(Meta)Inferential Levels of Entailment beyond the Tarskian Paradigm

Eduardo Alejandro Barrio* Federico Pailos† Damian Szmuc‡

University of Buenos Aires, Puan 480, C1420 Buenos Aires, Argentina
and
IIF-SADAF (CONICET), Bulnes 642, C1176ABL Buenos Aires, Argentina

Abstract

In this paper we discuss the extent to which the very existence of substructural logics puts the Tarskian conception of logical systems in jeopardy. In order to do this, we highlight the importance of the presence of different levels of entailment in a given logic, looking not only at inferences between collections of formulae but also at inferences between collections of inferences—and more. We discuss appropriate refinements or modifications of the usual Tarskian identity criterion for logical systems, and propose an alternative of our own. After that, we consider a number of objections to our account and evaluate a substantially different approach to the same problem.

1 Introduction

Since the Tarskian paradigm became the standard to characterise a logical system', there have been different attempts to generalize it appropriately. Multiple-conclusions, non-contractive, non-monotonic, non-reflexive and even non-transitive accounts of what a consequence relation might be, are only some of the directions in which these investigations have been headed. Carolina Blasio was especially interested in a number of them, and she produced substantial advances in understanding phenomena of these

* *Email*: eabarrio@gmail.com
† *Email*: fpailos@hotmail.com
‡ *Email*: szmucdamian@gmail.com

sorts. In particular, she proved that certain essentially four-valued structures called B-matrices are necessary to characterize logics whose consequence relations are both non-reflexive and non-transitive, despite being standard in all other relevant respects.[1]

Interestingly, Carolina did not think that unveiling discoveries of this sort about peculiar consequence relations should be taken as a mere curiosity, or as an investigation that pertained to (only) algebraic structures without having any connection to the main issues in logic and philosophical logic, but as incarnating

> a contribution to the discussion about the concept of entailment and, hence, also the understanding of logic as a discipline [6, p. 256]

In this vein, our aim in this article is to exemplify how we can put certain substructural logics—especially non-reflexive and non-transitive logics—to good use, in order to improve our understanding of the concept of entailment and of logic as a discipline. In this vein, we will be paying special attention to the logic **ST**, developed by Cobreros, Egré, Ripley and van Rooij's and presented e.g. in [10] and [11].[2]

Our particular goal is to highlight the extent to which the very existence of substructural logics necessitates reworking the answer to the question of how to differentiate between logics. This is so, we argue, because even though logics are usually associated with consequence relations between collections of formulae, this perspective needs to be supplemented with some additional considerations revolving around e.g. inferences between inferences, and so on.

In order to achieve this goal, the paper is structured as follows.

In Section 2 we present some thoughts on how the existence of substructural logics—in a qualified sense—calls for certain subtleties in the identification of logical systems. In Section 3 we generalize these worries, using some technical results from [1]. Section 4 discusses a novel identity criterion that we put forward, with the aim of overcoming the difficulties faced by the other approaches. A different identity criterion is evaluated, in Section 5,

[1] Essentially, B-matrices are algebraic structures which generalize regular logical matrices, in that they do not present an algebra with a distinguished set of designated elements, but rather with two distinguished sets—one of accepted and one of rejected elements. For a more detailed discussion of these results, see [5] and [6].

[2] For a more detailed explanation of why **ST** is usually referred to as Strict-Tolerant logic, in the context of three-valued logics defined using the strong Kleene algebra, see footnote 11 below.

with reasons being presented not to embrace it. In Section 6, we discuss an argument against the structural/substructural distinction presented in [42], and provide compelling reasons to reject it. Section 7 discusses the argument presented in [11] against the idea that a weaker logic in the metainferential level is necessary different from another one with more metainferential power. Finally, Section 8 includes some concluding remarks.

2 The Tarskian paradigm, and beyond

What identifies a logical system as such, making it possible to differentiate it from other logical systems? A first stab at this seemingly simple question could lead us to *identify a logical system with its set of valid formulae.* This answer might be motivated by taking a logic to be defined axiomatically, thereby putting the emphasis on the set of axioms and the set of statements derivable from them, with the help of the rules of that system—i.e. the theorems of the logic in question.

However fine this approximation might sound, it can be noticed that there are different logics which, nevertheless, share their set of theorems. A clear example is given by Classical Logic **CL** and Graham Priest's Logic of Paradox **LP**. That these systems have the same set of theorems is a well-known fact registered, e.g., in [29]. One may wonder, then, what motivates the claim that these logics are different. Additionally, one may wonder whether or not this motivation can be put to good use, in developing a full-fledged alternative identity criterion for logical systems? The answer is that the impression asking us to tell apart logics with the same set of theorems but different set of valid inferences, such as **CL** and Priest's logic, is driven by the idea that they do not validate the same set of inferences. As a matter of fact, it is a famous feature of Priest's system that it does not validate every instance of Modus Ponens[3] while Classical Logic is committed to its unrestricted validity, in each and every one of its instances.

Thus, the second stop in the quest for an identity criterion for logics suggests us to differentiate logics in virtue of their valid inferences. This alternative is, in fact, at the heart of the *status quo* in this debate. It is common knowledge that, according to the Tarskian lore, to describe a logic it is essential to understand its underlying consequence relation—as in [45]. This is no other than a relation between sets of formulae and single formulae, which respects the structural features, known as Reflexivity,

[3] To wit, think of a paradoxical sentence p and a false sentence q: then the inference from p and $p \supset q$ to q is not valid in **LP**.

Monotonicity, and Transitivity. Later, Scott, Smiley and Shoesmith thought it was a good idea to keep the symmetry between premises and conclusions, thereby allowing the latter to be sets too—e.g., in [41] and [44]—resulting in multiple-conclusion consequence relations, which we will discuss from now on. In this vein, it is customary to omit all other details and *identify a logic with its set of valid inferences* thus rendered by its underlying consequence relation.

Our aim in this article is to show that the identification of a logic with its set of valid inferences has its flaws. To show this, we will appeal to *substructural logics* which, we claim, pose some neglected difficulties that force us to go beyond the Tarskian paradigm. However, in order to discuss these frameworks, and in particular in order to understand the qualified sense in which we take logics to be substructural, it would be beneficial to stop for a moment in order to report on a few preliminary issues.

In what follows, we will be working with a propositional language $\mathcal{L}$ equipped with the connectives $\neg, \wedge, \vee$, to be interpreted as negation, conjunction and disjunction, letting $FOR(\mathcal{L})$ be the set of recursively generated well-formed formulae of $\mathcal{L}$. As usual, we will let Γ, Δ, and other Greek capital letters represent sets of formulae, or sets of inferences, and Roman capital letters A, B, C represent formulae or inference.[4] An inference $\Gamma \Rightarrow \Delta$ on $\mathcal{L}$ is an ordered pair $\langle \Gamma, \Delta \rangle$ where $\Gamma, \Delta \subseteq FOR(\mathcal{L})$, letting $SEQ^0(\mathcal{L})$ be the set of all inferences on $\mathcal{L}$.

Even after observing the obvious, that is, the fact that substructural logics are logics which are not structural, saying that a logic is structural could nevertheless mean a number of different things. Traditionally, this means that the set of valid inferences of the corresponding logic is closed under the *properties* of Reflexivity, Monotonicity, and Transitivity. In other words, that if certain inferences are valid according to a given logic, and such and such conditions are met, then some other inference is valid according to said logic.[5] In this vein, if a logic $\mathbf{L}$ has an underlying consequence relation $\vdash_{\mathbf{L}}$ we will say that an inference $\Gamma \Rightarrow \Delta$ is valid according to $\mathbf{L}$ just in case it belongs to $\vdash$, which we may record as $\vdash_{\mathbf{L}} \Gamma \Rightarrow \Delta$. As such, a logic is said to be *structural* if and only if, for all $A \in FOR(\mathcal{L})$, and all $\Gamma, \Delta \subseteq FOR(\mathcal{L})$:

1. $\vdash_{\mathbf{L}} A \Rightarrow A$
2. If $\vdash_{\mathbf{L}} \Gamma \Rightarrow \Delta$, $\Gamma \subseteq \Gamma'$, and $\Delta \subseteq \Delta'$, then $\vdash_{\mathbf{L}} \Gamma' \Rightarrow \Delta'$

[4] Though we hope the context will make things clear enough, we will always clearly state whether, for example, Roman capital letters A, B, C represent formulae, or inferences.

[5] For a comprehensive study of substructural logics understood in this way, see [28], [33], and [19].

3. If $\vdash_{\mathbf{L}} \Gamma, A \Rightarrow \Delta$ and $\vdash_{\mathbf{L}} \Gamma \Rightarrow A, \Delta$, then $\vdash_{\mathbf{L}} \Gamma \Rightarrow \Delta$

In relation with logics understood in this way, however, there is an interesting phenomenon that we want to analyze. It concerns to logics that are structural but, nevertheless, have substructural extensions. For instance, the logic **ST** which we discuss in this article is a structural— in particular, a transitive logic—which has non-transitive extensions.[6] This system can be presented either proof-theoretically, as Gentzen's calculus LK without the Cut rule, plus all the inverted rules for the connectives, or semantically in terms of Strong-Kleene valuations, as we will prefer to do in what follows.

It will be clear from the technicalities presented in Section 3 that, understood as a logical system over the usual language containing $\neg, \wedge$ and $\vee$, **ST** has the same valid inferences as **CL**. However, if the language is *extended* so as to include a unary predicate Tr meant to represent truth, and if additionally its valuations are restricted in a way that said predicate is transparent—i.e., that for all Strong Kleene valuations (SK-valuations) and all $A \in FOR(\mathcal{L})$, where $\langle A \rangle$ is a quotation name for A, we have that $v(Tr(\langle A \rangle)) = v(A)$[7]—the extended system is non-transitive and, therefore, *substructural*. The same can be said, *mutatis mutandis*, of the extension of **ST** with vague predicates and appropriate tolerance principles, or of its extension with Prior's problematic connective tonk, or of its extension with Read's paradoxical-like connective bullet—see, respectively, [10], [36], and [32]. All such extensions are non-transitive, even though the basic system that they are extending is perfectly structural, according to the traditional understanding of structurality.

In this respect, there appears to be some ground to the idea that if a logic was extended in a way that made it substructural, then the claim that said logic is structural is not as strong as it would be if all of its extensions were structural. This seems reasonable, more than ever, when the extensions—or theories based on them, if the notions added are not logical—in question involve, as in the case of of **ST** that we just discussed, notions that can be arguably regarded as logical—such as the truth predicate. When such notions are around, **ST**'s semantics do not render a structural but rather a

[6]This phenomenon, however, is not limited to **ST** or other convoluted systems, but actually pertains to many other logics. In this respect, see, e.g., [38].

[7]From a proof-theoretic point of view, this will amount to the addition of the so-called truth-rules, that is:

$$\frac{\Gamma, A \Rightarrow \Delta}{\Gamma, Tr(\langle A \rangle) \Rightarrow \Delta} \qquad \frac{\Gamma \Rightarrow A, \Delta}{\Gamma \Rightarrow Tr(\langle A \rangle), \Delta}$$

substructural logic. To us, this points towards the need to reflect upon the way in which we read the structural features of standard logical systems. If the usual way of understanding them does not shed a light on cases like the one we are considering, then perhaps this calls for a reconsideration of substructurality.

This line of reasoning can be further developed if, instead of understanding the structural features as closure properties of the set of valid inferences of a logic, we understand them as as *metainferences*, that is, as inferences between inferences themselves. More formally, a metainference $\Theta \Rightarrow_1 B$ on $\mathcal{L}$ is an ordered pair $\langle \Theta, B \rangle$, where $\Theta \subseteq SEQ^0(\mathcal{L})$ and $B \in SEQ^0(\mathcal{L})$, such that $SEQ^1(\mathcal{L})$ is the set of all metainferences on $\mathcal{L}$. Therefore, affirming that a logic is structural would amount to said metainferences being valid according to the logic in question.

At this point, it is important to stop and notice that this modification requires us to allow that when working with a logic there is more than one level at which consequence relations might be present. There is, of course, the level of the inferences between collections of formulae. But there is, additionally, a consequence relation holding between collections of inferences and inferences. That we usually focus on the former as a consequence relation, and only think of the latter as properties under which the set of valid inferences are closed does not preclude looking at connections, between collections of inferences and inferences as some kind of inferences themselves.

Be that as it may, this alternative point of view surely demands detailing a method, standard, or procedure to determine the validity of a certain metainference in the context of a given logic. In what follows, we will detail our own take on this issue, which does not mean that we discard other ways of answering this question. Thus, we will assume that a logic can be presented *semantically* and doing so requires detailing some kind of valuations over some sort of algebraic structure (matrices, models, or otherwise). Often the notion of validity for inferences is defined in a way that valuations over structures either represent, or do not represent a counterexample to the validity of the inference in question. In the former case, we will say that a given valuation v does not satisfy the inference $\Gamma \Rightarrow \Delta$ in the logic $\mathbf{L}$ (which we will symbolize as $v \nvDash_{\mathbf{L}} \Gamma \Rightarrow \Delta$) whereas in the latter we will say that v satisfies the inference in the logic $\mathbf{L}$ (which we will symbolize as $v \vDash_{\mathbf{L}} \Gamma \Rightarrow \Delta$). Of course, in this vein the fact that an inference is valid amounts to its satisfaction by all valuations over all structures (which we will symbolize as $\vDash_{\mathbf{L}} \Gamma \Rightarrow \Delta$), while its invalidity amounts to the existence of some valuation over some structure which does not satisfy it (which we will symbolize as $\nvDash_{\mathbf{L}} \Gamma \Rightarrow \Delta$). Whence, given a semantics for $\mathbf{L}$, $\vDash_{\mathbf{L}}$ is a consequence relation on $FOR(\mathcal{L})$.

In the context of semantics understood in this way, we could consider what we call a *local criterion of metainferential validity*.[8] According to this understanding, a metainference is valid if and only if it *preserves satisfaction by valuations for the logic in question*.

We are of the opinion that adopting this criterion to rule over metainferential validity allows for a great conceptual advantage. Namely, it helps provide a unified account of validity which can be applied to all kinds of inferences—that is, for inferences between formulae, between inferences, and so on, as we will see in what follows. In fact, when we assess the validity of regular inferences we ask that, if the premises are satisfied (according to the standard for premise-formulae to be satisfied in the logic in question), then at least one of the conclusions is satisfied (according to the standard for conclusion-formulae to be satisfied in that logic). This is precisely what the local criterion for metainferential validity asks for, when deciding about metainferences.

Therefore, if we are bound to accept that logic—as a sub-discipline of philosophy and mathematics—is concerned with the study and analysis of validity in general, then embracing something other than the local understanding of metainferential validity would mean that validity is read one way when formulae are involved (in inferences), and another way when inferences are involved (in metainferences). This, we claim, would be a rather unpleasant option. Much to the contrary, having a unified stance towards validity, regardless of its relata, appears to be a satisfying option. The local reading allows for this unification, and this is why we adopt it in the rest of this article.

Thus, the understanding of the structural features of inferential logical consequence as metainferences, and the adoption of the local criterion for metainferential validity allow us to present an alternative definition of structural logic. Thus, we could say that a given semantics for a logic $\mathbf{L}$ induces a *structural* consequence relation $\vDash_{\mathbf{L}}$ if and only if for all valuations v for $\mathbf{L}$, all $A \in FOR(\mathcal{L})$ and all $\Gamma, \Delta \subseteq FOR(\mathcal{L})$:

1. $v \vDash_{\mathbf{L}} A \Rightarrow A$
2. If $v \vDash_{\mathbf{L}} \Gamma \Rightarrow \Delta$ and $\Gamma \subseteq \Gamma'$, $\Delta \subseteq \Delta'$, then $v \vDash_{\mathbf{L}} \Gamma' \Rightarrow \Delta'$
3. If $v \vDash_{\mathbf{L}} \Gamma, A \Rightarrow \Delta$ and $v \vDash_{\mathbf{L}} \Gamma \Rightarrow A, \Delta$, then $v \vDash_{\mathbf{L}} \Gamma \Rightarrow \Delta$

[8] For more about this notion, and the difference between a *local* and a *global* notion of metainferential validity, see [14]. A similar distinction was previously introduced by Humberstone in [22].

As regards this new notion of structurality and substructurality, let us notice that it opens the door to a highly interesting philosophical and logical phenomenon.[9] Namely, *the existence of two different logical systems which have the same set of valid inference, but different valid metainferences.* We argue that this points to an issue that has been inadvertently floating in the literature and that, to the best of our knowledge, no one has called the attention to. It pertains to the unfavourable position in which the existence of substructural logics puts the received view, which identifies logical systems with sets of valid inferences.

Indeed, if two systems can have the same set of valid inferences, and at the same time it is reasonable to say that they represent different logics, then our identity criterion for logical systems ought to focus on something more than the inferences which are valid according to a logic. In the following section we show that there are logics, discussed nowadays in the philosophical literature, which have these features.

3 Concerns about the Tarskian paradigm

One may wonder whether the worry that there are different logical systems sharing the same valid inferences is abstract, or if there are examples of this sort out there in the literature. We think that the answer is affirmative. In this vein, we take the logic **ST** to be an exponent of this phenomenon.[10] For the purpose of this discussion, then, let us show how this logic is defined, how it coincides with **CL**, why it can be said to be different from it from a conceptual point of view, and how this philosophical difference can be cashed out formally.

Definition 1. The Strong Kleene algebra is the structure

$$\mathbf{K} = \langle \{1, \frac{1}{2}, 0\}, \{f^{\neg}_{\mathbf{K}}, f^{\wedge}_{\mathbf{K}}, f^{\vee}_{\mathbf{K}}\} \rangle$$

where the functions $f^{\neg}_{\mathbf{K}}, f^{\wedge}_{\mathbf{K}}, f^{\vee}_{\mathbf{K}}$ are as follows

[9] Before moving on to these issues, let us also stress that it could be possible to find out which is the proof-theoretic counterpart to local metainferential validity, for example, in the context of sequent calculi. Until now, we have not developed an account of this sort, but someone specially interested in looking at logical systems from a proof-theoretic point of view might be intrigued to look into these issues.

[10] Yet another example is Rosenblatt's works of recent appearance [38] and [39] where a non-contractive version of **CL** is presented.

	$f_{\mathbf{K}}^{\neg}$
1	0
$\frac{1}{2}$	$\frac{1}{2}$
0	1

$f_{\mathbf{K}}^{\wedge}$	1	$\frac{1}{2}$	0
1	1	$\frac{1}{2}$	0
$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	0
0	0	0	0

$f_{\mathbf{K}}^{\vee}$	1	$\frac{1}{2}$	0
1	1	1	1
$\frac{1}{2}$	1	$\frac{1}{2}$	$\frac{1}{2}$
0	1	$\frac{1}{2}$	0

Moreover, the connectives $\supset$ and $\leftrightarrow$ are definable via the usual definitions. That is, $A \supset B =_{def} \neg A \vee B$ and $A \leftrightarrow B =_{def} (A \supset B) \wedge (B \supset A)$.

Definition 2. A Strong Kleene valuation (SK-valuation, hereafter) is a mapping from $FOR(\mathcal{L})$ to $\{1, \frac{1}{2}, 0\}$ that respects the Strong Kleene truth-tables above. Similarly, a Boolean valuation is a SK-valuation whose range is $\{1, 0\}$.

Below, we define what is for a valuation to satisfy a given inference in **CL**, and what is for a valuation to satisfy a given inference in **ST**. As is common practice, we will define the validity of an inference (and, more generally below, of an inference of an arbitrary inferential level) in a certain logic as nothing more than satisfaction by all valuations—in our case, all SK-valuations.

Definition 3. A Boolean valuation v satisfies an inference $\Gamma \Rightarrow \Delta$ in **CL** ($v \vDash_{\mathbf{CL}} \Gamma \Rightarrow \Delta$) if and only if it is not the case that $v(A) = 1$ for all $A \in \Gamma$ and $v(B) = 0$ for all $B \in \Delta$. Similarly, a SK-valuation v satisfies an inference $\Gamma \Rightarrow \Delta$ in **ST** ($v \vDash_{\mathbf{ST}} \Gamma \Rightarrow \Delta$) if and only if it is not the case that $v(A) = 1$ for all $A \in \Gamma$ and $v(B) = 0$ for all $B \in \Delta$.[11]

[11] Another way to present **ST**'s validity requires talking about *strict* and *tolerant* satisfaction or truth—which explains why this is known as a *Strict-Tolerant* consequence relation. A valuation v satisfies *tolerantly* a formula A if and only if $v(A) \in \{1, \frac{1}{2}\}$, and satisfies it *strictly* if and only if $v(A) \in \{1\}$. Then, a valuation v satisfies an inference $\Gamma \Rightarrow \Delta$ in **ST** if and only if if v *strictly* satisfies every $B \in \Gamma$, then v *tolerantly* satisfies at least one $A \in \Delta$. Finally, an inference from Γ to Δ is valid in **ST** if and only if for every valuation v, if v satisfies *strictly* every $B \in \Gamma$, then v satisfies *tolerantly* some $A \in \Delta$. Nevertheless, it is worth mentioning that this is not the only way **ST**'s supporters explain their position. They prefer to talk about strict and tolerant *assertion* rather than talking about strict and tolerant *satisfaction*, or strict and tolerant *truth*. As [16] explains, the reason why they use the idea of strict and tolerant assertion instead of any of the last two (pair of) notions, is to avoid revenge paradoxes related to the notions of "strictly true" and "strictly false' in the context of truth-theories based on **ST**.

In a similar fashion, the logic **TS**—short for *Tolerant-Strict* logic—, that will also play a key role in this paper, is also usually presented in terms of the notions of *strict* and *tolerant* satisfaction. Specifically, a valuation v satisfies an inference $\Gamma \Rightarrow \Delta$ in **TS** if and only if if v *tolerantly* satisfies every $B \in \Gamma$, then v *strictly* satisfies at least one $A \in \Delta$. Thus, an inference from Γ to Δ is valid in **TS** if and only if for every valuation v, if v satisfies *tolerantly* every $B \in \Gamma$, then v satisfies *strictly* some $A \in \Delta$.

These definitions allow to prove the much commented result stating that the set of valid inferences of **CL** and **ST** coincide—as shown, e.g., in [20] and [10].

Fact 4. *For all* $\Gamma, \Delta \subseteq FOR(\mathcal{L})$*:*

$$\vDash_{\mathbf{ST}} \Gamma \Rightarrow \Delta \quad \textit{if and only if} \quad \vDash_{\mathbf{CL}} \Gamma \Rightarrow \Delta$$

With respect to SK-valuations, there are other two logics which can be straightforwardly presented by appeal to such functions. They are **LP** and the so-called strong three-valued logic $\mathbf{K_3}$ due to Stephen Cole Kleene in [24]. For the purpose of introducing these logics, let us detail how satisfaction by a SK-valuation is defined for them.

Definition 5. A SK-valuation v satisfies an inference $\Gamma \Rightarrow \Delta$ in **LP** ($v \vDash_{\mathbf{LP}} \Gamma \Rightarrow \Delta$) if and only if it is not the case that $v(A) \in \{1, \frac{1}{2}\}$ for all $A \in \Gamma$ and $v(B) = 0$ for all $B \in \Delta$. Similarly, a SK-valuation v satisfies an inference $\Gamma \Rightarrow \Delta$ in $\mathbf{K_3}$ ($v \vDash_{\mathbf{K_3}} \Gamma \Rightarrow \Delta$) if and only if it is not the case that $v(A) = 1$ for all $A \in \Gamma$ and $v(B) \in \{0, \frac{1}{2}\}$ for all $B \in \Delta$.

These differences can be alternatively discussed by saying that what differentiates **LP** from $\mathbf{K_3}$ is the choice of designated value, i.e., the set of values which are to be preserved from premises to conclusion. Thus, whereas in **LP** these are 1 and $\frac{1}{2}$, in $\mathbf{K_3}$ it is only 1.

Discussing these systems, moreover, allows us to look at **ST** in a different fashion. In fact, we can say that if satisfaction by an SK-valuation in a logic is defined such that a SK-valuation satisfies an inference $\Gamma \Rightarrow \Delta$ in a logic **L** if and only if it is not the case that $v \vDash_{\mathbf{L}'} \Gamma$ and $v \nvDash_{\mathbf{L}''} \Delta$, then **L** instantiates an approach to logical consequence that has been christened by Chemla, Egré and Spector in [9] as *mixed*. In this vein, we may rightfully refer to **L** as $\mathbf{L}'/L''$.

Our aim in this article is not to develop a definitive philosophical argument to establish that **CL** and **ST** are different logics. For all we know, there may not be such a thing out there, which does not involve some sort question-begging reasoning. Our work here is not intended to convince believers in the identity of **CL** and **ST** that they are mistaken. Instead, we aim at offering technical developments and precise arguments to sustain in a formal way the rather loose and informal feeling that some people had so far, to the extent that **CL** and **ST** are not identical.

From a philosophical point of view, one may argue for the difference of **ST** and **CL** using some of the following arguments.

A first reason for the difference between **ST** and **CL** lies in the way these two logics understand the notion of validity as pertains to inferences. While, on the one hand, **CL** can be depicted as a system whose underlying notion of logical consequence can be understood in terms of truth-preservation, **ST** cannot—as acknowledged in many works revolving around **ST** and other substructural logics, such as [35] and [9]. Thus, even though there is nothing intrinsically mistaken in reading logical consequence as something other than the preservation of truth from premises to conclusion, for us this suggests that **CL** and **ST** should be told apart.

A second reason for the difference between **ST** and **CL** is that their identification would require accepting that transitivity is not essential for a proper understanding of classical reasoning. However, as claimed in [4], transitivity seems to be rather central for classical reasoning, in that it allows us to proceed cumulatively, moving from lemmata to theorems and later to corollaries—regardless of the topic under discussion. Thus, if transitivity is is not granted, we cannot guarantee that lemmas are knowledge checkpoints. Given that this is a feature quite important to reasoning in **CL** but not in **ST**, for us this suggests that these systems should be told apart.

Thirdly, there are some widely shared conceptions about which logical principles are distinctively valid in **CL**. Indeed, Explosion, Excluded Middle, Modus Ponens, and Disjunctive Syllogism are only a few of these, which every proper presentation of **CL** should respect. As such, it is reasonable to expect them to hold in the context of all inference levels. That is, they should hold within regular inferences between formulae, but also within metainferences. However, as remarked in [4] and [48], **ST** invalidates what might be understood as certain metainferential forms of Explosion, Modus Ponens, Disjunctive Syllogism, and so on—in a way that it is intimately related to its invalidation of Transitivity. Once more, for us this points towards telling **CL** and **ST** apart.

Finally, there is the important fact that **CL** is widely believed to be prone to trivialization when faced with semantic paradoxes and vagueness, while **ST** does not suffer these flaws. Interestingly, advocates of the identity of these systems may respond to an argument of this kind by saying that **ST** is nothing but a different *mode of presenting* **CL**. Thus, instead of using two-valued valuations, **ST** presents Classical Logic by means of three-valued Strong Kleene valuations. In this regard, they would argue that the result of closing an arbitrary piece of content under **CL** crucially depends on the presentation of said logic that we are implementing. We disagree with such a line of counter-argumentation, precisely because we think that closing arbitrary pieces of content under numerically different logical systems and looking at the consequences of each attempt provides a way—albeit a rather *indirect* one—to compare the initial logics. In fact, if two logics render exactly the same consequences, we think it is fair to assume that, for all intents and purposes, they are the same logic. However, this is not the case of **CL** and **ST**, which for us justifies the need to differentiate them. That is, we think it is necessary to have a more *direct* way of identifying and differentiating them, by providing a formal identity criterion for logical systems.

Thus, the above referred conceptual difference between **CL** and **ST** can be cashed out in the following formal manner, by pointing out the divergence between those metainferences that are (locally) valid in **ST** and those that are valid in **CL**.

Fact 6. *There are some* $\Theta \subseteq SEQ^0(\mathcal{L})$ *and* $B \in SEQ^0(\mathcal{L})$ *such that:*

$$\not\models_{\mathbf{ST}} \Theta \Rightarrow_1 B \quad \textit{and} \quad \models_{\mathbf{CL}} \Theta \Rightarrow_1 B$$

In fact, we think that it is rather appropriate to state that the case of the previously discussed logics calls for a refinement of the traditional paradigm regarding the identity criterion for logical systems. A refinement, that is, which *identifies logics not only with their set of valid inferences, but additionally with their set of valid metainferences*. This account—in fact, the third attempt at providing an identity criterion for logics—could seem like a reasonable end of the road not only for those adhering to the traditional view, but also for those advocating for the substructural logics that forked out of it.

Nevertheless, news are not as good for them. The reason is that it is possible to devise logical systems which coincide with regard to their valid inferences and metainferences, despite of having different sets of valid inferences between metainferences—i.e. different valid metametainferences or metainferences of level 2. To put it intuitively, a metametainference is an inference between metainferences. Speaking more formally, a metametainference $\Xi \Rightarrow^2 C$ on $\mathcal{L}$ is an ordered pair $\langle \Xi, C \rangle$, where $\Xi \subseteq SEQ^1(\mathcal{L})$ and $C \in SEQ^1(\mathcal{L})$. We denote, analogously, by $SEQ^2(\mathcal{L})$ the set of all metametainferences on $\mathcal{L}$.[12]

In order to understand a logic related in this way to **CL**, we discuss the system **TSST**, as presented in [1]. For that purpose, it is important to

[12] The same criterion of validity applies to metametainferences as to metainference. Thus, when assessing the validity of metametainferences or metainferences of level 2, this feature is readily established in the same way in which metainferences of level 1 are ruled in or out in **ST**, i.e. in a local way.

understand how a different substructural logic, the non-reflexive logic **TS** discussed, e.g., in [11] and [18] works.

Definition 7. A SK-valuation v satisfies an inference $\Gamma \Rightarrow \Delta$ in **TS** if and only if it is not the case $v(A) \in \{1, \frac{1}{2}\}$ for all $A \in \Gamma$ and $v(B) \in \{\frac{1}{2}, 0\}$ for all $B \in \Delta$, which we denote by $v \vDash_{\mathbf{TS}} \Gamma \Rightarrow A$.[13]

Moving on to **TSST** now, it is important to notice that its notion of validity is targeted at metainferences, only later inducing a derived notion of validity for inferences, formulae and so on.

Definition 8. A SK-valuation v satisfies a metainference $\Theta \Rightarrow_1 B$ in **TSST** if and only if $v \nvDash_{\mathbf{TS}} \theta$ for some $\theta \in \Theta$, or $v \vDash_{\mathbf{ST}} B$, which we symbolize as $v \vDash_{\mathbf{TSST}} \Theta \Rightarrow_1 B$.

Using these definitions, the following witness our previous point, i.e., that there are logics with the same set of valid metainferences, but different metametainferences.

Fact 9. *For all $\Theta \subseteq SEQ^0(\mathcal{L})$ and $B \in SEQ^0(\mathcal{L})$:*

$$\vDash_{\mathbf{TSST}} \Theta \Rightarrow_1 B \quad \textit{if and only if} \quad \vDash_{\mathbf{CL}} \Theta \Rightarrow_1 B$$

Fact 10. *There are some $\Xi \subseteq SEQ^2(\mathcal{L})$ and $C \in SEQ^2(\mathcal{L})$ such that:*

$$\nvDash_{\mathbf{TSST}} \Xi \Rightarrow_2 C \quad \textit{and} \quad \vDash_{\mathbf{CL}} \Xi \Rightarrow_2 C$$

Facing a case of this sort, someone who thinks that **ST** and **CL** are not the same logic may devise a fairly obvious response. If these systems have different set of valid metametainferences, then tweak the identity criterion of a logic to rule that two logics are identical only when they have the same set of valid inferences, metainferences (of level 1) and metametainfenreces (or metainferences of level 2). This will, surely, help in telling **CL** and the aforementioned logic **TSST** apart, as well as any other logic that dares to defy this new criterion.

[13] We should note that validity for inferences both in **ST** and **TS** is not, and cannot be defined in terms of the usual account of preservation of distinguished values. For what is worth, preservation is both a transitive and a reflexive relation, and inferential validity is neither transitive in **ST** nor reflexive in **TS**. Notwithstanding this fact, validity in **ST** and in **TS** can be understood in terms of certain relations between distinguished values "for the premises" and distinguished values "for the conclusions"—as seen in, e.g., [25], [17], [9], and [35]. We would like to thank an anonymous reviewer for urging us to clarify these matters.

Once more, the sad news for someone arguing along these lines is not only that such an attempt will be fruitless, but that any similar attempt will be futile. Indeed, for any metainference of a finite level n—i.e. metainferences, metametainferences, metametametainferences, and so on and so forth—it is possible to present a logic $\mathbf{L}_n$ that coincides with **CL** up to that inferential level, as shown in [1]. This can be made precise in the following way. A generalized metainference $\Omega \Rightarrow^n D$ of level n on $\mathcal{L}$ (for $1 \leq n < \omega$) is an ordered pair $\langle \Omega, D \rangle$, where $\Omega \subseteq SEQ^{n-1}(\mathcal{L})$ and $D \in SEQ^{n-1}(\mathcal{L})$. $SEQ^n(\mathcal{L})$ is the set of all metainferences of level n on $\mathcal{L}$.

In [1], the authors prove that the systems in their *hierarchy* coincide with **CL** in incrementally many inferential levels. More formally, the hierarchy is presented through the following technicalities.

Definition 11. The collection $\mathbb{ST} = \{\mathbf{L}_i \mid i \in \mathbb{N}\}$ of logical systems is recursively defined so that $\mathbf{L}_0 = \mathbf{LP}$, $\mathbf{L}_1 = \mathbf{ST}$, and for $2 \leq j$, $\mathbf{L}_j = \overline{\mathbf{L}_{j-1}}/\mathbf{L}_{j-1}$ (where $\overline{\mathbf{ST}} = \mathbf{TS}$ and, in general, $\overline{\mathbf{L}_j} = \mathbf{L}_n/\mathbf{L}_m$ if $\mathbf{L}_j = \mathbf{L}_m/\mathbf{L}_n$).[14]

Definition 12. For $2 \leq j$ and $\mathbf{L}_j \in \mathbb{ST}$, a metainference $\Gamma \Rightarrow_{j-1} A$, where $\Gamma \subseteq SEQ^{j-2}(\mathcal{L})$ and $A \in SEQ^{j-2}(\mathcal{L})$, we say a SK-valuation v satisfies $\Gamma \Rightarrow_{j-1} A$ in $\mathbf{L}_j$ ($v \vDash_{\mathbf{L}_j} \Gamma \Rightarrow_{j-1} A$) if and only if $v \nvDash_{\overline{\mathbf{L}_{j-1}}} \gamma$ for some $\gamma \in \Gamma$, or $v \vDash_{\mathbf{L}_{j-1}} A$.

With the main result being the following, namely, that every logic $\mathbf{L}_{n+1}$ of said hierarchy coincides with **CL** up to its set of valid metainferences of level n, but later on it differs from it.

Theorem 13. *For all $n \geq 1$, for all $\Gamma \subseteq SEQ^{n-1}(\mathcal{L})$, $A \in SEQ^{n-1}(\mathcal{L})$*

$$\vDash_{\mathbf{L}_{n+1}} \Gamma \Rightarrow_n A \quad \textit{if and only if} \quad \vDash_{\mathbf{CL}} \Gamma \Rightarrow_n A$$

These facts, we claim, can be used to draw a number of conceptual consequences revolving around the question of what identifies a logical system as such, making it possible to differentiate it from other logics. In the next section, we touch on these issues in detail.

[14] Obviously, $\overline{\mathbf{ST}} = \mathbf{TS}$ because each one of these logics switches the standards for premises and conclusions that every valuation should meet in order for the inference to be sound in that logic. Thus, **ST**'s standard for premises is a *strict* one—i.e., $\{1\}$—, while its standard for conclusions is *tolerant* —i.e., $\{1, \frac{1}{2}\}$. Conversely, **TS**'s standard for premises is a *tolerant* one—i.e., $\{1, \frac{1}{2}\}$—, while its standard for conclusions is *strict* —i.e., $\{1\}$. More concretely, since we explained that **ST** can be understood as $\mathbf{K_3}/\mathbf{LP}$, the logic **TS** can thereby be understood as the system $\mathbf{LP}/\mathbf{K_3}$.

4 A new identity criterion for logics

We are convinced that these results support two related philosophical conclusions concerning identity criteria for logical systems.

The first conclusion is *negative*. We noticed that if an identity criterion rules two logics identical if and only if they have the same set of valid metainferences of some arbitrarily large metainferential level n, then it is possible to present a case of two logics having precisely these same features despite of being different in some other relevant respect. Therefore, *logics cannot be identified with their set of valid inferences or metainferences of any arbitrarily large metainferential level.*

The second conclusion is *positive*. Even though logical systems cannot be thus identified, for any two logical systems there will be in principle either an inference or a metainference of some level with regard to which they will disagree. Therefore, *there must be some identity criterion out there capable of cashing our intuitions out.*

In this vein, our proposal for such an identity criterion is the following. A logic is to be identified by an *infinite sequence* constituted by its set of valid inferences, its set of valid metainferences, its set of valid metainferences of level 2, and so on and so forth. Thus, let the set of valid inferences according to a logic $\mathbf{L}$ be denoted by $VSEQ^0_{\mathbf{L}}$, and the set of valid metainferences of level n according to $\mathbf{L}$ be denoted by $VSEQ^n_{\mathbf{L}}$. Then, a logic $\mathbf{L}$ is to be identified with an infinite sequence of the following form:[15]

$$\langle VSEQ^0_{\mathbf{L}}, VSEQ^1_{\mathbf{L}}, VSEQ^2_{\mathbf{L}}, VSEQ^3_{\mathbf{L}}, \ldots, VSEQ^n_{\mathbf{L}}, \ldots \rangle$$

Each of these sets can be understood as a consequence relation, albeit one holding between different relata. While the first one holds between collections of formulae, the second one holds between collections of inferences, the third one between collections of metainferences, and so the story goes.

Notice, importantly, that there is no natural endpoint for this sequence, since there always will be a more complex (although finite) inferential level that could be essential in telling apart one system from another logic that is extremely similar to it. Whence, instead of taking a logic to be identified with a single consequence relation, be it its inferential consequence relation or a metainferential consequence relation of some level, we identify a logic with a bundle or a sequence of consequence relations.

[15]We would like to thank Francesco Paoli for suggesting a terminology along these lines, and for discussion concerning these issues.

As pointed out by an anonymous reviewer this can be questioned by, e.g., taking incompatibility or invalidity to be as primitive a logical concept as consequence or validity. Thus, a logic should not be defined solely in terms of what it validates, but also in what it invalidates. In this respect, we would like to notice that such an objection will only hold some water if invalidity is not understood as the negation of validity—for, otherwise, deciding what is valid immediately determines what is not valid. If invalidity is understood in some other qualified sense, however, this objection might point towards an interesting issue that would be of great interest, and that we would like to analyze in future works.[16]

As we stressed a number of times before, applying this identity criterion has the effect of rendering **CL** different from the systems in the hierarchy and, particularly, from **ST**. The main advocates of this last logic, however, claim that the fact that certain extensions of **ST** invalidate key metainferences like Transitivity is not enough to argue that this logic is weaker or even not identical to **CL**.[17]. In this respect, the key to their particular understanding of logics as classes of consequence relations resides in pairing conse-

[16]An anonymous reviewer suggest to take into account Brandom's incompatibility semantics ([8]) into account for this purpose. Although that might be an interesting option that we would be eager to explore in the future, our clarification is mostly directed at highlighting the work that, e.g., Cobreros *et al* [12], Scambler [40], and Rosenblatt [37] have been doing on *anti*validity, as opposed to mere *in*validity. Moreover, there are some pieces of work that discriminate invalidities from the negation of validities. In particular, we are thinking about [26] and [27]. Nevertheless, those articles deals with the validity paradox, and not specifically with how a logic should be defined, which is what we are mostly interested in this work.

[17]See, for example, [11] However, these authors' disregard for metainferences is due to their taking metainferences to be properties under which the set of valid inferences of a given logic is closed under, whereas we think of them as logical inferences on their own right. This highlights the logical import of these objects for us, and the secondary role they have for these authors.

Moving on to the positive reasons to embrace our point of view concerning how to identify logical systems, we think one advantage thereof is its generality. Indeed, as pointed out previously, the fact that two logics can have the same set of valid inferences despite of having different valid metainferences is a general fact, and it is not restricted to the relationship that **CL** and **ST** have. This is witnessed by some recent investigations due to Melvin Fitting in [15], where it is shown that non-transitive versions of Kleene's three-valued logic $\mathbf{K_3}$, of **LP**, and of Belnap-Dunn's four valued logic $\mathbf{E_{fde}}$ can be studied. In the wake of such investigations, we wonder whether variations of this non-transitive kind of other logical systems—such as various many-valued logics, but also Intuitionistic Logic, and the normal modal logics—can be somehow introduced by semantic means.[18]

Another reason for adopting our identity criterion for logics is that it may lead to some new insights concerning two central issues associated to Logical Pluralism, that is, the view that there is more than one correct logic.

On the one hand, this identity criterion and the technical results which prompted its consideration have interesting effects on the success of the so-called Collapse Argument—as discussed, e.g., in [47], [30], [31], and [23]. The aim of such an argument is, in a nutshell, to show that someone adopting a plurality of logics might find herself being a Logical Monist, given the normative guidance that one is expected to have from a logic. In this vein, if a Pluralist believes the premises of an inference which is valid according to one of the logics she embraces but not to another, there can be no pluralism as to whether she should believe the conclusion or not. Thus, either she does or she does not, whence Logical Pluralism collapses into Logical Monism.

However, if in line with our identity criterion two different logics need not differ with regard to their valid inferences, then a Logical Pluralist embracing two logical systems with the same set of valid inferences need not be threatened by the current form of the Collapse Argument. As pointed out in [2], realizing this should suggest the opposing party that there is a need to refine such an argument, to cope with different logics sharing their set of valid inferences, but probably differing as regards their valid metainferences of some inferential level. This can be done, as shown in the previously referred article, more or less in the following way. Suppose a Pluralist embraces two different logics with the same inferences of level n but different inferences of level $n+1$. If she believes the premises of a metainference of level $n+1$ that is valid according to one of the logics she embraces, but not to another of such systems, there can be no pluralism as to whether or not to believe the conclusion of said metainference. Whence, Logical Pluralism collapses into Logical Monism, at last.

On the other hand, our novel perspective on the identity of logical systems also bears some consequences for the so-called Intra-theoretical Logical Pluralism. This species of Pluralism—defended, e.g., in Hjortland's work [21]—maintains that more than two logics can be hosted within the same theoretical framework. Hjortland himself identifies such a framework with a three-sided sequent calculus, showing that **LP** and Kleene's strong three-valued logic $\mathbf{K_3}$ can be thus understood. The same point, however, can be made by appealing to the fact that these systems are induced by a different selection of distinguished values, over the same truth-tables. Whence, the same semantics would host two different logics. These variants of intra-theoretical Logical Pluralism are interesting in their own right, but we want to point out how our technical results above suggest the existence of yet another incarnation of this idea.

In this respect, the issue of Logic Pluralism is usually the question of which are the correct logics and, therefore, of which are the correct inferences. The previously discussed results concerning the hierarchy presented in [1] show, for instance, that whether or not Explosion is valid, is a question that should be disambiguated. This is precisely because in the context of the same logic Explosion could be valid at one inferential level, but not at the other. Thus, for example, the logic **ST** coincides with **CL** up to its inferences and therefore validates Explosion at that level. Nevertheless, it does not coincide with **CL** concerning its metainferences, and in fact Explosion is invalid in its metainferential form.[19]

We know that this is a generalized phenomenon, one that, as discussed in the paragraphs above, can be generalized to other systems in the hierarchy, and to other logics besides Classical Logic. Therefore, it should be concluded that in some sense different logics can coexist within the same logical system, i.e., the hierarchy discussed presents a case whereby different logics can be allocated in different metainferential levels of the same system.

quence relations which hold between different kind of relata through certain appropriately constrained translation functions. For this purpose, their account makes essential use of a generalization of the Tarskian paradigm due to Blok and Jónsson, by means of which consequence relations are allowed to be fully abstract, holding between relata of all sorts.[20] More concretely, the following definitions are crucial in establishing their desired points.

Definition 14. An abstract consequence relation over the set A is a relation $\vdash \subseteq \wp(A) \times A$ obeying the following conditions for all $a \in A$ and for all $X, Y \subseteq A$:

- Reflexivity: $X \vdash a$ whenever $a \in A$
- Monotonicity: if $X \vdash a$, then $X \cup Y \vdash a$
- Transitivity: if $X \vdash b$ for all $b \in Y$ and $Y \vdash a$, then $X \vdash a$

It is interesting to note that, by definition, an abstract consequence relation is Tarskian, i.e. its enjoys all the usual structural properties, and moreover A is required to be a set granting Contraction, Exchange and Associativity *by default*. Furthermore, abstract consequence relations can be substitution-invariant just like Tarskian or substructural consequence relations can be, although the account of invariance for these newly introduced relations requires a considerable amount of subtleties—which is why we are not going to discuss it or any related issues in this paper.

Definition 15. Two abstract consequence relations $\vdash_1, \vdash_2$ over the sets A_1, A_2, respectively, are *similar* if there is a mapping $\tau : A_1 \to \wp(A_2)$ and a

5 Assessing an alternative

In this section we review Dicher and Paoli's attempt to provide an identity criteria for logics, which relies on an application of the Blok-Jónsson generalization of Tarskian consequence relations. Our argument against their view aims to show that from their perspective substructural logics, as well as multiple-conclusion logics, would be non-existent. Moreover, if they were existent, they would not be associated with the kind of systems that we usually refer to with these labels. Finally, we further conclude that applying their perspective renders the structural features like Reflexivity, Monotonicity and Transitivity ineffable—i.e., that there cannot be a logic where an axiom or rule of any sort represents any of them.

In [14] and [13], Dicher and Paoli take a different stance on the matter of which is the appropriate identity criterion for logical systems. In their opinion, a logic ought to be identified with a certain *equivalence* class of consequence relations

[20]For more on this, see [7].

mapping $\rho : A_2 \to \wp(A_1)$ such that the following conditions hold, for every $X \cup \{a\} \subseteq \{A_1\}$ and for every $b \in A_2$:

- S1: $X \vdash_1 a$ iff $\tau[X] \vdash_2 \tau(a)$
- S2: $b \dashv\vdash \tau(\rho(b))$

It should be noted that Blok and Jónsson reserve the stricter denomination that two abstract consequence relations are *equivalent* for relations satisfying substitution-invariance. For the present purposes, since we left aside the issue of invariance for the time being, these notions will be interchangeable. Be that as it may, it is important to notice that in light of the previous definitions it is straightforward to see how two consequence relations holding of different kind of relata—e.g. formulae and sequents, or formulae and equations—can be equivalent: we only need to require that there are appropriate translations functions, complying with the previously referred constraints.

The outlined account is not only important because it represents a different alternative to the criterion we advanced, that we might want to assess, but also because it implies a distinctive perspective on the alleged coincidence of the valid inferences of **CL** and **ST**. Precisely, Dicher and Paoli discuss whether or not it can be truly said that these two logics have the same set of valid inferences, to which they answer negatively.[21] To establish their point, they heavily draw from the fact that the logic **ST** is mainly presented as a sequent calculus—namely, the calculus they refer to as LK^-_{INV}, i.e. a variant of Gentzen's calculus LK for **CL** minus the Cut rule plus elimination rules for all the connectives.

Moreover, these authors claim that when looking at which is the candidate for the abstract consequence relation induced by a given sequent calculus, it is obvious that the target should be a relation between sets of sequents and sequents. This is, actually, no other than the calculus' *derivability* relation by means of which the sequent $\Gamma \Rightarrow \Delta$ follows from the set of sequents $\{\Gamma_1 \Rightarrow \Delta_1, \ldots, \Gamma_n \Rightarrow \Delta_n\}$ in the sequent calculus C if and only if the former is provable in the calculus that results from adding the latter sequents as axioms to C.

Thus, if **ST** is mainly a consequence relation relating sequents and **CL** is usually associated with a consequence relation holding between formulae, then appropriate translation functions are required in order to determine whether or not they share the same valid inferences. Interestingly

[21] For more details, see [14] and [13].

enough, when the corresponding transformers are at play it is shown that **ST**'s induced consequence relation among sequents corresponds to a relation between formulae that is nothing other than **LP**. Not **CL**, as expected. Therefore, given that **ST** and **CL** do not belong to the same class, they cannot be taken to represent the same logic or to have the same valid inferences, for that matter.

Now, before turning to an assessment of the advantages and disadvantages of Dicher and Paoli's implementation of Blok and Jónsson's abstract account of logical consequence, let us highlight that, contrary to what we have claimed, Dicher and Paoli make essential use of the fact that *a given logic embodies one and only one consequence relation*. In the case of sequent calculi, it corresponds to the one that holding between (collections of) sequents and sequents. The reader may object, at this point, that sequent calculi are intuitively built with the intention of establishing facts concerning inferences between collections of formulae, whence there must be a way of extracting such a thing from a given sequent calculus. Proof-theoretic folklore seems to indicate that this can be done by looking at the derivable or *provable sequents* of the said calculus.

Faced with such a challenge, these authors can only reply that they, at least implicitly, agree with these ideas. However, the way in which they propose to accommodate them is intriguingly different than the way in which one would expect them to be coped with. There is a way of extracting a consequence relation between formulae from a sequent calculus, but it is rather *indirect*. Instead of saying that an inference is valid in the logic associated with a given sequent calculus if the corresponding sequent is derivable, they have to answer that it is valid if a certain inference between sequents is derivable. That is, the inference between (collections of) sequents that results from translating the target inference between (collections of) formulae. Similarly, there is no *direct* way in which a proper consequence relation between inferences or sequents can be extracted from an axiomatic or natural deduction calculus, but only *indirectly* through an appropriate translation function to a sequent system.

Besides these issues, there is no doubt that Dicher and Paoli's utilization of the Blok–Jónsson approach has at least one advantage. It allows to clearly establish something that we already know from an intuitive point of view, i.e. that there can be different presentations of the same logic. We knew that the same logic could be presented with different notations, that it could be introduced by way of different proof-systems (Hilbert-style axiomatizations, natural deduction calculi, sequent calculi, etc.), and that it could hold of different relata (collections of formulae, equations, sequents, etc.) The Blok-

Jónsson point of view neatly paves the way to provide a formal confirmation of that idea. Thus, providing a clear formal account of equivalence between presentations of the same logic is definitely an advantage of it.

However useful this point of view may be, it nevertheless has many flaws, at least in its current form. Its main shortcoming is obvious from the outset, but deserves to be called out explicitly. It undergenerates, i.e. it does not qualify as logics systems which many people take to be logics.

Prime among these are substructural logics, the very kind of systems that motivated our investigations in the first place. As is easy to observe, the Blok-Jónsson definition of an abstract consequence relation satisfies the Tarskian axioms. Therefore, none of the equivalence classes induced by such an account is going to be substructural and, if logics are to be identified with equivalence classes rendered by the abstract definition, then this means that substructural logics do not exist—and, indeed, cannot exist—at all. A similar thing happens with multiple-conclusion logics. The Blok-Jónsson approach is single-conclusion by definition, whence no equivalence class induced by it is going to include a consequence relation that is multiple-conclusion. If logics are identified with these classes, then multiple-conclusion logics are rendered nonexistent—and, in fact, their existence is impossible in this context.

A defendant of the Blok-Jónsson account may argue that a straightforward modification of its definitions may indeed allow for an abstract characterization of logical systems which, in principle, could let substructural and multiple-conclusion logics in. Such a fully general modification would only require of an abstract consequence relation over an arbitrary set A to be just a relation without asking for any of the Tarskian axioms to be satisfied (except for substitution invariance, presumably). In other words, there might be room for a substructural version of the derivability relation where, e.g., repetitions of sequents count, and so on.[22] Notice, however, that this generalization will still give unintuitive results. As a matter of fact, given the way in which sequent calculi are associated to logical systems via their external consequence relation, this alternative route will classify some systems

[22]We would like to thank an anonymous reviewer for urging us to consider this option. If this option is considered, however, we may as well ask why a defendant of the Blok-Jónsson framework (like Dicher and Paoli) would opt for a generalization of said account but still claim that the derivability relation—i.e., the internal consequence, the set of derivable sequents—does not codify validity in the target calculus—because it might be a substructural consequence relation. From our point of view, the openness to this generalization undermines the reasons for rejecting the identification of the provable sequents of a calculus with the valid inferences of the logic associated with such a calculus.

as substructural or multiple-conclusion—just not the ones that we usually take to be either substructural or multiple-conclusion.

In fact, the failure of e.g. Reflexivity for the logic associated to a sequent calculi will not be reflected by the failure or non-derivability of the structural rule of Reflexivity, that is

$$\overline{\varphi \Rightarrow \varphi}$$

but of the metainference, which we may call meta-Reflexivity

$$\frac{\Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta}$$

Similar remarks can be made about substructural logics failing to be e.g. monotonic, contractive, or transitive. Logics are usually classified as non-monotonic if their allegedly corresponding sequent calculi are such that not all instances of the following rules are derivable.

$$\frac{\Gamma \Rightarrow \Delta}{\Gamma, \Sigma \Rightarrow \Delta} \qquad \frac{\Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Pi, \Delta}$$

Whereas, according to the utilization of the Blok-Jónsson account due to Dicher and Paoli, the logics induced by sequent calculi could be legitimately called non-monotonic only if the following transition was invalid.[23]

$$\mathit{from} \ \frac{\Gamma_1 \Rightarrow \Delta_1 \quad \ldots \quad \Gamma_n \Rightarrow \Delta_n}{\Gamma \Rightarrow \Delta} \ \mathit{to} \ \frac{\Gamma_1 \Rightarrow \Delta_1 \quad \ldots \quad \Gamma_n \Rightarrow \Delta_n \quad \Sigma \Rightarrow \Pi}{\Gamma \Rightarrow \Delta}$$

In turn, logics are counted as non-contractive if their corresponding sequent calculi are such that these rules are not derivable.

$$\frac{\Gamma, \varphi, \varphi \Rightarrow \Delta}{\Gamma, \varphi \Rightarrow \Delta} \qquad \frac{\Gamma \Rightarrow \psi, \psi, \Delta}{\Gamma \Rightarrow \psi, \Delta}$$

But under the current approach inspired by the abstract take on logical consequence, they could only be rightfully referred to as non-contractive if they invalidate the following

$$\mathit{from} \ \frac{\ldots \quad \Gamma_n \Rightarrow \Delta_n \quad \Gamma_n \Rightarrow \Delta_n \quad \ldots}{\Gamma \Rightarrow \Delta} \ \mathit{to} \ \frac{\ldots \quad \Gamma_n \Rightarrow \Delta_n \quad \ldots}{\Gamma \Rightarrow \Delta}$$

[23]We present this version of Monotonicity because we are working in a multi-conclusion framework. Nevertheless, we are aware that usually it is only Left-Weakening that is required to classify a logic as monotonic.

Likewise, while it is usual to say that a logic is non-transitive if the next rule is not derivable.

$$\frac{\Gamma, \varphi \Rightarrow \Delta \quad \Sigma \Rightarrow \varphi, \Pi}{\Gamma, \Sigma \Rightarrow \Pi, \Delta}$$

The present context demands that a properly speaking non-transitive logic rather invalidates the following

$$from \quad \frac{\Gamma_1 \Rightarrow \Delta_1 \quad \ldots \quad \Gamma_n \Rightarrow \Delta_n}{\Gamma \Rightarrow \Delta} \quad and \quad \frac{\Gamma \Rightarrow \Delta \quad \Gamma_1 \Rightarrow \Delta_1 \quad \ldots \quad \Gamma_n \Rightarrow \Delta_n}{\Sigma \Rightarrow \Pi}$$

$$to \quad \frac{\Gamma_1 \Rightarrow \Delta_1 \quad \ldots \quad \Gamma_n \Rightarrow \Delta_n}{\Sigma \Rightarrow \Pi}$$

Analogously, the presence of multiple-conclusion inferences in the logic associated to a given sequent calculus cannot be reflected by the presence of sequents whose succedents have greater-than-one cardinality. Instead, the extension of the Blok-Jónsson framework to a Scott-Shoesmith-Smiley multiple-conclusion framework should be represented by sequent rules whose conclusion is not a single sequent, but a set of sequents of the following form[24]

$$\frac{\Gamma_1 \Rightarrow \Delta_1 \quad \ldots \quad \Gamma_n \Rightarrow \Delta_n}{\Sigma_1 \Rightarrow \Pi_1 \quad \ldots \quad \Sigma_n \Rightarrow \Pi_n}$$

where, it might be assumed to keep the symmetry with the ongoing discussion, such a transition should be read as allowing to infer at least one of the conclusion sequents.[25]

One further consequence of this alternative account is that it renders the structural logical properties of Reflexivity, Contraction, Weakening, Transitivity and even those of Exchange and Associativity completely *ineffable*, i.e. they cannot be expressed in the object language of a given calculus. However, unlike other ineffability phenomena (such like those concerning truth and semantic notions), this is not a feature that we need to accept on pain

[24]We would like to than an anonymous reviewer for suggesting us to consider this option.

[25]An anonymous reviewer suggests that Dicher and Paoli might, in the end, not be as challenged by these objections as we think they are. The reason for this would be that they could have a cogent understanding (for instance, a bilateralist reading along the lines of [34]) both of the internal and of the external consequence relations associated with substructural sequent calculus. To this we reply that we do not doubt that such an option is available. What we do doubt is that their endorsement of the Blok-Jónsson framework allows them to place logic at both levels, whereas for them either the internal or the external consequence relation does not induce a proper consequence relation—philosophically speaking. On the other hand, since we do not embrace such a framework, we can so to say "see" proper consequence relations at more than one level, i.e., by looking at inferences, metainferences, metametainferences, and more.

of triviality. What renders these structural properties ineffable is the very fact that trying to express them with a rule immediately requires having a calculus whose main consequence relation handles objects of a degree of complexity higher than intended, whence the rule that allegedly represented the structural property now represents nothing but an inference of a certain form between entities of higher complexity.

This can be seen by reflecting on the case of the structural rules, discussed above. When introducing e.g. the Cut rule in a sequent calculus we are working—precisely—within a sequent calculus, i.e. a system whose main consequence relation is not between collections of formulae but between collections of sequents. Thus, the Cut rule only represents a certain rule for chaining sequents and never a property of inferences. Inferences, in sequent calculi, only happen at the level of sequent-to-sequent transitions and never at the level of the sequents themselves. That this phenomenon leads to said ineffability can be seen from the effects that moving to a calculus not having the previously referred Cut rule for sequents would have. Consequence, in such a calculus, will only happen at the level of metainference-to-metainference transitions—whence the target rule would, again, fail to express the intended structural property.

Finally, as a coda, let us assess a different point of view. If, as we claim, Dicher and Paoli's account implies the non-existence of substructural logic, it is worth saying a few words about Shapiro's works [42] and [43] where it is argued that the structural-substructural divide is more blurry than usually thought. Shapiro's argument is not diametrically opposed to Dicher and Paoli's (i.e., he does not argue that structural logics do not exist, and that all logics are substructural), but something weaker than that. His claim is that some logics than are paradigmatically seen as structural, like **LP** discussed above, can be presented in a way that makes it look like a substructural system. We will evaluate in detail Shapiro's argument in the next section.

6 A defense of the structural/substructural distinction

We have argued for the distinction between **CL** and **ST**. Basically, we think that, while the former codifies a structural consequence relation, the latter does not. What underlies this statement is the belief that the distinction between structural and substructural systems makes sense. This seems for us a pretty standard and obvious position. Nevertheless, in [42], Lionel Shapiro puts it into question. Specifically, Shapiro claims that such a difference does not make sense when applied to theories equipped with a transparent notion of truth. Although we have not talked extensively about theories of truth of this kind, we will discuss Shapiro's account, for we think that his point can be easily replicated with respect to the logics on top of which the theories of truth are built. Thus, in what follows, we will evaluate, and reply to, Shapiro's argument, which blurs the structural/substructural divide. The existence of such a distinction is a consequence of our identity criterion of logics, and Shapiro's argument, thus, puts a threat to our position. His point can be fairly summarized in the following quote:

> I'll argue that the division of such proposals—i.e., the truth theories presented by Shapiro as solutions to semantic paradoxes—into those that are fully structural and those that are substructural is less clear than it has seemed; indeed, I'll suggest that it embodies a confusion. In support of this conclusion, I'll present a case study where what deserves to count as the same logical response to the paradoxes can be represented in two different ways: (1) using a fully structural consequence relation that abandons standard rules for the language's conditional, as well as (2) using a consequence relation that preserves versions of these conditional rules by "going substructural." The question "Is paradox being blocked by invoking a consequence relation with non standard structure?" is thus ill-posed. [42, p. 2]

Thus, although he does not present a general proof of the claim that the distinction between structural and substructural theories of truth is ill-posed, he presents what he refers to as a *case study*. This consists of an analysis of the well-known truth theory **LPT**, that is, the transparent theory of truth based on **LP**. In this regard, he presents Beall's structural sequent calculi for **LPT**, which he refers to as $\mathbf{S_{LPT}}$, and which validate every classically valid metainference. Afterwards, Shapiro introduces his own version of a calculus for **LPT**, which he goes on to call $\mathbf{S'_{LPT}}$, allegedly representing a substructural calculus for **LPT**. From this, he concludes that **LPT**, a standard and well-known presumably structural truth-theory, can be presented either as a structural or a substructural system.

We think Shapiro's argument is wrong. The structural/substructural distinction makes perfect sense. We will not get into the details of the aforementioned systems, but for what is worth let us underscore that there is nothing essential in this argument about the theory of truth based on

LP. The argument, and the discussion, can equally be made as regards the base logic **LP**. For our purposes, it is enough to say that, while Beall's calculus is pretty standard, Shapiro's introduces a new way to bunch the premises of a sequent—with a semicolon instead of a comma, which is the traditional set-theoretic way to bunch premises. Thus, while every sequent in Beall's system—what Shapiro calls *ordinary sequents*—has a proof in Shapiro's calculus, the latter contains new sequents and, more importantly, new sequents that are derivable. The difference between the two kind of sequents can be detailed as follows:

> While a sequent's succedent will remain a set of sentences, its antecedent may now be either a set or an ordered pair of nonempty sets. Both possibilities can be represented using the notation $\Sigma; \Gamma \vdash \Delta$. When Σ is empty, this stands for $\Gamma \vdash \Delta$, and when Γ is empty, it stands for $\Sigma \vdash \Delta$. A sequent whose antecedent is a set will be called an ordinary sequent. As before, I use the comma for set union. [42, p. 9]

As a result of the pertinent definitions, Shapiro proves that in S'_{LPT} the general versions of the Exchange, Cut, and Contraction rules are all inadmissible, when applied to the semicolon structure. Nevertheless, the three of them are admissible when restricted to ordinary sequents. That is his main reason for claiming that S'_{LPT} is substructural and that, therefore, it induces a substructural logic. He then goes on to argue that:

> we would be right to wonder whether [the set of provable ordinary sequents of S'_{LPT}, that is to say, the set of valid inferences of **LPT′**] furnishes a novel response to paradox, distinct from Beall's own response using fully structural **LPT**. [42, p. 13]

In a nutshell, Shapiro claims that **LPT** can be viewed as a substructural theory of truth because it can be seen as induced by a substructural sequent calculus. Nevertheless, **LPT′** is just S'_{LPT}'s set of ordinary valid sequents. And not only **LPT′** matches **LPT** in all its validities, it is also fully structural—i.e., S'_{LPT} is fully structural when restricted to ordinary sequents, because every sound structural **LPT** rule is admissible in it, when limited to its ordinary sequents. Thus, **LPT′**—i.e., the truth theory itself—is not substructural. At most, the sequent calculus that fixes it is.

For us the situation is the following. Either (i) the truth-theory **LPT′** is just the inferences that correspond to the ordinary sequents that have a proof in S'_{LPT}, or (ii) the truth-theory **LPT′** corresponds to the whole set of inferences—i.e., ordinary and not-ordinary—that have a proof in S'_{LPT}. If (i) is true, then $\mathbf{LPT}'$ is just **LPT**, and therefore it is fully structural. If (ii) is true, then $\mathbf{LPT}'$ might be properly substructural. But then it furnishes a novel response to paradox, a very different one that **LPT**. In none of these situations it is the case that the same truth-theory have a fully structural but also a substructural version.

This is not the only threat that our claim that **ST**, being a substructural logic, and **CL**, being fully structural, are in fact two different logics. In the next section, we will assess another argument of this kind, due to Cobreros, Egré, Ripley, and van Rooij.

7 An argument against the relevance of metainferences

In their paper [11], Cobreros, Egré, Ripley, and van Rooij present an argument to supports the idea that **ST** is just **CL**, disguised in a three-valued suit. Their argument constitutes a reaction to the view that, all other things being equal, a logic $\mathbf{L}_1$ is weaker than a logic $\mathbf{L}_2$ if $\mathbf{L}_1$ has less valid metainferences as $\mathbf{L}_2$. To this extent, they present what they take to be some counterexamples to this idea.

The first of these pertains to modal logics. In this respect, they claim that **S4** is surely weaker than **S5**, because **S5** is an inferential strengthening of **S4**. In other words, because axioms need to be added to **S4** in order to obtain **S5**. This results in new theorems being provable in **S5**, and in new inferences being valid because of that. However, these authors point out, such a strengthening also causes the loss of some valid **S4**-metainferences. To wit this, they say, consider the following metainference:

$$\frac{\emptyset \Rightarrow \diamond p \rightarrow \Box \diamond p}{\emptyset \Rightarrow \bot}$$

It is easy to observe that **S4** is closed under this metainference, for the inference $\emptyset \Rightarrow \diamond p \rightarrow \Box \diamond p$ is not valid in it. However, such a metainference is not valid in **S5**, because although the inference $\emptyset \Rightarrow \diamond p \rightarrow \Box \diamond p$ is valid in it, the inference $\emptyset \Rightarrow \bot$ is not. With the intention of making a stronger point, these authors later prove that, given certain minimal background conditions, every time a logic is inferentially strengthened (i.e., every time axioms are added to a logic), then some previously valid metainferences are lost in the process—unless the logic is strengthened all the way up to the universal consequence relation.

The second counterexample that these authors mention pertains to the relation between propositional and first-order logic. In this case, to make their point they require some assumptions. In particular, they need to assume that both logics share the same language. Thus, quantified formulas like $\forall x Px$ are treated as atoms in the propositional version of **CL**, although they are appropriately considered as quantified formulas in the first-order version of **CL**. Given this, they consider the following metainference:

$$\frac{\forall x Px \Rightarrow Pa}{p \Rightarrow q}$$

Again, it is easy to observe that—given the previous assumptions—this metainference is invalid in the first-order version of **CL**, although it is valid in the propositional version of **CL**. But, surely, the authors claim, this is not enough to say that the former is weaker than the latter. Briefly, Cobreros, Egré, Ripley, and van Rooij think that these two examples are enough to prove their point that a metainferential weakness is not a substantial weakness after all.We are committed to the opposite view, as we think that two logics that coincide in the set of inferential validities, but have different metainferential validities, are two different logics. Moreover, we think that, at least, their arguments have a much narrower scope than the initially intended.

On the one hand, every time these authors refer to the validity or invalidity of a metainference, they are referring to it being *globally* valid or invalid, respectively. Thus, their examples does not affect our position, that is built in emphasizing the importance of *local* metainferential validity. Locally, the first metainference is also **S4**-invalid, as there are **S4**-models that satisfies $\emptyset \Rightarrow \diamond p \rightarrow \Box \diamond p$, but do not satisfy $\emptyset \Rightarrow \bot$. Pretty much the same happens with the second metainference, which is also *locally* invalid in the propositional version of **CL**, for there are valuations which satisfy $\forall x Px \Rightarrow Pa$—given these sentences should be treated as atoms—but do not satisfy $p \Rightarrow q$. Thus, none of the example affect our present discussion. In other words, they have not shown that an inferential strengthening of a logic necessary results in the loss of *locally* valid metainferences. Therefore, they have not justified their view that a metainferential weakness, in the local sense, is not a substantial weakness.

On the other hand, there are substantial differences between the aforementioned cases and the case of the **ST** phenomenon. We claim that **ST** and **CL** behave differently at the metainferential level, and this becomes clear—even from a *global* point of view—when we *add new vocabulary* to the language, specifically, when we add a transparent truth predicate. By this we mean that the effects of such an addition serve to fundamentally test whether two logics that coincide at the inferential level are really the same. In this respect, when we add a transparent truth predicate, **CL** and **ST** do not behave in the same way: as we discussed earlier, the former becomes trivial whereas the latter does not. However, the cases discussed by these authors are not of this kind, because the alleged counterexamples are devised in a way that no new vocabulary is added when changing from, e.g., the propositional to the first-order version of **CL**.

8 Conclusion

In this paper we discussed the extent to which the very existence of substructural logics questions the limited perspective that the Tarskian paradigm has about the identity criterion for logics. As a result of this, we showed that logics cannot be identified with a given set of valid inferences, nor with a given set of valid metainferences—of any arbitrary inferential level—thereby suggesting that a logic has to be identified with an infinite sequence of consequence relations, and not a single one. We evaluated some objections concerning the technicalities that led us to these results, showing that they cannot be reasonably held. Furthermore, we assessed a different identity criterion for logics coming from a generalization of the Tarskian paradigm due to Blok and Jónnson, discarding it on the basis of the fact that it precludes the existence of substructural and multiple-conclusion logics, together with the fact that it renders structural properties ineffable.

Acknowledgements We would like to thank Bogdan Dicher, Francesco Paoli, Greg Restall, Dave Ripley, and the members of the Buenos Aires Logic Group for discussing previous versions of this work with us. In addition, special thanks go to the anonymous reviewers for this journal whose suggestions helped to substantially improve the paper, and to the editors of the special issue for their patience and help throughout the editorial process.

References

[1] Eduardo Barrio, Federico Pailos, and Damian Szmuc. A Hierarchy of Classical and Paraconsistent Logics. *Journal of Philosophical Logic*. Forthcoming. DOI: 10.1007/s10992-019-09513-z.

[2] Eduardo Barrio, Federico Pailos, and Damian Szmuc. Substructural Logics, Pluralism and Collapse. *Synthese*. Forthcoming. DOI: 10.1007/s11229-018-01963-3.

[3] Eduardo Barrio, Federico Pailos, and Damian Szmuc. What is a Paraconsistent Logic? In Jacek Malinowski and Walter Carnielli, editors, *Contradictions, from Consistency to Inconsistency*, pages 89–108. Springer, 2018.

[4] Eduardo Barrio, Lucas Rosenblatt, and Diego Tajer. The logics of strict-tolerant logic. *Journal of Philosophical Logic*, 44(5):551–571, 2015.

[5] Carolina Blasio. Revisitando a Lógica de Dunn-Belnap. *Manuscrito*, 40(2):99–126, 2017.

[6] Carolina Blasio, João Marcos, and Heinrich Wansing. An Inferentially Many-Valued Two-Dimensional Notion of Entailment. *Bulletin of the Section of Logic*, 46:233–262, 2018.

[7] Willem Blok and Bjarni Jónsson. Equivalence of consequence operations. *Studia Logica*, 83(1):91–110, 2006.

[8] Robert Brandom. *Between Saying and Doing: Towards an Analytic Pragmatism*. Oxford University Press, Oxford, 2008.

[9] Emmanuel Chemla, Paul Egré, and Benjamin Spector. Characterizing logical consequence in many-valued logics. *Journal of Logic and Computation*, 27(7):2193–2226, 2017.

[10] Pablo Cobreros, Paul Egré, David Ripley, and Robert van Rooij. Tolerant, classical, strict. *Journal of Philosophical Logic*, 41(2):347–385, 2012.

[11] Pablo Cobreros, Paul Egré, David Ripley, and Robert van Rooij. Reaching transparent truth. *Mind*, 122(488):841–866, 2014.

[12] Pablo Cobreros, Elio La Rossa, and Lucas Tranchini. (I Can't Get No) Antisatisfaction. Unpublished manuscript.

[13] Bogdan Dicher and Francesco Paoli. The original sin of proof-theoretic semantics. *Synthese*. Forthcoming. DOI: 10.1007/s11229-018-02048-x.

[14] Bogdan Dicher and Francesco Paoli. ST, LP, and tolerant metainferences. In Can Başkent and Thomas Macaulay Ferguson, editors, *Graham Priest on Dialetheism and Paraconsistency*. Springer, Dordrecht. Forthcoming.

[15] Melvin Fitting. The Strict/Tolerant Idea and Bilattices. Unpublished manuscript.

[16] Andreas Fjellstad. Naive modus ponens and failure of transitivity. *Journal of Philosophical Logic*, 45(1):65–72, 2016.

[17] Szymon Frankowski. Formalization of a plausible inference. *Bulletin of*

the Section of Logic, 33(1):41–52, 2004.

[18] Rohan French. Structural reflexivity and the paradoxes of self-reference. *Ergo*, 3(5):113–131, 2016.

[19] Nikolaos Galatos, Peter Jipsen, Tomasz Kowalski, and Hiroakira Ono. *Residuated Lattices: An Algebraic Glimpse at Substructural Logics*. Studies in Logic and the Foundations of Mathematics. Elsevier, Amsterdam, 2007.

[20] Jean-Yves Girard. *Proof theory and logical complexity*. Bibliopolis, Napoli, 1987.

[21] Ole Thomassen Hjortland. Logical pluralism, meaning-variance, and verbaldisputes. *Australasian Journal of Philosophy*, 91(2):355–373, 2013.

[22] Lloyd Humberstone. Valuational semantics of rule derivability. *Journal of philosophical logic*, 25(5):451–461, 1996.

[23] Rosanna Keefe. What Logical Pluralism Cannot Be. *Synthese*, 191(7):1375–1390, 2014.

[24] Stephen Cole Kleene. *Introduction to metamathematics*. North-Holland, Amsterdam, 1952.

[25] Grzegorz Malinowski. Q-consequence operation. *Reports on Mathematical Logic*, 24(1):49–59, 1990.

[26] Toby Meadows. Fixed points for consequence relations. *Logique et Analyse*, 57(227), 2014.

[27] Federico Pailos and Diego Tajer. Validity in a dialetheist framework. *Logique et Analyse*, 70(238):191–202, 2017.

[28] Francesco Paoli. *Substructural Logics: a Primer*. Kluwer, Dordrecht, 2002.

[29] Graham Priest. The logic of paradox. *Journal of Philosophical logic*, 8(1):219–241, 1979.

[30] Graham Priest. Logic: One or many? In J. Woods and B. Brown, editors, *Logical Consequence: Rival Approaches. Proceedings of the 1999 Conference of the Society of Exact Philosophy*. Hermes, Stanmore, 2006.

[31] Stephen Read. Monism: The One True Logic. In D. de Vidi and T. Kenyon, editors, *A Logical Approach to Philosophy: Essays in Memory of Graham Solomon*. Springer, Dordrecht, 2006.

[32] Stephen Read. General-elimination harmony and the meaning of the logical constants. *Journal of Philosophical Logic*, 39(5):557–576, 2010.

[33] Greg Restall. *An Introduction to Substructural Logics*. Routledge, 2000.

[34] David Ripley. Paradoxes and failures of cut. *Australasian Journal of Philosophy*, 91(1):139–164, 2013.

[35] David Ripley. Revising Up: Strengthening Classical Logic in the Face

of Paradox. *Philosophers' Imprint*, 13(5):1–13, 2013.

[36] David Ripley. Anything goes. *Topoi*, 34(1):25–36, 2015.

[37] Lucas Rosenblatt. Bilateralism and Invalidities. Manuscript.

[38] Lucas Rosenblatt. Non-contractive classical logic. *Notre Dame Journal of Formal Logic*. Forthcoming.

[39] Lucas Rosenblatt. On structural contraction and why it fails. *Synthese*. Forthcoming. DOI: 10.1007/s11229-019-02241-6.

[40] Chris Scambler. Classical Logic and the Strict Tolerant Hierarchy. *Journal of Philosophical Logic*. Forthcoming. DOI: 10.1007/s10992-019-09520-0.

[41] Dana Scott. On engendering an illusion of understanding. *The Journal of Philosophy*, 68(21):787–807, 1971.

[42] Lionel Shapiro. The very idea of a substructural approach to paradox. *Synthese*. Forthcoming. DOI: 10.1007/s11229-016-1230-x.

[43] Lionel Shapiro. LP, K3 and FDE as Substructural Logics. In P. Arazim and T. Lavička, editors, *The Logica Yearbook 2014*, page 257–272. College Publications, London, 2016.

[44] David Shoesmith and Timothy Smiley. *Multiple-conclusion logic*. Cambridge University Press, Cambridge, 1978.

[45] Alfred Tarski. On the concept of logical consequence. In J. Corcoran, editor, *Logic, Semantics, Metamathematics: papers from 1923 to 1938*, page 409–420. Hackett, Indianapolis, IN., 2nd edition, 1983.

[46] Morgan Thomas. A Kripke-style semantics for paradox-tolerant, non-transitive intuitionistic logic. Unpublished manuscript.

[47] Timothy Williamson. Equivocation and existence. *Proceedings of the Aristotelian Society*, 88:109–127, 1987.

[48] Elia Zardini. Naive modus ponens. *Journal of Philosophical Logic*, 42(4):575–593, 2013.