Higher-Level Plural Logic

MSc Thesis *(Afstudeerscriptie)*

written by

Orestis Dimou Belegratis
(born June 22nd, 2000 in Athens, Greece)

under the supervision of **Dr. Luca Incurvati**, and submitted to the Examinations Board in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam*.

Date of the public defense:	**Members of the Thesis Committee:**
July 8, 2025	Dr. Balder ten Cate (Chair)
	Dr. Luca Incurvati (Supervisor)
	Dr. Thomas Schindler
	Dr. Simone Picenni

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

In the current thesis, I aim to explore numerous logical and philosophical questions regarding higher-level plurals. I begin by proposing my own conception of higher-level plural reference, dubbed *Combinatorial Reference*, in a way that: (a) supports the intelligibility of higher-level plurals, (b) enjoys several conceptual advantages over the alternatives in the literature, and (c) naturally motivates a specific logical framework. I then compare the formal aspects of that framework to what I consider to be the best alternative to higher-level plural logic, namely the *generalized cover* approach to plurals. I axiomatize both logics as many-sorted first-order theories and then prove that they are both *Morita Equivalent* and *Bi-interpretable*. Subsequently, I turn to an important question regarding the height of the hierarchy of higher-level plurals and how this hierarchy can motivate another one consisting of more and more expressive higher-level plural languages. I compare my strategy to others from the literature, most notably that of Linnebo and Rayo (2012), and I find that I can both avoid the criticisms voiced against them and reach a similar conclusion. Finally, I consider the arguments of Button and Trueman (2022) regarding the formalization of higher-level plural logic. While they purport to establish that the formalization should be carried out in a one-sorted framework, I argue for my preferred, typed formalization due to its intimate conceptual connection to *Combinatorial Reference*.

Acknowledgements

First, I want to thank Luca for being a great supervisor, especially for his support during all those moments when I felt like the project was going nowhere. Working with Luca has been a great privilege. Our discussions about logic and philosophy have shaped my thinking in ways that I will carry with me indefinitely. Moreover, Luca's help and support throughout the stressful perior of grad-school applications will be something I will never forget. I also would like to thank the members of my committee for taking the time to read my thesis and for asking me deep, thought-provoking questions.

And now I want to thank all the people that made Amsterdam feel like home. I want to thank Mees for giving me a home, when just finding a house seemed like an impossibility. Mees and Bruno, you were the best roommates anyone could've ever asked for, albeit not the best FIFA players. I want to thank the Greek set-theorists, Orestis and Spyros, for all the early nights, the late nights, the conversations, and for being there for me since day one of the MoL. I want to thank Umberto for teaching me how to take my insane ideas and give them substance, and for unapologetically being who he is. I want to thank Merlin, whose philosophical acumen amazed from the very beginning and without whom I would've been a far worse philosopher and person. I want to thank Marios for everything: the endless conversations, the endless walks, the endless coffee-drinking, the endless debauchery. More generally, thanks to everyone who was there for me and who allowed me to be there for them: Christos, Izzie, Pablo, Katherin, Emre, Giannis, Özgün, Mette, Clara, Lamarana, Matteo.

Finally, some think that it should go without saying, but these are exactly the things that need to be said. I want to thank my brother Christos and my parents, Giannis and Chrysoula, for their unwavering support through all of these years. I will be eternally grateful for all the opportunities they gave me and for believing in me. This thesis is dedicated to you. Last but not least, many thanks to all the people who have always been there for me: George, Alex, Giannis, Artemis, Giannis, and Aggeliki. You've all made me the person I am today. Don't forget: it's your fault.

Contents

Chapter 1

Introduction

1.1 Motivation and Outline

In some sense, this thesis is about *reference*. Not necessarily a form of reference that is actually used, but a form of reference that *could* be utilized. A form of reference that would allow us to make sense of the intended interpretation behind certain formal languages (logics), thus deeming them legitimate for philosophical applications and investigations. So in another sense, this thesis is about *formal languages (logics)* and what it takes for them to be considered legitimate.

This work is situated primarily in the literature regarding *pluralism*. I aim to convince someone who already accepts the legitimacy of *plural reference*, i.e., the ability to refer directly and simultaneously to multiple objects, that there is another legitimate form of reference that she can acquire a grasp of. I will dub this form as *Combinatorial Reference*, and argue that it is a comprehensible way of understanding the way in which higher-level plurals refer. Higher-level plural terms are going to be understood as the result of some form of *iteration* of the move from the singular to the plural. In my take on the matter, what is iterated is the *containment* relation between singulars and plurals - roughly corresponding to the locution "is among" - to create plural terms of higher levels that can contain terms of all lower levels.

The four main questions guiding this thesis were first discussed by Oliver and Smiley (2016, p. 276). They maintain that any logician who tackles higher-level plural logic needs to answer four questions. (1) Is the idea of higher-level pluralities even intelligible? Many do not think so, but I argue that it is. (2) If the idea of higher-level pluralities is intelligible, how is it properly expressed? I argue that on the formal side it should be expressed by allowing for higher-level terms, quantifiers, and predicates; the upshot of formalizing it thusly is that it provides insight into how higher-level pluralities refer. (3) Are the levels cumulative or somehow exclusive? I argue that they are cumulative and make this an essential aspect of the formalism. (4) How far do the levels go? Is there a natural upper bound to the levels of the plural hierarchy? I argue that there is none and that one can make sense of as many levels as there are ordinals to index them.

Each of these questions, along with other naturally arising ones, is dealt with throughout this thesis. The obvious starting point is the intelligibility of the concept of a higher-level plurality, since there is no consensus in the literature on this.

Making sense of higher-level plurals is a deeply contested matter, and has been deemed highly problematic for numerous reasons by multiple authors, such as Simons (1982), Lewis (1990), Rosen and Dorr (2002), Linnebo (2003), Rumfitt (2005), McKay (2006), and Ben-Yami (2013), Payton (2025).

Others, such as Rayo (2006), Oliver and Smiley (2016), Simons (2016), Wagner (2015), and Grimau (2018), have argued for their intelligibility and their theoretical significance.

My first order of business, in Chapter 2, is to sketch a conceptual map of the numerous alternative positions in the literature regarding higher-level plurals. Three main positions will be described: (a) a *skeptical* one expressing several reservations about the intelligibility of higher-level plurals, (b) one that argues that higher-level plurals can be *paraphrased away* or that are completely *dispensable*, and (c) one that takes them at *face-value*. There is a plethora of specific formulations of each of these positions, each with its own intricacies and nuances. Doing justice to all of these details cannot possibly be achieved in a part of a single chapter, but I believe that the important takeaways have been faithfully represented in a way that does not limit the scope of my arguments.

After sketching the competing attitudes towards higher-level plurals, I propose my own way of understanding higher-level plural terms, higher-level plural reference, and higher-level plural logic. All this revolves around the notion of *containment* - plurals of some level are *contained* in plurals of higher levels - and its utilization to grasp the notion of *Combinatorial Reference*. After offering an understanding of *Combinatorial Reference*, I compare it *conceptually* to the other ways of understanding (apparently) higher-level plural terms, both to the ones that ultimately deem them dispensable and also to those that take them at face-value.

A general theme throughout this thesis is the distinction between *conceptual* and *formal* comparisons. Conceptual comparisons are about the intended interpretations of formal frameworks and whether they offer a more accurate depiction of the thing they are trying to capture; sometimes, these comparisons are just simply about the informal ways in which a given notion is explicated. This is what we will talk about at the end of Chapter 2. Formal comparisons, on the other hand, have to do purely with the syntactic and semantic properties of logical frameworks. In Chapter 3, we investigate the formal relations between higher-level plural logic and plural logic endowed with *covers*, only to find that they are *theoretically equivalent* under two important senses of the term: *Morita Equivalence* and *Bi-interpretability*. The cover-based approach appears to be one of the main competitors to my approach. Therefore, the importance of this chapter lies in accentuating the conceptual differences by disregarding the formal ones.

Chapter 4 is then devoted to answering the question: How high does the plural hierarchy extend? The question is broken down into two parts. I first try to understand whether *Combinatorial Reference* can help us grasp how plurals of transfinite levels would refer, and if that is so, whether there is a level of plurals that serves as a natural upper bound for this hierarchy. The answer is that under some auxiliary assumptions, α-level plurals refer according to *Combinatorial Reference* for every ordinal α. This is what we will later refer to as the **Plural Axiom of Ordinals**.

The second part involves determining the connection between the Plural Axiom of Ordinals and the legitimacy of higher-level plural languages involving terms and quantifiers of at most level α for a given ordinal α. We first propose two different conceptions of language legitimacy, before arguing that under one of the two conceptions, α-level plural languages are legitimate for all ordinals α. Given that this understanding of plurals of transfinite levels rests heavily on a set-theoretic background, I explore the possibility of providing an autonomous way of countenancing an ever-extending hierarchy of plurals and plural languages. I do so by adapting a technique due to Gödel (1933) and arguing for its success, contra the claims of Button and Trueman (2022). The autonomy achieved through this process results in a self-supporting framework with more potential philosophical applications due to its minimal dependence on other concepts.

Subsequently, we turn to an alternative approach due to Linnebo and Rayo (2012), which purports to reach similar conclusions, and the criticisms it has sustained due to Studd (2021) and Button and Trueman (2022). I argue that, irrespective of whether these criticisms condemn Linnebo and Rayo's attempt, they do not impact my argument at all. Although my goal is slightly more modest than that of Linnebo and Rayo, the criticisms are avoided, and the goal is attained. I, finally, turn to a brief discussion on whether the logic of higher-level plurals should be formulated in a typed (many-sorted) or in an untyped framework. While Button and Trueman (2022) argue for the latter, I believe that there are benefits to be reaped from the former approach, which I ultimately take to be more in line with my specific goals in this thesis.

Chapter 2

Higher-Level Plurals

The overarching theme of this chapter is to outline a particularly novel way to conceive of the notion of higher-level plural reference, through what I will call *Combinatorial Reference*. The goal of this project is not to try to convince someone already skeptical of *plural* reference that she should also acknowledge its higher-level counterpart. It is rather the pluralist, she who already accepts the legitimacy of plural terms, reference, and quantification, that I aim to convince of the legitimacy of my conception of higher-level plural reference. Thus, I have to start at the very beginning by presenting, albeit briefly, pluralism.

Next on our agenda will be a survey of the existing approaches to higher-level plural reference. I will attempt to give an overview of each of the main ideas and the corresponding formal apparatuses, before embarking on criticizing them for their shortcomings. Then, I will present my proposal as to how higher-level plural reference should be understood, I will compare it with the existing alternatives, and I will argue for its superiority.

2.1 Pluralism

In natural language, the singular terms such as *John, the book, the dog, etc.,* are taken to refer, if at all, to specific, unique objects. Under at least one conception of quantification, standard first-order quantifiers are similarly taken to bind variables that range over a domain of objects. Singular terms, however, are not the only kind of terms one finds in natural language; similarly, this way of referring to objects does not seem to be the sole referential device we employ.

In natural language, one can additionally find *plural* terms, such as *the cats, the plates*, or *the Beatles*. A natural question to pose then is whether the existence of such terms sanctions the legitimacy of an alternative notion of reference as well. Boolos (1984, 1985) argues that this should be the case: plural terms, contra singular ones, refer, if at all, to multiple objects at once instead of a unique one. *The cats* refer to the cats, *the books* to the books, and so on.

It could hardly be contested that we have this ability, and a testament to this would be actual language usage. Based on this observation, Boolos, among other pluralists, argues that we can use these expressive resources to outline a logical framework that takes them to heart. And this formal choice enjoys a multitude of both formal and conceptual upshots as well.

Firstly, it allows us to capture our linguistic practices in a more direct way. Consider the alternative. What would a plural term like *the cats* refer to if not to the cats? An off-hand answer comes from standard set-theoretic semantics and answers immediately: the set containing all and only these cats.

Although nobody would contest the formal adequacy of the set-theoretic approach to these semantics or the ability of set theory to offer a truth-preserving paraphrase for that discourse, we would be remiss if we thought that these semantics capture precisely our linguistic practice or the commitments incurred from it. Of course, when talking of *the cats*, we are not inadvertently and without explicitly doing so referring to sets. Boolos' proposal is substantially more intuitive: when plural terms occur in sentences, they refer to multiple objects at once. It also gives a more immediate way to talk of things that cannot be adequately represented using sets, such as talk of plurals that involve class-many objects. We cannot regiment in set theory talk of *the ordinals*, precisely because no set containing them exists in pain of the Burali-Forti paradox.

On the formal front, if we want to capture these aspects, we outline a syntax allowing for plural terms, plural variables, and irreducibly plural quantification, which binds the plural variables. On the semantic level, we use the notion of *plural reference* to explain how plural terms refer to multiple objects at once; in other words, we use semantic resources which are *homophonic* to the language whose meaning we are investigating, precisely because we already possess a grasp of those resources. In that way, using these semantic resources seems to result in a framework whose intended interpretation is more in line with our linguistic practice. Additionally, this intended interpretation seems to capture our ontological commitments, in the sense that talking of *the cats* does not commit one to the existence of any other entity apart from all the individual cats (Boolos, 1984). We follow suit with Oliver and Smiley (2016) and use *plurals, pluralities*, and later on *superplurals* and *higher-level plurals*, as *pseudo-singular* terms; namely, terms that are singular, but function plurally in the semantics. Therefore, talk of the existence of a plurality will simply be shorthand for the existence of individuals plurally quantified over or plurally referred to.

Secondly, to press this point further and argue for the incorporation of plural terms, variables, and predicates in our logical language, Boolos invokes the infamous Geach-Kaplan sentence "Some critics admire only one another". Whereas this sentence cannot be formalized in first-order logic, one can formalize it in plural logic or monadic second-order logic.[1] Therefore, it seems that plurals afford a way to genuinely extend our expressive resources.

Furthermore, an additional reason to believe that plural terms, or the notion of plural reference, are irreducible is due to the different ways in which predicates with plural arguments function. The relevant distinction here is between *collective* and *distributive* predication. In the latter case, saying that for some objects the predicate F holds is simply to say that F holds of each of these objects individually. In the former case, however, F holds collectively of all of these objects; the truth of the sentence "Harry and Sally met" is surely not because "Harry met" and "Sally met", both of which are ungrammatical. If we do not accept the legitimacy of plural reference, then how are we supposed to account for the collective predication?[2]

To make a long story short, multiple philosophers have accepted the legitimacy of the notion of plural reference and have acknowledged that plurals should/could be part of our philosophical

[1] It turns out that these two share the same expressive power. There is a way to define a translation from each formula of plural logic to a formula of monadic second-order logic, and vice versa, so that theorems of the one are mapped to theorems of the other. In more technical parlance, this shows that plural logic is interpretable in monadic second-order logic, and vice versa.

[2] One may argue that an adequate paraphrase would be that "Harry met with Sally and Sally met with Harry", thus dispensing with the need for plurals. This worry can be silenced by considering more intricate examples involving reciprocal verb constructions, such as "John, Mary, and Daniel played against each other in a three-way game", which does not seem to lend itself to a paraphrase like the one above. This does not imply that there is no adequate paraphrase at all, but that *this* strategy fails.

and logical theorizing. Whether plural logic is *logic proper* is obviously a substantial question that I deliberately avoid getting into. Different logical systems have been proposed regarding plural logic, but we will focus on just two, which take the legitimacy of the aforementioned notions to heart.

The most prominent logical systems involving plurals are PFO and PFO+. The main difference between the two is that the latter allows for plural predication, while the former does not.[3]

Let us start with PFO+, which extends first-order logic with identity. The extension happens as follows by adding:

1. Plural terms by incorporating plural variables ($xx, yy, \ldots$), corresponding to natural language pronouns, e.g., "they", and plural constants ($cc, dd, \ldots$), corresponding to plural proper names, e.g., "the cats".

2. Quantifiers binding plural variables ($\forall xx, \exists yy, \ldots$).

3. A primitive, logical, binary predicate $\prec$ for plural membership (containment) roughly corresponding to "is one of" or "is among".

4. Symbols for *collective* plural predicates with superscripts indicating the arity of each predicate ($P^1, P^2, \ldots, Q^1, \ldots$).

The formula formation rules are the obvious ones. Let us note that we only need symbols for collective plural predicates, since distributive predication can be expressed as $P(xx) \leftrightarrow \forall x(x \prec xx \rightarrow P(x))$, where using the same symbol P is simply an abuse of notation. The less expressive language PFO is just the fragment of PFO+ that only satisfies (1)-(3).

The formal framework PFO+ is also supplemented with a list of axioms and rules of inference. In this case, the axioms and rules include those of first-order logic with identity, and the expected ones for the plural quantifiers. Furthermore, there are three additional axioms governing plurals, which in combination with the previously mentioned axioms and rules comprise what Florio and Linnebo call *traditional plural logic* (Florio and Linnebo, 2021, p.20).

First, we have the axiom telling us that any plurality must be *non-empty*. Intuitively, this follows from the fact that for some things to exist, there needs to be at least one thing among them. Second, the *indiscernibility* axiom, which tells us that coextensive pluralities satisfy the same formulas. Thirdly and finally, we have the axiom scheme of *unrestricted comprehension*. For any formula $\phi(x)$ containing x, but not xx, as free, if there is an object satisfying it, then there exists also the plurality of all the things that satisfy it. The axioms are (prefaced with the necessary universal quantifiers whenever free variables are involved):

Non-empty $\forall xx \exists y\ (y \prec xx)$
Indisc $\forall xx \forall yy (xx \approx yy \rightarrow (\phi(xx) \leftrightarrow \phi(yy))$[4]
P-Comp $\exists x \phi(x) \rightarrow \exists xx \forall x (x \prec xx \leftrightarrow \phi(x))$

As I briefly mentioned before, Boolos (1984) provides a way to provide semantics for plural logic by employing plural resources in the metalanguage. The domain is just *some objects* - not the set with

[3]The following outline draws significantly from Florio and Linnebo (2021), which, alongside the references therein, serves as a great reference for the logical, philosophical, and linguistic aspects of plural logic.

[4]$xx \approx yy$ means coextensive in the sense that $\forall z(z \prec xx \leftrightarrow z \prec yy)$. Later on, we will simply use a primitive identity symbol that holds exactly of coextensive plurals. This version of the axiom will then follow as an instance of substitution of identicals.

these objects as members - and singular terms denote specific, unique objects in the domain, while plural terms denote *multiple objects* in the domain directly. In this way, Boolos dispenses with any set-theoretic regimentation of the semantic values of plurals. This feature of plural terms, i.e., their ability to signify or denote multiple objects of our domain of discourse, is what we will be calling *plural reference*. This concept will be of major importance for what is to come, and especially how it relates to *higher-level plural reference* in general and specifically the notion of *Combinatorial Reference* I will be proposing.

Before moving on to higher-level plural logic, we should make some further remarks about traditional plural logic. The formal systems of PFO and PFO+ often figure as part of the philosopher's toolkit from the latter quarter of the 20^{th}-century onwards, due to their numerous philosophical advantages in applications. This is important because it will serve as a guide for part of our discussion of higher-level plural logic, as I would like to argue that there are more philosophical benefits to be reaped from employing higher-level plural logic.

The first alleged philosophical upshot of plural logic is that, at least according to some, it constitutes "pure logic". Plural logic is *ontologically innocent*, in that plural quantifiers incur no additional ontological commitments other than those incurred by the first-order ones. This view is supported both by our intuitions and by how Boolos provides semantics for traditional plural logic. It is also *topic neutral*, meaning that the validity of its rules and axioms does not depend on the subject matter they are applied to. Furthermore, a defining aspect of the logicality of plural logic is its *epistemic primacy*, in the sense that the logical notions can be grasped without appealing to non-logical ones, and similarly, the logical truths are true independently of any non-logical ones. Whether plural logic can indeed be classified as true logic remains an unsettled matter, but even if it is, the ideas required to make sense of higher-level plural logic could not afford logicality for *that* framework. However, one thing can be salvaged, which we will return to later, and that is its *ontological innocence*.

Particular philosophical applications include logicism (Hewitt, manuscript), set theory (Lewis, 1990), mathematical nominalism (Hellman, 1989, 1996; Burgess and Rosen, 1997), metaphysics (Cotnoir and Baxter, 2014), and the philosophy of language and formal semantics. Although we will not deal with any specific application of higher-level plural logic in this thesis, I hope that the way I argue for its legitimacy shows that there will be philosophical benefits to be reaped from it. A testament to this is its application to logicism by Grimau (2018).

In what's to come, one should always have in mind that our audience is comprised of *the pluralists*, i.e., those who believe in the legitimate usage of plural logic on both the syntactic and the semantic front. These are the ones whom I will be trying to convince of the legitimacy of higher-level plural logic. As it seems, there is already a plethora of opinions on it in the pluralist camp, which is why we turn to them in the following section.

2.2 Higher-Level Plural Landscape

Pluralists disagree on numerous fundamental issues regarding plurals themselves, let alone higher-level plurals. There is a great variation of opinions, just as one would expect from a philosophical issue, regarding the ontological status of plurals, the preferred framework to reason about them, whether a plural syntax needs to be accompanied by singular or plural semantics, and much more. Putting such issues aside for the time being, we will try and understand the reservations of pluralists of all kinds regarding the intelligibility of higher-level plurals. In due time, I will put my cards on the table and outline the motivations that guide my conception of plurals and how that informs their higher-level counterparts. This section will be broken down into three parts: first come the skeptics, then come ones who argue for dispensability, and finally come the believers, i.e., those who take higher-level plurals at face-value.[5]

The initial obstacle one encounters when trying to make sense of the debate on higher-level plurals (HLP) is understanding precisely what one means by HLP. Then, the problem becomes even thornier as people with conflicting methodologies and aims try to answer different questions, but often even the same question in conflicting ways. The main methodological difference concerns the role that natural language has to play in this investigation. The aims differ as to whether one tries to understand HLP as an *asset* the pluralist could take advantage of or whether one believes that HLP is a problematic notion, casting doubt on pluralism itself. I will make appropriate comments whenever such issues arise and try to clarify the reasoning of the authors raising those issues. With that in mind, let us turn to *skepticism*.

I. Skepticism The skeptic I will introduce here comes in many shapes and forms. Her opinions have been voiced by very prominent philosophers, as we will see shortly, and her thoughts will help us carve a clear path to an intelligible conception of higher-level pluralism. Her worries start immediately when trying to understand what one could mean by *higher-level plural*.[6] Whatever it may mean, it surely would be the result of some kind of *iteration*. The question she first posits then is the following (Linnebo, 2003; Uzquiano, 2004; Linnebo and Nicolas, 2008):

Iteration: Can the move from the singular to the plural be iterated?

Our skeptic winces at the vagueness of this question, but concurs that there is some merit in initially putting the question in vague terms and trying to gradually precisify it. Maybe an alternative formulation, of narrower scope, could prove to be enlightening:

Iteration*: Is there something that stands to the plural, the same way the plural stands to the singular?

The existential quantifier in the question should be taken with a grain of salt. We do not want to rush in and make heavy ontological claims from the get-go, but we also do not want to rule them out. Maybe there really exists, in a substantial ontological sense, something like that, or maybe there isn't; settling this matter is also part of our discussion on higher-level plurals. As I mentioned in the previous section, I will be using existential and singular language for plurals, but it should always be taken as a mere *façon de parler* unless indicated otherwise.

Nevertheless, this question, our skeptic thinks, is not a singular question, but rather a schema depending on what one substitutes for 'something'. The question can be cast in terms of plural quantification, plural terms, or plural reference, thus providing us with three instances of the schema:

Q-Iteration: Is there any *quantificational* device that stands to plural quantification as plural quantification stands to singular/first-order quantification?

[5] This part has been greatly influenced by the relevant discussions from Wagner (2015, Chapter 2) and Grimau (2018, Chapter 3).

[6] Henceforth, I will use HLP as a catch-all term referring to terms, quantifiers, and reference, when no distinction is important to the overarching point. Whenever it is, I will qualify HLP accordingly.

T-Iteration: Are there *terms* that stand to plural terms, as plural terms stand to singular terms?

R-Iteration: Is there any *referential* device that stands to plural reference as plural reference stands to ordinary, singular reference?

Our skeptic realizes that tackling any one of the questions cannot be carried out in isolation. Answering one of them often, but not always, indicates an answer to the others. For instance, **Q-Iteration** lends itself to an immediate negative answer in light of a negative answer to **T-Iteration**, but not necessarily *vice versa*. A positive answer to the latter, coupled with an account of semantic reduction of higher-level plurals to other singular and/or plural resources, is compatible with a negative answer to the former.

Now our skeptic turns her attention to a subtlety. There was a change in the way **Iteration** was initially formulated and how it was later clarified and split up into numerous questions. The initial question involves the modal *can*, while the latter four questions are stated in a way that investigates whether there is a fact of the matter that will serve as either truth- or falsity-maker for them.

This fact of the matter turns out to be intimately related to the empirical question of whether natural language and its usage involve such terms, quantificational, and referential devices as those indicated by the iteration-questions. The questions more relevant to my aims would be the more open-minded ones **MQ-Iteration, MT-Iteration,** and **MR-Iteration** asking whether such resources *can* be understood in a non-trivial way. Once again, whenever something important hinges on the differences between the actual and the M-variant of a question, I will mention it explicitly. The connection between the two variants of each question has been very eloquently put by Gabriel Uzquiano. He writes:

> If English contained quantifier phrases, which, one could argue, behaved like plurally plural quantifiers, then perhaps that would be some evidence, albeit inconclusive, for the coherence of plurally plural reference. In the absence of such evidence, however, advocates of plurally plural quantification must first make it plausible that plurally plural quantification could intelligibly be introduced into the language. (Uzquiano, 2004, p. 439)

The idea then is that a positive answer to an **Q-Iteration** provides defeasible evidence for the legitimacy of **R-Iteration**. In the absence of such evidence, then, one needs to turn to the modalized versions of the questions and try to make sense of the relevant notions. I want to add to this that even an outright, negative answer to such a question, not a mere *lack* of evidence in favor of HLP but the existence of evidence against them, would still not settle the matter; it simply makes a positive answer to the modalized questions a more difficult feat. Natural language evolution is an empirical phenomenon largely governed by conventions, so the unavailability of HLP does not necessarily mean that we could not make sense of HLP in a way that motivates its use in a logico-philosophical context.

Moreover, even if natural language includes some HLP resources, it will surely not employ plurals of all finite levels. So, for some higher-level plurals we could have defeasible evidence for the positive answer to *Iteration*, but for others we would not. The importance, however, should not be put on the levels themselves; the evidence we are looking for is one for an intelligible iterative move akin to the one from the singular to the plural. The hierarchy that may arise from such an iteration, if there is one at all, is a matter depending on additional assumptions (see Chapter 4).

Having said that, it is not entirely clear why we should think that natural language does *not* employ HLP resources. Surprisingly, even people who have expressed sympathy towards HLP, such as Rayo (2006), believe that natural language lacks such features. To counter this commonplace view, numerous scholars have offered linguistic evidence in favor of the existence of HLP terms in natural language.

Linnebo and Nicolas (2008) make a case about superplural terms in English, namely terms that stand to plural as plural ones stand to singular terms, but they express skepticism about a relevant notion of quantification. Lists of plurals, such as "The Beatles, the Led Zeppelin, and the Pink Floyd" or "the players and the coaches", have been invoked as candidates for superplurals in English. The syntactic fact of adjoining multiple pluralities to one another is not enough to establish that semantically these lists should be understood as superplurals, which is one of the reasons that Uzquiano would call this evidence *inconclusive*.

Similarly to the case of plurals, one compelling reason to think so stems from the way in which these lists get predicated over. Collective predicates, for instance, applying to these lists could indicate that the best way to regiment these lists semantically would be by using a superplurality. For instance, if I uttered "The Beatles, the Led Zeppelin, and the Pink Floyd gave a joint concert", then the predication does not distribute to each of the bands (pluralities), and lends some support to the list being genuinely superplural.

We will return to such points in the next part, where we will discuss in detail the (in)dispensability of HLP. The matter of (in)dispensability is also important to our skeptic. Consider the particular case of HLP terms. If we find apparently HLP terms in natural languages, but can dispense with them semantically in favor of already accepted singular and plural resources, there is no need for HLP semantics. However, our skeptic concurs that this is no knock-down argument against the intelligibility of HLP; it very well may be the case that HLP are dispensable in the context of natural language, but also that a coherent conception of them can be successfully outlined.

Returning to natural language, in Grimau (2018, 2019) one can find multiple reasons why one would think that natural language employs some higher-level pluralities. In Linnebo (2022) and Grimau (2018, 2019), one finds similar cases for the existence of such resources in other natural languages, such as Icelandic, Finnish, Estonian, Khamtanga, Breton, and Classical Arabic. In fact, in some of these languages one can find *lexicalized* HLP terms.

Having acquired a better grasp of the questions and their relations to natural language, our skeptic concludes: natural language is neither enough to *establish* nor to *refute* the intelligibility of HLP.

For these reasons, she turns to the more conceptual arguments regarding HLP. She immediately voices her reservations about higher-level plural quantification, and finds herself considering some of Ian Rumfitt's worries:

> We can all understand what is meant by saying that a speaker is quantifying plurally over some objects. But what is meant by saying that he is quantifying plurally plurally over objects? The only sense I can make of this is as saying that one is quantifying plurally over objects with many members. But again, that brings in ontological commitments not incurred by quantification over the first-order domain. (Rumfitt, 2005, p. 97)

One who finds oneself aligned with Rumfitt would deem HLP quantification as somehow problematic. The problem is not that it is unintelligible, but rather that it can only way to make sense of it is by an iterated form of plural quantification: plural quantification over pluralities. But if a pluralist agrees with the criterion of ontological commitment of Boolos (1984) positing that we are committed to the existence of the individuals over which we plurally quantify, then quantifying plurally over plurals would ultimately commit one to the existence of pluralities.

The principal worry voiced here, also one that is very eloquently discussed by Linnebo (2003), is that accepting higher-level plural quantification would ultimately distort our initial intuitions about pluralities. In the same spirit, McKay argues that a legitimate form of higher-level plural quantification would involve singularizing devices. His "perplural" corresponds to what I refer to as *superplural* or second-level plural. He writes:

> The language of perplurals can be understood if we build in singularizing assumptions. Given those singularizing assumptions, perplurals are expressible by ordinary plurals applied to the 'higher-level' objects that the singularizing process introduces. If we do not make the singularizing assumptions, then there is no evident way to understand perplurals [. . .]. (McKay, 2006, p. 138)

Linnebo has argued that this iteration, where the plural has to be singularized to be plurally quantified over, puts pressure on the ontological innocence of the plural. He continues by expressing that the pluralist can be motivated by similar arguments as for pluralism to accept this iteration, thus putting pressure on the Boolosian pluralist position (Linnebo, 2003). Especially given the correspondence between plural quantification and second-order quantification over monadic predicates (Boolos, 1985), it is but a stone's throw away to start thinking about whether this quantificational device can be iterated. But David Lewis voices similar reservations about that as well:

> [Boolos' view of the relation between second-order and plural quantification] hints that the third, fourth, and higher orders cannot be far behind but what might plurally plural quantification be? (Infinite blocks of plural quantifiers? - That will be only a skimpy third order, and no start at all on the fourth.) (Lewis, 1990, p. 70-71)

Our skeptic sympathizes with these worries. She feels that the straightforward answer to **Iteration** is problematic for the traditional conception of pluralism. This does not entail that HLP quantification is inherently faulty, but it seems that this conception of it would not be in line with pluralism as we know it. Plurally quantifying over plurals commits us to plurals, as would a notion of superplural reference understood as direct reference to plurals.

The skeptic finds herself at a crossroads. She believes that these arguments condemn the naïve answer to *Iteration* and does not see immediately any other way to understand HLP. She considers a point made by Rayo (2006), who agrees that *this* conception of superplural quantification is problematic, and proposes to think of the HLP quantifiers as *sui-generis* quantification over individuals.

The problem is that it is not clear what this form of quantification would amount to. If superplurals contain plurals, then a corresponding notion of superplural *reference* should take into account that these individuals form pluralities before being referred to. This idea will be important for this project later on and will be fleshed out further, but once again, our skeptic is not entirely certain that this would lead to a legitimate conception of reference. She finds herself in agreement with the old writings of Simons (1982), who later on changed his mind (Simons, 2016).

That skepticism is captured in the following passage, where Simons' "manifold" corresponds to our plurality:

> We might look upon different plural expressions as relating to the objects in a different way. Consider the case of the chairs against the wall. It might be that we should wish to look upon the role of different referring expressions like this: ’these chairs’, ’these pairs of chairs’, ’these pairs of pairs of chairs’. [. . .] It is one thing, however, to draw diagrams [...] showing how we may group and subgroup individuals into larger or smaller groups: it is quite another to think that we have made any semantic sense of these diagrams in terms of higher-order manifolds. (Simons, 1982, p. 192-193)

Thus, our skeptic concludes that this idea is surely not a non-starter, but the burden of explaining its merit is on the side of the higher-level pluralist. In the meantime, one is justified in one’s skepticism towards this notion.

The final problem our skeptic turns to has to do with the notion of intersubstitutivity *salva veritate*. She thinks, and finds herself in agreement with Linnebo and Nicolas (2008), Florio (2010), and Ben-Yami (2013), that a naïve approach to HLP quantification and reference as quantification over and reference to individuals suffers from a serious defect. To illustrate this we allude to a natural language example involving lists of plurals:

(1) My children and your children played against each other.

(2) The boys and the girls played against each other.

Suppose that we are in a scenario where both of us have at least two children, at least one boy and a girl, and that indeed my children played against yours. If the superplural terms in (1) and (2) are supposed to co-refer to the individual children, then it is clear that substituting the terms to move from (1) to (2) does not preserve truth-value. The takeaway: either HLP reference as direct reference to individuals is incoherent, or a better notion of (co-)reference is required. For this reason, I will later argue for a view of reference that is reference to *individuals*, but not directly; this results in a natural criterion of co-reference which avoids these drawbacks.

And with that, we conclude the presentation of our skeptic’s worries. She initially pondered over how to intuitively grasp what higher-level plurals are. This was captured by the three different questions of iteration - in terms of quantification, terms, and reference - and their modalized versions. She then considered the arguments from natural language and how they impact the answers to each of these questions, before turning to the conceptual underpinning of HLP. She voiced her uneasiness with numerous off-hand answers to the iteration questions, but did not immediately conclude that this settles the intelligibility of HLP. Contrarily, she concludes that there are numerous challenges that the proponent of HLP has to meet and that the burden of proof is on them.

Having set the stage, we will now focus on two alternative kinds of approaches to HLP. We will first see, both formally and conceptually, how some are trying to dispense with HLP in favor of other singular and plural resources, and then we will turn to the views that take HLP at face-value.

II. Dispensability: The idea of the (in)dispensability was entertained by our skeptic and was considered an important option to be explored. As it turns out, the dispensability of HLP is a well-trodden path with numerous staunch supporters whose work I will present in detail. Proponents of dispensability are willing to accept in part the legitimacy of HLP terms, but deny their essentiality. The proposals can be roughly divided into two categories: (a) the ones arguing that HLP terms can be consistently *paraphrased away*, and (b) the ones arguing for their *semantic needlessness*. In what follows, the formal frameworks accompanying a philosophical position will also be introduced.

The point of interest here is the presence of HLP terms in natural language and their function therein. In the previous part, I tried to make a case that matters of natural language should not restrict the intelligibility of HLP and are definitely not to be considered as decisive evidence for or against HLP.

Having said that, I do believe that a conception of HLP needs to be materially adequate, in the sense of being able to account for the natural language phenomena in which HLP apparently occur. To be slightly more precise, even if either of the two dispensability challenges could prove adequate, this does not preclude other ways to understand HLP. What it demands from any such understanding is to prove that it does the job of explaining HLP at least as well as the alternatives. The conceptual aspect of this comparison will be carried out in Section 2.7, but the next chapter is devoted to a specific formal comparison of the best conceptual alternative to my preferred view of HLP.

Paraphrase

Paraphrasing attempts often try to show that there is a systematic way to move from a certain kind of discourse to another in a way that ensures that some semantic facts, for instance meaning or truth-value, remain unaltered, but other needless notions are dispensed with. I am certain that the paraphrase is intended to preserve truth-value, but I am less certain about the preservation of meaning.

A paradigmatic case, and a very early one, of a paraphrase attempt was proposed by Black (1971). Black initially proposed a nominalistic conception of sets by roughly equating them to the pluralities of their elements; the idea was that when one invokes a set, one is only committed to the existence of the elements of that set instead of also being committed to the set itself. He then explains that set-talk serves a specific purpose; its main function is *convenience* by allowing us to indirectly talk about the members of a set through mention of the set itself, in cases where we may not know exactly the members of the set or when they may be too many to list.

Following Black, talk of superplurals would correspond to talk of sets of sets, which, due to similar ontological arguments as the ones considered above, would end up committing one to the existence of sets. However, Black believes that *in principle* there is a way to reduce talk of sets of sets to "assertions about sets simpliciter" (Black, 1971, p.633). Unfortunately, he offers no such recipe for carrying out this task, but also no reasons to believe that this would even in principle be possible.

Obviously, a pluralist need not follow Black as far as to take sets to function as shorthand for pluralities; after all, numerous arguments could show that plurals are conceptually quite different from sets.[7] However, the analogy can be drawn simply in terms of plurals. If superplural terms can always be paraphrased away by only using plural vocabulary, then the former would simply be redundant.

Other proposals have, rather implicitly, taken this line of thought to heart. For instance, McKay (2006) proposes understanding superplural reference, and in general HLP reference, as plural reference to singular entities that can have other objects as their members. McKay, for instance, would take groups, teams, or pairs to be such singular entities.

[7] See Incurvati (2020, p. 2-11) for a relevant discussion about the conceptual differences between sets and plurals.

There are numerous reasons why both paraphrase strategies prove to be inadequate.[8] I want to present only two arguments that purport to show why these strategies are inadequate, since they fit well within the dialectic of my proposal. The first one, which we alluded to after our brief discussion of Black and which will also be important later, is that we cannot offer a paraphrase in all cases. The second one purports to explain that their paraphrase strategies are not in line with the pluralist spirit.

Regarding the former, suppose that we want a paraphrase for a sentence such as[9]:

(3) The cardinals, the ordinals, and the transitive sets overlap.

According to our previous examination, this sentence involves the HLP term "The cardinals, the ordinals, and the transitive sets", which the advocate of paraphrazing would want to change into another term by using entities with members. Sets, on pain of paradox, are obviously out of the question. One could consider adding *proper classes* to one's ontology and try to talk about the class of the cardinals and the classes each containing one of the rest. However, we just pushed to problem to the next level, because we can then form statements regarding the proper classes and the non-proper ones, such as "The proper classes and the non-proper ones are disjoint".

Two approaches can be sketched here. One would either be led to adopt a hierarchy of classes to deal with this. This is a move that results in a substantial ontological proliferation, whose success hinges on the lack of a final level in the hierarchy. For if there were such a level n, then statements like the one above, involving apparently plural reference to n-level class would need $(n+1)$-level classes to be formalized. One then would need an indefinitely extensible hierarchy of classes to be able to do what one could accomplish with a hierarchy of plurals.[10] Another way would be to posit a type-free (self-applicable) theory of classes.[11] There are proper class-many proper classes and similarly for non-proper ones, so by only invoking classes, one would be able to offer a paraphrase.

I do not see any conclusive way of arguing for one over the other in general. I do believe, though, that for pluralists neither the choice of an indefinitely extensible hierarchy of classes nor of a type-free notion of class is the best course of action. The uneasiness with the former approach does not stem from the need for a hierarchy, but rather from the singularizing that takes place at each level. Similarly, in the latter approach, the problem is once again the singularizing. Consider, for instance, the pluralist insight that one can refer plurally to the cardinals without invoking the class including them. For pluralists, uttering (3) should be captured in a pluralist manner without the need to invoke singular terms and unique entities, just as plural terms and reference should not be paraphrased in terms of sets. Even if one can offer a consistently well-applicable paraphrazing method in this way, the pluralist demands a more straightforward regimentation of this discourse. For this reason, I think that such approaches are ultimately null and void for the pluralist.

Semantic Dispensability

Having argued that paraphrasing is deemed unfavorable for the pluralist, we now turn to the views regarding semantic dispensability. The proponents of such a view typically concede that there are

[8]See Grimau (2018, p. 114-119) and Wagner (2015, p. 31-32) for some relevant arguments.

[9]The set and class examples are from (Grimau, 2018, p. 128)

[10]Instead of invoking classes outright as entities, one can argue on potentialist set-theoretic grounds that talking of the cardinals is in fact talk of a set-sized portion of the cardinals. Although this could be an interesting point to pursue, the anti-pluralist worries mentioned later would apply to it as well, so I will not explore it any further.

[11]Maddy (1983) offers such an account.

terms that appear to be HLP terms, but they argue for some semantic reduction of them using already acceptable singular and plural resources. We will focus mainly on two views: (a) the *articulated reference* view of Ben-Yami (2013), and (b) (generalized) cover-based semantics. The former is a conceptual view about how to dispense with cases seemingly invoking HLP reference. The latter originated in the formal work of Gillon (1987, 1992) and Schwarzschild (1996) using only singularist-friendly machinery, which has recently been adapted to be more appealing to pluralists by Payton (2025), Nicolas and Payton (2025), and Nicolas and Payton (forthcoming).[12]

Let us start with Ben-Yami's work. Ben-Yami's conceptual framework to understand HLP reference rests on the notion of *Articulated Reference*. Ben-Yami believes that HLP reference is redundant; HLP terms refer plurally to individuals (plurality) simply because they "contain" other terms that refer to some of these individuals (subplurality). That notion is made precise as follows:

Articulated Reference: A referring expression can refer to a plurality by virtue of containing other referring expressions that refer to some of that plurality.

In his words, the plural term "Jack and Jill" refers to Jack and Jill in virtue of containing the name "Jack" referring to Jack and the name "Jill" referring to Jill. Hence, the way in which "Jack and Jill" refers is articulated. This is exactly what occurs in the case of higher-level plural terms as well. Consider the following example, due to Oliver and Smiley, discussed by Ben-Yami:

(4) The joint authors of multi-volume treatises on logic are Whitehead and Russell, and Hilbert and Bernays.

Since we are conjoining plural terms, it seems that on the level of syntax we are creating a higher-level plural term. This syntactic feature could be taken to lead to a genuinely semantic feature as well: the existence of a higher-level counterpart to plural reference. But Ben-Yami believes that this temptation should be resisted. We can understand examples like these by alluding to articulated reference once again. What seems to be referring to the *superplural* containing the pluralities of Whitehead and Russell, and Hilbert and Bernays, simply articulately refers to the plurality containing all of them. The articulation is due to the fact that the plurality contains the two referring subpluralities of Whitehead and Russell, and of Hilbert and Bernays.

There are numerous problems with this view. Grimau (2018, p. 123) maintains that it is not clear how other terms, not equivalent to lists, could be considered as articulated expressions. Wagner (2015, p. 19) makes a similar observation that rests on the vagueness of the notion of "containment", which seems to be a syntactic feature whose properties are not clear. In the case regarding lists of plurals, then this syntactic notion of containment is clearly understood through *Russell and Whitehead* clearly being a part of the referring expression "Whitehead and Russell, and Hilbert and Bernays". In other cases, where the syntactic structure of a linguistic item does not determine whether it has any referring parts, it is not clear how it can be considered a case of articulated reference.

Grimau also explains how, under this view, ordinary plural reference seems to be reducible to singular reference, at least in the cases where a plural term is a list of singular terms (Grimau, 2018,

[12] An alternative which we will *not* focus on is the *mereological* approach to plurals proposed by Link (1998, 2002), and more recently advocated for and employed by Champollion (2017). Although this approach enjoys a privileged status in the linguistic literature, it formally rests on a set-theoretic meta-theory, which ultimately goes against pluralism as we have outlined it. For that reason, we do not go into it at all. For a discussion on comparing the pluralist and the mereological approaches to plurals, see Florio and Nicolas (2020).

p. 124). To illustrate this point, consider the term "Jack and Jill", which does not refer directly and plurally to *Jack and Jill*, but does so in virtue of "Jack" referring to Jack and "Jill" referring to Jill. This, although not necessarily a non-starter, seems to also be at odds with the pluralist's conception of plural reference as a sui-generis referential device. This is not a knock-down argument in any case, since it is compatible with pluralism that there be *some* connection between plural reference and singular reference, so long as it is not an outright reduction of the former to the latter. It is simply on Ben-Yami to explain what his "in virtue of" means.

The reason why this proposal should not be too appealing to the pluralist is simply its material inadequacy. It accounts only for a very limited number of HPL terms (lists), a very anti-climactic result for a reductive approach to HLP reference.

With that in mind, we can focus on another approach: cover-based semantics. In a series of papers, Nicolas and Payton propose a view of higher-level plurals extending already existing work of Linnebo and Nicolas (2008), arguing that higher-level plural terms are licensed by the predication involving them, and of the cover-based semantics of Gillon (1987, 1992) and Schwarzschild (1996). As Nicolas and Payton put it:

> The intuitive idea behind cover-based semantics is that the application-conditions of a plural predicate can be sensitive to how the participants in a conversation 'divide' the referent(s) of a plural term.

The division happens based on a given *cover* determined by the context. Formally, a *cover* C of a non-empty set x is a non-empty set of subsets of x whose union is x. That means that every member of x is a member of a member of C. The empty set should also not belong to C, i.e., $\emptyset \notin C$. A plural predicate, then, does not apply directly to the referent of a plural phrase, but to each member of the cover.

To understand this idea better, and to also understand the reservations of Nicolas and Payton (forthcoming), and Grimau (2018, 2019) we move to some examples. Consider the one given in Nicolas and Payton (forthcoming):

(5) The students and their teachers met in adjacent rooms.

Under one reading of this sentence, we would have that all the students went to one room and all their teachers to an adjacent one. The cover would then be the set $\{\{the\ students\}, \{the\ teachers\}\}$. The predicate *met in adjacent rooms* would then have to be true of each of the two sets contained in the cover, which is obviously wrong, as the predication is supposed to be collective. Using a similar example (Nicolas and Payton, forthcoming, p. 10), which gets the predication correctly, shows that the relevant reading of the sentence is lost in the process. The example is:

(6) The philosophy students and their teachers, and the history students and their friends, met in adjacent rooms.

Now, the relevant reading is that the philosophy students and their teachers met in adjacent rooms, and so did the history students and their friends. According to the cover

$$\{\{the\ philosophy\ students,\ their\ teachers\}, \{the\ history\ students,\ their\ friends\}\}$$

the predication distributes correctly over the two elements of the cover. However, it does not capture the relevant reading of the sentence since *met in adjacent rooms* is now true of two sets of individuals, which does not respect the way in which we understand the predication.

For this reason, Nicolas and Payton offer these two desiderata for cover-based semantics:

> In short, in order to account for these varieties of apparently higher-level predication, we need a cover which does two things: ensure that the predicate if true *of* the right things (i.e., that we get the correct collective or distributive reading); and retain the structure which is intuitively relevant to the truth of the sentence. (Nicolas and Payton, forthcoming, p. 10)

Apart from the apparent inadequacy of standard cover-based semantics to provide an accurate regimentation of some sentences, pluralists, such as Nicolas and Payton, find them lacking on the conceptual level as well. On the cover-based approach, singularist resources such as set theory, are employed in the meta-theory to analyze plurals; but have already seen and argued extensively that this would not be preferred by the kind of pluralist that we are interested in.

This conceptual drawback leads to a formal obstacle as well, precisely due to understanding plurals as sets. This implies that plural terms denoting class-many individuals plurally will not be up for a formal treatment using cover-based semantics.

Nicolas and Payton propose an amendment to standard cover-based semantics in order to deal with these shortcomings: the notion of a cover be *generalized* so that instead of dividing a set into subsets, it divides a given plurality into subpluralities. More formally:

Definition 1 (Generalized Cover). *δ is a generalized cover of aa if there are some indices ii such that:*

(i) for every i among the ii, $\delta(aa, i)$ are among aa;

(ii) for every x among aa, there is some i such that x is among $\delta(aa, i)$.

To capture superplurals, one need only make sure that the cover is extended in a way that the subpluralities of the initial plurality are then themselves divided into further subpluralities. In their paper, Nicolas and Payton go into great detail to show how different kinds of collective and distributive predication can be captured with this machinery.

Precisely because they are using plurals instead of sets, the previous qualms about the size of the denotation have been dealt with. However, a new problem arises within their approach stemming from a plural variant of Cantor's Theorem (Florio and Linnebo, 2021), that the plurals are strictly more than the individuals, and their proposal to index the plurals with individuals. The individuals are simply not numerous enough to index every plurality.

They propose three ways to deal with this. The first two are disregarded by them (Nicolas and Payton, forthcoming, p. 26-27), and I agree that they provide unsatisfactory solutions to their problem. The third, and more promising solution, adopts a rather unconventional approach to indexing and adapts the definition of a generalized cover so that pluralities themselves can be considered as indexes. The indexes ii and subpluralities thereof are now both considered indexes. Therefore, individuals and plurals together will be large enough to index all the plurals, as long as there are as many individuals in ii as in the plurality aa for which the cover is defined.

Note, however, that this is not the only way to define the generalized covers. Payton (2025), for instance, defines them as plural relations that relate a plurality to its subpluralities, without any need for indexes in the way they were used above. In any case, for now, this seems like the most promising option for the proponent of the dispensability of HLP. In due time, we will look into the differences between this approach and mine in detail. Before doing so, we will go into some of the proposals that take HLP at face-value.

III. Face-Value:

This is the third and final part of my sketch of the HLP landscape. In this part, we will focus on the informal conception of HLP reference that has been outlined by Grimau (2018). I aim to understand her philosophical insights into HLP reference, in order to later compare them with mine. The choice of formal logic to accompany those insights is of less importance for now, because the general frameworks employed by Grimau and me are very similar; they only differ with respect to the semantics. The syntax of the logic we both posit is a cumulative extension of the logic developed in Rayo (2006). Other alternatives in the literature include the logics of Oliver and Smiley (2016) and Florio (2014a). On the formal side, the main comparison that will be of interest will be between HLP logics and the non-HLP ones; however, in Chapter 4, we will briefly carry out a limited comparison between typed and untyped approaches to HLP logic.

The main idea behind face-value approaches to HLP is that there are legitimate notions of HLP reference and quantification, both of which need not be reducible to any form of plural or singular reference and quantification. Rayo, for instance, explains that higher-level plural quantification has to be considered as a *sui-generis* form of quantification over *individuals*, rather than an iterated form of plural quantification or some sort of first-order quantification over special entities. Although this is more of a slight nudge in the right direction rather than a clear-cut proposal, a successful precisification of it would help assuage some of the skeptic's worries presented previously. For instance, an appropriate conception of quantification over individuals may not necessarily lead to the ontological commitments that an iterated form of plural quantification would lead to. Higher-level plural quantification for the cover-based semanticists can be understood as quantification over covers; existential superplural quantification gets reduced to existential quantification over a cover, meaning more or less that there is a way to split a plurality into subpluralities.

The matter of quantification has received little to no attention in the literature on HLP, with the exception of Linnebo and Nicolas (2008). The bulk of the attention has been on HLP reference since an understanding of it would provide us with the intended interpretation of the quantifiers. Existential superplural quantification could be understood by the locution "there are individuals superplurally referred to". The main task is to elucidate what it means precisely for individuals to be HLP-referred to.

Grimau maintains that HLP reference is a certain species of what she calls *restricted reference*. Restricted reference is a mode of reference that makes salient a certain *aspect* of the object(s) picked out. HLP reference then should be understood as follows (Grimau, 2018, p. 138):

> HLP terms denote some objects *under their aspect of being organised in a certain way.*

The aspect of being organized in a certain way is what Grimau calls a *cluster*. In fact, she advocates for a view of clusters as a certain kind of plural properties and relations. Consider once again example (6) above. The superplural term "the philosophy students and their teachers, and the history students and their friends" refers to all the individual philosophy students, their teachers, the history students, and *their* friends under the cluster that takes them to be organized in the two pluralities, the first of which contains the philosophy students and their teachers, while the latter contains the history students and their friends.

In this way, she is also able to argue against the claims of the unintelligibility of HLP reference based on a lack of intersubstitutivity *salva veritate*. Remember the example we mentioned before:

(1) My children and your children played against each other.

(2) The boys and the girls played against each other.

The two superplural terms are not co-referring under Grimau's proposal. In the first case, "my children and your children" refer to the individual kids *under the cluster* that divides them into my and your children; this obviously differs from the cluster that divides them into the boys and the girls. In general, Grimau has put forth a proposal that gives a profound elucidation of the notion of HLP reference.

However, this is not to say that Grimau's proposal is impeccable. Specifically, I am uncertain about her treatment of the **Iteration** questions we asked before. She argues that the proper way to understand how the move from the singular to the plural gets iterated is by considering limiting cases. Singular reference is a limit case of plural reference, and as such, n-level plural reference is a limit case of $(n+1)$-level plural reference.

I believe that this misses the point of the iteration that happens between each successive level. My point will become clearer when I argue for my preferred answer to the **Iteration** question, but for now, I want to present one of the arguments against Grimau, which has been hinted at by Wagner (2015) and is based on remarks of Oliver and Smiley (2016).

The first step is putting up clear-cut boundaries between singular and plural reference. Singular and plural reference differ with respect to the number of things they refer to: the former kind involves reference to unique things, while the latter to multiple things. When it comes to the number of things terms can refer to, plural terms and singular terms offer both an *exclusive* and an *exhaustive* treatment.

Two comments are in order here. Firstly, if an answer to **Iteration** is to offer any elucidation to HLP, then Grimau's answer to reference-iteration will simply not do. We are supposed to explain what superplural reference is by taking it to be something to which plural reference is a limit case. If it is to be a limit case with respect to the number of referents, then there is no such thing, at least insofar as the difference between plural and singular is exclusive and exhaustive.

Alternatively plural reference to individuals is superplural reference to individuals through the empty cluster (if that exists at all). The empty cluster is not to group the individuals in any way, and so allows for a direct reference to them. It should also be distinguished from the identity-cluster that would just group them into the plurality of all these individuals. This difference amounts to that between a (cumulative) superplural containing only some individuals and a superplural containing the plurality of all the individuals. Plural reference is a direct reference to individuals, which, as I understand it, according to Grimau, should be equated with superplural reference to the same individuals. This is a point which will be discussed in Chapter 2.4 as it regards the cumulativity of superplurals, but I do not think that it is an indicator of the iterative move behind the intelligibility of superplurals, but rather a by-product of a given conception of reference for them.

Secondly, as Wagner argues, the difference between how HLP terms of different levels refer should be cashed out in terms of *levels of reference*. She does not go into details to precisify this further, but I take her intuition to be correct. Even then, I will fail to see that Grimau's answer can be considered adequate for the purposes of the iteration question.

Insofar as we take the iteration question to be an important aspect of understanding HLP, then answers to it, even partial ones, should serve as a guide to pinning down the meaning of HLP. Grimau's answer fails to deliver. Another comment, which will be pursued further shortly, is of a methodological nature. Grimau's version of HLP reference invokes a wide assortment of additional tools, such as aspects, clusters, plural properties, etc.. This is by no means a substantial critique against her approach, but we will see later that my proposal will be able to deliver the goods without the need for such a diverse formal machinery.

Finally, a remark should be made about the similarity between Grimau's approach and the generalized cover-based approach. In at least one conception of the cover-based approach, that of Payton (2025), covers are to be understood as plural relations relating a plural to its subpluralities. This conception of a cover is very similar, if not the same, as that of a *cluster* as a plural property. This conceptual similarity stops abruptly because Grimau chooses to understand these plural relations as contributing to genuinely superplural reference, while the proponents of the cover approach use them to dispense with superplural reference. My conception of HLP reference, to which we will now turn, will be different enough from both of these conceptions. In light of the formal arguments in Chapter 3, formally, our preferred frameworks will not be too different from one another either.

2.3 Combinatorial Reference

Having sketched a by-and-large complete picture of the philosophical and linguistic landscape on higher-level plurals, it is now time to offer my proposal of how to understand HLP reference. Before going into the notion of reference I want to describe, it is important to make my desiderata here clear. Firstly, I do not presume that the kind of reference I will propose for HLP is the kind of reference that is *actually* employed by human beings. It is rather a notion we *can grasp* based on a series of relatively uncontroversial assumptions; grasping it can allow us to use a logical framework with a plethora of potential applications. In addition, my conception of HLP reference should be consistent with the use of HLP in natural language; for the levels for which HLP terms figure in English, this notion of reference should be able to explain how these terms are used in an attempt to refer to some objects.

Moreover, I want to propose a conception of HLP that inherits as much as possible from plurals on the philosophical side, e.g., if plurals are ontologically innocent, then so are HLP. As we will see when discussing the crucial notion of *containment*, if plurals are ontologically innocent, then so are HLP. Finally, following Rayo (2006), I believe that higher-level plural quantification should be viewed as a sui-generis form of quantification over individuals, and I believe that my conception of higher-level plural reference motivates a natural account for such a kind of quantification. In parallel, HLP terms will refer to many individuals *in different ways*, rather than considering them as referring to pluralities of the previous level. The formal details will be considered later on, so for now we will focus on the conceptual underpinning of higher-level plurals.

My proposal starts by assuming that, as pluralists, we already possess the concepts of singular and plural reference and quantification. The other notion of which we need an understanding is *containment*. This is one of the first points of departure between my proposal and other HLP ones: I take it that the answer to the **Iteration** question is to be provided in terms of containment. Let us focus for now on the case of superplurals before generalizing the results to encapsulate plurals of all finite levels.

My preferred answer to **Iteration** is the following:

Superplurals *contain* plurals, just like plurals *contain* individuals.

So, for the hierarchy of plurals to come to fruition, we need a process of successive iteration of this notion of *containment* along some well-ordering.[13] Ben-Yami also employed a similar notion in his discussion of his *Articulated Reference*, but his proposal suffered from being too focused on the particular syntactic features of some natural language occurrences of HLP terms. Wagner (2015) voiced the worry that Ben-Yami's notion of containment lacks a robust presentation that can explain its key characteristics, and I believe that my proposal will circumvent these criticisms.

I, on the other hand, take higher-level containment to mimic the plural notion of containment understood as "is among/is one of". One could argue, for instance, that it is part of the content of a superplurality that other pluralities are contained in it. This containment doesn't need to be explicit on the syntactic level (contra Ben-Yami), but a superplurality could not genuinely be considered as such if no pluralities were contained in it.

Hence, a superplurality contains specific pluralities; for instance, in the examples (1),(2) given above with the children, the superplural "my children and your children" contain both the plurality containing my children and that containing your children, but none of the pluralities containing the boys or the girls. For the same reason, it is inherently different from the plurality containing both my and your children.

The notion of containment is what structures the superplurality and makes salient how it is built up from the given individuals one starts with. This doesn't need to be explicit from the syntactic structure of the term; it is a feature of all superplurals nonetheless. That is, a superplural could not be one if it were not built up from individuals first grouped up in pluralities, which are then contained in it. This will become even clearer shortly when we discuss the logic of HLP, which will have specific axioms capturing these features formally.

For now, we take containment to be a fundamental notion. Then, this way of combining individuals into pluralities that are then contained in superpluralities is what motivates the notion of *Combinatorial Reference*:[14]

Combinatorial Reference: A superplural term refers to some individuals in virtue of the pluralities it *contains and how they refer* to individuals.

The goal of this principle is to make explicit that the combinatorial aspect of the structure of a superplural is that which contributes to the way in which the superplural term refers to individuals. If we have the individuals $1, 2, 3$, the pluralities $\{1, 2\}$ and $\{1, 3\}$[15], and the superplural $[\{1, 2\}, \{1, 3\}]$.

[13] A *well-ordering* is usually taken to be a set A endowed with an operation $\leq$, such that $\leq$ is a total order on A, and every non-empty subset of A has a $\leq$-minimal element.

[14] The name is inspired by how Kurt Gödel and Penelope Maddy distinguish between combinatorial and logical conceptions.

[15] Curly brackets are used for convenience here and not to be confused with sets.

The superplural refers to $1, 2, 3$ in virtue of containing $\{1, 2\}$ and $\{1, 3\}$ and the fact that they refer plurally to $1, 2$ and $1, 3$ respectively.

In this way, we solve the problem of co-reference and intersubstitutivity *salva veritate* that was raised before. The answer is simply that the superplural does not refer to $1, 2, 3$ simpliciter, but *superplurally* in a way that is made precise by the plurals it contains and their referents.

The next step is generalization to include HLP reference for all finite levels. The iteration of *containment* simply gives us that any $(n + 1)$-level plurality will contain n-level pluralities. Hence, we have that:

Combinatorial Reference: An $(n + 1)$-level plurality refers to some individuals in virtue of the n-level pluralities it *contains and how they refer* to individuals.

Then, the notion of higher-level plural reference is not necessarily an iteration similar to the move from the singular to the plural, but rather a recursion that generalizes the ways in which one can refer to some objects. That is how we are now able to understand the comments of Wagner (2015) that the difference between the way HLP terms refer is one of *level*. They all still refer to individuals, but in a different way that depends on their level, i.e., which lower-level plurals they contain and how they refer to individuals.

Some more comments are in order. Superplural reference, *qua Combinatorial Reference*, does not stand to plural reference as plural reference stands to singular. It is indeed reference of a different kind. So, once again, the answer to **Iteration** falls on containment rather than on reference. Similarly, superplural quantification, understood simply as quantification binding superplural variables, is a sui-generis form of quantification over individuals and not the result of iteration.

On another note, *Combinatorial Reference* can be grasped by anyone who understands the notions of singular reference, plural reference, containment, and can carry out the induction step for all the natural numbers. The base case, i.e., superplural *Combinatorial Reference*, has been dealt with, and then, assuming that we know how k-level pluralities refer, we know how $(k + 1)$-level pluralities refer as well.

Let us return to one of the previous examples to explain once again how superplural terms refer and give us the correct truth conditions. This lends support to the claim that the notion of *Combinatorial Reference* is materially adequate. Remember the example:

(5) The students and their teachers met in adjacent rooms.

The superplural term "The students and their teachers" refers to the individual students and the individual teachers in virtue of their being first grouped up in the pluralities *the students* and *the teachers*, and them plurally referring to the students and the teachers respectively. In this case, *met in adjacent rooms* is considered a plurality-collective superplural predicate, in the sense that it does not distribute over the pluralities contained in the superplurality. For this reason, we would want to say that this sentence is true of "the students and their teachers", and this is precisely the upshot of taking them to be a superplurality. The predicate is not true of the individual pluralities (as that would be problematic for the same reasons as discussed above related to cover semantics), but true of the superplurality. We will return to this point more specifically during the comparison with the generalized cover approach.

Finally, as indicated above, *Combinatorial Reference* gives a natural criterion for co-reference. If HLP terms refer to individuals in terms of the pluralities (of some level) they contain and how they refer, co-reference is guaranteed by two HLP containing the same pluralities. If we accept that pluralities of all levels are extensional, which I will accept and make it part of the formal framework later on, then having the same pluralities is a criterion of HLP identity. Then, even if two different HLP are built up from the same individuals, they will not co-refer unless they satisfy this criterion.

Thus, we have offered a way to, at least in principle, make sense of the plural hierarchy and how HLP refer according to *Combinatorial Reference*. The rest of the chapter will include the following: (a) a discussion on the cumulativity of the hierarchy, (b) another on the notion of *containment*, (c) the formal logic of HLP, and (d) the conceptual comparison with the alternatives to my proposal.

2.4 Cumulativity

A recent discussion on the adoption of specific type theories, i.e., logics involving some kind of hierarchy of entities with each entity having its own type, has revolved around the *cumulativity* of the types. Such a discussion was initiated by Linnebo and Rayo (2012), and created a long subsequent debate with answers from Florio and Shapiro (2014), and then Linnebo and Rayo (2014). The idea of cumulativity has been severely criticized in Button and Trueman (2022), but for reasons that have to do with the formalization of plural logic in untyped rather than cumulative typed logics. These objections will be discussed in detail in Chapter 4.

Until now, our exposition regarding HLP has focused on the non-cumulative reading of plurals, namely that a plural of some level can contain plurals of the immediately preceding one. In this section, I will argue for the legitimacy of the *cumulative* reading of plurals, where plurals can contain other ones of *all* preceding levels. After establishing the legitimacy of the cumulativity of plurals, I will incorporate it as a key feature of the formalism of plurals. For brevity's sake, I often refer to individuals as 0-level plurals.

First, a note on terminology. One may think that taking a (cumulative) superplural as something that can contain both pluralities and individuals changes the meaning of the term 'superplural'. My answer to that objector would simply be that if she agrees that it makes sense to consider something which can contain both plurals and individuals, she may call it a 'shmuperplural'. I believe that the cumulative conception of a superplural does not differ substantially from the traditional conception of it, and so I choose to keep the same name, with the traditional case serving as a special case of the cumulative one.

Some of the examples as to whether it makes sense for superplurals to be cumulative turn once again on our intuitions about natural language usage. To draw an analogy, if we can view lists of plurals as superplurals, then we can similarly think of mixed lists of plurals and individuals as superplurals. Consider the following examples:

(7) My children and her children played against each other.

(8) The professor and the students held presentations in adjacent rooms.

Similarly to how we argued above, in this case too, the term *my children and her children* is best regimented as a superplural term. Analogously, in sentence (7), the best way to understand the term *the professor and the students* would be by invoking a cumulative superplural term containing the individual *the professor* and the plurality of *the students*. One alternative would be to consider either the plurality of all the students and the professor, but for reasons akin to those given previously, this does not get the correct reading of the sentence because of the way that the predicate *held presentations in adjacent rooms* operates. Someone who accepts the legitimacy of higher-level plural terms and the corresponding notion of reference, i.e., not someone who believes that they should be analyzed away, should be able to accept the legitimacy of cumulative superplurals for more or less the same reasons that compelled her to accept superplurals.

Others, like Ben-Yami or Nicolas and Payton, may object on the grounds that the term is a plural term that exhibits a specific case of articulated reference or refers according to a generalized cover. In that case, both proposals would need to be slightly modified to allow for a cumulative reading, but there is a straightforward way to do so. In the case of Nicolas and Payton, a plurality will not only be divided into subpluralities, but also into individual objects.

If we accept the intuition behind higher-level plurals being able to contain plurals of all preceding levels, then a modification of the notion of *Combinatorial Reference* is in order. In its previous form it read:

Combinatorial Reference: An $n+1$-level plurality refers to some individuals in virtue of the n-level pluralities it contains and how they refer to individuals.

Under the cumulative reading of higher-level plurals, we offer the following modification:

Combinatorial Reference*: An $n+1$-level plurality refers to some individuals in virtue of the singular terms and k-level pluralities, for all $k \leq n$, it contains, and how they refer to individuals.

Additionally, the cumulativity of the plural hierarchy allows us to preserve one of the intuitions behind plurals. According to the standard conception of plurals, plurals need to contain at least two individuals. In other words, there are neither empty pluralities nor singleton pluralities.[16] Of course, there are formal ways in which one can force such concepts, for instance by talking of all the things that are not self-identical, or of all the things that are equal to a specific individual. In my opinion, these are distorting the plural intuitions a bit too much, and so I take that the plural hierarchy has no empty or singleton HLP terms.[17]

Without singleton HLP we would not be able to talk superplurally of *the professor and the students* as in (8) in a non-cumulative hierarchy. The alternatives are also not too attractive: (a) a plurality containing the professor and the students gives rise to wrong truth conditions, and (b) the predicate cannot be distributed, so it must either be a plural or a superplural.

In my opinion, cumulativity is both intuitively appealing on its own and allows for a direct extension of some of the plural intuitions. For that reason, we will incorporate cumulativity into our formal system. We will present that shortly after an interlude on the notion of *containment*, the cornerstone of our approach to HLP reference.

2.5 Containment

This section's importance can hardly be overstated, as its principal aim is to elucidate one of the fundamental notions on which this thesis is based. Understanding how *containment* works is motivated

[16] For formal reasons, the logic of HLP presented later will allow for singleton pluralities, but only as a formal simplification.
[17] Similar ideas have been advanced in (Wagner, 2015, Ch. 3.1.1).

by a multitude of reasons.

Firstly, as we will see later on, the way in which we understand the notion of containment will be crucial in differentiating our referential approach to HLP from the analyzing-away approach of Ben-Yami and simultaneously avoiding the problems associated with it that Wagner (2015) has pointed out.

An additional aim of mine is to uphold many of the intuitions behind plurals themselves, especially their ontological innocence. The ontological innocence of plural quantification rests partly on how we understand plural reference. Similarly, if HLP terms refer according to *Combinatorial Reference*, and that in turn depends on the notion of *containment*, it seems that this latter notion will be the key to understanding the ontological innocence of HLP.

Since the motivation behind the hierarchy of plurals arises from iterating the containment relation, we need to make the features of that relation as precise as possible. *Combinatorial Reference* can be grasped and can be a coherent notion of reference even under alternative conceptions of containment. My goal here is to make sure that I offer a conception of containment that: (a) shows how the difference in reference between plurals and HLP is one of *level* or of a *different way of referring*, and (b) retains some of the features of the individual-plural version of it without resulting in ontological proliferation.

Let us start by looking at how containment works between individuals and plurals. In this case, there is a very tight connection between containment, reference, and ontological commitment. A plural contains individuals, refers to multiple individuals (plurally), and one is ontologically committed only to the individuals contained in the plurals. Since plurals are not entities over and above the individuals that comprise them, it would be amiss to consider containment as a relation between individuals and plurals. As far as I understand it, containment should be represented minimally as a *syntactic* relation between singular terms and plural terms, governed by the relevant axioms of PFO(+). Its *semantic* role is then twofold: (a) it relates single individuals to multiple ones, whenever the individual is among the multiple ones, and (b) it explains how the referents of singular terms contribute to the way in which a plural containing them refers. The latter is supposed to capture the fact that the individuals that are plurally referred to by a plural term can each be referred to by a singular term that is meaningfully said to be *contained* in the plural one.

The containment relation, when generalized to apply to HLP terms, retains the syntactic features of its singular/plural version, and is also allowed to be cumulative. As we will see in the next section, this motivates a very natural axiomatization of higher-level plural logic and the containment relation invoked there. On the syntactic side, things are fairly straightforward and pose no real difficulties, but the question remains: how are we to understand the semantic features of the new containment relation?

If we allowed for a direct generalization of the semantic aspects of the singular/plural version of it, we would run head-first into unwanted ontological commitments. Directly iterating both semantic features would involve similarly confusing locutions as the ones discussed in Section 2.2, for instance by saying that it relates many individuals to many many individuals, or individuals to individual*es*. Analogously, it would possibly result in a form of HLP quantification as plurally plural quantification. In fact, the first semantic aspect of the relation is dropped altogether, precisely because it would lead to such peculiarities.

What we iterate is the second semantic aspect: *containment* is supposed to relate two ways of referring by explaining how one way of referring contributes to another. Using triplets of letters for superplurals, if we said $aa \prec aaa$ - meaning that the aa's are contained in aaa - then the plural aa refers to some (if not all) the individuals to which aaa refers. But since aa is "included" in aaa, the way in which it refers to the individuals contributes to how aaa refers to them. In that way, $\prec$ shows us how the way in which aaa refers to individuals is made salient by lower-level plurals, thus also respecting the insights of Wagner about HLP terms providing *different levels of reference*.

Another example that will help illustrate this: a superplural containing only the two pluralities cc and dd refers to all the individuals in cc and dd in virtue of their being first grouped into the cc and the dd. In some sense, the plural reference of the terms cc and dd to the individual c's and d's is contained in the reference of the superplural containing cc and dd to those same individuals. This is precisely what gives rise to a different level of reference, i.e., reference to individuals but *in another way*.

Prima facie, this manner of understanding containment does not incur any new ontological commitments; it rather extends our ideology by offering several novel expressive resources. The expressive resources are the different kinds of HLP terms, alongside the quantifiers that bind them, which simply explain how one is able to refer to individuals *in different ways*. Still, the only things that exist are the individuals; we just have increasingly complicated ways of referring to them. And with that intuition and understanding in place, we can provide an insight into the intended interpretation of Higher-Level Plural Logic. In the next section, we do so for the fragment of it involving plurals of all finite levels. Outlining a transfinite extension thereof will be the focus of Chapter 4.

2.6 Higher-Level Plural Logic

The final piece of the puzzle is to see how a formal logic of HLP will give us a good insight into how higher-level plural terms, predicates, and quantifiers function. What is more, this framework will also make salient the features of the notion of containment, whose role in *Combinatorial Reference* is of paramount importance.

The first choice to be made here would be between the two main formal contenders for a logic of HLP: the logics of Rayo (2006) and Oliver and Smiley (2016). I opt for the former, precisely because we have a good depiction of how higher-level plural terms function on the level of syntax that mimics their referential features.

That is because Rayo explicitly posits a syntax distinguishing between the different levels of the plural hierarchy, while Oliver and Smiley have a catch-all term for any plural in the hierarchy. Also, (a) Rayo's framework will generalize on PFO(+) which we already saw, (b) there is an immediate way to extend his logic to one that captures a cumulative hierarchy of higher-level plurals, and (c) has a straightforward interpretability result with simple type theory and the cumulative extension of HLPL with cumulative type theory (Rayo, 2006; Linnebo and Rayo, 2012), which will be of interest to us in the subsequent chapters.

This choice, however, is not without unwanted consequences. Having a strictly typed system and being structurally similar to the simple theory of types means that our hierarchy suffers from some similar defects. For instance, there is no cross-type quantification.

Having already argued for a cumulative take on the plural hierarchy, I will present a version of Rayo's system that involves cumulativity. Additionally, we will present the version of higher-level plural logic (HLPL) that extends PFO+, meaning that we will have plural predicates of all levels as well that are considered to be *collective*. If one disregards the plural predicates and the formula formation rules including them, then one gets the system of HLPL that extends only PFO. Let us start with the language $\mathcal{L}_{HLPL}$ before moving to the proof-system:

- **Logical Vocabulary:**
  1. Variables $x^k, y^k, \ldots$ of level $k \geq 0$ corresponding to k-level plural variables.
  2. The logical connectives $\neg, \rightarrow$.
  3. The universal quantifier $\forall$
  4. For each $n \in \mathbb{N}$, a two-place predicate $\prec^n$ for pluralities of level less than n being among pluralities of level n understood as "is/are among of" or "is one/ are some of the".[18]
  5. An identity symbol for individuals $=$.

- **Non-Logical Vocabulary:**
  1. Constants $c^k, d^k, \ldots$ of level $k \geq 0$ corresponding to k-level plurals.
  2. Singular and higher-level plural n-adic predicates $P_n, Q_n, \ldots$.[19]

- **Formation Rules:**
  1. Variables of level k and constants of level k are terms of level k:
  2. For t^0 and u^0, $t^0 = u^0$ is a formula.
  3. For t^m and u^n, $t^m \prec u^n$ is a formula.
  4. For $m_i \geq 0$ and $1 \leq i \leq n$ we that for $t^{m_1}, \ldots, t^{m_n}$ and P_n, $P_n(t^{m_1}, \ldots, t^{m_n})$ is a formula.
  5. If ϕ and ψ are formulas, then so are $\neg\phi$ and $\phi \rightarrow \psi$.
  6. If ϕ is a formula, then so is $\forall x^k \phi$.
  7. Nothing else is a term or a formula of $\mathcal{L}_{HLPL}$.

- **Defined Expressions:**
  1. Higher-Level Plural Identity $t^k = t^l := \forall x^m(x^m \prec t^k \leftrightarrow x^m \prec t^l)$, for $m = max(k, l) - 1$ and either $k \neq 0$ or $l \neq 0$.
  2. Existential Quantification $\exists x^k \phi := \neg\forall x^k \neg\phi$.
  3. The ϵ-operator $a^\alpha \epsilon\, b^\beta := \exists x^\gamma(x^\gamma = b^\beta \wedge a^\alpha \prec x^\gamma)$.

Now for the proof system: we extend a standard deductive system for first-order logic with identity without positing any rules for the quantifiers:

$$\phi \rightarrow (\psi \rightarrow \phi)$$

$$(\phi \rightarrow (\psi \rightarrow \chi)) \rightarrow ((\phi \rightarrow \psi) \rightarrow (\phi \rightarrow \chi))$$

$$(\neg\phi \rightarrow \neg\psi) \rightarrow (\psi \rightarrow \phi)$$

[18]We will omit the superscript where no confusion arises.

[19]Here we follow the lead of Grimau (2018) who argues that the predicates need not be syntactically typed; they can take any term as an argument in any position. Not much hangs on this choice.

(SRI) Singular Reflexivity of Identity
$\forall x_0(x_0 = x_0)$

(Ind) Indiscernibility of Identicals
$\forall x_0 \forall y_0(x_0 = y_0 \rightarrow (\phi(x_0) \rightarrow \phi(y_0)))$, where y_0 is free for x_0 in $\phi(x_0)$.

(MP) Modus Ponens
From ϕ and $\phi \rightarrow \psi$ infer ψ.

We proceed by stipulating axioms and rules for higher-level plurals. For all $k \geq 0$ we have that:

(HLP-UI) Higher-Level Plural Universal Instantiation
$\forall x^k \phi(x^k) \rightarrow \phi(t^k)$, where t^k is free for x^k in ϕ.

(HLP-UG) Higher-Level Plural Universal Generalization
From $\phi \rightarrow \psi(x^k)$ infer $\phi \rightarrow \forall x^k \psi(x^k)$, only if x^k does not occur free in ϕ or in any premise of the deduction.

(HLP-C) Higher-Level Plural Comprehension
$\exists x^k \phi(x^k) \rightarrow \exists x^{k+1} \forall x^k (x^k \prec x^{k+1} \leftrightarrow \phi(x^k))$, where ϕ is a formula containing x_k and possibly other variables free, but no occurence of x^{k+1}.

(HLP-NE) Higher-Level Plural Non-Emptiness
$\forall x^{k+1} \exists x^k (x^k \prec x^{k+1})$

(LR) Level-Raising
$\forall x^k \exists x^l (x^k = x^l)$, for $0 \leq k \leq l$.

(HLP-Ext) Higher-Level Plural Extensionality
$\forall x^k \forall x^l (\forall x^m (x^m \prec x^k \leftrightarrow x^m \prec x^l) \rightarrow (\phi(x^k) \leftrightarrow \phi(x^l)))$, for $m = max(k, l) - 1$ and either $k \neq 0$ or $l \neq 0$, where x^l is free for x^k in $\phi(x^k)$.

This is the framework that we choose to work with for higher-level plurals. This is what we will call the ω-level plural logic, since it contains all pluralities of finite levels. In Chapter 4, our aim will try to see whether this logic can be extended to also include transfinite levels and deal with the challenges that arise from such extensions.

Based on this framework, we are also able to acquire a better understanding of the containment relation $\prec$ by regimenting syntactically our intuitions that stem from plural logic. In the case of plurals, due to the full comprehension axiom, we were able to say that whenever we have a property satisfied by some individuals, then we could form the plurality containing precisely these individuals. The immediate next step tells us that whenever we have some pluralities (individuals plurally referred to) that satisfy a specific property, then we can form the superplural that contains precisely those pluralities. That superplurality will then refer to those same individuals *superplurally*, i.e., by taking into account their having initially been grouped in pluralities.

Our pretheoretic understanding of higher-level plurals containing plurals of lower levels constitutes an integral part of how we understand HLP reference, and as a result, the way in which we formalize HLPL. In that sense, on the formal side of things, we employ HLP quantifiers that bind variables of the corresponding level. The intuition behind an HLP quantifier, for instance in a sentence like $\exists x^2 \phi(x^2)$, will be that for some individuals there is a way to refer to them superplurally such that ϕ holds. Although in HLPL we chose to have the universal quantifier as primitive, we could have provided an equivalent formulation with a primitive existential one. Universal HLP quantification of the form $\forall x^2 \phi(x^2)$ can be glossed as: in any way in which we can refer to individuals superplurally, $\phi(x^2)$ holds. I find myself in agreement with Boolos when arguing that the existential quantifier is more easily grasped than the universal one; after all, it seems that the universal one in the HLP case is very similar to the powerset operation from set theory. In general, an understanding of *Combinatorial Reference* seems enough to understand how HLP quantifiers would function as well.

Before moving on to the next chapter, we will offer a comparison between our approach and that of Grimau (2018), and ours and that of Nicolas and Payton (forthcoming).

2.7 Conceptual Comparison

My *Combinatorial Reference* is very similar to three notions of reference found in the literature. Articulated reference proposed by Ben-Yami (2013), the generalized cover approach of Nicolas and Payton (forthcoming), and the restricted reference approach of Grimau (2018). To kick things off, we start by comparing *Combinatorial Reference* with Grimau's work.

Grimau:

On the formal side, both Grimau and I are availing ourselves to (almost) the same logic ω-level logic. We both take higher-level plurals to be cumulative and regiment them using a cumulative extension of the logic of Rayo (2006) by incorporating insights from Linnebo and Rayo (2012). On the formal side, there is little to compare. Grimau proposes two different ways of thinking of the semantics of HLPL and chooses the latter, because it reflects her preferred understanding of HLP reference. The semantics I will prefer, discussed in Chapter 4 in relation to language legitimacy, will simply use plural resources in the meta-language and are compatible with either way of conceptualizing HLP reference. We mostly focus on the conceptual differences between our competing proposals regarding reference.

One point where our work will be different is that Grimau is mostly interested in the best way to understand plurals in natural language, and that is one of her main goals. For that reason, the ω-level logic is more than enough, and there are also applications of it in other contexts, such as Neo-Logicism (Grimau, 2018, Chapter 7). I, on the other hand, will undertake the task of considering higher-level plural logic as a legitimate philosophical tool, which has as a first task to figure out whether the height of the cumulative plural hierarchy is extended in the transfinite.

The main conceptual point of departure between our proposals is the following. I already touched upon this issue previously, but I will repeat the main points. Grimau understands the notion of higher-level plural reference through what she calls *restricted reference*. For this reason, she employs the concepts of *aspect* and *cluster*. In her words:

> [...] HLP terms denote some objects *under their aspect of being organised in a certain way*. [...] this mode of reference is a species of the broader mode of reference which I call 'restricted reference' - i.e. reference to some objects *under a certain aspect thereof*.
>
> According to this view, higher-level plurals refer to some objects under what I call a 'cluster'. Clusters are plural properties (or perhaps relations [...]) which hold of some

objects in so far as they are organised in various groups. (Grimau, 2018, p. 138)

Up until this point, it seems that our proposals have much in common. And they do. The difference is that I do not employ a concept of cluster, because I rest on the notion of containment. The reason why she opts to introduce that concept has to do with the matters of co-referential higher-level plurals and their intersubstitutivity *salva veritate* that we discussed above. Let us showcase our differences through a familiar example:

(1) My children and your children played against each other.

(2) The boys and the girls played against each other.

For Grimau to explain the failure of intersubstitutivity here, the intensional notion of a *cluster* is employed to distinguish how the two superplural terms refer to the individuals. My proposal for co-reference (motivated by (Wagner, 2015, p. 66)) is simpler and only employs plural notions:

Co-Reference: Two higher-level plurals co-refer if and only if they contain the same lower-level plurals.

Since a higher-level plural refers to some individuals in virtue of the lower-level plurals it contains and how they refer to individuals, if two higher-level plurals contain the same lower-level plurals, then by definition, they will co-refer. Ultimately, I believe that both mine and Grimau's criteria are adequate, but I opt for the principle above precisely because of its simplicity and its direct relation to my principle of *Combinatorial Reference*.

Overall, the divergence between Grimau and me lies on the conceptual level and our goals. Our proposals need not be considered as rivals; the reason I opt for my regimentation of HLP reference is that it paints a more cohesive picture of HLP without needing to resort to additional machinery.

Articulated Reference:

As was mentioned previously, Ben-Yami believes that we need not invoke a notion of HLP reference, insofar as we accept that of *articulated reference*. The idea is that plural terms can refer to individuals in a special way by articulating some aspect of these referents.

On the conceptual level, this does not differ terribly from *Combinatorial Reference*. The articulation that happens when referring to the individual children involves grouping them into "my children and your children", rather than "the boys and the girls". This idea is very similar to the way in which I take the grouping of the individuals into pluralities to play a part in the way superplurals refer to those individuals.

Furthermore, we avoid Grimau's worry that articulated reference in some cases reduces plural to singular reference. Our approach is completely aligned with the pluralist insights and does not suffer from such a defect.

Additionally, both of us rest on a notion of containment to make our proposals precise, but I believe that I avoid some of his proposal's shortcomings. Specifically, his *intuitive* idea of containment is better captured by the way I argued that containment gets iterated and by the axioms of HLPL that govern the containment predicate.

The matter turns to whether there is any substantial difference between "my children and your children" articulating their reference and "my children and your children" referring combinatorially.

As far as I am concerned, *Combinatorial Reference* enjoys the same benefits as articulated reference and improves upon its drawbacks.

Generalized Cover-Approach:

Nicolas and Payton have proposed that we reduce the notion of higher-level plural reference to plural reference under a contextually determined generalized cover. They argue that their approach differs from Ben-Yami's in three substantial aspects (Nicolas and Payton, forthcoming, p. 28-33), which we will discuss to see whether they apply to *Combinatorial Reference* as well, due to its similarities to articulated reference. They criticize Ben-Yami's approach firstly on the fact that containment is suspiciously based on syntactic features of the notion of containment. As we saw before, such a criticism does not apply to my proposal.

They proceed then to argue that Ben-Yami distinguishes distributive predication from cases of articulated reference, while their approach offers a unified explanation in terms of covers. For us, a superplural predicate distributing over pluralities is not genuinely superplural and thus reduces to predication over pluralities. Thus, similarly, we have a distinction between that and collective predication that requires superplural resources.

In my opinion, this does not constitute a problem, no more than plural distributive predication gets reduced to predication over individuals. It is only a problem for Ben-Yami, who takes articulated reference to be a species of plural reference. What is more, one could argue that I offer a more unified approach by generalizing already acceptable plural resources, which Nicolas and Payton also accept, instead of adding another primitive notion, that of a cover.

Finally, the main difference between ordinary plural reference and apparently superplural reference lies in the way one interprets the predicate. The predicate determines the choice of a cover that splits the plurality into subpluralities. Granted, supporters of HLP do not tell a wholly different story. Superplural terms in the context of a subject-predicate sentence are genuinely superplural and refer as such only if the predicate is not distributive. In that sense, there is an intricate balance between superplural terms and the predicates licensing them being genuinely superplural.

Once again, the crux of the issue is whether there is any substantial difference between splitting a plurality into further subpluralities, thus creating a downwards extending hierarchy, and understanding these phenomena through the lens of an upwards extending hierarchy of plurals. As we will see in Chapter 3, formally no such difference arises. As I said before, my aim is a very specific one: if we *could* understand HLP as I proposed and this provides a materially adequate conception of them when it comes to natural language, then we have every reason to avail ourselves of these resources for logical and philosophical theorizing. The point here is to show that these supposed conceptual differences do not cut as deep as some would like to think.

These two hierarchies I just mentioned are so similar that they give rise to a very similar formal picture. (Wagner, 2015, Ch. 2.1.4) has given an argument that standard cover-based semantics and the plural hierarchy are *structurally isomorphic*. This isomorphism is realized only if one talks of set-sized pluralities. If one allows talk of class many objects, as I believe one should, then this isomorphism breaks down.

However, with generalized covers applying to pluralities instead of sets, one can regain this intuitive correspondence between that approach and the hierarchy posited by HLPL. As mentioned before, when we first discussed this approach, indexing the subpluralities by individuals is problematic due to a plural version of Cantor's Theorem. Thus, if Nicolas' and Payton's proposal remained that way, the isomorphism would break down once again, and the formal upper hand would be given to HLPL, where the plural version of Cantor's Theorem is taken seriously.

Afterwards, the authors propose an indexing using pluralities as well as individuals to be able to have as many subpluralities as they would want. In fact, they would have at least as many subpluralities as the cardinality of the powerset of the individuals with which they started. And now, finally, the differences between the two approaches do not seem as substantial as they initially did.

There are two important comments to be made here. Firstly, the fact that covers are somehow licensed by the way that plurals are predicated over seems puzzling. Context-determining factors make clear which cover is supposed to be used to split up the plural into its subpluralities. My uneasiness with this conception is that it conflates pragmatic and semantic aspects of plurals. Initially, it seems weird that through the context and the relevant predication, there is a clear way in which a plurality is to be carved. Subsequently, if that is supposed to be mediated somehow through the context and communication between speakers, then I am uncertain about how this sheds light on the semantic aspects of these linguistic items.

Secondly, as mentioned before, an alternative formulation of covers takes them to be plural relations, which is akin to the way in which Grimau understands *clusters*. Then the two corresponding forms of reference are very similar to one another, but somehow they differ with respect to whether they take HLP seriously.

Both of these conceptual aspects will prove increasingly peculiar after the formal result of the next chapter. Since the logics HLPL and the one from the cover-based approach (CPL) are going to be shown to be theoretically equivalent under two important notions of equivalence, it seems that the philosophical differences should motivate the choice of framework. But the philosophical differences between CPL and Grimau's understanding of HLPL are minimal, so if there are no substantial formal ones either, it is hard to figure out precisely the point of disagreement.

On the other hand, my proposal differs from both Grimau's and CPL, even if it shares a formal regimentation with the former one. The main disagreement here, in my opinion, clearly rests on the conceptual level. On the formal side of things, what is the difference between countenancing an ascending hierarchy of pluralities versus splitting an existing plurality into a descending hierarchy of sub-pluralities? The next result will show that no significant formal differences arise and that the battle will have to be won on philosophical grounds.

The divergence of opinion rests on whether we should take higher-level plurals at face-value or not; I believe that we *can*, not necessarily that we *must*, and that by choosing to do so we avail ourselves to a very powerful formal framework, whose potential philosophical applications are vast in number.

2.8 Conclusion

Before moving to the more in-depth formal comparison mentioned previously, I want to briefly summarize the findings of this chapter. Firstly, we began by describing the main tenets of pluralism that one should have in mind when trying to argue for the legitimacy of HLP resources *in addition* to plural ones. Then, we offered an overview of numerous different positions in the literature regarding HLP to show that the field of HLP is ripe for investigation.

Having outlined the alternatives, I advocated for my preferred notion of reference regarding higher-level plurals and how the notion of containment plays a crucial role in grasping it. I extended the initial formulation of *Combinatorial Reference* to apply to the cumulative reading of HLP, which I will henceforth assume, before delving deeper into the intricacies of the notion of containment. With this pre-theoretic understanding of the relevant notions, I talked of the logic of HLP and how one should think of its intended interpretation.

Finally, I made three different comparisons between my proposal, the one of Ben-Yami (2013), that of Nicolas and Payton (forthcoming, 2025), and that of Grimau (2018). I found the most substantial resistance to come from the generalized cover approach. In fact, the conceptual similarities are multiple, and as we will see shortly, so are the formal ones. I do not presume to have settled how apparently HLP terms function in natural language, but I believe that I have outlined a legitimate conception of HLP reference that can adequately motivate the usage of HLPL.

Chapter 3

Formal Comparison

3.1 Comparing the Logics

In this section, our aim is simple: we provide formal comparison results between HLPL and CPL. HLPL is the cumulative version of higher-order plural logic described above, and CPL is PFO+ extended accordingly to include *covers* and show how they interact with the predication over plural terms.

As mentioned above, we have good reasons to view higher-level plurals as cumulative. On the cover-based approach, this amounts to a cover partitioning a plurality into both individuals and subpluralities, instead of just into subpluralities; this partitioning corresponds to the cumulative version of the superplurals. Further partitioning the initial subpluralities into more subpluralities and/or individuals gives us third-level plurals, and by iterating this process, plurals of all finite levels. As the logic of HLPL we are interested in is an ω-level language, the corresponding version of CPL would be one where covers can partition pluralities up to any finite number of subpluralities of subpluralities (and/or individuals).

To illustrate how this would work based on an already familiar example:

(8) The professor and the students held presentations in adjacent rooms.

The plurality in question is that containing all the individual students and the professor, which then is carved by the context-determined cover into the individual *the professor* and subplurality *the students* of which the predicate *held presentations in adjacent rooms* holds true.

It seems that there are significant intuitions supporting the existence of an intimate connection between the plural hierarchy and the ability of covers to break down plurals into further subplurals. In what follows regarding the comparison of the logical frameworks, we will restrict ourselves to a modest goal: compare the non-cumulative version of the fragment of HLPL only containing plurals of levels $0, 1, 2$ (individuals, plurals, and superplurals) and the version of CPL that only allows for unique partitions of plurals into subplurals.

In the following versions of the logics, we will allow for singleton pluralities (and superplurals), i.e., plurals containing a unique individual. Although this often goes against some of the intuitions behind plurals, considering singleton pluralities on the formal level does not impact the scope of the arguments. In fact, it would also allow for capturing sentences like (8) without needing to go cumulative.

Another important aspect to point out is the choice of framework in which the comparison will be carried out. In the literature, there is a limited amount of work on comparing different logical systems and a methodology for doing so. Some opt for stating the debate in the framework of *category theory*, since it offers fertile ground to compare theories formulated in a more-or-less ambient logical background.

Our approach here will be a different one, since due to the nature of CPL and HLPL, there is a natural choice of framework where the comparison can be carried out: many-sorted logic. Many-sorted logic should be roughly understood as a logic allowing for multiple kinds of quantification in *term* position, where the quantifiers range over the elements of sets that are disjoint for different sorts. Relation and function symbols come with a specific arity, as in first-order logic, but also with a specification of which sorts can be used for which argument places.

What we aim for here is stating both CPL and HLPL as many-sorted theories over some signatures, and then proceeding to prove that they "are the same" in some appropriate sense. We will discuss that sense in more detail in Section 3.3, but for now it is good to know that we will consider two: (a) *Morita Equivalence*, and (b) *bi-interpretability*. In fact, we will show that the two theories are the same under both conceptions, thus bypassing debates about which of the two fits better our idea of "sameness of theory".

Another important caveat: the choice of carrying out the comparison in many-sorted logic is not devoid of consequences. If any many-sorted theory is understood as talking about different kinds of *objects*, then casting plural logic (and any extension thereof) into that framework would not respect the intended interpretation of plural quantification and reference. Since we have spent some time criticizing any attempts to capture HLP that are not genuinely pluralist, we need a reason to consider this approach as warranted. The answer is simple and relies on a severance between the formal semantics and their intended interpretations. Here we are simply interested in an adequate formal depiction of the logics, which obviously does not respect the intended interpretation, but captures all the relevant formal details. We are not trying to say anything about how plurals should be understood; we have already accepted the pluralist view, and that does not change, as the many-sorted approach functions as a mere heuristic to carry out the comparison using the plethora of tools coming from standard model theory.[1]

The second point, a formal one, regards the semantics. Plural logic as we have seen it, and various extensions thereof, are often used to talk about specific kinds of discourses. Two interrelated cases involve talk of Absolute Generality, the idea that one's first-order quantifiers can be meaningfully taken to range over absolutely every object, and quantification or reference to *class-sized* totalities.

Absolute Generality, a metaphysical thesis to which we will return later on, may be true of the world or it may not; I do not aim to argue for either stance on the matter, but if it is true, then standard model theoretic techniques will falter for fear of paradox as there cannot be a universal set (if absolute generality is true) or any class-sized model, if we want to capture that kind of discourse. Immediately, it becomes clear that so long as we use set theory as our meta-theory, we will not be able to capture all of the intended models at face value.

I will outline two ways in which one may choose to tackle this problem. In any case, I maintain that the best way to start this investigation is by using the model theory of many-sorted logic, even if

[1] In this case, since both logics are extending standard first-order logic, I believe that many-sorted indeed offers a level playing field. I am not sure whether that would be the case if one of the two logics to be compared were, for instance, intuitionistic.

it may not settle the score immediately. The first way to deal with that would be to investigate the many-sorted theory of a given plural logic in a plural meta-theory. This approach would be in line with that of Boolos (1984, 1985), Rayo and Uzquiano (1999), Rayo and Williamson (2003), Yi (2005, 2006), and Oliver and Smiley (2016) among others. This would entail a significant amount of work where the various interpretability and equivalence results between logical theories would need to be cast in a meta-language involving plural resources. Obviously, such work far exceeds the goals of this thesis, but I take it to be an unavoidable task for any proponent of higher-order semantics who would like to carry out meaningful comparisons between logics and logical theories.

Another option would be to use a meta-theory that invokes *proper classes*; for instance, we could provide a model theory in von Neumann-Bernays-Gödel (NBG) or Morse-Kelley (MK) class theory. This would at least allow for models of the appropriate size, e.g., with a domain of all the ordinals, and would minimize the loss incurred by using ZFC as a meta-theory. Another alternative would be to use a set theory that allows for something like a universal set, such as Quine's New Foundations (NF). If one opts for a set or class theory to regiment one's semantics, instead of going for the plural alternative, which is the conceptually more appropriate choice, one should avoid set theories like NF. The main reason is that, even though they allow for a universal set, they do so by severely restricting the separation axiom, telling us which subsets of given sets exist. I believe, then, that choosing NF (or a similar set theory) would solve a problem by creating another.

All in all, there is certainly an abundance of details that need to be ironed out, but I believe both that starting by putting the question of comparison in the context of many-sorted logic is a good first step and also that there is good reason to think that the details *can* be ironed out. Using many-sorted logic, we will be able to invoke a plethora of results that will help us prove that CPL and HLPL are the same in a formally important sense.

3.2 Many-Sorted Logic (MSL)

Many-sorted first-order logic, henceforth MSL, has been regarded by numerous logicians as a very natural logical framework for numerous logical and mathematical purposes. Reasons for that usually include the fact that in mathematics we very often make distinctions between different kinds of objects and work simultaneously with both; consider, for instance, the axiomatization of Euclidean Geometry that uses two kinds (sorts) of objects, namely points and lines. Similarly, in the theory of vector spaces, we often talk of vectors and a field comprised of scalars.

In philosophy, the appeal of MSL has been steadily increasing precisely because it is thought that it can work as a level playing field for carrying out comparisons between logical systems that, at first glance, look very different from one another. Manzano (1996) has advocated at great length in favor of this view. Florio (2023) carries out a comparison in many-sorted logic between typed and untyped systems regarding their expressibility (a matter to which we will return in the final chapter). Finally, Halvorson (2019, Chapters 4-5) uses MSL as the background theory to talk of the theoretical equivalence of theories in the contexts of philosophy, logic, and science.

In this section, we will present the basic formal definitions regarding MSL. In the next section, we will use those definitions to talk about two kinds of theoretical equivalence.[2]

[2]My presentation of MSL mainly follows McEldowney (2020), and in some parts Halvorson (2019) and Hodges (1993). Two extensive treatments of MSL, with slightly different notation, can be found in Manzano (1996) and Manzano and Aranda (2022).

A many-sorted first-order signature $\mathcal{L}$ is a quadruple $(\mathcal{S}, \mathcal{R}, \mathcal{F}, \mathcal{C})$. $\mathcal{S}$ is a non-empty set of sort symbols. $\mathcal{R}$ is a set of *sorted* relation symbols such that for every $R \in \mathcal{R}$ has an arity $s_1 \times s_2 \times \cdots \times s_n$ such that every $s_i \in \mathcal{S}$. Similarly, $\mathcal{F}$ is a set of function symbols such that each $f \in \mathcal{F}$ has arity $s_1 \times \cdots \times s_n \to s_{n+1}$, and $\mathcal{C}$ is a set of constant symbols each equipped with an arity s_c.

Terms and formulas are defined recursively in the usual manner as for first-order logic. A many-sorted first-order $\mathcal{L}$-language is the set of formulas generated from the signature $\mathcal{L}$. In what follows, we will not distinguish between an $\mathcal{L}$-language and the signature $\mathcal{L}$; we will denote both by $\mathcal{L}$. Each sort s comes with its own identity symbol $=_s$, where the superscript will be omitted when it is clear from context. Each sort also comes with its own quantifiers and a countable stock of variables, meaning that any variable and any quantifier binding a variable ranges over a unique sort. Quantifiers of sort s are often denoted by $\exists_s$ and $\forall_s$, but when we talk of HLPL and CPL, we will use the notational conventions regarding plurals by writing variables randing over individuals as $x, y, \ldots$, variables ranging over plurals as $xx, yy, \ldots$, variables ranging over superplurals as $xxx, yyy, \ldots$, and the ones ranging over covers just with capital letters as $I, J, \ldots$.

For variables, their sort will also be called their *arity*. The *arity* of an $\mathcal{L}$-formula $\phi(x_1, \ldots, x_n)$ with free variables $x_1, \ldots, x_n$ of sorts $s_1, \ldots, s_n$ will be denoted by $s_1 \times \cdots \times s_n$. An $\mathcal{L}$-sentence is an $\mathcal{L}$-formula without any free variables.

We define a $\mathcal{L}$-structure $\mathcal{M}$ as another quadruple (S, R, F, C) that assigns "meaning" to the symbols of $\mathcal{L}$. The assignment of meaning happens as follows:

1. For each sort symbol $s \in \mathcal{S}$, there is a non-empty set $M_s \in S$ which will be a sort of $\mathcal{M}$.
2. For each relation symbol $R \in \mathcal{R}$ of arity $s_1 \times \cdots \times s_n$, there is a relation $R^{\mathcal{M}} \subseteq M_{s_1} \times \cdots \times M_{s_n}$.
3. For each function symbol $f \in \mathcal{F}$ of arity $s_1 \times \cdots \times s_n \to s_{n+1}$, there is a total function $f^{\mathcal{M}}$ from $M_{s_1} \times \cdots \times M_{s_n}$ to $M_{s_{n+1}}$.

For every $s \in \mathcal{S}$, the identity symbol $=_s$ is interpreted as equality between objects in M_s. We denote finite tuples of objects and variables as $\bar{a}, \bar{b}, \ldots$ and $\bar{x}, \bar{y}, \ldots$ respectively. With the notion of structure in place, one can define truth in a model in the standard Tarskian way. Roughly, whenever we have an $\mathcal{L}$-formula $\phi(x_1, \ldots, x_n)$ of arity $s_1 \times \cdots \times s_n$, and the objects $a_1, \ldots, a_n$ such that $a_i \in M_{s_i}$, we define recursively $\mathcal{M} \vDash \phi(a_1, \ldots, a_n)$. We then say that $\mathcal{M}$ *satisfies* $\phi(a_1, \ldots, a_n)$ or that it is a *model* of it.

An $\mathcal{L}$-theory T is a set of $\mathcal{L}$-sentences. We call T satisfiable if there is an $\mathcal{L}$-structure that satisfies all of its sentences, i.e., that there is an $\mathcal{L}$-structure $\mathcal{M}$ such that for all $\phi \in T$ we have that $\mathcal{M} \vDash \phi$. We say that T *logically implies* ϕ (and we write $T \vdash \phi$) if every model of T satisfies ϕ. Similarly, we define $T \vdash T'$ for two $\mathcal{L}$-theories T and T'. We say that T and T', as before, are *logically equivalent* iff $T \vdash T'$ and $T' \vdash T$, i.e., whenever they have the same theorems. We denote logical equivalence of theories by $T \equiv T'$.

Whenever we look at our $\mathcal{L}$-structures, it will be in a meta-language of set theory (ZFC), meaning that the domain will be comprised of pairwise disjoint sets corresponding to each sort, terms pick out objects of the right sort, and relation symbols will be interpreted as relations of the appropriate arity on the domain specified before. Had we wanted to be even closer to how plural logics are used, we would take our meta-language to be set theory *with urelements* (like ZFCU). Formally, not much depends on this choice; it will merely simplify the proofs.

3.3 Theory Comparison in MSL

With the preliminaries of MSL out of the way, we now turn to two kinds of equivalence between MSL-theories: (a) *Morita Equivalence*, and (b) *bi-interpretability*. Before giving the relevant definitions, a few comments are in order.

When talking of equivalence between theories, we are trying to establish that certain formal aspects of two theories are shared in a manner that makes the two theories, at least intuitively, "the same". Barrett and Halvorson (2016) and Halvorson (2019) have proposed *Morita Equivalence* as the appropriate formal notion for that task, while others, such as McEldowney (2020) and Visser (2021), argue that *bi-interpretability* enjoys that privilege. The discussion often turns to matters regarding the philosophy of science and theoretical equivalence between theories in that realm, but we will avoid jumping into these murky waters.[3]

We aim to bypass this debate entirely. We will provide a proof that HLPL and CPL as MSL theories are Morita Equivalent; based on a result from McEldowney (2020), we infer that this proof also suffices for bi-interpretability. The choice to prove Morita Equivalence for the two theories here is merely formal and does not reflect my stance on the previous debate.

Morita Equivalence:

To show that two theories T_1 over some many-sorted signature Σ_1 and T_2 over some possibly different many-sorted signature Σ_2 are Morita Equivalent, we need to extend both theories appropriately, so that they are formulated over the same signature and they are logically equivalent.

The extensions, appropriately dubbed *Morita Extensions*, are ones where we define new sorts, according to some acceptable ways of doing so, and explicitly define new relation, function, and constant symbols. I introduce the relevant definitions below.

Definition 2. *Let $\mathcal{L} \subset \mathcal{L}^+$, and let $R \in \mathcal{L}^+$ be a relation symbol. Then, an* ***explicit definition of*** R ***in terms of*** $\mathcal{L}$ *is an $\mathcal{L}^+$-sentence of the form:*

$$\forall \bar{x}(R(\bar{x}) \leftrightarrow \phi(\bar{x})), \tag{3.1}$$

where ϕ is an $\mathcal{L}$-formula. Then, if $f \in \mathcal{L}^+$ is a function symbol, an ***explicit definition of*** f ***in terms of*** $\mathcal{L}$ *is an $\mathcal{L}^+$-sentence of the form:*

$$\forall \bar{x} \forall y(f(\bar{x}) = y \leftrightarrow \phi(\bar{x}, y)), \tag{3.2}$$

where $\phi(\bar{x}, y)$ is an $\mathcal{L}$-formula. Finally, if $c \in \mathcal{L}^+$ is a constant symbol, an ***explicit definition of*** c ***in terms of*** $\mathcal{L}$ *is an $\mathcal{L}^+$-sentence of the form:*

$$\forall x(c = x \leftrightarrow \phi(x)) \tag{3.3}$$

where $\phi(x)$ is an $\mathcal{L}$-formula.

We also use $\exists!$ so that $\exists! x\phi(x)$ means that there exists a *unique* x such that ϕ. Based on that, we have that (3.2) implies $\forall \bar{x} \exists! y \phi(\bar{x}, y)$ and that (3.3) implies $\exists! x \phi(x)$. These two sentences are called *admissibility conditions* for (3.2) and (3.3) respectively.

Through explicit definitions we can expand our theories to new signatures by defining new constant, function, and relation symbols using resources from the initial signatures. There are four ways in which one can extend a signature by defining new *sort symbols* outlined in Barrett and

[3]For a general discussion of bi-interpretability as a notion of equivalence, see Button and Walsh (2018).

Halvorson (2016). We will present one of them, defining new sorts as *quotient* sorts, as it will be the one we will use.

Definition 3. *Let $\mathcal{L} \subset \mathcal{L}^+$ such that s is a sort symbol in $\mathcal{L}^+ \setminus \mathcal{L}$, s_1 is a sort symbol in $\mathcal{L}$, and μ is a function symbol in $\mathcal{L}^+$ of arity $s_1 \to s$. Then, an explicit definition of s and μ as a* ***quotient sort*** *in terms of $\mathcal{L}$ is an $\mathcal{L}^+$-sentence of the form*

$$\forall_{s_1} x_1 \forall_{s_1} x_2 [\mu(x_1) = \mu(x_2) \leftrightarrow \phi(x_1, x_2)] \wedge \forall_s z \exists_{s_1} x [\mu(x) = z], \tag{3.4}$$

where $\phi(x_1, x_2)$ is an $\mathcal{L}$-formula.

Remark 4. *Note that (3.4) implies the following $\mathcal{L}$-sentences, which basically capture that $\phi(x_1, x_2)$ encodes an equivalence relation and are the admissibility conditions for (3.4):*

$\forall_{s_1} x(\phi(x, x))$

$\forall_{s_1} x_1 \forall_{s_1} x_2 [\phi(x_1, x_2) \to \phi(x_2, x_1)]$

$\forall_{s_1} x_1 \forall_{s_1} x_2 \forall_{s_1} x_3 [(\phi(x_1, x_2) \wedge \phi(x_2, x_3)) \to \phi(x_1, x_3)]$

Definition 5. *Let T be an $\mathcal{L}$-theory. Then, a* ***Morita extension*** *T^* of T to $\mathcal{L}^+$ is an $\mathcal{L}^+$-theory of the form $T \cup \{\delta_\sigma \mid \sigma \in \mathcal{L}^+ \setminus \mathcal{L}\}$ satisfying the following conditions:*

1. *For each symbol $\sigma \in \mathcal{L}^+ \setminus \mathcal{L}$, the sentence δ_σ is an explicit definition of σ in terms of $\mathcal{L}$.*
2. *If α_σ is an admissibility condition for a definition δ_σ, then $T \vdash \alpha_\sigma$.*
3. *For each sort symbol $s \in \mathcal{L}^+ \setminus \mathcal{L}$ and function symbol $f \in \mathcal{L}^+ \setminus \mathcal{L}$, if f appears in the sort definition $\delta_s \in T^*$, then $\delta_f = \delta_s$.*

By using the notion of a Morita extension, we finally define the notion of Morita Equivalence.

Definition 6. *Let T and T' be many-sorted first-order theories. The theories T and T' are* ***Morita Equivalent*** *if there are finite sequences of Morita extensions $T, T_1, \ldots, T_n$ and $T', T'_1, \ldots, T'_m$, such that $T_n \equiv T'_m$.*

The definitions for bi-interpretability will be avoided, because they will not figure in any of the proofs to come. The main idea is the following: if we have two theories T, T' over some signatures Σ, Σ', we call them *bi-interpretable* if we have definable functions f, g such that whenever $\mathcal{M} \vDash T$ we have that $f(\mathcal{M}) \vDash T'$, and whenever $\mathcal{N} \vDash T'$ we have that $g(\mathcal{N}) \vDash T$, and additionally the bijections $f \circ g$ and $g \circ f$ are definable.[4]

If one wishes to see how bi-interpretability works in the many-sorted setting, one can consult McEldowney (2020). The important result from that paper that will be important for our work is the following:

Proposition 7. *Let T and T' be Morita equivalent by way of finite sequences of Morita extensions which define no new coproduct sorts. Then T and T' are* ***bi-interpretable****.*

As it has been indicated by the preliminaries above, the only way in which we will be defining new sorts is by *quotient sorts*, thus avoiding coproducts. This will show later on that our theories T_{HLPL} and T_{CPL} are bi-interpretable, as well as Morita equivalent.

[4]The precise details can be found in Button and Walsh (2018).

3.4 HLPL and CPL in MSL

In this part, we will regiment both of these logics as specific many-sorted theories that we will investigate. A note on terminology first; although we could change the notation to be more in line with that of MSL, we choose to keep it in accordance with that of HLPL and PFO+: letters like $x, y, z, \ldots$ will be used for individuals, $xx, yy, zz, \ldots$ for plurals, $xxx, yyy, zzz, \ldots$ for superplurals, and capital letters like $I, J, K, \ldots$ to denote covers.

3.4.1 HLPL in MSL

As was mentioned at the beginning of the chapter, the version of HLPL we will focus on will be slightly different from the one we have presented in Chapter 1. The logic presented in Chapter 1 contains a countably infinite collection of sorts and is also cumulative with respect to the containment relation (terms of type n can contain terms of type k for all $k < n$). In this section, we will look at the non-cumulative variant of that language that includes terms of levels $0, 1, 2$. This requires slightly altering the axioms of HLPL.

The formulation of HLPL given in Chapter 1 can almost immediately be understood as a many-sorted theory. We start by defining the many-sorted first-order signature $\Sigma_1 = (\mathcal{S}_1, \mathcal{R}_1)$, which contains no function and no constant symbols. We will have three sorts σ_0, σ_1, and σ_2 corresponding to individuals, plurals, and superplurals, respectively. The only two two-place relation symbols we will need will be $\prec$ of arity $\sigma_0 \times \sigma_1$ and $\prec^1$ of arity $\sigma_1 \times \sigma_2$.

Then, the Σ_1-theory T_{HLPL} will contain the list of axioms (without the inference rules) from Section 2.1.6 with some exceptions. For instance, the Level-Raising axiom (LR) will not be included, since it is only relevant in the cumulative version. We restate those axioms here:

HLP-C Higher-Level Plural Comprehension
$\exists x^k \phi(x^k) \to \exists x^{k+1} \forall x^k (x^k \prec x^{k+1} \leftrightarrow \phi(x^k))$, where ϕ is a formula containing x_k and possibly other variables free, but no occurrence of x^{k+1}, for $k \in \{0, 1\}$.

HLP-NE Higher-Level Plural Non-Emptiness
$\forall x^{k+1} \exists x^k (x^k \prec x^{k+1})$, for $k \in \{0, 1\}$.

HLP-Ext Higher-Level Plural Extensionality
$\forall x^{k+1} \forall y^{k+1} (\forall x^k (x^k \prec^k x^{k+1} \leftrightarrow x^k \prec^1 y^{k+1}) \leftrightarrow x^{k+1} = y^{k+1}$, where $\prec^0 \equiv \prec$ and $k \in \{0, 1\}$.

The axioms (HLP-C) and (HLP-NE) are both exactly the same as previously. The axiom of (HLP-Ext) differs slightly because it is now expressed in terms of HLP identity, instead of the right-hand side specifying that the plurals agree on all formulae. This is because in HLPL, the notion of HLP identity is a defined one, while here we take identity as primitive. For that reason, (HLP-Ext) as expressed here in the context of MSL implies the axiom scheme stated in Section 2.1.6. It is easy to verify that the Σ_1-signature $(A, \mathcal{P}(A) \setminus \emptyset, \mathcal{P}(\mathcal{P}(A) \setminus \emptyset) \setminus \emptyset, \in, \in)$ for some $A \neq \emptyset$ is indeed a model of T_{HLPL}, meaning that T_{HLPL} is satisfiable.

3.4.2 CPL in MSL

CPL extends PFO+ so the many-sorted theory T_{CPL} will be a theory over a signature with at least two sorts, one for individuals and one for plurals. The important aspect will be regimenting the covers appropriately. As far as I know, three related but slightly different formulations of CPL have been proposed in Nicolas and Payton (forthcoming), Nicolas and Payton (2025), and Payton (2025). The main idea is that covers act either as multi-valued functions or as plural relations by mapping or relating plurals with some of their subpluralities in an exhaustive manner. For us, they will act as functions from plurals to some of their subpluralities; in set-theoretic parlance, as functions from a set to individuals to a set of its subsets.

In CPL variables need to be indexed by a cover to occur in formulas, meaning that a string of the form $\exists xx F(xx)$ is not well-formed, but $\exists xx^i F(xx^i)$ is, with the latter capturing the fact that the predicate applies to the subpluralities of xx as articulated by the cover i. In certain formulations of CPL, there is no quantification over covers, where indexing with respect to a cover is supposed to capture the application of a specific one. In other formulations it is preferred to make the existence of covers explicit, so quantification over covers is also permitted.

In the MSL formulation of CPL, we also have quantification over covers. We will also make some choices as to how to represent them, but hopefully in a way that does justice to the original formulations of CPL. The way we will attempt to capture covers will be through defining a sort σ_c for covers, and a relation ϵ of arity $\sigma_c \times \sigma_1 \times \sigma_1$ that stands for how a cover articulates a plural into its subpluralities. The point with T_{CPL} will be to sketch out how covers interact with plurals, so we will allow no other relation symbols apart from $\prec$ and ϵ. In Section 3.6, we will sketch how the proofs presented later on can be adapted to account for signatures with additional relation symbols.

To define CPL as a many-sorted theory, we start by defining the relevant signature $\Sigma_2 = (\mathcal{S}_2, \mathcal{R}_2)$. Note that once again, we have no function and no constant symbols. We have three different sort symbols σ_0, σ_1, and σ_c, where the first two are about individuals and plurals and coincide with those of Σ_1, while the third is a sort for covers. Then, we have the two-place relation symbol $\prec$ of arity $\sigma_0 \times \sigma_1$, and the three-place relation symbol ϵ of arity $\sigma_c \times \sigma_1 \times \sigma_1$.

Then, we define the axioms of the Σ_2-theory T_{CPL}. T_{CPL} has all the axioms of PFO+ relating the first two sorts (singular and plural terms) with one another and telling us how the containment relation $\prec$ behaves. Note that these axioms correspond to the axioms of T_{HLPL} for $k = 0$. The three axioms are the following:

Non-empty $\forall xx \exists y\, y \prec xx$
Indisc $\forall xx \forall yy (\forall z (z \prec xx \leftrightarrow z \prec yy) \rightarrow (xx = yy))$
P-Comp $\exists x \phi(x) \rightarrow \exists xx \forall x (x \prec xx \leftrightarrow \phi(x))$

Finally, we need to pin down the axioms for the relation symbol ϵ. In Chapter 1, we posited that "i covers aa" just in case:

(a) for all xx, if $i(aa, xx)$ then xx is a subplurality of aa, and

(b) for all x, if $x \prec aa$ then there are some xx such that $i(aa, xx)$ and $x \prec xx$.

The first axiom explains that covering a plurality means relating it to its subpluralities. The second tells us that this relation is exhaustive in the sense that each individual contained in the plurality will be contained in some sub-plurality of aa given by the cover. We will formalize these two axioms and add three more explaining how covers are supposed to function. The latter three axioms are not clear in any of the three papers about generalized covers (Nicolas and Payton, forthcoming, 2025; Payton, 2025), so we have added them to make the formal presentation clearer by explicitly capturing how covers are to behave formally. They are necessary because they talk about the existence of covers and about when two covers are the same.

Continuing with the meanings of the axioms, the third axiom scheme acts as a comprehension axiom that tells us that for every definable way to carve a plurality into subpluralities, there is a cover that witnesses that carving. The fourth axiom tells us that each cover applies to a unique plurality. The third axiom tells us that if two covers apply to the same plurality and carve it up in the same way, they are the same cover. By combining the axioms, we have that for each plurality there are covers that can divide it in any way possible into subpluralities, but each cover does that for one plurality and no others.

Furthermore, we will use the defined symbol $\leq$ to indicate that yy is a subplurality of xx. Formally: $yy \leq xx \leftrightarrow \forall x(x \prec yy \rightarrow x \prec xx)$. The axioms are:

(a) $\forall I \forall xx \forall yy(\epsilon(I, xx, yy) \rightarrow yy \leq xx)$,

(b) $\forall I \forall xx(\exists yy\, \epsilon(I, xx, yy) \rightarrow \forall x(x \prec xx \leftrightarrow \exists zz\, (\epsilon(I, xx, zz) \wedge x \prec zz)))$, and

(c) $\exists zz\phi(zz) \rightarrow \exists I \exists xx \forall yy(\phi(yy) \leftrightarrow \epsilon(I, xx, yy))$, where ϕ is a Σ_2-formula containing yy free, and also does not contain any terms of sort σ_c or the relation symbol ϵ.

(d) $\forall I \exists xx(\exists yy\epsilon(I, xx, yy) \wedge \forall zz \forall ww(\epsilon(I, zz, ww) \rightarrow xx = zz)$.

(e) $\forall I \forall J \forall xx(\forall yy(\epsilon(I, xx, yy) \leftrightarrow \epsilon(J, xx, yy)) \leftrightarrow I = J)$.

The Σ_2-theory T_{CPL} then is the theory containing the four axioms above and all of the (Non-empty), (Indisc), and (P-Comp).

3.5 T_{HLPL} and T_{CPL} are Morita Equivalent and Bi-interpretable

Having set up all the preliminaries, we can finally turn to the main result of this chapter. We will show that the theories T_{HLPL} and T_{CPL} are *Morita Equivalent*, and from that, based on a result proved by McEldowney (2020), we deduce that the two theories are also *bi-interpretable*. I believe it is worth remarking once more that with this proof, we bypass the debate regarding which of the two is the best notion of theoretical equivalence.

Before presenting the proofs in detail, it is important to understand, in general, a sketch of the idea driving these proofs. Both T_{CPL} and T_{HLPL} have a sort for individuals and a sort for plurals that are governed by the same axioms. The point of departure between the two theories is that one uses the notion of cover, while the other uses superplurals. However, we will show that we can extend the signatures of each of these two theories by equating superplurals with covers in both cases. Then, the tools of PFO+ present in both theories will allow us to define all the rest of the relation symbols and extend the signatures and theories in a way that gives us two logically equivalent theories, and thus the Morita Equivalence of T_{CPL} and T_{HLPL}.

The proof proceeds syntactically by using explicit definitions, but the same idea (and some of the same formulas) could be used to provide us with the definable isomorphisms between the models of the theories, which would prove that they are bi-interpretable. The choice to carry out the proof in this way has been made for two reasons. First, to avoid delving into the subtleties of the debate regarding theoretical equivalence as mentioned above. Specifically, from the way we prove Morita Equivalence, we deduce bi-interpretability. Bi-interpretability does not generally imply Morita Equivalence (McEldowney, 2020, p. 406), and as far as I know, there is no result telling us when bi-interpretable theories are generally Morita Equivalent.[5] With this approach we can get both results. The second reason is that, in my opinion, this proof shows how intimately connected the formal notions of a cover and a superplural are, not only in the models they define, but also in the way they behave syntactically.

Morita Extensions of T_{HLPL} and T_{CPL}

We will extend the language Σ_1 of T_{HLPL} in three steps. Firstly, we define covers as a *quotient sort* of superplurals by using the identity as our equivalence relation, thus identifying covers with superplurals. To do so, we use the $\Sigma_1 \cup \{\sigma_c, i_1\}$-formula δ_{σ_c} to simultaneously define the sort σ_c and the function symbol i_1 of arity $\sigma_2 \to \sigma_c$:

$$(\delta_{\sigma_c})\ \ \forall xxx \forall yyy (i_1(xxx) = i_1(yyy) \leftrightarrow xxx = yyy) \ \wedge\ \forall I \exists xxx (i_1(xxx) = I)$$

This is how we get our first Morita extension of T_{HLPL}:

Proposition 8. *The $\Sigma_1 \cup \{\sigma_c, i_1\}$-theory $T'_{HLPL} = T_{HLPL} \cup \{\delta_{\sigma_c}\}$ is a **Morita extension** of the Σ_1-theory T_{HLPL}.*

Proof. The proof is immediate by the definition of a Morita Extension. We extended the theory by defining a new quotient sort using $=$ as an equivalence relation, which trivially satisfies the admissibility requirements. □

Secondly, we explicitly define the ϵ relation by saying that a cover exhaustively carves a plurality into subpluralities whenever there is a superplural corresponding to that cover that contains exactly those subpluralities. We do so using the formula:

$$(\delta_\epsilon)\ \ \forall I \forall xx \forall yy (\epsilon(I, xx, yy) \leftrightarrow yy \leq xx \ \wedge\ \exists xxx (i_1(xxx) = I \wedge\ yy \prec^1 xxx \ \wedge \forall x (x \prec^0 xx \leftrightarrow \exists zz (x \prec^0 zz \ \wedge\ zz \prec^1 xxx \ \wedge\ zz \leq xx)))$$

Then, we have another Morita extension:

Proposition 9. *The $\Sigma_1 \cup \{\sigma_c, i_1, \delta_\epsilon\}$-theory $T''_{HLPL} = T'_{HLPL} \cup \{\delta_\epsilon\}$ is a **Morita extension** of the $\Sigma_1 \cup \{\sigma_c, i_1\}$-theory T'_{HLPL}.*

Proof. The proof is also immediate because δ_ϵ is an explicit definition of the relation symbol ϵ in terms of the signature $\Sigma_1 \cup \{\sigma_c, i_1\}$. □

Before seeing the third step that extends T''_{HLPL} to a theory T^*_{HLPL}, we will see how the theory T_{CPL} will be extended also in three steps. The first step is completely analogous, where superplurals are explicitly defined as a *quotient sort* of covers using the identity symbol. We do that and simultaneously define the function symbol i_2 or arity $\sigma_c \to \sigma_2$ by using the $\Sigma_2 \cup \{\sigma_2, i_2\}$-formula δ_{σ_2}:

[5] In Friedman and Visser (2014), Friedman and Visser prove that in some cases bi-interpretability implies *synonymy* and *definitional equivalence*. Definitional equivalence in turn implies Morita Equivalence when the sorts in the signatures of both theories are *the same*. Even if Friedman and Visser's arguments applied in this case, this would not give us Morita Equivalence since the initial theories include different sort symbols. Synonymy would require a different approach still, because Friedman and Visser's results are about one-sorted theories.

(δ_{σ_2}) $\forall I \forall J (i_2(I) = i_2(J) \leftrightarrow I = J) \ \wedge \ \forall xxx \exists I (i_2(I) = xxx)$

And analogously, we have the following proposition, whose proof is exactly as for Proposition 8:

Proposition 10. *The $\Sigma_2 \cup \{\sigma_2, i_2\}$-theory $T'_{CPL} = T_{CPL} \cup \{\delta_{\sigma_2}\}$ is a **Morita extension** of the Σ_2-theory T_{CPL}.*

Secondly, we explicitly define the relation symbol $\prec^1$ standing for plurals being contained in superplurals using the following $\Sigma_2 \ \cup \ \{\sigma_2, i_2\}$-formula, which tells us that a plural aa is contained to a superplural aaa whenever there is a cover corresponding to aaa and a plural bb, such that the cover applies to bb and aa is one of bb's subplurals carved out by the cover. Formally:

($\delta_{\prec^1}$) $\forall xxx \forall xx \ (xx \prec^1 xxx \leftrightarrow \exists I \exists yy (i_2(I) = xxx \ \wedge \ \epsilon(I, yy, xx)))$.

The second extension is defined as follows, and the proof here is analogous to that of Proposition 9:

Proposition 11. *The $\Sigma_2 \cup \{\sigma_2, i_2, \prec^1\}$-theory $T''_{CPL} = T'_{CPL} \cup \{\delta_{\prec^1}\}$ is a **Morita extension** of the $\Sigma_2 \cup \{\sigma_2, i_2\}$-theory T'_{CPL}.*

Finally, we need to extend both T''_{HLPL} and T''_{CPL} to theories defined over the same signature. That signature will be $\Sigma = \Sigma_1 \cup \Sigma_2 \cup \{i_1, i_2\}$. The last thing missing is an explicit definition of i_2 for T''_{HLPL} and one of i_1 for T''_{CPL}. We do so by using the same $\{\sigma_c, \sigma_2, i_1, i_2\}$ sentence:

($\delta_{i_1 i_2}$) $\forall xxx \forall I (i_2(I) = xxx \leftrightarrow i_1(xxx) = I)$

We finally have these two propositions, whose proofs are also immediate and therefore omitted:

Proposition 12. *The Σ-theory $T^*_{HLPL} = T''_{HLPL} \cup \{\delta_{i_1 i_2}\}$ is a **Morita extension** of the $\Sigma_1 \cup \{\sigma_c, i_1, \delta_\epsilon\}$-theory T''_{HLPL}.*

Proposition 13. *The Σ-theory $T^*_{CPL} = T''_{CPL} \cup \{\delta_{i_1 i_2}\}$ is a **Morita extension** of the $\Sigma_2 \cup \{\sigma_2, i_2, \delta_{\prec^1}\}$-theory T''_{CPL}.*

And this finally brings us to the main result of this chapter:

Theorem 14. *The Σ_1-theory T_{HLPL} and the Σ_2-theory T_{CPL} are **Morita Equivalent**.*

Proof. It suffices to show that the Σ-theories T^*_{HLPL} and T^*_{CPL} are *logically equivalent*.[6]

First, we show that $T^*_{HLPL} \vdash T^*_{CPL}$, i.e., that every model of the former theory is a model of the latter. Let $\mathcal{M}$ be an arbitrary model of T^*_{HLPL}:

– We want to show that $\mathcal{M} \vDash \forall I \forall xx \forall yy (\epsilon(I, xx, yy) \to yy \leq xx)$. Take an arbitrary I of sort σ_c and arbitrary aa, bb of sort σ_1, such that $\mathcal{M} \vDash \epsilon(I, aa, bb)$. Since $\mathcal{M} \vDash T^*_{HLPL}$, specifically we have that $\mathcal{M} \vDash \delta_\epsilon$, which means that for the given I, aa, bb: $\mathcal{M} \vDash \epsilon(I, aa, bb) \leftrightarrow bb \leq aa \ \wedge \ \exists xxx(i_1(xxx) = I \wedge bb \prec^1 xxx \wedge \forall x(x \prec^0 aa \ \leftrightarrow \ \exists zz(x \prec^0 zz \wedge zz \prec^1 xxx \wedge zz \leq aa))$. Because $\mathcal{M} \vDash \epsilon(I, aa, bb)$, we have that $\mathcal{M} \vDash bb \leq aa \ \wedge \ \exists xxx(i_1(xxx) = I \ \wedge \ bb \prec^1 xxx \ \wedge \ \forall x(x \prec^0 aa \ \leftrightarrow \ \exists zz(x \prec^0 zz \ \wedge \ zz \prec^1 xxx \ \wedge \ zz \leq aa))$, which implies that $\mathcal{M} \vDash bb \leq aa$. Therefore, we have that $\mathcal{M} \vDash \epsilon(I, aa, bb) \to bb \leq aa$.

  Because I, aa, bb were all chosen arbitrarily we have that $\mathcal{M} \vDash \forall I \forall xx \forall yy (\epsilon(I, xx, yy) \to yy \leq xx)$, and because $\mathcal{M}$ was chosen arbitrarily we have that $T^*_{HLPL} \vdash \forall I \forall xx \forall yy (\epsilon(I, xx, yy) \to yy \leq xx)$.

[6]In what follows, we restrict the discussion to non-empty models.

– We want to show that $\mathcal{M} \vDash \forall I \forall xx(\exists yy\, \epsilon(I, xx, yy) \rightarrow \forall x(x \prec xx \leftrightarrow \exists zz\, (\epsilon(I, xx, zz) \;\wedge\; x \prec zz)))$. Let us take arbitrary I and aa of sorts σ_c, σ_1 respectively, such that $\mathcal{M} \vDash \exists yy \epsilon(I, aa, yy)$. Therefore, we have /that there is bb of sort σ_1, such that $\mathcal{M} \vDash \epsilon(I, aa, bb)$. Once again, because $\mathcal{M} \vDash \delta_\epsilon$, for the given I, aa, bb and because $\mathcal{M} \vDash \epsilon(I, aa, bb)$ we have, as before, that: $\mathcal{M} \vDash bb \leq aa \;\wedge\; \exists xxx(i_1(xxx) = I \;\wedge\; \forall x(x \prec aa \;\leftrightarrow \exists zz(x \prec zz \wedge zz \prec^1 xxx \wedge zz \leq aa)) \wedge bb \prec^1 xxx)$.

  This implies, from the second conjunct, that there exists aaa of sort σ_2 such that $\mathcal{M} \vDash i_1(aaa) = I \;\wedge\; \forall x(x \prec aa \;\leftrightarrow \exists zz(x \prec zz \leftrightarrow zz \prec^1 aaa \wedge zz \leq aa)) \wedge bb \prec^1 aaa$. This further implies that $\mathcal{M} \vDash \forall x(x \prec aa \;\leftrightarrow \exists zz(x \prec zz \wedge zz \prec^1 aaa \wedge zz \leq aa))$. It suffices to show that for any cc of sort σ_1 it is the case that $\mathcal{M} \vDash (cc \prec^1 aaa \;\wedge\; cc \leq aa) \leftrightarrow \epsilon(I, aa, cc)$. This follows directly from the definition of δ_ϵ, so we can deduce that $\mathcal{M} \vDash \forall x(x \prec aa \;\leftrightarrow \exists zz(x \prec zz \wedge zz \prec^1 aaa \wedge zz \leq aa)) \rightarrow \forall x(x \prec aa \leftrightarrow \exists zz\, (\epsilon(I, aa, zz) \;\wedge\; x \prec zz))$, and because the antecedent is satisfied in $\mathcal{M}$ that $\mathcal{M} \vDash \forall x(x \prec aa \leftrightarrow \exists zz\, (\epsilon(I, aa, zz) \;\wedge\; x \prec zz))$. Because I and aa were arbitrary, this shows exactly that $\mathcal{M} \vDash \forall I \forall xx(\exists yy\, \epsilon(I, xx, yy) \rightarrow \forall x(x \prec xx \leftrightarrow \exists zz\, (\epsilon(I, xx, zz) \;\wedge\; x \prec zz)))$, and because $\mathcal{M}$ was arbitrary that $T^*_{HLPL} \vdash \forall I \forall xx(\exists yy\, \epsilon(I, xx, yy) \rightarrow \forall x(x \prec xx \leftrightarrow \exists zz\, (\epsilon(I, xx, zz) \wedge x \prec zz)))$.

– We want to show that $\mathcal{M} \vDash \exists zz \phi(zz) \rightarrow \exists I \exists xx \forall yy(\phi(yy) \leftrightarrow \epsilon(I, xx, yy))$. Suppose that $\mathcal{M} \vDash \exists zz \phi(zz)$. We know that for this specific ϕ it is the case that $\mathcal{M} \vDash \exists zz \phi(zz) \rightarrow \exists xxx \forall xx(xx \prec^1 xxx \leftrightarrow \phi(xx))$ from (HLP-C). Because we know that $\mathcal{M} \vDash \exists zz \phi(zz)$, we have that $\mathcal{M} \vDash \exists xxx \forall xx(xx \prec^1 xxx \leftrightarrow \phi(xx))$, or equivalently that there exists aaa of sort σ_2 such that $\mathcal{M} \vDash \forall xx(xx \prec^1 aaa \leftrightarrow \phi(xx))$.

  From an instance of the plural comprehension axiom scheme, we have that $\mathcal{M} \vDash \exists xx(\forall x(x \prec xx \leftrightarrow \exists zz(zz \prec^1 aaa \;\wedge\; x \prec zz)))$, i.e., there exists bb of sort σ_1 such that $\mathcal{M} \vDash \forall x(x \prec bb \leftrightarrow \exists zz(zz \prec^1 aaa \;\wedge\; x \prec zz))$. We also have that for the given aaa, there exists I of sort σ_c such that $\mathcal{M} \vDash i_1(aaa) = I$. It is easy to verify now, based on δ_ϵ that $\mathcal{M} \vDash \forall yy(yy \prec^1 aaa \leftrightarrow \epsilon(I, bb, yy))$. Finally, because of $\mathcal{M} \vDash \forall xx(xx \prec^1 aaa \leftrightarrow \phi(xx))$, we get that $\mathcal{M} \vDash \forall yy(\phi(yy) \leftrightarrow \epsilon(I, bb, yy))$. From that we deduce that $\mathcal{M} \vDash \exists I \exists xx \forall yy(\phi(yy) \leftrightarrow \epsilon(I, xx, yy))$, which implies that $\mathcal{M} \vDash \exists zz \phi(zz) \rightarrow \exists I \exists xx \forall yy(\phi(yy) \leftrightarrow \epsilon(I, xx, yy))$. Because $\mathcal{M}$ was arbitrary we have that $T^*_{HLPL} \vdash \exists zz \phi(zz) \rightarrow \exists I \exists xx \forall yy(\phi(yy) \leftrightarrow \epsilon(I, xx, yy))$.

– We want to show that $\mathcal{M} \vDash \forall I \exists xx(\exists yy \epsilon(I, xx, yy) \wedge \forall zz \forall ww(\epsilon(I, zz, ww) \rightarrow xx = zz)$. Take an arbitrary I of sort σ_c. We know that $\mathcal{M} \vDash \exists xxx(i_1(xxx) = I)$, or equivalently that there exists aaa of sort σ_2 such that $\mathcal{M} \vDash i_1(aaa) = I$. From the plural comprehension axiom, we have that $\mathcal{M} \vDash \exists xx(\forall x(x \prec xx \leftrightarrow \exists zz(zz \prec^1 aaa \;\wedge\; x \prec zz)))$, i.e., there exists bb of sort σ_1 such that $\mathcal{M} \vDash \forall x(x \prec bb \leftrightarrow \exists zz(zz \prec^1 aaa \;\wedge\; x \prec zz))$. We also know that the right-hand side is satisfied for some $x = a$ of sort σ_0 and some cc of sort σ_1, because of (HLP-NE) for both $k = 0$ and $k = 1$. We still need to show that $\mathcal{M} \vDash \epsilon(I, bb, cc)$, but this follows directly from the construction of bb and that $\mathcal{M} \vDash \delta_\epsilon$.

  The final part is to show that $\mathcal{M} \vDash \forall zz \forall ww(\epsilon(I, zz, ww) \rightarrow xx = zz)$. From the way we have explicitly defined δ_ϵ, from plural and superplural extensionality, we have that the bb above is uniquely specified. The proof is also omitted. We have proven that $\mathcal{M} \vDash \forall I(\exists xx \exists yy \epsilon(I, xx, yy) \wedge \forall zz \forall ww(\epsilon(I, zz, ww) \rightarrow xx = zz)$, and because $\mathcal{M}$ was arbitrary, that

$$T^*_{HLPL} \vdash \forall I \exists xx(\exists yy\, \epsilon(I, xx, yy) \;\wedge\; \forall zz \forall ww(\epsilon(I, zz, ww) \rightarrow xx = zz)$$

- We want to show that $\mathcal{M} \vDash \forall I \forall J \forall xx(\forall yy(\epsilon(I, xx, yy) \leftrightarrow \epsilon(J, xx, yy)) \leftrightarrow I = J)$. Let there be arbitrary I, J of sort σ_c and arbitrary aa of sort σ_1. Assume that $\mathcal{M} \vDash \forall yy(\epsilon(I, aa, yy) \leftrightarrow \epsilon(J, aa, yy))$. By the δ_ϵ axiom we know that the aaa and bbb such that $\mathcal{M} \vDash i_1(aaa) = I \wedge i_2(bbb) = J$ will contain exactly the bb for which $\mathcal{M} \vDash \epsilon(I, aa, bb)$, so by the extensionality of superplurals, $aaa = bbb$, which implies that $I = J$. The converse follows immediately because $I = J$ and they both would apply to the same aa of sort σ_1. Therefore, we have proven that $\mathcal{M} \vDash \forall I \forall J \forall xx(\forall yy(\epsilon(I, xx, yy) \leftrightarrow \epsilon(J, xx, yy)) \leftrightarrow I = J)$, and thus that $T^*_{HLPL} \vdash \forall I \forall J \forall xx(\forall yy(\epsilon(I, xx, yy) \leftrightarrow \epsilon(J, xx, yy)) \leftrightarrow I = J)$.

- We want to show that $\mathcal{M} \vDash \delta_{\prec_1}$, i.e., that $\mathcal{M} \vDash \forall xxx \forall xx(xx \prec^1 xxx \leftrightarrow \exists I \exists yy(i_2(I) = xxx \wedge \epsilon(I, yy, xx)))$. Let aaa, aa be arbitrary objects of sorts σ_2 and σ_1 respectively, such that $\mathcal{M} \vDash aa \prec^1 aaa$. For aaa, we have that $\mathcal{M} \vDash \exists I(i_1(aaa) = I)$, or equivalently that there exists I of sort σ_c that $\mathcal{M} \vDash i_1(aaa) = I$, and from the $\delta_{i_1 i_2}$ axiom that $\mathcal{M} \vDash i_2(I) = aaa$. Now, as before, we define bb of sort σ_1 such that $\mathcal{M} \vDash \forall x(x \prec bb \leftrightarrow \exists zz(zz \prec^1 aaa \wedge x \prec zz))$. It is easy to verify now based on $\mathcal{M} \vDash \delta_\epsilon$ that $\mathcal{M} \vDash \epsilon(I, bb, aa)$. So we have proven that $\mathcal{M} \vDash aa \prec^1 aaa \rightarrow i_2(I) = aaa \wedge \epsilon(I, bb, aa)$, or equivalently that $\mathcal{M} \vDash aa \prec^1 aaa \rightarrow \exists I \exists yy\ (i_2(I) = aaa \wedge \epsilon(I, yy, aa))$.

  The converse direction follows immediately by δ_ϵ and $\delta_{i_1 i_2}$. Because aa, aaa were arbitrary, we have proven that $\mathcal{M} \vDash \delta_{\prec^1}$, and by the arbitrariness of $\mathcal{M}$ that $T^*_{HLPL} \vdash \delta_{\prec^1}$.

- Finally, we will show that $\mathcal{M} \vDash \delta_{\sigma_2}$, i.e., that $\mathcal{M} \vDash \forall I \forall J(i_2(I) = i_2(J) \leftrightarrow I = J) \wedge \forall xxx \exists I(i_2(I) = xxx)$. For arbitrary I, J of sort σ_c, assuming that $\mathcal{M} \vDash i_2(I) = i_2(J)$, we have that $\mathcal{M} \vDash i_1(i_2(I)) = i_1(i_2(J))$, because i_1 is a function symbol interpreted as a total function. But there exists aaa of sort σ_2 such that $\mathcal{M} \vDash i_2(I) = aaa \wedge i_2(J) = aaa$, and from the explicit definition of the function symbols, we have that $\mathcal{M} \vDash i_1(aaa) = I \wedge i_1(aaa) = J$, and so that $\mathcal{M} \vDash I = J$. The converse follows directly from the functionality of i_2.

  Now we need to show that $\mathcal{M} \vDash \forall xxx \exists I(i_2(I) = xxx)$. Take an arbitrary aaa of sort σ_2. It follows from the non-emptiness of superplurals that there exists aa of sort σ_1, such that $\mathcal{M} \vDash aa \prec^1 aaa$. Since we have proven that $\mathcal{M} \vDash \delta_{\prec^1}$ we have that $\mathcal{M} \vDash \exists I \exists yy(i_2(I) = aaa \wedge \epsilon(I, yy, aa))$, which gives us precisely what we wanted.

  Because I and J were arbitrary, we have that $\mathcal{M} \vDash \delta_{\sigma_2}$, and because $\mathcal{M}$ was chosen arbitrarily, we also get that $T^*_{HLPL} \vdash \delta_{\sigma_2}$.

Because all the axioms of T_{CPL} regarding only sorts σ_0 and σ_1 are included in T_{HLPL} already, we have shown that $T^*_{HLPL} \vdash T^*_{CPL}$. Let us note here that because T_{HLPL} is satisfiable and T^*_{HLPL} results from T_{HLPL} from a finite sequence of Morita extensions, T^*_{HLPL} is satisfiable as well (McEldowney, 2020). Because of the first part of the proof, this also shows that T^*_{CPL} (and hence T_{CPL} as well) is satisfiable.

Now we will prove the other direction, i.e., that $T^*_{CPL} \vdash T^*_{HLPL}$. Let $\mathcal{M}$ be an arbitrary model of T^*_{CPL}:

- We will start by proving that $\mathcal{M} \vDash (HLP - C)$, i.e., that $\mathcal{M} \vDash \exists xx\phi(xx) \rightarrow \exists xxx \forall xx(xx \prec^1 xxx \leftrightarrow \phi(xx))$, for some formula ϕ that has no occurrence of xxx and contains at least xx as free. Suppose that $\mathcal{M} \vDash \exists xx\phi(xx)$. We know that $\mathcal{M} \vDash \exists zz\phi(zz) \rightarrow \exists I \exists xx \forall yy(\phi(yy) \leftrightarrow \epsilon(I, xx, yy))$, because $\mathcal{M} \vDash T^*_{CPL}$. Therefore, we have that $\mathcal{M} \vDash \exists I \exists xx \forall yy(\phi(yy) \leftrightarrow \epsilon(I, xx, yy))$, i.e., that there exist I of sort σ_c and aa of sort σ_1, such that $\mathcal{M} \vDash \forall yy(\phi(yy) \leftrightarrow \epsilon(I, aa, yy))$.

Additionally, we know that there exists aaa of sort σ_2 such that $\mathcal{M} \vDash i_2(I) = aaa$. We also know that $\mathcal{M} \vDash \delta_{\prec^1}$, i.e., that for the given I, aa, aaa we have $\mathcal{M} \vDash \forall yy(yy \prec^1 aaa \leftrightarrow i_2(I) = aaa \ \wedge \ \epsilon(I, aa, yy))$.

From this we immediately deduce that $\mathcal{M} \vDash \forall yy(yy \prec^1 aaa \leftrightarrow \phi(yy))$, which implies that $\mathcal{M} \vDash \exists xxx \forall yy(yy \prec^1 xxx \leftrightarrow \phi(yy))$, which gives us exactly that $\mathcal{M} \vDash (HLP - C)$ and, because $\mathcal{M}$ was arbitrary, that $T^*_{CPL} \vdash (HLP - C)$.

– Moreover, we will prove that $\mathcal{M} \vDash (HLP - NE)$, i.e., that $\mathcal{M} \vDash \forall xxx \exists xx(xx \prec^1 xxx)$. This follows immediately because for every superplural there is a corresponding cover, and a cover always has a plural to which it applies (from axiom 4 regarding covers). We then use the fact that $\mathcal{M} \vDash \delta_{\prec^1}$ to get the result. The proof here is omitted.

– Then, we will prove that $\mathcal{M} \vDash (HLP - Ext)$, i.e., that $\mathcal{M} \vDash \forall xxx \forall yyy(\forall xx(xx \prec^1 xxx \leftrightarrow xx \prec^1 yyy) \leftrightarrow xxx = yyy)$. Take arbitrary aaa, bbb of sort σ_2. From $\delta_{\prec^1}$ we have that $\mathcal{M} \vDash \forall xx(xx \prec^1 aaa \leftrightarrow xx \prec^1 bbb)$ is equivalent to $\mathcal{M} \vDash \forall xx(\exists I \exists yy(i_2(I) = aaa \wedge \epsilon(I, yy, xx)) \leftrightarrow \exists J \exists zz(i_2(J) = bbb \ \wedge \ \epsilon(J, zz, xx)))$. However, because any given cover applies to a unique plurality and i_2 is a function symbol, we have that $\mathcal{M} \vDash \exists I \exists J \exists ! yy \exists ! zz \forall xx(\epsilon(I, yy, xx) \leftrightarrow \epsilon(J, zz, xx))$. But we also know that the specific yy and zz, because of the second axiom for covers, have exactly the same elements, and by plural extensionality, they are identical. Because each cover applies to a unique plurality and because of the extensionality for covers (axiom 5), we have that $I = J$, and hence that $aaa = bbb$. The converse direction follows directly from $\delta_{\prec_1}$. Because aaa, bbb were arbitrary, we have that $\mathcal{M} \vDash (HLP - Ext)$ and because $\mathcal{M}$ was arbitrary we have that $T^*_{CPL} \vdash (HLP - Ext)$.

– Now turning to the axioms extending T_{HLPL} to T^*_{HLPL}, we want to show that $\mathcal{M} \vDash \delta_{\sigma_c}$, i.e., that $\mathcal{M} \vDash \forall xxx \forall yyy(i_1(xxx) = i_1(yyy) \leftrightarrow xxx = yyy) \ \wedge \ \forall I \exists xxx(i_1(xxx) = I)$. The proof here is completely analogous to the previous one where we proved that for each $\mathcal{N}$ such that $\mathcal{N} \vDash T^*_{HLPL}$, it is the case that $\mathcal{N} \vDash \forall I \forall J(i_2(I) = i_2(J) \leftrightarrow I = J) \wedge \forall xxx \exists I(i_2(I) = xxx)$.

– Finally, we want to show that $\mathcal{M} \vDash \delta_\epsilon$, i.e., that $\mathcal{M} \vDash \forall I \forall xx \forall yy(\epsilon(I, xx, yy) \leftrightarrow yy \leq xx \ \wedge \ \exists xxx(i_1(xxx) = I \wedge \ yy \prec^1 xxx \ \wedge \forall x(x \prec^0 xx \wedge \exists zz(x \prec^0 zz \ \wedge \ zz \prec^1 xxx \ \wedge \ zz \leq xx)))$. Let I be an arbitrary object of sort σ_c and aa, bb two arbitrary objects of sort σ_2. First, suppose that $\mathcal{M} \vDash \epsilon(I, aa, bb)$. From the axioms, we know that $\mathcal{M} \vDash bb \leq aa$ and that $\mathcal{M} \vDash \exists xxx \ i_2(I) = xxx$, or equivalently that there exists a aaa such that $\mathcal{M} \vDash bb \leq aa \ \wedge \ i_2(I) = aaa$, which by the axiom $\delta_{i_1 i_2}$ is equivalent to $\mathcal{M} \vDash bb \leq aa \ \wedge \ i_1(aaa) = I$. Also, because $\mathcal{M} \vDash i_2(I) = aaa \ \wedge \ \epsilon(I, aa, bb)$, by the axiom $\delta_{\prec^1}$ we have that $\mathcal{M} \vDash bb \prec^1 aaa$. Then, by the axioms governing ϵ, we have that $\mathcal{M} \vDash \forall x(x \prec xx \leftrightarrow \exists zz(x \prec zz \ \wedge \ \epsilon(I, aa, zz)))$. By applying the axiom δ_{σ_c}, we also get that $\mathcal{M} \vDash \forall x(x \prec xx \leftrightarrow \exists zz(x \prec zz \ \wedge \ \epsilon(I, aa, zz))) \rightarrow \forall x(x \prec xx \leftrightarrow \exists zz(x \prec zz \ \wedge \ zz \prec^1 aaa \ \wedge \ zz \leq aa)$ and also that $\mathcal{M} \vDash bb \prec^1 aaa$. Because I, aa, bb were arbitrary, we have proven that $\mathcal{M} \vDash \delta_\epsilon$, and because $\mathcal{M}$ was arbitrary we have proven that $T^*_{CPL} \vdash \delta_\epsilon$.

Therefore, we have proven also that $T^*_{CPL} \vdash T^*_{HLPL}$. This concludes the proof that the two theories are *logically equivalent*. □

And finally, from the work of McEldowney (2020), we have also the following corrolary:

Corollary 15. *The Σ_1-theory T_{HLPL} and the Σ_2-theory T_{CPL} are bi-interpretable.*

Proof. The proof is based on the fact that in extending the theories, we did not use any coproduct sorts. Then, from the work of McEldowney (2020) we infer that Morita equivalence, in this case, implies bi-interpretability. □

3.6 The extended result

Suppose now that we want to talk about plurals, superplurals, and covers using specific predicates. We are not going to focus on the case of adding relation symbols of sort σ_0 as this is dealt with in a straightforward manner. We have a very specific aim here. We want to see how CPL with predicates for plurals and HLPL with predicates for superplurals can come to say the same things.

The claim of the cover-based semanticists is that every predication needs to be indexed by a specific cover. Instead of writing $\exists xx Fxx$, one should make explicit according to which cover the predication is to happen by writing $\exists xx^i Fxx^i$, meaning that the predicate distributes over the subpluralities that i splits xx into. In our framework in MSL, this will be captured by saying that $\exists xx \exists I(\forall yy\, \epsilon(I, xx, yy) \rightarrow Fyy)$, where Fyy denotes standard collective plural predication. This will correspond in a natural way to how superplural predicates are supposed to work.

Before working on the formal details, I want to stress that now we are going to add relation symbols and prove the interpretability result without adding any new axioms governing those relation symbols. As the first proof was supposed to illustrate the formal proximity of the notions of a superplural and that of a cover, this proof is supposed to show the relevant proximity between superplural predication and predication with respect to a cover.

3.6.1 The correct signatures and theories

We want to compare two theories that try to capture the different higher-level plural phenomena. Regarding singular terms, both theories will share the same predicates, in the sense that they make the same claims about them, and that they can form the same plurals (due to the plural comprehension axiom). We will also need plural predicates for both HLPL and CPL. Then, we will add to CPL two-place predicate symbols for a cover and a plural to show how a cover distributes a *plural* predicate over the plural's subpluralities. In a similar vein, HLPL will be endowed with infinitely many *distributive* superplural predicates. Therefore, the only extra axioms we need will connect the two-place predicates with plural predications in the former case, and the distributive superplural predicates with plural predications in the latter one.

Let us proceed with the formalization. We define HLPL as the theory T_{HLPL} over the language generated by the signature Σ_1.[7] This signature has the three sort symbols σ_0, σ_1, and σ_2, the two-place special relation symbols $\prec$ of arity $\sigma_0 \times \sigma_1$ and $\prec^1$ of arity $\sigma_1 \times \sigma_2$, capturing the corresponding notions of containment. We have countably many predicates P^1_n, for all $n \in \omega$ of sort σ_0, countably many predicates P^2_n, for all $n \in \omega$ of sort σ_1, and countably many predicates P^3_n, for all $n \in \omega$ of sort σ_2.

The axioms of the Σ_1-theory T_{HLPL} as exactly those as before, containing all the axioms for plurals and superplurals, and additionally the axiom scheme $\forall xxx(P^3_n(xxx) \leftrightarrow \forall xx(xx \prec^1 xxx \rightarrow P^2_n(xx)))$ distributing the superplural predicate over the plurals contained in each superplural. The difference compared to the previous case is that in the axiom **HLP-C** there are more pluralities (of both levels)

[7] Note that Σ_1 now is different from the one used before. It is actually a larger signature.

to be defined, because our language is currently more expressive than it was. A similar thing will happen to CPL.

We define CPL as the theory T_{CPL} over the language generated by the signature Σ_2. This signature has the three sort symbols σ_0, σ_1, and σ_c, the two-place relation symbol $\prec$ of arity $\sigma_0 \times \sigma_1$, the three-place relation symbol ϵ of arity $\sigma_c \times \sigma_1 \times \sigma_1$, capturing how a cover applies to a plural and splits it into subplurals. We have countably many predicates P_n^1, for all $n \in \omega$ of sort σ_0, countably many predicates P_n^2, for all $n \in \omega$ of sort σ_1, and countably many predicates P_n^c, for all $n \in \omega$ of sort $\sigma_c \times \sigma_1$.[8]

The axioms of the Σ_2-theory T_{CPL} as exactly those as before, containing all the axioms for plurals and covers, and additionally the axiom scheme $\forall I \forall xx(P_n^c(I, xx) \leftrightarrow \forall yy(\epsilon(I, xx, yy) \rightarrow P_n^2(yy)))$ distributing the plural predicate over its subplurals given to us by a cover.

3.7 Logico-Philosophical Remarks

Before we conclude this section, I want to raise some final points about the equivalence just proven. Firstly, I want to point out that in regimenting CPL as the theory T_{CPL}, I chose to define covers as functions with a unique plurality in their domain; this choice corresponds directly to extending plural logic with superplurals. If we allowed for the covers to be functions with a domain containing a plural and some of its subpluralities, then we would need third-level plurals to capture them.

The idea for such a proof would be very similar in character; the reason we avoided doing that is that proving the theoretical equivalence of the two theories as regards capturing superplurals is enough to capture most (if not all) natural language phenomena. One could protest that not everything about natural language is saved in the process because we did not talk of the *cumulative* version of superplurals. While strictly speaking that is true, formally nothing of much significance is lost here because we allowed for singleton pluralities. I argued in the previous chapter about the conceptual inadequacy of singleton pluralities, but on the formal front, it allows for a formal regimentation of sentences involving cumulative superplurals, for instance, lists of pluralities and individuals, that outputs the right truth conditions. Therefore, when covers split up a plurality into subpluralities, this split in effect also involves individuals as singleton pluralities; similarly, when a superplural contains plurals, it contains individuals as singleton pluralities as well.

With this technically adequate machinery, we are able to capture all relevant natural language phenomena, albeit not in the most intuitively appealing manner. Although natural language is not my direct aim, it is the aim of the advocates of the cover-based approach. By proving that the two frameworks are Morita equivalent and bi-interpretable, I have secured that, on the formal side, our frameworks are capable of getting the right truth conditions for the same phenomena.

Obviously, this does not settle the matter on how HLP terms function in natural language, neither does it establish that both approaches are correct (if any of them is). The correctness of any analysis of HLP requires delving deep into the workings of natural language and investigating whether one of the two intended interpretations of the frameworks captures those workings more accurately. What I have established here is that neither framework is formally inferior to the other.

The second remark is about the choice of carrying out the comparison of the two logics in the framework of MSL. In doing so, formal choices about how to best formalize both HLPL and CPL were

[8] I should note at this point that I have used the same relation symbols for the common sorts of these two theories to showcase that they make the same claims regarding those sorts. In a more careful formulation of the problem, I would use different relation symbols.

made; I do not see any way in which these choices impact my argument negatively, as I believe that I respected the initial formalizations as much as possible. The only possible choice that may seem arbitrary is the choice to have covers apply to unique plurals and not multiple ones. This is a choice that simplifies the proofs and not one that impacts the overall scope of application of CPL, as there are enough covers to carve any plurality in any way one wishes to do so. Since any predication of a plural is indexed to a cover, the appropriate cover can always be chosen without any qualms as to its existence. Whenever, in HLP parlance, we have a genuinely superplural predicate applying to a superplural, we know that there is always going to be a plural term, a plural predicate, and a cover, so that the relevant predications both come out as true or as false. This point is also supported by the fact that from the way that the models of T_{CPL} and T_{HLPL} were regimented as Σ_2 and Σ_1 structures, the domains of the sorts σ_c and σ_2 have the same cardinalities, meaning that there is a bijection between them.

A third remark has to do with the formulation of the theories without any non-logical predicates. Basically, we only talked of the relation symbols capturing containment and the application of covers to plurals without any other kinds of plural or superplural predicates. In the following chapter, we will discuss the presence of non-logical predicates in HLPL, and as Linnebo and Rayo (2012) mention, avoiding them seems like an artificial restriction to impose. Therefore, it should be pointed out that when involving non-logical predicates in the language, we will have to make sure that the theoretical equivalence results go through just the same.

I have good reason to believe that this would be the case. We mentioned in Section 3.4.2 that in CPL predication over plurals always involved covers; for any non-logical predicate F, we would have $\exists xx^i Fxx^i$ instead of $\exists xx Fxx$, where the indexing shows the appropriate cover to be used. In the MSL version, this would be captured by saying that $\exists xx \exists I(\forall yy(\epsilon(I, xx, yy) \rightarrow Fyy))$, meaning that the predication is distributed to the plurals yy the cover associates to xx. Collective predication is simply captured by the cover I such that $\epsilon(I, xx, xx)$. Then, we know how to translate this way of predicating to the HLPL version, because for this cover I there is a unique superplural aaa to which it corresponds, and then F will be considered a superplural predicate; $\exists xx \exists I(\forall yy(\epsilon(I, xx, yy) \rightarrow Fyy))$ would then be roughly equivalent to $\exists xxx(i_1(xxx) = I \ \wedge \ Fxxx)$. Singular predication is the same in both CPL and HLPL, so that poses no problems.

The other direction is very similar in character. Every time we have a superplural predication of the form $\exists xxx Fxxx$, then we simply associate xxx with the corresponding cover I and the plural xx that contains only the individuals contained in plurals contained in xxx. Then $\exists xxx Fxxx$ is roughly equivalent to $\exists I \exists xx(i_2(I) = xxx \ \wedge \ \forall yy(\epsilon(I, xx, yy) \rightarrow Fyy))$. Predicates of higher levels all become plural predicates applying to pluralities indexed to a cover. This is but a mere sketch of how the proofs would be extended, but I believe that it shows that a more general result is but a stone's throw away.

The final, and most important point, which was also mentioned at the outset of this chapter, is the fact that the model theory may not capture precisely all the models we would associate with these logics. I take this to be a significant challenge to this approach, which is why I outlined alternative ways to deal with it previously.

3.8 Conclusion

I began by introducing many-sorted logic as the framework where the formal comparison between HLPL and CPL should be carried out. I outlined some of the limitations of this approach and ways in which it can be dealt with, ultimately concluding that presenting the formal comparison in this way is an important first step to any later one.

I, then, introduced *Morita Equivalence* as a notion of theoretical equivalence, which lends itself to natural comparisons between many-sorted first-order theories. Some remarks were also made about the rival notion of *Bi-interpretability*. I proceeded by formulating both HLPL and CPL as many-sorted first-order theories by precisely describing the axioms for both. In the case of CPL, where no precise enough formulation of the axioms existed, I attempted to formalize the commitments of the proponents of the cover-based approach in many-sorted logic. With the formalizations in place, I showed that the two theories are both Morita Equivalent and Bi-interpretable.

The main takeaway from this chapter, precisely due to the previous result, is that formally, there is little to be considered about the choice between CPL and HLPL. The divergence between the two frameworks happens on the conceptual level, and whether they are adequate for their own purposes. I do not presume that I have established any conceptual superiority of HLPL contra CPL in how they deal with natural language phenomena. This is the main motivation for CPL, but due to their theoretical equivalence, we know that both frameworks will be adequate in carrying out the role of a formal semantics. To understand and grasp a notion of reference for higher-level plurals, HLPL is clearly more suitable, which is why it is given priority in this thesis.

Concerning natural language, there are already several worries mentioned in the previous chapter. It is not entirely clear how covers are to be determined in each case, and, as I mentioned, they seem to be dependent on context in a manner more akin to pragmatics than to semantics. Additionally, the linguistic data from Grimau (2019), especially those about the lexicalized HLP terms in some languages, seem to show that humans employ some form of superplural reference; whether that is the same as *Combinatorial Reference*, or not, is yet to be decided.

Furthermore, precisely because of the formal equivalence between HLPL and CPL, someone who wants to genuinely discern between the two frameworks should provide a good story about how they differ with respect to their interpretation. Grimau's view of HLP reference and the cover-based approach share a multitude of similarities on the conceptual level as well, specifically due to the connections of their respective notions of reference to plural properties. In that regard, little is to be discerned between the two approaches; my approach differs from both substantially in that it rests on a different view of reference based on the notion of containment. My main critique of the cover-based approach is that the choice of a cover is heavily context and pragmatically dependent; both of the face-value approaches (mine and Grimau's) avoid that shortcoming. I believe then that there is some evidence for the semantic superiority of HLP reference contra the cover-based approach.

Chapter 4

Height of Hierarchy

Until this point, we have concentrated purely on plurals of finite levels. As soon as we allow the iteration from plurals to superplurals, it is fairly straightforward to keep that iteration going for all finite levels. In this chapter, I want to pursue that task even further. I wish to understand the limits of this iteration, if there are any, by extending it to transfinite levels indexed by ordinals.

This question will be broken down into two parts: (a) what is the largest ordinal α for which an α-level plural exists, and (b) what is the largest ordinal α for which an α-level plural *language* is considered *legitimate*. The former question will be dealt with in a manner akin to that of the second chapter; the main difference will be that a grasp of *Combinatorial Reference* will require *transfinite induction* instead of simple induction on the natural numbers. In some sense, we will be generalizing the intuitions coming from the plurals of finite levels in a straightforward manner, provided that we have access to a rich enough theory of ordinals.

The latter question requires more care. First, we will have to discuss the notion of *language legitimacy* to better understand what it means to allow for a language to be admitted to our philosophical theorizing. Then, we see how different kinds of higher-level plural languages can be deemed legitimate and which are the relevant requirements for their legitimacy; here, the connection between the two questions will become apparent, as countenancing a language with α-level terms without a grasp of how these terms refer would seem at best counterintuitive.

After having established that conclusion - the legitimacy of HLPL - we will turn to an alternative strategy due to Linnebo and Rayo (2012) that purports to establish a more ambitious claim that would entail the legitimacy of HLPL. I will argue that my proposal reaps some of the benefits without falling prey to the shortcomings of theirs. The main critiques of that strategy come from Studd (2019, 2021) and Button and Trueman (2022). We first discuss Studd's work to highlight the difficulty that Linnebo and Rayo's arguments face, before turning to Button and Trueman, whose points will impact how we view the formalization of higher-level plural logic.

4.1 Plural Hierarchy

Determining the height of the plural hierarchy is our first order of business. Once again, talk of the "height of the plural hierarchy" is not intended to be about a hierarchy of *entities*, but a mere *façon de parler*. It simply serves the purpose of putting more succinctly the question: for which ordinals α can we grasp how α-level plural terms would refer to individuals according to *Combinatorial Reference*?

This question should be decoupled from the legitimacy of α-level plural languages, although as we will see later on, a positive or negative answer to the former impacts the outcome of the latter.

My aim, more specifically, is to extend our conception of *Combinatorial Reference* accordingly to account for plurals of larger and larger levels into the transfinite. Let us start by restating the way in which plurals of finite levels refer:[1]

Combinatorial Reference*: An $(n+1)$-level plurality refers to some individuals in virtue of all the k-level pluralities for all $k \leq n$ it contains and how they refer to individuals.[2]

The role of a higher-level plural term in a sentence is to indicate how some individuals are referred to in virtue of the way in which they are combined. That way of referring is explained by the notion of *Combinatorial Reference*, which in turn relies on how individuals and plurals are contained in other plurals. When I first described this kind of reference in its non-cumulative form (Section 2.1.3), I mentioned the prerequisites to grasping that notion of reference in more detail. Firstly, one needs to know how singular terms and plural terms refer to individuals. Then, one needs to learn how to iterate the notion of containment so that higher-level plural terms contain plurals of the immediately preceding level. The last piece of the puzzle is the ability to do induction on the natural numbers: we first showed how superplurals refer (base case), and then given how n-level plurals refer to individuals, one knows how $(n+1)$-level plurals refer.

The components necessary to grasp the cumulative version of *Combinatorial Reference* are the same, apart from the inductive step, which requires *strong induction* on the natural numbers.[3] The base case is similar: superplurals containing plurals and singular terms and referring to individuals in virtue of them. The inductive step: suppose we know how k-level plurals refer to individuals for all $k < n$. Now we can also know how n-level plurals refer to individuals, because we know how all the plurals contained in them refer to individuals.

An equivalent way to view this would be the following. Instead of looking at all the plurals of previous levels, we simply look at the plurals of the immediately preceding level, which now also include plurals *of all previous levels*. Thus, when a plural shows up in one level of the hierarchy, it also shows up in all subsequent levels. Then, we could just stick to the original formulation of *Combinatorial Reference*, but essentially nothing of importance depends on this.

The important part is that as long as a pluralist knows how to iterate the notion of containment and has a grasp of a fragment of arithmetic strong enough to carry out the mathematical induction needed here, then that pluralist will have access to *Combinatorial Reference* (CR). The pluralist can avail herself of:

ω-level CR:[4] All n-level plurals, for $n < \omega$, refer according to *Combinatorial Reference*.

The question I now turn to is whether she can do the same with α-level CR for any ordinal α. That would mean that for all ordinals $\beta < \alpha$, the plurals of level β would be able to refer according to CR. As with all constructions involving transfinite ordinals, the first order of business is to check the case of limit ordinals.

[1] A brief reminder: higher-level plurals are cumulative with respect to the containment relation.

[2] Another reminder: singular terms are viewed as 0-level plurals.

[3] Formally, the two kinds of induction mentioned here are equivalent, but I take it that the move to strong induction should be mentioned nonetheless.

[4] Henceforth, I drop the * superscript as all higher-level plurals will be cumulative.

So, what would it take for an ω-level plural to refer? According to CR, an ω-level plural would refer to some individuals in virtue of the pluralities it contains and how they refer to individuals. According to our current conception of plurals, extending them to transfinite levels would mean that an ω-level plural should be able to contain plurals of all finite levels. Were we to extend ω**-level CR** to $(\omega + 1)$**-level CR** uniformly, we would ask that ω-level plurals refer according to the pluralities of finite levels they contain and how *they* refer to individuals. But we already know what it means for all n-level plurals to refer to individuals, so assuming that we grasp how a plural can contain plurals of all finite levels we have a way to make the move from ω**-level CR** to $(\omega + 1)$**-level CR**. In fact, for any limit ordinal λ, if we know how all higher-level plurals of level less than λ refer, then we know how λ-level plurals refer. So, once again, a move from λ**-level CR** to $(\lambda + 1)$**-level CR** is immediate.

Turning once again our attention to *containment*, we should stress the significance of it in making sense of plurals that correspond to limit ordinals. In set-theoretic language, containment has a characteristic of the operation of *union*; every time we have some pluralities of some levels, we can have a plurality that can contain all of them. In other words, we can countenance a plurality whose level is the supremum (or maximum) of all the levels the previous pluralities had. Expressing in a way more in line with how we understand plurals, suppose that we have some plural terms (of possibly different levels) and we know how each of them refers to individuals; then, we can make sense of a plural term that refers to individuals *in virtue of containing* all the previous pluralities. This could be considered as a version of Linnebo and Rayo's **Principle of Union** (Linnebo and Rayo, 2012), but one which does not suffer from the same shortcomings as theirs. We will return to these shortcomings as soon as we discuss Studd's critique of it (Studd, 2021).

Generally then, if we understand singular, plural, superplural reference, can iterate *containment*, and have access to *transfinite induction*, instead of just induction on the natural numbers, we have access to α**-level CR** for all ordinals α. With these prerequisites we can iterate the containment relation along any well-order[5] and get pluralities of ever increasing levels. The case for successor ordinals is dealt with exactly as in the finite case, and the case for limit ordinals is as described in the previous paragraph. We will be referring to this fact as follows:

Plural Axiom of Ordinals: For any ordinal α we have α**-level CR**.[6]

As it may already be evident, arriving at the axiom of ordinals makes some hefty demands on our meta-theory due to the requirement of transfinite induction and access to the whole class of ordinals. These demands are related to criticisms aimed at type-theoretic frameworks which require a theory of ordinals in the meta-theory to supply one with the relevant types.[7]

However, for us, a potential and more pressing problem does not lie in the availability of the ordinals, but rather in the *justification* for the Plural Axiom of Ordinals. More specifically, someone may have qualms as to why we should accept that plurals should be governed by a *union principle*, that legitimizes considering plurals of α-level, where α is a limit ordinal. The union principle basically tells us that the containment relation can be iterated into the transfinite as well.

[5]Reminder: a set A endowed with a relation $\leq$ is a well-order iff $\leq$ is a total order and every non-empty subset of A has a $\leq$-least element.

[6]The name is inspired by the set theoretic Axiom of Ordinals (Potter, 2004, Chapter 13.2).

[7]See for instance Linnebo and Rayo (2012) who make similar claim about the meta-theoretic demands for accepting their results, and Button and Trueman (2022) who compare such type-theoretic approaches to un-typed frameworks powerful enough to define the "types" internally.

Bear in mind, once again, that we are not trying to argue that any such plural terms exist in natural language. However, given our understanding of the containment relation, as a relation that is governed, by-and-large, by the same structural features as the singular-plural containment, we *could* keep iterating it to apply to HLP terms of transfinite levels. On the semantic level, this would rely on our ability to understand how the way in which things on the left side of the relation of containment refer contributes to the way in which things on the right side refer.

Then, *if* we were to countenance limit-level terms, e.g., ω-level terms, we would need them to figure on the right-hand side of *containment*. By our understanding of CR, ω-level terms would refer in virtue of the finite level plurals they contain, and we already possess a grasp of how these plurals would refer. Therefore, we would have a uniform way, provided by CR, to understand the reference of limit-level plurals. This does not impose that containment *must* be iterated into the transfinite, but it is enough to accept that we *could* iterate it and also grasp how the corresponding terms would refer by using a direct generalization of CR. A generalization that is directly supported by the transfinite induction mentioned above.

As a result, I take it that the proponent of CR, who also accepts enough set theory to have access to the class of ordinals and transfinite induction, has ample, but defeasible, reasons to use the **Plural Axiom of Ordinals**, unless important considerations undermine the union principle for containment.

In the next section, we will focus on fleshing out the relation between **α-level CR** and the legitimacy of an α-level plural language. We will show that under an appropriate sense of legitimacy, accepting the **Plural Axiom of Ordinals** suffices for countenancing α-level plural languages, for all ordinals α. In whatever follows, we assume that the axiom is justified and we make free use of it.

4.2 Language Legitimacy

The primary goal of this section is to figure out what it takes for a formal language to be considered *legitimate*. A formal language here is simply a collection of symbols along with some formation rules that explain what counts as a well-formed formula. We can even endow this language with an axiomatic system and an entailment relation, just as we did above when presenting the axiomatic system for HLPL.

Generally, when confronted with a formal language that includes some specific expressive resources, such as infinitely many different and irreducible kinds of quantification, we have to consider whether we are in a position to consider such a language as legitimate. This legitimacy is *prima facie* an obvious, necessary condition for its applicability in philosophy.

The specific question we want to tackle is whether HLPL, including plurals of level α for all ordinals α, can be considered legitimate. Here, the language we deal with is not simply a formal tool, but is accompanied by an intended interpretation. Terms of level α are supposed to function like α-level plural terms, which refer to individuals according to CR.

Using the terminology of the previous section, the even more specific question is whether a grasp of **α-level CR** suffices for countenancing an α-level plural language, and if so, what kind of level-language that will be. In the remainder of the section, we will introduce two kinds of legitimacy for a language and two distinctions as regards the kinds of typed languages one may be interested in. This will not only be done to clarify the goals of this proposal, but also to set the stage for a discussion regarding the available alternatives for countenancing higher-level languages. Let us deal with these matters in turn.

Legitimacy:

We will follow Studd (2021), who draws from Lewis (2010), to distinguish between two kinds of language legitimacy. David Lewis distinguishes between *languages* and *language*. The latter refers to language as we use and engage in it as "a form of rational, convention-governed human social activity" (Lewis, 2010, p. 7). Languages, on the other hand, are supposed to be (interpreted) languages in the logician's sense. In our case, a formal language will be considered as *thinly legitimate* if there is an interpretation function mapping the (non-logical) expressions of the language to their intended semantic values.

Thickly legitimate languages are thinly legitimate languages accompanied by a connection between them and *language* in Lewis' sense. So, we are talking about an interpreted language with an appropriate interpretation function, for which we can explain how it is connected to actual language usage.

The reason we introduce this distinction is that it is important both for the philosophical motivation behind HLPL and the range of applications that the framework could have. For instance, as we will see shortly, the version of HLPL that allows for transfinite levels incorporates a deductive system with infinitary rules of inference. Even if that language were *thinly* legitimate, for many it would not be considered as *thickly* legitimate due to the inability of humans to carry out an argumentation involving infinitely many premises.

However, there is a possible distinction that can be made even within the realm of thinly legitimate languages that is not directly discussed by Lewis or Studd. This distinction is about the grasp we have of the intended semantic values of a thinly legitimate language; as discussed in Chapter 2, Simons (1982) maintains that offering a *formal* account of the semantics of his manifolds, and our HLP, by the way in which they result from grouping individuals into plurals, etc., is immediate, but that does not ensure that we employ this notion of reference. Therefore, we may be able to offer an interpretation function to the intended semantic values of a language - ensuring thin legitimacy- but the semantic values may be in principle, be incomprehensible to us.

This is to be differentiated by thick legitimacy, simply because semantic values need not be understood only by relating them to actual language usage. The way I have argued about the meaningfulness of HLP rests on already acquired concepts, such as singular and plural reference, alongside a relevant principle of induction. Through these concepts, we can comprehend CR, thus having a good grasp of the semantic values we are interested in. An interpretation function mapping the non-logical vocabulary of an HLP language to these semantic values would ensure the thin legitimacy of that language; the grasp we have of those values would ensure a slightly further form of legitimacy firmly situated in-between thin and thick legitimacy.

In any case, our initial aim will be *thin* legitimacy, having already argued for the plausibility of the semantic values of plurals. Thick legitimacy, a prerequisite of which is thin legitimacy anyway, seems to be too strong a requirement of HLP, but also not necessarily a needed one for its philosophical applicability. Since we have a principled manner in which to understand HLP resources, the addition of the thin legitimacy of these languages should be enough to ensure their role in philosophical theorizing.

We turn to the second distinction, the one between languages involving terms of different levels, since they may differ with respect to the status of their legitimacy, and because of their interconnections regarding the ability to provide a model theory for them. Both of these aspects will be important for the next sections and especially for the exposition of the arguments found in Linnebo and Rayo (2012).

Kinds of Level Languages:

There are ways to separate the languages we will be interested in; one will be regarding the existence (or not) of non-logical predicates, while the other has to do with the constant symbols allowed. This point needs to be treated with care because Rayo and Studd have introduced relevant terminology corresponding to slightly different logical systems.

Whereas Rayo refers to higher-level plural languages, where the typed syntactic entities are all terms, Studd's main example is a simple type theory where the typed entities are predicates. From Rayo (2006) we know that there is an immediate way to provide a translation between HLPL and simple type theory, but they nevertheless lend themselves to different conceptual and formal treatments in some respects. I will return to this point later on, but for now, it would be good to have in mind that having constant symbols for higher-level plurals and allowing for non-logical predicates are two very different things.

We are now going to provide both classifications of languages that are invoked by Rayo (2006) and Studd (2021) by slightly altering their terminology to avoid any conflations. We start with Rayo's classification:

> *Basic α-level language:* An α-level language containing quantifiers and terms of level α that does not have predicates taking $(\alpha+1)$-level terms as arguments.
>
> *Complete α-level language:* An α-level language that results from extending a *basic* α-level one with non-logical predicates taking $(\alpha+1)$-level terms as arguments.[8]

The difference between a basic language and a complete one is analogous to that between PFO and PFO+, where the latter contains plural predicates while the former does not. Therefore, this distinction concerns *predication*. The following one due to Studd (2021) concerns *levels* and *constants*:

> *Generic α-level language:* An α-level language such that (a) each of its variables has level $\beta < \alpha$, (b) for all $\beta < \alpha$ it has countably many variables of level β, and (c) each of its constants has level $\gamma \leq \alpha$.
>
> *Full α-level language:* A generic α-level language, such that for all $\gamma \leq \alpha$ it has at least one constant of level γ.
>
> *Cofinally Full α-level language:* A generic α-level language such that for each $\gamma < \alpha$ it has at least one constant whose level exceeds γ.

Having explained these classifications of level languages, we now turn to the legitimacy of HLPL. Once again, the legitimacy of HLPL will not have to do with actual language usage, but with our ability to provide a good account of its intended interpretation. The different kinds of level languages were introduced mainly to investigate whether there is any significant difference in the way in which we legitimize them.

Afterwards, we will focus on how Linnebo and Rayo try to reach a similar conclusion through a different route, which will ultimately be much more problematic than initially thought to be. One

[8] Rayo calls Complete languages "full".

of the problems will come from the fact that their arguments for language legitimacy do not work uniformly on all the kinds of level languages we have considered.

4.3 HLPL Legitimacy

In the second chapter, we focused on the ω-level plural language containing plurals of all finite levels. In this case, we extend the language to include terms of transfinite levels. This can be done either by offering a basic or a complete language, depending on whether predication of the top level is allowed (assuming one exists).

In effect, apart from extending the syntax, the formation rules will remain the same. The previous axioms are simply adjusted to talk of any ordinal α, instead of some natural number. For instance, the higher-level plural comprehension axiom scheme will be:

> For any ordinal α, $\exists x^\alpha \phi(x^\alpha) \rightarrow \exists x^{\alpha+1} \forall x^\alpha (x^\alpha \prec x^{\alpha+1} \leftrightarrow \phi(x^\alpha))$, where ϕ is a formula containing x^α and possibly other variables free, but no occurence of $x^{\alpha+1}$.

The next step is the addition of some rules of inference. For all levels $\beta \geq \alpha$, we have the two inference rules provided that all formulae are well-formed and b^β does not occur in any undischarged assumption on which $\phi(b^\beta)$ depends:[9]

$$\forall E^\beta_\alpha \ \frac{\forall x^\beta \phi(x^\beta)}{\phi(a^\alpha)} \qquad\qquad \forall I^\beta_\alpha \ \frac{\phi(b^\beta)}{\forall x^\alpha \phi(x^\alpha)}$$

Additionally, for each limit level λ we have the infinitary rule, which captures how limit types conform to whatever was happening before them; that "nothing essentially 'new' happens at limit types" (Button and Trueman, 2022, p. 3):

$$\text{Limit}^\lambda \ \frac{\forall x^\alpha \phi(x^\alpha), \text{ for all } \alpha < \lambda}{\forall x^\lambda \phi(x^\lambda)}$$

Notice that these rules are justified based on the conception of reference we have associated with higher-level plurals. For instance, in the Limit$^\omega$ rule, we want to say that if something holds for all pluralities of finite levels, then it will hold of all ω-level pluralities. Remember that ω-level pluralities are those that contain pluralities of all finite levels and refer to individuals in virtue of the contained plurals and how *these* plurals refer to individuals. Therefore, if we know that for all $n \in \mathbb{N}$ we have that $\phi(x^n)$ for all the n-level plurals, then we can infer that $\phi(x^\omega)$, precisely because any ω-level plural refers to individuals in virtue of the finite level plurals it contains.

What will it take now for an α-level plural language to be considered legitimate? We would need to provide a semantics for it that explains how we are supposed to be interpreting the non-logical parts of the language in a way that respects the intended interpretation of HLP. Succeeding in doing so would result in the *thin legitimacy* of the α-level language.

A way to do so would be by extending the semantics given by Oliver and Smiley (2016). The basic idea behind their interpretation of plural logic invokes plural resources in the meta-language. Following suit, we want to invoke HLP resources in the meta-language, which, as I mentioned in

[9] This part on the additional rules of inference draws heavily from the presentation of Cumulative Type Theory in Button and Trueman (2022).

Chapter 2, can be interpreted according to different conceptions of HLP reference; I, however, interpret them using *Combinatorial Reference* and the **Plural Axiom of Ordinals**.

Having accepted simple (first-level) plural reference allows for a way to extend the notion of a *function* from something that outputs (at most) a unique value to something that can out multiple values simultaneously. Since we can refer to multiple individuals plurally, we can accept functions that have multiple outputs.[10]

With a background of the **Plural Axiom of Ordinals**, we want to allow functions that are α-level multivalued.[11] Let us first focus on second-level multivalued functions, which would serve the purpose of assigning a superplural term its semantic value, to present the main points of the argument.

Second-level multivalued functions are supposed to add a level of complexity to the way in which they map a term to some individuals, just like the way that superplural terms refer to individuals. In my opinion, second-level multivalued functions are supposed to generalize the way in which we already understand function composition. In ordinary calculus, having two functions $f : A \to B$ and $g : B \to C$ allows one to construct their composition $g \circ f : A \to C$.

In a similar way that the information that f carries is utilized by g, so we will think of a multivalued function - that depicts plural reference - as contributing to a second-level multivalued function - depicting superplural reference - precisely mimicking our conception of *Combinatorial Reference*. It is important to note here that this analogy should be taken with a grain of salt. I am not trying to say that second-level multivalued functions are in fact compositions of two multivalued functions, but that there is a similarity in the way that these functions give us their outputs given an (some) input(s).

Thus, if we wanted to map a superplural constant aaa to its semantic value, we would need a function f that not only maps aaa to some individuals, but also takes into account the fact that the individuals comprising aaa have been first combined into the pluralities (and individuals) contained in aaa. Then, f would not output some of the a's directly, but the a's grouped up appropriately. A manner that becomes clear when one has mastered the workings of *Combinatorial Reference*. This function will be denoted by val henceforth, and maps terms of a given level to the individuals they refer to by taking into account how they have been grouped.[12] As mentioned in Chapter 2, this does not depend on *Combinatorial Reference* specifically, but it can be interpreted similarly using different notions of HLP reference. This is precisely the reason why there is no direct comparison between these semantics and those preferred by Grimau (2018, Chapter 6).

The next step is to iterate this move about functions, which provides us with the conception of an α-level multivalued function, which then can be used for the semantics of an α-level plural language. And once again, the notion of *containment* plays a pivotal role in understanding this procedure, and our previously outlined conception of it, as outlined in Section 2.1.5, is what supplies us with a conception of HLP as an extension of ideology without incurring any ontological commitments. The heavy lifting for this iteration to go through for all ordinals is obviously done by the **Plural Axiom of Ordinals**, since it allows us to have a conception of α**-level CR** for all α, which in turn allows us to grasp how α-level multivalued functions work.

In that way, we are simply able to provide a semantics and an intended interpretation for α-level

[10]For a more detailed version of this argument, see (Oliver and Smiley, 2016, Ch. 9).

[11]Simple multivalued functions are multivalued ones of level 1.

[12]A more technically comprehensive presentation of these semantics can be found in (Oliver and Smiley, 2016, Chapter 15). For my purposes, the way to understand them according to CR is at issue and not their formal intricacies.

plural languages which basically talk about some individuals and all the ways in which one can refer to them according to α**-level CR**. Therefore, we have directly acquired the *thin legitimacy* of such languages.

The thin legitimacy of at least some kinds of α-level languages has been achieved; the question is, based on the two classifications provided in the previous section, do we get legitimacy for all of them with a swift move, or should we be more careful when it comes to these distinctions?

I think that we can rest assured that the thin legitimacy of *generic, full,* or *cofinally full* α-level languages has been attained. This is a three-way distinction regarding the kinds of constants the language involves, which, in our case, are plural terms of different levels. Since a grasp of the **Plural Axiom of Ordinals** gives us the relevant conception of an α-level multivalued function for all α, we can provide a suitable interpretation function for any α-level plural term for all α. Hence, we can consider *full* α-level languages as legitimate, and since that is a stronger requirement than demanding the same out of *generic* or *cofinally full* level languages, we are licensed to deem them all legitimate.

The point which needs to be handled with care is whether a different approach should be adopted with *basic* or *complete* languages. The reason is that in this case we are not only interested in the semantic values of HLP terms but also of predicates taking HLP terms as arguments. Firstly, let us note that our predicates are not typed, in the sense that the argument positions can be filled by HLP terms of any level. One of the reasons is to mimic the fact that natural languages contain predicates which may take HLP terms of different levels as arguments. Consider "John and Mary played against each other", which is about a first-level plurality, and "My children and your children played against each other", which is about a second-level plurality.

The semantic value of a predicate F would be given through an α-level multivalued function val of the appropriate level, considering F a HLP property of appropriate levels. Accordingly, for a predicate F applying to terms $a_1, \ldots, a_n$, where the a_i can be singular or HLP terms of any level, val satisfies $F(a_1, \ldots, a_n)$ iff $valF$ holds of $val\ a_1, \ldots, val\ a_n$. What we must understand now is whether the **Plural Axiom of Ordinals**, and the subsequent conception of α-level multivalued functions that it licenses, are enough to grasp what HLP predication is supposed to be about. In other words, are these tools enough to thinly legitimize the *complete* α-level languages for all levels or not?

The question can also be framed as follows: Is a grasp of HLP reference enough to grasp the notion of HLP properties and relations?[13] This is a matter that requires care because of the competing analyses of HLP properties. The two main contenders are the following: we either (a) make a type distinction between pluralities (individuals included) of different levels and properties, or (b) HLP properties that hold between some pluralities of some levels will be a plurality of a level that is the *maximum* of the levels of the other pluralities of which the property holds.[14]

I do not aim to settle this debate at this point, but I am sympathetic towards the former view, especially in light of the arguments presented in Florio (2014b). This sympathy is also reflected on the choice of semantics that I have presented above, since plural predicates denote plural properties through val. I will proceed by assuming this type distinction, but before doing so I want to outline what choosing the latter option would amount to.

[13]Henceforth, I will be talking of *HLP properties* for ease of exposition.

[14]There is a point to be made here about a choice regarding the formal framework. I have allowed for predicates to take finitely many arguments, but it is also *prima facie* plausible to allow for more arguments, even transfinitely many. For now, I restrict the discussion to predicates with finitely many arguments; the points I make can almost directly be translated to the liberalized version of the predicates by switching *maximum* for *supremum*.

If HLP properties are to be understood simply as pluralities of some level, then it is clear that we do not need much more than the **Plural Axiom of Ordinals** to deem *basic* and *complete* level languages as legitimate. To make this clearer, consider a predicate F and the plural terms aa^{k_i} for some finite $i \in \mathbb{N}$; then intuitively $F(aa^{k_1}, \ldots, aa^{k_i})$ is true iff for all $l \leq i$ it is the case that aa^{k_l} is *contained* in the m-level plurality containing all the pluralities that satisfy F, where $m = max\{k_1, \ldots, k_i\}$. Therefore, since we understand *containment* and have a grasp of how plurals of all levels refer, we can understand the semantic structure of plural predicates.

Matters become trickier when a type distinction between pluralities and properties is introduced. Trickier in the sense that pluralities and properties do not lend themselves to a uniform treatment as before. Fortunately, this does not give rise to any insurmountable obstacles. The reason is that I take it that we already possess an understanding of what it means for properties to hold of individuals referred to plurally or superplurally. Such an indication, although not a definitive one, is our ability to use and evaluate the relevant predications involving superplural terms in natural language. Generalizing those intuitions to pluralities of higher levels should not be too hard.

Additionally, much like our conception of singular properties, i.e., properties that hold of unique individuals, we already have a general grasp of plural properties, i.e., properties that hold of many individuals collectively. HLP properties are a species of plural properties, in the sense that they hold of multiple individuals, that are referred to in the appropriate way. When predicates apply to singular terms, they aim to say something of an object; plural predicates, in much the same way, aim to say something of objects. HLP predicates purport to say something of objects (individuals) referred to in a special way.

Since we have a grasp of that kind of reference, and the conception of higher-level multivalued functions is supposed to capture that, it seems that the semantics of predicates given above can capture adequately the intended interpretation of HLP predicates, namely HLP properties. Once again, then, it seems that there is not too big of a difference between considering *basic* or *complete* HLP languages as legitimate.

To summarize then we saw that the legitimacy of HLP languages of different levels is a matter that rests on our grasp of *containment*, the **Plural Axiom of Ordinals**, and (higher-level) plural properties. Alas, we have not provided a way towards *thick* legitimacy, as such a goal would be one that is too ambitious, but I believe that we have achieved an important feat: HLP languages of all levels are *thinly* legitimate.

4.4 Meta-theoretic Demands

In what we previously showed, the way to the thin legitimacy of all α-level plural languages goes through the acceptance of the **Plural Axiom of Ordinals**. The axiom, whose reasonableness and justification we already saw in Section 4.1, is the reason why we can allude to α-level multivalued functions as the valuation functions mapping the non-logical constants of our language to their intended interpretation. Therefore, just as the plural hierarchy encompasses a level for each ordinal, so do we have the legitimacy of all α-level plural languages. Looking at our previous argumentation more closely shows that $(\alpha + 1)$**-level CR** is enough for the thin legitimacy of all β-level languages for all $\beta \leq \alpha$. That is precisely because β-level plural terms require β-level multivalued functions to map them to their intended interpretations, which can be grasped insofar as (at least) $\beta + 1$**-level CR** is available to us.

Remember that an acceptance of the **Plural Axiom of Ordinals** rests (among others) on the acceptance of a meta-theory that has access to the whole class of ordinals and can carry out transfinite recursion. Although this is not a problem *per se*, it severely limits the philosophical applicability of HLPL. If making sense of the intended interpretation of the logical language, requires a meta-theory - i.e., set theory - strong enough to already model most of the things that HLPL would be applicable to, then one would wonder where the upshot of HLPL is.

Alluding once again to the conceptual and formal distinction mentioned in the Introduction, there may be conceptual reasons to prefer HLPL over the alternatives. For instance, the more linguistic applications may prefer HLPL, because it captures HLP terms at face-value, something which other approaches to HLP cannot do. However, more fundamental and metaphysical applications interested in the ontological commitments of the frameworks invoked would not be too eager to use HLPL.

Consider Hellman's allusion to plural logic due to its apparent ontological innocence in Hellman (1989, 1996). For his aims, HLPL would seem like a good candidate as it would extend the expressive resources of plural logic without incurring any additional ontological commitments. However, it would be argued that these ontological commitments have been simply shifted to the meta-level, because to make sense of the framework, we already need an over-inflated ontology that includes all the ordinals. If, on the other hand, we can dispense with the deep set-theoretic demands, we would widen the range of the possible philosophical applications of HLPL.

In order to achieve that, we will need to find an alternative to ordinals that will allow us to both keep the existing formulation of *Combinatorial Reference* and posit languages of higher levels. In other words, we need some entities that can function as ordinals, namely that are well-ordered by some relation so that: (a) we can make sense of comparing the levels of the languages and what it means to countenance languages of higher and higher levels, and (b) so that we can employ the principle of induction for well-orders which was essential in the presentation of *Combinatorial Reference*.

The strategy that will allow us to achieve that is inspired by Gödel (1933).[15] Let us present the strategy in some detail. Suppose that one accepts the existence of a countable ordinal $\alpha \geq \omega$ without any qualms, accepts $(\alpha + 1)$**-level CR**[16] and is ready to consider any α-level language as thinly legitimate. Then, suppose that with those resources we are able to provide HLP counterparts to the well-ordered set α within an α-level language; then we would have a way to talk of the ordinal α as an HLP entity.[17] Furthermore, if we can prove within an α-level language the existence of an HLP term functioning as an "ordinal" β greater than α, then we will be able to use "ordinals" less or equal to β, and the fact that they are well-ordered to grasp and accept $(\beta + 1)$**-level CR**. With that in place, by arguments similar to the ones presented above, we can consider β-level languages as thinly legitimate. We would then proceed to iterate this process for as long as we can in order to get our hierarchy of plurals and higher-level languages.

In a nutshell, if one is ready to accept a big enough countable ordinal, one can prove the existence of well-orders within one's theory, which can then be used to serve as the types for the entities invoked. It is now pertinent to show that this can, in principle, be carried out before talking about the least

[15] Gödel's remarks have been discussed at length by Tait (2001), Koellner (2003, Chapter 1.2), and Incurvati (2020, p. 78-81).

[16] Note that here $\alpha + 1$ is just a matter of notation, which should not be taken to necessarily commit one to the existence of the ordinal $\alpha + 1$. The reason is that $(\alpha + 1)$**-level CR** only mentions ordinals up until α.

[17] In what follows, I will be talking of set-theoretic ordinals and their HLP counterparts by only using the word "ordinal". Whenever the distinction is not clear from the context, I will qualify the term accordingly.

possible initial ordinal that can kick this process off, and finally about the number of iterations that one should carry out to get as many higher-level languages as possible.

Let us start by talking about the existence of large ordinals provable in an $(\omega+\omega+3)$-level language. The argument presented here is similar to that presented by Button and Trueman (2022), which also draws on results of Linnebo and Rayo (2012) and Degen and Johannsen (2000). The arguments of Linnebo and Rayo (2012) will be presented in detail in the next section; the only thing to note now is that an α-level plural language corresponds to a cumulative type theory with typed predicates for every ordinal $\beta \leq \alpha$. We will allude to the fact that certain versions of set theory are *interpretable* in appropriate HLP languages. The kind of set theory we need is the first-order theory Zr, which includes all the Zermelo axioms for set theory plus the axiom telling us that all the sets are arranged in well-ordered ranks.[18]

Button and Trueman (2022) prove the following theorem, which, in our terminology, states that Zr is interpretable in an appropriate α-level plural language T^α:

Theorem 16. $T^\alpha \vdash Zr^{(\kappa)}$*, where* $\kappa > \omega$ *a limit ordinal and* $\kappa + 2 < \alpha$*.*

$Zr^{(\kappa)}$ results from a translation that takes the variables of Zr and restricts them to the rank κ. Suppose now that one accepts the existence of the ordinal $\omega + \omega + 3$, the corresponding form of *Combinatorial Reference*, and the thin legitimacy of the corresponding HLP language $T^{\omega+\omega+3}$. Then, one can interpret $Zr^{(\omega+\omega)}$, wherein one can prove Hartog's Theorem that posits the existence of an uncountable, well-ordered set A.[19] By using that well-order A, we can get a grasp of how all higher-level plurals until level A would refer according to *Combinatorial Reference*. As a result, we can then get the thin legitimacy of A-level plural languages, where we can carry out this process again to ascend to higher and higher levels.

To summarize this process: we start with a plural language of an appropriate level, which interprets an appropriate set theory. Then, within that set theory, we can prove the existence of an ordinal that is bigger than the initial level with which we started. The very same thing, due to interpretability, is provable within the plural language. Therefore, within the plural language, we can prove the existence of larger ordinals than the ones we posited in the beginning.

From how we described *Combinatorial Reference* in Chapter 2 and the above sections of this Chapter, all we need is the existence of longer and longer well-orders along which we can iterate the containment relation, and thus grasp the corresponding versions of *Combinatorial Reference*. However, Button and Trueman (2022) mention that there may be a problem with the fact that we need to already accept the existence of the ordinal $\omega + \omega + 3$ for that process to go through. Maybe one has second thoughts about the existence of infinitely many objects.

In our case, I do not take this as a serious threat. Pluralists do not generally have a problem with the existence of infinities. The existence of a countable ordinal as low as $\omega+\omega+3$ should not be considered too steep a cost for the hierarchy of increasingly more expressive languages it affords us. Even one who only accepts arithmetic and the existence of the naturals could help oneself to larger countable ordinals by alluding to work in predicative foundations for mathematics and use a countable ordinal smaller than the Feferman-Schütte ordinal Γ_0, which is significantly larger than $\omega + \omega + 3$.[20] If one

[18]See (Button and Trueman, 2022, p. 32-33) and the references therein for a relevant exposition regarding that theory.

[19]For details, see (Potter, 2004, p. 185) and (Incurvati, 2020, p. 92). It is also important to note that this requires the notion of ordinal defined using the Scott-Tarski trick as explained in (Potter, 2004, Chapter 11).

[20]For more on Γ_0, see Schütte (1965) and Feferman (1964). A general overview of predicativity and the role of Γ_0 can be found in Feferman (2005) and references therein.

believes that ascending a hierarchy of ever more expressive languages affords a legitimate way to extend one's expressive resources, the predicativist strategy can give us large countable ordinals by only accepting arithmetic. Then, with these countable ordinals and the corresponding HLP languages, we can prove the existence of uncountable ordinals, and keep iterating this process to reach more expressive heights.

Alternatively, for the ones with qualms about the actual existence of infinitudes, one could allude to claims of *possibility* to allow for the process to be carried out. For instance, claims of possibility (logical or mathematical) akin to those invoked by Chihara (1990, 2003), Hellman (1989) or Parsons (2007) to talk of infinitudes without committing to their actual existence. Obviously, such a route requires a convincing argument about the connection between modal claims about infinity and accepting the usage of transfinite ordinals. I do not aim to offer such an argument right now, but I consider it a plausible avenue to pursue for those who want to invoke the tools of HLPL and have reservations about certain ontological claims.

The final detail to iron out has to do with the number of iterations that one is able to carry out when applying this strategy. Button and Trueman (2022) put the problem as follows:

> [Suppose that if one] has accepted the existence of an ω-sequence of theories [$T_{\alpha_1}, \ldots, T_{\alpha_n}, \ldots$], then [one] can bootstrap [one's] way to their limit, [T_{α_ω}]. This suggestion is spurious. If some [T_{α_i}] is sufficient to introduce an entity with order-type [α_ω], then we can simply take [T_{α_ω}] as the $i+1^{th}$ theory. The important case is when none of the theories [T_{α_i}] suffices to introduce anything with order-type [α_ω]. But in this case, [one] will worry whether 'taking the limit' is ontologically innocent; for, by assumption, [one] has not found any ontologically innocent theory which supplies [α_ω]. (Button and Trueman, 2022, p. 11)

Button and Trueman have good reasons to express their uncertainty, but it is important to be careful about the relevant details to make sense of exactly what's happening here. In their presentation, they assume that the limit-theory (limit-level language) of the ω-sequence $T_{\alpha_1}, \ldots, T_{\alpha_n}, \ldots$ is the theory T_{α_ω} which requires an entity of order-type α_ω to be typed accordingly. However, this is not entirely accurate. The limit theory $T = \bigcup_{n\in\omega} T_{\alpha_n}$ of the sequence should be the limit of the sequence of theories (defined as their union) and not necessarily the theory including the entities of level $\alpha_\omega = sup\{a_n \mid n \in \omega\}$.

The theory T is the one that pools together all the previously accepted resources. The relevant question to ask now is whether that theory is indeed different from T_{α_ω}, and *if* it is, whether it is powerful enough to prove the existence of an ordinal greater than all the α_n. If the two theories are not different, then T is perfectly acceptable because it only uses previously accepted resources.[21] This goes to show that there is nothing inherently problematic with taking limit theories; the problem is whether they will be able to deliver the goods.

Let us summarize before turning to the alternative approaches to establish the existence of open-ended hierarchies of higher-level languages. We started this section with the aim to provide an independent account of how one could understand *Combinatorial Reference* and the thin legitimacy of higher-level languages without having to rest on a significant amount of set theory. We outlined Gödel's technique applied to higher-level languages as it was presented by Button and Trueman (2022) and saw that there are good reasons to believe that pluralists can apply it to provide an autonomous

[21]This line of argument is reminiscent of the Principle of Union of Linnebo and Rayo (2012). Studd (2021) has criticized that principle, as we will see shortly, but for reasons which do not impact my argumentation here.

story in favor of HLPL. For the potential applications of HLPL to other philosophical areas, people will have to supply this story with additional reasons as to why this approach suffices for their purposes, especially when one of their aims is ontological innocence. However, it is a good first step to have shown that there are ways in which one can navigate the higher-level plural landscape that are independent of set theory.

4.5 Linnebo and Rayo

The achievement of the previous sections should not be underestimated. To illustrate its significance it is important to counterpose it with a similar attempt, that of Linnebo and Rayo in their Linnebo and Rayo (2012), to transform simple type theory to a cumulative one with transfinite levels. Higher-level plural logic, since the influential paper Rayo (2006), has been considered as one of the possible interpretations of simple type theory (STT), and a cumulative version of the former - such as that of the previous chapter and sections - as an interpretation of cumulative type theory (CTT).

That means that if Linnebo and Rayo's argumentation proves successful, then the legitimacy of HLPL would be an immediate consequence. However, as we will see, their arguments have come under scrutiny in Studd (2021) and Button and Trueman (2022). We will first present their argumentation to understand the main critique voiced by Studd. In the next section, I will talk a bit more about Linnebo and Rayo's three assumptions to show that they are not as self-evident as they make them to be, and that my arguments are somewhat independent of whether they hold or not. Then, we will turn to the critique voiced in Button and Trueman (2022) and discuss its impact on this work.

To kick this discussion off, we first need to talk about simple type theory (STT) and cumulative type theory (CTT). STT is in essence higher-order logic of order ω; this means that we extend standard first-order logic with identity with predicates (variables and/or constants) of all finite levels, which take as arguments predicates of the immediately previous level, and quantifiers of all finite orders binding predicate variables of the corresponding level. CTT, as we will see it here, relaxes the formation rules of STT by allowing predicates to take as arguments entities of any previous level (cumulativity) and by allowing transfinite levels.

The intimate connections between STT and the non-cumulative version of HLPL of Section 2.1.6, and between CTT and HLPL as it has been discussed in this chapter, should be fairly clear even now. The former connection was established by Rayo (2006), while the latter has been mentioned by Linnebo and Rayo (2012) and Button and Trueman (2022) among others. The connection is that predicates of different levels can be viewed as pluralities of the same levels.

However, that is not the only interpretation; one could consider a Fregean interpretation which takes predicates to stand for *concepts*. Linnebo and Rayo's arguments, alongside Studd's critique thereof, do not draw on the interpretation of STT and CTT; for that reason, in what follows, we will be neutral with respect to how the frameworks are interpreted.

Their arguments, although they would also suffice for the aim of this thesis to countenance HLP languages of all orders α, are more ambitious than that. Taking inspiration from Gödel, they want to show that there is some intimate connection between CTT and set theory, as CTT results from STT if some "superfluous restrictions are lifted". However, they want to motivate philosophically the lifting of the relevant restrictions.

Their argument rests on three theses (assumptions), which purport to establish that whoever accepts them and is ready to use first-order quantifiers in her theorizing, *should* also be ready to countenance full CTT (CTT with a level for every ordinal).[22] The assumptions are:

Absolute Generality: One's first-order quantifiers can meaningfully be taken to range over absolutely all objects.

Semantic Optimism: Given an arbitrary language, it should be possible to articulate a generalized semantic theory for that language.

Principle of Union: For λ a limit ordinal, suppose that one is prepared to countenance languages of order β for every $\beta < \lambda$. Then one should also countenance languages of order λ (i.e., languages containing variables of type β, for every $\beta < \lambda$), on the grounds that they would be made up entirely of vocabulary that had been previously deemed legitimate.

Let us take a closer look. **Absolute Generality** is a thesis advocated for by numerous authors, most notably by Williamson (2003). The idea is that in numerous areas, it seems that the possibility of absolutely general quantification is necessary for the expressibility of the main tenets of particular doctrines. In metaphysics, for instance, a realist claiming that everything is concrete would have to be able to generalize about absolutely everything for her claim to capture exactly what she purports to express. Contrary theses to **Absolute Generality** have been advanced by Studd (2019).[23]

Semantic Optimism is, in effect, a statement explaining that we should be able to provide a theory for all the different interpretations a language could be given. **Semantic Optimism** enjoys a certain kind of intuitive appeal; if we accept the usage of a language in our theorizing, we would be willing to carry out investigations regarding its semantics. If we fall short of an understanding of the different kinds of ways in which the language can be given its semantics, we will have only achieved a partial insight into that language itself. **Semantic Optimism** will be discussed in the next subsection.

Finally, the **Principle of Union** also has an intuitively true ring to it. Our understanding of a limit ordinal is that it is simply the union (or supremum) of all the ordinals that come before it. Then, if we have β-order languages for all $\beta < \alpha$, where α is a limit ordinal, an α-level language that pools together all the resources of the previous languages should be able to be countenanced as well.

With these assumptions in place, the argument purporting to establish the legitimacy of cumulative type theory is as follows. If one is willing to use first-order languages in one's theorizing, then, by **Semantic Optimism**, it should be possible to provide a generalized semantic theory for that first-order language. Due to **Absolute Generality**, this cannot be carried out in another first-order language; it can be done, however, in a second-order language. By induction, one is led to accepting languages of any finite order, and thus by the **Principle of Union** a language of ω-order.

In general, if α is a successor ordinal, then one can provide a generalized semantic theory for an α-order language in a language of order $\alpha + 1$, and if α is a limit ordinal in a language of order $\alpha + 2$. The three assumptions, along with an application of transfinite induction, allow one to conclude that they can countenance α-order languages for any ordinal α.

[22]In the discussion pertaining to Linnebo and Rayo's paper, we will be talking of the *order* of a language instead of its *level*. Keep in mind that in the context of HLPL, first-order logic is a 0-level plural language, so technically a first-level language would correspond to a second-order one.

[23]Several different positions on **Absolute Generality** can be found in Rayo and Uzquiano (2006).

Note here that we have been intentionally vague about the kind of languages that one is supposed to be countenancing. Let us make that a bit more precise in order to go through Studd's critique of the **Principle of Union**. In his critique, he assumes **Absolute Generality** and **Semantic Optimism**, in order to focus on the inadequacies of the **Principle of Union**.

Based on the distinctions mentioned above, Linnebo and Rayo are talking about *complete α-level languages*, the reason being that they believe that disallowing the existence of predicates taking variables of level α as arguments is an artificial restriction. Therefore, what matters now is whether we are talking of *generic*, *full*, or *cofinally full* languages of level α.

Studd's Critique:

Semantic Optimism can be split into two theses, a positive and a negative one, the second of which relies on **Absolute Generality**. The positive thesis tells us that it is possible to give a generalized semantic theory for an α-level language in a language of level $\alpha + 1$, or of level $\alpha + 2$ when α is a limit ordinal. This fact is independent of the *kind* of level language that we are dealing with.

The negative thesis posits that it is impossible to give a generalized semantic theory for an α-level language in another α-level language. Contrarily, the kind of level language we are dealing with here matters. The proof of the result that gives rise to the negative thesis relies on the fact that the language is a *full* level language, as the proof requires the use of a constant of level α.

Providing then a generalized semantic theory for this α-level language requires generalizing over interpretations. If this is to take place in a meta-language also of level α, then the variable ranging over the interpretations would need to be of level $\beta < \alpha$. The proof then proceeds by an analogue of Cantor's Theorem, where one proves that there are strictly more entities of type α than of type β, for $\beta < \alpha$.

The necessity of a constant of type α implies immediately that the same proof does not go through for languages of *generic* level. Additionally, the **Principle of Union** pulls towards languages of generic level, as "pooling together the resources" of all the β-level languages for all $\beta < \alpha$ would not result in a *full* α-level language as no previous language contained an α-level constant.

Whereas this shows that the argumentation of Linnebo and Rayo is not as ironclad as they initially thought it to be, there are ways to amend it. Studd provides such a way by considering languages of *cofinally full* level (Studd, 2021, p. 279). The more important point, which is equally important for our purposes, lies elsewhere.

Studd introduces the Lewisian distinction between thin and thick legitimacy as we did previously. He argues that **Semantic Optimism** demands thick legitimacy, as providing generalized semantic theories requires meta-languages that we can use and understand. However, the **Principle of Union** is not compatible with thick legitimacy. Studd gives an example of a sequence of languages $\mathcal{B}_n$ such that $\mathcal{B}_n$ is a thickly legitimate language of level n. The union language $\mathcal{B}_\omega$, he describes then, is one containing an infinite lexicon; languages we are able to comprehend and use are compositional languages with a finite lexicon, which are nothing like $\mathcal{B}_\omega$.

Another reason, which Studd does not explicitly mention, is that ordinally typed languages of transfinite levels usually employ *infinitary rules of inference*. This would be the case for the ω-level language that pools together the resources of all finite-level languages of higher-order logic. Even if all languages of finite level can be deemed thickly legitimate, the limit level language will not be as there is little reason to believe that humans would be able to use infinitary rules.

Therefore, the **Principle of Union** in its thick reading is clearly untenable. What about the thin reading? If we accept the thin (and even thick) legitimacy of a sequence of n-level languages, does that warrant the thin legitimacy of the ω-level language including all the previous resources? In other words, is the existence of suitable interpretation functions I_n for each n-level language for the existence of a suitable interpretation function I_ω for the ω-level language? Studd answers in the negative. Such an interpretation would be an entity of type $\omega + 1$, which we have not already accepted in our theorizing.

Linnebo and Rayo, and anyone sympathetic to their argumentation find themselves in a disadvantageous position. Accepting the legitimacy of an ω-level language when having access only to resources of type less than ω appears untenable. I believe that my argumentation regarding the legitimacy of α-level plural languages for all ordinals α does not suffer from the same defects, precisely because of my strategy. I have proposed that the height of the hierarchy can be determined by first extending the hierarchy of plurals and *only then* that of the languages. Because of the **Plural Axiom of Ordinals** or the Gödelian strategy, we first acquire a grasp of the plural hierarchy as a whole or until some level, before deeming the corresponding languages as legitimate. That legitimacy relies only on our grasp of the relevant part of the hierarchy and the semantics we have proposed. As a result, at least part of Linnebo and Rayo's goal can be salvaged. Let us take a closer look at the assumptions of their argument to show how my argumentation does not rest on their assumptions.

4.6 A Brief Discussion of the Three Assumptions

The literature on higher-order logic (simply type theory) is intimately connected to that on absolute generality and generalized semantic theories. It does not then strike one as odd that Linnebo and Rayo's argumentation employs results from both areas; the idea is that if someone accepts the two theses of **Absolute Generality** and **Semantic Optimism**, thus adopting second-order languages in their theorizing, they are but a stone's throw away from accepting languages of all finite orders.

After all, it is not merely a coincidence that Rayo employs both of these assumptions to argue that one should countenance higher-level plural languages of every finite order Rayo (2006). His argument is similar to the one outlined above. If one accepts first-order (0-level) languages and believes both in **Absolute Generality** and **Semantic Optimism**, then one can only provide a generalized semantic theory for that language in a *basic* first-level language.

The resources of the plural logic PFO are thus needed to do so. But for reasons that have to do with Tarski's Theorem regarding the undefinability of truth, one can only give a generalized semantic theory for a basic first-level language in a *full* first-level language. But, due to **Absolute Generality** and a plural version of Cantor's Theorem, the generalized semantics for that language require a *basic* second-level language, and so on, and so forth.

This argument is trying to establish something very specific: if you include first-order logic in your theorizing and accept the two theses, you *should* also accept logics of any finite level. The word "should" is important here, because it tries to impose a certain kind of normativity on this conclusion. You would be unjustified in considering that any of these languages is beyond reach for theorizing.

This goal differs from mine in a very important respect. I am not trying to establish a normative claim like this; I wish to convince the pluralist, who accepts additionally enough set theory to carry out transfinite induction or is convinced by the Gödelian technique presented above, that the resources of HLPL are within her reach. If **Absolute Generality** and **Semantic Optimism** happen to also be established beyond doubt, then anyone who accepts first-order logic would *have* to accept languages of all finite orders as well. Although that would suffice for my goal, I would certainly opt for an alternative route for the same destination that avoids these controversial assumptions.

Note also that Linnebo and Rayo's arguments, if successful, will be more consequential than mine. Whereas I am to convince the *pluralist*, their arguments also put pressure on anyone who simply accepts first-order logic.

It is important to highlight that I am not trying to argue *against* **Absolute Generality** or **Semantic Optimism**. I realize that there is little consensus on these matters, so I believe that doing without them would increase the scope of my argument.

For instance, multiple authors have contested the plausibility of absolutely general quantification.[24] Some, such as Grimau (2018), state **Absolute Generality** as a motivation behind pluralism. I still believe that, if we can, we should not found our argument on it. Personally, although **Absolute Generality** may be used as an argument for pluralism, I take it that it is not one of the main motivations for it. A subtler motivation would be our apparent ability to denote things that cannot be consistently taken to be the members of a set. I take it that the ability to quantify over absolutely everything and that of denoting non-set-sized totalities should be discerned, for the latter is a more immediate motivation for pluralism than the former. Whether one can consistently hold the latter without holding the former is also not something I aim to defend, but I believe that, nonetheless, the distinction is *prima facie* plausible.

Whereas the literature on **Absolute Generality** is vast, the same cannot be said of **Semantic Optimism**, so let us briefly turn our focus to it.[25]

Semantic Optimism, as we have already touched upon, constitutes an intuitively possible thesis, if we are to understand all the ways in which one could assign meaning to a language. To enhance that intuition consider the alternative, namely a kind of semantic *pessimism*. What is it about a language that makes it untenable that one provides a generalized semantic theory? Rayo admits that semantic pessimism would amount to a quite radical view:

> For one is forced to countenance the view that a language might have features whose investigation is ruled out by the nature of the language itself. (Rayo, 2006, p. 246)

This would indeed be a striking feature. However, bizarreness can hardly ever be deemed as conclusive evidence for falsity. **Semantic Optimism** is also an important aspect of the Barwise-Cooper approach to natural language semantics, especially those of generalized quantifiers. This constitutes a highly influential and successful approach to semantics, so whoever wants to avail herself of it would need additional ideological resources. Specifically, she would need irreducibly plural and superplural quantification.

One with roughly naturalistic tendencies then need not endorse **Semantic Optimism** all the way, but only to the extent that it is important for the Barwise-Cooper approach to semantics. My approach offers exactly that: the ideological resources required for that approach, but will do so without resting on the presupposition of **Semantic Optimism** as a general requirement of all legitimate languages.

The reason is that, even though **Semantic Optimism** requires a general investigation into the semantics of languages, legitimate languages often hand-in-hand with an intended interpretation. If

[24]See Footnote 23 in this Chapter.

[25]My discussion on **Semantic Optimism** draws heavily on Rayo (2006), Linnebo (2006), mainly Chapter 3.6 of Studd (2019), and Studd (2021).

we can make sense of that interpretation, then I do not see any pressing reason to accept **Semantic Optimism**, unless we wish to carry out a meta-semantic investigation. Since our aim here is not such, I take it that avoiding endorsing **Semantic Optimism** helps build my argument on secure foundations.

Finally, in light of Studd's critique, it has become clear that relying on the **Principle of Union** to countenance languages of transfinite level will result in severe drawbacks. The argument presented in Section 4.3 is not dependent on the **Principle of Union**, in the sense that we make no normative claim about which languages one should countenance given others that one already accepts. The argument relies on the **Plural Axiom of Ordinals** and our ability to provide the semantics and an intended interpretation for α-level languages based on it. Also, given those assumptions, differentiating among *generic*, *full*, or *cofinally full* becomes indeed superfluous, as their legitimacy is dealt with simultaneously in one fell swoop.

Whereas Linnebo and Rayo try to ascend a hierarchy of CTT-languages by *successively* accepting languages of higher and higher orders, my argument relies first on accepting the **Plural Axiom of Ordinals** and then *simultaneously* getting the thin legitimacy of all α-level plural languages.

With that in mind, it should be noted that I rely on the intelligibility of another union-like principle that does not have to do with the levels of the languages, but with the notion of *containment*. But as was mentioned in Section 4.1, the status of that feature of containment is not as controversial and can be considered at least defeasibly justified. As I noted at the beginning of this chapter, it generalizes on already accepted intuitions about containment and how CR works.

4.7 Pluralities Typed and Untyped

In the final section, I want to focus on a critique by Button and Trueman (2022) that involves the way in which we formalized HLPL, its semantics, and its meta-theoretic demands (at least when using transfinite induction). When setting up the formal details regarding HLP, we chose to proceed in a typed (many-sorted) setting, whereby every plural was accompanied by an ordinal indicating its level. In Chapter 2, I mentioned briefly that there are untyped alternatives to formulating HLPL, such as the one provided in (Oliver and Smiley, 2016, Chapter 15). I brushed over the details of this formal choice and attributed it to the fact that the typed system captures more accurately the syntactic behavior of HLP.

In this section, I want to reflect more on the significance of this choice, especially given the attempts of Button and Trueman (2022) to establish that the type restrictions of CTT in its plural interpretation are *unjustified*. The conclusion they draw from this criticism is that a logic of plurals understood cumulatively should be formulated in a single-sorted, untyped system, rather than one positing transfinitely many levels.

The main idea behind their argumentation is that the semantics for a given language have to justify or motivate the type restrictions on the syntactic front. If the intended interpretation and the semantics invite the introduction of untyped variables, then the type restrictions are considered superfluous and *should* be abandoned in favor of an untyped syntax. I will try to challenge exactly the normativity of the previous claim; in fact, I will try to argue that even if the introduction of untyped variables is in line with the intended interpretation of a typed framework, there may still be reasons for preferring a typed over an untyped framework.

Let us first see their argumentation, specifically as regards pluralities, in detail. First, remember the ϵ operator of Linnebo and Rayo (2012) capturing a certain kind of membership relation.[26] For $\gamma = max(\alpha, \beta) + 1$ we have the *defined* expression:[27]

$$a^\alpha \,\epsilon\, b^\beta \text{ iff } (\exists x^\gamma = b^\beta)\; a^\alpha \prec x^\gamma$$

Suppose that $\alpha \geq \beta$, meaning that the expression $a^\alpha \prec b^\beta$ is not well-formed. Even if that expression is not well-formed, it very well may be the case that there is a plurality of type $\gamma = max(\alpha, \beta) + 1$ that contains a^α and it is equal to b^β. In fact, from the axioms of Level Raising and Higher-level Plural Comprehension, we can show that a plurality $x^\gamma = b^\beta$ always exists; what matters is showing then that in this plurality a^α is contained. Also, precisely because of HLP identity, there is an important sense in which x^γ and b^β can be considered to be *the same*.

It is also not hard to see that in the cases where $a^\alpha \prec b^\beta$ is well-formed, then it is also equivalent to $a^\alpha \epsilon\, b^\beta$. Therefore, ϵ-membership provides us with a way to faithfully mimic $\prec$ when the entities are of the appropriate levels, and also simulate a form of membership intimately connected to $\prec$ even when the level restrictions are not respected.

For these reasons, Button and Trueman argue that we can consider $a^\alpha \prec b^\beta$ as well-formed irrespective of the levels of the pluralities involved. Thus, any model of CTT (in the plural interpretation) will end up making either $a^\alpha \prec b^\beta$ or its negation true independently of the relation between α and β. According to the plural semantics that Button and Trueman mention, we would have for all α and β that:

> $a^\alpha \prec b^\beta$ is true iff what $'\, b^\beta\, '$ refers to includes what $'\, a^\alpha\, '$ refers to.

Precisely because of the way in which the type restrictions for $\prec$ were lifted, it becomes clear that α and β do not play an essential role here. This can then be transformed into the following semantic clause, which dispenses with them and licenses the usage of untyped variables:

> $a \prec b$ is true iff what $'\, b\, '$ refers to includes what $'\, a\, '$ refers to.

The argument then simply runs as follows: CTT's type restrictions are *unstable*, because they invite the usage of untyped variables. Specifically, the plural semantics license the move from typed to untyped variables.[28] One of the conclusions that Button and Trueman draw from this result is the fact that the logic of higher-level plurals should be developed in an untyped framework like that presented in (Oliver and Smiley, 2016, Chapter 15). Precisely because we have opted for a cumulative and typed version of HLPL, we should not take this argumentation lightly.

Firstly, we should start by pointing out that, based on the way we have proposed to understand HLP, the semantic clause above is not accurate enough. As we mentioned, the way in which HLP terms refer differs from simple plurals in terms of *level* or *the way of reference*. So, containment is not understood as a plural referring to something which is included in the referent of another plural. The semantic clause should rather read:

> $a^\alpha \prec b^\beta$ is true iff the way in which $'\, a^\alpha\, '$ refers to some individuals contributes to the way in which $'\, b^\beta\, '$ refers to some (or all) of the same individuals.

[26]See Chapter 2 and the formulation of higher-level plural logic.

[27]The right-hand side of the iff corresponds to Button and Trueman's $(\exists x^\gamma = b^\beta)\; x^\gamma(a^\alpha)$. Also, $(\exists x^\gamma = b^\beta)\phi$ abbreviates $\exists x^\gamma(x^\gamma = b^\beta \wedge \phi)$.

[28]Button and Trueman continue by arguing that the type restrictions of STT, contrary to those of CTT, are justified in *some* semantics, the Fregean semantics, but this would be beside the point here.

This is important because it reflects more accurately our conception of HLP, but it does not immediately impact Button and Trueman's argument. The same clause can be formulated without any sort of mention of the levels invoked:

> $a \prec b$ is true iff the way in which $'a'$ refers to some individuals contributes to the way in which $'b'$ refers to some (or all) of the same individuals.

However, I believe that this clause would strike one as odd if one did not already possess a very thorough understanding of HLP reference. The reason I am not talking specifically about CR here is that I think that this line of argument would be similar under other conceptions of how HLP terms refer. If one chooses to take HLP reference at face-value, one should be clear about the way this is captured on the formal level.

In the clause offered by Button and Trueman, one should understand what it means for the referent of one term to be included in the referent of another. For instance, if a superplural term ddd refers to something, then a plural cc contained in ddd would need to refer to something *included* in the referent of ddd. This way of articulating that mode of reference is consistent with (at least) two competing interpretations.

Firstly, it could be that $'b'$ refers to individuals and $'a'$ refers to some of them, and thus the referents of $'b'$ "include" those of $'a'$. Then, plural and singular terms would refer to individuals directly, something deeply problematic, as we say in Chapter 2. Interpreted in this way, Button and Trueman's clause tells us nothing about the way in which HLP terms refer to individuals.

Secondly, it could be that $'b'$ refers to a superplural that includes (contains) the plural to which $'a'$ refers. This would be the most charitable way to interpret the semantic clause, as it would also be in line with the way that Oliver and Smiley (2016) conceive of the reference of HLP terms. My dissatisfaction with this semantic clause stems from its using the pseudo-singular[29] terms superplural and plural on the meta-level. The usage of "plural" is fairly unproblematic, assuming an already good grasp of how pluralities function semantically.

However, if one's aim is to explain how superplurals and other related terms are to function semantically, then one should not use these terms on the meta-level. The case of plurals is different, precisely because we have a long list of natural language examples that facilitate our grasp of them. If the way in which superplural terms function semantically is what we are trying to explicate, then I am not certain whether alluding to them on the meta-theoretic level is warranted. Apart from that, it runs the risk of creating confusion around the ontological commitments of pluralities. That is because we invoke a pseudo-singular meta-language, of which we have limited understanding, and that could be misinterpreted to mean that, in some way, plurals *qua entities* are included in superplurals *qua entities*. This way of construing the semantics, presupposing again our limited understanding of HLP reference, is dangerously similar to taking HLP quantification as a form of iterated plural quantification.

This does not mean, obviously, that the remarks of Button and Trueman miss their mark completely, but rather that the plausibility of such semantic clauses rests on a prior understanding of HLP reference. But this is exactly what I purported to provide by positing CR as the way to understand how plurals of higher levels refer. The notion of level in understanding CR has been essential and should not be dispensed with. Even if I invoke the amended semantic clause that is in line with CR, namely that:

[29] In the sense of Oliver and Smiley (2016), a pseudo-singular term is a singular term that functions semantically as plural.

> $a \prec b$ is true iff the way in which $'a'$ refers to some individuals contributes to the way in which $'b'$ refers to the some (or all) of the same individuals,

without a prior understanding of CR, talking of "the way" in which something refers is unintelligible.

My line of argument is therefore that there is a role to be played by the typed framework, and that is its ability to work symbiotically with the form of reference it helps explicate. We have the conception of CR that we formalize using HLPL to acquire a more precise grasp of it. Indeed, this is an indispensable first step without which neither the semantics nor the ontology would make sense. The ontological commitments of pluralists are respected, as has been argued in Chapter 2, which is made clear both by the concept of CR and the logic invoked to help us grasp it more thoroughly.

In general, I think that some concepts can be captured adequately in both a typed and an untyped language; it is a pragmatic matter to choose which of the two would be preferable given the specific aims that one may want to accomplish. Given our aim to outline and grasp a form of HLP reference, I take a typed framework to be a more useful starting point.

On the other hand, if we aim at a *thickly legitimate* language, or one that we would like humans to be capable of mastering its use, then an untyped framework, such as that of Oliver and Smiley, would be a better choice than the typed one we chose. The reason: no infinitary rules of inference are required. This does not imply thick legitimacy from the get-go, but at least it does not rule it out as the infinitary inference rules do.

It should also be noted that the stratification of plurals into a hierarchy of different levels is not lost on the untyped formulation of the theory. It is still the case that there is a notion of level that is not imposed from the outside in the form of types or sorts, but it is rather a *definable* notion based on the axioms of the theory. In fact, within that framework, one can show that every plural belongs to a level, for every plural there is an earliest level where it appears, and if it appears on one level, then it also does in any subsequent ones.

Although this is the case, abstracting away from these levels does not allow us to be as precise as we can concerning the way in which plurals refer to individuals. The obvious upshot is that this notion of level does not presuppose a heavy meta-theoretic machinery, precisely because it is definable within the theory. That circumvents an important limitation regarding the applicability of HLPL defined from a meta-theoretic perspective involving a significant amount of set theory. Comparing it to the Gödelian strategy that allows us to provide an infinite hierarchy of languages without requiring the ordinals to be supplied meta-theoretically, this perspective is in important respects more elegant. But even if that is the case, I still wish to maintain that the first step should be a typed formulation of HLPL that helps us understand CR. The question remaining is understanding precisely what the formal connections between the typed and the untyped framework are, and whether one would enjoy a status of formal superiority.[30]

To summarize, I started by explaining Button and Trueman's argument for the superiority of untyped frameworks for dealing with plural logic. I find that argument to be lacking on an important point: although it shows that we *can* introduce untyped variables, it does not establish that we *should*. In fact, in cases where two frameworks, typed and untyped, share the same intended interpretation and are about the same concepts, it very well may be the case that the two emphasize different aspects

[30] For a discussion on some differences between typed and untyped logics, see Florio (2023) and Linnebo (2006).

of these concepts. In this case, I believe that understanding the notion of CR requires the assistance of a typed framework whose syntax reflects how the variables of different sorts are connected.

Although it would require some results regarding the *theoretical equivalence* of the two frameworks to show that the differences lie on the conceptual and not on the formal side of things, I still think that I have supported adequately the claim that two different kinds of formalisms can be important for different reasons. In fact, I think I have argued adequately that, because of those reasons and given our aim of explicating HLP reference, the option for a typed framework over an untyped one is the favorable one.

4.8 Conclusion

We dealt with a multitude of problems within this chapter. We started by trying to understand whether our conception of CR can give rise to a hierarchy of plurals that exceeds the finite levels and tried to establish that it does. We can view CR as the way in which plurals of all transfinite levels refer to individuals by iterating the containment relation along the ordinals; this is what we referred to as the **Plural Axiom of Ordinals**. However, doing so presupposes that one can make use of the class of all ordinals and the fact that transfinite induction holds for them.

Much as the ability to refer to all finite-level plurals according to CR gave us a way to understand the ω-level language of HLPL, we then wanted to explore the possibility of deeming α-level plural languages as legitimate for all ordinals α. To understand these questions better, we first introduced two distinctions regarding the different kinds of level languages one can posit, and a main distinction between the two kinds of language legitimacy that we would focus on drawing on work from Lewis (2010) and Studd (2021).

Subsequently, we argued that the **Plural Axiom of Ordinals** sanctions the *thin legitimacy* of all α-level plural languages, and supplies a way to grasp the intended interpretation of all these different languages. After focusing on this, we saw a way to provide an autonomous path to an increasing hierarchy of plural languages that does not rest on set theory. This was done by a Gödelian technique that allowed us to iterate the containment relation along large well-orders provable from within already acceptable theories, without a need for a supply of ordinals from the meta-level.

Then, we discussed how Linnebo and Rayo (2012) try to establish a similar result concerning an open-ended hierarchy of languages through three main assumptions: **Absolute Generality**, **Semantic Optimism**, and the **Principle of Union**. We looked at critiques mainly from Studd (2021), offered our own on the three assumptions, and found our argumentation to be compatible with, but not resting on them. Specifically, we saw that **Plural Axiom of Ordinals** rests on a certain principle of union, which is different from that of Linnebo and Rayo in ways that lend support to our version but not theirs.

Finally, we delved deeper into the critiques of Button and Trueman (2022), who believe that a logic of higher-level plurals should be carried out within an untyped framework as the type restrictions of CTT are unstable under the plural interpretation. We argued against their claims, maintaining that our typed framework should be a first step in any formalization of HLPL because it lends important insight into the workings of HLP reference, which a typed framework fails to provide.

Chapter 5

Conclusion

5.1 A Summary

In the introduction, we set out two primary goals: (a) grasp higher-level plural reference, and (b) determine the height of the plural hierarchy it gives rise to. The first goal was tackled in the second chapter. I initially introduced the kind of pluralism that I am taking as a starting point for my discussion, before presenting alternative positions on the intelligibility of HLP. We saw three kinds of attitudes: skepticism towards HLP, dispensability, and face-value. After critically discussing them, I turned to my proposal by going through its main aspects, the logic it is associated with, and a conceptual comparison with the alternatives. The main conclusions we drew from that chapter were that *Combinatorial Reference* arises as a promising candidate for HLP reference, that is not equivalent to any of the existing alternatives, does not suffer from any of their main setbacks, and also bears significant conceptual similarities with some of them.

Chapter 3 was then devoted to an in-depth formal comparison of what I have taken to be the most compelling rival to HLPL, namely the generalized cover approach and its associated logic CPL. The pros and cons of each of the intended interpretations were discussed in Chapter 2, so the aim of Chapter 3 was to show that in some important sense of the word, the two logics make the same formal claims about their domain of application.

To be a bit more precise, I first formulated and axiomatized both HLPL and CPL as many-sorted first-order theories, talked about their consistency, and proved that they are *theoretically equivalent* in two very important ways: they are both *Morita Equivalent* and *bi-interpretable*. With this result, I purport to show that the differences between the two alternatives do not lie on the formal, but rather on the conceptual side. Settling the debate between the superiority of one over the other has to be based on philosophical and linguistic considerations, but not formal ones. For the purposes of this thesis, the choice of working with HLPL is warranted, and formally, it does not restrain us in any way that CPL does not.

Finally, Chapter 4 deals with the second primary goal: the height of the hierarchy. There, the question was split up into two different ones. First, we looked into the hierarchy of plural terms that one could, in principle, make sense of based on *Combinatorial Reference*. Second, given *this* hierarchy, if we can actually motivate a similar hierarchy of higher-level plural languages, how high does *that* hierarchy extend? The first question was answered by the **Plural Axiom of Ordinals**, where I argued that the HLP hierarchy extends as high as the ordinals. The second question was answered similarly due to the previous insight. Both questions received alternative answers with a similar conclusion, because of the attempt to make fewer meta-theoretic demands by alluding to the Gödelian technique of Section 4.4. Then, I embarked on comparing this approach to higher-level languages with that of Linnebo and Rayo (2012), while dealing with the relevant critique of Studd (2021). I offered my own critiques to Linnebo and Rayo's work as well, and argued that my proposals' assumptions are orthogonal to those of Linnebo and Rayo, thus not suffering from the same defects. Finally, some attention was given to the remarks of Button and Trueman (2022) who purport to establish through a series of arguments that higher-level plurals should be regimented in an untyped, rather than a typed, logical framework. I found their argumentation not to be as compelling as they take it to be, in the sense that they do not force us to adopt an untyped framework. I concluded by arguing that my choice of a typed logic gives rise to some significant conceptual differences, which, for my goal to unearth the connections between HLP and *Combinatorial Reference*, make this choice the preferred one.

5.2 Future Work

I take it that this thesis makes a compelling case about the intelligibility of higher-level plurals and the logic associated with them; a case substantial enough so that people undertake the task of applying it to several philosophical areas. In Chapter 2, I offered a list of areas of application for plural logic, and I believe that there are the first instances where one could investigate the potential applicability of HLPL. An example could be in the attempt to provide generalized semantic theories, which require resources in addition to plural logic. Studd (2019) notes that in some cases, especially when considering generalized quantifiers, superplural resources (at least those) are required. With a prior understanding of superplural logic, the logico-philosophical investigations in the semantics of formal languages can be carried out without any qualms as to the legitimacy of the formal and conceptual tools used.

More specifically, I would like to provide a more thorough understanding of the formal comparison between the alternative plural frameworks. Firstly, I would extend the interpretability results of Chapter 3 to be between the many-sorted first-order theories HLPL and CPL with a countably infinite assortment of primitive predicates and to the cumulative versions of these theories. Secondly, I would want to see if *stronger* interpretability results, such as *synonymy*, can be attained for these frameworks. Thirdly, an additional, and similar comparison, between HLPL, CPL, and the mereological approach (after extending it to HLP) would be in order, and of great significance to the comparison between the logico-philosophical and the logico-linguistic literature.

Additionally, given the ability to compare logical frameworks that use semantics other than set-theoretic ones, e.g., higher-order semantics, I would want to compare the same frameworks with respect to these semantics. However, the field of theory/logic comparison has not extended its reach to alternative semantic approaches as of yet.

Regarding the work in Chapter 4, I believe that there is work in understanding precisely what kind of language legitimacy is required for the philosophical applicability of a given framework. Attempts in the literature invoke either *thick* legitimacy, for instance by offering a paraphrase in natural language, or a fundamental sui-generis understanding of the framework, such as that argued by Williamson (2003, 2013). I believe that I carve a middle way by the notion of *thin* legitimacy accompanied with a grasp of the intended interpretation of the framework using already acquired concepts. The limits of this approach, I think, are worthy of exploration, as they provide interesting insights into grasping the intended interpretations of formal languages even if those languages cannot be used by humans due to their infinitary character. It is of large philosophical interest to compare different frameworks that share an intended interpretation, are in some sense theoretically equivalent, but one could be used by following its formal rules while the other could not. This is a matter intricately connected to the discussion between the choice of an untyped or a typed framework.

Furthermore, future work would attempt to make the Gödelian strategy even more formally precise. The idea of using higher-order entities to provide well-orders is in itself interesting and could have other applications in the vast area of higher-order metaphysics as well. Specifically, this strategy could provide an autonomous foundation for higher-order metaphysics since it would only employ higher-order resources.

A further application in metaphysics would concern the literature on "Composition as Identity", a thesis proposed by Baxter (1988) and endorsed by Lewis (1990) and Cotnoir and Baxter (2014). The thesis roughly states that the mereological sum of some parts is just those parts. It is not an uncontroversial take on the ontological status of sums, but its tenability rests heavily on the ontological innocence of plural quantification (Lewis, 1990, p. 87). In a mereological setting where atoms can be viewed as individuals, and sums created from them as pluralities, then further sums created from some sums and some atoms can be viewed as superplurals. For "Composition as Identity" to lead to the ontological innocence of mereology, as Lewis argues, the ontological innocence of superplural quantification is needed. *Combinatorial Reference* and its apparent ontological innocence can be utilized in this area so that proponents of "Composition as Identity" can afford mereology the ontological innocence they wish for it.

Finally, I would like to offer an axiomatization for an untyped theory of plurals based on the Level Theory of Button (2021), because it allows for an object-language definition of the notion of level. I believe that such an axiomatization, due to its significant similarities with other set theories that define levels within the theory, such as the theory Z of Potter (2004) or Z^+ of Incurvati (2020), can be a good first-order candidate for this task. This is a personal choice for the axioms of such a theory, but I would like to use them to offer three comparisons: (a) between it and the HLP logic Oliver and Smiley (2016) which is defined using a different predicate logic than first-order logic with identity, (b) between it and the HLP languages we have considered in this thesis, in a way very similar to that from Linnebo and Rayo (2012) and Button and Trueman (2022), and (c) to offer a comparison between it and set theory. I take it that all comparisons will be of philosophical and formal interest.

The third one, especially, I hope to use as a way to argue that an untyped plural logic and a set theory with an explicit notion of rank/level should roughly make the same claims about plurals and sets, respectively. The difference between them should lie in the adoption of *urelements*, and specifically whenever the urelements are class-many, plurals will act differently from sets because there can be no sets of class-size, but we can refer plurally to class-many individuals. In other words, whenever set-sized collections are invoked, sets and pluralities including the elements of those sets should make roughly the same claims *if* one believes, as I do, that set theory provides our best theory of collections.

All in all, there is still a lot of work to be done, and many more avenues to be pursued both for the foundations of plural logic and for its applications.

Bibliography

Barrett, T. W. and H. Halvorson (2016). Morita equivalence. *Review of Symbolic Logic 9*(3), 556–582.

Baxter, D. L. M. (1988). Identity in the loose and popular sense. *Mind 97*(388), 575–582.

Ben-Yami, H. (2013). Higher-level plurals versus articulated reference, and an elaboration of salva veritate. *Dialectica 67*(1), 81–102.

Black, M. (1971). The elusiveness of sets. *Review of Metaphysics 24*(4), 614–636.

Boolos, G. (1984). To be is to be a value of a variable (or to be some values of some variables). *Journal of Philosophy 81*(8), 430–449.

Boolos, G. (1985). Nominalist platonism. *Philosophical Review 94*(3), 327–344.

Burgess, J. P. and G. Rosen (1997). *A Subject with No Object: Strategies for Nominalistic Interpretation of Mathematics*. New York: Oxford University Press.

Button, T. (2021). Level theory, part 1: Axiomatizing the bare idea of a cumulative hierarchy of sets. *Bulletin of Symbolic Logic 27*(4), 436–460.

Button, T. and R. Trueman (2022). Against cumulative type theory. *Review of Symbolic Logic 15*(4), 907–49.

Button, T. and S. P. Walsh (2018). *Philosophy and Model Theory*. Oxford, UK: Oxford University Press.

Champollion, L. (2017). *Parts of a Whole: Distributivity as a Bridge Between Aspect and Measurement*. New York, NY.

Chihara, C. S. (1990). *Constructibility and Mathematical Existence*. New York: Oxford University Press.

Chihara, C. S. (2003). *A Structural Account of Mathematics*. Oxford and New York: Oxford University Press UK.

Cotnoir, A. J. and D. L. M. Baxter (Eds.) (2014). *Composition as Identity*. Oxford: Oxford University Press USA.

Degen, W. and J. Johannsen (2000). Cumulative higher-order logic as a foundation for set theory. *Mathematical Logic Quarterly 46*(2), 147–170.

Feferman, S. (1964). Systems of predicative analysis. *Journal of Symbolic Logic 29*(1), 1–30.

Feferman, S. (2005). Predicativity. In S. Shapiro (Ed.), *Oxford Handbook of Philosophy of Mathematics and Logic*, pp. 590–624. Oxford University Press.

Florio, S. (2010). The semantics of plurals: A defense of singularism.

Florio, S. (2014a). Semantics and the plural conception of reality. *Philosophers' Imprint 14*, 1–20.

Florio, S. (2014b). Untyped pluralism. *Mind 123*(490), 317–337.

Florio, S. (2023). On type distinctions and expressivity. *Proceedings of the Aristotelian Society 123*(2), 150–172.

Florio, S. and O. Linnebo (2021). *The Many and the One: A Philosophical Study of Plural Logic*. Oxford, England: Oxford University Press.

Florio, S. and D. Nicolas (2020). Plurals and mereology. *Journal of Philosophical Logic 50*(3), 415–445.

Florio, S. and S. Shapiro (2014). Set theory, type theory, and absolute generality. *Mind 123*(489), 157–174.

Friedman, H. M. and A. Visser (2014). When bi-interpretability implies synonymy. *Logic group preprint series 320*, 1–19.

Gillon, B. S. (1987). The readings of plural noun phrases in english. *Linguistics and Philosophy 10*(2), 199–219.

Gillon, B. S. (1992). Towards a common semantics for english count and mass nouns. *Linguistics and Philosophy 15*(6), 597–639.

Gödel, K. (1933). The present situation in the foundations of mathematics. *in* Gödel, 1990, 45-53.

Gödel, K. (1990). *Collected Works II*. Oxford: Oxford University Press.

Grimau, B. (2018). *From Plurals to Superplurals: In Defence of Higher-Level Plural Logic*. Ph. D. thesis, University of Glasgow.

Grimau, B. (2019). In defence of higher-level plural logic: Drawing conclusions from natural language. *Synthese 198*(6), 5253–5280.

Halvorson, H. (2019). *The Logic in Philosophy of Science*. Cambridge and New York: Cambridge University Press.

Hellman, G. (1989). *Mathematics Without Numbers: Towards a Modal-Structural Interpretation*. Oxford, England: Oxford University Press.

Hellman, G. (1996). Structuralism without structures. *Philosophia Mathematica 4*(2), 100–123.

Hewitt, S. (manuscript). Frege's theorem in plural logic.

Hodges, W. (1993). *Model Theory*. Encyclopedia of Mathematics and its Applications. Cambridge University Press.

Incurvati, L. (2020). *Conceptions of Set and the Foundations of Mathematics*. Cambridge University Press.

Koellner, P. (2003). *The Search for New Axioms*. Ph. D. thesis, Massachusetts Institute of Technology.

Lewis, D. K. (1990). *Parts of Classes*. Blackwell.

Lewis, D. K. (2010). Languages and language. In D. Byrne and M. Kölbel (Eds.), *Arguing about language*, pp. 3–35. Routledge.

Link, G. (1998). *Algebraic Semantics in Language and Philosophy*. CSLI Publications.

Link, G. (2002). The logical analysis of plurals and mass terms: A lattice-theoretic approach. In P. H. Portner and B. H. Partee (Eds.), *Formal Semantics - the Essential Readings*, pp. 127–147. Blackwell.

Linnebo, Ø. (2003). Plural quantification exposed. *Noûs 37*(1), 71–92.

Linnebo, O. (2006). Sets, properties, and unrestricted quantification. In A. Rayo and G. Uzquiano (Eds.), *Absolute generality*, pp. 149–178. Oxford University Press.

Linnebo, O. and D. Nicolas (2008). Superplurals in english. *Analysis 68*(3), 186–197.

Linnebo, O. and A. Rayo (2012). Hierarchies ontological and ideological. *Mind 121*(482), 269–308.

Linnebo, O. and A. Rayo (2014). Reply to florio and shapiro. *Mind 123*(489), 175–181.

Linnebo, (2022). Plural Quantification. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2022 ed.). Metaphysics Research Lab, Stanford University.

Maddy, P. (1983). Proper classes. *Journal of Symbolic Logic 48*(1), 113–139.

Manzano, M. (1996). *Extensions of First Order Logic*. New York: Cambridge University Press.

Manzano, M. and V. Aranda (2022). Many-Sorted Logic. In E. N. Zalta and U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Winter 2022 ed.). Metaphysics Research Lab, Stanford University.

McEldowney, P. A. (2020). On morita equivalence and interpretability. *Review of Symbolic Logic 13*(2), 388–415.

McKay, T. (2006). *Plural Predication*. New York: Oxford University Press.

Nicolas, D. and J. D. Payton (2025). Collective nouns and the distribution problem. *Synthese 205*(4), 1–29.

Nicolas, D. and J. D. Payton (forthcoming). Superplurals analyzed away. *Inquiry: An Interdisciplinary Journal of Philosophy*.

Oliver, A. and T. Smiley (2016). *Plural Logic: Second Edition, Revised and Enlarged*. Oxford University Press.

Parsons, C. (2007). *Mathematical Thought and its Objects*. New York: Cambridge University Press.

Payton, J. D. (2025). From singular to plural...and beyond? *Philosophy and Phenomenological Research 110*(3), 983–1012.

Potter, M. D. (2004). *Set Theory and its Philosophy: A Critical Introduction*. Oxford, England: Oxford University Press.

Rayo, A. (2006). Beyond plurals. In A. Rayo and G. Uzquiano (Eds.), *Absolute generality*, pp. 220–54. Oxford University Press.

Rayo, A. and G. Uzquiano (1999). Toward a theory of second-order consequence. *Notre Dame Journal of Formal Logic 40*(3), 315–325.

Rayo, A. and G. Uzquiano (Eds.) (2006). *Absolute Generality*. New York: Oxford University Press.

Rayo, A. and T. Williamson (2003). A completeness theorem for unrestricted first- order languages. In J. C. Beall (Ed.), *Liars and Heaps: New Essays on Paradox*. Oxford University Press UK.

Rosen, G. and C. Dorr (2002). Composition as a fiction. In R. Gale (Ed.), *The Blackwell Companion to Metaphysics*, pp. 151–174. Blackwell.

Rumfitt, I. (2005). Plural terms : Another variety of reference? In J. L. Bermu?dez (Ed.), *Thought, reference, and experience: themes from the philosophy of Gareth Evans*, pp. 84–123. Clarendon Press.

Schwarzschild, R. (1996). *Pluralities*. Springer.

Schütte, K. (1965). Predicative well-orderings. In J. Crossley and M. Dummett (Eds.), *Formal Systems and Recursive Functions*, Volume 40 of *Studies in Logic and the Foundations of Mathematics*, pp. 280–303. Elsevier.

Simons, P. (1982). Numbers and manifolds. In B. Smith (Ed.), *Parts and Moments. Studies in Logic and Formal Ontology*, pp. 160–197. Philosophia Verlag.

Simons, P. (2016, 03). *The Ontology and Logic of Higher-Order Multitudes: Logic, Philosophy, and Linguistics*, pp. 55–69. Oxford University Press.

Studd, J. P. (2019). *Everything, More or Less: A Defence of Generality Relativism*. Oxford, England: Oxford University Press.

Studd, J. P. (2021). *Infinite Types and the Principle of Union*, pp. 266–290. Routledge.

Tait, W. W. (2001). Godel's unpublished papers on foundations of mathematics. *Philosophia Mathematica 9*(1), 87–126.

Uzquiano, G. (2004). Plurals and simples. *The Monist 87*(3), 429–451.

Visser, A. (2021). Extension and interpretability. In M. Mojtahedi, S. Rahman, and M. Zarepour (Eds.), *Mathematics, Logic, and their Philosophies: Essays in Honour of Mohammad Ardeshir*, pp. 53–92. Springer.

Wagner, E. (2015, November). Superplural logic. Master's thesis, Universiteit van Amsterdam. Available at `https://eprints.illc.uva.nl/id/eprint/964/1/MoL-2015-23.text.pdf`.

Williamson, T. (2003). Everything. *Philosophical Perspectives 17*(1), 415–465.

Williamson, T. (2013). *Modal Logic as Metaphysics*. Oxford, England: Oxford University Press.

Yi, B. (2005). The logic and meaning of plurals. part i. *Journal of Philosophical Logic 34*(5-6), 459–506.

Yi, B. (2006). The logic and meaning of plurals. part ii. *Journal of Philosophical Logic 35*(3), 239–288.