SIMILATIVE PLURALS AND THE NATURE OF ALTERNATIVES

by

Ryan Smith

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF LINGUISTICS

In Partial Fulfillment of the Requirements

For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

2020

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation prepared by Ryan Walter Smith titled *Similative Plurality and the Nature of Alternatives,* and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy.

_________________________________ Date: May 12, 2020
Robert Henderson

_________________________________ Date: May 12, 2020
Heidi Harley

_________________________________ Date: May 12, 2020
Andrew Carnie

_________________________________ Date: May 12, 2020
Simin Karimi

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.

We hereby certify that we have read this dissertation prepared under our direction and recommend that it be accepted as fulfilling the dissertation requirement.

_________________________________ Date: May 12, 2020
Robert Henderson
Dissertation Committee Co-Chair
Department of Linguistics

_________________________________ Date: May 12, 2020
Heidi Harley
Dissertation Committee Co-Chair
Department of Linguistics

ACKNOWLEDGEMENTS

I've heard it said that the acknowledgements section of the dissertation is the hardest part of the entire thing, and I can see why: it's impossible to express in only a few paragraphs just how much this six-year journey has meant to me, and how much I've grown in the process. U of A, and the city of Tucson as a whole, will always hold a special place in my heart. Given the richness of my grad school experience, it will be impossible to thank everyone who has contributed to just how wonderful my time here was. That said, I'll give it my best shot.

I owe an enormous debt of thanks to my committee: my advisors Robert Henderson and Heidi Harley, and my committee members Simin Karimi and Andrew Carnie.

Robert played an absolutely pivotal role in my becoming a semanticist. I was thrilled when he joined the department in my second year, and it was thanks to his seminar on dependent plurals that I came upon the topic that would grow into this dissertation. It was a real pleasure working through analyses in his office, and he always encouraged me and helped me sharpen my thinking on a problem, even in the early days when I was nervous about showing him my formulas. If I'm even half the semanticist he is, I'd say I turned out pretty well.

It goes without saying that Heidi is an intellectual giant, one of the most brilliant people I have ever known. I knew from the summer before I started grad school that she would be my advisor, and that was a great choice. Heidi has always challenged me, pointing out fascinating connections that I hadn't thought of, while also being a great champion of my work. She's also an incredibly kind person, a joy to talk to. I'll miss coming to her office for meetings.

Despite not being my advisor, Simin is almost certainly the professor I've worked with most closely in the department, and someone I can call a dear friend. I owe her a huge debt of gratitude for funding me through her NSF grants on the grammars of Iranian languages, through which I gained invaluable experience in fieldwork and an immense knowledge of the Iranian language family. With her help, I also got the chance to travel to Europe and Singapore, and make connections in the field of Iranian linguistics. Besides the professional benefits, she has also always been there for me in my times of great personal need. I cannot thank her enough for all she has done for me.

Andrew was the first person I met in the Linguistics department, and he's really the whole reason I came to U of A in the first place. I still remember eating lunch at Sabor with him when I was a mere prospective student. We've had many engaging conversations since then, both in his office and walking to and from class. It was an immense pleasure taking syntax with him, as well as TAing for him for syntax at the undergrad and grad level.

The U of A linguistics department is filled with wonderful faculty, and I've had a number of great teachers in the department outside of my committee. Chief among these, I'd like to thank Mike Hammond, who instilled in me a great love for phonology that I'll carry with me into my career as a linguist, and Amy Fountain, who was a model teaching mentor early in grad school.

I met some of my best friends in grad school, and while I can't thank everyone, a few folks really deserve direct shout outs. A huge thanks to Kristina Mihajlovic, Jianrong Yu, Amber Lubera, Lucy Hall Hartley, Jian Gang Ngui, and Shin-Ichi Kitada, for being the best squad a guy could ask for. JR and Amber are great collaborators (not to mention great karaoke buddies), and have really been there for me when I was at my lowest. Lucy and I could always bond over books and music (not to mention the occasional gossip session), and Jian Gang (GANG!) has been a great pal, both when talking syntax/semantics and just goofing around. Shin, though with us for only a short time and now back in Japan, definitely solidified himself as part of the crew, and I look forward to hanging out with him sometime in Tokyo! Finally, I can't even begin to do justice to how much Kristina means to me: she is the best friend I've ever had, and so much more. She gets her own paragraph below.

I owe an enormous debt of gratitude to Mohsen Mahdavi Mazdeh. Without his Persian lessons early on in grad school, I never would have written this dissertation, or gotten involved in research on Persian or the Iranian languages in the first place. He is also a fantastic linguist, and a great friend. *xeyli mamnoon, dooste man!*. Similarly, I want to thank Ryoichiro Kobayashi, who collaborated with me on some of the earliest work on *-toka* and *-tachi*, as well as some other phenomena in Japanese. Without his judgments and great linguistic sense, much of chapter 4 of this dissertation would not exist. Ryo is also a great linguist, and I'm happy to have him as a collaborator and friend.

No dissertation acknowledgements are complete without a shoutout to the cohort: Colleen Patton, Bill Cotter, stanley donahoo (I know how you like your name all lower case!), Skye Anderson, Zechy Wong (see you in Singapore!), Shannon "Gotta Do It!" Grippando, Leah Rice, and especially Joe Dupris. Joe gets an even bigger shoutout for all the great times we had together, including being my roommate in our second year, being my gym buddy for years (it's thanks to you I finally got in shape!), and for all the late nights playing guitar and watching the Simpsons and Invader Zim

Thanks as well to my other friends in linguistics (in no particular order): Christina

Newhall (and Roo!), Jon Geary (one of the funniest and most competent people I've known, and a great conference co-organizer), Andrew Zupon (Big Z!), Masha Alexeeva, George-Michael Pescaru, Josh Meyer, Dave Medeiros, Patrick Thomas (who has incredible taste in music, I must say!), Luis Irizarry (pretty cool), Rana Nabors, Roya Kabiri (and Hamed!), Christian Ruvalcaba, Muriel Fisher (the best office neighbor!) and Bri Alexander, among many, many others. Thank you all!

A special thanks goes to my first linguistics professors at Baylor: Lydia Grebenyova, Jeanette Marsh, Clay Butler, and Melisa Dracos. Dr. Grebenyova in particular set me on the path to becoming a formal linguist. I wouldn't be doing this now if it weren't for her. Likewise, I thank Yuko Prefume for teaching me Japanese, and both her and Yoshiko Fujii Gaines for encouraging me in my interest in Japanese linguistics both at Baylor and beyond. *Arigatoo gozaimasu!*

This dissertation wouldn't exist without the judgments of many native Persian and Japanese speakers. For this I thank Mohsen, Rana, Roya, Simin, Masoud Jasbi, Nazila Shafiei, Vahideh Rasekhi, and audiences at NACIL2 for their feedback on the judgments reported here. Similarly, I thank Ryo, Michael Yoshitaka Erlewine, Kunihiko Yabushita, Yuko Prefume, Yoshiko Fujii Gaines, and audience members at a variety of conferences for their Japanese judgments.

The biggest open secret is that we academics would be completely lost without the support of our departments' business staff and coordinators. I want to include a special thanks to Marian Wiseley, Kimberley Young, and Stephanie Amado. Marian in particular has been a huge help throughout my grad school career, and has shown me great kindness despite the times where I definitely bungled up something financial. I'll miss coming to her office to talk with her.

Of course, I wouldn't even be here today if it weren't for my family: many thanks to Mom and Dad, for bringing me into this world, raising me, and supporting me every step of the way. Many thanks as well to my grandparents Ted and Maureen Douglas, who also had a big hand in raising me, and have always been there for me throughout my life. Thanks to my sister Megan and brother-in-law Logan, and my cousins Brittlea and Carley Douglas. And of course, thanks to Cookie, Taco, Bunny, Jameson, Day-Z and Dexter, as well as to Lilcee, Fred, Chloe, and Miley, for touching my life, and being so wonderful in your own ways.

Finally, I thank Kristina, my best friend and love, for everything: from the goofy memes to the long, emotionally deep conversations, from movies and karaoke to hikes and running a 10k, for being there for me at my best and at my lowest, for constantly encouraging and inspiring me to be the best version of myself (*semper ad meliora*). Of all the things I've done and experiences I've had in Tucson, the best was meeting you. *Volim te puno.*

DEDICATION

This dissertation is dedicated to my mother, Sharon Smith, who always encouraged me to pursue my interests and be my own person. It's been a pleasure sharing everything I've learned with you, Mom. I love you!

TABLE OF CONTENTS

TABLE OF CONTENTS – *Continued*

TABLE OF CONTENTS – *Continued*

ABSTRACT

Standard approaches to implicature hold that scalar implicatures arise when listeners enrich the basic meaning of a speakers utterance by reasoning over a constrained set of alternatives to that utterance (Grice 1975; Horn 1978; Geurts 1998, a.o.). For example, a sentence like Rostam ate some of the cookies implies that Rostam did not eat all of the cookies because, so the reasoning goes, the speaker could have uttered the stronger sentence Rostam ate all of the cookies, but chose not to. On this view, then, alternatives are 1) calculated on the basis of entire utterances, and 2) derived via replacement of certain expressions in the utterance with other expressions in the lexicon of the language (Katzir 2007).

In this dissertation, I challenge this view on the basis of the semantic and pragmatic properties of an understudied variety of plural, similative plurals, focusing on m-reduplication in Persian and the morphemes -toka and -tari in Japanese as case studies. These expressions are associated with a non-homogeneous plural inference: they refer to a plural entity composed of at least one entity in the denotation of the bare nominal to which the plural applies, and at least one entity in a set that is in some sense similar to the set denoted by the bare nominal. Using evidence from downward-entailing semantic contexts, as well as pragmatic contexts involving speaker ignorance, I demonstrate that this inference is not entailed by similative plurals, but is merely implicated.

My proposed analysis of the non-homogeneous plural implicature has major consequences for the theory of implicatures. First, in order to derive this implicature, calculation must occur within the syntactic structure prior to existential closure of

the event variable, and is thus a case of subsentential implicature calculation (Landman 2000; Chierchia 2004; Zweig 2009). Furthermore, the analysis of one set of speaker's judgments calls for the use of an abstract alternative, one that does not correspond to any lexical item of the languages in question (Chemla 2007; Buccola et al. 2018; Charlow 2019). This alternative cannot be derived via a series of deletion and lexical replacement operations on the structural representation of the sentence under evaluation, and thus poses a problem for theories requiring alternatives to be derived by such means (Katzir 2007). Instead, I propose that the required alternatives be derived from the conceptual representation of the expression. Following this, I extend the analysis to include other types of non-homogeneous plurality, such as associative plurals and personal pronoun constructions, and demonstrate the existence of a common core in the semantics of non-homogeneous plurals.

CHAPTER 1

INTRODUCTION: NON-HOMOGENEOUS PLURALITY

1.1 Non-homogeneous plurality

This dissertation is concerned with the formal semantics and pragmatics of *non-homogeneous plurality*, plurals that may refer to multiple individuals not all of which fall into the denotation of the expression to which the plural applies. Non-homogeneous plurals fall into two types, depending on the nature of the relationship between the individuals in the set the plural takes as an argument and those individuals in the plurality that fall outside that set. The first type, *similative plurals*, are plurals in which only one individual in the plurality needs to be in the set denoted by the plural's complement, while other individuals are from sets that are merely *similar* in the context to the base set from which the plural is derived. The second type, *associative plurals*, are those in which the plurality is composed of a core individual or group of individuals, along with those associated with the individual in some way, such as family members, friends, or colleagues. Similative plurals are exemplified by *m-reduplication* in Persian (1), and arises as well in Japanese with the morphemes *-toka* and *-tari* (2). A representative associative plural is Japanese *-tachi*, given in (3).

(1) Mohsen ketâb metâb xund
    Mohsen book RED read.PST
    'Mohsen read a book and other such things'

(2) a. Taro -toka -ga ki -ta
Taro -TOKA -NOM come -PST
'Taro and someone else like that came'

b. Taro -ga heya -o sooji si -tari si -ta
Taro -NOM room -ACC clean do -TARI do -PST
'Taro cleaned his room and did other such things.'

(3) Taro -tachi -ga ki -ta
Taro -ASSOC -NOM come -PST
'Taro and his associates came.'

The non-homogeneity of each of these plural expressions distinguishes them from English-type bare plurals like (4), which are interpreted homogeneously: every member of the plurality has the property named by the argument of the plural morpheme *-s*.

(4) John saw dogs

1.2 Goals of the Dissertation

Non-homogeneous plurals are understudied in the formal linguistics literature; similative plurals in particular, to my knowledge, have not analyzed in the formal semantics literature. This dissertation thus offers the first formal treatment of an understudied variety of plurality.

While a formal treatment of similative plurality and non-homogeneous plurality in general is an important contribution in and of itself, the study of non-homogeneous plurality also informs semantic and pragmatic theory. In particular, in chapter 3 and 4, I argue that the implicature associated with similative plurals, what I call the *non-homogeneous plural inference*, requires two radical departures from standard approaches to scalar implicature: first, rather than computing implicatures over entire *utterances* or *speech acts*, as in classical Gricean and more recent Neo-Gricean analyses (Horn (1972); Grice (1975); Horn (1984); Geurts (2010), a.o.), we must resort to *subsentential implicature calculation*, according to which implicatures are calculated within the syntactic structure as a compositional semantic mechanism (Landman (2000); Chierchia (2004); Chierchia (2006)). Second, and perhaps even more radical, we must allow alternatives to be *abstract*: they need not correspond to a lexical item in the language at issue (Chemla (2007); Buccola et al. (2018); Charlow (2019)). This not only poses a problem for traditional Gricean and Neo-Gricean approaches to implicature, due to their focus on what is or could be *uttered* as the unit of implicature computation, but is also problematic for approaches that derive alternatives by means of deletion and lexical replacement operations on syntactic structures (Katzir (2007)). The pragmatics of similative plurals thus lead to the development of a novel, non-Gricean theory to implicatures and alternative.

1.3 Outline of the dissertation

The dissertation is structured as follows. Chapter 2 develops a formal system for reasoning about similarity. In it, I derive a *similarity metric* based on a formalization of the intuition that similarity involves a contextually sufficient degree of shared properties. I then construct a *similarity relation* on the basis of this metric, prove that it is *reflexive*, *symmetric*, and *non-transitive*, all properties that similarity intuitively has, and use it to construct a notion of *similarity set*, to be used in the analysis of similative plurals in the following chapters. Chapter 3 introduces the first case study of similative plurals, Persian m-reduplication. It is shown that the non-homogeneous plural inference associated with m-reduplication in upward-entailing contexts disappears in downward-entailing and non-monotonic environments, as well as in pragmatic contexts establishing speaker ignorance. Drawing a comparison between the non-homogeneous plural inference and the multiplicity inference observed with English bare plurals (Krifka (2004); Spector (2007); Zweig (2009), a.o.), I analyze the non-homogeneous plural inference as a scalar implicature, and show that the derivation of this implicature requires local implicature calculation and abstract alternatives, as discussed above. This section also treats variation in the interpretation of m-reduplication in non-upward-entailing environments. I then show that this system makes a number of correct predictions, such as the incompatibility of m-reduplication with differential object marking in the absence of an additional definite plural marker, and the existence of *dependent readings* when an m-reduplicated nominal occurs in the scope of another plural expression (Zweig (2009)). Chapter 4 moves to the second case study of similative plurality, Japanese *-toka* and *-tari*. I show that, despite some syntactic differences from m-reduplication, these expressions possess remarkably similar semantic and pragmatic properties to m-reduplication, and develop a semantic and pragmatic analysis of these particles using the same strategy as that adopted in the preceding chapter. Chapter 5 moves to a different type of non-homogeneous plurality, *associative plurality*, and shows how the same general tools used for the analysis of similative plurality can be extended to that of associatives. This chapter also develops an analysis of *plural pronoun constructions* (PPCs), in which a plural pronoun with a comitative PP complement is seemingly interpreted as potentially singular in reference, as in the Russian example in (5) (Vassilieva (2005); Vassilieva and Larson (2005)).

(5) | My | s | Pet | -ej | pojdë | -m | domoj |
|---|---|---|---|---|---|---|
| 1.PL.NOM | with | Petya | -INSTR | go.PRF.FUT | -1.PL | home |

‘Peter and I will go home/Peter and we will go home.’

I show that associative plurals and PPCs may be given an identical analysis, deriving the strong and weak non-homogeneous behavior of both classes of phenomena under a unified framework. I further explore the connection between similatives on the one hand and associatives/PPCs on the other, showing that both involve mixture operations on a set of individuals and another set of individuals related the first set in some way. Chapter six concludes the dissertation with a discussion of the theoretical issues raised throughout the dissertation, the range of variation exhibited in the interpretation of similatives, focusing on constraints on alternative generation, predicted constraints on possible similative markers following from the formal system in chapter 2, and future directions for research prompted by the work reported in this dissertation.

CHAPTER 2

DEFINING SIMILARITY

The goal of this chapter is to develop a formal system to capture intuitions about *similarity*, which will play a role throughout this dissertation. The formalism aims to construct a *similarity relation*, with reflexive, symmetric, and non-transitive properties, derived from the proportion of properties two expressions share in the lexicon. This relation is then used to construct *similarity sets*, which will play a major role in the analysis of natural language phenomena in the next several chapters. Formal techniques from the theory of generalized quantifiers (Barwise and Cooper (1981): Keenan and Stavi (1986)) and type-shifting (Montague (1973); Partee (1987) will be used extensively throughout this chapter, establishing a strong connection between these domains and the formal theory of similarity I develop here.

The chapter is structured as follows. The first section details properties that a similarity relation should intuitively have. Next, I formalize a notion of similarity *between individuals*, deriving from a treatment of similarity as *sufficient degree of overlap in shared properties*. This is formalized by collecting the set of properties two individuals have and finding the *Jaccard Index* between these sets to define a similarity metric (Jaccard (1901), Levandowsky and Winter (1971)). I then derive a similarity relation between individuals and prove that it exhibits the desired formal properties. The third section takes the results of the previous section and extends them to similarity *between functions from individuals to truth values* or *predicates*, defining a similar metric between these functions as well as a similarity relation between functions with the same set of desired properties. I further prove a connection between individual and functional similarity, such that two individuals are individually similar whenever their IDENT-*shifted* interpretations are functionally similar. Finally, I demonstrate how these definitions can be used to define a notion of *similarity set*, the set of similar objects in the context. Two equivalent formalizations of these sets will be offered, one in terms of similarity relations, the other in terms of a topological notion known as the *closed r-ball around a point.* A method of providing contextual restrictions on these sets will also be introduced. These notions will be extensively used in the analysis of similative plurality and of prototype inferences in the next several chapters.

2.1 Properties of the similarity relation

Intuitively, the notion of similarity corresponds to a relation between objects: two objects are *similar* if they share a sufficient number of properties, and dissimilar if they do not share enough of them, with what counts as *sufficient* depending on the context. The similarity relation should possess three key properties: *reflexivity*, *symmetry*, and *non-transitivity.* A relation is *reflexive* if, for any individual x, the relation holds between x and itself. In this setting, this means that every individual should be similar to itself, regardless of the context, because any individual will share all of its properties with itself. In this sense, individuals are *maximally self-similar.*A relation is *symmetric* if, for any two individuals x and y, if the relation holds between x and y, then the relation holds between y and x. This intuitively applies to similarity as well: if an individual x counts as similar to another individual y, then y should also count as similar to x. This is because if x shares enough properties to count as similar to y, then y should have the same degree of overlap in properties with x.

The final property of the relation, non-transitivity, is perhaps less intuitive. A relation is *non-transitive* when it is not necessarily the case that for three individuals x, y, and z, if the relation holds between x and y, as well as between y and z, then we can conclude that it holds between x and z as well. Here, we should *not* be able to conclude that x is similar to z if x is similar to y and y is similar to z[1]. The reason for this is that, if the similarity relation were transitive, then we could end up with an undesirable situation in which everything is similar to everything else. This clashes with our intuitions about what it means for things to be similar. To see this, imagine a chain of individuals arranged according to how similar they are to one another. While each individual is similar to the person they are directly adjacent to, once we move far enough away from the individual we started with, we will almost certainly find that the individual we end up at has very little in common with the first individual.

Taken together, the properties of reflexivity, symmetry, and non-transitivity define what is known as a *tolerance relation*. Tolerance relations are quite similar to a better known class of relations, equivalence relations, in being reflexive and transitive. They differ in that while equivalence relations are transitive, tolerance relations need not be. In this case, all equivalence relations are tolerance relations, but the reverse does not hold. Tolerance relations are useful for comparing objects along a scale of

[1]Note that this is not the same as requiring our relation to be *intransitive*. This is a stronger condition, stating that there is *no* x, y, z, such that x is similar to y, y is similar to z, and x is similar to z. Intuitively this is quite clear: three people lined up according to similarity may all be similar to one another. It simply doesn't follow that only two people directly adjacent to one another count as similar.

distinguishability, and as such are perfect for modeling similarity between objects.

2.2 Individual similarity

Having discussed the properties a similarity relation should have, I will now proceed to construct such a similarity relation. This section is devoted specifically to similarity between elements of the domain of individuals: *individual similarity.* The core of this idea is very simple: individuals are compared to one another with respect to the *properties they have in common.* The more properties two individuals share, the more similar they are, and whether they count as similar *simpliciter* will depend on contextual factors.

The construction of the similarity relation between individuals will proceed in several steps. First, I will provide a rigorous definition of the otherwise vague term 'property' as well as a method for collecting these properties into a set. Next, I will construct a similarity function that makes reference to the property sets of the individuals whose degree of similarity is being evaluated and explore its formal properties. Finally, I will make use of this function to construct a similarity relation between individuals, and will prove that this relation has all of the desirable properties discussed above.

2.2.1 Properties and property sets

The definition of *property* I use is typical of formal semantics: I define it simply as a *predicate*, or function from individuals to truth values[2]. A property of an individual is then any predicate that maps that individual to true. Equivalently, we can think of this property as a set of which the individual in question is a member, with the

[2]In an intensional semantics, the term *property* is typically reserved for functions from possible worlds to predicates. For the sake of convenience I restrict my attention to extensional models.

predicate being the *characteristic function* of this set. We can always retrieve the underlying set from its characteristic function *f*. One way to do this is with Winter (2016)'s * function, as follows.

(6) $f^* = \{x \mid f(x) = 1\}$

Having defined the relevant notion of property, we need a way to collect the properties of an individual into a set. This is the *property set of an individual*, denoted $\textsc{prop}(x)$ for any individual c.

(7) Property set of an individual
$\textsc{prop}(c) = \{P \mid P(c) = 1\}$

It turns out that there is another, equivalent way to define the property set of an individual. Namely, we can take advantage of the theory of *type-shifting* and the ability to retrieve a set from its characteristic function. The first step is to take the *Montague Lift* of an individual, which takes an individual of type e and returns a function of type <<e,t>,t>, a function from predicates to truth values (Montague (1973), Keenan and Stavi (1986), Partee (1987)).

(8) Montague Lift of an individual
$\textsc{lift}(c) = \lambda P.P(c)$

The LIFT function was originally used by Montague (1973) to permit composition of an individual with verbs, which in his system took arguments of type <<e,t>,t>. Later, Partee and Rooth (1983) made use of LIFT to permit generalized conjunction and disjunction of ordinary individuals with quantificational noun phrases, and Partee (1987) included it in her study of type-shifting operations and their interconnections.

For our purposes, LIFT is useful precisely because it generates the *characteristic func-*

tion of the set of functions that map a given individual to true. In other words, its underlying set is the property set of an individual!

(9) $\textsc{lift}(c)^* = \{P \mid P(c) = 1\}$

We can therefore define the property set of an individual, using independently motivated tools from the formal semantics literature, as the underlying set of the Montague Lift of that individual.

(10) $\textsc{prop}(c) = \textsc{lift}(c)^*$

It is worth seeing this technique in action. Assume we have a model consisting of a domain of individuals {j, m}, the set of truth values {0, 1}, and three functions, given below.

(11) $f = [\mathrm{j} \mapsto 1,\ \mathrm{m} \mapsto 0]$, $g = [\mathrm{j} \mapsto 1,\ \mathrm{m} \mapsto 1]$, $h = [\mathrm{j} \mapsto 0,\ \mathrm{m} \mapsto 1]$

LIFT(j), then, is the function that maps any function to 1 if it maps j to 1, and to 0 otherwise.

(12) $\textsc{lift}(\mathrm{j}) = [f \mapsto 1,\ g \mapsto 1,\ h \mapsto 0]$

Finally, we retrieve j's property set by applying the * function. This delivers to us the set containing the functions f and g, exactly the result desired.

(13) $\textsc{lift}(\mathrm{j})^* = \{f, g\}$

2.2.2 The similarity metric

Armed with this technique for collecting up the properties of an individual, we can now formalize a notion of the degree of similarity between two individuals with respect to the properties they share. I formalize this as the ratio of the number of properties two individuals have in common over the number of properties they have overall. Set-theoretically, this is the cardinality of the *intersection* of two property sets divided by the cardinality of the *union* of those sets. This function is known as the *Jaccard Index* of two sets (Jaccard, 1901), used here to define the *similarity function* between the two *individuals*, and denoted throughout by sim(x,y).

(14) Individual similarity function

$$\text{sim(x,y)} = \frac{|\textsc{prop}(x) \cap \textsc{prop}(y)|}{|\textsc{prop}(x) \cup \textsc{prop}(y)|}$$

This way of measuring similarity has several important properties that will allow us to define a similarity relation with the appropriate properties. It is worth it to explore some of these properties in detail. First, the similarity function always returns a value between 0 and 1. This is because the cardinality of the union of two sets is always greater than or equal to that of the intersection of those sets. This in turn follows from the fact that the intersection of two sets is always a subset of their union. Therefore, the ratio of the two cardinalities will always be less than or equal to one. What is more, if two property sets have no elements in common, the similarity function will return a value of 0, the smallest possible value. Taken all together, on this definition of the similarity function, two objects x and y are *maximally similar* if sim(x,y) = 1, and *maximally dissimilar* if sim(x,y) = 0.

In exploring the properties of the similarity function, it is useful to provide a method for converting it into a *distance function* by subtracting the value of sim(x,y) from 1. We do this in order to show that the function we've developed is a *metric* on the set of individuals, which is usually done in terms of the properties of a distance metric. The analogue of sim defined in terms of distance is given below, where $\triangle$ denotes the *symmetric difference* of two sets.

(15) $$\text{d(x,y)} = 1 - \text{sim(x,y)} = \frac{|\textsc{prop}(x) \triangle \textsc{prop}(y)|}{|\textsc{prop}(x) \cup \textsc{prop}(y)|}$$

The distance function here and its similarity counterpart are not quite true metrics in all models[3]: it possesses the following characteristics that true metrics have:

(16) Properties of the distance/similarity function

1. Non-negativity: d(x,y) $\geq$ 0 (sim(x,y) $\geq$ 0)
2. Minimal self-difference/maximal self-similarity: d(x,x) = 0 (sim(x,x) = 1)
3. Symmetry: d(x,y) = d(y,x) (sim(x,y) = sim(y,x))
4. Triangle inequality: d(x,y) + d(y,z) $\geq$ d(x,z)

Each of these properties follows from the definition of the functions. First, non-negativity follows from the fact that the size of sets is always non-negative: the cardinality function only returns natural number values, and the smallest value it can return is that of the empty set, 0. Minimal self-difference/maximal self-similarity follows from the fact that the intersection and union of the property sets of x are equal, so sim(x,x) will always equal 1, and d(x,x) will always equal 0, the latter because the symmetric difference of identical sets is the empty set, which has cardinality 0. Symmetry follows from the commutativity of union, intersection, and symmetric difference): changing the order of the sets being operated on does not change the result of applying these operations[4].

In order for a similarity/distance function to qualify as a true metric, it must possess all of the properties we have shown sim/d to have above. However, they must pos-

[3]The Jaccard Index, used to define the similarity function over individuals, *is* a metric on the collection of finite sets. However, because I am using it indirectly to define functions over *individuals*, some of its properties get lost in the transition.

[4]The final property, the triangle inequality, is inherited from the Jaccard Index, and is a bit more difficult to prove. Because no properties of the similarity relations I develop using the similarity function hinge crucially on the triangle inequality, I will not prove it here. See Levandowsky and Winter (1971) for a proof of the triangle inequality for the Jaccard Index, which holds for this application as well.

sess an additional property as well. This is the property of *identity of indiscernables*, defined below.

(17) Identity of indiscernables
$d(x,y) = 0 \rightarrow x = y$ ($sim(x,y) = 1 \rightarrow x = y$)

What identity of indiscernables says is that two individuals can only be maximally self-similar or minimally self-different if they are the same individual. This property does *not* hold of the similarity and distance functions we have defined for all models. We can see that this fails if we consider two individuals that just happen to have identical property sets: by happy accident, they happen to be in the same sets named by the language, but they need not be the same object. Because of this failure to satisfy identity of indiscernables, sim and d are*pseudometrics*.

There is a circumstance in which these functions do qualify as full metrics: in models where an identity function for each individual is present: $\lambda x.x = c$, for given individual c. The identity function on an individual will only be present in the property set for that individual. This has the consequence that the only individual maximally similar/minimally different from an individual c will be c itself, thereby validating identity of indiscernables. The systematic availability of identity functions on individuals will play a crucial role later in this chapter, so for our purposes the similarity and distance function defined above is a true metric.

2.2.3 The similarity relation between individuals

The final ingredient we need to define a similarity relation is an idea borrowed from degree semantics, the concept of a *standard of comparison*, which is most likely contextually determined but can also be fixed. I'll denote this with Strd(sim), read as the standard of the similarity relation, or just the standard. This function returns a numerical value, which, like the similarity metric itself, can only take values between 0 and 1. The similarity relation between individuals $\sim_i$ can now be defined as follows.

(18) Similarity relation between individuals
x $\sim_i$ y iff sim(x,y) $\geq$ Strd(sim)

That is, x and y are similar iff their measure of similarity is greater than or equal to the standard.

Does $\sim_i$ have the properties we want a similarity relation to have? Recall that we want our relation to be reflexive and symmetric, but not transitive. Here I will prove that $\sim_i$ has these properties.

Reflexivity

THEOREM: $\sim_i$ is reflexive.

PROOF: This follows from the maximal self-similarity property of sim, which in turn follows from the fact that the intersection of the property set of x with itself is equal to the union of the same: they are both equal to the property set itself. As such, no matter how large the property set of x is, sim(x,x) = 1. Furthermore, because Strd(sim) can only be set between 0 and 1, and because the standard can never be greater than 1, x will always count as similar to itself regardless of the value of Strd(sim). QED

Symmetry

THEOREM: $\sim_i$ is symmetric

PROOF: This follows from the symmetric property of sim, which in turn follows from the commutativity of union and intersection in the definition of sim: PROP(x) $\cup$ PROP(y) = PROP(y) $\cup$ PROP(x), and, likewise, PROP(x) $\cap$ PROP(y) = PROP(y) $\cap$

PROP(x). Therefore, sim(x,y) = sim(y,x). Because of this, it follows that if sim(x,y) $\geq$ Strd(sim), then sim(y,x) $\geq$ Strd(sim), each of these just being expansions of the definition of $\sim_{\mathrm{i}}$. QED

Non-transitivity

THEOREM: $\sim_{\mathrm{i}}$ is *not* transitive. PROOF: It suffices to show that given some universe and some setting of Strd(sim) that transitivity need not hold of some three elements x y and z. Consider the model in (19).

(19) D = {w, x, y, z}
f = [x $\mapsto$ 1, y $\mapsto$ 1, z $\mapsto$ 0, w $\mapsto$ 0], g = [x $\mapsto$ 0, y $\mapsto$ 1, z $\mapsto$ 1, w $\mapsto$ 0], h = [x $\mapsto$ 1, y $\mapsto$ 1, z $\mapsto$ 0, w $\mapsto$ 1], i = [x $\mapsto$ 0, y $\mapsto$ 1, z $\mapsto$ 1, w $\mapsto$ 1]

Consider now the property sets of x, y, and z.

(20) PROP(x) = {f, h}
PROP(y) = {f, g, h, i}
PROP(z) = {g, i}

Calculating the similarity between each of x, y, and z, we get the following values.

(21) a. sim(x,y) = $\frac{|\{f,h\}|}{|\{f,g,h,i\}|}$ = 1/2
b. sim(y,z) = $\frac{|\{g,i\}|}{|\{f,g,h,i\}|}$ = 1/2
c. sim(x,z) = $\frac{|\{\}|}{|\{f,g,h,i\}|}$ = 0

If we set the value of the standard such that $1/2 \geq$ Strd(sim) > 0, it follows that x $\sim_{\mathrm{i}}$ y and y $\sim_{\mathrm{i}}$ z, but, crucially, x $\nsim_{\mathrm{i}}$ z. This is enough to show that transitivity need not hold for any similarity relation in any model. QED

2.2.4 Interim summary

We have shown that $\sim_{\mathrm{i}}$ is reflexive, symmetric, and non-transitive. Putting all of this together, we have proven that $\sim_{\mathrm{i}}$ is a tolerance relation, having exactly the sort of properties we want a similarity relation to have. This was not stipulated; we did not, for instance, merely introduce a relation and imbue it with the properties we desired by means of axioms. Instead, these properties were *derived* from more basic principles concerning how similarity is calculated.

2.3 Predicate similarity

The previous section dealt with similarity as it relates to members of the domain of individuals. The results of that section provide the foundations for this one, in which I extend the techniques introduced here to define a similarity relation between *functions from individuals to truth values*, or *predicates*. The intuition behind *predicate similarity* is very similar to the one for individual similarity: we will evaluate the degree of similarity between two predcates in terms of the degree to which they share *superterms*. The idea of a 'superterm' of some predicate P, much like the concept of 'property' from the previous section, will also be formulated as a predicate. However, rather than treating it as a predicate that maps P to true, it will be formulated as one that maps to true *every individual that P maps to true*. From a set-theoretic perspective, a superterm is simply a superset of P's underlying set.

The section is structured similarly to the previous one. First, I explain the reason for defining predicate similarity in terms of functions from individuals to truth values rather than sets. Next, I define the property sets of a predicate, and construct a similarity function using exactly the same method as for individual similarity, in terms of the Jaccard Index of the property sets of two predicates. After examining the properties of this function, I then construct a predicate similarity relation, and demonstrate that it, like the individual similarity relation, is a tolerance relation. I then conclude the section by demonstrating a connection between individual and predicate similarity: two individuals are i-similar iff their IDENT-lifts are p-similar.

2.3.1 Why functions?

Before we proceed, a question immediately arises that needs to be addressed: why formulate similarity in terms of type <e,t> functions? Why not, for example, deal directly with sets of individuals? Property sets that amount to collecting up supersets of a given set are relatively easy to define, as done in (22).

(22) Property set of a set PROP(S) = {S' | S $\subset$ S'}

The biggest reason for defining predicate similarity in terms of functions rather than sets is practical, and has to do with integration with the compositional semantics to be used later in the dissertation. Essentially, the sorts of natural language expressions we will be wanting to construct similarity sets for, such as *book*, do not denote, say, the set of books directly, but rather the *characteristic function* of the set of books. As such, the input to the function generating similarity sets will not be a set, but its characteristic function, so it is more convenient to deal with such functions directly. Predicates, then, will be not sets but functions of type <e,t>, which map individuals to the truth values 0 and 1.

2.3.2 Property sets for predicates

Much like in the case of individual similarity, predicate similarity will require a notion of property set. This notion needs to be defined slightly differently, however, as the Montague Lift is a type-shifter specific to individuals. Rather, we will define the property set of a predicate P as the set of functions Q such that, for any individual that P maps to true, Q also maps that individual to true. The formal definition is given in (23).

(23) $\textsc{prop}(\mathrm{f}) = \lambda \mathrm{g}.\forall \mathrm{x}[\mathrm{f}(\mathrm{x}) \rightarrow \mathrm{g}(\mathrm{x})]^*$

(23) is essentially a functional equivalent of the set-based definition from (22) with the universal quantification over individuals playing the role of the subset relation. We can now define a similarity function between predicates exactly the same way we did with individuals, in terms of the Jaccard Index of their property sets.

(24) Predicate similarity function

$$\mathrm{sim(f,g)} = \frac{|\textsc{prop}(f) \cap \textsc{prop}(g)|}{|\textsc{prop}(f) \cup \textsc{prop}(g)|}$$

As before, it's also possible to derive a distance function from the similarity function by subtracting the similarity of two functions from 1.

(25) Predicate distance function

$\mathrm{d(f,g)} = 1 - \mathrm{sim(f,g)}$

It's pretty clear that predicate similarity has the same properties that individual similarity does, for all the same reasons: non-negativity follows from the non-negativity of the cardinality function, maximal self-similarity follows from the equality of the intersection and union of two identical sets, and symmetry follows from the commutativity of intersection and union. It also lacks identity of indiscernables: two sets may happen to be subsets of the same sets, and therefore possess identical property sets. Once again, we are dealing with a pseudo-metric. In fact, it is not clear we can force the predicate similarity function to qualify as a full metric, since the presence of identity functions over *functions* in the model is notnecessary, or even guaranteed, in the same way that those over individuals are. In this way, predicate similarity differs from individual similarity, though this difference will not matter for our purposes.

2.3.3 The predicate similarity relation

Once again, we can define a similarity relation, this time between type <e,t> functions, on the basis of the predicate similarity function:

(26) $\mathrm{f} \sim_{\mathrm{p}} \mathrm{g}$ iff $\mathrm{sim(f,g)} \geq \mathrm{Strd(sim)}$

This relation has the same properties as the individual similarity relation, and this can be shown as we did with that relation.

THEOREM: $\sim_{\mathrm{p}}$ is reflexive.

PROOF: As before, this follows from the maximal self-similarity property of sim, which in turn follows from the fact that the intersection and union of the property set of a function f with itself are equal. They therefore have the same cardinality, and division of the cardinality, regardless of its actual value, by itself is equal to 1. Hence sim(f,f) = 1. Furthermore, because Strd(sim) can only be set between 0 and 1, f will always count as similar to itself regardless of the value of Strd(sim). QED

THEOREM: $\sim_{\mathrm{p}}$ is symmetric.

PROOF: This follows from the symmetric property of sim, which in turn follows from the commutativity of union and intersection in the definition of sim: $\mathrm{PROP(f)} \cup \mathrm{PROP(g)} = \mathrm{PROP(f)} \cup \mathrm{PROP(g)}$, and, likewise, $\mathrm{PROP(f)} \cap \mathrm{PROP(g)} = \mathrm{PROP(g)} \cap \mathrm{PROP(f)}$. Therefore, sim(f,g) = sim(g,f). Because of this, it follows that if $\mathrm{sim(f,g)} \geq \mathrm{Strd(sim)}$, then $\mathrm{sim(f,g)} \geq \mathrm{Strd(sim)}$, each of these just being expansions of the definition of $\sim_{\mathrm{p}}$. QED

THEOREM: $\sim_{\mathrm{p}}$ is non-transitive.

PROOF: $\sim_{\mathrm{p}}$ will not be transitive in every model. This can be demonstrated using a method similar to the one we used to demonstrate the non-transitivity of $\sim_{\mathrm{i}}$, though in this case we will be concerned with how functions relate to other functions rather than with individuals. We can use a similar toy model to the one considered in (19). This slightly modified model is given in (27).

(27) D = {w, x, y, z}
f = [x $\mapsto$ 0, y $\mapsto$ 1, z $\mapsto$ 1, w $\mapsto$ 0],
g = [x $\mapsto$ 0, y $\mapsto$ 1, z $\mapsto$ 1, w $\mapsto$ 1],
h = [x $\mapsto$ 1, y $\mapsto$ 0, z $\mapsto$ 1, w $\mapsto$ 1]
i = [x $\mapsto$ 1, y $\mapsto$ 1, z $\mapsto$ 1, w $\mapsto$ 1]

(28) gives the property sets for the functions f, g, and h. These are the sets of functions that map to true every individual the function of interest maps to true. Note that each function is an element of its own property set.

(28) PROP(f) = {f, g, i}
PROP(g) = {g, i}
PROP(h) = {h, i}

Now we calculate the similarity of f, g, and h to each other.

(29) a. sim(f,g) $= \dfrac{|\{g,i\}|}{|\{f,g,i\}|} = 2/3$
b. sim(g,h) $= \dfrac{|\{i\}|}{|\{g,h,i\}|} = 1/3$
c. sim(f,h) $= \dfrac{|\{i\}|}{|\{f,g,h,i\}|} = 1/4$

Setting the standard of sim to 1/3, it now holds that f $\sim_{\mathrm{p}}$ g, g $\sim_{\mathrm{p}}$ h, but f $\nsim_{\mathrm{p}}$ h. Therefore, transitivity does not hold for $\sim_{\mathrm{p}}$ in this model. QED

We have thus shown that $\sim_{\mathrm{p}}$, like $\sim_{\mathrm{i}}$ before it, is a reflexive, symmetric, non-transitive relation. Taking all of this together, we have proven that it too is a tolerance relation, just as desired.

2.4 Relating individual and predicate similarity

Thus far, I have developed two similarity relations, $\sim_{\mathrm{i}}$ for individuals and $\sim_{\mathrm{p}}$ for type <e,t> functions. A question we can ask at this stage is whether there is any relation between these two relations; that is, can we draw a connection between individual and predicate similarity?

The answer to this question is yes. The connection hinges on the availability of an independently motivated type-shifter, Partee (1987)'s IDENT type-shifter, which takes individuals of type e and returns a function of type <e,t>.

(30) $\textsc{ident}(\mathrm{c}) = \lambda \mathrm{x}.\mathrm{x} = \mathrm{c}$

By virtue of being predicates themselves, IDENT-shifted individuals can be compared via predicate similarity. It can then be shown that if any two individuals satisfy the individual similarity relation, then their IDENT-shifted counterparts satisfy the predicate similarity relation[5]. This is stated formally below.

[5] A word of caution is in order. Technically, this result will only hold if the standard of the individual similarity function is equal to the standard of the predicate similarity function. This need not be the case: if two sim functions are defined, they could in principle be assigned two non-equal standards, which could result in situations in which two individuals count as similar but their IDENT-shifted counterparts do not, and vice versa.

One way to enforce the constraint that the standard of both similarity functions be identical without merely stipulating that it be so is to define a single, *polymorphic* similarity function: a single function is defined over both individuals and predicates, computing the similarity between expressions of either type, as in (i).

(i) $\mathrm{sim}(\alpha,\beta) = \dfrac{|\textsc{prop}(\alpha) \cap \textsc{prop}(\beta)|}{|\textsc{prop}(\alpha) \cup \textsc{prop}(\beta)|}$ for $\alpha,\ \beta \in D$ or $\alpha,\ \beta \in D_{\mathrm{t}}{}^{D}$

This definition collapses the two sim functions defined above into a single function, one defined over both individuals and functions of type <e,t>. This is a variety of *parametric polymorphism*, in which a function defined over more than one type of expression performs the same operation on

THEOREM: x $\sim_i$ y $\Leftrightarrow$ IDENT(x) $\sim_p$ IDENT(y)

To prove this theorem, it is useful to first prove an additional fact, namely the following *lemma*: for any individual, the property set of that individual is identical to the property set of its IDENT-shifted counterpart.

LEMMA: PROP(x) = PROP(IDENT(x))

PROOF: The proof is by contradiction. Recall that PROP(x) is the set of functions that maps x to true. Likewise, PROP(IDENT(x)) is the set of functions that map to true everything that λy.y = x maps to true. In order to falsify this lemma, we would need to show that there is some element of PROP(x) that is not an element of PROP(IDENT(x)), or vice versa. In other words, there needs to be a function that maps x to true that does not map to true everything that everything that λy.y = x does, or, vice versa, there must be a function that maps everything to true that λy.y = x does but which does not map x to true. However, by definition, anything functions that maps x to true is one that will map to true whatever IDENT(x) maps to true, because the only thing that IDENT(x) maps to true is x itself. Conversely, any function that maps everything to true that IDENT(x) does will necessarily be one that maps x to true, for the same reason. We have thus run into a contradiction, and therefore it must be the case that the property sets of x and IDENT(x) are one and the same. QED

A slight variation on the proof of this lemma may prove easier to understand. The proof is still by contradiction, but substitutes set talk for function talk. First, PROP(x), the set of functions that map x to true, may equivalently be considered as the set of sets of which x is a member. Likewise, PROP(IDENT(x)) may be thought

---

expressions of every type for which it is defined. Here, the sim function returns the Jaccard Index of the property sets of individuals and functions. A single standard can then be provided for this function, which will be the same regardless of whether the function is operating on individuals or functions.

of as the set of sets of which {x}, the *singleton set with x as its sole member*, is a *subset*. In this case, we would have to show that there is a set containing x as a member that does not contain {x} as a subset, or, vice versa, that there is a set with {x} as a subset that does not contain x as a member. We immediately see that there can be no such set: any set A such that x $\in$ A is necessarily one that contains {x} as a subset, and vice versa. We again arrive at a contradiction, and must therefore conclude that PROP(x) = PROP(IDENT(x)). QED

Armed with this lemma, we can now turn to the proof of the main theorem we wish to prove, stated again below.

THEOREM: x $\sim_{\mathrm{i}}$ y $\Leftrightarrow$ IDENT(x) $\sim_{\mathrm{p}}$ IDENT(y)

PROOF: As the above lemma shows, PROP(x) and PROP(y) are identical to PROP(IDENT(x)) and PROP(IDENT(y)), respectively. It immediately follows from this that sim(x,y) = sim(IDENT(x),IDENT(y)), because the Jaccard Indexes of the two pairs of property sets will necessarily be the same. Assuming there is a single standard defined for the sim function regardless of the type of expression it operates on, if sim(x,y) $\geq$ Strd(sim), then it is also the case that sim(IDENT(x),IDENT(y)) $\geq$ Strd(sim). The last statements are simply expansions of the definition of each similarity relation, and therefore we have established the truth of the claim. QED

The upshot of this theorem is that it turns out that we don't actually need two similarity relations: we can make use of predicate similarity for individuals by simply shifting them using the IDENT type-shifter, without losing anything from the individual similarity relation. This also permits easy evaluation of the similarity of a predicate and an individual, which will prove useful in one of the applications to be explored later in this dissertation.

2.5 Similarity sets

Having constructed similarity relations between individuals and predicates, I now proceed to apply them to construct the set of objects similar to a given individual or predicate, *similarity sets*, which will play a crucial role in the analysis of similative plurals in the coming chapters. This will be accomplished by a pair of functions, notated $\simeq$, which will take an individual or predicate as an argument and return the set of objects similar to it in the context. These will be sets of individuals regardless of whether the input to $\simeq$ is an individual or a predicate of individuals. This is because in the analysis of similative plurals, we will need to take the individuals that a given function maps to true and "mix" them with the individuals in the similarity set, rather than perform an operation on functions.

We start with the individual case. The similarity set of an individual x is given below.

(31) Similarity set of an individual

$$\simeq(\mathrm{x}) = \{\mathrm{y} \mid \mathrm{y} \sim_{\mathrm{i}} \mathrm{x}\}$$

As one can say, this setup is very straightforward: $\simeq$ takes an individual argument and returns the set of individuals similar to x.

The definition of the similarity set of a predicate, on the other hand, is a bit more involved. We need $\simeq$ to take a predicate of individuals as an argument and return a set of individuals. A definition that accomplishes this is given in (74).

(32) Similarity set of a predicate

$$\simeq(\mathrm{f}) = \bigcup\{\mathrm{g}^* \mid \mathrm{g} \sim_{\mathrm{p}} \mathrm{f}\}$$

We can unpack this definition by conceiving of the operation of $\simeq$ as involving a number of steps. First, $\simeq$ collects up the set of functions g that count as similar to

f. Then, Winter's * function is applied to each function in the set, extracting from each g a set g*, the set of individuals that g maps to true. This transforms our set of functions into a family of sets of individuals. Finally, the $\bigcup$ operator performs iterated union over this set, flattening the family of sets into a single set of individuals. Due to the reflexive property of the similarity relation, the similarity set of an individual will always contain the individual itself. Likewise, the similarity set of a function will always contain those individuals the function itself maps to true. However, we may run into cases in which we want a similarity set that does not contain such individuals; I refer to this notion as the *proper similarity set* of an individual or function, denoted $\sim$(x) or $\sim$(f), defined as follows.

(33) Proper similarity set of an individual
$\sim(\mathrm{x}) = \simeq(\mathrm{x})\backslash\{\mathrm{x}\}$

(34) Proper similarity set of a function
$\sim(\mathrm{f}) = \simeq(\mathrm{f})\backslash\mathrm{f}^*$

Here, proper similarity sets are simply defined as similarity sets of a given individual or function with that individual or set of individuals mapped to true by the function removed. This notion, along with similarity sets *simpliciter*, will also play a role in the analysis of similative plurals to come.

2.6 Contextual restrictions on similarity sets

Thus far, we have considered an unrestricted notion of similarity sets: we construct such sets on the basis of whichever individuals or functions count as similar with respect to the contextual standard, without regard for any other restrictions the objects might have. However, such unrestricted similarity sets don't seem to be very frequently used in natural language. Ordinarily, some additional, contextually determined property is held in mind by speakers when considering what counts as similar to another object. For example, when considering what counts as similar to apples, speakers consider similar kinds of *fruits*, rather than just anything that might be reminiscent of an apple. In this case, speakers might group apples with oranges or plums, but not with small red balls. In another context, the relevant property may not be fruits but red things, such that apples would be grouped with small red balls but not with oranges.

We can model this phenomenon as a species of *contextual domain restriction* (Barwise and Cooper (1981); von Fintel (1994), *inter alia*). This is familiar from the study of natural language quantification, in which a quantificational noun phrase is interpreted as having some additional restriction beyond what is said on the surface. For example, (220) is interpreted as referring not to every bottle in the entire universe, but to those bottles relevant to the current context (Stanley and Gendler Szabó, 2000).

(35) Every bottle is empty

For present purposes, it suffices to define a contextual domain restriction as intersection of the similarity set with some contextually determined set. I will therefore define a version of the functions $\simeq$ and $\sim$ that are relativized to a contextual parameter C, as in (36) and (37) below.

(36) a. $\simeq_C(\mathrm{x}) = \simeq(\mathrm{x}) \cap C$

b. $\sim_C(\mathrm{x}) = \sim(\mathrm{x}) \cap C$

(37) a. $\simeq_C(\mathrm{f}) = \simeq(\mathrm{f}) \cap C$

b. $\sim_C(\mathrm{f}) = \sim(\mathrm{f}) \cap C$

As will be shown in the Chapter 3, some languages, such as Persian, permit only an implicit domain restriction. In this case, the parametrized similarity sets as in (36-37) are appropriate. However, other languages, such as Japanese, to be discussed in Chapter 4, allow for the overt expression of the domain restriction. In these cases, a function will be defined which takes another function as an argument and returns a parametrized similarity set, one restricted by the set of which the functional argument is the characteristic function. In this case, C is replaced by a function, and the definitions above are slightly modified, as in (38) and (39).

(38) a. $\simeq_{\mathrm{g}}(\mathrm{x}) = \simeq(\mathrm{x}) \cap \mathrm{g}^{*}$

b. $\sim_{\mathrm{g}}(\mathrm{x}) = \sim(\mathrm{x}) \cap \mathrm{g}^{*}$

(39) a. $\simeq_{\mathrm{g}}(\mathrm{f}) = \simeq(\mathrm{f}) \cap \mathrm{g}^{*}$

b. $\sim_{\mathrm{g}}(\mathrm{f}) = \sim(\mathrm{f}) \cap \mathrm{g}^{*}$

2.7 The topology of similarity: similarity sets as closed balls

It turns out that the concept of similarity set developed here has close connections to a concept from topology known as an *r-ball*, which is a set of objects whose distance from a given point p falls within a certain radius r. In particular, it maps onto the *closed r-ball*, which includes the points at its boundary[6]. This is defined as in (40).

(40) Closed r-ball

$B_r[p] = \{\mathrm{x} \mid \mathrm{d(x,p)} \leq r\}$

This can also be defined in terms of the similarity function, as follows.

[6] To be distinguished from an *open* ball, which excludes the points at its boundary i.e. only includes the points that are strictly less than its radius. The choice to use a closed ball is largely arbitrary, and depends on whether we want to include individuals/sets at the similarity/distance cut-off.

(41) Closed r-ball (similarity-based definition)

$B_r[p] = \{\text{x} \mid \text{sim(x,p)} \geq r\}$

As one can see, the closed ball around a point is equivalent to the similarity set of that point: we can think of r as the standard of the distance/similarity function, and p as the individual or function we require objects to be similar to. This is simply the definition of the similarity set with the similarity relation $\sim$ replaced with its definition in terms of the similarity function and its standard. As such, we can define $\simeq$(x) as an object language predicate that denotes the closed ball around x.

(42) $[\![\simeq(\text{x})]\!] = B_{\mathit{Strd}}[\text{x}]$

Defining $\simeq$(f) in terms of closed balls is similar to how it was defined above: first we define the closed ball around f as a set of sets as follows.

(43) $B_{\mathit{Strd}}[[\![\text{f}]\!]] = \{\text{g*} \mid \text{sim(g,f)} \geq \mathit{Strd}\}$

But we aren't done yet. We want $\simeq$(f) to be the set of *individuals* in one of the sets in the ball around P. We can achieve this by having $\simeq$(f) denote the union of the sets in the ball.

(44) $[\![\simeq(\text{f})]\!] = \bigcup B_{\mathit{Strd}}[[\![\text{f}]\!]]$

We can also define the proper similarity set of an individual/function on this definition. This amounts to essentially the same approach taken in the previous section: removing the individual, or set of individuals mapped to true by the function, from the similarity set.

(45) $[\![\sim(\text{x})]\!] = B_{\mathit{Strd}}[\text{x}]\backslash\{\text{x}\}$

(46) $[\![\sim(\text{f})]\!] = \bigcup B_{\mathit{Strd}}[[\![\text{f}]\!]]\backslash\text{f*}$

2.8 Conclusion

In this chapter, I've constructed similarity and distance functions, two defined for individuals and another two for sets, in terms of the Jaccard Index of their property sets. I've shown that these functions are pseudometrics, given that they satisfy non-negativity, symmetry, the triangle inequality, and maximal self-similarity, but not necessarily identity of indiscernables. From these functions, I've defined similarity relations over individuals and functions of type <e,t>, and have shown both relations to be reflexive, symmetric, and non-transitive, making them tolerance relations as desired. I further demonstrated a connection between these two relations, showing that a similarity relation holds between two individuals iff one holds between the predicate of individuals derived by type-shifting both individuals using Partee (1987)'s IDENT type-shifter. Finally, I defined the notion of similarity sets for individuals and functions in two different but equivalent ways, one in terms of the relations I've constructed and the other in terms of the closed ball around the individual/function being evaluated. A method for parametrizing these sets to another set, whether implicit or explicit, was also introduced. These notions can now be fully and directly incorporated into the analysis of similative plurals to be developed in the next two chapters of this dissertation.

CHAPTER 3

SIMILATIVE PLURALITY: PERSIAN M-REDUPLICATION

This chapter applies the formal tools developed in the previous chapter to the analysis of *similative plurals*, focusing on *m-reduplication* in Persian as the first case study. The major empirical result of this chapter is that the non-homogeneous plural inference associated with m-reduplication—that the expression refers to a plural entity composed of at least one entity in the denotation of the bare nominal to which reduplication is applied, and at least one entity that is in some sense similar to that kind of object—is not *entailed* by reduplication, but merely *implicated.* I propose an analysis of m-reduplication as denoting a *mereological mixture*–the set of objects derived by summing the objects in two sets (Heycock and Zamparelli (1999), Heycock and Zamparelli (2000), Champollion (2015))–of the set denoted by the bare nominal with its similarity set, and further take into account some interspeaker variation in judgments about the meaning of m-reduplication in downward-entailing and question contexts. I then derive the non-homogeneous plural reading as a scalar implicature. This analysis has implications for the theory of implicature calculation, as well as for the theory of alternatives. In particular, the analysis of one set of speakers' judgments calls for the calculation of the implicature *below* the site of existential closure of the event variable, and is thus a case of *local* or *subsentential implicature calculation* (Landman (2000); Chierchia (2004); Zweig (2009). Furthermore, the analysis of the other set of speaker's judgments calls for the use of an *abstract alternative*, one that does not correspond to any lexical item of the language (Chemla (2007);

Buccola et al. (2018); Charlow (2019)). This alternative cannot be derived via a series of deletion and replacement operations on the structural representation of the sentence under evaluation, and thus poses a problem for the structural theory of alternatives (Katzir, 2007). I propose that the required alternatives be derived from a *conceptual* representation of the expression in a *language of thought* (Fodor (1975); Buccola et al. (2018)).

The chapter is structured as follows. First, I provide background on the Persian language, discussing its genetic classification, geographic distribution, and linguistic properties. Second, I detail the properties of Persian m-reduplication, discussing its behavior in upward-entailing contexts and some previous descriptive work on it and related phenomena. Third, I reveal the sensitivity of the interpretation of m-reduplication to the direction of entailment, showing that it possesses inclusive readings in downward-entailing and question contexts. I further demonstrate that establishing speaker ignorance is enough to eliminate the non-homogeneous reading even in upward-entailing contexts. Fourth, I develop the mixture semantic analysis of m-reduplicated nominals, as well as a pragmatic analysis of the derivation of the non-homogeneity inference as a scalar implicature. I also discuss here how the phenomena pose a challenge for global calculation of the implicature and for structural approaches to alternatives. Finally, I discuss some implications of the current analysis. as well as areas for future research.

3.1 Background on Persian

Persian, also known as Farsi, is a Southwestern Iranian language spoken throughout the Greater Iran region of western Asia. Varieties of Persian are the official language of Iran, Afghanistan, and Tajikistan, and it is also spoken by a significant number of people in Uzbekistan, Iraq, and Azerbaijan, not to mention by members of the Iranian diaspora throughout the world. Windfuhr (2009) estimates that there are around 70 million native speakers of Persian, with a total of 110 million speakers if one counts L2 speakers. There are three major varieties of Persian: 1) Iranian Persian, spoken, as one would expect, in Iran, 2) Dari, spoken in Afghanistan, and 3) Tajik, spoken in Tajikistan and, as a minority language, in Uzbekistan. The judgments reported in this chapter are all from Iranian Persian, primarily the colloquial variety of the capital of Iran, Tehran. As I will not be discussing other varieties of Persian in this dissertation, I will henceforth use the term Persian to refer to the Tehrani dialect reported below.

The unmarked word order in Persian is generally Subject-Object-Verb. Verbs agree with their subject in person and number, and may be inflected for tense, aspect, and mood. Much of this can be seen in (47).

(47) man ketâb -o mi- xun -am
1.SG book -DOM ASP- read.PRS -1.SG
'I am reading the book.'

Despite its neutral SOV word order, however, Persian is otherwise a head-initial language: adjectives follow the nouns they modify (48), the language has prepositions rather than postpositions (49), and clausal complements follow clause-selecting verbs (50). The *-e* morpheme in (48) is known as *ezâfe*, and occurs between the noun and the phrase modifying it, as well as between any other APs and NPs modifying the head noun.

(48) ketâb -e bozorg -e man
book -EZ big -EZ 1.SG
'My big book'

(49) bâ Royâ
with Roya
'with Roya'

(50) Simin goft (ke) Royâ raft xune
Simin say.PST (that) Roya go.PST house
'Simin said that Roya went home.'

Persian also possesses differential object marking (DOM), marked by the enclitic -*râ*, which in colloquial speech is normally pronounced -*ro* or -*o*, depending on whether it follows a vowel or consonant, respectively. NPs with DOM differ distributionally from nouns without DOM: the former precede indirect objects and certain kinds of adverbials in neutral word order, while the latter are usually left-adjacent to the verb. (51) illustrates this difference.

(51) a. Ranâ ketâb -o be Mohsen dâd
Rana book -DOM to Mohsen give.PST
'Rana gave the book to Mohsen.'

b. Ranâ be Mohsen ye ketâb dâd
Rana to Mohsen a book give.PST
'Rana gave a book to Mohsen.'

The exact contribution of -*râ* is debated, but researchers generally agree that it contributes a type of specificity (Ghomeshi (1997); Karimi (1999); Jasbi (2015); Karimi and Smith (2019)). I will return to a specific semantic analysis of -*râ* later in this chapter when I discuss its interaction with m-reduplication.

A final aspect of Persian grammar relevant to the current discussion is bare nominals. Persian permits bare nouns in multiple environments. In object position, bare nouns are interpreted existentially, and receive a number-neutral interpretation.

(52) man ketâb xarid -am
1.SG book buy.PST -1.SG

'I bought one or more books.'

Bare nouns in subject position, however, are usually interpreted as *definite* and *singular*.

(53) ketâb xeyli bozorg -e
book very big -COP.PRS.3.SG
'The book is very big.'

One might treat this difference by postulating that bare nouns denote properties, and undergo pseudoincorporation (Massam (2001); Dayal (2011); Fereshteh (2015)). Bare subjects, then, might be analyzed as singular predicates of individuals that have been type-shifted to a definite singular reading via the IOTA type-shifter (Partee, 1987). It should be noted, though, that bare subjects may also appear with existential number-neutral interpretations as well in certain contexts, as in (**??**).

(54) hanuz ketâb ru miz -e!
still book on table COP.PRS.3.SG
'Books are still on the table!'

As I will mostly be concerned with bare nouns in object position in this chapter, I will analyze bare nouns in object position as existentially quantified and number-neutral, without committing to a specific analysis of the phenomenon. I will discuss this further in the analysis section later in this chapter.
This concludes the background on the Persian language. In the next section, I turn to the descriptive properties of Persian's similative plural: *m-reduplication*.

3.2 M-reduplication in Persian

Persian possesses a type of full root reduplication, termed m-reduplication in other languages with a similar construction, and further termed a *similative plural* by

Armoskaite and Kutlu (2013), which applies to nouns to create a *non-homogeneous* plural: that is, the plurality is understood to include objects with a property distinct from that of the overtly mentioned object (Nakanishi and Tomioka, 2004). (55), for instance, is judged true if Mohsen read at least one book, as well as something else similar to a book in the context, such as a magazine.

(55) Mohsen ketâb metâb xund
Mohsen book RED read.PST
'Mohsen read a book and other such things'

(55) is judged infelicitous if only one book was read, if only books were read, or if only a magazine or something else similar to a book was read (with one important exception, to be discussed later). For the first two situations, the non-reduplicated bare plural is used, as exemplified in (56).

(56) Mohsen ketâb xund
Mohsen book read.PST
'Mohsen read (one or more) books.'

The interpretation of an m-reduplicated noun is to a certain extent context-dependent: (55) may be interpreted as referring to a set of reading material. In this case, *ketâb metâb* will refer to the set of reading materials similar to a book, such as a magazine or comic book, but will exclude, for instance, online reading material such as Wikipedia pages. It could also be interpreted as types of entertainment: *ketâb metâb* might then be taken to refer to books, movies, and other pastimes involving fiction, but might exclude such things as board games. Whatever general category is selected in the context, it will always be centered on the sorts of objects denoted by the head in the reduplication.

In addition to appearing bare and receiving an existential interpretation, m-reduplicated nominals can also be quantified (57) and made definite (58). (57) is interpreted as meaning that Mohsen read two things, one of which was a book and the other of which was something similar[1]. (58) means that Mohsen read the set of objects known to both speaker and hearer, which is composed of one or more books and one or more book-like things.

(57) Mohsen do tâ ketâb metâb xund
Mohsen two CL book RED read.PST
'Mohsen read two things, one a book, the other book-like.'

(58) Mohsen ketâb metâb -hâ -ro xund
Mohsen book RED -DEF.PL -DOM read.PST
'Mohsen read the book and book-like things.'

NPs headed by m-reduplicated nouns have several properties that suggest that they behave like plural expressions. First, they are compatible with collective predicates. This is shown in (59), which is true as long as Mohsen collects at least one flower as well as at least one flower-like thing, in this context other types of plant matter such as leaves or sticks.

(59) Mohsen gol mol jam' kard
Mohsen flower RED collection do.PST
'Mohsen collected flowers and other such things.'

Second, NPs with m-reduplicated heads cannot be marked with the differential object marker *-ro* without the presence of the specific plural marker *-hâ*[2].

[1] Some speakers find modification of m-reduplicated nouns by small numerals to be a bit odd. This effect is not specific to m-reduplication, but has been observed with Japanese *-tachi* plurals as well. Typically these sentences are completely natural if a large or imprecise number is used instead.

[2] The contribution of *-hâ* has been proposed to be a kind of definiteness or specificity marking on plural or number-neutral nominals. See Jasbi (2014) for an analysis along these lines.

(60) *Mohsen ketâb metâb -ro xund
Mohsen book RED -DOM read.PST
Intended: 'Mohsen read the one book and bookish thing.'

(61) Mohsen ketâb metâb -hâ -ro xund
Mohsen book RED -SP.PL -DOM read.PST
'Mohsen read the book and book-like things.'

A similar effect is seen when NPs with m-reduplication are in subject position.

(62) *mard pard umad -an mehmuni
man RED come.PST -3.PL party
Intended 'Men and other such people came to the party.'

(63) mard pard -HA umad -an mehmuni
man RED SP.PL come.PST -3.PL party
'The men and such came to the party.'

Intuitively, this is because bare nominals marked by *-ro*, as well as those in subject position, are definite and *singular*; this effect can be observed in (64). M-reduplicated nominals, being plural, are incompatible with this construction, and require the presence of an additional plural specificity marker.

(64) Mohsen ketâb -o xund
Mohsen book -DOM read.PST
'Mohsen read the book.'

Thus far, it appears reasonable to treat NPs involving m-reduplication as involving reference to non-homogeneous plural entities; for example, on an algebraic approach to the semantics of plurals like that of Link (1983), a noun like *ketâb metâb* would denote a set of sum individuals such that each individual contains at least one book as a part, as well as at least one thing similar to a book in the context, and nothing else, as in (65), where 'b' represents an individual book, 'm' represents an individual magazine, and 'c' represents an individual comic book.

(65) $[\![\textit{ketâb metâb}]\!] = \{\mathrm{b}\oplus\mathrm{m},\ \mathrm{b}\oplus\mathrm{c},\ \mathrm{b}\oplus\mathrm{m}\oplus\mathrm{c},\ ...\}$

This analysis is simple, straightforward, and accounts for all the observations above: it correctly predicts that m-reduplicated expressions can be quantified and are compatible with collective predicates, while ruling out the possibility of their cooccurence with *-ro*. It seems, then, that this is all that needs to be said. However, in the next section I will provide evidence that things are not as simple as they seem.

3.3 M-reduplication does not entail non-homogeneity or plurality

There is evidence that suggests that m-reduplicated nominals do *not* exclusively denote sets of non-homogeneous sum individuals, as (65) suggests. In fact, in many semantic environments, m-reduplicated nominals are compatible not only with non-homogeneous plural interpretations, but also with homogeneous plural and singleton interpretations as well. These environments are generally downward-entailing or non-monotone, and include negation, conditionals, polar questions, and imperatives. Furthermore, even in upward-entailing environments, m-reduplicated nominals may have singular or homogeneous plural readings if speaker ignorance is established pragmatically. Finally, there is a certain amount of speaker variation in these judgments: some speakers also permit m-reduplication to denote something from the set of objects similar, but not identical, to those denoted by the bare noun. I go through these cases in turn.

3.3.1 Non-upward-entailing environments

We start with non-upward-entailing environments. First, consider the negation in (66). Here, the speaker is understood to mean not that they simply did not read books *and* similar things, but rather that they did not read *any* books or books and similar things. For some speakers the interpretation is stronger: the speaker did not read *any* books *or any other things* like that[3].

(66)

man	ketâb	metâb	na-	xund	-am
I	book	-RED	NEG-	read.PST	-1.SG

'I didn't read books or anything like that'

Next, we consider conditionals. In (67), the addressee may tell the speaker even if she ate only one or more apples. For some speakers, the addressee may tell the speaker even if she did not eat an apple, but a similar fruit. She need not have eaten, say, an apple *and* an orange for the use of m-reduplication to be felicitous.

(67)

age	sib	mib	mi-	xor	-i,	be	man	be-	gu!
if	apple	RED	IMP-	eat.PRS	-2.SG	to	1.SG	SUBJ-	say

'If you eat an apple or something, tell me!'

We find the same sort of behavior with polar questions. In (68), an affirmative response to the question is possible even if the one answering the question ate only one or more apples and nothing else. Again, for some speakers, an affirmative response is possible even when the answerer ate just something similar to an apple, such as an orange.

(68) a.

emruz	sib	mib	xord	-i?
today	apple	RED	eat.PST	-2.SG

'Did you eat an apple or something today?'

[3] For the sake of simplicity of presentation I gloss examples with the more inclusive interpretations.

b. Âre, ye sib / do tâ sib / ye porteqâl xord -am
yes one apple / two CL apple / one orange eat.PST -1.SG
'Yes, I ate an apple/two apples/an orange.'

The final example we consider involves imperatives. As can be seen in (69), a felicitous way to comply with an imperative containing an m-reduplicated nominal like *sib mib* would be eating an apple, eating two apples, or, for some speakers, eating a similar type of fruit, such as an orange.

(69) sib mib bo- xor!
apple RED SUBJ- eat
'Eat an apple or something!'

All of this suggests that the non-homogeneity inference associated with m-reduplication in Persian is not entailed, but is actually derived via implicature. In this way, m-reduplication exhibits properties in common with English bare plurals, which have readings that exclude non-atomic individuals in upward-entailing contexts, as can be seen in (70), but readings that *include* them in non-upward-entailing contexts, as in (71) (Krifka (2004); Spector (2007); Zweig (2009), *inter alia*).

(70) John saw dogs yesterday
$\rightsquigarrow$ John saw more than one dog yesterday

(71) a. John didn't see dogs yesterday
$\rightsquigarrow$ John didn't see any dogs yesterday
b. Q: Did John see dogs yesterday?
A: Yes, John saw one dog yesterday
A': No, John didn't see any dogs yesterday

Although the multiplicity inference is attested in (70), it vanishes in the scope of negation and in polar questions, as can be seen in (71). We therefore find a parallel between these kinds of two plurality-denoting expressions across languages.

3.3.2 Ignorance contexts

Even in upward-entailing contexts, it is possible to interpret m-reduplication as not being restricted to non-homogeneous plurals, as long as speaker ignorance is established by the context.

(72) CONTEXT: You see Roya carrying a small lunchbox, in which she usually keeps an apple for an afternoon snack, but sometimes brings some other kind of fruit. You don't know exactly how many she has in the box (and are not entirely sure what kind of fruit it is).

Royâ	sib	mib	dâr	-e
Roya	apple	RED	have.PRS	-3.SG

'Roya has an apple or something.'

Here, the speaker is not committed to Roya having more than one apple, nor are they committed to her having anything but apples. In fact, some speakers are not even committed to Roya having an apple in the first place. For these speakers, she could just have some other kind of small round fruit, such as an orange.

We find a parallel with the behavior of English bare plurals here too. de Swart and Farkas (2010) note that bare plurals can be used in contexts in which some number of entities is known to exist, but for which there is not enough evidence to establish exactly how many entities there are. In these cases involving ignorance, a plural may be used when the speaker does not know whether or not there is more than one object, just like in the m-reduplicated example above.

(73) Inclusive reading of bare plurals with ignorance (deSwart & Farkas 2010:30)

a. Context: Speaker walks into basement and notices mouse droppings]
Aghhh, we have mice!

b. Context: Speaker walks into unknown house and notices toys littering the floor
There are children in this house.

All of this goes to show that the non-homogeneous plural inference observed in Persian m-reduplication is much like the multiplicity inference associated with English bare plurals: it is sensitive to the monotonicity of the semantic environment it is in, and also vanishes in more global pragmatic contexts establishing speaker ignorance, even in semantic contexts in which the inference would otherwise be expected to arise. In the following section, I pursue an analysis further connecting m-reduplication to English bare plurals involving scalar implicature.

3.4 Analysis

My analysis consists of two components. The first component is semantic, and is the proposal that m-reduplication denotes a *mereological mixture* of the set denoted by the bare nominal and a set of objects that are similar to that bare nominal in the context. The exact type of mixture denoted by m-reduplication depends will depend on the speaker: for some speakers, the mixture is *fully inclusive*, while for others it is only *partially* so. I will elaborate on the precise formulation of these terms in the following subsection.

The second component of the analysis is pragmatic. Taking the unenriched meaning of an m-reduplicated nominal, the non-homogeneous plural reading of m-reduplication is derived from a partially or fully inclusive interpretation via scalar implicature, yielding an *exclusive mixture* interpretation corresponding to the non-homogeneous plural interpretation.

In what follows, I elaborate on the formal ingredients of the analysis, making precise the notion of *similarity sets*, *mereology*, and *mixtures*.

3.4.1 Ingredients

Similarity sets

The notion of similative plural makes crucial reference to sets of objects similar to other objects. This calls for a *similarity relation* $\sim_C$, which will hold between two predicates. The similarity relation is a *tolerance relation*, meaning it is reflexive and symmetric, but not transitive. In the analysis, I'll be dealing with what I call the *similarity set/predicate* of a given set/predicate, notated $\mathrm{P}^{\simeq}$ for given P. This is the set of objects x for which there is some predicate Q such that Q is similar to P in the context and each x is a Q. A formal definition is given in (74).

(74) $[\![\simeq(\mathrm{P})]\!] = \{\mathrm{x} \mid \exists \mathrm{Q}\text{: } \mathrm{Q} \sim_C \mathrm{P}\ \&\ \mathrm{Q}(\mathrm{x})\}$

Because the similarity relation is reflexive, P is similar to itself, so P is a subset of the similarity set. It is possible to define a similarity set that does not include P, as in (75). This is the *proper similarity set* of P, and I denote it as $\mathrm{P}^{\sim}$.

(75) $[\![\sim(\mathrm{P})]\!] = \{\mathrm{x} \mid \exists \mathrm{Q}\text{: } \mathrm{Q} \sim_C \mathrm{P}\ \&\ \mathrm{Q} \neq \mathrm{P}\ \&\ \mathrm{Q}(\mathrm{x})\}$

Mereology

Throughout this work I assume a mereological approach to the semantics of plurals (Link (1983), Krifka (2004), Champollion (2010), a.o.). In addition to atomic individuals used to model singulars, we will also have sum individuals: the sum of two individuals x and y is denoted by x⊕y. I assume, following standard practice (Champollion, 2010), that one-place predicates denote sets of atomic individuals. In order to derive number-neutral/plural predicates, I make use of the algebraic (or cumulative) closure operator * (Link, 1983), which generates the set of all sums of individuals in a given set. An example of this is given in (76).

(76) a. $[\![P]\!] = \{a, b, c\}$

b. $[\![{*P}]\!] = {*}[\![P]\!] = \{a, b, c, a\oplus b, a\oplus c, b\oplus c, a\oplus b\oplus c\}$

In this work, I will use capital X and Y to denote variables over both atomic and sum individuals. Furthermore, following Landman (2000), I will treat predicates as sets of *atomic* individuals, using the cumuative closure operator to derive predicates containing sum individuals.

Mixtures

The analysis I develop for m-reduplication relies heavily on the notion of a *mixture*, an idea that has been alluded to in various places throughout this paper. To motivate this concept, consider (57), repeated in (77).

(77) Mohsen do tâ ketâb metâb xund
Mohsen two CL book RED read.PST
'Mohsen read two things, one a book, the other book-like.'

In upward-entailing contexts, this sentence is true when there are two things, one of which is a book, and the other something similar. As such, there must be sum individuals in the denotation of *ketâb metâb* with two atomic parts, one of which is a book and the other of which is a similar object. This does not come for free however; one could not simply take the union of the two sets to get the desired effect. Instead, we need to be able to sum individuals from one set with individuals of the other. I will refer to such a set of sums as a *mixture*(Heycock and Zamparelli (1999), Heycock and Zamparelli (2000), Champollion (2015)).[4].

(78) Mereological mixture
$[\![\textsc{Mix}(P,Q)]\!] = \{X \oplus Y \mid X \in {*P}, Y \in {*Q}\}$

To make the effect of a mixture clearer, let us consider an example. Consider two sets M = {m1, m2} and W = {w1, w2}. The mixture of these two sets is the set of sums of each element of the algebraic closure of M with each element of the algebraic closure of W.

(79) $\textsc{Mix}(M,W)$
$= \{X \oplus Y \mid X \in {*M}, Y \in {*W}\}$
$= \{X \oplus Y \mid X \in \{m1, m2, m1 \oplus m2\}, Y \in \{w1, w2, w1 \oplus w2\}\}$
$= \{m1 \oplus w1, m1 \oplus w2, m1 \oplus w1 \oplus w2, m2 \oplus w1, m2 \oplus w2, m2 \oplus w1 \oplus w2, m1 \oplus m2 \oplus w1, m1 \oplus m2 \oplus w2, m1 \oplus m2 \oplus w1 \oplus w2\}$

The mixture of two sets denotes another set. I treat the set formed from the mixture of two others as a predicate, such that it is possible to evaluate whether or not a particular object is a member of the mixture. I distinguish the object language predicate Mix from the mixture operation Mɪx.

(80) $[\![\mathrm{Mix}(P,Q)(t)]\!] = 1$ iff $[\![t]\!] \in [\![\mathrm{Mix}(P,Q)]\!]$

(81) $[\![\mathrm{Mix}(P,Q)]\!] = \textsc{Mix}([\![P]\!],[\![Q]\!])$

Mixtures have been invoked in the analysis of conjunctions of plural nouns, particularly when they are quantified (Heycock and Zamparelli, 2000). For instance, in

[4]My terminology is most similar to Champollion's, who refers to these sets as *P/Q mixtures*. Heycock and Zamparelli (1999), with whom the concept originates, refer to the operation of mixing sets in this way as *Set Product*.

(82), there is a reading in which five people came, some of which were men and some of which were women (Champollion, 2015).

(82) Five men and women came

For this reading to be possible, 'men and women' must contain in its denotation sum individuals of cardinality five with parts from the set of men and parts from the set of women. A mixture of men and women is just the sort of set needed for a translation of (82), such as in (83).

(83) $\exists X[|X| = 5 \wedge \text{Mix}(M,W)(X) \wedge {}^*\text{Came}(X)]$

In previous applications, mixtures are generally *exclusive*: the sets used to form the mixture are not themselves subsets of the mixture.[5] For the analyses I develop for m-reduplication below, mixtures that include at least one of (the algebraic closure of) the sets in the mixture, as well as ones that include (the algebraic closure of) both sets. I refer to the former variety of mixture as a *partially inclusive* mixture, and to the latter variety as a *fully inclusive*, or simply *inclusive*, mixture.

(84) A mixture $\textsc{Mix}(P,Q)$ is partially inclusive iff $^*P \subset \textsc{Mix}(P,Q)$ or $^*Q \subset \textsc{Mix}(P,Q)$

(85) A mixture $\textsc{Mix}(P,Q)$ is fully inclusive iff $^*P \subset \textsc{Mix}(P,Q)$ and $^*Q \subset \textsc{Mix}(P,Q)$

The sort of mixture derived via the MIX operation depends on the relation between the two sets being mixed. If one set is a subset of the other, this will result in a partially inclusive mixture. If the two sets are identical, a fully inclusive mixture will result. It is also possible to define new mixture operations to derive these varieties of mixture. I will make use of both strategies to derive the required mixtures in the

[5]By the definition of mixture I've provided, mixtures are exclusive if the sets used to form them are disjoint.

analysis of m-reduplication below.

Alternatives and exhaustification

The final ingredient required for the analysis is a notion of alternatives and pragmatic enrichment. I will discuss the alternatives for m-reduplication in the analysis section, but for the sake of explicitness I will make use of an exhaustification operator, particularly *Exh* due to Fox (2007)[6] to derive scalar implicatures, defined as in (86).

(86) $Exh(\mathrm{A})(\mathrm{p}) = \mathrm{p}\ \&\ \forall \mathrm{q} \in \mathrm{IE}(\mathrm{A})(\mathrm{p})\text{: } \neg \mathrm{q}$

Essentially, what *Exh* does is negate all of the innocently excludable alternatives of p. An alternative is *innocently excludable* iff its negation does not contradict what p asserts.

This concludes the introduction of the formal tools to be used in this paper. In the next two sections, I apply the ingredients developed here to the analysis of m-reduplication. Due to the observed variation in speakers' judgments concerning the meaning of m-reduplication in non-upward-entailing environments and contexts involving speaker ignorance, I divide the analysis into two separate, but very closely related analyses, starting with the set of speakers I refer to as *the more restrictive speakers*, who do not permit m-reduplicated nominals to denote objects that are simply similar to the bare nominal, followed by the analysis for the ones I call *the less restrictive speakers*, whose denotation for m-reduplication does include atomic individuals in the proper similarity set of the bare nominal. Each of these analyses call for an elaboration of a basic approach to the calculation of the observed impli-

[6]See also the *O* operator of Chierchia (2004), Chierchia (2006), which is defined similarly.

cature.The first calls for calculation of the implicature at a subsentential level, while the second calls for an abstract alternative.

3.4.2 More restrictive speakers and local implicature calculation

For more restrictive speakers, I propose that m-reduplication denotes a mereological mixture of the predicate denoted by the bare noun and a set of objects similar to that bare noun, its similarity set. Given that m-reduplication only targets the head noun of an NP, (87) provides a structure for m-reduplication that treats RED as a categorizing head n, along the lines of an analysis in Distributed Morphology (Halle and Marantz (1993), Harley and Noyer (1999)), taking the root as an argument, and (88) gives a logical translation for the reduplication morpheme.

(87) [nP Root [n RED]]

(88) RED $\rightsquigarrow$ λP.λX.Mix(P,≃(P))(X)[7]

Because *P is a subset of *≃(P), the mixture of the two sets is a *partially* inclusive mixture: summing the elements of *P with the elements of the subset of *∼(P) equal

[7] This will technically put a function from individuals to truth values in the first argument of MIX, which, being a set of pairs of individuals and truth values, is not what MIX operates on. We need a function that takes this function into the set it characterizes. This can be accomplished by using a function like the one defined in Winter (2016), defined below.

(88) $f^{*} = \{x \mid f(x) = 1\}$

We can then slightly modify those expressions making reference to mixtures to reflect this move, writing instead Mix(dom(f),dom(≃(f))), the mixture of the underlying set of the characteristic function f with its similarity set. Where confusion does not arise, I will continue to write Mix(P,Q) throughout the paper.

to *P will produce that same set, due to the idempotence of the sum operation. It will also produce sum individuals that are composed of at least one member of *P and at least one member of the cumulative closure of P's proper similarity set, $*\sim$(P). Crucially, however, no member of $*\sim$(P) is a member of this mixture; this is because every member of that set has been summed with an element of *P. Mix(P,$\simeq$(P)) is therefore not a fully inclusive mixture. This exactly captures the meaning of m-reduplication for more restrictive speakers in non-upward-entailing and ignorance contexts: for these speakers, *ketâb metâb* could mean one or more books, or a sum of at least one book and at least one book-like thing, but not simply a single book-like thing or sum containing only book-like things.

A translation of (55), repeated as (89a), is given below in (89b). I assume existential closure of the variable in the translation of m-reduplication, which can be accomplished by Partee (1987)'s $\exists$ type shifter. Composition with the verb can take place either via additional type shifting or by just having verbs take type <<e,t>,t> arguments as in Montague (1973), among other options.

(89) a. Mohsen ketâb metâb xund
Mohsen book RED read.PST
'Mohsen read a book and other such things'

b. $\exists$X[Mix(Book,$\simeq$(Book))(X) $\wedge$ *Read(X)(m)]

(89b) is true if Mohsen read something, and that thing is in the mixture of Book and $\simeq$(Book), that is, it is either one or more books or at least one book summed with at least one bookish thing. In order to derive the exclusive mixture reading, we need to eliminate the singletons and sums composed of nothing but books.

To that end, I propose that the alternative to an m-reduplicated noun is its bare counterpart. This makes sense from a variety of standpoints. First of all, bare nouns in Persian receive number-neutral existential readings in object position, denoting a set of one or more things (Ghomeshi (2003); Jasbi (2015)).[8] In effect, they denote exactly the set of things we want to exclude.

(90) Mohsen ketâb xund
Mohsen book read.PST
'Mohsen read one or more books.'

Second, the bare noun is at least as complex as its m-reduplicated counterpart, derivable via either the deletion of the reduplication morpheme or substitution of the * operator in place of the reduplicant, and therefore is predicted to be available as an alternative by approaches such as Katzir (2007), which derive alternatives structurally. A possible tree for the bare nominal in which a head introducing the * operator is projected in the syntax is given in (91)[9].

(91) a.
nP
Root n

b. $* \rightsquigarrow \lambda P.\ \lambda x.{}^{*}P(x)$

Finally, the bare nominal, at least at first glance, appears to be *logically stronger* than the m-reduplicated nominal: because the denotation of *P is a subset of the set denoted by Mix(P,≃(P)), the former *entails* the latter in upward-entailing contexts. The bare nominal, then, seems to be a viable candidate as an alternative to m-reduplication.

Adopting a structural approach to deriving alternatives for the moment, we derive the

[8] Bare nouns often receive definite singular readings in subject position, but this is not a hard restriction; they can also receive a number-neutral reading as well, though definite readings may be more common.

[9] A simple alternative would be to introduce * lexically. This has no bearing on the availability of the bare noun as an alternative on the structural approach to alternatives

bare noun as an alternative via deletion of the node corresponding to reduplication in the syntax. Alternatively, if the * operator is present on a node in the syntax, we can derive the bare nominal via replacement of the node corresponding to the reduplicant with one containing *. The representation of this, which corresponds to the translation of (90), is given in (92).

(92) $\exists X[{}^*\text{Book}(X) \wedge {}^*\text{Read}(X)(m)]$

We can then apply *Exh* to the sentence in (55), resulting in (93).

(93) $\textit{Exh}(A)(55) = \exists X[\textsc{Mix}(\text{Book}, \sim(\text{Book}))(X) \wedge {}^*\text{Read}(X)(m)] \wedge \neg\exists X[{}^*\text{Book}(X) \wedge {}^*\text{Read}(X)(m)]$

Conjoining the negation of the alternative containing the bare nominal with the translation of the sentence containing the m-reduplicated nominal amounts to the statement that Mohsen read either a book, multiple books, or at least one book and one book-like thing, and Mohsen did not read one or more books, i.e. Mohsen read at least one book and at least one book-like thing. At first glance, it appears that we needed only one simple step to derive the exclusive mixture reading, or the non-homogeneous plural reading, from the partially inclusive mixture reading.

The astute reader will notice a problem with this account: technically speaking, (89b) is true whenever (92) is, but (92) is also true whenever (89b) is. This is certainly the case if (89a) is understood distributively, as made explicit in (94).

(94) $\exists X[\text{Mix}(\text{Book}, \simeq(\text{Book}))(X) \wedge \forall y[y <_{\text{AT}} X \rightarrow {}^*\text{Read}(y)(m)]]$

A similar problem is found in implicature analyses of English bare plurals, where the sentences with the bare plural and those with its desired alternative, a singular indefinite, entail each other (Spector (2007), Zweig (2009)).

In order to solve this problem, I follow Zweig (2009)[10] in, first, moving to a Neo-Davidsonian event semantics, and, second, calculating the implicature below the existential closure of the event variable, as originally proposed by Landman (2000) and developed by Chierchia (2004). For the first step, we need to provide Neo-Davidsonian translations for (89a) and (90) prior to existential closure of the event variable. These are given in (95) and (96), respectively.

(95) λe.$\exists$X[Mix(Book,$\simeq$(Book))(X) $\wedge$ *READ(e) $\wedge$ AGENT(e) = m $\wedge$ THEME(e) = X]

(96) λe.$\exists$X[*Book(X) $\wedge$ *READ(e) $\wedge$ AGENT(e) = m $\wedge$ THEME(e) = X]

Now we can discuss how implicatures are calculated in this system. As Zweig (2009) notes, because the expressions in competition with each other are *predicates*, not *propositions*, the scalar relation is not entailment, but *set containment*: a predicate A is stronger than a predicate B if A is a proper subset of B. The question we need to ask, then, is if (96) is stronger than (95) in this setup. We can do this as follows[11]. Suppose there are two events: e_1, an event of Mohsen reading a book (say, *Moby Dick*), and e_2, an event of Mohsen reading a magazine (say, the latest issue of *Time*). Additionally, imagine we have $e_3 = e_1 \oplus e_2$, the sum of e_1 and e_2. This is an event of Mohsen reading *Moby Dick* and *Time*. e_1 and e_3 are both in the set of events denoted by (95), as they are both events in which Mohsen reads something in the mixture of the set of books with its similarity set. However, of these events, only e_1 is in the set of events denoted by (96); this is because this set of events contains only those events with themes in the set of one or more books. e_1 meets this requirement,

[10] See Spector (2007) for an alternative solution making use of the notion of *higher-order implicature*.

[11] This scenario follows Zweig (2009)'s scenario for evaluating the scalar relationship between events with plural and singular arguments.

as its theme is the single book *Moby Dick*, but e_3 does not, as its theme is a sum individual composed of a book and a magazine. As such, there is a scalar relationship between (96) and (95): the former denotes a proper subset of the latter. As such, (96) is a stronger alternative of (95), as desired.

We are now in a position to calculate the implicature. Applying *Exh* to (95) results in the enrichment in (98), in which the alternative corresponding to (96) is negated[12].

(98) $Exh(\mathrm{A})(95) = \lambda e.\exists X[\mathrm{Mix}(\mathrm{Book},\simeq(\mathrm{Book}))(X) \wedge {}^{*}\textsc{read}(e) \wedge \textsc{agent}(e) = m \wedge \textsc{theme}(e) = X] \wedge \neg\exists X[{}^{*}\mathrm{Book}(X) \wedge {}^{*}\textsc{read}(e) \wedge \textsc{agent}(e) = m \wedge \textsc{theme}(e) = X]$

Existentially closing the event variable then leads to the final enriched meaning in (99).

(99) $\exists e[\exists X[\mathrm{Mix}(\mathrm{Book},\simeq(\mathrm{Book}))(X) \wedge {}^{*}\textsc{read}(e) \wedge \textsc{agent}(e) = m \wedge \textsc{theme}(e) = X] \wedge \neg\exists X[{}^{*}\mathrm{Book}(X) \wedge {}^{*}\textsc{read}(e) \wedge \textsc{agent}(e) = m \wedge \textsc{theme}(e) = X]]$

Note that calculating the implicature below the site of existential closure of the event variable results in the negation of the stronger alternative ultimately scoping below the event variable: (99) asserts that there is a reading event with Mohsen as agent and something in the mixture of books and book-like things as theme, and it is not the case that *this event* has something in the set of just books as its theme. That is, while subevents that are part of this event may have a (sum) individual book as theme, the event itself does not. Given this, we can conclude that the theme of the asserted event must be something in the mixture of books and book-like objects that is *not*

[12]This requires a suitable reformulation of *Exh* so that it can apply to predicates of events. Chierchia (2004) accomplishes this by generalizing all Boolean operators to functional types Partee and Rooth (1983). One possible formulation in this context is the following.

(97) $Exh(\mathrm{A})(\mathrm{P}) = \lambda e.\mathrm{P}(e) \wedge \forall Q \in \mathrm{IE}(\mathrm{A})(p){:}\ \neg\, Q(e)$

itself in the set of one or more books. This can only be a sum individual composed of at least one book and one object in the similarity set of books. This formulation therefore successfully derives the exclusive mixture reading of m-reduplication while avoiding the problem of calculating the implicature at the propositional level.

3.4.3 Less restrictive speakers and abstract alternatives

Recall that some speakers, while having the same interpretation of m-reduplication in upward-entailing contexts as the speakers discussed in the previous section, permit a more inclusive reading of m-reduplicated nominals in non-upward-entailing and ignorance contexts: for them, an m-reduplicated nominal may denote something that is just similar to a book, in addition to any number of books and book-book-like sums. These speakers, therefore, treat m-reduplicated nouns as *fully inclusive* mixtures, rather than partially inclusive ones as the more restrictive speakers do.
To generate fully inclusive mixtures, we could propose a variation on the original mixture function, an inclusive mixture function as in (100).

(100) Inclusive mixture
$[\![\text{I-M}\textsc{ix}(\text{P,Q})]\!] = \{\text{X} \oplus \text{Y} \mid \text{X} \in {}^*\text{P} \cup {}^*\text{Q},\ \text{y} \in {}^*\text{P} \cup {}^*\text{Q}\}$

This sums elements of the union of (the algebraic closure of) both sets with elements from the same set, guaranteeing the presence of both *P and *Q as subsets of the mixture. Using this, we could give the meaning of m-reduplication for less restrictive speakers as (101), which produces an inclusive mixture of P and its proper similarity set.

(101) $\textsc{red} \rightsquigarrow \lambda\text{P}.\lambda\text{X}.\text{I-Mix}(\text{P},\sim(\text{P}))(\text{X})$

This will produce a fully inclusive mixture. However, the reader will notice that, since the inclusive mixture function draws individuals from the union of both sets, that the inclusive mixture of these two sets is the same as mixing $\simeq$(P) with itself. We could therefore represent m-reduplication with (102).

(102) RED $\rightsquigarrow$ λP.λX.Mix($\simeq$(P),$\simeq$(P))(X)

We can take this further: it turns out that for any P, mixing P with itself is just *P. This can be immediately appreciated given another fact about mixtures, namely, that the inclusive mixture of two sets is equivalent to applying algebraic closure to the union of those sets. That is:

(103) Inclusive mixture as the algebraic closure of the union of two sets
I-MIX(P,Q) = *(P $\cup$ Q)

It immediately follows that I-MIX(P,P) = *P, since *(P $\cup$ P) = *P. Because of this equivalence, we can also define the other types of mixture in terms of *, $\cup$, and set difference.

(104) Exclusive mixture
E-MIX(P,Q) = *(P $\cup$ Q) $\setminus$ (*P $\cup$ *Q)

(105) Left-inclusive mixture
L-MIX(P,Q) = *(P $\cup$ Q) $\setminus$ *Q

(106) Right-inclusive mixture
R-MIX(P,Q) = *(P $\cup$ Q) $\setminus$ *P

Given that this is the case, we could simply translate the reduplicative morpheme as in (107).

(107) RED $\rightsquigarrow$ λP.λX.*$\simeq$(P)(X)

In other words, m-reduplication for less restrictive speakers simply denotes the algebraic closure of the similarity set of the predicate it takes as an argument. We can therefore provide the basic, unenriched translation of (55) for less restrictive speakers as (108).

(108) $\lambda e.\exists X[{}^{*}{\simeq}(\text{Book})(X) \wedge {}^{*}\text{Read}(e) \wedge \textsc{agent} = m \wedge \textsc{theme} = X]$

This captures the unenriched meaning of m-reduplication for less restrictive speakers, but it raises an issue for the derivation of the exclusive mixture reading via scalar implicature. In particular, while we are still able to derive the bare nominal as an alternative via the same process of structural deletion/replacement, it is no longer enough to generate the exclusive mixture reading; if the bare noun is the only alternative to m-reduplication, we instead predict a partially inclusive reading in upward-entailing contexts without speaker ignorance, one in which m-reduplication refers to a set of either book-bookish sums or to one or more book-like things, but not books. This is incorrect, as the two groups of speakers agree on the interpretation of m-reduplication in upward-entailing contexts as an exclusive mixture.

In order to derive an exclusive mixture reading from the fully inclusive reading, two alternatives are required: one corresponding to the bare noun, *P, and the other corresponding to the set of things similar but not identical to the bare noun. This is the *proper similarity set*, *∼(P). This alternative is *abstract*: it does not correspond to any lexical item in the language. It therefore cannot be derived via lexical replacements or deletions, as a structural approach to the generation of alternatives would require.

This point is worth elaborating on. Suppose we want to derive an appropriate set of alternatives from a syntactic tree containing as a subtree (87), repeated in (109) immediately below for convenience.

(109) Subtree corresponding to *ketâb metâbb*

nP

$\sqrt{\text{BOOK}}$ n

RED

One strategy for deriving alternatives would replace the head containing RED morpheme with one containing the * operator, deriving the necessary alternative denoting sets of one or more books (110)[13]. Additionally, one may derive alternatives containing other subsets of the similarity set by replacing the root $\sqrt{\text{BOOK}}$ with other roots in the lexicon. (111) shows the result of replacing $\sqrt{\text{BOOK}}$ with $\sqrt{\text{MAGAZINE}}$.

(110) Subtree derived from (109) via replacement of RED with *

nP

$\sqrt{\text{BOOK}}$ n

*

(111) Subtree derived from (109) via replacement of RED with * and $\sqrt{\text{BOOK}}$ with $\sqrt{\text{MAGAZINE}}$

nP

$\sqrt{\text{MAGAZINE}}$ n

*

While each of these sets is a subset of the similarity set of *Book*, they are not enough to generate the exclusive mixture reading. This is because the similarity set also contains *sums* of the atoms in each of these sets: sum individuals consisting of books and magazines, magazines and newspapers, newspapers and comic books, and so on. Since expressions like those in (110) and (111) do not denote sets containing such mixed sums, they cannot be the only alternatives to an m-reduplicated nominal. We seem to also need *mixtures* of the appropriate alternatives. An obvious way to

[13]On an approach that does not project * in the syntax, this could be treated by deleting the node containing RED or by replacing it with one that is semantically vacuous.

attempt to derive these is by merely replacing the root with contextually appropriate lexical alternatives while maintaining the RED morpheme in the syntactic structure. (112) shows one possible result of this type of replacement.

(112) Subtree derived from (109) via replacement of $\sqrt{\text{BOOK}}$ with $\sqrt{\text{MAGAZINE}}$

nP

$\sqrt{\text{MAGAZINE}}$ n

RED

This will allow the generation of mixtures, which may contain the sorts of sum individuals that we would like to exclude, such as sums of magazines and newspapers that do not contain books. Unfortunately, however, this does not solve the problem. The reason has to do with the nature of the similarity relation, namely, the fact that it is *symmetric*: $\text{P} \sim \text{Q} \iff \text{P} \sim \text{Q}$. Therefore, if magazines count as similar to books, then books count as similar to magazines. This further implies that if the set of magazines is in the similarity set of books, then the set of books is in the similarity set of magazines. Given this fact, the set denoted by the expression in the subtree in (112) will contain sum individuals composed of books and magazines, the sorts of individuals we do *not* want to exclude in our analysis. At best, this will simply result in incorrect predictions: excluding alternatives to an m-reduplicated nominal generated by substituting the root with roots that are similar in the context will result in unattested interpretations of m-reduplication in upward-entailing contexts. At worst, these alternatives won't be innocently excludable in the first place, as excluding them may result in contradicting the assertion of a sentence containing an m-reduplicated nominal. Their presence would therefore play no role in the calculation of the desired implicature.

The only option that guarantees the derivation of the exclusive mixture reading from the fully inclusive one is making use of the alternative corresponding to the property similarity set, which, as discussed above, is an abstract alternative. Because this alternative does not correspond to a lexical item in the language, it will not be possible to derive the alternative using the structural replacement/deletion approach sketched above.

There has been some work suggesting that abstract alternatives are possible, and even necessary for the analysis of some phenomena (Chemla (2007); Buccola et al. (2018); Charlow (2019)). For example, Chemla (2007) and Buccola et al. (2018) discuss the fact that the French sentence in (113), like its English counterpart, is odd, despite the fact that French, unlike English, lacks a word corresponding to *both*, competition with which would explain the oddness of the sentence in English.

(113) #Jean s'est cassé tous les bras
Jean REFL=be.PRS.3.SG broken all DEF.PL arms
'#Jean broke all his arms.' (implies: Jean has more than two arms

Perhaps more strikingly, Charlow (2019) discusses the case of exceptionally scoping indefinites, as in (114), in which an indefinite takes widest scope despite being deeply embedded within another phrase.

(114) John overheard the rumor that *a student of mine* was expelled

As with other indefinites, exceptionally scoping indefinites are associated with an implicature: (114) means that not all of my students were expelled. The problem is that, on either a choice-functional (Reinhart, 1997b) or an alternative semantics (Kratzer and Shimoyama, 2002) approach to exceptional scope, the required alternative, a universal quantification over choice functions or alternatives, does not correspond to any overt or covert lexical item in English. This alternative, therefore, is abstract.

Let us return to the m-reduplication case at hand. Although the presence of an abstract alternative is not completely unprecedented, it is not entirely clear how to derive such alternatives. One could merely stipulate that the proper similarity set of an expression is lexically associated with m-reduplicated nominals via a Horn scale. A more principled approach is suggested by Buccola et al. (2018): alternatives are calculated on the basis of the *conceptual representation* of an expression. This is accomplished by operations on representations in the *language of thought* (Fodor, 1975), making use of a notion of conceptual complexity, rather than structural complexity as argued by Katzir (2007). For the sake of concreteness, I suggest that the alternatives for m-reduplication be derived from the representation in the *logical language*, which in this context serves as a proxy for the language of thought.[14] If this is the case, then despite the semantic equivalence of (101), (102), and (107), for the purpose of deriving alternatives from the conceptual representation of the expression, (101) is to be preferred. In this case, the representation of (55) for less restrictive speakers is as in (115).

(115) λe.$\exists$X[I-Mix(Book, $\sim$(Book))(X) $\wedge$ *Read(e) $\wedge$ AGENT = m $\wedge$ THEME = X]

Because *Book and *$\sim$(Book) are part of the representation of the mixture in the logical language, and, more generally, components in the complex concept of their mixture, we can derive sentences containing them as alternatives to (115). Upon applying *Exh* to (115) and existentially closing the event variable, we can now derive (116).

(116) $\exists$e$\exists$X[I-Mix(Book, $\sim$(Book))(X) $\wedge$ *Read(e) $\wedge$ AGENT = m $\wedge$ THEME = X] $\wedge$ $\neg\exists$X[*Book(X) $\wedge$ *Read(e) $\wedge$ AGENT = m $\wedge$ THEME = X] $\wedge$

[14]I return to the issue of constraining alternatives derived from a non-linguistic level of representation in the final chapter of this dissertation.

$$\neg\exists X[{}^*{\sim}(\text{Book})(X) \wedge {}^*\text{Read}(e) \wedge \textsc{agent} = m \wedge \textsc{theme} = X]$$

(116) means that there is an event of Mohsen reading something made up of either one or more books, one or more book-like things, or sums of books and book-like things, and it is not the case that he read one or more books and it is not the case that he read one or more book-like things. This is equivalent to saying that Mohsen read a sum of books and book-like things. This is equivalent to the non-homogeneous plural interpretation, and can be represented as an exclusive mixture of the set of books and its proper similarity set, given in (117).

(117) $\exists e \exists X[\text{E-Mix}(\text{Book}, {\sim}(\text{Book}))(X) \wedge {}^*\text{Read}(e) \wedge \textsc{agent} = m \wedge \textsc{theme} = X]]$

We have thus derived the exclusive mixture reading of m-reduplication for less restrictive speakers.

A brief look back at the more restrictive speakers

Recall that for the analysis of more restrictive speakers, I used a structural approach to deriving the alternative to m-reduplication. However, for less restrictive speakers, I used a conceptualist approach to deriving alternatives. We therefore have two entirely distinct analyses of the computation of the implicature for each set of speakers. Ideally, we would prefer to be able to unify the two analyses using the same mechanism for deriving alternatives for both groups of speakers.

My proposal is to adopt the conceptualist approach for the more restrictive speakers as well. Recall that for those speakers, m-reduplication denotes a partially inclusive mixture, Mix(P,$\simeq$(P)). This means that we will have two alternatives, *P and *$\simeq$(P). However, only one of these alternatives is innocently excludable: *P. If we were to negate an alternative containing *$\simeq$(P), we would contradict the main assertion, because the partially inclusive mixture denoted by m-reduplication is a subset of the inclusive mixture denoted by $^*\simeq(P)$. $^*P$, on the other hand, corresponds to the bare noun, which denotes a subset of the partially inclusive mixture, and is therefore logically stronger and thus innocently excludable. Negating this alternative will deliver the exclusive mixture reading, as desired.

The takeaway here is that a unified approach to m-reduplication for both sets of speakers is possible, one making exclusive use of a conceptualist approach to alternatives. Fitting this with the notion of innocent exclusion[15], we guarantee that the enrichment of the underlying meaning ends up making use of the same alternative as the previously proposed structural approach.

3.5 Prediction of the analysis: Dependent readings of m-reduplicated nominals

Recall that I have adopted a local implicature analysis of the exclusive mixture reading of m-reduplication for both sets of speakers. While this is motivated for the more restrictive speakers, due to the fact that global calculation of the implicature is incapable of deriving the correct reading for these speakers. However, this is technically not required for the analysis of the less restrictive speakers' judgments. This is because m-reduplication denotes an inclusive mixture for these speakers, and therefore the alternatives to m-reduplication in this context are strictly stronger than the unenriched interpretation of m-reduplication. Certainly, one motivation for analyzing both sets of speakers' judgments as involving local implicature calculation is simple uniformity of analysis, and nothing is lost by analyzing the more restrictive speakers using local implicatures. That said, it would be ideal to find an additional empirical motivation for using local implicature calculation for both groups of speakers, where an alternative that calculates the implicature globally for less restrictive speakers

[15]Or monotonicity more generally.

would fail.

This evidence comes from *dependent readings* of plural expressions (Chomsky (1975); de Mey (1981); Partee (1985); Zweig (2009)). As noted previously, English bare plurals are associated with a multiplicity inference in upward-entailing contexts. However, in contexts involving quantified subjects, a bare plural may be interpreted as distributed across the individuals associated with the quantified subject. Consider the example in (118).

(118) Three boys flew kites

In this sentence, it need not be the case that each of the three boys flew more than one kite. Rather, (118) is true in a context in which each boy flew only one kite, as long as more than one kite was flown overall: it would be false if the boys shared a single kite and took turns flying it. This is the *dependent reading* of the bare plural. In order to derive dependent readings of English bare plurals, Zweig (2009) makes use of the local implicature mechanism developed above. Essentially, (118) receives the following unenriched translation, prior to existential closure of the event variable. Capitalized variables range over both sums and atoms, while lowercase variables range only over atoms.

(119) $\lambda e.\exists X \exists Y[{}^*\textsc{boy}(X) \wedge |X| = 3 \wedge {}^*\textsc{kite}(Y) \wedge \textsc{fly}(e) \wedge \textsc{ag}(e) = X \wedge \textsc{th}(e) = Y]$

This event description competes with an alternative containing a singular rather than a bare plural. The English sentence corresponding to this alternative is given in (120), and its translation is given in (121).

(120) Three boys flew a kite

(121) $\lambda e.\exists X \exists y[{}^*\textsc{boy}(X) \wedge |X| = 3 \wedge \textsc{kite}(y) \wedge \textsc{fly}(e) \wedge \textsc{ag}(e) = X \wedge \textsc{th}(e) = y]$

The event description in (121) is stronger than that in (119): (121) is a set of events with an atomic kite as theme, while (119) is a set of events with atomic or sum kites as theme. (121) is therefore a subset of (119), and can be innocently excluded. Doing so leads to strenghtening of (119), and existential closure of the event variable gives the final truth conditions in (122).

(122) $\exists e.\exists X\exists Y[{}^*\textsc{boy}(X) \wedge |X| = 3 \wedge {}^*\textsc{kite}(Y) \wedge |Y| > 1 \wedge \textsc{fly}(e) \wedge \textsc{ag}(e) = X \wedge \textsc{th}(e) = Y]$

This formula means that there is a kite flying event with three boys as agent and more than one kite as theme. However, it does *not* require that each of the three boys flies more than one kite: rather, the subevents of this event may involve only an atomic boy as agent and an atomic kite as theme. (122) is true as long as there are multiple kites flown overall. This is exactly the dependent plural reading.

Given that I have made use of the local implicature calculation system for both sets of Persian speakers, we would expect that m-reduplicated nominals may also show dependent readings. Moreover, we predict that, since both sets of speakers agree on the interpretation of m-reduplication in upward-entailing contexts, that both sets of speakers will have the same dependent reading of m-reduplication.

This prediction is borne out: dependent readings do arise with m-reduplicated nominals. Consider the sentence in (123).

(123) CONTEXT: three children, Ali, Simin, and Forough, are eating fruits. Ali eats an apple, Simin eats an orange, and Forough eats a pear.

Se	tâ	bacche	sib	mib	xord	-an
three	CL	child	apple	RED	eat.PST	-3.PL

'Three children ate apples and the like.'

Both sets of speakers agree that (123) is true in the context provided. As one can see, then, there is no requirement that each child eat both an apple and something else for (123) to be true. What is more, both sets of speakers agree that (123) is infelicitous if each child ate only one or more apples or just something similar to an apple[16]. This, therefore, is the m-reduplication analogue to dependent readings of bare plurals in English: (123) is true as long as at least one apple is eaten by one of the children, and more than one kind of fruit is eaten overall, but there is no requirement that every child eat an apple and at least one other kind of fruit.

3.5.1 Dependent readings with more restrictive speakers

Let us see how the approach making use of local implicature calculation derives the dependent reading of m-reduplication for both sets of speakers. We will start, as before, with the more restrictive speakers. Recall that for these speakers, m-reduplication denotes a partial mixture of a set with its similarity set. The translation of the reduplicant for these speakers was given in (88), which I repeat in (124) below for convenience.

(124) $\textsc{red} \rightsquigarrow \lambda\mathrm{P}.\lambda\mathrm{X}.\mathrm{Mix}(\mathrm{P},\simeq(\mathrm{P}))(\mathrm{X})$

Following Landman (2000) and Zweig (2009) among others, I assume that weak quantificational DPs like "three boys" take scope below existential closure of the event variable by default[17]. A translation of (123) prior to closure of the event

[16]Assuming of course that the context does not involve speaker ignorance or indifference.

[17]Technically, it makes no *semantic* difference whether the existential quantification over events is scoped above or below those over individuals: since they are all existentially quantified, the interpretations will be equivalent. For the present purposes, scoping the existential quantifiers over individuals below the site of existential closure of the event variable matters is crucial for the calculation of the right implicature.
Note that scope for numerals *does* matter if we invoke an additional scope mechanism, such as

variable for more restrictive speakers is given in (125).

(125) λe.$\exists$X$\exists$Y[*Child(X) $\wedge$ |X| = 3 $\wedge$ Mix(Apple,$\simeq$(Apple))(Y) $\wedge$ *EAT(e) $\wedge$ AGENT(e) = X $\wedge$ THEME(e) = Y]

More restrictive speakers have as the sole innocently excludable alternative to an m-reduplicated nominal its non-reduplicated bare counterpart. In this case, the alternative to *sib mib* is simply the predicate corresponding to *sib* 'apple'. This is composed into the alternative event description in (126).

(126) λe.$\exists$X$\exists$Y[*Child(X) $\wedge$ |X| = 3 $\wedge$ *Apple(Y) $\wedge$ *EAT(e) $\wedge$ AGENT(e) = X $\wedge$ THEME(e) = Y]

Excluding (126) amounts to negating the part of the expression in (126) below the abstraction over events and conjoining it below the site of existential closure of the event in (125). Once the event variable is closed, this delivers the enriched formula in (127).

(127) $\exists$e[$\exists$X$\exists$Y[*Child(X) $\wedge$ |X| = 3 $\wedge$ Mix(Apple,$\simeq$(Apple))(Y) $\wedge$ *EAT(e) $\wedge$ AGENT(e) = X $\wedge$ THEME(e) = Y] $\wedge$ $\neg\exists$X$\exists$Y[*Child(X) $\wedge$ |X| = 3 $\wedge$ *Apple(Y) $\wedge$ *EAT(e) $\wedge$ AGENT(e) = X $\wedge$ THEME(e) = Y]]

Let's break this down: (127) is true iff there is an event of eating with three children as agent and something in partial mixture of apples and applie-like things as theme, and *this event* is not an event that with all three of the following properties: 1. being an event of eating, 2. being an event with three children as agent, and 3. being an

via a distributivity operator. This is accomplished by scoping one of the phrases containing a numeral over a universal quantification over the parts of the variable existentially quantified over, and interpreting the other quantified expressions within the scope of this universal quantifier. This will require that each numerical expression be evaluated with respect to each part of the expression to which the distributivity operator applies.

event with something in the set of one or more apples as theme. Since we know, by virtue of the first few conjuncts, that this event *is* an eating event with three children as agent, exclusion of the alternative amounts to denying that the theme of the event is in the set of one or more apples. We can therefore rewrite (127) as (128).

(128) $\exists e[\exists X \exists Y[{}^{*}\text{Child}(X) \wedge |X| = 3 \wedge \text{Mix}(\text{Apple},\simeq(\text{Apple}))(Y) \wedge {}^{*}\textsc{eat}(e) \wedge \textsc{agent}(e) = X \wedge \textsc{theme}(e) = Y] \wedge \neg\exists Y[{}^{*}\text{Apple}(Y) \wedge \textsc{theme}(e) = Y]]$

Here, we can see that three children are the agent of an eating event with something in the partial mixture of apples and apple-like objects as theme, but that this object is *not* an atomic or sum individual in the set of apples. We must therefore conclude that the theme of this event is in the *exclusive mixture* of apples and apple-like things: it must be a sum individual containing at least one apple and one thing similar to an apple. Therefore, (128) is equivalent to (129), which provides a final representation for the analysis.

(129) $\exists e \exists X \exists Y[{}^{*}\text{Child}(X) \wedge |X| = 3 \wedge \text{E-Mix}(\text{Apple},\sim(\text{Apple}))(Y) \wedge {}^{*}\textsc{eat}(e) \wedge \textsc{agent}(e) = X \wedge \textsc{theme}(e) = Y]$

Let us now consider the truth conditions of (129). This formula is true iff there is an eating event with three children as agent and a sum of apples and apple like fruits as theme. Much like the English dependent plural example in (122), (129) does not require that each of the three children eat an apple and some other fruit; rather, the subevents of the event whose existence is asserted may have a single child as agent and either an apple, orange, or pear as theme. It requires only that an apple and some other fruit be eaten overall. This is exactly the dependent reading we want to derive!

3.5.2 Dependent readings with less restrictive speakers

Having derived the dependent reading of m-reduplication for more restrictive speakers, let us turn now to deriving this reading for less restrictive speakers. Recall that these speakers treat m-reduplicated nominals as fully inclusive mixtures prior to enrichment, but assign them the same exclusive reading in upward-entailing contexts and get the same dependent reading that more restrictive speakers do. The translation of the reduplicant for less restrictive speakers was given above in (101), repeated for convenience in (130)

(130) $\textsc{red} \rightsquigarrow \lambda$P.$\lambda$X.I-Mix(P,$\sim$(P))(X)

The unenriched meaning of (123) for these speakers is then (131). Once again, we leave the event variable open.

(131) λe.$\exists$X$\exists$Y[*Child(X) $\wedge$ |X| = 3 $\wedge$ I-Mix(Apple,$\sim$(Apple))(Y) $\wedge$ *$\textsc{eat}$(e) $\wedge$ $\textsc{agent}$(e) = X $\wedge$ $\textsc{theme}$(e) = Y]

Recall now that less restrictive speakers have two alternatives for m-reduplication: the alternative corresponding to the non-reduplicated bare noun, and the one corresponding to the proper similarity set to the set denoted by the bare nominal. The event descriptions derived by replacing the inclusive mixture in the formula in (131) with the bare nominal and the proper similarity set are given in (132) and (133), respectively.

(132) λe.$\exists$X$\exists$Y[*Child(X) $\wedge$ |X| = 3 $\wedge$ *Apple(Y) $\wedge$ *$\textsc{eat}$(e) $\wedge$ $\textsc{agent}$(e) = X $\wedge$ $\textsc{theme}$(e) = Y]

(133) λe.$\exists$X$\exists$Y[*Child(X) $\wedge$ |X| = 3 $\wedge$ *$\sim$(Apple)(Y) $\wedge$ *$\textsc{eat}$(e) $\wedge$ $\textsc{agent}$(e) = X $\wedge$ $\textsc{theme}$(e) = Y]

As before, we exclude these alternatives by conjoining their negations with the event description in (131), producing a new event description. Applying existential closure to this event description results in the enriched meaning of (123) for less restrictive speakers, and is given (134).

(134) $\exists$e[$\exists$X$\exists$Y[*Child(X) $\wedge$ |X| = 3 $\wedge$ I-Mix(Apple,$\sim$(Apple))(Y) $\wedge$ *EAT(e) $\wedge$ AGENT(e) = X $\wedge$ THEME(e) = Y] $\wedge$ $\neg\exists$X$\exists$Y[*Child(X) $\wedge$ |X| = 3 $\wedge$ *Apple(Y) $\wedge$ *EAT(e) $\wedge$ AGENT(e) = X $\wedge$ THEME(e) = Y] $\wedge$ $\neg\exists$X$\exists$Y[*Child(X) $\wedge$ |X| = 3 $\wedge$ *$\sim$(Apple)(Y) $\wedge$ *EAT(e) $\wedge$ AGENT(e) = X $\wedge$ THEME(e) = Y]]

As with (127), the event variable in each of the conjuncts corresponding to the unenriched assertion and negated alternatives is bound by the same existential quantifier. This means that innocent exclusion of the alternatives is evaluated with respect to a single event. We can therefore breakdown (134) as follows: there is an event of eating with three children as agent and something in the *inclusive mixture* of one or more apples and the proper similarity set of apples as theme. Furthermore, just as with (127), this event is not one with all three of the following properties: 1. being an eating event, 2. being an event with three children as agent, and 3. being an event with one or more apples as theme. What's more, since the proper similarity set is also an excluded alternative for less restrictive speakers, this is also not an event with all three of the following properties: 1. being an eating event, 2. being an event with three children as agent, and 3. being an event with one or more objects in the proper similarity set of apples. As before, because this event *is* an eating event with three children as theme, innocent exclusion of the alternatives amounts to denying only that the event in question has an element of the inclusive mixture of apples and the proper similarity set of apples as theme, but that this is object is neither an atomic apple nor a sum consisting only of one or more apples, nor is this object an atom or sum individual in the cumulative closure of the proper similarity set of apples. All of this means that (134) is equivalent to (135).

(135) $\exists e[\exists X \exists Y[{}^*\text{Child}(X) \wedge |X| = 3 \wedge \text{I-Mix}(\text{Apple},\sim(\text{Apple}))(Y) \wedge {}^*\textsc{eat}(e) \wedge \textsc{agent}(e) = X \wedge \textsc{theme}(e) = Y] \wedge \neg\exists Y[{}^*\text{Apple}(Y) \wedge \textsc{theme}(e) = Y] \wedge \neg\exists Y[{}^*{\sim}(\text{Apple})(Y) \wedge \textsc{theme}(e) = Y]]$

Because (135) requires that the theme of the event be something in the inclusive mixture of apples and the proper similarity set of apples, but denies that this object is a sum of one or more apples or a sum of one or more objects in the cumulative closure of the proper similarity set of apples, this leaves us with only one possibility: the theme of this event must be a sum individual containing at least one apple and at least one thing from the proper similarity set of apples. In other words, the theme of the event in question is in the exclusive mixture of Apple and ∼(Apple). We can therefore provide a final representation as (136), which is exactly the same as (129) above.

(136) $\exists e \exists X \exists Y[{}^*\text{Child}(X) \wedge |X| = 3 \wedge \text{E-Mix}(\text{Apple},\sim(\text{Apple}))(Y) \wedge {}^*\textsc{eat}(e) \wedge \textsc{agent}(e) = X \wedge \textsc{theme}(e) = Y]$

Since we are dealing with the exact same formula as in (129), it is no surprise that (136) is true in a context in which three children ate some kind of apple-like fruit, that it is not necessary that each child eat an apple in addition to another kind of fruit, and that the only requirement is that an apple and some other kind of fruit similar to an apple be eaten overall. In other words, we have successfully derived the exact same dependent reading for less restrictive speakers that we had previously derived for more restrictive speakers!

3.6 Explaining other properties of m-reduplicated nominals

The current analysis of m-reduplication does not posit that m-reduplicated nominals denote exclusive mixtures. As such, it does not treat such nominals as referring exclusively to sum individuals. In this section, I demonstrate that the analysis nevertheless still accounts for the original observations motivating treating m-reduplication as involving plural reference in the first place.

First, recall that m-reduplicated nominals are compatible with collective predicates, as in (59), repeated below as (137).

(137) Mohsen gol mol jam' kard
Mohsen flower RED collection do.PST
'Mohsen collected flowers and other such things.'

Of course, collective predicates like *jam' kardan* 'collect' only contain sum individuals in their denotation. M-reduplicated nominals, just like bare plurals in English, do contain sum individuals in their denotation, so we expect them to be acceptable with predicates like *jam' kardan*.

The issue is more complex when we consider the unacceptability of m-reduplication with differential object marking without an additional plural marker.

(138) *Mohsen ketâb metâb -ro xund
Mohsen book RED -DOM read.PST
Intended: 'Mohsen read the one book and bookish thing.'

As noted previously, this receives a straightforward explanation if m-reduplication simply denotes a set of sum individuals, since without the plural marker differential object marking forces a singular interpretation. However, I have already demonstrated that m-reduplication denotes a partial or fully inclusive mixture. As such, the unacceptability of (138) appears at first glance to be a bit of a puzzle.

This puzzle can be resolved by being more specific about the contribution of differential object marking in Persian. Building on an approach due to Jasbi (2015), we could treat the syntactic structure of a *-ro*-marked object as containing a type-shifter, Partee (1987)'s *iota*, which takes a predicate and returns the unique individual satisfying that predicate. The definition is given in (139).

(139) $\textsc{iota}(\mathrm{P}) = \iota \mathrm{x}[\mathrm{P}(\mathrm{x})]$

IOTA also introduces a presupposition that the set denoted by the predicate to which it applies is *singleton*: there can be only one member of the set. Let us consider how this would work with a bare noun like *ketâb*. Assume that bare nouns in Persian denote the cumulative closure of a set. As such, *ketâb* would denote the set of all possible sums of books.

(140) $[\![\textit{ketâb}]\!] = {}^*\{\mathrm{x} \mid \mathrm{x} \text{ is a book}\}$

The cumulative closure operator * can also apply to singleton sets. In this case, the * operator returns the same set. In other words, a singleton set is its own cumulative closure. The reason for this is that, as the set of all sums formable from a set, the cumulative closure of a singleton set can only sum the sole member of that set with itself. Because of the idempotence of the sum operation $\oplus$, this simply returns the sole member of the set.

(141) ${}^*\{\mathrm{x}\} = \{\mathrm{x}\}$, for any x.

Given this fact, the only context in which a bare noun in Persian will be compatible with *-ro* is if it is singleton. This successfully predicts that *-ro*-marked bare nouns are acceptable only if they denote singleton sets.

Now consider the case of an m-reduplicated noun, which denotes a mixture of a set with its similarity set. Assume a context in which there is a single book and a single magazine, the latter of which is in the similarity set of book[18]. The mixture of these two sets is given in (**??**).

(142) MIX({b},{m}) = {b, m, b⊕m}

As one can see, even a mixture of two singleton sets is necessarily non-singleton. Since m-reduplicated nominals denote such mixtures, they will always denote non-singleton sets, and will therefore be uncompatible with *-ro* without further modification. Now consider an example of m-reduplication with *-hâ*, which is completely acceptable with *-ro*.

(143)

Mohsen	ketâb	metâb	-hâ	-ro	xund
Mohsen	book	RED	-SP.PL	-DOM	read.PST

'Mohsen read the books and stuff.'

Just as we were more specific about the contribution of *-ro*, we can account for this case by proposing an explicit semantics for *-hâ*. I propose, building on a proposal due to Jasbi (2014), that *-hâ* reduces a set to its maximal member.

(144) *-hâ* $\rightsquigarrow$ $\lambda P.\lambda x.x = \oplus P$

One could rightfully ask why I don't simply define P-hâ as an individual, the maximal member of *P. There are two reasons. First, I want to maintain a uniform analysis of -*râ*-marked NPs, including those marked by *-hâ*. More importantly, though, it appears empirically necessary, as NPs marked by *-hâ* can be quantified, as the following shows (Ghomeshi, 2003).

[18]One could in principle have a context in which there is only a single book. Here, though, it would be odd to use m-reduplication, as the m-reduplicated nominal would mean the same thing as the bare nominal. We could then appeal to a pragmatic principle that requires that a more specific form, in this case the bare nominal, be used when it is equivalent to the m-reduplicated one. Alternatively, we could place a definedness condition on m-reduplication, stating that the result of m-reduplication is only defined when there is at least one member of the similarity set that is not equal to the members of the set under evaluation.

(145) ye ketâb -hâ -i
a book -SP.PL -INDEF
'Some books'

In the case of (143) above, the addition of *-hâ* reduces the set to its single maximal member. Because this set is now a singleton, though containing a single sum individual, it satisfies the uniqueness requirement of IOTA, and is thus successfully predicted to be acceptable with *-ro*.

3.7 Conclusion

In this chapter, I've developed an analysis of non-homogeneous plural inferences in m-reduplication as scalar implicature, based on an analysis of the unenriched meaning of similative expressions as mereological mixtures of sets with their similarity sets. The analysis explains the behavior of m-reduplication in upward-entailing as well as non-upward-entailing and ignorance contexts, and their interaction with quantificational expressions, collective predication, and other properties of the grammars of the languages under discussion, while also explaining interspeaker variation in the unenriched meaning of similative plurals.

The analysis makes several empirical and theoretical contributions. Empirically, it presents a formal treatment of the properties of similative plurals, which has not been done previously. It further reveals connections with the theory of bare plurals, such as those found in English, and thereby facilitates comparison with other types of plurals cross-linguistically. Theoretically, the analysis has implications for the theory of implicature calculation and of alternatives. In particular, I have argued that the derivation of the non-homogeneous plural, or exclusive mixture interpretation of similative plurals from their underlying (partially) inclusive mixture semantics requires the use of both local implicature calculation (Landman (2000); Chierchia (2004); Zweig (2009)) and abstract alternatives (Chemla (2007); Buccola et al. (2018); Charlow (2019)). Consequently, the phenomena discussed here prove problematic for structural approaches to alternatives (Katzir, 2007).

It is important to note that although the analysis casts doubt on the *structural* theory of *alternatives*, it does not exclude a *grammatical* approach to *implicatures*. On such theories, implicatures are derived by semantic operators in the syntax, rather than via global pragmatic inferences operating on an unenriched semantic representation. Such theories are compatible with a variety of approaches to the derivation of alternatives, including the conceptualist approach advocated for here. Likewise, different approaches to the nature of alternatives are compatible with a range of approaches to the calculation of implicatures. The conceptualist approach I adopted here was combined with a grammatical approach to implicatures, but the two ideas are logically independent.

This concludes the analysis of m-reduplication in Persian. In the next chapter, I turn my attention to similative plurals in Japanese, focusing on *-toka* and *-tari*, making use of the tools developed here and in the previous chapter and extending them where necessary.

CHAPTER 4

SIMILATIVE PLURALS IN JAPANESE

This chapter provides the second case study on similative plurality, examining the semantic and pragmatic behavior of the morphemes *-toka* and *-tari* in Japanese. Empirically, it is found that *-toka* and *-tari* behave very similarly to Persian m-reduplication: they exhibit the same sensitivity to the monotonicity of the environment in which they are embedded, display exclusive readings in upward-entailing contexts but inclusive readings in downward-entailing and non-monotone contexts. Also like m-reduplication, they permit inclusive readings when speaker ignorance is established in upward-entailing contexts. They also exhibit the same kind of interspeaker variation as was found with m-reduplication among Persian speakers: there is a set of more restrictive speakers and a set of speakers with less restrictive judgments. *-toka* and *-tari* do differ from m-reduplication in two ways, however: first, the complements of *-toka* and *-tari* can be far more syntactically complex than that of m-reduplication. Second, unlike m-reduplication, *-toka* and *-tari* allow restrictions on their domains to be overtly expressed. I propose an analysis of these expressions that is very similar to that proposed for m-reduplication in the previous chapter, with modifications to that analysis to allow for a treatment of properties specific to *-toka* and *-tari* that m-reduplication does not display.

The chapter is structured as follows. First, I provide background on the Japanese language, focusing on its grammatical properties at a descriptive level. Second, I introduce *-toka* and *-tari* and discuss their behavior in upward-entailing contexts. I also discuss grammatical differences between these morphemes and m-reduplication, revealing the greater syntactic complexity of their complements as compared to m-reduplication, their compatibility with an overt domain restriction, and their ability to be coordinated. Third, I reveal the sensitivity of the interpretation of m-reduplication to the direction of entailment, showing that it possesses inclusive readings in downward-entailing and question contexts. I further demonstrate that establishing speaker ignorance is enough to eliminate the non-homogeneous reading even in upward-entailing contexts. Fourth, I develop a mixture analysis for *-toka* and *-tari* along lines similar to Persian m-reduplication, factoring in differences between the two phenomena. As with m-reduplication, I supplement the semantic analysis with a pragmatic analysis to derive the non-homogeneity inference associated with *-toka* and *-tari* as a scalar implicature, making use of local implicature calculation and abstract alternatives as was done for m-reduplication. I then turn to two previous analyses of *-toka* and *-tari*, Smith and Kobayashi (2018) and Smith (2019), which adopt two distinct approaches to the phenomena, the first in terms of Hamblin-style alternative semantics (Hamblin (1973); Kratzer and Shimoyama (2002)), and the second in terms of higher-order scalar implicatures (Spector, 2007). I discuss the merits of these analyses, but ultimately reveal several inadequacies in them, and demonstrate how the analysis presented in this chapter improves on them. Finally, I extend the analysis to the similative coordinator *ya*, and discuss how the analysis compares to previous approaches to this morpheme (Hayashishita and Bekki (2011); Sudo (2014)). This analysis is in turn extended to the zero coordinator, allowing for a treatment of the coordinating uses of *-toka* and *-tari*.

4.1 Background on Japanese

Japanese is a Japonic language spoken in Japan, as well as by members of the Japanese diaspora around the world. It is a rigidly head-final language, with basic SOV word order, postpositions, and modifier-modified order within the noun phrase. The language exhibits nominative-accusative alignment, and sentences exhibit topic-comment structure. Noun phrases are marked by case markers, which follow the noun phrase they mark. Verbs are conjugated for tense, aspect, negation, and mood. The sentence in (146) illustrates many of these properties.

(146)

Taro	-wa	Hanako	-ga	Makudonarudo	-kara	kat	-ta	oishi	-soo		
Taro	-TOP	Hanako	-NOM	McDonald's	from	buy	-PST	delicious	-SEEM		
na	hambaagaa	-o	tabe	-nakat	-tara	boku	-wa	kanasi	-ku	nar	-u
PRED	hamburger	-ACC	eat	-NEG	-if	1.SG	-TOP	sad	-PRED	become	-PRS

'If Taro doesn't eat the delicious-looking hamburger that Hanako bought from McDonald's, I will get sad.'

Like Persian, Japanese noun phrases are generally number-neutral in reference. Unlike Persian, they are also generally underspecified for definiteness, as Japanese lacks definite and indefinite articles, as well as differential object marking. As an example, *kyooju* 'professor' in (147) is compatible with both an indefinite and definite interpretation, and may be either singular or plural.

(147)

Taro	-wa	kyooju	-o	mi	-ta
Taro	-TOP	professor	-ACC	see	-PST

'Taro saw a/the professor(s)'

This wide variability in the interpretation of noun phrases in Japanese can be accounted for by treating nouns as lexically number-neutral predicates of individuals (as in analyses like that of Chierchia (1998)), and by assuming the availability of type-shifting operations like IOTA and $\exists$ (Partee (1987)).

4.2 *-toka* and *-tari*

Japanese possesses two morphemes for expressing similative plural meaning: *-toka* and *-tari*. Both of these morphemes exhibit a strong resemblance to m-reduplication in Persian. This section details these similarities, while also discussing key differences between the Persian and Japanese phenomena. I argue that these differences are ultimately syntactic in nature, and elaborate on the syntactic structure of complements to *-toka* and *-tari* while proposing a mixture-based semantics for both morphemes analogous to that previously proposed for m-reduplication.

4.2.1 Properties of *-toka* and *-tari* in upward-entailing contexts

-toka is a suffix that attaches to nominal contituents, while *-tari* is a suffix that attaches to extended projections of the verb phrase. in that they are associated with non-homogeneous plural inferences. For example, (148) is true if Taro and at least one other person similar to him in the context comes, while (149) is judged true if Taro cleans his room and does other similar actions, such as other chores. (148) is infelicitous if only Taro comes, or if only one person other than Taro comes. Likewise, (149) is infelicitous if Taro cleaned his room and did nothing else, or if he did something other than clean his room.

(148) Taro -toka -ga ki -ta
Taro -TOKA -NOM come -PST
'Taro and someone else like that came'

(149) Taro -ga heya -o sooji si -tari si -ta
Taro -NOM room -ACC clean do -TARI do -PST
'Taro cleaned his room and did other such things.'

Like m-reduplicated nominals, *-toka* NPs are compatible with collective predicates.

(150) Taro -toka -ga kooen -de atsumat -ta
Taro -TOKA -NOM park AT gather -PST
'Taro and others gathered at the park.'

4.2.2 *-toka* and *-tari* in non-upward-entailing contexts

-toka and *-tari* display another characteristic in common with Persian m-reduplication in that they too show a sensitivity to the monotonicity of the semantic environments in which they are found: their meanings become inclusive in downward-entailing and non-monotone contexts. What's more, just like with m-reduplication, there is variation among Japanese speakers with regard to whether *-toka* and *-tari* behave partially inclusively or fully inclusively. I will note this variation throughout this section.

Let us start with negation. For less restrictive speakers, (151) is true as long as Taro studies neither English, nor anything similar to it in the context, such as French or German. For more restrictive speakers, (151) will be true as Taro doesn't study English, or English in addition to some other similar thing, but is compatible with a situation in which they study just something similar.

(151) Taro -wa eigo -toka -o benkyoo si -nakat -ta
Taro -TOP english -TOKA -ACC study do -NEG -PST
Less restrictive: 'Taro didn't study English or anything like that.'
More restrictive: 'Taro didn't study at least English.'

The situation with (152) is similar: less restrictive speakers interpret this as meaning that Taro didn't clean his room or do any other chores, but more restrictive speakers permit this to mean that Taro at the very least did not clean his room. For these speakers, Taro may have done some other chore, such as doing the dishes.

(152) Taro -ga heya -o sooji si -tari si -nakat -ta
Taro -NOM room -ACC clean do -TARI do -NEG -PST
Less restrictive: 'Taro didn't clean his room or do anything like that'
More restrictive: 'Taro didn't clean his room, and may not have done other things too.'

For conditionals, less restrictive speakers interpret (153) as meaning that Yosuke will serve tea as long as someone like Taro shows up, while more restrictive speakers require that Taro show up, possibly accompanied by someone else, for Yosuke to serve tea. Similar effects are observed for (154): for less restrictive speakers, Taro's mom will be happy as long as Taro performs some chore similar to cleaning his room, while for more restrictive speakers, Taro's mom will only be happy if he at least cleans his room.

(153) Taro -toka -ga ki -tara Yosuke -wa ocha -o das -u
Taro -TOKA -NOM come -COND Yosuke -TOP tea -ACC serve -PRS
Less restrictive: 'If Taro or someone like that comes, Yosuke will serve tea.'
More restrictive: 'If Taro (and possibly someone else like that) comes, Yosuke will serve tea.'

(154) Taro -ga heya -o sooji si -tari si -tara, mama -wa yorokob -u
Taro -NOM room -ACC clean do -TARI do COND mom -TOP become.happy -PRS
Less restrictive: 'If Taro cleans his room or does something like that, his mom will be happy'

More restrictive: 'If Taro cleans his room (and possibly does something else similar), his mom will be happy.'

Polar questions exhibit the same effect. Less restrictive speakers may respond affirmatively to (155) as long as someone similar to Taro came, while more restrictive speakers will only respond affirmatively as long as at least Taro came. Likewise, less restrictive speakers will respond affirmatively to (149) as long as Taro did one of his chores, while more restrictive speakers will only answer affirmatively if Taro at the very least cleaned his room.

(155) Taro -toka -ga ki -ta no?
Taro -TOKA -NOM come -PST Q
Less restrictive: 'Did Taro or someone like that come?'
More restrictive: 'Did at least Taro come?'

(156) Taro -ga heya -o sooji si -tari si -ta no?
Taro -NOM room -ACC clean do -TARI do -PST Q
Less restrictive speakers: 'Did Taro clean his room or do other such things?'
More restrictive speakers: 'Did Taro at least clean his room?'

The last case we consider is imperatives. For less restrictive speakers, one can comply with (157) by bringing food or drink, while for more restrictive speakers food is required, but additional things may be brought. Similarly, for less restrictive speakers, the imperative in the addressee can comply with (158) by dancing or providing a similar form of entertainment, while for more restrictive speakers the addressee must at least dance in order to comply with the imperative, though they may perform additional entertaining acts.

(157) tabemono -toka motteko -i!
food -TOKA bring -IMP

Less restrictive: 'Bring me food or something like that!'

More restrictive: 'Bring me food at least!'

(158) tsumaranai. Odot -tari si -ro!
boring dance -TARI do -IMP

Less restrictive: 'I'm bored. Dance or something!'

More restrictive: 'I'm bored. Dance at the very least!'

4.2.3 Ignorance contexts

What's more, just like with Persian m-reduplication, Japanese sentences with *-toka* and *-tari* are sensitive to pragmatic aspects of the context. As (159) shows, both *-toka* and *-tari* permit inclusive readings when speaker ignorance is established, despite the fact that the semantic environment in which the *-toka/tari* phrase is found is upward-entailing.

(159) a. CONTEXT: Hanako has a lunch box in which she usually carries a few apples, but sometimes brings other kinds of fruit. You don't know exactly what's in the box, nor do you know exactly how many things are in it.

Hanako -wa ringo -toka -o mot -te i -ru
Hanako TOP apple -TOKA -ACC carry -PROG be -PRS

'Hanako has an apple or something.'

b. CONTEXT: You and a friend have made plans to meet up with Taro. He seems to be running late, but then you remember he had a few chores to do today. You don't know exactly what he's doing today, but you know it's something around the house. Your friends asks you why Taro's running late. You say:

sentaku -o si -tari si -te i -ru
laundry -ACC do -TARI do PROG be -PRS
'He's doing laundry or something like that.'

As such, we find that *-toka* and *-tari* exhibit semantic and pragmatic behavior exactly like that of Persian m-reduplication.

4.3 Structural differences between *-toka/-tari* and m-reduplication

A few differences exist between Persian m-reduplication on the one hand and *-toka* and *-tari* on the other. For one, both *-toka* and *-tari* may be coordinated by juxtaposing two *-toka/-tari* phrases.

(160) Taro -toka Hanako -toka -ga ki -ta
Taro -TOKA Hanako -TOKA -NOM come -PST
'Taro, Hanako, and someone else like that came'

(161) Taro -ga heya -o sooji si -tari sentaku si -tari si -ta
Taro -NOM room -ACC clean do -TARI laundry do -TARI do -PST
'Taro cleaned his room, did laundry, and did other such things.'

Furthermore, although m-reduplication may only target the head noun, *-toka* and *-tari* may attach to significantly larger constituents. *-toka*, for instance, may scope over adjectives (162) and relative clauses (163): (162) is taken to mean that Hanako reads things similar to scientific articles, such as other kinds of intellectually stimulating literature; the other things she read need not be scientific in nature. Similarly, (163) is not necessarily interpreted as meaning Hanako eats fish-like things that were bought at Walmart, but that Hanako eats things similar to fish bought at Walmart, such as other things from department stores, or for those with negative associations with Walmart's goods, things of low quality.

(162) Hanako -wa yoku kagakutekina kiji -toka -o yom -u
Hanako -TOP often scientific article -TOKA -ACC read -PST
'Hanako often reads scientific articles and the like.'

(163) Hanako -wa Walmart -de kat -ta sakana -toka -o tabe -ru
Hanako -TOP Walmart -AT buy -PST food -TOKA -O eat PRS
'Hanako eats fish that was bought from Walmart and the like.'

The complement of *-tari* is verbal, and may contain aspect markers and negation.

(164) Hanako -wa akachan -to ason -de -i -tari si -ta
Hanako -TOP baby -WITH play -PROG be -TARI do -PST
'Hanako was playing with the baby, among other things.'

(165) Taro -wa sentaku si -nakat -tari si -ta
Taro -TOP laundry do -NEG -TARI do -PST
'Taro did things like not doing the laundry'

-tari phrases may not, however, contain modal adverbs or suffixes.

(166) a. *Taro -wa sakana -o tabe soo- dat -tari si -ta
Taroo -TOP fish -ACC seem COP -TARI do -PST
Intended: 'Taroo seems to have eaten fish, among other thing.'

b. *Hanako -wa hasir -u yoo dat -tari si -ta
Hanako -TOP run -PRS alleged COP -TARI do -PST
Intended: 'Hanako allegedly ran, among other things.'

c. *Taro -ga saihu -o tabun otosi -tari si -ta
Taro -NOM wallet -ACC maybe drop -TARI do -PST
Intended: 'Taro may have dropped his wallet among other things.'

The structure of the complement of *-tari* can be explored further by coordinating two *-tari* phrases. Two *-tari* phrases may share a single object, as in (167), but may also be large enough to contain distinct subjects, as (168) shows.

(167) Taro -wa booru -o ket -tari nage -tari si -ta
Taro -TOP fish -ACC kick -TARI throw -TARI do -PST
'Taro kicked and threw the ball, among other things.'

(168) sono ronbun -wa Ken -ga home -tari Ryoo -ga seme -tari si -ta
that thesis -TOP Ken -NOM praise -TARI Ryoo -NOM criticize -TARI do PST
'As for that thesis, Ken praised it and criticized it, among other things.'

However, each *-tari* phrase may not contain a separate topic phrase.

(169) *sono ronbun -wa Ken -ga home -tari sono happyoo -wa Ryoo -ga seme -tari si -ta
that thesis -TOP Ken -NOM praise -TARI that presentation -TOP Ryoo -NOM criticize -TARI do PST
'As for that thesis, Ken praised it and criticized it, among other things.'

These facts suggest that the complements of *-toka* and *-tari* are phrasal; *-toka* takes a full noun phrase as its complement, while the complement of *-tari* may be at least as large as an AspP, but not as large as a CP.

What's more, Japanese allows *-toka/-tari* phrases to modify another expression that overtly manifests the contextual restriction on the similative, whereas the contextual restriction of Persian m-reduplication cannot be similarly expressed.

(170) Taro -wa mikan -toka kudamono -o tabe -ta
Taro -TOP orange -TOKA -ACC eat -PST
'Taro ate oranges and other such **fruits**'

(171) otera -ni it -tari ryokoo si -ta -i
temple -DAT go -TARI travel do -DES -PRS

‘I want to go to temples and do other such **travel-type things**

(172) *Mohsen mive -ye sib mib xord
Mohsen fruit -EZ apple RED eat.PST
Intended: ‘Mohsen ate apples and other such **fruits**

Despite these syntactic differences between *-toka/-tari* and m-reduplication, there is reason to believe that they do not greatly differ in their basic semantics. In particular, I maintain that, like m-reduplication, *-toka* and *-tari* both take arguments of predicative type: *-toka* takes type <e,t> arguments, and that *-tari* takes type <v,t> arguments, with v the type of events. The fact that these expressions cannot take arguments of a higher type can be seen by their interaction with quantifiers. Quantificational expressions, such as *subete* ‘all’, must take scope over the entire similative plural: (173) can only mean that everyone in the combined set of professors and similar things, such as grad students, came. It could not be used to describe a situation in which, for instance, every professor came, in addition to something that counts as similar to every professor, such as five graduate students.

(173) subete no kyooju -toka -ga ki -ta
all -GEN professor -TOKA -NOM CL come -PST
‘Everyone in the set of professors and similar things came’/*’Every professor came, and so did something similar to every professor’

Such a result is expected if *-toka* denotes something of type <e,t>: a quantificational DP, of type <<e,t>,t>, could not be an argument of *-toka*, but it could take the *-toka* phrase as an argument.

4.4 A mixture analysis of *-toka* and *-tari*

In line with the discussion above, I propose that *-toka* and *-tari* be analyzed as (partially) inclusive mixtures. I begin my analysis with *-toka*. As with m-reduplication, I analyze *-toka* as a function that takes a type <e,t> argument[1]. Due to the fact that *-toka*, unlike m-reduplication, allows for the overt expression of the contextual restriction on the domain in which similarity is being evaluated, I also propose that *-toka* takes an additional type <e,t> argument. (175a) shows the translation of *-toka* for more more restrictive speakers, while (175b) provides the one for less restrictive speakers. The result is another predicate of individuals for both sets of speakers.

(175) a. $\textit{-toka}_{\text{more restrictive}} \rightsquigarrow \lambda P.\lambda Q.\lambda X.\text{Mix}(P,P^{\simeq}{}_{Q})(X)$

b. $\textit{-toka}_{\text{less restrictive}} \rightsquigarrow \lambda P.\lambda Q.\lambda X.\text{I-Mix}(P,P^{\sim}{}_{Q})(X)$

Syntactically, I propose that *-toka* takes an NP complement. It may then adjoin either to another NP corresponding to its second type <e,t> semantic argument, or to a null expression I give here as *C*.

(176) Syntactic structure of *-toka* phrases

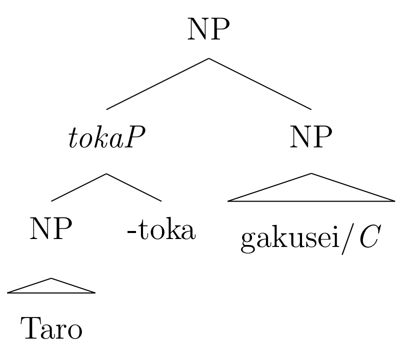

[1]Names like *Taro* can be lifted to type <e,t> via the IDENT type-shifter of Partee (1987).

(174) IDENT t $\Rightarrow$ λx.x = t

Deriving the exclusive reading proceeds exactly as in the m-reduplication case, where the *-toka* sentence competes with the alternative corresponding to the bare noun for more restrictive speakers, and with the alternatives corresponding to the bare noun and its proper similarity set for less restrictive speakers. Let us begin with the more restrictive speakers. The translation of a sentence like (148) prior to existential closure of the event variable and pragmatic strengthening is as in (177).

(177) λe.$\exists$X[Mix(λx.x = t,$\simeq$(λx.x = t)(X) $\wedge$ *COME(e) $\wedge$ AGENT(e) = X]

As with the more restrictive speakers of Persian with respect to m-reduplication, there are two alternatives to consider: the one corresponding to the bare noun, in this case *Taro*, and the one corresponding to the similarity set of this noun, in this case the set of individuals that are considered similar to Taro in the context, but are not identical to Taro. However, as with the analogous m-reduplication case, excluding both alternatives would contradict what is asserted by the sentence prior to strengthening. Here, the only innocently excludable alternative is the bare noun alternative, *Taro*. Negating this alternative and conjoining it within the scope of existential closure of the event variable yields the result in (178)

(178) $\exists$e[$\exists$X[Mix(λx.x = t,$\sim$(λx.x = t)(X) $\wedge$ *COME(e) $\wedge$ AGENT(e) = X] $\wedge$ $\neg\exists$X[X = t $\wedge$ *COME(e) $\wedge$ AGENT(e) = X]]

This says that there is an event in which some individual in the partially inclusive mixture of Taro with other individuals like him came, but it is not the case that only some individual identical to Taro came. In other words, (178) will be true iff some individual came, and that individual is a sum of Taro and someone like him, as desired.

A similar anaylysis holds for the less restrictive speakers as well. As with the analysis of less restrictive speakers with regard to m-reduplication in Persian, these speakers differ in from the more restrictive speakers in having a fully inclusive mixture representation for *-toka*. This means the translation of (148) will be as in (179) for these speakers.

(179) $\lambda e.\exists X[\text{I-Mix}(\lambda x.x = t,\sim(\lambda x.x = t)(X) \wedge {}^*\textsc{come}(e) \wedge \textsc{agent}(e) = X]$

The set of alternatives to (179) are the same as those for (177). However, in this case, both alternatives are innocently excludable: the set containing only Taro and the set of individuals similar but not identical to Taro are proper subsets of their mixture. Excluding these alternatives and closing the event variable results in the following sentence.

(180) $\exists e[\exists X[\text{I-Mix}(\lambda x.x = t,\sim(\lambda x.x = t)(X) \wedge {}^*\textsc{come}(e) \wedge \textsc{agent}(e) = X] \wedge \neg\exists X[X = t \wedge {}^*\textsc{come}(e) \wedge \textsc{agent}(e) = X] \wedge \neg\exists X[{}^*\sim(\lambda x.x = t)(X) \wedge {}^*\textsc{come}(e) \wedge \textsc{agent}(e) = X]]$

This says that there is an event of coming whose agent is something in the inclusive mixture of Taro and those similar to him, but the agent of that event is neither identical to Taro himself nor is it just some number of people similar to Taro. This leaves but one option for the agent of this event: it must be a sum of Taro and at least one indivdual similar, but not identical to him. This is the desired exclusive mixture reading, identical to the reading of (178) above.

Having shown how the analysis correctly derives the exclusive mixture reading of *-toka* sentences for both sets of speakers, let us turn now to the analysis of *-tari*. Because *-tari* involves verbal material, I propose that it takes as arguments predicates of *events*, and mixes a set of events with its similarity set or proper similarity set. Once again, we make a distinction between the analysis of *-tari* for more and less restrictive speakers, given in (181a) and (181b), respectively.

(181) a. -*tari*$_{\text{more restrictive}}$ $\rightsquigarrow$ λV.λV'.λe.Mix(V,V$^{\simeq}{}_{\text{V'}}$)(e)

b. -*tari*$_{\text{less restrictive}}$ $\rightsquigarrow$ λV.λV'.λe.I-Mix(V,V$^{\sim}{}_{\text{V'}}$)(e)

Syntactically, I propose that -*tari* may take any verbal projection as complement, such as *v*P or AspP. This corresponds to -*tari*'s first semantic argument, which denotes a predicate of events. As with -*toka*, I further propose that the phrase headed by -*tari* adjoins to another phrase, in this case a *v*P corresponding to -*tari*'s second semantic argument. This phrase may be headed by a lexical verb itself, or may be simply a light verb *si* 'do', to which tense inflections may attach. This analysis is sketched in (182).

(182) Syntactic structure of -*tari* phrases

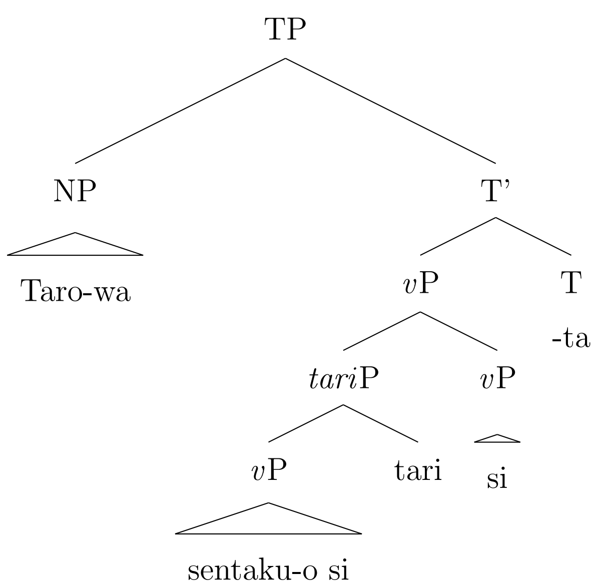

Here again, the derivation of the exclusive mixture reading from the partial and fully inclusive readings respectively proceeds via calculating and negating the innocently excludable alternatives. In the case of more restrictive speakers, the translation of a sentence like (149) is as in (183)[2].

(183) λe.Mix(λe.clean(e),$\sim$(λe.clean(e)))(e) $\wedge$ AGENT(e) = t $\wedge$ THEME(e) = ιx[room(x)]

As with the more restrictive analysis of *-toka*, for these speakers the only innocently excludable alternative is the one corresponding to the verb phrase that *-tari* takes as complement. Negating this alternative derives the exclusive mixture interpretation from the underlying partially inclusive semantics: (184) is true iff there is an event in the partially inclusive mixture of cleaning events and similar sorts of events, but is not itself in the set of cleaning events. The event must therefore be a sum of a cleaning event and an event of a different, but similar, sort in the context, such as an event of organizing the room.

(184) $\exists$e[Mix(λe.clean(e),$\sim$(λe.clean(e)))(e) $\wedge$ AGENT(e) = t $\wedge$ THEME(e) = ιx[room(x)] $\wedge$ $\neg$(clean(e) $\wedge$ AGENT(e) = t $\wedge$ THEME(e) = ιx[room(x)])]

In the case of less restrictive speakers, the translation of (149) involves an inclusive mixture, as in (185).

(185) λe.I-Mix(λe.clean(e),$\sim$(λe.clean(e)))(e) $\wedge$ AGENT(e) = t $\wedge$ THEME(e) = ιx[room(x)]

Here, both alternatives derived from the arguments of the mixture function in the translation of *-tari*, the predicate of events corresponding to the overt verb phrase and

[2] For the sake of legibility I have chosen to give *-tari* scope only over the bare predicate of events, λe.clean(e), rather than over the entire VP, which would include the theme. This parse is in fact independently attested, and can be observed in sentences involving coordination of *-tari* phrases, as observed in (167). This could be derived by attaching *-tari* directly to a constituent containing the verb and excluding the object, assuming this constituent is also of type <v,t>. Alternatively, the complement of *-tari* may be a much larger constituent in the extended projection of the verb, and the object may undergo across the board movement from both conjuncts, leaving a variable in each conjunct. The analysis presented below is compatible with both possibilities.

its proper similarity set, are innocently excludable, and negating these derives from the underlying fully inclusive semantics an exclusive mixture reading, once again as desired.

(186) $\exists$e[I-Mix(λe.clean(e),$\sim$(λe.clean(e)))(e) $\wedge$ $\textsc{agent}$(e) = t $\wedge$ $\textsc{theme}$(e) = ιx[room(x)] $\wedge$ $\neg$(clean(e) $\wedge$ $\textsc{agent}$(e) = t $\wedge$ $\textsc{theme}$(e) = ιx[room(x)]) $\wedge$ $\neg$($\sim$(λe.clean(e))(e) $\wedge$ $\textsc{agent}$(e) = t $\wedge$ $\textsc{theme}$(e) = ιx[room(x)])]

Having developed my analysis of *-toka* and *-tari* in terms of the mixture semantics developed for m-reduplication, I turn now to previous analyses of these particles.

4.5 Previous analyses of *-toka* and *-tari*

Two previous analyses of *-toka* and *-tari* exist in the formal semantics literature: Smith and Kobayashi (2018) and Smith (2019). I will discuss these approaches in turn, demonstrating the inadequacies of each analysis and how the present proposal improves on both.

4.5.1 Smith & Kobayashi 2018

In order to derive the polarity sensitivity of these particles, Smith and Kobayashi (2018) analyze the semantics of *-toka* and *-tari* using a Hamblin-style alternative semantics (Hamblin (1973), Kratzer and Shimoyama (2002)), according to which *-toka* and *-tari* are generators of sets of individuals and sets of predicates of individuals, respectively.

(187) $[\![\alpha\text{-}toka]\!] = \{\text{x} \mid \text{x} \sim \alpha\}$

(188) $[\![\beta\text{-}tari]\!] = \{\text{P} \mid \text{P} \sim \beta\}$

The sets of propositions generated by composing these sets with the rest of the sentence via Pointwise Function Application need to be flattened into singleton sets in order to be interpretable as declaratives, so two propositional quantifiers serve to collapse these sets: in upward-entailing contexts, this is accomplished by a default universal quantifier, defined as in (189. In other contexts, the alternatives are effectively existentially quantified, either by the semantics of operators in conditionals or modals, or explicitly via the propositional existential quantifier, defined in (190).

(189) $[\![\forall]\!](\mathrm{A}) = \{\forall \mathrm{p}[\mathrm{p}{\in}\mathrm{A} \rightarrow \mathrm{p}\ \bar{1}]\}$

(190) $[\![\exists]\!](\mathrm{A}) = \{\exists \mathrm{p}[\mathrm{p}{\in}\mathrm{A} \wedge \mathrm{p}\ \bar{1}]\}$

This analysis is able to capture the sensitivity of *-toka* and *-tari* to their semantic environment, and is also able to capture the non-homogeneity effect associated with these expressions. There are at least two problems with this approach, however. First, use of the universal propositional quantifier results in overly strong truth conditions: an expression with *-toka* or *-tari* is predicted to be true only if *all* the alternatives are true. In reality, they are judged true/felicitous in upward-entailing contexts when at least one other alternative is true. As such, (148) is felicitous if Taro *and at least one similar person* comes, not just if Taro and *all* of those similar to him in the context come. Second, the analysis predicts that *-toka/-tari* should receive something like a disjunctive interpretation in downward-entailing contexts; for example, the conditional in (153) should mean that Yosuke will serve tea if Taro *or* someone similar comes. This judgment is attested, but not all Japanese speakers agree with it: the interpretation other speakers get is that, if *at least* Taro comes (possibly along with those similar to him in the context), Yosuke will serve tea. The Hamblin alternatives-based analysis has no way of accounting for this set of speakers' judgments.

A third problem with the analysis arises from the interaction of alternatives with collective predicates. In a Hamblin-style alternative semantics, composition of alternative individuals with predicates is *pointwise*: each individual composes with a function, ultimately delivering a set of results corresponding to the result of applying the function to each argument. Because some alternative individuals generated by a *-toka* phrase like *Taro-toka* will be singular, and collective predicates require that their argument be *plural*, we expect that *-toka* phrases containing singular alternatives should be unacceptable with collective predicates. This prediction is not borne out: *-toka* phrases with singulars are acceptable with collective predicates like *atsumaru* 'gather', while the counterpart of such a sentence without *-toka* is unacceptable. This contrast is not explained on Smith & Kobayashi's analysis.

(191) Taro -toka -ga kooen -de atsumat -ta
Taro -TOKA -NOM park LOC gather -PST
'Taro and others like that gathered at the park.'

(192) #Taro -ga kooen -de atsumat -ta
Taro -NOM park LOC gather -PST
'Taro gathered at the park.'

Finally, the analysis is somewhat unprincipled: although the analysis derives an empirically unproblematic interpretation for less restrictive speakers in downward-entailing contexts, the analysis of polar questions requires a stipulated use of the existential propositional quantifier, and the fact that weaker readings are attested in certain contexts is not connected to their downward-entailing nature. Because of this, the analysis also has nothing to say about *-toka* and *-tari*'s sensitivity not just to the semantic environment in which they appear but also to more pragmatic concerns such as speaker ignorance.

4.5.2 Smith 2019

Smith (2019) proposes a solution to these issues using a *higher-order implicature* analysis Spector (2007). On this analysis, expressions suffixed with *-toka/-tari* are identical in their unenriched meanings to the expression without these morphemes. *-toka* and *-tari* phrases then compete with exhaustified variants of the sentence without the morphemes. This amounts to negating an alternative with the meaning of, say, *only Taro* came in the case of *-toka*, or, in the case of *-tari*, Taro *only did laundry*. The exclusion of these alternatives leads to an enrichment of the basic meaning, such that the interpretation of a *-toka* sentence ends up being, essentially, 'Taro came, and it's not the case that only Taro came', and the interpretation of a *-tari* sentence is 'Taro did laundry, and it's not the case that Taro only did laundry.' These are both non-homogeneous readings, though derived in a way different from the analysis proposed in this chapter.

This analysis makes many correct predictions where Smith & Kobayashi's approach failed: it correctly accounts for the fact that only one alternative in addition to the one asserted needs to be true, rather than all of them, it naturally accounts for the partially inclusive readings that many Japanese speakers get with such sentences in downward-entailing, non-monotone and ignorance contexts rather than only permitting the fully inclusive reading, it is less stipulative than the previous analysis, and it is able to tie the sensitivity of these expressions to certain types of contexts to well-understood properties of implicature.

Unfortunately, this analysis runs into additional problems. Empirically, it suffers from the same problem with collective predicates as Smith & Kobayashi's analysis: because the constituent modified by *-toka* can be singular, and the implicature is calculated at the sentence level, the analysis predicts that a collective predicate like 'gather' predicated of a *-toka* phrase should be semantically anomalous, contrary to fact. It also struggles to provide a natural account of the judgments of less restrictive speakers, and the analysis of these judgments is essentially stipulated. On a more conceptual level, if sentences with *-toka/-tari* are equivalent to those without such a morpheme, it isn't clear why speakers would use the expression in non-upward-entailing contexts in the first place. The reason is clear independent of the analysis: *-toka/tari* expressions are weaker than their bare counterparts. This, however, is not reflected in an analysis based on higher-order scalar implicature.

4.5.3 Solving problems with previous analyses

The analysis proposed in this chapter straightforwardly solves the problems with both Smith and Kobayashi (2018) and Smith (2019)'s analyses. Lke Smith (2019), the current analysis solves several problems with the analysis of Smith and Kobayashi (2018): *-toka/tari* sentences are not predicted to have universal readings, and the account explains the connection between the different interpretations of *-toka/tari* to the properties of implicature in non-upward-entailing environments and ignorance contexts. Beyond this overlap with Smith (2019), the analysis solves the problems that plagued both analyses. First, recall that Smith and Kobayashi (2018) could handle the judgments of less restrictive but not those of more restrictive speakers, and Smith (2019) had a principled analysis of the more restrictive speakers but an *ad hoc* analysis of the less restrictive speakers. Second, both analyses incorrectly predict that *-toka* phrases should be incompatible with collective predicates when the complement of *-toka* denotes an atomic individual. On the current analysis, it is possible to account for the judgments of both less and more restrictive speakers in the same general mixture-based framework. What's more, because mixtures contain sum individuals in their denotation, the analysis correctly predicts that *-toka* phrases are compatible with collective predicates. Finally, the analysis improves on a problem specific to the analysis of Smith (2019) as well. In particular, we now have an explanation for why speakers would use *-toka/tari* in non-upward-entailing contexts: they are logically weaker than their counterparts without *-toka/-tari*.

4.6 Conclusion and Future research

In this chapter, I examined the similative plural morphemes *-toka* and *-tari*, showing that their semantic and pragmatic behavior mirrors that of m-reduplication in Persian, despite some small ultimately syntactic differences between the two classes of expression. I proposed a mixture-based semantic analysis in tandem with an implicature analysis to explain the behavior of these expressions in a range of contexts, in addition to proposing a slightly different syntactic analysis for them in order to account for their differences from m-reduplication. I further demonstrated how the current analysis improves on previous analyses of these morphemes (Smith and Kobayashi (2018); Smith (2019)).

This is not the end of the story for the study of similative expressions in Japanese, and their analysis will be the focus of future work. There are two other similative expressions in Japanese: *ya* and the null coordinator, the latter of which I will denote with $\varnothing$. These expressions differ from *-toka* and *-tari* in being coordinative: *ya* requires two DP arguments, and the null coordinator requires two arguments of the same type[3]. Like *-toka* and *-tari*, they are associated with both a conjunc-

[3]Generally, these will also be DPs. I remain intentionally vague about the types of the coordinands of $\varnothing$ at this juncture, and will make a more concrete proposal in the analysis of coordinative uses of *-toka* and *-tari* below.

tive and a non-exhaustive inference in upward-entailing environments. To illustrate, (193) below is judged felicitous if Taro, Hanako, *and* someone else like them comes (Hayashishita and Bekki (2011); Sudo (2014)).

(193) Taro ya/∅ Hanako -ga ki -ta
Taro YA/∅ Hanako -NOM come -PST
'Taro, Hanako, and someone else came.'

The coordinands of *ya* and ∅ exhibit a similarity restriction: coordinations of DPs whose referents have very little in common are judged odd, unlike similar coordinations using *to* 'and'.

(194) #neko ya ringo ya amerika
cat YA apple YA America
'Cats, apples, America, and things like that.'

(195) neko to ringo to amerika
cat and apple and America
'Cats, apples, and America.'

Singular DPs coordinated by *ya* or ∅ are compatible with collective predicates, as (196) shows (Sudo, 2014). This is also true of ordinary DP conjunction with *to* 'and' (197), but not of disjunctions with *ka* 'or' (198).

(196) a. John ya Martin -ga isshoni kenkyuu -o si -ta
John YA Martin -NOM together research -ACC do -PST
'John, Martin, and others did research together.'

b. Taroo -wa kono hon ya kono zasshi -o tabane -ta
Taro -TOP book YA magazine -ACC bundle -PST
'Taro bundled up this book, this magazine, and other such things.'

(197) Taroo -wa kono hon to kono zasshi -o tabane -ta
Taro -TOP book and magazine -ACC bundle -PST

'Taro bundled up this book and this magazine.'

(198) #Taroo -wa kono hon ka kono zasshi -o tabane -ta
Taroo -TOP this book or this magazine -ACC bundle -PST
'Taro bundled up this book or this magazine'

Much like the non-uniform plural inference observed with m-reduplication and *-toka* and *-tari*, the non-exhaustive conjunctive reading of *ya* and ∅ vanishes in non-upward-entailing environments. It is replaced by a *disjunctive* reading, where the disjunction includes both overt coordinands as well as the unmentioned individuals similar to the overt coordinands.

(199) Negation

Taro -wa Hanako ya Jiro -o mi -nakat -ta
Taro -TOP Hanako YA Jiro -ACC see -NEG PST

'Taro didn't see Hanako, Jiro, or anyone like that.'

(200) Antecedent of a conditional

Taro ya/∅ Hanako -ga paatii -ni ki- tara, Yoosuke -wa ocha -o
Taro YA/∅ Hanako -GA party -DAT come -COND Yosuke -TOP tea -ACC
das -u
serve -PRS

'If Taro, Hanako, or someone like that comes to the party, Yosuke will serve tea.'

(201) Polar question

Taro ya/∅ Hanako -ga ki -ta no?
Taro YA/∅ Hanako -NOM come -PST Q

'Did Taro, Hanako, or someone like that come?'

(202) Imperative

paatii -ni Taro ya/∅ Hanako -o tsureteko -i!
party -DAT Taro YA/∅ Hanako -ACC bring -IMP

'Bring Taro, Hanako, or someone like that to the party!'

Taken all together, *ya* and ∅ exhibit many of the same properties of m-reduplication and *-toka* and *-tari*. That said, extending the mixture-based analysis to these expressions is not completely straightforward. While capturing the disjunctive interpretation of these expressions in non-upward-entailing contexts is straightforward, the non-exhaustive conjunctive reading is not as easy to derive as it is for *-toka* and *-tari*. The reason is that we need to exclude sums of one of the overt coordinands with an unmentioned member of their similarity set that does not include the other coordinand, all while also ensuring that a sum of both coordinands with at least one member of their similarity set is *not* excluded. This turns out to be a non-trivial task. For the moment, I leave this to future research.

CHAPTER 5

ASSOCIATIVE PLURALITY AND PLURAL PRONOUN CONSTRUCTIONS

This chapter marks a shift in focus away from similative plurality, the topic of the previous two chapters, and toward *associative plurality*, another type of non-homogeneous plural found in many languages. It also introduces *plural pronoun constructions* (PPCs), a type of coordinating construction involving a plural personal pronoun and a comitative adpositional phrase in which the complement of the comitative head is interpreted as part of the plurality expressed by the pronoun. Following a connection between these two seemingly different phenomena often suggested in the typological literature, but rarely if ever formally implemented, I develop a unified analysis of the two phenomena in terms of operations on the argument of an associative/PPC and its *associate set*. The analysis not only successfully captures the semantic behavior of associative and PPC constructions under a unified framework, but also reveals a common core between associative and similative plurality: both types of plurals are derived via operations on sets constructed from relations between elements of the domain of individuals.

The chapter is structured as follows. Section 1 provides an overview of the properties of associative plurals cross-linguistically, focusing on their semantic behavior with different types of arguments and restrictions on the kinds of arguments they can take. Section 2 then introduces Plural Pronoun Constructions and their properties. Section 3 provides a formal analysis of both associatives and PPCs using a common toolkit couched in terms of operations on individuals and a set constructed from salient *social relations*, and shows how these tools derive the semantic behavior of these expressions. The fourth section discusses a novel way of understanding the underlying structure that associatives and pronouns share. Section 5 reviews previous analyses of both phenomena investigated in this chapter, and demonstrates how the unified analysis improves on previous approaches.

5.1 Associative plurality

Associative plurals are attested in a number of languages, ranging from Indo-European languages like Persian and Afrikaans (den Besten (1996); Ghomeshi (2018)) to Uralic languages like Hungarian, from Mandarin to Japanese (Moravcsik (2003); Nakanishi and Tomioka (2004)). Examples of associative plurals from each of these languages are given below in (203).

(203) a. Japanese

Taro -tachi
Taro -ASSOC

'Taro and his associates'

b. Chinese

Zhangsan tamen
Zhangsan 3.PL

'Zhangsan and company'

c. Hungarian

Péter -ék
Peter -ASSOC

'Peter and his associates'

d. Persian

Farhâd -inâ
Farhad -ASSOC

'Farhad and his family/close friends'

e. Afrikaans

Pa -hulle
dad ASSOC

'Dad and mom'/'Dad and his associates'

Associative plurals differ from more familiar English-like plurals, but pattern with similative plurals, in their *non-homogeneity*: while in an English plural, every part of the plurality must have the property named by the argument of the plural morpheme, the parts of an associative plural need not all be describable by the overt argument of the plural. They are merely required to be *associated* with the overt argument. For example, (204) is true as long as Taro and some other person(s) associated with him in some way came. This group of people may be one or more family members, coworkers, friends, or any other people with a contextually salient association with Taro.

(204) Taro -tachi -ga ki -ta
Taro -ASSOC -NOM come -PST
'Taro and his associates came.'

Associative plurals and similative plurals pattern alike with respect to their non-homogeneity. The distinction between these two types of plurals lies primarily in the *relationship* between two sets in the denotation of the plural: 1) the base set falling in the denotation of the noun to which the plural is applied and 2) the set of objects in the non-homogeneous plural that is not in the denotation of the bare noun. For similative plurals, the relation is one of similarity: sets of objects similar to the base set in some contextually salient sense are included in the denotation of the plural. For associative plurals, this relation is *social* in nature: the members of the two sets stand in some contextually determined association with one another. The difference can be illustrated with a minimal pair in Japanese: consider (205), with the similative morpheme -*toka*, and (206), with the associative plural -*tachi*.

(205) Kyoo -wa *pro* kyooju -toka -o mi -ta
today -TOP *pro* professor -TOKA -ACC see -PST
'Today I saw professors and such.'

(206) Kyoo -wa *pro* kyooju -tachi -o mi -ta
today TOP *pro* professor -ASSOC -ACC see -PST
'Today I saw (the) professor(s) and their associates.'

(205) is true and felicitous in a context in which the speaker saw at least one professor as well as similar such people, such as other university staff, like graduate students and university board members. What's more, each member of this set need not have any connection to one another: the professors, graduate students and other members of the plurality may simply be random individuals meeting that description, with no association with one another. In this kind of context, -*tachi* is infelicitous. On the other hand, (206) is true and felicitous in a context in which the speaker sees some professor and his family, a context in which -*toka* is unlikely to be used.

5.1.1 Distributional restrictions on associative plurals

Associative plurals exhibit two common restrictions cross-linguistically. First, they exhibit an *animacy* requirement: associative plurals can only attach to expressions denoting animate entities. In fact, in the majority of languages that have them, associative plurals are restricted to human beings (Moravcsik, 2003)[1]. (207) displays

[1] Exceptions may include other social animals, such as dogs.

this requirement for Japanese: while *-tachi* may attach to names or expressions denoting people, it may not attach to expressions denoting inanimate objects such as books.

(207) a. Sensei -tachi
teacher -ASSOC
'(The) teacher and associates'

b. #hon -tachi
book -ASSOC
Intended: '(the) book and associates'

The second widely attested restriction on associative plurals is a *definiteness* requirement: in most languages with associative plurals, the argument of the associative must be definite (Moravcsik, 2003). This is true for most of the languages discussed above, but not of Japanese, as Nakanishi and Tomioka (2004) demonstrate. For example, a *-tachi* phrase may contain numerals (208), and may occur in existential constructions (209), in both cases with indefinite interpretations.

(208) 200 -nin -ijyoo -no gakusei -tachi
200 -CL -OR.MORE -GEN student -ASSOC
'200 or more students'

(209) kooen -ni kodomo -tachi -ga i -ta
park -LOC child -ASSOC -NOM be -PST
'There are children in the park'

5.1.2 Non-atomicity

Like all attested varieties of plurals, associative plurals are associated with a multiplicity inference in upward-entailing contexts. However, unlike similative plurals and English bare plurals, associative plurals are *not* number-neutral in non-upward-entailing contexts: they retain their non-atomic reference in these contexts. This can be shown clearly with the antecedent of a conditional (210) and polar question (211).

(210) Taro -tachi -ga ki -tara boku -wa yorokob -u
Taro -ASSOC -NOM come -COND I -TOP become.happy -PRS
'If Taro *and* his associates come, I'll be happy.'

(211) a. Taro -tachi -ga ki -ta no?
Taro -ASSOC -NOM come -PST Q
'Did Taro and his associatives come?'

b. Hai, Taro to Hanako -ga ki -ta / #Taro -ga ki -ta
Yes, Taro and Hanako -NOM come -PST / Taro -NOM come -PST
'Yes, Taro and Hanako came/#Taro came'

While in a similar sentence to (210) making use of *-toka* I will be happy if Taro comes alone, in (210) itself my happiness is only guaranteed if Taro and one of his associates come; all bets are off if only Taro comes. Likewise, it is not possible to answer (211) in the affirmative if only Taro comes. This once again contrasts with the sentence's equivalent with *-toka*, which can be answered affirmatively if only Taro or someone similar to him comes.

5.1.3 Strong and weak non-homogeneity

Associative plurals display two different degrees of non-homogeneity, depending on properties of the argument of the associative plural morpheme. When applied to singular expressions, such as names, associatives are *strongly non-homogeneous*: they necessarily refer to a plural individual consisting of the individual denoted by the phrase to which the plural is attached and some other individual. For example, (204) is only true if Taro and one of his associates come. It is false if only Taro comes, or if only his associate comes. This holds as well for associatives in other languages, such as Afrikaans *-hulle*. This is shown in (212).

(212) a. Japanese

Taro -tachi -ga ki -ta
Taro -ASSOC -NOM come -PST

'Taro and his associates came'/#Taro came'

b. Afrikaans

Pa -hulle
Dad -ASSOC

'Dad and his associates/#Dad'

However, if the argument of the associative morpheme is plural or number-neutral, the plurality is instead *weakly non-homogeneous*: the plurality may contain only individuals describable by the argument of the associative morpheme, or it may additionally contain those individuals' associates. In other words, the non-homogeneity typically seen with associatives seems to be *weakened* with these types of arguments. This can be seen when associatives take a conjunction of names as an argument. For example, (213) is true even if only Taro and Hanako come (Tatsumi, 2017). Likewise, the Akrifaans example in (214) may refer to a group consisting of only Piet and Koos, though it could refer to a larger group of people containing their associates as well (den Besten, 1996).

(213) Taro to Hanako -tachi -ga ki -ta
Taro and Hanako -ASSOC -NOM come -PST

'Taro, Hanako, and their associates came/Taro and Hanako came.'

(214) Piet en Koos -hulle
Piet and Koos -ASSOC

'Piet, Koos, and associates/Piet and Koos'

The same weak non-homogeneity holds of plural definites, as in (215), as well as of number-neutral indefinites (216). Tatsumi (2017) refers to the latter case as an "additive reading" of the associative.

(215) die kinders -hulle
the children -ASSOC
'The children, plus or minus a non-child'

(216) Kyooju -tachi -ga ki -ta
professor -ASSOC -NOM come -PST
'Professors and their associates came/Professors came.'

Cross-linguistically, many associative plural markers have the same form as third person plural pronouns. This is true of the Chinese, Persian, and Afrikaans examples given above. For this reason, they are often compared to certain plural pronominal constructions in other languages, to which I turn next.

5.2 Plural Pronoun Constructions

Plural Pronoun Constructions (henceforth PPCs) are a type of construction occurring in many languages, especially Slavic languages like Russian and Polish, in which a plural personal pronoun appears to have singular reference when it cooccurs with a comitative phrase (Dyla (1988); den Dikken (2001); Vassilieva (2005); Vassilieva and Larson (2005)). (217) gives two examples from Russian. In the first example, the first person plural pronoun *my* 'we' receives what has been termed an *exclusive* interpretation (Moravcsik, 2003): it refers to the speaker and some unspecified associate. In the second example, on the other hand, the pronoun coocurs with a comitative prepositional phrase *s Petej* 'with Petya,' and receives an *inclusive* interpretation:

the plurality denoted by the pronoun *includes* the individual denoted by the complement of the preposition. As such, (217b) is true in a situation in which only the speaker and Petya will go home; it is not necessary that two people, including the speaker, along with Petya go home, with three people in total going[2].

(217) Interpretation of Russian plural pronoun and PPC (Vassilieva (2005); Vassilieva and Larson (2005))

a. My pojdë -m domoj
1.PL.NOM go.PRF.FUT -1.PL home
'We will go home.'

b. My s Pet -ej pojdë -m domoj
1.PL.NOM with Petya -INSTR go.PRF.FUT -1.PL home
'Peter and I will go home/Peter and we will go home.'

A recurring idea in the typological and syntactic literature is that associatives and inclusive readings of PPCs are closely connected (den Besten (1996); Moravcsik (2003); Vassilieva (2005), a.o.). I follow this literature in adopting this perspective and proposing my own unified analysis of associative plurals and PPCs. I turn to this unified analysis in the next section.

5.3 A unified analysis of associative plurals and PPCs

In this section, I develop a formal framework for the analysis of associative plurals and PPCs. The section proceeds as follows. The first subsection introduces the formal ingredients required for the analysis. These tools are then applied to the analysis

[2]Vassilieva (2005) notes that the latter is also a possible interpretation. This suggests that the interpretation of a PPC is not limited to the speaker and the complement of the comitative preposition, but that the set is potentially larger. This will be cashed out in the analysis to come in section 3.

of associative plurals and PPCs in sections 2.2 and 2.3, respectively, and it is demonstrated that the analysis not only correctly accounts for the empirical observations made earlier in the chapter, but also makes additional correct predictions about the behavior of associatives and PPCs.

5.3.1 Ingredients

The basis for the analysis is a set of relations SOCREL $\subset \wp(D \times D)$, with each relation conceived standardly as a set of ordered pairs. Intuitively, SOCREL is a set of *social relations* between two individuals, and we can think of this as the basis for a notion of association between people. We can then use a *choice function* f over this set of relations, which will pick out a particular relation in this set depending on the context[3]. An illustration of SOCREL and the effect of a choice function on SOCREL is given below.

(218) a. SOCREL = {FAMILY MEMBER, PARTNER, FRIEND, COLLEAGUE, ...}

b. f(SOCREL) $\in$ SOCREL

Next, I define two similar operations on relations. The first operation involves restricting a relation R to a particular individual x. A *relation restricted to an individual*, henceforth R_x, is defined as in (**??**), where where p is a variable over ordered pairs and π_1 is a function that returns the first projection of an ordered pair.

(219) Relation restricted to x

$R_x = \{p \mid p \in R \; \& \; \pi_1(p) = \mathrm{x}\}$

[3]Choice functions on sets have been used extensively in formal semantics in the analysis of exceptionally scoping indefinites (Reinhart (1997a); Winter (1997), Kratzer (1998). In this context, a choice function is a function from sets to individuals that returns some member of the set it takes as an argument. Choice functions in this context are either existentially quantified over or left open as a free variable to be supplied by context, much like in the analysis here.

Intuitively speaking, this operation restricts our attention to relationships holding only between the individual x and others. This will allow us to restrict our selected social relations to the individual whose associates we would like to construct a plural from.

I further define a method for restricting relations to a particular *property* P, where a property is simply a set of individuals[4]. A *relation restricted to P*, henceforth R_P, is defined as in (220).

(220) Relation restricted to P

$$R_P = \{p \mid p \in R \ \& \ \mathrm{P}(\pi(p))\}$$

Breaking this down, this operation restricts the first projection of each pair in the relation to those with the property P. For instance, if the property by which R was restricted were the property of being a professor (λx.PROFESSOR(x)), then $R_{\lambda x.professor(x)}$ will be a relation between professors and other individuals.

In addition to a means for restricting relations to particular individuals or properties, we need a way to flatten relations into simple sets. I accomplish this using the function FLAT, which takes a relation and returns a set of sum individuals, where each individual is a sum of the first and second projections of the pairs in the relation.

(221) $\textsc{flat}(R) = \{\mathrm{x}\oplus\mathrm{y} \mid \langle \mathrm{x},\mathrm{y}\rangle \in R\}$

Finally, we need a way to extract a set from its characteristic function. This operation was used in the second chapter in order to construct similarity sets out of functions, and was written $f*$. In order to avoid confusion with the cumulative closure operator * of Link (1983), of which use will also be made, I will name this function EXTRACT.

(222) $\textsc{extract}(\mathrm{f}) = \{\mathrm{x} \mid \mathrm{f(x)} = 1\}$

[4] Or, equivalently, a function from individuals to truth values

With all of the above formal tools in place, it is now possible to define a notion of *associate set*, notated ASET. The *associate set of an individual* is defined in (223).

(223) $\textsc{aset}(\text{x}) = \text{I-Mix}(\{\text{x}\}, {}^*\textsc{flat}(f(\textsc{socrel})_x))$

Let's break this down. The ASET of an individual x is an inclusive mixture of the singleton set containing x and the set of sums derived from flattening some contextually selected social relation restricted to only those pairs containing x as its first projection. For instance, a possible ASET for x may be a set containing the atom x as well as the set of sums of x with one or more coworkers, family members, or friends.

We will also define a notion of the *association set of a predicate*, defined as in (224)[5].

(224) $\textsc{aset}(\text{P}) = \text{I-Mix}(\textsc{extract}(\text{P}), {}^*\textsc{flat}(f(\textsc{socrel})_P))$

In this case, the ASET is the mixture of not a singleton set, but of a potentially non-singleton set with a flattened social relation restricted to members of this set. In other words, this set will contain all members of P as well as every possible sum of the members of P with those standing in a contextually salient social relationship with some member of P. A possible ASET for a predicate, then, might be a set of professors and sums of professors and their advisees. Furthermore, because the set of sums is constructed from pairs of individuals standing in a particular relation, we restrict the set to those sums containing sum individuals whose atomic parts are associated with each other in the relevant way, and sums of such sums. As such, this system will not generate sums of individuals that are not themselves associated, even if one individual *is* associated with some element of the set out of which the

[5] I am overloading the ASET notation here by using the same name for functions with a similar effect but which operate over different types of arguments.

ASET is constructed[6].

5.3.2 Analysis of associative plurals

Having constructed all the necessary formal tools, I present my analysis of associative plurals here. I present two very similar analyses, the first for associatives that only occur with definites, as in Afrikaans, and the second for associatives that may combine with predicates and receive indefinite as well as definite interpretations, as in Japanese. A lexical entry for the former type, with Afrikaans *-hulle*, is given in (225).

(225) *-hulle* $\rightsquigarrow \lambda \mathrm{x}.f(\lambda \mathrm{y}.\textsc{aset}(\mathrm{x})(\mathrm{y}) \wedge \neg\textsc{atom}(\mathrm{y}))$

Breaking this down a bit, *-hulle* takes an individual argument x and returns some *non-atomic* individual in the ASET of x. The choice of individual is contextually determined, as reflected in the use of a choice function *f*. *f* applies to the set of non-atomic individuals in x's ASET and returns one of those individuals.

When applied to an expression denoting an atomic individual, the ASET of that individual will contain the individual and all sums of that individual with its associates. However, because of the non-atomic condition, the atomic individual out of which the ASET is constructed will never be selected by *f*. If, on the other hand, (225) is applied to an argument denoting a sum individual, the non-atomic condition will not exclude this individual from its ASET. This means that the individual selected by *f* from the set could very well be the argument of *-hulle* itself, along with any other sum containing that argument as a proper part. This means that a single

[6] The mixture analysis presented here generates set that also contains sum individuals composed of a professor, their associate, and one or more professors not associated with the first professor's associate. While this may appear excessively inclusive, I do not see any empirical disadvantages to this move.

denotation for *-hulle* successfully derives the *strong* homogeneity of associatives with atom-denoting arguments as well as the *weak* non-homogeneity observed with sum-denoting arguments.

Let us turn now to the other type of associative plural, represented by Japanese *-tachi*. This type of associative differs from Afrikaans *-hulle* in permitting indefinite interpretations. I therefore analyze *-tachi* as taking a type <e,t> argument and returning a predicate (also of type <e,t>). I give its translation in (226) below.

(226) *-tachi* $\rightsquigarrow \lambda P.\lambda X.\textsc{aset}(P)(X) \wedge \neg\textsc{atom}(X)$

It is straightforward for *-tachi* to compose with predicates like *kyooju* 'professor(s)' on this analysis: on the assumption that Japanese bare nouns are number-neutral predicates (Chierchia (1998); Nakanishi and Tomioka (2004)), *-tachi* will take this argument and return a predicate of individuals in the set of one or more professors and sums with professors and their associates as parts. The result can then be either existentially-closed by the ∃ type-shifter or made definite by IOTA (Partee, 1987).

(227) *kyooju-tachi* $\rightsquigarrow \lambda X.\textsc{aset}(\lambda Y.{}^*\textsc{professor}(Y))(X) \wedge \neg\textsc{atom}(X)$

What about type e arguments, such as names? Here we invoke the IDENT type-shifter of Partee (1987), familiar from preceding chapters, which shifts an individual-denoting expression to a predicate-denoting one. Upon shifting the individual to type <e,t>, *-tachi* can compose with this, generating that individuals ASET. The result can then be type-shifted via IOTA[7].

[7] In principle, it is also possible for the result to be *existentially closed* as well, despite the fact that *-tachi* is generally felt to be definite by native speakers when it applies to a name. The analysis of Nakanishi and Tomioka (2004) runs into the same conundrum, though they claim that only the IOTA shifter is "pragmatically appropriate" for the use of *-tachi* with definites.

Another way to resolve this apparent issue is to argue that *name* + *-tachi may*, in fact, receive an indefinite reading, and that the apparent definiteness of such constructions is an illusion brought about by the definiteness associated with the name. That is, it is the *name* that presupposes

(228) $[\![\textit{Taro-tachi}]\!] = [\![\textit{-tachi}]\!](\textsc{ident}(t) = \lambda X.\textsc{aset}(\lambda y.y = t) \wedge \neg\textsc{atom}(X)$

Because IDENT, when applied to an individual, returns the singleton set containing that individual, the ASET of an IDENT-shifted individual x will consist of just x and sums containing x and its associates as parts. In other words, ASET(IDENT(x)) = ASET(x), and we achieve the same results in this case as we did with *-hulle* above, modulo the need for type-shifting in the case of *-tachi*.

Given the near-equivalence of the analysis of *-tachi* with names and that of *-hulle* generally, we correctly predict that *-tachi* will be strongly non-homogeneous when combined with an atom-denoting expression like *Taro*. This is due to the fact that the non-atomic condition in the translation of *-tachi* rules out the one atomic individual in Taro's AST, namely, Taro himself, and otherwise leaves only sums of Taro and his associates. With conjunctions of names, such as *Taro to Hanako* 'Taro and Hanako,' we once again derive *weak* non-homogeneity. In the case of conjunctions of names, the reason is exactly the same as for the case with *-hulle* above: the non-atomic condition does not exclude the sum individual denoted by this conjunction, so the *-tachi* phrase may denote a set containing the individual it takes as an argument, along with other larger sum individuals.

We have an additional case, unique to *-tachi*, to consider here: that of *-tachi* applying to number-neutral predicates, such as *kyooju* 'professor(s).' These too exhibit a type of weak non-homogeneity; (216), repeated as (229) below, is true if either at least one professor and their associates came, or if at least two professors came.

(229) Kyooju -tachi -ga ki -ta
professor -ASSOC -NOM come -PST

familiarity on the part of the listener, while the actual group of individuals denoted by the *-tachi* plural need not be similarly familiar. This seems to accord with the intuitions of the Japanese speakers I have consulted, but I will leave more detailed investigation of this idea to future research.

'Professors and their associates came/Professors came.'

Tatsumi (2017) calls this latter reading an "additive" reading, but it is clear that it has much in common with other weak readings of associative plurals, and in fact both readings are compatible (and expected) with the unified analysis of *-tachi* presented above. Consider again the result of applying *-tachi* to the meaning of *kyooju*, as in (230).

(230) $\lambda X.\textsc{aset}(\lambda Y.{}^{*}\textsc{professor}(Y))(X) \wedge \neg\textsc{atom}(X)$

The weak non-homogeneity of (230) once again comes about via an interaction between the non-atomic condition contributed by *-tachi* and the potential plurality of the predicate to which it applies, *PROFESSOR. The ASET of such a predicate will contain every individual in the extension of that predicate, which in this case includes both atomic and sum professors, as well as all sums of those professors with their associates. The non-atomic condition will eliminate the atomic professors, guaranteeing that *-tachi* may not refer to a single professor. However, sums consisting only of professors, as well as sums of professors and their associates, will be left untouched. This derives weak non-homogeneity for these expressions: they may, but need not, denote more than a sum of professors.

5.3.3 Analysis of PPCs

Having analyzed associative plurals, I now turn to the analysis of PPCs. Essentially, I directly extend the analysis of weak and strong non-homogeneity in associative plurals to the same phenomena noted for PPCs, by decomposing a pronoun into two distinct heads: one head for the person feature of the pronoun, and another for the associative. The comitative phrase is serves as the complement of the person head. The following tree in (231) illustrates the intended syntactic decomposition

(231) Decompositional analysis of plural personal pronoun

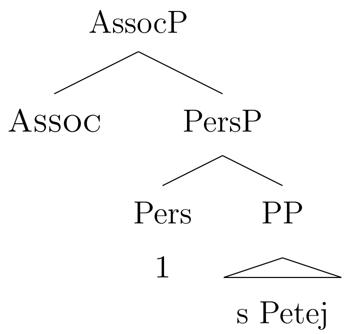

As is standard, I analyze first and second person features as denoting the speaker and the addressee, both of type e, in the context.

(232) $[\![1]\!] = \textit{speaker}$

I analyze the associative morpheme exactly as I analyzed *-hulle*, that is, as selecting an individual from the ASET of its individual argument via a choice function.

(233) $\text{ASSOC} \rightsquigarrow \lambda \text{x}.f(\lambda \text{y}.\text{ASET}(\text{x})(\text{y}) \wedge \neg\text{ATOM}(\text{y}))$

Finally, I analyze the comitative phrase in a PPC as involving sum formation, providing a semantic formalization from the mereological literature (Link (1983), a.o.) of an often suggested analysis of comitative coordination in the literature (Dyla (1988); Progovac (1997)). The example below gives this analysis for the Russian preposition *s*.

(234) $s \rightsquigarrow \lambda \text{y}.\lambda \text{x}.\text{x}\oplus\text{y}$

Putting all of this together, I now present the complete analysis of PPCs. In the absence of a comitative, a PPC denotes an individual selected from the ASET of the individual denoted by the person feature in the syntactic structure of the pronoun. (235) gives the analysis of the Russian first person plural pronoun *my* in the absence of a comitative.

(235) $[\![my]\!] = [\![\textsc{Assoc}(1)]\!] = f(\lambda \text{x}.\textsc{aset}(\textit{speaker})(\text{x}) \wedge \neg\textsc{atom}(\text{x}))$

Here, we derive strong non-homogeneity for the interpretation of the pronoun, which in this case looks like the typical interpretation of a first person pronoun: it will necessarily refer to a sum individual with the speaker as a proper part.

Now consider the case with a comitative phrase. The comitative preposition takes two individuals as arguments and returns their sum. In this case, it takes its complement DP and the person feature as its arguments, returning the sum of the individuals these two phrases denote. (**??**) provides a translation for the PersP in (231) above following this analysis.

(236) $[\![\text{PersP}]\!] = [\![\text{s}(1)(\text{Petej})]\!] = \textit{speaker} \oplus \text{p}$

To complete the analysis, we factor in the role of the ASSOC head. The full structure of *my s Petej* then receives the analysis in (237).

(237) $[\![\textit{my s Petej}]\!] = [\![\textsc{assoc}(\text{PersP})]\!] = f(\lambda \text{x}.\textsc{aset}(\textit{speaker} \oplus \text{p})(\text{x}) \wedge \neg\textsc{atom}(\text{x}))$

As with the analysis of coordinated names with *-hulle*, the non-atomic condition will not exclude the sum of the speaker and Petya. We therefore expect *weak* non-homogeneous behavior: the PPC may refer to the speaker and Petya alone, rather than referring to at least three people.

In this way, we have not only succeeded in deriving the variable behavior of PPCs depending on the presence or absence of a comitative phrase, but we have also achieved a formally explicit theoretical unification of PPCs and associative plurals: the analysis of exclusive and inclusive readings of plural personal pronouns in these constructions is exactly identical to that of strong and weak non-homogeneity in associative plurals, respectively.

5.3.4 Additional predictions of the analysis

The analysis I have proposed for associative plurals and PPCs derives their behavior as either strongly or weakly non-homogeneous plural expressions depending on the nature of the complement of the associative. In addition to deriving the facts it was designed for, the analysis makes two additional predictions. First, it falls out from the analysis that associatives may only take animate arguments. In particular, we predict, in accordance with the typology of associative plurals, that associatives can only be well-formed with arguments that are capable of standing in social relations with other entities. This is because, obviously, the analysis is constructed out of social relations themselves. Any (set of) entities that do not stand in social relations will simply have a vacuous ASET consisting only of itself.

The second prediction is more formal in character: we predict that associative plurals will differ from similative plurals in necessarily denoting sum individuals regardless of the monotonicity of their semantic environment. This follows from the non-atomic predicate in the translation of *-tachi*, which also plays a key role in the derivation of strong and weak non-homogeneity. This prediction is borne out: associative plurals that permit indefinite readings do not have inclusive readings in downward-entailing or non-monotonic contexts. This is demonstrated for *-tachi* in a conditional (238) and polar question (239).

(238)

sensei	-tachi	-ga	ki	-tara,	Yosuke	-wa	ocha	-o	das	-u
teacher	-ASSOC	-NOM	come	-COND	Yosuke	TOP	tea	-ACC	serve	-PRS

'If teachers and their associates come, Yosuke will serve tea.'

(239)

sensei	-tachi	-ga	ku	-ru	no?
teacher	-ASSOC	-NOM	come	-PRS	Q

'Are teachers and their associates coming?'

For (238), Yosuke will only serve tea if at least one teacher and some other individual (either another teacher or an associate, such as the teacher's student, colleague or spouse) comes. Yosuke will not serve tea if only a single teacher comes. Likewise, an affirmative response to (239) is only felicitous if at least two people come, one of whom must be a teacher and the other of whom may be that teacher's associate. If only a single teacher comes, the addressee is expected to answer in the negative.

5.4 On pronouns and associatives

The previous section presented what amounts to an identical semantic analysis for associative plurals and PPCs. Having done this, we can take the analysis a step further by asking how it explains the shared form of associative markers and pronouns cross-linguistically. Some previous analyses, to be discussed further in the next section, propose that associatives are actually pronominal constructions themselves: the associative amounts to a plural pronoun, and its complement is interpreted as *included* in the group denoted by the pronoun (den Besten (1996); Moravcsik (2003); Tatsumi (2017)). For example, den Besten (1996) proposes the following analysis in (240), in which a DP headed by a pronoun contains another DP out of which the associative is constructed.

(240)

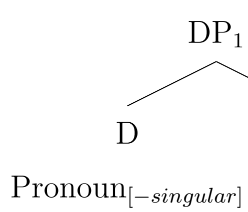

My own analysis allows for a novel perspective on the connection between pronouns and associatives that flips the aforementioned picture on its head: rather than positing that associatives are pronominal, we posit that plural pronouns contain a head encoding associative semantics. More concretely, I propose that plural pronouns spell out associative heads. This can be modeled in a framework like Distributed Morphology (Halle and Marantz (1993); Harley and Noyer (1999), a.o), in which Vocabulary Items are inserted according to context-sensitive rules. The idea here is that associative heads can be Spelled Out as first or second person plural pronouns in the context of a first or second person feature, with the third person plural being the *elsewhere* form. I give an example of this style of analysis is given for the Afrikaans plural pronoun series[8].

(241) [ASSOC] $\Longleftrightarrow$ *ons* / __[1]
[ASSOC] $\Longleftrightarrow$ *julle* / __[2]
[ASSOC] $\Longleftrightarrow$ *hulle* Elsewhere

On this analysis, what associative plurals and pronouns have in common is an underlying associative structure, not an underlying pronoun (whether overt or null). In Afrikaans, then, the ASSOC head will be spelled out as *hulle* in contexts other than when it occurs with a first or second person feature. These will include both third person pronominal contexts as well as in the associative plural construction. Other languages, then, might have an associative marker that is not pronominal in form. This is the case for Japanese *-tachi*, which is straightforwardly analyzed as the exponent of *assoc* in Japanese.

(242) [ASSOC] $\Longleftrightarrow$ *-tachi*

What is crucial here is that we no longer have to posit a pronominal core to associative plurals in order to capture the relation between associatives and PPCs.

[8] There are a few ways to implement this analysis. For instance, one could also have the Pers head undergo head movement to ASSOC, and insert Vocabulary Items according to the features on a single head.

5.5 Previous analyses

In this section, I discuss previous analyses of associative plurals and PPCs. The majority of previous analyses are largely syntactic in nature, with little explicit semantic analysis. This is especially true of analyses of PPCs. Given this state of affairs, I will mostly restrict the discussion to semantic analyses. Furthermore, many of the analyses discussed below deal most explicitly with Japanese *-tachi*, for no reason other than the most explicit semantic proposals are those focused primarily on *-tachi*. That said, I will discuss some more syntactic proposals along with the semantic intuitions behind them in the context of the connection between PPCs and associative plurals, and the analyses of *-tachi* are applicable to associative plurals more generally, as noted by the authors of the works to be discussed below.

5.5.1 Previous approaches to the semantics of associative plurals

The majority of analyses of associative plurals tend to follow their cross-linguistically common restriction to definite complements and an overall definite interpretation. For example, Kurafuji (2004) analyzes Japanese *-tachi* and Mandarin *(ta)men* as inherently definite, combining with a property and returning the maximal individual (notated with the σ operator) satisfying that property.

(243) Kurafuji's (2004) analysis of associative plurals

$$[\![\mathrm{PL}]\!] = \lambda \mathrm{P}.\sigma \mathrm{x}[\mathrm{P(x)}]$$

This analysis works for Mandarin *(ta)men*, which is restricted to definites, and could be extended just as well to other associative plurals, such as Afrikaans *-hulle*. However, it does not fully capture the behavior of Japanese *-tachi*, which permits indefinite interpretations in addition to definite interpretations, as Nakanishi and Tomioka

(2004) show with examples like (244).

(244) kooen -ni kodomo -tachi -ga i -ta
park LOC child -ASSOC NOM be -PST
'There were children in the park.'

Another issue with Kurafuji's analysis is that it fails to account for the non-homogeneity of associative plurals: (243) generates the maximal sum individual of the set it takes as an argument, but that set is not non-homogeneous. For instance, when combined with a predicate like *kyooju* 'professor,' this analysis of *-tachi* will only generate a sum individual in the set of professors. It would not permit a sum individual composed of, say, a professor and their non-professor associate.

Nakanishi and Tomioka (2004) propose a different approach, which aims both to capture the possibility of non-honogeneous plural interpretations of associatives and their compatibility with indefinite interpretations. On their analysis, associatives like *-tachi* receive two distinct, but related translations, depending on whether their argument is of type e or type <e, t>, given in (245).

(245) a. *-tachi* $\rightsquigarrow$ λx.λY.x$\leq$Y & $|$Y$|$ $\geq$ 2 & x represents Y

b. *-tachi* $\rightsquigarrow$ λP.λY.x$|$Y$|$ $\geq$ 2 & P represents Y

On this treatment, associative plurals like *-tachi* are ambiguous between a function that takes an individual as an argument and one that takes a predicate as an argument. Both *-tachi*s share the property of taking an argument and returning a predicate of sum individuals that is in some sense "represented" by the argument of *-tachi*. The "represents" relation is used to encode close association, and allows for the denotation of *-tachi* to be non-homogeneous. Nakanishi & Tomioka provide the following as an example.

(246) CONTEXT: Martians conquered every part of the earth, except for Canada,

and Canadians now await the last assault. The army storming towards Canada actually consists mainly of earthlings, led by a handful of Martians. The Canadians say:

Kaseijin -tachi -ga seme -te ki -ta
Martian -ASSOC -NOM attack -GER come -PST

'Martians came to attack'

Here, the idea is that the plurality doing the attacking is represented by the Martians, but need not be made up of Martians entirely, and in fact my consists mostly of non-Martians.

On this analysis, since *-tachi* denotes a predicate after composing with its argument, its individual argument needs to be closed. This is accomplished by type-shifting, as in my own analysis developed in section 3, using IOTA or $\exists$ to shift the predicate of individuals to an individual or existentially quantified NP, respectively.

This analysis is very close to the approach developed in section 3. Both analyses allow for non-homogeneous interpretations of the plurality, make use of type-shifting operations, and make use of a relation to encode a notion of association. There are three problems with Nakanishi & Tomioka's analysis, in my view. First, Nakanishi & Tomioka's analysis requires two distinct homophonous lexical entries for *-tachi*, one to compose with individuals and another to compose with predicates. Second, the *represents* relation is very vague, and does not clearly place any restrictions on what sorts of relations are at issue with a *-tachi* plural. This means that it is not clear that Nakanishi & Tomioka's analysis captures the nature of associativity as involving social relations. Finally, and related to the last point, the *represents* relation does not clearly place any restrictions on what sorts of individuals or predicates may "represent" the plurality denoted by a *-tachi* phrase. That is, it is not clear how the restriction of *-tachi* to human referents follows from the analysis; why can't a plurality of businessmen be represented by, say, the building in which they work?

The analysis developed in section 3 overcomes each of these three issues. First, it posits only a single *-tachi* that combines with predicates and individuals, the latter by an independently motivated type-shifting operation. Second, the association relation is explicitly defined in terms of a set of contextually determined social relations, thus placing restrictions on what the ASET of the argument of *-tachi* may be. Finally, the arguments of *-tachi* are restricted to those that may stand in social relations in the first place. Given that standing in a social relation implies animacy, it is clear that a building could not appear in an associative plural, due to the fact that buildings do not have relations of the relevant sort.

A final set of analyses, represented by den Besten (1996) and Tatsumi (2017), attempt to analyze associative plurals as involving either overt, as in Afrikaans, or covert, as in Japanese, pronominal structure, as discussed briefly in the previous section. den Besten's proposed structure is as in (240), repeated for convenience in (247), while Tatsumi's analysis is given in (248).

(247)

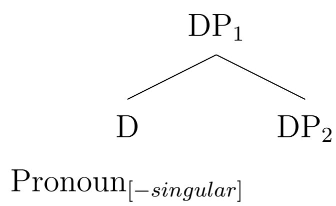

(248)

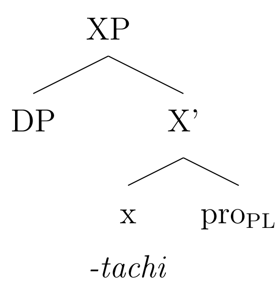

The intuition behind both analyses is that the DP in the specifier of XP is interpreted as a part of the meaning of the pronoun. Tatsumi formulates this semantics explicitly, positing that -*tachi*, when combined with pro_{PL}, takes an individual and returns a predicate of individuals, where the individual is interpreted as a contextually salient part of a group.

(249) $[\![\text{-}\textit{tachi}\ \text{pro}_{\text{PL}}]\!] = \lambda\text{y}.\lambda\text{z}.\sigma\text{v.z} \leq \sigma\text{v[group(v)(c)]} \wedge \text{y} \leq_c \sigma\text{v[group(v)(c)]}$

Both den Besten and Tatsumi assume that associatives may only take definite arguments, though Tatsumi allows the open variable of the translation of the constituent containing -*tachi* and pro to be existentially quantified. The fact that two variables y and z are used in Tatsumi's analysis permits non-homogeneity of the plurality: with a singular proper name, the analysis derives strong non-homogeneity, while with a plural noun phrase, such as 'the professors,' the analysis will permit for weak non-homogeneity.

Although den Besten's analysis is not explicit enough to permit for detailed semantic critique, Tatsumi's proposal is, and it runs into a few problems that the analysis developed in this chapter does not. First, his analysis, much like Nakanishi & Tomioka's, is vague about the restrictions on individuals that -*tachi* takes as an argument: in principle, there is nothing stopping -*tachi* from taking an inanimate individual argument. The notion of a group, utilized in the analysis in (**??**) does not provide such a restriction, as groups may be composed of inanimate objects, nor does the notion of "contextually salient part" place any restriction on the argument of -*tachi*. Second, and also like Nakanishi & Tomioka's analysis, Tatsumi's approach is not explicit about how the atomic parts of the plurality are related to each other: although groups are typically associated with each other in some way, one can also pick out arbitrary groups, such as individuals who happen to be standing next to each other in a line, with no relation to each other at all. The ASET notion developed in my own analysis, on the other hand, solves both problems, as it is constructed on the basis of a restriction to individuals standing in social relations, and thereby predicts a restriction to animate individuals.

A more serious issue with Tatsumi's analysis is that it predicts that associative plurals with apparently predicative arguments must be type-shifted to individuals. For example, Tatsumi proposes that a predicate like *gakusei* 'student(s)' is shifted by a silent 'def', which returns the largest sum individual in the set of students. The idea, then, is that *gakusei-tachi* will refer to an individual of which all the students are a salient part.

(250) a. $[\![\textit{gakusei}]\!] = \lambda$x.student(x)

b. $[\![\text{def } \textit{gakusei}]\!] = \sigma$x[student(x)]

This, however, does not appear to be a requirement on *-tachi*: as (Nakanishi and Tomioka, 2004) note, *-tachi* plurals formed from common nouns may be indefinite, and need not refer to any salient set of individuals. This is compatible with both Nakanishi & Tomioka's analysis, as well as my own analysis. Tatsumi's only support for his proposal that common noun arguments of *-tachi* need to be definite comes from the fact that *-tachi*, as well as other plurals in Japanese, is incompatible with wh-words like *dare* 'who', but a plural can be formed from *doitsu* 'which person'.

(251) a.

*Mary	-wa	dare	-ra/-tachi	-o	kirat	-teiru	no
Mary	TOP	who	WHO	-PL/-ASSOC	hate	-ASP	Q

Intended: 'Who does Mary hate?'

b.

?Mary	-wa	doitsu	-ra	-o	kirat	-teiru	no
Mary	-TOP	which.person	-PL	-ACC	hate	-ASP	Q

'Which people does Mary hate?'

The problem with this argument is twofold. First, the two constrasting sentences are not fully minimal pairs: while *dare* is shown to be incompatible with plural markers generally, *doitsu* is shown only with *-ra*, a plural marker which may not be identical in meaning to *-tachi*. Second, definiteness is not at issue in this example: both *dare* and *doitsu* are **wh-indefinites**. The contrast may instead be related to the fact that, even in a language like English, *who* cannot be pluralized, while D-linked wh-phrases like *which person/which people* can.

A final point relating to analyses like den Besten's and Tatsumi's concerns the appeal to pronominal structure in associative plurals. As noted in section 4, the connection between pronominal constructions like PPCs and associatives need not be thought about in terms of underlying pronominal structure, but may instead be thought of in terms of a shared underlying *associative* structure. This removes the motivation for a covert pronoun in *-tachi* plurals, and permits a rethinking of den Besten's syntactic proposal for Afrikaans associatives.

5.5.2 Previous analyses of PPCs

Previous formal approaches to PPCs are primarily syntactic in nature (Dyla (1988); den Dikken et al. (2001); Vassilieva (2005)), and are largely compatible with the semantic analysis I have developed for them here. Previous semantic proposals start with the intuition that plural personal pronouns are composed of two parts: a part identical to the meaning of the plural pronoun's singular counterpart, and a contextually specified "remainder". On this view, then, plural personal pronouns are "incomplete expressions" (Vassilieva and Larson (2005); Vassilieva (2005), Siggurdson and Wood (2020)), as schematized in (252).

(252) $we = I +$ others

you(pl) = *you*(sg) + others

they = *he/she/it* + others

Vassilieva and Larson (2005) assume this general approach, and then adopt a special set of interpretive rules for plural pronouns, making use of the interpretive schema from Larson and Segal (1995). The same analysis is proposed for Icelandic PPCs by Siggurdson and Wood (2020).

(253) $\mathrm{Val}(<\mathrm{X},\mathrm{Y}>,[\mathit{my}],\sigma)$ iff $|\sigma(\mathrm{a}) \cup \mathrm{Y} - \mathrm{X}| = 0$

Here, σ(a) picks out a set containing the speaker of the sentence, X represents a set of individuals, and Y is a free variable for a set. The idea behind this is that the speaker and whatever Y contains is all that is in the set. This will derive the 'speaker + others' reading of the Russian first person plural pronoun *my.* In order to analyze the contribution of the comitative, Vassilieva & Larson introduce two additional rules, one permitting a comitative to be analyzed as a single individual, and another for analyzing a DP containing a pronoun and a PP.

(254) Val(Y, [s Petej] iff Y = {Peter}

(255) $\mathrm{Val}(\mathrm{X}, [\mathrm{D\ PP}],\sigma)$ iff $\mathrm{Val}(<\mathrm{X},\mathrm{Y}>,\mathrm{D}],\sigma)$ & $\mathrm{Val}(\mathrm{Y}, \mathrm{PP}, \sigma)$

With these rules, it is possible to interpret a plural personal pronoun with a comitative complement inclusively; (**??**) derives an interpretation where *my s Petej* contains only the speaker and Peter.

(256) Val(X, *my s Petej*, σ) iff $|(\sigma(\mathrm{a}) \cup \{\mathrm{Peter}\}) - \mathrm{X}| = 0$

While this style of approach is able to derive the strong and weak non-homogeneity of PPCs, it does so using very *ad hoc* semantic interpretation rules: rules are invoked that are particular to comitative phrases and constituents containing pronouns *qua*

D heads and prepositional phrases. Moreover, the analysis is wholly unconnected to associative plurality. The analysis I developed in section 3, on the other hand, makes no use of such special interpretation rules; the analysis uses only standard compositional rules from the formal semantics literature, and is able to account for strong and weak non-homogeneity just as well. Furthermore, my own analysis ties together associative plurals and PPCs using an identical analysis in a unified theoretical framework.

5.6 Conclusion: The common core of associative and similative plurality

In this chapter, I proposed a unified analysis of associative plurals and PPCs, constructing appropriate associate sets for each from underlying social relations and deriving the strong and weak non-homogeneity they display depending on the nature of the expression they take as an argument. The analysis was shown to make additional correct predictions about distributional restrictions on associatives and the behavior of associative plurals in non-upward-entailing contexts, and allows for a novel approach to the connection between associative markers and pronouns that does not require associative plurals to themselves be or contain pronouns in their syntactic structure. Finally, the analysis improves on previous approaches by allowing for a single analysis of associative markers like *-tachi* where previous analyses have posited ambiguities, and proposing a fully explicit unified semantics for associatives and PPCs where others have left the details vague.

CHAPTER 6

CONCLUSION

This dissertation has developed a general framework for reasoning about non-homogeneous plurality. Starting with similative plurality, it developed a formal model of the calculation of similarity in context, out of which a similarity relation was constructed. I then showed how to put these tools to use in the analyses of Persian and Japanese similative plurals, and further showed that the semantic and pragmatic properties of these plurals require recourse to local implicature calculation and abstract alternatives. Finally, I applied the mixture semantics developed for similative plurals to the analysis of associative plurals.

Having accomplished this much, I conclude with an elaboration of the connection between the phenomena investigated throughout this dissertation, discussion on how alternatives in the analysis of similatives plurals may be constrained, and a discussion of directions for future research.

6.1 The common core of non-homogeneous plurality

In concluding this chapter, I would like to elaborate on the connections between the analysis of associatives and PPCs presented here and that of similative plurals developed in the preceding two chapters. Although the analyses of the two phenomena seem superficially distinct–for instance, one is an inclusive plural that receives exclusive readings in upward-entailing contexts, while the other is semantically exclusive–there is a deep underlying core to both types of non-homogeneous plural on the view

I have proposed here: both similative and associative plurality involve a *mereological mixture* of two sets, where one set is constructed out of individuals bearing a contextually determined *relation* with (individuals in) the other set. For similative plurals, the relation at issue is a *similarity relation* between either individuals or functions from individuals to truth values, and the plural is derived from a mixture of a set and the *similarity set* constructed on the basis of the similarity relation. For associative plurals, the relation at issue is a contextually determined *social relation*, and the mixture is the *associate set* derived by mixing a set with another set constructed on the basis of the selected social relation. This means, then, that the analyses proposed throughout this dissertation provide a unified framework for reasoning about non-homogeneous plurality as a phenomenon, with similative plurals forming one class of non-homogeneous plurals and associative plurals and PPCs forming another class.

6.2 Symmetry and contraining abstract alternatives

The analysis I have developed makes crucial use of *abstract* alternatives, which are derived from a non-linguistic conceptual level of representation, rather than by deletion and replacement operations over syntactic structure restricted by the lexicon of the language in question. Since these alternatives are not constrained by lexicalized resources, they must be constrained in some other way. This is crucial because there are plausible alternatives to m-reduplication that, if derivable, would halt the calculation of any implicature.

As an example, consider the case of Mix(P,$\sim$(P)). On the analysis I proposed in chapters 3 and 4 of this dissertation, the alternatives to such a mixture are each of the (cumulative closures of the) arguments of the mixture function, *P and *$\sim$(P). Excluding these alternatives is sufficient to generate the exclusive mixture reading of m-reduplication in Persian and *-toka* and *-toka* in Japanese. However, if an expression that corresponds to the reading we would like to derive via implicature, E-Mix(P,∼(P)) in this case, is available as an alternative via a replacement operation on conceptual representations, the analysis will fail: at worst, we fail to generate any implicature at all, since the exclusion of such an alternative in combination with the other alternatives will result in a contradiction of the assertion. This is an instance of the *symmetry problem*, whereby an alternative corresponding to a strengthened reading of a sentence cannot be excluded from competition with that sentence (Katzir (2007); Fox and Katzir (2011)).

The structural approach to alternatives advocated by Katzir (2007) avoids the symmetry problem by restricting operations on syntactic structures to replacement and deletion operations. This means that alternatives can only be *at most* as structurally complex as the structure from which they are derived. For a sentence like *John ate some of the cookies*, then, alternatives of equal complexity like *John ate all of the cookies* are ruled in, while the symmetry-inducing alternative *John ate some but not all of the cookies* is underivable, due to its structural complexity exceeding that of the sentence under consideration.

As I have argued against the structural approach to alternatives in this dissertation, this solution to the symmetry problem is not available to my analysis. However, a version of it, a sort of *conceptual complexity*, may be applicable. On such an approach, alternatives to ϕ must be *at most as complex as* ϕ in some sense. One way to operationalize this in this domain is to consider the relative complexity of different types of mixture operations defined in chapter 3. Recall that the inclusive mixture of two sets is equivalent to the cumulative closure of the union of those sets, as in (257).

(257) $\text{I-Mix}(P,Q) = {}^*(P \cup Q)$

Exclusive mixtures, on the other hand, are more complex, being equivalent inclusive mixtures with the cumulative closure of each set used to form the inclusive mixture removed.

(258) $\text{E-Mix}(P,Q) = {}^*(P \cup Q) \setminus ({}^*P \cup {}^*Q)$

In this way, we can think of E-Mix as more complex than I-mix, in that its definition is of a greater set-theoretic complexity than that of I-Mix. In this way, we could exclude E-Mix as a viable alternative to I-Mix.

In a related vein, Buccola et al. (2018) offer *primitiveness* as a constraint on alternatives. On this view, the language of thought contains *primitive concepts*, out of which more complex, composite concepts may be constructed. Solving the symmetry problem here amounts to stating what is a primitive: for instance, the conceptual counterparts of *some* and *all* are primitives, while the counterpart of *some but not all* is not, and must be constructed from other primitives, like $\exists$, $\forall$, $\wedge$, and $\neg$. The symmetry problem may be solved on this approach by proposing that alternatives must be drawn from lexical items of the *language of thought*. In this way, $\exists$ and $\forall$, being primitives, are available alternatives to $\exists$, while the LOT counterpart of *some but not all*, being non-primitive, is not available. Alternatively, Buccola et al. (2018) propose a graded view of alternativehood, where alternatives drawn from non-primitive concepts are accompanied by a greater cost than those drawn from primitive concepts.

Applying these ideas to the domain of mixture semantics, one possibility is that inclusive mixtures are primitive, while exclusive mixtures are not. We might then expect, in line with the discussion of Buccola et al. (2018), that exclusive mixtures are rarely, if ever lexicalized.[1]

6.3 Future Directions

This dissertation serves as an initial investigation into a number of phenomena, which may be elaborated on in future research. For one, there is the question of how similative expressions may be constrained cross-linguistically. The formalism in chapter 2 restricts evaluation of similarity to individual and predicative expressions, but is not defined over more complex objects, such as quantifiers. One might expect, then, to find languages that permit similatives to be able to apply to predicative expressions like nouns/noun phrases and verbs/verb phrases, but not to, say, quantificational determiners or quantificational NPs.

A particularly interesting area for future study is the *structure* of similative plurals cross-linguistically. The languages I have investigated in this dissertation form similative plurals in different ways, with Persian and other languages of Western and Southern Asia forming them via reduplication, and Japanese forming them with dedicated suffixes. Because the similarity set does not correspond to an overt expression in the language, I argued that an analysis appealing to abstract alternatives is necessary. However, one could imagine a logically possible language that forms similatives by means of expressions corresponding to the similarity set. If such a language exists, it may be amenable to an analysis avoiding appeal to abstract alternatives. Such an analysis could then be extended to the cases I have considered in this dissertation.

[1]In fact, the context in which exclusive mixture was originally proposed, as *Set Product* in Heycock and Zamparelli (1999), has since been reanalyzed by Champollion (2015), who does not rely on an exclusive mixture operation to derive the interpretation of conjoined NPs

REFERENCES

Armoskaite, S. and A. Kutlu (2013). Category neutral similative: Evidence from Turkish. Ms.

Barwise, J. and R. Cooper (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy*, **4**, pp. 159–219.

Buccola, B., M. Kriz, and E. Chemla (2018). Conceptual alternatives: competition in language and beyond. Ms., retrieved from Lingbuzz: lingbuzz/003208.

Champollion, L. (2010). *Parts of a whole: Distributivity as a bridge between aspect and measurement.* Ph.D. thesis, The University of Pennsylvania.

Champollion, L. (2015). Ten men and women got married today. *Journal of Semantics*, **33(3)**, pp. 561–622.

Charlow, S. (2019). Scalar implicature and exceptional scope. Ms., retrieved from Lingbuzz: lingbuzz/003181.

Chemla, E. (2007). French both: a gap in the theory of antipresupposition. *Snippets*, **15(1)**, pp. 4–5.

Chierchia, G. (1998). Reference to kinds across language. *Natural Language Semantics*, **6(4)**, pp. 339–405.

Chierchia, G. (2004). Scalar implicatures, polarity phenomena, and the syntax/pragmatics interface. In Belletti, A. (ed.) *Structures and beyond.* Oxford University Press.

Chierchia, G. (2006). Broaden your views: implicatures of domain widening and the "logicality" of language. *Linguistic Inquiry*, **37(4)**, pp. 535–590.

Chomsky, N. (1975). Questions of form and interpretations. In Austerlitz, R. (ed.) *Scope of American Linguistics*, pp. 158–196. Lisse: The Peter De Ridder Press.

Dayal, V. (2011). Hindi pseudoincorporation. *Natural Language and Linguistics Theory*, **29(1)**, pp. 123–167.

de Mey, S. (1981). The dependent plural and the analysis of tense. In Burke, V. A. and J. Pustejowsky (eds.) *Proceedings of NELS 11*, pp. 58–78.

de Swart, H. and D. Farkas (2010). The semantics and pragmatics of plurals. *Semantics and Pragmatics*, **3(6)**, pp. 1–54.

den Besten, H. (1996). Associative DPs. *Linguistics in the Netherlands*, **13(1)**, pp. 13–24.

den Dikken, M. (2001). Pluringulars, pronouns and quirky agreement. *The Linguistic Review*, **18(1)**, pp. 19–41.

den Dikken, M., A. Liptak, and Z. Zsolensky (2001). On inclusive reference anaphora: new perspectives from Hungarian. In Megerdoomian, K. and L. Bar-el (eds.) *Proceedings of WCCFL 20*, pp. 137–149. Somerville: Cascadilla Press.

Dyla, S. (1988). Quasi-comitative coordination in Polish. *Linguistics*, **26**, pp. 383–414.

Fereshteh, M. (2015). *Bare nouns in Persian: Interpretation, Grammar, and Prosody*. Ph.D. thesis, University of Ottawa and Humboldt-Universitat zu Berlin.

Fodor, J. (1975). *The language of thought*. Harvard University Press.

Fox, D. (2007). Free choice and the theory of scalar implcatures. In *Presupposition and implicature in compositional semantics*, pp. 71–112. Palgrave Macmillan UK.

Fox, D. and R. Katzir (2011). On the characterization of alternatives. *Natural language semantics*, **19(1)**, pp. 87–107.

Geurts, B. (2010). *Quantity implicatures*. Cambridge University Press.

Ghomeshi, J. (1997). Topics in Persian VPs. *Lingua*, **102(2-3)**, pp. 133–167.

Ghomeshi, J. (2003). Plural marking, indefiniteness, and the noun phrase. *Studia linguistica*, **57(2)**, pp. 47–74.

Ghomeshi, J. (2018). The associative plural and related constructions in Persian. In Korangy, A. and C. Miller (eds.) *Trends in Iranian and Persian Linguistics*, pp. 233–262. Berlin: Mouton de Gruyter.

Grice, H. (1975). Logic and conversation. In et al., C. (ed.) *Syntax and semantics 3: speech act*, pp. 41–58. Elsevier.

Halle, M. and A. Marantz (1993). Distributed morphology and the pieces of inflection. In Hale, K. and S. J. Keyser (eds.) *The view from Building 20*, pp. 111–176. Cambridge, Mass: MIT Press.

Hamblin, C. (1973). Questions in Montague English. *Foundations of language*, **10(1)**, pp. 41–53.

Harley, H. and R. Noyer (1999). Distributed morphology. *Glot international*, **4(4)**, pp. 3–9.

Hayashishita, J. and D. Bekki (2011). Conjoined nominal expressions in Japanese. In *JSAI International Symposium on Artificial Intelligence*, pp. 54–67. Berlin, Heidelberg: Springer.

Heycock, C. and R. Zamparelli (1999). Toward a unified analysis of DP conjunction. In Dekker, P. (ed.) *Proceedings of the Twelfth Amsterdam Colloquium*, pp. 127–132.

Heycock, C. and R. Zamparelli (2000). Friends and colleagues: Plurality, coordination, and the structure of DP. *Natural Language Semantics*, **13(3)**, pp. 201–270.

Horn, L. (1972). *On the semantics properties of the logical operators in English*. Ph.D. thesis, University of California, Los Angeles.

Horn, L. (1984). Towards a new taxonomy for pragmatic inference: Q-based and R-based implicature. In Shiffrin, D. (ed.) *Meaning, form, and use in context: linguistic applications*, pp. 11–42. Washington: Georgetown University Press.

Jaccard, P. (1901). Étude comparative de la distribution florale dans une portion des Alpes et des Jura. *Bulletin de la Société vaudoise des sciences naturelles*, **37**, pp. 547–579.

Jasbi, M. (2014). The semantics of differential object marking in Persian. Qualifying paper, Stanford University.

Jasbi, M. (2015). Semantics of differential object marking in Persian. In *Proceedings of the 89th meeting of the Linguistic Society of America*.

Karimi, S. (1999). Specificity Effect: Evidence from Persian. *The Linguistic Review*, **16**, pp. 125–142.

Karimi, S. and R. W. Smith (2019). Another look at Persian -*râ*: a single formal analysis of a multifunctional morpheme. In Larson, R., S. Moradi, and V. Samiian (eds.) *Advances in Iranian Linguistics Vol. 1*. Amsterdam: John Benjamins.

Katzir, R. (2007). Structurally-defined alternatives. *Linguistics and Philosophy*, **30(6)**, pp. 669–690.

Keenan, E. and J. Stavi (1986). A semantic characterization of natural language determiners. *Linguistics and Philosophy*, **9(3)**, pp. 253–326.

Kratzer, A. (1998). Scope or pseudoscope? Are there wide scope indefinites? In Rothstein, S. (ed.) *Events and grammar*, pp. 163–196. Dordrecht: Kluwer Academic Publishers.

Kratzer, A. and J. Shimoyama (2002). Indeterminate Pronouns: The view from Japanese. In *3rd Tokyo conference on psycholinguistics*. Yukio Otsu.

Krifka, M. (2004). Bare NPs: Kind-referring, Indefinites, Both, or Neither? In Bonami, O. and P. Hofherr (eds.) *Empirical Issues in Formal Syntax and Semantics*, volume 5, pp. 111–132.

Kurafuji, T. (2004). Plural morphemes, definiteness, and the notion of semantic parameter. *Language and Linguistics*, **4**, pp. 211–242.

Landman, F. (2000). *Events and Plurality: the Jerusalem Lectures*. Kluwer Academic Publications.

Larson, R. and G. Segal (1995). *Knowledge of meaning: an introduction to semantic theory*. Cambridge, MA: MIT Press.

Levandowsky, M. and D. Winter (1971). Distance between sets. *Nature*, **234(5323)**, pp. 34–35.

Link, G. (1983). The logical analysis of plurals and mass terms: a lattice-theoretical approach. In Bauerle, R., C. Schwarze, and A. von Stechow (eds.) *Meaning, use and interpretation of language*, pp. 303–323. de Gruyter.

Massam, D. (2001). Pseudo noun incorporation in Niuean. *Natural Language and Linguistics Theory*, **19(1)**, pp. 153–197.

Montague, R. (1973). The proper treatment of quantification in ordinary English. In *Approaches to natural language*, pp. 221–242. Springer, Netherlands.

Moravcsik, E. (2003). A semantics analysis of associative plurals. *Studies in Languages*, **27(3)**, pp. 469–503.

Nakanishi, K. and S. Tomioka (2004). Japanese Plurals are Exceptional. *Journal of East Asian Studies*, **13.2**, pp. 113–140.

Partee, B. (1985). "Dependent plurals" are distinct from bare plurals". Ms., University of Massachusetts, Amherst.

Partee, B. (1987). Noun Phrase Interpretation and Type-Shifting Principles. In *Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers*, pp. 115–143. Foris, Dordrecht.

Partee, B. and M. Rooth (1983). Generalized conjunction and type ambiguity. In Bauerle, R., C. Schwarze, and A. von Stechow (eds.) *Meaning, use and interpretation of language*, pp. 334–356. de Gruyter.

Progovac, L. (1997). Structure for coordination. *Glot International*, **3(7)**, pp. 3–6.

Reinhart, T. (1997a). Quantifier scope: How labor is divided between QR and choice fuctions. *Lingustics and Philosophy*, **20(4)**, pp. 335–397.

Reinhart, T. (1997b). Quantifier scope: how labor is divided between QR and choice functions. *Linguistics and Philosophy*, **20(4)**, pp. 335–397.

Siggurdson, H. and J. Wood (2020). We Olaf: Pro x-NP constructions in Icelandic and beyond. *Glossa: A Journal of General Linguistics*, **5(1)**, p. 16.

Smith, R. W. (2019). Non-uniform plural inferences as higher-order scalar implicature. In Espinal, M. T., E. Castroviejo, M. Leonetti, L. McNally, and C. Real-Puigdollers (eds.) *Proceedings of Sinn und Bedeutung 23*, pp. 305–316.

Smith, R. W. and R. Kobayashi (2018). Alternating conj/disjunctions: the case of Japanese -toka and -tari. In Sauerland, U. and S. Solt (eds.) *Proceedings of Sinn und Bedeutung 22*, pp. 393–406.

Spector, B. (2007). Aspects of the pragmatics of plural morphology: on higher-order implicatures. In *Presupposition and implicature in compositional semantics*, pp. 243–281. Palgrave Macmillan UK.

Stanley, J. and Z. Gendler Szabó (2000). On quantifier domain restriction. *Mind & Language*, **15(2-3)**, pp. 219–261.

Sudo, Y. (2014). Higher order implicatures of -ya in Japanese. Handout from TEAL 9.

Tatsumi, Y. (2017). A compositional analysis of plural morphemes in Japanese. In Erlewine, M. (ed.) *Proceedings of GLOW in Asia XI*, pp. 233–241. Cambridge, MA: MIT Working Papers in Linguistics.

Vassilieva, M. (2005). *Associative and pronominal plurality*. Ph.D. thesis, Stony Brook University.

Vassilieva, M. and R. Larson (2005). The semantics of the plural pronoun construction. *Natural Language Semantics*, **13(2)**, pp. 101–124.

von Fintel, K. (1994). *Restrictions on Quantifier Domains*. Ph.D. thesis, University of Massachusetts, Amherst.

Windfuhr, G. (ed.) (2009). *The Iranian Languages*. Psychology Press.

Winter, Y. (1997). Choice functions and the scopal semantics of indefinites. *Linguistics and Philosophy*, **20(4)**, pp. 399–467.

Winter, Y. (2016). *Elements of formal semantics: an introduction to the mathematical theory of meaning in natural language*. Edinburgh University Press.

Zweig, E. (2009). Number-neutral bare plurals and the multiplicity implicature. *Linguistics and Philosophy*, **7(2)**, pp. 353–407.