

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Bicontextualism

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1889221> since 2023-12-11T16:06:47Z

Published version:

DOI:10.1215/00294527-2023-0004

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Bicontextualism

Lorenzo Rossi

Abstract Can one quantify over absolutely everything? Absolutists answer positively, while relativists answer negatively. Here, I focus on the absolutism vs. relativism debate in the framework of theories of truth, where relativism becomes a form of *contextualism* about truth predications. Contextualist theories of truth provide elegant and uniform solutions to the semantic paradoxes while preserving classical logic. However, they interpret harmless generalizations (such as «everything is self-identical») in less than absolutely comprehensive domains, thus systematically misconstruing them. In this paper, I show that contextualism is broadly compatible with absolute generality. More specifically, I develop a *bipartite* contextualist semantics, or «bicontextualism», on which sentences are split in two groups: the *unproblematic* sentences, which are compatible with absolute generality, and the *problematic* ones, which are given a relativist semantics. I then argue that bicontextualism retains the advantages of (orthodox) contextualism, and does not give rise to new revenge paradoxes.

1 Introduction

Can one quantify over absolutely everything? *Generality absolutists* (henceforth: «absolutists») answer positively, while *generality relativists* (henceforth: «relativists») answer negatively.¹ Relativists typically motivate their rejection of absolute generality via an *argument from paradox*: given any alleged maximal domain of quantification D , some reasoning typically along the lines of set-theoretic, property-theoretic, or truth-theoretic paradoxes shows that something is not in D —hence D does not contain everything. But D was arbitrary, and the argument fully general: therefore, no domain D can contain absolutely everything.

A rich tradition in the philosophy of mathematics, going back at least to Ernst Zermelo [89], takes the arguments from paradoxes to afford a conclusive lesson on how to interpret the talk of domains and quantifiers.² In a nutshell, the *Zermellian* view envisages a never-ending succession of set-theoretic models to interpret domains of quantification. Importantly, such a succession never provides an «ultimate» model.

2010 Mathematics Subject Classification: Primary 03A05, 03B16; Secondary 03B65
Keywords: quantifiers, absolute generality, truth, paradoxes

As James Studd [83, p. 53] emphasizes, one can interpret the Zermellian picture as a form of relativism: quantification is always relative to an always expandable domain. The Zermellian picture is naturally opposed to a *Cantorian* view,³ on which, in addition to set-theoretical models, there is an absolute, non-set-sized domain, which includes absolutely everything, and which is beyond the reach of any set-theoretical interpretation. It is, in other words, the ultimate, all-encompassing domain.

The absolutism vs. relativism debate is not confined to set theory. In the framework of *theories of truth*—the focus of this paper—generality relativism morphs into a form of *contextualism about truth predications* (henceforth: «contextualism»), and the argument from paradox typically employs a contextualist version of the Liar Paradox. Here's a very brief outline.⁴ Consider a sentence λ equivalent to « $\ulcorner \lambda \urcorner$ is not true», where $\ulcorner \lambda \urcorner$ is a name of the sentence λ . The Liar reasoning seemingly establishes that λ is both true and untrue. In order to avoid the contradiction, contextualists postulate a covert *contextual* element, which changes or shifts in the course of the Liar reasoning. In this way, one establishes that λ is not true in the original context, but it is true in another context (which is perfectly consistent). However, in order for the context shift to be possible, the interpretation provided in the first context could not have been maximal: something must have not been available in the first context that becomes available in the second (and makes λ true in it). And since one can always run a Liar reasoning, the argument concludes, no context is absolutely comprehensive.^{5,6}

Contextualist theories of truth have a number of virtues: they provide elegant and uniform solutions to the semantic paradoxes, and preserve full classical logic. However, they reject absolute generality, and such rejection is not to be taken lightly. By the relativist's lights, even generalizations such as «everything is self-identical» and «everything that is possible is necessarily possible» can express propositions only in less than absolutely comprehensive domains. But this is a misconstrual: in scientific and metaphysical theorizing, such generalizations are to be read unrestrictedly (Williamson [85]). Or, at least, such a reading is clearly possible. What is more, such a misconstrual is unnecessary: interpreting the above generalizations unrestrictedly doesn't give rise to any paradox. To be sure, absolutism can recover such maximally general interpretations, but cannot provide the elegant and appealing solution to paradoxes that relativism has to offer. We have an inevitable trade-off, or so it seems.

The purpose of this paper is to show that this trade-off is illusory. There is no need to choose between an appealing solution to the paradoxes and an unrestricted interpretation of harmless generalizations: one can simply have both. This is achieved by curtailing absolute generality only when it is strictly necessary. More precisely, absolute generality needs to be restricted *only when* interpreting sentences such as λ . Sentences such as «everything is self-identical» can harmlessly be interpreted unrestrictedly. The theory developed here is therefore a *bipartite* contextualist semantics (henceforth: «bicontextualism»), in which the «unproblematic» sentences are given an absolutist semantics, while the remaining ones are given a relativist semantics.⁷

Finally, a few words on the scope of this paper. First, the purpose of this paper is *not* to adjudicate the absolutism vs. relativism debate *tout court*. This work assumes that, in their truth-theoretical clothes, both relativism and absolutism have appealing traits, and presents a theory that combines them. However, the intransigent absolutist who denies that relativism has *any* appeal is not going to be interested in the view on

offer, simply because she doesn't accept half of its motivation. *Ditto* for the intransigent relativist. However, such extreme positions are too strict to be taken into serious consideration. Second, this paper is mostly concerned with the *truth-theoretic* incarnation of the absolutism vs. relativism debate. Versions of this debate appear in many different areas (including set theory, higher-order logics, natural language semantics, the metaphysics of properties, and theories of truth), so it would be overambitious to think that a single proposal can simultaneously address it in all these areas. That being said, I think that bicontextualism has some interest beyond theories of truth. For one thing, theories of truth are (also) developed with the objective of providing truth-conditions, and thus theories of meaning, for (fragments of) natural languages.⁸ For another, bicontextualism employs both set-theoretical and higher-order resources, in both relativist and absolutist fashions. So, while bicontextualism is a theory of truth, it is closely connected to the absolutism vs. relativism debate in neighboring areas, and the paper includes a brief discussion of its implications for natural language semantics, set theory, and higher-order languages (§5).

The remainder of the paper is structured as follows. In §2, I present «orthodox» (that is, fully relativist) contextualism. In §3, I develop bicontextualism, and in §4 I explore its main philosophical implications. §5 addresses some objections, and §6 concludes.

2 Orthodox contextualism

Here I outline the contextualist approach to truth and paradox originally conceived by Charles Parsons [63] and substantially expanded by Michael Glanzberg [24, 25, 26, 27].⁹

Let me begin with a clarification (and a simplification). Contextualists often work under the assumption that propositions are the primary truth-bearers. As Parsons [63] points out, however, nothing crucial hinges on this: the talk of propositions is not necessary, and contextualism can be equivalently formulated taking sentences as truth-bearers.¹⁰ As it will become clear later, the core of the contextualist approach to truth and paradox consists in the idea that truth predications are to be interpreted in a suitable *hierarchy*. But the hierarchical structure of truth predications can be spelled out for sentences and propositions alike.¹¹ For ease of exposition, I will therefore take sentences as truth-bearers, and talk about sentences being true «in a context c », or more simply «in c ». However, the talk of propositions can always be recovered by replacing « $\ulcorner \varphi \urcorner$ expresses a true proposition in c » for « $\ulcorner \varphi \urcorner$ is true in c », where $\ulcorner \varphi \urcorner$ is a name of the sentence φ . Moreover, the idea that a sentence φ expresses a proposition in c can be interpreted as « $\ulcorner \varphi \urcorner$ is true in c or $\ulcorner \varphi \urcorner$ is false in c », where the latter means « $\ulcorner \neg \varphi \urcorner$ is true in c ».¹²

Assume that naïve rules for *truth introduction* and *truth elimination* hold unrestrictedly, so that one can always infer « $\ulcorner \varphi \urcorner$ is true in c » from φ and *vice versa*, when reasoning in c . Now a contextualist version of the Liar Paradox can be given.¹³

LIAR IN CONTEXT:

Let c be the context of reasoning, and λ_c be equivalent to « $\ulcorner \lambda_c \urcorner$ is not true in c ».

- | | |
|---|--------------------------|
| (1) $\ulcorner \lambda_c \urcorner$ is true in c or $\ulcorner \neg \lambda_c \urcorner$ is true in c | [Assumption] |
| (i-a) $\ulcorner \lambda_c \urcorner$ is true in c | [Assumption] |
| (ii-a) λ_c | [Truth elimination, i-a] |

- | | | |
|---------|---|---|
| (iii-a) | $\lceil \lambda_c \rceil$ is not true in c | [Definition of λ_c] |
| (iv-a) | Contradiction | [i-a and iii-a] |
| (i-b) | $\lceil \neg \lambda_c \rceil$ is true in c | [Assumption] |
| (ii-b) | $\neg \lambda_c$ | [Truth elimination, i-b] |
| (iii-b) | $\neg \lceil \lambda_c \rceil$ is true in c | [Definition of λ_c] |
| (iv-b) | $\lceil \lambda_c \rceil$ is true in c | [Double negation elimination, iii-b] |
| (v-b) | λ_c | [Truth elimination, iv-b] |
| (vi-b) | Contradiction | [ii-b and v-b] |
| (2) | Contradiction | [Disjunction elimination, 1, iv-a, vi-b] |
| (3) | It is not the case that ($\lceil \lambda_c \rceil$ is true in c or $\lceil \neg \lambda_c \rceil$ is true in c) | [Negation introduction 1, 2 (discharge of 1)] |
| (4) | $\lceil \lambda_c \rceil$ is not true in c and $\lceil \neg \lambda_c \rceil$ is not true in c | [De Morgan Equivalence, 3] |
| (5) | $\lceil \lambda_c \rceil$ is not true in c | [Conjunction elimination, 4] |
| (6) | λ_c | [Definition of λ_c] |
| (7) | $\lceil \lambda_c \rceil$ is true in c | [Truth introduction, 6] |
| (8) | Contradiction. | [5, 7] |

The LIAR IN CONTEXT is then used by orthodox contextualists to provide an argument against absolute generality.

CONTEXTUALIST ARGUMENT FROM PARADOX:

The LIAR IN CONTEXT is not sound, because the derivation (1)-(8) involves a covert context shift.¹⁴ First, one proves λ_c , i.e. that λ_c is not true in c , and then that λ_c is true *in a context c' different from c* . A context shift takes place between (6) and (7). Hence, the LIAR IN CONTEXT does not establish (7), but rather:

(7*): $\lceil \lambda_c \rceil$ is true *in c'* (for a context c' different from c).

Since (5) and (7*) are consistent, the contradiction is blocked. The reasoning from (1) to (7*) is sound, and it shows that the interpretation of the truth predicate in c is not maximally general: there are sentences that are neither true nor false in c , but true in c' .¹⁵

Note finally that there's nothing special about the LIAR IN CONTEXT: the CONTEXTUALIST ARGUMENT FROM PARADOX can be based on any other paradoxical reasoning involving unrestricted truth introduction and truth elimination. Having rejected absolute generality, orthodox contextualists develop a relativist view of quantification and truth, to which I now turn.

2.1 Contextualism à la Glanzberg In an orthodox contextualist framework, λ_c receives *two* interpretations: one for the context c where it is neither true nor false, and one for the context c' where it is true. In order to formally model this result, Glanzberg [25] employs iterations of Kripke's (1975) construction (strong Kleene version). I will now outline Glanzberg's construction in a simplified form, which will also clarify what is the connection between restrictions on generality and Kripke's theory (or iterations thereof).¹⁶

Kripke's construction is a model-theoretic construction which determines an interpretation for the truth predicate.¹⁷ More precisely, it yields a set E which serves as the extension of the truth predicate, namely the set of (names of) true sentences.¹⁸

In Kripke's construction, the set E is built in stages. More specifically, E is a certain stage occurring in a transfinite sequence of sets $E_0, E_1, \dots, E_\alpha, \dots$ indexed by ordinals. Here is, in a nutshell, how the sequence is defined. A «base model» \mathcal{M} is selected that interprets the truth-free part of the language. At stage 0, nothing is in the extension of the truth predicate, so that E_0 is empty. At stage 1, truth-free atomic sentences that are satisfied by the base model \mathcal{M} (such as « $0 = 0$ ») are declared true (so, they are added to E_1), and truth-free atomic sentences that are not satisfied by \mathcal{M} (such as «grass is red») are declared false (so, their negation is added to E_1). At stage 2, E_2 is obtained by applying the strong Kleene evaluation schema to the sentences in E_1 . So, for example, if both φ and ψ are in E_1 , then their conjunction $\varphi \wedge \psi$ is in E_2 ; if χ is in E_1 , then its double negation $\neg\neg\chi$ is in E_2 (similar clauses hold for the remaining connectives and quantifiers; the latter can be interpreted both substitutionally or *à la* Tarski). Finally, truth predications work similarly: if φ is in E_1 , then « $\ulcorner \varphi \urcorner$ is true» is in E_2 , and if $\neg\varphi$ is in E_1 , then « $\ulcorner \varphi \urcorner$ is not true» is in E_2 .¹⁹ All subsequent successor stages are constructed in a similar fashion, while one takes unions at limit stages. In this way, one builds a monotonic sequence:

$$E_0, E_1, \dots, E_\alpha, \dots$$

For cardinality reasons,²⁰ there is a (limit) ordinal β s.t. the sequence stops including new sentences from β onwards:

$$E_0, E_1, \dots, E_\alpha, \dots, E_\beta = E_{\beta+1} = \dots$$

E_β is called a *Kripkean fixed point*, and is identified with the final interpretation E .

The sketch I just provided describes the *least* Kripkean fixed point, i.e. the fixed point obtained setting $E_0 = \emptyset$. *Non-minimal* Kripkean fixed points (i.e. fixed points that include the least one) result by letting E_0 be non-empty. Every Kripkean fixed point E satisfies the so-called *transparency* of truth: for every sentence φ , the sentence « $\ulcorner \varphi \urcorner$ is true» is in E if and only if φ is in E . A Kripkean fixed point E is *consistent* if no φ is s.t. both φ and $\neg\varphi$ are in E . The least Kripkean fixed point, together with several others, is consistent (I will essentially only consider consistent Kripkean fixed points). Finally, if a Kripkean fixed point E is consistent, then there are sentences φ s.t. neither φ nor $\neg\varphi$ is in E : the sentence λ equivalent to « $\ulcorner \lambda \urcorner$ is not true» is a case in point. In other words, consistent Kripkean fixed points provide an essentially *partial* interpretation of the truth predicate. Indeed, the inner logic of a consistent Kripkean fixed point is strong Kleene logic, and hence it is a partial, non-classical logic (where, e.g., the law of excluded middle is not valid: $\lambda \vee \neg\lambda$ is a counterexample). Sentences that are neither in the extension nor in the anti-extension of a Kripkean fixed point are called *gappy* in it.

On Glanzberg's approach, a consistent fixed point E is used to determine the extension of the predicate «is true in the context c ». Glanzberg's theory, however, is classical, and hence a consistent fixed point E is not enough, since we've just seen that it provides a partial, non-classical interpretation of the language. As we have seen in the CONTEXTUALIST ARGUMENT FROM PARADOX, in a contextualist framework λ_c is declared to be not true in c . But λ_c precisely says that λ_c is not true in c . Therefore, a faithful contextualist model should satisfy λ_c . To accomplish this result, Glanzberg takes the so-called *closing-off* of E . That is, the sentences that are validated in his models are not merely those in the extension of the predicate «true in c » (viz. E), but rather the set of (names of) sentences that are classically satisfied by

$\langle \mathcal{M}, E \rangle$. Now, λ_c is not in E . Therefore, $\langle \mathcal{M}, E \rangle$ does not satisfy « $\ulcorner \lambda_c \urcorner$ is true in c ». Hence, $\langle \mathcal{M}, E \rangle$ satisfies « $\ulcorner \lambda_c \urcorner$ is not true in c ». But by definition of λ_c , this means that $\langle \mathcal{M}, E \rangle$ satisfies λ_c , as required. More compactly:

$$\begin{aligned} \langle \mathcal{M}, E \rangle &\not\models \ulcorner \lambda_c \urcorner \text{ is true in } c \text{ iff} \\ \langle \mathcal{M}, E \rangle &\models \ulcorner \lambda_c \urcorner \text{ is not true in } c \text{ iff} \\ \langle \mathcal{M}, E \rangle &\models \lambda_c \end{aligned}$$

While Kripkean fixed points unrestrictedly satisfy transparency, closed-off fixed points don't: if E is consistent, λ_c is in the closing-off of E (as just seen), but « $\ulcorner \lambda_c \urcorner$ is true in c » clearly isn't.

Now, in order to model the context shift, and interpret the predicate «is true in c' », Glanzberg constructs a *new* fixed point, but rather than starting from the empty set, he takes as starting point the sentences in the closing-off of E . And since λ_c is in the closing-off of E , the sentence « $\ulcorner \lambda_c \urcorner$ is true in c' » is in the fixed point built over it. Recall that, in the construction of a Kripkean fixed point, if $\varphi \in E_\alpha$, then « $\ulcorner \varphi \urcorner$ is true» $\in E_{\alpha+1}$. So, « $\ulcorner \lambda_c \urcorner$ is true in c' » is in the extension of the new fixed point. And clearly, it is also in the closing-off of the new fixed point (a Kripkean fixed point is always a subset of its own closing-off, by definition).

Putting things together, in Glanzberg's theory λ_c receives a formal treatment that matches the informal account of the LIAR IN CONTEXT. More specifically, λ_c is not true in the original context c (modeled by the closing-off of the first Kripkean fixed point), but is true in a subsequent context c' (modeled by the closing-off of the second fixed point). This succession of interpretations goes on indefinitely, for we can now formulate a new Liar sentence $\lambda_{c'}$ equivalent to « $\ulcorner \lambda_{c'} \urcorner$ is not true in c' », which is then declared true in a subsequent closed-off fixed point which models «true in c'' ». And so on. As in the conclusion of the CONTEXTUALIST ARGUMENT FROM PARADOX, orthodox contextualists take this line of reasoning to substantiate and formalize the idea that the interpretation of the truth predicate in c is not maximally general, as there are sentences that are neither true nor false in c , but true in c' . And since the LIAR IN CONTEXT reasoning can be run in any context, and c is completely arbitrary, no interpretation is maximally general.

2.2 Advantages of orthodox contextualism Orthodox contextualism has a number of theoretical virtues. First, it preserves classical logic. As a result, unlike non-classical approaches, mathematical and scientific theories can be combined with contextualist theories of truth without losing any content.²¹

Second, contextualist theories can non-trivially model paradoxical reasonings such as the Liar. Paradoxical reasonings are intuitively compelling—that is arguably why they are so bewildering, and why identifying where they go wrong is so controversial. However, virtually all non-contextualist theories limit themselves to blocking paradoxical arguments in one place or another (to avoid triviality), without providing *reconstructions* that do justice to the intuitive soundness of paradoxical arguments. Contextualist theories, in addition to blocking paradoxical reasonings as originally formulated, also provide sound re-interpretations of them (as shown in the CONTEXTUALIST ARGUMENT FROM PARADOX). So, contextualists have a simple explanation why paradoxical reasonings seem sound: because, once correctly interpreted, they *are* sound.²² As a consequence, contextualist theories manage to explain away the apparent awkwardness of accepting a sentence (namely λ) which,

by the theory's own lights, is not true.²³ In a contextualist's construal, this is not surprising: since λ is established from no premises and it «says» that λ is not true in c , it actually isn't true in c . However, such awkwardness is «made up for» in the next context c' , where λ is shown to be true.

Third, contextualist theories offer a uniform solution to all semantic paradoxes, including *revenge* paradoxes. Contextualist theories essentially trade on a restriction of absolute generality, and revenge paradoxes for contextualism try to reinstate absolute generality, in a form or another [e.g. 63]. However, absolute generality is the price contextualists *already pay*: if a paradox results from re-instating absolute generality, it hardly counts as a new, «revenge» paradox. Because of their ban on absolute generality, contextualist approaches are not free from expressive limitations. But they can offer a *uniform* solution to semantic paradoxes, including attempts at revenge (more on this in §5.4).²⁴

Finally, contextualist approaches improve on some of our best classical theories of truth. Consider the theory KF (for Kripke-Feferman), an axiomatization of Kripke's theory of truth in classical logic developed by Solomon Feferman [15].²⁵ KF has a number of virtues, including a very good approximation of the Tarskian compositional clauses for (type-free) truth, an elegant simplicity, and a high proof-theoretical power. However, KF cannot be consistently supplemented with all the instances of the following rules:

(Tr-Intro) from φ , infer $\text{Tr}(\ulcorner \varphi \urcorner)$ (Tr-Elim) from $\text{Tr}(\ulcorner \varphi \urcorner)$, infer φ

Now, KF can be closed under one of Tr-Intro and Tr-Elim but not both.²⁶ The contextualist theory developed by Glanzberg [25] recovers essentially all the positive features of KF (formulated in a contextualist setting) but, unlike KF, it can also be closed under Tr-Elim *and* a context-shifting version of Tr-Intro, according to which if φ holds in a context c , then φ is true in a suitably more extended context c' (this is a formal counterpart to the rule informally employed in the CONTEXTUALIST ARGUMENT FROM PARADOX). So, unlike in KF, in a contextualist theory one can always declare that, if φ has been established, then it is true in a suitable context.

2.3 Disadvantages of orthodox contextualism The main disadvantage of orthodox contextualism is its ban on absolute generality. More specifically, orthodox contextualism disallows maximally general interpretations of the truth predicate and, therefore, generally speaking disallows absolute generality. But how can it be that sentences such as «everything is self-identical» or «everything that is possible is necessarily possible» do not talk about absolutely everything? At the very least, it seems clear that one *can* use these sentences to talk about absolutely everything, nor is there any danger of paradox in doing so. Yet, orthodox contextualists reject this possibility, thereby imposing a highly revisionary semantics on general truths.

Such a revisionary semantics has high costs. For one thing, absolute generality is required for theoretical purposes: semantic, scientific, and philosophical generalizations arguably require the utmost level of generality [85]. For another, the blanket ban that orthodox contextualism puts on absolute generality seems to overshoot: even if one agrees with orthodox contextualists that the truth predicate should be re-interpreted when it comes to modeling the semantics of Liar sentences (and relevantly similar cases), no such re-interpretation is required for « $\ulcorner 0 = 0 \urcorner$ is true», nor is any domain restriction required for $\forall x(x = x)$. So, why can't such interpretations be maximally general?

We seem to have reached an impasse. On the one hand, orthodox contextualism (*à la* Glanzberg) strikes an excellent balance as a theory of truth, for its scientific applicability, uniformity, and strength. On the other, its ban on absolute generality cripples its interpretation of simple generalizations and unproblematic truth predications. In the next section, I will argue that the impasse is illusory: contextualism is broadly compatible with absolute generality. It follows, then, that the CONTEXTUALIST ARGUMENT FROM PARADOX has to be rejected.

3 Bicontextualism

3.1 Heuristics The basic idea of bicontextualism is that whether the interpretation of a sentence can be maximally general depends on that very sentence: in unproblematic cases (e.g. unproblematic truths such as «everything is self-identical» and unproblematic falsities such as «something is not is self-identical»), it can; in problematic cases (such as a Liar sentence), it cannot.²⁷ Therefore, the semantics developed here is bipartite: it provides an absolutist interpretation for the former sentences, and a relativist interpretation for the latter. The idea of curtailing the semantics only when it comes to problematic sentences is not new. Most non-classical theories of truth «recapture» classical logic when they deal with intuitively unproblematic sentences, and classical theories satisfy instances of Tr-Intro and Tr-Elim for truth in unproblematic cases.²⁸ In a similar fashion, bicontextualism recovers absolute generality whenever possible.

There are various ways to draw the line between unproblematic and problematic sentences.²⁹ In the approach I follow here, I do *not* proceed by first distinguishing the two cases, and then providing a semantics for them. Rather, the distinction between problematic and unproblematic sentences is hardwired in the semantics, as it were. In a nutshell, the unproblematic sentences are identified with the sentences in the least Kripkean fixed point, built for suitable languages. Therefore, the first task of this section is to introduce such languages (§3.2) and to provide an absolutist-friendly version of Kripke's construction (§3.3).

There are good reasons, I believe, to identify the unproblematic sentences with those in the least Kripkean fixed point. The least fixed point contains only *grounded* sentences, i.e. sentences whose value ultimately depends on the value of atomic sentences of the truth-free fragment of the language.³⁰ In particular, the least fixed point does not contain sentences which are *paradoxical* and *ungrounded* [in the now-standard terminology from Kripke 45] such as λ , or sentences that are unparadoxical and ungrounded. An example of ungrounded but unparadoxical sentence (in Kripke's sense) is given by a *truth-teller* sentence τ equivalent to « $\ulcorner \tau \urcorner$ is true». τ is ungrounded in the sense just sketched, but it is unparadoxical because, while neither τ nor $\neg\tau$ are in the extension of the least Kripkean fixed point, either one of them can be in the extension of some consistent non-minimal fixed point. For these reasons, the least Kripkean fixed point constitutes a natural (and rather strict) option to filter out sentences which are not completely unproblematic. However, other options are possible; in particular, the theory to be developed in §§3.3-3.4 can be straightforwardly adjusted to identify the unproblematic sentences with the sentences in some non-minimal consistent Kripkean fixed point.

Having characterized the unproblematic sentences, I now turn to the problematic ones. As seen in §2.1, sentences such as λ are outside of the extension of the

least Kripkean fixed point together with their negation, but they are evaluated in its closing-off. So, the problematic sentences are, roughly, those which are in the closing-off *without* also being in the (now absolutist) least fixed point (the distinction will require more care, as we shall see). A relativist semantics for such sentences is provided in §3.4. Finally, having developed both the absolutist and the relativist components of the semantics, these two «halves» are combined together in a single, bipartite interpretation, and a unified notion of consequence is provided (§3.5).

3.2 The object-language(s) Strictly speaking, I employ many object-languages. This is merely a matter of convenience: I employ different languages in order to straightforwardly model different interpretations of «true in c », as it will be clear from the following (see also §5.1).³¹

Definition 3.1 Let \mathcal{L} be a collection of first-order languages with identity, indexed by ordinals:

$$\mathcal{L} := \mathcal{L}_0, \mathcal{L}_1, \dots, \mathcal{L}_\alpha, \dots$$

such that every language in \mathcal{L} satisfies the following requirements (where $\alpha < \beta$):

- (a) \mathcal{L}_α includes the membership relation \in , and has no function constants,
- (b) \mathcal{L}_α contains a fresh unary predicate Tr_α (for truth),
- (c) \mathcal{L}_α and \mathcal{L}_β have the same individual variables, individual constants, and relation constants (with the exception of Tr_α),
- (d) \mathcal{L}_α has at least one *acceptable* structure \mathcal{M}_α (more on this shortly),
- (e) For every acceptable structure \mathcal{M}_α , for every \mathcal{L}_α -expression e , there is a closed term $\ulcorner e \urcorner$ that denotes the code of e in \mathcal{M}_α (*ditto*),
- (f) For every acceptable structure \mathcal{M}_α , for every open \mathcal{L}_α -formula $\varphi(x)$ there is a closed \mathcal{L}_α -term t_φ such that $(t_\varphi)^{\mathcal{M}_\alpha} = (\ulcorner \varphi(t_\varphi/x) \urcorner)^{\mathcal{M}_\alpha}$.

Some explanations are in order: I collect them in a single remark—the reader uninterested in (or already familiar with) the technical details can skim or skip it.

Remark 3.2

- By (a), every \mathcal{L}_α includes the language of set theory, so mathematical notions can be represented in the usual way. Also, no \mathcal{L}_α includes function constants, which simplifies the definitions in §3.3, but is of no consequence (functions can be defined as special cases of relations). (b) and (c) entail that if $\alpha < \beta$, $\mathcal{L}_\alpha \subseteq \mathcal{L}_\beta$. The same holds for the sets of constants ($\text{Con}_{\mathcal{L}_\alpha}$), variables ($\text{Var}_{\mathcal{L}_\alpha}$), formulae ($\text{For}_{\mathcal{L}_\alpha}$), and sentences ($\text{Sent}_{\mathcal{L}_\alpha}$).
- (b) allows us to formulate truth predications and, together with (c), it ensures that languages higher up in \mathcal{L} only differ from languages lower down in \mathcal{L} in the new truth predicates. So, if φ is a sentence of \mathcal{L}_α , then $\text{Tr}_{\alpha+1}(\ulcorner \varphi \urcorner)$ is a sentence of $\mathcal{L}_{\alpha+1}$ (more about $\ulcorner \cdot \urcorner$ shortly). Ordinal indices are used to identify contexts. The informal « φ is true in c » of §2 is now formalized as $\text{Tr}_\alpha(\ulcorner \varphi \urcorner)$, and will be interpreted by a suitable structure for \mathcal{L}_α .
- An acceptable structure (required in (d)) is a model \mathcal{M}_α s.t. a well-behaved machinery to code and decode \mathcal{L}_α -expressions is definable in \mathcal{M}_α . A coding is a function that associates \mathcal{L}_α -expressions with elements of the domain of \mathcal{M}_α .³² (e) ensures that each \mathcal{L}_α has sufficiently many terms for codes of \mathcal{L}_α -expressions. This provides a more precise formal counterpart for the notion

- of a «name» of φ employed in §2. I will often identify \mathcal{L}_α -expressions with their codes (no confusion will arise from this).
- Requirement (f) ensures that intuitively self-referential sentences, such as the sentences λ and τ mentioned above, are part of the languages in \mathcal{L} . So, for instance, the formal counterpart of a Liar sentence in \mathcal{L}_α is a sentence $\neg\text{Tr}_\alpha(t_{\lambda_\alpha})$, where t_{λ_α} denotes (in \mathcal{M}_α) the same element as $\ulcorner\neg\text{Tr}_\alpha(t_{\lambda_\alpha})\urcorner$. Abbreviate $\neg\text{Tr}_\alpha(t_{\lambda_\alpha})$ as λ_α . A Liar sentence is thus a sentence λ_α equivalent to $\neg\text{Tr}_\alpha(\ulcorner\lambda_\alpha\urcorner)$, i.e. « $\ulcorner\lambda_\alpha\urcorner$ is not true in context α », formalizing the informally defined λ_c I used in §2.³³
 - I did not specify whether \mathcal{L} is a proper class or a set (and, in the latter case, which is supremum of the set of ordinals used as indices). This is on purpose: many countable transfinite limit ordinals will do.³⁴ However, if α is sufficiently large, it becomes impossible to define codings as well-behaved as is needed by the acceptability requirement. Nevertheless, weaker notions of codings can be defined also for languages $\mathcal{L}_{\alpha_\kappa}$ (for α_κ an ordinal of cardinality κ , and κ large enough), and the semantic construction to be developed in §§3.3-3.4 can be carried on employing these weaker codings as well.³⁵

Summing up, our object-languages are given by any collection \mathcal{L} that respects requirements (a)-(f). (Many such collections exist.) This is for the sake of generality: a bicontextualist semantics can be developed for a wide variety of (minimally expressive) languages, and is not restricted to a specific vocabulary. I now turn to the tasks outlined above: formulate an absolutist-friendly Kripkean construction to interpret the unproblematic sentences of each $\mathcal{L}_\alpha \in \mathcal{L}$, and a relativist-friendly version of the Kripkean construction to model the problematic ones, and their context-shift.

3.3 Absolutist semantics Rayo and Uzquiano [66] have shown how to construct a second-order, Tarskian semantics that interprets a first-order object-language in an absolutist way. In §3.3.1, I sketch the basics of the Rayo-Uzquiano approach. In §3.3.2, I adapt it to construct a second-order, *Kripkean* semantics that interprets the languages in \mathcal{L} in an absolutist way. This construction is then used to isolate the unproblematic sentences, which are interpreted unrestrictedly.

3.3.1 The Rayo-Uzquiano approach to absolutist models Rayo and Uzquiano take the language of first-order set theory as their object-language, and the corresponding second-order language as their meta-language. However, their proposal applies to any first-order language. Therefore, I will take an arbitrary $\mathcal{L}_\alpha \in \mathcal{L}$ as object-language, and the corresponding second-order language (\mathcal{L}_α^2) as meta-language.

In model theory, models are sets. More specifically, a model \mathcal{M} of a language \mathcal{L} is a set $\langle M, I \rangle$, where M is a non-empty set of individuals (the domain of the model), and I is a set of tuples of elements of M which provides the extension of the individual and relation constants of \mathcal{L} . That a model is a set, as is well-known, makes it impossible for it to interpret quantifiers unrestrictedly: no set includes absolutely everything. Therefore, «everything is self-identical» ($\forall x(x = x)$) or «nothing belongs to the empty set» ($\neg\exists x(x \in \emptyset)$) are not interpreted as being about absolutely everything (or absolutely every set): they are about the (set-many) things which are in M . As Rayo and Uzquiano [66, pp. 316-317] argue, this is no minor incident: it entails that «no standard model provides the language of set theory with its intended interpretation». To solve the problem, Rayo and Uzquiano propose an alternative notion of model:

The core of our proposal is that we conceive of a model, not as a single set-theoretic object, but rather as given by the values of a second-order variable X . Accordingly, we take satisfaction to be a relation that a formula φ bears, not to a certain structured set, but to the values of X . These objects will encode a specification of the individuals over which our first-order quantifiers are to range and a specification of the ordered pairs that are to be assigned to $\langle\langle\rangle\rangle$.

[66, pp. 318-319, notation adapted]

So, rather than using a set as a model, they define a second-order formula $\mathbb{M}(X)$, with a free second-order variable X , which encodes what it takes for X to work as a model. To distinguish it from the usual notion of model, we will read $\mathbb{M}(X)$ as $\langle\langle X$ is an RU-model $\rangle\rangle$. The values of the variable X then provide instances of the newly defined RU-models.³⁶ Here is the definition.³⁷

Definition 3.3 (RU-model) For every unary second-order variable X , the formula $\langle\langle X$ is an RU-model $\rangle\rangle$, in symbols $\mathbb{M}(X)$, is defined as:

$$\begin{aligned} & \exists x X(\langle\langle \ulcorner \forall \urcorner, x \rangle\rangle) \wedge \\ & \forall x [X(x) \rightarrow \exists y (x = \langle\langle \ulcorner \forall \urcorner, y \rangle\rangle) \vee \exists y \exists z_1, \dots, \exists z_n (\text{Rel}_{\mathcal{L}_\alpha}(y) \wedge x = \langle\langle \ulcorner y \urcorner, \langle z_1, \dots, z_n \rangle \rangle)] \wedge \\ & \forall x [\text{Con}_{\mathcal{L}_\alpha}(x) \vee \text{Var}_{\mathcal{L}_\alpha}(x) \rightarrow (X(\langle\langle \ulcorner \forall \urcorner, \ulcorner x \urcorner \rangle\rangle))] \wedge \\ & \forall x \forall y_1, \dots, y_n [\text{Rel}_{\mathcal{L}_\alpha}(x) \wedge X(\langle\langle \ulcorner x \urcorner, \langle y_1, \dots, y_n \rangle \rangle) \rightarrow X(\langle\langle \ulcorner \forall \urcorner, y_1 \rangle\rangle) \wedge \dots \wedge X(\langle\langle \ulcorner \forall \urcorner, y_n \rangle\rangle)] \end{aligned}$$

Let's unpack the definition.

Remark 3.4

- I follow Rayo and Uzquiano [66, p. 319] in taking an RU-model $\langle\langle$ to be given by ordered pairs of two different types: (1) ordered pairs of the form $\langle\langle \ulcorner \forall \urcorner, x \rangle\rangle$, which [...] encode the fact that x is to be within the range of our quantifiers», and (2) ordered pairs of the form $\langle\langle \ulcorner R \urcorner, \langle y_1, \dots, y_n \rangle \rangle$, for R a relation constant of \mathcal{L}_α , which encode the fact that the tuple $\langle y_1, \dots, y_n \rangle$ is part of the interpretation of the relation R . This is what the second conjunct of the definition says, i.e. that every x to which X applies (informally: $\langle\langle$ everything in the RU-model $\rangle\rangle$) is either a pair of the first kind or a pair of the second kind.
- The first conjunct of the definition says that X applies to something, i.e. something is in the range of the quantifiers. In standard model-theoretic terms, this corresponds to the assumption that a model has a non-empty domain.
- The third conjunct says that every \mathcal{L}_α -individual constant or variable is associated with one element in the range of the quantifiers (together with the second conjunct, this ensures that each constant or variable is associated with exactly one element). In standard model-theoretic terms, this corresponds to the assumption that a model specifies a denotation for individual constants and individual variables.³⁸
- The last conjunct says that every \mathcal{L}_α -relation constant is associated with a tuple of elements in the range of the quantifiers. In standard model-theoretic terms, this corresponds to the assumption that a model specifies the extension of the relation constants.

RU-models are given by values of X in $\mathbb{M}(X)$.³⁹ Thus, certain values of X will yield standard, set-theoretic models, those where the collections forming the $\langle\langle$ domain $\rangle\rangle$ in $\mathbb{M}(X)$ and the $\langle\langle$ denotation $\rangle\rangle$ and $\langle\langle$ extension $\rangle\rangle$ of the vocabulary of \mathcal{L}_α

in $\mathbb{M}(X)$ form sets. Some others will not be set-theoretical models. An absolutist model is one in which the value corresponding to the «domain» in $\mathbb{M}(X)$ is given by absolutely everything, and the «denotations» of constants and variables, and the «extension» of relations are defined on absolutely everything.⁴⁰

Once one has the absolutist notion of RU-model, one can use it to provide a semantics—that is, a notion of satisfaction—for each \mathcal{L}_α . To this end, it will be useful to first isolate the «domain» component of the formula $\mathbb{M}(X)$.

Definition 3.5 (RU-domain) For all second-order variables X and Y , the formula « Y is the RU-domain of the RU-model X », in symbols $\mathbb{D}(X, Y)$, is defined as:

$$\mathbb{D}(X, Y) \text{ iff } \mathbb{M}(X) \wedge \forall x (Yx \leftrightarrow \exists y (Xy \wedge y = \langle \ulcorner \forall \urcorner, x \rangle))$$

So, $\mathbb{D}(X, Y)$ holds just in case X is an RU-model, and Y applies exactly to the things which constitute the second element of the first kind of pairs in X (as explained in the first item in Remark 3.4), namely its «domain»—henceforth, its «RU-domain».

Next, we need the analogue of a y -variant of a variable assignment, namely an RU-model Y which differs from the RU-model X at most in the value it assigns to the \mathcal{L}_α -variable y .

Definition 3.6 (RU-variant) For all first-order variables y , and unary second-order variables X and Y , the formula « Y is an RU- y -variant of X », in symbols $\mathbb{V}(y, Y, X)$, is defined as:

$$\begin{aligned} \text{Var}_{\mathcal{L}_\alpha}(y) \wedge \mathbb{M}(X) \wedge \mathbb{M}(Y) \wedge \exists! Z (\mathbb{D}(X, Z) \wedge \mathbb{D}(Y, Z)) \wedge \\ \forall x [x \neq y \rightarrow \forall z (X(\langle z, x \rangle) \leftrightarrow Y(\langle z, x \rangle))] \end{aligned}$$

So, X and Y are RU- y -variants of each other if they (i) are both RU-models, (ii) have the same RU-domain, and (iii) disagree at most on their interpretation of the \mathcal{L}_α -variable y .⁴¹

Definitions 3.3-3.6 provide us with all the ingredients to formulate absolutist-friendly satisfaction conditions for the languages in \mathcal{L} . This is, indeed, what Rayo and Uzquiano proceed to do, offering a Tarskian semantics. However, the Rayo-Uzquiano method does not only apply to Tarskian satisfaction: it applies to any inductive definition of satisfaction. In particular, as anticipated, I now use it to formulate Kripkean, absolutist-friendly satisfaction conditions for each language \mathcal{L}_α , which I then use to interpret its unproblematic sentences.

3.3.2 A Kripkean absolutist semantics for unproblematic sentences We now define a predicate which can be informally (if imprecisely) rendered as «the RU-model X with RU-domain Y Kripke-satisfies x relative to the accepted Z » (more on «the accepted Z » in a moment).⁴² I also assume that the RU-model X we use as base model is acceptable, i.e. that its definition includes a conjunct which translates the acceptability conditions (see Remark 3.2) into our official meta-language \mathcal{L}_α^2 —as this is obviously possible, I skip the lengthy definition for space reasons.

Definition 3.7 Let x be a first-order variable, and X , Y , and Z be second-order variables. $\text{KSat}_\alpha(x, X, Y, Z)$ if and only if:⁴³ $\mathbb{M}(X)$ and $\mathbb{D}(X, Y)$ and

- (i) Zx , or

- (ii) x is $\ulcorner R(t_1, \dots, t_n) \urcorner$ and there are y_1, \dots, y_n s.t. $X(\langle \ulcorner \forall \urcorner, y_1 \rangle) \wedge \dots \wedge X(\langle \ulcorner \forall \urcorner, y_n \rangle)$ and $X(\langle \ulcorner R \urcorner, \langle y_1, \dots, y_n \rangle \rangle)$, or
- (iii) x is $\ulcorner \neg R(t_1, \dots, t_n) \urcorner$ and there are y_1, \dots, y_n s.t. $X(\langle \ulcorner \forall \urcorner, y_1 \rangle) \wedge \dots \wedge X(\langle \ulcorner \forall \urcorner, y_n \rangle)$ and $\neg X(\langle \ulcorner R \urcorner, \langle y_1, \dots, y_n \rangle \rangle)$, or
- (iv) x is $\neg \neg y$ and $\text{KSat}_\alpha(y, X, Y, Z)$, or
- (v) x is $y \wedge z$ and $\text{KSat}_\alpha(y, X, Y, Z)$ and $\text{KSat}_\alpha(z, X, Y, Z)$, or
- (vi) x is $\neg(y \wedge z)$ and $\text{KSat}_\alpha(\neg y, X, Y, Z)$ or $\text{KSat}_\alpha(\neg z, X, Y, Z)$, or
- (vii) x is $\forall yz(y)$ and for every W s.t. $\forall(y, W, X)$, $\text{KSat}_\alpha(z(y), W, Y, Z)$, or
- (viii) x is $\neg \forall yz(y)$ and for some W s.t. $\forall(y, W, X)$, $\text{KSat}_\alpha(\neg z(y), W, Y, Z)$, or
- (ix) x is $\text{Tr}_\alpha(y)$ and y is $\ulcorner z \urcorner$ for $z \in \mathcal{L}_\alpha$ and $\text{KSat}_\alpha(z, X, Y, Z)$, or
- (x) x is $\neg \text{Tr}_\alpha(y)$ and either y does not code an \mathcal{L}_α -sentence, or y is $\ulcorner z \urcorner$ for $z \in \mathcal{L}_\alpha$ and $\text{KSat}_\alpha(\neg z, X, Y, Z)$.

This definition formalizes the inductive construction of the extension of a Kripkean fixed point for languages in \mathfrak{L} in a higher-order framework. Let's highlight its main features.

Remark 3.8

- In §2.1, a Kripkean fixed point was defined as an element occurring at certain stages in a succession indexed by ordinals. By contrast, and following Halbach [32] (Ch. 15), here I directly define the fixed point (rather than a succession which *reaches* a fixed point), as shown by the «if and only if» at the beginning of Definition 3.7. Nothing crucial hinges on these differences.
- The second-order variables X and Y appear in the conjuncts $\mathbb{M}(X)$ and $\mathbb{D}(X, Y)$. So, the value of X is the base, acceptable RU-model, and the value of Y is the RU-domain where the quantifiers of \mathcal{L}_α range. I do not specify a single value for X and Y , so that they can be assigned both set-theoretic and absolutist models.
- The variable Z appears in (i), as the disjunct « Zx »: this guarantees that, if Z applies to x , then $\text{KSat}_\alpha(x, X, Y, Z)$ holds, i.e. x is in the fixed point defined by $\text{KSat}_\alpha(x, X, Y, Z)$. This clause allows us to define both the least and non-minimal Kripkean fixed points in our framework. If Z takes an empty value, the clause Zx is never satisfied, and one obtains the least fixed point. I will write $\text{KSat}_\alpha(x, X, Y, \emptyset)$ for the least fixed point whose RU-domain is Y . If Z takes a non-empty value, then $\text{KSat}_\alpha(x, X, Y, Z)$ yields the non-minimal fixed point built *over* that value.⁴⁴

Summing up: the second-order formula $\text{KSat}_\alpha(x, X, Y, Z)$ characterizes the notion of «(possibly) absolute Kripkean satisfaction» for each \mathcal{L}_α . Using it, we can finally define the set of *absolutely general truths* of each \mathcal{L}_α .

Definition 3.9 For every $\mathcal{L}_\alpha \in \mathfrak{L}$, the set of *absolutely general truths* of \mathcal{L}_α , in symbols Abs_α , is defined as follows:

$$\text{Abs}_\alpha := \{ \varphi \in \mathcal{L}_\alpha \mid \text{KSat}_\alpha(\ulcorner \varphi \urcorner, X, Y, \emptyset) \}$$

Simply put, Abs_α contains the sentences that are in the least Kripkean fixed point for the language \mathcal{L}_α , over a base RU-model X and its RU-domain Y . Clearly, here I am mostly interested in the case in which the value of Y is given by absolutely

everything, but, again, I leave it open for the sake of generality. Since Abs_α is given by the *least* fixed point of the (absolutist-friendly) Kripke construction, these sentences match the intuition spelled out in §3.1 that such fixed point provides an appropriate way to isolate the unproblematic sentences. However, variants of Abs_α can easily be defined that correspond to consistent non-minimal fixed points.

3.4 Relativist semantics I now turn to the problematic sentences, and their relativist semantics. Recall (from §2 and §3.1) that, in the bicontextualist picture, the problematic sentences of each language \mathcal{L}_α must:

- (a) first be declared to be neither true_α nor false_α , and then be declared either $\text{true}_{\alpha+1}$ or $\text{false}_{\alpha+1}$, in order to model the context shift;
- (b) always be interpreted over domains that are not absolutely general.

Concerning (b), for the sake of simplicity, I will always employ set-size domains in the case of problematic sentences, but this restriction can also be relaxed.⁴⁵ As discussed in §2 and §3.1, problematic sentences are (roughly) identified as those sentences that are gappy in the least Kripkean fixed point. Therefore, if φ is problematic, $\text{KSat}_\alpha(x, X, Y, \emptyset)$ does not apply to φ , nor to $\neg\varphi$, but the closing-off of the fixed point defined by $\text{KSat}_\alpha(x, X, Y, \emptyset)$ applies to exactly one of them. In order to accomplish this, I reproduce the closing-off in my higher-order framework, making sure the domain is always set-sized and using different set-sized closed-off fixed points to interpret problematic sentences in different contexts.⁴⁶

3.4.1 Closing-off

Definition 3.10 Let \mathcal{M}_0 be an acceptable model of the Tr_0 -free fragment of \mathcal{L}_0 (namely $\mathcal{L}_0 \setminus \{\text{Tr}_0\}$), and M_0 its domain. The relativistic closing-off for \mathcal{L}_0 (relative to \mathcal{M}_0) is:

$$\text{C-Off}_0 := \{ \varphi \in \text{Sent}_{\mathcal{L}_0} \mid \langle \mathcal{M}_0, \{x \in M_0 \mid \text{KSat}_0(x, \mathcal{M}_0, M_0, \emptyset)\} \rangle \models \varphi \}$$

Let $\alpha > 0$, and let \mathcal{M}_α be an acceptable $\mathcal{L}_\alpha \setminus \{\text{Tr}_\alpha\}$ -structure s.t. for every $\beta < \alpha$, $\mathcal{M}_\beta \prec_{\bigcup_{\beta < \alpha} \text{C-Off}_\beta} \mathcal{M}_\alpha$, i.e. \mathcal{M}_α is an elementary extension of \mathcal{M}_β with respect to the sentences in $\bigcup_{\beta < \alpha} \text{C-Off}_\beta$.⁴⁷ Then, the relativistic closing-off of \mathcal{L}_α (relative to \mathcal{M}_α) is given by:

$$\text{C-Off}_\alpha := \{ \varphi \in \text{Sent}_{\mathcal{L}_\alpha} \mid \langle \mathcal{M}_\alpha, \{x \in M_\alpha \mid \text{KSat}_\alpha(x, \mathcal{M}_\alpha, M_\alpha, \bigcup_{\beta < \alpha} \text{C-Off}_\beta)\} \rangle \models \varphi \}$$

Let me unpack the above definition.

Remark 3.11

- First, one constructs the set of sentences defined by $\text{KSat}_0(x, \mathcal{M}_0, M_0, \emptyset)$, i.e. the extension of the least Kripkean fixed point for \mathcal{L}_0 . By construction, the RU-domain of $\text{KSat}_0(x, \mathcal{M}_0, M_0, \emptyset)$ is a set, M_0 , and hence cannot contain absolutely everything—as desired. Then, one considers the set of sentences that are classically true in $\langle \mathcal{M}_0, \{x \in M_0 \mid \text{KSat}_0(x, \mathcal{M}_0, M_0, \emptyset)\} \rangle$, where \mathcal{M}_0 interprets the Tr_α -free fragment of \mathcal{L}_0 and $\{x \in M_0 \mid \text{KSat}_0(x, \mathcal{M}_0, M_0, \emptyset)\}$ is the extension of Tr_0 . This is the closing off proper, C-Off_0 .
- Then, one builds a new fixed point over C-Off_0 , defined by the formula $\text{KSat}_1(x, \mathcal{M}_1, M_1, \text{C-Off}_0)$. Once *that* fixed point is closed-off, we have the

set C-Off₁, which interprets \mathcal{L}_1 , and agrees with the interpretation of \mathcal{L}_0 offered by C-Off₀. Again, the RU-domain of C-Off₁ is a set, namely M_1 . And so on: the process goes on, unaltered, for every language in \mathcal{L} .

- The requirement that \mathcal{M}_α is an elementary extension of every \mathcal{M}_β for $\beta < \alpha$ with respect to the sentences in $\bigcup_{\beta < \alpha} \text{C-Off}_\beta$ ensures that \mathcal{M}_α agrees with all the previous \mathcal{M}_β s concerning the sentences in the closing-offs of the fixed-points defined over them. Therefore, all the fixed points that are built along the sequence are consistent (and so are the corresponding closing-offs), and extend each other. So, taking $\bigcup_{\beta < \alpha} \text{C-Off}_\beta$ as ‘the accepted Z ’ (as per Definition 3.7) does not result in inconsistent sets of \mathcal{L}_α -sentences. Finally, the acceptability requirement ensures that there are definable coding functions for all the \mathcal{L}_β s up to \mathcal{L}_α . Note that the existence of such an \mathcal{M}_α is an immediate consequence of Kripke’s original construction.

What happens to the intuitively problematic sentences? Definition 3.10 yields a succession of set-sized closed-off fixed points, so it works exactly as Glanzberg’s construction reviewed in §2.1. Let λ_0 be a Liar sentence in \mathcal{L}_0 . Now, λ_0 is not in the extension of the fixed point defined by $\text{KSat}_0(x, \mathcal{M}_0, M_0, \emptyset)$. Clearly, then, since $\text{KSat}_0(x, \mathcal{M}_0, M_0, \emptyset)$ yields the extension of Tr_0 , λ_0 is not in it. Therefore (cf. §2.1):

$$\begin{aligned} \langle \mathcal{M}_0, \{x \in M_0 \mid \text{KSat}_0(x, \mathcal{M}_0, M_0, \emptyset)\} \rangle &\not\models \text{Tr}_0(\ulcorner \lambda_0 \urcorner) \\ &\models \neg \text{Tr}_0(\ulcorner \lambda_0 \urcorner) \\ &\models \lambda_0 \end{aligned}$$

In words: λ_0 is not true in \mathcal{L}_0 . The sentence that says of itself that it is not true in \mathcal{L}_0 is indeed untrue in \mathcal{L}_0 —as desired.

Things change when one moves to \mathcal{L}_1 and C-Off₁. C-Off₁ agrees with C-Off₀ on \mathcal{L}_0 , but it interprets \mathcal{L}_1 , which also features a new, more encompassing notion of truth, namely Tr_1 . Since $\lambda_0 \in \text{C-Off}_0$, λ_0 goes into the extension of Tr_1 given by the fixed point for \mathcal{L}_1 built over C-Off₀. In other words, $\text{KSat}_1(x, \mathcal{M}_1, M_1, \text{C-Off}_0)$ applies to (the code of) λ_0 , and therefore:

$$\langle \mathcal{M}_1, \{x \in M_1 \mid \text{KSat}_1(x, \mathcal{M}_1, M_1, \text{C-Off}_0)\} \rangle \models \text{Tr}_1(\ulcorner \lambda_0 \urcorner)$$

Namely, λ_0 is true in \mathcal{L}_1 —again, as desired.

In turn, \mathcal{L}_1 has its own problematic sentences, e.g. a sentence λ_1 equivalent to $\neg \text{Tr}_1(\ulcorner \lambda_1 \urcorner)$, and they are treated exactly as sketched above in C-Off₂. And so on for every $\mathcal{L}_\alpha \in \mathcal{L}$.

3.4.2 A Kripkean relativist semantics for problematic sentences Having constructed the succession of closed-off fixed points for languages in \mathcal{L} , one can finally use them to single out the problematic sentences, which are to be interpreted over set-sized domains, separating them from the absolutely general ones.

Definition 3.12 For every $\mathcal{L}_\alpha \in \mathcal{L}$, the set of *relatively general truths* of \mathcal{L}_α , in symbols Rel_α , is defined as follows:

$$\text{Rel}_\alpha := \text{C-Off}_\alpha \setminus (\text{Abs}_\alpha \cup \{\varphi \in \mathcal{L}_\alpha \mid \neg \varphi \in \text{Abs}_\alpha\})$$

Again, let’s unpack the above definition, to make sure Rel_α captures its intended extension.

Remark 3.13

- The problematic sentences are obtained as a complement: since the closing-off of a Kripkean fixed point includes that fixed point, $C\text{-Off}_\alpha$ includes all the sentences in the least Kripkean fixed point over base model \mathcal{M}_α —including $\forall x(x = x)$ and many other unproblematic sentences which are also in Abs_α . This explains the first complement: we need to remove Abs_α from $C\text{-Off}_\alpha$.
- Second, we need to remove from $C\text{-Off}_\alpha$ all the \mathcal{L}_α -sentences φ s.t. that $C\text{-Off}_\alpha$ *disagrees about* φ with Abs_α , i.e. the φ s s.t. $\varphi \in C\text{-Off}_\alpha$ but $\neg\varphi \in \text{Abs}_\alpha$. $C\text{-Off}_\alpha$ and Abs_α might disagree over sentences which have nothing to do with semantic paradoxes, merely because their domains are different—e.g. if one is absolutely general, and the other one is set-sized. Suppose Abs_α is defined with reference to absolutely all sets, and suppose κ is a cardinal s.t. $|M_\alpha| < \kappa$. Let φ be the sentence «there is no set of cardinality κ ». Clearly, $C\text{-Off}_\alpha$ and Abs_α disagree about φ , since $\varphi \in C\text{-Off}_\alpha$ but $\neg\varphi \in \text{Abs}_\alpha$. In this case, we should remove φ from Rel_α , since we accept Abs_α for unproblematic sentences. This explains the second complement in the definition of Rel_α .
- Finally, note that problematic sentences that are validated by closed-off fixed points, such as λ_α , are in Rel_α (while, of course, $\neg\lambda_\alpha$ is not), as required.

3.5 Bicontextualism: the full story With the notions of absolute and relative truth in place, one can finally provide a bicontextualist notion of consequence, and thus a theory of truth proper.

Definition 3.14 For every $\mathcal{L}_\alpha \in \mathfrak{L}$, and every $\{\Gamma, \varphi\} \subseteq \text{Sent}_{\mathcal{L}_\alpha}$, the argument from Γ to φ is *bicontextually valid*, in symbols $\Gamma \models_\alpha^{\text{bc}} \varphi$, if and only if:

if all the sentences in Γ are in $\text{Abs}_\alpha \cup \text{Rel}_\alpha$, so is φ .

In a nutshell, for every language \mathcal{L}_α , its bicontextualist semantics interprets all the unproblematic sentences over a possibly absolutely unrestricted domain, and the problematic ones over a restricted, always extendable domain. Bicontextual validity is then defined as preservation of either absolute truth in a model (Abs_α) or relative truth in a model (Rel_α). Given that absolute and relative truth in a model are exhaustive and exclusive (as far as the true sentences of \mathcal{L}_α are concerned), Definition 3.14 is in effect a bicontextualist version of preservation of truth in a model, as it is to be expected for a classical theory.⁴⁸

Bicontextualism validates several desirable truth-theoretical principles. First, it validates all the axioms of the theory KF (discussed in §2.2) for every language \mathcal{L}_α . This means that, for every $\mathcal{L}_\alpha \in \mathfrak{L}$, the axioms of KF formulated for \mathcal{L}_α are made either absolutely or relatively true in the bicontextualist semantics for \mathcal{L}_α . For example, in a bicontextualist construal, the axiom KF4 reads:

$$\forall x[\text{Sent}_{\mathcal{L}_\alpha}(x \wedge y) \rightarrow (\text{Tr}_\alpha(x \wedge y) \leftrightarrow \text{Tr}_\alpha(x) \wedge \text{Tr}_\alpha(y))].$$

The bicontextualist can therefore argue that the notion of truth we find in natural language satisfies the axioms of KF, but the latter are subject to a sort of typical ambiguity, since the semantics has to determine the appropriate context (i.e. level in the hierarchy of languages and models) to interpret them.⁴⁹

Moreover, bicontextualism is closed under the following truth-elimination rule, for $\varphi \in \mathcal{L}_\alpha$:

$$\text{Tr}_\alpha(\ulcorner \varphi \urcorner) \models_\alpha^{\text{bc}} \varphi$$

and the following truth-introduction rule, for $\varphi \in \mathcal{L}_\alpha$:

$$\text{if } \Gamma \models_\alpha^{\text{bc}} \varphi, \text{ then } \Gamma \models_{\alpha+1}^{\text{bc}} \text{Tr}_{\alpha+1}(\ulcorner \varphi \urcorner)$$

(The limit case reduces to the successor one, since limit stages are defined as unions). The proof is immediate, and follows from the construction of the Kripkean fixed points and their closing-off explained in §3.3.2 and §3.4.1. These rules are closely related to the principles Tr-Intro and Tr-Elim discussed in §2.2. However, while classical theories of truth cannot consistently feature both truth-elimination and truth-introduction principles, bicontextualism can unproblematically have both, provided that the truth-introduction rule is in general context-shifting. However, if $\varphi \in \text{Abs}_\alpha$, then $\text{Tr}_\alpha(\ulcorner \varphi \urcorner) \in \text{Abs}_\alpha$ as well. In other words, applying the truth predicate to an absolutely true sentence does not require a context shift. The context-shifting nature of truth-introduction *in general* is uniquely due to the presence of problematic sentences, and their relativist interpretation.

Clearly, the converse rules hold as well, for $\varphi \in \mathcal{L}_\alpha$:

$$\neg\varphi \models_\alpha^{\text{bc}} \neg\text{Tr}_\alpha(\ulcorner \varphi \urcorner)$$

$$\text{if } \Gamma \models_{\alpha+1}^{\text{bc}} \neg\text{Tr}_{\alpha+1}(\ulcorner \varphi \urcorner), \text{ then } \Gamma \models_\alpha^{\text{bc}} \neg\varphi$$

And while the truth introduction rule in general induces a shift to a language higher up in the hierarchy, the $\neg\text{Tr}$ elimination rule induces a shift to a language *lower down* in the hierarchy.

In conclusion, bicontextualism improves on the main shortcoming of orthodox contextualism: the lack of an absolutely general interpretation for arguably general truths and falsities. At the same time, bicontextualism retains the attractive features of orthodox contextualism (§2.2): the preservation of classical logic, the ability to reconstruct paradoxical reasonings and explain their intuitive soundness, and a uniform solution to standard and revenge paradoxes (more on this in §5.4). Finally, bicontextualism validates appealing truth-theoretical principles, such as the ones of KF, as well as versions of both truth-introduction and truth-elimination principles. The latter principles arguably bring the semantics of the truth predicate fairly close to a naïve interpretation, while remaining within the boundaries imposed by classical logic.

4 Revisiting the CONTEXTUALIST ARGUMENT FROM PARADOX

In addition to its consequences for formal theories of truth, the development of bicontextualism has some consequences for the foundations of semantics more generally. Both absolutist and relativist semantics get some things right, and some things wrong. Absolutist semantics is well-suited to interpret absolutely general claims such as «everything is self-identical», but cannot model paradoxical arguments as arguments that *expand* a given interpretation of the truth predicate, or a given domain of quantification. Conversely, relativist semantics offers an adequate account of paradoxical reasonings, but systematically misinterprets absolutely general claims. Both problems have a common source: the assumption, shared by absolutists and

relativists alike, that every language—even one as simple as \mathcal{L}_α —has a *unified* interpretation:

UNIFIED INTERPRETATION: All the sentences of a given languages are given a single interpretation (be it first- or higher-order, relativist or absolutist).

This assumption lies at the heart of contemporary semantics. In model-theoretic semantics, for example, models are ordered pairs of the form $\langle M, I \rangle$, where M is a set and I is a function (also a set) that maps linguistic expressions to appropriate extensions in M (see §3.3.1). In a higher-order absolutist semantics, the situation is similar: even though here RU-domains are not sets, the semantics still defines a unique satisfaction relation (as in Rayo and Uzquiano [66]) and/or a unique interpretation (as in Rayo and Williamson [68]) that relates syntactic objects of the appropriate kind (such as terms or formulae) and things that belong to an RU-domain.

Bicontextualism challenges UNIFIED INTERPRETATION. By bicontextualist lights, every minimally expressive language such as \mathcal{L}_α (and, *a fortiori*, natural language) has two kinds of sentences: the problematic and the unproblematic ones. The appropriate interpretation for a sentence φ depends on whether φ is one of the former or one of the latter. In order to do justice to the semantics of both kinds of sentences, UNIFIED INTERPRETATION has to be abandoned in favor of a bipartite one. Such a bipartite interpretation is clearly visible in Definition 3.14, which employs both Abs_α (with its possibly absolutely general RU-domain given by the value of X), and Rel_α (with its necessarily restricted domain M_α). This also shows that the revision to the framework of standard semantics provided by bicontextualism is not particularly radical: one kind of interpretation is not sufficient, while more than two are not necessary. Two kinds of interpretations—one for problematic and one for unproblematic sentences—are just right.

We are now in a position to provide a diagnosis of what goes wrong in the CONTEXTUALIST ARGUMENT FROM PARADOX. The final part of the argument (see §2) can now be recast as follows.

[The LIAR IN CONTEXT] shows that the interpretation of the language in context c is not maximally general: there are sentences which are true or false in c' but neither true nor false in c . But since the LIAR IN CONTEXT can be performed in any context, and by UNIFIED INTERPRETATION there is a single interpretation for all sentences, no interpretation can be maximally general.

The argument breaks down at the second step. The fallacy: UNIFIED INTERPRETATION. By bicontextualist lights, problematic sentences such as Liar sentences show that there just isn't a single interpretation for every sentence of the language. All that follows from the CONTEXTUALIST ARGUMENT FROM PARADOX, properly understood, is that the initial interpretation of the truth predicate in *problematic sentences*, like Liar sentences, can always be expanded. Consequently, all that can be recovered from the CONTEXTUALIST ARGUMENT FROM PARADOX is the following:

[The LIAR IN CONTEXT] shows that the interpretation of *some sentences* in context c is not maximally general: there are sentences which are true or false in c' but neither true nor false in c . Therefore, one cannot provide a maximal interpretation for the truth predicate *when it is applied to such sentences* (e.g. λ_c).

Orthodox contextualists—and relativists working in the tradition of Russell and Zermelo—are right to see problematic sentences as sentences that can be used to *expand* any interpretation of the truth predicate *as it applies to such sentences*. But it

simply does not follow from this that the interpretation of *every sentence* (including, e.g., $\text{Tr}_\alpha(\ulcorner 0 = 0 \urcorner)$ and $\forall x(x = x)$) can similarly be expanded.

5 Objections and replies

I will now discuss some potential objections. I will address the worry that bicontextualism suffers from severe expressive limitations (§5.1), the alleged non-uniqueness of the interpretation for unproblematic sentences (§5.2), worries related to the object-language/meta-language distinction (§5.3), and revenge paradoxes (§5.4).

5.1 Expressive limitations A first objection to bicontextualism claims that it is a *typed* theory of truth, as it provides interpretations for a hierarchy of languages. Worse still, the objection continues, in the bicontextualist setting we have a new truth predicate Tr_α for each new language \mathcal{L}_α . But, the objection concludes, natural languages feature a type-free truth predicate, and therefore no typed theory is applicable to or relevant for them.

However, such an objection would be misguided. The hierarchy of languages in \mathcal{L} is a mere technical expedient aimed at offering a simple presentation of the context shift that takes place in the interpretation of problematic sentences. First, all the languages in \mathcal{L} share all the non-truth-theoretic vocabulary. Second, every language in \mathcal{L} allows for iterated truth-predications and self-referential sentences, and is interpreted in a completely uniform way—the construction of Abs_α and Rel_α is always the same, across the hierarchy of languages.

A related worry is that infinitary generalizations and blind ascriptions cannot in general be modeled in a hierarchical theory, as they might be paradoxical. Suppose Bach and Telemann only utter, respectively, the following sentences:

- (9) Everything Telemann says about Händel is true;
- (10) Everything Bach says about Händel is not true.

Kripke [45, p. 695-696] used sentences such as (9) and (10) to argue for the inadequacy of Tarskian hierarchical approaches, and one might worry that they show contextualist and bicontextualist approaches to be similarly inadequate. However, Glanzberg [28, p. 233] correctly observes that Kripke's objection does not apply to his version of contextualism, and his defense covers bicontextualism as well. The reason, in short, is that both versions of contextualism employ iterations of Kripkean fixed points, and thus can easily interpret (9) and (10) in any such fixed point. More explicitly, applying the construction of §3.4 to (9) and (10), it's clear that neither of them is true or false in \mathcal{L}_α , while (10) is true in $\mathcal{L}_{\alpha+1}$ and (9) is false, as expected.

Finally, one could object that the distinction between problematic and unproblematic sentences is not robust, since *contingent Liars* (such as the pair (9)-(10) above) show that it might be contingent whether a sentence is problematic (see Kripke [45, p. 696 and following]). The point is not well-taken, however. Essentially every semantic theory of truth has to interpret contingent Liars depending on contingent facts, but this does not make the problematic/unproblematic distinction ill-defined or unusable. Given any relevant collection of contingent facts that determines the extension of the non-semantic fragment of the language, bicontextualism, just like any other semantic theory of truth, categorizes the resulting sentence as problematic or unproblematic.

5.2 Non-uniqueness One might worry that bicontextualism doesn't provide the treatment of unproblematic sentences we were originally after. In §2.3, the worry continues, orthodox contextualism was criticized for not providing a unique, maximal interpretation of unproblematic sentences, but bicontextualism also provides an unending succession of distinct interpretations $\text{Abs}_0, \text{Abs}_1, \text{Abs}_2, \dots$, and not a single one. Where is, therefore, the advantage of bicontextualism with respect to orthodox contextualism?

The worry is legitimate but misguided, for at least two reasons:

- (i) Unproblematic truth-predications receive a maximal interpretation: once they are declared to be absolutely true or untrue (or false), they remain so, and their interpretation is never altered (as required in §2.3). To see this more precisely, assume $\beta < \alpha$. Therefore, since $\text{Abs}_\beta \subseteq \text{Abs}_\alpha$, by the fact that every Abs_β is the extension of a consistent Kripkean fixed point, we have that:
 - (a) If $\varphi \in \text{Abs}_\beta$, then $\text{Tr}_\beta(\ulcorner \varphi \urcorner) \in \text{Abs}_\beta \subseteq \text{Abs}_\alpha$, and $\text{Tr}_\alpha(\ulcorner \varphi \urcorner) \in \text{Abs}_\alpha$.
 - (b) If $\varphi \in \text{Abs}_\beta$, then $\neg \text{Tr}_\beta(\ulcorner \varphi \urcorner) \notin \text{Abs}_\beta$, $\neg \text{Tr}_\beta(\ulcorner \varphi \urcorner) \notin \text{Abs}_\alpha$, and $\neg \text{Tr}_\alpha(\ulcorner \varphi \urcorner) \notin \text{Abs}_\alpha$.⁵⁰
 (a) and (b) entail that absolute truth is preserved upwards across the languages in \mathcal{L} , and no unproblematic sentence can be declared to be true $_\beta$ first and untrue $_\alpha$ (or false $_\alpha$) later, or untrue $_\beta$ (or false $_\beta$) first and true $_\alpha$ later.
- (ii) Unproblematic quantified sentences can always be interpreted over the same, maximally general RU-domain. Setting X in the definition of every Abs_α to be absolutely general, every Abs_α is defined over the same RU-domain, i.e. absolutely everything (again, as required in §2.3).

The situation is very different for the problematic, relatively general sentences: as we have seen, truth-predications in $\text{Rel}_0, \text{Rel}_1, \text{Rel}_2, \dots$ receive two distinct interpretations. For $\beta < \alpha$, λ_β is declared untrue $_\beta$ first, and true $_\alpha$ later (against (i) above). Moreover, quantified sentences are interpreted over always different, never-maximal, ever-growing domains (against (ii) above). This picture is in line with the core idea of bicontextualism: unproblematic sentences are interpreted «once and for all», in a maximally general way (a maximally general interpretation for the truth predicate, and a maximally general domain of quantification), while problematic sentences undergo the contextualist treatment described in §2.1 and §3.4.

On a diagnostic level, this worry is structurally similar to the ones in §5.1: it takes a technical (and presentational) aspect of bicontextualism—namely that it features more than one language and, consequently, more than one collection Abs_α —to distort and invalidate its basic intuition concerning absolutely general and relatively general truths. We have seen, however, that this is not the case.

5.3 The object-language/meta-language distinction Bicontextualism employs a higher-order meta-theory to provide a semantics for first-order languages. But the second-order language employed in the meta-theory is not given a bicontextualist interpretation. Indeed, it is given no interpretation at all. Does this show that bicontextualist semantics is expressively incomplete?⁵¹

A first avenue of reply for the bicontextualist is to insist that:

- (i) natural languages do not feature *genuine* higher-order quantification, and therefore a semantics for an (idealized version of a) natural language should not be concerned with higher-order sentences;⁵²

- (ii) the semantics can legitimately make a *purely instrumental* use of higher-order quantification in the meta-theory, to provide an absolutist interpretation for a fragment of the target object-language, which is first-order.

Relatively to (i), by *genuine* higher-order quantification I mean a quantification into predicate position that is not paraphrased away as a first-order quantification. For instance, consider $\forall X \forall x (X(x))$. A literal reading of $\forall X \forall x (X(x))$ would be nonsensical, as it would amount to something like «every X X s every x », which is not an English sentence. That is, the X would need to simultaneously be the syntactic object to which the quantifier applies *and* the predicate in the expression following the quantifier. Yet this doesn't seem possible in languages such as English. To be sure, plenty of English paraphrases of $\forall X \forall x (X(x))$ are available, e.g. «everything has every property». However, this paraphrase is not genuine, as it effectively treats X as a first-order variable, and not as a predicate.⁵³

One might object that, even conceding (i), (ii) is not satisfactory: how can one plausibly make use, even instrumentally, of a meta-theory one cannot express in English, and which has no interpretation? To address this further objection, the bicontextualist might follow Williamson [85] and argue that even if higher-order quantification cannot be expressed in natural languages, it can still be understood by the «direct method». As Williamson puts it:

Perhaps no reading in a natural language of quantification into predicate position is wholly satisfactory. If so, that does not show that something is wrong with quantification into predicate position, for it may reflect an expressive inadequacy in natural languages. We may have to learn second-order languages by the direct method, not by translating them into a language with which we are already familiar. [85, p. 459]

In conclusion, the objection that bicontextualism is an incomplete semantics because it does not interpret the higher-order meta-language is misplaced: bicontextualism is an intelligible and workable semantics for (fragments of) natural languages even in absence of such an interpretation.⁵⁴

5.4 Revenge Can one formulate revenge paradoxes for bicontextualism? The answer, in a nutshell, is: yes, but revenge sentences are treated exactly as the standard problematic sentences. Therefore, bicontextualism is not subject to «new» revenge paradoxes: it stands or falls on its treatment of standard paradoxes.

Consider Parsons's «Superliar» (1974): a sentence λ_A which, roughly, says that for every α , it is not true $_{\alpha}$. Then, λ_A intuitively says that λ_A is not true in absolutely any context. Therefore, one can argue that λ_A cannot be treated as bicontextualism prescribes, i.e. as being neither true nor false in one context, and being true in another because, when one establishes λ_A in the course of the LIAR IN CONTEXT, one proves that λ_A is neither true nor false in *absolutely any* context. How can then λ_A turn out to be true in some context after all?

However, on a bicontextualist view, one cannot simply take for granted that λ_A is interpreted as actually saying that λ_A is neither true nor false in absolutely any context. On a bicontextualist view, whether λ_A can be interpreted unrestrictedly depends on whether λ_A is problematic or not. But it quite clearly is. Indeed, λ_A is in the gap of the least Kripkean fixed point for its language. Appearances to the contrary, λ_A effectively fails to talk about absolutely every context, and is interpreted

in a restricted way. Bicontextualism treats λ_A in its relativistic, context-shifting part, exactly as it does with the original Liar sentences.

To be sure, the would-be revenger will object that this simply misconstrues λ_A . After all, λ_A is designed to say precisely that it is neither true nor false in *absolutely any* context. However, the objection is misplaced: what a sentence effectively says is determined by its interpretation, not merely by its superficial form. And, on a bicontextualist semantics, λ_A is simply *not* interpreted in an absolutely unrestricted way. To be sure, one can *insist* that λ_A be interpreted unrestrictedly. But this would be like insisting, in the face of Tarski's Theorem, that truth be naïve even if logic is classical.

The above arguments and replies clearly do not show that bicontextualism is free from expressive limitations. Quite the contrary: the cases of λ_α and λ_A show that interpreting some sentences as absolutely unrestricted is precluded to a bicontextualist semantics. But what ultimately matters for bicontextualism as a theory of truth is that it does not suffer from «new» expressive limitations, i.e. revenge paradoxes which would undermine its applicability. No theory of truth is free from expressive limitations. The best a theory can hope for is to address all semantic paradoxes, standard and revenge alike, *uniformly*. Bicontextualism arguably satisfies this requirement.

6 Concluding remarks

I opened this paper asking whether one can quantify over absolutely everything. Absolutists answer «yes», relativists answer «no». In the framework of semantic theories of truth, bicontextualists offer a more nuanced view, and answer «sometimes».

Absolutists and relativists disagree on fundamental issues about quantification, yet the strengths of their respective theories concern distinct kinds of sentences: absolutists can do justice to the semantics of utterly general truths, while relativists can accurately capture the more elusive meaning of paradoxical sentences, and model paradoxical reasonings. Absolutism and relativism are incompatible, but their crucial insights are not in tension: they can be harmonized in a comprehensive view, without generating «new» revenge paradoxes. At the same time, bicontextualism is not merely the result of putting the best of absolutism and relativism together. It is a distinct and coherent view, with implications for the notions of interpretation and domain, and therefore for the foundations of semantics more generally.

Notes

1. This quick gloss only is intended to informally introduce absolutism and relativism, not to fully characterize them. For absolutism, see e.g. Lewis [50], McGee [57], Rayo and Uzquiano [66], Rayo and Williamson [68], Williamson [85]; for relativism, see e.g. Hellman [38], Linnebo [52], Button [6], Studd [83]. For intermediate or hybrid positions between absolutism and relativism, see e.g. Williamson [87], Shapiro [76], Fine [18], Linnebo [53, 54], Uzquiano [84]. The view I defend here is also a hybrid one, albeit of a quite different kind.
2. See also Russell [74], Dummett [12], Simmons [78], Dummett [13], Simmons [79], Giaquinto [23, Ch. 6.2], Studd [83, Ch. 2.2-2.3].

3. Cantor [9]; see also Hallett [36], Lavine [46, Chs. III.4, IV, and V], and Studd [83, Ch. 7.3].
4. See Parsons [63], Glanzberg [24, 25, 26, 27, 28]. More details, and references, are provided in §2.
5. I intentionally omitted to specify what such contextual element consists in, and how are contexts identified, since different contextualist theories adopt different proposals. However, the structural features of contextualism can be presented without committing oneself to any specific view of the contextual elements in the Liar reasoning (more on this in §2). Moreover, contextualism is compatible with essentially every formal representation of contexts [e.g. 49, 44, 81, 82]. For example, if contexts are thought to include propositions, and propositions are taken as truth-bearers, one could argue (as Glanzberg [24] does) that it is the proposition expressed by λ which is lacking in the first context but not in the second.
6. There is a clear analogy between the contextualist interpretation of semantic paradoxes and the relativist interpretation of set-theoretic paradoxes (see e.g. Dummett [12, 13] and Studd [83, Ch. 4.5 and 7]). The analogy becomes even more evident if, following Glanzberg [24, 25], a Liar sentence is construed as a sentence λ_e which says that there is no true proposition expressed by λ_e , and the context shift is interpreted as an expansion in the domain of the existential quantifier in λ_e . In this way, one establishes that λ_e does not express a true proposition when the quantifier in λ_e is interpreted over the starting domain, and that λ_e expresses a true proposition when the quantifier in λ_e is interpreted over another domain. I did not follow Glanzberg's presentation here because it departs from the usual formulation of Liar sentences (by treating them as quantified statements). However, as it will become clear in §2, such differences are inessential: the contextualist construal of the Liar reasoning can be rendered equally well using λ or λ_e , and (more generally) taking sentences or propositions as truth-bearers.
7. A precise characterization of «problematic» and «unproblematic» sentences is offered in §§3.3-3.4.
8. See Glanzberg [28] and Glanzberg and Rossi [29].
9. Other varieties of contextualism include Burge [5], Barwise and Etchemendy [2], Simmons [79], Gaifman [21], Simmons [80]. I focus on the Parsons-Glanzberg tradition because it lends itself easily to the modifications I propose here.
10. «The view about the Liar paradox here presented has to meet two objections: first, that it presupposes the dubious notion of sentences as expressing propositions, and of propositions as the primary bearers of truth or falsity; [...]. [...] the approach can also be formulated in the situation where truth-values are attributed to sentences as I shall do in the next section» [63, p. 391].
11. For more on this, see Glanzberg [28], especially footnote 27 and *infra*.
12. If one further assumes that a sentence expresses at most one proposition in a given context, then « $\ulcorner \varphi \urcorner$ is true in c » is interpreted as « $\ulcorner \varphi \urcorner$ expresses exactly one true proposition in c ». See Parsons [63, p. 392 and following] for more details on the switch between the sentential and the propositional settings.

13. The Liar Paradox is often presented making use of the so-called T-SCHEMA, i.e.: $\varphi \leftrightarrow \langle \ulcorner \varphi \urcorner \text{ is true} \rangle$, but it is more convenient to use the corresponding inference rules in the present setting. See Field [16, Ch. 13] and Murzi and Rossi [61] for more details on the differences between the T-SCHEMA, the truth rules, and other principles for naïve truth. This presentation of the LIAR IN CONTEXT follows Glanzberg [25].
14. The question of what triggers and explains the context-shift, and where exactly such a shift takes place in paradoxical derivations, is a matter of debate in the literature. I follow Murzi and Rossi [59] in identifying the truth-introduction rule as the context-shifting principle. For more details and discussion, see Glanzberg [25, especially pp. 33-4], Gauker [22], and Mankowitz [55].
15. Again, the point can also be made in terms of propositions. In terms of propositions, (5) means that λ_c does not express a proposition in c , and (7) means that λ_c expresses a proposition in c' . So, there are propositions not quantified over in c but quantified over in c' . Therefore, the quantification in c was not maximally general. But c is arbitrary and the LIAR IN CONTEXT can be performed in any given context, so one cannot quantify over absolutely all propositions. *A fortiori*, one cannot quantify over absolutely everything.
16. Glanzberg's theory is a rich and complex one, and I must limit myself to sketching it. For reasons of space, I will also *not* explicitly advocate Glanzberg's own brand of contextualism (for that, I refer the reader to Glanzberg's own papers), and simply take it as the starting point to articulate my own view.
17. Kripke's construction does not provide, by itself, a context-relative interpretation of the truth predicate. For this reason, when presenting Kripke's theory, I simply write $\langle \ulcorner \varphi \urcorner \text{ is true} \rangle$, rather than $\langle \ulcorner \varphi \urcorner \text{ is true in } c \rangle$. Later, *iterations* of Kripke's theory will be used to model context-relative truth predications.
18. In Kripke's original construction, not a single set E but a pair $\langle E, A \rangle$ is constructed, where A is the anti-extension of the truth predicate, i.e. the set of (names of) false sentences, i.e. sentences whose negation is true. By contrast, here I only define E, but its definition also specifies conditions under which negations are true. Thus, A can be immediately obtained as the set of (names of) sentences φ s.t. $\neg\varphi$ is in E. Nothing crucial hinges on this, but this presentation considerably simplifies the construction in §3.3.2. For more details on this and other variants of Kripke's construction, see Halbach [32, Ch. 15].
19. For a presentation of the strong Kleene scheme, and a fuller presentation of Kripke's theory, see, e.g., McGee [56, Ch. 4], Field [16, Ch. 3], Horsten [41, Ch. 9], and Halbach [32, Ch. 15].
20. I.e., because there are more ordinal stages than sets of sentences that can be interpreted in this way.
21. For the criterion of scientific applicability of theories of truth, see Leitgeb [48]. One widely discussed example of «loss of content» concerns theories of truth formulated in a non-classical logic over a mathematical base theory. In some such theories, one can show that less theorems are provable than in the same base theory formulated in classical logic. The theory Partial Kripke-Feferman (PKF) formulated over Peano Arithmetic (PA) in a non-classical logic is a case in point [35]. In the arithmetical context, a precise measure

of the loss of mathematical content is available, in terms of proof-theoretic ordinals, i.e. ordinals that measure the amount of transfinite induction that a theory can prove. Halbach and Horsten [35] have shown that the proof-theoretic ordinal of PKF over PA in the target non-classical logic is much smaller than the proof-theoretic ordinal of PA in classical logic. In other words, the former theory proves less arithmetical theorems than the latter and therefore, one can argue, some «mathematical content» is lost. For more on the loss of mathematical content in non-classical theories of truth, see [33, 17]. Thanks to an anonymous referee for helpful comments on this point.

22. Delia Graff Fara [30, p. 50] has famously argued that a convincing solution to a paradox should explain why the paradox is intuitively compelling. Graff Fara's discussion concerns the Sorites Paradox, but there is no reason not to extend her requirement of *psychological adequacy* to semantic paradoxes.
23. See Horsten [41, Ch. 9.3], Field [16, Ch. 3.3-3.4], and Sagi [75] for similar arguments.
24. See also Juhl [43], Bacon [1, pp. 312-313] and Murzi and Rossi [59] (§5).
25. In order to present the truth-theoretical axioms of KF, I follow Halbach [32, Chs. 5 and 7]. \mathcal{L} is some (suitable) language for which the axioms are given. «Sent» and «CTer» are formulae of \mathcal{L} denoting the class of sentences and closed terms, respectively; a dot under a logical operator represents the corresponding primitive recursive operations (on names of formulae). For details on the encoding of syntax, see Feferman [14]. $\forall t \varphi(t)$ abbreviates $\forall x(\text{CTer}_{\mathcal{L}}(x) \rightarrow \varphi(x))$. Moreover, $\text{Tr}(\ulcorner \varphi(t/x) \urcorner)$ indicates that the result of substituting all free occurrences of x with t in $\varphi(x)$ is true, where x is the only free variable in $\varphi(x)$. Finally, t° indicates the value of the term t , i.e. the element denoted by t (a number, if t is a numeral, and \mathcal{L} is the language of arithmetic). As Picollo [65] points out, the function symbol $^\circ$ cannot be part of any language that satisfies the premises of Strong Diagonalization, on pain of triviality, but (as customary) I write it as an object-linguistic function symbol to preserve readability. In its original formulation, KF is not a contextualist theory, so its axioms concern truth *simpliciter* rather than truth in context. (Notice that one instance of the first two axioms is needed for every predicate of the base language \mathcal{L}):
 - (KF1) $\forall t_1, \dots, \forall t_n [\text{Tr}(R(t_1, \dots, t_n)) \leftrightarrow R(t_1^\circ, \dots, t_n^\circ)]$
 - (KF2) $\forall t_1, \dots, \forall t_n [\text{Tr}(\neg R(t_1, \dots, t_n)) \leftrightarrow \neg R(t_1^\circ, \dots, t_n^\circ)]$
 - (KF3) $\forall x [\text{Sent}_{\mathcal{L}_{\text{Tr}}}(x) \rightarrow (\text{Tr}(\neg \neg x) \leftrightarrow \text{Tr}(x))]$
 - (KF4) $\forall x [\text{Sent}_{\mathcal{L}_{\text{Tr}}}(x \wedge y) \rightarrow (\text{Tr}(x \wedge y) \leftrightarrow \text{Tr}(x) \wedge \text{Tr}(y))]$
 - (KF5) $\forall x [\text{Sent}_{\mathcal{L}_{\text{Tr}}}(x \wedge y) \rightarrow (\text{Tr}(\neg(x \wedge y)) \leftrightarrow \text{Tr}(\neg x) \vee \text{Tr}(\neg y))]$
 - (KF6) $\forall x \forall y [\text{Sent}_{\mathcal{L}_{\text{Tr}}}(\forall xy) \rightarrow (\text{Tr}(\forall xy) \leftrightarrow \forall t (\text{Tr}(y(t/x))))]$
 - (KF7) $\forall x \forall y [\text{Sent}_{\mathcal{L}_{\text{Tr}}}(\forall xy) \rightarrow (\text{Tr}(\neg \forall xy) \leftrightarrow \exists t (\text{Tr}(\neg y(t/x))))]$
 - (KF8) $\forall t [\text{Tr}(\ulcorner t \urcorner) \leftrightarrow \text{Tr}(t^\circ)]$
 - (KF9) $\forall t [\text{Tr}(\neg \ulcorner t \urcorner) \leftrightarrow \text{Tr}(\neg t^\circ) \vee \neg \text{Sent}_{\mathcal{L}_{\text{Tr}}}(t^\circ)]$

Similar axioms can be stated for \vee , \rightarrow , \leftrightarrow , and \exists .

26. See Halbach [32, Ch. 15], and Field [16, Chs. 7, 13] for more on Tr-Intro and Tr-Elim in classical theories.
27. Bicontextualism is developed in detail for mathematical languages (which however can model paradoxical phenomena). In §5.1, I sketch how bicontextualism can address sentences whose (un)problematic status depends on contingent factors («contingent Liars»).

28. See Murzi and Rossi [60] and Rosenblatt [71, 72] for classical recapture (and its problems) in non-classical logics, and Reinhardt [70], Bacon [1] for the classical case.
29. For a theory which identifies «problematic» sentences and their properties, see Rossi [73].
30. For more details on dependence and grounding, see Yablo [88], Leitgeb [47], Beringer and Schindler [3], Rossi [73], Picollo [64]. Which sentences count as grounded also depends on the evaluation scheme employed in Kripke's construction. Under the strong Kleene scheme (which we have been adopting), λ is ungrounded, while $\lambda \vee 0 = 0$ and $\neg(\lambda \wedge 0 = 1)$ are grounded, as they are in the extension of the least Kripkean fixed point. This is because a true disjunct and a false conjunct suffice to make a disjunction true and a conjunction false (respectively) in strong Kleene logic. By contrast, $\lambda \vee 0 = 1$ and $\lambda \wedge 0 = 0$ are ungrounded. And this is because a disjunction with a false and an undefined disjunct is undefined in strong Kleene logic, as is a conjunction with a true and an undefined conjunct. However, Kripke's construction can be developed with other evaluation schemes too. In the weak Kleene scheme, for instance, all of λ , $\lambda \vee 0 = 0$, $\lambda \wedge 0 = 1$, $\lambda \vee 0 = 1$, and $\lambda \wedge 0 = 0$ would count as ungrounded [8]. I do not discuss these variants in the interest of space.
31. Much of the mathematical complexity of Glanzberg's original construction derives from the need to re-interpret the same language. Availing oneself to multiple languages helps avoiding some of that complexity.
32. A full definition of acceptable structure is provided in Moschovakis [58, Ch. 5]. There's a minor complication here. Moschovakis's original notion of acceptability requires \mathcal{M}_α to have an isomorphic copy of the natural numbers in the standard signature $\{0, S, +, \cdot\}$ as a substructure. Since isomorphisms are only defined for structures in the same signature, the signature of \mathcal{M}_α (i.e. \mathcal{L}_α) has to include $\{0, S, +, \cdot\}$. However, since we want to avoid languages with function constants, we have to employ a variant of Moschovakis's original notion, as spelled out, for example, in McGee [56], Ch. 1, where the natural number structure is only required to be *interpretable* in \mathcal{M}_α . See Hodges [40, Ch. 5] and Button and Walsh [7, Ch. 5] for more details the notion of interpretability. As such variant notion is easily specified model-theoretically, I will implicitly assume it.
33. Working axiomatically, the form of diagonalization involving the provable identity between the terms t_φ and $\ulcorner \varphi(t_\varphi) \urcorner$ is sometimes referred to as «strong diagonalization» (and contrasted with «weak diagonalization», which involves the material equivalence between ψ and $\varphi(\ulcorner \psi \urcorner)$). The corresponding Strong Diagonalization Lemma is due to Jeroslow [42]. For more discussion, see Heck [37], Picollo [64]. The approach I adopt here is model-theoretic: rather than adopting a base theory and proving a diagonalization lemma for it, I consider general requirements on languages and structures, which include the availability of strong diagonalization.
34. Hierarchies of languages similar to \mathfrak{L} are studied in Halbach [34, 31] and Halbach [32, Ch. 9]
35. Essentially, such codings are ZFC-definable functions (that are absolute w.r.t. transitive models) from \mathcal{L}_α -expressions to structures of the form $\langle H(\kappa), \in|_{H(\kappa)} \rangle$, where $H(\kappa)$ is the set of all sets hereditarily of cardinality $< \kappa$, and $\in|_{H(\kappa)}$ is the membership relation restricted to it. Constructions of codings for languages of high cardinality can be found

in Dickmann [11]. Such codings are employed in the construction of semantic theories of truth in Glanzberg and Rossi [29].

36. Rayo and Uzquiano use the informal paraphrase « X is a model», but its surface form should be taken with some care. X is a second-order variable, whereas in « X is a model» it occurs in the scope of the predicate «is a model», i.e. as taking the place of an individual. But only first-order variables take the place of individuals: second-order variables take the place of predicates. I follow them (e.g. when writing « X is an RU-model»), with the proviso that this grammatical form should not be given much weight.
37. The following definition follows Rayo and Williamson [68], which adapt the RU-models of Rayo and Uzquiano [66] to an arbitrary first-order language, such as \mathcal{L}_α .
38. Standardly, models do not specify the denotation of individual variables—*assignments* (over a given model) do that. Here I follow Rayo and Williamson [68] and take models to provide the values of variables too. This minor change comes at no cost, but substantially simplifies the formulation of some clauses in Definition 3.7.
39. X is a second-order variable. So, its values are not single individuals. By analogy, consider the formula « $Y(\text{John})$ ». Two possible values of Y are the individuals who eat, and the individuals who run, so that « $Y(\text{John})$ » will be interpreted as «John eats» and «John runs» respectively.
40. The above paraphrases, with their reifying talk of «domain», «extension», and the like, improperly (if inevitably) gloss higher-order quantification as first-order quantification. Taking the talk about, say, domains literally in an absolutist higher-order setting is problematic, for it seemingly presupposes that such «domains» are *themselves objects*—be they sets, proper classes, what have you. See Cartwright [10] and Studd [83, §1.4]. Rayo and Williamson [68, p. 3] make a related point: «Such informal explanations are a kind of useful nonsense, a ladder to be thrown away once climbed, because they use the second-order (predicate) variable [« X »] in first-order (name) positions in sentences of natural language; nevertheless, they draw attention to helpful analogies between [second-order]-interpretations and [model-theoretic]-interpretations.»
41. I require that the RU-domain Z is unique in order to ensure that the variants defined on the RU-models X and Y are defined on the same domain. However, since RU-domains are extensional, one could define identity between them as is standardly done in second-order logic (i.e. letting $Z_1 = Z_2 \leftrightarrow \forall x(Z_1(x) \leftrightarrow Z_2(x))$), and require that the RU-domains of the RU-models X and Y are identical.
42. Again, we are restricting ourselves to the strong Kleene version of Kripke’s construction, so «Kripke-satisfies» should be understood as «Kripke-Kleene satisfies».
43. For the sake of readability, I omit the requirement that x is a sentence of \mathcal{L}_α .
44. Non-minimal fixed points will be used in §3.4 to interpret the problematic \mathcal{L}_α -sentences, as outlined in §2.1.
45. I employ set-sized domains in order to make the contrast between the quantifier domains of problematic and unproblematic sentences vivid, but one could employ the very apparatus of bicontextualism to provide a Glanzberg-style contextualist theory where

the interpretation of the truth predicate shifts (when it comes to problematic sentences), and yet the quantifiers are still ranging over non-set-sized collections, including proper classes. Thanks to an anonymous referee for helpful comments on this point.

46. Recall that unproblematic sentences are interpreted in non-closed-off fixed points because they require no context shift (and the context shift, as per Glanzberg's original treatment, is modeled by the closing off).
47. For the notion of elementary extension, see Hodges [40, Ch. 2].
48. That the consequence relation of Definition 3.14 is classical follows from the construction of Abs_α and Rel_α : the inner logic of both sets of sentences is classical logic. One might worry that, since the base acceptable RU-model X and the base acceptable model \mathcal{M}_α employed in defining Abs_α and Rel_α are kept fixed, Definition 3.14 does not provide a notion of validity. This is largely a terminological issue, but can easily be remedied by considering a variant of Definition 3.14 where X is kept fixed (and is maximally general), and one quantifies over all the acceptable \mathcal{M}_α s used in defining Rel_α .
49. I owe this suggestion to Hannes Leitgeb.
50. The same holds replacing claims formalizing untruth (e.g. $\neg \text{Tr}_\beta(\ulcorner \varphi \urcorner)$) with corresponding claims formalizing falsity ($\text{Tr}_\beta(\ulcorner \neg \varphi \urcorner)$).
51. The question how to interpret higher-order languages is a deep and ramified issue, and cannot be exhaustively addressed here. For discussion, see Boolos [4], Shapiro [77], Linnebo [51], Williamson [85, 86], Rayo and Uzquiano [67], Rayo [69], Oliver and Smiley [62], Florio [19], Florio and Linnebo [20].
52. Of course, one could further object that the bicontextualist's meta-theory can be given using *plural logic* rather than second-order logic, which would put more pressure on the need to interpret the meta-language.
53. See Higginbotham [39] for further discussion. Thanks to Salvatore Florio, Michael Glanzberg, Julien Murzi, Simon Schmitt, and Brett Topey for fruitful discussion on this issue.
54. An alternative line of defense for the bicontextualist would be to insist that the second-order meta-language needs no semantics at all, because the meta-theoretic notions employed are sufficiently well characterized *inferentially*, i.e. by the axioms and rules of second-order logic and second-order ZF implicitly adopted in the meta-theory. A version of this position in the context of arithmetic, set theory, and model theory itself is articulated by Button and Walsh [7, Chs. 10-12].

References

- [1] Bacon, A., "Can the classical logician avoid the revenge paradoxes?," *Philosophical Review*, vol. 124 (2015), pp. 299–352. [25](#), [26](#)
- [2] Barwise, J., and J. Etchemendy, *The Liar: An Essay on Truth and Circularity*, Oxford University Press, Oxford, 1987. [23](#)

- [3] Beringer, T., and T. Schindler, “A graph-theoretic analysis of the semantic paradoxes,” *The Bulletin of Symbolic Logic*, vol. 23 (2017), pp. 442–492. [26](#)
- [4] Boolos, G., “To be is to be a value of a variable (or to be some values of some variables),” *Journal of Philosophy*, vol. 81 (1984), pp. 430–50. [28](#)
- [5] Burge, T., “Semantical paradox,” *Journal of Philosophy*, vol. 76 (1979), pp. 169–98. [23](#)
- [6] Button, T., “Dadaism: restrictivism as militant quietism,” *Proceedings of the Aristotelian Society*, vol. 110 (2010), pp. 387–98. [22](#)
- [7] Button, T., and S. Walsh, *Philosophy and Model Theory*, Oxford University Press, Oxford, 2018. [26](#), [28](#)
- [8] Cain, J., and Z. Damjanovic, “On the weak Kleene scheme in Kripke’s theory of truth,” *The Journal of symbolic logic*, vol. 56 (1991), pp. 1452–1468. [26](#)
- [9] Cantor, G., *Grundlagen einer allgemeinen Mannigfaltigkeitslehre. Ein mathematischer-philosophischer Versuch in der Lehre des Unendlichen*, Teubner, Leipzig, 1883. [23](#)
- [10] Cartwright, R., “Speaking of everything,” *Nous*, vol. 28 (1994), pp. 1–20. [27](#)
- [11] Dickmann, M. A., *Large infinitary languages: model theory*, volume 83, North-Holland Amsterdam, 1975. [27](#)
- [12] Dummett, M., *Frege: Philosophy of Language*, second edition, Duckworth, London, 1983. [22](#), [23](#)
- [13] Dummett, M., *Frege: Philosophy of Mathematics*, Duckworth, London, 1991. [22](#), [23](#)
- [14] Feferman, S., “Arithmetization of metamathematics in a general setting,” *Fundamenta Mathematicae*, vol. 49 (1960), pp. 35–92. [25](#)
- [15] Feferman, S., “Reflecting on Incompleteness,” *Journal of Symbolic Logic*, vol. 56 (1991), pp. 1–49. [7](#)
- [16] Field, H., *Saving Truth from Paradox*, Oxford University Press, Oxford, 2008. [24](#), [25](#)
- [17] Field, H., “The power of naive truth,” *The Review of Symbolic Logic*, (2020), pp. 1–34. [25](#)
- [18] Fine, K., “Relatively unrestricted quantification,” pp. 20–44 in *Absolute Generality*, edited by A. Rayo and G. Uzquiano, Oxford University Press, 2006. [22](#)
- [19] Florio, S., “Unrestricted quantification,” *Philosophy Compass*, vol. 9 (2014), pp. 441–454. [28](#)
- [20] Florio, S., and O. Linnebo, “On the innocence and determinacy of plural quantification,” *Noûs*, vol. 50 (2016), pp. 565–583. [28](#)
- [21] Gaifman, H., “Pointers to truth,” *Journal of Philosophy*, vol. 89 (1992), pp. 223–261. [23](#)
- [22] Gauker, C., “Against stepping back: A critique of contextualist approaches to the semantic paradoxes,” *Journal of Philosophical Logic*, vol. 35 (2006), pp. 393–422. [24](#)
- [23] Giaquinto, M., *The search for certainty*, Clarendon Press, Oxford, 2002. [22](#)

- [24] Glanzberg, M., “The liar in context,” *Philosophical Studies*, vol. 103 (2001), pp. 217–51. [3](#), [23](#)
- [25] Glanzberg, M., “A contextual-hierarchical approach to truth and the liar paradox,” *Journal of Philosophical Logic*, vol. 33 (2004), pp. 27–88. [3](#), [4](#), [7](#), [23](#), [24](#)
- [26] Glanzberg, M., “Truth, reflection, and hierarchies,” *Synthese*, vol. 142 (2004), pp. 289–315. [3](#), [23](#)
- [27] Glanzberg, M., “Context and unrestricted quantification,” pp. 45–74 in *Absolute Generality*, edited by Rayo, A. and Uzquiano, G., Oxford University Press, Oxford, 2006. [3](#), [23](#)
- [28] Glanzberg, M., “Complexity and Hierarchy in Truth Predicates,” in *Unifying the Philosophy of Truth*, edited by Achourioti, T. and Galinon, H. and Martínez Fernández, J. and Fujimoto, K., volume 36 of *Logic, Epistemology, and the Unity of Science*, Springer, 2015. [19](#), [23](#)
- [29] Glanzberg, M. and Rossi, L., “Truth and Quantification,” in preparation. [23](#), [27](#)
- [30] Graff Fara, D., “Shifting sands: an interest-relative theory of vagueness,” *Philosophical Topics*, vol. 28 (2000), pp. 44–81. [25](#)
- [31] Halbach, V., “Tarskian and Kripkean Truth,” *Journal of Philosophical Logic*, vol. 26 (1997), pp. 69–80. [26](#)
- [32] Halbach, V., *Axiomatic Theories of Truth*, Cambridge University Press, Cambridge, 2011. [13](#), [24](#), [25](#), [26](#)
- [33] Halbach, V., and C. Nicolai, “On the costs of nonclassical logic,” *Journal of Philosophical Logic*, vol. 47 (2018), pp. 227–257. [25](#)
- [34] Halbach, V., “Tarski hierarchies,” *Erkenntnis*, vol. 43 (1995), pp. 339–367. [26](#)
- [35] Halbach, V., and L. Horsten, “Axiomatizing Kripke’s theory of truth,” *Journal of Symbolic Logic*, vol. 71 (2006), pp. 677–712. [24](#), [25](#)
- [36] Hallett, M., *Cantorian set theory and limitation of size*, Oxford University Press, Oxford, 1984. [23](#)
- [37] Heck, R. K., “Self-reference and the languages of arithmetic,” *Philosophia Mathematica*, vol. 15 (2007), pp. 1–29. (originally published under the name ‘Richard G. Heck, Jr’). [26](#)
- [38] Hellman, G., “Against ‘absolutely everything’!,” pp. 75–97 in *Absolute Generality*, edited by Agustín Rayo and Gabriel Uzquiano, Oxford University Press, 2006. [22](#)
- [39] Higginbotham, J., “Higher-order logic and natural language,” *Proceedings of the British Academy*, vol. 95 (1998). [28](#)
- [40] Hodges, W., *Model theory*, Cambridge University Press, 1993. [26](#), [28](#)
- [41] Horsten, L., *The Tarskian Turn. Deflationism and axiomatic truth*, MIT Press, Cambridge, (Mass.), 2012. [24](#), [25](#)
- [42] Jeroslow, R. G., “Redundancies in the Hilbert-Bernays derivability conditions for Gödel’s Second Incompleteness Theorem,” *Journal of Symbolic Logic*, vol. 38 (1973),

- pp. 359–367. [26](#)
- [43] Juhl, C. F., “A context-sensitive liar,” *Analysis*, vol. 57 (1997), pp. 202–204. [25](#)
- [44] Kaplan, D., “Demonstratives,” in *Themes from Kaplan*, edited by Almog, J. and Perry, J. and Wettstein, H., Oxford: Oxford University Press, 1989. [23](#)
- [45] Kripke, S., “Outline of a theory of truth,” *Journal of Philosophy*, vol. 72 (1975), pp. 690–716. [4](#), [8](#), [19](#)
- [46] Lavine, S., *Understanding the infinite*, Harvard University Press, Cambridge (MA), 1994. [23](#)
- [47] Leitgeb, H., “What truth depends on,” *Journal of Philosophical Logic*, vol. 34 (2005), pp. 155–92. [26](#)
- [48] Leitgeb, H., “What theories of truth should be like and cannot be,” *Philosophy Compass*, vol. 2 (2007), pp. 276–290. [24](#)
- [49] Lewis, D., “Scorekeeping in a language game,” pp. 172–187 in *Semantics from different points of view*, Springer, 1979. [23](#)
- [50] Lewis, D., *Parts of Classes*, Basil Blackwell, Oxford, 1991. [22](#)
- [51] Linnebo, O., “Plural quantification exposed,” *Noûs*, vol. 37 (2003), pp. 71–92. [28](#)
- [52] Linnebo, O., “Sets, properties, and unrestricted quantification,” pp. 149–178 in *Absolute Generality*, edited by A. Rayo and G. Uzquiano, Oxford University Press, 2006. [22](#)
- [53] Linnebo, O., “Pluralities and sets,” *The Journal of Philosophy*, vol. 107 (2010), pp. 371–391. [22](#)
- [54] Linnebo, O., “The potential hierarchy of sets,” *The Review of Symbolic Logic*, vol. 6 (2013), pp. 205–228. [22](#)
- [55] Mankowitz, P., “The Liar without relativism,” *Erkenntnis*, (2021), pp. 1–22. [24](#)
- [56] McGee, V., *Truth, Vagueness, and Paradox*, Hackett Publishing Company, Indianapolis, 1991. [24](#), [26](#)
- [57] McGee, V., “Everything,” pp. 54–78 in *Between Logic and Intuition. Essays in Honor of Charles Parsons*, edited by G. Sher and R. Tieszen, Cambridge University Press, Cambridge, 2000. [22](#)
- [58] Moschovakis, Y., *Elementary Induction on Abstract Structures*, North-Holland and Elsevier, Amsterdam, London and New York, 1974. [26](#)
- [59] Murzi, J. and Rossi, L., “Reflection principles and the Liar in context,” *Philosophers’ Imprint*, vol. 15 (2018), pp. 1–18. [24](#), [25](#)
- [60] Murzi, J. and Rossi, L., “Generalized Revenge,” *Australasian Journal of Philosophy*, vol. 1 (2020), pp. 153–177. [26](#)
- [61] Murzi, J. and Rossi, L., “The expressive power of contextualist truth,” pp. 88–114 in *Modes of Truth. A unified approach to truth, modality, and paradox*, edited by C. Nicolai and J. Stern, Routledge, 2021. [24](#)

- [62] Oliver, A., and T. Smiley, *Plural Logic*, Oxford University Press, 2013. [28](#)
- [63] Parsons, C., “The Liar Paradox,” *Journal of Philosophical Logic*, vol. 3 (1974), pp. 381–412. [3](#), [7](#), [21](#), [23](#)
- [64] Picollo, L., “Alethic reference,” *Journal of Philosophical Logic*, vol. 49 (2020), pp. 417–438. [26](#)
- [65] Picollo, L., “Reference and truth,” *Journal of Philosophical Logic*, vol. 49 (2020), pp. 439–474. [25](#)
- [66] Rayo, A., and G. Uzquiano, “Toward a theory of second-order consequence,” *Notre Dame Journal of Formal Logic*, vol. 40 (1999), pp. 315–325. [10](#), [11](#), [18](#), [22](#), [27](#)
- [67] Rayo, A., and G. Uzquiano, editors, *Absolute Generality*, edited by Rayo, A. and Uzquiano, G., Oxford University Press, Oxford, 2006. [28](#)
- [68] Rayo, A., and T. Williamson, “A completeness theorem for unrestricted first-order languages,” pp. 331–356 in *Liars and Heaps*, edited by Jc Beall, Oxford University Press, 2003. [18](#), [22](#), [27](#)
- [69] Rayo, A., “Beyond plurals,” pp. 220–254 in *Absolute Generality*, edited by Agustín Rayo and Gabriel Uzquiano, Oxford University Press, Oxford, 2006. [28](#)
- [70] Reinhardt, W. N., “Some remarks on extending and interpreting theories with a partial predicate for truth,” *Journal of Philosophical Logic*, vol. 15 (1986), pp. 219–251. [26](#)
- [71] Rosenblatt, L., “Classical recapture and maximality,” *Philosophical Studies*, (2020), pp. 1–20. [26](#)
- [72] Rosenblatt, L., “Maximal non-trivial sets of instances of your least favorite logical principle,” *The Journal of Philosophy*, vol. 117 (2020), pp. 30–54. [26](#)
- [73] Rossi, L., “A unified theory of truth and paradox,” *The Review of Symbolic Logic*, vol. 12 (2019), pp. 209–254. [26](#)
- [74] Russell, B., “On some difficulties in the theory of transfinite numbers and order types,” *Proceedings of the London Mathematical Society*, vol. 4 (1906), pp. 29–53. [22](#)
- [75] Sagi, G., “Contextualism, relativism and the liar,” *Erkenntnis*, vol. 82 (2017), pp. 913–928. [25](#)
- [76] Shapiro, S., “All sets great and small: And I do mean all,” *Philosophical Perspectives*, vol. 17 (2003), pp. 467–490. [22](#)
- [77] Shapiro, S., *Foundations without Foundationalism: A Case for Second-Order Logic*, Oxford University Press, Oxford, 1991. [28](#)
- [78] Simmons, K., “The diagonal argument and the Liar,” *Journal of Philosophical Logic*, vol. 19 (1990), pp. 277–303. [22](#)
- [79] Simmons, K., *Universality and the Liar: An Essay on Truth and the Diagonal Argument*, Cambridge University Press, Cambridge, 1993. [22](#), [23](#)
- [80] Simmons, K., *Semantic Singularities: Paradoxes of Reference, Predication, and Truth*, Oxford University Press, Oxford, 2018. [23](#)

- [81] Stalnaker, R., “On the representation of context,” *Journal of logic, language and information*, vol. 7 (1998), pp. 3–19. [23](#)
- [82] Stalnaker, R., “Common ground,” *Linguistics and philosophy*, vol. 25 (2002), pp. 701–721. [23](#)
- [83] Studd, J., *Everything, more or less*, Oxford University Press, Oxford, 2019. [2](#), [22](#), [23](#), [27](#)
- [84] Uzquiano, G., “Varieties of indefinite extensibility,” *Notre Dame Journal of Formal Logic*, vol. 56 (2015), pp. 147–166. [22](#)
- [85] Williamson, T., “Everything,” *Philosophical Perspectives*, vol. 17 (2003), pp. 415–465. [2](#), [7](#), [21](#), [22](#), [28](#)
- [86] Williamson, T., *Modal logic as metaphysics*, Oxford University Press, Oxford, 2013. [28](#)
- [87] Williamson, T., “Indefinite extensibility,” *Grazer Philosophische Studien*, vol. 55 (1998), pp. 1–24. [22](#)
- [88] Yablo, S., “Grounding, dependence, and paradox,” *Journal of Philosophical Logic*, vol. 11 (1982), pp. 117–137. [26](#)
- [89] Zermelo, E., “Über Grenzzahlen und Mengenbereiche: Neue Untersuchungen über die Grundlagen der Mengenlehre,” *Fundamenta Mathematicae*, vol. 16 (1930), pp. 29–47. [1](#)

Acknowledgments

Thanks to Salvatore Florio, Michael Glanzberg, Hannes Leitgeb, Julien Murzi, Chris Scambler, Simon Schmitt, and Brett Topey for helpful comments, discussion, and suggestions on this work.

Department of Philosophy and Education
Center for Logic, Language, and Cognition (LLC)
University of Turin
via Sant’Ottavio 20
Torino 10124
Italia
lo.rossi@unito.it
<http://lorenzorossi.org/>