THE REVIEW OF SYMBOLIC LOGIC
Volume 13, Number 2, June 2020

A FULLY CLASSICAL TRUTH THEORY CHARACTERIZED BY SUBSTRUCTURAL MEANS

FEDERICO MATÍAS PAILOS
Philosophy Department, University of Buenos Aires

Abstract. We will present a three-valued consequence relation for *metainferences*, called **CM**, defined through **ST** and **TS**, two well known *substructural* consequence relations for inferences. While **ST** recovers every classically valid inference, it invalidates some classically valid metainferences. While **CM** works as **ST** at the inferential level, it also recovers every classically valid metainference. Moreover, **CM** can be safely expanded with a transparent truth predicate. Nevertheless, **CM** cannot recapture every classically valid meta-metainference. We will afterwards develop a hierarchy of consequence relations $\mathbf{CM_n}$ for metainferences of level n (for $1 \leq n < \omega$). Each $\mathbf{CM_n}$ recovers every metainference of level n or less, and can be nontrivially expanded with a transparent truth predicate, but cannot recapture every classically valid metainferences of higher levels. Finally, we will present a logic $\mathbf{CM_\omega}$, based on the hierarchy of logics $\mathbf{CM_n}$, that is *fully classical*, in the sense that every classically valid metainference of any level is valid in it. Moreover, $\mathbf{CM_\omega}$ can be nontrivially expanded with a transparent truth predicate.

§1. Introduction. Nonclassical theories of truth pursue two conflicting desiderata. On the one hand, they search for a paradox-free transparent truth predicate. On the other hand, they want to retain as much classical logic as possible. This conflict is recently examined in [19]. There, Hjortland claimed that "nonclassical theories try to recapture classical principles in special cases. This is a form of damage control". Hjortland calls this desideratum "the *maxim of minimal mutilation*."

Within the family of substructural solutions to semantic paradoxes, the one that seems to have the best credentials regarding the "damage-control policy", is **ST+**. **ST+** is a truth-theory that has the nontransitive theory **ST** as its base logic.[1] What is special about it is that its supporters have a strong case for the claim that **ST** is just classical logic in disguise. The main reason for this assertion is that an inference is valid in classical logic if and only if it is valid in **ST**. A logic is defined in terms of its valid (and invalid) inferences. Therefore, it is not possible to discriminate between classical logic and **ST**.

Nevertheless, **ST** and classical logic seem to behave very different in other, and very relevant at that, respects. In particular, **ST** does not validate some key classically valid metainferences. A metainference can be roughly characterized as an inference that has a

Received: December 27, 2017.
2010 *Mathematics Subject Classification*: 03BXX, 03B47, 03B50, 03B60, 03B53.
Key words and phrases: logic and metainferences and metainferential validity and substructural logics and empty logic.

[1] A logic is usually understood as an order pair of a language and a consequence relation. A consequence relation is, informally, a *standard* that an inference may, or may not, live up to (or, extensionally, a set or collection of inferences). Nevertheless, in what follows we will freely use both terms as synonyms.

© Association for Symbolic Logic, 2019

doi:10.1017/S1755020318000485

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

set of inferences as premises, and a particular inference as the conclusion. An important metainference that **ST** gives up is Cut. But is not the only one. Metainferential forms of Modus Ponens and Explosion are also invalid in **ST**. Thus, **ST** is not *optimal* in keeping as much classical logic as possible. Moreover, a theory of truth based on a nonclassical logic that recovers not only every classically valid inference, but also every classically valid metainference, seems to better achieve this goal. In the following sections, we will present both a logic and a truth theory with those features. We will called them **TS/ST+** and **TS/ST**, respectively. One special feature of **TS/ST** is that two well known substructural consequence relations—namely, **ST** and **TS**—are used to characterize them in a pretty straightforward way.

This paper is structured as follows. In §2 we introduce in a more precise way the distinction between inferences and metainferences, along with a suitable notion of metainferential validity. We will also distinguish between different—finite—levels that a metainference can take. These ideas are essential to understand exactly what **TS/ST** accomplishes. We will also present two substructural logics, **ST** and **TS**, that will be used to define **TS/ST**. In §3 we present **TS/ST**, a metainferential consequence relation that not only validates every classically valid inference, as **ST**, but also every classically valid metainference (of level 1). Moreover, we will show how **TS/ST**—or, as we will also call it, **CM**—can be expanded with a transparent truth predicate. That new truth theory will be called **TS/ST+**. In §4 we introduced $\mathbf{CM}_2$, a consequence relation for metainferences of level 2. $\mathbf{CM}_2$ recovers not only the same things as **TS/ST**, but also recaptures every classically valid meta-metainference. This result is pushed one step forward in §5. In this section, we define a hierarchy of logics $\mathbf{CM_n}$. Every $\mathbf{CM_n}$ validates every classically valid metainference of level n or less, and can be safely expanded with a transparent truth predicate. In §6 we introduce the logic $\mathbf{CM}_\omega$. $\mathbf{CM}_\omega$ is the union of the theories in the hierarchy of logics $\mathbf{CM_n}$. Every classically valid metainference of any level is valid in $\mathbf{CM}_\omega$. Moreover, it can also be expanded with a transparent truth predicate. Finally, in §8 we offer some concluding remarks, and point out some possible future developments of this project.

§2. Metainferences, metainferential validity and two substructural logics. In the next section, we will introduce the target consequence relation for metainferences, **TS/ST** (that will later be called **CM**, to emphasize the relation that it has with the other logics that we will introduce). **TS/ST** will validate every classically valid sentence, inference and metainference of level 1. But it is not a standard consequence relation: it does not relate sets of *formulas* of a language $\mathcal{L}$, but sets of *inferences* (understood as pairs of sets of formulas of $\mathcal{L}$). But before moving on, it will be necessary to make some clarifications. In particular, we will need to make precise the distinction between *standard* inferences (as relations between sets of sentences) and metainferences. Moreover, we will need to clarify what we mean by *metainferential validity*. Finally, we will introduce in some detail two crucial substructural logics: **ST** and **TS**. They will play a key part in specifying how **TS/ST** works.

2.1. Inferences and metainferences. To understand and carry out our investigation, it will be essential to have an accurate grasp of the received view about inferences, metainferences and consequence relations.

To analyze these matters, it will be useful to fix some terminology. Let $\mathcal{L}$ be a propositional language, such that $\mathbf{FOR}(\mathcal{L})$ is the absolutely free algebra of formulae of $\mathcal{L}$, whose

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

universe we denote by $FOR(\mathcal{L})$. As usual, we will let Γ, Δ, and other Greek capital letters represent *finite* sets of formulae, while γ, δ, and other Greek minuscules represent formulae such that $\gamma \in \Gamma$, $\delta \in \Delta$. Finally, A, B, C and D represent formulae that does not belong to sets of formulae represented by Greek capital letters.

DEFINITION 2.1. *A Tarskian consequence relation over a propositional language $\mathcal{L}$ is a relation $\vDash \subseteq \mathcal{P}(FOR(\mathcal{L})) \times FOR(\mathcal{L})$ obeying the following conditions for all $A \in FOR(\mathcal{L})$ and for all $\Gamma, \Delta \subseteq FOR(\mathcal{L})$:*

(1) $\Gamma \vDash A$ *if* $A \in \Gamma$ (*Reflexivity*)
(2) *If* $\Gamma \vDash A$ *and* $\Gamma \subseteq \Gamma'$, *then* $\Gamma' \vDash A$ (*Monotonicity*)
(3) *If* $\Delta \vDash A$ *and* $\Gamma \vDash B$ *for every* $B \in \Delta$, *then* $\Gamma \vDash A$ (*Cut*).

Additionally, a (Tarskian) consequence relation $\vDash$ is substitution-invariant whenever if $\Gamma \vDash \Delta$, and σ is a substitution on $\mathbf{FOR}(\mathcal{L})$, then $\{\sigma(A) \mid A \in \Gamma\} \vDash \{\sigma(B) \mid B \in \Delta\}$.

DEFINITION 2.2. *A Tarskian logic over a propositional language $\mathcal{L}$ is an ordered pair $(\boldsymbol{FOR}(\boldsymbol{L}), \vDash)$, where $\vDash$ is a substitution-invariant Tarskian consequence relation.*

Notwithstanding, many philosophers have argued against the Tarskian conception of logic. For example, Shoesmith and Smiley [31], Avron [2], and Scott [30] argue for the generalization of the Tarskian account multiple consequences. Avron [2] and Gabbay [17] give reasons for relaxing the Monotonicity condition. Moreover, Malinowski [23] and Frankowski [14] argued for a generalization or liberalization that allows logics to drop Reflexivity, Cut, or both of them. It is this context where logics like the nonreflexive **TS** of the nontransitive **ST** where introduced.

These modifications, in turn, bring on the possibility of a shift in the nature of the collection of *formulae* featured in the consequence relation. Thus, instead of treating logical consequence to hold between (sets of) formulae, it may hold between labeled formulae, multisets of formulae (where repetition matters), sequences of formulae (where order makes a difference), with multi-conclusions instead of single-conclusions, etc. Our target logic, **TS/ST**, explicitly moves from understanding logical consequence as a relation between collections of *formulae*, to conceiving it as a relation between collections of *inferences*. Here, we should understand *inferences* as an ordered pair whose first member is a set (or multiset, or sequence, etc.) of formulae, and a set of formulae as its second term. A *metainference*, then, is an ordered pair whose first member is a set (or multiset, etc.) of *inferences*, and a single inference as its second term. This is a relatively standard move in the literature. Avron, in [1], first, and Blok and Jónsson in [6], second, discuss a generalization of the Tarksian account that allows moving to logical consequence relations that do not hold only between collections of formulae, but between objects of other nature.[2] Thus, more precisely:

DEFINITION 2.3. *An inference on $\mathcal{L}$ is an ordered pair (Γ, Δ), where $\Gamma, \Delta \subseteq FOR(\mathcal{L})$ (written $\Gamma \vDash \Delta$). $SEQ^0(\mathcal{L})$ is the set of all inferences or sequents on $\mathcal{L}$.*

[2] Moreover, Avron suggested in [1] that the idea that logical consequence can be said to hold of *relata* other than formulae is very reasonable to those used to *sequent calculus* systems—and, most prominently, with substructural sequent calculi.

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

DEFINITION 2.4 ([11]). *A metainference on $\mathcal{L}$ is an ordered pair (Γ, A), where $\Gamma \subseteq SEQ^0(\mathcal{L})$ and $A \in SEQ^0(\mathcal{L})$ (written $\Gamma \vDash^1 A$). $SEQ^1(\mathcal{L})$ is the set of all metainferences on $\mathcal{L}$.*[3]

We will say, accordingly, that the following one on the left is an inference, whereas the one on the right is a metainference

$$\Gamma \vDash A \qquad\qquad \frac{\Delta \vDash B \quad \Sigma \vDash C}{\Gamma \vDash A}.$$

Therefore, we will distinguish two different kinds of consequence relations: the *inferential* and the *metainferential*. **TS/ST**—our target logic—is a consequence relation of this last type. In the next subsection we will give a precise characterization of this notion.

2.2. Metainferential validity. A metainference, thus, is a pair (Γ, δ), where Γ is a set of inferences and δ is a particular inference. Every metainference will have the following structure:

$$\frac{\Gamma_1 \vDash \Delta_1, \ldots, \Gamma_n \vDash \Delta_n}{\Sigma \vDash \Pi}$$

We will call every $\Gamma_i \vDash \Delta_i$, a *premise* of the metainference, while $\Sigma \vDash \Pi$ will be its *conclusion*.

Let $\Gamma_i \vDash \Delta_i$ (for $1 \leq i \leq n$) and $\Sigma \vDash \Pi$ be inferences, and let $\vDash_1$ stand for the (meta)inferential relation between premises and conclusion of a metainferences. An immediate question that pops up is the following: when is a metainference $\Gamma_1 \vDash \Delta_1, \ldots, \Gamma_n \vDash \Delta_n \vDash_1 \Sigma \vDash \Pi$ valid in a logic **L**? Here is a plausible answer:

DEFINITION 2.5. *A metainference $\Gamma_1 \vDash \Delta_1, \ldots, \Gamma_n \vDash \Delta_n \vDash_1 \Sigma \vDash \Pi$ is valid in (a propositional logic) **L** if and only if, for every valuation v, if v satisfies every $\Gamma_i \vDash \Delta_i$ according to **L**, then v satisfies $\Sigma \vDash \Pi$ according to **L**.*

A valuation v satisfies an inference $\Gamma \vDash \Delta$ in a specific logic if and only if v is not a counterexample of $\Gamma \vDash \Delta$'s validity in that logic.

This way to characterized the notion of metainferential validity is known as a "local conception of metainferential validity" ([11]).[4] As it stands, the definition specifies what it takes for a particular metainference to be valid in a particular logic. Nevertheless, it can—and will—be used to specified when a metainferential *scheme* is valid in a logic **L**. In a nutshell, a scheme is valid in **L** if and only if every *instance* of it is valid.

TS/ST, our target logic, is a consequence relation for metainferences. To understand exactly how it works, we need to introduce first two *inferential* but *non-Tarskian* consequence relations: the non-reflexive **TS** and the non-transitive **ST**.[5]

2.3. ST and TS: Two substructural consequence relations. **ST** and **TS** are labelled as *substructural* because at least one *structural* feature of a Tarskian consequence relation

[3] We will not consider multi-conclusions metainferences. This is not because we think there is something wrong with them, but because it will unnecessarily complicate the presentation of the logics we will further introduce.

[4] For more about this notion, and the difference between a *local* and a *global* notion of metainferential validity, see [11]. A similar distinction was previously introduced by Humberstone, in [20].

[5] Those readers that are already familiar with these logics may safely skip the following subsection.

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

is given up by them. **ST** abandons Cut, while **TS** drops *Reflexivity*. We will present propositional versions of **ST** and **TS**.

The logic **ST** can be portrayed as a *p-logic*, as devised by Frankowski in [14] as a means to characterize logical systems where valid derivations are such that the degree of strength of the conclusions can be smaller than that of the premises.[6]

DEFINITION 2.6 ([14]). *A p-consequence relation over a propositional language $\mathcal{L}$ is a relation $\vDash \subseteq \mathcal{P}(FOR(\mathcal{L})) \times \mathcal{P}(FOR(\mathcal{L}))$ obeying the following conditions for all $A \in FOR(\mathcal{L})$ and for all $\Gamma, \Delta \subseteq FOR(\mathcal{L})$:*

(1) $\Gamma \vDash \Delta$ *if for some* $A \in \Delta$, $A \in \Gamma$ (*Reflexivity*)

(2) *If* $\Gamma \vDash \Delta$ *and* $\Gamma \subseteq \Gamma'$, *then* $\Gamma' \vDash \Delta$ (*Monotonicity*).

Additionally, a p-consequence relation $\vDash$ is substitution-invariant whenever if $\Gamma \vDash \Delta$, and σ is a substitution on $\boldsymbol{FOR}(\mathcal{L})$, then $\{\sigma(B) \mid B \in \Gamma\} \vDash \sigma(A)$ -for some $A \in \Delta$.

DEFINITION 2.7 ([14]). *A p-logic over a propositional language $\mathcal{L}$ is an ordered pair $(\boldsymbol{FOR}(\boldsymbol{L}), \vDash)$, where $\vDash$ is a substitution-invariant p-consequence relation.*

In general, *p-logics* can be connected to *p*-matrices. **ST** can be represented as a *p-matrix* logic associated to the 3-element Kleene algebra.

DEFINITION 2.8 ([15]). *For $\mathcal{L}$ a propositional language, an $\mathcal{L}$-p-matrix is a structure $\langle \mathcal{V}, \mathcal{D}^+, \mathcal{D}^-, \mathcal{O} \rangle$, such that $\langle \mathcal{V}, \mathcal{O} \rangle$ is an algebra of the same similarity type as $\mathcal{L}$, with universe $\mathcal{V}$ and a set of operations $\mathcal{O}$, where $\mathcal{D}^+, \mathcal{D}^- \subseteq \mathcal{V}$ and $\mathcal{D}^+ \subseteq \mathcal{D}^-$.*

DEFINITION 2.9. *The 3-element Kleene algebra is the structure*

$$\boldsymbol{K} = \langle \{1, \frac{1}{2}, 0\}, \{f_{\boldsymbol{K}}^{\neg}, f_{\boldsymbol{K}}^{\wedge}, f_{\boldsymbol{K}}^{\vee}\} \rangle$$

where the functions $f_{\boldsymbol{K}}^{\neg}, f_{\boldsymbol{K}}^{\wedge}, f_{\boldsymbol{K}}^{\vee}$ are as follows

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

$f_{\boldsymbol{K}}^{\wedge}$	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

$f_{\boldsymbol{K}}^{\vee}$	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 functions $\vDash$ and are definable via the usual definitions.

DEFINITION 2.10 ([8]). *A 3-valued **ST**-matrix is a p-matrix*

$$\mathcal{M}_{\boldsymbol{ST}} = \langle \{1, \frac{1}{2}, 0\}, \{1\}, \{1, \frac{1}{2}\}, \{f_{\boldsymbol{K}}^{\neg}, f_{\boldsymbol{K}}^{\wedge}, f_{\boldsymbol{K}}^{\vee}\} \rangle$$

such that $\langle \{1, \frac{1}{2}, 0\}, \{f_{\boldsymbol{K}}^{\neg}, f_{\boldsymbol{K}}^{\wedge}, f_{\boldsymbol{K}}^{\vee}\} \rangle$ is the 3-element Kleene algebra.

Now, as is common practice, semantic structures such as *p*-matrices induce consequence relations and, therefore, *logics* with the help of valuation functions, e.g., homomorphisms from $\mathbf{FOR}(\mathcal{L})$ to the set of truth-values of the semantic structure in question—in this case, the set $\{1, \frac{1}{2}, 0\}$. Valuations are extended from propositional variables to complex formulae with the help of the truth-functions for the connectives; in this case the functions given by

[6] For an extensive presentation of **ST**, see also [8], [27], [29], and [9].

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

the 3-element Kleene algebra. In our particular case, we can define what a valid inference or sequent is in any *p-matrix logic*—and, therefore, in **ST**—in the following straightforward manner. Notice that, below, $\models_{\mathcal{M}}$ is a substitution-invariant *p-consequence relation*, whence $(\mathbf{FOR}(\mathcal{L}), \models_{\mathcal{M}})$ is a *p-logic*. In addition to that, when a *p-* matrix $\mathcal{M}$ induce some *p-logic* **L**, we may interchangeably refer to $\models_{\mathcal{M}}$ as $\models_{\mathbf{L}}$.

DEFINITION 2.11. *For $\mathcal{M}$ a p-matrix, an $\mathcal{M}$-valuation v satisfies a sequent or inference $\Gamma \models \Delta$ (written $v \models_{\mathcal{M}} \Gamma \models \Delta$) if and only if if $v[\Gamma] \subseteq \mathcal{D}^+$ (e.g., if $v(\gamma) \in \mathcal{D}^+$, for every $\gamma \in \Gamma$), then $v(A) \in \mathcal{D}^-$, for some $A \in \Delta$. A sequent or inference $\Gamma \models \Delta$ is $\mathcal{M}$-valid (written $\models_{\mathcal{M}} \Gamma \models \Delta$) if and only if $v \models_{\mathcal{M}} \Gamma \models \Delta$, for all $\mathcal{M}$-valuations v.*

But a more straightforward characterization of **ST**'s validity is the following one:[7]

$$\models_{\mathbf{ST}} \Gamma \models \Delta \text{ if and only if } \text{ for every valuation } v, \begin{cases} \text{if } v(B) \in \{1\} & \text{for all } B \in \Gamma \\ \text{then } v(A) \in \{1, \frac{1}{2}\} & \text{for some } A \in \Delta. \end{cases}$$

Another interesting generalization of Tarskian consequence relations is the notion of *q*-consequence relation, due to Malinowski [23].[8]

DEFINITION 2.12 ([23]). *A q-consequence relation over a propositional language $\mathcal{L}$ is a relation $\models \subseteq \mathcal{P}(FOR(\mathcal{L}))\times \subseteq \mathcal{P}(FOR(\mathcal{L}))$ obeying the following conditions for all $A \in FOR(\mathcal{L})$ and for all $\Gamma, \Delta \subseteq FOR(\mathcal{L})$:*

(1) *If $\Gamma \models \Delta$ and $\Gamma \subseteq \Gamma'$, then $\Gamma' \models \Delta$ (Monotonicity)*
(2) *$\Gamma \cup \{A \mid \Gamma \models A\} \models \Delta$ if and only if $\Gamma \models \Delta$ (Quasi-closure).*

DEFINITION 2.13 ([23]). *A q-logic over a propositional language $\mathcal{L}$ is an ordered pair $(\boldsymbol{FOR}(\boldsymbol{L}), \models)$, where $\models$ is a substitution-invariant q-consequence relation.*

DEFINITION 2.14 ([23]). *For $\mathcal{L}$ a propositional language, an $\mathcal{L}$-q-matrix is a structure $\langle \mathcal{V}, \mathcal{D}^+, \mathcal{D}^-, \mathcal{O} \rangle$, such that $\langle \mathcal{V}, \mathcal{O} \rangle$ is an algebra of the same similarity type as $\mathcal{L}$, with universe $\mathcal{V}$ and a set of operations $\mathcal{O}$, where $\mathcal{D}^+, \mathcal{D}^- \subseteq \mathcal{V}$ and $\mathcal{D}^+ \cap \mathcal{D}^- = \emptyset$.*

The 3-valued *q*-matrix logics associated to the 3-element Kleene algebra that will be dealing with in our ongoing investigation is the logic **TS**.

---

[7] Another way to present **ST**'s validity requires talking about *strict* and *tolerant* satisfaction or truth. 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 $\Gamma \models \Delta$ 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 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 [13] 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**.

[8] Wansing and Shramko offer in [32] a way to read these two kinds of logics. While a *p-logic* is devised to qualify as valid derivations of conclusions whose degree of strength (understood as the conviction in its truth) is smaller than that of the premises, the relation of *q*-logic is devised to qualify as valid derivations of true sentences from nonrefuted premises (understood as hypotheses).

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

DEFINITION 2.15 ([8], [24]). *A 3-valued **TS**-matrix is a q-matrix*

$$\mathcal{M}_{\boldsymbol{TS}} = \langle \{1, \frac{1}{2}, 0\}, \{1\}, \{0\}, \{f_{\boldsymbol{K}}^{\neg}, f_{\boldsymbol{K}}^{\wedge}, f_{\boldsymbol{K}}^{\vee}\}\rangle$$

such that $\langle \{1, \frac{1}{2}, 0\}, \{f_{\boldsymbol{K}}^{\neg}, f_{\boldsymbol{K}}^{\wedge}, f_{\boldsymbol{K}}^{\vee}\}\rangle$ *is the 3-element Kleene algebra.*

TS is discussed by e.g., Cobreros, Ripley, Egré and van Rooij in [8], and also by Chemla, Egré, and Spector in [7] in the context of the more general discussion of what represents a "proper" consequence relation between formulae. Moreover, it was also discussed by Malinowski in [24] as a tool to model empirical inference with the aid of the 3-valued Kleene algebra, and more recently was stressed by Rohan French in [16], in connection with the paradoxes of self-reference.

Now we need to clarify how q-matrix logics validate or invalidate inferences —following, e.g., [33, p. 196]. Notice that, below, $\vDash_{\mathcal{M}}$ is a substitution-invariant q-consequence relation, whence $(\mathbf{FOR}(\mathcal{L}), \vDash_{\mathcal{M}})$ is a q-logic. In addition to that, when a q-matrix $\mathcal{M}$ induce some q-logic $\mathbf{L}$, we may interchangeably refer to $\vDash_{\mathcal{M}}$ as $\vDash_{\mathbf{L}}$.

DEFINITION 2.16. *For $\mathcal{M}$ a q-matrix, an $\mathcal{M}$-valuation v satisfies a sequent or inference $\Gamma \vDash A$ (written $v \vDash_{\mathcal{M}} \Gamma \vDash A$) if and only if if $v[\Gamma] \cap \mathcal{D}^- = \emptyset$, then $v(A) \in \mathcal{D}^+$.*

For $\mathcal{M}$ a q-matrix, an inference $\Gamma \vDash \Delta$ is $\mathcal{M}$-valid (written $\vDash_{\mathcal{M}} \Gamma \vDash \Delta$) if and only if $v \vDash_{\mathcal{M}} \Gamma \vDash \Delta$, for all $\mathcal{M}$-valuations v.

The following is a more straightforward characterization of **TS**'s inferential validity:

$$\vDash_{\mathbf{TS}} \Gamma \vDash \Delta \text{ if and only if } \text{for every valuation } v, \begin{cases} \text{if } v(B) \in \{1, \frac{1}{2}\} & \text{for all } B \in \Gamma \\ \text{then } v(A) \in \{1\} & \text{for some } A \in \Delta. \end{cases}$$

From these definitions of **ST** and **TS**'s validities, the following facts follow.

FACT 2.17 ([8]). ***TS** is a nonreflexive, and thus a substructural, logic.*

FACT 2.18 ([8]). ***ST** is a nontransitive, and thus a substructural, logic.*

FACT 2.19 ([16]). ***TS** has no valid inferences or sequents.*

FACT 2.20 ([18], [27]). ***ST** and classical propositional logic **CL** have the same set of valid inferences or sequents.*

In the next section we will present a notion of consequence relation for *metainferences* based on both **ST** and **TS**.[9] Finally, we must stress that, though we have distinguished between *consequence relations* and *logics*, in the rest of the article we will use them interchangeably.

§3. TS/ST, a logic for classically valid metainferences. We will now present **TS/ST**, a logic that not only validates every classically valid *inference*—as **ST** does—, but also validates every classically valid *metainference* (of level 1).[10] This is possible because

[9] [3], [11], and [26] have shown that—through some suitable translation—the set of valid inferences in **LP** coincides with the set of valid meta-inferences in **ST**. Moreover, [16] have conjectured that —again, through some suitable translation—the set of valid inferences in **SK**, e.g., Strong Kleene logic, coincides with the set of valid meta-inferences in **TS**.

[10] As we already have mentioned, we will present single-conclusion *metainferential*logics, **TS/ST** being the first of them. Moreover, the consequence relation will be defined for inferences with a

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

TS/ST's consequence relation embraces a feature of the *inferential* consequence relation **ST**, but applies it to the metainferential level. Just to remember, an inference is valid in **ST** if and only if, for every valuation v, if the premises satisfy certain—demanding—standard, then the conclusion meet some *less demanding* code. In particular, an inference is valid in **ST** if and only if, if for every premise A, $v(A) = 1$—e.g., if v *strictly* satisfies every premise—, then for some conclusion B, $v(B) = 1, \frac{1}{2}$—e.g., v *tolerantly* satisfies a conclusion.[11] Conversely, if v *strictly* satisfies every premise, but does not *tolerantly* satisfies some conclusion, then v is a *counterexample* to the validity of that inference in **ST**.

We will adopt a similar norm for **TS/ST**'s validity notion, but for the metainferential level. Thus, **TS/ST**'s standard for the premises will be *more demanding* than its standard for the conclusion.

DEFINITION 3.1. *A metainference* $\Gamma_1 \vDash \Delta_1, \ldots, \Gamma_n \vDash \Delta_n \vDash_1 \Sigma \vDash \Pi$ *is valid in* ***TS/ST*** *if and only if, for every valuation v, if every* $\Gamma_i \vDash \Delta_i$ *is satisfied by v according to* ***TS****, then v satisfies* $\Sigma \vDash \Pi$ *in* ***ST****.*

Before proving our main result, it is worth noticing two facts. The first one relates **ST** and the semantic consequence relation of classical propositional logic, **CL**. The second one relates **TS** and **CL**. So now we should pause to clarify how—the propositional—**CL** behaves.

The valuations that define **CL**'s consequence relation are bivalent, exclusive and exhaustive: for every valuation v and every formula A, either $v(A) = 1$ or $v(A) = 0$, but not both. An inference $\Gamma \vDash \Delta$ is valid in **CL** if and only if, for every valuation v, either $v(\gamma) = 0$ (for some $\gamma \in \Gamma$), or $v(\delta) = 1$ (for some $\delta \in \Delta$). Similarly, a valuation v is a counterexample to $\Gamma \vDash \Delta$ in **CL** if and only if for every $\gamma \in \Gamma$, $v(\gamma) = 1$, and for every $\delta \in \Delta$, $v(\delta) = 0$.

We would like to stress the following fact: a valuation v is a counterexample in **ST** to an inference $\Gamma \vDash \Delta$ if and only if v is a counterexample to that inference in **CL**. Remember that v is a counterexample to $\Gamma \vDash \Delta$ in either of those two logics if and only if, for every $\gamma \in \Gamma$, $v(\gamma) = 1$, and for every $\delta \in \Delta$, $v(\delta) = 0$. Similarly, v satisfies an inference $\Gamma \vDash \Delta$ in **TS** if and only if v satisfies that inference in **CL**. Recall that v satisfies $\Gamma \vDash \Delta$ in either of those two logics if and only if either $v(\gamma) = 0$, for some $\gamma \in \Gamma$, or, for some $\delta \in \Delta$, $v(\delta) = 1$.

Now we can introduce the main result. It establishes that a metainference (of level 1) is valid in **CL** if and only if it is valid in **TS/ST**.[12]

THEOREM 3.2. *The First Collapse Result* (*For every metainference* $\Gamma_1 \vDash \Delta_1, \ldots, \Gamma_n \vDash \Delta_n \vDash_1 \Sigma \vDash \Pi$) $\Gamma_1 \vDash \Delta_1, \ldots, \Gamma_n \vDash \Delta_n \vDash_1 \Sigma \vDash \Pi$ *is valid in* ***CL*** *if and only if* $\Gamma_1 \vDash \Delta_1, \ldots, \Gamma_n \vDash \Delta_n \vDash_1 \Sigma \vDash \Pi$ *is valid in* ***TS/ST****.*

finite set of premises. Nevertheless, we hope that the results that we will introduce may be safely extended to multi-conclusion versions of these logics, that also admit arguments with infinite premises.

[11] Here, $v(B) = 1, \frac{1}{2}$ just means that $v(B) \in \{1, \frac{1}{2}\}$. For simplicity, a similar notation will be used in the rest of the paper.

[12] As both an anonymous referee and Thomas Ferguson (in private conversation) describes it, **TS** is providing a strict standard on premises of metainferences, while **ST** is playing the role of a tolerant criteria. We would like to thank the anonymous referee (and Thomas Ferguson) for this observation.

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

Proof. $\Gamma_1 \vDash \Delta_1, \ldots, \Gamma_n \vDash \Delta_n \vDash_1 \Sigma \vDash \Pi$ is *invalid* in **CL** if and only if for some v and every i ($1 \leq i \leq n$), if v satisfies every $\Gamma_i \vDash \Delta_i$ *according to* ***CL***, then v *does not* satisfies $\Sigma \vDash \Pi$ *in* ***CL***, if and only if, if for every $\Gamma_i \vDash \Delta_i$, $v(\gamma) = 0$ or $v(\delta) = 1$, for some $\gamma \in \Gamma_i$ or some $\delta \in \Delta_i$, then for every $\sigma \in \Sigma$ and every $\pi \in \Pi$, $v(\sigma) = 1$ and $v(\pi) = 0$, if and only if if v satisfies every $\Gamma_i \vDash \Delta_i$ *according to* ***TS***, then v *does not* satisfies $\Sigma \vDash \Pi$ *according to* ***ST***, if and only if $\Gamma_1 \vDash \Delta_1, \ldots, \Gamma_n \vDash \Delta_n \vDash_1 \Sigma \vDash \Pi$ is *invalid* in **TS/ST**. □

Another way to understand this result is the following: a metainference (of level 1) is invalid in **CL** if and only if there is a valuation v that satisfies every premise and does not satisfy the conclusion. But a valuation v satisfies a premise of a metainference of level 1—e.g., an inference—according to **CL** if and only if, either v gives value 0 to some premise of the inference, or gives value 1 to some conclusion. And that is precisely what it takes for v to satisfy that premise according to *TS*. Moreover, a valuation v does not satisfy the conclusion of the metainference according to *CL* if and only if v gives value 1 to every premise of the inference that is the conclusion of the metainference, and value 0 to every conclusion of the inference that is the conclusion of the metainference.[13] And that is just what it takes for v to be a counterexample to the validity of that inference in **ST**.

One last thing about metainferences before moving on. It is usual to consider sentences as *degenerate* cases of inferences (e.g., with an empty set of premises). In a similar fashion, inferences can be interpreted as *degenerate* cases of metainferences of level 1 (e.g., with an empty set of premises). Thus, every classically valid inference will be valid in **TS/ST**, and vice-versa. A degenerate case of a metainference of level 1—e.g., an inference—is valid in **TS/ST** if and only if every valuation satisfies the conclusion according to the standard for conclusions—e.g., according to **ST**. And **ST** recovers every classically valid inference.

3.1. A classical metainferential theory of truth. A truth predicate *Tr* is *transparent* if it satisfies the *transparency* condition, that states that, for every valuation v, every model $\mathcal{M}$ and every sentence A, $v_M(A) = v_M(Tr(\langle A \rangle))$. It is worth noticing that **TS/ST** can be safely expanded with a transparent truth predicate. We will called the resulting theory, **TS/ST+**.

As Kripke proves in [22], *Tr* can be understood as a fixed-point of a jump operator over the **SK**'s scheme. But the **SK**'s scheme is the one used to specify how **ST**, **TS**, and, moreover, **TS/ST**'s valuations work. Thus, **TS/ST+**'s truth predicate *Tr* can also be understood as a fixed-point, proving that the theory itself is nontrivial.

Of course, the details should be handled with care. As *Tr* is a predicate, valuations are no longer enough to interpret **TS/ST+**'s language. Models, then, should be included in the picture.

Let $\mathcal{L}^+$ be the result of expanding **TS/ST**'s language $\mathcal{L}$ with a predicate *Tr* and an infinite number of individual constants.[14] $\mathcal{L}^+$'s interpretation will be partially constrained to ensure that paradoxical sentences will be around.[15] In particular, some individual constants will be treated as *distinguished names*. Thus, we will fix a $1 - 1$ function τ from names

[13] Remember that even if metainferences are single-conclusion, the inference that is the conclusion of a metainference might be multi-conclusion.

[14] To keep things as simple as possible, **TS/ST**'s language will not include quantifiers nor variables for names.

[15] We follow here the strategy used in [28]. This way of handling self-reference is also similar to the one used in [5].

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

to formulas of $\mathcal{L}^+$, and require, for every *distinguished name n*, that n should denote $\tau(n)$ in every model. This means that set $FOR(\mathcal{L}^+)$ of formulas of $\mathcal{L}^+$ must be a subset of the domain of every model, and that only infinite models will be admitted.[16]

DEFINITION 3.3. *An **TS/ST**-model for the language $\mathcal{L}^+$ is a structure $< D, I >$ such that*:

- *D is a domain such that $FOR(\mathcal{L}^+) \subseteq D$, and*
- *I is an interpretation function such that*:
  - *For an ordinary name a, $I(a) \in D$*
  - *For a distinguished name n, $I(n) = \tau(n)$*
  - *For an n-ary predicate P, $I(P) \in \{0, \frac{1}{2}, 1\}^{D^n}$*
  - *For propositional letters p, $I(p) \in \{0, \frac{1}{2}, 1\}$*
  - *For atomic formulas (that are not propositional letters) $A = P(t_1, t_2, \ldots, t_n)$, $I(A) = I(P)(I(t_1), I(t_2), \ldots, I(t_n))$*
  - $I(\neg A) = 1 - I(A)$
  - $I(A \wedge B) = min(I(A), I(B))$
  - $I(A \vee B) = max(I(A), I(B))$
  - $I(A \vDash B) = max(1 - I(A), I(B))$.

A **TS/ST+**-model is a **TS/ST**-model that obeys the following restriction: for every pair of formulas A, $Tr(\langle A \rangle)$, $I(A) = I(Tr(\langle A \rangle))$. This last condition ensures the transparency of *Tr*.

With these notions at hand, it is possible to prove that **TS/ST+** is nontrivial—or, more accurately, satisfiable.

THEOREM 3.4. ***TS/ST+** is satisfiable.*

Proof. **TS/ST+**'s satisfiability follows from the fact that the truth predicate can be interpreted as a fixed-point of a jump operator over the *Strong Kleene scheme*, that is the ones used by **TS/ST**'s—and **TS/ST+**—models. In fact, **TS/ST**'s models are just simplifications of Ripley's **ST**'s models ([28].[17] And **ST**'s models are just the Strong Kleene logic **SK**'s (or **LP**'s) models expanded with distinguished names, following the lines of [22].

TS/ST+-models are the subset of **TS/ST**-models (for $\mathcal{L}^+$) that satisfy transparency. Moreover, **TS/ST+**'s models are only (simplified versions of) **ST**'s models for a first-order language with a transparent truth predicate. And, as Ripley shows in [28], these models are *transparent* not only because the truth predicate is transparent, but also because "no amount of adding *Ts* [e.g., truth predicates like *Tr*] or removing them can make a valid argument invalid, or vice versa" ([28], page 6). □

3.2. Why TS/ST is not fully classical. So far, we have shown that **TS/ST** recovers not only every classically valid inference, but also recaptures every classically valid metainference (of level 1). Nevertheless, **TS/ST** is not *fully* classical. Many classically valid *meta-metainferences*—e.g., metainferences of level 2—are invalid in **TS/ST**. In particular, the meta-metainference that we will call **Meta-Cut**, turns out to be invalid in **TS/ST**.

[16] This maneuver is standard in works on theories of truth. See, for example, [21] and [28].
[17] This is because $\mathcal{L}^+$ does not include quantifiers nor an identity predicate.

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

$$\text{Meta-Cut}\ \dfrac{\dfrac{\Gamma_1^1 \vDash \Delta_1^1, \ldots, \Gamma_j^1 \vDash \Delta_j^1}{\Sigma_1^1 \vDash \Pi_1^1, A} \qquad \dfrac{\Gamma_1^2 \vDash \Delta_1^2, \ldots, \Gamma_k^2 \vDash \Delta_k^2}{A, \Sigma_1^2 \vDash \Pi_1^2}}{\dfrac{\Gamma_1^1 \vDash \Delta_1^1, \ldots, \Gamma_j^1 \vDash \Delta_j^1, \Gamma_1^2 \vDash \Delta_1^2, \ldots, \Gamma_k^2 \vDash \Delta_k^2}{\Sigma_1^1, \Sigma_1^2 \vDash \Pi_1^1, \Pi_1^2}}$$

A metainferential consequence relation not only determines a validity standard for metainferences (and, furthermore, for inferences and sentences): it also fixes a standard for metainferences of higher levels. In particular, if C is a propositional metainferential consequence relation, a meta-metainference is valid in C if and only if every valuation v that satisfies every premise according to C, satisfies the conclusion in C. Moreover, for any metainference $\Gamma_1, \Gamma_j \vDash_n \Delta$ of any level n (where Δ and each Γ_i are metainferences of level $n-1$, and $\vDash_n$ stands for the metainferential relation of level n between premises and conclusion), $\Gamma_1, \Gamma_j \vDash_n \Delta$ is valid in C if and only if, for every valuation v, if v satisfies every premise Γ_i according to C, then v satisfies the conclusion Δ according to C.

It is not hard to realize why **Meta-Cut** is invalid in **TS/ST**. Take any instance of it that involves only propositional letters, where A is different from every other sentence in the premises and the conclusion. The valuation v such that $v(A) = \frac{1}{2}$, $v(\gamma) = v(\delta) = v(\sigma) = 1$—for every $\gamma \in \Gamma_i^j$, every $\delta \in \Delta_m^n$ and every $\sigma \in \Sigma_k^l$–, and $v(\pi) = 0$—for every $\pi \in \Pi_r^s$—, satisfies every premise according to **TS/ST**, but does not confirm the conclusion in **TS/ST**.

Nevertheless, it is possible to design a new consequence relation that recovers not only every classically valid inference—as **ST** and **TS/ST** do– and every classically valid metainference –as **TS/ST** does—, but also every classically valid meta-metainference. The next section will be devoted to present such logic.

§4. A consequence relation for meta-metainferences: $\mathbf{CM}_2$. To see how this new consequence relation works, we will need first to introduce a new metainferential consequence relation: **ST/TS**.

DEFINITION 4.1. *A metainference* $\Gamma_1 \vDash \Delta_1, \ldots, \Gamma_k \vDash \Delta_n \vDash_1 \Sigma \vDash \Pi$ *is valid in* ***ST/TS*** *if and only if, for every valuation v, if v satisfies every* $\Gamma_i \vDash \Delta_i$ *according to* ***ST****, then v confirms* $\Sigma \vDash \Pi$ *in* ***TS****.*

Now we are in a position to define the consequence relation that recovers every classically valid inference, metainference and meta-metainference—including *Meta-Cut*. We will call it $\mathbf{CM}_2$ (short for "consequence relation for metainferences of level 2").

DEFINITION 4.2. *A meta-metainference* $\Gamma_1, \ldots, \Gamma_n \vDash_2 \Delta$ *is valid in* $\boldsymbol{CM}_2$ *if and only if, for every valuation v, if every* Γ_i *is satisfied by v according to* ***ST/TS****, then v confirms* Δ *according to* ***TS/ST****.*

As we have already anticipated, the main result about $\mathbf{CM}_2$—or **STTS/TSST**, if one prefers a more informative name—says that a meta-metainference (or metainference of level 2) is valid in **CL** if and only if it is valid in $\mathbf{CM}_2$.

THEOREM 4.3. *The Second Collapse Result.*

For every meta-metainference $\Gamma_1, \ldots, \Gamma_n \vDash_2 \Delta$, $\Gamma_1, \ldots, \Gamma_n \vDash_2 \Delta$ *is valid in* ***CL*** *if and only if it is valid in* $\boldsymbol{CM}_2$.

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

Proof. For every valuation v, v is a counterexample to $\Gamma_1, \ldots, \Gamma_n \models_2 \Delta$ in **CL** if and only if it is a counterexample in $\mathbf{CM}_2$. A valuation is a counterexample in **CL** to the validity of a meta-metainference if and only if it satisfies every premise and does *not* satisfy the conclusion. But for every valuation v, v satisfies a metainference in **CL** if and only if v satisfies the metainference in **ST/TS**. And, moreover, v is a counterexample to a metainference in **CL** if and only if v satisfies it in **TS/ST**. Thus, the counterexamples for meta-metainferences in **CL** are exactly the same as the ones in $\mathbf{CM}_2$. Therefore, a meta-metainference is valid in **CL** if and only if it is valid in $\mathbf{CM}_2$. □

$\mathbf{CM}_2$ can also be expanded with a transparent truth predicate. The resulting theory is called $\mathbf{CM}_2$**+**, and is satisfiable.

THEOREM 4.4. ***CM***$_2$***+*** *is satisfiable.*

Proof. $\mathbf{CM}_2$**+**'s language is $\mathcal{L}^+$, e.g., **TS/ST+**'s language. Moreover, $\mathbf{CM}_2$'s models are the same as **TS/ST**'s models, and the same relation exist between $\mathbf{CM}_2$**+**'s and **TS/ST+**'s models. Thus, the proof that $\mathbf{CM}_2$**+** is satisfiable is the same as the one for **TS/ST+**. □

4.1. More problems. $\mathbf{CM}_2$'s recovers not only every classically valid meta-metainference, but also every classically valid metainference, inference and sentence. But metainferential consequence relations of level n not only determine a set of valid metainferences of level n, but a set of valid metainferences of level higher than n. In the case of $\mathbf{CM}_2$, for any metainference $\Gamma \models_n A$ of any level n higher than 2, $\Gamma \models_n A$ is valid in $\mathbf{CM}_2$ if and only if, for every valuation v, if v confirms every premise $\gamma \in \Gamma$ according to $\mathbf{CM}_2$, the v satisfies A in $\mathbf{CM}_2$. Thus, though $\mathbf{CM}_2$ recovers more classically valid meta-metainferences than **CM**, it is not fully classical either. Many classically valid *meta-meta-metainferences*—or *metainferences of level 3*—are invalid in $\mathbf{CM}_2$. In particular, what what we will call **Meta-Meta-Cut**—a classically valid meta-meta-metainference—turns out to be not valid. At this point, it might be useful to remind us the structure of **Meta-Cut**.

$$\text{Meta-Cut}\ \dfrac{\dfrac{\dfrac{\Gamma_1^1 \vDash \Delta_1^1, \ldots, \Gamma_j^1 \vDash \Delta_j^1}{\Sigma_1^1 \vDash \Pi_1^1, \mathrm{A}} \quad \dfrac{\Gamma_1^2 \vDash \Delta_1^2, \ldots, \Gamma_k^2 \vDash \Delta_k^2}{\mathrm{A}, \Sigma_1^2 \vDash \Pi_1^2}}{\Gamma_1^1 \vDash \Delta_1^1, \ldots, \Gamma_j^1 \vDash \Delta_j^1, \Gamma_1^2 \vDash \Delta_1^2, \ldots, \Gamma_k^2 \vDash \Delta_k^2}}{\Sigma_1^1, \Sigma_1^2 \vDash \Pi_1^1, \Pi_1^2}$$

Meta-Meta-Cut is a metainferential scheme of level 3 that has, as a conclusion, a metainference of level 2, that has, as a conclusion, **Meta-Cut**'s conclusion. The premises of this conclusion are a finite set of metainferences. Each premise has the following structure (for $1 \leq j\omega$, $1 \leq k\omega$):

$$PREMS\ \dfrac{\Theta_1^k \vDash \Xi_1^k, \ldots, \Theta_j^k \vDash \Xi_j^k}{\Phi^k \vDash \Psi^k}$$

Meta-Meta-Cut has two premises. The conclusion of each premise will be, on one hand, the first premise of **Meta-Cut**, and, on the other hand, the second **Meta-Cut**'s premise. **Meta-Meta-Cut**'s premises have the same set of premises as the premises of **Meta-Meta-Cut**'s conclusion.

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

This is how the **Meta-Meta-Cut** scheme looks like:

$$\dfrac{\dfrac{\dfrac{\textit{PREMS}}{\Gamma_1^1 \vDash \Delta_1^1, \ldots, \Gamma_j^1 \vDash \Delta_j^1}}{\Sigma_1^1 \vDash \Pi_1^1, A} \quad \dfrac{\dfrac{\textit{PREMS}}{\Gamma_1^2 \vDash \Delta_1^2, \ldots, \Gamma_k^2 \vDash \Delta_k^2}}{A, \Sigma_1^2 \vDash \Pi_1^2}}{\dfrac{\textit{PREMS}}{\Gamma_1^1 \vDash \Delta_1^1, \ldots, \Gamma_j^1 \vDash \Delta_j^1, \Gamma_1^2 \vDash \Delta_1^2, \ldots, \Gamma_k^2 \vDash \Delta_k^2}}{\Sigma_1^1, \Sigma_1^2 \vDash \Pi_1^1, \Pi_1^2}$$

Meta-Meta-Cut is an invalid $\mathbf{CM}_2$ meta-meta-metainference. The valuation v such that $v(A) = \frac{1}{2}$, $v(\gamma) = v(\delta) = v(\sigma) = 1$—for every $\gamma \in \Gamma_i^j$, every $\delta \in \delta_m^n$ and every $\sigma \in \Sigma_k^l$—, $v(\pi) = 0$—for every $\pi \in \Pi_r^s$—, and $v(\theta) = v(\xi) = v(\phi) = v(\phi) = 1$—for every $\theta \in \Theta_d^e$, every $\xi \in \Xi_d^e$, every $\phi \in \Phi^e$, and every $\psi \in \Psi^e$—satisfies every premise, but does *not* confirms its conclusion, both things according to $\mathbf{CM}_2$.

Now, it can be shown (though we will not try it here) that $\mathbf{CM}_2$ cannot recover every classically valid metainference of a level higher than 4. Nevertheless, it is possible to design a hierarchy of metainferential consequence relations n—one for each n ($1 \leq n < \omega$)—such that for every consequence relation for metainferences of level n, that consequence relation recovers every classically valid metainferences of level n or less. Thus, for any classically valid metainference of some level j, there is a consequence relation characterized by substructural means that recovers it. The next section will devoted to present in detail the construction of this hierarchy.

§5. A general result. Before presenting in detail the hierarchy, we would like to emphasize what we are going to show:

FACT 5.1. *General Result*: *For any level n, there is a consequence relation* $\boldsymbol{CM_n}$ *for metainferences of level n, such that every metainference of level n* $\Gamma_1, \ldots, \Gamma_k \vDash_n \Delta$, $\Gamma_1, \ldots, \Gamma_k \vDash_n \Delta$ *is valid in* **CL** *if and only if* $\Gamma_1, \ldots, \Gamma_k \vDash_n \Delta$ *is valid in* $\boldsymbol{CM_n}$.

To prove this general result, we must first define an operation $*$ on consequence relations.

DEFINITION 5.2. *For any metainferential consequence relation* $\boldsymbol{L_j/L_k}$, $\boldsymbol{(L_j/L_k)^* = L_k/L_j}$.

When applied to a metainferential consequence relation, the star operation $*$ gives another consequence relation. This new consequence relation switches the original consequence relation's standards for premises and conclusion of a sound argument (e.g., a metainference).

With the help of the star operation $*$, it is now possible to define the desired hierarchy of metainferential consequence relations.

DEFINITION 5.3. *The first step of the hierarchy is, by this time, well understood. They are the following*:

$\boldsymbol{CM_1 = TS/ST}$.

And for every n such that $2 \leq n < \omega$, $\boldsymbol{CM_n}$ *is the following consequence relation*:

Let $\Gamma_1, \ldots, \Gamma_k$ *and* Δ *be metainferences of level* $n - 1$. *A metainference of level n* $\Gamma_1, \ldots, \Gamma_k \vDash_n \Delta$ *is valid in* $\boldsymbol{CM_n}$ *if and only if, for every valuation v, if every* $\Gamma_1, \ldots, \Gamma_k$ *is satisfied by v according to* $(\boldsymbol{CM_{n-1}})^*$, *then v satisfies* Δ *according to* $\boldsymbol{CM_{n-1}}$.

Thus, for every consequence relation $\boldsymbol{CM_n}$, the following will be true:

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

FACT 5.4. *Third Collapse Result.*
For every level n ($1 \leq n < \omega$)*, a metainference of level n* $\Gamma_1, \ldots, \Gamma_k \models_n \Delta$ *is valid in* ***CL*** *if and only if it is valid in* $\mathbf{CM_n}$.

Proof. We will, in fact, a prove a more general claim. For every valuation v, v is a counterexample to $\Gamma_1, \ldots, \Gamma_k \models_n \Delta$ in **CL** if and only if it is a counterexample in $\mathbf{CM_n}$, and v satisfies $\Gamma_1, \ldots, \Gamma_k \models_n \Delta$ in **CL** if and only if v satisfies $\Gamma_1, \ldots, \Gamma_k \models_n \Delta$ in $(\mathbf{CM_n})^*$.

The proof is an induction on the levels of metainferences.

Base cases. We have already proven the theorem for $\mathbf{CM}_1$ and $(\mathbf{CM}_1)^*$—e.g., **TS/ST** and **ST/TS**.

Inductive case. We will need to prove the theorem for $\mathbf{CM_n}$ and $\mathbf{CM_n^*}$, for any $n > 1$.

By inductive hypothesis, for every valuation v, v satisfies a metainference of level $n-1$ ($n-2$) in **CL** if and only if v satisfies it in $(\mathbf{CM_{n-1}})^*$, and v is a counterexample to a metainference of level $n-1$ ($n-2$) in **CL** if and only if v disconfirms it in $\mathbf{CM_{n-1}}$.

Now, a valuation v satisfies a metainference of level n $\Gamma_1, \ldots, \Gamma_k \models_n \Delta$ in **CL**, if and only if v does not satisfy a premise or if it satisfies the conclusion—both according to **CL**—, if and only if there is a Γ_i such that v does not satisfy Γ_i or v satisfy Δ in **CL**, if and only if (*by inductive hypothesis*) there is a Γ_i such that v does not satisfy Γ_i in $\mathbf{CM_{n-1}}$ or v satisfies Δ in $(\mathbf{CM_{n-1}})^*$, if and only if v satisfies $\Gamma_1, \ldots, \Gamma_k \models_n \Delta$ in $(\mathbf{CM_n})^*$.

Moreover, a valuation v is a counterexample to a metainference of level n $\Gamma_1, \ldots, \Gamma_k \models_n \Delta$ in **CL** if and only if v satisfies every premise and does not confirms its conclusion—both according to **CL**—, if and only if every Γ_i is such that v satisfies Γ_i and v does not confirm Δ in **CL**, if and only if (*by inductive hypothesis*) every Γ_i is such that v satisfies Γ_i in $(\mathbf{CM_{n-1}})^*$ and v does not confirms Δ in $(\mathbf{CM_{n-1}})$, if and only if v is a counterexample to $\Gamma_1, \ldots, \Gamma_k \models_n \Delta$ in $\mathbf{CM_n}$. □

Before moving on, we would like to stress a general fact, that explains why each step in the hierarchy of logics $\mathbf{CM_n}$ recovers increasingly more metainferences. We have already mentioned that what it takes for a valuation to be a counterexample in **ST** to an inference $\Gamma \models \Delta$ is exactly what it takes for it to be counterexample to that inference in **CL**, and that what it takes for it to satisfy an inference $\Gamma \models \Delta$ in **TS** is exactly what it takes for it to satisfy that inference in **CL**. Similarly, for any level n, what it takes for a valuation to be a counterexample in $\mathbf{CM_n}$ to a metainference of level n is exactly what it takes for it to be counterexample to that inference in **CL**. And, moreover, for any level n, what it takes for a valuation to be counterexample to an inference in **CL**, and that what it takes for a valuation to satisfy a metainference of level n in * is exactly what it takes for it to satisfy that inference in **CL**.

5.1. Still more problems. Nevertheless, none of the consequences relations $\mathbf{CM_n}$ is *fully* classical. In fact, every $\mathbf{CM_n}$ invalidates a version of *Meta-Cut* of level $n-1$.

THEOREM 5.5. *For every n* ($1 \leq n < \omega$)*,* ***CM****n invalidates Meta-Cut*$_n$.

Proof. The proof is inductive. We have already present the structure of both *Meta-Cut* and *Meta-Meta-Cut*, and show that *Meta-Cut* is invalid in $\mathbf{CM_i}$—e.g., **TS/ST**. Furthermore, it is possible to specify the structure of each *Meta-Cut* of level j ($2 \leq j < \omega$) constructively. Each *Meta-Cut*$_j$—e.g., *Meta-Cut* of level j—is a two-premise metainference of level $j+1$. Each premise is a metainference of level j that has, as its conclusion, one of the premises of an instance of *Meta-Cut* of level $j-1$. The conclusion of *Meta-Cut* of level j is also a metainference of level j that has, as a conclusion, the conclusion of an instance of

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

Meta-Cut of level $j-1$. But the two premises and the conclusion of the particular version of *Meta-Cut*$_j$ that we will be talking about, has the same (and single) premise. For the sake of simplicity, it will be a metainference of level $j-1$ such that every propositional letter in it does not appear in the instance of the metainference *Meta-Cut*$_{j-1}$ used to design this instance of *Meta-Cut*$_j$. Moreover, every inference that occurs in it will be an instance of $p \vDash q$. Here, p and q are propositional letters that appear only once in the metainference. Thus, the metainference that is the premise of both premises, and also the premise of the conclusion of *Meta-Cut*$_j$, can be considered a *contingent* metainference, in the sense that some valuations will satisfy it, and some others will not.

Now we can present a counterexample to *Meta-Cut*$_j$. It will be a valuation j such that, for every p, q in the premise of both premises of *Meta-Cut*$_j$—e.g., two metainferences of level j—, and also in the premise of the conclusion, $v(p) = v(q) = 1$. The rest of the formulas—e.g., the conclusion of each premise, and also the conclusion of the conclusion—will be part of *Meta-Cut*$_{j-1}$. By inductive hypothesis, there is a valuation $v*$ that is a counterexample to *Meta-Cut*$_{j-1}$. For all those formula B in *Meta-Cut*$_{j-1}$—e.g., the conclusion of each premise, and also the conclusion of the conclusion—, let $v(B) = v*(B)$. In particular, let the *meta-cut-formula* A be such that $v(A) = v*(A) = \frac{1}{2}$. The valuation v satisfies every premise and does not confirm the conclusion of *Meta-Cut*$_j$ in $\mathbf{CM_j}$. Remember that *Meta-Cut*$_j$ is a metainference of level $j+1$. A valuation v does not satisfy an instance of *Meta-Cut*$_j$ in $\mathbf{CM_j}$ if and only if it satisfies every premise and does not satisfy the conclusion, both according to $\mathbf{CM_j}$. The conclusion and each premise are metainferences of level j. Thus, v satisfies any of them in $\mathbf{CM_j}$ if and if either if v does not satisfy some premise according to $\mathbf{CM^*_{j\text{-}1}}$, or it confirms the conclusion according to $\mathbf{CM_{j-1}}$. As $v(B) = v*(B)$, for every B in *Meta-Cut*$_{j-1}$, and $v*$ is a counterexample to it, then $v*$—and thus v—satisfies every *Meta-Cut*$_{j-1}$'s premise according to $\mathbf{CM^*_j}$. Therefore, it satisfies the conclusion of every premise of *Meta-Cut*$_j$, because it satisfies each conclusion.[18] Moreover, v does not satisfy *Meta-Cut*$_j$'s conclusion. Remember that for every p, q in the premise of the conclusion, $v(p) = v(q) = 1$. Then, the premise of the conclusion is satisfied by v according to $\mathbf{CM^*_j}$. As $v(B) = v*(B)$, for every B in *Meta-Cut*$_{j-1}$, and $v*$ is a counterexample to it, then v will not satisfy *Meta-Cut*$_{j-1}$'s either. But *Meta-Cut*$_{j-1}$'s conclusion is the conclusion of the conclusion of *Meta-Cut*$_j$. Then, v will satisfy each premise of *Meta-Cut*$_j$, but will not confirm its conclusion. Therefore, v is a counterexample to it. □

Thus, though no logic in the hierarchy recovers every classically valid metainference, each metainference is recaptured by some logic $\mathbf{CM_n}$. Nevertheless, no logic $\mathbf{CM_n}$ recaptures *every* classically valid metainference. In the next section we will present a logic with this desired feature. To define it, we will use the hierarchy that we have presented.

§6. A fully classical logic characterized by substructural means. The goal of this section is to define a consequence relation that recovers every classically valid metainference. We will call it $\mathbf{CM_\omega}$, and it could be understood as "the union of every $\mathbf{CM_n}$".

DEFINITION 6.1. *A meta-inference of level n—for any level n—*$\Gamma_1, \ldots, \Gamma_k \vDash_n \Delta$ *is valid in* $\boldsymbol{CM_\omega}$*, if and only if* $\Gamma_1, \ldots, \Gamma_k \vDash_n \Delta$ *is valid in some* $\boldsymbol{CM_j}$ $(1 \leq j < \omega)$.

[18] For every metainference of every level j, the set of counterexamples to it in $\mathbf{CM^*_j}$ is included in the set of counterexamples to it in $\mathbf{CM_j}$, as it is not hard to see. We leave the proof as an exercise to the reader.

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

FACT 6.2. *The General Recovery Result.*
For every level n ($1 \leq n < \omega$)*, a metainference of level n* $\Gamma_1, \ldots, \Gamma_k \vDash_n \Delta$ *is valid in* ***CL*** *if and only if it is valid in* $\boldsymbol{CM_\omega}$.

Proof. We know, by the *Third Collapse Result*, that for every level n ($1 \leq n < \omega$), a metainference of level n $\Gamma_1, \ldots, \Gamma_k \vDash_n \Delta$ is valid in **CL** if and only if it is valid in $\mathbf{CM_n}$. But if it is valid in $\mathbf{CM_n}$, then, by the definition of $\mathbf{CM_\omega}$, it is valid in $\mathbf{CM_\omega}$. Therefore, every classically valid metainference of any level n will be valid in $\mathbf{CM_\omega}$. □

What is, then, the exact relationship between **CL** and $\mathbf{CM_\omega}$? Dave Ripley (in [28], [29]) maintains that **ST** is just classical logic. His point is that **ST** recovers every classically valid inference. Others, like [4] and [11], argue against Ripley's claim. They claim that there are some central features of **CL** that **ST** cannot recover. Those features are some classically valid metainferences, mainly of level 1. But **ST** loses classically valid metainferences in every level. **CM**—e.g., **TS/ST**—recovers every classically valid metainferences, but still loses some classically valid meta-metainferences, like *Meta-Cut*. $\mathbf{CM_2}$ recovers them, but still cannot recover every classically valid meta-meta-metainference. Moreover, for every n, though $\mathbf{CM_n}$ recovers every classically valid metainference of level n, it cannot recover every classically valid metainference of higher levels. In this regard, $\mathbf{CM_\omega}$ works better than every $\mathbf{CM_n}$. $\mathbf{CM_\omega}$ recovers every classically valid metainference of every level. If a classical logic is defined through the inferences and metainferences of every level that turns out valid in it, then $\mathbf{CM_\omega}$ is a *fully* classical logic. The only special thing about it is that it is defined through substructural means.

One more attractive feature of $\mathbf{CM_\omega}$ is that not only it is a *fully* classical logic, but it also supports a transparent truth predicate. We will call the theory obtained from $\mathbf{CM_\omega}$ by expanding the language with a transparent truth predicate *Tr*, $\mathbf{CM_\omega}+$. The following is an important result about this theory:

THEOREM 6.3. $\boldsymbol{CM_\omega}$**+** *is satisfiable.*

Proof. $\mathbf{CM_\omega}$**+**'s satisfiability proof runs just as the corresponding proof for **TS/ST+**. Notice that the language and the models of both theories are exactly the same. □

Therefore, $\mathbf{CM_\omega}$**+** is the truth theory we were looking for. It not only validates every classically valid metainference of any level, but also supports a transparent truth predicate.

§7. Objections and replies. Before ending, we would like to address two objections that have been raised to us by an anonymous referee. The first one is the following. She has suggested that **TS/ST** does not actually recovers Cut, because what we have called *Cut* does not give the correct metainferential rendering of the unrestricted Cut rule. As far as we know, the discussion around Cut as always been focuses in rule-versions (additive or multiplicative, or even mixed-forms) of the metainferential schema that we have called *Cut*. This is exactly how Cut is presented, for example, in [28] and [29]. Surely, the different metainferential schemes that are instances of what we have called $Meta - Cut_n$ are related to Cut in some way. Nevertheless, we are not inclined to claim that they are, as the referee suggest, the same rule—or instances of a possible unrestricted Cut rule. Each rule -or, more precisely, each instance of $Meta - Cut_n$- has its own unique level. We can't tell how we are meant to identify rules (e.g., metainferences) across levels. In any case, this is not an easy task. But until that issue is settled, we cannot justifiably claim that there is such an "unrestricted Cut rule". In any case, if, as the referee claims, *Meta-Cut* (e.g., $Meta - Cut_1$)

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

is the rule that truly express the unrestricted Cut rule, then, as $\mathbf{CM}_2$ validates it, $\mathbf{CM}_2$ truly recovers Cut. Nevertheless, certainly the referee will claim that the metainference that really expresses Cut (or, as she claims, gives the correct metainferential rendering of the unrestricted Cut rule) in $\mathbf{CM}_2$ is $Meta - Cut_2$. And, more generally, as no logic $\mathbf{CM_n}$ validates $Meta - Cut_n$, then no logic in the hierarchy really recovers Cut. But notice that none of these objections can be raised against $\mathbf{CM}_\omega$. As $\mathbf{CM}_\omega$ validates any metainference that is valid in at least one of the logics $\mathbf{CM_n}$ in the hierarchy, then, as every metainferential version of Cut is recovered at some point in the hierarchy, $\mathbf{CM}_\omega$ itself validates every single one of them. Moreover, as it was proved, $\mathbf{CM}_\omega$ can be safely expanded with a transparent truth predicate. The resulting theory, $\mathbf{CM}_\omega\mathbf{+}$, is satisfiable. But this result has nothing to do with the nontransitivite of $\mathbf{CM}_\omega\mathbf{+}$. In fact, $\mathbf{CM}_\omega\mathbf{+}$ is as transitive as $\mathbf{CM}_\omega$, because both of them validate every metainferential version of Cut, including (what we have called) Cut itself. The main reason that explains $\mathbf{CM}_\omega\mathbf{+}$'s satisfiability is that it is defined through *SK-models*, and its logical constants are the usual ones. In particular, no constant like a strong negation (e.g., a strong negation, in this context, might be understood as a unary constant $\sim$ such that for every formula A and valuation v, $v(\sim A) = 0$ if and only if $v(A) = 1$, and in any other case, $v(\sim A) = 1$).[19] Every paradoxical sentence that can be expressed with the vocabulary of the theory, like the Liar or any Curry-sentence, might (or will) receive the value $\frac{1}{2}$ in every valuation.

The second objection points directly to $\mathbf{CM}_\omega$. It claims that it is not correct to say that $\mathbf{CM}_\omega$ is fully classical just because it not only recovers every classical inference, but also every metainference of any level. If $\mathbf{CM}_\omega$ were a real fully classical logic, then it would behave just like any presentation of classical logic. For example, it would not be satisfiable when a transparent truth predicate is added to it. But, as any standard presentation of classical logic, such as a Hilbert-style axiomatic calculus for classical logic, cannot be expanded with a transparent truth predicate, this shows that classical logic and $\mathbf{CM}_\omega$ are two different things.

This point is contentious, for at least three reasons. The first one is the following. Some authors, such as Ripley and his co-authors in [28], defend that different presentations of the same logics may behave differently with respect to the same language. Then, it is perfectly possible that, for example, a two-valued presentation of classical logic—e.g., **CL**—cannot be safely expanded with a transparent truth predicate, while that kind of predicate can be added to a three-valued presentation of it—e.g., **ST**—without the resulting theory becoming trivial. We do not endorse this position. Nevertheless, we do think that it is an open debate whether or not a logic is sensible to its different presentation—e.g., in a sequent calculus, in a natural deduction system, in an axiomatic one, in a two-valued semantics, in a three-valued semantics, etc.

The second reason is that in order to establish whether or not two logic are in fact the same, a criteria to distinguish between logics is required. Here, we have implicitly assume an *extensional* criteria. This norm proposes the following: for any two logics L_1 and L_2, L_1 is the same logic as L_2 if and only if the validates exactly the same sentences, inferences and metainferences of level n, for any level n. If this is the right way to distinguish between logics, then it is correct to claim that **CL** and $\mathbf{CM}_\omega$ are the same logic. But there are other criteria that have been proposed. For example, Dicher and Paoli, in [12], claim that a logic is a similarity class of abstract consequence relations. And according to that criteria, for example, **ST** and **CL** are two different logics—in fact, **ST** is just **LP**. If Dicher and Paoli's

[19] For more about strong or *classical* negations, see [10].

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

criteria is the right one, then every logic in the hierarchy is also equivalent to **LP**. We are not sure whether or not $\mathbf{CM}_\omega$ and **CL** belong the the same similarity class of abstract consequence relation. In any case, whether or not Dicher and Paoli's criteria is the right way to distinguish between logics, or whether the *extensional* criteria is the right one, or if none of them get things right, is a complex issue. Thus, what we want to defend is the following, weaker version of our original claim. If the *extensional* criteria is the right way to distinguish logics, then it is certainly true that **CL** and $\mathbf{CM}_\omega$ are the same logic—and it would be right to claim that $\mathbf{CM}_\omega$ is *fully classical*. But if the antecedent of this conditional is false, then **CL** and $\mathbf{CM}_\omega$ may be different logics. We think that Dicher and Paoli's is not a good way to identify logics for many reasons. Just to mention one: according to them, there cannot be substructural logics, because every abstract consequence relation needs to be reflexive, monotone and transitive. Thus, if one thinks that substructural logics *are* logics, then one needs to reject Dicher and Paoli's criteria.

Finally, even if it is true that $\mathbf{CM}_\omega$ is not equivalent to **CL**, it still can be true that $\mathbf{CM}_\omega$ is fully classical. The point, here, is what characterized a logic as classical. If this means to be equivalent to **CL**, then $\mathbf{CM}_\omega$ might not be classical—at least, for example, if two logics cannot be equivalent if they behave differently when they are expanded with the say vocabulary. But if what defines a logic as classical is to validate the same set of metainferences (of every level) as **CL**, then $\mathbf{CM}_\omega$ is certainly classical—even if it turns out to be not equivalent to **CL**.

§8. Conclusion. We have present a three-valued consequence relation for *metainferences*, called **TS/ST**, or (latter) **CM**, defined with the help of the substructural logics **ST** and **TS**. Even though every classically valid inference is valid in **ST**, some metainferences—most importantly, Cut—are lost. **CM** not only recovers every classically valid inference, but also every classically valid metainference. Moreover, it is possible to expand **CM** with a transparent truth predicate. Thus, it seems to work better than a truth theory based on **ST**, in the sense that—in Hjortland's words—it "mutilates" less classical logic in the process. Nevertheless, **CM** is not entirely classical either. Some classically valid meta-metainferences—most importantly, *Meta-Cut*—are lost. But the meta-metainferential consequence relation that we have called $\mathbf{CM}_2$ can recapture every classically valid metainference. Despite everything it recovers, $\mathbf{CM}_2$ still cannot recapture every classically valid meta-meta-metainference. Nevertheless, the consequence relation $\mathbf{CM}_3$ can recover every classically valid metainference. In fact, these three logics are the initial steps of an infinite hierarchy of logics, the $\mathbf{CM_n}$ logics. Each $\mathbf{CM_n}$ specifies a consequence relation for metainferences of level n, that recovers every classically valid metainference of level n or less, but still cannot recapture every classically valid metainferences of higher levels. Nevertheless, the logic $\mathbf{CM}_\omega$ do recover every classically valid metainference. Therefore, $\mathbf{CM}_\omega$ is a *fully classical* logic, in the sense that every classically valid inference or metainference of any level is valid according to it. Moreover, $\mathbf{CM}_\omega$ can also be safely expanded with a transparent truth predicate. Thus, it seems like a very attractive truth theory.

Nevertheless, there still is plenty work to do in relation to these logics and truth theories. For example, it seems not easy to imagine a proof theory for them. One option, somewhat in the line of the sequent calculi developed in [25], is to present a meta-sequent calculi for **TS/ST+**.[20] Paradoxes might be avoided by given up *Meta-Cut*. Presumably, if that

[20] We are aware that Rohan French is currently working on a similar sequent calculi for metainferences, but for nonreflexive theories.

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

approach works, it can be adopted for theories based on logics and truth theories that are higher in the hierarchy. Nevertheless, that cannot be what a sequent calculi for $\mathbf{CM}_{\omega}$ with a transparent truth predicate might look like. Notice that that proof theory should make valid every classically valid metainference (e.g., metarule) of any level, while avoiding paradoxes. We leave the exploration of these routes for future work.

§9. Acknowledgments. I would like to thank the members of the Buenos Aires Logic Group, Dave Ripley, Rohan French, Bogdan Dicher, Andreas Fjellstad, and Thomas Ferguson for helpful feedback concerning previous versions of this paper. Thanks also go to anonymous reviewers for their comments and suggestions to improve the paper. Finally, I would like to thank the National Scientific and Technical Research Council of Argentina (CONICET) for the economic support.

BIBLIOGRAPHY

[1] Avron, A. (1988). The semantics and proof theory of linear logic. *Theoretical Computer Science*, **57**(2–3), 161–184.

[2] Avron, A. (1991). Simple consequence relations. *Information and Computation*, **92**(1), 105–139.

[3] Barrio, E., Rosenblatt, L., & Tajer, D. (2015). The logics of strict-tolerant logic. *Journal of Philosophical Logic*, **44**(5), 551–571.

[4] Barrio, E., Rosenblatt, L., & Tajer, D. (2015). The logics of strict-tolerant logic. *Journal of Philosophical Logic*, **44**(5), 551–571.

[5] Barwise, J. & Etchemendy, J. (1987). *The Liar*: *An Essay on Truth and Circularity*. Oxford, UK: Oxford University Press.

[6] Blok, W. & Jónsson, B. (2006). Equivalence of consequence operations. *Studia Logica*, **83**(1), 91–110.

[7] Chemla, E., Egré, P., & Spector, B. (2017). Characterizing logical consequence in many-valued logics. *Journal of Logic and Computation*, **27**(7), 2193–2226.

[8] Cobreros, P., Egré, P., Ripley, D., & van Rooij, R. (2012). Tolerant, classical, strict. *Journal of Philosophical Logic*, **41**(2), 347–385.

[9] Cobreros, P., Egré, P., Ripley, D., & Van Rooij, R. (2014). Reaching transparent truth. *Mind*, **122**(488), 841–866.

[10] De, M. & Omori, H. (2015). Classical negation and expansions of belnap–dunn logic. *Studia Logica*, **103**(4), 825–851.

[11] Dicher, B. & Paoli, F. (2017). St, lp, and tolerant metainferences. Unpublished manuscript.

[12] Dicher, B. & Paoli, F. (2018). The original sin of proof-theoretic semantics. Unpublished manuscript.

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

[14] Frankowski, S. (2004). Formalization of a plausible inference. *Bulletin of the Section of Logic*, **33**(1), 41–52.

[15] Frankowski, S. (2004). P-consequence versus q-consequence operations. *Bulletin of the Section of Logic*, **33**(4), 197–207.

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

[17] Gabbay, D. (1996). *Labelled Deductive Systems*. Oxford: Oxford University Press.

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press

[18] Girard, J.-Y. (1987). *Proof Theory and Logical Complexity*. Napoli: Bibliopolis.

[19] Hjortland, O. Theories of truth and the maxim of minimal mutilation. Manuscript.

[20] Humberstone, L. (1996). Valuational semantics of rule derivability. *Journal of Philosophical Logic*, **25**, 451–461.

[21] Kremer, M. (1988). Kripke and the logic of truth. *Journal of Philosophical Logic*, **17**, 225–278.

[22] Kripke, S. (1975). Outline of a theory of truth. *Journal of Philosophy*, **72**(19), 690–716.

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

[24] Malinowski, G. (2014). Kleene logic and inference. *Bulletin of the Section of Logic*, **43**(1/2), 43–52.

[25] Priest, G. & Wansing, H. (2015). External curries. *Journal of Philosophical logic*, **44**(4), 453–471.

[26] Pynko, A. (2010). Gentzen's cut-free calculus versus the logic of paradox. *Bulletin of the Section of Logic*, **39**(1/2), 35–42.

[27] Ripley, D. (2012). Conservatively extending classical logic with transparent truth. *The Review of Symbolic Logic*, **5**(2), 354–378.

[28] Ripley, D. (2012). Conservatively extending classical logic with transparent truth. *Review of Symbolic Logic*, **5**(2), 354–378.

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

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

[31] Shoesmith, D. & Smiley, T. (1978). *Multiple-Conclusion Logic*. Cambridge: Cambridge University Press.

[32] Shramko, Y. & Wansing, H. (2010). Truth values. In Zalta, E., editor. *The Stanford Encyclopedia of Philosophy*. Stanford Uniersity, Summer 2010 edition. Available at: `http://plato.stanford.edu/archives/sum2010/entries/truth-values/`.

[33] Shramko, Y. & Wansing, H. (2011). *Truth and Falsehood*: *An Inquiry into Generalized Logical Values*. Springer Science & Business Media.

NATIONAL SCIENTIFIC AND TECHNICAL RESEARCH COUNCIL
SARMIENTO 440
C104AA7 BUENOS AIRES, ARGENTINA
and
DEPARTMENT OF PHILOSOPHY
UNIVERSITY OF BUENOS AIRES
PUÁN 480
C1406CQJ BUENOS AIRES, ARGENTINA
E-mail: federico.pailos@gmail.com

https://doi.org/10.1017/S1755020318000485 Published online by Cambridge University Press