UNIVERSITY OF CALIFORNIA

Los Angeles

Truth in the World

A dissertation submitted in partial satisfaction

of the requirements for the degree of Doctor of Philosophy

in Philosophy

by

Torsten Odland

ABSTRACT OF THE DISSERTATION

Truth in the World

by

Torsten Odland

Doctor of Philosophy in Philosophy

University of California, Los Angeles, 2024

Professor Samuel John Cumming, Chair

For the last century, the dominant philosophical approach in the Anglophone world to the traditional question “What is truth?” has been the Deflationist answer: truth has no substantive nature beyond what is given by disquotational principles. This dissertation clears the way for a non-Deflationist account of truth by focusing on questions that seem to me more tractable: questions about truth-bearers. I motivate and defend the view I call Particularism, according to which the fundamental truth-bearers are concrete particular representations—paradigmatically, mental states like token beliefs. The dissertation is composed of relatively independent papers that develop the Particularist view and draw out its implications for debates in logic, semantics, and theories of representational content.

In Chapter 1, I present the strongest consideration in favor of Particularism: that it offers an attractive solution to the semantic paradoxes. For the Particularist, truth is properly predicated of sentence tokens, and semantic pathology arises as an illegitimate dependency structure among tokens, rather than simply being a function of an interpretation associated with a sentence type. Particularism offers a straightforward account what has been called "The Chrysippus Intuition"—the fact that it is apparently sensible to judge Liar-paradoxical objects to be not true—it is compatible with classical logic, and it preserves a conception of truth as a unified property. Crucially, unlike other broadly "Contextualist" responses to the Liar, I argue that Particularism is not threatened by Revenge Paradoxes.

In Chapter 2, I address what I take to be the deepest objection to Particularism: the view apparently conflicts with other platitudes that are part of our common sense understanding of truth. For instance, it is plausible that there are truths about the Milky Way that have and never will never be thought or stated by anyone. Whatever these truths are, it looks like they cannot be concrete particular representations. I argue that the tension is merely apparent: in addition to particular representations, we derivatively judge representational kinds to be true or false, whether they have been instantiated or not. Following Hanks, I suggest that a representational kind is true if and only if its instantiations would be true if it were instantiated, and I defend this analysis against putative counterexamples.

In Chapter 3, I offer reasons for skepticism regarding Particularism's main rival—the view that the fundamental truth-bearers are not concrete representations but, rather, the propositional content expressed by those representations. The traditional view of propositional content presupposes that there is a privileged classification of token representations that partitions them with respect to sameness of content. I argue that there is no such thing: we utilize

multiple distinct standards of sameness of content, none of which is privileged. My argument is based on the existence of semantic underdetermination. Plausibly, we use expressions (e.g. "sandwich") that do not have determinate extensions and that are open to precisification in multiple directions. This gives rise to indefinitely many possible cases of content fission, where a semantically underdetermined token is identified as having the same content as two tokens that are precisified in opposite ways. Content fission requires recognizing a huge variety of legitimate notions of sameness of content. This implies a pluralism about "propositional contents" that sits uneasily with the suggestion that propositional contents are the fundamental truth-bearers.

The dissertation of Torsten Odland is approved.

Joshua David Armstrong

David B. Kaplan

Sean Walsh

Una Stojnić

Samuel John Cumming, Committee Chair

University of California, Los Angeles

2024

For Anna and Sophie

CONTENTS

LIST OF FIGURES

ACKNOWLEDGEMENTS

A multitude of people contributed to this dissertation through their support, friendship, inspiration, questions, and criticism. (The faults remain, largely, my own.)

I must begin by thanking my committee members. Sam Cumming was a generous and thoughtful chair, and I thank him for his guidance. Sam's critical probing always pushed me to investigate things more deeply—a far cry from "throwing sand in my eyes," as he once described it. Sean Walsh was an indispensable interlocutor throughout my research, and I appreciate his willingness to talk about anything, with rigor, but also with patience and kindness. I thank Josh Armstrong for always encouraging me to think about where my projects fit into a bigger philosophical picture. Una Stojnić made many arguments in this dissertation sharper and more interesting, and I thank her for her penetrating questions. And many thanks to David Kaplan, who will always be a model for me of how to do philosophy of language and logic well.

Since the beginning of my graduate studies, Calvin Normore has been a continuous source of encouragement and insight regarding all my projects. I hope to emulate Calvin's philosophical temperament—cheerful, open-minded, and inexhaustibly curious. I am also grateful to have known and worked with Tyler Burge, John Carriero, Pamela Hieronymi, Katie Elliot, Andrew Hsu, and Adam Crager, each of whom has taught me so much.

I was lucky to have a community of terrific philosophers as friends and colleagues among the UCLA graduate students. Talking with Gabe Dupré decisively shaped the sorts of questions I think are important in the study of language and mind, and I am extremely grateful for his guidance and example. Christian De Leon provided important feedback on several ideas that eventually became dissertation chapters, and I have appreciated coming to know his perspective

on language and communication. I thank Ally Peabody for being a caring friend and for continually exposing my philosophical prejudices. Thanks as well to Gabby Johnson, Bill Kowalsky, John Kardosh, Ian Boon, Kim Johnston, Tim Buchanan, Tristen Cardwell, Catherine Hochman, Nefeli Ralli, Ekin Zeytinoglu, Paul Talma, James Gu, Seungtaek Lee, Jürgen Lipps, and Jonah Ragir for so many delightful and instructive conversations over the years.

Chapters 2 and 3 of this dissertation were dramatically improved thanks to audiences at the Central States Philosophical Association 2022 Conference, the Texas Tech 2023 Graduate Conference, and the UT Austin 2023 Graduate Conference. In particular, I would like to thank Bill Melanson, Jason Lemmon, Ryan Simonelli, David Boylan, Paul Forrester, Chad Hall, Harvey Lederman, Mark Sainsbury, Hans Kamp, and Sam Wolf Kantor for their comments. Chapter 2 reuses material published as "The problem of unarticulated truths" in *Philosophical Studies*, and it is reused here by the permission of Springer Nature.

My graduate studies would have been far bleaker were it not for Brian Thorn, my friend and one-time roommate. Brian has made me laugh more than anyone has a right to in their life. I also thank my parents Anne and Robin and my brother Lars for supporting and encouraging all my philosophical projects, even if my attempts to explain them have been pitiful.

To those I have forgotten—I thank you too.

Finally, I owe an enormous debt of gratitude to my partner Anna-Claire Siena for helping me find confidence in my work and for making my life outside philosophy so wonderful.

VITA

2020 C. Phil in Philosophy, University of California, Los Angeles

2017 M.A. in Philosophy, University of California, Los Angeles

2015 B.A. in Philosophy and Creative Writing, Columbia University

PUBLICATIONS AND PRESENTATIONS

"The Problem of Unarticulated Truths," 2023. *Philosophical Studies*

"Fine-grained Content and Open Texture," UT Austin Philosophy Graduate Conference; April 2023

"Fine-grained Content and Open Texture," Texas Tech Philosophy Graduate Conference; April 2023

"The Problem of Unarticulated Truths," 2022 Central States Philosophical Association Conference; University of Nebraska-Lincoln: April 2022

"Explanations in Semantics and the Theorist's Natural Language Competence," Mind and Language Workshop, University of California, Los Angeles; December 2018

"What do we learn from Truth-Conditional Semantic Explanations?" 10th California Universities Semantics and Pragmatics Workshop; University of California, Irvine; Oct 2017

CHAPTER 1

Introduction

The thread that connects these chapters is a picture about the basic function of the concept *truth*. According to that picture, truth is the standard by which we access individual cognitive states for correctness. In the most basic application of the concept, we are assessing *particular* mental happenings for correctness—one person's judgment, idea, or belief. The truth-bearers, on this understanding, are concrete parts of the natural world—they are products of mental agents and can stand in causal relations with other objects and events. I call this view about truth-bearers *Particularism.* The Particularist picture is an extremely old one; it seems to lie behind any account of truth as a correspondence between the mind and the world. Nonetheless, Particularism is widely rejected in the contemporary philosophical landscape.

A number of considerations suggest that truth must be predicated of objects that are independent of particular psychologies. Some of these derive from linguistic data—for instance, the fact that truth is very naturally predicated of clausal complements—but the most serious have a basis in substantive metaphysical and logical commitments. It is very natural to think that even if there were no mental agents, there might still be truths—for instance, it would be *true* that there were no mental agents and *true* that 2 + 2 =4. And it would seem that logical relations attach to thoughts a way that is completely independent of whether those thoughts are ever entertained. This influence of these considerations, especially as expressed so forcefully by Frege, has helped to solidify an anti-psychologistic perspective on truth.

The essays in these chapters are meant to point a way back to a form of psychologistic view of the truth-bearers. In my view, we abstract away some important facts if we treat truth as

property of pure contents, rather than a property of individual representations that possess those contents. The main contention of the first chapter is that understanding truth as a property of non-repeatable particulars is crucial for making sense of the semantic paradoxes. The fact that the truth-value of representation can vary across tokens of the same semantic type is what allows us to reflect on paradoxical utterances and coherently judge them untrue. In the second chapter, I deflect the main metaphysical objections against Particularism by appealing to representational types as derivative truth-bearers. In the third chapter, I draw attention to the ways in which appeals to abstract contents obscure the fact that there are multiple legitimate notions of sameness of content. Together, these chapter sketch a psychologistic picture of truth's place in the world that I think is coherent and plausible.

Perhaps the dominant question in the philosophy of truth, for the last century, has been the question of Deflationism: is there anything to say about the nature of truth beyond what is given by trivial disquotational principles? Though I rarely discuss Deflationism directly in this dissertation, it has been sort of foil that has shaped the direction of my research. If you accept Particularism, Deflationism has very little motivation. It is not easy to formulate standard disquotational principles in a Particularist setting, because an individual token cannot be "disquoted" the way a sentence type can. And when you do provide such principles with a Particularist friendly formulation, as I do on p. 14, they end up not looking so trivial. Even if, as I guess is probably right, there is no hope of providing an illuminating, reductive explanation of the nature of truth, I hope this dissertation demonstrates that, contra the Deflationist, there are a variety of non-trivial and explanatory things to be said about it.

CHAPTER 2

Particularism about Truth and the Significance of Revenge

0. Introduction

It is a familiar refrain in the literature on the semantic paradoxes that any theory that aims to diagnose all cases of semantic pathology and preserve classical logic is doomed to face Revenge problems (see Sharp 2008; Bacon 2015, 2018; Priest 2008; Armour-Garb 2008). These Revenge problems are often presented in the form of a dilemma: either the theory of truth on offer is expressively limited or it refutes itself. One cannot rationally accept the second horn of this dilemma: any theory that genuinely refutes itself is not rational to believe. But the story is a bit more complex regarding the first horn. Whether or not an expressive limitation is problematic depends on what the theory is meant to accomplish.

In my view, there is a sort of response to the semantic paradoxes that respects classical logic and avoids revenge problems: the *Particularist* response. The Particularist response is distinguished by identifying the truth-bearers with non-repeatable particular representations, the semantic properties of which are not fully determined by their representational type. Although, the standard revenge arguments do show that the Particularist needs to accept a kind of expressive limitation—namely, that no linguistic agent can fully describe the semantic properties of their own linguistic products—this is not a limitation that undermines the Particularist's claim to give a general diagnosis of Liar-like semantic pathology.

In Section 1, I introduce the Particularist response to the semantic paradoxes, and in Section 2, I illustrate how a revenge argument might be raised against it. In Section 3, I clarify the notion of *interpreted language* that is implicit in the revenge argument and suggest that the

Particularist should treat interpreted languages as *stable semantic states* of linguistic agents. In Section 4, I argue that the standard revenge argument raised for the Particularist shows that no agent in a sufficiently expressive stable semantic state can express a complete semantic theory for tokens produced in the state they occupy. I conclude by arguing that this conclusion is not particularly surprising, and it compatible with the key virtues the Particularist can claim for themselves. Particularism is Revenge proof.

1. Particularism

Let me sketch some reasoning regarding the Liar paradox as it might arise in natural language. Let "L1" be the name of the sentence token on the next line.

L1 L1 is not true.

L1 is a meaningful declarative English sentence token—I take myself and my reader to understand it. Here is an argument that it cannot be true. Suppose it is true. Something is true if and only if what it says is the case, and what L1 says is that L1 is not true. It follows that L1 is not true. But this contradicts our supposition. Suppose, then, that L1 is not true. Since what L1 says is that L1 is not true, then what L1 says is indeed the case. And, therefore, it is true.

What are we to make of this? Let us hold fixed that L1 is either true or not true (and not both)[1]. If that is so, then the intuitive principle we appealed to when we were reasoning must have led us astray, namely: something is true if and only if what it says is the case. This principle is a one of a family of "disquotational" principles and inference rules that are often taken as central to the concept of truth, the most famous of which is Tarski's T-schema. If we do not have

[1] Of course, many will want to get off the bus here, and there are well known accounts of the Liar that turn on rejecting the Law of Excluded Middle or the Law of Non-Contradiction. The merits of such views have been debated in many places (see Ripely, Beall, and Glanzberg 2018 Chapter 5 for an overview). I will simply say that I find the costs of giving up these laws to outweigh the benefits.

more systematic ambitions, the story of the semantic paradoxes might end here: the intuitive disquotational principles associated with truth are not unrestrictedly valid.

But those with systematic ambitions will be left with pressing questions: why do these principles break down when they do? How do we understand truth if not via such principles? And, to go back where we started, is L1 true or not? In my view, the last of these questions (in contrast to the former two) seems to have an intuitive, theory independent answer. Even if we bracket the intuitive T-schema and avoid being led into contraction, there is an asymmetry between the suggestion that L1 is true and the suggestion that it is not. It is difficult to see how it could be rationally coherent to judge that L1 is true. By judging something to be true, you endorse it. But if you judge L1 to be true, you endorse something that contradicts your very judgment. If you really understand L1, this is absurd, whether or not you are careful to hold off from explicitly contradicting yourself by drawing the further inference that L1 is not true. By contrast, one apparently *can* rationally judge that L1 is not true. The fact that the intuitive disquotational principles fail for L1 suggests that there is something *wrong* with it—wrong it in a way that is incompatible with its being true. Compelled by this, you might make the following judgment:

L2 L1 is not true.

This judgement is not incoherent in the way the former was. L2 is a different token than L1, so you do not, in the act of judging L2, ascribe untruth to your own judgment. And although your judgment instantiates a sentence that is *of the same type* as L1, this does not imply that you endorse L1. Two people assertively uttering the same sentence (e.g. “I am the world’s best chef”) do not automatically express endorsement of each other’s assertions.

L2 looks like it expresses a reasonable result that follows from diagnosing L1 with semantic failure. Gupta (2001) calls this the "Chryssipus Intuition": when the dust of confusion has settled, one wants to say, whatever else is going on with the Liar-paradoxical object, it cannot be true[2]. The puzzle at the heart of the Liar, in my view, is to give an account of how L2 could be semantically different from L1, given that they are tokens of the same sentence and apparently say the same thing. We cannot give them the same treatment: if we want to say that L2 is true, we should also say that L1 is true and thereby contradict ourselves; if we want to say that L1 is unfit for truth, then L2, our apparently sound diagnosis, is not true either[3].

Assuming that the truth-value of a natural language sentence token is determined compositionally, we have two standard sorts of explanations for why two tokens of the same sentence type might differ in truth-value: ambiguity and context sensitivity. The ambiguity suggestion raises more puzzles than it solves[4], and, as far as I know, has not been defended by anyone. A number of philosophers (Burge 1979; Parsons 1974; Simmons 1993, 2018; Glanzberg 2004; Murzi and Rossi 2018) have, however, offered Contextualist explanations of the difference

[2] It is telling, in my view, that responses to the Liar that do not straightforwardly imply a statement of the Chrysippus Intuition (e.g. Field's (2008) paracomplete view), are usually accompanied by a story according to which the Liar can be legitimately *rejected* (see Field 2008 p.73-78) .

[3] This phenomenon—that when we try to offer a diagnosis of the Liar, we are driven to either contradict ourselves or assert something we judge to be untrue—reemerges in all Revenge paradoxes (see Beall 2007, Sharp 2007).)

[4] The only real candidate for an ambiguous expression in "L1 is not true" would seem to be "true," and "true" does not pass ordinary linguistic tests for ambiguity (e.g. the Zwicky-Sadock test). And even supposing "true" were ambiguous, we are still left with many questions. For instance, how many truth-predicates does our language have? And are they all associated with their own disquotational principles? If so, then we are going to have as many Liar paradoxes as there are truth-predicates. The resulting picture seems to undermine the import of whatever it is that L2 says.

between L1 and L2. Common to all these accounts is the view that: (a) there is a contextually variable element in a sentence like "L1 is not true," and (b) there is a way of resolving this contextually variable element such that "L1 is not true" semantically fails and lacks a truth value, and *other* ways of resolving it such that it correctly reports this semantic failure. Probably the deepest objection to such accounts is that, in order to prevent the generation of revenge paradoxes (e.g., "This very English sentence is true in no context"), they must restrict our ability to quantify over linguistic contexts[5]. In addition to striking many as *ad hoc*, this conflicts with the sort of generality that is aspired to in empirical semantics. (Ordinarily, when natural language semanticists assign a semantic value to a context sensitive expression, like "now," they take themselves to be giving an account of its denotation relative to *any* context[6].) But a more basic issue with Contextualist responses to the Liar is that "linguistic context" in the familiar sense seems irrelevant to the difference between L1 and L2[7].

The most salient contrast between L1 and L2 is that L1 ascribes untruth to *itself*, whereas L2 ascribes untruth to L1. This difference seems to be context invariant: wherever and however L1 and L2 are produced, L1 will predicate untruth of itself and L2 will predicate untruth of L1. Moreover, we can be sure that any token of "L1 is not true" (if it is an English token and it follows our referential stipulation about "L1") will be true whatever context it is produced in, so

[5] For some examples of this objection, see Williamson 1998, Bacon 2015, Gauker 2006; for responses see Glanzberg 2006, Murzi and Rossi 2018, Simmons 2018 Chapter 9.

[6] Similarly, it is difficult to see how you could give an account of entailment between context sensitive sentences without considering the truth value of those sentences relative to *all* contexts.

[7] Glanzberg (2004) does acknowledge this by characterizing the context sensitivity as "extraordinary": the variation in contexts expands the range of propositions available to be expressed, not simply varying which sentences express which propositions. But extraordinary or not, his notion of context shift strikes me as the wrong category for describing what, I argue in the next paragraph, is better characterized in terms of relations between tokens.

long as that token is not L1. The key feature that distinguishes paradoxical and non-paradoxical tokens of "L1 is not true" seems to be *whether or not that token is identical with L1.* This feature is quite different from other features that are modeled as standard elements of linguistic contexts, e.g. agent, addressee, time, world, conversational common ground. Linguistic contexts, however they are modeled, are meant to capture speaker's *general* knowledge about how to interpret sentence types in various circumstances. By contrast, in order for a semantic theory to describe the salient difference between L1 and L2, it needs to make semantic predictions directly about tokens and distinguish between them *qua* tokens.

Such a semantic theory will embody the metaphysical commitments about truth-bearers that I am calling *Paricularism.* Particularism is the conjunction of two claims:

Token Truth-bearers

The fundamental truth-bearers are concrete, non-repeatable representations.

Particularist Semantics

> For some token representations, the factors that determine its truth-conditions involve features that individuate it as a token.

The first commitment distinguishes it from theories that treat truth as a property of sentence types, sentence-types-in-context, or abstract propositions. The second commitment is the lesson we ought to learn from L1 and L2. A token can have its truth-conditions influenced by a variety of features that are in principle sharable—what type it instantiates, time and location of utterance—but none of these features seems to be decisive in accounting for the difference between L1 and L2. In some cases, being *the very token* is it can be a factor that influences the determination of truth-conditions for a token representation.

Although most work in natural language semantics treats sentence-types-in-context as the main unit of analysis, there is reason to think that, with regard to some phenomena, semantics

should concern itself with token expressions and relations of dependence among tokens. It is widely held that expression tokens play a fundamental role in *metasemantics*[8]—the account of *why* linguistic expressions have the meanings they do—since most believe that linguistic expression types mean what they do in virtue of how agents use particular tokens of those expression. And there are areas where it is not obvious that the metasemantic and semantic projects can be pursued independently. For instance, many have suggested that the lexical semantics of vague or open textured words involve constraints set by canonical token utterances that provide precedents for subsequent use (Ludlow 2014; Cumming 2023). Given the generality of their concerns, it is often appropriate for linguists and philosophers frame theories at a level of description that abstracts away from reference to particular networks of token expressions. But a token-based approach to semantics is not unmotivated or new[9].

To turn Particularism into a real solution to the semantic paradoxes, we require a systematic theory that predicts which tokens are paradoxical and accounts for the difference between, e.g., L1 and L2 in a principled way. Haim Gaifman (1988, 1992, 2000) has shown how this can be done. (See Appendix 1 for a formal presentation of a Gaifman-style semantics). Here I will outline the mechanics of Gaifman's approach.

---

[8] The distinction between semantics—an account that describes the meanings of—and metasemantics—an account that explains *why* expressions have the meanings they do—goes back to Kaplan 1989, though similar distinctions were made by Stalnaker (1997) and Dummett (1975). Kaplan 1989 also presents arguments, in the context of arguing against Reichenbach's 1947 account of token-reflexive expressions, that sentences-in-context, rather than tokens or utterances, are the appropriate objects of semantic and logical analysis. Building on this account, Braun (2018) argues that, while token expressions have a role to play in metasemantics, they have no role to play in semantics proper. For further discussion and debate regarding token-reflexivity, see among, others Garcia-Carpintero 1998, Predelli 2006, Simchen 2013.

[9] It is adopted by Reichenbach 1947, Garcia-Carpintero 1998, Perry 2001, Korta and Perry 2011, Crimmins 1995, Schiffer 2005, and Ismael 2011, among others. It is also the perspective held by virtually all logicians of Medieval Latin Christendom (Brower-Toland 2022).

A Gaifman-style semantics takes a domain *D* of tokens of expressions in some language *L* and constructs a valuation function on *D* that assigns each member TRUE, FALSE, or GAP. We assume that *L* has two semantic predicates, "Tr" and "Fa," and that the non-semantic fragment of *L* comes with a background interpretation, associating each sentence type containing no semantic vocabulary with a truth value. (This can be thought of as specifying the default truth value of any token of the respective sentence type.) The construction proceeds by extending some initial valuation (canonically, the empty valuation that is defined on no tokens) according to three rules until all members of *D* are evaluated. The *Standard Value Rule* operates on non-pathological tokens and assigns them the default classical value associated with the sentence type they instantiate[10]; the other two rules operate on pathological tokens and assign them GAP. The construction proceeds on two levels—a token level and a type level. As tokens are evaluated, the background interpretation associating default truth values with sentence types is updated to reflect this change. So, for instance, if a token named "a" is evaluated as GAP, the background interpretation is extended to associate the sentence type "Tr(a)" with the value FALSE.

[10] This procedure might seem to stand in tension with the suggestion that tokens rather than types are the fundamental truth-bearers. Formally, sentence *types* get evaluated with classical values first in Gaifman's semantics, and the truth-value of non-pathological tokens is determined by the truth value assigned to their type. One might think this indicates that the truth-values of tokens *depend on* the truth-values of types. This is a mistake, in my view. The assignment of truth-values to sentence types is a formal convenience of Gaifman's presentation that can be eliminated without changing the resulting valuation of tokens. For instance, instead of keeping track of an induced valuation, we could successfully add clauses to our theory that make *universal generalizations* about all the tokens of a given type that have not been previously assigned GAP. The Standard Value Rule reflects the uncontroversial point that the truth-value of a token expression is determined in part by what type it instantiates—a point which does not imply that the types themselves are truth-bearers. Ultimately, the interest of the Particularist semantics offered by Gaifman is considerably undermined if we take "fundamental" truth to be the property attributed to sentence types in the metalanguage. If that is what is fundamental, it is unclear why we should bother giving a semantics for a language with a truth predicate restricted to tokens.

According to Gaifman, a token is pathological if it is a member of a defective dependence structure. This analysis is broadly of a piece with other "grounding" based accounts (Herzberger 1970, Kripe 1975, Yablo 1982, Maudlin 2004, Leitgeb 2005). Predications of truth or falsity establish dependence relations among tokens: if *x* predicates truth of *y*, then assessing the truth value of *x* depends on first assessing the truth value of *y*. Pathology arises when these chains of dependence do not terminate: when the dependence relations form closed loops or infinite descending chains. The core idea here is that, if evaluating the truth value of some object depends on the results of evaluating others, this chain of dependence needs to terminate in objects that can be evaluated independently[11].

To show Gaifman's account in action, consider L1 and L2. The lines in the diagram below indicate the dependence relations established by the tokens of the truth predicate:

[11] A full defense of the grounding-based account of semantic pathology is beyond the scope of the present paper, which is narrowly aimed at answering Revenge objections. Revenge arguments are often presented as a problem that arises for *any* account of semantic pathology that accepts classical logic, whatever the details of how it analyses semantic pathology (see Sharp 2008, Bacon 2015, Priest 2008, Armour-Garb 2008). That said, there are significant worries to be raised about whether or not the grounding-based account is overbroad. On standard developments, it will count innocuous generalizations like "Nothing true is false" as ungrounded. Most seriously, since "grounded" is itself a semantic predicate, Gaifman's account ought to predict that any token of "Nothing ungrounded is true" is itself ungrounded. These are real difficulties, but they are distinct from the problems raised by Revenge arguments, and I think they can be answered. A rough suggestion (one I am developing in another paper) is the following: some semantic generalizations are *instance-based* generalizations that establish dependence relations with individual tokens, while others are *generic* generalizations that depend only on the nature of truth (c.f. Linnebo 2022). In my view, tokens of "Nothing true is false" or "Nothing ungrounded is true" do not stand in dependence relations to any tokens, because they are made true, not by individual tokens, but by something like the *laws* of truth. A distinction like this is already present implicitly, I think, in standard accounts of semantic grounding. All accounts of grounding recognize sentences involving semantic vocabulary that are essentially counted as grounded by courtesy, e.g. principles about truth-aptness like "Only sentences are true."

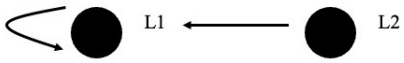

Figure 2.1

L2 is dependent on L1 and nothing else, whereas L1 is dependent only on itself. L1 constitutes a closed-loop—if we try to evaluate it, the dependence relations lead us in a (one membered) circle. In Gaifman's semantics, this means that the Closed-Loop Rule is enabled on L1, and at some stage in our construction we can assign L1 the value GAP. Once this happens, the type-level track of the theory is updated so that "L1 is not true" is associated with the default value TRUE. This allows us to evaluate L2 as TRUE in the next step with the Standard Value Rule.

If a token is semantically pathological it *lacks* a truth-value[12]. The value GAP signifies *semantic failure*. The general idea of semantic failure is familiar from philosophy of language: there are some sentence tokens that, despite being syntactically well-formed and purportedly truth-apt, malfunction in such a way that they cannot be evaluated for truth or falsity. The paradigmatic case is *reference failure*. To take Strawson's (1950) case: suppose I extend cupped hands towards someone, saying, as I do so, "This is a fine red one," when in fact there is nothing in my hands. The expression type "this" functions to refer, on an occasion of use, to some contextually salient object. But, in the imagined case, no object is provided, so the token "this" apparently does not refer to anything. Plausibly, we cannot assess this utterance as either true or

[12]One could develop a Particularist semantics along Gaifman's lines that—following Bradwardine, Albert of Saxony, and Buridan (see Read 2002)—evaluates pathological tokens as false. I won't discuss this alternative except to point out that it seems to me somewhat quixotic to strive to preserve bivalence when it is accompanied by other significant changes to classical semantics (e.g. allowing that tokens of p and $\neg p$ can both be false). The characterization I give below, of Liar-tokens as failing in their semantic function, strikes me as more faithful to the phenomenon.

false. The utterance would be true just in case the object referred to is a fine red one and it would be false if the object referred to is not a fine red one—but since no object is in fact referred to, neither of these conditions can obtain[13]. Though they fail for a different reason, Liar tokens fail in the same sense: a precondition for their being true or false fails to hold.

Gaifman's Particularism is sometimes classified among "No Proposition" or "Meaningless" responses to the Liar paradox (Bacon 2018). These labels are sometimes applied to any view that characterizes Liar-paradoxical objects as semantically malfunctioning, and in that sense, the classification is apt. But it is worth emphasizing that *propositions* as such play no role in Gaifman's semantics as I am developing it here.[14] The explanation of whether a given token is pathological depends only on its place in a network of semantic dependence—these networks can be characterized without mentioning propositions, and without a general theory relating propositions to truth-aptness. I am happy to say that L1 does not "say anything," but this is meant merely to report the semantic pathology—it does not explain it.

What does Particularism tell us about the T-schema or related disquotational principles? Standard formulations of such principles cannot be straightforwardly applied in a Particularist setting. Take Tarski's T-schema:

---

[13] This understanding of semantic failure implies that more is required for a token utterance to be false than for its truth-conditions to be unfulfilled (cf. Dummett 1959). To my mind, examples like Strawson's show there is no reason to suppose that bivalence holds in general for utterance tokens. This is compatible with supposing that bivalence might hold at another level of linguistic analysis, e.g. the level of assertions (see Glanzberg 2003). In my view, however, questions of truth and falsity do arise fundamentally for utterance tokens, so there is nothing "insubstantial" about a truth-value gap at the utterance level.

[14] And although "meaningless" in the present context is obviously, to some extent, a term of art, I regard it as very misleading. It is precisely *because* Liar sentences have a specific meaning that they generate a puzzle—and, in my view, no small part of responding to the puzzle involves characterizing the sort of meaning they have in a coherent way.

s is true if and only if *p*

One instantiates the T-schema by replacing "*s*" with an expression that refers to a sentence and replacing "*p*" with the sentence referred to (or a translation of it); and the resulting instantiation is itself a sentence type. But truth, for the Particularist, is properly predicated of tokens, so "*s*" will need to be a replaced with an expression referring to a token, and, if we want the substitution for *"p*" to be the object the substitution of "*s*" refers to, then "*p*" needs to be replaced by a token. But if we demand that the substitution of "*p*" must be the *very* token that the substitution of "s" refers to, then applications of the T-schema will be massively restricted—token instantiations of T-schema will only make semantic predications about objects they contain as subtokens. But the T-schema is meant to explain the validity of arguments involving "long distance" disquotation, like the following:

Tarski's utterance of "Snow is white" is true. Snow is white.

Snow is white. Tarksi's utterance of "Snow is white" is true.

In these arguments, the "disquoted" token is not literally Tarski's utterance—it is mine. What it has in common with Tarski's utterance is that it instantiates the same sentence type. This indicates the sort of formulation of the T-schema that the Particularist ought to take seriously. Let's say that, in the following two-way inference schema, "*s*" is to be replaced by a name of a token *x* and "*p*" is to be replaced by a token of the same type as *x*.

s is true ⊣⊢ *p*

Clearly, this schema will only preserve truth in restricted class of cases. The token referred to by the substitution for "*s*" needs to be part of the same language as the substitution for "*p*"; the schema has to be limited to applications where both tokens occur in the same context; and

neither the token referred to by the substitution for "*s*" or the token substituted for "*p*" can be part of a pathological dependence structure.

When conditions are right, one can state the truth-conditions of a token by producing another token that instantiates the same type. But whether these conditions happen to hold is not a matter of logic, or something that can be guaranteed by the meaning of the truth-predicate. Nor do they seem to be obviously constitutive of truth. Intuitively, whether a token is true does not depend on the existence (or possibility) of other tokens of the same type stably preserving its semantic features[15].

In a moment, I am going to turn to a Revenge objection against Particularism. But first let me summarize the points that, to my mind, recommend a Particularist response to the semantic paradoxes. First, Particularism can be developed in a manner that preserves classical logic. Although Gaifman's semantics for tokens is not bivalent, the system preserves classical logic in the sense that, within the realm of non-pathological tokens, any instantiation of a classically valid schema is true and all instantiations of classically valid inferences preserve truth[16]. Second, it explains the Chrysippus Intuition. Third, it has a quality I'll call **Unity**: it is compatible with the view that there is one property—*truth*—that all truth-predicates function to express.

---

[15] This does not contradict the (plausible) idea that in order for there to be representations *at all* there need to be repeatable types, tokens of which reliably have the same semantic properties.

[16] For this to really amount to "preserving classical logic," we have to suppose that pathological tokens are outside logic's domain—the rules and laws logic is concerned with do not bear on them. (This is the sort of view that Kripke 1975, somewhat infamously, advocates regarding his own non-classical semantics.) Warren 2023 draws an illuminating comparison between pathological sentences and other sorts of sentences that logic (in one good sense of the word) can legitimately ignore, e.g. imperatives, and he gives a a sketch of a natural deduction system that incorporates rules for excluding pathological sentences from arguments. With Warren, I think that serious sorts of departures from classical logic involve proposing non-classical rules of inference that are meant to govern reasoning with, among other things, pathological objects.

Many extent responses to the semantic paradoxes make available multiple candidates for the property of being true. For instance, in Tarski's hierarchy, there are only truth-predicates *for particular languages*. He does not attempt to characterize a universal truth-predicate applicable to truths in whatever language, and his work has suggested to many that such an idea is incoherent (see Field 2008 for this interpretation). This fragmentation of truth into a plurality of different properties is preserved in many responses to the Liar that, contra Tarski, treat languages with a self-applicable truth predicate[17]. The Particularist, by contrast, can maintain that there is one property—roughly, *representing things as they are*—that tokens of the object language truth-predicate function to attribute. In the case of some tokens, like L1, this function may be stifled, but the theory does not provide any *alternative* property that it attributes on those occasions. *Prima facie*, it is a good thing if we can coherently respond to the semantic paradoxes while preserving the pretheoretical picture that there is a unique property, *truth,* serving as a target for our cognitive activities. Our cognitive lives seem to be oriented toward truth as a fundamental value: we want our beliefs to be true; we take good evidence to lead reliably to true beliefs. It is not obvious what these foundational platitudes mean if there are multiple properties that are equally good candidates for being the property *truth*.

2. Revenge

[17] In general, this sort of fragmentation arises as a distinction between a truth-predicate in the object-language and a truth-predicate in the metalanguage. For instance, in Kripke's theory of truth, the Liar sentence takes the value GAP in the minimal fixed-point, and given that GAP is meant to signify "expresses no proposition," it is natural for the theorist to conclude that the Liar sentence is not true. But the object language truth-predicate cannot be used to make this claim—the Liar sentence is not in the anti-extension of that predicate. So Kripke's theory provides us with two non-equivalent truth-predicates. Whether this genuinely provides us with two different truth *properties* is a matter we might debate, but the existence of two non-equivalent predicates raises the question.

Bacon (2015, 2018) has developed revenge arguments as they apply to Particularist accounts in an especially forceful way, and my version of it draws heavily on his. Bacon suggests that, in order for a theory to succeed in the project of distinguishing which objects are semantically pathological and which are not, the "Diagnostic Project", a semantic theory ought to prove each instance of a *restricted* version of the T-schema, which in a Particularist setting looks like this:

RTS For all *x,* if *x* is a non-pathological token of $s \rightarrow (x \text{ is true} \leftrightarrow p)$

where "*s*" is replaced by a term referring to sentence and *p* is replaced by the same sentence. The intuitive idea here is that failures of the T-schema are sign of Liar-like semantic malfunction. So, any theory that is compatible with there being non-pathological tokens for which the T-schema does not hold must not have carried out the diagnosis correctly. Note: Bacon is assuming that the sentences in question are not context sensitive[18]. Since, I do not think that context sensitivity is relevant to generating the semantic paradoxes, I will grant this. From here on out, I am going to bracket all ordinary context sensitivity and assume that we are discussing languages purged of any context sensitive vocabulary.

Any Particularist semantics should also include some axiom saying that nothing pathological is true:

PATH For all *x,* if *x* is pathological, *x* is not true.

A standard Revenge Argument for the Particularist focusses on a sentence type like *u,* on the next line.

u No token of *u* is true

[18] Otherwise, it is pretty clear how there could be failures of the T-schema without any semantic pathology.

Let's suppose we have a Particularist semantics for tokens of English sentences, call it *E*, that proves every instance of RTS. In particular, it will prove RTS-*u,* the instance of RTS where "*s*" is replaced with a term referring to *u* and "*p*" is replaced with *u*:

RTS-u For all *x,* if *x* is a non-pathological token of *u* → (*x* is true ↔ No token of *u* is true)

RTS-*u,* with PATH, classically implies that all tokens of *u* are pathological. Suppose there is a non-pathological token of *u*, call it *c*. From RTS-*u*, we can infer that *c* is true if and only if no token of *u* is true. Since *c* is itself a token of *u*, this result will lead to contradiction[19]. Since *c* was an arbitrary instance, we can infer that all tokens of *u* are pathological. If we accept the account of semantic pathology as ungroundedness, this verdict looks reasonable. Any token of *u* will create a closed circular network of dependence consisting of itself and other tokens of *u.* In a Gaifman-style semantics, one can prove that any token of a sentence analogous to *u* will receive the value GAP (see Appendix 2). But the result is also quite puzzling, because in combination with PATH, it implies *u:*

u No token of *u* is true

Since RTS-*u* is a theorem of *E,* PATH is one of its axioms, and these together imply *u*, *u* itself is a theorem of *E.* So *E* has a theorem, *u*, such that it implies all tokens of that theorem are pathological.

[19] Here is an informal proof. *c* is either true or not true. Contradiction follows from each of these disjuncts. Suppose that *c* is true. Therefore, no token of *u* is true. This contradicts our assumptions that *c* is a token of *u* and that *c* is true. Suppose, on the other hand, that *c* is not true. Since, *c* is not true if and only if it is not the case that no token of *u* is true, we can infer that some token of *u* is true. Call that token *d*. *d* is a true token of *u*. By PATH, *d* is non-pathological. From RTS-*u* we can infer that *d* is true if and only if no token of *u* is true. Therefore, no token of *u* is true. This contradicts the claim that *d* is a true token of *u*.

In Bacon's (2015) account of this argument, he suggests that it shows that any classical response to the Liar (Gaifman's Particularist account included) either fails in the Diagnostic Project or refutes itself by having theorems it implies are pathological. This is a bit imprecise as applied to *E*: *E* consists of *sentence types,* and it does not imply that its theorem *u*, the sentence type, is pathological, just that all tokens of *u* are. But this result still shows *E* to be radically self-undermining. By its own lights, *E* cannot be articulated truly—someone who tried to express its theorems would end up producing tokens that are not true. If nothing else, this makes the theory self-refuting *in practice*. And although the idea that *E* is true but cannot be truly tokened might be intelligible, this is not something that a Particularist, insofar as they take tokens to be the fundamental truth-bearers, can comfortably say.

The Particularist, at this point, might appeal to appeal to a distinction between object-language and metalanguage. After all, the puzzle only arises because *u* is both an element of the language in which the theory is stated and a sentence whose tokens the theory is meant to describe. There will be no risk of self-refutation if, say, the theory consists of Japanese sentences but is addressed to tokens of English sentences. And, although Bacon's RTS does not make *explicit* reference to an object language, we have to presuppose some implicit reference to an object language for the schema to be plausible. Consider this instantiation of RTS:

> For all *x,* if *x* is a token of "Many dogs have awful breath," then *x* is true if and only if many dogs have awful breath.

Unless we suppose that "Many dogs have awful breath" picks out an *interpreted sentence,* and that the sentence substituted for "*p*" is the *same* interpreted sentence, there is no reason to think this instantiation is true. If "Many dogs have awful breath" picks out an uninterpreted string of shapes or sounds, no semantic theory should be in the business of predicting the truth-conditions of *all* its tokens, because those tokens might be produced with different meanings. And if it picks

out a sentence type *as* interpreted a specific way, that had better be the way my token of "Many dogs have awful breath" is interpreted. For instance, if "Many dogs have awful breath" picks out an interpreted sentence of, say English as spoken in the 17th Century, when "awful" meant *awe-inspiring*, then the instantiation is likely false—many dogs do have awful breath, but, presumably, not all tokens of "Many dogs have awful breath" *qua* expression of 17th Century English are true[20].

If we make the relativization of RTS to a language explicit, the schema looks like this:

Relative-RTS For all *x,* if *x* is a non-pathological token of *s* produced in language *L* → (*x* is true ↔ *p*)

where "*s*" is replaced with a term referring to a sentence, "*L*" with a term referring to a language, and "*p*" replaced with the sentence substituted for "s" or a translation of that sentence as interpreted in the language referred to into the metalanguage. Let's suppose that Bacon is right to suggest that, in order to succeed in the Diagnostic project, a semantic theory must imply every instance of Relative-RTS. This effectively implies that, if a semantic theory is stated in the language *K*, in order to fulfill the Diagnostic Project it must characterize the truth-conditions of all tokens produced in the language *K.*

We can frame a Revenge argument in terms of Relative-RTS on the basis of a different revenge sentence—one that refers to the language the theory is stated in[21]. So, if our semantic theory is stated in English, our revenge sentence can be:

---

[20] Perhaps it is vacuously true because there are no such tokens. But the point is that this is a sort of generalization a semantic theory has no business making.

[21] *u* cannot be used to derive the desired result. If we fix English as the language of our theory, then one instance of Relative-RTS will be "For all *x*, if *x* is a non-pathological token of *u* produced in English, then *x* is true if and only if no token of *u* is true." But this does not show that the theory implies that that no token of *u* is true (and thereby has *u* has one of its theorems).

relative-u No token of *u* produced in English is true.

If our semantic theory implies every instance of Relative-RTS, and it includes PATH as an axiom, then it will also have *relative-u* as a theorem—and, thereby, it will have a theorem that it claims cannot be truly tokened.

But what should we make of this? The Particularist may want to protest that Bacon's Revenge argument should really be interpreted as a *reductio* of the demand that, in order to succeed in the diagnostic project, a Particularist semantic theory stated in *K* must characterize the truth-conditions of all tokens produced in the language *K.* If a Particularist insists on a separation between the object-language, tokens of which they intend to describe, and the metalanguage in which they develop their theory, they run no risk of developing a theory that refutes itself. So perhaps the lesson of the Revenge argument is simply that no Particularist semantic theory for all tokens produced in a language *L* can be stated in the language *L.*

Without further elaboration, though, this response looks quite lame. First, it offers no reason to think that Bacon is *wrong* to suppose that succeeding the Diagnostic Project involves producing a theory that implies every instance of Relative-RTS. To reject this requirement simply because it dooms any Particularist semantic theory to self-refutation would clearly be special pleading. And, by insisting on an object-language/metalanguage distinction, this response also seems to undermine the Particularist's claim to preserve Unity. If a Particularist semantic theory stated in English can ascribe truth-conditions to all tokens of Spanish sentences containing "verdadero" but cannot ascribe truth-conditions to all tokens of English sentences

---

The theory is compatible with there being true tokens of *u*, so long as those tokens are produced in a language other than English.

containing "true," this would seem to suggest "true" and "verdadero" do not express exactly the same thing.

If these objections are cogent, then Bacon's Revenge argument seriously damages any Particularist account of the semantic paradoxes. Ultimately, I will argue in Section 4 that the response given above—that no Particularist semantic theory for all tokens produced in a language *L* can be stated in the language *L*—is basically correct and the objections misfire. But, before I can make that case, I need to clarify the notion of an interpreted language that the Particularist is appealing to when they distinguish between object-language and metalanguage.

3. Languages, Agents, and States

The Revenge argument crucially depends on making generalizations about tokens "produced in" a given language, e.g. English. This is because, unless we specify a language, it would be foolish for any Particularist semantics to make predictions about the truth value of a token based simply on the type it instantiates[22]. If we know something is a token of the sentence "1 + 1 = 10," but we do not know whether it was produced in decimal notation or binary notation, we will not be in a position to assign it a truth value. But what exactly does this "produced in" locution amount to? What makes it the case that a given token of "1 + 1 = 10" is produced as an expression of decimal vs. binary notation?

This is a general question about how to apply truth theories that make generalizations about linguistic tokens based on the expression types they instantiate. Commonly, it's assumed that linguistic expression types cannot be assessed for truth-conditions *tout court*; they are true or

[22] Gaifman (2000) recognizes this and accounts for it by stipulating that the objects of semantic evaluation are tokens (or, in his terms, pointers) *as they are used* in some particular language. He does not, however, give an account of what this means, as I propose to do here. To the extent that it is unclear what it means for a token be *used as an element of a specific language* it is not obvious how Gaifman's semantics can be applied empirically.

false only relative to some way of interpreting them, and they are not intrinsically determined to have some specific interpretation. This holds of natural languages as well as artificial languages. To be sure, natural language expression types are not individuated, like artificial languages, as strings of abstract shapes or sounds[23]. And in fact, there is good reason to suspect that semantic features *do* play a part in the individuation of natural language expressions—for instance, it is plausible to think that verbs are essentially associated with theta-grids, which encode abstract aspects of meaning. But, for all that, it would be radical to suppose that natural language sentence types are *essentially* such that all of their tokens can be associated with rich semantic properties that are sufficient to fix their truth-value (even in a specified context)[24]. So, with regard to natural language too, if one wants to make generalizations about the truth-conditions of tokens based on the expression type they instantiate, it is necessary to restrict the generalization to a group of tokens for which it can be assumed that, if tokens instantiate the same type, then they have the same meaning, i.e. a group of tokens that all belong to some uniform language.

[23] See the literature on words, e.g. Kaplan 1990**,** Hawthorne and Lepore 2011, Bromberger 2011, and Stojnić 2022, among others.

[24] There are a variety of arguments to be made in favor of this point. Perhaps the simplest derives from common sense: we commonly say that natural language expressions *change* their meaning. As I mentioned in the last section, it is often maintained that, before the 19th century, the English word "awful" meant *awe-inspiring* rather than *reprehensible*. If this is correct, then "awful" cannot be individuated in such a way that would fix its extension—tokens of the word produced in the 18th Century have a different extension from tokens produced today. Though I won't argue for it here for reasons of space, I think modern theoretical linguistics is perfectly compatible with, and provides some support for, the claim that linguistic expressions are not individuated with respect to rich semantic properties. (The perspective on the metaphysics of linguistic expressions I am adopting here is potentially in tension with the account of word individuation offered by Stojnić 2022. Stojnić's proposal is that words are individuated as finely as the neologizing events that associate expressions with meanings. But, so long as the *meanings* Stojnić has in mind do not determine rich truth-conditional properties like characters or extensions for all tokens of a given word, her view will be compatible with mine.)

The question still stands: what does it mean to say that a token expression is produced in one language rather than another?

I want to suggest that, in answering this question, it is helpful to think of a language not as an abstract object, e.g. a pairing of expression types and meanings (cf. Lewis 1975), but as a naturally occurring *capacity.* As part of the process of normal maturation, humans develop an ability to produce and interpret token expressions of novel linguistic types, and they use this ability to communicate with one another. This ability would be miraculous if the repeatable expression types figuring in token utterances did not systematically determine the meaning of those utterances. The normal functioning of language as we know it depends on conditions of semantic *stability* across tokens of the same expression type.

Put another way, we encounter phenomena that must be explained by reference to linguistic agents occupying what I will call a *stable semantic state*:

Stable Semantic State

A linguistic agent *a* occupies a stable semantic state if and only if there is some set of expression types D such that, for all $e \in D$, if *a* were to produce a token *x* of *e* then, necessarily, any token of *e* produced by *a* would have the same meaning as *x*[25].

[25] Where does polysemy show up on this picture? One might object: normal linguistic agents are not really in "semantic stable states" with respect to all the expression types in their language, because many of those expression types have *multiple* meanings. The precise way polysemy is realized in the lexicon is still a matter of debate (see Lupukhina et al. 2018), but it may well be that some natural language expressions have multiple meanings. Nonetheless, there are fruitful research programs in semantics that abstract away from this by treating the objects of interpretation as disambiguated. (Ultimately, disambiguation *needs* to happen somewhere if we want to assess linguistic products as categorically true or false.) My conception of stable semantic states takes up this idealization. Although it is clearly a simplification, I do not think remedying it will affect the basic idea that normal linguistic agency requires conditions of semantic stability across tokens of the same expression type.

The notion of "necessity" here is supposed to be *linguistic necessity*: whatever sort of natural necessity generalizations in semantics are supposed to have[26]. And the counterfactual conditional can be understood on the strict analysis with necessity accordingly restricted: it is linguistically necessary that if *p* then *q*. The characterization above tells us that whenever there is a linguistically necessary regularity in meaning that holds across tokens of the same type produced by a given agent, that there is a stable semantic state that agent occupies. If *S* is a stable state, I will use the term *S*'s "domain" to refer to the set of expression types over which *S* is a regularity. I will leave the notion of "meaning" that figures in this characterization as a primitive. (I don't think this is any more objectionable than Tarksi's helping himself to a primitive notion of "translation.") The use I make of the notion will be compatible with a variety of more substantive accounts of meaning, e.g. explications of meaning in terms of rules of use, inferential roles, links with mental representations, or default truth-conditions.

There are two reasons why I suggest that an interpreted language should be thought of as a stable capacity of an expression producer, rather than an abstract mapping between expressions and meanings. First, for natural languages, is not obvious how this mapping should be characterized. There is a long tradition of thinking of an interpreted language in *truth-conditional* terms, i.e. as association between sentences and truth-conditions (Davidson 1967, 1973; Lewis 1969, 1975; Partee 1980). But it is difficult to see how an interpreted language, so conceived, could contain a truth-predicate. As Tarski (1936) showed, the natural way of developing a theory of truth for such a language results in a (classically) inconsistent theory. This pushes us either

[26] It is not intended to rule out, for instance, an agent moving from one stable state into another. If a word changes its meaning (contra Stojnic 2022), then an agent with that word in their lexicon will transition to a different state when the change takes place.

toward either thinking that no interpreted language contains its own truth predicate[27] or toward thinking that some interpreted languages somehow have their proper interpretation characterized by a theory of truth that is logically false (c.f. Tarski 1936, Patterson 2009, Eklund 2002 on the suggestions that there are *inconsistent* languages). Both of these positions, I think, are inadequate for describing natural language.

Second, however we characterize the mapping between expressions and meanings, if we identify a language as an abstract interpretation function, we are still left with the question we began with at the start of this section: in virtue of what does a *token* expression count as belonging to one interpreted language rather than another? It seems that any plausible answer to this question will have to talk about the circumstances the token was produced in and the state of the agent who produced it[28]. By simply identifying interpreted languages with stable semantic states, we remove this further explanatory step linking abstract languages to concrete linguistic activities.

---

[27] We can prove for instance, that no Lewisian-language contains a Liar sentence. For Lewis, a language *L* is a function that maps sentences to sets of possible worlds. A sentence is true-in-*L* if and if *L* maps that sentence to a set of worlds that contains the actual world. Suppose that *L* contains devices for self-reference and a predicate "T" expressing *truth-in-L*, and let λ be a Liar sentence, i.e. λ = "¬T(<λ>)". The intension of "T" should map every world *w* to the set {*x* | *L* maps *x* to a set of worlds that contains *w*}. (Or alternatively, {*x* | *L* maps *x* to a set of worlds that includes the actual world}. The proof goes through either way.) This implies that *L* will map the Liar sentence "¬T(<λ>)" to the set of worlds {*w* | *L* maps λ to a set of worlds that does not include *w*}. If this is the proposition expressed by λ in *L*, then it is true-in-*L* if and only if it is not true-in-*L*, because the actual world is a member of this set if and only if it is not. This is a reductio of our assumption that *L* contains the resources to form the sentence "¬T(<λ>)," where "T" expresses true-in-*L*. For a rich discussion of the issues raised by Liar sentences in a Davidsonian semantic framework, see Lepore and Ludwig 2005 p. 133-138.

[28] c.f. Lewis 1969, 1975. After characterizing the sort of interpretation function he thinks constitutes a language, Lewis must take up the question of what it is for a person or community to *speak* a language.

To summarize: in order for any reasonable semantic theory to draw inferences about the truth value of tokens based on the type they instantiate, such inferences need to be mediated by assumptions about what interpreted language the tokens in question belong to, and I propose that we understand interpreted languages as *stable states of expression producers*. (At the end of Appendix 1, I show how a Gaifman-style semantic can be modified to make explicit the mediating role of stable semantic states.) In the next section, I will argue that what Revenge arguments, like Bacon's, show is that no agent can articulate a complete semantic theory for tokens produced in the very state they occupy.

4. State Limitation

The Revenge argument was meant to show that any Particularist semantic theory that aims to fulfill the Diagnostic Project will have theorems that it implies cannot be truly tokened. As I pointed out in Section 2, this result can only obtain if we assume that the theory in question makes predictions about tokens that are produced in the very language that the theory is stated in. Now that we have identified languages and stable semantic states, we need to clarify what it means to say that a *theory*—a collection of abstract sentence types—is stated in this or that language. To say that a collection of sentences T is "a theory stated in *L,*" where *L* is a stable semantic state, is to ascribe a canonical interpretation to the members of T according to which each sentence *s* in T has the meaning that a token of *s* would have when produced by an agent in state *L.* In terms of this notion, we can define the idea of a *self-articulating semantic theory* for a stable semantic state *L*.

> A set of sentence types H is a *self-articulating semantic theory* for a stable semantic state *L* if and only if, for every member *p* of H,
> (i) *p* is an expression type in *L*'s domain;
> (ii) tokens of *p* produced by an agent in state *L* make true predictions about the truth-values of tokens produced by agents in *L*
> (iii) H is closed under (classical) logical consequence

The Revenge argument arises from the assumption that, in order to succeed in the Diagnostic Project, any Particularist semantics must imply every instance a restricted T-schema. Given our identification of languages with stable semantic states, the restricted T-schema ought to take the following form:

State-Relative-RTS For all *x,* if *x* is a non-pathological token of *s* produced in state *S* → (*x* is true ↔ *p*)

where "*s*" is replaced by a term referring to a sentence type, "*S*" is replaced with a term referring to a stable semantic state, and *p* is replaced with the sentence referred to by the substitution for "*s*" or a translation of that sentence (as interpreted in the stable semantic state referred to by "*S"*) into the language of the theory.

I will argue in Section 4.2 that Bacon is wrong to think, for Particularism to succeed in the Diagnostic Project, it must present a semantic theory that implies every instance of State-Relative-RTS. Nonetheless, I think the schema is an important benchmark of success in semantics. In particular, if a Particularist semantics is to be a *complete* semantic theory for tokens produced in a state *J*, then it must imply every instance of State-Relative-RTS where "*S*" is replaced with a term referring to *J*. So, for instance, if I want to develop a Particularist semantic theory (in English as spoken by me) for tokens produced in (say) Raphael Nadal's current state, the theory ought to imply every instance of the schema:

> For all *x,* if *x* is a non-pathological token of *s* produced in Nadal's current stable semantic state → (*x* is true ↔ *p*)

where "*s*" is replaced by an expression in my language that refers to a sentence in Nadal's language, and *p* is replaced by a sentence in my language that translates Nadal's. Suppose my semantic theory did not imply every instance of this schema. In that case, there would be a sentence type, *a*, in Nadal's language for which, although I have a sentence in my language, *b*,

with the same meaning as *a*, I am not prepared to rule out there being cases in which he produces non-pathological token of *a* and it has a different truth-value than tokens of *b* produced by me. (This, of course, is easily explicable if *a* and *b* are context sensitive, but, as usual, we are bracketing all context sensitivity.) Either my theory is incomplete or we have a case where, for no obvious reason, sameness of meaning fails to preserve truth.

In my view, Revenge arguments aimed at Particularism really show that no semantic theory for a state *S* can be *both* self-articulating and complete. For suppose that T is a complete self-articulating semantic theory for the state *S*. If *S* has the resources to frame a semantic theory for tokens produced in *S*, it must be able to refer to *S* and express a theory of its own syntax. So *S*'s domain will include a sentence like *State-Relative-u*:

State-Relative-u No token of *State-Relative-u* produced in *S* is true.

If T is complete semantic theory for *S*, then it will have the State-Relative-*u* instance of State-Relative-RTS as a theorem. The State-Relative-*u* instance of State-Relative-RTS, combined with PATH ("Nothing pathological is true"), classically implies State-Relative-*u*. Since any Particularist semantic theory has PATH as an axiom, T will have State-Relative-*u* as a theorem. But from this, we can show that T is not a self-articulating semantic theory for *S*. Any token of State-Relative-*u* tokened by an agent in S is guaranteed to be pathological[29]. Therefore, T does not consist of sentence types that are true as tokened by an agent in *S*—and therefore, it cannot be a self-articulating semantic theory for *S*.

[29] See end of Appendix Two for a proof that any token of State-Relative-*u* produced by an agent in state *S* will belong to a pathological dependence structure. A Gaifman-style semantic for tokens produced by an agent in state S will predict that any token of State-Relative-*u* receives the value GAP.

The upshot of the revenge argument, in my view, is that for any stable semantic state *L* there is no set of sentences that counts as a complete, self-articulating semantic theory for *L.* I will call this result *State Limitation*.

State Limitation is strange. It rules out as impossible a descriptive semantic project that is intuitively worth pursuing: formulating a complete semantic theory for tokens produced in one's own stable state. But, on reflection, I do not think State Limitation is particularly problematic. It is one of a variety of ways in which our ability to express truths is constrained. For instance, we are familiar with limitations that derive from our spatiotemporal location. Given the causal constrains on reference and the development of concepts, it is doubtful that ancient Persians could say anything true about iPads or quarks. It is not shocking that our expressive capacities should come with *in principle* constraints either. It is a familiar fact in biology that the structures that sustain abilities also engender limitations: the skeletal structure that gives us the capacity to walk on two legs inhibits our ability to slither like a snake[30]. Stable states are powers to produce linguistic objects the meaning of which are stable across tokens of the same type. And this very stability ensures, for sophisticated states, that some types are associated with meanings that guarantee that they can only be tokened pathologically.

It is useful to compare State Limitation with the expressive limitations occasioned by "contingent Liars." For any circumstance in which a sentence can be tokened, we can construct a sentence type that (relative to a state) will get trapped in a pathological network when tokened in that circumstance. For my current state, we can use the following schema to construct such sentences, where *C* is the circumstance of tokening:

Every token of this very sentence produced in my current state in *C* is not true.

[30] To use a favorite example of Chomsky's.

One instantiation of this sentence is "Every token of this very sentence produced in my current state painted on the door of the White House is not true." I can produce many true tokens of this sentence. But I cannot produce a true token of it by painting it on the door of the White House. This is a kind of expressive limitation, but there is nothing objectionable about it—it is a limitation we should expect to obtain, given what the sentence means in my current state.

Revenge sentences like *State-Relative-u* illustrate a limit case of this phenomenon: the relevant circumstance is just *being in a certain stable state*. This sort of circumstance is quite special, though, because it is the very circumstance that enables an agent in that state to produce tokens whose truth-conditions depend systematically on their linguistic type. State Limitation emerges from the fact that the very conditions that allow us to produce indefinitely many truths in a systematic way also constitute circumstances that guarantee pathology for certain expression types.

Now that we have a fuller understanding of the sort of expressive limitation revealed by the Revenge argument, I will address the objections raised at the end of Section 2: does State Limitation threaten the Particularist's claims to preserve Unity and succeed in the Diagnostic Project?

4.1 Unity

State Limitation is reminiscent of Tarski's claim (1936, 1946) that no language can be semantically closed, and it is illustrative to compare them. For Tarski, a language is semantically closed if it contains its own truth-predicate. Tarski's explicit goal was to investigate the "semantic conception of truth," and he characterized this idea by saying that a predicate T_L expresses the semantic conception of truth, for a given language L, if it "makes assertable" every

instance of the T-schema, where "s" is to be replaced by term referring to a sentence of L and "p" by that sentence or a translation of it:

s is T_L if and only if p

Tarski's Undefinability Theorem shows that, for any language *L* rich enough to express a theory of its own syntax, for any open sentence *P(v*) with only the variable *v* free, *L* will contain a sentence, λ, for which the syntactic theory for the language proves ¬(P(<λ>) ↔ λ), where "<λ>" is a code of λ. In the case where "P" = "T_L," λ is effectively a Liar sentence, and the syntactic theory proves an exception to the T-schema: ¬(T_L(<λ>) ↔ λ). Since, according to the semantic conception of truth, something only counts as a truth-predicate for given language if it makes assertable every instance of the T-schema, then the Undefinability theorem not only shows, for a sufficiently rich language *L*, that T_L is not definable in *L*—it shows that it is inconsistent to suppose that any predicate in *L is* a truth-predicate for *L*. And the same reasoning shows that no expressively rich language could contain a *universal* truth-predicate[31].

[31] Tarski holds that any characterization of a sentence as true or false implicitly refers to some language that gives the sentence an interpretation. So, we can represent a universal truth-predicate as a predicate taking (sentence, language) pairs as arguments. A predicate "T" would express the universal semantic conception of truth if it made assertable every instance of the following schema:

(s, *L*) is T if and only if p

where "s" is replaced by the name of a sentence, "*L*" with the name of a language, and "p" with a sentence that translates the former sentence as it is interpreted in the language referred to. Since any sentence interpreted relative to a language *G* counts as a translation of itself relative to *G*, if *L* is the metalanguage as well as the object language, the schema should hold whenever "p" is replaced by the very sentence referred to be the substitution for "s". However, any language that can express its own syntax and refer to itself will contain a sentence, λ, such that the syntactic theory for the language proves ¬(T(<λ>, *L*) ↔ λ). This is an exception to the Universal T-schema. Since a predicate only counts as a Universal truth-predicate, on the semantic conception, if it implies every instance of the Universal T-schema, it follows that no language rich enough to describe its own syntax contains such a predicate.

The Particularist departs from Tarksi's approach in two crucial ways. First, for the Particularist, predications of truth do not involve any implicit or explicit reference to languages or schemes of interpretation. This is forced on Tarski, because a sentence type can only be assessed for a truth value *relative to* some interpretation or other. But a token sentence, in virtue of being produced by an agent in a particular stable state, will simply *have* a meaning and, often enough, *have* a truth value. As we saw, if we want to formulate a systematic Particularist semantic theory, it will have to recognize stable semantic states as mediating inferences from types to tokens, but it does not treat truth as a *relation* between tokens and states. A truth predicate in a Particularist semantics (whether in the object language or the metalanguage) expresses a monadic property of tokens.

The Particularist also differs from Tarski regarding what it means for a linguistic expression to be a truth predicate. By stipulating that he is interested in predicates that express the semantic conception of truth, Tarski guarantees that a truth predicate (for a language *L*) is *defined* by a theory that implies every instance of the T-schema. The denial of semantic closure, then, implies that, for any language *L,* there is no predicate in *L* that expresses what a truth-predicate for *L* expresses[32]. By contrast, I have simply started with the naïve assumption that there is a property of token representations—roughly, representing things correctly—that words like "true" and "verdadero" function, relative to certain states, to express. I do not require, that in order to express this property, these expressions must be suitable for stating a particular kind of semantic theory. The distinctive contribution of the Particularist response to the paradoxes is the claim that predicates that function to express truth will *malfunction* when tokened in particular

[32] Tarski might say "no *consistent* language." I think the idea of an "inconsistent" language is a confusion that derives from thinking of interpreted languages as individuated by reference to a theory of truth (c.f. Burge 1979).

instances. For example, although L1 is pathological and L2 is true, the predicate "true" in each token functions to express the same property. This is crucial to the analysis of these tokens—it is only because "true" in L1 expresses *truth* that it initiates a dependence chain that is ungrounded.

These two points allow the Particularist to consistently maintain that, although there cannot be a universal semantic theory, the truth-predicates we find in natural language are *universal* truth-predicates. We address semantic theories to tokens produced in a specific state not because arises, as in Tarski, the notion of truth is interpretation relative, but because we only have sound empirical grounds for making systematic predictions about tokens based on their form if we restrict our predictions to a specific state. And State Limitation does not imply that the truth-predicate for a given state S expresses some *limited version* of the property truth. It just reflects the fact that some tokens of the truth predicate are systematically guaranteed to malfunction as tokened in that state.

Since Tarski's work, a great deal of the literature on the Liar paradox has centered on projects that involve adding a truth-predicate to a formal language, e.g., the language of arithmetic, whereas account I have offered depends crucially on thinking of truth-bearers as the concrete outputs of a naturally occurring capacity. It is natural, then, to wonder what light Particularism sheds on semantic paradoxes in formal languages. Formal languages differ from natural languages in a way that, in my view, is crucial: we give formal language expressions whatever meanings they have *by stipulation*. Tarski's theorem shows that us that a certain kind of stipulation is (classically) inconsistent: if *L* is a formal language capable of expressing its own syntax, then it will be inconsistent for us to suppose that there is any predicate *P* in *L* that universally satisfies the schema $P(\langle s \rangle) \leftrightarrow s$. This means that we cannot stipulate, for such a language, that any of its predicates unrestrictedly validate the T-schema. In my view, Tarski was

essentially right about the upshot of this: such languages cannot contain a predicate that expresses truth. For, insofar as we believe that anything we can prove is true and nothing we can prove is untrue, we cannot accept that any predicate *P* in a language *L* expresses the property of being true if our syntactic theory for *L* allows us to prove, for some sentence s, $(\neg P(\langle s \rangle) \wedge s) \vee (P(\langle s \rangle) \wedge \neg s)$[33]. This does not imply that we cannot ascribe truth to sentences of an interpreted formal language. We do this by using tokens of our natural language truth predicates. We can also define predicates in formal languages that express properties approximating truth. We just cannot add a predicate to an expressively rich formal language that, in every one of its occurrences, expresses the universal property *truth* that we express with natural language tokens. To summarize: when we are semantically characterizing a formal language, we do not, as in natural language, confront the "the Liar paradox" as an empirical problem[34]; our task is the creative one of defining predicates that suffice to approximate truth for whatever our mathematical or metalogical purposes happen to be.

4.2 The Diagnostic Project

Why should we think that State Limitation conflicts with success in the Diagnostic Project? Why, in other words, should we think that, in order to succeed in the Diagnostic Project, we need to frame a semantic theory, within a state *S*, that includes every instance of RTS for sentences of in *S'*s domain? Bacon suggests that any theory that doesn't satisfy this condition

[33] Which follows classically from Tarski's Theorem. For reasons of space, I am presenting this argument in a highly compressed form.

[34] Whatever contradictions we derive in such a setting are the result of an inconsistent stipulation for the meaning of the supposed truth-predicate—not a consequence of the predicate's antecedently given meaning.

will be formally compatible with *false positives.* That is, we could consistently add to such a theory a sentence of the form:

x is a non-pathological token of *s* produced in *S* and it is not the case that *x* is true iff *s*

where *S* is the state in which we are developing our theory. This sentence would say of some token that it is healthy but nonetheless fails to preserve truth across disquotation within the state *S*. But since our "pathology" predicate is meant to apply to all tokens that fail to preserve truth across disquotation within a single state, this shows that our theory does not rule out the existence of Liar-like objects that are *non*-pathological. So, our pathology predicate is not fit to diagnose what goes wrong with the Liar.

There are two points I would make in response. First, the Revenge argument shows that no set of sentences that is closed under classical consequence can be a self-articulating semantic theory for a given state. But it has not provided, for an agent in a state S, some instance of the RTS for S that they will not be able to truly articulate. Consider the *u* instance of State-Relative-RTS:

State-Relative-RTS-u For all *x,* if *x* is a non-pathological token of *u* produced in state *S* → (*x* is true ↔ no token of *u* produced in *S* is true)

On a Gaifman-style semantics for tokens produced in *S*, tokens of this sentence will come out true—an agent in *S* can produce such tokens truly. So, the Revenge argument does not show that the Particularist cannot assert (or must deny) tokens of the RTS instance for *u.* The puzzling thing is that this sentence type classically implies *u*, which can only be tokened pathologically by an agent in *S*. The thrust of the Revenge argument, as Bacon presents it, is that an agent in *S* who endorses a token of RTS-*u* is logically committed to endorsing tokens of *u*, thereby endorsing something pathological. But this is too quick. Logical consequence is typically understood as a relation between (sets of) sentence types, and it is not a trivial matter to extend the notion to

collections of tokens, given the Particularist account of semantic pathology. For instance, any Particularist will want to deny a simple account according to which, if the sentence type *p* classically implies the sentence type *q,* then all tokens of *p* imply all tokens of *q*. Otherwise they would be committed to saying that, for instance, L2 implies L1, by Repetition. My view (similar to that defended in Warren 2023) is that semantically pathological object do not stand in entailment relations to anything. The onus, admittedly, is on the Particularist to offer a fully developed account of proof and entailment among tokens, which I am not providing in this paper. But it is not fair to judge the issue in advance and simply assume that any agent in *S* that asserts an instance of RTS-*u* thereby logically committed to any token of *u.*

The second point is that a Particularist can have good reasons for failing to assert every instance of RTS, without it being fair to say that their "pathology" predicate has false negatives. On the natural extension of the Gaifman semantics that includes "pathological" as a semantic predicate, some token instantiations of RTS will *malfunction.* Consider this instance, where "*t*" is the name of the token on the next line:

t For all *x,* if *x* is non-pathological token of "*t* is not true" produced in *S,* then *x* is true iff *t* is not true.

A Particularist in state S can justifiably decline to endorse this token on the grounds that, in so doing, they would endorse a token that malfunctions[35]. And, if so, they will not endorse every token instance of RTS. But this does not show that they are open to their pathology predicate having false negatives. A pathology predicate has false negatives if it is *false* of some object that exhibits Liar-like semantic malfunction, and no such object has yet been shown to exist.

[35] *t* gets trapped in a closed loop consisting of itself and all tokens of "*t* is not true" produced in *S*.

For both of these reasons, it is wrong to think that, unless the Particularist can produce a self-articulating semantic theory that includes every instance of RTS, their theory of semantic pathology is open to false negatives. So, we should reject Bacon's characterization of what it takes to succeed in the diagnostic project. But, how then should we understand it?

The diagnostic project is to identify a feature that is common to all objects that exhibit Liar-like pathology and explains why they are pathological. The general account of semantic pathology offered by the Particularist is this: tokens of semantic predicates, like "truth," give rise to dependence relations among tokens, and any token that is part of a network of dependence relations that constitutes a closed-loop or infinite descending chain is pathological[36]. This account, to be sure, only goes part of the way in yielding predictions about which tokens are pathological. To make such predictions, one would need to know, in addition, what tokens there are and what states they were produced in, and one would need a detailed semantic theory for tokens produced in those states. But these two projects—giving a general account of semantic pathology and giving a systematic semantics for tokens produced in some specific state—are, in principle, separable. State Limitation puts restrictions on the latter project. For the Particularist, given the semantic state an agent occupies, there will be restrictions regarding which states they

[36] As I acknowledge in footnote eight, there are questions to be raised about whether this is something the Particularist can say without saying something ungrounded. Is this just the Revenge problem rearing its head in a new place? There is a certain similarity: these questions raise doubts about whether the Particularist can present their theory without saying something that, by their own lights, is not true. Of course, the Particularist has to address these questions (I sketch my preferred answer in footnote 8), but there are a number of reasons to think that they raise issues that are separate from those raised by traditional "Revenge" paradoxes. For one thing, expressions generating Revenge are supposed to replicate Liar-like reasoning. The apparently ungrounded generalizations that a Particularist wants to make (e.g. tokens of "Nothing true is false," "Nothing ungrounded is true.") are intuitively unproblematic and can be counted as grounded and true by fiat without generating contradiction or absurdity. It seems to be a problem that can be localized to the analysis of semantic pathology in terms of grounding.

can express a complete semantic theory for. But that does not imply any restriction on an agent's ability to express the general account of semantic pathology offered by the Particularist, or imply that the account is mistaken.

One might object here that the Particularist, in their general account of semantic pathology, is helping themselves to language-transcendent concepts that they are not obviously entitled to. For instance, the account given above helps itself to the idea that some predicates (as tokened in a given state) are *semantic predicates*. But this notion is basically introduced by example, using "truth" in the mouths of English speakers as a paradigm. How, it might be pressed, do we generalize from this case to a general concept of a "semantic predicate" applicable to expressions that might give rise to pathological tokens in any state? Have we really succeeded in giving a general diagnosis of semantic pathology if we do not have a definition of what it is to be a "semantic predicate" relative to a given state?

These questions should not give rise to skeptical worries. For one thing, there may be a general, non-circular definition of semantic predicates[37]. But even if there is not, I do not think that the Particularist who uses "semantic predicate" as a primitive is worse off than the biologist who starts investigating cows without defining "cow." We encounter semantic predicates in our own languages, and we generalize to other cases based on their similarity to paradigms like "true," "applies to," and "refers to." One need not be in possession of a definition to think that "semantic predicate" picks out a natural grouping that can be used to give a fully general account of semantic pathology.

[37] Dalglish (2020) claims that a notion is semantic if it is constitutively associated with principles of default inference in which expressions are both used and mentioned. I think this is informative and probably correct, though I am not sure if the use/mention distinction or the relevant notion of "constitutive principles" can be explicated without going in a circle that involves truth and other obviously semantic notions.

5. Conclusion

Many authors have argued that Particularist responses to the semantic paradoxes, like any other response that aims to preserve classical logic, face Revenge objections: either they refute themselves or they incur an objectionable kind of expressive limitation. I have argued that the sort of expressive limitation the Particularist must accept—State Limitation—is really quite innocuous. It is compatible with thinking that the Particularism gives a correct and general explanation of semantic pathology, and it is also compatible with Unity—the idea that all truth-predicates function to predicate a single property. Since Revenge arguments are often thought to be *the* central problem faced by classical responses to the semantic paradoxes, it counts strongly in Particularism's favor that it can accept the results of Revenge arguments unscathed.

CHAPTER 2

The Problem of Unarticulated Truths

There is a tension implicit in the common-sense platitudes that characterize truth. On the one hand, we have ideas like the following:

(1) Truth is the correspondence of a thought with its object. (See Aquinas 1952, Descartes 1991 AT II:597, Kant 1998 A58/B82.)

(2) To say of what is that it is, is true. (Aristotle Metaphysics 2016 1011b25)

The picture suggested by these claims is that truth is fundamentally a property of representational vehicles, e.g. thoughts or acts of saying. By "representational vehicles" ("RVs," for short) I mean concrete particulars with semantic properties, produced by some cognitive agent at a particular time.[38] I will leave the notion some- what vague, hoping that examples—utterances, belief states—are a sufficient guide for now. I will call the thesis that the fundamental truth-bearers are token representational vehicles the RV-First view.

On the other hand, it is quite natural to talk as if truth is utterly independent of any representing vehicle. This is the perspective we take when we say something like (3):

(3) There are truths about the Milky Way that will never be articulated by anyone.

(3) is *prima facie* incompatible with the view that truth is fundamentally a property of representational vehicles. No representational vehicles could be the "truths" that count as witnesses for the generalization in (3). On the most natural understanding of "articulation" applicable to representational vehicles, a representational vehicle is articulated just in case it is

[38] I mean concrete in the sense isolated by Hoffman and Rosenkrantz (2005). Even if mental states are modes of Cartesian souls, they will count as concrete on this understanding.

produced or brought into existence. But according to this reading, (3) would incoherently suppose there are representational vehicles that are unproduced. Rather (3) seems to be talking about a class of truths that are not representational vehicles, and, moreover, what it says is intuitively plausible. Probably there are objects in the Milky way that no one has or will ever think or talk about— but how could there fail to be *truths* about what those objects are like?[39]

How should this tension be resolved?[40] One promising line of thought goes as follows: the fundamental truth-bearers are particular RVs, but other things can have truth-values in a derivative sense; in particular, we can group RVs into kinds and derivatively characterize those kinds as true or false. This line of thought meshes well with the recent trend in the propositions literature to analyze propositions as *kinds or abstract groupings* of more fundamental, concrete representational things. I take Soames (2014), Hanks (2011, 2013), and Grzankowski and Buchanan (2019) as representatives of this approach.[41] If we posit *kinds* of RVs and suppose that they are truth- bearers in a derivative sense, then it is plausible that we can accept both (3) and

[39] One finds perhaps the purest expression of this perspective in Frege: "What I acknowledge as true, I judge to be true quite apart from my acknowledging it's truth or even thinking about it. That someone thinks it has nothing to do with the truth of a thought" (1997 p. 342). One plausible reaction to (3), championed by Frege among others, is to see it as supporting the view that the fundamental truth-bearers are members of a sui generis ontological category—propositions or Gedanken—that are "abstract" Hoffman and Rosenkrantz's (2005) sense.

[40] Historically, many philosophers have favored a theological solution: there is a divine mind that necessarily exists and supplies vehicles for the truths that would otherwise appear to lack them Augustine (2010 p. 395) and Leibniz (1714) even offer arguments for the existence of God on the basis that necessary truths must subsist in a necessarily existing substance.

[41] Tyler Burge expresses a similar point of view, though not in the service of developing a detailed theory of propositional content (see 2010 p. 37–42 and 2019 p. 45). Jeff King's (2007) account of propositions is also a core example of an attempt to ground the existence and truth-conditions of propositions in the activities of real cognitive agents. However, on King's view propositions are not types that are instantiated in particular representations; they are a special sort of existentially quantified fact.

the RV-First View. For there is good reason to think that there are many true RV-kinds that are never instantiated, and these can serve as witnesses for the generalization in (3).[42]

Of course, the merits of this line of thought really rest on the *account* it gives of what it is for an RV-kind to be true. In this paper, I will be considering and the exploring the consequences of one such account:

> *D* For an RV-kind *x* to be true is for *x* to be such that, were it instantiated, its instances would be true.

I take this to be the most intuitive and simple way of characterizing what it is for an RV-kind to be true.

In Section 1, I introduce D and explain the sense in which I take it to define truth as a *derivative* property of RV-kinds. In Section 2, I introduce a crucial distinction for RV-kinds between being true-at a possible world vs. being true-in a possible world. In Section 3, I develop what I take to be the strongest putative counterexamples to D and argue that they fail.

1. Truth and RV-Kinds

It will be helpful as a preliminary to situate my project with respect to debates about the nature of propositions. If we take a broadly "abstractionist" approach to propositions, according to which they are essentially kinds, types, or reified equivalence classes of token RV's, then it is

[42] I should mention that Soames makes some remarks that suggest an alternative strategy for dissolving this puzzle. In the context of considering whether there really are enough cognitive-act types to play the role of all the propositions we are prepared to accept as true, he claims that propositions do not need to exist in order to be true (2014 p. 103). For some properties, he suggests, an individual can have that property without existing—for instance, Socrates has the property of being admired, though he does not exist (see also Salmon 1987). Truth is a property like this. In the present context, we could adopt Soames' strategy and say that the truths about the Milky Way quantified over in (3) are particular RVs that are true but do not exist. Frankly, I find this to be a non-starter—on the understanding of "exist" at issue, I feel compelled to deny that *there is* anything that does not exist. There are more ontologically plausible interpretations of the examples Soames gives, in my view.

natural to think that what I am calling RV-kinds just are propositions. I do not object to this interpretation but neither do I insist on it—for two reasons. First, the puzzle raised by (3) is a problem for any RV-First view, and I aim to provide a defense of the RV-First perspective that is maximally general. So, I do not want to make assumptions about RV-kinds that would privilege one particular account of propositions over another. Second, my own view is that there is no partition of token RV's into equivalence classes that *uniquely* deserves to be called the partition according to *sameness of content*. So, for any RV token, there may be multiple kinds that are equally good candidates for the proposition expressed by that token.

That said—for the purposes of making things concrete, it will by helpful to adopt Hanks' view when discussing particular examples. For Hanks, the relevant RV-kinds are complex action types, the instances of which are dated particular acts of predication. Particular assertions are paradigm cases of "predications" in Hank's sense. Since, for Hanks, declarative mood conventionally encodes predication, by uttering literally a declarative natural language sentence one thereby makes an assertion. So, my producing the following token sentence constitutes a token predication:

(4) Magnus Carlsen is clever.

By producing this token, I assert that Magnus Carlsen is clever. This is an instantiation of an action type that Hanks represents in the following way:

(5) |-<**Magnus Carlsen**, *clever*>

the This is the sort of action that consists in (a) referring to Magnus Carlsen (via the name "Magnus Carlsen"), (b) expressing the property of being clever, and (c) predi- cating the property expressed of the object referred to. For the purposes of thinking through my examples of RV-tokens and RV–kinds, it will be fine to understand them on the model of (4) and (5).

However, I will often (sloppily, from Hanks' perspective) simply treat token natural language sentences themselves as RV-tokens, rather than the token actions that consist in uttering them.

One key consequence of Hanks' account is that RV-kinds essentially mark out token representations that have the same truth-conditions: necessarily, any token predication that instantiates (5) will be true if and only if Magnus Carlsen is clever, because any such token is a predication of cleverness of Magnus Carlsen. This consequence—that RV-kinds *rigidly* pick out tokens that form an equivalence class with respect to their truth-conditions—is a necessary feature of anything that counts as an RV-kind, as I am using the term.

In the arguments I will help myself to the following schematic principles, where "p" is to be replaced by any declarative English sentence:

Referring to RV-Kinds

Expressions of the form "The claim that *p*" refer to RV-kinds.

Rule of Thumb Condition for Kind Membership

> A token of a sentence *p* will generally be a member of the kind referred to by tokens of the expression "the claim that *p*."

The *Referring* principle expresses my view that, when we speak about "claims" or "statements" in abstraction from any particular event of claiming or stating, we are referring to RV-kinds. The *Rule of Thumb Condition* provides a bridge between talking about particular RVs and talking about the kinds they instantiate. There are obviously cases in which this Rule of Thumb fails. For instance, the claim that I am hungry is not a claim that LeBron James might make by uttering the sentence "I am hungry."[43] Since claims are meant to track (at minimum) sameness of truth-conditions, to the extent that a sentence type has context dependent truth-conditions, not all

[43] At any rate, the sense of "claim" according to which all tokens of "I am hungry" express the same claim does not out an RV-kind that has a truth-value.

tokens of that sentence will be instances of the same claim. I will not offer a general theory of the conditions under which two RV's should be counted as instances of the same claim.[44] However, all the cases in which I apply the Rule of Thumb— jumping from an RV-token to the kind it instantiates—are in conformity with ordinary-language usages of "the claim that *p*" and should be uncontroversial.

As I announced in the introduction, I propose to define *truth* as a property of RV- kinds, in the following way:

> *D* For an RV-kind *x* to be true is for *x* to be such that, were it instantiated, its instances would be true.

I intend this as a real definition of truth as property of RV-kinds: a statement of *what it is* for an RV-kind to be true. Hanks explicitly endorses *D* in his explanation of the manner in which cognitive act-types inherit truth-conditions from cognitive act tokens (2013 p. 568); a good case can be made that it follows from Soames' commitments as well.[45] The counterfactual conditional in *D* should be understood in terms of the Lewis-Stalnaker semantics: in all of the words that are relevantly simi- lar to our own, if *x* is instantiated at those worlds, then *x*'s instances are true in

[44] Indeed, I am skeptical there is a single theory adequate to all purposes for which we use words like "claim" or "statement."

[45] Strictly, Soames says that a true proposition (cognitive act-type) is true in virtue of the fact that it represents things as being certain way and things are that way (Soames 2014 p. 96). But a proposition *p* represents things as being a certain way virtue of the fact that "all conceivable instances" of *p* represent things as being that way (*ibid* p. 96). In lieu of positing an ontology of "merely conceivable instances," this claim seems appropriately paraphrased as: a proposition *p* represents things as being a certain way in virtue of the fact that, were it to be instantiated, its instantiations would represent things as being that way. If this paraphrase is fair, then, since a token cognitive act is likewise true in virtue of representing things as they are, Soames' characterization of truth for propositions, plausibly, can be partially reduced to the following: a proposition (cognitive act-type) is true in virtue of the fact that, were it to be instantiated, its instantiations would be true.

those worlds. A world is "relevantly similar" to our own, roughly, if and only if it differs at most in the existence of a token expression of the relevant RV-kind.[46]

I propose that the "truths" that serve as witnesses for the generalization in (3) are RV-kinds. For, it is highly plausible that (3)' is true, and (3)' posits RV-kinds that, given *D*, make the generalization in (3) true:

> (3)' There are RV-kinds that (a) are about the Milky Way that (b) have never and will never be articulated by anyone and (c) are such that were they tokened, their tokens would be true.

For instance, I think the claim consisting of the claim that the Milky Way is a galaxy conjoined with itself twenty thousand times is (a) about the Milky Way, (b) has never and will never be tokened, and (c) is such that, were it to be tokened, its tokens would be true. I take it that (a) and (b) are *prima facie* plausible, but I will spell out the reasoning for (c). Let's call the claim consisting of the claim that the Milky Way is a galaxy conjoined with itself twenty thousand times *c.* Take an arbitrary world *x* that differs from the actual world at most in that *c* is instantiated. In *x* the Milky Way is a galaxy, for otherwise *x* would differ drastically from the actual world, contrary to hypothesis. Now suppose that *c* is tokened in *x*. Necessarily, if

[46] One might doubt that this is coherent: surely two worlds cannot differ *solely* in the fact that in one world includes a token representation that the other does not. This is correct, but I do not think it raises any special problems for understanding *D*. We can say that a world *v* differs from the actual world at most in the existence of token *t* if all the facts that distinguish *v* from our world are: (a) constitutive of the existence of *t* or (b) entirely grounded by the fact that *t* exists combined with facts hold in the actual world. So, for example, if we consider a world *z* in which a token representation *s* exists that does not exist in the actual world, there are some differences between *z* and the actual world that are compatible with *z* differing from the actual world *at most* in the existence of *s*: for instance, *z* will contain all *s*'s parts, which the actual world may not contain; and *z* will be such that *s* exists and there are horses, whereas that conjunctive fact does not hold of the actual world. These latter facts are, respectively, (a) constitutive of the existence of *s* and (b) entirely grounded by the fact that *s* exists combined with a fact that holds in the actual world (there are horses).

something is a token of *c* it is true if and only if the Milky Way is a galaxy—because all such tokens consist of a conjunction each of the elements of which is true if and only if the Milky Way is a galaxy. Therefore, since *x* is a world in which the Milky Way is a galaxy, *x* is a world in which *c*'s token is true. Since *x* was an arbitrary world, we can conclude that *c* is such that, in all the worlds that differ at most from the actual world in that *c* is instantiated, *c*'s tokens are true. This proves (c), on the intended Lewis- Stalnaker understanding of the counterfactual. (3)' is consistent with the RV-first view. Since (3)' in combination with *D* implies (3), and *D* is a principle RV-first theorists should be happy to accept, this amounts to showing that (3) is consistent with the RV-first view.

Does *D* characterize RV-kinds as true in a *derivative*, rather than fundamental, sense? One might object to this along the following lines. Trenton Merricks character- izes a "fundamental truth bearer" as an object that (a) has a truth-value and (b) is such that its possessing that truth-value is not (even partially) analyzed in terms of its being related to another object that has a truth value (2015 p. 22). This definition suggests that some RV-kinds might be fundamental truth-bearers. Consider one of the RV-kinds about the Milky Way that serves as a witness for (3)' above: call it *a*. By hypothesis, *a* is true and has never been tokened. The analysis I would give of *a*'s being true is that *a* is such that, if it were tokened, its tokens would be true. But this analysis does not mention some distinct object, standing in a relation to *a*, that is actually true—it just ascribes a modal property to *a.* So it seems that, according to Merrick's characteriza- tion of what makes a truth-bearer fundamental, *a* should count as a fundamental truth bearer. Since *a* is not an RV-token, this result is inconsistent with the RV-first view.

However, I think that my definition of truth for RV-kinds still marks them as derivative truth-bearers in a recognizable sense, even if that sense is not Merricks'. They are derivative

truth-bearers in the sense that, in the statement of what it is for an RV-kind to be true, truth is predicated of another sort of thing. Truth as property of RV-tokens is "prior in definition" to truth as a property of RV-kinds.[47] A strong way of putting it is that, on my view, a truth-predicate for RV-kinds is *eliminable*— we could just as well talk about RV-kinds that could have true tokens, rather than true RV-kinds. The fact stated by (6) reduces to that stated by (7):

(6) The claim that Shakespeare wrote plays is true.

(7) If the claim that Shakespeare wrote plays were tokened, its tokens would be true.

Does this mean that *strictly* speaking the property of truth does not apply to RV- kinds? Grzankowsi and Buchanan (2019) recommend this position. But if we accept that no RV-kind is really true then my proposed analysis of (3) fails—and, anyway, I do not find the view well

[47] I have heard the following objection about *D*: *D* explains the truth of RV-kinds in terms of the truth of *possible RV tokens*. But surely this is problematic from the RV-first perspective: merely possible RV- tokens are not actual concrete RV's. If we need to posit merely possible RV-tokens to explain full story about truth, then it is just false that concrete RVs are the fundamental truth-bearers. This line of thought is misleading. Plausibly, the sense in which *D* quantifies over "possible tokens" is ontologically noncommittal. With many philosophers, I reject the Barcan Formula, and so reject the inference from, e.g., "There might have existed a cake I baked yesterday" to "There exists something that could have been a cake I baked yesterday." For the same reasons, I would object to an interpretation of *D* according to which it is committed to the existence of things that could be RV-tokens. Or to the existence of non- actual "merely possible things," which I regard as a confused idea (see Williamson 2013 p.22–23). (Of course, some philosophers have argued for the Barcan formula and the view that ontology is necessary (see Williamson 2010, 2013; Fritz and Goodman 2017). My sense is that this remains a minority position, and the issues it raises are mostly orthogonal to my present concerns.) *D* ascribes a modal property to an RV-kind; the property will be modeled in a possible-worlds semantics framework by means of individuals inhabiting alternative possible worlds, but I am agnostic about the ontological commitments involved in such modeling (see Stalnaker 2012 p.22–51 for an "actualist" take on these matters). So it is somewhat misleading to say, as I do, that, in *D*, "truth is predicated of another sort of thing"—there are no *things* (possible tokens) such that *D* ascribes truth to them. The important point is that truth as a property of RV-kinds is a complex modal property, the nature of which is made intelligible in terms of truth as property of RV-tokens. (Thanks to Bill Melanson for pressing me on these issues.).

motivated. There is an intelligible reason why we extend our concept of truth to RV-kinds, even if what it is for a RV-token to be true is not exactly what is for an RV-kind to be true. This is an instance of a pattern that is utterly general regarding the ascription of properties to kinds. For an individual car to be red is (very roughly) for the parts on its exterior to reflect certain wavelengths of light. A *type* of car, e.g. the 1998 Dodge Viper Red, can also be red, but it cannot be red in exactly the same way. (In order to reflect light an object has to be located somewhere in space, but a part of a *type* of car, e.g. the hood of the 1998 Dodge Viper Red, is not located anywhere). Ascribing truth to RV-kinds is no more illegitimate than ascribing colors to kinds of material objects.[48]

Perhaps from a metaphysically rigorous point of view, we should say that truth as property of RV-kinds is distinct from truth as a property of RV-tokens, because the real definition of the former cannot apply to the latter. But all the same, natural language truth-predicates can be correctly applied (without straightforward ambiguity) to *both* RV-tokens and RV-kinds. This is what matters, from my perspective, because this is what allows me to say that (3) says something correct. In ordinary usage, some RV-kinds really are true.

2. Truth-in a World vs. Truth-at a World

Is *D* correct about what it is for an RV-kind to be true? If we limit our attention to RV-kinds that actually are instantiated, *D* predicts that these RV-kinds are true if and only if all of their tokens are true. This seems to track our intuitions about what claims are true. For RV-kinds as well as RV-tokens, one of the primary functions of the concept of truth is to license

[48] See Hanks 2016 for a similar argument.

"disquotational" inferences. For instance, consider the following schemata where "*p*" is to be replaced by some English sentence.

(8) The claim that *p* is true	(8)' *p*
(9) *p*	(9)' The claim that *p* is true

In many ordinary cases, the token instantiated in line (9) or (8)' will be an instance of the claim that that is referred to in line (8) or (9)'. When this holds, the inference is a good one.

(10) The claim that there are no unicorns is true.	(10)' There are no unicorns.
(11) There are no unicorns.	(11)' The claim that there are no unicorns is true

When we know that a given RV-kind is true, we can be confident in endorsing a particular instantiation of it. If these sorts of inferences were not in general valid, it is difficult to see what the purpose of extending the concept of truth to RV-kinds would be. But in order for these inferences to be valid, it must be the case that, for any instantiated RV-kind *x, x* is true if and only if all its instantiations are true. Otherwise the argument schemata above might take one from a truth to a falsehood. So, if there are counterexamples to *D*, they will not be found among actually instantiated RV-kinds: there are no true claims that are instantiated in untrue acts of claiming, or true acts of claiming untrue claims.[49] The potentially controversial cases concern untokened RV-kinds.

[49] One might think that there are examples of claims for which this does not hold. In my view this reflects the fact that, in ordinary language, we are promiscuous about what sorts of equivalence classes we want to pick out with the word "claim." There is an intuitive sense in which two people who say "I'm hungry" have made the same claim, though their particular acts of claiming might have different truth values. One the other hand, it would be strange to talk of *the* single claim that both of these people made as having *one* truth-value—which tends to confirm *D.* Some cases are a bit trickier, for instance, cases where "claim" seems to pick out a temporal proposition: e.g. "the claim that China's growth is slowing is true now, but it was false when economists put it forward in 1990." This example suggests that the claim that China's growth is slowing is true but it was tokened falsely in 1990. This is not inconsistent with *D* on its own, but inconsistency will arise if we add some natural assumptions, i.e. that the tokens

Before we turn to that however, it will be instructive to consider an objection to the effect that the definition does not hold of necessity, even if it is extensionally correct. Consider the RV-kind instantiated by (12):

(12) No RV-tokens exist.

the (12) suggests the following argument. Consider a possible world *w* in which no RV- tokens exist. In *w*, the RV-kind instantiated by (12)—the claim that no RV-tokens exist—is true. But, in *w*, it is not the case that if the RV-kind instantiated by (12) were tokened, those tokens would be true. In all the closest possible worlds to *w* in which an RV-token *x* instantiates the claim that no RV-tokens exist, *x* is false—it falsifies itself. Therefore *w* is a world in which the RV-kind instantiated by (12) is true, but it is not the case that, were that kind tokened, its tokens would be true. This implies that the proposed definition does not hold of necessity.

To my mind, this argument goes wrong in its first step: inferring from the fact that, in *w*, no RV-tokens exist, that, in *w*, the claim that no RV-tokens exist is true. We need to introduce a distinction, for RV-kinds, between being *true-at* a world versus being *true-in* a world.
The *truth-at/truth-in* a world distinction distinguishes between what can truly be said, in our actual circumstances, of various counterfactual circumstances, versus what could be truly said in those counterfactual circumstances, if they obtained.[50] If we predicate truth of contingent objects

---

expressions from 1990 have not changed their truth-value. There are a variety of ways a defender of *D* could deal with this issue, but, since the details are not entirely pertinent to my present concerns, I will leave them for another time. (It seems promising to me to suggest that, since timeless truth is more fundamental than time relative truth, the sort of RV-kinds that count as the paradigm case for the application of the truth-predicate to RV-kinds will be timelessly true.)

[50] This follows Williamson's formulation in (2013 p. 297), though he speaks of "truth of" not "truth-at". One can find similar presentations of the distinction in Mates (1970), Fine (1985), Davies and Humber- stone (1980) and Stalnaker (2011). The distinction is sometimes credited to the fourteenth Century philosopher Jean Buridan, who gives a lucid presentation of it in his Sophismata (1966).

like RV-tokens, there is no avoiding such a distinction, and (12) illustrates this nicely. Assuming that RV-tokens are contingent beings, there are possible worlds in which no RV-tokens exist. Call the class of such worlds G. (12) correctly describes the worlds in G—it is "true-at" those worlds. But none of the worlds in G is such that, were they actual (12) would have the property of being true. For (12) can only be true if it exists (and is an RV-token), but all of the worlds in G are such that, were they actual, no RV-tokens would exist at all. This is the sense in which (12) fails to be *true-in* the members of G.

If we extend the truth-at/truth-in distinction to RV-kinds we will be able to block the argument above right at the beginning. For the argument proceeded in the following way:

(a) *w* is a world in which there are no RV-tokens.

(b) *w* is a world in which the claim that there are no RV-tokens is true.

But this in effect assumes there is no substantive truth-at/truth-in distinction for RV-kinds. If there is such a distinction, there is no guarantee that a world that can be characterized by means of an RV-kind *p* will be a world in which *p* is true. This distinction can be used to explain the invalidity of the analogous inference for the case of sentence tokens:

(a)' *w* is a world in which there are no RV-tokens.

(b)' *w* is a world in which tokens of the English sentence "There are no RV-tokens" are true.

Is there a substantive truth-at/truth-in distinction applicable to *claims*, in addition to sentences types or concrete tokens? Here is an argument that there is not: in general, it seems that an object is true-*in* a world if and only if it is true-*of* that world and it exists there (with its actual semantic properties). So, for truth-bearers that exist contingently, it makes sense to mark the truth-at/truth-in distinction, because they may not exist in every world where their descriptive conditions are

met. But if we assume (for the moment) that RV-kinds exist necessarily, the distinction would be idle because RV-kinds would be true-in all and only the worlds they are true-at.

I am not sure if RV-kinds really are necessary existents, but, even assuming they are, I think this argument should be resisted. We do not arrive at the *truth-in* relation via deriving it from the *truth-at* relation plus existence. Rather, *truth-in* is the result of translating talk of "what would be true" into a framework where we explicitly refer to and quantify over possible situations. To say that *x* is true-in a possible situation *y* is just to say that, were *y* realized, *x* would have the property *truth*. So, the question at hand is: if *w* were realized, would the claim that there are no RV-kinds be true? I say no, on the grounds that, even if *w* were realized, that claim would not be such that, were it to be tokened, its tokens would be true. This is a somewhat theory-driven response, but in the present context I do not think it begs the question.[51]

I propose then, that there is a substantive truth-at/truth-in distinction for RV-kinds that we can characterize as follows:

> An RV-kind *x* is true-at a world *w* iff *x* is such that, were it tokened, its tokens would be true-at *w*.
>
> An RV-kind *x* is true-in a world *w* iff, were *w* realized, then *x* would be true.

With this distinction in hand we can unravel the problem for *D* raised by (12). When we consider the world *w* in which there are no RV-tokens, we are considering a world *at which* the claim that there are no RV-tokens is true. But if we take seriously the truth-at/truth-in distinction, this does

[51] I do not think there is much hope of giving a non-theory-driven account of the inference from (a) to (b)—commonsense thinking about truth does not seem to provide us with a neutral verdict that is stable. The transition between (a) and (b) is natural enough, but I expect many would be uneasy endorsing structurally similar inferences, i.e. there are no claims, therefore, the claim that there are no claims is true.

not necessarily imply that, *in w,* the claim that there are no RV-tokens is true. If this inference is blocked, then *w* cannot be put forward as an example in which *D* fails to hold.

3. Counterexamples?

The strongest arguments that might be marshaled to show that *D* has extensional counterexamples also fail, I think, for similar reasons. The arguments I have in mind try to show that there is an RV-kind that accurately describes our world but would not have true tokens if tokened (or misdescribes our world but would have true tokens if tokened). For instance, take some true RV-kind *p* that is never tokened, and consider the following schematically described RV-kinds:

Never-*p* The RV-kind consisting of the conjunction of *p* and the claim that *p* is not tokened.

Sometime-*p* The RV-kind consisting of the conjunction of *p* and the claim that *p* is tokened sometime.

Here is an argument that Never-*p* is true. By hypothesis *p* is true and untokened. (Since I believe there are infinitely many untokened true RV-kinds, I have no reason to object here.) Never-*p* is a conjunction of two claims. The first is *p*, which is true by hypothesis; the second is the claim that *p* is not tokened, which is also true by hypothesis. So Never-*p* is a conjunction of two true claims, and is, therefore, true. But Never-*p* is not such that, were it tokened, its tokens would be true. Any token of Never-*p* has a token of *p* as a part, so any world in which Never-*p* is tokened is a world in which *p* is tokened. And therefore, in any world in which Never-*p* is tokened, its tokens are false, because their second conjunct is false. So Never-*p* is a counterexample to *D*. The argument is parallel with Sometime-*p,* though the conclusion is perhaps more disturbing: Sometime-*p* is false, but *D* counts it as true.

Never-*p* actually raises two different (though related) problems for *D.* First, these cases show that *D* is inconsistent with supposing that that RV-kinds form a Boolean algebra. It can be that two claims *x* and *y* are individually such that, were they tokened, their tokens would be true, but the conjunction [*x* and *y*] is not such that, were it tokened, its tokens would be true. So if *D* is correct, not every conjunction of true claims is itself true. Second, Never-*p* shows us an example of an RV-kind that says of itself, in effect, that it is not instantiated, and that *D* therefore must count as not true even if what it says is, apparently, the case.

The second problem is, in my view, the deeper one. The first problem can be ade- quately addressed by replacing *D* with a recursive definition of truth for RV-kinds based on their composition.[52] For example:

> An atomic RV-kind *x* is true iff *x* is such that, were it instantiated, its tokens would be true.
>
> An RV-kind of the form ~*x* is true iff *x* is not true[53].
>
> An RV-kind of the form *x* ^ *y* is true iff *x* is true and *y* is true.
>
> …

I do not pursue this sort of response in detail for a couple of reasons. First, to develop such an account I would need to endorse a particular theory of RV-kind *structure*. This would violate my aim to give a defense of the RV-First view that is maximally general and depends on no special assumptions about RV-kinds—including that they are structured. I am not sure there is a unique notion of "structure" that applies uniformly to all schemes of classification that count as marking RV-kinds. More importantly, although this solution would disarm Never-*p* and Sometime-*p*, it

[52] Thanks to Ryan Simonelli for suggesting this approach to me.

[53] Of course, this will need to be adjusted if we want to include RV-kinds that are neither true nor false.

would leave the general issue they raise—the second problem—untouched. There are other plausible cases of claims that characterize themselves as uninstantiated. In fact, if we develop a theory of RV-kind structure that recapitulates the structure of First-Order Logic, we will likely be able to *prove* the existence of such RV-kinds via diagonalization.

So let us take the second problem head on. I think that are grounds from within common sense for taking Never-*p* to be false (or at the very least not true). For the following seems to me to be a platitude: a true claim is something that can be truly claimed. And, on this score, Never-*p* fails—it is a claim that, in principle, can only ever be claimed *falsely*. As a representation, it is self-undermining. Looking at Never-*p* from this perspective, it does not seem to stretch ordinary intuitions to count it as false, and I propose we accept this verdict. The case of Never-*p* is a bit odd, but I do not think that common sense provides us with dispositive reasons for thinking that it is true.

The really significant reason in favor of judging Never-*p* to be true derives not from common sense but from logical theory. One might plausibly object: to suppose that a conjunction of two truths can be false involves a radical revision of classical semantics and would seem to imperil the validity of rules like Conjunction Introduction. On reflection, though, I think accepting that RV-kinds do not form a Boolean algebra with respect to their truth-values raises no deep logical problems. This would be problematic if we wanted a *logic* of RV-kinds, i.e. if we thought that RV-kinds were the sort of thing that fundamentally stand in relations like logical implication. (If we are thinking of RV-kinds as propositions, and take propositions to have a central role in logic, this point of view is understandable.) However, from the RV-first perspective it is not obvious that a logic of RV-kinds is something we need.

It is natural to think that the properties of interest to the logician—roughly, properties that underlie truth-preservation due to logical form—are instantiated by the *fundamental* truth-bearers. For the RV-First Proponent, then, it is natural to think that logic is about token RVs (studied at a level of abstraction), since these are what fundamentally have truth-values. So the fact that RV-kinds do not form a Boolean algebra, does not imply that the RV-first theorist denies, e.g., Conjunction Introduction. From the RV-First perspective, it is not obvious that rules like Conjunction Introduction are properly applied to RV-kinds, because these are not fundamentally the constituents of arguments—RV-tokens are.[54] And it is consistent with *D* to suppose that, necessarily, the set of RV-tokens is a Boolean algebra: necessarily, every token conjunction is true iff each of its conjuncts is true, and so on. This is the natural way to understand clauses for conjunction, disjunction, and negation in truth-conditional semantic theories if we understand them as issuing lawful universal generalizations *about token RVs* based on their structure. So the fundamental truth-bearers can still conform to Boolean operations on the RV-first view I am sketching.

Moreover, even though RV-kinds do not form a Boolean algebra with respect to their truth-values, there is a property in the vicinity with respect to which they *do* form a Boolean algebra: whether or not they are true-at the actual world. Although a conjunctive claim composed of two true conjuncts may not always be true, a conjunctive claim composed of two claims that are both true at *w* will always be true-at *w*.[55] Indeed, what is particularly strange about Never-*p* is that, although it is false, it is true-at the actual world. It stands in relation to the actual word that

[54] There are substantial questions about what a logic for token expressions would actually look like. For some suggestions see Klima (2004), Radulescu (2018).

[55] Supposing, of course, that truth-at a world is defined in the ordinary way.

(12) stands to *w* in the example above. Its descriptive conditions are satisfied by the actual world, but that does not—I am suggesting—make it true. With this in mind, we can see that the argument above goes wrong by ignoring the truth-at/truth-in distinction. The actual world satisfies the descriptive condition encoded by Never-*p* (i.e. it is true-at the actual world), but, given the truth-at/truth-in distinction, that does not imply that it is true. When we compute a truth-value for Never-*p* on the basis of Boolean conjunction, we are applying a rule that applies *not to its* truth-value, but rather to its status as truth-at or not true-at the actual world.[56] So the argument goes wrong in concluding that Never-*p* is true based on the fact that it is a conjunction of true claims.

Admittedly, these examples like Never-*p* do bring out a feature of *D* that is genuinely odd. Since, according to *D* an RV-kind counts as true in virtue of being such that, were it tokened, its tokens would be true, untokened RV-kinds count as true, not in virtue of describing the world as it is, but in virtue of having tokens that would describe the world as it is *if the world were otherwise*. Since being otherwise in this way (i.e. having tokens) can be part of what the RV-kind describes, we can potentially get RV-kinds that are false despite describing the world correctly or are true despite misdescribing it. This, I submit, is just a consequence we should accept if we take seriously that truth is fundamentally a property of token representations. RV-kinds are only means of categorizing tokens according to their semantic similarity, and it is only

[56] One might object: this property being true at the actual world is an artifact of possible worlds model theory which is itself of dubious metaphysical import; so it is no real consolation to be told that RV-kinds form a Boolean algebra with respect to this artificial property. In fact, although it is convenient for expository purposes to work with a possible worlds ontology, the property can be identified without talking about possible worlds at all. (Buridan draws the truth-at/truth-in distinction—perhaps invented it—without any possible worlds apparatus.) We could instead talk about these RV-kinds having fulfilled truth-conditions—keeping in mind that this fails to make them true.

in an extended sense that kinds themselves have truth-values or represent anything. Although we can recognize a sense in which Never-*p* expresses a condition that holds of our world, it is an abstraction from (possible) token representations that can only ever be false.

4. Conclusion

I began with a puzzle: truth seems to be a property of particular representations, but there is pressure within common sense to countenance truths that cannot be identified with any particular representation—truths that have not and never will be articulated. I have argued that it is compatible with the RV-first view that there are untokened RV-kinds and that these RV-kinds can be assessed for truth-value in a derivative sense, that stated by *D*: for an RV-kind to be true is for it to be such that, were it instantiated, its tokens would be true. I considered putative counterexamples to *D,* but found them to rest on ignoring the truth-at/truth-in distinction and, for that reason, argued that they fail. *D* provides us with a sound principle that allows us to maintain that token representational vehicles are the fundamental truth-bearers while also recognizing a domain of truths that do not owe their truth to being articulated by anyone.

CHAPTER 3

Open Texture and Capacious Pluralism about Content

1. Introduction

It is common to suppose that propositions, whatever else they are good for, ought to provide us with a privileged way of categorizing representations with respect to their semantic similarity: the sort of semantic similarity that constitutes *sameness in content*. For two representations to have the same content is for them to express one and the same proposition. Indeed, a number of recent authors have suggested that propositions should be thought of as *reifications* of the semantic properties of token representations that establish relations of sameness of content.[57] Like W.V.O. Quine and Hartry Field, I think this picture obscures something deep. We utilize *many* distinct notions of sameness of content—none of which is privileged as suitable for all purposes. We should be *Capacious Pluralists* about classifying contents.

Other philosophers have arrived at similar conclusions—that there are multiple legitimate notions of sameness of content, or that sameness of content only holds relative to some implicit parameter (Cumming 2013; Dickie and Rattan 2010; Bjerring and Schwarz 2017; Haze 2022; Grzankowski and Buchanan 2021). Often such views are motivated by an attempt to account for *hyperintensional* differences in content, as a means to explaining Frege-cases or Kripke's puzzle

[57] See Hanks 2011, 2015; Soames 2014, 2019; Grzankowski and Buchanan 2019.

about belief. I am sympathetic to these accounts, but I think that Capacious Pluralism can be given independent motivation. In particular, we are forced to utilize many distinct standards of sameness of content because we are often coordinating between representations that are semantically underdetermined. Semantically underdetermined expressions, I will argue, can become more precise without breaking content preserving links with past tokens; and because expressions can be precisified in contrary directions, we require multiple separate standards of sameness of content to relate tokens that are precisified in contrary ways.

I begin, in Section 2, by explicating the notion of content that I aim to investigate, which I characterize in terms of semantic properties that ground the application of intersubjective rational norms. In Section 3, I describe a specific kind of semantic underdetermination—what Friedrich Waismann called "open texture"—and give an argument that the correct application of intersubjective rational norms to tokens of open textured expressions requires using a large number of different standards of sameness of content. In Section 4, I close by discussing the deflationary upshot this argument has for some traditional debates about content.

2. Content Classification Schemes

What are propositional contents supposed to be or do? Philosophers have appealed to them for a variety of purposes[58], filling theoretical roles determined by a mix of logical, metaphysical, epistemological, and psychological considerations. I will be concerned with one

[58] They have been identified with: the primary bearers of truth and falsity; the com- positional semantic values of natural language sentences; the "entities" corresponding to type *t* in a higher-order logic; the fundamental bearers of modal properties like necessity and contingency; the objects of mental attitudes like belief; the objects of assertion and agreement; the objects of choice. There are, of course, connections between these domains that make it somewhat plausible to think that one sort of thing can fill all these roles, but this is a substantial piece of theory, not a definitional matter.

important strand in this mix: the role of contents as *communicable objects of thought*. It is a banal piece of folk psychology that, via communication, we come to have shared beliefs. I can believe something, say it to you, and you can, thereby, come to believe the same thing. This seems to be an ordinary way to describe communicative exchanges that are a ubiquitous part of our social life. I am going to accept this folk-psychological banality as true, and my discussion will be concerned with the notion of content that underwrites it.

Rather than focusing on contents *qua* (possibly higher-order) entities, I am going to examining the conditions under which two token representations can be said to have the *same content.* Appealing to an independent domain of content-objects to settle questions about which representations agree in content has been less helpful than one might expect—because, for over a century there have been heated debates about what these content-objects are and how to individuate them. Although I do not presuppose it in what follows, in my view, the core function of contents is to categorize token mental states and other concrete representational things into equivalent semantic types, and a proper metaphysics of contents will treat them as abstractions from the semantic *properties* of token representations (in this I follow Hanks 2011 (2015); Soames 2014 (2019); Grzankowski and Buchanan 2019; Burge 2010; Rescorla 2020). Even if one does subscribe to independent metaphysics of contents qua special objects, though, there is nothing objectionable about proceeding as if *content-sameness* is the comparatively primitive notion.

What kind of semantic similarity suffices for two token representations, e.g. beliefs, utterances, to have the same content? One suggestion might be: two representations have the same content if and only if they have all the same semantic properties. But on traditional understanding of the content role, it is not demanded *a priori* that, in order for two

representations to have the same contents, they need to be semantic duplicates of each other. Usually, theorists who utilize the framework of propositional content have not thought that propositional content is exhaustive of meaning/semantic significance more broadly. Frege, for instance, thought that pairs of sentences like the following, expressed the same proposition (*Gedanke*):

(1) John Lennon's sunglasses are cheap.

(2) John Lennon's sunglasses are inexpensive.

"Cheap" and "inexpensive" are arguably not *completely* synonymous, although they may apply to all the same objects (in all possible worlds). "Cheap" has a slightly negative affective resonance that "inexpensive" lacks. Frege would have said that, although they make the same contribution to expressed thought, they differ in "tone" (Frege 1956).[59] To take another example, many direct reference theorists have supposed that the following two token sentence express the same proposition.

(3) This author is hungry.

(4) Torsten Odland is hungry.

But it cannot really be said that there are *no* semantic differences between (3) and (4). The former contains a complex demonstrative with an NP constituent, "author," that applies to authors; this is a semantic property that the latter sentence token lacks. Although both sentences

[59] Frege's own examples were "Roß" and "Pferd" (roughly, "Steed" and "Horse"). The idea that tone is a feature associated with an expression as part of its meaning is widespread but not uncontroversial. See Ernie Lepore and Stone 2018 and Anderson and Ernie Lepore 2011 for arguments against this view. My disagreement with these authors is partially a verbal issue regarding which features one wants to call "semantic," but not entirely. It seems highly plausible that the difference in tone between, e.g., "dog" and "doggie" (Gazdar 1979), amounts to a robust public convention that English speakers understand as part of their knowledge of the language.

have argument DPs that contribute me as a referent, they do so in different ways, and a complete description of their semantic properties will reveal such a difference.

If content sameness does not require complete sameness of semantic features, what sorts of semantic features are relevant to making comparisons of content? Here, I will follow Frege—who articulated in a particularly vivid way the notion of propositional content that has shaped 20th and 21st Century philosophical debates—in taking two sorts of features as central.[60] First, contents track the properties of a representation that determine that representation's *truth-value.* Second, contents track those properties of a representation in virtue of which it can be assessed with respect to *rational norms*. In ordinary life, we manifest our sensitivity to these features in the way that we apply what I'll call *Content Targeting Principles*. So, for instance:

1. It is irrational to believe *p* and also believe ¬*p*.
2. In order to successfully communicate that *p*, one's audience must come to believe that *p*.
3. If two agents believe that *p*, they agree with each other.
4. Don't believe *p* if *p* is false.

Principles like these state, in schematic form, norms and generalizations that apply to agents in virtue of the content of their representational states. These principles arise from our common-sense picture of rationality, but they can be sharpened and altered by theoretical investigation in

[60] There are, of course, significant antecedents in Stoic discussions of lekta, "sayables," (Bobzien 2003; Bronowski 2019) and Bolzano's Wahrheiten an sich (Morscher 2018). Given the two features I take to be central to the idea of propositional content, it not surprising that the thinkers who contributed most, historically, to articulating this idea were deeply interested in characterizing the subject matter of logic.

logic, psychology, and epistemology.[61] Being extracted from common-sense, there are a variety of ways these principles might be regimented. For instance, in some of my examples above the schematic letters *p* and *q* are replaceable by singular terms, while in others they must be replaced by sentences. It is revealing, in my view, that our intuitive understanding of these principles does not decide between either of these sorts of regimentation. We apply CTPs by associating the schematic letters with semantic properties that we take to ground comparisons of content. So for instance, consider an instantiation of CTP (3):

> If two agents believe that Venus is visible in the evening sky, then they agree with each other.

My suggestion is that, when we apply CTP like this, the material substituted for the schematic letters (here, a token of "Venus is visible in the evening sky") instantiates a semantic property that provides a basis for content comparisons. For instance, here is one semantic property of the substituted token: it predicates the property *being visible in the evening sky* of Venus. Let's call this semantic property *Ref-Venus*. If we take Ref-Venus to be the property that determines the content schematically represented as *p*, then our application of the above principle provides us with a rule for assessing agreement: if two agents have belief states that predicate being visible in the evening sky of Venus, then they agree with each other. To make this more concrete, let's apply this principle to the belief states of, say, a 7th Century BCE Babylonian astronomer and an 11th Century CE Mayan astronomer. Individuals in both communities made detailed observations of Venus and its movements. Suppose there was a Babylonian astronomer and a Mayan astronomer who each had a belief state that (a) referred to Venus and (b) predicated of Venus the

[61] For instance, Bayesian principles governing the credence function and belief update are Content Targeting Principles developed in a more formal theoretical setting.

property *being visible in the night sky*. The instantiation above tells us that, that the Mayan and the Babylonian thereby agree.

This seems to be a reasonable standard for comparing the content of these two beliefs. It is not obvious, based on what I have said, that the Mayan and Babylonian astronomer have a common "way" of thinking about Venus—that there is a Fregean sense by which they each think of Venus; nor have we specified anything very specific about the internal structure of the token beliefs states in the Mayan's and Babylonian's head, respectively. But, none the less, a historian of astronomy who wanted to compare the scientific cultures of the Mayans and Babylonians would count them, in virtue of having token beliefs that instantiate Ref-Venus, as believing something in common—as agreeing. And this is not an idle matter: if we take them to have beliefs with the same content (rather than simply being semantically similar), we will apply further CTPs to them in a uniform manner. For instance, we can assess what inferences would be rational for them to make on the basis of *common* evidence. We can say: if they come to believe that the moon is the smallest object visible in the evening sky, then they would have a reason to infer that the moon is smaller than Venus. By taking the Ref-Venus property as establishing sameness of content, the historian is able to put the Babylonian and the Mayan into rational engagement[62] with each other, despite a large spaciotemporal gap and idiosyncrasies of culture and language[63]. On a finer-grained standard, this may not be possible.

[62] I pick up the notion of "rational engagement" from Dickie and Rattan 2010.

[63] To wit: "Unlike the Babylonians, the Mayans apparently made no attempt to correct for the varying lengths of the Venus synodic period but took the 584-day Venus Round as constant" (McCluskey 1983). Here I think it is natural to read McCluskey as saying that there is something that the Mayans believed and the Babylonians rejected.

I will call a property that can be used to correctly apply a CTP (e.g. Ref-Venus) a *content grounding* property:

> A semantic property *F* is *content grounding* if and only if it is possible that there are token representations *x* and *y* such that: *x* and *y* both instantiate *F*, and, in virtue of instantiating *F*, the correct application of CTPs to *x* and *y* will treat them indistinguishably.[64]

(I am going to assume in what follows that necessarily coextensive properties are identical. Nothing of substance hangs on this, but it will make my exposition simpler if we have a principle of individuation for properties ready to hand.)

Before proceeding, I should make a note about my methodology. When we apply CTPs, we are produce some sentence that amounts to an attitude report or an indirect speech report. That might suggest I am naively supposing that, if a natural language attitude report of the form ⌜ A believes that φ ⌝ is true, then the token expression φ must have the same content as the representation being reported on. But this assumption is not at all tenable. For instance (to use an example from Abreu Zavaleta 2019), if Anna assertively uttered, "Carla has been seeing a new lover from LA," we can truly report her assertion in the following ways:

(5) Anna said that someone has been seeing a new lover from LA.

(6) Anna said that Carla has been seeing a new lover from California.

No one would think that the embedded sentences in these reports have the same content as Anna's utterance—its easy to think of situations in which Anna's utterance and the embedded sentence have different truth-values— and there are a variety of other cases that illustrate the same point (see Blumberg and Lederman 2020, Sosa 1970, Cappelen and Ernest Lepore 2005

[64] I.e. as associated with the same schematic letter.

among others.) But my method does not depend on the assumption that, in all true speech/attitude reports, the embedded clause has the same content as the target reported on. In fact, my argument in this paper will not presuppose any particular analysis of attitude reports or the denotations of embedded CPs.[65] Rather, my method is to consider token representations and their semantic properties, and to ask if any of those properties provide a legitimate basis for applying CTPs to those states in a uniform manner. True attitude reports in natural language do provide defeasible evidence for how CTPs can legitimately be applied—because, at least in non-technical settings, we express our application of CTPs by producing natural language attitude reports. The evidence is defeasible because, sometimes (as in Abreu Zavaleta's cases) natural language attitude reports characterize the semantic properties in of a representation in ways that, on reflection, aren't appropriate for applying CTPs.[66] True reports are evidence nonetheless, because there are cases in which we really are interested in applying CTPs to ourselves and others, and in those cases it is important that we make attitude reports in which the content of the report matches the content of the target.

Let me summarize what I've said so far about the content role. On this picture, contents are supposed to be those aspects of representational tokens in virtue of which we can *jointly* carry on cognitive activities where we take ourselves to be subject to intersubjective rational norms. That there are such properties, and that such properties are routinely passed from one

[65] Although, it fits naturally with the view that embedded CPs denote properties of contentful individuals (see Kratzer 2006, Moltmann 2013, Moulton 2015, Elliott 2020, Bondarenko 2022).

[66] So for instance, it might be the case that Anna's utterance is false (Carla has not been seeing anyone), but the token embedded clause in (5) is true (someone or other is seeing a new lover from LA). In that case, there are rational norms (e.g. CTP 4) that demand different attitudes towards these tokens.

agent to another in communication, are assumptions fundamental to our ordinary understanding of joint inquiry and decision-making. In particular, it is part of the ordinary understanding of *science* that independent investigators transmit beliefs to each other and can be held to public standards of rationality. I will not try to enumerate a canonical set of CTPs—nor do I mean to imply that it is self-evident or incontestable what should be included on that list. For instance, some CTPs (like 1 above) are grounded in logical laws, and I think that there are substantive debates to be had about what generalizations are laws of logic. So, I do not mean to suggest that common-sense supplies us with a fixed list of CTPs that we could, say, convert into a Ramsey-sentence that would provide us with an implicit definition of what semantic properties count as contents. But the CTPs that I will make use of in this paper, like those above, should not be objectionable, and they are good candidates to appear on any canonical list of CTPs.

My characterization of content-grounding property is weak in a notable sense. It leaves open the possibility that the property that suffices for sameness of content with regard to one pair of representations may be insufficient with regard to different pair. There is some *prima facie* reason to think content comparisons are actually like this. For instance, consider an Ancient Greek who has two ways of thinking about Venus, as Hesperus and as Phosphorus, and who is ignorant of the fact that Hesperus is Phosphorus. Suppose that this Greek has a token belief, *a*, that predicates *being visible in the evening* of the object thought of as Phosphorus, and a token belief, *b* that negates the predication of *being visible in the evening* to the object thought of as Hesperus. Can we apply CTPs to *a* and *b* taking the property *predicates being visible of Venus* as content-grounding, as we did in the case of the case of the Mayan and the Babylonian? No—this will lead us to *misapply* CTPs to *a* and *b*. For if *a*'s content is determined by that property, then *b* will count as *a*'s negation, and our Greek will count as violating CTP 1 in our list: there is some

p such that he believes both *p* and ¬*p*. But this must be a misapplication of CTP 1, because our Greek is not irrational in virtue of having the beliefs *a* and *b*.

The standard Fregean reaction to this sort of case is to infer that the property we started with—which tracked merely tracked sameness of referential and predicational structure—was not content determining in the first place: *a* and *b* show that, whatever properties determine content, they must be finer-grained that these. This ultimately may be the correct reaction. But we should not define "content grounding" property to as to rule out the possibility that a property can be content grounding with respect to one network of tokens while failing to be content grounding with respect to another. The fact that *a* and *b* do not have the same content does nothing to undermine our initial judgement that the Mayan and Babylonian agree with each other, and our basis for thinking this seems to rest entirely on the fact that they both have beliefs that predicate *being visible in the evening sky of Venus*. As we said, if there is some further "mode of presentation" common to the Mayan's thoughts about Venus and the Babylonian's thoughts it isn't obvious what it is, and our judgment about their agreement seems well supported without it. I do not rule out the possibility that, as the Fregean suggests, any property that grounds sameness of content for some network of tokens must ground sameness of content with respect to any network of tokens. But this is a claim that must be argued for—it ought not be stipulated.

Individual content comparisons reflect more general principles for classifying representations according to their content. For instance, whatever reasons we have for taking the Babylonian and the Mayan astronomer to believe the same thing apply equally to two individuals who have beliefs predicating *being the US president in 2023* of Joe Biden. The notion of sameness of content at work in the Ref-Venus example is, roughly, same property predicated of

same subject matter. And the application of many CTPs requires that content grounding properties stand in formal-logical relations to other content grounding properties. If a property is to serve as the value of *p*, in an application of e.g. CTP 1, there must be some other property that counts as its negation, to serve as the value of ¬*p*[67]. Content grounding properties come in natural families. Some of those families are of particular interest, and I will call them *Content Classification Schemes* (CCs):

> A set of content grounding properties is a Content Classification Scheme if and only if, necessarily, every truth-apt representation instantiates a single property in that set.

A Content Classification Scheme is a complete partition of truth-apt token representations (actual and possible) into equivalence classes according to mutually exclusive content grounding properties. As such, every CC determines an equivalence relation that amounts to a *general* conception of sameness of content: two token representations have the same content with respect to CC_x if and only if they both instantiate some property in CC_x. A CC, defined in this way, will only ascribe contents to truth-apt representations—not the compositional parts of those representations (if they have them). This will be apt for views of content according to which contents do not themselves have compositional structure. But we can extend the notion of a CC to cover schemes that classify the compositional parts of representations. A *Compositional* CC would be a CC supplemented by a set of semantic properties of compositional parts of truth-apt representations, constrained by some theory that indicates the contribution those properties make to determining the content possessed by truth-

[67] It is plausible that every content grounding property is such that the Boolean operations ¬, ∧, and ∨ are defined on them. Plausible, but I will not argue for it. I think it is just barely conceivable that there might be creatures with simple cognitive capacities that have truth-evaluable states such that nothing would count as that state's negation

apt representations containing parts that possess those properties.[68] A Compositional CC is what we would need if we wanted to say that, e.g. a token of "rot" by a German speaker had the same content as a token of "red" by an English speaker.

Traditionally, philosophers have specified Content Classification Schemes by developing a theory of propositions—all traditional theories of propositions determine Content Classification Schemes[69]. To illustrate this, consider the view of propositions as sets of metaphysically possible worlds (see, e.g., Stalnaker 1976, Stalnaker 1984). According to this conception, the proposition expressed by a representation *x* is the set of possible worlds *x* is true of. Consider this set of properties: $G = \{F \mid$ there is a set of possible worlds W such that, for all x, necessarily, x instantiates F if and only if x expresses $W\}$. Each one of the properties in this set corresponds to a possible-worlds proposition, and, necessarily, is instantiated by just the objects that express that proposition. So *G* is the Content Classification that counts two representations as having the same content just in case they express the same possible-worlds proposition. Other familiar conceptions of "propositional content" can be mapped to CCs in an analogous way[70].

---

[68] The sort of compositional account of propositions as cognitive act types developed in Hanks 2011, 2015 and Soames 2014, 2019 can be seen as developing such a theory of how the semantic features of sub-truth-apt representations (e.g. predicating redness) contribute to determining the content grounding properties of truth-apt representations.

[69] In the succeeding discussion I am effectively assuming that every truth-apt representation expresses at most one proposition, and that its truth-value is determined by that proposition. For those who deny this, the picture will have to be adjusted somewhat. (The simplest approach would be to say that two representations have the same content if and only if they express exactly the same proposition or propositions.)

[70] Roughly, just take your favored account of propositions and consider the set of properties $\{F \mid$ there is a proposition p s.t., F is the property of expressing p and only $p\}$. This characterization is rough because, depending on your theory of propositions it may be provable that there are more propositions than can be collected in any set—and that therefore, there is no one-one mapping from the propositions to the members of any set. However, the characterization of CC's in terms of sets rather than, say, in plural terms, is purely a matter convenience in exposition. So by itself,

What is distinctive of the Content Classification Schemes associated with traditional theories of propositions is that they are meant to be, in some sense, privileged. The distinctions among the various propositions ought to be fine-grained enough to apply Content Targeting Principles correctly to all representations, and ought not to draw any distinctions that are irrelevant to the correct application of Content Targeting Principles. To make this idea precise, let me define the sense in which one CC can *fine-grain* or *coarse-grain* another:

> CC_x *fine-grains* CC_y if and only if, possibly, there exist two token representations *a* and *b* such that CC_y classifies *a* and *b* as having the same content, CC_x classifies *a* and *b* has having distinct contents, and the correct application of CTPs to *a* and *b* involves counting them as having distinct contents.
>
> CC_x *coarse-grains* CC_y if and only if, possibly, there exist two token representations *a* and *b* such that CC_y classifies *a* and *b* as having distinct contents, CC_x classifies *a* and *b* has having the same content, and the correct application of CTPs to *a* and *b* involves counting them as having the same content.

These relations are not converses, and neither relation is guaranteed to be asymmetric– there may be cases where one CC fine-grains another CC that in turn fine-grains the former. Now, in terms of these notions, I'll define what it is for a Content Classification Scheme to be *authoritative*:

> CC_x is authoritative if and only if there is no CC_y that fine-grains or coarse-grains it.

---

the fact the propositions may be more numerous than the cardinality of any set is no obstacle. There are, however, related objections to be raised regarding Russell-Myhill-like paradoxes of propositions. It is true that, if contents are to be associated one-one with semantic properties, any theory of content will have to restrict any comprehension principles that would generate a violation of (some analogue of) Cantor's Theorem. But there is no obvious reason to think this is objectionable as such—every account of propositions has to find some way of avoiding Russell-Myhill—and it raises issues orthogonal to my concerns here. For a discussion of the sorts of constraints Russell-Myhill places on theories of propositions see, e.g., Uzquiano 2015 and Walsh 2016.

Traditional theories of propositions are meant to determine a Content Classification Scheme that is authoritative in this sense. It is easy to show that, if there is an authoritative CC, it is the only CC.[71]

The main question I want to address is whether or not there is an authoritative CC. In the next section, I will argue that there is not: any legitimate way of classifying representations according to sameness of content is either fine-grained or coarse-grained by some other Content Classification Scheme. Indeed, I will argue for the view I am calling Capacious Pluralism about Content: there are many legitimate Content Classification Schemes and no obvious limits on how many there might be.

3. Open Texture

It is widely held that many natural language predicates (and corresponding concepts) are semantically indeterminate. A predicate *F* is semantically indeterminate if, possibly, there is an object *o* such that it not determined, given the meaning of *F* and the features of *o*, whether or not *F* applies to *o* (see McGee and McLaughlin 1995; Shapiro 2006, Shapiro 2010; Heck 2003; Weatherson 2010 among others). If *F* is a predicate and *o* and object, and it is indeterminate whether *F* applies to *o*, then *o* is a *borderline case* for *F*. I will assume that, if *o* is borderline case for a predicate *F*, then the meaning of a sentence of the form ⌜ *Fo* ⌝ does not determine whether tokens of that sentence are true nor false.[72] If we are confronted with a borderline case, it seems that *applying* the predicate or *rejecting* its application are both legitimate options, as far as the

[71] Suppose CC_x is an authoritative CC. Given that we are individuating properties in terms of necessary coextensiveness, then any CC that neither fine-grains nor coarse-grains CC_x is composed of properties with the same modal profile as CC_x, and, therefore, is identical with CC_x.

[72] This is not uncontroversial, but it is widely accepted by those who accept the existence of semantic indeterminacy in the first place.

predicate's meaning is concerned. This openness to contrary elaborations, I will argue, requires us to apply multiple standards of sameness of content to uses of semantically underdetermined predicates.

As I understand it, semantic indeterminacy includes the sort vagueness that gives rise to the sorites paradox, but my argument will be focused on a slightly different phenomenon: indeterminacy due to what Friedrich Waismann (Waismann 1947) called "open texture." He developed this idea in the context of arguing that many natural language predicates cannot be strictly defined.[73] Let's say that a predicate G is strictly definable if and only if there is some distinct complex predicate H such that, *as a matter of semantic rule*, G applies to an object if and only if H applies to that object. This fits with the conception of definitions as rules of abbreviation. In our post-positivistic age, many will be apt to believe that *few* or *no* natural language predicates are definable in this sense, but Waismann's discussion contains an insight that remains valuable.

Waismann points out that, for most linguistic expressions, there is no definition of that expression that we would accept as a strict rule determining its application conditions in all possible circumstances. Take "suitcase." Suppose we follow Merriam Webster in saying that "suitcase" is defined as "a portable case designed to hold a traveler's clothes and personal articles" ("Suitcase" 2024). Now, further suppose that an inventor designs a portable case for travelers that holds their clothes and other personal articles and transforms them into an edible slurry. Waismann suggests that, presented with such an object—which differs in a radical way from any suit- case we have previously encountered—we would be unsure whether "suitcase,"

[73] In Waismann's original discussion, he is indiscriminate in applying "open texture" to concepts as well as linguistic items. He likely saw this distinction as of little import.

applies to it, despite the fact that it satisfies the Merriam Webster definition. This demonstrates that we do not take it to be a semantic rule that "suitcase" applies to whatever satisfies "a portable case designed to hold a traveler's clothes and personal items." Waismann's argues that, for *any* proposed definition of "suitcase," it is possible to construct a case in which some object satisfies that definition but we are not compelled as a matter of semantic rule to apply "suitcase" to that object.[74] So "suitcase" is indefinable.

Waismann's insight is to suggest that some words, e.g. "suitcase," are indefinable because the features that bear on their application are *open-ended*. This is what he what he means by "open texture." (Here I follow Cumming 2023.[75]) At any given time, an open textured expression will be semantically associated with a collection of features that bear positively on its application and a collection that bears negatively on its application, but these collections are open to elaboration. Although the slurry-case may have all the features we've taken, in the past, to count decisively in favor of applying "suitcase," the meaning of "suitcase" leaves it open whether additional features may bear decisively against (or for) its application. If an expression

[74] One might have a different reaction to Waismann's cases: perhaps the Merriam Web- ster definition fails but some other definition will succeed; or perhaps our hesitancy regarding applying "suitcase" in the imagined cases is not due to the fact that there is no definition that determines the application of "suitcase," but, rather, to our *ignorance* of the semantic rules governing "suitcase." In principle, these responses are open, but neither has much plausibility on reflection. If anyone can be said to be an expert on these matters, surely it would be professional lexicographers. The fact that their proposed definitions of "suitcase" do not provide a strict rule governing the expression's application strongly suggest that no such definition can be given. C.f. Elbourne 2011 p.1-6.

[75] Sometimes "open texture" is simply identified with what I'm calling semantic indeterminacy: an expression is open textured if and only if there are possible objects such that, the meaning of the expression and the nature of those objects do not determine whether the expression applies to those objects (See, e.g., Shapiro 2006 and Gauker 2017.) It is worth keeping these ideas separate because not all semantically indeterminate expressions need be indefinable in the sense Waismann discusses.

is open textured, it will be semantically indeterminate, because its meaning will not determine fixed application conditions that settle its extension.

(One might ask: what about the feature *suitcasehood*? Doesn't "suitcase" as a matter of semantic rule apply to everything that has that feature? In the case of some indefinable predicates, this might be the right response. For instance, Putnam (1962; 1975) and Kripke 1980 famously argued that the extension of natural kinds terms is not determined by an object's satisfying a certain description, but rather by being causally linked to a certain real essence via an initial baptism. We can develop Waismann-style cases to show that a natural kind term, e.g. "gold," is indefinable.[76] Currently, our best definition of "gold" might be "substance with the atomic number 79." But if we were to encounter a substance with an atomic number of 79 that emitted a heretofore never observed form of radiation, we would not think that "gold" must apply to the substance in virtue of the definition of "gold"— in principle, we might adjust our theory in light of the new observation so as to no longer accept that all gold has the atomic number 79. But even if "gold" is not definable (in Waismann's sense), we may want to say that it isn't open textured either. There is a feature that strictly determines whether or not "gold" applies to an object: namely, whether it is *gold*, e.g. has the real essence of the substance we (actually) baptised as "gold."

[76] "Gold" is in fact one of Waismann's own examples, though he considers a more verificationist friendly candidate definition that is specified in terms of gold's characteristic spectrometer readings. It is a bit hermeneutically awkward to suggest, as I do here, that one of Waismann's main examples of open texture is not really open textured. But the suggestion depends on taking seriously the "natural kind" metasemantics, proposed by Putnam and Kripke, that Waismann never entertained and may well have rejected.

This may be the correct account of natural kind terms[77], but it would be misleading to give an analogous account for "suitcase." There does not seem to be a underlying essence of *suitcasehood* that our word could latch onto. To the extent there is a property *suitcasehood*, plausibly, it is made what it is by virtue of its role in a cluster of human activities that includes the classificatory practice of calling things "suitcases." So *suitcasehood* should not be understood as a fully independent feature of objects that *governs* how "suitcase" is to be applied. Although it is true that "suitcase" applies to all and only things that are suitcases, this is compatible with the claim that "suitcase" is open textured—that the features that bear positively or negatively on its application are open-ended.)

I will base my argument for Capacious Pluralism on a not-so-outlandish imagined history[78] of the evolution of an open textured expression:

> Suppose at time t_1 there is a community of English speakers living on an island in the North Sea who have the word "sandwich" in their vocabulary. As of t_1, the hotdog has never been introduced to their community. Eventually it is, and the islanders encounter the hotdog. Initially, the question "Does 'sandwich' apply to hotdogs?" has no clear answer for the islanders. At that time, hotdogs are borderline cases for "Sandwich." Let *a* be an utterance of "sandwich" produced at t_1. The island has a northern and a southern half, and over time the two regions develop different linguistic practices. In the north, a regional court rules early on that, for the purposes of taxing and regulating restaurants, hotdogs

[77] For some criticisms, see Dupré 1981 and Wilson 2006.

[78] I owe a debt to Mark Wilson 1982, 2006 for my thoughts about this sort of case. The view I will argue for is also similar to the account of content Hartry Field offers in his 2017 paper. For reasons of space, I have left out a detailed comparison of my view with Field's. The most important point of divergence is that Field takes cases like the one I am about to discuss to show that the common-sense concept of synonymy is not transitive. I think this ignores something important about the formal structure of content comparisons. Sameness of content needs to be an equivalence relation in order to play its role as determining a uniform standard for applying rational norms. So where Field sees failures of transitivity, I see multiple ways of tracking contents.

> count as sandwiches. In the south, there is no relevant regulatory decision, and the population tends to associate hotdogs with other sausage-related paradigmatic non- sandwiches——e.g. sausage rolls, toad in the hole. By t_2, there is a well-established pattern of usage that distinguishes Northerners from Southerners: Northerners apply the word "sandwich" to hotdogs, whereas Southerners deny its application. Let *b* and *c* be utterances of "sandwich" produced at t_2 in the North and South respectively.

I will argue for three claims regarding this story. First, there is a CC according to which *b* and *c* have distinct contents—call it CC_1. Second, there is a CC according to which *a* and *b* have the same content—CC_2. Third, there a CC according to which *a* and *c* have the same content—CC_3. Since the same-content relation for every CC is transitive, this implies that either $CC_1 \neq CC_2$ or $CC_1 \neq CC_3$. Either way, the situation involves at least two CCs.

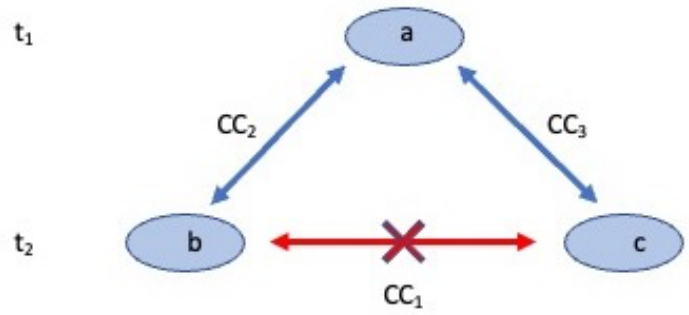

Figure 4.1 Diagram of the CCs in the island story

Hopefully it is uncontroversial that some CC distinguishes *b* from *c*. Given the patterns of usage described, a Northern utterance at t_2 of "A hotdog is a sandwich" will be true and a Southern utterance of the same sentence at the same time will be false. This will make a difference to the application of many content targeting principles. Take the principle: don't assert *p* if *p* is false. There needs to be a CC according to which a Southerner who asserts "A hotdog is a sandwich" violates this norm, whereas a Northerner does not. And in order for that to be the

case, there has to be some CC that assigns utterances of "sandwich" produced in the North a different content than utterances of "sandwich" produced in the South.

It is perhaps more surprising that there are CCs that group tokens *a* and *b* and *a* and *c* together, respectively. The core intuition here is captured in Henryk Mehlberg's claim that semantically indeterminate expressions can be "understood in many ways without being misunderstood" (Mehlberg 1958). At t_1, it is admissible both to apply "sandwich" to hotdogs and to reject its application, so both populations are extending the use of "sandwich" in a manner is compatible with its meaning at t_1.

As one piece of evidence, consider the pattern of attitude reports within this community. Members of either of the descendant populations will be able to make true (de dicto) speech and attitude reports about members of the ancestral populations using the word "sandwich." For instance, if a speaker the uttered "I want a sandwich" at t_1, speakers in the North or South at t_2 can report this truly by uttering "She says she wanted a sandwich" (see Cappelen 2018 p. 109-110, Francén 2022, Dorr and Hawthorne 2014 for further evidence in favor of this point).[79] This is not conclusive. It could be that, although such reports are true, the ancestral token does not have the same content as either of the token embedded clauses in the future attitude reports. But it does suggest that sameness of content is preserved within each branch of the developing usage of the word "sandwich."

[79] Is it plausible to suggest that these reports are felicitous although false—that they're just "loose talk"? I don't think so. As I will argue a bit later, I think that we routinely make homophonic belief and speech reports using semantically indeterminate vocabulary and it is doubtful that when we make these reports, our words are indeterminate in exactly the same way as the speakers we are reporting on. If, in order for such homophonic reports to be true, all aspects of semantic indeterminacy in the target had to be preserved in the embedded material, we would have to reject a huge number of intuitively true reports as false. Thanks to David Boylan for pressing me on this.

From a more theoretical perspective, the correct application of CTPs to our islanders will require taking recognizing chains of sameness of content that link past tokens to descendent tokens in either community.[80] Inhabitants of the island can engage in joint inquiries and projects spanning the time between t_1 and t_2, and these projects can be assessed according to common rational norms. For instance, suppose Omar, a dietitian, writes the following in an article at t_1: "All sandwiches contain gluten." Liv, a Northerner, then reads this article at t_2, takes its claims to be true, and comes to hold a belief she would express by saying "All sandwiches contain gluten." Since Liv, who is gluten-intolerant, already holds a belief that she would express by saying "If all sandwiches contain gluten, I will get sick eating sandwiches," she draws a conclusion she expresses by saying "I will get sick eating sandwiches." I think that we can correctly describe this as a case in which Omar asserted *p*, Liv took *p* on authority from Omar, and, applying modus ponens, inferred *q* on the basis of her prior belief in $p \rightarrow q$. If this is correct, then there must be a CC that counts Omar's "sandwich" and each of Liv's "sandwich"s as having the same content. This story can be told with the same plausibility about someone living in the *South* at t_2, so there must also be a CC that groups ancestral tokens of "sandwich" with Southern descendant tokens.

What basis would we have for denying this account? One might point out that that, insofar as Liv uses "sandwich" in a way that determinately applies to hotdogs, her "sandwich" is semantically different from Omar's. If we suppose that, for any legitimate conception of "content," *x* and *y* can have the same content only if they determinately apply to all and only the

[80] There has been a great deal of discussion of similar points in the literature on conceptual engineering, where many have been concerned to articulate a sense in which one and the same concept (or subject matter) can persist through changes in extension. C.f. Sawyer 2018, Cappelen 2018, Prinzing 2018, Richard 2019 among others.

same objects, then we should deny that there is any CC that groups together Liv's "sandwich" and Omar's.

However, in order for representations *x* and *y* to have the same content, we do not ordinarily insist that they have to be semantic *duplicates*. As I pointed to out before, it is traditional to think that some semantic differences, e.g. differences in tone, are compatible with sameness in content. The notion of content I am interested in picks out whatever semantic properties are involved in applying Content Targeting Principles. So, the important question is whether we can correctly apply CTPs in a way that links tokens of "sandwich" produced by Omar and Liv. I suggested above that there must be, because we can describe the process above as an instance of someone coming to believe that *p* on the basis of testimony that *p*.

It is true that the belief Liv expresses with "All sandwiches contain gluten" is more committal than Omar's token of that sentence. Liv takes it to imply that all hotdogs contain gluten, while Omar's sentence does not commit him to this. But the meaning of Omar's sentence doesn't *rule out* Liv's interpretation either—at t_1 it is legitimate to apply "sandwich" to hot-dogs. Liv understands Omar's sentence in one of the ways it is open to being understood. I take this as grounds to think that there is a notion of content that links Liv's belief and Omar's sentence.

Moreover, in this case, the communicative goals of the exchange seem to be fulfilled. Omar intended to get his audience to come to share the belief he expressed by uttering "All sandwiches contain gluten." By his lights, he appears to succeed: he can truly report that, on the basis of his testimony, Liv "believes that all sandwiches contain gluten." To come at the same point another way, consider two possible reception histories for Omar's article: in scenario A, Liv is as we have described her; in scenario B, Liv lives in a community in which "sandwich" is

used in a way that preserves all the indeterminacy of Omar's usage, and on the basis of Omar's article comes to have a belief she expresses by saying "All sandwiches contain gluten." Is Omar's original intention—the intention he would express by saying "I want my audience to believe that all sandwiches contain gluten"—any better satisfied by scenario B than by scenario A? I think not. Omar's conception of the mental state that counts as a "belief that all sandwiches contain gluten" is just as indeterminate as his conception of "sandwich," and so fails to distinguish B rather than A as satisfying it. In sum: Omar apparently succeeds in communicating his belief to Liv[81]; if that is so, we should count the semantic properties that are preserved across this communicative exchange as content grounding.

If it seems odd to characterize "sandwich" in Liv and Omar's mouths as having the same content, this is likely because we are characterizing two tokens as sharing content despite making different contributions to truth- value. Liv's "sandwich" determinately applies to hotdogs, whereas Omar's does not, and this difference can be relevant to the determining the truth- values of token sentences containing "sandwich" (e.g. "This is a sandwich," where a hotdog is demonstrated).[82] Since content is meant to track the semantic features that are determinative of

[81] And for the same reason, Liv succeeds in her goal of coming to believe what Omar intended to communicate.

[82] There are complicated issues raised here that I do not have space to discuss. For instance, if we want to assess the validity of an intersubjective chain of reasoning carried on by Omar and Liv, should we understand being truth-valueless as a designated or undesignated value? The most basic question, to put it vaguely, is: how do representations without truth-values relate logically to representations with truth-values? On one extreme, we might think that representations without truth-values are not fit to stand in logical relations with anything—logic should simply ignore them. (See, e.g., Kripke's understanding of the intermediate truth-value in the Kleene truth-table underlying his theory of truth Kripke 1975, or Ludlow's claim that a utterance of a semantically underdetermined sentence will be inadmissible unless its meaning is sharpened enough, in context, to settle its truth-value Ludlow 2014.) On the other, we might adopt the perspective that validity holds in virtue of purely structural features of representations that hold

truth-value, this seems to amount to a content distinction. But the matter is subtle. A token of "sandwich" in Omar's mouth does not have a *different* extension than a token in Liv's. Due to open texture, *neither* token has a meaning determines an extension.

The question is: what is the appropriate way to assess two expression tokens that lack extensions for sameness of content? The principle suggested by the objection under consideration is that the two tokens need to be indeterminate about exactly the same objects. This standard is intelligible, but it does not seem to track the semantic properties we are usually interested in when we apply CTPs. Open textured terms are ubiquitous in natural language. In the ordinary case, when we use these terms in conversation, we do not have strong reasons to believe that our conversational partners use them in a manner that preserves all the indeterminacy of our own usage. If sameness of content always requires preserving *all* indeterminacy, we would have to accept that much of our ordinary communication may not involve the transfer of content. I take this to be a *reductio* of the proposal. In most of our ordinary assessments of communication and joint reasoning, we use a standard of sameness of content that is less exacting than the one envisioned.

So, let us grant my three claims: there is a CC, CC_1, according to which *b* and *c* have distinct contents; there is a CC, CC_2, according to which *a* and *b* have the same content; and there is a CC, CC_3, according to which *a* and *c* have the same content. The history of this island involves a sort of content fission. The usage of "sandwich" evolves in such a way that *a* can stand in same-content relations with two tokens that have distinct contents. Given that any same-content relation is transitive, to apply content targeting principles to this populations we need (at

---

independently of whether those representations have truth-values (Camp 2002 takes up this perspective, for example).

least) two notions of sameness of content. If we accept this story as a case of content fission, it is a very short step to accepting Capacious Pluralism about content. For the fission-case we have described is just one of countless ways that "sandwich," as used at t_1 by the islanders, can be precisified.

Intuitively, there is no obvious limit on the sort of possible objects that would strike islanders as borderline cases for applying "sandwich." There are all sorts of possible sandwich-like objects they have never considered that raise problems similar to that of the hotdog, and many of these can be settled independently of one another. For instance, we might continue the story of our island further by allowing the forces of globalization to bring the McDonald's McGriddle—a hand-held breakfast consisting of egg, meat, and cheese between two pancakes—to their shores. Since settling the question about hotdogs does not provide any clear guidance about whether an object with pancakes as casing satisfies "sandwich", the islanders can further precisify their usage in opposite ways.

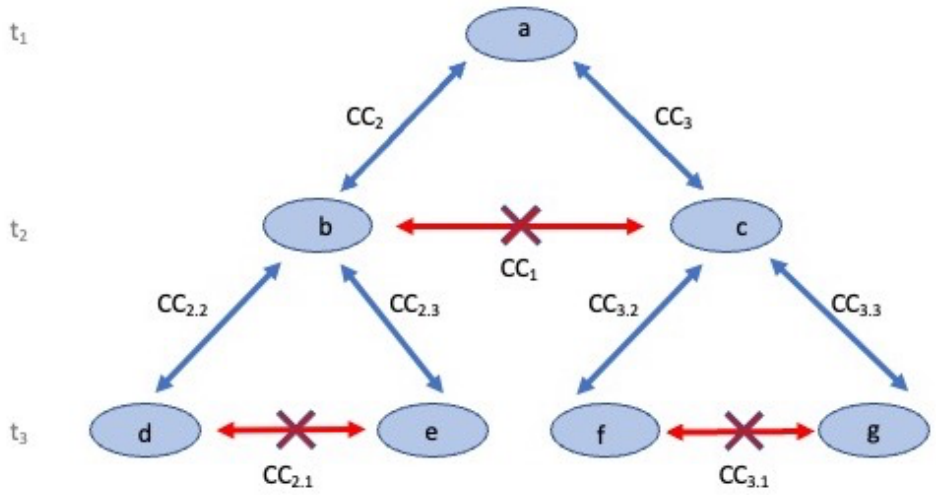

Figure 4.2 An imagined continuation of the island story with another level of fission.

Since the processes involved in these further fission case are exactly analogous, we should expect the same patterns regarding the application of CTPs. This makes available new patterns of semantic similarity and difference that call for additional CCs, beyond those we had recognized previously.[83] Since each token of "sandwich" in this tree is open textured, this process can be continued on without any obvious limit. "Sandwich" at each successive stage remains open regarding what features count for or against its application, so it remains open to elaboration in contrary ways.[84] If this is correct, the number of CCs that might be required to distinguish the different relations of sameness of content is extremely vast.

---

[83] Formally, the pattern illustrated in Figure 4.2 is compatible with there being only two CCs: for example, one CC according to which every node has the same content and another CC that distinguishes every node. But these CCs do not draw all the relevant distinctions—for instance, there is an intelligible sense in which *b*, *d*, and *e* share a content that should be distinguished from that possessed by *c*, *f*, and *g*. We can prove that there are more than two CCs involved in the scenario by adding three plausible premises. First, for every node, there is a CC according to which it has the same content as *a*. Second, the CCs involved are *symmetrical* in the sense that: for every branching node *x* and its two daughters *y* and *z*, there is a CC according to which *x* has the same content as *y* if and only if there is a CC according to which *x* has the same content as *z*, and there is a CC according to which *x* has a distinct content from *y* if and only if there is a CC according to which *x* has a distinct content from *z*. In the fission cases we are describing, the relationship between an indeterminate parent node and its comparatively determinate daughters are parallel, so any departure from symmetry would be arbitrary. Third, there are CC's that distinguish clusters of descendent nodes: for instance, there is a CC according to which $b = d \neq c = g$. I leave the proof to the reader.

[84] Open texture makes this open-endedness especially vivid, but an argument with the same structure can be given for a sorites-susceptible vague predicate like "tall." Plausibly, the only feature that is relevant to whether or not an object satisfies "tall" (relative to a comparison class) is its height; so the factors that bear on its application are not open- ended, as in the case of "sandwich." "Tall" is semantically indeterminate because its meaning does not specify a threshold on the height scale that distinguishes its extension from its anti-extension. But the argument in favor of tracking sameness of content across contrary precisifications of "sandwich" holds just as well for tracking sameness of content across precisifications of "tall." And, for any decision about a borderline case for "tall," there are infinitely many potential borderline cases that can be settled independently. So as with "sandwich," the process of distinguishing finer and finer precisifications of "tall" can go on indefinitely. This argument will not work, however, for a vague predicate if the scale with respect to which that predicate is vague is not *dense*. For instance, "has a lot of kids for a 21st Century Canadian" is vague, but, given that it does not apply

This, then, establishes that there is no authoritative CC. It follows that it is misleading to talk of two representations "having the same content" *simpliciter*. Assessments of content sameness are always made relative to some CC that determines what sorts of similarity count as content sameness, because it may be that two beliefs or statements have the same content with regard to one CC but differ with respect to another.

I will end this section by defusing a worry that might be raised by the last point. The suggestion that there are many standards of sameness of content seems to sit uneasily with the idea that, when we apply rational norms and other CTPS, we do so absolutely. If two representations have the same content according to CC_A but differ according to CC_B, then we do not have a single way to apply content targeting norms to those representations. It might be that, according to CC_A the representations realize a valid inference pattern, while according to CC_B it involves equivocation. When a situation like this obtains, what should we say about the rationality of the resulting argument? Are we left simply saying that it is rationally permissible relative to CC_A but rationally impermissible relative to CC_B?

No: although there are many different CCs that can be legitimately used to apply CTPs to some tokens or other, if we focus on a particular network of tokens, it is often an objective matter what standards are appropriate for applying CTPs within that network. This is what we saw with Liv and Omar. In that case, we said that, since Omar successfully communicates his belief to Liv, it is only appropriate to apply CTPs to that network that count Omar's belief and Liv's belief as sharing content. To be sure, there are finer-grained CC's that would distinguish between

---

to Canadians with two children and does apply to Canadians with five children, there is a small number of ways it can be precisified.

Omar's belief and Liv's belief.[85] But, applied to the network containing just Omar and Liv, such a CC would be illegitimate—it would count a successful belief-transfer as a case of equivocation. The mere fact that there are multiple standards of sameness of content is no threat to objectivity in the applications of rational norms to particular networks. There are different norms of polite greeting around the world, but, usually, when we step into someone's home, a particular norm is operative and the others are irrelevant.

However, allowing for a *certain* amount of non-objectivity in the application of CTPs might be seen as a feature of my account and not a bug. I do not have a fully general argument that would show that, for any given collection of representations, if there are two CCs that differ in the way they apply CTPs to that collection, at most one of those CCs applies CTPs to those representations correctly. I am not sure such an argument can be given. There are cases where it is not obvious how CTPs should be applied, in which more than one conflicting CC appears legitimate; the literature on Frege-cases provides ample illustrations of this. Does Kripke's (Kripke 1979) puzzled Pierre agree with an Englishman who utters "London is pretty"? Does everyone who assents to "Woodchucks are mammals" thereby agree with everyone who assents

[85] If our conclusion above, regarding there being at least three CCs in our initial island scenario, is correct, this implies that there must be a CC according to which *either a* has a distinct content from *b*, or *a* has a distinct content from *c*. Since *b* and *c* are parallel in their differences from *a*, we should conclude that each differs from *a* according to some CC (possibly the same or possibly different). So there must be a CC that distinguishes "sandwich" as used in Omar's community from "sandwich" as used in Liv's. More concretely, it is easy to imagine networks of tokens where the property *applying determinately to all hotdogs*, possessed by *b* but not *a*, might be relevant to content comparisons. For instance, consider a network consisting of three tokens of "This hotdog is a sandwich" uttered in each of the three communities—plausibly, it would be appropriate to count none of these tokens as expressing agreement.

to “Groundhogs are mammals” (Mates 1954)? There may be principled univocal answers to these questions, but it is not completely obvious that there are.

4. Debates about Propositions

Capacious Pluralism is interesting largely because of the import it has for debates about propositions. The most direct upshot of Capacious Pluralism is that it rules out any view of propositions that determines an *authoritative* Content Classification Scheme.

A CC would be authoritative if it could be used in every circumstance to correctly apply content targeting principles to every truth-apt representation. Naively, one might have thought that what it is for propositions to be “the objects of thought and talk” is for them to determine an authoritative CC: two representations have the same content if and only if they express the same proposition(s). But our story of the islanders shows there cannot be any such thing: there is no content classification scheme that is identical with CC_1, CC_2, and CC_3, and we need each of them to apply content targeting principles correctly. This also implies that one tool by which theories of propositions are sometimes compared to each other is faulty. Consider a standard sort of argument that Fregeans make against Russellian propositions. Here is a content targeting principle: it is irrational for someone to believe both p and $\neg p$. Suppose there is an agent with beliefs they would express by means of the sentences “Mark Twain was clever” and “Samuel Clemens was not clever.” On the standard Russellian account of propositions, there is a p such that this agent believes both p and $\neg p$, so if we apply the CTP above, we’ll judge the agent to be irrational. But these beliefs *don’t* make the agent irrational, just uninformed. Therefore, the objects of thought have to be finer-grained than Russellian propositions. The problem with this sort of argument is that it holds Russellian propositions to a standard that nothing can meet. If

there is no authoritative CC, then *every* CC will apply content targeting principles incorrectly in some cases. So we can't rule out a view as illegitimate on just those grounds.

In general, since we know there are multiple distinct standards of sameness of content, we don't necessarily need to interpret distinct accounts of propositions as rivals. It could be that the Russellian is describing a semantic property that counts as sameness of content with respect to some networks of representations, while the Fregean is describing a property that counts as sameness of content with respect to others. The interesting questions, then, won't be about whether the Russellian propositions are *really* contents, but whether there are situations regarding which to apply CTPs correctly, you need to use a Russellian rather than Fregean standard of sameness.

There are a variety of other philosophical projects related to propositions with regard to which Capacious Pluralism encourages caution. If there are multiple systems of objects that are equally well "propositions," then certain traditional questions—"What is the nature of propositions?"; "Are propositions structured or non-structured?"—are not obviously well posed. Perhaps there is a univocal answer to these questions, but perhaps not. The sorts of semantic similarity that count as sameness of content might be heterogeneous in a way that makes such questions inappropriate.

Of course, the import of Capacious Pluralism for wider debate about contents depends on whether the content role that I have investigated—that which grounds the intersubjective application of rational principles—matches the role of content as it appears in those debates. And this is a live issue, because philosophers have appealed to "contents" or "propositions" for a

variety of purposes.[86] In Section 2, I argued that my notion of content picks up the core of the traditional conception articulated by Frege. But it is worth pausing to consider an alternative conception: contents that classify representations according to the semantic features that figure in predictions, generalizations, and explanations in scientific psychology. This is the conception of contents as *psychological natural kinds*. One might object that, for all I have argued, it may be that, in the psychological natural-kind-sense there is a single authoritative way of classifying representations according to same- ness of content. And there is some plausibility to the idea that *this* is the notion of content that deserves serious philosophical and scientific attention.

This objection is legitimate, as far it goes. I have not provided any evidence for thinking that, if a semantic property is suitable for grounding the application of intersubjective rational principles, then it is the sort of property that will be privileged in detailed psychological prediction and explanation, or vice versa. Although these systems of categorization surely do interact at some level of abstraction, it easy to imagine how they might come apart when we are giving a detailed account of psychological mechanisms removed from everyday observation.[87] But even if *no* property that counts as a CTP-content constitutes a psychological natural kind, I don't think it is fair to say that natural-kind-content is the only sort of content worthy of

[86] If, say, one is mainly interested in characterizing the "entities" corresponding to the syntactic category of sentences in a higher-order logic (as in Goodman 2017) my arguments put essentially no constraints on one's project.

[87] I say "surely" for a few reasons. First, one might think that the folk-psychological theory framed using CTP-contents is, within a limited domain, predictively successful enough that it likely captures some distinctions that cognitive science will recognize as real, if superficial. Secondly, once we know what sorts of semantic properties are useful for in giving psychological explanations, it would be natural to *adopt them* for the purposes of applying CTPs, at least in some cases. So it seems quite likely that the properties constituting CTP-contents and natural-kind-contents overlap each other.

philosophical attention. Not everything we care about is a natural kind. Insofar as the psychological natural kinds pick out semantic properties that aren't CTP-contents, those kinds will not track the semantic properties we regularly transmit to each other in communication and by which we hold each other to intersubjective rational standards. So, to the extent that natural-kind-contents diverge from CTP-contents, they leave out what we might call the *public* content of joint inquiry and deliberation. Although CTP-contents may not capture all the properties and distinctions that are relevant to psychological explanation, the role they play in *epistemology* for social creatures like ourselves is not something we can easily replace or ignore.

5. Conclusion

Philosophical debates about propositions and their nature typically proceed from the assumption that propositions determine a preeminent content classification. I suggested we should understand this as a content classification that will draw all and only the distinctions relevant for assessing representations with respect to Content Targeting Principles. But I have argued that there isn't any such thing. Since some representations are semantically open textured, and their semantic properties can be extended in a variety of ways without severing content preserving links, the evolution of an open textured representation within a community can involve content fissions that might, in principle, go on indefinitely. To properly apply Content Targeting Principles in such a setting requires us to multiple standards of sameness of content, and there is no obvious limit on the number of distinct standards that may be required.

APPENDIX ONE

Gaifman has developed two versions of his "Pointer Semantics"—the version with operations on pointers in 1988 and 1992, and the version in 2000 without. In this exposition, I blend some elements from each. I follow Gaifman's earlier version in including compositional operations on tokens because it allows for a more intuitive description of the "direct call" relationship. The two systems do make some subtly different predictions, but the core Particularist idea is preserved in both versions, and the response I offer to the revenge argument can be adapted to either. Throughout this exposition I will assume that $\mathcal{L}$ is a first-order language (without function symbols) which is interpreted except for the semantic predicates "Tr( )" and "F( )." As a simplifying assumption, I will treat quantification substitutionally, so I will assume that every object in the universe of the background interpretation of $\mathcal{L}$ has a name. (I will treat these names as elements of the metalanguage as well as $\mathcal{L}$, so I will say, e.g. that "*a*" refers to *a*.) I use the following as metavariables ranging over tokens: $p, r, p_1, p_2 \ldots$ And use the following as metavariables ranging over valuations: *v, v*'…

Token Networks

A *token-system* for a language $\mathcal{L}$ consists of:

1. A set *P* of tokens.
2. A mapping ↓ from *P* onto the set of wffs of $\mathcal{L}$, such that every $p \in P$ is associated with at wff $p\downarrow$. We interpret ↓ as expressing the instantiation relation: $p\downarrow = y$ iff *p* is a token of *y*.
3. Two sorts of operations on tokens:
    a. Two functions, ( )1 and ( )2, associating every $p \in P$ with tokens *p1* and *p2* such that: if $p\downarrow$ = A * B, where * is a binary connective, then $p1\downarrow$= A and $p2\downarrow$= B; if $p\downarrow$

= ¬A, then $p1\downarrow$= A and $p1$= $p2$; in all other cases $p = p1 = p2$. We interpret these functions as mapping tokens onto their subtokens.

b. A function (|) taking a token p and a term t of ℒ as arguments such that: if Q is a quantifier, and $p\downarrow$ = QxA(x), then $(p|t)\downarrow$ = A(t); if $p\downarrow$ is not a quantified formula then $(p|t) = p$. We interpret this function as mapping tokens of quantified formulae and terms onto tokens of substitution instances.

Although I am treating quantification substitutionally here, the system can be modified to include objectual quantification by defining a satisfaction relation between tokens and variable assignments.

A token p *directly calls* a token r if and only if one of the following conditions holds:

1. $p\downarrow$ = ¬A or A * B, and $r = (p)2$ or $(p)2$
2. $p\downarrow$ = a quantified formula QxA(x) and r = $(p \mid t)$ for some term t
3. $p\downarrow$ = Tr(r) or Fa(r)

A *calling path* from p to r is a sequence of tokens $p_1 \ldots p_n$, with $n > 1$, $p_1 = p$, $p_n = r$, such that every p_i calls p_{i+1} directly. A token p *calls* a token r if and only if there is a calling path from p to r.

Building a Total Evaluation

A valuation v for a system of tokens is a (possibly partial) function from members of that to the values TRUE, FALSE, or GAP. TRUE and FALSE we will call *standard values*; GAP is a non-standard value, signifying semantic failure. A token p is evaluated by v if and only if $v(p)$ is defined. A valuation v' extends a valuation v if and only if, for all p that are evaluated by v, $v(p) = v'(p)$. Any valuation v determines a two-valued function $\mathfrak{v}$ from sentence types of ℒ to standard

values which we will call the "induced valuation of *v*." An induced valuation $\mathfrak{v}$ is recursively defined, relative to a given valuation *v,* as follows:

1. If α is an atomic sentence not containing "Tr" or "Fa," then $\mathfrak{v}(\alpha)$ = the valuation of α in the background interpretation.
2. If $v(p)$ = TRUE, then $\mathfrak{v}$(Tr(p)) = TRUE and $\mathfrak{v}$(Fa(p)) = FALSE
3. If $v(p)$ = FALSE, then $\mathfrak{v}$(Tr(p)) = FALSE and $\mathfrak{v}$(Fa(p)) = TRUE
4. If $v(p)$ = GAP, then $\mathfrak{v}$(Tr(p)) = $\mathfrak{v}$(Fa(p)) = FALSE
5. For non-atomic sentences $\mathfrak{v}$ is determined in accordance with the Weak-Kleene truth-tables, (where the third value is "undefined" rather than "GAP") with $\forall$ and $\exists$ treated as (possibly infinite) conjunction and disjunction.

Gaifman shows how to construct a *total* evaluation for a system of tokens of $\mathcal{L}$ that is capable of modeling the Particularist response to the semantic paradox and preserves a variety of truth-theoretic desiderata. The construction proceeds in steps by applying rules to an initial valuation *v,* which yields a new valuation *v'* in which further tokens are evaluated; and so on, until we reach a fixed-point in which all the tokens are evaluated. I will describe the construction assuming that our initial valuation is the *empty valuation* $\emptyset$—the valuation that is defined on no tokens. (Strictly, we can arrive at a fixed-point with the desired characteristics using other initial valuations, so long as these valuations are *self-supporting* in a sense that Gaifman defines. One could have a debate about which initial valuation leads to the construction of a fixed-point that better models natural language, but the choices between them turn on issues orthogonal to my purposes here.)

There are three rules that we use to construct new valuations: the Standard Value Rule, the Closed Loop Rule, and the Groundless Tokens Rule. The first rule, as it sounds, assigns standard values to tokens, while the others deal with failures and assign GAP.

Standard Value Rule: if $p\downarrow = \alpha$, $\mathfrak{v}(\alpha)$ is defined, and $v(p) \neq$ GAP, then assign to p the value $\mathfrak{v}(\alpha)$.

The antecedent of this conditional we call the *enabling condition* for the rule. It is necessary and sufficient for applying this rule to a token p that it meets these conditions. If it does, we say that the rule is *enabled on p*.

If we start with a valuation v on which p is unevaluated, applying the Standard Value rule builds a new valuation, v', extending v, on which p gets the value of $\mathfrak{v}(p\downarrow)$. So for instance, if we start with $\emptyset$, the Standard Value Rule will be enabled on any p that instantiates a sentence not containing "Tr" or "Fa," and applying it will result in assigning p the value $p\downarrow$ receives in the background interpretation. If p already gets a standard value, applying this rule will result in an unchanged valuation. Since the rule is only enabled on a token if it has not been evaluated as GAP, the rule cannot be used to revise a GAP.

To state the Closed Loop Rule, we first need to define what counts as closed loop of tokens. A set of tokens G is *closed* on v if and only if every member of G is unevaluated by v, and for every $p \in G$, for every token r unevaluated by v, if there is a calling path from p to r consisting only tokens unevaluated by v, then $r \in G$. If G is closed on v and, in addition, every member of G calls some member of G, then G is a *closed non-terminating* set on v. A set of tokens G is a *closed loop* for a valuation v if and only if G is closed on v and every member of G calls every member of G.

Closed Loop Rule: If a set of tokens *G* is a closed loop for *v*, assign GAP to all the members of *G*.

The Groundless Tokens rule, in turn, depends on the definition of a groundless set of tokens. A set of tokens *G* is *groundless* for a valuation *v* if and only if *G* is a closed non-terminating set on *v* that does not have any non-empty subset that is a closed-loop for *v*. Groundless sets all involve, in one way or another, infinite descending chains of calls among tokens. The simplest example of a groundless set (under the empty evaluation) would be a set consisting of p_i for every natural number *i,* such that: $p_0\downarrow = \mathrm{Tr}(p_1)$, $p_1\downarrow = \mathrm{Tr}(p_2)$, … $p_n\downarrow = \mathrm{Tr}(p_{n+1})$,…

Groundless Tokens Rule: If a set of tokens *G* is groundless for a valuation *v,* assign GAP to all the members of G.

Gaifman proves that, for any set of tokens of $\mathcal{L}$, there is a unique total valuation for that set that can be reached from ∅ by applying his three rules (and this is true for any self-supporting initial valuation). So, the total evaluation one arrives at does not depend on the order in which the rules are applied.

Incorporating Stable Semantic States

In Section 2 of the paper, I introduce the notion of a *stable semantic state* and argue that, if a Gaifman-style Particularist semantics is to make empirically correct predictions about tokens, those predictions need to be sensitive to which state those tokens were produced in. Here, I will outline what sorts of adjustments need to be made to a Gaifman-style semantics to make that sensitivity explicit. I will suppose that theory is intended to build a total valuation for a set consisting of tokens of $\mathcal{L}$ produced in some state *G*. *G* will determine what expressions of $\mathcal{L}$ are

semantic v.s. non-semantic (relative to *G*) and provide a background interpretation for the non-semantic fragment.

First, in order for the GAP rules to deliver the right results, we have to suppose that, not only are pointer networks closed under the ( )1, ( )2, and (|) operations, but further that every operation on a token produced in *G* yields a token produced in *G*. (This is a technicality we can avoid by opting for Gaifman's 2000 semantics, which dispenses with operations on tokens/pointers. Though, if we think of operations like ( )1 and ( )2 as mapping tokens to their subtokens, this additional supposition is not very substantive: if a token conjunction was produced in S, then, presumably, so were its token conjuncts.)

Second, the definition of "*r* directly calls *q*" needs to be sensitive to what state *r* and *g* are produced in. (From the mere fact that *r* instantiates "Tr(*a*)" we cannot draw any conclusions about what it directly calls—this will depend on what "Tr" and "*a*" mean relative to the state *r* was produced in.) Since *G* is the only state we are building a valuation for, we can accomplish this by explicitly limiting the scope of the definition: if *r* is a produced in *G*, then *r* directly calls *q* if and only if one of the following conditions holds, etc.

Third, we need to explicitly relativize the notion of an "induced" valuation to a specific state. So, strictly, instead of saying that $\mathfrak{v}$ is the induced valuation of *v*, we should say that it is the induced valuation of *v* relative to the state *G*—though for theories that are only concerned with one state, this can be omitted. The recursive definition of the induced valuation of *v* relative to *G* should be updated to specify that the "background interpretation" that determines the induced value for atomic sentences is the background interpretation specified by *G*.

Fourth, and finally, we must update the Standard Value Rule so that its enabling condition includes facts about what state the token in question was produced in. As I suggest in

the body of the paper, at the end of Section Two, we can revise the rule this way (where $\mathfrak{v}$ is the induced valuation of *v* relative to *G*):

Standard Value Rule

> If $p\downarrow = \alpha$, *p* is produced in state *G*, $\mathfrak{v}(\alpha)$ is defined, and $v(p) \neq$ GAP, then assign to *p* the value $\mathfrak{v}(\alpha)$

With these updates in place, applications of Gaifman's three rules to the initial $\emptyset$ valuation will yield successive valuations in which additional tokens produced in *G* get evaluated. If there are tokens in the domain P that are not produced in *G*, none of the rules will be enabled on them at any stage in the evaluation process. Successive applications of the rules will result in a *total valuation* of all the tokens in P produced in state *G* if and only if no token produced in *G* directly calls a token that is not produced in *G*. This is a principled limitation—the theory only provides instructions for interpreting tokens that are produced in *G*, so if there are direct call relations that point beyond that domain, the theory needs supplemental information regarding how to interpret those tokens.

APPENDIX TWO

Here is a proof that every token of *u*, $\forall x((x \downarrow = \text{'}u\text{'}) \rightarrow \neg \text{Tr}(x))$, receives the value GAP in a Gaifman-style semantics, whatever is in domain of tokens the theory is concerned with.

Let *p* be an arbitrary token of *u*. *p* directly calls all the token substitution instances of $(x \downarrow = \text{'}u\text{'}) \rightarrow \neg \text{Tr}(x)$, i.e. all $(p \mid t)$ for every term *t*. Each substitution instance initiates a series of calls with the following structure:

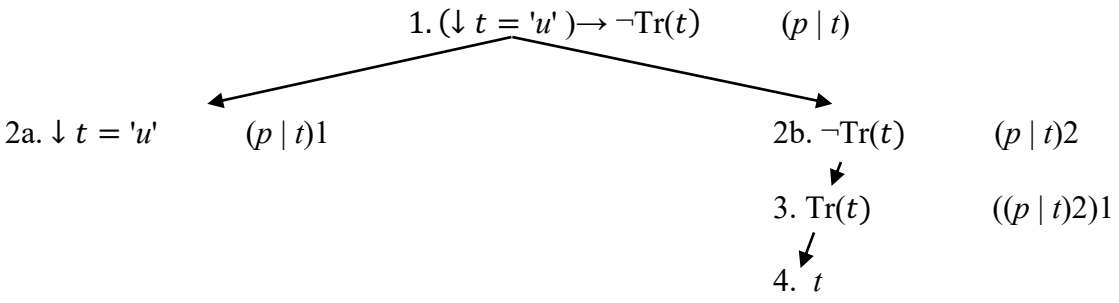

Figure Appendix 2.1

For all cases in which *t* refers to an object that is not a token of *u*, the 2a node will be a token of a sentence "$t \downarrow = \text{'}u\text{'}$," which receives the value FALSE in the background interpretation (given the intended interpretation for "↓"). By the definition of the induced valuation, the induced valuation of ∅ assigns $(p \mid t) \downarrow$ TRUE whenever *t* refers to an object that is not a token of *u*. So, in all substitution instances $(p \mid t)$ where *t* does not refer to a token of *u*, the Standard Value Rule is enabled on the 1 and 2a nodes of this tree relative to ∅. Since any token enabled relative to ∅ can be evaluated independently of any other token, there will be a valuation extending ∅ in which every token $(p \mid t)$, where *t* does not refer to a token of *u,* is evaluated but no other tokens are. Since *p* was an arbitrary token of *u,* every token of *u* initiates the same structure of direct calls.

There is, then a valuation v^* extending ∅ in which, for every token x of u, every token substitution (x | t), where t does not refer to a token of u, is evaluated but no other token is.

Consider some substitution instance (p | c), where c *does* refer to a token of u. p directly calls (p | c), node 1, which is a conditional calling the tokens in node 2a, (p | c)1, and node 2b, (p | c)2. The token in node 2a instantiates the sentence "c↓ = 'u'," which receives the value TRUE in the background interpretation. (p | c)2↓ = the sentence ¬Tr(c), so (p | c)2 in turn calls the token in node 3, ((p | c)2)1, which in turn calls the token in node 4, c. Since every token of u has a name in our semantic framework, for every token of u, there is some substitution instance (p | t), where t refers to that token. Since, p was an arbitrary token of u we can conclude that every token of u calls every token of u, via the analogous structure of calls.

Since every token occupying the 2a node for a substitution instance (x | t), where t refers to a token of u is enabled relative to ∅, there is a valuation v^{**}, extending v^*, which evaluates each of these tokens as FALSE, but evaluates nothing else. The following set, CL, is a closed loop relative to v^{**}: the set containing every token of u and every token that occupies nodes 1, 2b, or 3 for some substitution instance (x | t), where x is a token of u and t is a term referring to a token of u. First, I will show the set is a loop. Every token that occupies nodes 1, 2b, or 3 for some substitution instance (x | t), where x is a token of u and t is a term referring to a token of u, calls some token of u and is called by some token of u. Since every token of u calls every other token of u, and calling is transitive, it follows that every member of CL calls every other. Now, I will show that CL is closed on v^{**}: every member of CL is unevaluated by v^{**}, and for every o ∈ CL and every token r unevaluated by v^{**}, if there is a calling path from o to r consisting only tokens unevaluated by v^{**}, then r ∈ CL. By hypothesis, no member of CL is evaluated by v^{**}. Suppose that o is an arbitrary member of CL and r is an arbitrary token unevaluated by v^{**} and

there is a calling path from *o* to *r* consisting of tokens unevaluated by *v***. If there is a calling path consisting of tokens unevaluated by *v*** from *o* to *r* then there is a calling path consisting of tokens unevaluated by *v*** from some token of *u*, call it *k,* to *r*—since every member of CL is unevaluated and every token of *u* calls every member of CL. Every calling path from a token of *u* starts with its substitution instances, so if *r* is called by *k* it must be through some substitution instance (*k* | *t*). If (*k* | *t*) is a substitution instance where *t* does not refer to a token of *u*, there is no calling path from *k* to *r* consisting of tokens that are unevaluated by *v*** because every such substitution instance is evaluated as TRUE by *v***. Therefore, whatever calling path of unevaluated tokens leads from *k* to *r* goes via a substitution instance (*k* | *t*) where *t* does refer to a token of *u*. In other words, *r* must be a token occupying node 2a, 2b, 3, or 4 for some substitution instance (*k* | *t*) where *t* refers to a token of *u. r* cannot be a token occupying node 2a for such a substitution instance, because all 2a nodes are evaluated TRUE by *v*** whereas *r* is unevaluated. Therefore, *r* is a token occupying nodes 2b, 3 or 4 for a substitution instance (*k* | *t*), where *t* refers to a token of *u*. It follows that either *r* is itself a token of *u* (because anything occupying node 4 for such a substitution instance is a token of *u*) or it is a token occupying nodes 2b or 3 for a substitution instance (*x* | *t*), where *t* refers to a token of *u*. Therefore, *r* is a member of CL. Since *o* and *r* were arbitrary tokens, it follows that, for every $o \in CL$ and every token *r* unevaluated by *v***, if there is a calling path from *o* to *r* consisting only tokens unevaluated by *v***, then $r \in CL$. So, CL is closed on *v***. Therefore, CL is a closed loop relative to *v***. Therefore, in any fixed point extending ∅, every member of CL, and in particular every token of *u*, receives the value GAP. Our construction of *v*** depended on no assumptions about the domain besides what is required if it is to contain any tokens of *u*. Therefore, for any domain of tokens of L+, any token of *u* will receive the value GAP in the fixed point extending ∅.

In section 3 of the paper, I claim that the same sort of proof can be given for the sentence *State-Relative-u* = "$\forall x((x \downarrow= \textit{State-Relative-u} \wedge \text{Produced}(x, S)) \rightarrow \neg \text{Tr}(x))$." The only significant difference between *u* and *weakened-u* is the addition the conjunct "Produced(*x, S*)" in the restrictor of the universal generalization. "Produced" and "*S*" are non-semantic expressions the interpretation of which must specified by the background interpretation for L++: "*S*" names a semantic state and "Produced" holds of an ordered pair (*x*, *y*) if *x* is a token produced in state *y*. The proof that every token of *weakened-u* will receive GAP in an evaluation extending ∅ will proceed in exactly the same way as above, except that the tree of calls under each substitution instance for *weakened-u* will contain an additional nodes for substitutions of the form "Produced(*t, S*)" each of which gets an evaluation in the background interpretation, and is therefore enabled relative to ∅.

BIBLIOGRAPHY

Abreu Zavaleta, Martín. "Weak Speech Reports." *Philosophical Studies* 176, no. 8 (2019): 2139–2166.

Anderson, Luvell, and Ernie Lepore. "Slurring Words." *Noûs* 47, no. 1 (2011): 25–48.

Aquinas, Thomas. *Truth: A Translation of Quaestiones Disputatae De Veritate, Volume 1, Questions I-IX*. Translated by Mulligan, Robert W. Chicago: Henry Regnery Company, 1952.

Aristotle. *Metaphysics*. Translated by Reeve, C.D.C. Hacket Publishing Company, Inc., 2016.

Armour-Garb, Bradley. "Read and Indirect Revenge." In *Unity, Truth and the Liar*, edited by Shahid Rahman, Tero Tulenheimo, and Emmanuel Genot. Vol. 8. Logic, Epistemology, and the Unity of Science. Springer, 2008.

Augustine of Hippo. *On the Free Choice of the Will, on Grace and Free Choice, and Other Writings*. Translated by Peter King. Cambridge University Press, 2010.

Bacon, Andrew. "Can the Classical Logician Avoid the Revenge Paradoxes?" *Philosophical Review* 124, no. 3 (2015): 299–352.

———. "Radical Anti-disquotationalism." *Philosophical Perspectives* 32, no. 1 (2018): 41–107.

Ball, Derek. "Semantics as Measurement." In *The Science of Meaning: Essays on the Metatheory of Natural Language Semantics*, edited by Derek Ball and Brian Rabern. Oxford University Press, 2018.

Barker, Chris. "The Dynamics of Vagueness." *Linguistics and Philosophy* 25, no. 1 (2002): 1–36.

Bjerring, Jens Christian, and Wolfgang Schwarz. “Granularity Problems.” *Philosophical Quarterly* 67, no. 266 (2017): 22–37.

Blumberg, Kyle, and Harvey Lederman. “Revisionist Reporting.” *Philosophical Studies* 178, no. 3 (2020): 755–783.

Bobzien, Susanne. “Stoic Logic.” In *The Cambridge Companion to Stoic Philosophy*, edited by Brad Inwood. Cambridge University Press, 2003.

Bondarenko, Tatiana Igorevna. “Anatomy of An Attitude”. Ph.D. Thesis, Massachusetts Institute of Technology, 2022.

Bromberger, Sylvain. “What Are Words? Comments on Kaplan , on Hawthorne and Lepore, and on the Issue.” *Journal of Philosophy* 108, no. 9 (2011): 486–503.

Bronowski, Ada. *The Stoics on Lekta: All There Is to Say*. Oxford University Press, 2019.

Brower-Toland, Susan. “Medieval Theories of Propositions: Ockham and the Later Medieval Debate.” In *The Routledge Handbook of Propositions*, edited by Chris Tillman and Adam Murray. Routledge, 2022.

Burge, Tyler. *Origins of Objectivity*. Oxford University Press, 2010.

———. “Semantical Paradox.” *Journal of Philosophy* 76, no. 4 (1979): 169–198.

———. “Psychological Content and Egocentric Indexes.” In *Blockheads! Essays on Ned Block’s Philosophy of Mind and Consciousness*, edited by Pautz, Adam and Stoljar, Daniel, 41–69. MIT Press, 2019.

Buridan, Jean. *Sophisms on Meaning and Truth*. Appleton-Century-Crofts, 1966.

Camp, Joseph L. *Confusion: A Study in the Theory of Knowledge*. Harvard University Press, 2002.

Cappelen, Herman. *Fixing Language: An Essay on Conceptual Engineering*. Oxford University Press, 2018.

Cappelen, Herman, and Ernest Lepore. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Wiley-Blackwell, 2005.

Choi, Shungho, and Michael Fara. "Dispositions." *The Standford Encyclopedia of Philosophy* (2012).

Crimmins, Mark. "Contextuality, Reflexivity, Iteration, Logic." *Philosophical Perspectives* 9 (1995): 381–399.

Cumming, Sam. "Creatures of Darkness." *Analytic Philosophy* 54, no. 4 (2013): 379–400.

———. "Semantic Reasons." *Noûs* 57, no. 3 (2023): 641–666.

Dalglish, Steven. "Accepting Defeat: A Solution to Semantic Paradox with Defeasible Principles for Truth". Ohio State University, 2020. http://rave.ohiolink.edu/etdc/view?acc_num=osu1597757494987204.

Davidson, Donald. "Radical Interpretation." *Dialectica* 27, no. 1 (1973): 314–328.

———. "Truth and Meaning." *Synthese* 17, no. 3 (1967): 304–323.

Descartes, René. *The Philosophical Writings of Descartes, Vol. 3: Correspondence, Trans. by John G. Cottingham, Robert Stoothof, Dugald Murdoch, and Anthony Kenny.* Cambridge University Press, 1991.

Dickie, Imogen, and Gurpreet Rattan. "Sense, Communication, and Rational Engagement." *Dialectica* 64, no. 2 (2010): 131–151.

Dorr, Cian, and John Hawthorne. "Semantic Plasticity and Speech Reports." *Philosophical Review* 123, no. 3 (2014): 281–338.

Dummett, Michael A. E. “What Is a Theory of Meaning?” In *Mind and Language*, edited by Samuel D Guttenplan. Clarendon Press, 1975.

Dummett, Michael. “Truth.” *Proceedings of the Aristotelian Society* 59, no. 1 (1959): 141–62.

Dupré, Gabe. “Idealisation in Natural Language Semantics: Truth-Conditions for Radical Contextualists.” *Inquiry: An Interdisciplinary Journal of Philosophy* (2020).

Dupré, John. “Natural Kinds and Biological Taxa.” *Philosophical Review* 90, no. 1 (1981): 66–90.

Eklund, Matti. “Inconsistent Languages.” *Philosophy and Phenomenological Research* 64, no. 2 (2002): 251–275.

Elbourne, Paul. *Meaning: A Slim Guide to Semantics*. Oxford University Press, 2011.

Elliott, P.D. “Elements of Clausal Embedding”. Ph.D. Thesis, University College London, n.d.

Field, Hartry. “Egocentric Content.” *Noûs* 51, no. 3 (2017): 521–546.

———. *Saving Truth From Paradox*. Oxford University Press, 2008.

Fine, Kit. “Plantinga on The Reduction of Possibilist Discourse.” In *Alvin Plantinga*, edited by Tomberlin, J.E. and van Inwagen, P. Vol. 5. Profiles (An International Series on Contemporary Philosophers and Logicians). Dordrecht: Springer, 1985.

Francén, Ragnar. “Mananas, Flusses and Jartles: Belief Ascriptions in Light of Peripheral Concept Variation.” *Philosophical Studies* 179, no. 12 (2022): 3635–3651.

Frege, Gottlob. “The Thought: A Logical Inquiry.” *Mind* 65, no. 259 (1956): 289–311.

———. *The Frege Reader*. Edited by Beaney, Michael. Blackwell Publishing Ltd, 1997.

Fritz, Peter, and Jeremy Goodman. “Counting Incompossibles.” *Mind* 126, no. 504 (2017): 1063–1108.

Gaifman, Haim. “Operational Pointer Semantics: Solution to Self-referential Puzzles I.” In *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge*, edited by M. Y Vardi, 43–60. Morgan Kaufman, 1988.

———. “Pointers to Propositions.” In *Circularity, Definition, and Truth*, edited by Chapuis, André and Gupta, Anil, 2000.

———. “Pointers to Truth.” *Journal of Philosophy* 89, no. 5 (1992): 223.

Garcia-Carpintero, Manuel. “Indexicals as Token-reflexives.” *Mind* 107, no. 427 (1998): 529–564.

Gauker, Christopher. “Against Stepping Back: A Critique of Contextualist Approaches to the Semantic Paradoxes.” *Journal of Philosophical Logic* 35, no. 4 (2006): 393–422.

———. “Open Texture and Schematicity as Arguments for Non-referential Semantics.” In *Meaning, Context, and Methodology*, edited by Sarah-Jane Conrad and Klaus Petrus, 13–30. De Gruyter, 2017.

Gazdar, Gerald. “Pragmatics, Implicature, Presuposition and Lógical Form.” *Critica* 12, no. 35 (1979): 113–122.

Glanzberg, Michael. “A Contextual–hierarchical Approach to Truth and the Liar Paradox.” *Journal of Philosophical Logic* 33, no. 1 (2004): 27–88.

———. “Against Truth-Value Gaps.” In *Liars and Heaps: New Essays on Paradox*, edited by J. C Beall. Clarendon Press, 2004.

———. “Context and Unrestricted Quantification.” In *Absolute Generality*, edited by A. Rayo and G. Uzquiano, 45–74. Oxford University Press, 2006.

Goodman, Jeremy. “Reality Is Not Structured.” *Analysis* 77, no. 1 (2017): 43–53.

Grzankowski, Alex, and Ray Buchanan. “Content Pluralism.” *Inquiry: An Interdisciplinary Journal of Philosophy* (2021).

———. “Propositions on the Cheap.” *Philosophical Studies* 176, no. 12 (2019): 3159–3178.

Gupta, Anil. “Truth.” In *The Blackwell Guide to Philosophical Logic*, edited by L. Goble. Blackwell, 2001.

Hanks, Peter. *Propositional Content*. Oxford University Press, 2015.

———. “Structured Propositions as Types.” *Mind* 120, no. 477 (2011): 11–52.

———. “What Are the Primary Bearers of Truth?” *Canadian Journal of Philosophy* 43, no. 5 (2013): 558–574.

Hawthorne, John, and Ernie Lepore. “On Words.” *Journal of Philosophy* 108, no. 9 (2011): 447–485.

Haze, Tristan Grotvedt. *Meaning and Metaphysical Necessity*. Routledge, 2022.

Heck, Richard. “Semantic Accounts of Vagueness.” In *Liars and Heaps: New Essays on Paradox*, edited by J. C Beall, 106–27. Oxford University Press, 2003.

Herzberger, Hans. “Paradoxes of Grounding in Semantics.” *Journal of Philosophy* 67, no. 6 (1970): 145–167.

Hoffman, Joshua, and Rosenkrantz, Gary S. “Platonistic Theories of Universals.” In *The Oxford Handbook of Metaphysics*, edited by Loux, Michael J. and Zimmerman, Dean W. Oxford University Press, 2005.

Horwich, Paul. “A Use Theory of Meaning.” *Philosophy and Phenomenological Research* 68, no. 2 (2004): 351–372.

Ismael, Jenann. “Reflexivity, Fixed Points, and Semantic Descent; How I Learned to Stop Worrying and Love Reflexivity.” *Acta Analytica* 26, no. 4 (2011): 295–310.

Kant, Immanuel. *Critique of Pure Reason*. Edited by Guyer, Paul and Wood, Allen. Cambridge University Press, 1998.

Kaplan, David. "Demonstratives: An Essay on the Semantics, Logic, Metaphysics and Epistemology of Demonstratives and Other Indexicals." In *Themes From Kaplan*, edited by Joseph Almog, John Perry, and Howard Wettstein, 481–563. Oxford University Press, 1989.

———. "Words." *Aristotelian Society Supplementary Volume* 64, no. 1 (1990): 93–119.

King, Jeffrey C. (2007). The nature and structure of content. New York: Oxford University Press.

Klima, G. "Consequences of a Closed, Token-Based Semantics: The Case of John Buridan." *Bulletin of Symbolic Logic* 10, no. 4 (2004): 592–593.

Kratzer, Angelika. "Decomposing Attitude Verbs", 2006.

Kripke, Saul A. "A Puzzle About Belief." In *Meaning and Use*, edited by A. Margalit, 239–83. Reidel, 1979.

———. *Naming and Necessity: Lectures Given to the Princeton University Philosophy Colloquium*. Vol. 217. Harvard University Press, 1980.

Kripke, Saul. "Outline of a Theory of Truth." *Journal of Philosophy* 72, no. 19 (1975): 690–716.

Laurence, Stephen. "A Chomskian Alternative to Convention-based Semantics." *Mind* 105, no. 418 (1996): 269–301.

Leibniz, G.W. "The Principles of Philosophy, or, the Monadology (1714)." In *Philosophical Essays*, by Leibniz, G.W., 213–225. edited by Ariew, Roger and Garber, Daniel. Hacket Publishing Company, Inc., 1989.

Leitgeb, Hannes. "What Truth Depends On." *Journal of Philosophical Logic* 34, no. 2 (2005): 155–192.

Lepore, Ernie, and Kirk Ludwig. *Donald Davidson: Meaning, Truth, Language, and Reality*. Oxford University Press, 2005.

Lepore, Ernie, and Matthew Stone. "Pejorative Tone." In *Bad Words: Philosophical Perspectives on Slurs*. Oxford University Press, 2018.

Lewis, David K. "Languages and Language." In *Minnesota Studies in the Philosophy of Science*, edited by Keith Gunderson, 7:3–35. University of Minnesota Press, 1975.

Lewis, David Kellogg. *Convention: A Philosophical Study*. Vol. 20. Wiley-Blackwell, 1969.

Linnebo, Øystein. "Generality Explained." *Journal of Philosophy* 119, no. 7 (2022): 349–379.

Lopukhina, A, A. Laurinavichyute, K. Lopukhin, and O. Dragoy. "The Mental Representation of Polysemy Across Word Classes." *Frontiers in Psychology* 9, no. 192 (2018).

Ludlow, Peter. *Living Words: Meaning Underdetermination and the Dynamic Lexicon*. Oxford University Press, 2014.

Mates, Benson. "Leibniz on Possible Worlds." *Critica* 4, no. 10 (1970): 123–127.

———. "Synonymity." *Journal of Symbolic Logic* 19, no. 3 (1954): 223–223.

Maudlin, Tim. *Truth and Paradox: Solving the Riddles*. Vol. 116. Oxford University Press, 2004.

McCluskey, S.C. "Maya Observations of Very Long Periods of Venus." *Journal for the History of Astronomy* 14, no. 2 (1983).

McGee, Vann, and Brian McLaughlin. "Distinctions Without a Difference." *Southern Journal of Philosophy* 33, no. S1 (1995): 203–251.

Mehlberg, Henry. *The Reach of Science*. Vol. 10. University of Toronto Press, 1958.

Merricks, Trenton. *Propositions*. Oxford University Press, 2015.

Moltmann, Friederike. “Propositions, Attitudinal Objects, and the Distinction Between Actions and Products.” *Canadian Journal of Philosophy, Supplementary Volume on Propositions, Edited by G. Rattan and D. Hunter* 43, no. 5–6 (2013): 679–701.

Morscher, Edgar. “Bernard Bolzano.” In *Stanford Encyclopedia of Philosophy*, 2008.

Moulton, Keir. “CPs: Copies and Compositionality.” *Linguistic Inquiry* 46, no. 2 (2015): 305–342.

Murzi, Julien, and Lorenzo Rossi. “Reflection Principles and the Liar in Context.” *Philosophers’ Imprint* 18 (2018).

Parsons, Charles. “The Liar Paradox.” *Journal of Philosophical Logic* 3, no. 4 (1974): 381 – 412.

Patterson, Douglas. “Inconsistency Theories of Semantic Paradox.” *Philosophy and Phenomenological Research* 79, no. 2 (2009): 387 – 422.

Perry, John. *Reference and Reflexivity*. Center for the Study of Language and Inf, 2001.

Predelli, Stefano. “The Problem with Token-reflexivity.” *Synthese* 148, no. 1 (2006): 5–29.

Priest, Graham. “Revenge, Field, and ZF.” In *Revenge of the Liar*, edited by J.C. Beall. Oxford University Press, 2008.

Prinzing, Michael. “The Revisionist’s Rubric: Conceptual Engineering and the Discontinuity Objection.” *Inquiry: An Interdisciplinary Journal of Philosophy* 61, no. 8 (2018): 854–880.

Radulescu, Alexandru. “Token-Reflexivity and Repetition.” *Ergo: An Open Access Journal of Philosophy* 5 (2018): 745–763.

Read, Stephen. “The Liar Paradox from John Buridan Back to Thomas Bradwardine.” *Vivarium* 40, no. 2 (2002): 189–218.

Reichenbach, Hans. *Elements of Symbolic Logic*. Dover Publications, 1947.

Rescorla, Michael. “Reifying Representations.” 135–177, 2020.

Richard, Mark. *Meanings as Species*. Oxford University Press, 2019.

Salmon, Nathan. "Existence." *Philosophical Perspectives* 1 (1987): 49–108.

Sawyer, Sarah. "The Importance of Concepts." *Proceedings of the Aristotelian Society* 118, no. 2 (2018): 127–147.

Schiffer, Stephen. "Russell's Theory of Definite Descriptions." *Mind* 114, no. 456 (2005): 1135–1183.

Shapiro, Stewart. *Vagueness in Context*. Oxford University Press, 2006.

———. "Vagueness, Metaphysics, and Objectivity." In *Cuts and Clouds: Vagueness, Its Nature, and Its Logic*, edited by Richard Dietz and Sebastiano Moruzzi. Oxford University Press, 2010.

Simchen, Ori. "Token-Reflexivity." *Journal of Philosophy* 110, no. 4 (2013): 173–193.

Simmons, Keith. *Semantic Singularities: Paradoxes of Reference, Predication, and Truth*. Oxford University Press, 2018.

———. *Universality and the Liar: An Essay on Truth and the Diagonal Argument*. Cambridge University Press, 1993.

Soames, Scott. "Propositions as Cognitive Acts." *Synthese* 196, no. 4 (2019): 1369–1383.

———. "Cognitive Propositions." In *New Thinking About Propositions*, by King et al., 91–124. Oxford University Press, 2014.

Sosa, Ernest. "Propositional Attitudes De Dicto and De Re." *Journal of Philosophy* 67, no. 21 (1970): 883–896.

Stalnaker, Robert C. *Inquiry*. Cambridge University Press, 1984.

———. “Propositions.” In *Issues in the Philosophy of Language: Proceedings of the 1972 Oberlin Colloquium in Philosophy*, edited by Alfred F Mackay and Daniel Davy Merrill, 79–91. Yale University Press, 1976.

Stalnaker, Robert. *Mere Possibilities: Metaphysical Foundations of Modal Semantics*. Princeton University Press, 2011.

———. “Reference and Necessity.” In *A Companion to the Philosophy of Language*, edited by Bob Hale, Crispin Wright, and Alexander Miller, 902–919. Wiley-Blackwell, 1997.

Stojnić, Una. “Just Words: Intentions, Tolerance and Lexical Selection.” *Philosophy and Phenomenological Research* 105, no. 1 (2022): 3–17.

Strawson, Peter F. “On Referring.” *Mind* 59, no. 235 (1950): 320–344.

“Suitcase.” *Merriam Webster*, n.d. https://www.merriam-webster.com/dictionary/suitcase#:~:text=noun,traveler’s%20clothing%20and%20personal% 20articles.

Tarski, Alfred. “Der Wahrheitsbegriff in Den Formalisierten Sprachen.” *Studia Philosophica* 1 (1936): 261–405.

———. “The Semantic Conception of Truth and the Foundations of Semantics.” *Philosophy and Phenomenological Research* 4, no. 3 (1943): 341–376.

Travis, Charles. “Pragmatics.” In *A Companion to the Philosophy of Language*, edited by Bob Hale and Crispin Wright, 87–107. Blackwell, 1997.

Uzquiano, Gabriel. “A Neglected Resolution of Russell’s Paradox of Propositions.” *Review of Symbolic Logic* 8, no. 2 (2015): 328–344.

Waismann, F. “Verifiability.” *Journal of Symbolic Logic* 19, no. 1 (1951): 117–44.

Walsh, Sean. “Predicativity, the Russell-Myhill Paradox, and Church’s Intensional Logic.” *Journal of Philosophical Logic* 45, no. 3 (2016): 277–326.

Warren, Jared. “Restricting the T-schema to Solve the Liar.” *Philosophy and Phenomenological Research* 108, no. 1 (2023): 238–258.

Weatherson, Brian. “Vagueness as Indeterminacy.” In *Cuts and Clouds: Vagueness, Its Nature, and Its Logic*, edited by Richard Dietz and Sebastiano Moruzzi. Oxford University Press, 2010.

Williamson, Timothy. “Indefinite Extensibility.” *Grazer Philosophische Studien* 55, no. 1 (1999): 1–24.

———. *Modal Logic as Metaphysics*. Oxford University Press, 2013.

———. “Necessitism, Contingentism, and Plural Quantification.” *Mind* 119, no. 475 (2010): 657–748.

Wilson, Mark. “Predicate Meets Property.” *Philosophical Review* 91, no. 4 (1982): 549–589.

———. *Wandering Significance: An Essay on Conceptual Behavior.* Vol. 57. Oxford: Clarendon Press, 2006.

Yablo, Steve. “Grounding, Dependence, and Paradox.” *Journal of Philosophical Logic* 11, no. 1 (1982): 117–137.