Closed Structure

Peter Fritz[1,2] · Harvey Lederman[3] · Gabriel Uzquiano[4]

Abstract

According to the structured theory of propositions, if two sentences express the same proposition, then they have the same syntactic structure, with corresponding syntactic constituents expressing the same entities. A number of philosophers have recently focused attention on a powerful argument against this theory, based on a result by Bertrand Russell, which shows that the theory of structured propositions is inconsistent in higher order-logic. This paper explores a response to this argument, which involves restricting the scope of the claim that propositions are structured, so that it does not hold for all propositions whatsoever, but only for those which are expressible using closed sentences of a given formal language. We call this restricted principle CLOSED STRUCTURE, and show that it is consistent in classical higher-order logic. As a schematic principle, the strength of CLOSED STRUCTURE is dependent on the chosen language. For its consistency to be philosophically significant, it also needs to be consistent in every extension of the language which the theorist of structured propositions is apt to accept. But, we go on to show, CLOSED STRUCTURE is in fact inconsistent in a very natural extension of the standard language of higher-order logic, which adds resources for plural talk of propositions. We conclude that this particular strategy of restricting the scope of the claim that propositions are structured is not a compelling response to the argument based on Russell's result, though we note that for some applications, for instance to propositional attitudes, a restricted thesis in the vicinity may hold some promise.

✉ Peter Fritz
   peter.fritz@acu.edu.au

1 Dianoia Institute of Philosophy, Australian Catholic University, Melbourne, Australia

2 Department of Philosophy, Classics, History of Art and Ideas, University of Oslo, Oslo, Norway

3 Department of Philosophy, Princeton University, Princeton, NJ, USA

4 School of Philosophy, University of Southern California, Los Angeles, CA, USA

1 Introduction

According to the structured theory of propositions, if two sentences express the same proposition, then they have the same syntactic structure, with corresponding syntactic constituents expressing the same entities. A consequence of this view is that if the propositions expressed by simple predications Fa and Gb are identical, then the denotation of F is identical to that of G and the denotation of a is identical to that of b. Views along these lines go back to Russell [35]; more recent endorsements can be found in Kaplan [20], Salmon [36], Soames [37] and King [22, 23].

A number of philosophers, including Dorr [10] and Goodman [14], have recently focused attention on a powerful argument against this theory. Given that the theory concerns identity among propositions and properties, it is natural to formalize it using higher-order logic. In the language of higher-order logic, quantification over propositions is regimented as quantification binding variables occupying the syntactic position of sentences, and similarly for properties or relations and predicates. Using higher-order logic, we can formalize some of the commitments of the structured proposition view. We do so below using a principle we call STRUCTURE, which includes among its instances sentences of the following form:

$$\forall X \forall Y \forall x \forall y (Xx = Yy \rightarrow X = Y \wedge x = y)$$

A result by Russell [35, Appendix B] shows that the instances of this schema in which x and y occupy the syntactic position of sentences are inconsistent; *a fortiori*, STRUCTURE is inconsistent as well. (This result was rediscovered by Myhill [29], which is why it is sometimes named after Russell and Myhill.) So, it has been argued, the theory of structured propositions must be rejected.

How might a proponent of structured propositions respond? Broadly speaking, there are two strategies they might adopt, separately or in tandem. First, they might reject aspects of the formal framework of higher-order logic, whether on the grounds that the syntax of standard higher-order languages is somehow problematic, or on the grounds that standard principles of such logics are not valid. Second, they might formulate and endorse a theory which gives voice to the idea that propositions are in some sense structured, without however implying the problematic instances of STRUCTURE. An example of the former strategy is a predicativist response, which weakens the logical principles of quantification so as to restrict higher-order comprehension principles to predicative instances. In such instances, the defining condition of an entity of some type involves no bound variables of that type or higher, or free variables of higher type. Walsh [39] shows that STRUCTURE is consistent in such a system. In this paper we will explore a response of the second type, which retains the standard logical principles of higher-order quantification, but replaces STRUCTURE with a weaker principle we call CLOSED STRUCTURE.

One way of motivating this weaker principle is via the role structured propositions play in the philosophy of language, where they are taken to serve as the meanings of sentences of natural languages like English. But complete sentences of English do not have free variables (setting aside demonstratives and anaphora for the moment; we return to them in Section 5.1). Faced with the inconsistency of STRUCTURE, a philosopher of language might thus claim that all that matters for

their theorizing is that complete sentences of languages like English which do not have free variables, express structured propositions. Formally, they might thus propose to replace the universally quantified formulas of STRUCTURE by their instances formed using closed terms. The resulting schematic principle of CLOSED STRUCTURE thus includes among its instances sentences (formulas without free variables) of the following form:

$$\xi\varepsilon = \zeta\eta \to \xi = \zeta \wedge \varepsilon = \eta$$

Whereas STRUCTURE imposes conditions on all propositions, CLOSED STRUCTURE imposes conditions only on propositions that are expressible by closed formulas.

Can this move from STRUCTURE to CLOSED STRUCTURE help the proponent of structured propositions? A first question about the approach is whether the restriction in CLOSED STRUCTURE is sufficiently restrictive, i.e. whether this principle, unlike STRUCTURE, is in fact consistent in standard classical higher-order logic. The first main result of this paper shows that it is: no inconsistency can be derived from the instances of CLOSED STRUCTURE in this language.

From a formal perspective, this result sheds light on an interesting aspect of the Russell-Myhill theorem itself. The standard proof of the Russell-Myhill theorem does not explicitly specify counterexamples to STRUCTURE: for example, for a certain argument a, the proof specifies a property F for which it can be proven that there exists a property $X \neq F$ such that $Xa = Fa$. This shows that a counterexample to STRUCTURE must exist, but it doesn't explicitly specify one, since it doesn't explicitly specify X. Our result shows that this aspect of the proof of the Russell-Myhill theorem is in fact essential: all failures of STRUCTURE which can be established in our background logic involve free variables.

This first result seems to be good news for theorists of structured propositions, opening up a potentially interesting avenue of response to the argument against structured propositions based on the Russell-Myhill theorem. But our second main result reveals stark limitations on the extent of the good news. CLOSED STRUCTURE is a schematic principle, so its strength depends on the expressive power of the language in which it is stated. It is only philosophically interesting that CLOSED STRUCTURE is consistent in the basic higher-order language if it is also consistent in every extension of the language which the theorist of structured propositions is apt to accept. But our second main result shows that CLOSED STRUCTURE is in fact inconsistent in a very natural extension of the standard language of higher-order logic.

This extension adds a way of speaking *plurally* about propositions. Boolos [4] argued that we can not just quantify singularly over things, such as sets, but also plurally, using phrases such as 'there are some sets such that . . . '. His observations have motivated what we can think of as an extension of first-order logic in which such plural quantification over individuals can be regimented. We adapt these ideas to our type-theoretic setting, and introduce an extension in which we can quantify plurally over propositions and form closed plural terms. Such resources are novel, but they can be motivated as a natural combination of the ideas of higher-order languages and plural quantification.

We show that once the plural resources are added, there are instances of CLOSED STRUCTURE which are inconsistent. The proof of this result makes use of a tight

connection between the Russell-Myhill theorem and Cantor's theorem. In fact, Russell's proof of his result is completely parallel to the standard proof of Cantor's theorem. The latter shows that for every function f from a powerset $\mathcal{P}(A)$ to the set A, there are distinct sets $X, Y \in \mathcal{P}(A)$ such that $f(X) = f(Y)$. Boolos [5] and Kanamori [17] describe proofs of Cantor's theorem which show how such sets can be defined explicitly, for any given function f. Kanamori [17] derives the result as a corollary of a lemma for the proof of Zermelo's well-ordering theorem, which Bell [3] had adapted to the study of type-reducing correspondences between second- and first-order entities. These proofs effectively adapt the so-called paradox of Burali-Forti [6]. We show how the proof of Cantor's theorem by Boolos and Kanamori can be adapted to the Russell-Myhill theorem, which gives us a proof of the inconsistency of CLOSED STRUCTURE. Our result relies on plural talk of propositions, as it crucially depends on the extensionality of such plurals, i.e., a principle which roughly states that if every proposition is one of these propositions just in case it is one of those propositions, then these propositions are those propositions.

We argue that if one has already accepted the basic higher-order language used here, one should accept its extension by our plural resources as well. Because of this, we conclude that CLOSED STRUCTURE must be given up. In spite of our positive first result, this second result shows that the response to the Russell-Myhill theorem which restricts STRUCTURE to CLOSED STRUCTURE is unsuccessful.

The main focus of this paper is on responses to the Russell-Myhill theorem which uphold classical logic and endorse CLOSED STRUCTURE as a weakening of STRUCTURE. But at the end of the paper we turn briefly to the implications of our discussion for one style of non-classical theory, which rejects the classical principle that if $a = b$ then any sentence φ is equivalent to the sentence which results from substituting a for b in φ. One important application of the theory of propositions is to theorizing about propositional attitudes as well as propositional attitude reports. But arguably identity does not license substitution within the context of attitude reports. For instance, although Hesperus is Phosphorus, it is arguably not true that necessarily anyone who believes Hesperus is Hesperus believes that Hesperus is Phosphorus. So if we want to develop a systematic theory of identity which covers propositional attitudes we may reject the idea that coreferring names can be everywhere intersubstituted. And, if we reject this idea for names, it is natural to reject it for sentences as well, to hold that two expressions which denote identical propositions may fail to be intersubstitutable inside the scope of an operator corresponding to an attitude verb. For some philosophers these considerations might suggest that it is not identity, but some finergrained notion of 'cognitive equivalence' which is of most interest for the study of propositional attitudes, and, therefore, to the study of propositions in application to these attitudes.

To the extent that STRUCTURE was motivated in the first place by reflection on propositions as the objects of attitudes, philosophers who hold that identity does not license intersubstitution within attitude reports may be interested not in STRUCTURE and CLOSED STRUCTURE, which govern the identity relation, but variants of these principles which govern cognitive equivalence instead. If two expressions express cognitively equivalent entities, then they are intersubstitutable in every context, and it follows from this that the relevant variant of STRUCTURE will still be inconsistent by the Russell-Myhill theorem. But such a theorist may endorse the relevant

variant of CLOSED STRUCTURE, which can still be shown to be consistent. And, crucially, they should not be concerned about our result concerning plural propositional quantification. For that result relies on the claim that, when two plural terms express coextensive and hence identical pluralities, they are everywhere intersubstitutable. Someone who already rejects the claim that terms for identical entities are intersubstitutable will naturally reject the claim that terms for identical plural terms are intersubstitutable, and rejecting this assumption of the theorem is sufficient to escape a version of our second result. So, while our results show that CLOSED STRUCTURE is untenable, they leave open whether a variant of CLOSED STRUCTURE, which concerns cognitive equivalence rather than identity, is tenable.

The plan of the paper is as follows. Section 2 introduces the background formalism of standard higher-order logic. Section 3 proves the consistency of CLOSED STRUC-TURE in this setting, alongside a couple of strengthenings of this result. Section 4 presents the extension of the logic by plural propositional quantification, and proves the inconsistency of CLOSED STRUCTURE in this setting. Section 5 provides more extended conceptual discussion: we present motivations for restricting STRUCTURE to CLOSED STRUCTURE (5.1), defend the intelligibility of plural propositional quantification (5.2), and finally consider in detail the considerations about attitude reports just described (5.3). The reader primarily interested in conceptual issues related to our results may wish to skip to Section 5, turning back to previous sections as necessary for relevant definitions and statements. Section 6 concludes, and two appendices include some technical lemmas appealed to in the presentation of the main results.

2 Background

2.1 Pure Relational Type Theory

This section introduces the basic logical framework, which will be a form of type theory. Type theories codify different grammatical categories as types. For example, many type theories include a type for individual terms, a type for sentences, and types for various predicates, such as the type of predicates which take individual terms as arguments.

There are two common forms of type theories, a functional and a relational one. Here, we adopt a simple relational typed language, following Gallin [13, ch. 3] who presented a simplified version of the type theory used by Montague [27]. Since we focus on propositions, we don't essentially need a type of individual terms. Everything below would still go through with such an addition, but for simplicity, we leave it out. This allows us to work with a very simple type theory, which is defined by a single recursive rule, according to which every finite sequence of types $\langle t_1, \ldots, t_n \rangle$ is a type, the type of relational expressions which take n expressions, of types $t_1, \ldots, t_n$, respectively, as arguments.

Definition 2.1 Let the set T of types be the smallest set such that $\langle t_1, \ldots, t_n \rangle \in T$ whenever $t_1, \ldots, t_n \in T$.

This requires some unpacking: Note that n may be 0; thus $\langle\rangle$ is a type. This is the type of relational expressions which take no arguments, which we can consider to be formulas. From this first type, we can recursively generate higher types. For example, $\langle\langle\rangle\rangle$ is the type of predicates which take formulas as arguments, i.e., of unary sentential operators. Analogously, $\langle\langle\rangle, \langle\rangle\rangle$ is the type of binary sentential operators, and $\langle\langle\langle\rangle\rangle\rangle$ is the type of expressions taking a single sentential operator as an argument.

Let a *(T-indexed) family of sets* be a function X mapping every $t \in T$ to a set X_t. To define the formal language, we fix families of sets V and C, which specify the variables and constants of each type, respectively. We assume that all sets V_t are countably infinite, and that any two sets in these families are disjoint. Further, we assume that the following constants are included, which we call *logical*:

$$\rightarrow \; \in C_{\langle\langle\rangle, \langle\rangle\rangle}$$
$$\forall_t \in C_{\langle\langle t\rangle\rangle} \text{ for all } t \in T$$

The remaining constants will be called *non-logical*. $\rightarrow$ is intended to be interpreted as the material conditional. $\forall_t$ is intended to be interpreted as the relevant form of universal quantification. Note that $\forall_t$ is not a variable-binder; rather, $\forall_t$ combines with a predicate F of type $\langle t\rangle$ to form a sentence $\forall_t F$. The customary uses of quantifiers as variable-binders will be recaptured by introducing a way of constructing a relation term $(\lambda x.\varphi)$ from any variable x and formula φ, expressing the property of being an x such that φ. With this, $\forall(\lambda x.\varphi)$ states that every x is such that φ.

In general, we will allow λ to bind any finite non-empty non-repeating sequence of variables, thus producing complex expressions of all types apart from $\langle\rangle$. The only other syntactic constructions in the formal language are the use of variables and constants as expressions of the relevant type, and the application of relational expressions to suitably typed arguments to produce formulas.

Definition 2.2 Let L be the family of sets defined inductively as follows, where any $\varepsilon \in L_t$ is called an *expression of type t*, and expressions of type $\langle\rangle$ are also called *formulas*:

ATOMIC: Any variable or constant of type t is an expression of type t.
APPLICATION: $(\eta\varepsilon_1 \ldots \varepsilon_n)$ is a formula,
    if $n > 0$, $\varepsilon_1, \ldots, \varepsilon_n$ are expressions of types $t_1, \ldots, t_n$, respectively, and η is an expression of type $\langle t_1, \ldots, t_n\rangle$.
ABSTRACTION: $(\lambda x_1 \ldots x_n.\varphi)$ is an expression of type $\langle t_1, \ldots, t_n\rangle$,
    if $n > 0$, φ is a formula and $x_1, \ldots, x_n$ are pairwise distinct variables of types $t_1, \ldots, t_n$, respectively.

We adopt a number of conventions to simplify notation. First, recall that we have not specified what the elements of V_t, i.e., the variables, are. Any terms such as 'x' are thus meta-variables – variables in the language of logical English in which this article is written, possibly ranging over object-language variables. Similarly, '$\rightarrow$' is treated as a meta-language expression denoting an object language expression. By convention, a term such as 'Fa' denotes the result of concatenating F and a.

We conventionally restrict the range of meta-variables via the kind of letter used. Meta-variables ranging over object-language variables will generally be taken from the Roman alphabet. Restrictions to specific types will sometimes be notated by specifying the type in the upper index of the first occurrence, for example, writing '$X^{\langle\langle\rangle\rangle}$' for a meta-variable ranging over variables of type $\langle\langle\rangle\rangle$. Furthermore, '$p$', '$q$', '$r$', '$s$' customarily range over variables of type $\langle\rangle$, also called *propositional variables*, and uppercase Roman letters customarily range over variables of higher types. Greek letters will generally be used to range over any expressions (of a given type), whether simple or complex, or sets thereof. In particular, 'φ', 'ψ', 'χ' customarily range over formulas, whereas 'Γ', 'Δ' customarily range over sets of formulas. 'ε' and 'η' will often be used to range over expressions of arbitrary type. Meta-variables standing for specific object-language expressions will be taken from the Roman alphabet, but notated using a sans-serif font, such as 'F'.

Next, we indicate sequences by adding a bar on top of the relevant expression. E.g., '$\bar{t}$' usually indicates a sequence $\langle t_1, \ldots, t_n \rangle$ of types, for some $n < \omega$. '$\forall \bar{x}$' abbreviates '$\forall x_1 \ldots \forall x_n$'. We omit parentheses when this does not lead to ambiguity. $\rightarrow$ will be written in infix notation; e.g., we use '$p \rightarrow q$' to denote $\rightarrow pq$. Likewise, an expression such as $\forall x^t \varphi$ stands for $\forall_t \lambda x^t. \varphi$. Similar conventions will be used for the following abbreviations:

$$\bot := \forall p\ p \qquad\qquad \neg := \lambda p.p \rightarrow \bot$$
$$\top := \neg\bot \qquad\qquad \vee := \lambda pq. \neg p \rightarrow q$$
$$\wedge := \lambda pq. \neg(p \rightarrow \neg q) \qquad\qquad \leftrightarrow := \lambda pq.(p \rightarrow q) \wedge (q \rightarrow p)$$
$$\exists_t := \lambda X^{\langle t \rangle}. \neg \forall y^t \neg Xy \qquad\qquad =_t := \lambda x^t y^t. \forall Z^{\langle t \rangle}(Zx \leftrightarrow Zy)$$
$$\sqsubseteq_{\bar{t}} := \lambda X^{\bar{t}} Y^{\bar{t}}. \forall \bar{z}^{\bar{t}}(X\bar{z} \rightarrow Y\bar{z}) \qquad\qquad \equiv_{\bar{t}} := \lambda X^{\bar{t}} Y^{\bar{t}}. X \sqsubseteq Y \wedge Y \sqsubseteq X$$

Further, we indicate negations of applications of binary sentential operators by overlaying the operator, written infix, with a slash. E.g., $p \neq q$ abbreviates $\neg(p = q)$.

Despite the use of common logical symbols as abbreviations, it should not be assumed that these stand for the relevant logical notions. It should not be assumed, for example, that $\neg$ expresses negation. Doing so would be to make assumptions about the relationship between logical connectives which should be treated as controversial in the present context. Instead, we will only rely on the truth-functional adequacy of the abbreviations. For example, if $C_{\langle\langle\rangle\rangle}$ includes a constant n for negation, we only assume the truth of:

$$\forall p(\neg p \leftrightarrow np)$$

We don't assume the stronger claim that the relevant propositions are identical:

$$\forall p(\neg p = np)$$

The reason why we do not simply treat $\neg$, $\vee$ etc. as constants is that this would render some of the constructions below tediously complicated. No difficulties would arise from doing so, however.

We define as usual the notion of a *free* (occurrence of) a variable, and the notion of an expression being *free for* a variable in an expression. If expressions $\bar{\eta}$ are free for variables $\bar{x}$ in ε, we write $\varepsilon[\bar{\eta}/\bar{x}]$ for the result of simultaneously replacing every free occurrence of x_i in ε by η_i, and $\varepsilon(\bar{\eta})$ for $\varepsilon[\bar{\eta}/\bar{x}]$, given contextually salient free variables $\bar{x}$.

Definition 2.3 Let $\vdash$ be the proof system in L with the following axiom schemas and rules:

TAUT:	tautologies	MP:	$\varphi, \varphi \to \psi/\psi$
UI:	$\forall x \varphi \to \varphi[\varepsilon/x]$	UG:	$\varphi \to \psi/\varphi \to \forall x \psi$ (x not free in φ)
$e\beta 1:$	$(\lambda \bar{x}.\varphi)\bar{x} \to \varphi$	$e\beta 2:$	$\varphi \to (\lambda \bar{x}.\varphi)\bar{x}$

We write $\vdash \varphi$ for φ being derivable in $\vdash$.

Here, a formula is considered a tautology just in case it is built up from propositional variables using $\to$ and the defined connectives $\bot, \neg, \top, \vee, \wedge, \leftrightarrow$, corresponding to a theorem of classical propositional logic. TAUT and MP can thus be seen as contributing the principles of classical propositional logic, and analogously, UI and UG can be seen as contributing the principles of elementary quantification theory without identity. $e\beta 1$ and $e\beta 2$ codify the natural idea that some things stand in the relation of being so-and-so just in case they are so-and-so. Note that using the other axiomatic principles, we can derive from $e\beta 1$ and $e\beta 2$ any instance of the schema $(\lambda \bar{x}.\varphi)\bar{\varepsilon} \leftrightarrow \varphi[\bar{\varepsilon}/\bar{x}]$.

From these six principles, we can derive all the principles of elementary quantification theory with identity (the principles of first-order logic, extended to higher-order quantifiers), including the reflexivity of $=$ and Leibniz's Law. Additionally, we can derive an extensional comprehension principle which states for any formula that there is a relation which relates some things just in case they satisfy the formula:

$$\vdash \varepsilon = \varepsilon$$
$$\vdash \varepsilon = \eta \to (\varphi(\varepsilon) \to \varphi(\eta))$$
$$\vdash \exists X \forall \bar{y}(X\bar{y} \leftrightarrow \varphi) \quad (X \text{ not free in } \varphi)$$

Of course, we cannot derive the claim that co-extensiveness entails identity, i.e., that if $\varepsilon \equiv \eta$, then $\varepsilon = \eta$. These observations can be established along familiar lines, so we omit the details. When presenting evidence for claims about derivability in the following, we give indications as to how the proof would proceed, relying on any standard inferences in elementary quantification theory with identity.

Because we don't treat $\forall$ as a variable-binder, there is one highly plausible quantificational principle which does not correspond to any principle of first-order logic. This is the schema $\forall_t \varepsilon \leftrightarrow \forall x^t \varepsilon x$. Whether it is added to the present axiom system is immaterial for all of the following, as the schema is valid on the relevant models, but not required for the relevant deductions. We have therefore omitted the axiom schema for simplicity, but it could be added without changing any of the results.

As noted above, there is an intended interpretation for $\to$ and $\forall_t$, on which $\to$, for example, is interpreted as the material conditional. More generally, we assume that any interpretation of non-logical constants induces a corresponding intended interpretation of all expressions. We further assume that on any such interpretation, every sentence which is derivable in $\vdash$ is true.

Define the *universal closure* of a formula φ, written $\forall(\varphi)$, as $\forall \bar{x} \varphi$, where $\bar{x}$ are the free variables in φ (ordered according to some fixed well-ordering of the variables). We say that Γ *entails* Δ, written $\Gamma \vdash \Delta$, if for every $\delta \in \Delta$, there are $\gamma_1, \ldots, \gamma_n \in \Gamma$ such that $\vdash \forall(\gamma_1) \wedge \cdots \wedge \forall(\gamma_n) \to \forall(\delta)$. In such a context, we write Γ, Δ for $\Gamma \cup \Delta$, and use the name of a schematic principle for the set of its instances. Γ is *inconsistent*

if $\Gamma \vdash \bot$, and *consistent* otherwise. Note that according to these conventions, p is inconsistent, since $\forall(p)$ is $\forall p\ p$, i.e., $\bot$. The reason for defining $\vdash$ in this way is that it simplifies the statement of schematic principles, allowing us to consider them as implicitly universally generalized.

2.2 Structure

According to the structured view of propositions, if $Fa = Gb$ then $F = G$ and $a = b$. This should hold not just for specific examples F, G, a, b, but in general. Furthermore, such a principle should hold not only for unary applications of a specific type, but for applications of any type $\bar{t}$. Thus we arrive at the following principle, which is schematic in the type, in the sense that X, Y may be of any type $\langle t_1, \ldots, t_n \rangle$ (with $n > 0$), and $\bar{x}, \bar{y}$ accordingly of types $t_1, \ldots, t_n$, and where i is any number $1 \le i \le n$:

STRUCTURE: $\forall X \forall Y \forall \bar{x} \forall \bar{y}(X\bar{x} = Y\bar{y} \rightarrow X = Y \wedge x_i = y_i)$

STRUCTURE is inconsistent, as the Russell-Myhill theorem shows.

Theorem 2.4 (Russell [35, Appendix B], Myhill [29]) *For any formula χ with a free variable of type $\langle\langle\rangle\rangle$, there is an expression $\varepsilon^{\langle\langle\rangle\rangle}$ such that*

$$\vdash \exists X^{\langle\langle\rangle\rangle}((\chi(X) = \chi(\varepsilon)) \wedge X \not\equiv \varepsilon).$$

Proof Define an expression R of type $\langle\langle\rangle\rangle$ as follows:

$$\mathsf{R} := \lambda p.\neg \forall X((\chi(X) = p) \rightarrow Xp)$$

Assume for contradiction that $\neg\mathsf{R}\chi(\mathsf{R})$. Then $\forall X((\chi(X) = \chi(\mathsf{R})) \rightarrow X\chi(\mathsf{R}))$. So as $\chi(\mathsf{R}) = \chi(\mathsf{R})$, it follows that $\mathsf{R}\chi(\mathsf{R})$, contradicting the assumption. Thus $\mathsf{R}\chi(\mathsf{R})$, whence there is an X such that $\chi(X) = \chi(\mathsf{R})$ and $\neg X\chi(\mathsf{R})$. Since $\mathsf{R}\chi(\mathsf{R})$, $\chi(\mathsf{R})$ witnesses $X \not\equiv \mathsf{R}$. We conclude $\exists X((\chi(X) = \chi(\mathsf{R})) \wedge X \not\equiv \mathsf{R})$. $\qquad \square$

Corollary 2.5 STRUCTURE *is inconsistent in* $\vdash$.

Proof Letting χ be $X\bot$, Theorem 2.4 constructs an expression ε for which:

$$\vdash \exists X((X\bot = \varepsilon\bot) \wedge X \not\equiv \varepsilon).$$

Since $X \ne \varepsilon$ follows from $X \not\equiv \varepsilon$, this contradicts the instance of STRUCTURE for type $\langle\langle\rangle\rangle$. $\qquad \square$

The argument against STRUCTURE does not give any concrete example of a failure of STRUCTURE. For example, in the proof of Corollary 2.5, X has not been specified explicitly, and there is no obvious way of doing so. Thus, although we know that failures of STRUCTURE exist among simple predications of type $\langle\langle\rangle\rangle$, we have not explicitly specified any concrete examples of this.

This raises the question whether any concrete examples of a failure of STRUCTURE can be specified. If not, one might hope that in the course of any applied philosophical enterprise, e.g., in the philosophy of language, we will never encounter

such counterexamples to SMALLCAPS STRUCTURE. We might thus wonder whether the results of instantiating the universal quantifiers in STRUCTURE using closed expressions are consistent:

CLOSED STRUCTURE: $\xi\bar{\varepsilon} = \zeta\bar{\eta} \rightarrow \xi = \zeta \wedge \varepsilon_i = \eta_i$ (where $\xi, \zeta, \bar{\varepsilon}, \bar{\eta}$ are closed)

In the next section, we show that the instances of CLOSED STRUCTURE in L are in fact consistent.

3 Consistency

To show the consistency of CLOSED STRUCTURE, we construct models for it. Simple intensional models for $\vdash$ start from a set of worlds W, and recursively define domains for different types, by letting the domain of a type $\bar{t}$ be the set of functions from sequences of elements of domains of types $t_1, \ldots, t_n$ to $\mathcal{P}(W)$, the powerset of W. The idea behind such models is that relations can be modeled using a function mapping any suitable arguments to the set of worlds in which they are related by the relation. Note that the case of the general definition of domains for $n = 0$ yields that elements of the domain of type $\langle\rangle$, i.e., propositions, are functions from the set of sequences of length zero, i.e., $\{\langle\rangle\}$, to $\mathcal{P}(W)$. Such functions stand in a natural one-to-one correspondence with sets of worlds.

These simple models clearly can't be used to show the consistency of CLOSED STRUCTURE, since they individuate propositions in a rather coarse-grained manner. But we can adapt them, by adjoining syntactic components, letting elements of the domains be pairs of coarse-grained "logical" contents and fine-grained "syntactic" contents. In the case of type $\langle\rangle$, this takes the straightforward form of letting elements of the domain of type $\langle\rangle$ be pairs whose first element is a function from $\{\langle\rangle\}$ to a powerset $\mathcal{P}(W)$ and whose second element is a formula. For higher types, the modification of lower-level domains has to be taken into account recursively, so that elements of the domain of type $\bar{t}$ are pairs whose first element is a function from sequences of elements of domains of types $t_1, \ldots, t_n$ to $\mathcal{P}(W)$, and whose second element is an expression of the relevant type.

Such models suffice to show that CLOSED STRUCTURE is consistent. But as we will see, they do so by enforcing a level of fineness of grain which even adherents of CLOSED STRUCTURE might find excessive. For example, these models require sentences which differ only in relabeling bound variables to express distinct propositions. We therefore generalize the construction, and allow for syntactic components to be, instead of individual expressions, equivalence classes of expressions under some suitable equivalence relation on terms.

Before defining these specific models more rigorously, it is worth noting that the consistency of CLOSED STRUCTURE can already be derived relatively easily from a result of Muskens [28, Lemma 6.1, p. 108]. However, this proof method does not obviously extend to some of the strengthened consistency result we will establish below. Furthermore, the models we will use provide a more intuitive picture of how propositions might be individuated in a way which is consistent with CLOSED STRUCTURE. The idea of constructing models by pairing coarse-grained contents

with more fine-grained (syntactic) contents has also proven fruitful in other contexts; see, e.g., Dorr [10, Appendix].

3.1 Models

The equivalence relation on terms just mentioned must respect types, and satisfy three coherence constraints. The first of these prevents distinct variables from being related. The second and third require, recursively, that any complex expressions are equivalent if they have the same structure and corresponding components are equivalent. We call such (type-indexed) equivalence relations *coarsenings*:

Definition 3.1 Let a *coarsening* be a function mapping every $t \in T$ to an equivalence relation $\sim_t$ on L_t such that:

(1) For any variables of type t, if $x \sim_t y$ then $x = y$.
(2) If $\varepsilon_1 \sim_{t_1} \eta_1, \ldots, \varepsilon_n \sim_{t_n} \eta_n$ and $\xi \sim_{\bar{t}} \zeta$, then $\xi \bar{\varepsilon} \sim_{\langle \rangle} \zeta \bar{\eta}$.
(3) If $\varphi \sim_{\langle \rangle} \psi$, then $\lambda \bar{x}.\varphi \sim_{\bar{t}} \lambda \bar{x}.\psi$.

We usually omit the index of $\sim_t$, and write $[\varepsilon]_\sim$ for the equivalence class of ε under $\sim$, i.e., $\{\eta : \eta \sim \varepsilon\}$, and $L_t/\sim$ for the set of equivalence classes under $\sim$, i.e., $\{[\varepsilon]_\sim : \varepsilon \in L_t\}$. Given the constraints on a coarsening, we can lift the syntactic operations of application and λ-abstraction to equivalence classes, which we notate as follows:

$$[\varepsilon]_\sim + [\eta_1]_\sim + \cdots + [\eta_n]_\sim := [\varepsilon \eta_1 \ldots \eta_n]_\sim$$
$$\lambda \bar{x} + [\varphi]_\sim := [\lambda \bar{x}.\varphi]_\sim$$

Relative to a coarsening, domains can be defined as outlined, writing Y^X for the set of functions from X to Y, $X \times Y$ for the set of pairs $\langle x, y \rangle$ with $x \in X$ and $y \in Y$, and $\Pi_{1 \leq i \leq n} X_i$ for the set of sequences $\langle x_1, \ldots, x_n \rangle$ with $x_1 \in X_1, \ldots, x_n \in X_n$.

Definition 3.2 For any set W and coarsening $\sim$, define a set $D_{\bar{t}}$, for every type $\bar{t} = \langle t_1, \ldots, t_n \rangle$, as follows:

$$D_{\bar{t}} := \mathcal{P}(W)^{\Pi_{1 \leq i \leq n} D_{t_i}} \times L_{\bar{t}}/\sim$$

In addition to providing a set of worlds and a coarsening, a model will have to specify how constants are to be interpreted, and which propositions are true. The interpretation of constants is specified using a function mapping any constant to its logical content; the syntactic content is determined by the constant itself. The true propositions are determined by specifying a designated actual world.

Definition 3.3 A *closed structure model* (CSM) is a structure $\langle W, \sim, I, @ \rangle$, such that W is a set, $\sim$ is a coarsening, I is a function mapping every non-logical constant c of type $\bar{t}$ to a function $I(c) : \Pi_{1 \leq i \leq n} D_{t_i} \to \mathcal{P}(W)$, and $@ \in W$. For such a model,

an *assignment function* is a function mapping every variable of a type t to an element of D_t.

We can now define how CSMs interpret expressions of the language, by defining a function $[\![\cdot]\!]^{\cdot}$ which maps every expression ε to an element $[\![\varepsilon]\!]^a$ of the relevant domain, relative to an assignment function a. The interpretation of a variable is determined, as usual, by the assignment function. The interpretation of a constant contains, as its syntactic component, the equivalence class of the constant. The logical component is specified by the interpretation function in the case of non-logical constants; in the case of logical constants, this is given by the usual set-theoretic operations, suitably lifted.

Some notational conventions are useful to specify these lifts. First, we specify functions using '$\mapsto$', writing '$x \mapsto f(x)$' for the function mapping any relevant x to $f(x)$. Effectively, '$\mapsto$' performs the function of λ-abstraction in the meta-language. Second, we have to strip off the syntactic content of elements of the domain for type $\langle\rangle$, since these are not elements of $\mathcal{P}(W)$ itself, but pairs $\langle x, y\rangle$. We notate this by using π_1 and π_2 for the projection functions, where $\pi_1(\langle x, y\rangle) = x$ and $\pi_2(\langle x, y\rangle) = y$. For brevity, we also write $[\![\varepsilon]\!]^a_i$ for $\pi_i([\![\varepsilon]\!]^a)$ (with i being 1 or 2). Further, note that if $\langle x, y\rangle$ is an element of the domain of type $\langle\rangle$, then x is not a set of worlds, but a function from $\{\langle\rangle\}$ to sets of worlds. However, such functions naturally correspond one-to-one to sets of worlds; we thus treat them as indistinguishable, abusing notation somewhat for brevity. Similarly, we don't distinguish between x and the singleton sequence $\langle x\rangle$.

For the interpretation of predications, we obtain the logical content by applying the logical content of the predicate to the arguments. We obtain the syntactic content by concatenating the syntactic contents of the predicate and arguments; the constraints on the coarsening guarantee that this is well defined. Finally, the logical content of a λ-abstract is the function on the relevant domains mapping every sequence to the logical content of the complement clause, on the assignment function interpreting the bound variables using the relevant sequence. To notate this, we write $a[\bar{o}/\bar{x}]$ for the function which differs from a only in mapping any x_i among $\bar{x}$ to o_i. The syntactic content of a λ-abstract binding variables $\bar{x}$ is obtained by prefixing the syntactic content of the complement clause with $\lambda\bar{x}$. However, here, we need to factor out the interpretation of the bound variables $\bar{x}$ by the assignment function. We do so by interpreting the complement clause relative to an assignment function which enforces the syntactic content of any of these variables to be just the equivalence class of this variable. To notate this, we write $a[\bar{x}]$ for the function which differs from a only in mapping any x_i among $\bar{x}$ to $\langle\pi_1(a(x_i)), [x_i]_\sim\rangle$.

Truth can be defined as containing the designated actual world @, and validity can be defined as truth under every assignment function.

Definition 3.4 Given a CSM $\mathfrak{M} = \langle W, \sim, I, @\rangle$, define $[\![\cdot]\!]^{\cdot}$ to be the function mapping every expression ε of type t and assignment function a to $[\![\varepsilon]\!]^a \in D_t$ such that:

$$[\![x]\!]^a = a(x), \text{ for every variable } x$$

$$
\begin{aligned}
\llbracket c \rrbracket^a &= \quad \langle I(c), [c]_\sim \rangle, \text{ for every non-logical constant } c \\
\llbracket \to \rrbracket^a &= \quad \langle \langle d, e \rangle \mapsto W \backslash \pi_1(d) \cup \pi_1(e), [\to]_\sim \rangle \\
\llbracket \forall_t \rrbracket^a &= \quad \langle d \mapsto \bigcap_{e \in D_t} \pi_1(d)(e), [\forall_t]_\sim \rangle \\
\llbracket \varepsilon \eta_1 \ldots \eta_n \rrbracket^a &= \quad \langle \llbracket \varepsilon \rrbracket_1^a (\langle \llbracket \eta_1 \rrbracket^a, \ldots, \llbracket \eta_n \rrbracket^a \rangle), \llbracket \varepsilon \rrbracket_2^a + \llbracket \eta_1 \rrbracket_2^a + \cdots + \llbracket \eta_n \rrbracket_2^a \rangle \\
\llbracket \lambda \bar{x} \varphi \rrbracket^a &= \quad \langle \bar{o} \mapsto \llbracket \varphi \rrbracket_1^{a[\bar{o}/\bar{x}]}, \lambda \bar{x} + \llbracket \varphi \rrbracket_2^{a[\bar{x}]} \rangle
\end{aligned}
$$

A formula φ is *true* relative to $\mathfrak{M}$ and an assignment function a, written $\mathfrak{M}, a \vDash \varphi$, if $@ \in \llbracket \varphi \rrbracket_1^a$. φ is *valid in* $\mathfrak{M}$ if $\mathfrak{M}, a \vDash \varphi$ for every assignment function a.

As usual, given soundness of $\vdash$ with respect to CSMs, i.e., the fact that every theorem of $\vdash$ is valid in every CSM, we can establish the consistency of some formulas by constructing a CSM in which they are valid: if $\mathfrak{M}$ validates some formulas, then it validates their universal closures; thus by soundness, $\mathfrak{M}$ witnesses that $\bot$ is not entailed by the universal closures of the formulas, as required for consistency. And soundness of $\vdash$ with respect to CSMs can be established along the usual lines; the details can be found in Appendix A.1, with soundness stated in Lemma A.5.

It is worth noting that the definition of CSMs could be varied in various ways, without affecting their usefulness for our consistency proofs. For example, we could generalize them by replacing $\mathcal{P}(W)$ by a complete Boolean algebra and @ by an ultrafilter of this algebra; cf. Gallin [13, ch. 4]. We could also make them more restrictive by requiring W to be a singleton $\{@\}$, which corresponds to using a two-element Boolean algebra. We opt for possible worlds models here since they illustrate the flexibility of the construction without requiring any appeals to abstract algebraic concepts.

3.2 Closed Structure

We now show that there are CSMs which validate CLOSED STRUCTURE. Thus, assuming the truth of $\xi \bar{\varepsilon} = \zeta \bar{\eta}$ in a CSM, we aim to show the truth of $\xi = \zeta$ and $\varepsilon_i = \eta_i$. We first show that $\xi \bar{\varepsilon} \sim \zeta \bar{\eta}$, by establishing that the syntactic content of any closed term is its equivalence class under $\sim$. That is, we show that in CSMs, $\llbracket \varepsilon \rrbracket_2 = [\varepsilon]_\sim$ for every closed term ε; the proof, by induction on the complexity of ε, can be found in Appendix A.2, under Lemma A.6. From $\xi \bar{\varepsilon} \sim \zeta \bar{\eta}$, we aim to conclude that $\xi \sim \zeta$ and $\varepsilon_i \sim \eta_i$; we therefore restrict attention to coarsenings supporting this inference, which we call *structural*. From $\xi \sim \zeta$ and $\varepsilon_i \sim \eta_i$, we aim to conclude the truth of $\xi = \zeta$ and $\varepsilon_i = \eta_i$, so we restrict ourselves to certain CSMs guaranteeing this, which we call *safe*. This line of argument shows that CLOSED STRUCTURE is valid on any safe CSM based on a structural coarsening. The consistency of CLOSED STRUCTURE can now be established by producing a safe CSM based on a structural coarsening. And this is easily done, as identity – the trivial coarsening relating every term just to itself – is structural and such that any CSM based on it is safe.

Definition 3.5 A coarsening $\sim$ is *structural* if, in general,

if $\xi \bar{\varepsilon} \sim \zeta \bar{\eta}$, then $\xi \sim \zeta$ and $\varepsilon_i \sim \eta_i$.

A CSM is *safe* if $[\![\varepsilon]\!]^a = [\![\eta]\!]^a$ whenever $\varepsilon \sim \eta$. A coarsening is *safe* if every CSM based on it is safe.

Proposition 3.6 *Any safe* CSM *based on a structural coarsening validates* CLOSED STRUCTURE.

Proof Let $\mathfrak{M}$ be safe and based on a structural coarsening $\sim$. Assume $\mathfrak{M}, a \vDash \xi\bar{\varepsilon} = \zeta\bar{\eta}$, for closed $\xi, \zeta, \bar{\varepsilon}, \bar{\eta}$. Then $[\![\xi\bar{\varepsilon}]\!]_2 = [\![\zeta\bar{\eta}]\!]_2$, so with Lemma A.6, $[\xi\bar{\varepsilon}]_\sim = [\zeta\bar{\eta}]_\sim$, i.e., $\xi\bar{\varepsilon} \sim \zeta\bar{\eta}$. Since $\sim$ is structural, $\xi \sim \zeta$ and $\varepsilon_i \sim \eta_i$. And as $\mathfrak{M}$ is safe, $[\![\xi]\!]^a = [\![\zeta]\!]^a$ and $[\![\varepsilon_i]\!]^a = [\![\eta_i]\!]^a$. Thus $\mathfrak{M}, a \vDash \xi = \zeta \wedge \varepsilon_i = \eta_i$, as required. $\square$

Theorem 3.7 CLOSED STRUCTURE *is consistent in* $\vdash$.

Proof Identity is trivially a safe and structural coarsening. So by Proposition 3.6, any CSM based on identity validates CLOSED STRUCTURE. Consistency follows with Proposition A.5. $\square$

3.3 First Strengthening

Initially, we characterized structured views of propositions as requiring that for sentences to express the same proposition, they must, first, have the same syntactic structure and, second, be such that corresponding terms express the same entities. So far, we have targeted the second of these informal requirements: In L, every sentence is either syntactically simple, or an application of the form $\xi\bar{\varepsilon}$. CLOSED STRUCTURE states that in the case of sentences $\xi\bar{\varepsilon}$ and $\zeta\bar{\eta}$ which are of the same structure in the sense of being applications of the same type, if the propositions expressed are the same then the components are the same. However, this does not address the first informal aspect of structured views of propositions, that sentences expressing the same propositions have to have the same structure. It is therefore worth verifying that CSMs also model this aspect, at least for closed terms. That is, we would minimally like to verify that for closed terms, $\xi\bar{\varepsilon} = \varphi$ can only be true if φ is an application $\zeta\bar{\eta}$ with ζ the same type as ξ.

Using the models used in the proof of Theorem 3.7, this is easily done: Choosing identity as the coarsening, $\xi\bar{\varepsilon} = \varphi$ can only be true if $[\![\xi\bar{\varepsilon}]\!]_2 = [\![\varphi]\!]_2$, which by Lemma A.6, entails $\xi\bar{\varepsilon} \sim \varphi$, i.e., $\xi\bar{\varepsilon} = \varphi$. More generally, we can show that for any safe CSM based on a coarsening $\sim$, an equation $\varepsilon = \eta$ involving two closed terms is true just in case $\varepsilon \sim \eta$. This is captured by the following two principles:

DISTINCTNESS$_\sim$: $\varepsilon \neq \eta$ (where ε, η closed and $\varepsilon \nsim \eta$)
IDENTITY$_\sim$: $\varepsilon = \eta$ (where ε, η closed and $\varepsilon \sim \eta$)

Proposition 3.8 *For every coarsening* $\sim$, *every instance of* DISTINCTNESS$_\sim$ *and* IDENTITY$_\sim$ *is valid in any safe* CSM *based on* $\sim$.

Proof For DISTINCTNESS$_\sim$, consider any CSM based on $\sim$ and closed expressions $\varepsilon \nsim \eta$. Then $[\varepsilon]_\sim \neq [\eta]_\sim$, so by Lemma A.6, $[\![\varepsilon]\!] \neq [\![\eta]\!]$, whence $\varepsilon \neq \eta$ is valid. IDENTITY$_\sim$ is immediate by safety. $\square$

Since the models we have constructed so far are based on $\sim$ being identity, they show that CLOSED STRUCTURE is consistent with DISTINCTNESS$_=$ and IDENTITY$_=$. These principles capture the view that closed terms are co-denoting if and only if they are the same. This corresponds to a very fine-grained notion of structure, where the structure of an expression is identified with the expression itself. This may in fact be considered too fine-grained, since it means that sentences differing only in alphabetic variation of bound variables will have distinct structure, and so will express distinct propositions. For example, the models we have constructed so far validate the extreme principle that $\forall p \; p \neq \forall q \; q$. This raises the question whether CLOSED STRUCTURE is consistent with versions of DISTINCTNESS$_\sim$ and IDENTITY$_\sim$ which count alphabetic variation in bound variables as preserving structure, and so require $\forall p \; p = \forall q \; q$ to be true, rather than false. For the remainder of this section, we explore some such more coarse-grained views.

The idea that expressions which differ merely in alphabetic variation of bound variables express the same entity is one of the three main principles of identification in standard systems governing λ, usually labeled with the Greek letter α. The second is labeled with β, and strengthens our principles $e\beta 1$ and $e\beta 2$ from a material equivalence to an identity. The third is labeled with η, and states that for any higher type, a term ε expresses the same as the result of applying it to variables $\bar{x}$ not free in ε and then abstracting them away again using λ. In the present setting, we can formulate three corresponding principles in the form of schematic identities:

$\alpha :$ $(\lambda \bar{x}.\varphi) = (\lambda \bar{y}.\varphi[\bar{y}/\bar{x}])$
$\beta :$ $(\lambda \bar{x}.\varphi)\bar{x} = \varphi$
$\eta :$ $(\lambda \bar{x}.\varepsilon \bar{x}) = \varepsilon$ ($\bar{x}$ not free in ε)

These common principles give rise to a natural question: Which of them are consistent with CLOSED STRUCTURE? We will show that while β is incompatible with CLOSED STRUCTURE, both α and η are jointly satisfiable with CLOSED STRUCTURE. The first of these follows by a well-known line of argument, a version of which can be found in Dorr [10, p. 58]:

Proposition 3.9 CLOSED STRUCTURE *and β are jointly inconsistent in* $\vdash$.

Proof By β, $(\lambda p.p)\bot = \bot = (\lambda p.\bot)\bot$. So with CLOSED STRUCTURE, $(\lambda p.p) = (\lambda p.\bot)$, whence $(\lambda p.p)\top = (\lambda p.\bot)\top$. So with β again, $\top = \bot$, from which $\top \leftrightarrow \bot$ and so $\bot$ is derivable. $\qquad \square$

We now construct CSMs of α, η and CLOSED STRUCTURE. We first use α and η to define a coarsening $\sim^{\alpha,\eta}$. To do so, we show more generally how any set of equations T naturally gives rise to a corresponding coarsening $\sim^T$, namely the coarsening which relates terms if they can be obtained from each other by successive replacements of sub-terms corresponding to equations in T.

Definition 3.10 For any consistent set of equations T, let:

$\varepsilon \sim^{iT} \eta$ iff for some $\chi = \kappa \in T$ and term ϑ, $\vartheta(\chi) = \varepsilon$ and $\vartheta(\kappa) = \eta$.
$\varepsilon \sim^{sT} \eta$ iff $\varepsilon \sim^{iT} \eta$ or $\eta \sim^{iT} \varepsilon$.

$\varepsilon \sim^T \eta$ iff there are expressions $\bar{\chi}$ such that $\varepsilon = \chi_1 \sim^{sT} \cdots \sim^{sT} \chi_n = \eta$.

We can show that in this way, α, η give rise to a coarsening $\sim^{\alpha,\eta}$ which is structural and safe. This is shown in Appendix A.2 by Lemmas A.9–A.11. These results ensure, by previous results, that any CSM based on $\sim^{\alpha,\eta}$ validates CLOSED STRUCTURE. Moreover, the safety of $\sim^{\alpha,\eta}$ ensures that α and η are validated, since $\sim^{\alpha,\eta}$ will relate any pair of terms corresponding to an instance of α or η.

Theorem 3.11 CLOSED STRUCTURE*, α and η are jointly consistent in $\vdash$.*

Proof By Lemmas A.9–A.11, $\sim^{\alpha,\eta}$ is a safe and structural coarsening, so by Proposition 3.6, any CSM based on $\sim^{\alpha,\eta}$ validates CLOSED STRUCTURE. Likewise, since $\sim^{\alpha,\eta}$ is safe, it follows that any such model validates α and η. Consistency follows with Proposition A.5. □

Another very restrictive aspect of the models constructed so far concerns the identification of constants. Let c and d be distinct non-logical constants of the same type. Since $c \not\sim^{\alpha,\eta} d$, it follows with Proposition 3.8 that CSMs based on identity or $\sim^{\alpha,\eta}$ validate $c \neq d$. But plausibly, some constants may co-denote, so it is instructive to verify that CLOSED STRUCTURE is consistent with identifications of the form $c = d$. This is easily done: Consider any set E of equations of the form $c = d$, with c, d constants. It is easily seen that $\sim^E$ is a structural coarsening such that a CSM based on it is safe if $I(c) = I(d)$ whenever $c = d \in E$. Such models evidently exist. By Proposition 3.6, they validate CLOSED STRUCTURE, and by safety, they validate all equations of E.

The models constructed here might also be of interest for other theories of intermediate fineness of grain of propositions. For example, consider a view which endorses α, β and η. As noted, this is inconsistent with CLOSED STRUCTURE. But one might still want to endorse as fine-grained an individuation of propositions as possible, given these constraints. That is, one might want to endorse IDENTITY$_{\sim\alpha,\beta,\eta}$ and DISTINCTNESS$_{\sim\alpha,\beta,\eta}$. Analogous to Lemma A.11, we can show that $\sim^{\alpha,\beta,\eta}$ is safe, and with this, the joint consistency of IDENTITY$_{\sim\alpha,\beta,\eta}$ and DISTINCTNESS$_{\sim\alpha,\beta,\eta}$ follows from Proposition 3.8.

3.4 Second Strengthening

CLOSED STRUCTURE captures, for closed terms, the idea that if sentences with the same structure express the same propositions then corresponding components are co-denoting. However, it does so in a rather shallow way: all it captures of the structure of a complex sentence, which must be an application, is what the predicate is, and what the arguments are – it does not capture the syntactic structure of these components themselves. To illustrate this, consider an equation of the form $(\lambda x.Rxx)a = (\lambda x.Sxx)a$. According to CLOSED STRUCTURE, we can infer from this that $(\lambda x.Rxx) = (\lambda x.Sxx)$. But there is no obvious way of inferring that $R = S$. However, on a natural construal of the idea of structured propositions, this would be

required as well. We will now show how CLOSED STRUCTURE can be strengthened to capture this idea, and prove that this is consistent as well.

The identity $(\lambda x.Rxx)a = (\lambda x.Sxx)a$ is of the form $\xi(\varepsilon) = \xi(\eta)$, with ξ being $(\lambda x.Yxx)a$, and ε and η being R and S, respectively. A couple of generalizations suggest themselves naturally. First, admitting multiple parameters, we might state that for any closed formulas, if $\xi[\bar{\varepsilon}/\bar{x}] = \xi[\bar{\eta}/\bar{x}]$ and all $\bar{x}$ have free occurrences in ξ, then $\varepsilon_i = \eta_i$. Second, it seems plausible to hold that an intermediate step may be inserted in the first motivating example: if $(\lambda x.Rxx)a = (\lambda x.Sxx)a$, then $(\lambda x.Rxx) = (\lambda x.Sxx)$, which in turn is only true if $R = S$. Thus a naturally general formulation applies to all non-vacuous substitutions of the form $\xi[\bar{\varepsilon}/\bar{x}] = \xi[\bar{\eta}/\bar{x}]$, whatever type ξ may be. Finally, it turns out that for consistency, we need not require $\xi[\bar{\varepsilon}/\bar{x}]$ and $\xi[\bar{\eta}/\bar{x}]$ to be closed; it suffices to require that the parameters $\bar{\varepsilon}$ and $\bar{\eta}$ are closed. This leads us to the following strengthened form of CLOSED STRUCTURE:

STRONG CLOSED STRUCTURE: $\xi[\bar{\varepsilon}/\bar{x}] = \xi[\bar{\eta}/\bar{x}] \to \varepsilon_i = \eta_i$ ($\bar{\varepsilon}, \bar{\eta}$ closed, all $\bar{x}$ free in ξ)

The instances of CLOSED STRUCTURE can be obtained from the special cases in which ξ is of the form $X\bar{y}$. The requirement of all $\bar{x}$ to have free occurrences in ξ is needed to ensure that STRONG CLOSED STRUCTURE is not trivially inconsistent. Without this restriction, various obvious inconsistencies, such as $\bot = \bot \to p = q$, would count as instances.

We build models of this strengthened principle by imposing additional constraints on coarsenings. What is needed is that different replacements of the same variables in one term only produce equivalent outcomes if the replacements are equivalent. Call this being *strongly structural*:

Definition 3.12 A coarsening $\sim$ is *strongly structural* if, in general,

if $\xi[\bar{\varepsilon}/\bar{x}] \sim \xi[\bar{\eta}/\bar{x}]$ and all $\bar{x}$ are free in ξ, then $\varepsilon_i \sim \eta_i$.

Using a variation on the proof of Proposition 3.6, we can show that any safe CSM based on a strongly structural coarsening validates STRONG CLOSED STRUCTURE; the details can be found in Appendix A.2 under Proposition A.12. With this, we can show STRONG CLOSED STRUCTURE to be consistent:

Theorem 3.13 STRONG CLOSED STRUCTURE *is consistent in* $\vdash$.

Proof Using Proposition A.12, we can reason as for Theorem 3.7; thus, it suffices to show that identity is strongly structural. This can be established by a straightforward induction on the complexity of ξ. □

We leave open whether this strengthening may be combined with that of the previous section, i.e., whether STRONG CLOSED STRUCTURE is consistent with α and η. We conjecture that the answer is positive, and that this can be established by the proof method used here. It would suffice to show that $\sim^{\alpha,\eta}$ is strongly structural, but doing so is not completely straightforward. For example, a natural attempt at a proof would employ an induction on the complexity of terms which requires that if abstractions

$\lambda \bar{x}.\varphi$ and $\lambda \bar{x}.\psi$ are equivalent then φ and ψ are equivalent. But this is not the case in general, as $\lambda x.(\lambda x.Rxx)x \sim^{\alpha,\eta} \lambda x.Rxx$, but $(\lambda x.Rxx)x \not\sim^{\alpha,\eta} Rxx$. Due to complications like this, proving that $\sim^{\alpha,\eta}$ is strongly structural requires a more detailed analysis of this coarsening than we will provide here.

4 Inconsistency

The consistency of CLOSED STRUCTURE (with modest principles of identification, like α and η) may give proponents of structured propositions hope that any failures of STRUCTURE may be restricted to cases which will in an important sense never be encountered. We now show that these hopes are in vain, as CLOSED STRUCTURE is inconsistent in a modest extension of the type theory we have used until now, namely the one which adds a way of talking *plurally* about propositions.

4.1 Plural Propositional Quantification

Plural quantification roughly captures what is expressed by English phrases such as 'Some things are such that ...', which may be followed by statements which attribute something to the relevant things collectively, rather than individually. Claims of this form need not say anything about sets, classes, groups, mereological sums, or any similar entities, so they are not naturally captured by singular quantification over such entities. A more direct way of regimenting such quantification extends first-order logic by adding plural variables. Following Burgess and Rosen [8], these are often written as doubled first-order variables $xx, yy, zz, \ldots$. Such variables may be bound by quantifiers $\exists$ and $\forall$, and used in statements such as $x \prec yy$, which states that x is one of the yy.

Here, we will make use of such plural quantification over propositions. In order to fit this into our type system, we add a new type into our type hierarchy, written $[\langle\rangle]$, the type of plural propositional terms. Keeping the recursive rules of generating types, this gives rise to further new types, such as the type $\langle[\langle\rangle]\rangle$ of unary predicates applying to plural propositional terms.

Definition 4.1 Let the set T' of types be the smallest set such that $[\langle\rangle] \in T'$, and $\langle t_1, \ldots, t_n \rangle \in T'$ whenever $t_1, \ldots, t_n \in T'$.

As in the case of L, we construct a language L' of expressions of each type based on a choice V of variables of each type and a choice C of constants of each type. Again, C is assumed to contain logical constants $\rightarrow$ and $\forall_t$ for each type t (now of T'). Additionally, we now assume that a logical constant $\prec$ of type $\langle\langle\rangle, [\langle\rangle]\rangle$ is included. Using infix notation for $\prec$, we can use $p \prec qq$ to state that p is one of, or among, qq. Finally, we add an analog of λ for plural terms: $\pi p.\varphi$ will be used to stand for the propositions (plurally) which satisfy φ. An analogous operator for forming plural terms in the standard setting of plural extensions of first-order logic can be found in Oliver and Smiley [30, ch. 8]; see also Whitehead and Russell [40, p. 30] for a plural reading of their notation for classes. We adapt earlier conventions on

formulas, doubling the relevant letters; for example, we let doubled lowercase Greek letters like $\varphi\varphi$ range over plural propositional terms.

Definition 4.2 Define the family of sets L', based on V' and C', inductively using ATOMIC, APPLICATION, ABSTRACTION for the types of T' and:

PLURAL PROPOSITIONAL ABSTRACTION: If p is a propositional variable and φ is a formula, then $(\pi p.\varphi)$ is an expression of type $[\langle\rangle]$.

It is worth noting that we can think of the present plural extension as a fragment of a more comprehensive type theory, according to which any types $t_1, \ldots, t_n$ give rise to two types: the type $\langle t_1, \ldots, t_n \rangle$ of an n-ary relational expression taking as arguments n expressions of types $t_1, \ldots, t_n$, respectively, and the type $[t_1, \ldots, t_n]$ of an n-ary (generalized) plural term $\eta\eta$ which can be used in formulas like $\varepsilon_1 \ldots \varepsilon_n \prec \eta\eta$, where $\varepsilon_1, \ldots, \varepsilon_n$ are of types $t_1, \ldots, t_n$, respectively. Closely related type theories are discussed by Myhill [29] and Fine [11]. Since such a rich extension of the type theory will likely be considered more controversial than just adding plural propositional resources, we restrict ourselves to the latter. Conversely, one may think of plural propositional quantifiers as quantifiers ranging over properties of propositions restricted by a distinguished condition, chosen so as to validate the axiomatic principles given below. However, we consider the best motivation for the existence of such a condition to proceed via plural propositional quantification, so we will simply consider plural propositional quantification as *sui generis*.

We continue to use the earlier abbreviating conventions, and extend them to the new types whenever this is defined. This includes $\exists_{[\langle\rangle]}$ and $=_{[\langle\rangle]}$, but not $\sqsubseteq_{[\langle\rangle]}$ and $\equiv_{[\langle\rangle]}$, since plural propositional terms cannot be applied to any arguments. Instead, we can naturally define the latter as follows:

$$\varphi\varphi \sqsubseteq_{[\langle\rangle]} \psi\psi := \forall p(p \prec \varphi\varphi \to p \prec \psi\psi)$$
$$\varphi\varphi \equiv_{[\langle\rangle]} \psi\psi := \varphi\varphi \sqsubseteq \psi\psi \wedge \psi\psi \sqsubseteq \varphi\varphi$$

Here, the definition of '$\sqsubseteq$' is intended to be read such that p is the first propositional variable (according to some fixed ordering) not free in $\varphi\varphi$ or $\psi\psi$. We extend notation about negation, e.g., writing $\not\prec$ for a negated application of $\prec$, and introduce shortening notation for quantifiers as familiar from set theory; e.g., '$\forall p \prec qq(\ldots)$' abbreviates '$\forall p(p \prec qq \to \ldots)$' and so on.

$\vdash$ is extended by instances of the quantificational principles for the new types, and an analog of the principles governing λ for π. Finally, we add a principle of extensionality, according to which pluralities are individuated extensionally: if $\varphi\varphi$ and $\psi\psi$ have the same propositions among them, then they are the same:

Definition 4.3 Let $\vdash'$ be the proof system in L' with the principles of $\vdash$ (which may be instantiated using expressions of any types in T'), and

$$e\beta1': \quad p \prec (\pi p.\varphi) \to \varphi \qquad\qquad e\beta2': \quad \varphi \to p \prec (\pi p.\varphi)$$
$$\text{EXT}: \quad \varphi\varphi \equiv \psi\psi \to \varphi\varphi = \psi\psi$$

EXT is the crucial axiomatic principle in which plural propositional terms differ from the resources of L'. We will therefore record any appeals to it in deductive

arguments explicitly. EXT is standardly included in axiomatizations of plural quantification; see Burgess [7, p. 197] and Linnebo [26]. The axiom is clearly motivated on a common way of thinking about pluralities, according to which they are nothing over and above the things they comprise; see Roberts [34]. Going beyond extensionality, it is sometimes even said that for pluralities to be coextensive *is* for them to be identical; see Williamson [42, p. 700] and Oliver and Smiley [30, p. 109]. This motivates a strengthening of EXT where the conditional is replaced by identity, but no such stronger principle will be required here. Finally, the notion of plural propositional quantification is introduced here in analogy with certain uses of English phrases, but not necessarily by requiring it to match English usage precisely. (We discuss this point in more detail in Section 5.2 below.) As part of our introduction of these new formal quantifiers, we may well simply stipulate their meaning in such a way that EXT comes out true.

As above, we can derive the principles of elementary quantification theory from the principles of $\vdash'$, including the following plural comprehension principle:

$$\vdash' \exists pp \forall q (q \prec pp \leftrightarrow \varphi) \ (pp \text{ not free in } \varphi)$$

Note that this includes the case in which φ is a contradiction; thus according to $\vdash'$, there are pp among which there is no proposition. We call this the *empty* plurality. This inclusion may be at odds with some uses of English phrases which motivate the introduction of plurals, since such uses seem to carry existential import: it seems that if some propositions jointly satisfy a certain condition, then there is at least one proposition among them. Yet, the inclusion of the empty plurality is unproblematic in the present context, for two reasons. First, following Burgess and Rosen [8] and Burgess [7], we may stipulate that plural quantification captures what is expressed by phrases like 'there are zero or more . . . ', which explicitly rule out such existential import. Second, as noted already in connection with EXT, we make no claim that plural propositional quantifiers correspond directly to any expressions of English.

$\vdash'$ is easily seen to be consistent: As noted at the beginning of Section 3, simple models for $\vdash$ can be constructed by starting from a set of worlds W, and using as the domain of any type $\bar{t}$ the set of functions from sequences of elements of domains of types $t_1, \ldots, t_n$ to $\mathcal{P}(W)$. Such models are straightforwardly extended to interpret L', in a way which guarantees that the theorems of $\vdash'$ are validated: In brief, let the domain of type $[\langle\rangle]$ be the powerset of the domain of type $\langle\rangle$; $\prec$ is interpreted as the function mapping any arguments to W or $\emptyset$, depending on whether the first argument is a member of the second; and $\pi p.\varphi$ is interpreted as the set of propositions which satisfy φ.

4.2 Closed Structure and Plurals

In L', there are additional instances of the schematic principle CLOSED STRUCTURE involving plural propositional terms, such as the following:

$$\bot \prec (\pi p.\varphi) = \bot \prec (\pi p.\psi) \rightarrow (\pi p.\varphi) = (\pi p.\psi)$$

where no variable other than p is free in φ or ψ. Further, $\vdash'$ extends $\vdash$ in a nontrivial way, by adding EXT. This raises the question whether the schema CLOSED

STRUCTURE is consistent in $\vdash'$ as well. We show that this is not the case, by establishing the following result:

Theorem 4.4 *For every formula χ with a free plural propositional variable, there are closed plural terms $\varphi\varphi$ and $\psi\psi$ such that*

$$\vdash' \chi(\varphi\varphi) = \chi(\psi\psi) \wedge \varphi\varphi \not\equiv \psi\psi.$$

The remainder of this section establishes this result, along the lines of Burali-Forti's so-called paradox. Sketching the idea behind the proof, consider, for a given formula χ with a free variable pp, the sequence of propositions defined recursively by letting every element be $\chi(pp)$, where pp is the plurality of propositions preceding it in the sequence. The first element of this sequence is $\chi(\pi p.\bot)$ (since nothing precedes it); the second is $\chi(\pi p.p = \chi(\pi p.\bot))$; and so on. We first show how to capture this sequence in L' using an expression M of type $\langle\langle\rangle, \langle\rangle\rangle$. We then consider the proposition ∞, defined as $\chi(\varphi\varphi)$, where $\varphi\varphi$ is the field of M (the propositions ordered by M). It turns out that ∞ must be an element of the sequence M as well; indeed, it must be the final element. But this means that ∞ is identical to $\chi(\psi\psi)$, where $\psi\psi$ are the propositions among $\varphi\varphi$ excluding ∞ itself, since otherwise, ∞ would not be an element of the sequence M. Therefore, $\varphi\varphi$ and $\psi\psi$ are the required plural terms, as $\chi(\varphi\varphi) = \infty = \chi(\psi\psi)$ while $\varphi\varphi \not\equiv \psi\psi$.

To make this idea precise, we first introduce the required order-theoretic notions in L'. We use R and S as variables of type $\langle\langle\rangle, \langle\rangle\rangle$, and Θ for an arbitrary expression of the same type. Let the *field* of a binary propositional relation be the propositions it relates – one way or the other – to some proposition:

$$\mathsf{ff}(\Theta) := \pi p.\exists q(\Theta pq \vee \Theta qp)$$

Let a binary propositional relation be a *well-order* if it is reflexive, anti-symmetric, transitive, and well-founded (on the propositions in its field):

$$\mathsf{W} := \lambda R.$$
$\forall p \prec \mathsf{ff}(R)(Rpp) \wedge$	(reflexive)
$\forall p \forall q \forall r(Rpq \wedge Rqr \rightarrow Rpr) \wedge$	(transitive)
$\forall p \forall q(Rpq \wedge Rqp \rightarrow p = q) \wedge$	(antisymmetric)
$\forall qq \sqsubseteq \mathsf{ff}(R)(\exists p(p \prec qq) \rightarrow \exists p \prec qq \forall q \prec qq(Rpq))$	(well-founded)

Note that if R is a well-order, then it is *total*:

$$\forall p \prec \mathsf{ff}(R)\forall q \prec \mathsf{ff}(R)(Rpq \vee Rqp).$$

Next, we define the restriction of a well-order Θ to the strict predecessors of φ:

$$\Theta|_\varphi := \lambda pq.\Theta pq \wedge \Theta q\varphi \wedge q \neq \varphi$$

Since we have no extensionality principle for relations, relations may behave in unexpected ways, compared to the familiar extensional treatment of relations in set theory. We therefore spell deductive proofs out in somewhat fine detail.

With these general notions on the table, pick a formula χ with a free plural propositional variable. The remaining development in this section will be relative to this

choice of χ, but we don't record this relativity explicitly, to simplify notation. We aim to define the sequence M as sketched above, and prove that it is a well-order. This can be done by first identifying the initial sub-sequences of the intended sequence M, and then defining M as their union. These sub-sequences can be identified as the well-orders R which satisfy the condition that every proposition in the field of R is the result of applying χ to its strict predecessors in R. We call them *chains*. In symbols, this can be defined as follows:

$$\mathsf{C} := \lambda R.\mathsf{W} R \wedge \forall p \prec \mathsf{ff}(R)(p = \chi(\mathsf{ff}(R|_p)))$$

With this, we can define M as relating p to q just in case p is related to q by some chain, and ∞ as the result of applying condition χ to the field of M:

$$\mathsf{M} := \lambda pq.\exists R(\mathsf{C}R \wedge Rpq)$$
$$\infty := \chi(\mathsf{ff}(\mathsf{M}))$$

We now need to establish two claims: First, that M is a chain, and second, that ∞ is in the field of M. The proofs are relegated to Appendix B. The first is established in Lemma B.2, using the fact that chains are totally ordered by the relation of being an initial sub-chain of one another, which we show in Lemma B.1. The second is established in Lemma B.3, by noting that if ∞ were not in the field of M, then appending ∞ to the end of M would produce a chain M^+ containing ∞, whence ∞ would be in the field of M after all. With these lemmas, we can prove the main result of this section:

Proof of Theorem 4.4 For any formula χ with a free plural propositional variable, we have defined closed plural terms $\mathsf{ff}(\mathsf{M})$ and $\mathsf{ff}(\mathsf{M}|_\infty)$. We show that these witness the claim, by establishing:

(a) $\vdash' \chi(\mathsf{ff}(\mathsf{M})) = \chi(\mathsf{ff}(\mathsf{M}|_\infty))$
(b) $\vdash' \mathsf{ff}(\mathsf{M}) \not\equiv \mathsf{ff}(\mathsf{M}|_\infty)$

For (a), note that by Lemmas B.2 and B.3, M is a chain and $\infty \prec \mathsf{ff}(\mathsf{M})$. Thus $\infty = \chi(\mathsf{ff}(\mathsf{M})|_\infty)$. So by definition of ∞, $\chi(\mathsf{ff}(\mathsf{M})) = \infty = \chi(\mathsf{ff}(\mathsf{M})|_\infty)$.

For (b), $\infty \prec \mathsf{ff}(\mathsf{M})$ by Lemma B.3, but $\infty \not\prec \mathsf{ff}(\mathsf{M}|_\infty)$ by construction of $\mathsf{M}|_\infty$. Thus ∞ witnesses $\mathsf{ff}(\mathsf{M}) \not\equiv \mathsf{ff}(\mathsf{M}|_\infty)$. $\qquad\square$

Theorem 4.4 immediately entails:

Corollary 4.5 CLOSED STRUCTURE *is inconsistent in* $\vdash'$.

Furthermore, Theorem 4.4 can be strengthened in two non-trivial ways, by noting that the witnessing plural parameters may be chosen to be fields of well-orders, one of which includes the other. All that needs to be shown is that $\mathsf{M}|_\infty$ is a well-order, which is routine to derive from the fact that M is a well-order. Thus, we obtain:

Corollary 4.6 *For every formula χ with a free plural propositional variable, there are closed terms Θ and Θ' such that*

$$\vdash' \mathsf{W}\Theta \wedge \mathsf{W}\Theta' \wedge \mathsf{ff}(\Theta) \sqsubseteq \mathsf{ff}(\Theta') \wedge \chi(\mathsf{ff}(\Theta)) = \chi(\mathsf{ff}(\Theta')) \wedge \mathsf{ff}(\Theta) \not\equiv \mathsf{ff}(\Theta')$$

This concludes the main technical discussion in this paper. Two aspects are worth highlighting, before we move on to discuss their philosophical significance. First, the inconsistency of CLOSED STRUCTURE in $\vdash'$ entails the inconsistency of STRONG CLOSED STRUCTURE in $\vdash'$. But the inconsistency of the latter is in fact already an immediate consequence of EXT: according to STRONG CLOSED STRUCTURE, $(\pi p.\bot \wedge \top) \neq (\pi p.\bot \wedge \bot)$, which contradicts EXT.

Second, we could have used a functional type theory instead of a relational one. However, since in the present dialectical context, we cannot assume α, β and η, we cannot – as is often possible when these identities are endorsed; see Dorr [10, Appendix] – treat functional and relational type theories as mere syntactic variants. And some subtleties arise in the functional setting, since quantification over binary relations plays an important role in the results of this section. Functional type theory only provides quantification over relations by simulating relations as higher-order functions, via a process known as Currying or Schönfinkeling, according to which a binary relation R is treated as a function mapping a first argument x to a function mapping a second argument y to $(Rx)y$. Above, we relied on a version of $e\beta 1/2$ according to which $(\lambda xy.\varphi)\varepsilon\eta \leftrightarrow \varphi[\varepsilon/x, \eta/y]$. The corresponding equivalence in a functional setting may be written as $(\lambda x.(\lambda y.\varphi))\varepsilon\eta \leftrightarrow \varphi[\varepsilon/x, \eta/y]$. This does not obviously follow from the most straightforward way of capturing the idea of material β equivalence in a functional setting, which is the schematic principle $(\lambda x.\varphi)\varepsilon \leftrightarrow \varphi[\varepsilon/x]$. These issues can be sidestepped if one can avail oneself of a way of pairing propositions, i.e., a formula ρ with two free propositional variables such that $\rho(\varphi, \psi) = \rho(\varphi', \psi') \leftrightarrow (\varphi = \varphi' \wedge \psi = \psi')$. In this case, quantification over binary relations among propositions can be simulated using quantification over properties of propositions. And on some views of structured propositions, such pairing contexts may be easy to come by, as even the binary Boolean connectives may provide examples. But if one retreats from STRUCTURE to CLOSED STRUCTURE, one might equally restrict the recoverability of conjuncts from conjunctions to closed sentences. All of these complications simply do not arise in the relational setting, and since we see no reason to consider relational languages more problematic than the functional ones, we have opted for a relational formulation.

5 Discussion

The results established here show that retreating from STRUCTURE to CLOSED STRUCTURE (both defined in Section 2.2) holds some initial promise: CLOSED STRUCTURE, unlike STRUCTURE, is consistent in the basic higher-order logic $\vdash$ (introduced in Section 2.1). The prospects of such a retreat are, however, limited, as this consistency result is not preserved under the addition of plural propositional resources: CLOSED STRUCTURE is inconsistent in the logic $\vdash'$ (defined in Section 4.1). In this section, we discuss the wider ramifications of these findings.

One immediate observation which can be made from these results concerns the source of the inconsistency of structured proposition views. The Russell-Myhill argument is naturally given a cardinality-theoretic gloss: it can be seen as using the fact that there are more properties of propositions than propositions to conclude that

there cannot be a distinct proposition for every property of propositions, which is inconsistent with STRUCTURE. The inconsistency of CLOSED STRUCTURE cannot be given such a cardinality-theoretic gloss: CLOSED STRUCTURE is entailed by the claim that distinct closed sentences express distinct propositions; in a countable language, it therefore does not require there to be more than countably infinitely many propositions.

In the remainder of this section, we consider how one could motivate CLOSED STRUCTURE without also being committed to STRUCTURE, scrutinize the plural propositional resources used here, and discuss some related principles motivated by considerations concerning propositional attitudes.

5.1 Restricted Structure

What does it take to endorse CLOSED STRUCTURE but not STRUCTURE? To answer this question, it is important to note that the viability of CLOSED STRUCTURE depends on the language under consideration. The main results of this paper already illustrate this, as CLOSED STRUCTURE is consistent without, but inconsistent with, plural propositional quantifiers and the π operator. But even when we restrict ourselves to the simple relational language L without such plural resources, the choice of language matters. This is because L is parametric with respect to a choice of non-logical constants.

The inconsistency of STRUCTURE shows that even in L, CLOSED STRUCTURE is only viable on certain restrictions on the interpretation of non-logical constants. This follows from the Russell-Myhill theorem (Theorem 2.4 above): As we have seen, there is some X such that $X\perp = \mathsf{R}\perp$, even though $X \neq \mathsf{R}$. So, in a language with a constant C such that $C = X$, $C\perp = \mathsf{R}\perp$, but $C \neq \mathsf{R}$. Since R is closed, this contradicts an instance of CLOSED STRUCTURE. So, even when we consider only instances of CLOSED STRUCTURE in L, we must restrict ourselves to certain well-behaved languages. This poses the question how one might motivate instances of CLOSED STRUCTURE in such a language without motivating the problematic instances as well. We start with this issue, setting aside for the moment the inconsistency of CLOSED STRUCTURE in $\vdash'$.

One option is to postulate a distinction between logical and non-logical entities, and to endorse CLOSED STRUCTURE for languages in which all constants express logical entities. Call this *logical structuralism*. On this view, a logical constant like $\rightarrow$ would express a logical entity, but likewise, a non-logical constant N might be used to stand for a logical entity such as negation. Furthermore, expressing a logical entity would plausibly be held to be closed under APPLICATION and ABSTRACTION, so that every closed expression of L would express a logical entity. Endorsing CLOSED STRUCTURE for any such language would be a way of giving voice to the idea that reality is not completely structured, but contains a purely logical core which is structured – a core of logical entities which are structured just like a language.

Logical structuralism has the advantage of starting from an intuitive conception of logicality (however vague) which serves to single out a non-trivial class of interpreted languages in which CLOSED STRUCTURE may be endorsed. However, this view is a very substantial retreat from STRUCTURE: according to logical structuralism, we

can presumably not even conclude that the proposition that the moon is red is distinct from the proposition that the moon is blue, even though this case, like many similar ones, poses no danger of leading to inconsistency. It is therefore interesting to explore the possibility of alternatives, which trade the intuitiveness of the logical/non-logical distinction for a much more inclusive class of interpreted languages.

As a starting point, a natural view would endorse CLOSED STRUCTURE for any choice of constants corresponding to a word of any extant spoken language. To avoid arbitrariness, the view should further endorse CLOSED STRUCTURE for any extension of such a language we might introduce. One might wonder if this wouldn't immediately lead to problems: why couldn't we simply introduce a new constant C for a witness to the existence of an X such that $X\bot = \mathsf{R}\bot$ and $X \neq \mathsf{R}$? But note that although on a suitably abstract conception of languages, there clearly are languages with such a constant, it is a substantial claim that we can introduce and speak such a language. It may be that successfully introducing such a constant C requires not only the existence of an X satisfying the relevant condition, but also a unique way of singling it out, since otherwise, the meaning of the new constant would not be unambiguously determined. The view would thus depend on there not being any way of uniquely singling out such an X.

This suggests that all failures of STRUCTURE may in some sense be inaccessible or ineffable, so we may call this idea *effable structuralism*. This formulation of the view also takes care of a worry we raised earlier about CLOSED STRUCTURE: In formal semantics, demonstratives and anaphora are often formalized using free variables, which are interpreted relative to a contextually determined assignment function. Motivating structured propositions via their role in the philosophy of language, it might thus seem that this requires the full strength of STRUCTURE. But demonstratives and anaphora can only be used to denote entities which can be made contextually salient, and there is no reason to think that this includes everything. Rather, if there are ineffable entities, then they should not be available as referents of demonstratives or anaphora. English sentences containing demonstratives or anaphora can thus be argued to motivate nothing stronger than *effable structuralism*.

The idea of effable structuralism can be illustrated using the idea of dark matter in physics. All of the matter which we can see is ordinary matter. However, from gravitational effects, we can infer that there must be another form of matter – dark matter. Corresponding to the matter we can see are all of the individuals, propositions, properties and relations which we can name and think of. According to effable structuralism, any instance of CLOSED STRUCTURE with constants for such ordinary entities is true. Corresponding to the indirect evidence for the existence of dark matter is the Russell-Myhill theorem, which entails that there are some counterexamples to STRUCTURE. Indeed, the analogy even extends to cardinality considerations: Dark matter is thought to make up most of the matter in the universe; correspondingly, there are set-theoretic reasons to think that violations to STRUCTURE are pervasive rather than being isolated cases (see Uzquiano [38, p. 339]).

The way we have sketched the idea of effable structuralism suggests that we can formulate the view more precisely using a metaphysical distinction using a (type-indexed) term for being *effable*. This, however, renders the view self-defeating: once we have a term for being effable, we can run the Russell-Myhill theorem with all

quantifiers restricted to effable entities, and so produce an argument for the existence of an effable counterexample to STRUCTURE. It is worth spelling this out in detail. So assume that for every type t, there is a closed term E_t of type $\langle t \rangle$ expressing the property of being effable. E_t might be a constant, or it might be a complex term. According to the envisaged position, every closed term expresses something effable:

CLOSED-EFFABLE: $\mathsf{E}\varepsilon$ (ε closed)

From this, we can derive the existence of effable counterexamples to STRUCTURE.

Proposition 5.1 *For any formula χ with a single free variable of type $\langle\langle\rangle\rangle$, there is a closed expression $\varepsilon^{\langle\langle\rangle\rangle}$ such that*

$$\text{CLOSED-EFFABLE} \vdash \mathsf{E}\varepsilon \wedge \exists X^{\langle\langle\rangle\rangle}(\mathsf{E}X \wedge (\chi(X) = \chi(\varepsilon)) \wedge X \not\equiv \varepsilon).$$

Proof Define an expression R of type $\langle\langle\rangle\rangle$ as follows:

$$\mathsf{R} := \lambda p. \neg \forall X((\mathsf{E}X \wedge \chi(X) = p) \to Xp)$$

Assume for contradiction that $\neg\mathsf{R}\chi(\mathsf{R})$. Then $\forall X((\mathsf{E}X \wedge \chi(X) = \chi(\mathsf{R})) \to X\chi(\mathsf{R}))$. Since χ has only a single free variable, R is closed, whence with CLOSED-EFFABLE, $\mathsf{E}\mathsf{R}$. So as $\chi(\mathsf{R}) = \chi(\mathsf{R})$, it follows that $\mathsf{R}\chi(\mathsf{R})$, contradicting the assumption. Thus $\mathsf{R}\chi(\mathsf{R})$, whence there is an X such that $\mathsf{E}X$, $\chi(X) = \chi(\mathsf{R})$ and $\neg X\chi(\mathsf{R})$. Since $\mathsf{R}\chi(\mathsf{R})$, $\chi(\mathsf{R})$ witnesses $X \not\equiv \mathsf{R}$. We conclude $\mathsf{E}\mathsf{R} \wedge \exists X(\mathsf{E}X \wedge (\chi(X) = \chi(\mathsf{R})) \wedge X \not\equiv \mathsf{R})$. □

This argument shows that whatever 'effable' means, it is incoherent to assume both that STRUCTURE holds for all effable entities and that every closed expression of the language expresses something effable. But those were the two core components of effable structuralism as sketched above. So, although our sketch of effable structuralism may initially have sounded compelling, it is not clear that there even is a coherent view to be entertained. Note that these challenges to articulating effable structuralism do not apply here with equal force to logical structuralism: restricting quantifiers in the Russell-Myhill theorem to logical entities only provides us with a logical counterexample to STRUCTURE if logicality is logical, but it is coherent – and perhaps even well-motivated – for a proponent of logical structuralism to deny this.

Instead of trying to articulate a view of effable structuralism using the notion of effability, one can of course simply endorse all instances of CLOSED STRUCTURE for a given language. But this provides no motivation for the instances of CLOSED STRUCTURE. It also does not provide any way of articulating the idea that CLOSED STRUCTURE holds for any language we are in a position to speak.

This second limitation can to some extent be overcome. At least, this can be done on the plausible but not uncontroversial assumption that to be able to introduce a new constant for some entity, it is necessary and sufficient to be able to single the entity out uniquely. What exactly it takes to single something out is a difficult matter, but arguably, if there is a formula φ with a single free variable x of some type t which is uniquely satisfied, then the unique satisfier can be singled out with the help of this formula. In such a case, we say that the unique satisfier is *defined by φ*. In symbols:

$$\varphi \triangleright x := \forall y^t (\varphi(y) \leftrightarrow y = x)$$

The informal sketch of effable structuralism thus motivates not just CLOSED STRUCTURE, i.e., STRUCTURE restricted to what is expressed by closed terms, but STRUCTURE restricted to what is definable. This can be captured by the following schema:

DEFINABLE STRUCTURE: $\forall \bar{x} \bar{y} (\varphi_0 \triangleright X \wedge \psi_0 \triangleright Y \wedge \varphi_1 \triangleright x_1 \wedge \cdots \wedge \varphi_n \triangleright x_n \wedge \psi_1 \triangleright y_1 \wedge \cdots \wedge \psi_n \triangleright y_n \rightarrow (X\bar{x} = Y\bar{y} \rightarrow X = Y \wedge x_i = y_i))$

DEFINABLE STRUCTURE immediately entails CLOSED STRUCTURE, since anything expressed by a closed term ε is defined by the formula $x = \varepsilon$. But DEFINABLE STRUCTURE is not entailed by CLOSED STRUCTURE, since we can show that DEFINABLE STRUCTURE is invalid in some of the models of CLOSED STRUCTURE defined above. The relevant models are those in which we have a term $\square$ expressing the property of being a proposition with trivial logical content, i.e., whose first coordinate is W, the set of all worlds. Such an operator, and similar ones expressing identity of logical content or entailment among logical contents, which could be used interchangeably, do not just arise as an artifact of the model theory, but are natural notions for advocates of structured propositions to want to appeal to. For instance, they might wish to have a way of expressing coarse logical equivalences among finely individuated propositions, of the kind that holds between, say, a proposition p and its self-conjunction $p \wedge p$. But once we have $\square$, we can define the unique operator X which maps every proposition $\langle l, s \rangle$ to $\langle W, Is \rangle$, where I is $\lambda p.p$. Then $X\top = I\top$, even though $X \neq I$.

Proposition 5.2 DEFINABLE STRUCTURE *is invalid in any* CSM *such that there is a closed term* $\square$ *of type* $\langle \langle \rangle \rangle$ *for which:*

$$[\![\square]\!]_1(\langle l, s \rangle) = \begin{cases} W & \text{if } l = W \\ \emptyset & \text{otherwise} \end{cases}$$

Proof Letting $I := \lambda p.p$, and X a variable of type $\langle \langle \rangle \rangle$, define:

$$\varphi := \forall p \square Xp \wedge X\top = I\top$$

Any $\langle L, S \rangle$ satisfies $\forall p \square Xp$ iff L is the constant function $o \mapsto W$, and $\langle L, S \rangle$ satisfies $X\top = I\top$ iff $LW = W$ and S is I. Thus φ is uniquely satisfied by $\langle o \mapsto W, I \rangle$. Letting X stand for this element, $X\top = I\top$ by definition. And $X \neq I$ since $X\bot$ while $\neg I\bot$. So DEFINABLE STRUCTURE is not valid. $\square$

One way of seeing that this includes some safe CSMs based on a structural coarsening is to note that in a CSM with a one-element set of worlds, I satisfies the condition on $\square$. Thus we obtain:

Corollary 5.3 CLOSED STRUCTURE *does not entail* DEFINABLE STRUCTURE *in* $\vdash$.

Indeed, we have not been able to rule out that for every CSM, there is some closed term $\square$ satisfying the condition. More generally, we have not been able to determine whether DEFINABLE STRUCTURE is consistent. As we noted in the beginning, CLOSED STRUCTURE is arguably only plausible as a general schema when it is endorsed for all languages we might come to speak. The same line of thought extends to DEFINABLE STRUCTURE as well, so the consistency of this principle is crucial. We therefore consider this the most important question left open by our discussion:

Open Question 5.4 *Is* DEFINABLE STRUCTURE *consistent in* $\vdash$?

5.2 Plural Propositional Quantifiers

The views sketched in the previous section motivate ways of endorsing instances of CLOSED STRUCTURE without endorsing the full strength of STRUCTURE. But they motivate not only the instances of CLOSED STRUCTURE in a language L without plural resources, but also instances of the corresponding richer language L' which includes such resources. Since some instances of CLOSED STRUCTURE in this richer language are inconsistent, this poses a problem for these views. To illustrate this more concretely, consider a formula χ with a free plural propositional variable pp which involves no non-logical constants. For example, χ may be $\bot \prec pp$. We have shown that there are plural terms $\varphi\varphi$ and $\psi\psi$ such that $\vdash' \chi(\varphi\varphi) = \chi(\psi\psi) \wedge \varphi\varphi \neq \psi\psi$. These terms are constructed from χ without invoking any non-logical constants, so if χ contains no non-logical constants, then $\varphi\varphi$ and $\psi\psi$ also don't contain any non-logical constants. Thus, in L', there are *purely logical* counterexamples to CLOSED STRUCTURE. These examples contradict logical structuralism, and *a fortiori* the stronger effable structuralism.

The only reasonable option we see for defending such restricted structure views against the inconsistency of purely logical instances of CLOSED STRUCTURE in L' is to deny the coherence of plural propositional quantification. After all, the only alternative is to reject some of the logical principles of $\vdash'$, which we take to be an unpromising avenue: we consider the principles of $\vdash$ to be far more plausible than CLOSED STRUCTURE, and the additional principles of $\vdash'$ to be constitutive of the notion of plural propositional quantification. This raises the question whether one could motivate denying the coherence of plural propositional quantification. The answer to this depends arguably on one's general understanding of how we should conceive of the higher-order languages used here. We are sympathetic to a view along the lines suggested by Prior [33] and Williamson [41]. According to Prior and Williamson, we generally become competent speakers (or rather, writers) of higher-order languages by the same kinds of methods we use to become competent speakers of natural languages. Furthermore, what sentences of higher-order languages mean is not generally determined by a translation into any natural language. This is of course not to say that the facts in virtue of which higher-order languages are meaningful are completely inscrutable. A number of metasemantic factors can easily be identified which constrain the meanings of such languages. In the case of L, the following seem to us especially important:

First, there are general logical postulates which constrain the meanings of logical connectives. For example, the axiomatic principles UI and UG of ⊢ highly constrain the meanings of universal quantifiers. These principles can therefore partly be seen as definitional. Standard model theories serve to show that these constraints are jointly satisfiable, and so contribute to an overall argument for the good standing of the ideology. This first constraint should not be misunderstood as an inferentialist theory of meaning of logical connectives. Most importantly, we have not committed ourselves to claiming that such logical principles pin down the meanings of these connectives uniquely.

Second, natural languages like English contain phrases which are in many ways similar to the relevant logical connectives. For example, the meaning of the English conditional construction 'if . . . then . . . ' may serve to narrow down the meaning of →. Similarly, the meaning of phrases like 'every proposition' may serve to narrow down the meaning of $\forall_{()}$. Again, such constraints are not meant to give the full story. For example, → and 'if . . . then . . . ' may come apart in meaning in crucial respects. Maybe 'if . . . then . . . ' is ambiguous, allowing for a material conditional reading and other readings. On a very simple model of the metasemantics of →, the axiomatic principles involving → may serve to weed out the unwanted meanings of 'if . . . then . . . ', and thus the two first constraints would work in tandem to single out a unique meaning for →. Analogously, English phrases like 'every proposition' may be highly ambiguous, but axiomatic principles like UI and UG may go some way towards disambiguating them.

Third, the meanings of higher-order expressions may to some extent be identified by a kind of triangulation: In English, there are singular terms, predicates which may be applied to such terms, and quantifiers binding variables in the position of these terms. The meanings of higher-order quantifiers may be constrained by stating that the meaning of a quantifier binding variables in the position of predicates stands to predicates as quantifiers binding variables in the position of singular terms stands to singular terms.

Together, such matters provide significant grounds for the metasemantics of higher-order languages. It is of course still a substantial claim that higher-order ideology is in good standing. As in science, when introducing terms like 'phlogiston' and 'electron' for entities which are not directly observable, the success of the relevant terminological introduction must at least partly be judged on the basis of the success of theorizing in such terms. Thus, whether the recent use of higher-order languages in metaphysics is successful must at least partly be judged on whether such metaphysical theorizing gives rise to successful theories. And as in the sciences, the success of such theories will have to be judged on general criteria of theory choice such as simplicity, strength, and fit with the phenomena. There is of course much to be filled out in this sketch of an account of higher-order languages, but the present basics will suffice to evaluate the additional resources of plural propositional quantification.

The case for the plural propositional resources of L' is in fact very similar to that of the more standard higher-order resources of L. First, our axiomatic principles of quantification apply to plural propositional quantifiers just as they apply to ordinary higher-order quantifiers. In fact, this constraint is substantially stronger in the case of plural quantifiers, since we also impose the extensionality axiom EXT. And as before,

the satisfiability of these constraints can be established model-theoretically. Second, natural languages contain something akin to plural quantification, using phrases such as 'some thing*s*'. Just as 'some proposition' may be a rough guide to the meaning of $\forall p$, 'some proposition*s*' may be a rough guide to the meaning of $\forall pp$. And such plural propositional talk in English is highly natural: for example, when considering matters of modality, it is natural to speak of some proposition*s* being individually possible but jointly impossible. Third, the meanings of plural propositional quantifiers may also be constrained by a form of triangulation, now with respect to the previously established languages of higher-order logic and a plural extension of first-order logic: Using x for a singular individual variable and xx as a plural individual variable, the meaning of $\forall pp$ may be constrained by positing that it stands to $\forall p$ as $\forall xx$ stands to $\forall x$. That is, plural propositional quantification is what stands to propositional quantification as plural individual quantification stands to singular individual quantification.

Insofar as we have reason to think that standard higher-order languages are in good standing – and we think that we do – we have equally good reason to think that the extension by plural propositional resources is in good standing as well. Of course, there are differences in detail between the two accounts, and these differences may account for a difference in the meaningfulness of the relevant connectives. Furthermore, it may simply be that in the first case, the introduction of new logical ideology is successful whereas it is unsuccessful in the second case. After all, there is a substantial external component to metasemantics: however well-constrained the introduction of new ideology is, whether it is successful depends on whether reality provides the required structure for it to latch onto. Again, the success of such introductions must ultimately be judged on whether the relevant ideology leads to successful theorizing. Given the novelty of plural propositional quantification, any assessment we can undertake here will have to be highly preliminary. But from what we can see so far, the case is promising: given the failure of STRUCTURE, the motivations for CLOSED STRUCTURE are limited. In our estimation, they do not outweigh the intuitive appeal of plural propositional quantification.

In this defense of plural propositional resources, we have focused on plural propositional *quantifiers*. But the intelligibility of the plural operator π can be motivated similarly, by analogy to the standard variable-binder λ of higher-order logic. The use of π is in one sense crucial, since without it, there are no closed plural terms, and so no instances of CLOSED STRUCTURE with plural parameters. Our argument against CLOSED STRUCTURE therefore depends on π. However, recall from the previous section that many of the motivations for CLOSED STRUCTURE extend as well to DEFINABLE STRUCTURE. And in the presence of plural propositional quantifiers, the proof of the inconsistency of CLOSED STRUCTURE using π is straightforwardly adapted to a proof of the inconsistency of DEFINABLE STRUCTURE which does not rely on π. Without π, plural comprehension must be added as an axiom, as it can no longer be derived. One can then replace the closed plural terms $\mathsf{ff}(\mathsf{M})$ and $\mathsf{ff}(\mathsf{M}|_\infty)$ in the relevant instance of CLOSED STRUCTURE by plural variables which are bound by universal quantifiers restricted to the unique pluralities which are coextensive with the defining conditions of $\mathsf{ff}(\mathsf{M})$ and $\mathsf{ff}(\mathsf{M}|_\infty)$, respectively.

Finally, it is worth noting that plural propositional resources give rise to a range of interesting further questions. For example, one may wonder whether we can perform a speech act of plural assertion using plural terms. Just as we can use the sentence $\forall p(p = \neg\neg p)$ to assert a certain (coarse-grained) view about the individuation of propositions, can we use the term $\pi p.p = \neg\neg p$ to assert the instances of this universal claim? What, if anything, would be the difference between these speech acts? We think that these are great questions, but not ones which need to be answered for present purposes.

5.3 Attitudes

One of the reasons some might initially be attracted to STRUCTURE relates to attitude reports in English. Here is how one might reason: It is plausible that for any two sentences s and s' of English, if s and s' have different syntactic structure, or corresponding syntactic constituents which denote distinct entities, then $\ulcorner$it is possible that someone believes s but does not believe $s'\urcorner$ will be true, i.e. s and s' are not invariably intersubstitutable in the context of belief reports. Contraposing, if s and s' are intersubstituable in the context of attitude reports then they have the same syntactic form, and each of the corresponding constituents denote identical entities. Identity licenses intersubstitution, so if $\ulcorner s = s'\urcorner$ is true then s and s' have the same syntactic form, and each of the corresponding constituents denote identical entities. STRUCTURE generalizes this idea in the setting of higher-order logic.

So one might reason, but in fact, this argument does little to motivate STRUCTURE. This is because the argument relies on the intersubstitutibility of identicals, which can be stated as the following schema in our higher-order languages:

SUBSTITUTION: $\varepsilon = \eta \rightarrow (\varphi(\varepsilon) \rightarrow \varphi(\eta))$

This principle is controversial in the context of attitude reports, for reasons related to familiar Frege cases deriving from [12]: it is natural to think that although Hesperus is Phosphorus, the Greeks knew that Hesperus was visible in the evening, but did not know that Phosphorus was visible in the evening. Theories of attitude reports are usefully divided into two camps, depending on whether, in a language enriched by constants which intuitively stand for attitude verbs, they reject SUBSTITUTION on the basis of such examples. Let us consider the two options in turn.

First, consider those who take Frege cases to show that SUBSTITUTION fails in the presence of attitude contexts (as argued by Goodman and Lederman [16], this includes Fregeans, though see Pickel and Rabern [31]). On this view, the above argument cannot be used to motivate STRUCTURE, since it uses SUBSTITUTION as a premise, and SUBSTITUTION fails. It's worth noting that in the language used above, the connective $=$ is defined, and SUBSTITUTION is derivable from principles governing the primitive connectives. Those who want to reject SUBSTITUTION might therefore prefer to reject the claim that $=$, as defined above, is even coextensive with identity. For the purposes of this section, we will therefore consider $=$ as a primitive identity connective.

Consider now those who hold that Frege cases do *not* show that SUBSTITUTION fails in the presence of attitude contexts. This includes Millians as the position is traditionally understood (although Bacon and Russell [2, section 2] dispute this con-

sensus; cf. Goodman and Lederman [15] for discussion). On this view, there must be some way of accounting for our intuitive judgements concerning Frege cases. Depending on the details of this story, it may well be possible to use it also to account for the intuitions which motivated STRUCTURE, i.e., the intuitions according to which s and s' are not intersubstitutable in the context of attitude reports if they have different syntactic forms or corresponding constituents which denote distinct entities, without actually endorsing STRUCTURE.

Thus, whatever one's position on SUBSTITUTION, attitude ascriptions do not on their own motivate STRUCTURE. For those who reject SUBSTITUTION, there is nevertheless an interesting principle in the vicinity of STRUCTURE. For they might want to theorize using a notion of *cognitive equivalence*, which is finer-grained than identity. (Sometimes the label "cognitive equivalence" is used for a relation among linguistic expressions, but we will be using it as a relation among worldly entities.) Notating this $\approx$, they might then want to endorse the following principle:

SUBSTITUTION$_{\approx}$: $\quad \varepsilon \approx \eta \rightarrow (\varphi(\varepsilon) \rightarrow \varphi(\eta))$

Indeed, they might consider SUBSTITUTION$_{\approx}$ to be definitional of $\approx$. Some will dispute whether this notion of cognitive equivalence makes sense for types other than t; we won't argue for this idea here, but note that at type t, it can plausibly be defined along the lines of 'λpq. necessarily a person believes p if and only if the person believes q'.

The argument sketched at the beginning of this section may not convincingly motivate STRUCTURE. But for those who replace SUBSTITUTION with SUBSTITUTION$_{\approx}$, an obvious variant of the argument makes a convincing case for a corresponding variant of STRUCTURE which replaces identity with cognitive equivalence:

STRUCTURE$_{\approx}$: $\quad \forall X \forall Y \forall \bar{x} \forall \bar{y} (X\bar{x} \approx Y\bar{y} \rightarrow X \approx Y \wedge x_i \approx y_i)$

Using the quantificational principles of $\vdash$, this is inconsistent as well, by a corresponding variant of the Russell-Myhill theorem. This is not necessarily a problem: In the context of attitude reports, in response to the so-called problems of "quantifying in", some authors have rejected UI (e.g. Bacon and Russell [2], Lederman [24]) and many have denied $e\beta 1/2$ (this is probably the best understanding of Kaplan [18], cf. Kaplan [19]; see Yalcin [43], Lederman [25]). We think those who reject UI or $e\beta 1/2$ and who are also attracted to STRUCTURE$_{\approx}$ will most naturally seek to respond to the Russell-Myhill theorem using their weaker quantificational logic or weaker principles governing abstraction. But some have suggested that both UI and $e\beta 1/2$ should be upheld even when dealing with attitude reports, and such "classical opacists" will hold that STRUCTURE$_{\approx}$ is inconsistent; cf. Caie et al. [9]. On this view, the following variant of CLOSED STRUCTURE provides an interesting fallback position:

CLOSED STRUCTURE$_{\approx}$: $\quad \xi\bar{\varepsilon} \approx \zeta\bar{\eta} \rightarrow \xi \approx \zeta \wedge \varepsilon_i \approx \eta_i$ (where $\xi, \zeta, \bar{\varepsilon}, \bar{\eta}$ are closed)

In fact, one might argue that CLOSED STRUCTURE$_{\approx}$ was all that was ever motivated by cases of attitude ascriptions, for reasons similar to those which motivated effable structuralism: There may be ineffable entities to which no person can stand

in substantial cognitive relations, so that they would not be differentiable in terms of attitude ascriptions. But arguably, every entity expressed by a closed term of a spoken language is accessible in the relevant sense, so that at least STRUCTURE$_\approx$ restricted to what is expressed by closed terms is motivated – and that is just what is captured by CLOSED STRUCTURE$_\approx$. Given the inconsistency of STRUCTURE$_\approx$, one might therefore consider CLOSED STRUCTURE$_\approx$ an independently motivated fallback position.

Admittedly, one might worry about cases of attitude ascriptions with free variables in the complement clause which are bound by a quantifier; these cases may well motivate something stronger than CLOSED STRUCTURE$_\approx$. It would be worth exploring appropriate strengthenings, but here, we want only to note that whatever principle is motivated, it need not be as strong as STRUCTURE$_\approx$. CLOSED STRUCTURE$_\approx$ is a natural starting point for someone interested in exploring this style of view, and an examination of its consistency will provide insight into what strengthenings can consistently be endorsed.

The results established above show that CLOSED STRUCTURE$_\approx$ is consistent in $\vdash$, since we may – for the purposes of an abstract consistency proof – interpret $\approx$ and $=$ alike. The question for the classical opacitist endorsing CLOSED STRUCTURE$_\approx$ is now whether the inconsistency of CLOSED STRUCTURE in $\vdash'$ poses a problem. Just as the Russell-Myhill theorem can be transposed to show the inconsistency of STRUCTURE$_\approx$, the argument of Section 4 can be transposed to produce an argument for the inconsistency of CLOSED STRUCTURE$_\approx$. But this argument will rely on a principle corresponding to EXT (the extensionality of plural terms), according to which coextensiveness suffices not only for identity but also cognitive equivalence. And this is an assumption the classical opacitist will have antecedent reasons to reject, similar to their reasons for rejecting of SUBSTITUTION. So the inconsistency results in this paper do not give the classical opacitist reason to abandon CLOSED STRUCTURE$_\approx$.

Principles of structured propositions restricted to closed instances thus have something to offer in the context of attitude ascriptions. In closing, we want to pose a question about whether they might offer even more than we've suggested so far, asking whether a restriction to closed instances can help with a different limitative result on propositions. The theorem is due to Prior [32], and related to a model-theoretic observation by Kaplan [21]; see Bacon et al. [1] for discussion. To sketch it, let Q be a sentential operator expressing the attitude of being entertained (by a certain agent at a certain time). *Prima facie*, it would seem that one could uniquely entertain the thought that everything one entertains is false. *A fortiori*, it should be possible to entertain this thought, and nothing of a different truth value. That is, letting r abbreviate $\forall q(Qq \to \neg q)$, the following would seem to be possible:

P : $\quad Q\mathrm{r} \land \forall p(Qp \to (p \leftrightarrow \mathrm{r}))$

But this, as Prior shows, is inconsistent. Formally, we can treat Q as a variable of type $\langle\langle\rangle\rangle$; then $\vdash \neg$P. So surprisingly, if one entertains r, then one also entertains some p which differs in truth-value from r. But as in the case of the Russell-Myhill result, the proof of Prior's theorem does not explicitly specify any example of such a p.

This raises the question whether it is consistent to suppose that *among propositions expressed by closed sentences*, only propositions with the same truth value as r are being entertained. Generalizing from r, one might thus ask whether for every closed formula φ, the following set is consistent in $\vdash$:

$$\{Q\varphi\} \cup \{Q\psi \rightarrow (\psi \leftrightarrow \varphi) : \psi \text{ closed}\}$$

And, if all such sets are consistent in $\vdash$, then given the results established here, it is natural to follow up and ask whether they are also consistent in $\vdash'$. We leave this line of inquiry for another occasion.

6 Conclusion

The Russell-Myhill theorem shows the inconsistency of the view that propositions are structured, as captured by the principle STRUCTURE in relational type theory. Yet, intriguingly, it provides us with no concrete examples of the failure of propositions to be structured. Here, we have explored views which claim, in some form or another, that the failures of propositions to be structured can be contained, allowing propositions to be structured whenever it matters. Centrally, we have shown that such views show some initial promise, by proving the consistency of the principle of CLOSED STRUCTURE, which restricts STRUCTURE to closed sentences. Some consistency questions on strengthenings of this principle were answered as well, and others were left open, most importantly Question 5.4, which asks whether the principle of DEFINABLE STRUCTURE is consistent.

While interesting from a formal perspective, these open questions are in a sense superseded by our results on an extension of the relational type theory by plural propositional resources. First, we have shown that in such an extended language, CLOSED STRUCTURE is inconsistent. This result uses plural propositional quantifiers and the π operator to provide an explicit, and purely logical, counterexample to propositions being structured. Thus, even the most restrictive view considered here, logical structuralism, which restricts the claim that propositions are structured to propositions specified using purely logical vocabulary, is untenable. Second, we have shown that there are no compelling reasons to think that the resources of standard relational type theory are meaningful but plural propositional resources are not.

In the end, we therefore conclude that CLOSED STRUCTURE must be rejected alongside STRUCTURE. Despite some initial promise, even the restricted views of structured propositions considered here run into problems. Our conclusions are, however, not completely negative: Due to the opacity of attitude reports, so-called classical opacitists may want to endorse, instead of CLOSED STRUCTURE, a variant principle we call CLOSED STRUCTURE$_\approx$, which replaces identity with a relation $\approx$ of cognitive equivalence. In the setting without plural propositional resources, the consistency of this principle follows from the consistency of CLOSED STRUCTURE. And in the setting with plural propositional resources, the inconsistency of CLOSED STRUCTURE does not entail the inconsistency of CLOSED STRUCTURE$_\approx$, due to the opacity of $\approx$.

Appendix A: Consistency

A.1 Soundness

To show soundness, we require two standard lemmas, showing that the interpretation of free variables is well-behaved.

Lemma A.1 *In any* CSM, *if a and b agree on the free variables in ε, then $[\![\varepsilon]\!]^a = [\![\varepsilon]\!]^b$.*

Proof If ε is a variable x, then $a(x) = b(x)$, whence $[\![x]\!]^a = [\![x]\!]^b$. If ε is a constant c, then $[\![c]\!]^a = [\![c]\!]^b$ by construction. The remaining cases follow routinely by induction hypothesis. $\square$

This means that the assignment function is inert in the interpretation of any closed term ε. Consequently, we will omit it, and write simply $[\![\varepsilon]\!]$.

Lemma A.2 *In any* CSM, $[\![\eta[\bar{\varepsilon}/\bar{x}]]\!]^a = [\![\eta]\!]^{a[[\![\varepsilon_1]\!]^a/x_1,...,[\![\varepsilon_n]\!]^a/x_n]}$.

Proof By induction on the structure of η.

ATOMIC: If η is x, then $[\![x[\varepsilon/x]]\!]^a = [\![\varepsilon]\!]^a = [\![x]\!]^{a[[\![\varepsilon]\!]^a/x]}$. The case of constants is immediate by Lemma A.1.

APPLICATION: Routine by induction hypothesis.

ABSTRACTION: Let η be $\lambda\bar{y}.\varphi$. Since every ε_i is free for x_i, $[\![\lambda\bar{y}.\varphi[\bar{\varepsilon}/\bar{x}]]\!]^a = [\![\lambda\bar{y}.(\varphi[\bar{\varepsilon}/\bar{x}])]\!]^a$. This is:

$$\langle \bar{o} \mapsto [\![\varphi[\bar{\varepsilon}/\bar{x}]]\!]_1^{a[\bar{o}/\bar{y}]}, \lambda\bar{y} + [\![\varphi[\bar{\varepsilon}/\bar{x}]]\!]_2^{a[\bar{y}]}\rangle$$

By induction hypothesis and the fact that $\bar{x}$ and $\bar{y}$ are disjoint, this is:

$$\langle \bar{o} \mapsto [\![\varphi]\!]_1^{a[[\![\varepsilon_1]\!]^a/x_1,...,[\![\varepsilon_n]\!]^a/x_n][\bar{o}/\bar{y}]}, \lambda\bar{y} + [\![\varphi]\!]_2^{a[[\![\varepsilon_1]\!]^a/x_1,...,[\![\varepsilon_n]\!]^a/x_n][\bar{y}]}\rangle$$

which is $[\![\lambda\bar{y}.\varphi]\!]^{a[[\![\varepsilon_1]\!]^a/x_1,...,[\![\varepsilon_n]\!]^a/x_n]}$, as required. $\square$

Further, we note that the classical truth-conditions hold for primitive logical resources.

Lemma A.3 *For every* CSM $\mathfrak{M}$ *and assignment function a:*

(1) $\mathfrak{M}, a \vDash (\lambda\bar{x}.\varphi)\bar{\varepsilon}$ *iff* $\mathfrak{M}, a[[\![\varepsilon_1]\!]^a/x_1, \ldots, [\![\varepsilon_n]\!]^a/x_n] \vDash \varphi$

(2) $\mathfrak{M}, a \vDash \varphi \to \psi$ *iff* $\mathfrak{M}, a \vDash \varphi$ *only if* $\mathfrak{M}, a \vDash \psi$

(3) $\mathfrak{M}, a \vDash \forall x^t\varphi$ *iff* $\mathfrak{M}, a[o/x] \vDash \varphi$ *for all* $o \in D_t$

Proof (1) $\mathfrak{M}, a \vDash (\lambda\bar{x}.\varphi)\bar{\varepsilon}$
iff $@ \in [\![\lambda\bar{x}.\varphi]\!]_1^a(\langle[\![\varepsilon_1]\!]^a, \ldots, [\![\varepsilon_n]\!]^a\rangle)$
iff $@ \in [\![\varphi]\!]_1^{a[[\![\varepsilon_1]\!]^a/x_1,...,[\![\varepsilon_n]\!]^a/x_n]}$

iff $\mathfrak{M}, a[\llbracket \varepsilon_1 \rrbracket^a/x_1, \ldots, \llbracket \varepsilon_n \rrbracket^a/x_n] \vDash \varphi$.

(2) and (3) are routine. □

Similar conditions hold for the defined logical connectives; we note three representative cases.

Lemma A.4 *For every* CSM $\mathfrak{M}$ *and assignment function a:*

(1) $\mathfrak{M}, a \nvDash \bot$
(2) $\mathfrak{M}, a \vDash \neg\varphi$ *iff* $\mathfrak{M}, a \nvDash \varphi$
(3) $\mathfrak{M}, a \vDash \varepsilon = \eta$ *iff* $\llbracket \varepsilon \rrbracket^a = \llbracket \eta \rrbracket^a$

Proof Routine, using Lemma A.3. □

We can now establish soundness:

Proposition A.5 *If* $\vdash \varphi$ *then* φ *is valid in every* CSM.

Proof By a routine induction on the length of proofs, using the previous lemmas. □

A.2 Coarsenings

Given a coarsening $\sim$, fix a type-indexed choice function γ_t on $L_t/\sim_t$. That is, for each type t, let $\gamma_t : L_t/\sim_t \to L_t$ such that $\gamma_t(X) \in X$, for all $X \in L_t/\sim_t$. Assume further that for each variable x, $\gamma([x]_\sim) = x$; there exist such functions since being a coarsening guarantees that $x = y$ whenever $x \sim y$. As usual, we omit the type indices. γ maps every equivalence class of terms under $\sim$ to a designated representative. With this, we can define the result of replacing, in a given expression ε, any free variable x by the designated representative of the syntactic component of $a(x)$. We notate this as follows:

$$\varepsilon[a] := \varepsilon[\gamma(\pi_2(a(x_1)))/x_1, \ldots, \gamma(\pi_2(a(x_n)))/x_n] \quad (\bar{x} \text{ the free variables in } \varepsilon)$$

Lemma A.6 *In any* CSM, $\llbracket \varepsilon \rrbracket_2^a = [\varepsilon[a]]_\sim$.

Proof By induction on the complexity of ε.

ATOMIC: If ε is a variable x, then $\llbracket x \rrbracket_2^a = \pi_2(a(x)) = [\gamma(\pi_2(a(x)))]_\sim = [x[a]]_\sim$. If ε is a constant c, then $\llbracket c \rrbracket_2^a = [c]_\sim = [c[a]]_\sim$.

APPLICATION: If ε is an application $\xi\bar{\eta}$, then $\llbracket \xi\bar{\eta} \rrbracket_2^a = \llbracket \xi \rrbracket_2^a + \llbracket \eta_1 \rrbracket_2^a + \cdots + \llbracket \eta_n \rrbracket_2^a$. By IH, this is $[\xi[a]]_\sim + [\eta_1[a]]_\sim + \cdots + [\eta_n[a]]_\sim = [\xi\bar{\eta}[a]]_\sim$.

ABSTRACTION: If ε is an abstraction $\lambda\bar{y}.\varphi$, then $\llbracket \lambda\bar{y}.\varphi \rrbracket_2^a = \lambda\bar{y} + \llbracket \varphi \rrbracket_2^{a[\bar{y}]}$. By IH, this is $\lambda\bar{y} + [\varphi[a[\bar{y}]]]_\sim$. Let $\bar{x}$ be the free variables of $\lambda\bar{y}.\varphi$. Then $[\varphi[a[\bar{y}]]]_\sim$ may be specified as follows:

$$[\varphi[\gamma(\pi_2(a(x_i)))/x_1, \ldots, \gamma(\pi_2(a[\bar{y}](y_i)))/y_1, \ldots]_\sim$$

Note that $\pi_2(a[\bar{y}](y_i)) = [y_i]_\sim$, so $\gamma(\pi_2(a[\bar{y}](y_i))) = y_i$. Thus the replacements on $\bar{y}$ are vacuous; so $\lambda\bar{y} + [\varphi[a[\bar{y}]]]_\sim = \lambda\bar{y} + [\varphi[\gamma(\pi_2(a(x_i)))/x_1, \ldots]_\sim = [(\lambda\bar{y}.\varphi)[a]]_\sim$. $\qquad\square$

Note that in the case of closed terms ε, the assignment function can be omitted, in which case we obtain $[\![\varepsilon]\!]_2 = [\varepsilon]_\sim$.

Lemma A.7 *For any consistent set of equations T:*

(1)   If $\xi \sim^{sT} \zeta$ then $\xi\bar{\varepsilon} \sim^{sT} \zeta\bar{\varepsilon}$.
(2)   If $\varepsilon_i \sim^{sT} \eta_i$ then $\xi\varepsilon_1 \ldots \varepsilon_i \ldots \varepsilon_n \sim^{sT} \xi\varepsilon_1 \ldots \eta_i \ldots \varepsilon_n$.
(3)   If $\varphi \sim^{sT} \psi$ then $\lambda\bar{x}.\varphi \sim^{sT} \lambda\bar{x}.\psi$.

Proof (1) Assume $\xi \sim^{sT} \zeta$ is witnessed by expressions ϑ, χ and κ such that $\xi = \vartheta(\chi)$ and $\zeta = \vartheta(\kappa)$. Without loss of generality, the designated free variable of ϑ may be assumed not to occur in $\bar{\varepsilon}$. Then $\xi\bar{\varepsilon} = \vartheta\bar{\varepsilon}(\chi)$ and $\zeta\bar{\varepsilon} = \vartheta\bar{\varepsilon}(\kappa)$, whence $\xi\bar{\varepsilon} \sim^{sT} \zeta\bar{\varepsilon}$.

The cases of (2) and (3) are analogous, where in the case of (3), we assume the designated free variable of ϑ not to be among $\bar{x}$. $\qquad\square$

Lemma A.8 *For any consistent set of equations T, $\sim^T$ is a coarsening.*

Proof It is immediate by construction that $\sim_t$ is an equivalence relation. We check the three defining conditions.

For (1), we consider any distinct variables $x \neq y$, and show that $x \not\sim^T y$. Assume $x \sim^T y$ for contradiction. Then $T \vdash x = y$: In general, if $\varepsilon = \eta \in T$, then $T \vdash \vartheta(\varepsilon) = \vartheta(\eta)$. So if $\xi \sim^{iT} \zeta$, then $T \vdash \xi = \zeta$. By the reflexivity, symmetry and transitivity of $=$ in $\vdash$, it follows also that if $\xi \sim^T \zeta$, then $T \vdash \xi = \zeta$. So, as claimed, it follows from $x \sim^T y$ that $T \vdash x = y$, whence $T \vdash \forall x \forall y(x = y)$ by the definition of entailment. But $\vdash \neg\forall x\forall y(x = y)$: for type $\langle\rangle$, this is witnessed by $\vdash \bot \neq \top$; for any other type, it is witnessed by $\vdash (\lambda\bar{x}.\bot) \neq (\lambda\bar{x}.\top)$. Thus T is inconsistent, in contradiction to the assumption. So $x \not\sim^T y$.

For (2), assume $\varepsilon_1 \sim^T \eta_1, \ldots, \varepsilon_n \sim^T \eta_n$ and $\xi \sim^T \zeta$. Then there are $\bar{\vartheta}, \bar{\chi}$ such that

$$\xi = \vartheta^1 \sim^{sT} \cdots \sim^{sT} \vartheta^m = \zeta, \text{ and}$$
$$\varepsilon_i = \chi_i^1 \sim^{sT} \cdots \sim^{sT} \chi_i^{l_i} = \eta \text{ for all } i < n.$$

Then by Lemma A.7 (1 & 2),

$$\xi\varepsilon_1 \ldots \varepsilon_n = \vartheta^1\chi_1^1 \ldots \chi_n^1 \sim^{sT} \cdots \sim^{sT} \vartheta^m\chi_1^1 \ldots \chi_n^1$$
$$\cdots \sim^{sT} \vartheta^m\chi_1^{l_1} \ldots \chi_n^1$$
$$\cdots$$
$$\cdots \sim^{sT} \vartheta^m\chi_1^{l_1} \ldots \chi_n^{l_n} = \zeta\eta_1 \ldots \eta_n$$

So $\xi\bar{\varepsilon} \sim^T \zeta\bar{\eta}$, as required.

For (3), assume $\varphi \sim^T \psi$. Then there are $\bar{\chi}$ such that

$$\varphi = \chi_1 \sim^{sT} \cdots \sim^{sT} \chi_n = \psi.$$

Then by Lemma A.7 (3),

$$\lambda\bar{x}.\varphi = \lambda\bar{x}.\chi_1 \sim^{sT} \cdots \sim^{sT} \lambda\bar{x}.\chi_n = \lambda\bar{x}.\psi.$$

So $\lambda\bar{x}.\varphi \sim^T \lambda\bar{x}.\psi$, as required. □

Lemma A.9 $\sim^{\alpha,\eta}$ *is a coarsening.*

Proof Standard models show that α, η is consistent. The claim therefore follows by Lemma A.8. □

Lemma A.10 $\sim^{\alpha,\eta}$ *is structural.*

Proof We show that if $\xi\bar{\varepsilon} \sim^{\alpha,\eta} \zeta\bar{\chi}$, then $\xi \sim^{\alpha,\eta} \zeta$ and $\varepsilon_i \sim^{\alpha,\eta} \chi_i$ by induction on the length of the witnessing sequence. The case of a sequence of length 1 is immediate. So assume $\xi\bar{\varepsilon} \sim^{\alpha,\eta} \zeta\bar{\chi}$ is witnessed by a sequence of length $n+1$, with the nth element being ϑ. Then $\vartheta \sim^{s\alpha,\eta} \zeta\bar{\chi}$. Since no instance of α or η is an identity flanked by an application on either side, ϑ must be an application $\nu\bar{\mu}$, with $\nu \sim^{s\alpha,\eta} \zeta$ or $\nu = \zeta$, and $\mu_i \sim^{s\alpha,\eta} \chi_i$ or $\mu_i = \chi_i$. Thus $\nu \sim^{\alpha,\eta} \zeta$ and $\mu_i \sim^{\alpha,\eta} \chi_i$. Further, by IH, $\xi \sim^{\alpha,\eta} \nu$ and $\varepsilon_i \sim^{\alpha,\eta} \mu_i$. The claim follows by transitivity of $\sim^{\alpha,\eta}$. □

Lemma A.11 $\sim^{\alpha,\eta}$ *is safe.*

Proof We first show that for every instance $\chi = \kappa$ of α and η, $[\![\chi]\!]^a = [\![\kappa]\!]^a$, considering the two coordinates separately.

To show $[\![\chi]\!]_1^a = [\![\kappa]\!]_1^a$, we distinguish two cases. Consider first any instance of α: $[\![\lambda\bar{y}.\varphi[\bar{y}/\bar{x}]]\!]_1^a = \bar{o} \mapsto [\![\varphi[\bar{y}/\bar{x}]]\!]_1^{a[\bar{o}/\bar{y}]}$; by Lemma A.2, this is $\bar{o} \mapsto [\![\varphi]\!]_1^{a[\bar{o}/\bar{y}][\bar{o}/\bar{x}]}$, which by Lemma A.1 is $\bar{o} \mapsto [\![\varphi]\!]_1^{a[\bar{o}/\bar{x}]} = [\![\lambda\bar{x}.\varphi]\!]_1^a$. Consider now an instance of η: $[\![\lambda\bar{x}.\varepsilon\bar{x}]\!]_1^a = \bar{o} \mapsto [\![\varepsilon\bar{x}]\!]_1^{a[\bar{o}/\bar{x}]}$, which by Lemma A.1 is $\bar{o} \mapsto [\![\varepsilon]\!]_1^a(\bar{o}) = [\![\varepsilon]\!]_1^a$.

We now show $[\![\chi]\!]_2^a = [\![\kappa]\!]_2^a$. Since $\chi = \kappa$ is an instance of α or η, $\chi[a] = \kappa[a]$ is an instance of α or η as well. Thus $\chi[a] \sim^{\alpha,\eta} \kappa[a]$, whence $[\chi[a]]_{\sim^{\alpha,\eta}} = [\kappa[a]]_{\sim^{\alpha,\eta}}$. By Lemma A.6, it follows that $[\![\chi]\!]_2^a = [\![\kappa]\!]_2^a$.

So, for every instance $\chi = \kappa$ of α and η, $[\![\chi]\!]^a = [\![\kappa]\!]^a$. Therefore $[\![\vartheta]\!]^{a[[\![\chi]\!]^a/x]} = [\![\vartheta]\!]^{a[[\![\kappa]\!]^a/x]}$, whence with Lemma A.2, $[\![\vartheta(\chi)]\!]^a = [\![\vartheta(\kappa)]\!]^a$. It follows that if $\xi \sim^{i\alpha,\eta} \zeta$, then $[\![\xi]\!]^a = [\![\zeta]\!]^a$. Using the fact that identity is reflexive, symmetric and transitive, it follows by induction also that if $\xi \sim^{\alpha,\eta} \zeta$, then $[\![\xi]\!]^a = [\![\zeta]\!]^a$. □

Proposition A.12 *Any safe* CSM *based on a strongly structural coarsening validates* STRONG CLOSED STRUCTURE.

Proof Let $\mathfrak{M}$ be safe and based on a strongly structural coarsening $\sim$. Assume $\mathfrak{M}, a \models \xi[\bar{\varepsilon}/\bar{x}] = \xi[\bar{\eta}/\bar{x}]$, for closed $\bar{\varepsilon}, \bar{\eta}$, with all $\bar{x}$ free in ξ. Then $[\![\xi[\bar{\varepsilon}/\bar{x}]]\!]_2^a = [\![\xi[\bar{\eta}/\bar{x}]]\!]_2^a$, so with Lemma A.6, $[\xi[\bar{\varepsilon}/\bar{x}][a]]_\sim = [\xi[\bar{\eta}/\bar{x}][a]]_\sim$. Let $\bar{y}$ be the free variables distinct from those in $\bar{x}$. Let $\bar{z}$ be variables of the same types as $\bar{x}$, which do not occur in $\gamma(\pi_2(a(y_i)))$, for any $i \leq n$. Then $\xi[\bar{\varepsilon}/\bar{x}][a]$ is

$$\xi[\bar{z}/\bar{x}][\bar{\varepsilon}/\bar{z}][\gamma(\pi_2(a(y_1)))/y_1, \ldots, \gamma(\pi_2(a(y_n)))/y_n]$$

By the choice of $\bar{z}$, this is $\xi'[\bar{\varepsilon}/\bar{z}]$, where ξ' is

$$\xi[\bar{z}/\bar{x}][\gamma(\pi_2(a(y_1)))/y_1, \ldots, \gamma(\pi_2(a(y_n)))/y_n]$$

Analogously, $\xi(\bar{\eta})[a]$ is $\xi'[\bar{\eta}/\bar{z}]$. So $\xi'[\bar{\varepsilon}/\bar{z}] \sim \xi'[\bar{\eta}/\bar{z}]$. Since $\sim$ is strongly structural, it follows that $\varepsilon_i \sim \eta_i$. And as $\mathfrak{M}$ is safe, it follows that $[\![\varepsilon_i]\!]^a = [\![\eta_i]\!]^a$, as required. □

Appendix B: Inconsistency

We use '$\trianglelefteq$' for the relation of being an initial sub-chain:

$$\trianglelefteq := \lambda RS.R \sqsubseteq S \wedge \forall p \forall q (Spq \rightarrow (q \prec \mathrm{ff}(R) \rightarrow p \prec \mathrm{ff}(R)))$$

We start by showing that $\trianglelefteq$ totally orders chains:

Lemma A.13 $\vdash' CR \wedge CS \rightarrow R \trianglelefteq S \vee S \trianglelefteq R.$

Proof Assume CR and CS, i.e., that R and S are chains. Let qq be such that

$$qq = \pi p.p \prec \mathrm{ff}(R) \wedge p \prec \mathrm{ff}(S) \wedge R|_p \equiv S|_p.$$

We first show that qq is downward closed along R and S, in the following sense:

($\downarrow$)   If Rpq or Spq, then $q \prec qq$ only if $p \prec qq$.

Without loss of generality, consider any p, q such that Rpq and $q \prec qq$. The claim is immediate if $p = q$, so assume $p \neq q$. Since $q \prec qq$, $R|_q \equiv S|_q$. So as $R|_q pp$, also $S|_q pp$, whence Spq. Note that this means that our assumptions are symmetric with respect to R and S. We establish the three conjuncts required for $p \prec qq$:

(i & ii)   Since Rpq, $p \prec \mathrm{ff}(R)$. By symmetry, $p \prec \mathrm{ff}(S)$.

(iii)   We show that $R|_p \sqsubseteq S|_p$. So assume $R|_p rs$. Then Rrs, Rsp and $s \neq p$. By transitivity and antisymmetry of R, Rsq and $s \neq q$. Thus $R|_q rs$ and $R|_q sp$, and so $S|_q rs$ and $S|_q sp$. Therefore Srs and Srq. Since $r \neq q$, it follows that $S|_p rs$. Thus $R|_p \sqsubseteq S|_p$, and by symmetry, $S|_p \sqsubseteq R|_p$ as well. So $R|_p \equiv S|_p$, as required.

Thus $p \prec qq$, which establishes ($\downarrow$).

Next, we show the following claim:

($*$) $\mathrm{ff}(R) \sqsubseteq qq$ or $\mathrm{ff}(S) \sqsubseteq qq$.

Assume for contradiction that ($*$) fails to hold. Then $\mathrm{ff}(R) \not\sqsubseteq qq$, whence $\pi p.p \prec \mathrm{ff}(R) \wedge p \not\prec qq$ is non-empty. So as R is well-founded, there is an R-minimal $p \prec \mathrm{ff}(R)$ not among qq. Symmetrically, there is an S-minimal $q \prec \mathrm{ff}(R)$ not among qq.

We show that $\mathrm{ff}(R|_p) \equiv qq$: Consider first any $r \prec \mathrm{ff}(R|_p)$. Then Rrp and $r \neq p$, so by the minimality of p, it follows that $r \prec qq$. Conversely, consider any $r \prec qq$. Since $p \not\prec qq$, $r \neq p$. And by ($\downarrow$), Rpr fails to hold. But as both p and r are among $\mathrm{ff}(R)$, Rrp follows with the totality of R. Thus $r \prec \mathrm{ff}(R|_p)$.

By symmetry, $\mathrm{ff}(S|_q) \equiv qq$ as well, and so

$\mathsf{ff}(R|_p) \equiv \mathsf{ff}(S|_q)$.

Using EXT, we conclude:

$\chi(\mathsf{ff}(R|_p)) = \chi(\mathsf{ff}(S|_q))$.

Since $p \prec \mathsf{ff}(R)$, $q \prec \mathsf{ff}(S)$, and R and S are chains, we have:

$p = \chi(\mathsf{ff}(R|_p)) = \chi(\mathsf{ff}(S|_q)) = q$.

We show that $R|_p \equiv S|_q$ follows as a consequence. Without loss of generality, assume for contradiction that there are r, s such that $R|_p rs$ and not $S|_q rs$. Then Rrs, Rsp and $s \neq p$. So $r, s \prec \mathsf{ff}(R_p) \equiv \mathsf{ff}(S_q)$. Thus Srq and Ssq. Since not $S|_q rs$ was assumed it follows that not Srs. By the totality of S, this means that Ssr. So with transitivity and antisymmetry of S, we obtain Ssq and thus $s = q$. But this contradicts that $p = q$ and $s \neq p$. Thus, $R|_p \equiv S|_q$.

Therefore p (which is q) is such that $p \prec \mathsf{ff}(R)$, $p \prec \mathsf{ff}(S)$, and $R|_p \equiv S|_p$; therefore $p \prec qq$. But this contradicts the choice of p as not among qq. This establishes (∗), i.e., that $\mathsf{ff}(R) \sqsubseteq qq$ or $\mathsf{ff}(S) \sqsubseteq qq$.

Assuming $\mathsf{ff}(R) \sqsubseteq qq$, we show that $R \trianglelefteq S$. To show $R \sqsubseteq S$, assume Rpq. Then $q \prec qq$, so $q \prec \mathsf{ff}(S)$. Therefore, if $p = q$, then Spq. This leaves only the case of $p \neq q$ to be considered. In this case, $R|_q pp$. And since $q \prec qq$, $R|_q \equiv S|_q$; thus $S|_q pp$. So Spq, as required. Thus, $R \sqsubseteq S$. So, to show $R \trianglelefteq S$, assume Spq and $q \prec \mathsf{ff}(R)$. Then $q \prec qq$, so by (↓), $p \prec qq$, whence $p \prec \mathsf{ff}(R)$, as required. Thus if $\mathsf{ff}(R) \sqsubseteq qq$, then $R \trianglelefteq S$.

By symmetry, if $\mathsf{ff}(S) \sqsubseteq qq$, then $S \trianglelefteq R$. Since we have established (∗), i.e., that $\mathsf{ff}(R) \sqsubseteq qq$ or $\mathsf{ff}(S) \sqsubseteq qq$, it follows as claimed that $R \trianglelefteq S$ or $S \trianglelefteq R$.   □

With this, we can show that M is itself a chain.

Lemma A.14 $\vdash'$ CM.

Proof We first verify WM, i.e., that M is a well-order, by checking the four defining conditions.

Reflexivity: Immediate by reflexivity of chains.

Transitivity: If $\mathsf{M}pq$ and $\mathsf{M}qr$, then Rpq and Sqr for chains R and S. By Lemma B.1, $R \sqsubseteq S$ or $S \sqsubseteq R$. Thus, by the transitivity of chains, Rpr or Spr, whence $\mathsf{M}pr$.

Antisymmetry: Analogous to the case of transitivity, if $\mathsf{M}pq$ and $\mathsf{M}qp$, then by Lemma B.1, there is a chain Q such that Qpq and Qqp, whence with the antisymmetry of chains, $p = q$.

Well-foundedness: Consider any $pp \sqsubseteq \mathsf{ff}(\mathsf{M})$ such that there is some $p \prec pp$. By definition of M, $p \prec \mathsf{ff}(R)$ for some chain R. Let qq be such that

$qq = \pi q.q \prec pp \wedge q \prec \mathsf{ff}(R)$

$p \prec qq$, so qq is non-empty. Therefore, by the well-foundedness of R, there is an R-minimal $q \prec qq$. We show that q is M-minimal among pp. So consider any $r \prec pp$. If $r \prec \mathsf{ff}(R)$, then Rqr by the R-minimality of q, so $\mathsf{M}qr$. Otherwise, i.e.,

if $r \nprec \mathsf{ff}(R)$, there is a chain S such that $r \prec \mathsf{ff}(S)$. Then $S \not\sqsubseteq R$, so by Lemma B.1, $R \trianglelefteq S$. Assume for contradiction that $\mathsf{M}qr$ is not the case. Then Sqr also fails to be the case, whence by the totality of S, Srq. But $q \prec \mathsf{ff}(R)$, so with $R \trianglelefteq S$, $r \prec \mathsf{ff}(R)$, contradicting the assumption otherwise. Thus $\mathsf{M}qr$, establishing that the M-minimality of q. Therefore M is well-founded.

Finally, we show that if $p \prec \mathsf{ff}(\mathsf{M})$, then $p = \chi(\mathsf{ff}(\mathsf{M}|_p))$. If $p \prec \mathsf{ff}(\mathsf{M})$, then $p \prec \mathsf{ff}(R)$ for some chain R. Then $p = \chi(\mathsf{ff}(R|_p))$. We establish the following:

$$M|_p \equiv R|_p$$

$R|_p \sqsubseteq \mathsf{M}|_p$ is immediate from the fact that R is a chain. So assume $\mathsf{M}|_p qr$. Then $\mathsf{M}qr$, $\mathsf{M}rp$ and $r \neq p$. By Lemma B.1, $R \trianglelefteq \mathsf{M}$, since if $\mathsf{M} \trianglelefteq R$ then $\mathsf{M} \equiv R$. Thus $q, r \prec \mathsf{ff}(R)$. We know that $\mathsf{M}pr$ cannot hold, whence Rpr cannot hold either. By totality of R, Rrp follows. It remains to establish Rqr: If $q = r$, this follows by reflexivity, so assume otherwise. Then $\mathsf{M}rq$ fails, whence Rrq fails as well; by totality, Rqr follows. Thus $R|_p qr$ follows. This establishes $\mathsf{M}|_p \equiv R|_p$.

If follows that $\mathsf{ff}(M|_p) \equiv \mathsf{ff}(R|_p)$, and so with EXT that $p = \chi(\mathsf{ff}(R|_p)) = \chi(\mathsf{ff}(M|_p))$ as required. □

We can now show that ∞ is among the field of M.

Lemma A.15 $\vdash' \infty \prec \mathsf{ff}(\mathsf{M})$

Proof Assume for contradiction that $\infty \nprec \mathsf{ff}(\mathsf{M})$. Define:

$$\mathsf{M}^+ := \lambda pq.\mathsf{M}pq \vee (p \prec \mathsf{ff}(\mathsf{M}) \wedge q = \infty) \vee (p = \infty \wedge q = \infty)$$

We show that M^+ is a chain. First, we establish that it is a well-order:

Reflexivity: Immediate by construction.
Transitivity: Assume $\mathsf{M}^+ pq$ and $\mathsf{M}^+ qr$. If $r = \infty$, then $\mathsf{M}^+ pr$, so assume otherwise, i.e., $r \neq \infty$. Then $\mathsf{M}qr$, and as $\infty \nprec \mathsf{ff}(\mathsf{M})$, $q \neq \infty$ as well. Then $\mathsf{M}pq$, so $\mathsf{M}pr$ by transitivity of M.
Antisymmetry: Assume $\mathsf{M}^+ pq$ and $\mathsf{M}^+ qp$. Since $\infty \nprec \mathsf{ff}(\mathsf{M})$, $p = \infty$ iff $q = \infty$. So assume neither p nor q is ∞. Then $\mathsf{M}pq$ and $\mathsf{M}qp$, whence $p = q$ by antisymmetry of M.
Well-foundedness: Consider any non-empty $pp \sqsubseteq \mathsf{ff}(\mathsf{M}^+)$. If there is no $p \prec pp$ such that $p \prec \mathsf{ff}(\mathsf{M})$, then $p \prec pp$ iff $p = \infty$. Then trivially, ∞ is M^+-minimal among pp. So assume otherwise, i.e., that there is some $p \prec pp$ such that $p \prec \mathsf{ff}(\mathsf{M})$. Let qq such that:

$$qq = \pi q.q \prec pp \wedge q \prec \mathsf{ff}(\mathsf{M})$$

Since $p \prec qq$, qq is non-empty, so it follows from the well-foundedness of M that there is a $q \prec qq$ which is M-minimal among qq. We show that q is M^+-minimal among pp. So consider any $r \prec pp$. If $r \prec \mathsf{M}$, then $r \prec qq$, so $\mathsf{M}qr$ by M-minimality of q, whence $\mathsf{M}^+ qr$. Otherwise, $r = \infty$, in which case, $\mathsf{M}^+ qr$ as well. Thus q is M^+-minimal among pp as required.

Finally, we show that for all $p \prec \mathsf{ff}(\mathsf{M}^+)$, $p = \chi(\mathsf{ff}(\mathsf{M}^+|_p))$. We distinguish two cases:

Case 1: $p \prec \mathsf{ff}(\mathsf{M})$. Then as M is a chain, $p = \chi(\mathsf{ff}(\mathsf{M}|_p))$. We show that $\mathsf{ff}(\mathsf{M}|_p) \equiv \mathsf{ff}(\mathsf{M}^+|_p)$. $\mathsf{M} \sqsubseteq \mathsf{M}^+$, so $\mathsf{ff}(\mathsf{M}|_p) \sqsubseteq \mathsf{ff}(\mathsf{M}^+|_p)$. Conversely, consider any $q \prec \mathsf{ff}(\mathsf{M}^+|_p)$. Then M^+qp and $q \neq p$, and as $p \prec \mathsf{ff}(\mathsf{M})$ and $\infty \nprec \mathsf{ff}\mathsf{M})$, $\mathsf{M}qp$. So $q \prec \mathsf{ff}(\mathsf{M}|_p)$. Therefore, as claimed, $\mathsf{ff}(\mathsf{M}|_p) \equiv \mathsf{ff}(\mathsf{M}^+|_p)$. By EXT, $p = \chi(\mathsf{ff}(\mathsf{M}|_p)) = \chi(\mathsf{ff}(\mathsf{M}^+|_p))$, as required.

Case 2: $p \nprec \mathsf{ff}(\mathsf{M})$; then $p = \infty$. For any p, $p \prec \mathsf{ff}(\mathsf{M})$ iff $\mathsf{M}^+p\infty$ and $p \neq \infty$, which is the case iff $p \prec \mathsf{ff}(\mathsf{M}^+|_\infty)$. Thus $\mathsf{ff}(\mathsf{M}) \equiv \mathsf{ff}(\mathsf{M}^+|_\infty)$. By EXT, $\infty = \chi(\mathsf{ff}(\mathsf{M})) = \chi(\mathsf{ff}(\mathsf{M}^+|_\infty))$, as required.

By construction, $\mathsf{M}^+\infty\infty$. We have just shown M^+ to be a chain, so it witnesses $\mathsf{M}\infty\infty$. Thus $\infty \prec \mathsf{ff}(\mathsf{M})$, contradicting our assumption to the contrary. □

References

1. Bacon, A., Hawthorne, J., Uzquiano, G. (2016). Higher-order free logic and the Prior-Kaplan paradox. *Canadian Journal of Philosophy*, *46*, 493–541.
2. Bacon, A., & Russell, J.S. (2019). The logic of opacity. *Philosophy and Phenomenological Research*, *99*, 81–114.
3. Bell, J.L. (1995). Type reducing correspondences and well-orderings: Frege's and Zermelo's constructions re-examined. *The Journal of Symbolic Logic*, *60*, 209–221.
4. Boolos, G. (1984). To be is to be a value of a variable (or to be some values of some variables). *The Journal of Philosophy*, *81*, 430–449.
5. Boolos, G. (1997). Constructing Cantorian counterexamples. *Journal of Philosophical Logic*, *26*, 237–239.
6. Burali-Forti, C. (1897). Una questione sui numeri transfiniti. *Rendiconti del Circolo Matematico di Palermo*, *11*, 154–164.
7. Burgess, J.P. (2004). E Pluribus Unum: plural logic and set theory. *Philosophia Mathematica*, *12*, 193–221.
8. Burgess, J.P., & Rosen, G. (1997). *A subject with no object*. Oxford: Clarendon Press.
9. Caie, M., Goodman, J., Lederman, H. (forthcoming). Classical opacity. *Philosophy and Phenomenological Research*.
10. Dorr, C. (2016). To be F is to be G. *Philosophical Perspectives*, *30*, 39–134.

11. Fine, K. (1977). Properties, propositions and sets. *Journal of Philosophical Logic*, *6*, 135–191.
12. Frege, G. (1892). Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik*, *NF 100*, 25–50.
13. Gallin, D. (1975). *Intensional and higher-order modal logic*. Amsterdam: North-Holland.
14. Goodman, J. (2017). Reality is not structured. *Analysis*, *77*, 43–53.
15. Goodman, J., & Lederman, H. (forthcoming a). Perspectivism. *Noûs*.
16. Goodman, J., & Lederman, H. (forthcoming b). Sense, reference and substitution. *Philosophical Studies*.
17. Kanamori, A. (1997). The mathematical import of Zermelo's well-ordering theorem. *Bulletin of Symbolic Logic*, *3*, 281–311.
18. Kaplan, D. (1968). Quantifying in. *Synthese*, *19*, 178–214.
19. Kaplan, D., & Schilpp, P.A. (1986). Opacity. In Hahn, L.E. (Ed.) *The Philosophy of W. V. Quine* (pp. 229–289). La Salle: Open Court.
20. Kaplan, D. (1989). Demonstratives. In Almog, J., Perry, J., Wettstein, H. (Eds.) *Themes from Kaplan* (pp. 481–563). Oxford: Oxford University Press. Completed and circulated in mimeograph in the published form in 1977.
21. Kaplan, D. (1995). A problem in possible-world semantics. In Walter Sinnott-Armstrong, Diana Raffman, and Nicholas Asher, editors, *Modality, Morality and Belief: Essays in Honor of Ruth Barcan Marcus*, Cambridge University Press, Cambridge.
22. King, J. (1996). Structured propositions and sentence structure. *Journal of Philosophical Logic*, *25*, 495–521.
23. King, J.C. (2007). *The nature and structure of content*. Oxford: Oxford University Press.
24. Lederman, H. (forthcoming). Fine-grained semantics for attitude reports. *Semantics and Pragmatics*.
25. Lederman, H. Fregeanism, sententialism, and scope. unpublished.
26. Linnebo, Ø. (2017). Plural quantification. In Zalta, E.N. (Ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University summer 2017 edition.
27. Montague, R. (1974). *Formal Philosophy: Selected papers of Richard Montague*. New Haven: Yale University Press. Edited and with an introduction by Richmond H. Thomason.
28. Muskens, R. (2007). Intensional models for the theory of types. *The Journal of Symbolic Logic*, *72*, 98–118.
29. Myhill, J. (1958). Problems arising in the formalization of intensional logic. *Logique et Analyse*, *1*, 78–83.
30. Oliver, A., & Smiley, T. (2016). *Plural logic*. Oxford: Oxford University Press, second, revised and enlarged edition.
31. Pickel, B., & Rabern, B. (forthcoming). Frege and saving substitution. *Philosophical Studies*.
32. Prior, A.N. (1961). On a family of paradoxes. *Notre Dame Journal of Formal Logic*, *2*, 16–32.
33. Prior, A.N. (1971). Platonism and quantification. In *Objects of Thought* (pp. 31–47). Oxford: Oxford University Press.
34. Roberts, S. (forthcoming). Pluralities as nothing over and above. *Journal of Philosophy*.
35. Russell, B. (1903). *The principles of mathematics*. Cambridge: Cambridge University Press.
36. Salmon, N. (1986). *Frege's Puzzle*. Cambridge: MIT press.
37. Soames, S. (1987). Direct reference, propositional attitudes, and semantic content. *Philosophical Topics*, *15*, 47–87.
38. Uzquiano, G. (2015). A neglected resolution of Russell's paradox of propositions. *The Review of Symbolic Logic*, *8*, 328–344.
39. Walsh, S. (2016). Predicativity, the Russell-Myhill paradox, and Church's intensional logic. *Journal of Philosophical Logic*, *45*, 277–326.
40. Whitehead, A.N., & Russell, B. *Principia Mathematica*, 1–3, Cambridge: Cambridge University Press, 1910–1913.
41. Williamson, T. (2003). Everything. *Philosophical Perspectives*, *17*, 415–465.
42. Williamson, T. (2010). Necessitism, contingentism, and plural quantification. *Mind*, *475*, 657–748.
43. Yalcin, S. (2015). Quantifying in from a Fregean perspective. *Philosophical Review*, *124*, 207–253.