i-1}$ and $\mathfrak{a}_i$ satisfy the first conjunct and the second conjunct of (T)'s antecedent respectively. Consider a non-transitivist according to whom the resulting instance of (T) fails, and who accepts on non-logical grounds all the premises Γ of $\mathfrak{a}_0, \mathfrak{a}_1, \mathfrak{a}_2 \ldots \mathfrak{a}_{i-1}$ and accepts (on non-logical or logical grounds) the premises Λ of $\mathfrak{a}_i$ which are not conclusions of any of $\mathfrak{a}_0, \mathfrak{a}_1, \mathfrak{a}_2 \ldots \mathfrak{a}_{i-1}$. Such non-transitivist would then have to accept the premises Λ, Θ of $\mathfrak{a}_i$ (this seems extremely plausible; I will justify it more fully in section 4.3). Yet, given the foregoing explanation of what the non-transitivist regards as a correct pattern of reasoning, she will not regard herself as committed to accepting the conclusions Ξ of $\mathfrak{a}_i$. How could she then still maintain that $\mathfrak{a}_i$ is *valid*, given that *she accepts its premises but refuses to infer its conclusions*? Is her refusal to infer the conclusion not an implicit admission that she does not regard $\mathfrak{a}_i$ as valid (and thus that she does not regard $\mathbb{F}_i$, which is instantiated by $\mathfrak{a}_i$, as a universally valid argument form)? Notice how, in a certain respect, these questions are particularly pressing for a *non-dualist* non-transitivist, since she cannot help herself to the dualist rejection that, if every premise of $\mathfrak{a}_i$ is true, so is the conclusion of $\mathfrak{a}_i$, rejection which would certainly go some way towards explaining the refusal to infer the conclusion of $\mathfrak{a}_i$.

Before addressing these urgent questions, notice that logical nihilism, as against non-logical/logical dualism, is actually not a possible option for a non-transitivist, at least in the following sense. It might well be that all of the logical nihilist, the non-logical/logical dualist and the non-dualist non-transitivist accept the premises Γ of a classically valid argument $\mathfrak{a}$ of form $\mathbb{F}$ while refusing to infer its conclusions Δ. The non-logical/logical dualist might do this because, even though she recognizes $\mathfrak{a}$ as valid, she regards some coordinate of Γ as describing a logical fact, and so she rejects that, if every coordinate of Γ is true, so is some coordinate of Δ (and on these grounds refuses to infer Δ from Γ). The non-dualist non-transitivist might do this because, even though she recognizes $\mathfrak{a}$ as valid and accepts that, if every coordinate of Γ is true, so is some coordinate of Δ, she still refuses to infer Δ from Γ (on grounds which we will explore shortly). The logical nihilist, however, refuses to infer Δ on the very simple grounds that she does not regard $\mathbb{F}$ as a universally valid argument form and that, in particular, she regards its instance $\mathfrak{a}$

[44] Of course, even in a non-transitive logic not *every* argument form is usually such, and so the label 'nihilist' might in the usual cases be a bit of an exaggeration; still, I will use it indiscriminately as it is even literally correct in some other cases, and I think conveys the right tones even in those cases in which it is literally speaking an exaggeration.

as invalid. She thus cannot regard $\mathfrak{a}$ as involved in a possible failure of the transitivity of consequence, as the dualist and the non-dualist non-transitivist do. Logical nihilism is an *alternative* to non-transitivism, not one of its species.

To come back to the question as to how a (non-dualist) non-transitivist can recognize as valid an argument whose premises she accepts and whose conclusions she does not accept, we must observe that the connection, presupposed by this question, between recognizing an argument as valid and inferring the conclusions if one also accepts its premises is much less straightforward than it might seem at first glance. One can recognize an argument as valid and accept its premises while still not inferring its conclusions because one is somehow prevented by external circumstances from doing so (by a threat, a psychological breakdown, a sudden death etc.). Or because one fails to recognize, maybe on account of their syntactic complexity, that the premises (conclusions) are indeed premises (conclusions) of an argument one recognizes as valid. Or because one has a general policy of not inferring conclusions, maybe for the reason that one has been told by one's guru that every inference is sacrilegious. Or because one simply cannot be required to infer all the conclusions of all the arguments one recognizes as valid and whose premises one accepts—this is arguably not a requirement on resource-bounded rationality, and, unless each and every single truth is an aim of belief, it is not clear why it should even be a requirement on resource-unbounded rationality.

The previous counterexamples may seem to trade on "*deviant*" cases. Still, unless a plausible independent characterization of "deviancy" is provided (a highly non-trivial task), the objection against the non-transitivist would seem to lose much of its force—why should non-transitivism be itself classified as one of the "deviant" cases (which we know from the foregoing counterexamples generally to exist)? Be that as it may, stronger, if more controversial counterexamples can be given where there is actually *epistemic force against the inference's being drawn*:

MO Mo may be told by a source she is justified to trust that, if Mo's initial is 'M', then Mo is a horribly bad *modus-ponens* inferrer (which does not imply that Mo is horribly bad at recognizing the validity of *modus-ponens* arguments). Mo may also know that her initial is 'M'. Mo recognizes the validity of the relevant instance of the rule of *modus ponens*, yet, given that it would be unwarranted to believe that one is a horribly bad *modus-ponens* inferrer exactly *via* a *modus-ponens* inference, there is epistemic force against Mo's inferring the conclusion that she is a horribly bad *modus-ponens* inferrer.[45]

DAVE Sincere Dave might believe *of* each of the 1,000,000 substantial and independent statements of his new history book that *that* statement is true—if Dave didn't really believe a statement to be true, why would he have put it in the book in the first place? Together with the certainly true assumption that those are all the statements in his book, it follows that all the statements in Dave's book are true.

[45]Thanks to Daniele Sgaravatti, Martin Smith and Crispin Wright for discussion of this and similar examples.

Dave recognizes the argument as valid, yet, given Dave's fallibility as a historian, there is epistemic force against modest Dave's inferring the conclusion that all the statements in his book are true.[46,47]

As the reader will have spotted, these two candidate counterexamples can also be turned into candidate counterexamples (a new one and an old one respectively) to many *closure* principles for knowledge and justification. And, in fact, unsurprisingly all the main candidate counterexamples to closure principles I know of can conversely be turned into candidate cases in which one recognizes an argument as valid and accepts its premises while still not inferring its conclusions and in which there is actually epistemic force against the inference's being drawn.[48]

[46]This counterexample is of course a version of the preface paradox (see Makinson [1965]; see Christensen [2004] for a recent congenial discussion). The counterexample is even more telling given the relationships between probabilistic structure and some kinds of non-transitive deductive consequence relations remarked upon in fn 20. As Crispin Wright has emphasized to me in conversation, there is an important asymmetry between the counterexample, which crucially relies on the fact that there is epistemic force against accepting the conjunction of all the premises, and some application of non-transitivity, where one would wish not just to accept all the premises, but also their conjunction (more or less equivalently, one would wish not only to accept of every premise that it is true, but also to accept that every premise is true). The point of the counterexample is however only to show that there can be epistemic force against an inference's being drawn even if the argument is recognized as valid and all its premises are accepted. The *source* of this epistemic force in **DAVE** is such that it only applies to multiple-premise arguments, whereas the source of the force in some application of non-transitivity will evidently not carry this restriction, as it also happens in **MO** and in the other cases mentioned in fn 48 (Harman [1986], pp. 11–24 is the *locus classicus* for the problematization of the connection between validity and inference).

[47]To connect with the issue mentioned in fn 19, the argument from β (a sentence expressing Dave's overall body of evidence relevant to the topics treated in his book) to γ_1 (where γ_i expresses the proposition that the ith statement in Dave's book is true) is defeasibly valid, as is the argument from β, γ_1 to 'γ_1 and γ_2'. Thus, by (CT), the argument from β to 'γ_1 and γ_2' would also be defeasibly valid. Yet, the argument from β, 'γ_1 and γ_2' to 'γ_1, γ_2 and γ_3' is also defeasibly valid. Thus, by (CT), the argument from β to 'γ_1, γ_2 and γ_3' would also be defeasibly valid. With another 999,997 structurally identical reasonings, we reach the conclusion that the argument from β to 'γ_1, γ_2, $\gamma_3 \ldots$ and $\gamma_{1,000,000}$' is also defeasibly valid. Since the last argument is not defeasibly valid, we can assume that, in at least one of those 999,999 reasonings, (CT) (and not just (T)) fails for the defeasible consequence relation in question. Notice that an interesting feature of cases like this in which it is also (CT) (and not just (T)) that fails is that, contrary to the cases like that presented in section 2.3 in which it is only (T) that fails, there is no interesting sense in which, in any of the 999,999 reasonings in question, the state of information represented by the initial premise β *defeats* (or even makes less reasonable) the inference from the initial conclusion 'γ_1, γ_2, $\gamma_3 \ldots$ and γ_i' to the final conclusion 'γ_1, γ_2, $\gamma_3 \ldots$, γ_i and γ_{i+1}' (and so, in this respect, it is a bit misleading to call this kind of non-deductive consequence relations 'defeasible'). The problem seems rather that, in at least one of the 999,999 reasonings in question, *the way in which* Dave has arrived at the initial conclusion 'γ_1, γ_2, $\gamma_3 \ldots$ and γ_i' makes it unreasonable for him to draw the further inference to the final conclusion 'γ_1, γ_2, $\gamma_3 \ldots$, γ_i and γ_{i+1}' (contrast with the case in which Dave, who still has β, has arrived at 'γ_1, γ_2, $\gamma_3 \ldots$ and γ_i' not *via* successive defeasible inferences, but simply because a trustworthy source has revealed to him that that sentence is true: then, the further inference from β,'γ_1, γ_2, $\gamma_3 \ldots$ and γ_i' to 'γ_1, γ_2, $\gamma_3 \ldots$, γ_i and γ_{i+1}' would seem eminently reasonable).

[48]Thus, at some point in the first performance of a 1,000,000-step proof of Fermat's Last Theorem (from φ_0 to $\varphi_{999,999}$), for a few i René accepts φ_i and recognizes as valid the argument from φ_i to φ_{i+1}, yet a case can be made that, given René's fallibility as an inferrer, for some such i there is epistemic force against René's drawing the relevant inference (see Hume [1739], Book I, Part IV, Section I); judging

These counterexamples are certainly sufficient to open up conceptual space for a genuinely non-transitivist (rather than logical nihilist) position. Still, more needs to be said by way of a positive explanation of the refusal to draw an inference when an argument is recognized as valid and all its premises are accepted (at least, more needs to be said if one wants to avoid giving the explanation that the non-logical/logical dualist gives). Moreover, a doubt now arises as to the very point of non-transitivism: if a story needs to be told anyway in order to vindicate the rationality of accepting the premises of an argument recognized as valid while refusing to infer its conclusions, could such a story not be applied in a transitive framework, recognizing, say, that, since, [for every $\varphi \in \mathrm{ran}(\Theta)$, Γ entails Δ, φ and Λ, Θ entails Ξ], Λ, Γ does entail Δ, Ξ, and accepting Λ, Γ, while refusing to infer Δ, Ξ? Such a doubt would certainly be pressing at least for the distinctive punch that intermediate and strong applications of non-transitivity are supposed to have.

The situation which seems to be emerging is this. The non-transitivist needs to differentiate between the *acceptance-related normative property* of being N_0 (which triggers the normative force of consequence once all the premises of an argument recognized as valid are N_0) and the acceptance-related normative property of being N_1 (which, while not sufficient for triggering such force, is nevertheless sufficient for generating a commitment to accepting a sentence that is N_1). Such a distinction between acceptance-related normative properties would allow the non-transitivist to insist that, in the cases in which she accepts the premises of an argument recognized as valid while refusing to infer its conclusions, this is so because the premises are only N_1 and not N_0. And, together with the additional claim that, if the premises of an argument recognized as valid are N_0, its conclusions are N_1, such a distinction between acceptance-related normative properties would also allow the non-transitivist to reply to the awkward question of the last paragraph by saying that, since Λ, Γ is N_0, if consequence were transitive, Δ, Ξ would have to be N_1, and so she would after all be committed to accepting it. To the identification of the acceptance-related normative properties of being N_0 and of being N_1 we must now turn.

whether, in a series of 1,000,000 pairwise indiscriminable trees going from 100 ft tall to 1 ft tall she's looking at from the distance, the ith tree is at least 50 ft tall, for a few i and for some j Sherry accepts 'I know that I know that I know (j times) that the ith tree is at least 50 ft tall' and, given her limited powers of discrimination, recognizes as valid the argument from 'I know that I know that I know (j times) that the ith tree is at least 50 ft tall' to 'I don't know that I don't know that I know that I know that I know ($j - 2$ times) that the $i + 1$th tree is at least 50 ft tall', yet a case can be made that, given what Sherry knows about the first tree and the last tree in the series, for some such i there is epistemic force against Sherry's drawing the relevant inference (see Zardini [2013k]); looking at the zebras in the zoo, Fred accepts 'Those animals are zebras' and recognizes as valid the argument from 'Those animals are zebras' to 'Those animals are not cleverly disguised mules', yet, given Fred's ignorance of zoos' policies, a case can be made that there is epistemic force against Fred's drawing the relevant inference (see Dretske [1970]).

4.3 The Normativity of Consequence

Once again, let us go back without loss of generality to the situation where the non-transitivist accepts φ for non-inferential reasons,[49] recognizes both that ψ follows from φ and that χ follows from ψ but does not accept that χ follows from φ. I say that, in such a situation, her non-transitivism is sufficient to save her from a commitment to accepting χ. To see this, consider the highly plausible general normative principle to the effect that one has inferential reasons[50] to accept all *but no more* than the consequences of what one has *non-inferential* reasons to accept (insofar as one does have such non-inferential reasons).[51] Now, our non-transitivist only has non-inferential reasons to accept φ. Since ψ is indeed a consequence of φ, and since she has non-inferential reasons to accept φ, by the general normative principle just introduced she has indeed inferential reasons to accept ψ, and so she is indeed committed to accepting ψ.[52] However, on her view, χ is not a consequence of φ, and so she is not committed to accepting χ simply because she has non-inferential reasons to accept φ—indeed, since in a non-transitive framework χ need not follow from anything she has non-inferential reasons to accept, by the general normative principle just introduced she has no reasons at all to accept χ (even though she is committed to accepting ψ and recognizes that χ follows from ψ!).

Drawing on the foregoing distinction between *having non-inferential reasons to accept* and *having inferential reasons to accept*, I propose that, roughly, we identify the acceptance-related normative properties of x-is-N_0 and x-is-N_1 with one-has-non-inferential-reasons-to-accept-x and one-has-inferential-reasons-to-accept-x respectively.[53]

[49]Henceforth, by 'having *non-inferential reasons* to accept φ' and its like I really just mean 'having a subjective basis for accepting or in fact accepting φ not because it is the conclusion of a valid argument all of whose premises one has reason to accept or accepts'—acceptance for non-inferential "reasons" need not be objectively well-grounded at all. My use of 'reason' and its like can thus be understood as roughly synonymous with '*subjective* reason' and its like.

[50]Henceforth, by 'having *inferential reasons* to accept φ' and its like I really just mean 'having a subjective basis for accepting φ because it is the conclusion of a valid argument all of whose premises one has reasons to accept'. Notice that, contrary to my own use up to section 4.2, this notion of an inferential reason is in some important respect *external*, as the possession of an inferential reason in this sense does not depend on one's recognition of the validity of the argument. Having said that, I stress that I am switching to the more external notion only to argue that the non-dualist non-transitivist can escape unwanted commitments even in such external sense: points analogous to the ones to be made will apply for a more internal notion that requires one's recognition of the validity of the argument.

[51]The 'insofar'-qualification will henceforth be implicitly understood; it is supposed to take care of the point forcefully made by Harman [1986], pp. 11–12 to the effect that the fact that the conclusion is unacceptable may *defeat* the reasons one had for accepting the premises.

[52]This last step should make clear that, to simplify the discussion, I am assuming that the relevant (non-inferential or inferential) reasons to accept a certain sentence are always so strong as to imply a commitment to accepting that sentence. Conversely, given my understanding of 'reason' (see fns 49 and 50), commitment to accepting a certain sentence is a reason (non-inferential or inferential, as the case may be) to accept that sentence. In view of this equivalence, I will put points concerning one's having reasons of some kind or other (i.e. non-inferential or inferential) to accept a sentence in terms of one's being committed to accepting that sentence.

[53]'Roughly' because **DAVE** (and possibly other counterexamples too) shows that, in some cases, even if one has non-inferential reasons to accept the premises these are not N_0 (i.e. they do not trigger the

I take it that, as it has been introduced at the end of section 4.2, the distinction between x-is-N_0 and x-is-N_1 has some very intuitive appeal and import in our *ordinary evaluation of reasons*: we would ordinarily distinguish between one's reasons to accept a certain sentence being so strong as to permit (or even mandate) acceptance of whatever follows from that sentence and one's reasons to accept a certain sentence being simply strong enough as to permit (or even mandate) acceptance of that sentence. In the former case, one's reasons allow (or even mandate) one to take the sentence as a *initial point for further reasoning*, whereas in the latter case they only allow (or even mandate) one to take the sentence as a *terminal point of acceptance* (see Smith [2004], pp. 196–199 for a defence of this distinction within a transitive framework). In the lights of the remarks already made in fn 20 and in connection with **DAVE**, it should go without saying that the distinction also makes perfectly good *probabilistic* and, more generally, *defeasible* sense. Moreover, I also take it that the proposed identification of that distinction with the distinction between one-has-non-inferential-reasons-to-accept-x and one-has-inferential-reasons-to-accept-x has some very intuitive appeal and import in our ordinary evaluation of reasons as well: we would ordinarily think that *the strength of our reasons decreases with further inferences* (and, contrary to standard assumptions in formal models of uncertainty, that this is so even if the arguments relied on in such inferences are deductively valid, single-premise arguments, as evidenced most clearly by the first two cases mentioned in fn 48).[54]

In such a distinction between acceptance-related normative properties, we see how a trace of non-logical/logical dualism does necessarily remain in the non-dualist non-transitivist's position: only, the distinction is not between two different kinds of *facts*, but between two different kinds of *reasons for accepting a sentence*—either non-inferential or inferential. In view of this distinction, the non-dualist non-transitivist can be seen not as endorsing the rather exotic restriction to non-logical facts of the application of the laws of logic, but as adhering unswervingly both to the verdicts of validity and invalidity issued by her non-transitive logic and to the general normative principle introduced two paragraphs back that the commitments generated by the laws of logic on a certain *position* (a sequence of sentences accepted for non-inferential reasons) coincide with the logical consequences of that position. That is, if one has non-inferential reasons to accept a certain sequence of sentences, one is indeed committed by logic and those very same reasons to accepting each and every consequence of them, but also committed only to that (at least by logic and those very same reasons), so that, if logic is non-transitive, one is not committed by logic and those very same reasons to accepting a consequence of a consequence of one's position which is not already a consequence of one's position, even if it is a consequence of something one is committed to accepting and even accepts (such, as it were, "consequences at one remove" of course do not exist if logic is transitive, but they do if it isn't).

normative force of consequence, at least with respect to certain conclusions). In what follows, I will ignore this further complexity.

[54]It is perhaps worth noting that neither acceptance-related normative property implies actual acceptance, whilst being actually accepted implies both being N_0 (see fn 49) and being N_1 (from fn 49, reflexivity of consequence and (NI) (see below)—indeed, by reflexivity of consequence and (NI), being N_0 itself implies being N_1). This is so because both acceptance-related normative properties pertain to what (subjectively) *ought to be* the case rather than to what *is* the case.

It is crucial to see that this general principle governing consequence can be accepted *by non-transitivists and transitivists alike as exhausting its normative force* (at least as far as the aspects we are concerned with here go). For our purposes, we can focus on its positive component, which can be fleshed out a bit more generally and precisely as the following principle of *connection between having non-inferential reasons to accept and having inferential reasons to accept*:

(NI) If one has non-inferential reasons to accept a sequence of sentences Γ, and this entails a sequence of sentences Δ, one has inferential reasons to accept Δ.

Notice that, *pending any further specification of the properties of the consequence relation*, (NI) does *not* imply that consequence in turn maps inferential reasons to accept a sequence of sentences onto inferential reasons to accept a sequence of sentences. In other words, pending any further specification of the properties of the consequence relation, while (NI) does imply that the normative force of consequence applies to premises one has non-inferential reasons to accept, producing inferential reasons to accept (and so commitments to accepting) the conclusions, it does not imply that such force applies to premises one has simply inferential reasons to accept. Indeed, the non-transitivist can be seen as exploiting exactly the fact that, by itself, (NI) does not imply the stronger principle of *preservation of inferential reasons to accept*:

(II) If one has inferential reasons to accept a sequence of sentences Γ, and this entails a sequence of sentences Δ, one has inferential reasons to accept Δ.

In particular, the non-transitivist can be seen as accepting (NI) while rejecting (II): on her view, one need not be committed to accepting consequences of commitments generated by logic.

It seems to me that, in her joint acceptance of (NI) and rejection of (II), the non-transitivist is occupying a reasonable position, given that the following theorem holds for every reflexive consequence relation (under certain very plausible additional assumptions which I'll make explicit at the relevant stages in the proof):

Theorem 1. (NI) *implies* (II) *iff the consequence relation is transitive.*

Proof.

- *Left-to-right.* We prove the contrapositive. Take a reflexive non-transitive consequence relation $\mathbf{L}$ such that, for every $\varphi \in \text{ran}(\Theta)$, $\Gamma \vdash_{\mathbf{L}} \Delta, \varphi$ and $\Lambda, \Theta \vdash_{\mathbf{L}} \Xi$, but $\Lambda, \Gamma \nvdash_{\mathbf{L}} \Delta, \Xi$, and consider an intermediate application of $\mathbf{L}$ by a subject s having non-inferential reasons only to accept Γ and Λ.[55] s's intermediate application

[55] Throughout this proof, in order to avoid excessive verbal clutter, I will sometimes let context disambiguate whether a sequence is accepted "conjunctively" (in the fashion of premises) or "disjunctively" (in the fashion of conclusions).

of **L** is such that s does not accept Δ, Ξ, even though she does accept, for every $\varphi \in \mathrm{ran}(\Theta)$, Δ, φ. In such a situation, (NI) only requires from s that she accept, for every $\varphi \in \mathrm{ran}(\Theta)$, Δ, φ (since, for every $\varphi \in \mathrm{ran}(\Theta)$, $\Gamma \vdash_{\mathbf{L}} \Delta, \varphi$)—it does not require from s that she accept Δ, Ξ (since $\Lambda, \Gamma \nvdash_{\mathbf{L}} \Delta, \Xi$). Therefore, s satisfies (NI). However, in such a situation, given plausible additional principles (II) does require from s that she accept Δ, Ξ. For, having non-inferential reasons to accept Γ, by (II) s has inferential reasons to accept (and so is committed to accepting), for every $\varphi \in \mathrm{ran}(\Theta)$, Δ, φ (since, for every $\varphi \in \mathrm{ran}(\Theta)$, $\Gamma \vdash_{\mathbf{L}} \Delta, \varphi$). By the additional principle of *semicolon-agglomeration of commitments to accepting "disjunctively"*:

(SACAD) If, for some function seq from ordinals to sequences, for every $\alpha \in \mathrm{dom}(\Pi)$ (with $\Pi = \psi_0, \psi_1, \psi_2 \ldots$),[56] one is committed to accepting "disjunctively" $\mathrm{seq}(\alpha), \psi_\alpha$, then one is committed to accepting "disjunctively" $\mathrm{seq}(0), \mathrm{seq}(1), \mathrm{seq}(2) \ldots, \Pi^{;}$

(where ';', unlike ',', denotes a right-conjunctive structural punctuation mark[57] and $\Pi^{\dagger}$ the result of substituting $\dagger$ throughout as the structural punctuation mark of Π), s is committed to accepting "disjunctively" $\Delta, \Delta, \Delta \ldots, \Theta^{;}$. Hence, by contraction, s is committed to accepting "disjunctively" $\Delta, \Theta^{;}$.[58] Since s is also committed to accepting "conjunctively" Λ (having non-inferential reasons to accept "conjunctively" Λ), by the additional principle of *quasi-monotonicity of commitment to accepting "disjunctively" over implication of commitment to accepting*:

(QCADICA) If commitment to accepting ("conjunctively") Π_0, Π_1 implies commitment to accepting ("disjunctively") Σ_0, then commitment to accepting "conjunctively" Π_0 implies that commitment to accepting "disjunctively" $\Sigma_1, \Pi_1^{;}$ implies commitment to accepting "disjunctively" Σ_1, Σ_0,

we have that s is committed to accepting "disjunctively" Δ, Ξ (since the fact that $\Lambda, \Theta \vdash_{\mathbf{L}} \Xi$ together with the reflexivity of **L**, (NI) and (II) allows us to detach the main consequent of the relevant instance of (QCADICA)). Therefore, (II) has a consequence (i.e. that s is committed to accepting "disjunctively" Δ, Ξ) that (NI) does not have (and, if s were not committed to accepting "disjunctively" Δ, Ξ, (II) would not hold while (NI) might still hold).

- *Right-to-left.* By cases. Take a reflexive transitive consequence relation **L** and suppose that, for every $\varphi \in \mathrm{ran}(\Theta_0)$, a subject s has inferential reasons to accept φ. This can be so:

[56] Throughout, $\mathrm{dom}(\Gamma)$ is the domain of Γ.

[57] For our purposes, a structural punctuation mark $\dagger$ is *right-conjunctive* iff $[\Gamma \vdash \Delta \dagger \varphi$ iff $[\Gamma \vdash \Delta$ and $\Gamma \vdash \varphi]]$. Needless to say, the usefulness of a right-conjunctive structural punctuation mark derives from its ability to allow us to mimic conjunctive operations over sentences of a language which may well lack a conjunctive operator.

[58] The use of contraction is only needed in the proof because, as observed in fn 10, (T) is not a pure transitivity principle and does in fact imply contraction. Contraction would not be needed in the proof if we worked with one of the more convoluted transitivity principles mentioned in fn 10.

(i) Either because $\Gamma \vdash_{\mathbf{L}} \varphi$, where $\Gamma \neq \langle . \rangle$ and s has non-inferential reasons to accept Γ (so that, by (NI), s has inferential reasons to accept φ);

(ii) Or because $\langle . \rangle \vdash_{\mathbf{L}} \varphi$ (so that, by (NI), s has inferential reasons to accept φ).

Notice that, by reflexivity of **L** and (NI), the case in which s has non-inferential reasons to accept φ is reduced to case (i). Notice also that, by the additional principle of *well-foundedness of inferential reasons over logically contingent sentences*:

(WIRLCS) The relation x-is-a-coordinate-of-a-sequence-which-gives-s-inferential-reasons-for-y-and-x-is-not-y is well-founded on the field of logically contingent sentences,

if s has inferential reasons to accept φ because $\Gamma \vdash_{\mathbf{L}} \varphi$, where $\Gamma \neq \langle . \rangle$, it is not possible that [Γ is such that, for some of its logically contingent coordinates φ_0, s has only inferential reasons to accept φ_0, namely that φ_0 follows from premises Γ_0, and Γ_0 is in turn such that, for some of its logically contingent coordinates φ_1, s has only inferential reasons to accept φ_1, namely that φ_1 follows from premises Γ_1, and Γ_1 is in turn such that, for some of its logically contingent coordinates φ_2, s has only inferential reasons to accept φ_2, namely that φ_2 follows from premises $\Gamma_2 \ldots$]. This yields that, if $\langle . \rangle \nvdash_{\mathbf{L}} \varphi$, each inferential reason s has to accept φ must ultimately be traceable back to a combination Γ_* of (possibly infinitely many, possibly infinitely long) sequences each of which can be reached in a finite number of steps and every coordinate of which s has non-inferential reasons to accept. Hence, by (monotonicity and) (T), we have that $\Gamma_* \vdash_{\mathbf{L}} \varphi$.[59] Letting $\Gamma = \Gamma_*$, this finite-chain case is also reduced to case (i). Now, suppose also that $\Theta_0 \vdash_{\mathbf{L}} \Xi$. Let Θ_1 be the sequence obtained from Θ_0 by replacing each coordinate falling under case (i) with its associated Γ and by deleting each coordinate falling under case (ii). Then, s has non-inferential reasons to accept all the coordinates of Θ_1 and, by (monotonicity and) (T), $\Theta_1 \vdash_{\mathbf{L}} \Xi$ just as well. Hence, by (NI), s has inferential reasons to accept Ξ.

□

Indeed, given that (II) straightforwardly implies (NI) no matter whether the consequence relation is transitive or not as long as it is reflexive (since, as we have seen in fn 54, non-inferential reasons to accept will then imply inferential reasons to accept),[60] theorem 1 can be reformulated to the effect that [(NI) is *equivalent* with (II) iff the consequence relation is transitive]. In view of this, it seems to me that the non-transitivist can reasonably insist that the *pure principle*, free of any relevant assumption concerning formal properties of consequence, which encodes its normativity (principle which is endorsed by

[59] As noted in fn 10, (T) in fact implies monotonicity. Monotonicity would not even *prima facie* be needed in the proof if we worked with one of the more convoluted transitivity principles mentioned in fn 10.

[60] The assumption of reflexivity will henceforth be implicitly understood.

non-transitivists and transitivists alike) is (NI) rather than (II), (II) being associated with such normativity only because equivalent with (NI) under the (rejected) assumption of transitivity. Hence, even though a non-transitive logic is naturally hospitable to a certain "softening" of the normative force of consequence, this does not mean that the core of that force is not preserved (let alone that no important requirement is placed by non-transitive consequence on rational beings).

Before proceeding further, an absolutely essential feature of this dialectic must be made clear. (II) is in effect a principle of *closure of having inferential reasons to accept* under logical consequence. Why then cannot one apply in this case a strategy analogous to the one we used in order to uphold in a non-transitive framework the necessity of truth preservation for consequence, which is in effect a principle of *closure of truth* under logical consequence? In the case of truth and of the other properties we considered in section 3.2, it was observed that there seem to be bridge principles linking the languages talking about these properties and the original language $\mathscr{L}_0$ claimed to be non-classical, principles which force the logic of the former languages to be itself non-classical. No such principle seems to govern properties like one-has-inferential-reasons-to-accept-x and, more generally, there does not seem to be any reason why the language $\mathscr{L}_1$ talking about this property should exhibit the same problematic features which motivated a deviance from classical logic for $\mathscr{L}_0$. To give one example, for an intuitionist, the language talking about which sentences of a standard quantified arithmetical language a mathematician has inferential reasons to accept may well lack the *unsurveyability* characteristic of the standard quantified arithmetical language, and so may well be classical (it may well be a surveyable matter which sentences the mathematician has a subjective basis for non-inferentially accepting, and even an intuitionist would typically take it to be a classical matter which sentences follow from these in intuitionist logic). Or, to give another example more germane to our concerns, for a tolerant logician, the language talking about which sentences of an ordinary language a competent speaker has inferential reasons to accept may well lack the *vagueness* characteristic of the ordinary language, and so may well be classical (it may well be a precise matter which sentences the speaker has a subjective basis for non-inferentially accepting, and even a tolerant logician would typically take it to be a classical matter which sentences follow from these in a tolerant logic). In all such cases, it is the respects in which the application of 'Subject s has inferential reasons to accept 'φ' ' (as against the application of φ, ' 'φ' is knowable', ' 'φ' is justifiedly believable' etc.) depends on *brutely factual subjective features* (see fns 49 and 50) that bring in a crucial *mismatch* between the truth conditions of the former sentence and those of the latter sentences, so that the problematic features which motivated a deviance from classical logic for the latter sentences need not be present in the former sentence.

Moreover, even if, contrary to what I have just argued, the logic of $\mathscr{L}_1$ were non-transitive, (II) would still be unacceptable for most non-transitivists. Suppose that a subject s has non-inferential reasons to accept φ, and that ψ follows from φ and χ from ψ. Suppose also that the argument from φ to χ is non-transitivistically invalid, so that, *qua* non-transitivists, we would wish to avoid imputing to s any commitment to accepting χ, despite her having (non-inferential) reasons to accept φ. Now, by (NI) and ψ's following from φ, we can infer that s has inferential reasons to accept (and so is committed to accepting) ψ. But reflect that the satisfaction of s's commitment to accepting ψ requires s to *treat* ψ *in a certain way* (at the least, very roughly, to assent to ψ on the basis of an inference if queried under normal circumstances). Crucially, such a way is also sufficient to establish s's having inferential reasons to accept ψ independently of s's having non-inferential reasons to accept φ together with ψ's following from φ and (NI). This means that there will typically be not only inferential reasons to accept that s has inferential reasons to accept ψ (namely, those provided by s's having non-inferential reasons to accept φ together with ψ's following from φ and (NI)), but also non-inferential reasons to accept that (namely, those provided by the observation of s's behaviour). More precisely, there will be such reasons whenever s in fact satisfies her commitment to accepting ψ. Having non-inferential reasons both to accept that s has inferential reasons to accept ψ and to accept (II), we would thus be committed to accepting that s has inferential reasons to accept (and so is committed to accepting) χ! (Of course, the argument propagates forward to any conclusion connected with χ through a chain of valid arguments.)

A very general lesson can be extracted from the main turn of the previous argument. Say that a proposition is *non-inferentially inaccessible* for a subject s at a time t iff its truth does not imply that s at t has any non-inferential reason to believe it. What the previous argument shows is that an at least intermediate or strong application of non-transitivity with respect to two arguments $\mathfrak{a}_0$ and $\mathfrak{a}_1$ made by a subject s at a time t requires the propositions expressed by the conclusions of $\mathfrak{a}_0$ to be non-inferentially inaccessible for s at t. This squares nicely with the joint acceptance of (NI) and rejection of (II) typical of at least intermediate and strong applications of non-transitivity: for what these imply is that whether or not the normative force of consequence applies to a subject's acceptance of a sentence should depend on the sentence's *pedigree* (non-inferential or inferential) in the subject's cognitive history, whereas lack of non-inferential inaccessibility precisely obliterates any distinction which might be drawn at that level.

Thus, sentences accepted *qua* conclusions of valid arguments all of whose premises are accepted may just not have the right pedigree to enter in turn as premises into further valid arguments possessing normative force. It is in this *relevance of the premises' pedigree to the normative force of a valid argument* that deductive non-transitive reasoning comes close to defeasible reasoning. For example, consider the defeasible validity of the argument from 'Al is a native speaker of Italian' to 'Al was born in Italy' (see section 2.3) and suppose that one accepts its premise. Is one committed to accepting its conclusion? *That will depend on the reason why one accepts its premise*: if one accepts the premise because one has heard Al's fluent Italian speech (and one has no evidence against the conclusion), one is committed to accepting the conclusion, but, if one accepts the premise because one has in turn inferred it from 'Al was born in Little Italy', one will not be so committed. Contrast this relevance of the premises' pedigree to the normative force of a valid argument in the case of deductive non-transitive reasoning and defeasible reasoning with the irrelevance of the premises' pedigree to the normative force of a valid argument in the case of deductive transitive reasoning. For example, consider the transitive validity of the argument from 'Al is a native speaker of Italian' to 'Al is a native speaker of Italian or

Al was born in Italy' and suppose that one accepts its premise. Then, one is committed to accepting its conclusion, *no matter for what reason one accepts its premise.* Non-dualist non-transitivism can thus be seen as an interesting hybrid joining important aspects of the *metaphysics* characteristic of deductive consequence relations as opposed to defeasible ones (i.e. the guarantee of truth preservation) with important aspects of the *normativity* characteristic of defeasible consequence relations as opposed to deductive ones (i.e. the relevance of the premises' pedigree to the normative force of a valid argument).

We can go a bit deeper. What underlies the relevant aspects of the normativity characteristic of defeasible consequence relations is the fact that the conclusion of a defeasibly valid argument can (in constrained ways) be *stronger* than its premise. Here, 'x is stronger than y' is of course not meant in the rather definitional sense of y's following from x but not *vice versa* (according to the relevant consequence relation); rather, it is meant in the intuitive sense of x "*being weightier than*" y, with the more precise and specific consequences that x *is (epistemically) less likely than* y and that x *has consequences that* y *does not have* (according to the relevant consequence relation). Now, *exactly the same holds for deductive non-transitive consequence relations*: the only relevant difference with defeasible consequence relations is that, while in the latter consequence relations the conclusion being stronger than the premise goes together with (and, one may wish to add, is grounded in) the possibility of the premise being true while the conclusion is not, this is not so in the former consequence relations. Hence, even for deductive non-transitive consequence relations, it is the case that, if χ is stronger than ψ, it might be that one has reasons to accept ψ but no reasons to accept χ even if χ follows from ψ—whether one does have the latter reasons will depend on the specifics of one's reasons to accept ψ (and on the specifics of the constrained ways in which χ is stronger than ψ).[61] In terms of the insightful taxonomy of Salmon [1967], pp. 5–11, a non-transitivistically deductively valid inference can thus be "*ampliative*" while remaining "*demonstrative*", and be so in a way that goes beyond what Salmon and many others seem to have been able to conceive, since such demonstrative inference can be ampliative not only in the broadly *semantic* sense of the conclusion not being implicit in the premise, but also in the broadly *epistemic* sense of the conclusion being less likely than the premise. Coming back to a theme emerged already in section 3.1, we can then see how non-transitive logics vindicate a radical sense in which *logic* can be *informative.*

[61]Of course, it will be particularly problematic to suppose that one has reasons to accept χ if one's reasons for accepting ψ consist in reasons for accepting φ and in ψ's following from φ (for ψ might in turn be stronger than φ). Notice that, in addition to explaining the relevant aspects of the *normativity* characteristic of deductive non-transitive consequence relations, this conception also accounts for the very *non-transitivity* of such consequence relations. Suppose that ψ follows from φ: if ψ is less likely than φ and χ is less likely than ψ, the difference in likelihood between χ and ψ might be small enough as to be compatible with χ's following from ψ, while the difference in likelihood between χ and φ might be large enough as to be incompatible with χ's following from φ; even more straightforwardly, let χ be one of the consequences that ψ has but φ does not have.

4.4 Non-Transitivity and Asymmetries between Premises and Conclusions

Let us draw together some of the contrasts explored so far. Exploiting the distinction just unearthed, the left-hand and right-hand sides of a sequent can then be interpreted in a non-transitive framework as expressing a *connection* between what we have non-inferential reasons to accept and what we have inferential reasons to accept (which does not of course rule out that, independently, we also have non-inferential reasons to accept it: reasons-based acceptance can be overdetermined). This distinction has substance *for transitivists and non-transitivists alike*: there is a perfectly good sense in which I have non-inferential reasons to accept that snow is white but have simply inferential reasons to accept that either snow is white or grass is blue. It should thus actually be *common ground* that consequence can indeed fail to preserve non-inferential reasons to accept, and that it only guarantees that, if one has non-inferential reasons to accept the premises, one has inferential reasons to accept the conclusions.

Within this common ground, the debate on transitivity can then be understood as follows. The non-transitivist can be seen as thinking that consequence can also fail to preserve inferential reasons to accept. She can thus be seen as focussing on the *common-ground* property of consequence of *guaranteeing* inferential reasons to accept *given* non-inferential reasons to accept, whilst the transitivist can be seen as focussing on the *further* (alleged) property of consequence of *preserving* inferential reasons to accept (which, as shown by theorem 1, follows from the common ground-property iff consequence is transitive). The non-transitivist thinks that consequence *only guarantees a connection* between the *two different* acceptance-related normative properties of being N_0 and of being N_1, whilst the transitivist thinks that it *also preserves* the *single* acceptance-related normative property of being N_1.

It is worth noting that, parallel to the *normative* distinction between having non-inferential reasons to accept and having inferential reasons to accept—which can be used to mark a major point of disagreement between the *non-transitivist* and the *transitivist*—a similar distinction can be drawn at the *metaphysical* level between truth simply in virtue of how the world is and truth in virtue of how the world is and the laws of logic—which can be used to mark a major point of disagreement between the *dualist* and the *non-dualist*. The left-hand and right-hand sides of a sequent can equally legitimately be interpreted in a non-transitive framework as expressing a *connection* between what is true simply in virtue of how the world is and what is true in virtue of how the world is and the laws of logic (which does not of course rule out that, independently, it is also true simply in virtue of how the world is: truth can be overdetermined). This distinction too has substance *for dualists and non-dualists alike*: there is a perfectly good sense in which it is true simply in virtue of how the world is that snow is white, but it is true in virtue of how the world is and the laws of logic that either snow is white or grass is blue. It should thus actually be *common ground* that consequence can indeed fail to preserve truth simply in virtue of how the world is, and that it only guarantees that, if all the premises are true simply in virtue of how the world is, some conclusion is true in virtue of how the world is and the laws of logic.

The dualist can then be seen as thinking that consequence can also fail to preserve truth in virtue of how the world is and the laws of logic. She can thus be seen as focussing on the *common-ground* property of consequence of *guaranteeing* truth in virtue of how the world is and the laws of logic *given* truth simply in virtue of how the world is, whilst the non-dualist (either transitivist or non-transitivist) can be seen as focussing on the *further* (alleged) property of consequence of *preserving* truth in virtue of how the world is and the laws of logic (which, as indicated in fn 43, does follow if the deduction theorem holds). The dualist thinks that consequence *only guarantees a connection* between *two different* metaphysical properties, whilst the non-dualist thinks that it *also preserves* a *single* metaphysical property.

4.5 The Locality of Non-Transitive Consequence

In the peculiar joint acceptance of (NI) and rejection of (II) we see one sense in which the requirements placed by a non-transitive logic might typically be *local*, in this case extending only so far as the consequences of sentences accepted for non-inferential reasons go rather than stretching to cover every consequence of any sentences one is committed to accepting for whichever reasons. Indeed, the idea that non-transitive consequence has distinctively "*localist*" features has been lurking behind much of our discussion in sections 4.1–4.4, and it is now time to make it emerge more clearly.

To start with a suggestive spatial metaphor,[62] the non-transitivist's picture of logical space is the rather unusual one which sees the space of consequences of a given point (that is, of a given sequence of sentences) as being bounded by a *horizon*: there are bounds to what is logically necessary in relation to a given point which are not generally such in relation to themselves and to some other points included in the same original horizon, so that *a movement inside the horizon can result in a movement of the horizon itself*. Needless to say, this picture goes precisely against the more traditional picture of logical space (inspiring the view of consequence as a *closure* operation introduced in fn 37) which sees the space of consequences of a given point as being bounded by a *frame*: there are bounds to what is logically necessary in relation to a given point which stretch so far as to be such also in relation to themselves and to all other points included in the same original frame, so that *no movement inside the frame can result in a movement of the frame itself*.

Throughout sections 4.1–4.4, we have explored at some length two alternative ways (dualist and non-dualist) in which this horizon-like picture of logical space can be made better sense of with what has in fact been the general idea that the effects of consequence are peculiarly local—that *consequence manages to constrain the truth value of sentences only at one remove from where one starts*. As we have seen, this may be because consequence in turn fails to impose a truth-preservation constraint on the sentences whose truth values have been so constrained (if one is a dualist), or it may be instead because such a constraint can still be rationally not accepted to have certain effects (if one is a

[62]Suggested to me in another context by Josh Parsons.

non-dualist). Either way, the upshot is that, although consequence does constrain the truth value of sentences at one remove from where one starts, it does not constrain the truth value of sentences at further removes—on longer distances, its constraining force gives out.

In conformity with the focus of our discussion in sections 4.1–4.4 (see fn 26), this point about the sensitivity of the constraining force of consequence to the distance from where one starts applies straightforwardly only to intermediate and strong applications of non-transitivity. But such sensitivity can in turn be seen to arise as a special case from the more general non-transitivist thought that a set of sentences may exhibit a *structure of merely local connections*, with consequence playing the natural role of a "one-stop inference ticket" that reflects such structure. Thus, to go back to the examples of applications of non-transitivity of section 2, in the case of relevance consequence is supposed to reflect merely locally holding *meaning* connections; in the case of tolerance, consequence is supposed to reflect merely locally holding *indifference* connections; in the case of evidence, consequence is supposed to reflect merely locally holding *probabilification* connections.

Accordingly, on this localist picture, transitivity can be seen to fail because, generally speaking, it would create *spurious connections*. To go back again to the specific examples of applications of non-transitivity of section 2, in the case of relevance, transitivity would create a spurious meaning connection between sentences which have none (even though they are the opposite extremes of a chain of genuinely connected sentences); in the case of tolerance, transitivity would create a spurious indifference connection between sentences which have none (even though they are the opposite extremes of a chain of genuinely connected sentences); in the case of evidence, transitivity would create a spurious probabilification connection between sentences which have none (even though they are the opposite extremes of a chain of genuinely connected sentences). In all these cases, transitivity would obliterate a *non-trivial distance structure* which non-transitivists think is exhibited by a set of sentences, inflating local connections between these into global ones.

4.6 Non-Transitivist Theories, Situations and Worlds

It is in view of what has been said that, in a non-transitive framework and under a certain assumption to be introduced shortly, two different readings of theoretical notions defined using the notion of *closure under logical consequence* must be sharply distinguished as being non-equivalent (in this section, for simplicity's sake, we assume monotonicity and return to taking sets to be the terms of the consequence relation and to restricting our attention to single-conclusion arguments). Notions so defined can be seen as having at their core the notion of a *theory*, straightforwardly defined as being any set of sentences closed under logical consequence. The theoretical interest of a notion which, in virtue of the closure clause, outruns that of a mere set of sentences should be evident in view of the normativity of consequence. One has non-inferential reasons to accept a set of sentences X. By (NI), one has thereby inferential reasons to accept (and so is committed to accepting) not only X, but also the set of all the consequences of X. The *objects of commitment* are thus always closed under logical consequence in the specified sense. And such objects well deserve to be called 'theories', for theories are traditionally thought precisely to be the objects of commitment: theories are what people are traditionally thought to hold, defend, attack, revise, try to confirm etc.

Under the simplifying assumptions made at the beginning of this section, we can identify positions (introduced in section 4.3) with arbitrary sets of sentences; the theory of a position is then the closure under logical consequence of that position (let us denote by '$\mathrm{thr}_{\mathbf{L}}$' the function from positions to theories under the consequence relation $\mathbf{L}$). A theory $\mathcal{T}$ is *prime* iff, whenever 'φ or ψ' belongs to $\mathcal{T}$, either φ or ψ belongs to $\mathcal{T}$: primeness is the property of a theory to provide a witness for each of its assertions. A theory $\mathcal{T}$ is *maximal* iff, for every φ, either φ or 'It is not the case that φ' belongs to $\mathcal{T}$: maximality is the property of a theory to settle every question. A prime theory represents a *situation*, a maximal theory a *world*; φ is *true* in a situation* (*true* in a world*) iff φ belongs to the theory representing that situation (world). Call the logic under which a theory is closed '*target logic*' (given our purposes, this will almost always be a non-transitive logic), the logic of the language talking about theories '*background logic*'.

What do non-transitivist theories, situations and worlds look like? There is no reason to think that the language talking about the logical consequences of sets of sentences of a language $\mathscr{L}$ should exhibit the same problematic features which motivate a deviance from classical logic for $\mathscr{L}$. It behoves us then to consider the case where the background logic is classical. Under this assumption, two readings of the definition 'The theory of a position is the closure under logical consequence of that position' must be sharply distinguished as being non-equivalent:

(i) The theory of a position P is the set of the logical consequences of P.

(ii) The theory of a position P is the smallest set $\mathcal{T}$ such that:

  (a) $\mathcal{T}$ contains all the logical consequences of P;
  (b) If φ is a consequence of $\mathcal{T}$, $\varphi \in \mathcal{T}$.[63]

Under the assumption of transitivity of the background logic, the notion delivered by reading (ii) is clearly too strong at least for intermediate and strong applications of non-transitivity, since it will force the theory of a position—that to which one is committed—to contain sentences which, according to the non-transitive target logic, are not logical consequences of the position.

Again, as in the case of (NI) and (II), it seems to me that, in her use of reading (i) rather than reading (ii), the non-transitivist is occupying a reasonable position, given that the following theorem holds:

[63] That is, the greatest lower bound under $\subseteq$ of the class of sets satisfying (a) and (b); that is, the set $\mathcal{T}_0$ such that $\varphi \in \mathcal{T}_0$ iff, for every $\mathcal{T}_1$ satisfying (a) and (b), $\varphi \in \mathcal{T}_1$.

Theorem 2. *Let* $\mathrm{thr(i)}_{\mathbf{L}}$ *and* $\mathrm{thr(ii)}_{\mathbf{L}}$ *be the functions corresponding respectively to readings* (i) *and* (ii)*. Then, [for every* P*,* $\varphi \in \mathrm{thr(ii)}_{\mathbf{L}}(P)$ *only if* $\varphi \in \mathrm{thr(i)}_{\mathbf{L}}(P)$*] iff* $\mathbf{L}$ *is transitive.*

Proof.

- *Left-to-right.* We prove the contrapositive. Take a non-transitive consequence relation $\mathbf{L}$ such that, for every $\varphi \in Y$, $X \vdash_{\mathbf{L}} \varphi$ and $Z \cup Y \vdash_{\mathbf{L}} \psi$, but $Z \cup X \nvdash_{\mathbf{L}} \psi$, and consider a position $P = Z \cup X$. Then $\psi \notin \mathrm{thr(i)}_{\mathbf{L}}(P)$ (since $Z \cup X \nvdash_{\mathbf{L}} \psi$), whereas $\psi \in \mathrm{thr(ii)}_{\mathbf{L}}(P)$ (since $Z \subseteq \mathrm{thr(ii)}_{\mathbf{L}}(P)$ by reflexivity of $\mathbf{L}$ and $Y \subseteq \mathrm{thr(ii)}_{\mathbf{L}}(P)$, while $Z \cup Y \vdash_{\mathbf{L}} \psi$).

- *Right-to-left.* Take a transitive consequence relation $\mathbf{L}$ and consider an arbitrary position P. Suppose that $\varphi \in \mathrm{thr(ii)}_{\mathbf{L}}(P)$. We aim to prove that $P \vdash_{\mathbf{L}} \varphi$. We do so by first defining by transfinite recursion the following hierarchy of positions:

  (i′) $P_0 = \{\varphi : P \vdash_{\mathbf{L}} \varphi\}$;

  (ii′) $P_{\alpha+1} = P_\alpha \cup \{\varphi : P_\alpha \vdash_{\mathbf{L}} \varphi\}$;

  (iii′) $P_\lambda = \bigcup(\{P_\alpha : \alpha < \lambda\})$.

  Notice that, since the language is finitary, by the well-ordering of the ordinals there will be a (not very big) first ordinal κ at which the process stabilizes and no new sentences are admitted as consequences—that is, for every $\alpha \geq \kappa$, $P_\alpha = P_\kappa$. Consider then the set P_κ. P_κ satisfies (a), since $P_0 \subseteq P_\kappa$. P_κ also satisfies (b), since $P_{\kappa+1} = P_\kappa$. Moreover, P_κ is the smallest set to do so. For suppose that $\mathcal{T}$ satisfies (a) and (b). We prove by transfinite induction that, for every $\alpha \leq \kappa$, $P_\alpha \subseteq \mathcal{T}$:

  (i″) Since, by (a), $\{\varphi : P \vdash_{\mathbf{L}} \varphi\} \subseteq \mathcal{T}$, $P_0 \subseteq \mathcal{T}$;

  (ii″) If $P_\alpha \subseteq \mathcal{T}$, since, by (b), $\{\varphi : P_\alpha \vdash_{\mathbf{L}} \varphi\} \subseteq \mathcal{T}$ as well, $P_{\alpha+1} \subseteq \mathcal{T}$;

  (iii″) If, for every $\alpha < \lambda$, $P_\alpha \subseteq \mathcal{T}$, then $\bigcup(\{P_\alpha : \alpha < \lambda\}) \subseteq \mathcal{T}$ as well.

  Thus, if $\varphi \in P_\kappa$, $\varphi \in \mathcal{T}$, and so P_κ is the smallest set to satisfy (a) and (b). It then follows that $P_\kappa = \mathrm{thr(ii)}_{\mathbf{L}}(P)$. We can now prove by transfinite induction that, given the transitivity of $\mathbf{L}$, if $\varphi \in \mathrm{thr(ii)}_{\mathbf{L}}(P)$ ($\varphi \in P_\kappa$), then $P \vdash_{\mathbf{L}} \varphi$. We do so by proving by transfinite induction that, for every $\alpha \leq \kappa$, if $\varphi \in P_\alpha$, then $P \vdash_{\mathbf{L}} \varphi$:

  (i‴) If $\varphi \in P_0$, then $P \vdash_{\mathbf{L}} \varphi$;

  (ii‴) If $\varphi \in P_{\alpha+1}$, then $\varphi \in P_\alpha \cup \{\varphi : P_\alpha \vdash_{\mathbf{L}} \varphi\}$, and so:

    (a′) Either $\varphi \in P_\alpha$, in which case, by the induction hypothesis, $P \vdash_{\mathbf{L}} \varphi$;

    (b′) Or $\varphi \in \{\varphi : P_\alpha \vdash_{\mathbf{L}} \varphi\}$, in which case, since, by the induction hypothesis, for every $\psi \in P_\alpha$, $P \vdash_{\mathbf{L}} \psi$, by (T) $P \vdash_{\mathbf{L}} \varphi$ as well;

  (iii‴) If $\varphi \in P_\lambda$, then, for some $\alpha < \lambda$, $\varphi \in P_\alpha$, and so, by the induction hypothesis, $P \vdash_{\mathbf{L}} \varphi$.

The proof is completed by observing that, since $P \vdash_{\mathbf{L}} \varphi$, $\varphi \in \mathrm{thr(i)}_{\mathbf{L}}(P)$ as well.

$\square$

Indeed, given that it is straightforward that, [for every P, $\varphi \in \mathrm{thr(i)}_{\mathbf{L}}(P)$ only if $\varphi \in \mathrm{thr(ii)}_{\mathbf{L}}(P)$] no matter whether **L** is transitive or not (because of (a)), theorem 2 can be reformulated to the effect that, [[for every P, $\varphi \in \mathrm{thr(i)}_{\mathbf{L}}(P)$ *iff* $\varphi \in \mathrm{thr(ii)}_{\mathbf{L}}(P)$] iff **L** is transitive]. In view of this, it seems to me that the non-transitivist can reasonably insist that the *pure notion*, free of any relevant assumption concerning formal properties of consequence, which plays the complex role usually assigned to the notion of a theory (role which is the same for non-transitivists and transitivists alike) is the one expressed by reading (i) rather than the one expressed by reading (ii), the latter being associated with such role only because reading (ii) is equivalent with (i) under the (rejected) assumption of transitivity.

Once reading (i) is distinguished as the appropriate notion to use when the target logic is non-transitive, the theory of theories can proceed very much as before (see e.g. Barwise and Perry [1983], pp. 49–116). For our purposes, we only need to note the following concerning situations. A prime theory $\mathcal{N}$ with a non-transitive target logic **L** representing a certain situation may be such that $\psi \in \mathcal{N}$, $\psi \rightarrow \chi \in \mathcal{N}$ but $\chi \notin \mathcal{N}$ (such will be the case for example in tolerant logics if only φ, $\varphi \rightarrow \psi$ and $\psi \rightarrow \chi$ belong to the position $\mathcal{N}$ is a theory of). The situation thus represented would be one where a conditional ($\psi \rightarrow \chi$) and its antecedent (ψ) are true*, but the consequent (χ) is not. At a glance, this may of course look like an unwelcome consequence, since **L** may actually be such that the rule of *modus ponens* is unrestrictedly valid in it (as it is in tolerant logics). Even worse, if the law of excluded middle also holds unrestrictedly in **L** (as it does in the tolerant logics advocated in my works on vagueness referenced in section 1), given $\mathcal{N}$'s primeness we would have that $\neg\chi \in \mathcal{N}$, and so that $\neg\chi$ is true* in the situation represented by $\mathcal{N}$! However, it should by now be clear that, analogously to the cases discussed in section 3.2, these consequences are due to the choice of adopting a transitive background logic in the theory of theories. This choice imposes that the link between the technical notion of truth* in a situation (and of truth* in a world)—which is then governed by a transitive logic—and the informal and philosophical notion of truth in a situation (and of truth in a world)—which, just as the notion of truth, is sensitive to the failures of transitivity in the logic of the object language—be at best very complex and mediated.

Again, the generality of the point can be illustrated with reference to more well-known deviations from classical logic. Consider a classical theory $\mathcal{C}$ of a prime intuitionist theory $\mathcal{I}$ formulated in language $\mathscr{I}$. $\mathcal{C}$ entails that the situation i represented by $\mathcal{I}$ is actually such that, for every $\varphi \in \mathscr{I}$, either φ is true* in i or φ is not true* in i (even though, of course, being the target logic intuitionist it need not be the case that, for every $\varphi \in \mathscr{I}$, either $\varphi \in \mathcal{I}$ or $\neg\varphi \in \mathcal{I}$). Such a conclusion would be repugnant given what would seem to be the most natural theory of truth in a situation (in the informal and philosophical sense) available to an intuitionist (namely one such that φ's being not true in a situation implies that φ is false in that situation and so that $\neg\varphi$ is true in that situation).

If a non-transitive logic is adopted as background logic, however, things change drastically and reading (ii) becomes again a viable option for intermediate and strong applications of non-transitivity. For example, going back to the theory $\mathcal{N}$ discussed two paragraphs back, reading (ii) no longer has the unwelcome consequence of declaring that $\chi \in \mathcal{N}$. For, in a non-transitive background logic, that $\chi \in \mathcal{N}$ does not follow from (set theory and) reading (ii), φ's belonging to $\mathcal{N}$, $\varphi \rightarrow \psi$'s belonging to $\mathcal{N}$ and $\psi \rightarrow \chi$'s belonging to $\mathcal{N}$. Indeed, there is now at least one reason to prefer reading (ii) over reading (i), since, if the background logic $\mathbf{L_0}$ is non-transitive (but not so as to affect definitional reasoning), reading (ii) allows for the construction of a theory $\mathcal{T}$ with a non-transitive target logic $\mathbf{L_1}$ representing a situation s such that, if $X \vdash_{\mathbf{L_1}} \varphi$, then truth* in s is preserved from X to φ (as befits the informal and philosophical notion of truth in a situation). For suppose that $X \vdash_{\mathbf{L_1}} \varphi$ and that every member of X is true* in s. Then, by definition of truth* in s, it follows that $X \subseteq \mathcal{T}$. From this, $X \vdash_{\mathbf{L_1}} \varphi$ and reading (ii), it follows that $\varphi \in \mathcal{T}$, which in turn entails, by definition of truth* in s, that φ is true* in s as well (notice how this argument only requires transitivity for definitional reasoning in $\mathbf{L_0}$).

5 Conclusion

Appealing as they may be, applications of non-transitivity require an extended philosophical discussion of what sense there is to be made of a rational being who reasons using a non-transitive logic. This paper has no doubt been only a first stab at meeting that pressing request in its generality. We have seen how this task involves dealing with some of the hardest problems at the interface between the philosophy of logic and the theory of rationality: the connection between structural rules and the very nature of the premises and conclusions of an argument (and of their acceptance or rejection); the relation between deductive and defeasible consequence relations; the relation between consequence and inference; the relation between consequence, preservation of truth and preservation of other epistemic properties; the normativity of consequence; the relation between premises and conclusion of a valid argument; the nature of the objects of logical commitment; the problem of theorizing about how all this behaves in a certain logic by using a different one etc. From a wider perspective, I hope that the foregoing investigations about how non-transitivity impinges on these and other issues have helped to see them in a new, more general light, and that the conceptualizations made and the distinctions drawn will prove fruitful also in the examination of the philosophical foundations of other logics.

References

Arnon Avron. Simple consequence relations. *Information and Computation*, 92:105–140, 1991.

Jon Barwise and John Perry. *Situations and Attitudes*. MIT Press, Cambridge MA, 1983.

Jean-Yves Béziau. Transitivity and paradoxes. In Jurğis Šķilters, Matti Eklund, Ólafur Jónsson, and Olav Wiegand, editors, *Paradox: Logical, Cognitive and Communicative Aspects*, volume I of *The Baltic International Yearbook of Cognition, Logic and Communication*, pages 87–92. University of Latvia Press, Riga, 2006.

Bernhard Bolzano. *Wissenschaftslehre*. Seidel, Sulzbach, 1837.

David Christensen. *Putting Logic in its Place*. Oxford University Press, Oxford, 2004.

Pablo Cobreros, Paul Égré, David Ripley, and Robert van Rooij. Tolerant, classical, strict. *Journal of Philosophical Logic*, 41:347–385, 2012.

Eduard Čech. *Topological Spaces*. Academia, Prague, revised edition, 1966.

Fred Dretske. Epistemic operators. *The Journal of Philosophy*, 67:1007–1023, 1970.

Richard Epstein. Relatedness and implication. *Philosophical Studies*, 36:137–173, 1979.

Gareth Evans. Semantic structure and logical form. In Gareth Evans and John McDowell, editors, *Truth and Meaning*, pages 199–222. Oxford University Press, Oxford, 1978.

Hartry Field. Truth and the unprovability of consistency. *Mind*, 115:567–605, 2006.

Hartry Field. *Saving Truth from Paradox*. Oxford University Press, Oxford, 2008.

Manuel García-Carpintero. The grounds for the model-theoretic account of the logical properties. *Notre Dame Journal of Formal Logic*, 34:107–131, 1993.

Peter Geach. Entailment. *Proceedings of the Aristotelian Society Supplementary Volume*, 32:157–172, 1958.

Rolf George. Bolzano's consequence, relevance, and enthymemes. *Journal of Philosophical Logic*, 12:299–318, 1983.

Rolf George. Bolzano's concept of consequence. *The Journal of Philosophy*, 83:558–564, 1986.

Gilbert Harman. *Change in View*. MIT Press, Cambridge MA, 1986.

David Hume. *A Treatise of Human Nature*. Noon and Longman, London, 1739.

Kazimierz Kuratowski. Sur l'operation $\overline{A}$ de l'analysis situs. *Fundamenta Mathematicae*, 3:182–199, 1922.

Casimir Lewy. Entailment. *Proceedings of the Aristotelian Society Supplementary Volume*, 32:123–142, 1958.

Casimir Lewy. *Meaning and Modality*. Cambridge University Press, Cambridge, 1976.

David Makinson. The paradox of the preface. *Analysis*, 25:205–207, 1965.

David Makinson. *Bridges from Classical to Nonmonotonic Logic*. King's College London Publications, London, 2005.

Erroll Martin and Robert Meyer. S (for syllogism). ms, 1982.

Vann McGee. Two problems with Tarski's theory of consequence. *Proceedings of the Aristotelian Society*, 92:273–292, 1992.

Sebastiano Moruzzi and Elia Zardini. Conseguenza logica. In Annalisa Coliva, editor, *Filosofia analitica*, pages 157–194. Carocci, Roma, 2007.

Francesco Paoli. *Substructural Logics: A Primer*. Kluwer, Dordrecht, 2002.

Rohit Parikh. The problem of vague predicates. In Robert Cohen and Marx Wartofsky, editors, *Language, Logic, and Method*, volume 31 of *Boston Studies in the Philosophy of Science*, pages 241–261. Reidel, Dordrecht, 1983.

Karl Popper. Two autonomous axiom systems for the calculus of probabilities. *The British Journal for the Philosophy of Science*, 6:51–57, 1955.

Graham Priest. *In Contradiction*. Oxford University Press, Oxford, 2nd edition, 2006.

Willard Quine. *Word and Object*. MIT Press, Cambridge MA, 1960.

Willard Quine. *Philosophy of Logic*. Harvard University Press, Cambridge MA, 2nd edition, 1986.

Stephen Read. Validity and the intensional sense of 'and'. *Australasian Journal of Philosophy*, 59:301–307, 1981.

Stephen Read. *Relevant Logic*. Blackwell, Oxford, 1988.

Stephen Read. Logical consequence as truth-preservation. *Logique et Analyse*, 183:479–493, 2003.

Alfréd Rényi. On a new axiomatic theory of probability. *Acta Mathematica Academiae Scientiarum Hungaricae*, 6:286–335, 1955.

Greg Restall. *An Introduction to Substructural Logics*. Routledge, London, 2000.

Richard Routley, Val Plumwood, Robert Meyer, and Ross Brady. *Relevant Logics and Their Rivals*, volume I. Ridgeview, Atascadero, 1982.

Mark Sainsbury. *Logical Forms*. Blackwell, Oxford, 2nd edition, 2001.

Wesley Salmon. *The Foundations of Scientific Inference*. University of Pittsburgh Press, Pittsburgh, 1967.

Dana Scott. Completeness and axiomatizability in many-valued logic. In Leon Henkin, editor, *Proceedings of the Tarski Symposium*, pages 411–435. American Mathematical Society, Providence RI, 1974.

Timothy Smiley. Entailment and deducibility. *Proceedings of the Aristotelian Society*, 59: 233–254, 1959.

Nicholas Smith. Vagueness and blurry sets. *Journal of Philosophical Logic*, 33:165–235, 2004.

Richard Sylvan. *Sociative Logics and their Applications.* Aldershot, 2000.

Alfred Tarski. Fundamentale Begriffe der Methodologie der deduktiven Wissenschaften. I. *Monatshefte für Mathematik und Physik*, 37:361–404, 1930.

Alfred Tarski. O pojęciu wynikania logicznego. *Przegląd filozoficzny*, 39:58–68, 1936.

Neil Tennant. *Anti-Realism and Logic.* Clarendon Press, Oxford, 1987.

Achille Varzi. On logical relativity. *Philosophical Issues*, 10:197–219, 2002.

Georg von Wright. The concept of entailment. In *Logical Studies*, pages 166–191. Routledge, London, 1957.

Douglas Walton. Philosophical basis of relatedness logic. *Philosophical Studies*, 36:115–136, 1979.

Alan Weir. Naïve set theory is innocent! *Mind*, 107:763–798, 1998.

Ryszard Wójcicki. *Theory of Logical Calculi. Basic Theory of Consequence Operations.* Kluwer, Dordrecht, 1988.

Elia Zardini. A model of tolerance. *Studia Logica*, 90:337–368, 2008a.

Elia Zardini. *Living on the Slippery Slope. The Nature, Sources and Logic of Vagueness.* PhD thesis, Department of Logic and Metaphysics, University of St Andrews, 2008b.

Elia Zardini. Towards first-order tolerant logics. In Oleg Prozorov, editor, *Philosophy, Mathematics, Linguistics: Aspects of Interaction*, pages 35–38. Russian Academy of Sciences Press, St Petersburg, 2009.

Elia Zardini. Truth without contra(di)ction. *The Review of Symbolic Logic*, 4:498–535, 2011.

Elia Zardini. Truth preservation in context and in its place. In Catarina Dutilh-Novaes and Ole Hjortland, editors, *Insolubles and Consequences*, pages 249–271. College Publications, London, 2012.

Elia Zardini. Higher-order sorites paradox. *Journal of Philosophical Logic*, 42:25–48, 2013a.

Elia Zardini. It is not the case that [P and 'It is not the case that P' is true] nor is it the case that [P and 'P' is not true]. *Thought*, 1:309–319, 2013b.

Elia Zardini. Evans tolerated. In Kensuke Akiba and Ali Abasnezhad, editors, *Vague Objects and Vague Identity*. Springer, Berlin, 2013c. Forthcoming.

Elia Zardini. First-order tolerant logics. *The Review of Symbolic Logic*, 2013d. Forthcoming.

Elia Zardini. Getting one for two, or the contractors' bad deal. Towards a unified solution to the semantic paradoxes. In Theodora Achourioti, Kentaro Fujimoto, Henri Galinon, and José Martínez, editors, *Unifying the Philosophy of Truth*. Springer, Berlin, 2013e. Forthcoming.

Elia Zardini. Naive *modus ponens*. *Journal of Philosophical Logic*, 2013f. Forthcoming.

Elia Zardini. Naive truth and naive logical properties. *The Review of Symbolic Logic*, 2013g. Forthcoming.

Elia Zardini. Borderline cases. ms, 2013h.

Elia Zardini. Closed without boundaries. ms, 2013i.

Elia Zardini. Context and consequence. An intercontextual substructural logic. ms, 2013j.

Elia Zardini. Open knowledge of one's inexact knowledge. ms, 2013k.

Elia Zardini. Sequent calculi for tolerant logics. ms, 2013l.

Elia Zardini. The bearers of logical consequence. ms, 2013m.