ACCESSIBILITY, LANGUAGE PRODUCTION, AND LANGUAGE CHANGE

by

ZARA HARMON

A DISSERTATION

Presented to the Department of Linguistics
and the Graduate School of the University of Oregon
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

June 2019

DISSERTATION APPROVAL PAGE

Student: Zara Harmon

Title: Accessibility, Language Production, and Language Change

This dissertation has been accepted and approved in partial fulfillment of the requirements for the Doctor of Philosophy degree in the Department of Linguistics by:

Vsevolod Kapatsinski	Chairperson
Melissa M. Baese-Berk	Core Member
Eric Pederson	Core Member
Dare Baldwin	Institutional Representative

and

Janet Woodruff-Borden	Vice Provost and Dean of the Graduate School

Original approval signatures are on file with the University of Oregon Graduate School.

Degree awarded June 2019.

DISSERTATION ABSTRACT

Zara Harmon

Doctor of Philosophy

Department of Linguistics

May 2019

Title: Accessibility, Language Production, and Language Change

This dissertation explores the effects of frequency on the learning and use of linguistic constructions. The work examines the influence of frequency on form choice in production and meaning inference in comprehension and discusses the effect of each modality on diachronic patterns of change in language. In production, high frequency of a form increases its accessibility given its meaning, and other related meanings. Under the pressures of online real-time speech production, greater accessibility makes a frequent form more likely to be selected over its competitors. Consequently, frequent forms are extended to novel meanings in production, resulting in a synchronic correlation between frequency and polysemy. At the same time, frequency in comprehension results in entrenchment—the more often a form is experienced with a meaning, the more confident the learner becomes that the form is unlikely to be used to express other meanings. The findings reconcile two seemingly contradictory effects of frequency in language change and language acquisition. While frequency results in extension of a frequent form to other meanings in production, it can, at the same time, cause entrenchment in comprehension, which curbs over-extension. The struggle between the pressures from production to extend and from comprehension to entrench molds language. I further provide

experimental evidence demonstrating that frequent forms push their infrequent competitors out of their shared meanings, and that infrequent forms competing with frequent forms tend to be assigned to novel related meanings in comprehension. This result suggests a mechanism for the survival of infrequent forms in specific niches and the existence of push chains in semantic change.

This dissertation includes previously published co-authored material.

CURRICULUM VITAE

NAME OF AUTHOR: Zara Harmon

GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:

University of Oregon, Eugene OR
University of Tehran, Tehran, Iran

DEGREES AWARDED:

Doctor of Philosophy, Linguistics, University of Oregon
Master of Arts, Linguistics, University of Oregon
Master of Arts, Linguistics, University of Tehran

AREAS OF SPECIAL INTEREST:

Planning, retrieval, and execution in language production, disfluencies, learning theory, Associative and Bayesian models of learning, usage-based approaches to language change, language acquisition

PROFESSIONAL EXPERIENCE:

Graduate Teaching Fellow, Department of Linguistics, University of Oregon
Adjunct ESL Instructor, American English Institute, University of Oregon
Assessment Coordinator and Consultant, Center for Applied Second Language Studies (CASLS), University of Oregon

GRANTS, AWARDS, AND HONORS:

Betty Foster McCue Scholarship for Human Development and Performance
Institute of Cognitive and Decision Sciences Dissertation Award
Graduate School Special OPPS Travel and Research Award

Linguistic Society of America Summer Institute Fellowship, 2017

Gary E. Smith Summer Professional Development Award

Jacqueline Schachter Outstanding Conference Presentation Award

Institute of International Education Fulbright FLTA Scholarship

PUBLICATIONS:

Harmon, Z., & Kapatsinski, V. (under revision). A theory of repetition and retrieval in sequential behavior.

Harmon, Z., Idemaru, K., & Kapatsinski, V. (2019). Mechanisms of learning in cue reweighting. *Cognition*, *189*, 76–88.

Harmon, Z., & Kapatsinski, V. (2017). Putting old tools to novel uses: The role of form accessibility in semantic extension. *Cognitive Psychology*, *98*, 22–44.

Kapatsinski, V. & Harmon, Z. (2017). A Hebbian account of entrenchment and (over)-extension in language learning. *Proceedings of the Annual Conference of the Cognitive Science Society, 39,* 2366–2371.

Harmon, Z., & Kapatsinski, V. (2016). Fuse to be used: A weak cue's guide to attracting attention. *Proceedings of the Annual Conference of the Cognitive Science Society, 38,* 520–525.

Harmon, Z., & Kapatsinski, V. (2016). Determinants of lengths of repetition disfluencies: Probabilistic syntactic constituency in speech production. *Chicago Linguistic Society*, *50*, 237–248.

Harmon, Z., & Kapatsinski, V. (2015). Studying the dynamics of lexical access using disfluencies. *Proceedings of Disfluencies in Spontaneous Speech*, *7*, 41–44.

ACKNOWLEDGEMENTS

This dissertation has benefitted immensely from the guidance and scholarship of my advisor, Volya Kapatsinski. I have interacted with few with more sense of curiosity, desire for exploring interdisciplinary ideas, and elegant approaches to constructing an argument. But, the most relevant to the work in this dissertation is Volya's breadth and depth of knowledge as a linguist and a cognitive scientist. It was from Volya that I learned that investigating the dynamic processes of development, learning, and change is the only way to study language. Throughout the years, he has generously shared his insight in all these areas and has always pointed me to relevant literature. Volya introduced me to learning theory, and the importance of computational modeling in spelling out theoretical proposals. I regret that these few words will not do justice to his tremendous contribution to my training in the past few years. I feel deeply fortunate to have had him as a mentor, but even more so to call him a friend. I am and will always be inspired by his inquisitiveness and generosity.

I am indebted to Eric Pederson for his support since the beginning of my graduate program. Eric has so generously shared ideas and fostered a learning environment for everyone in the department. He provided valuable feedback both on this dissertation and other work. I sincerely thank him for his support, insight, and above all for always being a trusted counsel. I would like to also thank Melissa Baese-Berk for helpful comments and discussion of the ideas in the dissertation and the first draft of the paper which turned into this dissertation. Finally, I would like to express gratitude to Dare Baldwin for insightful comments and discussion of ideas in the dissertation and for always being encouraging no matter how brief our conversation. I appreciate her feedback and advice.

Special thanks to Amy Smolek, Paul Olejarczuk, Matt Stave, and Manuel Otero for all the pleasant memories in the past few years inside and outside the lab. I have learned with you and from you on many occasions, and I appreciate that. I am indebted to Amy for patiently recording all the stimuli for the experiments reported in this dissertation. I would also like to thank everyone who shared ideas during my presentations at Cognitive Linguistics Workgroup, but specially Charlotte Vaughn for constructive comments on the first draft of the paper and Ted Bell for helpful discussions of my ideas.

Parts of this work have been presented at Interdisciplinary Advances in Statistical Learning, the Boston University Conference on Language Development, the Meaning in Flux Workshop, the International Cognitive Linguistics Conference, the American International Morphology Meeting, the Quantitative Morphology Meeting, and the

annual meetings of the Linguistic Society of America, Societas Linguistica Europaea, and the Cognitive Science Society. I would like to thank the audience of these conferences for their valuable feedback. This work has received funding from the Betty Foster McCue Scholarship for Human Development and Performance from the University of Oregon Graduate School and the Institute of Cognitive and Decision Sciences Dissertation Award.

I thank my family, especially mom, dad, and Sarah, as well as Jan and Mark Harmon for their unconditional love and support throughout my graduate studies.

Finally, I would like to say thank you to Russell Harmon. Russell encouraged me to do what I love and made sure I prioritized what contributed most to my success above anything else. Having him by my side has filled every day of my life with an ineffable sense of joy and comfort. I would not be where I am without his unending reassurance and encouragement to focus on accomplishing what I set out to do when I started my program.

To *Russell*

TABLE OF CONTENTS

Chapter Page

Chapter Page

Chapter Page

LIST OF FIGURES

Figure Page

Figure Page

Figure Page

LIST OF TABLES

Table Page

CHAPTER I

INTRODUCTION

This chapter provides an overview of the dissertation. I briefly review the areas of concern and the problem of interest and spell out how the current work contributes to solving that problem. Next, the theoretical framework and assumptions that shape this work are discussed. Research questions and hypotheses corresponding to each are also detailed. The chapter concludes with an outline of the chapters in the dissertation.

1.1 Areas of concern

In language, forms that are frequent have more meanings (Piantadosi et al., 2012; Reder, Anderson, & Bjork, 1974; Schnorr & Atkinson, 1970; Zipf, 1949). Excluding 2 phrasal verbs and 14 idioms, the American Heritage Dictionary lists 15 independent meanings for the verb *have* but only five independent meanings for the verb *possess*, two of which are considered archaic. This is one of many examples that demonstrates a correlation between frequency and number of meanings, highlighted in Figure 1.1. Providing a mechanistic account of this relationship is one of the major aims of the current work.

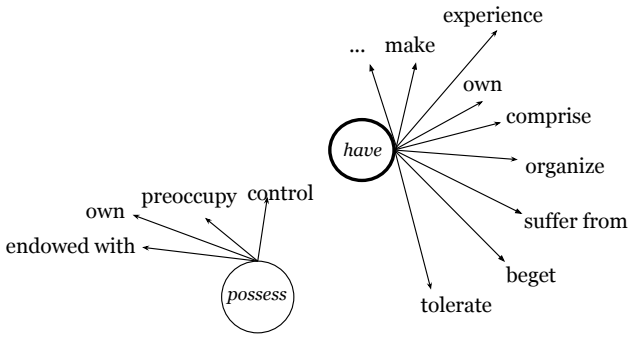

Figure 1.1. The correlation between frequency and polysemy, with *have* representing a frequent highly polysemous form.

The correlation can be attributed to two possible diachronic pathways. The first is that forms become frequent because they are used in many contexts by virtue of having different meanings (Haspelmath, 1999). The second is that forms are extended to new meanings because they are frequent (Bybee, 2003; Zipf, 1949). These pathways are not mutually exclusive, but the second pathway is controversial.

Furthermore, if frequency can cause extension to new uses, the mechanism by which it does so is a matter of controversy. First, frequency can cause semantic broadening, perhaps, via the mechanism of habituation (Bybee, 2003; see also Lambert & Jakobovits, 1960; Smith & Klein, 1990). This broadening would then make the form compatible with more contexts, allowing for new uses to emerge. Second, frequency can encourage the use of a form in a new context by making the frequent form more accessible in that new context. Experimentation can be used to evaluate hypothetical causal mechanisms linking frequency and polysemy. In this dissertation, I show that frequency can cause semantic extension and that it does so by affecting form accessibility. To do so, I manipulate frequency and accessibility, and examine the effects of these manipulations on semantic extension.

The hypothesis that frequency causes semantic extension, and that it does so through its effect on form accessibility has its roots in Zipf's (1949) proposal on the relationship between word frequency and number of meanings. Zipf likened linguistic expressions to tools, and noted that it would be functional for an artisan to reuse the most accessible tools to perform new tasks, absorbing the uses of less accessible tools. In the same way, it would be functional for a speaker to extend easily accessible expressions to new uses. A main goal of this dissertation is to develop Zipf's metaphor into a

mechanistic account of how frequency can cause forms to be more accessible in a new context.

An additional goal of this dissertation is to reconcile the literatures on semantic extension in language change and on over-extension in language acquisition. While high frequency has been proposed to result in extension of a form to novel meanings in language change, it has been argued to help retreat from over-extension in child language acquisition. How can frequency have seemingly opposite influences on language use? The current work aims to provide a solution, by showing that extension of a frequent form to new uses in production can co-exist, for principled reasons, with entrenchment of frequent forms to familiar uses in comprehension.

The broader goal of this dissertation is to show that processes influencing form selection during language production have the potential to influence linguistic structure by triggering language change. Furthermore, it demonstrates how the dynamic interactions between comprehension and production influence a form's productivity and range of uses.

1.2 Significance and contribution

Understanding the role of frequency in language change informs us about the ways in which language use and linguistic experience influence language structure. This is the central goal of usage-based linguistics (Bybee, 2006, 2010). The contribution of the present dissertation is three-fold.

First, it reports on experimental work that elucidates the mechanisms behind semantic extension in diachronic change by showing that frequency causes extension in the moment of production through its effect on form accessibility. The resulting

innovations can then be taken up by other members of the speech community, on the condition that the innovation is judged acceptable.

Second, by bringing together research on language acquisition and language change, this dissertation reconciles seemingly conflicting results in change and acquisition, by showing that frequency can simultaneously cause entrenchment in comprehension and (over-)extension in production.

A third contribution is demonstrating that the mechanisms that produce over-extension in children continue operating in adults. Extension has often been studied in child language, where it is considered *over*-extension because it does not match the adult community norms. The present study is focused on extension in adults. Because young children are not considered competent speakers, child innovations tend not to be taken up by other members of the speech community (Aitchison, 1991; Bybee, 2010; Bybee & Slobin, 1982; Cheshire 1982; Kapatsinski, Easterday & Bybee, 2019; Slobin, 1994, 2002; Vihman 1980). Demonstrating that frequency causes adult extension and not only child over-extension is therefore essential for explaining semantic change.

1.3 Theoretical basis of the study

I argue that the effect of frequency on form extension is mediated by its effect on accessibility. The influence of frequency on accessibility is well documented in picture naming: frequent words are more accessible from their meanings (Oldfield & Wingfield, 1965). The current work demonstrates that frequent forms are also more accessible given *related* meanings. The motivation for this prediction is the widely shared belief that semantics is distributed (e.g., Landauer & Dumais, 1997) so that related meanings activate the same features or attributes.

The work defends the idea that frequency causes both extension and entrenchment, with extension occurring in production and entrenchment in comprehension. Entrenchment refers to the notion that experiencing a form in a fixed semantic or structural context makes it unacceptable in new contexts (Ambridge, Pine, Rowland, & Young, 2008; Braine & Brooks, 1995; Theakston, 2004). Following Xu and Tenenbaum (2007), I ground the notion of entrenchment in Bayesian inference. Experiencing a form paired with meanings from a narrow region of semantic space provides evidence that the form is not used with meanings outside of that region. At the same time, under the pressures of real-time production, when a meaning that is close to the region is activated, a frequent form may nonetheless be the first one that comes to mind.

I take a usage-based constructionist approach to linguistic representation. From this perspective, there is no distinction between the grammar and the lexicon. Both consist of constructions, form–meaning mappings. It is "constructions all the way down" (Goldberg, 2006, p. 18) so that both syntactic structures and morphemes are constructions, as are words and idioms. A constructionist approach allows me to use data from different levels of language structure and predicts that they will be acquired and used in fundamentally similar ways. In particular, words, morphemes and syntactic constructions alike are productively extended to new uses, and can be over-extended. Thus, both over-extension of words to related meanings and over-generalization of regular patterns may occur as results of the same underlying processes or cognitive constraints. On the surface, word learning may seem very different from syntax acquisition, but the similarities are important and both are relevant to the question of

productivity. This flexibility is a hallmark of functionalist approaches to studying language, where levels of analysis that traditionally belonged to different subfields of linguistics interact and intersect (Bates & MacWhinney, 1989).

1.4 Research questions and hypotheses

Five major research questions guide the organization of the current work. The first research question focuses on the effect of frequency on form extension: Does frequency result in extension of forms to related meanings, as suggested in diachronic research? In accord with Zipf's proposal, I hypothesize that it does.

Research Question 2 is concerned with the mechanisms behind the effect of frequency on extension. If frequency results in extension, what is the cognitive mechanism behind this effect? I hypothesize that frequency increases accessibility of a form by strengthening the associations between the semantic features experienced with the form and the form itself. An alternative hypothesis maintains that frequency results in bleaching, so that frequent forms lose associations with some of the semantic features they used to possess (Bybee, 2003; Meillet, 1912).

Research Question 3 focuses on entrenchment. Does frequency result in entrenchment, as suggested in language acquisition literature (Braine & Brooks, 1995; Xu & Tenenbaum, 2007)? I hypothesize that it does, and that the locus of the entrenchment is in comprehension. This results in a dissociation between production and comprehension that is not predicted by bleaching.

Research Question 4 examines the relationship between entrenchment and extension: How do entrenchment and extension interact? I hypothesize that entrenched

form–meaning mappings acquired in comprehension can be reused in production, leading to retreat from over-extension.

Research Question 5 examines competition and the notion of pre-emption closely related to entrenchment. What is the effect of frequency of a form on the use and semantic extension of competing semantically-similar forms? I hypothesize that frequent forms pre-empt their semantic competitors (Boyd & Goldberg, 2011; Brooks & Tomasello, 1999; Brooks & Zizak, 2002; Clark & Clark, 1979; Goldberg, 1995; 2006; 2019, Perek & Goldberg, 2014; Robenalt & Goldberg, 2015), pushing them out of shared meanings.

1.5 Structure of this dissertation

Chapter II is an overview of previous work in the area, including entrenchment and its motivation in Bayesian inference, and the phenomena of extension and over-extension in children and adults. In Chapters III–IV, I report on three experiments testing the hypotheses outlined above in response to Research Questions 1–4. Chapter V focuses on the issue of pre-emption (Research Question 5) and its relevance to language acquisition and change. Chapter VI reports on an experiment that replicates Experiment I reported in Chapter III with a more natural language and addresses some alternative explanations for the results of Chapter V. Chapter VII discusses the major findings of the research so far, fleshes out the implications, and puts forward ideas for future directions. Chapter III and IV have previously appeared in a journal article co-authored with Volya Kapatsinski.

CHAPTER II

LITERATURE REVIEW

"But *glory* doesn't mean *a nice knock-down argument*," Alice objected.
"When I use a word," Humpty Dumpty said, in rather a scornful tone, "it means just what I choose it to mean—neither more nor less."
"The question is," said Alice, "whether you can make words mean so many different things."
"The question is," said Humpty Dumpty, "which is to be master—that's all."
Lewis Carroll: *Through the Looking-glass*

This chapter synthesizes the literature on overextension in word learning and morphosyntax, and the mechanisms by which learners retreat from overextension. I argue that frequent forms are the ones that tend to be extended to related uses in production. On this account, frequent forms inevitably accumulate meanings over diachronic time, resulting in a positive synchronic correlation between frequency and polysemy. Entrenchment and statistical pre-emption effect a retreat from *over*extension, which is faster for frequent forms. Computational simulations show that learning theory predicts the co-existence of entrenchment for frequent forms in comprehension and extension of these forms to new uses in production. Furthermore, competition with a strongly entrenched form can push another form out of the shared meaning. Throughout, I situate this dissertation relative to prior work on frequency effects and extension in language acquisition and language change.

2.1 Language acquisition

2.1.1 Over-regularization and over-extension[1]

Productivity is an essential aspect of language. In order to use language productively, children need to extend forms to new meanings and grammatical contexts. This includes both extension of a lexical form to related concepts and generalization of a grammatical form to related constructions. For example, children must learn that sneakers and flats could both be labeled as *shoes*. They need to learn that just as they use *-ed* to express the meaning PAST in the context of the verb *call*, they should also use it to express the meaning PAST in the context of the verb *try*. When children start producing sentences with more complicated syntax, they need to learn that the verb *break* can be used both in a transitive construction (*mommy broke the cup*), and in an unaccusative construction (*the cup broke*). Without extension, their language would be so limited that it would hardly be useful for communication (Bloom, 1973; Nelson, Rescorla, Gruendel & Benedict, 1978).

However, children's attempts at using language productively come at a price: As children extend forms to new objects, they sometimes overextend them to label objects that are not labeled that way by adults. That is, they apply a label more broadly than adults along at least one semantic dimension (Bloom, 1973; Clark, 1973; Gershkoff-Stowe & Smith, 1997; Fremgen & Fay, 1980; Huttenlocher, 1974; Mervis, 1987; Naigles & Gelman, 1995; Nelson et al., 1978; Thomson & Chapman, 1987; Rescorla, 1980). This phenomenon has been documented both in real language use and in experimental settings. For example, data on children's production during object naming activities at

[1] All of these terms are used interchangeably in the literature, but over-regularization is mostly rooted in morphology, over-extension in word learning, and overgeneralization in syntax.

home show that children often label objects erroneously[2]. For example, a child may label a BEAR as a *duck* (Gershkoff-Stowe & Smith, 1997) or label a STRAWBERRY as an *apple*, or even refer to both parents as *daddy* (Rescorla, 1980, p. 328). Clark (1973, pp. 79–83) documents a large number of extensions in difference semantic domains such as movement, shape, size, sound, taste, and texture. Some examples are provided in Table 2.1 below:

Table 2.1. Children's overextension errors, adapted from Clark (1973).

Label	First referent	(Over-)extensions in order	
candy	*candy*	*candy > anything sweet*	Leopold (1949)
fly	*fly*	*specks of dirt > dust > all small insects > his own toes > crumbs of bread > a toad*	Moore (1896)
tick-tock	*watch*	*clocks > all clocks and watches > gas-meter > fire hose wound on spool > bath scale with round dial >*	

In morphology, the most commonly discussed type of overgeneralization is over-regularization. The choice of this term is based on the fact that regulars typically outnumber irregulars (Bybee, 1995). For example, in English most overgeneralization errors involve generalizing a regular suffix such as past tense *-ed* or noun plural *-s* to contexts where their use is not appropriate (e.g., *goed* and *breaked* instead of *went* and

[2] Children's labeling also involves under-extensions and mismatches (see Clark, 1973), but here we are interested in cases where the child is clearly over-extending the label to related meanings for which the label should not be used. Over-extensions in general are more noticeable to adults compared to under-extensions (Anglin, 1977; Kay & Anglin, 1982). As Naigles & Gelman (1995, p. 20) point out: "it is more striking when a child calls an elephant doggie than when she FAILS to call a chihuahua doggie." They are also theoretically interesting because retreating from overextension seems to require entrenchment to the more specific meaning.

broke; *sheeps* and *foots* instead of *sheep* and *feet*). However, overgeneralization rates are not cross-linguistically associated with regularity. For example, overgeneralizations of ostensibly irregular German plural patterns outnumber overgeneralizations of the ostensibly regular *-s* plural (Köpcke, 1988; Marcus et al., 1992). While some researchers have argued that special mechanisms are needed to explain over-regularization (Pinker & Prince, 1988; Yang, 2016), in this dissertation, I consider it simply a species of overgeneralization (see also Bybee, 1995; Kapatsinski, 2018; MacWhinney, 2004).

While children's lexical over-extensions are mostly noticed and studied between the ages of 1;0 and 2;0 (Naigles & Gelman, 1995), children's morphological and syntactic over-generalizations have been documented up to age 7 (e.g., Theakston, 2004), 9, or 10 (Goldberg, 2019, p. 138, citing Hao, 2015) and even into adulthood (Bybee & Slobin, 1982; Dabrowska, 2012). It is natural then that this area of research has received much attention. As the theoretical arguments are relevant to the conceptualization of the hypotheses in this dissertation, I will take some time to discuss them here.

Verb argument structure over-generalization involves using a verb in a context where its subcategorization frame is violated. For example, knowledge of syntax requires children to know that despite the fact that the verb *break* can occur in both intransitive and transitive constructions as in examples (1)–(2), *giggle* is limited to only the intransitive construction, making (3) ungrammatical according to adult grammar (Bowerman, 1982, 1988; Pinker, 1989).

(1) *The cup broke.*

(2) *Daddy broke the cup.*

(3) **Don't giggle me.*

Other constructions that are often over-generalized are the dative construction, the passive construction, and the periphrastic causative construction. Below are some examples from Bowerman (1982):

Double object dative construction

(4) I said her no. (Age: 3;1)

(5) I do what my horsie says me to do. (Age 3;9)

(6) Shall I whisper you something? (Age 7;8)

Causatives

(7) Mommy can you stay this open? (Age 2;6)

(8) She came it over there. (Age 3;4)

(9) You cried her. (Age 5;3)

Locatives

(10) I'm gonna cover a screen over me. (Age 4;5)

(11) Feel your hand to that. (Age 6;10)

One might imagine that in a "perfect" world, parents would respond to all of the child's grammatical and ungrammatical uses of an inflectional morpheme or a verb and its argument structure construction with some sort of approval of grammatical uses and disapproval of ungrammatical uses (Gold, 1967). The child could then use both the positive evidence—existence of the grammatical utterances—and negative evidence—disapproval of ungrammatical ones—to guide production. However, in the real world, negative evidence is insufficient to guide the child's production (Braine, 1971; McNeill, 1966; Cazden, 1972; Zwicky, 1970; MacWhinney & Snow, 1985; Brown & Hanlon,

1970; see Marcus et al., 1992, Marcus, 1993 for a review; but cf. Chouinard & Clark, 2003; Schoneberger, 2010). The problem is exacerbated by the fact that many differences between regular and irregular verbs or verbs with different subcategorization frames are arbitrary (Pinker, 1989). The arbitrariness is true at least from a synchronic point of view, and the synchronic input is all that is available to the child. One solution for the child is to be conservative in his or her extension of forms (Baker, 1979), but that comes at the expense of productivity. Thus, the conservatism hypothesis is inconsistent with the existence of overextensions. Furthermore, at some point children need to abandon conservatism for their language to become adult-like. It seems as if to find out where the line is one must cross it.

2.1.2 Accounting for overextension errors in word learning

Overextension and overgeneralization errors are inherent to language learning. They are most numerous when the child's semantic and syntactic knowledge undergoes significant changes (e.g., Gershkoff-Stowe & Smith, 1997). During the times of rapid growth, many new forms or constructions are added to the constructicon before the learner is able to properly differentiate their meanings. Several proposals have been made to account for overextension.

In the domain of word learning, Clark (1973) proposed the Semantic Feature Hypothesis, which attributes over-extension to underspecified semantic representations: instead of *ball* meaning a round toy object, it might mean any round object (see also Mervis, 1987). To Clark, overextensions are largely based on perceptual similarity, e.g., in shape, size, etc. However, similarity in function may also be important (Lewis, 1957; Dewey, 1894). This is similar to Pinker's account of over-generalization errors in

morphology and syntax where the child's over-generalization of a construction to some verbs is attributed to his or her incomplete grasp of the verbs' semantics (Pinker, 1989).

As more words are added to the child's vocabulary, some words would take over parts of the semantic space initially mapped onto the over-extended item. These new items then push the over-extended item out of the shared areas, restricting its semantic space. The reduction in semantic space in turn renders the semantic representation of the word more specific. This process of restructuring and narrowing proposed by Clark is similar to accounts of pre-emption discussed later in the chapter. In a sense, the new words block the extension of the old word by taking over the space of semantic that does not belong to that word in adult's language. For this account to work, the child must have a tendency to avoid synonymy, which Clark (1987) calls the Principle of Contrast.

A second proposal is that the child is well aware that the label she is extending to label a certain object is not appropriate for that object but it is the best option available to her (Bloom, 1973; Nelson, Rescorla, Gruendel & Benedict, 1978). Analogy may play an important role here as the choice of the over-extended form is not entirely random (Hudson & Nelson, 1984; Nelson et al., 1978). That is, by using a word with a novel object, the child is not intending to say that the new object is the *same* as the object the word refers to. Rather, the child simply indicates that it is *like* that other object in certain respects.

Finally, a related but distinct proposal is that a word can be used to express meanings with which it has not been encountered simply because those meanings activate the word (Gershkoff-Stowe & Smith, 1997; Naigles & Gelman, 1995). This is the account of overextension I develop further in this dissertation. Because a word is

extended only to related meanings, the novel and familiar meanings of a word share attributes. When a novel meaning is encountered, these shared attributes are activated. Activating features that are shared with a familiar meaning of a word will in turn activate the word. The more frequent a word, the more strongly it will be associated with these shared features, and the more strongly it will be activated by them, leading to its production.

For example, consider the words *dog* and *cow*. In an experimental study by Naigles and Gelman (1995), the word *dog* was commonly overextended to a cow puppet in production. The two puppets shared many features: they were of a similar size, and were both black and white. These shared features would activate both the word *dog* and the word *cow*. However, due to the greater frequency of *dog*, the associations between the shared features and the form *dog* would be stronger than their associations with the form *cow*. Unless the non-shared semantic features of COW strongly activate the form *cow*, the form *dog* will be activated by COW more strongly than the form *cow*.

Three factors are argued to contribute to extension from this perspective. First, a word is likely to be chosen when its meaning is similar to the intended meaning (Naigles & Gelman, 1995; Nelson et al., 1978). Second, a word is likely to be extended to an object when the appropriate label for the object is inaccessible. This could be because the child has not learned a name for the object (Bloom, 1973; Nelson et al., 1978) or because the name is inaccessible at the moment of production, i.e., the child is faced with a retrieval problem (Huttenlocher, 1974; Naigles & Gelman, 1995; Thomson & Chapman, 1977). Third, a word is more likely to be extended to novel uses related to its known uses if it is experienced frequently (Gershkoff-Stowe & Smith, 1997; Naigles & Gelman,

1995). Accessibility of a form can also be increased, temporarily, by form priming, leading speakers to name a picture of a priest *nun* having just heard or said the homophone *none* (Ferreira & Griffin, 2003; see also Burke, Kester Locantore, Austin, & Chae, 2004; Gershkoff-Stowe, Connell, & Smith, 2006).

An important source of evidence for the position that overextensions are caused by a frequent form outcompeting a rare form during retrieval is that over-extensions in production need not be reflected in comprehension (Fremgen & Fay, 1980; Naigles & Gelman, 1995; Rescorla, 1980; Thomson & Chapman, 1977). When presented with *doggie*, the child may select a picture of a dog over that of a cow, and may even state that the cow is not a dog (Dromi, 1987). Yet, having not yet learned the form *cow*, the same child may nonetheless use the word *doggie* to refer to a large horned bovine (e.g., Bloom, 1973; Clark, 1973; Gershkoff-Stowe, 2001; Gershkoff-Stowe & Smith, 1997; Huttenlocher, 1974; Naigles & Gelman, 1995; Thomson & Chapman, 1977). When you have no access to *cow*, a *doggie*—presumably the semantically closest form you *can* access—will do.

In this dissertation, I will also contrast selection of a meaning based on a form in comprehension and selection of a form based on a meaning in production, and—like Naigles and Gelman (1995)—demonstrate that the two can be in apparent conflict. In addition, I compare production to a forced choice task that—like production—involves choosing a form on the basis of a meaning. Through this comparison, I show that the reason that frequent forms are over-extended in production is indeed the difference in accessibility between frequent and infrequent forms. Thus, the effect of frequency on form choice disappears if frequent and infrequent forms are made equally accessible.

However, whereas Naigles and Gelman (1995) proposed that comprehension tasks provide a better window on semantic representations than production tasks do, I will argue that the results of both tasks are consistent with mapping mature distributed semantic representations onto forms using bidirectional associations. While formal computational modeling must wait until later, a brief illustration can be provided by considering the results of Naigles and Gelman (1995).

Naigles and Gelman (1995) write (pp. 41–42): "Recall that when labelling the cow puppet, the most common error was to call it a dog; hence, *dog* would seem to refer to the larger, overextended category. However, in comprehension, it is now the *cow* label which appears to map onto the larger category, as it is overextended to include the dog puppet." They further argue that the comprehension data reflects the children's semantic representations, whereas production data reflect difficulties with retrieval. I argue instead that such dissociations are only apparent, and in fact consistent with bidirectional associations between forms and distributed semantic representations.

Consider the representation of dog and cow in Figure 2.1. Here, DOG and COW are both associated with the correct forms. The semantic representations are distributed and overlap in some features. The form–meaning associations involving *dog* are stronger because *dog* is more frequent than *cow*. The connection strengths are represented by line widths, which are 3 for *dog* ~ DOG and .5 for *cow* ~ COW. The activation of *dog* from DOG or DOG from *dog* is $3 \times 4 = 12$, while the activation of COW from *cow* or vice versa is $0.5 \times 4 = 2$. This difference reflects the well-known fact that frequent words are easier to recognize and produce than rare words (Howes, 1951; Oldfield & Wingfield, 1965).

Now consider what happens when the meaning COW is activated. Because it shares some features with DOG, it will activate the form *dog* more strongly than it will activate the form *cow*, leading to overextension of *dog* in production. When the form *dog* is presented, on the other hand, it will activate DOG much more strongly than it activates COW. In contrast, the form *cow* will activate COW more than DOG, but the difference in activations will be much weaker.

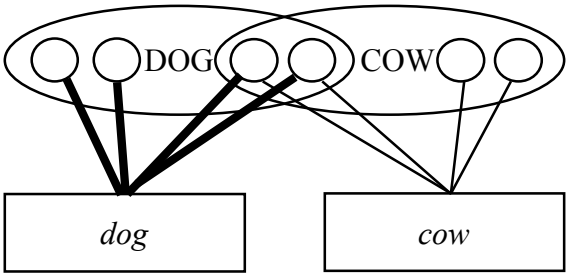

Figure 2.1. A schematic representation of *dog* and *cow*. Line width represents frequency.

The evidence for overextension of rare forms in comprehension is quite limited. One piece of evidence argued to support overextension in comprehension is Mervis and Canada's (1983) finding that children incorrectly consider a toy leopard a *kitty*. However, it is not at all clear that a toy leopard should not in fact be considered a kitty. Another piece of evidence is Naigles and Gelman's (1995) finding that children hearing the word *cow* would look at the dog puppet more than at the cat puppet when the cow puppet was absent. They argue that children presented with *dog* do not consistently look more at the cow puppet than at the cat puppet. However, children do in fact look significantly more at the cow puppet than at the cat puppet on some measures, and Naigles and Gelman do not

statistically test whether the interaction between the presented word and the puppet is significant.

Furthermore, research in the visual world paradigm has since shown that adults too look at pictures that are visually similar to the referent of a word more than they look at unrelated distractors. For example, adult participants in Dahan and Tanenhaus (2005) look at a picture of a rope when presented with the word *snake* and no snake is presented. Since adults presumably know that snakes are not ropes, this result questions the interpretation of preferential looking data in Naigles and Gelman as indicating an overly broad COW concept. Rather, looks can be drawn to referents merely similar to the referent of a word that the listener is hearing. While there is good evidence that repeated exposure to a form–meaning pairing results in a growing difference in how strongly the form activates that meaning vs. other similar meanings (Xu & Tenernbaum, 2007), there is little evidence that the original mapping is overly broad. For these reasons, I do not consider comprehension tasks to necessarily be more efficacious in elucidating semantic representations than production tasks.

2.1.3 Overgeneralization in morphosyntax

Another area where overgeneralization has received considerable attention is morphology. A huge body of research on inflectional morphology in English has discussed the underlying cause of overgeneralization errors. The influence of morphology on the field has been mainly due to the major impact morphological data has had on the theoretical debates surrounding rules versus analogy as the mechanism behind productivity (Albright & Hayes, 2003; Bybee, 1985, 1995, 2001; MacWhinney, 1978, 2004; Pinker, 1991; Pinker & Prince, 1988; Rumelhart & McClelland, 1986; Skousen,

1989). Generative accounts of morphological learning attribute these errors to over-generalizing a morphological rule such as VERB + *-ed* for past tense, while usage-based accounts attribute it to analogical processes involving whole words.

Despite the diversity of approaches, there is a consensus in this literature that probability of extending a morpheme to new words (whether correct or not) depends on type frequency—the number of words the morpheme is observed with (Bybee, 1995; O'Donnell, 2015), and on similarity between the new context to which the morpheme might be extended and the contexts in which it has been observed (see esp. Albright & Hayes, 2003; Hare & Elman, 1995; Kapatsinski, 2005; Suttle & Goldberg, 2011).

Type frequency is often confounded with how evenly the words featuring a particular morpheme are distributed in semantic and phonological space. However, it has been argued that type frequency alone is insufficient to account for the likelihood of extending a morpheme to highly unusual new words (e.g., the nonce verb *ploamph*, Prasada & Pinker, 1993). For example, the German plural *-s* is extended to new borrowings far more often than its low type frequency in the native lexicon would predict (Clahsen, 1999). To explain extension to highly unusual contexts, breadth of distribution needs to be taken into account (Hare & Elman, 1995). Morphemes that attach to a wide variety of existing words are more likely to be extended to very different new words because their context of occurrence is highly underspecified.

The role of token frequency of a morpheme is ambiguous in the literature on morphological productivity (see Kapatsinski, 2018b, 2018c, for a review). Some researchers have argued that high token frequency restricts morphemes to the stems they co-occurred with (Bybee, 2001; Hay, 2001). Others have argued that high token

frequency helps extension (Kirov & Cotterell, 2018), or that it can either help or hurt depending on type frequency (Barðdal, 2008; Madlener, 2016). Still others have argued that it has no effect (Albright & Hayes, 2003; Perfors, Ransom, & Navarro, 2014).

While the role of morpheme token frequency remains controversial, token frequency of a whole word clearly plays a role in paradigm leveling, a process involving extension of a form to new cells of a morphological paradigm, usually the ones that are semantically closest to the form's original use (Bybee & Brewer, 1980). For example, Tiersma (1982) has shown that Frisian generally restructured singular-plural paradigms on the basis of singulars, as if *geese* became *gooses* in English; with the exception of a small number of words in which the plural form is more frequent than the singular, such as *louse ~ lice*. In these words, the singular was reshaped on the basis of the plural. This is likely to occur in English too, if not prevented by formal education. On the TV show *Are you smarter than a fifthgrader* (Season 2, Episode 14), 0/5 adults could generate the singular form of the word *lice*, with 4/5 generating *lice* as the singular, and one generating *lie*. Although 5/5 fifthgraders got it right, this is only because they learned in third grade English class. It is generally the case that paradigm restructuring uses the more accessible, frequent forms as bases to restructure less frequent semantically similar forms (Bybee & Brewer, 1980). Overextension of frequent forms to other cells in the paradigm is also observed in children; especially those learning morphologically complex languages with rich paradigm structure (Aguado-Orea & Pine, 2015; Dąbrowska and Szczerbiński, 2006; Engelmann et al., 2019; Leonard et al., 2002; Marchman, 1997; Matthews & Theakston, 2006; Räsänen et al., 2016; Rubino & Pine, 1998; Theakston et al., 2003; Theakston & Rowland, 2009).

In morphologically poor languages like English, overextension of forms to new cells in the paradigm usually looks like omission of a suffix (Freudenthal, Pine, Aguado-Orea, & Gobet, 2007). For example, if one overextends *walk* to use with a third person singular subject, producing *She walk her dog every day*, this would appear as an omission of *-s* even if it is really an overextension of *walk* into a context that demands *walks*. The similarity of this process to omission has led generative syntacticians to consider it an omission process, a so-called 'root infinitive' stage (e.g., Radford, 1996; Rizzi, 1993; Wexler, 1998).

A usage-based model of syntax acquisition in children, MOSAIC, instead attributes these errors in child language to frequency of exposure to sentences with compound finites, which root infinitives are over-extensions of (Freudenthal, Pine, Aguado-Orea, & Gobet, 2007). Compound finites are observed when another verb in the sentence carries tense and agreement as in the English sentence *Can he play?*. According to MOSAIC, **he play* in child language is motivated by *Can he play?* in the input. Similarly, *Peut-il jouer?* in a French-speaking child's input and *Wil hij spelen?* in a Dutch-speaking child's input may results in **Il jouer* and **Hij spelen* errors in child language respectively (Ambridge & Lieven, 2011, pp. 152). Freudenthal et al. (2010) add support to this proposal by demonstrating a cross-linguistic correlation between the occurrence of compound finites and the rate of omission errors so that languages with high rate of compound finites have higher proportion of root-infinitive errors.

The factors that influence the productivity of morphological constructions may also be at work when it comes to syntactic constructions. While overextension of both morphological and syntactic constructions has been observed (e.g., in that children might

say *She said me Hi* instead of *She said Hi to me*), the syntactic literature has mostly focused on how such errors are curbed (e.g., Ambridge et al., 2008; Boyd & Goldberg, 2011; Braine & Brooks, 1995; Brooks et al., 1999; Goldberg, 2006, 2019) rather than on explaining the process that leads to their occurrence in the first place. In essence, the overextensions are attributed to immature knowledge about the particular verbs to which the constructions are overgeneralized (Pinker, 1989). In comparison, the characteristics of syntactic constructions being overgeneralized remain underexplored (though see below). To a large extent, this is because all syntactic constructions have often been assumed to be equally and fully productive. Quantitative corpus measures of productivity suggest that this is not the case: just as some morphological constructions often generate new words, and others do not, some syntactic constructions often generate new constructs while others do not (Zeldes, 2012). The existence of productivity differences among syntactic constructions suggests that the factors that affect the productivity of morphological constructions may also affect syntactic productivity. These factors include type and token frequency as well as the way that the verbs to which the construction applies are distributed in semantic space.

Some evidence for the importance of semantic breadth is presented by Suttle & Goldberg (2011). Given the same semantic distance—as measured by Latent Semantic Analysis (Landauer & Dumais,1997)—between a new verb and the center of the cluster of verbs that have already been observed to occur in a construction, extension of the construction to the verb was more likely when the previously observed verbs were more diverse. This result is closely related to Hare, Elman and Daugherty's (1996) results for diachronic changes in the English past tense system. The *-ed* suffix was able to overtake

the Old English strong verb changes despite initially having low type frequency because the verbs it could associate with were very diverse. Diversity may also explain the "minority default" behavior of German *-s* plurals and Semitic sound plurals (Boudelaa & Gaskell, 2002; Hare & Elman, 1995).

Several studies have explored the effects of how the token frequency of a construction is distributed over the instantiating types (Ellis & Ferreira-Junior, 2009; Goldberg, Casenhiser & Sethuraman, 2004; Madlener, 2016; McDonough & Nekrasova-Becker, 2014; McDonough & Trofimovich, 2013; Nakamura, 2012). The most consistent effect in this literature is that a skewed token frequency distribution over types reduces overextension, and a more balanced token frequency distribution promotes it (as first documented by Goldberg et al., 2004). However, the opposite effect has also been observed (Madlener, 2016), as have null effects (McDonough & Nekrasova-Becker, 2009; McDonough & Trofimovich, 2013; Nakamura, 2012).

As Madlener (2016) points out, interactions with type frequency and token frequency are likely responsible for these inconsistencies. A skewed token frequency distribution may promote generalization when both type and token frequency are high because enough types will be learned for the construction to be productive regardless of skew. Skew then increases the number of hapax legomena, which have been associated with productivity (Baayen, 1992; Zeldes, 2012). When type and token frequency are low, a skewed token frequency distribution may instead mean that no generalization occurs because only one type is token-frequent enough to be learned. The roles of type and token frequency remain underexplored in the literature on productivity of syntactic constructions.

2.1.4 Retreat from overgeneralization

After a period of overgeneralization, the child's use of a form becomes more adultlike. As a result, the overall accuracy often follows a U-shaped trajectory[3]: accuracy dips as a result of overextension/overgeneralization and then recovers. For example, in morphology, irregular forms are produced correctly at first but this period is followed by over-regularization, from which the frequent irregulars eventually recover (Ervin & Miller, 1963; see also Bybee & Slobin, 1982; Cazden, 1968; Ervin & Miller, 1964; Kuczaj, 1977; Marcus et al., 1992). The regular suffix *-ed* is first used only with the verbs with which it is observed, before being overextended to irregular verbs. Periods of lexical over-extension (Gershkoff-Stowe & Smith, 1997) also appear after the child correctly uses the words to refer to some appropriate referent(s) (Nelson et al., 1978). Just as periods of over-regularization (Marcus et al., 1992), periods of lexical over-extension tend to be brief and limited in length to around only a few weeks (Gershkoff-Stowe & Smith, 1997). Because forms are always learned in context, the experienced uses of a form must be appropriate as these uses are likely to be limited to these experienced contexts (Tomasello, 2003). Thus, the earliest uses of a form are appropriate. However, at some point the form can be over-extended beyond the range of contexts that adults find appropriate. The child then needs to retreat from this over-extension.

As mentioned earlier, explicit negative evidence is generally agreed to be insufficient to retreat from overextension. Consequently, the accounts of constraining an

[3] This is very likely a micro U-shape development in which some forms exhibit this trajectory and they exhibit it at different times as opposed to macro U-shape development in which all verbs or nouns are influenced by patterns of over-regularization.

overgeneralized grammar have limited themselves to explanations based on positive or indirect negative evidence. These include distributional and semantic evidence.

The most successful accounts of curbing over-generalization errors have surfaced in the constructionist tradition of language acquisition research as these accounts have established a strong connection between the child's input and his or her language use. In what follows, we review these accounts. Many of these explanations demonstrate that not only there is ample data for the child to reach adult grammar, hence reducing the need to resort to innate principles (e.g., Pinker, 1989), but have also provided information about potential mechanisms through which over-generalization errors are curbed.

2.1.4.1 Entrenchment

Overgeneralization errors are eventually curbed as the child's linguistic knowledge approximates that of adults. The mechanisms behind curbing over-generalization errors have been a focus of research on syntax acquisition for decades, partly because of strong claims in the generative literature that retreat from an overgeneral grammar is logically impossible (Gold, 1967; Baker, 1979; Yang, 2016). The argument is that all grammars that generate all of the observed utterances are equally likely and cannot be rejected. However, this argument ignores differences in how probable the observed data are according to the different grammars and is inconsistent with normative Bayesian inference. Suppose that there are two hypotheses H_1 and H_2, both of which state that the observed data (D) can occur, but H_1 is more restrictive than H_2 in that H_2 says that other, unobserved data can also occur. Because H_2 assigns some probability mass to unobserved data, the probability of the observed data D given H_1 (= 1) is inevitably larger than the probability of D given H_2 (< 1). The less restrictive

hypothesis grows increasingly unlikely as more data is observed (Kapatsinski, 2018d; Regier & Gahl, 2004; Xu & Tenenbaum, 2007).

This intuition is consistent with the notion of *entrenchment*, first proposed by Braine and Brooks (1995), which states that an increase in the frequency of a form in a certain construction prevents that form from being extended to use in other constructions in which it has not been attested. Starting with Theakston (2004), ample evidence for entrenchment in syntax has been provided by grammaticality judgement tasks. Theakston exposed three age groups, 5-year-olds, 8-year-olds and adults, to sentences containing an over-generalization error on the verb argument structure. The sentences were composed using semantically similar verbs but were different in that one group contained high-frequency verbs (e.g., *pour*; mean frequency = 790.7) and the other low-frequency verbs (e.g., *dribble*; mean frequency = 19.8). For all three groups, sentences containing low-frequency verbs were judged as more acceptable than sentences containing high-frequency verbs. Furthermore, the results indicated an increase in entrenchment with age, consistently with entrenchment curbing children's over-generalization errors: compared to 5-year-olds, 8-year-olds were significantly less likely to accept over-generalization errors for both low and high frequency verbs.

That said, Theakston did not find a relationship between the frequency of verb and the acceptability of overgeneralization errors involving the verb in adults, possibly due to a floor effect. The results of judgment for certain verbs also did not fit with the overall trend for overgeneralizations of frequent verbs to be less acceptable. For example, adults exhibited a preference for *come* and not *arrive* in a sentence such as *she came/arrive me to school* despite the fact that frequency of *come* is 282 times more than

frequency of *arrive* (p. 27). Theakston attributed this to a preference for the co-occurring verb and preposition as *at* is more likely than *to* to occur with *arrive* and *to* is more likely than *at* to occur with *come*. The preference for *come to* over *arrive to* could therefore override the lower acceptability of *came me* compared to *arrived me*.

Matthews, Lieven, Theakston, & Tomasello (2005) exposed children ages 2;9 and 3;9 to high, medium, and low frequency English verbs (e.g., *shove, push,* and *ram* respectively) in a novel "weird" SOV construction (see Akhtar, 1999). The younger group adopted the SOV construction only for low frequency verbs, while the older group failed to use it altogether. They concluded that not only children's knowledge of grammatical structures is lexical and item-based, but also that high frequency of exposure results in the entrenchment of these items in their original constructions.

Theakston and colleagues' work was followed by several studies by Ambridge and colleagues demonstrating the effect of entrenchment on retreat from over-generalization in both children and adults. Ambridge, Pine, Rowland, & Young (2008) tested children and adults on grammaticality judgement of overgeneralization of intransitive verbs to transitive constructions. Ambridge and colleagues were concerned with the limitations that blocking and competition face in accounting for certain cases of retreat from overgeneralization errors. For example, *the magician disappeared the rabbit* does not have a competitor that can block it (there is a gap that is filled with *vanish*), except for indirect competitors like periphrastic causative (*made disappear*). So, how does the child learn to retreat from overgeneralization errors in these cases? The concern is rooted in the claim that certain periphrastic causatives do not block their corresponding intransitive uses (e.g., *John made the baby stand up* vs. *John stood the baby up*;

Bowerman, 1988). Ambridge and colleagues focused on the causative alternation, and its overgeneralization, as in *the rabbit disappeared* → **The magician disappeared the rabbit.* They also improved on the previous studies by using similar rating scales for children and adults and added grammatical sentences as a baseline control (cf. Theakston, 2004).

Ambridge and colleagues tested the entrenchment hypothesis in a grammaticality judgement task. The prediction, as with the previous studies testing the effect of entrenchment, is that participants are more accepting of an overgeneralization error of a verb that has occurred with lower frequency with a particular construction compare to a verb that has occurred with high frequency with that construction. They exposed 5–6- and 9–10-year-olds, and adults to a graded grammaticality judgment task involving high frequency, low frequency, and novel verbs (e.g., *fall, tumble, meek*). Like Theakston (2004), Ambridge et al found a significant effect of frequency. However, unlike Theakston, they found no interaction between frequency and age. For all three age groups, the difference in judgments between grammatical over ungrammatical uses were greater for high-frequency verbs than for low-frequency verbs and greater for low-frequency verbs than for novel verbs.

2.1.4.2 Entrenchment as a suspicious coincidence effect

Entrenchment rests on the fact that exposure provides information regarding whether the form F can be used in contexts in which it has *not* yet been encountered. If one keeps encountering a word in a particular context X, and there are other frequent contexts Y in which the form is never encountered, then there is a growing probability that the form's absence from Y is not accidental. That is, there is some constraint that

prevents FY from occurring (Regier & Gahl, 2004; Stefanowitsch, 2008; Xu & Tenenbaum, 2007). In the terminology of Xu and Tenenbaum (2007), it is a *suspicious coincidence* that every time the form occurs it occurs in the context X.

In Xu and Tenebaum's (2007) study, children and adults learn artificial labels for objects in a subordinate category level in a taxonomic hierarchy. They are then tested on their tendency to generalize that label to the basic and superordinate category in the same taxonomic hierarchy. For example, they learn a label for Dalmatian and then they are tested on whether the same label could express the concept of dog, i.e., the basic level in the hierarchy encompassing Dalmatian (subordinate), dog (basic), and animal (superordinate). This was tested by asking the learners if other types of dog could be labeled with the same label. If a child refers to Dalmatian, Terrier, and German Shepherd with the same label, the child considers the label to refer to DOG.

Xu and Tenenbaum (2007) demonstrate that when the child experiences the label *fep* with Dalmatian, *fep* is perceived by the learner to mean DOG as the learner often accepts the use of *fep* to refer to other types of dog as well. However, when *fep* is observed three times with Dalmatians, its meaning gets restricted to DALMATIAN and is significantly less likely to be accepted as a label for DOG.

Xu and Tenenbaum (2007) attribute this effect to Bayesian inference. According to their account, the child implicitly thinks along the following lines: "if *fep* means DOG, what are the chances that every time I see a dog it is a Dalmatian? Not very likely!" Put differently, seeing a Dalmatian paired with *fep* once, the probability of p(Dalmatian & *fep* co-occurrence | *fep* = DOG) is, say, .1 if 10% of dogs are Dalmatians. Not too improbable. But when Dalmatian is paired with *fep* three times that probability is reduced to $.1^3$, or

.001, exceedingly unlikely. Therefore, children end up restricting the form *fep* to its original meaning of a subordinate category, DALMATIAN.

Regier and Gahl (2004) use Bayesian inference to account for the fact that the anaphoric one refers to a noun phrase rather than an N': In *He found a yellow balloon. I want one too*., *one* is interpreted as meaning another yellow balloon rather than another balloon of any color. Stefanowitsch (2008) shows that a similar line of reasoning can be followed for whether the verb *disappear* can occur in both transitive and intransitive constructions. As one keeps encountering *disappear* in intransitive contexts, it becomes less and less likely that it can also be transitive.

In both entrenchment and suspicious coincidence, the learner takes advantage of only the frequency of occurrence of the form with its construction or meaning. The effect is independent of whether there is a competitor form that would block the use of Dalmatian to refer to other types of dog, or the extension of the transitive construction to the verb *disappear*.

2.1.4.3 Pre-emption

An alternative mechanism for curbing overgeneralization errors is preemption. According to Clark and Clark's (1979) Principle of Pre-emption by Synonymy[4], "If a potential innovative denominal verb would be precisely synonymous with a well-

4 Clark & Clark (1979) also discuss the Principle of Pre-emption by Homonymy: "If a potential innovative denominal verb is homonymous with a well-established verb and could be confused with it, the innovative verb is normally pre-empted, and therefore is considered unacceptable." (Clark & Clark, 1979, p. 800). For example, to *summer, autumn, and winter in France* is acceptable, but to *spring* and *fall in France* is not, being pre-empted by the homonymous common verbs spring and fall.

established verb, the innovative verb is normally pre-empted by the well-established verb, and is therefore considered unacceptable." Clark and Clark give the example of the verb *hospitalize* preempting the verb *hospital* which could potentially be an appropriate verb formed from the noun *hospital*. In the Principle of Pre-emption by Synonymy, the locus of preemption is in comprehension (Clark & Clark, 1979, p. 798; see also MacWhinney, 1987):

> Thus the listener would 'reason' as follows: Suppose my interlocutor had intended to convey the sense 'put into a hospital'. If he had, he would have used the well-established verb hospitalize, which means precisely 'put into a hospital', because then he would have had good reason to think I would compute the intended sense uniquely. Since he used hospital, he must have meant something distinct from 'put into a hospital'. Yet the only reasonable sense I can come up with is 'put into a hospital', which I already know to be impossible. Thus I find hospital to be uninterpretable, and therefore unacceptable.

Pre-emption is well-documented in morphology, where it is called "blocking" (Aronoff, 1976; Kiparsky, 1982; MacWhinney, 1987, 1993; Rainer, 1988). For example, the occurrence of *went* in contexts where the child uses and expects *goed* can discourage the production of *goed*. Importantly, the irregular past tense form is thought to block the regular past tense form only if the child realizes that they both have the same meaning (Kuczaj, 1977, 1981). If the child fails to make this semantic connection, and treats the irregular and regular versions of the verb as two independent verbs, she will continue to over-regularize. It is only when the expected form and the observed form are identical in meaning that the observed blocks the expected. Thus, *brethren* can coexist with *brothers* and *hanged* could co-exist with *hung*.

Goldberg (1995) extended the idea of blocking to constraining the overextended use of a verb in a verb argument structure construction. According to Goldberg (1995), when a child fails to observe the most preferred argument structure construction for a meaning with a particular verb, the child suspects that the use of the verb in that construction is not allowed. As this experience gets repeated, the child becomes certain that the verb is not suitable for that construction. The preference for the construction stems from lack of synonymy in language:

> Since we have assumed that no two constructions are entirely synonymous both semantically and pragmatically (cf. chapter 3), it should be possible to find contexts in which a given construction is the most preferred. If the preferred form is *not* used, then the child is able to tentatively infer that that form is disallowed.

Evidence for this hypothesis was first presented by Brooks and Tomasello (1999). Brooks and Tomasello (1999) exposed children to two novel verbs, one with and one without an alternative pre-empting construction. This alternative construction was passive for the transitive novel verb and periphrastic causative for the intransitive novel verb. Children's over-generalization errors were assessed using an elicitation task. Brooks and Tomasello found an effect of pre-emption for the older group of children, ages 6 and 7, demonstrating that in the presence of an alternative construction, children were less likely to overgeneralize a verb to new contexts. Brooks and Zizak (2002) revisited and replicated these results in a similar study, again, showing that an alternative construction can discourage older children from overextending a verb to new contexts.

In more recent work, Goldberg and colleagues highlight the effect of frequency (of the pre-empting form) in preemption, referring to this effect as statistical pre-emption

(Boyd & Goldberg 2011, Goldberg, 2006, 2011; Robenalt & Goldberg, 2015, 2016; Perek & Goldberg, 2015). Statistical pre-emption refers to the idea that forms compete for expressing the same meaning. Whenever one expects to encounter one form as an expression of a particular meaning but encounters another, the encountered form pre-empts the expected form, weakening its association with the meaning. Thus, if the child is expecting to hear a verb in a certain construction, such as the double object construction (*Explain me this*), and instead consistently hears an alternative construction such as the prepositional object construction (*Explain this to me*), the child will learn not to use the verb in the double object construction.

There are two crucial requirements for pre-emption. First, the violation of expectations and blocking should happen consistently, in more than one instance, hence *statistical*. One instance can be ineffective as it may be viewed by the learner to have been an error on the part of the speaker (Goldberg, 2006). Second, that the two constructions should express the same meaning. If the meanings are different, the learner may think that the two constructions can co-exist, as they express (slightly) different meanings.

Boyd and Goldberg (2011) demonstrate that the use of a novel adjective such as *ablim* from the class of English a-adjectives (e.g., *asleep*, *acute*) is avoided in attributive constructions (e.g., **the asleep boy*) only when speakers experience that adjective used in an alternative construction with the same meaning (e.g., *The boy who is sleeping*), but not if the alternative formulation can be attributed to other reasons such as conjoining with a prepositional phrase in a complex modifier (e.g. *ablim and proud of himself*). Boyd and Goldberg argue that this is because pre-empting constructions are necessary to curb

overgeneralization and that these constructions must share the meaning of the construction to which a word is overgeneralized.

Robenault and Goldberg (2015) found that high frequency of a verb reduces the acceptability of the verb in a novel construction only if a pre-empting construction is available. In the absence of an alternative construction that would express the intended meaning of the speaker, high frequency of the verb fails to entrench that verb and prevent its over-extension to novel constructions. When an alternative construction is available, high frequency verbs are less likely to be extended to novel uses (see also Robenault & Goldberg, 2016).

2.1.4.4 Is entrenchment necessary?

While no researcher denies the existence of pre-emption, Goldberg and colleagues have argued that pre-emption is the only mechanism curbing overgeneralization, and that entrenchment is unnecessary (e.g., Goldberg, 2019). In this thesis, I argue instead that both entrenchment and pre-emption are necessary to account for retreat from overgeneralization (Ambridge et al., 2008, 2012, 2015, 2018; Braine & Brooks, 1995; Brooks et al., 1999; see also Clark & Clark, 1979).

One problem with a preemption-only account is that entrenchment is inherent to the learning mechanisms thought to give rise to pre-emption. As described above, Xu and Tenenbaum (2007) and Regier and Gahl (2004) show that entrenchment is inherent to Bayesian inference. Boyd and Goldberg (2011) also frame their preemption-only account of retreat from overgeneralization within a Bayesian inference framework. It is unclear how Bayesian inference could be used to learn constructions without giving rise to suspicious coincidence.

The suspicious coincidence effect involves an increase in the frequency of a form in a particular meaning that results in that form not being extended to other meanings. According to statistical pre-emption, extension is prevented by strengthening pre-emptors. Increasing the frequency of *fep* ~ DALMATIAN should therefore strengthen some pre-emptor. It cannot be *fep* because *fep* cannot be expected to pre-empt itself. It is therefore possible to account for Xu and Tenenbaum's (2007) suspicious coincidence effect with pre-emption alone but only if *meanings* can pre-empt each other. Under this account, encountering *fep* with DALMATIAN would strengthen the *fep* ~ DALMATIAN association, and DALMATIAN would then pre-empt *fep* from referring to GERMAN SHEPHERD and other breeds. However, if meanings pre-empted each other, languages would have a tendency to avoid homonymy and polysemy. Goldberg (2019) argues, convincingly, that encountering a form in a novel meaning should not lead the child to unlearn the other meaning of the word (see also Clark, 1987). If meanings pre-empted each other, frequent words would also not accumulate meanings over diachronic time (Goldberg, 2019), contrary to Zipf (1949).

Another difficulty for a preemption-only model is that pre-emption requires the child to expect to hear a form s/he does not observe based on a meaning. It is not clear that a child who uses *kitty* to refer to cows would in fact expect an adult to use the word *kitty* to refer to a cow. As noted above, children often know that a cow is not a cat despite using the word *kitty* to refer to both (Naigles & Gelman, 1995; Nelson et al., 1978). The child may simply be using *kitty* to refer to cows because it is more accessible than the word *cow*.

2.1.4.5 Entrenchment and prototype effects

The account of word learning in Xu and Tenenbaum (2007) assumes a strict hierarchical taxonomic hierarchy, where the alternative hypotheses the learner considers are in a set-superset relationship. Thus, *fep* is considered to either apply to all dogs, or all Dalmatians. However, the boundaries of semantic categories are often fuzzy (e.g., Bybee & Eddington, 2006; Kapatsinski & Janda, 2011; Labov, 1973; Lakoff, 1987; Rosch, 1978). Rather than having all-or-none membership, natural categories display prototype effects: some birds are birdier than other birds (Rosch, 1978). Similarly, lexical categories associated with constructions have prototype structure (Hopper & Thompson, 1980). Words display degrees of association with a specific syntactic construction (Bybee & Eddington, 2006; Stefanowitsch & Gries, 2003), and generalization of a construction to new words is predictable from the degree to which the new word is semantically similar to the words known to occur in the construction (Bybee & Eddington. 2006; Suttle & Goldberg, 2011).

The applicability of a form gradiently decreases around the encountered exemplars of its use. For example, a child that uses *kitty* to refer to cows is nonetheless generally able to say that a cat is a better example of a *kitty* than a cow is (Huttenlocher, 1974; Naigles & Gelman, 1995; Nelson et al., 1978; see also Bybee & Eddington, 2006; Kapatsinski & Janda, 2011; Labov, 1973; *inter alia* for adults). As the child does not know the word *cow*, there is nothing to pre-empt the extension of *kitty* to COW. The fact that COW is nonetheless not considered as good of a semantic match to *kitty* as CAT therefore provides evidence that experience with the form *kitty* itself constrains the extension of *kitty*. Findings of this kind are consistent with the idea that a form is

associated with a *distribution* in semantic space, where some parts of the space are associated with the form more strongly, and others less strongly.

The strongest associations tend to involve the areas of the space that have been paired with the form most frequently, and extension beyond the areas of the space paired with the form follows a gradient of similarity, as illustrated in Figure 2.2. Here, *kitty* has been mapped onto an area of semantic space with CAT in the center. Concepts like CAT themselves correspond to areas of semantic space: some cats are more *kitty*-ish than others. A COW is hovering at the edge of the semantic space *kitty* is mapped onto and is not as much of a *kitty* as CAT is. In fact, some COWs might not be *kitties* at all: calling a cow a kitty likely depends on how the cow is construed. There may or may not be a strict boundary associated with the area of space mapped onto *kitty* (Langacker, 1987).

In the realm of sound perception, Kleinschmidt and Jaeger (2016) account for the selective adaptation effect by showing that category narrowing with repeated exposure to the prototype is inherent to a parametric Bayesian model of phonetic categorization. Kleinschmidt and Jaeger (2016) represent phonetic categories as normal distributions describable by their mean and variance. Additional exemplars of the prototype presented to the subject of a selective adaptation experiment necessarily reduce variance, causing the category to become narrower. Selective adaptation in this model is therefore equivalent to entrenchment to the prototype. Figure 2.2 illustrates that entrenchment is equally inherent to a distributional representation of semantic categories. Any encounter with the prototype narrows the category by reducing its variance. As the child continues experiencing the word *kitty* paired with cats and not cows, the semantic distribution mapped onto *kitty* grows peakier, retreating from COW space.

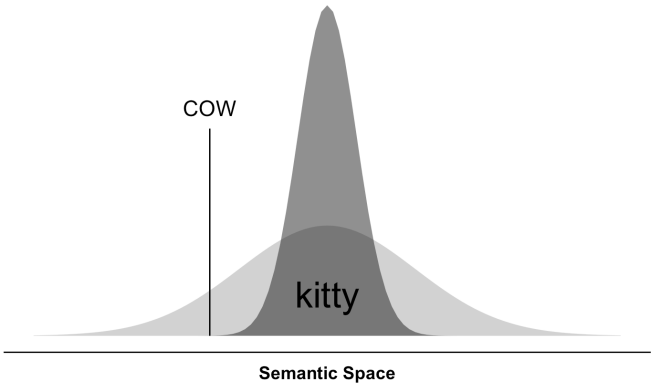

Figure 2.2. Entrenchment as a result of exposure to the prototype of a category. The semantic space is simplified into one dimension.

If entrenchment results from an increase in the peakiness of the distribution of the word's referents in semantic space, then an increase in the frequency of a form will not always result in entrenchment (as shown by Suttle & Goldberg, 2011). Any increase in frequency that increases the variance of the distribution will instead increase the likelihood of extending the form to new referents. Thus, adding token frequency to infrequent types exemplifying a construction will generally make the distribution of the construction's tokens more variable, increasing the likelihood of extending the construction to new types. This is consistent with the effects of token frequency distribution skew in Goldberg et al. (2004): decreasing the skew means distributing tokens over types more evenly, and results in an increased likelihood of generalizing the construction from motion events to events of appearance. Similarly, if one adds *kitty* exemplars paired with animals other than domestic cats, e.g. by referring to a LION as a *kitty*, then this increase in the token frequency of *kitty* may in fact increase the likelihood of extending *kitty* to COW.

Both proponents and opponents of entrenchment in syntax have assumed that

high-frequency verbs should necessarily be more entrenched in the constructions they are observed in, and therefore less likely to be extended to, and judged acceptable in novel constructions. Thus, *come* should be judged as less acceptable than *arrive* in the frame *He ____ me to school* (Theakston, 2004). However, this ignores how tokens of a verb are distributed over the semantic and constructional contexts. Frequent verbs like *come* are more polysemous than rare verbs like *arrive*, their meanings spread over a much wider area of semantic space (Zipf, 1949). To the extent that some of these meanings and uses are similar to the novel use participants are asked to judge in an acceptability judgment experiment, novel uses of frequent verbs should be judged to be *more* acceptable than novel uses of rare verbs (cf. Blything, Ambridge, & Lieven, 2014; Theakston, 2004). Taking into account the distribution of a verb's uses over semantic space may therefore help account for the occasional failure to observe entrenchment effects by comparing frequent and infrequent verbs (e.g., Robenault & Goldberg, 2015; and *come* vs. *arrive* in Theakston, 2004).

2.2 Language Change

2.2.1 Extension in diachrony

While the literature on child language acquisition has focused on how (over)extension is curbed, the literature on diachronic change has documented innumerable examples in which extension was *not* curbed, resulting in an increase in the semantic breadth or polysemy of the extended expression. In English, the word *dog* has been extended from a particular breed of dog to all dogs. Latin *tenere* meaning TO HOLD has been extended to mean TO HAVE in Spanish (*tener*).

As noted in Chapter I, Zipf (1949) documented that the token frequency of a word correlated with polysemy: a frequent word usually has many more meanings than a rare word. Zipf argued that this correlation is caused by the speaker's tendency to extend frequent words to new uses. This hypothesis, however, remains controversial. For example, Haspelmath (1999) argues that extension of a form to new uses causes the frequency of the form to rise, but not vice versa. The controversy is due in part due to the fact that Zipf did not propose a mechanism by which accessibility could cause extension. He simply stated that such an extension would make the language more functional. A major aim of this chapter is to place this hypothesis on a firmer foundation, by grounding it in learning theory and modern psycholinguistics.

Unlike Haspelmath (1999), Bybee (2003) does consider frequent use to also result in semantic extension, but through a process different from the one proposed by Zipf. Whereas Zipf attributed extension of frequent forms to their greater accessibility, Bybee (2003) attributes extension to an increase in semantic breadth, which itself results from frequency-caused habituation. In habituation, as a stimulus is repeated and grows less and less surprising, it loses its ability to elicit an associated response: A shock becomes less and less shocking; A piece of chow less and less tempting; And a construction less and less prone to activate its semantic associations. It is argued that frequently repeated forms lose the ability to evoke all of the associated semantic features (Bybee, 2003).

While both Bybee and Zipf propose that an increase in frequency of use causes semantic change, their views on the nature of the ensuing change differ. Bybee sees frequency-driven semantic change as a kind of semantic reduction, a decrease in specificity, whereas Zipf saw it as an extension of the frequent form to specific new

uses—particularly uses for which the speaker may not have a well-practiced, highly accessible expression.

In this dissertation, I argue that many well-documented instances of semantic broadening cannot be explained by habituation and can only be attributed to gradual extension from context to context. Whereas habituation is thought to cause extension by broadening the meaning of a word, I argue that extension often causes rather than is caused by semantic change. The use of a form is extended to a new context because the form is accessible in the new context. When a listener is then exposed to the form being used in a broader range of contexts, she reinterprets the form as having a broader meaning.

For example, Bybee and Eddington (2006) document how a variety of Spanish verbs have been extended to mean *become* in the context of different adjectives, including *hacerse* ~ MAKE ONESELF, *ponerse* ~ PUT ONESELF, *volverse* ~ TO RETURN, and *quedarse* ~ TO BE LEFT. Bybee and Eddington show that the verbs took on the meaning BECOME in specific adjective contexts and were then extended to other, semantically-similar adjectives. For example, *quedarse* was extended from *quedarse solo* ~ BE LEFT ALONE to *quedarse aislado* ~ ISOLATED, *soltera* ~ SINGLE and the antonymous *emparejado con* ~ PAIRED WITH. *Quedarse quieto* ~ QUIET/SPEECHLESS was extended to *tranquilo* ~ QUIET/PEACEFUL and then to *a gusto* ~ PLEASED, but also to *inmóvil* ~ MOTIONLESS, and *dormido* ~ ASLEEP and *muerto* ~ DEAD. This example shows how extension can eventually result in a form being associated with a synchronically arbitrary class of referents. There is little shared between PLEASED/CONTENTED and DEAD. Yet,

both adjectives take the same verb because of a chain of semantic extensions to specific adjectival contexts.

2.2.2 Entrenchment and pre-emption in diachrony

The positive correlation between token frequency and polysemy (Zipf, 1949) indicates that extension of frequent forms is much more common than their entrenchment in their frequent contexts. This point is underscored by the existence of grammaticalization, in which a word is extended to enough contexts to be considered grammatical. A classic example is *going to* VERB, which used to mean simply WALKING SOMEWHERE IN ORDER TO VERB, but has developed into a future marker by first being extended to non-motion contexts and finally to contexts that do not even involve volition (a cross-linguistically common development; Bybee et al., 1994).

Some examples of narrowing can be plausibly attributed to an influence of entrenchment. In particular, these are cases that Blank (1999) calls *restriction to the prototype*. For example, in English, *hound* has become restricted from any dog to a hunting dog. Blank argues that this change occurred in hunting context, where *dog* always meant HUNTING DOG. In Latin, *frumentum*, CEREAL, has become restricted to mean WHEAT. Similarly, *corn*, which is used in British English to refer to all kinds of grains, has become restricted to mean MAIZE in the US. In these cases of restriction to the prototype, a form that (nearly) invariably occurs in a specific semantic context becomes restricted to that context. However, other cases of restriction to the prototype are more difficult to explain this way. For example, in Latin, *homo* became restricted from PERSON to MALE PERSON (Blank, 1999). While Blank argues that this development is due to the prototypicality of men in a patriarchal society, some explanation is needed for why the

form *homo* would become restricted to men if it was continuing to be used to also refer to women. Entrenchment does not help with cases like these.

Most examples of narrowing are likely the result of pre-emption during form competition. Here, *homo* likely became restricted to men because of pre-emption by *femina*. Extension of a frequent form to a new context is a common cause of form competition, which can result the death of a form already associated with that context. Thus, the regular *-ed* suffix drove many irregular past tense forms into obsolescence. In other cases, the extended form can take over only a part of the semantic space occupied by another form, restricting that form to a narrow niche. For example, as *going to* was extended to express the meaning of FUTURE, it took over the semantic space of proximal future leaving *will* the meaning of DISTAL FUTURE (Poplack & Tagliamonte 2000). Older forms often survive in more conservative, rare constructional contexts associated with more formal styles (e.g., subordinate clauses; Bybee, 2001; Bybee et al., 1994).

In some cases, competition can be seen to result in a semantic push chain, whereby the extension of a form to a new context causes the form that used to occupy it to shift out of its prototypical meaning, developing a new ecological niche (Aronoff, 2016; MacWhinney, 1987, 1993; see also Lindsay & Aronoff, 2013). In the process, it may push another form out of the semantic space corresponding to its new prototype. As MacWhinney (1987, p. 292) explains:

> The situation is much like that in population genetics. If two species of birds are competing for exactly the same ecological niche, one of the two species will win out and the other species will move into another niche or die out altogether. The niche of the losing species may overlap partly with that of the winning species, but it cannot be an exact overlap.

A good example is presented by a variationist study of English future expressions (Torres Cacoullos & Walker, 2009). When *will* primarily competed with *shall* rather than *going to*, the prototypical uses of *will* involved *non*-first-person subjects. As *going to* expanded into the domain of future, it took this niche over, causing the prototype of *will* to shift to first person subjects. Thus, at present *will* is favored by first-person subjects even though it used to be disfavored by them. *Shall* has been almost pushed into obsolescence, holding on only in the context of offers, where the semantic distinctions between *shall*, *will* and *going to* are retained (cf. *Shall I call you a cab*? vs. *Will I call you a cab?* and *Am I going to call you a cab?*) Note that, in this context, *going to* has pushed *will* out of its original meaning of volition: *Will I call you a cab?* does not question the speaker's willingness to do so as much as *Am I going to call you a cab?* does. In Chapter V, I show how semantic push chains can result from competition and pre-emption.

2.3 The relationship between acquisition and change

In contrast to generative linguistics, which attributes language change to imperfect acquisition and reanalysis of the input by children (e.g., Lightfoot, 1999; see also Hudson Kam & Newport, 2005), usage-based linguistics has been generally wary of attributing language change to young children (Bybee, 2010; Kapatsinski, 2018; Kapatsinski, Easterday & Bybee, 2019; Slobin, 1994, 2002; Vihman, 1980). This skepticism comes from the fact that young children are generally not in the social position to spread their innovations, and are shaped into conformity with the norms of the surrounding speech community as instantiated by their parents. It is also motivated by empirical comparisons between the typology of sound change and the typology of pronunciation errors in young children, which bear little resemblance to each other. For

example, unstressed vowel reduction is a late-acquired process that is ubiquitous in language change whereas major place consonant harmony is ubiquitous in child speech yet unattested in language change (Bybee, 2010; Kapatsinski et al., 2019; Vihman, 1980). Instead, the sound changes we observe tend to parallel rapid-speech phenomena in adults (Browman & Goldstein, 1992; Kapatsinski et al., 2019; Mowrey & Pagliuca, 1995).

We must therefore be careful in drawing parallels between overextension errors observed with young children and extensions seen in language change. A major aim of this dissertation is to show that frequent words are in fact preferentially extended to novel related uses by adults because they are more accessible than their semantic competitors. That is, the extension of a word to new uses by an adult can be driven by the same factors driving over-extensions in children. The only difference between extension and over-extension is that over-extensions, being made by young children, are perceived to be unacceptable and squashed by the speech community. Overextensions are extensions that children grow out of.

2.4 The relationship between frequency, extension and entrenchment

The effects of word frequency in acquisition and change are at first glance contradictory. In acquisition, frequent words are the first to be restricted to acceptable uses, with rarer words being over-extended for a longer time. In language change, frequent words are the ones that tend to be extended to new uses. What resolves this contradiction?

First, the contradiction is likely only apparent because the studies of retreat from overgeneralization are focused on explaining the trajectory of retreat rather than the trajectory of the onset of overgeneralization. While children correct over-extensions of

frequent forms before they correct over-extensions of rare forms, they may also be more likely to over-extend frequent forms prior to curbing the over-extension. For example, in morphology, few children over-extend the rare irregular patterns compared to the frequent regular ones. Just as frequent forms replace rare forms in the diachronic process of paradigm leveling (Bybee & Brewer, 1980; Tiersma, 1982), children over-extend frequent members of a morphological paradigm to replace ones that are less frequent (Aguado-Orea & Pine, 2015; Dąbrowska & Szczerbiński, 2006; Engelmann et al., 2019; Leonard et al., 2002; Marchman, 1997; Matthews & Theakston, 2006; Räsänen et al., 2016; Rubino & Pine, 1998; Theakston et al., 2003; Theakston & Rowland, 2009; see also Rescorla, 1980).

Second, in this dissertation, I argue that high frequency favors extension in production while leading to entrenchment in comprehension and judgment. As discussed earlier in this chapter, Zipf's artisan both knows more about the ways in which frequent tools are and are *not* used and is at the same time more likely to reuse frequent tools for novel purposes. The artisan is likely to be confident that a novel use of a familiar tool is in fact novel, and at the same time be likely to use a familiar tool when the novel task is at hand. In the next chapter, I show that this dissociation is predicted by all models of learning form–meaning mappings.

As discussed earlier in this chapter, most evidence for entrenchment has come from judgment studies rather than production. Two studies examined entrenchment in production. One of these is in fact the first study that provided direct evidence for entrenchment (Brooks, Tomasello, Dodson, & Lewis, 1999). Brooks and colleagues tested the predictions of the entrenchment hypothesis in children ages 3, 4–5, and 8.

Children participated in an elicited production task involving intransitive and transitive verb pairs that differed in frequency, as measured by Age of Acquisition (AOA). These included examples such as early AOA *come* and late AOA *arrive* for intransitive pairs and early AOA *take* and late AOA *remove* for transitive. The child's task was to watch an enactment of actions involving puppets and objects and then describe the actions in response to the experimenter's question. (Brooks et al., 1999).

This procedure induced overgeneralization errors of intransitive verbs in transitive constructions such as **The cow is gonna arrive it* and overgeneralization errors of transitive verbs in intransitive constructions such as **It's hitting*. Overall, children, even as young as 3 years old, were more likely to overgeneralize late AOA verbs to constructions with opposite transitivity status than early AOA verbs. Ruling out other confounding factors, Brooks and colleagues attributed this effect to the frequency of the verbs, providing support for entrenchment in production.

However, a close reading of Brooks et al. (1999) shows that their methodology purposely reduced accessibility differences between verbs. Brooks et al. (1999, pp. 1328–1329) write that "Because the late AOA verbs may not have been in the productive vocabularies of all of the children, it was necessary to provide children with a great deal of exposure to the verbs to ensure that children would use both early and late AOA verbs. The experimenter also encouraged the children to use the target verbs by asking questions at regular intervals (e.g., Can you say *disappear*? Say *disappear*). Unfortunately, even with the experimenter modeling each target verb in over 50 utterances, occasionally a child would fail to produce even a single utterance containing the target verb." Because repetition priming boosts the activation of a rare word more than it boosts the activation

of a rare word (e.g., Plaut & Booth, 2000), this procedure reduces or eliminates activation differences between the two groups of words in the study. The same criticism also holds for the only other study reporting entrenchment in production (Blything et al., 2012). Even though I do not dispute that these studies have observed entrenchment, these findings do not contradict the hypothesis that frequent words should be preferentially extended to novel uses *because they are more accessible than rare words*. When accessibility differences between frequent and rare forms are leveled, I predict that speakers will not prefer to choose frequent forms to express novel meanings.

One way to eliminate differences between frequent and rare forms is to ask participants to choose among these forms in a two-alternative forced-choice (2AFC) task (Luce & Pisoni, 1998; Sommers, Kirk & Pisoni, 1997). In Chapters III–VI, I therefore compare a production task to a forced form choice task. If high token frequency of a form makes the form likely to be chosen to express a novel meaning because it increases the form's accessibility, high token frequency of a form should make the form likely to be chosen for production but should not increase its likelihood of selection in a forced choice task. Instead, in a forced choice task, frequent forms should either be as likely to be selected as rare forms or, if entrenchment occurs in production, they should be disfavored compared to rare forms.

2.5 Conclusion

Studies of comprehension suggest that frequently encountering a form–meaning pairing convinces the learner than the form cannot be used in any other way (Braine & Brooks, 1995; Brooks & Tomasello, 1999; Regier & Gahl, 2004; Stefanowitch, 2008; Xu & Tenenbaum, 2007). Nonetheless, frequent forms are the ones most likely to be

extended to new uses over historical time. Using a frequent form in a novel way seeds the process of language change because that novel use can then be picked up by others, spreading through the speech community. As the novel use diffuses through the community, it becomes conventional. Over historical time, extension of frequent forms results in the well-documented correlation between frequency of use and number of senses: in every language, it is the most frequent forms that are most polysemous (Piantadosi et al., 2012; Zipf, 1949).

Conventionalization of extensions is the primary mechanism behind the diachronic process of grammaticalization (Bybee, 2010; Heine, 2011). The importance of this diachronic process can hardly be overstated as it is the primary source of grammar: almost all grammatical morphemes, whether bound affixes or independent functors like prepositions, determiners or auxiliaries are former lexical words that have been gradually extended to more and more uses (Bybee, 2003; 2010; Christiansen & Chater, 2016).

Despite the correlation between frequency and semantic extension, the causal mechanisms behind grammaticalization remain controversial. For example, Haspelmath (1999) has argued that increases in frequency seen in grammaticalization are caused by the extension of the grammaticalizing form to new uses, which are in turn caused by semantic broadening. Bybee (2003) agrees that semantic broadening causes extension but suggests that high frequency causes semantic broadening. Like Haspelmath (1999), Heine (2011) does not allocate frequency a causal role in the process but suggests that extensions result in broadening.

This chapter has argued that entrenchment in comprehension can co-exist with extension in production. That is, a speaker can believe that they are using a word in a

novel way, and yet use the word that way because it is the most accessible word given the meaning to be expressed. Furthermore, dissociations between production and comprehension are expected even if form–meaning associations are bidirectional, and semantic representations are fully mature. A frequent word may be the most accessible form given a novel meaning even if the listener is unlikely to activate that meaning when presented with the form.

In the next chapter, I demonstrate that entrenchment and extension of frequent forms can indeed co-exist within the same speakers, and show that this is predicted by Hebbian learning of associations between semantic cues and form-based outcomes (Chapter III). Furthermore, entrenchment and extension also co-exist with statistical pre-emption. Pre-emption means that forms compete to express a meaning. If pre-emption coexists with entrenchment and extension, extension of a form can result in a push chain in semantic space. That is, an extended form can push another form out of its functional niche. This process is demonstrated in Chapter IV.

CHAPTER III

ACCESSIBILITY DRIVES EXTENSION

3.1 Experiment I[5]

The primary goal of this experiment was to test whether—as predicted by Zipf's tool analogy—entrenchment of frequent forms in comprehension can co-exist with extension of frequent forms in production, and whether the effect of frequency on form choice in production is mediated by the effect of frequency on accessibility. To this end, participants were exposed to a miniature artificial language comprised of four morphological constructions, and each participant's knowledge was tested using comprehension, production and forced form choice tasks.

I expect entrenchment effects in comprehension: increasing the frequency of a form–meaning pairing should make participants more likely to select the meaning when given the form, and less likely to select a related meaning with which the form has not been paired.

I expect frequent forms to have an advantage in production, being preferentially selected even when the meaning is novel. This means divergent effects of frequency in comprehension and production: the forms that are least likely to be mapped onto a novel meaning in comprehension should be the forms most likely to be selected for expressing the novel meaning in production. In other words, infrequent forms should be mapped onto the novel meaning in comprehension but not used to express that meaning in

[5] The experiment reported in this chapter has previously been published in Harmon, Z., & Kapatsinski, V. (2017). Putting old tools to novel uses: The role of form accessibility in semantic extension. *Cognitive Psychology*, *98*, 22–44.

production; frequent forms should not be mapped onto the novel meaning in comprehension but should be used to express it in production.

Finally, a forced form choice task allows us to evaluate whether the effect of frequency on form selection is mediated by its effect on accessibility. In the form choice task, as in production, the participant is given a meaning to express and asked to choose a form but the options to choose amongst are made available. Instead of retrieving a form from memory, the participant needs only to choose between the forms s/he hears. In this way, infrequent forms are made as (or almost as) accessible as their frequent competitors in the moment of selection. Indeed, previous studies on picture naming show that restricting the context, which results in pre-activation of related word forms, reduces frequency effects on naming latencies (Gollan et al., 2011; Griffin & Bock, 1998). Absence of frequency effects with closed-set tests of word recognition support this point as well (Clopper, Pisoni, & Tierney, 2006; Sommers, Kirk, & Pisoni, 1997). Thus we expect the effect of frequency on form choice seen in production to diminish or disappear in this task, as the effect crucially relies on accessibility differences between frequent and infrequent forms.

3.1.1 Languages and predictions

Two artificial languages were used in the experiment: Dan and Nem (Figure 3.1). Dan and Nem had the same morphological constructions, and shared the same set of nouns in their vocabulary. The nouns were 30 one- and two-syllable nonce nouns (e.g., *chool* and *osto*).

Dan and Nem consisted of the same four suffixes: -*sil*$_{\text{PL}}$, -*dan*$_{\text{PL}}$, -*nem*$_{\text{DIM}}$, and -*shoon*$_{\text{DIM}}$. The suffixes -*sil*$_{\text{PL}}$ and -*dan*$_{\text{PL}}$ are plural while -*shoon*$_{\text{DIM}}$ and -*nem*$_{\text{DIM}}$ are

diminutive. During the exposure phase, -*sil*$_{PL}$ and -*dan*$_{PL}$ were always paired with a picture of multiple large creatures, while -*nem*$_{DIM}$ and -*shoon*$_{DIM}$ were paired with a picture of one small creature, the identity of the creature varying with the stem. An unaffixed stem was always paired with a single large creature. As a result, -*sil*$_{PL}$ and -*dan*$_{PL}$ could be thought of as either simply plural (PL) or plural non-diminutive (BIG.PL). Likewise, -*shoon*$_{DIM}$ and -*nem*$_{DIM}$ could be thought of as either diminutive (DIM) or diminutive singular (DIM.SG). The interpretations participants actually infer from this experience were probed in the test phase.

In each language, one suffix is more frequent than the others: -*dan*$_{PL}$ is the frequent suffix in the Dan language and -*nem*$_{DIM}$ is the frequent suffix in the Nem language. The thicker lines in Figure 1 highlight this difference in frequency. The identity of the frequent suffix is the only difference between the two languages.

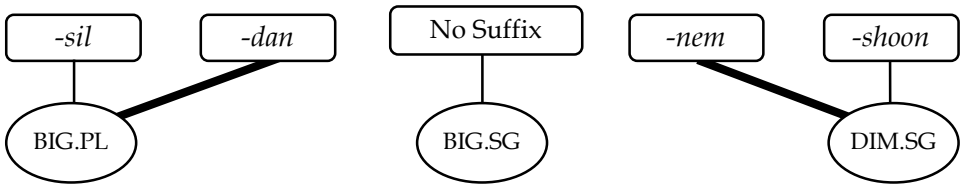

Figure 3.1. Structure of the two languages Dan (where -*dan*$_{PL}$ is frequent) and Nem (where -*nem*$_{DIM}$ is frequent) during exposure.
Thick lines represent the frequent form–meaning mappings: the plural -*dan*$_{PL}$ in the Dan language and the diminutive -*nem*$_{DIM}$ in the Nem language.

Crucially, participants exposed to either language are tested on the diminutive plural meaning (DIM.PL), which is never presented to learners during the exposure phase. Use of a suffix with the DIM.PL meaning constitutes extension of the suffix to a novel, though related, meaning. Presenting participants with the diminutive plural

meaning during the comprehension test further allows us to distinguish these possible interpretations: a participant who has mapped a suffix onto, say, simply DIM should be equally likely to choose plural and singular diminutive referents to map the suffix onto. A participant who has mapped it onto DIM.SG should choose the single small creature over multiple small creatures. We hypothesized that frequent forms will be used to refer to DIM.PL referents in production but would not be mapped onto the DIM.PL meaning in comprehension. Conversely, infrequent forms would be more equally mapped onto their original meanings and DIM.PL in comprehension but would not be used to express DIM.PL in production.

3.1.2 Methods

3.1.2.1 Participants

Seventy adult native speakers of American English, all undergraduate students at the University of Oregon, participated in the experiment. Participants received course credit for their participation. Each participant was exposed to only one of the two languages (35 participants experienced Dan, the language where -*dan*$_{\text{PL}}$ was the frequent suffix, while 35 others experienced Nem, the language where -*nem*$_{\text{DIM}}$ was the frequent suffix.).

3.1.2.2 Tasks

3.1.2.2.1 Exposure

Each exposure trial consisted of the presentation of a picture on the computer screen. One hundred milliseconds later, the name was presented auditorily over headphones, while the picture stayed on screen. Once the sound finished playing, the

picture was removed and replaced with the next picture and its corresponding name. The pictures were of three types: a single large creature, paired with the bare stem, multiple (five to nine) large creatures paired with the stem suffixed with -*sil*$_{PL}$ or -*dan*$_{PL}$, or a miniature version of the creature paired with the stem suffixed with -*nem*$_{DIM}$ or -*shoon*$_{DIM}$. The time lapse between trials was four hundred milliseconds. Trial order was randomized for each participant.

3.1.2.2.2 Production

Exposure was followed by an elicited production test. In this test, participants were presented with novel non-diminutive singulars paired with novel creatures. Each trial began with the presentation of a single large novel creature on the computer screen (Figure 3.2, panel A). One hundred milliseconds later, the name of the novel object was presented auditorily over headphones, as in the training stage. Once the sound finished playing, the picture was removed and replaced with a display of four pictures representing the four different meanings. Two hundred milliseconds later, three of these pictures disappeared, leaving the participants with one target picture to name (i.e., one meaning to express). The reason for presenting all the pictures at first was to make sure that the participants differentiated between the four meanings (particularly, with respect to realizing whether the target picture was small or large. Participants were asked to generate the corresponding form for the target meaning using the presented stem and their knowledge of the language and to say the form aloud. They had five seconds to do so. After this time period, the test automatically continued to the next trial.

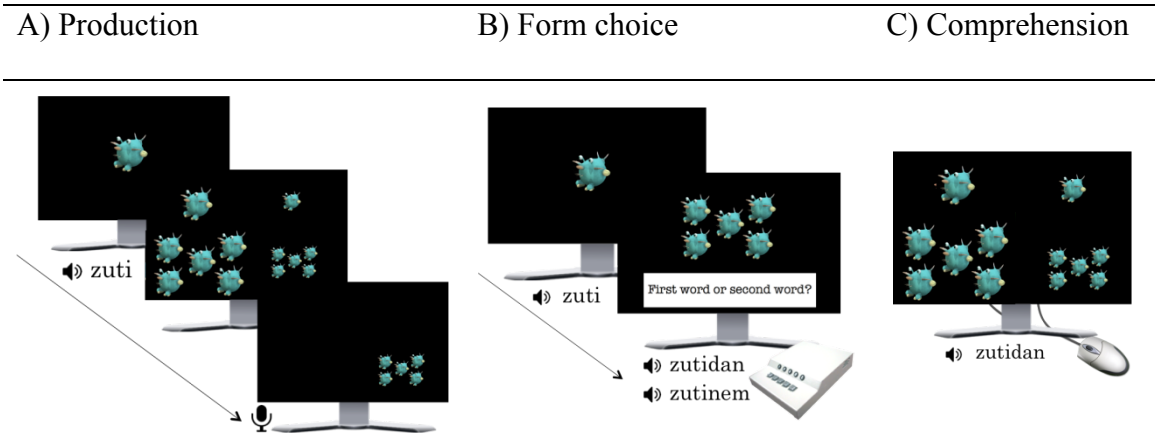

Figure 3.2. The three test tasks in the experiment, in order of presentation: Production (Panel A), 2AFC form choice task (Panel B), and form-to-meaning mapping task (Panel C).

3.1.2.2.3 Form choice

This task was a two-alternative forced form choice (2AFC) task: participants were asked to map a meaning onto one of two forms by pressing one of two buttons on a button box. As in production, participants were presented with the base, unaffixed form and the corresponding picture. Participants were then asked to choose between two forms to express one of the other three meanings. The onset of the first form was one hundred milliseconds after the corresponding picture appeared, with an ISI of one hundred milliseconds between the two forms. Immediately after the two forms were played, a prompt appeared on the screen under the picture of the object: "First word or second word?" (Figure 3.2, panel B). Participants pressed the corresponding button to choose between the two forms. They had three seconds to do so. As soon as they pressed a button, the experiment continued to the next trial. The goal of this test was to assess whether there is an effect of frequency on form choice when accessibility differences between frequent and rare forms have been attenuated. Because this task followed the

production task in this experiment, we did not consider it necessary to remind the participants of the semantic space by briefly flashing the four possible meanings on every trial. The results suggest that they remembered the space, as they rarely chose a form for a meaning that did not share features with the meaning of the form during training, e.g. choosing -*nem*$_{\text{DIM}}$ or -*shoon*$_{\text{DIM}}$ to refer to large creatures. In Experiment III, this task comes first, immediately after exposure, and does feature the four-picture flash.

3.1.2.2.4 Form-to-meaning mapping

In this task, participants were required to map a form onto one of four meanings, by clicking on a picture on the computer screen. At the beginning of each trial, four pictures representing the four possible meanings in the languages were presented on the computer screen (Figure 3.2, panel C). A hundred milliseconds later a form was auditorily presented to the participant, and the mouse cursor was activated to allow the participants to click on one of the four pictures. Participants had three seconds to do so. The experiment advanced to the next trial upon response.

3.1.2.3 Stimuli

The same test trials were used for all participants, but the training varied by language. The training comprised 180 trials (45 items that were repeated in random order in 4 cycles). Ten stems were presented to learners in each cycle. Five stems were used once with each suffix (e.g., *ostosil, ostodan, ostonem, ostoshoon*; 20 items). The other five were used without a suffix once (5 items) and were used four other times with the frequent suffix (20 items).

The same elicited production test was used for all subjects. This test comprised 75 items (25 stems by 3 suffix meanings). Of the 25 stems, 15 were novel and 10 were the same stems learners were exposed to during training.

The form choice task comprised 105 trials. On each trial, participants were presented with two alternative suffixed forms of the same stem. Order of presentation of the suffixes was counterbalanced for each suffix pair. The crucial trials involved the choice between -*dan*$_{\text{PL}}$ and -*nem*$_{\text{DIM}}$ (30 trials, 10 each for PL, DIM and DIM.PL meaning). The rest of the trials involved -*sil*$_{\text{PL}}$, -*shoon*$_{\text{DIM}}$, and the compositional suffix -*dannem*. Of the 30 stems in the language, 10 were used in this task: 5 novel stems that learners did not experience during either exposure or production and 5 stems that learners were exposed to during production only.

The form-to-meaning mapping task comprised 60 trials. Six forms were tested during this task, 10 times each: The unaffixed stem, -*sil*$_{\text{PL}}$, -*dan*$_{\text{PL}}$, -*nem*$_{\text{DIM}}$, -*shoon*$_{\text{DIM}}$ and finally a compositional suffix, -*dannem*. All four possible meanings of BIG.PL, BIG.SG, DIM.PL, and DIM.SG were available for mapping to each suffix. Of the 10 stems used in this task, 5 were shared with both production and exposure and 5 were shared only with production.

The visual stimuli were a set of novel creature images retrieved from the website sporepedia.com (http://www.sporepedia.com). All pictures were presented on a black background. The auditory stimuli were recorded by a female native speaker of American English. The stimuli were then edited in Praat (Boersma & Weenink, 2009) to remove the silence at the beginning and end of each word.

3.1.2.4 Analysis

All analyses were performed using mixed effect logistic regression with maximal random effects structure supported by the design. For each task, significance was then determined using nested model comparison with log likelihood tests (Barr, Levy, Scheepers, & Tily, 2013). For all tests, the dependent variable was a binary one indicating the participants' choice of form or meaning. All analyses were performed in R version 3.2.1 (R Core Team, 2015) using the lme4 package (Bates, Maechler, Bolker, & Walker, 2015). Support for the null was calculated using the BIC approximation to the Bayes factor (Wagenmakers, 2007). Graphs were produced using the ggplot2 (Wickham, 2009), mgcv (Wood, 2011) and itsadug (van Rij et al., 2016) packages.

Frequency effects were evaluated by examining the effect of Language (Dan vs. Nem) on form selection in production and forced form choice tasks and on meaning choice in the comprehension task. Language identifies the frequent form: -dan_{PL} is the frequent form in Dan language and -nem_{DIM} is the frequent form in the Nem language. The effect of form frequency is then evaluated by comparing -dan_{PL} in Dan and -nem_{DIM} in Nem on the one hand to -dan_{PL} in Nem and -nem_{DIM} in Dan on the other.

By using Language as our operationalization of frequency, we are controlling for any effects that might be due to a preference for a particular form or meaning: the form/meaning combination that is frequent in Dan is infrequent in Nem and vice versa. Therefore, a preference for the frequent form over the infrequent form cannot be due to the identity of the form or the meaning with which it was paired in training. By restricting analysis to -dan_{PL} and -nem_{DIM}, we ensure that frequent and infrequent forms have

equally infrequent competitors. Therefore, the form frequency effect cannot be attributed to a competitor frequency effect.

3.1.3 Results

We discuss tasks in the order in which participants encountered them. For this experiment, the participants encountered the production task immediately after exposure, followed by the form choice task, followed by comprehension. This order was chosen to prevent carry-over of training during test to the production task. The tasks are re-ordered in Experiments II and III.

3.1.3.1 Production

The analysis was performed on 70.86% of responses, where the response conformed to either one of the four suffixes. Of the excluded responses, 6.28% were missing; 9.35% were unsuffixed forms; 2.02% featured compositional suffixes; 10.46% were mispronunciations of the existing suffixes to the point of un-recoverability (e.g., *-dem, -shil*); and finally, 1.03% were responses where more than one response was provided.

The Production panel in Figure 3.3 shows, for each meaning, the proportion of productions that bore frequent -dan_{PL} and -nem_{DIM} suffixes. In the figure, frequent refers to -dan_{PL} in Dan and -nem_{DIM} in Nem, and infrequent refers to -dan_{PL} in Nem and -nem_{DIM} in Dan. Frequent -dan_{PL} and -nem_{DIM} suffixes are chosen on average 86% of the time in response to their original meaning (PL for -dan_{PL} and DIM for -nem_{DIM}), which is somewhat higher than the probability of that suffix given the meaning in training. When -dan_{PL} and -nem_{DIM} are infrequent, they are as frequent as their synonyms in training, and accordingly participants choose the two synonyms approximately equally often.

Therefore, the results for the original meanings show that the probability of choosing a form given a meaning resembles the probability of experiencing the same form given the same meaning in training.

The diminutive suffix -*nem*$_{\text{DIM}}$ was rarely chosen to express the non-diminutive plural meaning and the plural suffix -*dan*$_{\text{PL}}$ was rarely chosen to express diminutive singular (as shown by the very low bars in the Wrong meaning in the Production panel of Figure 3.3). Thus, participants learned the meanings of the suffixes well, even for the infrequent suffixes like -*nem*$_{\text{DIM}}$ in Dan. They chose a form only when cued by at least one semantic feature with which the form was paired in training. With this background, we can turn to investigate form choice given the novel DIM.PL meaning.

Accessibility differences predict that a frequent suffix should be more likely than an infrequent suffix to be used to express the novel DIM.PL meaning. In contrast, entrenchment predicts that a frequent suffix should be *less* likely than an infrequent suffix to express the DIM.PL meaning. We therefore focused on suffix choice for the DIM.PL meaning. The data were analyzed using a mixed-effects logistic regression model with language (Dan vs. Nem) as the fixed effect of interest and the binary dependent variable of suffix (-*dan*$_{\text{PL}}$ vs. -*nem*$_{\text{DIM}}$). Random intercepts for subjects and noun stems and random slopes for language within stem were also included in the model.

The results support the accessibility hypothesis. A form was significantly more likely to be chosen to express DIM.PL when it was frequently encountered during exposure ($\beta = 8.46$, $z = 2.89$; $\chi^2 (1) = 21$, $p < .0001$).

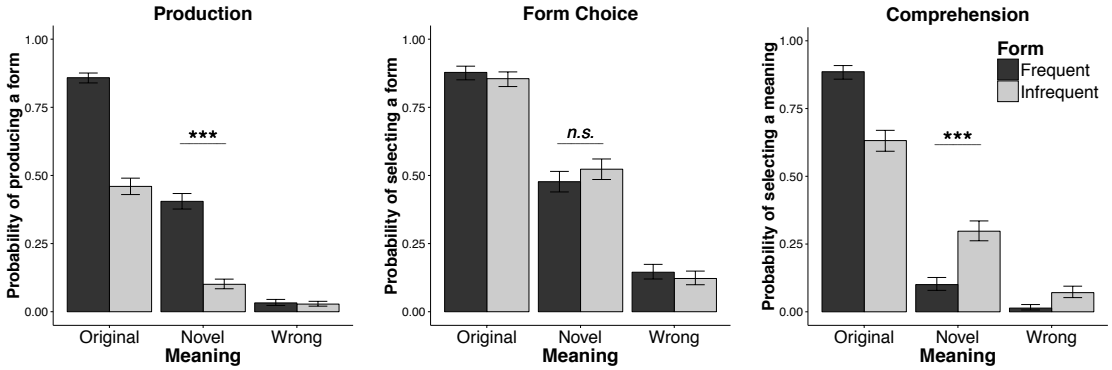

Figure 3.3. The results of Experiment I for the two suffixes -*dan*$_{PL}$ and -*nem*$_{DIM}$. Original meaning = meaning paired with a suffix during training, e.g. BIG.PL for -*dan*$_{PL}$; Novel meaning = DIM.PL, which shares a feature with each of the original meanings; Wrong meaning = meaning paired with a different suffix in training, e.g. DIM.SG for -*dan*$_{PL}$. Frequent suffix = -*dan*$_{PL}$ in Dan language and -*nem*$_{DIM}$ in Nem language. Infrequent suffix = -*dan*$_{PL}$ in Nem and -*nem*$_{DIM}$ in Dan. Left panel: Form choice probabilities in the production task. The distance between the tops of the bars and 100% for the original meaning in the Production panel is composed largely of probability of choosing the synonymous suffix: -*shoon*$_{DIM}$ over -*nem*$_{DIM}$ or -*sil*$_{PL}$ over -*dan*$_{PL}$. Middle panel: Form choice probabilities in the forced form choice task on the -*dan*$_{PL}$ vs. -*nem*$_{DIM}$ trials. Right panel: Meaning choice probabilities in the comprehension task. In this experiment, production preceded form choice, which preceded comprehension. (*** means $p < .001$).

An examination of the suffixes that comprised the compositional suffix in participant's productions adds support to the accessibility hypothesis. While only 72 compositional tokens were produced in reference to DIM.PL, they were produced by 9 different subjects with 25 different stems, thus both stem and subject random intercepts could be fit. Even with these random intercepts, the data strongly support an influence of accessibility: ($\beta = 45.04$, $z = 5.088$; $\chi^2(1) = 25.35$, $p < .0001$): 93% (67/72) of compositional responses included a frequent suffix.

3.1.3.2 Form choice

The analysis was conducted on 99.46% of the data, where participants provided a response. These results are presented in the Form Choice panel of Figure 3.3. The data

were analyzed using a mixed-effects logistic regression model with language as the fixed effect of interest and the binary dependent variable of suffix (*-dan*$_{PL}$ vs. *-nem*$_{DIM}$) chosen to express the novel meaning DIM.PL. Random intercepts for subjects and noun stems and random slopes for language within stem were included in the model.

As the aim of this task was to reduce accessibility differences between frequent and infrequent forms, we expected the effect of language to be reduced or eliminated. As predicted, the analysis revealed no significant effect of language on suffix choice for the DIM.PL meaning ($\beta = -0.2317$, $z = -0.423$; $\chi^2(1) = 0.1741$, $p = .68$). The BIC approximation to the Bayes factor suggests that we have sufficient power to accept the null, as the probability of the null hypothesis given the data is 96%. Figure 3.3 suggests that the lack of frequency effects in this task is not due to participants behaving randomly: Participants seldom select forms that are not appropriate for the presented meaning (witness the low Wrong bars). That is, *-dan*$_{PL}$ is seldom chosen to express DIM and *-nem*$_{DIM}$ is seldom chosen to express PL. Participants appear to know that *-dan*$_{PL}$ is plural, and *-nem*$_{DIM}$ is diminutive but have no preference to map DIM.PL onto one or the other.

To assess whether the effect of form frequency in production differs significantly from its effect in the form choice task, we ran a mixed effects logistic regression model that included task and language, as well as, crucially, an interaction between the two as fixed effects. The binary dependent variable was the choice between the two forms *-dan*$_{PL}$ and *-nem*$_{DIM}$ in response to the novel DIM.PL meaning. Random slopes for task within subject and language within stem were included in the model as well. Results revealed a significant interaction between task and language ($\beta = -6.56$, $z = -3.17$; χ^2 (1)

= 19.3, $p < .0001$), providing support for the claim that the effect of frequency in the two tasks differed.

3.1.3.3 Comprehension

The analysis was conducted on all trials for which participants provided responses (99.29%). The results are shown in the Comprehension panel of Figure 3.3. The data were analyzed using a mixed-effects logistic regression model with frequency as the fixed effect of interest and the binary dependent variable of mapping a suffix onto its original meaning (-*dan*$_{PL}$ to BIG.PL and -*nem*$_{DIM}$ to DIM.SG) versus the novel meaning of DIM.PL. Random slopes for frequent suffix within both subject and stem were included in the model.

The entrenchment hypothesis predicts that mapping onto DIM.PL should become less likely as frequency of co-occurrence with the original meaning increases. As participants repeatedly encounter -*dan*$_{PL}$ paired with non-diminutive plurals, they should become more and more confident that -*dan*$_{PL}$ is not only plural but also non-diminutive. (Unaffixed BIG.SG forms ensure that the DIM forms are more DIM than SG and that PL forms are more PL than BIG.) As participants repeatedly encounter -*nem*$_{DIM}$ paired with diminutive singulars, they should become more and more confident that -*nem*$_{DIM}$ is not only diminutive but also singular. The results revealed that a form was mapped onto DIM.PL significantly less often when it was frequent during exposure ($\beta = 5.883$, $z = 4.00$; χ^2 (1) = 17, $p < .0001$). In other words, increasing the frequency of a form–meaning mapping during exposure made participants more likely to map the frequent form onto its original meaning.

3.1.4 Discussion

In this experiment, we manipulated suffix frequency. In each language, one of the suffixes presented to participants was more frequent than the others. The frequent suffix was -*dan*$_{PL}$ in the Dan language and -*nem*$_{PL}$ in the Nem language. We examined the effects of frequency on form choice in production and meaning choice in comprehension.

On the basis of Zipf's tool analogy—supported by the empirical correlation between frequency and polysemy (Bybee, 2003; Piantadosi et al., 2012; Zipf, 1949)—and data on overextensions in child language (e.g., Bloom, 1973; Clark, 1973; Gershkoff-Stowe & Smith, 1997; Huttenlocher, 1974; Naigles & Gelman, 1995), we predicted that frequent forms are especially likely to be extended to new meanings in production. This prediction was confirmed in the present experiment: both -*dan*$_{PL}$ and -*nem*$_{DIM}$ were more likely to be used for expressing the novel DIM.PL meaning when they were frequent (-*dan*$_{PL}$ was used more in the Dan language, and -*nem*$_{DIM}$ was used more in the Nem language).

We hypothesized that extension of frequent forms in production can co-exist with entrenchment of frequent forms in comprehension documented by several studies in first language acquisition (e.g., Ambridge et al., 2008; Brooks & Tomasello, 1999; Wonnacott, Newport, & Tanenhaus, 2008). The results of the form-to-meaning mapping task were consistent with this prediction: forms were mapped onto DIM.PL less often when they were frequent than when they were infrequent.

Note that frequency has divergent effects in production and comprehension. High frequency of a form makes participants *less* likely to select the unfamiliar DIM.PL as a possible meaning of the form in comprehension, but at the same time *more* likely to use

the form to express DIM.PL in production. In other words, the forms participants are most likely to use to express DIM.PL are the forms that they are least likely to extract the DIM.PL meaning from. Figure 3.4 presents the individual subject data in the production and comprehension tasks of experiment I. The dissociation between these two tasks is also evident from the switch in subjects' preferences of the two suffixes -*dan*$_{PL}$ and -*nem*$_{DIM}$ to the meaning of DIM.PL in the two languages. Although this effect is stronger in Dan, it is apparent in both languages. I will further discuss individual subject data and relevant analyses in the next chapter in a comparison between the three experiments, but am including the individual data points here for the purpose of transparency lacking in bar graphs.

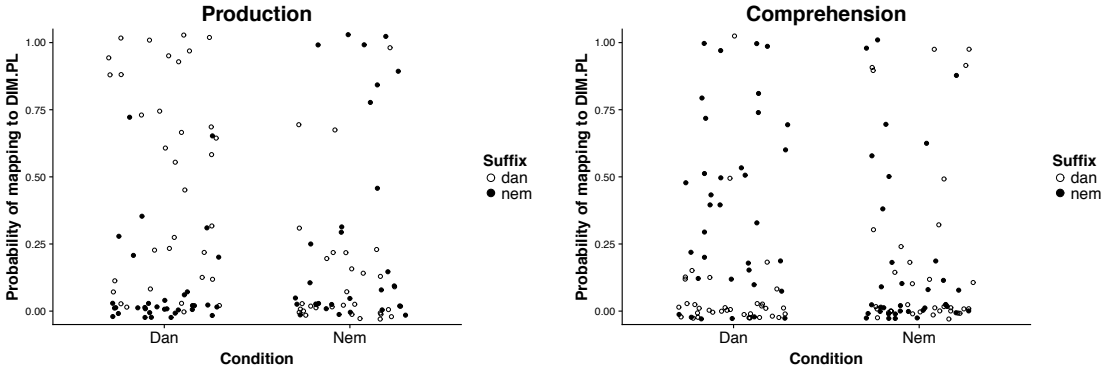

Figure 3.4. Individual participants' mappings of the suffixes -*dan*$_{PL}$ and -*nem*$_{DIM}$ to DIM.PL meaning in the production and comprehension tasks in Experiment I. Each white dot corresponds to an individual subject's use of -*dan*$_{PL}$ and each black dot corresponds to an individual subject's use of -*nem*$_{DIM}$. In production, the dots represent probabilities of choosing a suffix in response to DIM.PL. In comprehension, the dots represent probabilities of choosing DIM.PL in response to a suffix.

The results of the form choice task showed no effect of frequency on the learner's choice of form to express the novel meaning. The production and form choice tasks both involve selecting a form when cued with meaning. The notable difference between the

two tasks is that forms are presented to the learner in the form choice task, which renders infrequent forms accessible, leaving less room for frequency to influence accessibility (Clopper et al., 2006; Griffin & Bock, 1998; Gollan et al., 2011; Luce & Pisoni, 1998; Sommers et al., 1997).

Making infrequent forms (almost) as accessible as frequent forms eliminates the effect of frequency on form choice seen in production. Therefore, it appears that the effect of frequency on form choice in production is indeed due to frequent forms being more accessible. The retrieval demands inherent to production can push the speaker to produce the most accessible form to express a meaning, even when a less accessible form would be a better cue to the meaning for the same speaker. While others have made similar arguments to explain children's overextension errors (Gershkoff-Stowe, 2001; Naigles & Gelman, 1995), this is—to our knowledge—the first demonstration of this effect in unprimed adult speech production.

I suggest that the effect of accessibility on form choice underlies the diachronic process of semantic extension of frequent forms to new uses (Zipf, 1949; see also Kapatsinski, 2009). Whenever a speaker wishes to express a meaning, a conventionalized expression for that meaning may be inaccessible to the speaker in the moment of production—or at least it may be less accessible than a semantic competitor. The speaker will then use that competitor to express the meaning. The speaker's innovative use of the form can then be taken up by the listener. If and when the innovative use propagates through the speech community, a language change will have occurred. Note that the speaker who uses, say, -*dan*$_{\text{PL}}$ to refer to DIM.PL need not *believe* that -*dan*$_{\text{PL}}$ means DIM.PL or even that it is an appropriate form to use for this purpose. Indeed, the speaker

can be quite certain that the form is not used that way. The speaker's interlocutor has no access to these beliefs. All s/he knows is that the speaker has *used* the form -*dan*$_{\text{PL}}$ to express DIM.PL, providing positive evidence that the form can be used to express this meaning. That interlocutor can then use the form to mean DIM.PL with the firm belief that such use is acceptable. By simply using the form in a novel context, the speaker seeds a change in its meaning. In language change, changes in use often lead changes in belief.

The next chapter examines two alternative orders of test tasks designed to evaluate possible effects of learning during test and transfer across tasks. The primary aim of these follow-up experiments is to ensure that the differences between the three tasks described above could not be attributed to where these tasks appear in the task sequence. To this aim, the present experiment—in which the production task came first, immediately following exposure—was supplemented with two other experiments, each of which presents a different task first. However, before continuing to the next experiments, I show that the production–comprehension dissociation observed above is expected under almost any Hebbian account of learning form–meaning associations.

3.2 A Hebbian account of the production–comprehension dissociation[6]

The evidence reviewed so far suggests a production–comprehension dissociation at least at some point during learning: forms entrench in comprehension but not production. This section reports a simple computational model to demonstrate that this

[6] Parts of this section have appeared in Kapatsinski, V. & Harmon, Z. (2017). A Hebbian account of entrenchment and (over)-extension in language learning. *Proceedings of the Annual Conference of the Cognitive Science Society, 39,* 2366–2371.

dissociation is expected in Hebbian accounts of learning where forms and meanings are linked with bidirectional connections.

The choice of a Hebbian learning framework is motivated by the absence of entrenchment in production, which inevitably arises in error-driven learning (Gluck & Bower, 1988; Rescorla & Wagner, 1972). In brief, pre-emption and entrenchment correspond to outcome competition and cue competition respectively. In production, semantic features serve as cues to select a form, which serves as the outcome (e.g., Dell, 1986; Levelt, 1989). According to pre-emption, a form blocks another form from expressing the same meaning, because the meaning cues the pre-emptor more strongly than it cues the pre-empted form. This is, of course, inevitable because the speaker needs to choose one form to produce. According to entrenchment, the activated semantic cues do not cue a form that frequently co-occurs with *other* semantic cues as strongly as they cue a form that is rare in other contexts. Entrenchment in production therefore corresponds to a well-established effect in associative learning, namely the *base rate* or *associative interference* effect (Rescorla, 1968; see also Gluck & Bower, 1988; Medin & Edelson, 1988).

The base rate effect arises in error-driven learning models, whereby outcomes co-occurring with a certain cue are associated with that cue less strongly if they would be expected to occur just as often in the absence of the cue. The same effect does not arise in Hebbian learning models, due to their insensitivity to prediction error (e.g., Bush & Mosteller, 1951; Hebb, 1949). Absence of entrenchment in production therefore suggests that co-occurring meanings and forms wire together in a Hebbian manner, independently of prediction error.

3.2.1 Computational methods

According to Hebb (1949), neurons that fire together wire together. We assume a distinction between cues and outcomes, where outcomes follow cues. In the experiment reported in this chapter, learners see a word's referent before they hear the word. This order of presentation was chosen to reflect the temporal dynamics of real-life word learning (Pereira et al., 2014). As a result of the pictures being presented first, the meanings serve as cues while forms serve as outcomes. Specifically, suffix forms (*-dan*, *-nem*, etc.) are cued by the semantic features of the referents (BIG, DIM, PL and SG) plus a context cue, present on every trial.

All cue–outcome associations start at zero. On every trial, associations between the cues present on that trial and the following outcomes strengthened by an increment. Unlike error-driven models such as Rescorla & Wagner (1972), we did not multiply the increment in association strength by prediction error. This is what makes the model Hebbian: it does not learn less on trials with unsurprising (or no-longer-surprising) outcomes.

In essence, this base model is simply counting frequencies of form–meaning mappings. When it encounters a cue (meaning) followed by an outcome (form), it simply increases the weight of the link between them by a constant number, which we set to 1 in order to emphasize the model's nature as a simple frequency counter. The results do not change depending on what the number is.

In order to connect the model's knowledge to the experimental results, we need a set of linking hypotheses connecting the weights and activations of the model to the participants' responses in the experimental tasks. The linking hypothesis for production is

straightforward. Production involves activating forms given the semantic features present on that test trial and a context cue. The activation of a form is simply the sum of connection strengths from the semantic and context cues present on the test trial to that form. The choice of the form is then determined stochastically (Luce, 1963): the form is chosen in proportion to its activation value relative to the sum of all forms' activation values given the cues present.

The linking hypothesis for comprehension is more controversial. Note that the model, like the subjects, was trained only in the meaning-to-form direction. However, comprehension requires the listener to choose a meaning given a form, reversing the cue–outcome mappings they were trained on. The model must be able to do the same. We propose that the associations participants learn obey the Symmetry Principle: a cue–outcome association is as strong as the corresponding outcome–cue association (Asch & Ebenholtz, 1962; Kahana, 2002; Miller & Matzel, 1988). The symmetry allows us to use the mappings learned during exposure to training trials in which a meaning precedes the corresponding form to be used during comprehension, where meanings need to be activated given forms.

Here, a choice between two meanings depends on the difference in activations between the two meanings' contrasting features. For example, the probability of clicking on DIM.PL rather than BIG.PL when presented with -dan_{PL} is proportional to the difference in association strengths between -dan_{PL} ~ DIM (or DIM ~ -dan_{PL}) and -dan_{PL} ~ BIG (or BIG ~ -dan_{PL}). The bigger this difference, the more likely participants are to click on the meaning that was in fact paired with the form cue in training, i.e., its original meaning.

Besides the connections between the cues and outcomes present on a particular trial, there are three other sets of connections that could potentially be updated. Alternative theories of associative learning differ in their claims about whether these connections are indeed updated.

First, there are connections from the cues present on a trial to the outcomes absent from that trial. It is usually thought that these connections' weights are reduced, so that cues that are consistently paired with the *absence* of a certain outcome develop inhibitory connections to that outcome, with the subject learning the negative contingency present in the environment. In the simulations below, we reduce the weights of connections between present cues and absent outcomes by 0.5 on each trial. This reflects the generally lower salience of absent outcomes relative to present ones, which is inferred from slower learning of inhibitory associations compared to excitatory associations (Tassoni, 1995; van Hamme & Wasserman, 1994). The results below hold as long as absent outcomes are no more salient than present ones.

Second, there are connections from the absent cues to the present outcomes. These connections are assumed to *not* be updated by Rescorla & Wagner (1972). However, van Hamme and Wasserman (1994) as well as Tassoni (1995) argued that—if participants know the set of cues that *could* occur on every trial—the absence of a cue can be salient. In other words, learners may notice the consistent absence of a cue on trials containing a certain outcome and develop a negative association between that cue and the outcome. Once again, absences are less salient than presences, which means that even if connections from absent cues to present outcomes are learned, they are learned at half the speed of connections involving present cues and outcomes.

Finally, one could argue that connections from absent cues to absent outcomes may also be updated, gaining strength: when a cue and an outcome are *absent together,* the learner is in a position to learn that absence of the cue predicts absence of the outcome (Tassoni, 1995). Because these connections involve two absent stimuli, they are updated at 1/4 the speed of connections between present cues and outcomes.

Thus, models of learning can be arranged from simplest (wiring together present cues and outcomes only) and least veridical—least able to faithfully reproduce environmental contingencies—to most complex and most veridical (updating all connections on every trial). In what follows, we examine what kinds of updating are needed to capture the experimental results by independently varying whether each distinct set of connections undergoes updating. Table 3.1 summarizes the possible models from a simple frequency counter that updates only the connections between present cues and present outcomes to a fully veridical contingency tracker that updates all four sets of connections (in the normative direction). I will refer to the models we evaluate with the abbreviations shown on the left sides of the table cells. For example, the Rescorla–Wagner model updates only the sets of connections in the top row and can therefore be abbreviated as (p_c).

Table 3.1. The four distinct sets of cue-outcome connections.

	Outcome Present		Outcome Absent	
Cue Present	($p_{cue}p_{out}$)	+	($p_{cue}a_{out}$)	−
Cue Absent	($a_{cue}p_{out}$)	−	($a_{cue}a_{out}$)	+

Note. The signs show whether connection weights should become more positive (+) or more negative (−) in a model that is able to capture environmental contingencies veridically.

Table 3.2 shows predicted activations of the frequent suffix, its synonym, and the two other suffixes (which are always activated equally) by the semantic features of the novel meaning (DIM & PL) under all logically possible models of associative learning. The $p_{cue}p_{out}$ column represents the simplest possible model, a frequency counter (Bybee, 2010). Columns $p_{cue}a_{out}$, $a_{cue}p_{out}$, $a_{cue}a_{out}$ represent association sets that can be added to the frequency counter in order to make contingency learning more veridical, incorporating learning of connections involving absent cues and/or outcomes. Column p_{cue} is the model that learns only from present cues (a Hebbian version of Rescorla & Wagner, 1972). The last column is the full model that learns even about associations between absent cues and absent outcomes (a Hebbian version of Tassoni, 1995).

Table 3.2. Activation of the four suffixes given the DIM.PL meaning under alternative models.

Suffix	$p_{cue}p_{out}$	$p_{cue}a_{out}$	$a_{cue}p_{out}$	$a_{cue}a_{out}$	p_{cue}	all
Frequent	72	−42	−18	15	30	24
Synonym	24	−66	−6	21	−42	−12
Other	24	−66	−6	21	−42	−12

Note. Synonym denotes the synonym of the frequent suffixes. Other suffixes are the infrequent suffixes that are not the synonym of the frequent suffixes.

3.2.2 Extension in production

Extension of frequent forms to novel meanings is predicted if the activation of the frequent form exceeds that of all other forms, including the frequent form's synonym. In other words, a preference to extend the frequent form to novel meanings is predicted whenever the largest number in Table 3.2 is in the top row.

As seen in Table 3.2, extension of the frequent form is predicted by increasing the weights of connections from present cues to present outcomes, as well as by decreasing

the weights of connections from present cues to absent outcomes. Updating connections from absent cues acts against extension.

For the simulations reported in Table 3.2, it was assumed that an absence of a cue or outcome is noticed only half the time while its presence is always noticed. Associative learning in conditioning paradigms tends to be slower when reinforcement is signaled by the absence of a cue than when it is signaled by the presence of a cue (e.g. Wasserman, Dorner, & Kao, 1990). However, one might question whether absences are missed or ignored that often, and wonder whether noticing absences more would eliminate extension. It turns out not to matter much: a_{cue} does not overpower p_{cue} even if absences are as salient as presences. All extant models of learning agree that absent stimuli are no more salient than stimuli that are presented and therefore all predict extension of frequent forms to related meanings in production.

3.2.3 Entrenchment in comprehension

Table 3.3 reports activation differences between features that distinguish the novel meaning from the original meaning. Because of the Symmetry Principle, the activation differences correspond to meaning-to-form connection weights involving the semantic features in question. For example, the activation difference between the non-diminutive and diminutive plural for -dan_{PL} is the weight of the connection between -dan_{PL} and BIG minus the weight of the connection between -dan_{PL} and DIM. The activation difference between the singular and plural diminutive for -nem_{DIM} is the weight of the connection between -nem_{DIM} and SG minus the weight of the connection between -nem_{DIM} and PL (cf. Miller & Matzel, 1988). Entrenchment is observed if this difference is

larger (more positive) for a frequent form compared to the other forms, i.e. if the value in the top row in Table 3.3 is larger than the value in the bottom row.

Table 3.3 shows that entrenchment is favored by strengthening connections between present cues and present outcomes ($p_{cue}p_{out}$), weakening connections between absent cues and present outcomes ($a_{cue}p_{out}$), and strengthening connections between absent cues and outcomes ($a_{cue}a_{out}$). Because updating $p_{cue}p_{out}$ and $p_{cue}a_{out}$ weights pull in different directions, entrenchment only occurs if absent outcomes are less salient than present outcomes. In other words, weights of connections to absent outcomes must change less than the weights of connections to present outcomes. This appears to be a reasonable assumption (e.g. Tassoni, 1995), though not all extant models make it. For example, the Naïve Discriminative Learner (Baayen et al., 2011), which uses equilibrium equations for the Rescorla & Wagner (1972) model from Danks (2003), does not show entrenchment because the learning rates for present and absent outcomes in Danks' equations are equal. However, this lack of difference in salience between present and absent outcomes is merely a simplifying assumption (Danks, 2003, pp. 115–116).

Table 3.3. Activation differences between the original and novel meanings.

Suffix	$p_{cue}p_{out}$	$p_{cue}a_{out}$	$a_{cue}p_{out}$	$a_{cue}a_{out}$	p_{cue}	all
Frequent	36	0	36	6	36	78
Synonym	12	−12	12	−6	0	6
Other	12	12	0	0	24	24

Note. Each cell contains activation difference between the meaning paired with a form in training and the novel, diminutive plural, meaning. Activations of shared features of the competing meanings cancel out. Therefore, for plural suffixes this is the difference in activations between BIG and DIM, and for diminutive suffixes it is the difference between SG vs. PL. Entrenchment is predicted if Frequent > Other.

3.3 Conclusion

A production-comprehension dissociation therefore falls out of simple, Hebbian associative learning models, which acquire symmetrical form–meaning associations based on cue-outcome co-occurrence (Hebb, 1949; Miller & Matzel, 1988; see also McMurray, Horst, & Samuelson, 2012; Yu & Smith, 2012). Remarkably, all that is required to obtain the divergence between frequency effects in production and comprehension—entrenchment of the frequent in comprehension, and extension of the frequent in production—is the highly uncontroversial assumption that cue and outcome absences are less salient than present cues and outcomes (Tassoni, 1995; Wasserman et al., 1990), an assumption that is also normatively justified: almost every stimulus is absent more often than it is present, hence the presence of a stimulus is typically more informative about the contingencies in the learner's environment than its absence (McKenzie & Mikkelsen, 2007). Despite being somewhat surprising to human theorists, frequency-driven semantic extension in production is predicted to co-exist with entrenchment in comprehension by basic associative learning theory.

CHAPTER IV

ENTRENCHMENT CONSTRAINS EXTENSION

4.1 Experiment II[7]

Whereas the production test preceded the form choice and comprehension tasks in Experiment I, the comprehension task was the first test in Experiment II, followed by production, and concluding with the form choice task. The main purpose of this experiment was to ensure that the comprehension results in Experiment I were not due to carry-over effects from preceding test tasks.

Experiments II and III are also intended to determine what carry-over effects are present between the test tasks and, specifically, whether the effects of frequency in production and comprehension can be influenced by learning during test. In Experiment I, we observed a dissociation between the production test and the comprehension test: the forms chosen to express DIM.PL in production were seldom mapped onto DIM.PL in comprehension. This suggests that the form–meaning mappings observed during the production test were not reused in a subsequent comprehension test, indicating incomplete transfer of knowledge (in this case, of a system of form–meaning association weights) from production to comprehension. In Experiment II, the comprehension test precedes production, affording an opportunity to observe transfer in the opposite direction, from comprehension to production. As discussed in Chapter II, comprehension-to-production transfer of form–meaning mappings is thought to be crucial to eliminating accessibility-driven overextension in child language (Brooks & Braine, 1995; Regier &

[7] The experiments reported in this chapter have previously been published in Harmon, Z., & Kapatsinski, V. (2017). Putting old tools to novel uses: The role of form accessibility in semantic extension. *Cognitive Psychology*, *98*, 22–44.

Gahl, 2004). If this kind of comprehension-to-production transfer occurs, we expect that extension of frequent forms in production observed in Experiment I may be replaced by entrenchment in Experiment II, where production follows comprehension.

4.1.1 Methods

Sixty-six adult native speakers of American English participated in this experiment, with half assigned to each language (Dan vs. Nem). Each participant took part in only one experiment. Learners were tested using the same procedure utilized in Experiment I.

In addition to reordering the tasks, we removed the twenty trials involving the compositional suffix *-dannem* from the form choice task and the ten trials featuring *-dannem* from the form-to-meaning mapping task. The presence of *-dannem* could not have influenced production results in Experiment I because the production test preceded the tests that presented participants with *-dannem*. However, we were concerned that the presence of this compositional option could have discouraged participants from extending the other suffixes to DIM.PL in the form choice and form-to-meaning mapping tasks. Therefore, a secondary aim of Experiments II and III is to ensure that the results of form choice and form-to-meaning tasks observed in Experiment I cannot be attributed to the presence of a compositional option in these tasks.

4.1.2 Results

As in Experiment I, we discuss tasks in the order that participants encountered them. The order of tasks is the primary difference between Experiment II and Experiment I. The order of task presentation in the figures is kept consistent in all experiments for easy comparison of results within tasks across experiments.

4.1.2.1 Comprehension

As in Experiment I, a form was mapped onto the meaning it was paired with during exposure (original rather than the novel meaning) significantly more often when it was the frequent form ($\beta = 2.9949$, $z = 3.063$, χ^2 (1) = 12.80, $p = .0003$). These results are presented in the Comprehension panel of Figure 4.1 below.

To assess whether the effect of form frequency in the form-to-meaning mapping task differed between the two experiments, we tested for an experiment-by-language interaction effect on the likelihood of choosing the original vs. the novel DIM.PL meaning. This interaction was not significant, χ^2 (1) = 0.65, $p = .72$, indicating that the entrenching effect of frequency in comprehension did not differ across experiments (i.e. depending on whether the comprehension task preceded or followed production).

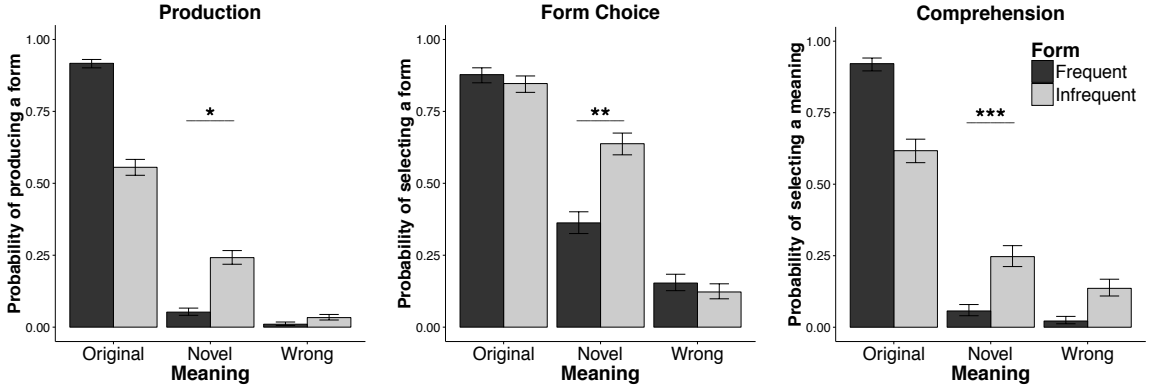

Figure 4.1. The results of Experiment II for the two suffixes -*dan*$_{PL}$ and -*nem*$_{DIM}$. In this experiment, comprehension preceded production, which preceded form choice. Original meaning = meaning paired with a suffix during training, e.g. BIG.PL for -*dan*$_{PL}$; Novel meaning = DIM.PL, which shares a feature with each of the original meanings; Wrong meaning = meaning paired with a different suffix in training, e.g. DIM.SG for -*dan*$_{PL}$. Frequent suffix = -*dan*$_{PL}$ in Dan language and -*nem*$_{DIM}$ in Nem language. Infrequent suffix = -*dan*$_{PL}$ in Nem and -*nem*$_{DIM}$ in Dan.

4.1.2.2 Production

The analysis was conducted on 80.91% of the data where participants' responses matched one of the four suffixes in Dan and Nem. As in Experiment I, participants were close to the probabilities of the suffixes in training given the original meanings. Unlike in Experiment I, *-dan*$_{\text{PL}}$ and *-nem*$_{\text{DIM}}$ were significantly *less* likely to be chosen for expressing the novel meaning of DIM.PL when they were frequently encountered during exposure ($\beta = -8.509$, $z = -2.633$, $\chi^2\,(1) = 5.43$, $p = .02$). These results are consistent with the entrenching effects of frequency in the comprehension task observed in both experiments.

Reliability of the difference between the two experiments was assessed by testing an interaction between experiment and language for predicting choice of form to express the novel DIM.PL meaning. The interaction was highly significant, $\chi^2\,(1) = 22.9$, $p < .0001$, confirming that the effect of frequency (language) differs across the two experiments with high form frequency significantly favoring extension in Experiment I and significantly disfavoring it in Experiment II. These results are presented in the Production panel of Figure 4.1.

4.1.2.3 Form choice

The analysis was conducted on 97.66% of the data where participants provided responses. Unlike in Experiment I, the effects of the form choice task paralleled the effects of the production and comprehension tasks in this experiment (Figure 4, Form Choice panel). There was a significant effect of language on suffix choice. As in production, forms were less likely to be selected to express the novel DIM.PL meaning when they were frequent in training ($\beta = -2.068$, $z = -3.32$, $\chi^2\,(1) = 10.5$, $p = .0012$).

Comparable to the production results, there was a significant interaction between experiment and language in the form choice data, $\chi^2(1) = 8.63$, $p = .013$, indicating that the effect of frequency on form choice in Experiment II differs reliably from the (null) effect of frequency on form choice in Experiment I. As in production, the direction of the effect of frequency is consistent with entrenchment, observed in comprehension in both experiments.

4.1.3 Discussion

The three tasks were re-ordered in Experiment II so that the test phase started with comprehension, which was followed by production, and finally form choice. One purpose of this re-ordering was to show that the entrenchment pattern observed in the comprehension task in Experiment I did not depend on the comprehension task occurring late in the testing sequence. Experiment II replicated the comprehension results of Experiment I, confirming that frequency leads to entrenchment in comprehension. It therefore appears that these results do not depend on where the comprehension task appears in the test sequence. The presence of entrenchment in the comprehension task in the absence of trials featuring the compositional form *-dannem* in Experiment II also confirms that the entrenchment effect does not crucially depend on their presence. These results therefore provide additional support for the dissociation between production and comprehension observed in Experiment I: a comprehension task that immediately follows the exposure phase shows entrenchment of frequent forms to their original meanings, whereas a production task that immediately follows exposure shows extension of frequent forms to a novel meaning.

Another purpose of reordering the test tasks was to test the hypothesis that entrenchment in comprehension results in the development of strong form–meaning mappings that can then guide production (Braine & Brooks, 1995; Regier & Gahl, 2004). This hypothesis was confirmed by the finding that the entrenchment effect was observed in both production and form choice tasks when they followed the comprehension task. It appears that exposing participants to the form-to-meaning task has led the participants to use the frequent forms in their original meanings only and to use the semantic competitors of these forms to express DIM.PL. Figure 4.2 presents the individual subject data in the production and comprehension tasks of Experiment II. The patterns of mappings in this experiment are very similar across the two tasks as subjects have a tendency to avoid mapping a frequent suffix to DIM.PL in both. Comparing across the two languages, -*nem*$_{\text{DIM}}$ is more likely to be mapped onto DIM.PL in Dan and -*dan*$_{\text{PL}}$ is more likely to be mapped onto DIM.PL in Nem.

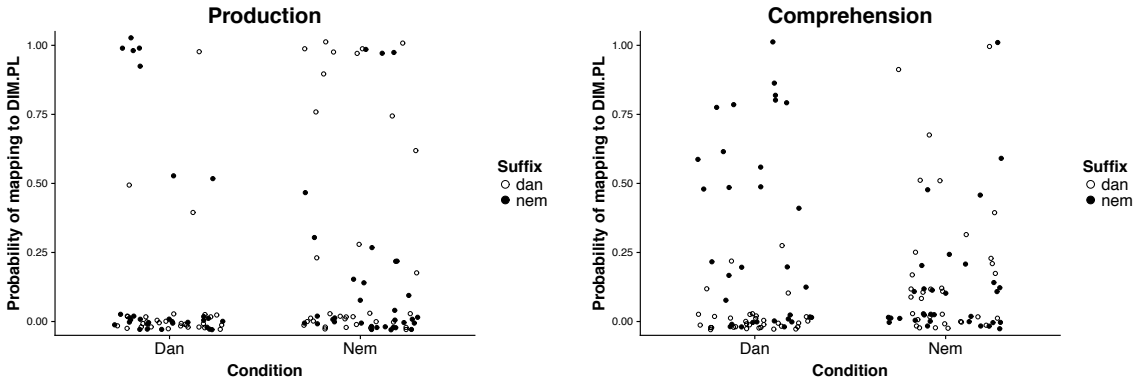

Figure 4.2. Individual participants' mappings of -*dan*$_{\text{PL}}$ and -*nem*$_{\text{DIM}}$ to DIM.PL meaning in the production and comprehension tasks in Experiment II.
Each white dot corresponds to an individual subject's use of -*dan*$_{\text{PL}}$ and each black dot corresponds to an individual subject's use of -*nem*$_{\text{DIM}}$. In production, the dots represent probabilities of choosing a suffix in response to DIM.PL. In comprehension, the dots represent probabilities of choosing DIM.PL in response to a suffix.

In order to further evaluate the effects of the comprehension task on the participants' system of form–meaning mappings, we examined the response patterns of individual participants in the two experiments. We suspected that completing the comprehension task first would cause participants to settle on a unique form for each meaning. If these mappings are reused in production, a participant in Experiment II should be unlikely to use more than one form to express a particular meaning—even a novel meaning—whereas participants in Experiment I should be more likely to do so.

To investigate this issue further, we assessed individual participants' form–meaning mapping systems by asking whether each participant extended the suffix they used most for either DIM or PL to express DIM.PL. For example, if a subject uses -dan_{PL} more than any other suffix for PL and DIM.PL, that individual would be coded as an extender, extending -dan_{PL} to all plurals, whether diminutive or not. If a subject uses -nem_{DIM} for DIM and DIM.PL, that individual would also be coded as an extender, using -nem_{DIM} for all diminutives. To be a non-extender, a participant would have to demonstrate a mutually-exclusive system, preferring not to use either their dominant DIM suffix or their dominant PL suffix to express DIM.PL, e.g., favoring -dan_{PL} for PL, -$shoon_{\mathrm{DIM}}$ for DIM but -nem_{DIM} or -sil_{PL} for DIM.PL.

Of the 70 subjects in Experiment I, only 21 (30%) used the suffixes in a mutually exclusive way. In other words, most participants in Experiment I extended a frequently used suffix to the novel meaning. In contrast, 52 out of 66 participants in Experiment II (79%) showed a tendency towards mutual exclusivity. The proportion of extenders differed significantly across the two experiments, $\chi^2(1) = 30$, $p < .0001$. Participants in

Experiment II tended not to use the same form to express two meanings (or, equivalently, a general, one-feature meaning), whereas participants in Experiment I did.

Participants in Experiment I also tended not to use the preferred form for DIM.PL in comprehension to express DIM.PL in production. For 28% of the participants, the form used most to express DIM.PL in production was the form mapped most onto DIM.PL in comprehension. In contrast, one form ~ BIG.PL mapping was dominant in both tasks for 59% of participants in Experiment II (a significant difference; $\chi^2(1) = 6.2, p = .01$), suggesting transfer from comprehension to production.

By the time the subjects get to choose forms to express meanings, their system of form–meaning mappings has settled enough for form choice to be no longer a choice at all: the novel meaning now has a form associated with it (usually, the synonym of the frequent form). Participants in Experiment I tended not to settle on a single form for DIM.PL in production: only 28% of participants used a single form to express DIM.PL 90% of the time or more. In contrast, most (57%) of participants in Experiment II consistently used a single form, a significant difference, $\chi^2(1) = 9.8, p = .002$. Overall, the results indicate that participants settled on a system of one-to-one form–meaning mappings during the comprehension task, which they then reused in production. In contrast, the production task did not result in settling on mutually exclusive form–meaning mappings that could then be reused in comprehension: most participants used more than one suffix to refer to DIM.PL in the production task in Experiment I.

4.2 Experiment III

In Experiment III, the form choice task precedes the comprehension task, which precedes the production task. The primary aim of this experiment is to determine whether

the lack of word frequency effects seen in the form choice task in Experiment I is still observed when the task occupies the same location in the task sequence as the production task in Experiment I.

We did not have strong predictions for the effects the form choice task might have on subsequent tasks prior to conducting Experiment I. Like production, the form choice task involves selecting a form to produce a meaning. Furthermore, the form choice task discourages settling on a particular form for each meaning. Given binary choices between four suffixes, participants *must* choose multiple forms when cued with the novel meaning over the course of the task. For example, on a -*nem*$_{\text{DIM}}$ vs. -*dan*$_{\text{PL}}$ trial one might choose -*nem*$_{\text{DIM}}$ to express DIM.PL while on a -*shoon*$_{\text{DIM}}$ vs. -*sil*$_{\text{PL}}$ trial one might choose -*sil*$_{\text{PL}}$ for the same meaning. (This is not the case for meanings seen during exposure because there is a *right* answer on every trial for these meanings.) Thus, the form choice task may leave even more room for the form–meaning mappings to shift during the comprehension task and then be reused in production. On the other hand, the mappings chosen during the form choice task may also influence form choice in production, competing with the mappings formed in comprehension, and therefore weakening the alignment between comprehension and production observed in Experiment II.

4.2.1 Methods

Sixty new adult native speakers of American English participated in the experiment, with thirty assigned to each language. Apart from a change in order compared to Experiment I, we revised the form choice task trial procedure to include a brief presentation of the four alternative meanings prior to form choice. This was similar to the presentation of meanings in the production test in Experiment I and aimed to

increase awareness of the difference between the four meanings. As in the production test, the four pictures appeared for 200 milliseconds and then disappeared, leaving the participants with the one target picture. The comparison between the form choice task results in Experiments I and III allows us to evaluate whether this change in procedure has any effect on the pattern of results in the task. The stimuli were identical to that of Experiment II.

4.2.2 Results

4.2.2.1 Form choice

The analysis was conducted on 98.61% of the data, where participants provided responses. The results are presented in the Form Choice panel of Figure 4.3. As in Experiment I, we found no significant effect of frequency on form choice ($\beta = 0.251$, $z = 0.513$, $\chi^2\ (1) = 0.26$, $p = .61$; $p(H_0|D)_{BIC} = 96\%$). The results of the form choice task are therefore robust to whether this task is preceded by the production task, whether the compositional form *-dannem* is included as an option to express DIM.PL and whether pictures representing the four possible meanings are briefly shown between presentation of the base and the form choice.

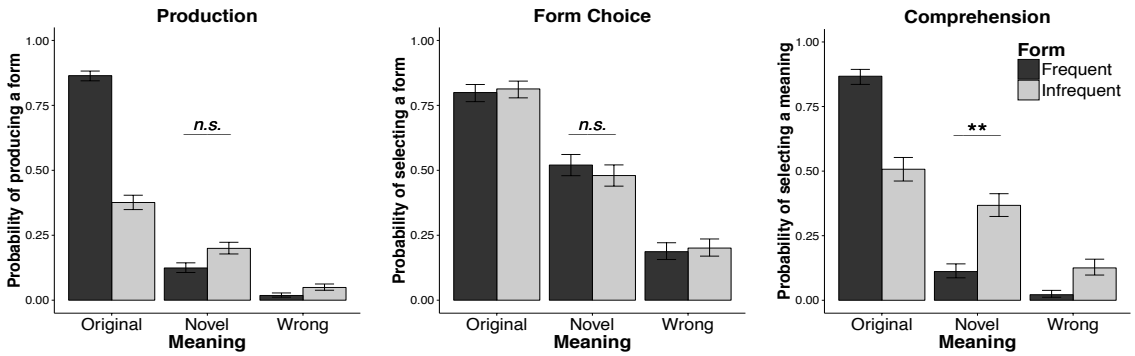

Figure 4.3. The results of Experiment III for the two suffixes -*dan*$_{PL}$ and -*nem*$_{DIM}$. In this experiment, the form choice task preceded comprehension, which preceded production. Original meaning = meaning paired with a suffix during training, e.g. BIG.PL for -*dan*$_{PL}$; Novel meaning = DIM.PL, which shares a feature with each of the original meanings; Wrong meaning = meaning paired with a different suffix in training, e.g. DIM.SG for -*dan*$_{PL}$. Frequent suffix = -*dan*$_{PL}$ in Dan language and -*nem*$_{DIM}$ in Nem language. Infrequent suffix = -*dan*$_{PL}$ in Nem and -*nem*$_{DIM}$ in Dan.

4.2.2.2 Comprehension

The analysis was conducted on 99.57% of the data where participants provided responses. The results are presented in the Comprehension panel of Figure 4.3. As in Experiments I and II, forms were mapped onto the meanings they were paired with during exposure rather than the novel meaning of DIM.PL significantly more often when they were frequent during exposure (an entrenchment effect; $\beta = 3.391$, $z = 3.46$; χ^2 (1) = 8.47, $p = .0036$).

4.2.2.3 Production

The analysis was conducted on 84.33% of the data where participants' responses matched one of the four suffixes in Dan and Nem. The results are presented in the Production panel of Figure 4.3. As in Experiments I and II, form probabilities given the original meanings were close to those experienced in training. However, unlike in previous experiments, there was no significant effect of frequency on the choice of form

to express the novel meaning of DIM.PL ($\beta = 0.732$, $z = 0.23$, χ^2 (1) = 0.05, $p = .82$; providing strong evidence for the null; $p(H_0|D)_{BIC} = 95.3\%$). Thus, we observed neither the significant entrenchment effect seen in Experiment II nor the extension of frequent forms observed in Experiment I. A significant interaction between frequency and experiment reveals that the effect of form frequency on choosing a form to express a novel meaning is affected by the test tasks that precede the production test (χ^2 (2) = 22.9, $p < .0001$). However, there is no significant difference between Experiments II and III, in which comprehension precedes production ($\beta = -13.82$, $z = -2.42$, χ^2 (1) = 3.57, $p = .059$; $p(H_0|D)_{BIC} = 82\%$).

4.2.3 Discussion

The three tasks in Experiment III were re-ordered so that the test phase started with the form choice task, which was followed by the form-to-meaning mapping task and then production. The results of the form choice task in Experiment III were the same as in Experiment I: form frequency did not significantly influence form choice in this task. Comparing the initial production task from Experiment I to the initial form choice task from Experiment III between participants—which controls for the possibility of carry-over effects—reveals that the favorable effect of frequency on semantic extension observed in an initial production task (Experiment I) reliably differs from the lack of a frequency effect observed in an initial form choice task (Experiment III; $\beta = -12.652$, $z = -7.20$, χ^2 (1) = 114, $p < .0001$). This result bolsters our confidence in the conclusion that frequent forms are extended to novel meanings in production because they are more accessible than infrequent forms: when accessibility differences between frequent and

infrequent forms are leveled—as they are in the form choice task—high frequency no longer favors semantic extension.

As in Experiments I and II, we found an entrenchment effect in the form-to-meaning mapping task, whereby the suffixes -*dan*$_{\text{PL}}$ and -*nem*$_{\text{DIM}}$ were less likely to be mapped onto DIM.PL when they were frequently encountered in another related meaning. As in Experiment II, participants finish the comprehension task in possession of entrenched one-to-one form–meaning mappings that are then reused in production: 78% of the participants used a mutually exclusive system in production (vs. 30% in Experiment I and 79% in Experiment II). There was a significant difference in subjects' generalization patterns between Experiment I and Experiment III, $\chi^2\,(1) = 25$, $p < .0001$, but no difference between Experiment III and Experiment II, $\chi^2\,(1) = 0.031$, $p = .9$. By the production task, most participants in both Experiment II and Experiment III have settled into a system of one-to-one form–meaning mappings, eliminating opportunities for semantic extension of frequent forms. Figure 4.4 presents the individual subjects preference for mapping -*dan* and -*nem* onto DIM.PL in the production and comprehension tasks of Experiment III. The preference for the non-frequent form in both languages for DIM.PL is parallel in both tasks.

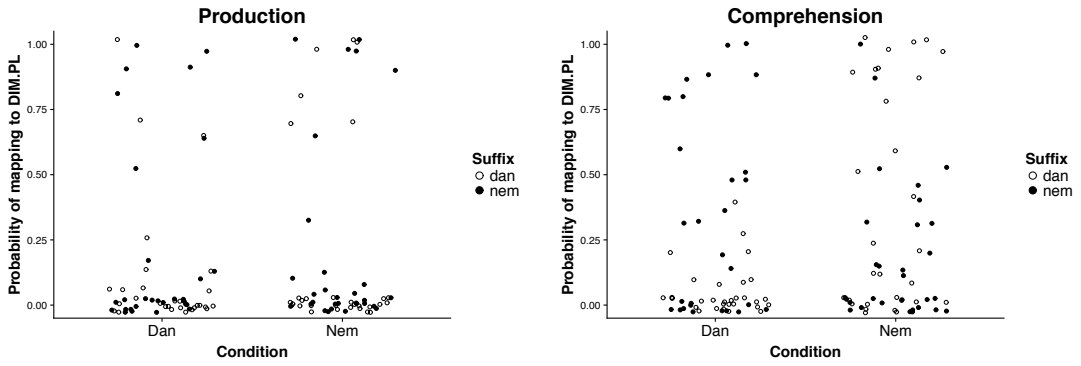

Figure 4.4. Individual participants' mappings of -*dan*$_{PL}$ and -*nem*$_{DIM}$ to DIM.PL meaning in the production and comprehension tasks in Experiment III.
Each white dot corresponds to an individual subject's use of -*dan*$_{PL}$ and each black dot corresponds to an individual subject's use of -*nem*$_{DIM}$. In production, the dots represent probabilities of choosing a suffix in response to DIM.PL. In comprehension, the dots represent probabilities of choosing DIM.PL in response to a suffix.

We examined the development of entrenchment over the course of the comprehension task by combining the comprehension results across experiments and examining the likelihood of mapping a form onto its original meaning vs. the novel (DIM.PL) meaning as a function of the form's frequency during exposure (the two lines in Figure 4.5) and the number of times one has experienced the form during the comprehension test (the horizontal axis in Figure 4.5). There was a significant effect of frequency during exposure ($\beta = 2.73$, $z = 17.92$, $p < .001$) and a significant interaction between frequency of a form and the number of times one has made a decision on the basis of that form in the comprehension task ($\beta = 0.41$, $z = 2.74$, $p = .006$). As shown in Figure 4.5, frequent suffixes are more likely to be mapped onto the original meaning from the start of the comprehension test but this advantage increases over the course of the comprehension test as a participant entrenches in the chosen mapping.

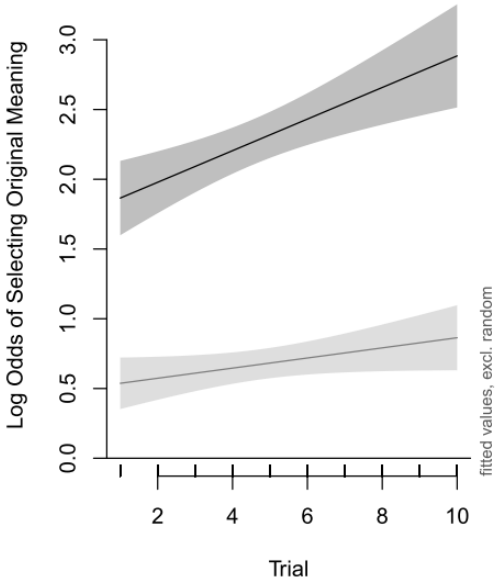

Figure 4.5. Entrenchment in comprehension in the three experiments
The effect of number of choices on choosing the original vs. novel meaning when a suffix is frequent (dark line) vs. infrequent (light line). As participants repeatedly map a suffix onto its original meaning in the comprehension test, they become more and more likely to map it onto that same meaning in future trials. This increase is faster for a form–meaning mapping that is frequently experienced in training.

The results of Experiments II and III suggest that form selection in production is rather easily influenced by learning during a preceding comprehension task. Form–meaning mappings appear to crystalize over the course of the comprehension task, and these calcified mappings are then reused in production.

Table 4.1 summarizes the results of the three experiments reported so far (see also §6.3.3 for individual subject data and §7.3 for a discussion of these results). As can be seen, the differences between the tasks observed within subjects in Experiment I using the task order intended to minimize carry-over effects are replicated between subjects for sequence-initial tasks across experiments (black font). Furthermore, the qualitative pattern of results in the comprehension task (entrenchment) is unaffected by preceding

tasks. What happens in production stays in production. However, what happens in comprehension does not stay in comprehension. Form choice is heavily affected by settling on a mutually exclusive system of form–meaning mappings in comprehension. Whereas prior to this settling form choice is heavily affected by form accessibility, choice after settling is driven by the form–meaning mappings settled upon.

Table 4.1. Results of the three experiments

Experiment	Task		
I	**Production** Extension	**Form Choice** Neither	**Comprehension** Entrenchment
III	**Form Choice** Neither	**Comprehension** Entrenchment	**Production** Neither
II	**Comprehension** Entrenchment	**Production** Entrenchment	**Form Choice** Entrenchment

Note. A summary of the effect of frequency on form and meaning choice in the three experiments. Black font illustrates the dissociation between production, comprehension and form choice tasks observed when the form choice and production tasks do not follow the comprehension task. The order of initial tasks top to bottom parallels the left-to-right order of tasks in Experiment I to facilitate comparison.

CHAPTER V

COMPETITOR FREQUENCY EFFECTS

In this chapter, I test two hypotheses focused on the effect of frequency of a suffix on the use of its synonymous competitor. In other words, I look at what happens to a -*sil*$_{\text{PL}}$ or -*shoon*$_{\text{DIM}}$ when it competes with a frequent -*dan*$_{\text{PL}}$ or -*nem*$_{\text{DIM}}$ versus when it has the same frequency as its competitor. The first is the hypothesis that competitors of frequent forms are pre-empted by them from mapping onto the shared meaning and are pushed out into the related novel meaning. In other words, entrenchment of the frequent forms in their original meanings accompanied with competition between forms encourages a one-to-one system of form–meaning mappings. Frequent forms are the first to settle into a specific meaning, pushing their competitors out of this meaning, even when the two are perfectly synonymous. This is a kind of preemption as the frequent forms preempt their competitors from expressing the shared meaning. If there are opportunities for the preempted competitors to map onto other similar areas of semantic space, they have a chance to remain extant. Otherwise, they will not stay in the system and become extinct.

Second, I test the hypothesis that preemption of forms by their frequent competitors occurs in both production and comprehension, but the locus of niche finding is in comprehension. While in production a preempted form may be eliminated due to retrieval difficulty, in comprehension the form is provided, so the listener has to interpret it. Finally, this effect may not be present in the form choice task as this task levels accessibility and may result in randomly choosing suffixes that have been paired with a meaning.

Third, I explore the possibility that frequency of a meaning influences how the novel meaning of DIM.PL is construed. The prediction is that DIM.PL would be construed as DIM.**PL** in Dan where the frequency of the PL meaning is high. Conversely, it would be construed as **DIM**.PL in Nem where the frequency of the DIM meaning is high. As a result, participants will select plural suffixes to express DIM.PL in Dan and diminutive suffixes in Nem.

The results of suffix competition in all three experiments are discussed in this chapter. As before, the tasks are discussed in the order of presentation to the participants in each experiment. The figures, however, are all matched for easy comparison. I discuss the implications of the findings for language acquisition and language change at the end of the chapter.

5.1 Experiment 1

5.1.1 Production

The production results for the competitors of the frequent suffixes (-*sil*$_{\text{PL}}$ and -*shoon*$_{\text{DIM}}$) are presented in the Production panel of Figure 5.1. An increase in the frequency of a form causes the synonymous competitor of that form to not be produced in response to the shared meaning with which both suffixes were paired during training. A mixed-effects logistic regression model with maximal random effect structure testing the mapping of the competitors to their original meaning revealed that after exposure, the competitors of the frequent suffixes (-*sil*$_{\text{PL}}$ in Dan and -*shoon*$_{\text{DIM}}$ in Nem) were less likely to be mapped onto their original meaning ($\beta = 6.436$, $z = 2.708$, $p = 0.007$). However, the results revealed no significant effect of competitor frequency on the choice of suffix in expressing the novel meaning of DIM.PL ($\beta = -0.356$, $z = -0.123$, $p = .902$).

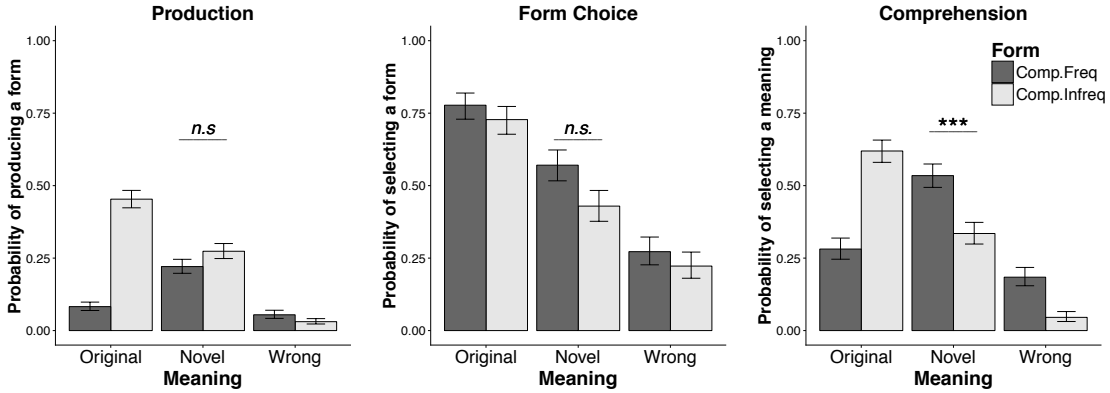

Figure 5.1. The results of Experiment I for the two suffixes -*sil*$_{PL}$ and -*shoon*$_{DIM}$. Original meaning = meaning paired with a suffix during training, e.g. BIG.PL for -*sil*$_{PL}$; Novel meaning = DIM.PL, which shares a feature with each of the original meanings; Wrong meaning = meaning paired with a different suffix in training, e.g. DIM.SG for -*sil*$_{PL}$. Competitor of the frequent suffix = -*sil*$_{PL}$ in Dan language and -*shoon*$_{DIM}$ in Nem language. Competitor of the infrequent suffix = -*sil*$_{PL}$ in Nem and -*shoon*$_{DIM}$ in Dan. Left panel: Form choice probabilities in the production task. The distance between the tops of the bars and 100% for the original meaning in the Production panel is composed largely of probability of choosing the synonymous suffix: -*nem*$_{DIM}$ over -*shoon*$_{DIM}$ or -*dan*$_{PL}$ over -*sil*$_{PL}$. Middle panel: Form choice probabilities in the forced form choice task on the -*sil*$_{PL}$ vs. -*shoon*$_{DIM}$ trials. Right panel: Meaning choice probabilities in the comprehension task. In this experiment, production preceded form choice, which preceded comprehension. (*** means $p < .001$).

5.1.2 Form choice

Results are presented in the Form Choice panel of Figure 5.1. Choice of the competitor suffixes was analyzed using a mixed-effects logistic regression model with language as the fixed effect of interest and the binary dependent variable of suffix (-*sil*$_{PL}$ or -*shoon*$_{DIM}$) chosen to express the novel meaning DIM.PL. Random intercepts for subjects and noun stems and random slopes for language within stem were included in the model. The analysis revealed no significant effect of language on suffix choice ($\beta = -0.823$, $z = -1.79$, $p = .073$). To ensure that the effect of frequency on the mapping of the

competitor suffixes to the novel meaning of DIM.PL in the same in both production and form choice task, a model tested the interaction between the two tasks. The model detected no significant difference between the two tasks ($\beta = -1.384$, $z = -0.473$, $p = .636$). A second model tested the effect of frequency on the mapping of the competitor suffixes to original meaning for each suffix. Competing with a frequent suffix did not influence how often the competitor suffixes were mapped less onto their original meanings ($\beta = -0.190$, $z = -1.253$, $p = .21$).

5.1.3 Comprehension

A mixed-effects logistic regression model (with random intercepts for subjects and random slopes for language within stem) probed the effect of competing with a frequent form on the mappings of -*sil*$_{\text{PL}}$ and -*shoon*$_{\text{DIM}}$, the competitors of -*dan*$_{\text{PL}}$ and -*nem*$_{\text{DIM}}$ (respectively). Once again, we focused on the likelihood of mapping the forms to the novel meaning (DIM.PL) vs. the meaning they were paired with during exposure (PL for -*sil*$_{\text{PL}}$ and DIM for -*shoon*$_{\text{DIM}}$) as the dependent variable (novel meaning vs. original meaning). Language served as the fixed-effect predictor, indicating whether the competitor of -*sil*$_{\text{PL}}$ or -*shoon*$_{\text{DIM}}$ was the frequent suffix (in Dan vs. Nem respectively). The combination of pre-emption with entrenchment predicted that a form should be more likely to be mapped onto a novel meaning when it competes with a frequent form for its original meaning, i.e. when it is a synonym of a frequent form.

The results are presented in the comprehension panel of Figure 5.1. As expected, these suffixes were more likely to be extended to the novel meaning of DIM.PL when their competitors were frequent ($\beta = -11.24$, $z = -6.7$, $p < .0001$).

5.1.4 Discussion

The comprehension data presented in Chapter III has shown that frequency causes forms to entrench in the specific meanings with which they co-occur. Thus, when -*dan*$_{PL}$ is infrequent, it bears the basic-level single-feature meaning PL and is mapped onto BIG.PL and DIM.PL equally often in comprehension. When -*dan*$_{PL}$ is frequent, it is instead mapped onto BIG.PL, and participants no longer select the DIM.PL meaning when presented with -*dan*$_{PL}$. In this chapter, we saw that, as the frequent forms entrench in their original meaning, their synonymous competitors are pushed out of that meaning. Thus, when -*dan*$_{PL}$ and -*sil*$_{PL}$ are equally infrequent, they are both mapped onto PL. However, -*sil*$_{PL}$ is pushed out of BIG.PL when -*dan*$_{PL}$ is frequent. As a result, participants tend to select DIM.PL when presented with -*sil*$_{PL}$ that competed with a frequent -*dan*$_{PL}$. This pattern is observed in production (see also Figure 5.2 presenting subject data on the mapping of these suffixes to DIM.PL in production and comprehension). The number of subjects who map -*sil*$_{PL}$ at least once onto DIM.PL increases from 15 in Nem language to 25 in Dan language as a result of an increase in frequency of -*dan*$_{PL}$, a significant difference, χ^2 (1) = 4, p = .04.

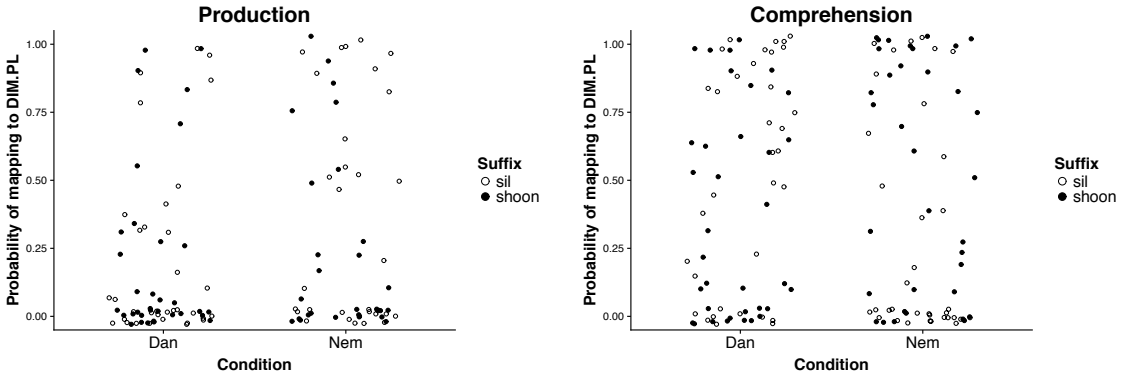

Figure 5.2. Individual participants' mappings of -*sil*$_{PL}$ and -*shoon*$_{DIM}$ to DIM.PL meaning in the production and comprehension tasks in Experiment I.
Each white dot corresponds to an individual subject's use of -*sil*$_{PL}$ and each black dot corresponds to an individual subject's use of -*shoon*$_{DIM}$. In production, the dots represent probabilities of choosing a suffix in response to DIM.PL. In comprehension, the dots represent probabilities of choosing DIM.PL in response to a suffix.

As a form–meaning pairing gains strength through frequency, it causes other forms to retreat from the shared area of the semantic space. This process appears to be crucially reliant on mutual exclusivity (Markman & Wachtel, 1988; Merriman & Bowman, 1989). A form is pushed out of its original meaning by a frequent competitor because the comprehender prefers to have only one form map onto a meaning. Interestingly, these effects of mutual exclusivity are seen in the comprehension, form-to-meaning mapping task—which has been the mainstay of work on mutual exclusivity—but not in production. One reason for the difference between the two tasks may be that in production learner face form retrieval difficulty, especially if production immediately follows exposure. Forms that are not accessible are much less likely to be used by speaker, a problem that is alleviated in comprehension when forms are provided. Overall, this result suggests that the semantic push chains observed in natural language change (Aronoff, 2016; MacWhinney, 1987, 1993; see also Lindsay & Aronoff, 2013) may likewise have their source in comprehension.

5.2 Experiment 2

5.2.1 Comprehension

The results of the comprehension task are presented in the Comprehension panel of Figure 5.3. As in Experiment I, suffixes were mapped onto the novel meaning of DIM.PL significantly more often when they competed with frequent synonyms (-*sil*$_{PL}$ in Dan and -*shoon*$_{DIM}$ in Nem) than when their synonyms were as infrequent as themselves (-*sil*$_{PL}$ in Nem and -*shoon*$_{DIM}$ in Dan; $\beta = -2.088$, $z = -3.504$, $p < .0005$).

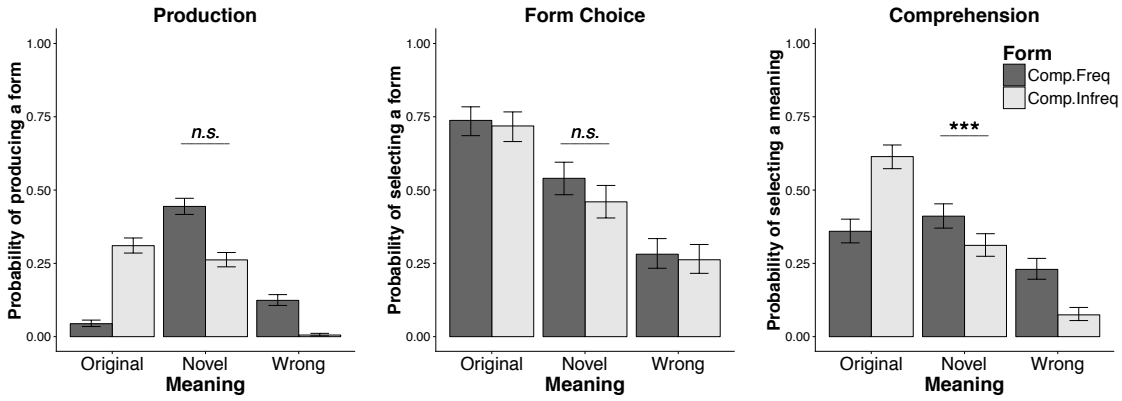

Figure 5.3. The probability of choosing Original, Novel, and Wrong meanings for -*sil*$_{PL}$ and -*shoon*$_{DIM}$, the competitors of the boosted suffixes -*dan*$_{PL}$ and -*nem*$_{DIM}$ in the three tasks in Experiment II.
In this experiment, comprehension preceded production, which preceded form choice. Original meaning is the meaning paired with a suffix during training, e.g. BIG.PL for -*sil*$_{PL}$; Novel meaning is DIM.PL; Wrong meaning is the meaning paired with a different suffix in training, e.g. DIM.SG for -*sil*$_{PL}$. Competitor of the frequent suffix = -*sil*$_{PL}$ in Dan and -*shoon*$_{DIM}$ in Nem. Competitor of the infrequent suffix = -*sil*$_{PL}$ in Nem and -*shoon*$_{DIM}$ in Dan.

5.2.2 Production

Competitor suffixes were pushed out of their original meanings significantly more often when they competed with a frequent suffix (-*sil*$_{PL}$ in Dan and -*shoon*$_{DIM}$ in Nem) than when they competed with an infrequent suffix (-*sil*$_{PL}$ in Nem and -*shoon*$_{DIM}$ in Dan; $\beta = 21.01$, $z = 3.037$, $p = .0024$). However, frequency had no effect on the mapping of the competitor suffixes to the novel meaning of DIM.PL ($\beta = -7.752$, $z = -1.955$, $p = .0505$). This unexpected null result (compare the difference between the dark and light gray bars in the Novel meaning in the production panel of Figure 5.3) is driven by the fact that the two competitor suffixes had very different behavior in the two languages. Looking at each suffix individually, it becomes clear that the plural suffix -*sil*$_{PL}$ was mapped onto the novel meaning of DIM.PL significantly more often when it did not compete with a frequent suffix (i.e., in the Nem language, $\beta = -20.46$, $z = -7.33$, $p < .0001$). However, language had no significant effect on the diminutive suffix -*shoon*$_{DIM}$ ($\beta = .2738$, $z = .176$, $p = .86$). In fact, a model testing the effect of frequency on the mapping of the competitors when they are collapsed does indicate a significant difference in the mapping of these suffixes onto DIM.PL in languages where they are frequent versus when they are infrequent. The results show that competitor suffixes are mapped more onto DIM.PL when they competed with a frequent suffix suffix (-*sil*$_{PL}$ in Dan and -*shoon*$_{DIM}$ in Nem) than when they competed with an infrequent suffix (-*sil*$_{PL}$ in Nem and -*shoon*$_{DIM}$ in Dan; $\beta = 9.982$, $z = 5.979$, $p < .0001$). Here, we will stick with the more conservative model but it is informative to know that at least for one suffix, entrenchment did occur.

The difference between the two suffixes was also observed for selecting suffixes to express their original meanings, where -*sil*$_{PL}$ was used to express PL more in the Nem

language ($\beta = 6.35$, $z = 3.56$, $p = .0004$) but language has no significant effect on selecting -*shoon*$_{\text{DIM}}$ to express DIM ($\beta = -0.998$, $z = 1.477$, $p = .5$).

5.2.3 Form Choice

As in Experiment I, there was no significant effect of competitor frequency on the choice of the competitor suffixes to express the Novel meaning of DIM.PL in the form choice task ($\beta = -0.523$, $z = -0.99$, $p = .321$). To check whether the effect of frequency on the mapping of competitor suffixes in this task was different from or the same as production, a model with an interaction between language and task was run. The results revealed no difference in the effect of frequency between the two tasks ($\beta = 1.297$, $z = 0.525$, $p = 0.6$). These results are presented in the Form Choice panel of Figure 5.3. Again, competing with a frequent suffix did not influence how often the competitor suffixes were mapped less onto their original meanings ($\beta = -$, $z = -$, $p =$).

5.2.4 Discussion

Similar results were found in Experiment II, where the comprehension task preceded production. The competitors of the frequent suffixes were pushed out of the original shared meaning and assigned to a novel meaning in comprehension. The same trend seemed to exist in production although not reliably. As such, we do not have enough information to generalize these findings to production. It is possible that the degree to which forms in a language were mutually exclusive in comprehension influences the degree to which these forms are mutually exclusive in production. Once again, the individual subject data in the production and comprehension tasks are presented (see Figure 5.4). These represent subjects' preferences for the two suffixes -

sil$_{PL}$ and -*shoon*$_{DIM}$ to the meaning of DIM.PL in the two languages. It is difficult to clearly detect a pattern of difference between the two tasks.

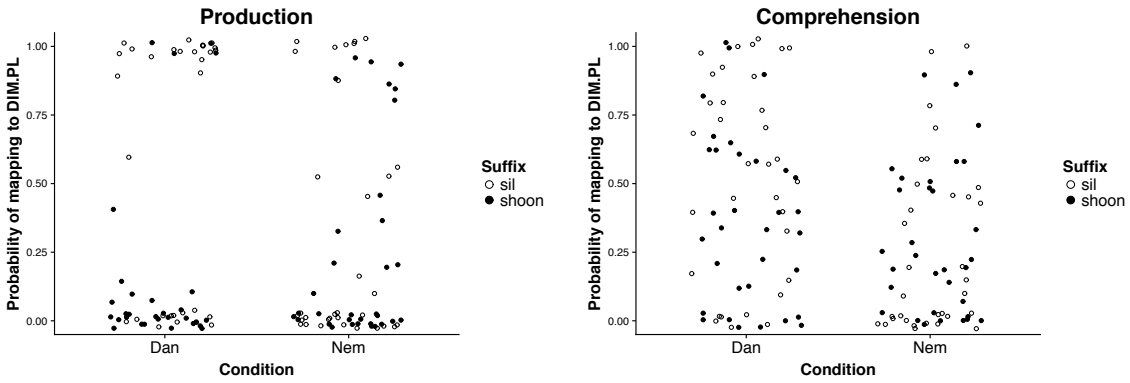

Figure 5.4. Individual participants' mappings of -*sil*$_{PL}$ and -*shoon*$_{DIM}$ to DIM.PL meaning in the production and comprehension tasks in Experiment II.
Each white dot corresponds to an individual subject's use of -*sil*$_{PL}$ and each black dot corresponds to an individual subject's use of -*shoon*$_{DIM}$. In production, the dots represent probabilities of choosing a suffix in response to DIM.PL. In comprehension, the dots represent probabilities of choosing DIM.PL in response to a suffix.

5.3 Experiment 3

5.3.1 Form Choice

The results are presented in the Form Choice panel of Figure 5.5. There was no effect of language on choosing -*sil*$_{PL}$ vs. -*shoon*$_{DIM}$, replicating the results of Experiments I and II, where no effect of form frequency was observed in this task ($\beta = -0.694$, $z = -1.443$, $p = .149$ for mapping onto novel meaning; $\beta = -0.694$, $z = -1.443$, $p = .149$ for mapping onto the original meaning).

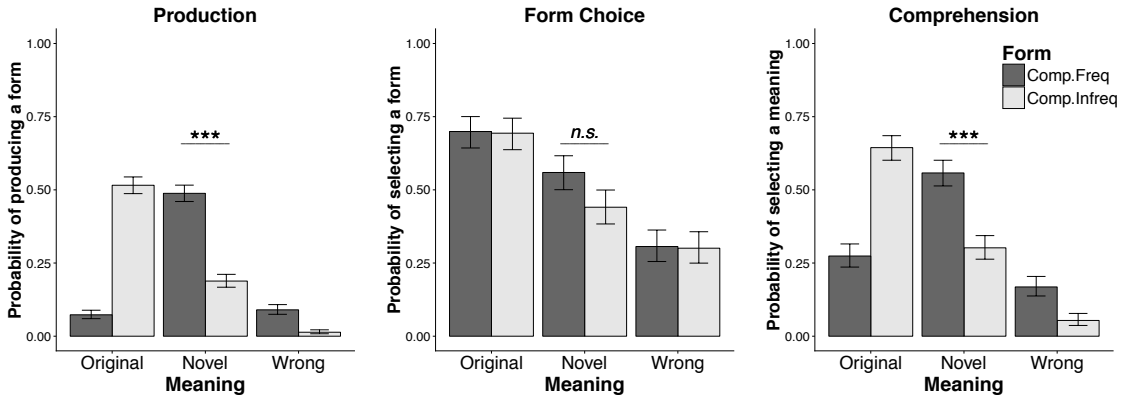

Figure 5.5. Probability of choosing original, Novel, and Wrong meanings for -*sil*$_{PL}$ and -*shoon*$_{DIM}$, the competitor suffixes in production, form choice, and comprehension tasks in Experiment III.
In this experiment, the form choice task preceded comprehension, which preceded production. Original meaning is the meaning paired with a suffix during training, e.g. BIG.PL for -*sil*$_{PL}$; Novel meaning is DIM.PL; Wrong meaning is the meaning paired with a different suffix in training, e.g. DIM.SG for -*sil*$_{PL}$. Competitor of the frequent suffix = -*sil*$_{PL}$ in Dan language and -*shoon*$_{DIM}$ in Nem language. Competitor of the infrequent suffix = -*sil*$_{PL}$ in Nem and -*shoon*$_{DIM}$ in Dan.

5.3.2 Comprehension

The results are presented in the Comprehension panel of Figure 5.5. As in Experiments I and II, suffixes were mapped onto the novel meaning significantly more often when competing with frequent suffixes for the original meaning (-*sil*$_{PL}$ in Dan and -*shoon*$_{DIM}$ in Nem vs. -*sil*$_{PL}$ in Nem and -*shoon*$_{DIM}$ in Dan; $\beta = -4.255$, $z = -4.88$, $p < .0001$).

5.3.3 Production

Results of the production task paralleled the results of the comprehension task (Figure 5.5. Production panel). Competitor suffixes were pushed out of their original meanings significantly more often when they competed with a frequent suffix (-*sil*$_{PL}$ in Dan and -*shoon*$_{DIM}$ in Nem) than when they competed with an infrequent suffix (-*sil*$_{PL}$ in

Nem and -*shoon*$_{\text{DIM}}$ in Dan; $\beta = 15.662$, $z = 3.574$, $p < .0004$). These suffixes were mapped onto the novel meaning significantly more often when they competed with frequent suffixes for their original meaning ($\beta = -21.09$, $z = -6.93$, $p < .0001$).

5.3.4 Discussion

When production follows comprehension, the competing suffixes settle into a mutually exclusive system. If participants learn a mutually exclusive system in comprehension, they can then use it in production. Figure 5.6 presents the individual subject data in the production and comprehension tasks in. These represent subjects' preferences for mapping the two suffixes -*sil*$_{\text{PL}}$ and -*shoon*$_{\text{DIM}}$ to the meaning of DIM.PL in the two languages. A higher number of participants mapped -*sil*$_{\text{PL}}$ onto DIM.PL in Dan in both production and comprehension and a higher number of participants mapped -*shoon*$_{\text{DIM}}$ onto DIM.PL in Nem in both tasks, in line with the niche-seeking behavior for these competing suffixes.

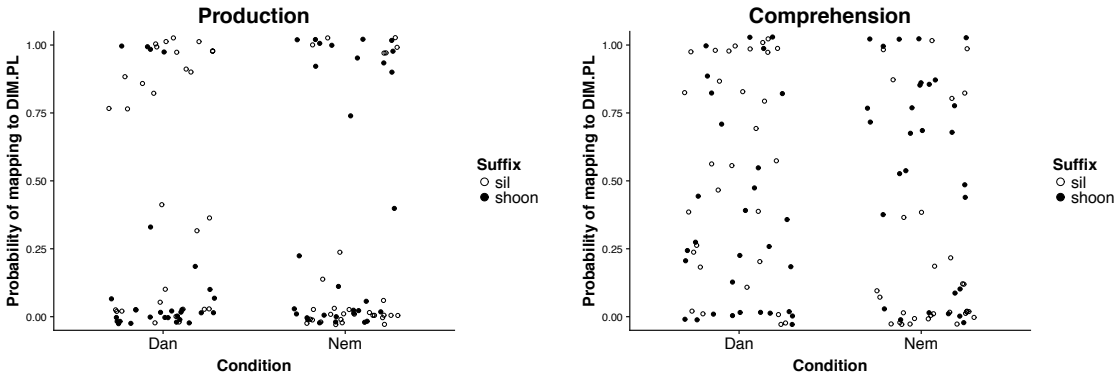

Figure 5.6. Individual participants' mappings of the competitor suffixes -*sil*$_{\text{PL}}$ and -*shoon*$_{\text{DIM}}$ to DIM.PL meaning in production and comprehension tasks in Experiment III. Each white dot corresponds to an individual subject's use of -*sil*$_{\text{PL}}$ and each black dot corresponds to an individual subject's use of -*shoon*$_{\text{DIM}}$. In production, the dots represent

probabilities of choosing a suffix in response to DIM.PL. In comprehension, the dots represent probabilities of choosing DIM.PL in response to a suffix.

5.4 Construal

To evaluate the possibility that the construal of the novel meaning is influenced by frequencies of its component features, I examined form choice in the two tasks where the direction of processing is from meaning to form in the three experiments reported so far. These are the production task and the forced form choice task. The prediction is that if DIM.PL is construed with PL as the more prominent component, participants will profile that component by choosing a plural suffix to express DIM.PL. Furthermore, the frequency of form should play minimal role in this process. More specifically, I expect that participants will choose -*sil*$_{\text{PL}}$ and -*dan*$_{\text{PL}}$ to express DIM.PL in Dan where the frequency of the PL meaning is high. Similarly, if DIM.PL is construed with DIM as the more prominent component, participants will profile that component by choosing a plural suffix to express DIM.PL. More specifically, they will choose -*nem*$_{\text{DIM}}$ and -*shoon*$_{\text{DIM}}$ to express DIM.PL in Nem where the frequency of the DIM meaning is high.

5.4.1 Production

Figure 5.7, Figure 5.8, and Figure 5.9 present data from each suffix in the production task in experiments I, II, and III, respectively. If frequency changed the construal of the DIM.PL, then *both* plural suffixes would be used to express DIM.PL in Dan more than in Nem, and both diminutive suffixes would be used to express DIM.PL in Nem more than in Dan.

To test whether construction frequency influenced construal of DIM.PL, a logistic regression model tested for an effect of Language on the choice of a diminutive vs. a plural suffix to express DIM.PL in production. Language had a significant effect on DIM vs. PL choice in the expected direction in Experiment I ($\beta = -2.538$, $z = -2.482$, $p = .0131$) and Experiment III ($\beta = -17.69$, $z = -5.512$, $p < .0001$) but not in Experiment II ($\beta = -0.352$, $z = -0.177$, $p = .859$). Across experiments, there was a preference to express PL over DIM in Experiment I ($\beta = 1.294$, $z = 2.702$, $p = .0069$) and Experiment II ($\beta = 2.444$, $z = 2.232$, $p = .0256$) but not in Experiment III ($\beta = -0.089$, $z = -0.098$, $p = .922$), where the choice of DIM vs. PL was strongly influenced by Language.

Overall, these results suggest that PL was a more salient feature of DIM.PL than DIM, or that participants considered it more important to express plurality. It also suggests that construction frequency can influence both the choice of form to express a meaning and the choice of what feature of the meaning to profile. That is, high frequency of -*dan*$_{PL}$ not only increases the likelihood of choosing -*dan*$_{PL}$ over -*sil*$_{PL}$, the other plural suffix, but also increases the likelihood of choosing *a* plural suffix over a diminutive one.

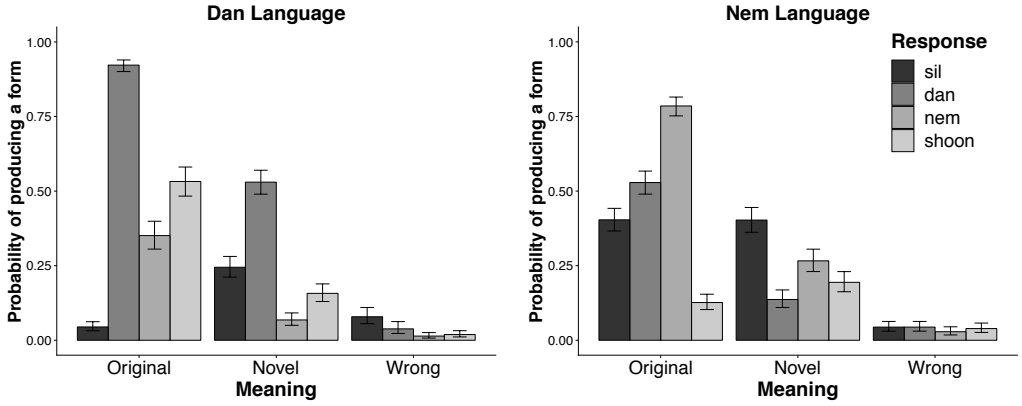

Figure 5.7. Results of the production task in Experiment I.

In this experiment, production preceded form choice which preceded comprehension. Original meaning is the meaning paired with a suffix during training; Novel meaning is DIM.PL; Wrong meaning is the meaning paired with a different suffix in training.

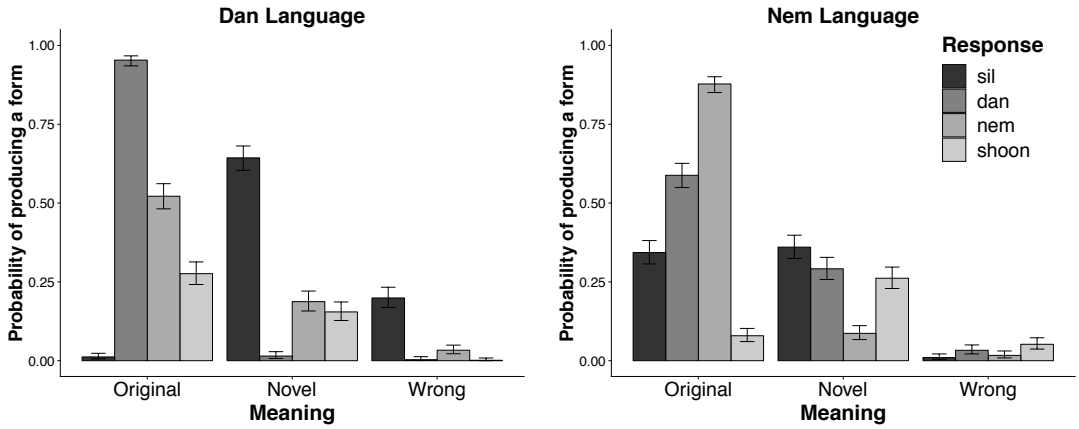

Figure 5.8. Results of the production task in Experiment II.
In this experiment, comprehension preceded production, which preceded form choice. Original meaning is the meaning paired with a suffix during training; Novel meaning is DIM.PL; Wrong meaning is the meaning paired with a different suffix in training.

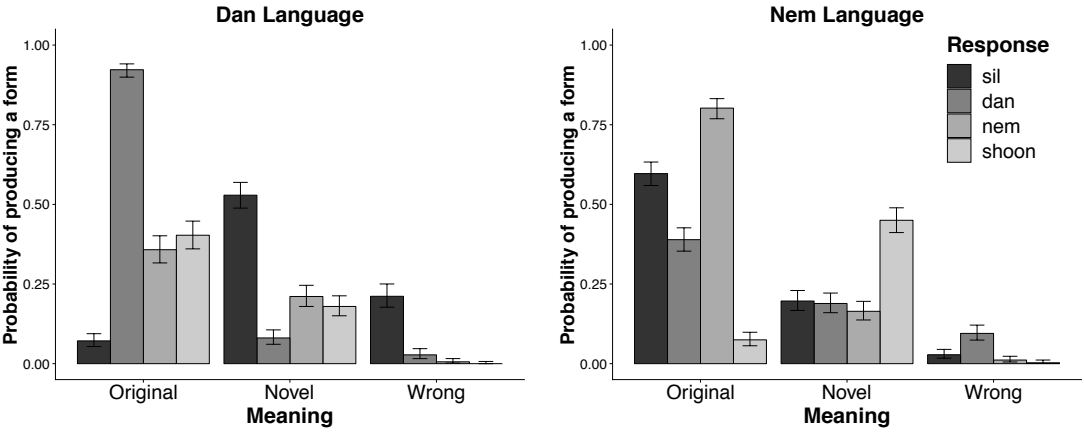

Figure 5.9. Results of the production task in Experiment III.
In this experiment, form choice preceded comprehension, which preceded production. Original meaning is the meaning paired with a suffix during training; Novel meaning is DIM.PL; Wrong meaning is the meaning paired with a different suffix in training.

5.4.2 Form Choice

The second task where the direction of processing is from meaning to form is forced form choice. The forced form choice task is often argued to provide a better window on semantics compared to production because it frees participants from accessibility issues (Ambridge et al., 2018; Schwab et al., 2018). Freed from the constraints of trying to access forms in real time, the speaker can choose the optimal form to highlight whatever aspect of the meaning they wish to profile. For example, the speaker may choose to say *intelligent* instead of *knowledgeable* to highlight certain features of the meaning.

Figure 5.10, Figure 5.11, and Figure 5.12 present data from each suffix in the Form Choice task in Experiments I, II, and III, respectively. From these data, there seems to be no evidence for frequency of meaning influencing construal of DIM.PL. A logistic regression model tested the effect of Language on the choice of a plural suffix vs. a diminutive suffix to express DIM.PL. Language had no significant influence on the choice in Experiment I ($\beta = -0.113$, $z = -0.291$, $p = .771$) or Experiment III ($\beta = -0.324$, $z = -1.068$, $p = .285$). There was an effect of Language in Experiment II, but in the unexpected direction: participants used plural suffixes to express DIM.PL in Nem more than in Dan ($\beta = 0.883$, $z = 2.264$, $p = .024$). Closer examination of the data shows this effect to be driven by the entrenchment effect reported in Chapter III: the frequent plural suffix *-dan*$_{\text{PL}}$ is entrenched to BIG.PL in Dan and the frequent diminutive *-nem*$_{\text{DIM}}$ is entrenched to DIM.SG in Nem during comprehension. These mappings are then reused in the subsequent form choice task. As a result, the frequent feature of DIM.PL is expressed less often than the infrequent feature in Experiment II.

To test whether DIM.PL was more likely to be construed as DIM or PL independently of Language, Language was excluded from the model, leaving an intercept-only model. There was no significant preference for one construal over another in any experiment (Experiment I: $\beta = 0.052$, $z = 0.265$, $p = .791$; Experiment II: $\beta = 0.108$, $z = 0.559$, $p = .576$; Experiment III: $\beta = -0.075$, $z = -0.485$, $p = .628$).

The lack of a frequency effect in the form choice tasks indicates that the effect of frequency in production is mediated by form accessibility. That is, the more salient PL feature of DIM.PL activates the associated forms more strongly than the less salient DIM feature, and a frequent feature activates the associated forms more strongly than an infrequent one. When form accessibilities are leveled, frequency and salience no longer have an effect on form choice.

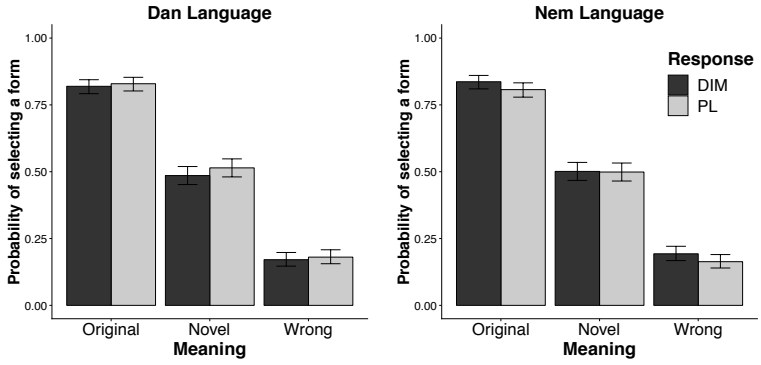

Figure 5.10. Results of the form choice task in Experiment I.
In this experiment, production preceded form choice, which preceded comprehension. Original meaning is the meaning paired with a suffix during training; Novel meaning is DIM.PL; Wrong meaning is the meaning paired with a different suffix in training. DIM suffixes are -*nem*$_{\text{DIM}}$, and -*shoon*$_{\text{DIM}}$ and PL suffixes are -*dan*$_{\text{PL}}$ and -*sil*$_{\text{PL}}$.

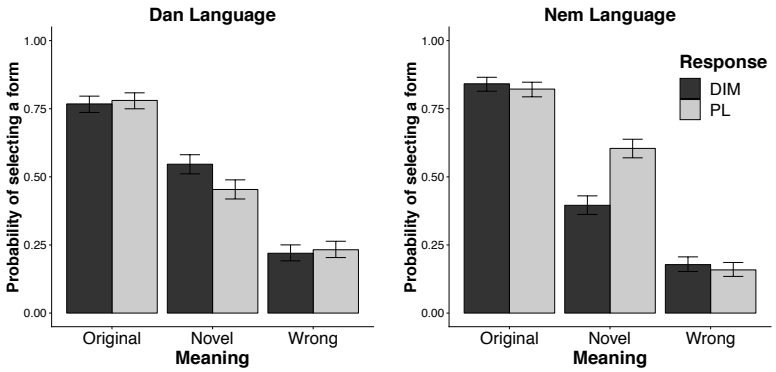

Figure 5.11. Results of the form choice task in Experiment II.
In this experiment, comprehension preceded production, which preceded form choice. Original meaning is the meaning paired with a suffix during training; Novel meaning is DIM.PL; Wrong meaning is the meaning paired with a different suffix in training. DIM suffixes are -*nem*$_{\text{DIM}}$, and -*shoon*$_{\text{DIM}}$ and PL suffixes are -*dan*$_{\text{PL}}$ and -*sil*$_{\text{PL}}$.

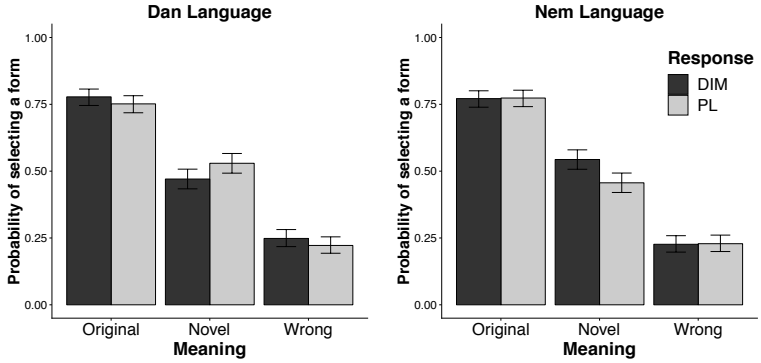

Figure 5.12. Results of the form choice task in Experiment III.
In this experiment, form choice preceded comprehension, which preceded production. Original meaning is the meaning paired with a suffix during training; Novel meaning is DIM.PL; Wrong meaning is the meaning paired with a different suffix in training. DIM suffixes are -*nem*$_{\text{DIM}}$, and -*shoon*$_{\text{DIM}}$ and PL suffixes are -*dan*$_{\text{PL}}$ and -*sil*$_{\text{PL}}$.

5.5 General Discussion: Niche-seeking

The effect of frequency on competitors of the frequent suffixes was explored in three experiments. The findings suggest that when a frequent and an infrequent form share a meaning, they compete for the shared meaning. Over time, the frequent form

strengthens its association with the meaning and pushes the competing form out of the shared meaning. In comprehension, the meaning of a frequent form co-occurring with a specific meaning narrows to that meaning. The meaning of its semantic competitor also narrows, to occupy the part of semantic space abandoned by the frequent form during the process of entrenchment.

While semantic broadening resulting from extension is perhaps more common than semantic narrowing (Bybee, 2003), both are attested. Our results suggest that narrowing can result from entrenchment or from pre-emption by a highly accessible form, which provides a mechanistic account of the process of niche-seeking (Aronoff, 2016; MacWhinney, 1987, 1993; see also Lindsay & Aronoff, 2013). For example, when *dog* was borrowed into English from Scandinavian, *hound* became specialized to a kind of hunting dog, surviving in its new niche (Traugott & Dasher, 2002, pp. 52–53). As *going to* is growing in strength as a future marker, it is pushing *will* out of the proximal future (Poplack & Tagliamonte, 1999) and 2nd person subject (Torres Cacoullos & Walker, 2009) contexts.

An alternative account of the pushing out phenomenon is that competition leaves the competitor of a frequent form meaningless, i.e., unassociated with a meaning. Because the learner assumes that all forms have meanings, she associates the meaningless form with the relatively novel meaning. This effect is reminiscent of the Novel Name Nameless Category Principle (Mervis & Bertrand, 1994). However, instead of a novel name mapping onto a nameless category, (newly) meaningless form would map onto a novel meaning.

Left unassociated with a meaning, the forms with frequent competitors can then map onto DIM.PL in comprehension because—in contrast to the meaning the form has appeared with—the DIM.PL meaning is novel and novelty draws attention. In addition—unlike the familiar meanings that appeared with other suffixes—the novel meaning has not had the opportunity to develop an inhibitory association with the form or to become associated with other suffixes. Under this account, the learner maps the form whose meaning she does not know onto the meaning that has not been *dis*associated from any form (McMurray, Horst, & Samuelson, 2012). Intriguingly, it is then not the frequent competing form that blocks its infrequent competitor from associating with the shared meaning. If the shared meaning occurred in the absence of any form for the same number of trials, the same effect would be expected to occur (see Ramscar, Dye, & Klein, 2013, for supportive data).

I leave it to future work to examine the roles of form competition vs. meaning frequency in this process of niche seeking. However, the underlying assumption of this alternative account that competitors of frequent forms are left meaningless by the competition is inconsistent with the results of forced form choice. When participants were presented with a meaning shared between a frequent form and an infrequent form, and asked to choose between the frequent form and its competitor, participants were equally likely to choose either form. In contrast, participants were far less likely to choose the competitor of a frequent form when presented with a meaning it was not paired with in training (the wrong meaning). Thus, participants generally knew to choose the competitor of a frequent form only in response to a meaning with which it has been paired. This result is inconsistent with the idea that competitors of frequent forms are

meaningless. Overall, the results therefore suggest that competition with a frequent form causes the prototype of a form's meaning to shift out of the semantic space the forms share.

An additional way to disentangle the predictions of the two accounts is to teach participants an artificial language in which the DIM.PL meaning is not novel, and is presented during training. This can be accomplished by explicitly mapping -*dan*$_{\text{PL}}$ onto PL and -*nem*$_{\text{DIM}}$ onto DIM, presenting both with the DIM.PL meaning. One of these forms can then be made to appear more frequently than the other in a specific meaning other than DIM.PL. The increasing frequency outside of the DIM.PL meaning should then shift the semantic prototype of the form away from DIM.PL. Thus, as -*dan*$_{\text{PL}}$ appears more and more often in BIG.PL, it should entrench into that meaning, as in Experiments I−III, leaving the DIM.PL meaning available for -*sil*$_{\text{PL}}$ to occupy. In contrast, if -*sil*$_{\text{PL}}$ is mapped onto DIM.PL when -*dan*$_{\text{PL}}$ is frequent only because novelty makes DIM.PL the most salient meaning (McMurray et al., 2012), then -*sil*$_{\text{PL}}$ should no longer be mapped onto DIM.PL when DIM.PL is not novel. This experiment is presented in the next chapter.

CHAPTER VI

NO SYNONYMY

In this chapter, I report on a fourth experiment testing the effect of frequency on extension versus entrenchment of forms. This experiment follows the task order of Experiment I, and is geared towards replicating the production-comprehension dissociation observed in that experiment. However, the form–meaning contingencies of the present experiment make mapping -*dan*$_{\text{PL}}$ and -*nem*$_{\text{DIM}}$ onto the basic-level single-feature categories PL and DIM consistent with training (see Figure 6.1). I have argued that the participants have also done this in Experiments I−III, and a bias in favor of single-feature meanings is consistent with both morphological typology (Baerman, Brown & Corbett, 2005; Pertsova, 2011) and experimental data on category learning (e.g., Goodman, Tenenbaum, Feldman, & Griffiths, 2008; Kurtz, Levering, Stanton, Romero, & Morris, 2013). However, the present experiment provides an important direct verification of this assumption in the semantic domains of number and size.

In the present experiment, mapping an infrequent -*dan*$_{\text{PL}}$ onto PL and an infrequent -*nem*$_{\text{DIM}}$ onto DIM is what the participants should do even if they have no bias for single-feature categories. If the results are similar to Experiment I, we can then confirm that they treated the training data of the first experiment as equivalent to the training data of this one and inferred single-feature meanings for infrequent -*dan*$_{\text{PL}}$ and -*nem*$_{\text{DIM}}$ whether or not they were explicitly paired with DIM.PL in training.

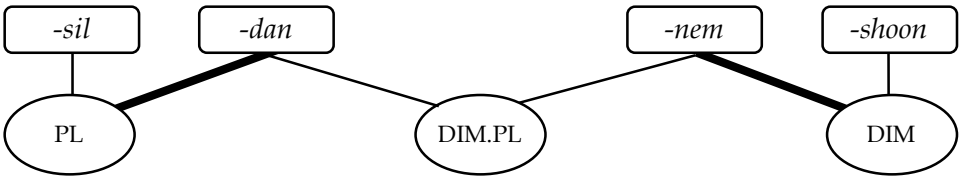

Figure 6.1. Experiment IV exposure
Structure of the two artificial languages Dan (where -*dan*$_{PL}$ is frequent in the PL meaning) and Nem (where -*nem*$_{DIM}$ is frequent in the DIM meaning) during exposure. Thick lines represent the frequent form–meaning mappings: the plural -*dan*$_{PL}$ in the Dan language and the diminutive -*nem*$_{DIM}$ in the Nem language.

An additional goal of this experiment is to disentangle the reasons for mapping competitors of frequent suffixes onto DIM.PL. As discussed in Chapter V, this result could occur 1) because frequent forms entrench into specific meanings that exclude DIM.PL and therefore do not preempt their infrequent competitors from remaining associated with that meaning; or 2) because the competitors of frequent forms are rendered meaningless by pre-emption, and are mapped onto the only novel meaning, presumably because of its salience (a version of the N3C Principle; Mervis & Bertrand, 1994). While I have argued that the form choice data of Chapter V are inconsistent with the N3C, the data of this chapter provide an additional way to distinguish among the two hypotheses by rendering the DIM.PL meaning familiar. If participants still choose it for competitors of frequent forms, this is then most likely due to its consistency with the meaning of the competitors of frequent forms.

The present experiment changes the structure of the input language in a way that, arguably, makes it more like the structure of natural languages. According to many researchers, natural languages abhor exact synonymy (e.g., Clark, 1987; Givon, 1979; Goldberg, 1995, 2006, 2019). The language used for Experiments I−III had completely synonymous suffixes, which may have led the learners to reanalyze it, assigning

competing suffixes to different meanings (Goldberg, p.c.). The language presented to learners in the present chapter was therefore designed to avoid complete synonymy. Thus, while in Experiments I−III, -*dan*$_{PL}$ and -*sil*$_{PL}$ were both paired with BIG.PL in training, and -*nem*$_{DIM}$ and -*shoon*$_{DIM}$ were paired with DIM.SG, here -*dan*$_{PL}$ is PL, -*sil*$_{PL}$ is BIG.PL, -*nem*$_{DIM}$ is DIM, and -*shoon*$_{DIM}$ is DIM.SG. While every meaning is shared by two suffixes, the choice of the suffix in this language can be seen as reflecting a *construal* of the referent (Langacker, 1987; see also Croft & Cruse, 2004). Thus, the choice between -*dan*$_{PL}$ and -*sil*$_{PL}$ for BIG.PL is equivalent to choosing between calling a creature a dog vs. a Dalmatian. The choice between -*dan*$_{PL}$ and -*nem*$_{DIM}$ for expressing DIM.PL can be seen as profiling diminutiveness versus number. DIM in particular is a category that is usually optional in languages of the world, in that the speaker can choose whether to express it or not. Number can also be optional (e.g., Müller, Storto, & Coutinho-Silva, 2006). There is therefore little reason for the learner to reanalyze the language in this experiment.

An additional difference between the language in Experiments I−III and that in Experiment IV is that entrenchment in Experiment IV cannot be interpreted as narrowing the extent of the region of semantic space associated with a form. A frequent -*dan*$_{PL}$ in Experiments I−III is never paired with DIM.PL in training. DIM.PL may therefore be excluded from the range of meanings -*dan*$_{PL}$ is thought to express as -*dan*$_{PL}$ grows in frequency, via the suspicious coincidence logic (Xu & Tenenbaum, 2007). In contrast, a frequent -*dan*$_{PL}$ in the present experiment is still mapped onto both DIM.PL and BIG.PL in training. Rather than learning that a frequent -*dan*$_{PL}$ means only BIG.PL, learners in

this experiment should therefore consider BIG.PL to be the *prototype* of the meaning of -*dan*$_{\text{PL}}$ when it is frequent (Bybee & Eddington, 2006; Lakoff, 1987; Rosch, 1978).

The fact that -*dan*$_{\text{PL}}$ and -*nem*$_{\text{DIM}}$ are presented with DIM.PL during training also means that the use of a form with the DIM.PL meaning cannot be considered extension of that form to a novel related meaning, unlike in Experiments I–III. Rather, it constitutes the use of a form to express a non-prototypical meaning that is part of its semantic range. From my perspective, however, extension is effected by precisely the same factors that affect the choice of a form in production when multiple near-synonyms are competing to express the meaning. If the results of the present experiment mirror those of Experiment I, this would provide additional support for this hypothesis.

Finally, the present experiment makes a potentially pre-empting form available for the DIM.PL meaning. In the morphological literature, it is assumed that forms whose meanings are most similar to the meaning-to-be-expressed block other forms, whose meanings are less similar (e.g., Caballero & Inkelas, 2013; Kapatsinski, 2018a, 2018b; Kiparsky, 1973; MacWhinney, 1978, 1987; Ramscar, Dye & McCauley, 2013; Yang, 2016). For example, *went* blocks *go* + *-ed* when the speaker wants to express GO and PAST because it can express both GO and PAST, whereas *go* can express only GO and *-ed* can express only PAST. The increasing frequency of -*dan*$_{\text{PL}}$ in the BIG.PL meaning is expected to make its meaning less compatible with DIM.PL. One might therefore expect that increasing the frequency of a form in a meaning other than DIM.PL would allow the other form paired with DIM.PL in training to pre-empt it. For example, an infrequent -*nem*$_{\text{DIM}}$ (which is mapped onto DIM) would block the frequent -*dan*$_{\text{PL}}$ from expressing the shared DIM.PL meaning. An infrequent -*dan*$_{\text{PL}}$ (mapped onto PL) would block the

frequent -*nem*$_{\text{DIM}}$ (mapped onto DIM.SG) from being used to express DIM.PL. However, this prediction is not made by the Hebbian learning models developed in Chapter II unless frequently encountering a form in the absence of a semantic feature cue leads to the development of strong inhibitory associations between the feature and the form. As discussed in Chapter II, most models of associative learning either do not learn associations involving absent cues (Rescorla & Wagner, 1972) or consider this kind of learning to be slow relative to learning associations involving present cues (Tassoni, 1995). I therefore did not expect infrequent forms to block frequent forms in this experiment.

Despite the differences between Experiments I and IV, I expected the results to be similar. That is, I expected divergent effects of frequency in production and comprehension and no effect of frequency in the form choice task, where both frequent and infrequent forms are equally accessible. More specifically, I expected the same dissociation between comprehension and production so that an entrenchment effect in comprehension and an extension effect in production would be observed. Experiment I demonstrated that increasing the frequency of a form–meaning pairing during training makes participants more likely to select the meaning when they are presented with the form. Frequent forms should have an advantage in production because they are more accessible when cued with DIM or PL. The prediction of the current experiment is that increasing the frequency of a form–meaning pairing renders the form more accessible, leading not only to the strengthening of its connection to its original meaning but also to increasing use of that form to express related meanings in which it has not been so frequently encountered. That is, despite increased frequency in BIG.PL shifting the

semantic prototype of -*dan*$_{PL}$ to BIG.PL, I expect a frequent -*dan*$_{PL}$ to be increasingly used to express DIM.PL. The increased use of -*dan*$_{PL}$ should occur despite the fact that there is another form that could express the meaning. I now describe the specifics of this design.

6.1 Languages in Experiment IV

Two artificial languages were used in the experiment: Dan and Nem (Figure 6.1). Dan and Nem had the same morphological constructions, and shared the same set of nouns in their vocabulary. The nouns were 30 one- and two-syllable nonce nouns (e.g., *chool* and *osto*).

Dan and Nem consisted of the same four suffixes: -*sil*$_{PL}$, -*dan*$_{PL}$, -*nem*$_{DIM}$, and -*shoon*$_{DIM}$. The suffixes -*sil*$_{PL}$ and -*dan*$_{PL}$ are plural while -*shoon*$_{DIM}$ and -*nem*$_{DIM}$ are diminutive. However, whereas -*sil*$_{PL}$ was always paired with a picture of multiple large creatures during the exposure phase, -*dan*$_{PL}$ was paired with pictures of both multiple large creatures and multiple small creatures. Similarly, -*shoon*$_{DIM}$ was paired with a picture of one small creature, while -*nem*$_{DIM}$ was paired with both a picture of one small creature and multiple small creatures. The identity of the creature varied with the stem. An unaffixed stem was always paired with a single large creature. As a result, -*sil*$_{PL}$ could be thought of as either simply plural (PL) or plural non-diminutive (BIG.PL) while -*dan*$_{PL}$ was probably more strongly conceived as PL. Likewise, -*shoon*$_{DIM}$ could be thought of as either diminutive (DIM) or diminutive singular (DIM.SG) while and -*nem*$_{DIM}$ was likely conceived more as a general DIM meaning. These possible interpretations were probed in the test phase.

In each language, one suffix is more frequent than the other four suffixes: -*dan*$_{PL}$ is the frequent suffix in the Dan language and -*nem*$_{DIM}$ is the frequent suffix in the Nem language. The thicker lines in Figure 6.1 highlight this difference in frequency. The identity of the frequent suffix is the only difference between the two languages. Importantly, the frequency was not increased in the DIM.PL meaning.

The main difference between this design and the design of the last three experiments was the fact that not only was DIM.PL present during the training phase, but that it was paired with both suffixes -*dan*$_{PL}$ and -*nem*$_{DIM}$ with equal frequency. An increase in the use of one of these two suffixes with the DIM.PL meaning during the production test phase of the experiment constitutes extension of the frequent suffix to a meaning that is not its prototypical meaning. Use of the two suffixes -*sil*$_{PL}$ and -*shoon*$_{DIM}$ with the DIM.PL meaning constitutes extension of a form to a novel but related meaning.

6.2 Methods

6.2.1 Participants

Eighty adult native speakers of American English, all undergraduate students at the University of Oregon, participated in the experiment. As before, each participant was exposed to only one of the two languages (40 participants experienced Dan, the language where -*dan*$_{PL}$ was the frequent suffix, while 40 others experienced Nem, the language where -*nem*$_{DIM}$ was the frequent suffix).

6.2.2 Tasks

The experiment started with an exposure phase, followed by production, followed by the forced forms choice task followed by comprehension. The procedure for all the

tasks in the current experiment were the same as in Experiment I. The visual stimuli and pictures were the same as the previous experiments.

6.2.3 Stimuli

The same test trials were used for all participants, but the training varied by language. The training comprised 220 trials (55 trials repeated in random order in 4 cycles). Ten stems were presented to learners in each cycle. Five stems were used once with each suffix in each meaning (e.g., *korasil*, *koradan*, *koradan*, *koranem*, *koranem*, *korashoon*; 30 items). The other five (*bani*, *chool*, *kudom*, *osto*, *vorke*) were used without a suffix once (5 items) and were used four other times with the frequent suffix (20 items).

The stimuli for production and forced choice tasks were the same as Experiment I. The comprehension task was different from Experiment I in that it did not include the compositional option, therefore comprising 50 trials in which the unaffixed stem, -*sil*$_{PL}$, -*dan*$_{PL}$, -*nem*$_{DIM}$, and -*shoon*$_{DIM}$ were presented to the subjects.

6.3 Results

6.3.1 Main suffixes

As in the previous chapters, I will present the results of the tasks in the orderof presentation to the subjects. The order of tasks in experiment IV was the same as Experiment I: participants encountered the production task immediately after exposure, followed by the form choice task, followed by comprehension. I will first report on the effect of frequency on the mapping of the main suffixes -*dan*$_{PL}$ and -*nem*$_{DIM}$ to the original and novel meaning. The effect of competition with these frequent suffixes on the mappings of the competing suffixes, -*sil*$_{PL}$ an and -*shoon*$_{DIM}$ is reported next.

6.3.1.1 Production

The analysis was performed on 70.50% of responses that included one of the four suffixes -*dan*$_{PL}$, -*nem*$_{DIM}$, -*sil*$_{PL}$ an and -*shoon*$_{DIM}$. Of the excluded suffixes, 1.45% were missing, 4.13% were stems, and less than 1% of the responses were responses in which either more than one suffix was provided or that could potentially be considered compositional. The rest of the responses were mispronunciations of the existing suffixes, blends of suffixes, blends of one or more suffix(es) with the stem, or nonce words.

The results are presented in the Production panel of Figure 6.2. Frequent suffixes were chosen on average 91% of the time in response to their frequent meaning (PL for -*dan*$_{PL}$ and DIM for -*nem*$_{DIM}$). However, when the same suffixes were infrequent (and therefore as frequent as their competitors), they were only chosen 39% of the time. The height of the bars in the Wrong meaning demonstrate that participants rarely used these suffixes in response to meanings with which they were not presented during training. Therefore, we can conclude that participants learned the meanings of the suffixes sufficiently well.

The accessibility hypothesis predicts that a frequent suffix should be more likely than an infrequent suffix to be used to express the meaning DIM.PL, despite the fact that it matches only one feature of a frequent form's semantic prototype. In contrast, entrenchment to the prototype predicts that a frequent suffix should be *less* likely than an infrequent suffix to be chosen to express the DIM.PL meaning. A mixed-effects logistic regression model was used to test these predictions. The model included random intercepts for subjects and noun stems and random slopes for language within stem. A form was significantly more likely to be chosen to express DIM.PL when it was

frequently encountered during exposure ($\beta = 2.97$, $z = 2.34$, $p = 0.02$). These results replicate the results of the first experiment and support the accessibility hypothesis.

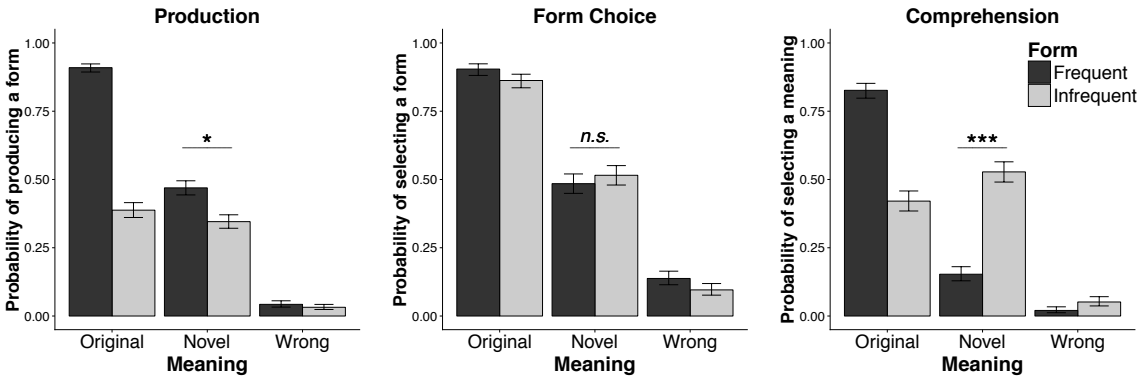

Figure 6.2. The results of Experiment IV for the two suffixes -*dan*$_{PL}$ and -*nem*$_{DIM}$. Original meaning is PL for -*dan*$_{PL}$ and DIM for -*nem*$_{DIM}$; Novel meaning is DIM.PL; Wrong meaning is meaning paired with a different suffix in training, e.g. DIM.SG for -dan$_{PL}$. Frequent suffix = -*dan*$_{PL}$ in Dan and -*nem*$_{DIM}$ in Nem. Infrequent suffix = -*dan*$_{PL}$ in Nem and -*nem*$_{DIM}$ in Dan. Left panel: Form choice probabilities in the production task. The distance between the tops of the bars and 100% for the original meaning in the Production panel is composed largely of probability of choosing the synonymous suffix: -*shoon*$_{DIM}$ over -*nem*$_{DIM}$ or -*sil*$_{PL}$ over -*dan*$_{PL}$. Middle panel: Form choice probabilities in the forced form choice task on the -*dan*$_{PL}$ vs. -*nem*$_{DIM}$ trials. Right panel: Meaning choice probabilities in the comprehension task. (*** means $p < .001$ and * means $p < .05$).

6.3.1.2 Form Choice

The analysis was conducted on 97.53% of the data, corresponding to trials on which participants provided responses. The results are presented in the Form Choice panel of Figure 6.2. A mixed-effects logistic regression model with maximal random-effects structure and Language as the fixed predictor indicated no significant effect of frequency on the choice of suffix to express DIM.PL ($\beta = -0.388$, $z = -0.739$, $p = .46$). To ensure that the null effect of frequency in the form choice task is reliably different from its significant effect in the production, a model tested the interaction between the two tasks. The model detected a significant difference between the two tasks with respect

to the effect of frequency; i.e., a task-by-frequency interaction ($\beta = -2.578$, $z = -3.299$, $p < 0.001$). Overall, the results of the form choice task in the current experiment are parallel to the results of the Form Choice task in Experiment I: The effect of frequency on form choice is mediated by accessibility whether that choice is made in response to a novel meaning or an extant meaning.

6.3.1.3 Comprehension

The analysis was conducted on all trials for which participants provided responses (99.29%). The results are shown in the Comprehension panel of Figure 6.2. The data were analyzed using a mixed-effects logistic regression model with frequency as the fixed effect of interest and the binary dependent variable of mapping a suffix onto its original meaning (-dan_{PL} to BIG.PL and -nem_{DIM} to DIM.SG) versus the meaning of DIM.PL. Random slopes for frequent suffix within both subject and stem were also included.

Mapping onto DIM.PL should become less likely as the form's frequency of co-occurrence with another meaning increases. As participants repeatedly encounter a frequent form paired with BIG.PL or DIM.SG, they should become less likely to select DIM.PL in response to the form. The results revealed that this is indeed the case. A form was mapped onto DIM.PL significantly less often when it was frequent during exposure ($\beta = 6.613$, $z = 4.037$, $p < .0001$). This replicates the results of the comprehension task in Experiment I. The parallelism between the results of the comprehension task in Experiment I and Experiment IV suggests that participants in both tasks initially mapped -dan_{PL} and -nem_{DIM} onto basic single-feature meanings in both experiments. Individual subject data in Figure 6.3 provide more detailed information on the subject's behavior.

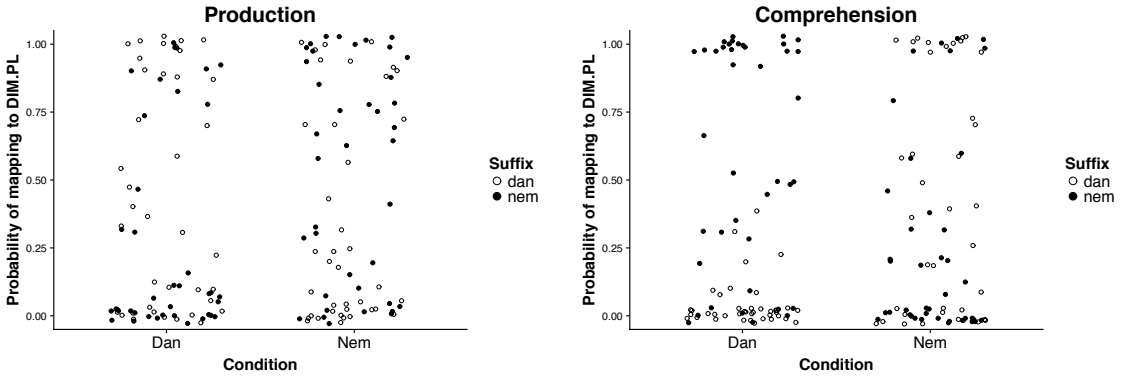

Figure 6.3. Individual participants' mappings of -*dan*$_{PL}$ and -*nem*$_{DIM}$ to DIM.PL meaning in the production and comprehension tasks in Experiment IV.
Each white dot corresponds to an individual subject's use of -*dan*$_{PL}$ and each black dot corresponds to an individual subject's use of -*nem*$_{DIM}$. In production, the dots represent probabilities of choosing a suffix in response to DIM.PL. In comprehension, the dots represent probabilities of choosing DIM.PL in response to a suffix.

6.3.2 Competitor frequency effects

As in Chapter V, we test the hypothesis that competitors of frequent forms are pre-empted by those forms from mapping onto the meaning that constitutes a frequent form's semantic prototype. As in Chapter V, the prediction is that the competitors will be pushed out of the prototypical meaning of the frequent form and retain only the part of the meaning that does not correspond to the frequent form's prototype. Unlike in Chapter V, however, this meaning is still part of the semantic range of the frequent form. There is therefore a possibility that retention of a semantic niche is harder for the competitors in the current design. They may therefore be pushed out of existence in production, and not reliably mapped onto any meaning in comprehension.

6.3.2.1 Production

The results for the competitors (-*sil*$_{PL}$ and -*shoon*$_{DIM}$) of the frequent suffixes are presented in the Production panel of Figure 6.4 (see Figure 6.5 for individual subject

data). An increase in the frequency of a form causes the synonymous competitor of that form to not be produced in response to the shared meaning with which both suffixes were paired during training. That is, high frequency of a form pushes its semantic competitors into obsolescence. This may be especially likely for forms whose semantic range is fully subsumed by the semantic range of a frequent form, and corresponds to the frequent form's semantic prototype.

A mixed-effects logistic regression model with maximal random-effects structure revealed that competitors of the frequent suffixes were less likely to be produced in response to their original meanings than competitors of infrequent suffixes ($\beta = 15.203$, $z = 3.84$, $p = .0001$). However, the results revealed no significant effect of language on whether these competing suffixes was more likely to be chosen to express the novel meaning of DIM.PL ($\beta = -3.325$, $z = -1.01$, $p = .313$).

Compared to Experiments I−III, where the competing forms -*dan*$_{\text{PL}}$ and -*nem*$_{\text{DIM}}$ were not paired with DIM.PL in training (Chapters III–IV), the present results show a dramatically lower rate of use of -*sil*$_{\text{PL}}$ and -*shoon*$_{\text{DIM}}$ in the DIM.PL meaning. When the range of a frequent form subsumes the range of a competing infrequent form, the infrequent form has no niche in which it can escape competition. It is therefore likely to be pushed into obsolescence rather than undergoing a prototype shift and semantic narrowing to the new prototype.

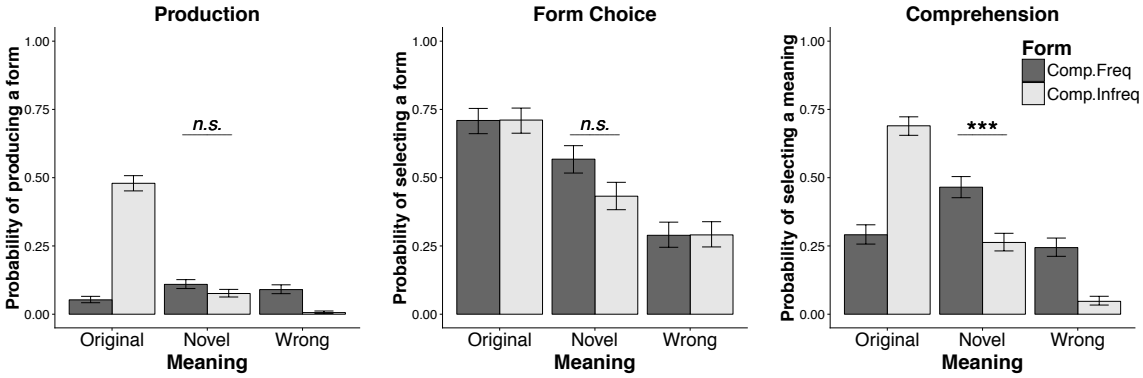

Figure 6.4. The results of Experiment IV for the competitor suffixes -*sil*$_{PL}$ and -*shoon*$_{DIM}$. Original meaning is the meaning paired with a suffix during training, e.g. BIG.PL for -*sil*$_{PL}$; Novel meaning is DIM.PL; Wrong meaning is the meaning paired with a different suffix in training, e.g. DIM.SG for -*sil*$_{PL}$. Competitor of the frequent suffix = -*sil*$_{PL}$ in Dan and -*shoon*$_{DIM}$ in Nem. Competitor of the infrequent suffix = -*sil*$_{PL}$ in Nem and -*shoon*$_{DIM}$ in Dan. Left panel: Form choice probabilities in the production task. Middle panel: Form choice probabilities in the forced form choice task on the -*sil*$_{PL}$ vs. -*shoon*$_{DIM}$ trials. Right panel: Meaning choice probabilities in the comprehension task.

6.3.2.2 Form choice task

Choice of the competitor suffixes in the form choice task was analyzed using a mixed-effects logistic regression model with language as the fixed effect of interest and the binary dependent variable of suffix (-*sil*$_{PL}$ or -*shoon*$_{DIM}$) chosen to express the novel meaning DIM.PL. Random intercepts for subjects and noun stems and random slopes for language within stem were included in the model. The analysis revealed no significant effect of language on suffix choice ($\beta = -0.847$, $z = -1.875$, $p = .061$). That is, the effects of competitor frequency on form choice in production were largely mediated by the effect of frequency on form accessibility.

Unlike the results of the production task, the results of the form choice task are parallel to those in Chapter III, where -*dan*$_{PL}$ and -*nem*$_{DIM}$ were not paired with DIM.PL in training. Thus, participants in both experiments select the competitor suffixes when presented with their original meanings ~75% of the time, and ~25% of the time when

presented with a meaning that share no features with the suffix's actual meaning. Accuracy in this task is unaffected by frequency of the competitor suffix, or by whether the competitor suffix is paired with DIM.PL in training. Thus, having a frequent competitor does not render a form meaningless. Rather, it simply makes it less likely to be selected for production pushing it into obsolescence.

6.3.2.3 Comprehension

A mixed-effects logistic regression model (with random intercepts for subjects and random slopes for language within stem) probed the effect of competing with a frequent form on the meanings chosen in response to -*sil*$_{PL}$ and -*shoon*$_{DIM}$. Once again, we focused on the likelihood of mapping the forms onto DIM.PL vs. the meaning they were paired with during exposure (PL for -*sil*$_{PL}$ and DIM for -*shoon*$_{DIM}$) as the dependent variable (Original meaning vs. Novel meaning). Language served as the fixed-effects predictor, indicating whether these competing suffixes competed with a frequent or infrequent suffix. I hypothesized that a form should be more likely to be mapped onto DIM.PL when the meaning it was paired with in training constituted the semantic prototype of a frequent form.

The results are presented in Figure 6.4 (see Figure 6.5 for individual subject data). As expected, these suffixes were more likely to be mapped onto the novel meaning of DIM.PL when their competitors were frequent ($\beta = -10.13$, $z = -4.45$, $p < .0001$). These results parallel the results of Experiments I–III. Replication of this results with a familiar DIM.PL means that the mapping of competitors of frequent forms onto DIM.PL in comprehension is not due to the novelty and salience of DIM.PL. Rather, this result can only be attributed to competition with a frequent form shifting the semantic prototype of

the infrequent form away from the semantic prototype of its frequent competitor. This happens even if the prototype shift does not allow the competitor to escape the semantic range of the frequent form.

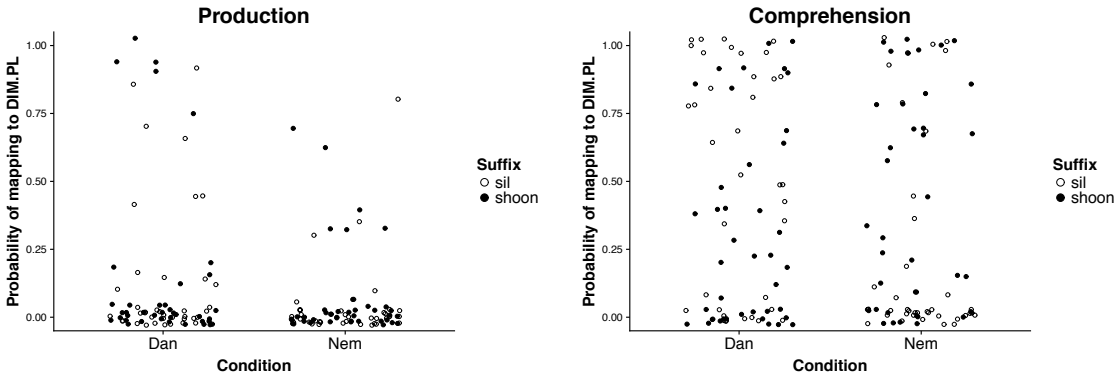

Figure 6.5. Individual participants' mappings of the competitor suffixes -*sil*$_{PL}$ and -*shoon*$_{DIM}$ to DIM.PL meaning in production and comprehension tasks in Experiment IV. Each white dot corresponds to an individual subject's use of -*sil*$_{PL}$ and each black dot corresponds to an individual subject's use of -*shoon*$_{DIM}$. In production, the dots represent probabilities of choosing a suffix in response to DIM.PL. In comprehension, the dots represent probabilities of choosing DIM.PL in response to a suffix.

6.3.3 A comparison of the individual subject behavior across the four experiments

Figure 6.6 further explores the correlation between production and comprehension choices of individual participants across experiments. Each dot in Figure 6.6 shows the probability of a form given the DIM.PL meaning in an individual participant's data. All graphs show a concentration of dots in the bottom left because a participant tends to map only one or two forms onto DIM.PL. Dots in the top right quadrant of the space correspond to participants for whom the form that is most likely to be used to express DIM.PL in production is the form most likely to be mapped onto DIM.PL in comprehension. Many of the participants in Experiment II and Experiment III are of this

type, resulting in a strong positive correlation between production and comprehension choices. In contrast, there is no correlation between production and comprehension choices in Experiment I because many participants fall into the top left and bottom right quadrants of the space. These are participants who use one form to express DIM.PL in production and tend to map a different form onto DIM.PL in comprehension. There are relatively few participants in the top right quadrant. Experiment IV falls in between. However, the positive correlation in Experiment IV does not arise for the same reason as the positive correlations in Experiments II and III.

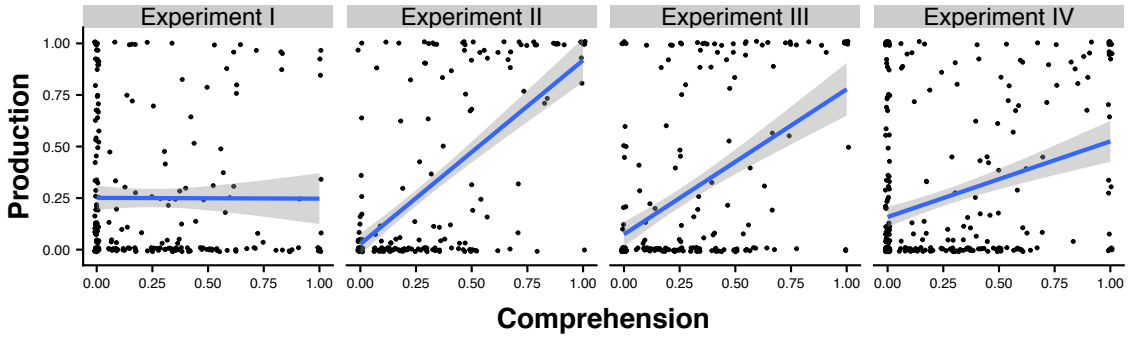

Figure 6.6. The correlation between participants' production and comprehension in the four experiments.
Each dot represents the probability of a suffix given the DIM.PL meaning in the production and comprehension data of an individual participant.

Figure 6.7 and Figure 6.8 show that the dots in the top quadrant of the Experiment IV panel in Figure 6.6 come from the fact that participants in Experiment IV are more likely to map *-dan* and *-nem* onto DIM.PL in comprehension compared to participants in Experiments I, II and III. This is in turn due to the contingency structure of the language: in Experiments I–III, *-dan* and *-nem* are never paired with DIM.PL in training, whereas in Experiment IV they are. As a result, participants in Experiments I–III learn that the

frequent form does not map onto DIM.PL, as seen from the absence of dots in the entire right half of the Experiment I–III panels in Figure 6.7. This is not the case in Figure 6.8, where the same form is infrequent, demonstrating the entrenchment effect. Participants in Experiment IV do not learn to restrict the frequent form from mapping onto DIM.PL because it is explicitly paired with DIM.PL in training.

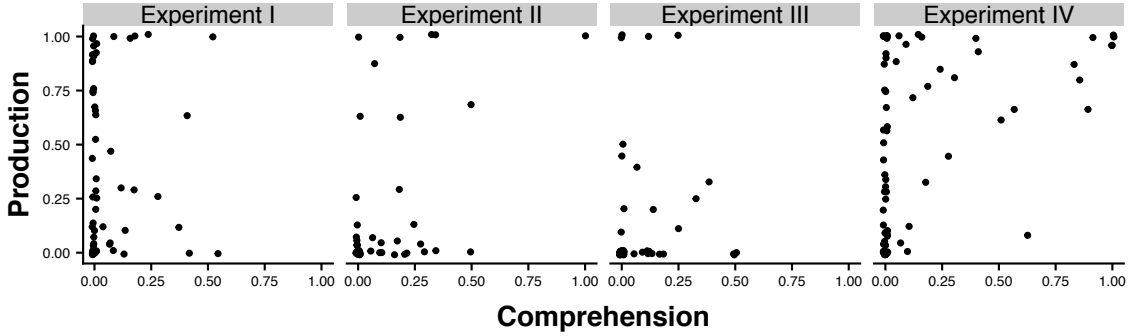

Figure 6.7. The correlation between participants' production and comprehension of frequent forms (-*dan*$_{PL}$ in Dan and -*nem*$_{DIM}$ in Nem) in the four experiments.
Each dot represents the probability of the frequent suffix given the DIM.PL meaning in the production and comprehension data of an individual participant.

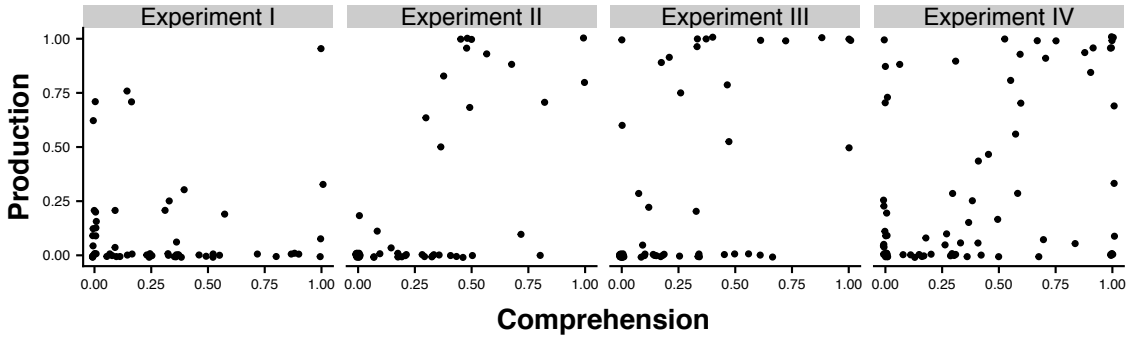

Figure 6.8. The correlation between participants' production and comprehension of infrequent forms (-*dan*$_{PL}$ in Nem and -*nem*$_{DIM}$ in Dan) in the four experiments.
Each dot represents the probability of the infrequent suffix given the DIM.PL meaning in the production and comprehension data of an individual participant.

In contrast, Figure 6.9 shows that most of the dots in the upper right quadrant in the Experiment II and III panels of Figure 6.6 come from participants in these experiments mapping the competitor of the frequent form onto the novel, DIM.PL meaning. That is, the positive correlation between production and comprehension responses of participants in Experiments II and III results from the frequent form pushing its competitor out into the novel meaning. The (weaker) positive correlation in Experiment IV results from learning the form–meaning contingencies of the language.

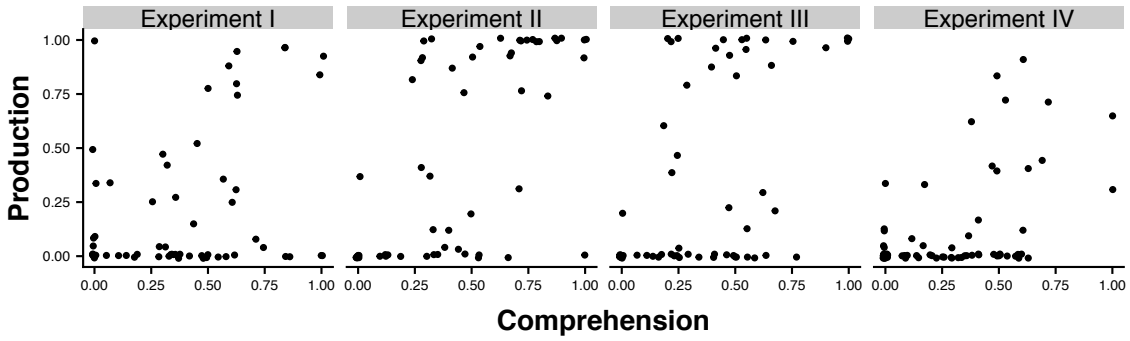

Figure 6.9. The correlation between participants' production and comprehension of competitors of frequent forms (-*sil*$_{PL}$ in Dan and -*shoon*$_{DIM}$ in Nem) in the four experiments.
Each dot represents the probability of the competitor of the frequent suffix given the DIM.PL meaning in the production and comprehension data of an individual participant.

6.3.4 Construal

To evaluate whether the construal of the novel meaning is influenced by frequency of its component features in Experiment IV, I examined the choice of the DIM suffix -*nem*$_{DIM}$ vs. the PL suffix -*sil*$_{PL}$ in both production and forced form choice. Participants in this experiment learned that -*dan*$_{PL}$ and -*nem*$_{DIM}$ map onto DIM.PL, as indicated by the higher bars for -*dan*$_{PL}$ and -*nem*$_{DIM}$ compared to -*sil*$_{PL}$ and -*shoon*$_{DIM}$ in

Figure 6.10, with -*dan*$_{PL}$ mapping onto PL and -*nem*$_{DIM}$ onto DIM. The use of one of these two suffixes in reference to DIM.PL would therefore constitute profiling one of the features of DIM.PL. However, an effect of Language on the use of -*dan*$_{PL}$ or -*nem*$_{DIM}$ in the production task is not diagnostic of construal because it can also be attributed to form accessibility. I therefore focus on the forced form choice task where accessibility differences between forms are leveled.

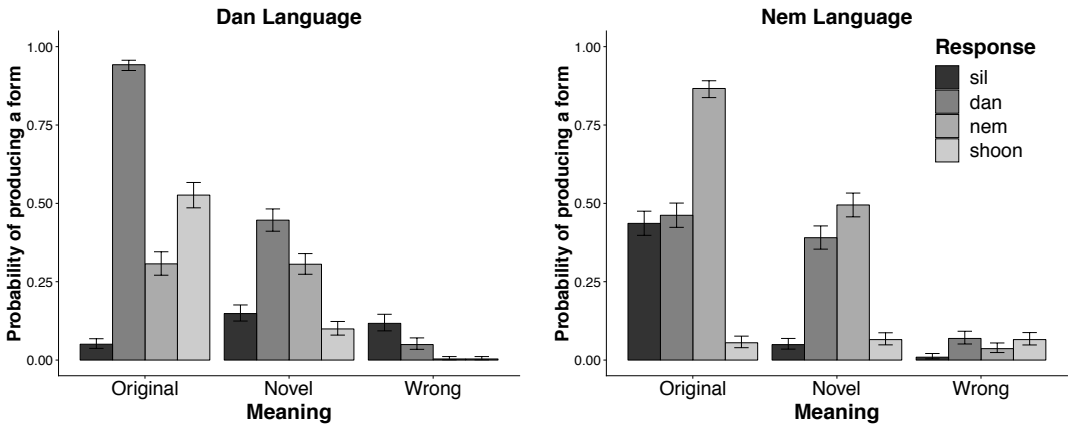

Figure 6.10. Results of the production task in Experiment IV.
Original meaning is PL for -*dan* and -*sil* and DIM for -*nem* and -*shoon*. Novel meaning is DIM.PL. Wrong meaning is the meaning each suffix was not paired with during exposure.

Across the two experiments, there was a significant preference for using the diminutive suffix to express DIM.PL, suggesting that participants preferred to profile the diminutive feature ($\beta = -0.448$, $z = -2.996$, $p = .0027$; Figure 6.11). However, there was no effect of Language on the choice of the DIM suffix -*nem*$_{DIM}$ vs. the PL suffix -*dan*$_{PL}$ ($\beta = 0.006$, $z = 0.021$, $p = .984$), indicating that form or feature frequency did not have a significant influence on the construal of DIM.PL. It is not clear why DIM appears to have been more salient to participants in Experiment IV, while PL was more salient to

participants in Experiments I–III. One possible post-hoc explanation is that DIM might have been reduced in salience in Experiments I–III compared to Experiment IV because it was redundant with SG during training. English-speaking participants therefore may not have noticed that it was also important. In Experiment IV, DIM occurs with both singular and plural referents, making it more likely to be identified as criterial. DIM therefore might be less likely to be noticed than PL, but more salient once noticed. I return to this issue in Chapter VII.

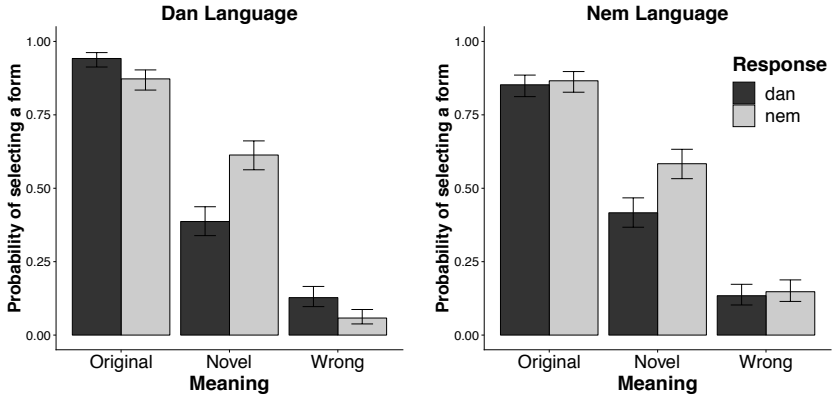

Figure 6.11. Results of the form choice task in Experiment IV.
Original meaning is PL for -*dan*$_{PL}$ and DIM for -*nem*$_{DIM}$. Novel meaning is DIM.PL. Wrong meaning is DIM for -*dan*$_{PL}$ and PL for -*nem*$_{DIM}$.

6.4 Discussion

This chapter reported on an experiment in which the order of the tasks replicated Experiment I, but the suffixes whose frequencies were manipulated (-*dan*$_{PL}$ and -*nem*$_{DIM}$), occurred with the DIM.PL meaning during training, as well as with the meanings they co-occurred with in Experiment I. Unlike in Experiment I, the language presented to participants in Experiment IV has no exact synonyms, allows for pre-emption to occur in

the DIM.PL meaning, and makes the DIM.PL meaning not novel. This change means that, in this experiment, training was consistent with these suffixes having the meanings PL and DIM rather than BIG.PL and DIM.SG. Increasing frequency of one of these forms' occurrence in the BIG.PL and DIM.SG meanings (the "original" meanings in Experiment I) now means learning that DIM.PL is not the prototypical / most common meaning of the form. Despite the differences between the experiments, the qualitative results of Experiment I were replicated, indicating that they are robust to the changes in design.

As in Experiment I, forms were more likely to be used to express a meaning when they appeared frequently, either in that meaning or in a meaning that shared semantic features with it. Forms did not entrench to their semantic prototypes in production: increased frequency in the BIG.PL meaning made -dan_{PL} more likely to be selected to express both the BIG.PL and the DIM.PL meanings. Similarly, increased frequency in the DIM.SG meaning made -nem_{DIM} more likely to be selected to express both DIM.SG and DIM.PL. These results therefore replicate the effect of frequency on form selection in production observed in Experiment I. As in Experiment I, the effect of frequency on form selection disappears when frequent and infrequent competitor forms are made equally accessible, by being presented to the participant as response options in a two-alternative forced choice task. Thus, the effect of frequency on form selection is mediated by its effect on form accessibility.

As in Experiment I, the form participants tend to use to express DIM.PL in production is not the form they map onto DIM.PL in comprehension. Whereas the frequent form is used to express DIM.PL, the infrequent form that has co-occurred with

DIM.PL and the semantic competitor of the frequent form are much more likely to be mapped onto DIM.PL in comprehension. For the competitor of the frequent form, DIM.PL in fact appears to be the semantic prototype, being chosen more frequently than the meaning with which the form was explicitly paired in training. Unlike in Experiment 1, this result cannot be due to the novelty of the DIM.PL meaning. Rather, the most likely interpretation is that a form abandons the semantic space that constitutes the semantic prototype of a frequent form.

The effect of frequency on the choice of form to express DIM.PL in production is smaller in Experiment IV than in Experiment I ($\beta = -4.4816$, $z = -2.697$ $p = .007$). This difference comes from the fact that an infrequent -dan_{PL} or -nem_{DIM} is more likely to be chosen to express DIM.PL in production in Experiment IV compared to Experiment I. This is a direct consequence of the difference between the languages presented to participants in the two experiments. In Experiment I, no form is paired with DIM.PL in training. In Experiment IV, -dan_{PL} and -nem_{DIM} are paired with DIM.PL during training even when they are infrequent. Therefore, even when they are infrequent, they are still relatively likely to be chosen to express the DIM.PL meaning, at the expense of the competitor suffixes -sil_{PL} and -$shoon_{\text{DIM}}$, which are never paired with DIM.PL in either experiment.

The implication of this difference for language change is that a competitor of a frequent form can survive only if its semantic niche is not subsumed completely within the semantic niche of a frequent form. If the frequent form occupies all of the space with which an infrequent form is associated, the infrequent form has 'nowhere to run', and is pushed out of existence (see also MacWhinney, 1987). While the comprehension and

forced form choice results of Experiment IV are almost identical to those of Experiment I, the production results differ in that the competitors of frequent forms are chosen for production very rarely in Experiment IV. In contrast, in Experiment I, these forms are frequently chosen to express the novel DIM.PL meaning that was not presented during training. A logistic regression model testing the effect of Experiment on the mapping of the main suffixes versus their competitors confirms this observation ($\beta = -4.409$, $z = -5.796$, $p < .0001$). In Experiment I, entrenchment leads the frequent form to abandon the DIM.PL meaning with which it does not co-occur. Its infrequent competitor does not entrench, remaining associated with either DIM or PL, and is therefore the form most compatible with DIM.PL. When DIM.PL is to be expressed, it can therefore compete with the frequent form within this niche. In Experiment IV, the frequent form is sometimes paired with the DIM.PL meaning during training, making it more difficult for the competitor to survive.An example of a frequent form driving an infrequent competitor into obsolescence is presented by Russian adversative conjunctions as described in Kapatsinski (2009); cf. the English *but* and *however*. These conjunctions are associated with preventive and denial of expectation meanings. The preventive meaning is illustrated by *He was going to flee the country but was detained at the airport*. The denial of expectation meaning is illustrated by *He was going to flee the country but didn't even think of saying goodbye*. According to Kapatsinski, in Russian, three conjunctions, *da*, *no* and *odnako* compete to express these meanings, with *no* being the most frequent. Kapatsinski shows that *da* is a better cue to the preventive meaning than *no* but, because of its high frequency, *no* is more likely than *da* to be chosen for production when the

preventive meaning is to be expressed. Accordingly, *da* is losing productivity, yielding to *no*.

While the present experiment replicated the qualitative pattern of results in Experiment I, the language in Experiment I maps onto the diachronic scenario of extension of familiar forms to novel uses more clearly. The production data in Experiment IV do not provide evidence for extension of a form to a *novel* meaning. It is only when the meaning is novel that we can speak of extension occurring. However, the similarity in the results between the two experiments suggests that the same mechanisms affect form selection in a novel meaning and in a familiar meaning.

The production data for the main suffixes -*dan*$_{\text{PL}}$ and -*nem*$_{\text{DIM}}$ show that shifting the prototype away from a meaning does not make the form less accessible in that meaning. As long as the frequency of a form is growing, its accessibility grows across its semantic range. As argued by Clark (1987), languages might abhor synonymy, but they do not mind polysemy (see also Goldberg, 2006, 2019). When a form develops a new meaning, and grows frequent in that meaning, rendering the old meaning non-prototypical, the form does not become less accessible when the newly peripheral meaning is to be expressed. Because of this, frequent forms accumulate meanings, becoming increasingly polysemous (Zipf, 1949), taking over semantic space rather than merely floating across it (Goldberg, 2019).

Just as production data of Experiment IV do not show extension, the comprehension data of Experiment IV do not necessarily show entrenchment. Rather, they are compatible with matching the probabilities of the various meanings conditional on the presented form. Thus, BIG.PL and DIM.PL are chosen equally often in response to

an infrequent *-dan*, whereas BIG.PL is chosen four times more often than DIM.PL in response to a frequent *-dan*. When *-dan*$_{PL}$ is infrequent, it is paired with BIG.PL and DIM.PL equally often in training. When it is frequent, it is paired with BIG.PL four times more often than with DIM.PL. Participants match these probabilities well.

Interestingly, probability matching is not observed with form choice. When participants face a forced form choice task, they appear to choose among forms compatible with a meaning randomly, without preference for the more frequent form. When they face an open-set form choice task, i.e., production, their responses fall between probability matching and maximizing. When two forms are equally probable given the meaning, the participants choose each about 50% of the time. However, when one of the forms is four times more probable than the other, it is chosen for production more often than its probability given the meaning to be expressed would suggest. I return to the question of probability matching in the general discussion. For now, I note that this is not a unique outlier, as similar results are obtained for form choice by Schumacher and Pierrehumbert (2017), who observe that response probabilities are attracted to 50% and 100% in a forced form choice task, rather than indicating probability matching.

Chapter VII

DISCUSSION

7.1 Summary of the findings

The results of Experiments I and IV demonstrated that frequent forms are more likely, compared to infrequent forms, to be extended to related meanings during online production, whether the meaning is novel as in Experiment I or familiar, as in Experiment IV. The effect of frequency on form choice is mediated by its effect on form accessibility. Under the pressure of real time processing to produce a meaning, speakers resort to forms that are easily accessible given the semantic features of the meaning they wish to express. These forms may or may not have been paired with the exact meaning to be expressed in the speaker's experience in the past: having co-occurred with other meanings sharing features with the meaning to be expressed is sufficient.

Forms compete for selection. Therefore, the synonyms of frequent forms are especially unlikely to be selected for production when expressing a meaning in which the frequent form is highly accessible. Interestingly, this is not a zero-sum game, so increasing the frequency of a form in a particular meaning makes the semantic features of that meaning more likely to be profiled. Thus, when *-dan*$_{\text{PL}}$ is frequent, it is more likely to be selected for production than when it is infrequent; its synonym, *-sil*$_{\text{PL}}$ is less likely to be selected for production compared to when it competes with an infrequent *-dan*$_{\text{PL}}$; but a PL suffix is more likely to be selected to express the DIM.PL meaning when one of the PL suffixes is frequent. That is, the frequency of a form–meaning mapping makes the meaning more likely to be expressed, representing a "thinking-for-speaking" effect (Slobin, 1987). If expressed, it is also most likely to be expressed using the frequent

form. Both effects of form frequency disappear in a binary forced form choice task where the alternative forms to be selected amongst are made equally accessible, indicating that the effects of frequency on form choice are mediated by its effect on form accessibility.

Despite being more accessible in production, and therefore more likely to be extended to similar novel uses, frequent forms entrench in the meanings with which they are experienced. The locus of this process is in comprehension, where upon hearing a form, the speaker maps it onto a meaning. As speakers repeatedly experience a form in a particular meaning, they become more confident that it does not map onto any other meanings or, at least, that the familiar meaning corresponds to the form's semantic prototype. The diachronic result of this process is *narrowing to the prototype* (Blank, 1999), which occurs when the entrenched form–meaning mappings learned through comprehension experience are entrenched enough to be the main determinant of form accessibility in production.

There is no contradiction between extension of frequent forms to new uses in production and entrenchment of the same forms in comprehension. Zipf (1949) likened words to tools, and the lexicon to a workbench, on which tools are arranged in descending order of frequency, so that frequently used tools are most accessible. An artisan sits by the workbench, deciding which tool to use for any particular task. Faced with a new task and an array of tools that have been used to perform similar tasks in the past, the artisan is likely to reach for a frequently used and therefore highly accessible tool. Note that this can happen even if the artisan is confident that the frequently used tool has never been used to perform the new task, and is less confident regarding the

previous uses of less accessible tools. In Chapter III, I have shown that this prediction holds for a wide array of Hebbian models of learning form–meaning associations.

Pre-emption pushes the competitors of a frequent form out of the area of semantic space corresponding to the form's prototype. Pre-emption shifts the infrequent competitor's semantic prototype away from the prototype of the frequent form, even if its only experienced uses correspond to the frequent form's prototype. Thus, when -*sil*$_{\text{PL}}$ competes with a frequent -*dan*$_{\text{PL}}$, a listener hearing -*dan*$_{\text{PL}}$ is likely to map onto the BIG.PL meaning rather than mapping it equally onto any PL meaning. This is the (negative) entrenchment effect. At the same time, -*sil*$_{\text{PL}}$ is likely to be mapped onto the DIM.PL meaning, even though it only co-occurs with the BIG.PL meaning in training. The reason that it *can* be mapped onto the DIM.PL meaning is because that meaning does not correspond to the frequent competitor's semantic prototype. This is the statistical pre-emption effect in comprehension. Together, the two mechanisms implement Clark's (1987) Principle of Contrast, leading languages to avoid exact synonymy. However, prototype shift does not ensure the survival of a competitor of a frequent form. As shown by Experiment IV, when the shifted prototype is still within the semantic range of the frequent form, the competitor is unlikely to be selected to express its prototypical meaning in production. It is only when the shifted prototype is outside of the frequent form's semantic range that the shift allows the competitor to survive in a specific semantic niche.

The processes of semantic extension and prototype shift resulting from entrenchment and pre-emption can result in chain shifts in semantic space (Aronoff, 2016; MacWhinney, 1987, 1993; see also Lindsay & Aronoff, 2013), in the following

way. Suppose that a form is extended to a new meaning. For example, *going to* is extended to mean a volitional future. As the frequency of *going to* in the volitional future increases, it pre-empts the older form *will* and therefore shifts its semantic prototype from volitional future to non-volitional future. When *going to* is then extended to non-volitional future as well, *will* is forced to shift further or face obsolescence. This shift is in progress in many dialects of English (e.g, Torres Cacoullos & Walker, 2009; see also Aaron, 2006, for Spanish).

7.2 The causes and consequences of extension[8]

Semantic extension has been a major recent focus of attention in grammaticalization theory. In grammaticalization, lexical items increase in frequency and are extended to more and more contexts as their meanings broaden while forms reduce and fuse with the forms around them. Although at some point in this process, the item is labeled "grammatical", the same processes occur before and after this point (e.g. Bybee, 2003; Bybee et al., 1994). Much research has focused on elucidating the relationships between these co-occurring processes, with little agreement. Observational diachronic data rarely allows one to identify causal relationships (though see Moscoso del Prado Martin & Brendel, 2016, for an exception). A primary contribution of the present dissertation is to evaluate a causal link from frequency of use to semantic extension, unmediated by semantic broadening.

While proposed as early as Zipf (1949), this causal link has not been seriously considered in recent theorizing on semantic change and grammaticalization. Rather, the

[8] Parts of this section have previously been published in Harmon, Z., & Kapatsinski, V. (2017). Putting old tools to novel uses: The role of form accessibility in semantic extension. *Cognitive Psychology*, *98*, 22–44.

debate has focused on whether frequency increases can directly cause semantic broadening—the loss of semantic specificity also known as bleaching or desemanticization—or are instead caused by it. In particular, Haspelmath (1999, p.1062) considers semantic broadening a prerequisite to extending the form to new uses, which in turn causes an increase in the frequency of a form. In contrast, Schuchardt (1885, p. 27) proposed that repetition by itself can lead to semantic bleaching: "If I say *g'Morgen* for *guten Morgen*, the adjective is deprived almost completely of its meaning, but only in consequence of the incessant repetitions".

While Schuchardt (1885) did not specify the mechanism by which frequency causes broadening, Bybee (2003) has argued that this mechanism is habituation. In habituation, as a stimulus is repeated and grows less and less surprising, it loses its ability to elicit an associated response: A shock becomes less and less shocking; A piece of chow less and less tempting; And a construction less and less prone to activate its semantic associations. Bybee exemplifies this process by the phrase *reiterate again*, as in *Mr. Speaker, let me reiterate again my great respect for the chairman, the gentleman from Illinois* (Blagojevich, 1999). According to Bybee (2003), *iterate*, which once expressed the meaning of "repeatedly" well enough on its own, has evolved into *reiterate* and then *reiterate again* as its ability to evoke repetition waned. Habituation results in semantic bleaching, a loss of specificity, rendering the bleached form suitable for a greater variety of contexts.

While Bybee (2003) and Haspelmath (1999) disagree on the causal relationship between frequency and semantic breadth, they agree that it is the increase in semantic breadth that causes the grammaticalizing form to be used in new contexts. According to

this view, grammaticalizing forms change in meaning by losing semantic features, whether through abrupt reanalysis or gradual weakening of form–meaning associations (see Campbell, 2000, p. 118−121, for a review). In contrast, Zipf (1949) saw grammaticalization as an extension of the frequent form to specific new uses—particularly uses the speaker has not associated with an alternative expression. Similarly, Heine (2011, p. 697) suggests that grammaticalization usually starts with extension, and it is extension that in turn triggers semantic broadening (see also Traugott & Trousdale, 2013). The present results are in agreement with this proposal, and provide a specific cognitive explanation for this ubiquitous trajectory of language change.

Bear in mind that the locus of habituation is in comprehension, where forms serve as cues to meanings. Habituation to a frequent form may make that form less likely to evoke its semantic outcomes, leading to a weakening of the form–meaning association. However, the comprehension results of Experiments I–IV consistently show that high frequency of a form–meaning pairing—at least in the range of frequencies examined in the current work—leads to a *strengthening* of form–meaning mappings.

In Experiments I–III, the participants experience suffixes paired with either multiple large creatures or a single small creature. This experience is intentionally ambiguous: participants could take the suffixes paired with multiple large creatures to be either plural or, specifically, non-diminutive plural. Similarly, they could take the suffixes paired with a single small creature to be either diminutive or specifically singular diminutive. If high frequency broadened a form's meaning, then frequent forms would be more likely to have single-feature meanings (DIM and PL rather than DIM.SG and BIG.PL). As a result, they would be more likely to be chosen to express the DIM.PL

meaning when all forms are equally accessible, i.e. in the form choice task. However, Experiments I and III provide strong positive evidence in favor of the null hypothesis that frequency does not influence form choice in this task.[9] These findings suggest that the participants believe that all four forms can be used to express the novel DIM.PL meaning, but are likely to choose the frequent form to express the meaning when the other forms are relatively inaccessible (see also Gershkoff-Stowe & Smith, 1997; Gershkoff-Stowe, 2001; Naigles & Gelman, 1995).

The results of the comprehension task are the opposite of those predicted by frequency-driven semantic broadening. Participants are much more likely to judge that the DIM.PL meaning is an appropriate meaning for a particular form when that form is infrequent in training. In particular, participants are about equally likely to click on DIM.PL and DIM.SG when presented with an infrequent -*nem*$_{\text{DIM}}$ in the comprehension task, suggesting that it is interpreted as simply DIM rather than DIM.SG. When the same form is frequent during training, participants are highly unlikely to click on the novel DIM.PL meaning, restricting their clicks to the specific meaning that has been paired with the form in training. We believe that the additional exposures to -*nem*$_{\text{DIM}}$ paired with DIM.SG referents increase the strength of the connection between -*nem*$_{\text{DIM}}$ and SG, making it less likely for participants to click on the DIM.PL referent when faced with that form. If frequency led to semantic broadening, we would expect exactly the opposite pattern of results.

[9] The effect of frequency on form choice in Experiment II is in the opposite direction from that predicted by frequency-driven semantic broadening. However, it is likely caused by transfer of form–meaning mappings from the preceding comprehension test and therefore does not provide strong evidence against this hypothesis.

The finding that participants are more likely to map frequent forms onto their original meanings is not a meaning frequency effect, but an effect of the frequency of the form–meaning pairing. If participants simply preferred to click on frequent meanings, then those meanings would be favored no matter the form cue. However, increasing the frequency of a form–meaning pairing *decreases* the likelihood that the form's synonym is mapped onto that meaning: -*sil*$_{\text{PL}}$ is mapped onto BIG.PL less often when -*dan*$_{\text{PL}}$ is frequent; and -*shoon*$_{\text{DIM}}$ is mapped onto DIM.SG less often when -*nem*$_{\text{DIM}}$ is frequent in the comprehension test. While frequent forms are more likely to be mapped onto their original meanings in comprehension, the synonyms of frequent forms (-*sil*$_{\text{PL}}$ in Dan and -*shoon*$_{\text{DIM}}$ in Nem) are more likely to be mapped onto the novel meaning than synonyms of infrequent forms (-*sil*$_{\text{PL}}$ in Nem and -*shoon*$_{\text{DIM}}$ in Dan).

The results suggest that a speaker's use of a frequent form to express a novel meaning may not reflect the speaker's beliefs about what forms can be used to express the meaning, or which forms best express that meaning. It is driven simply by differences in accessibility between forms activated by the set of semantic features that the speaker wishes to express. Several results in first language acquisition are also consistent with this idea and not explainable by habituation to frequent forms. For example, children may be well aware that a cow is not a dog and nonetheless use the form *doggy* to refer to cows when *cow* is not accessible (Gershkoff-Stowe, 2001; Gershkoff-Stowe & Smith, 1997; Kuczaj, 1982; Naigles & Gelman, 1995; Nelson et al., 1978). These results suggest that frequent forms are extended to novel uses because they are more accessible than their competitors (Zipf, 1949) rather than being bleached of semantic content.

Indeed, Bybee's (2003) example of *iterate* > *reiterate* > *reiterate again* appears unlikely to be driven by high frequency of *iterate* and then *reiterate* causing semantic broadening: neither form is particularly frequent. Rather, this may well be a case of extension, where the highly accessible constructions *re*-VERB and VERB *again* were extended to a lexical context in which they were previously unacceptable (and remain unacceptable for many). The meaning of REPETITION is strongly associated with the form *again*. This association makes *again* highly accessible from the meaning of REPETITION and triggers its production, even if there is no need for it, i.e. even if the form *iterate* still activates the meaning of REPETITION in comprehension. *Reiterate again* is then conventionalized. Once *again* is consistently produced with *reiterate*, *reiterate* can lose its association with repetition through cue competition, because it is overshadowed by a stronger cue (Pavlov, 1927; see also Harmon & Kapatsinski, 2016). A similar account can also be offered for the development of *irregardless*, by extension of *ir-* to a new stem because of its strong association with NEGATION.

By using the frequent form in a new context, one in which it has never been experienced, the speaker extends the form's range of uses. In this way, semantic extension can occur without semantic broadening. Yet, semantic broadening is sure to follow if the extension catches on and persists until the language is experienced by a new generation of learners. The extension of -*nem*$_{\text{DIM}}$ to DIM.PL would provide a learner of the language with positive evidence that -*nem*$_{\text{DIM}}$ does not mean specifically DIM.SG. Instead, its range of uses in the learner's experience would be fully consistent with a general diminutive marker. When that learner uses -*nem*$_{\text{DIM}}$ to refer to a DIM.PL referent, it would be with the belief that -*nem*$_{\text{DIM}}$ is a general diminutive marker. As previously

argued by Bybee (2010), in language change, **use leads and belief follows**. For example, while participants in grammaticality judgment tasks dislike *I'm gonna disappear it* more than they dislike *I'm gonna vanish it* (e.g., Theakston, 2004), it is *disappear* and not *vanish* that is being used transitively in production, as in the following examples: *Jamal Khashoggi was on the cusp of starting a democracy group to provide "a counter narrative in the Arab world and the West to Arab Spring skeptics" when the Saudis disappeared him.* (Urban Dictionary), or *Jamal Khashoggi Wanted to Launch a Pro-Democracy Group. Then the Saudis Disappeared Him* (The Daily Beast).

This is not to say that accessibility in production is the only route to semantic extension. Indeed, habituation seems to be a plausible explanation for certain developments, such as increasing acceptability of expletives (Jay, 1992), borrowings, or discourse particles, i.e. the gradual loss of markedness. However, I would like to suggest that accessibility-driven extension in production often feeds into semantic broadening, including the broadening we see in grammaticalization.

One important advantage of an extension account of grammaticalization is that it allows for grammatical forms to acquire a broad range of specific uses that cannot be described with a conjunction of necessary and sufficient semantic features and therefore cannot result from feature loss. Often, the range of uses of a grammatical morpheme is best described as a radial category, a network of family resemblances (Lakoff, 1987).

Plaster and Polinsky (2011) revisit Lakoff's (1987) classic example of a classifier in the Australian language Dyirbal was extended from use with human females to many other referents, resulting in the apparent class of *women, fire and dangerous things* (Lakoff, 1987). Plaster and Polinsky show that this change did not come about because

Dyirbal speakers have over time discovered semantic similarities between women, fire and dangerous things. Rather, many of the extensions were based on phonological rather than semantic similarities between nouns co-occurring with the classifier and novel nouns. For example, all animate nouns beginning with the syllable *yi* took on the classifier, presumably because the same syllable begins the word *yibi* ~ WOMAN. This example shows that extension of a form to use in a new context occurs because the new context shares features with older contexts in which the form occurs. These features do not have to be semantic, but rather simply need to be effective cues for the form in question to be retrieved and used in a new context.

To summarize, the meanings of linguistic forms are in constant flux, both on the timescale of language acquisition and the timescale of language change. As they change, meanings of individual forms can both broaden—becoming more general—and narrow—becoming more specific. I argue that broadening results from semantic extension, which can itself be driven by frequency. Extension of a form to a new use occurs when a speaker selects that form for the new use over a less accessible semantic competitor, despite never having heard the word used this way. These kinds of extensions are common in child language (Gershkoff-Stowe, 2001, 2002; Gershkoff-Stowe & Smith, 1997; Naigles & Gelman, 1995), where they are usually called *over*-extensions. Every instance of production has the potential to produce an extension, and every extension observable to a listener has the potential to seed language change (see also Bybee, 2010, and MacDonald, 2013, for related arguments). If accepted and later re-produced by a listener, the extension ceases to be an *over*-extension. In this way, novel uses of forms can gradually spread through the speech community.

7.3 The causes and consequences of entrenchment[10]

As the frequency of a form–meaning mapping increases, learners become more and more likely to map it onto the original meaning and less and less likely to map it onto the novel meaning in comprehension. Yet, the form least likely to be mapped onto the novel meaning in comprehension is the form most likely to be used to express this meaning in production. As shown in Chapter III, this production–comprehension dissociation is an almost inevitable outcome of associative learning even if bidirectional form–meaning connections are assumed to be formed and used for both comprehension and production. Therefore, production–comprehension dissociations of the type documented here are virtually certain to arise in the course of language acquisition.

To briefly summarize the argument, suppose that every exposure to a form accompanied by a referent with certain semantic features strengthens the connection between the form and the semantic features of the referent. These connections will be stronger for a frequent form than for an infrequent form. Thus, a frequent -*nem*$_{\text{DIM}}$ will have stronger connections to DIM and SG than an infrequent -*nem*$_{\text{DIM}}$. The results of the production task immediately follow. Note that the novel meaning, DIM.PL, shares one feature with the original meaning of every suffix, DIM.SG and BIG.PL. Therefore, all suffixes are eligible to be extended to this meaning: they will all be activated by it in proportion to their frequency in training. If form activation drives form choice in production, an uncontroversial assumption (e.g. Dell, 1986), the speaker will be likely to select the frequent forms over the infrequent ones to express the novel meaning.

[10] Parts of this section have previously been published in Harmon, Z., & Kapatsinski, V. (2017). Putting old tools to novel uses: The role of form accessibility in semantic extension. *Cognitive Psychology*, *98*, 22–44.

The entrenching effect of frequency in comprehension follows as well. Suppose that selection of a meaning in response to a form depends on how strongly its features are activated when cued by the form. A frequent form will have stronger associations to the features that distinguish its original meaning from the novel meaning. For example, a frequent -*nem*$_{\text{DIM}}$ will activate SG much more strongly than an infrequent -*nem*$_{\text{DIM}}$, boosting the activation of DIM.SG over DIM.PL and therefore making it less likely that the speaker will respond by clicking on DIM.PL.

Thus, even simple Hebbian learning of bidirectional form–meaning associations can predict the dissociation between the frequency effects observed in production and comprehension following the exposure phase. Indeed, almost any learning model would predict it, suggesting that such dissociations will inevitably arise in the course of language learning. Accordingly, many studies have observed such dissociations in child language (e.g., Bloom, 1973; Gerschkoff-Stowe & Smith, 1997; Kuczaj, 1982; Naigles & Gelman, 1995; Nelson et al., 1978).

The production–comprehension dissociations observed in child language usually disappear by adulthood instead of successfully seeding language change, and there is widespread agreement that this disappearance is driven by comprehension experience (e.g. Braine & Brooks, 1995; Regier & Gahl, 2004; Stefanowitsch, 2008), which curbs over-extension in production just as it curbs many other innovations. For this to happen, entrenchment in comprehension must modify form–meaning mappings that are used in production.

In our comprehension task, participants are required to decide on one of four possible meanings when presented with a form. The meanings therefore constitute

competing responses to be selected using the form as a cue. In the absence of corrective feedback, the more one chooses a particular response given a cue, the more likely one becomes to choose the same response when presented with that cue in the future (a Zipfian positive feedback loop; e.g. Martin, 2007; Oppenheim, Dell & Schwarz, 2010; Zipf, 1949). Every time a participant maps a form onto a meaning in the comprehension task, they have to overcome competition from other possible meanings / alternative responses. Since the mapping is always successful, in the absence of error feedback, the participant may adjust their system of form–meaning connections to make that response easier to produce on future trials. This can be accomplished by developing inhibitory connections between the form in question and other meanings, or the semantic features that distinguish those meanings from the chosen one. We can think of this as the simplest possible reinforcement learning process (Sutton & Barto, 1998), one in which the chosen response is always reinforced.

Entrenchment makes meanings more specific. In Xu & Tenenbaum's (2007) study, many participants enter the experiment with the (implicit) belief that a novel term like *fep* is more likely to refer to the basic level category DOG than to the subordinate category, WHITE SPOTTED DOG (i.e. Dalmatian). In our experiment, when a form is infrequent, participants are quite likely to click on the DIM.PL meaning, with which the form has never been explicitly paired in training. These responses are not simply errors because participants rarely click on the wrong meaning, which shares no features with the meaning that was paired with a suffix in training. Rather, they suggest that participants likely enter the experiment with the belief that suffixes refer to DIM and PL, rather than DIM.SG and BIG.PL.

However, for the frequent form, the original meaning has an advantage: it is strongly activated by the form due to the strong form–meaning association developed during exposure. When -*nem*$_{\text{DIM}}$ is frequent, participants no longer click on DIM.PL as much as on DIM.SG, since -*nem*$_{\text{DIM}}$ has built up a strong connection to SG during training. As the participant repeatedly chooses the original meaning in response to the frequent form in the comprehension task, s/he develops inhibitory connections between that form and the features that distinguish the novel meaning from the original one. Thus, a frequent -*nem*$_{\text{DIM}}$ would develop a strong inhibitory connection to PL. In a subsequent production task, those inhibitory connections then prevent the speaker from activating the frequent form when cued with the novel meaning. Under this proposal, entrenchment in comprehension can constrain production, as long as the inhibitory form–meaning connections developed in comprehension can be re-used in production. This requires the connections in question to be bidirectional.

While traditionally assumed by linguists of all persuasions, bidirectionality of form–meaning mappings has recently been questioned by Ramscar et al. (2010), who argued that learners acquire only associations whose direction matches the flow of time, from present to future but not from present to past. In our study, referents preceded forms during the exposure phase, which means that the exposure phase should have resulted in meaning-to-form and not form-to-meaning associations. These associations could be used for accurate production, which requires retrieving a form based on semantic cues, but not for comprehension. However, this hypothesis is contradicted by the fact that comprehension was highly accurate in our experiment, even when it immediately followed exposure: the plural suffixes were seldom mapped onto non-plural referents and

diminutive suffixes were almost never mapped onto non-diminutive ones. We should note that adults in Ramscar et al,'s study were required to perform mathematical operations upon presentation of the second stimulus in an exposure trial (form or referent), preventing them from attempting to retrodict the first stimulus and form a meaning-to-form association. This was not the case in the present experiment. We believe that, at least by the time we are adults, we do subconsciously attempt to form bidirectional form–meaning associations, though our efforts may be thwarted if attention is diverted (see also Arcediano, Escobar & Miller, 2005; Kahana, 2002; Matzel, Held & Miller, 1988; O'Reilly & Munakata, 2000; for other arguments for bidirectionality in associative frameworks).

An alternative to the proposed associative account of entrenchment is the Bayesian account of Regier & Gahl (2004) and Xu & Tenenbaum (2007). Under their proposal, entrenchment occurs because, as one experiences examples from a subcategory, it becomes increasingly unlikely that the experienced examples are a random sample from a larger category. In other words, the child thinks: "if *fep* means DOG, what are the chances that every time I see a dog it is a Dalmatian? Not very likely!" The probability of hearing *fep* referring to a Dalmatian once, given that *fep* means DOG, is .1 if 10% of dogs are Dalmatians. Not too improbable. But when Dalmatian is paired with *fep* three times, that probability is reduced to $.1^3$, or .001, exceedingly unlikely. Therefore, children restrict the form *fep* to the subordinate category in which it was experienced, Dalmatian, rather than extending it to all dogs (see also Regier & Gahl, 2004). In our study, as one experiences more and more -*nem*$_{\text{DIM}}$'s paired with single creatures, it becomes

increasingly unlikely that the experienced referents of -*nem*$_{\text{DIM}}$ are sampled from a population that contains equal numbers of singular and plural referents.

An advantage of the associative account is that it successfully predicts that the head-start enjoyed by the original meaning of the frequent form will grow into an insurmountable advantage through the comprehension test. The Bayesian account predicts that the comprehension test experience is irrelevant to entrenchment: it provides no information about the population from which the form's referents are sampled. Rather, the magnitude of the entrenchment effect should be determined entirely by frequency differences during exposure. In order to discriminate between these accounts, we examined the effect of trial during the comprehension test on the likelihood of mapping a suffix to the familiar vs. the novel meaning. As shown in Figure 4.5 (Chapter IV), form–meaning mappings are entrenched over the course of the comprehension test: frequent forms are more likely than infrequent forms to be mapped onto their original meaning from the start of the test, but this difference increases as the test progresses.

In the context of real-life language acquisition, the associative account demands that the learner predict and rule out the specific meanings to which a form is initially over-extended (e.g. DIM.PL). The comprehension test forces participants to discriminate between all four maximally specific meanings and should therefore be particularly conducive to entrenchment, However, we suspect that participants also engage in unprompted discriminative learning during passive exposure (e.g., Ramscar et al., 2010, 2013), and that therefore extended exposure without forced discrimination among the original and novel subordinate meanings would also eventually produce entrenchment to the experienced set of meanings.

7.4 Extension and entrenchment in continuous space

A limitation of the representation of semantics in the discussion above, and the computational modeling presented in Chapter III, is the assumption that semantic space is defined by a set of discrete dimensions, or features. It is also possible to conceptualize the semantic space as continuous, and some semantic dimensions are, perhaps, better conceptualized in this way; particularly those affective in nature (Osgood & Sebeok, 1954). In fact, Koranda, Zettersten, and MacDonald (2018) have demonstrated the effect of form accessibility on form choice in expressing directions within a continuous angle space. Furthermore, modern distributional semantic models treat the semantic space as composed exclusively of continuous dimensions (e.g., Landauer & Dumais, 1997). In this section, I show that entrenchment in comprehension is expected to coexist with extension of frequent forms in production even if we treat semantic space as continuous. In addition, I show that this dissociation can also be justified on the basis of normative Bayesian inference, which is well suited to dealing with distributions in a continuous space.

Kleinschmidt and Jaeger (2016) have argued that the selective adaptation effect in perceptual learning corresponds to entrenchment to the prototype: when participants are repeatedly presented with typical examples of a sound category, the category boundary shifts closer towards the category center. Here, I extend their reasoning to entrenchment in semantics, and show that shifting the prototype, as in Experiment IV, can also produce entrenchment.

Consider three semantic categories mapping onto three distinct forms in the unidimensional continuous semantic space represented in Figure 7.1 below. In a

continuous semantic space, forms correspond to distributions of meanings (in the same way that they correspond to distributions of sounds in the phonetic space). The lines in Figure 7.1 represent probability densities of the distributions corresponding to semantic categories. The categories represented by solid lines correspond to infrequent forms, while the category represented by the dashed black line corresponds to the frequent form. The peak of the density distribution for a semantic category associated with a form corresponds to its prototype, the point in semantic space most strongly activated by the form (see also Buz, Tanenhaus & Jaeger, 2016, for phonetic space).

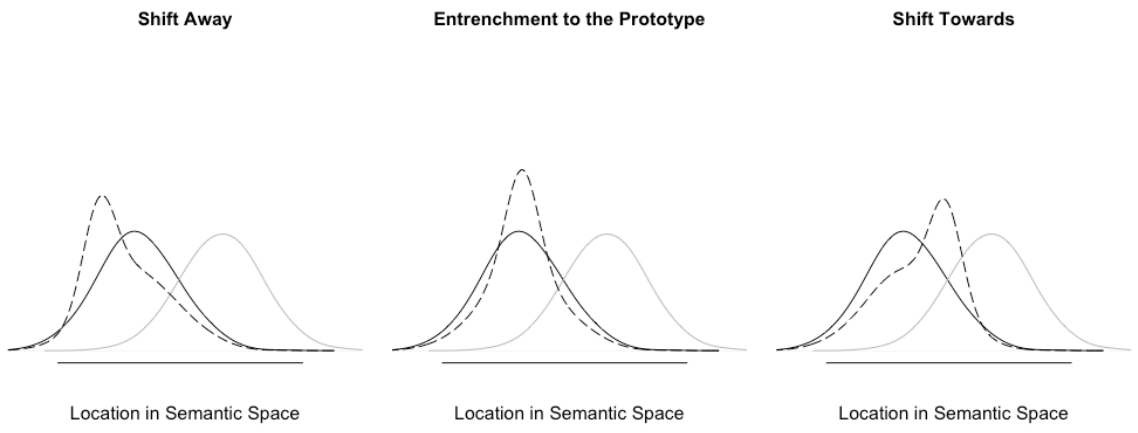

Figure 7.1. The effect of entrenchment in continuous space.

Experiments I–III are most closely represented by the middle panel of Figure 7.1, in which a frequent form is presented with additional examples of the meaning that constitutes its semantic prototype. Experiment IV is most closely represented by the left panel, where the additional presentations of a frequent form occur together with a meaning that is farther away from the prototype of the competing form than the prototype of the same form without the additional presentations. In the simulations underlying

Figure 7.1, the two low-frequency distributions are samples from normal distributions, while the high-frequency distribution is augmented with an additional set of tokens that correspond to the mean of the distribution in the middle panel, or are one standard deviation below (left panel) or above (right panel) the mean. The right panel would correspond to a version of Experiment IV in which the 'extra' examples of a frequent form would be paired with DIM.PL rather than DIM.SG or BIG.PL in training.

Figure 7.1 shows entrenchment in continuous space: the curve for the frequent form is below the curve for the (near-)synonymous infrequent form in regions of the space that are somewhat removed from the frequent form's prototype. Furthermore, entrenchment occurs whether additional examples reinforce the existing prototype or shift it.

The shift away from the competing category effected in Experiment IV does not have a clear equivalent in the experimental literature on sound categorization. However, in all experiments, participants are expected to activate the region of semantic space shared by the gray and black categories (Figure 7.1) less when presented with a frequent form (dashed line) than when presented with an infrequent form (solid lines). The only situation in which this prediction would not be made is if the additional examples that make a form frequent were presented with the intermediate, DIM.PL meaning.

While entrenchment in comprehension is expected, the probabilistic model in Figure 7.1 does not provide a way to model frequency-driven extension: in the left and middle panels of Figure 7.1, the dashed black line crosses the gray line further to the left than the solid black line does. Indeed, Kleinschmidt and Jaeger (2016) argue that entrenchment to the prototype narrows the region of the phonetic space that would be

mapped onto a form in a form choice task. In contrast, the present results indicate that, due to its higher accessibility, a frequent form is activated by off-prototype regions of semantic space more strongly than an infrequent form. This influence of accessibility is missing from Figure 7.1 because the area under each curve sums to 1.

Areas under the curve in Figure 7.1 are conditioned on the form: probabilities of all meanings given a form sum to 1. These probabilities of meanings given forms are useful for modeling comprehension, where the meaning is chosen and the form is provided. For production and form choice tasks, we instead need probabilities of forms given meanings. In order to derive the probability of a form given a meaning in a Bayesian model, these probabilities need to be multiplied by unconditioned form probabilities, i.e., their frequencies: according to Bayes' theorem, $p(\text{form} \mid \text{meaning}) \sim p(\text{meaning} \mid \text{form}) \times p(\text{form})$.

The consequences are shown in Figure 7.2 below. The curves corresponding to frequent forms rise above the curves corresponding to infrequent forms, predicting that frequent forms will be extended in production even as they entrench in comprehension. The dissociation between the effects of frequency in production and comprehension is therefore predicted not only by Hebbian associative models learning associations between forms and discrete semantic features, but also by Bayesian models mapping forms onto distributions in a continuous semantic space.

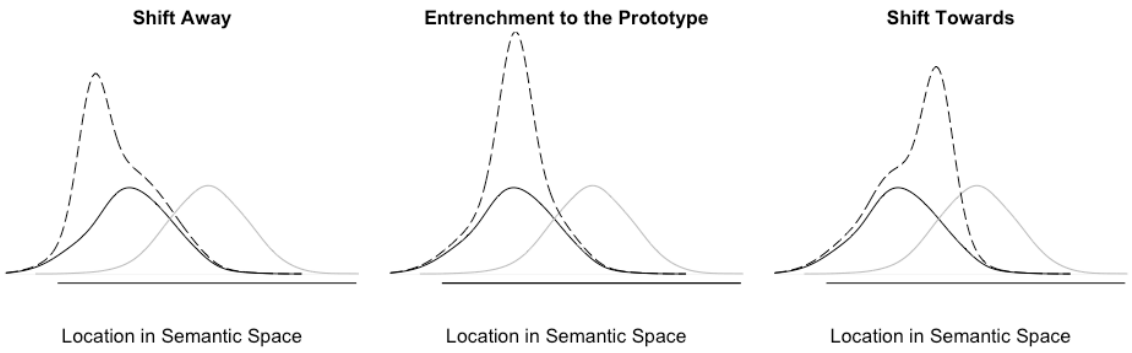

Figure 7.2. Extension of frequent forms in continuous space via Bayesian reasoning.

7.5 Production vs. forced choice

7.5.1 Bayesian vs. associative models

As pointed out by Norris and McQueen (Norris, 2006; Norris & McQueen, 2008), a major difference between associative/connectionist and Bayesian models of cognition is that the Bayesian mind deals in probability while the associative mind deals in activation. Norris has argued that this is an important advantage of Bayesian models because activation levels need to be somehow transformed into response probabilities, whereas Bayesian models directly predict response probabilities. However, form activations can be manipulated without manipulating form probabilities (Luce & Pisoni, 1998). In particular, when response choices are provided, response probabilities are much closer to random, compared to an open-set task in which responses have to be retrieved from memory or generated. Thus, *p*(form) appears not to have much of an effect on form choice in the forced form choice task in Experiments I–IV. This kind of ignorance of base rates is also observed in many other judgment tasks (Kahneman, 2011). From an associative perspective, bottom-up activation of the forms causes their activation levels in the forced choice task to be near ceiling, leaving little room for an influence of the form's

resting activation levels representing the forms' base rates (Luce & Pisoni, 1998) or of the strengths of associations connecting the presented meaning and the alternative forms. Thus, forms consistent with a meaning are chosen with equal probability in the forced choice task despite the frequent form having a higher probability given the meaning, as well as a higher probability overall. It is difficult to see how these task effects can be accounted for in a probabilistic framework.

7.5.2 Accessibility matters: Production tasks as a window on the grammar

Several researchers have used the results in Chapter III to argue that forced choice tasks provide a better window on the learned grammar than production tasks by reducing task demands (Ambridge et al., 2018; Goldberg, 2019; Schwab et al., 2018). I disagree with this interpretation. Forced choice tasks are not necessarily easier than production or open-set tasks. While it is easier to generate an acceptable response in the forced-choice task, it is harder to be faithful to learned probabilities in one's response choices, by either matching or maximizing the probabilities one has acquired.

Forced choice tasks reduce accessibility differences between probable and improbable responses. According to associative models of the mind, participants do not have direct access to response probabilities, and can only access them via activation differences (Luce & Pisoni, 1998). By leveling these differences, forced choice tasks often produce results that are more random and less reflective of the learned grammar than responses in production (see also Olejarczuk, 2018, for another example).

7.6 Why there is no entrenchment within the production test

I have argued that the comprehension test encourages settling on a system of mutually exclusive form–meaning mappings because it requires participants to decide

between the original meanings and the novel meaning. Note that the production task also requires participants to make a decision; in that case, choosing a form to produce. However, a mutually-exclusive system of form–meaning mappings does not appear to arise in the course of the production test.

One possible explanation for the lack of entrenchment in production is that participants in a form choice task consider that they need to reproduce the form variation in the input, i.e. to match probabilities of forms given each meaning. Adult participants in form choice tasks often attempt to match input probabilities; e.g., after training on an artificial language (Hudson Kam & Newport, 2005; Kapatsinski, 2010, 2013; Perfors, 2016; Wonnacott et al., 2008), after training on a simple non-linguistic pattern (Ferdinand et al., 2013) and when choosing a native-language pattern to apply to novel words (e.g. Coetzee & Pater, 2011; Kapatsinski, 2010). Sociolinguists have long argued that matching the variation in the input is part of sounding like a native speaker of a particular language variety (Labov, 1969; see Coetzee & Pater, 2011, for a review). Recently, Perfors (2016) has also shown that participants in a miniature artificial language learning experiment are more likely to regularize the input if they are convinced that they do not need to match the variation.

Note that individual participants in the production task were largely unable to *actually* match probabilities in Experiments I–IV. Participants choose the frequent form over a synonymous infrequent form more frequently than the ratio of their probabilities would suggest. Individual subjects data show that a large proportion of participants maximize probabilities while others display 50/50 responding. These results are consistent with those reported by Schumacher and Pierrehumbert (2017) and Schwab et

al. (2018). However, in unpublished work, I have asked participants to report their strategy in a simple binary form choice task, and most reported attempting to match the frequencies of the two forms used in the experiment. It is therefore likely that participants in form choice tasks try to match the variation in the productions they have experienced rather than consistently producing a single form for a given meaning, but task demands can prevent them from accomplishing this goal. This interpretation is supported by the results of Ferdinand et al. (2013), who presented their participants with one or more urns emitting colored balls, different urns being associated with different color probabilities. They found that participants matched color probabilities when faced with only one urn, but maximized when faced with multiple urns, guessing each urn's most common ball color when presented with the urn. It is possible that the lexical retrieval difficulties associated with production encourage producing the most accessible form every time one is faced with form choice. That is, whereas the forced choice task encourages random choice, the production task encourages regularization; always choosing the most likely response. Despite this, participants know that more than one response is associated with every meaning, and do not appear to entrench on the responses they produce.

The production test may also not be as conducive to entrenchment as the comprehension test because the production responses in our experiments are inherently more variable than responses in the comprehension task. Even if a participant always chooses the same suffix for a particular meaning, the actual form of their response will vary from trial to trial, because every trial involves a different stem. Therefore, choosing the same suffix from trial to trial does not mean choosing the same motor program, whereas choosing the same meaning in the comprehension task does involve the same

motor program for moving the mouse to a particular quadrant of the screen and clicking on it. One may not be able to entrench on a particular type of response if the actual form of that response is very different. Under this hypothesis, entrenchment may occur as a result of production practice when the alternative production responses constitute alternative production plans, e.g. when they are forms that can stand on their own (words or utterances).

Finally, language learning may also be biased against entrenchment of form–meaning mappings based on production because such entrenchment is counter-productive. In language acquisition, one needs to settle on the form–meaning mappings of the community despite initially producing forms in reference to the meanings they do not have for others. It would not be a good idea for a speaker to settle on his own production choices too fast if they want to grow out of overextension. Settling on over-extended forms in production does appear to be common in second language acquisition (Harley & King, 1989; Treffers-Daller & Calude, 2015). Both first and second language learners face the danger of settling on an over-extension. However, unlike first language learners, second language learners have acquired many concepts prior to starting to learn their second language. Therefore, they often need to express a concept for which they do not yet have a word. In such cases, they may choose a more easily accessible form and are then in danger of settling on that form–meaning mapping unless perceptual input prevents it. While the same danger is present for first-language learners (Gershkoff-Stowe & Smith, 1997), the perceptual input to second language learners is more limited, and therefore less likely to be able to curb over-extensions. Furthermore, second language

learners are experts in their native language and therefore may implicitly trust their own productions more and comprehension input less than relatively inexperienced children.

7.7 Pre-emption and entrenchment[11]

Several authors have credited entrenchment with the disappearance of overextension errors (Braine & Brooks, 1995; Regier & Gahl, 2004; Stefanowitsch, 2008). According to these proposals, entrenchment produced by additional perceptual experience curbs overextension in production. Recently, Boyd and Goldberg (2011) have suggested that this developmental process can instead be attributed to statistical pre-emption, where use of a construction to express a particular meaning is pre-empted by experiencing another construction paired with that meaning. For example, *He disappeared it* may be pre-empted by experiencing *He made it disappear*, instead of being unacceptable due to the high frequency of intransitive *disappear*, as in *He disappeared*. Pre-emption is closely related to the idea of accessibility-driven extension. This is most evident in a closed set, where accessibility of a form guarantees the inaccessibility of its competitor, i.e., its potential pre-emptor.

Entrenchment is implicated in the participants' reluctance to map the frequent form onto a novel meaning. When a form is infrequent, it tends to be mapped onto both the specific meaning it was paired with (the 'original meaning') and the DIM.PL meaning, which shares a feature with the original meaning. When a form is frequent, it tends to be mapped only onto the specific meaning it was paired with in training, either DIM.SG or BIG.PL. On a pure pre-emption account, this must be because it has been

[11] Parts of this section have previously been published in Harmon, Z., & Kapatsinski, V. (2017). Putting old tools to novel uses: The role of form accessibility in semantic extension. *Cognitive Psychology*, *98*, 22–44.

pre-empted from mapping onto either the more general meaning (DIM for -*nem*$_{DIM}$) or the novel meaning (DIM.PL) by another form.

Nothing except its own high frequency can prevent a frequent -*nem*$_{DIM}$ from mapping onto the single feature DIM: its -*shoon*$_{DIM}$ competitor is just as infrequent in Nem as in Dan—where both competitors map onto DIM—and should therefore be just as impotent a pre-emptor. This suggests that its high frequency in DIM.SG caused -*nem*$_{DIM}$ to map onto DIM.SG rather than DIM.

A statistical pre-emption account may suggest that -*shoon*$_{DIM}$ maps onto DIM.PL before pre-empting -*nem*$_{DIM}$ from mapping onto it. The problem with this proposal is that the mapping of -*shoon*$_{DIM}$ onto DIM.PL must be caused by the high frequency of -*nem*$_{DIM}$ in training: -*shoon*$_{DIM}$ is equally infrequent, whether -*dan*$_{PL}$ or -*nem*$_{DIM}$ is the frequent form. Why would -*shoon*$_{DIM}$ map onto DIM.PL instead of mapping onto DIM, as it does when -*nem*$_{DIM}$ is infrequent? On a pre-emption account, the answer is that it must be pre-empted from mapping onto DIM by -*nem*$_{DIM}$. However, -*nem*$_{DIM}$ is *less* likely to map onto DIM when it is frequent, and therefore should be less likely to pre-empt -*shoon*$_{DIM}$ from mapping onto DIM. Entrenchment appears to be a necessary mechanism for constraining extension in the present data.

The fact that synonyms of frequent forms tend not to be mapped onto their original meanings in comprehension suggests but does not necessarily entail that the frequent form is pre-empting its infrequent synonym. Instead, this may be a meaning frequency effect. Ramscar, Dye & Klein (2013) show that children are reluctant to map forms onto frequently encountered referents that have not been associated with a form, by virtue of occurring on every trial of training independently of the form present on that

trial. Suppose that every time participants encounter -*nem*$_{DIM}$ paired with the DIM.SG meaning, the *shoon* ~ SG and –*shoon* ~ DIM associations weaken (e.g. van Hamme & Wasserman, 1994). In other words, the more often a meaning is encountered without a form in training, the more the form is dissociated from that meaning (see also McMurray, Horst & Samuelson, 2012). When a form has a frequent semantic competitor, the shared meaning often occurs without the form, becoming dissociated from the form. On this account, -*shoon*$_{DIM}$ is seldom mapped onto DIM.SG when -*nem*$_{DIM}$ is frequent not because it has a frequent synonym but because DIM.SG has frequently occurred without -*shoon*$_{DIM}$. Accordingly, if DIM.SG were to occur without any form whatsoever on the same number of trials, the same reluctance to map -*shoon*$_{DIM}$ onto it is expected to occur. We leave this prediction for future work to evaluate. For now, we note that pre-emption can explain this result but is not uniquely supported by it.

Pre-emption is suggested by the finding that frequent forms are less likely to be mapped onto the novel meaning when production follows comprehension (e.g., Experiment II) than when it precedes it, a scenario that could account for comprehension experience constraining overgeneralization errors in production during development. When the competitors are pushed out of their shared meaning by the entrenched frequent form, they map onto the novel meaning, creating a mutually exclusive system of form–meaning mappings. Importantly, this mapping is then used in production, where the newly mapped competitors pre-empt the frequent forms from mapping onto the novel meaning. For frequent forms to push their competitors out of the shared meaning and not out of DIM.PL, an entrenching effect of frequency is necessary. However, pre-emption

appears necessary to account for the difference between experiments in which production is first and those in which it follows comprehension.

7.8 The role of children and adults in semantic change

A major contribution of the present thesis is to show that extension is immanent in every instance of language use, and is not limited to children. Bybee (2010; esp. pp. 114−135) has argued extensively that language changes through use, and that use rather than imperfect acquisition by children is the major instigator of language change in general and extension in particular (see also Aitchison, 1991; Bybee & Slobin, 1982; Cheshire 1982; Kapatsinski, Easterday & Bybee, 2019; Slobin, 1994, 2002; Vihman 1980). She argues that pre-adolescent children are usually not in a social position to spread their innovations to others. Indeed, the documented cases of children driving language change are all cases in which adult linguistic norms are lacking because the adults are second-language speakers or late acquirers of the language variety the children are acquiring (see Hudson Kam & Newport, 2005, for a review). In speech communities of fully entrenched adult speakers, children have little opportunity to bend the community speech norms to their will. Accordingly, many childhood innovations never catch on to become part of adult grammar of a language. For example, major place consonant harmony (*take* > *cake*) is ubiquitous in child speech but has not been incorporated into the adult phonology of any language (Bybee, 2010, p. 115; Vihman, 1980).

Many semantic changes that are the focus of research in the literature on language change are also difficult to attribute to young children because they require significant sophistication (Slobin, 1994, 2002). In particular, in the process of grammaticalization, words tend to be extended from observable entities and events to mental, relational and

intersubjective concepts (Traugott, 1988). Thus, *going to* is extended from the physical act of walking to the mental act of intending to do something, to an expression of predictions about the future. Child over-extensions in contrast tend to involve extending terms to concrete objects to which they do not refer in adult language, e.g., extending *doggie* to refer to a cow (Clark, 1973). As Slobin (1994, pp. 129−130) writes, "New meanings of grammatical forms arise in adult language use on the basis of pragmatic inferences drawn from existing referential or propositional meanings. Preschool-age children are not yet able to draw most of these inferences." As an example, he shows that the more grammaticalized uses of tense/aspect markers are acquired late, precluding the possibility that they result from semantic extensions made by young children.

Dale and Lupyan (2011) have also argued that languages with more non-native speakers tend to have simpler morphology, and attribute this to morphological paradigm leveling during second language learning. I have argued in Chapter II that paradigm leveling results from extending a frequent form of a word to semantically related paradigm cells, which occurs in both child and adult language learning. However, children are shaped out of such uses by the speech community whereas adults may not be (as easily). The correlation reported by Dale and Lupyan may therefore be the result of either adult learners' inferior language skills, but rather of the higher sociolinguistic status that allows them to spread their innovations.

Recent studies have emphasized that the mental lexicon grows throughout one's lifetime (Keuleers, Stevens, Mandera & Brysbaert, 2015; Ramscar, Hendrix, Shaoul & Baayen, 2014). Based on a megastudy of over 300,000 Dutch participants, Keuleers et al. (2015) estimate that the lexicon grows by ~48% between age 13 and age 70, with more

than half of this increase occurring after age 20. Thus, adults have many opportunities to extend their morphological and lexical tools to novel uses. Gershkoff-Stowe & Smith (1997, p.65) have argued that extensions may be especially likely to occur when the lexicon is expanding rapidly, whether one is a child or an adult.

While the growth of the lexicon looks gradual when averaged across speakers, lexical growth within an individual likely experiences repeated periods of acceleration and deceleration as one enters new social environments and therefore encounters whole new lexical fields (school, college, a new city, a new job, a new partner, etc.) These periods of rapid change in the lexicon may be especially likely to give rise to novel uses of words that the adult has previously learned. Indeed, sociolinguists have argued that various kinds of linguistic innovations are particularly likely to originate in individuals with high social mobility (Labov, 1966). Once initiated, the novel use of a form must spread through the community to become the community norm. Children may be especially important to this process of actuation (see also Clark & Clark, 1979). Whereas adults may often initiate language change, children may be particularly likely to take up innovative uses instead of rejecting them as unacceptable because, for children, the form is not entrenched in its original meaning.

As suggested by Heine (2011), extension tends to trigger semantic broadening. Unlike extension, semantic broadening is likely to be initiated by children. When a child learns the language from an adult who uses the form in both its novel and original meanings, the child may infer a novel, more general meaning for the form, turning extension into broadening. However, as mentioned before, semantic broadening is unlikely to be limited to children. The meanings of sublexical units are widely believed to

be learned by generalization over the lexical units that contain them (e.g. Aronoff, 1976; Baayen et al., 2011; Bybee, 1985, 2010). As new words are learned in adulthood, the meanings of morphemes inside these words are likely to continue to change. For this reason, exposure to new uses of a particular morpheme in adulthood may well modify one's beliefs about the meaning of the morpheme.

In general, I expect the results of the present experiments to be replicated with children. Indeed, Schwab et al. (2018) have recently reported that children too extend frequent forms to related meanings in production, but do not prefer frequent forms over infrequent competitors in forced form choice. Several researchers have also reported dissociations between production and comprehension in children (Gershkoff-Stowe, 2001; Gershkoff-Stowe & Smith, 1997; Kuczaj, 1982; Nagles & Gelman, 1995; Nelson et al., 1978). As discussed in Chapter 3 and § 7.4, the results of the present experiments are predicted by simple, associative learning as well as Bayesian inference, mechanisms that have been argued to be within the purview of children. However, to the extent that such differences may be found (see Ramscar et al., 2010, 2013, for convincing examples), the behavior of both adults and children needs to be considered in attempting to explain the recurrent patterns of language change. It is when both children and adults favor a particular language change that the change is likely to be actuated.

7.9 Semantic Effects

7.9.1 The effect of meaning frequency on form choice

The languages Dan and Nem differ not only in the frequencies of forms, but also in the frequencies with which the semantic features DIM and PL are expressed in the input. One might therefore wonder whether the Dan language leads learners to be more

likely to profile the PL aspect of the DIM.PL meaning, while the Nem language leads them to profile DIM. This hypothesis is borne out to some extent in the production data of Experiments I-III: the use of plural suffixes to express DIM.PL is more common in Dan, and the use of diminutive suffixes more common in Nem. The hypothesis cannot be tested in the same way in Experiment IV because the languages in that experiment do not contain pairs of synonymous PL or DIM suffixes.

The production data in Experiments I–III are particularly striking because the effect of meaning frequency persists across all three experiments, even though the effect of form frequency reverses in production from Experiment I to Experiments II-III. Thus, even though Dan favors the use of *-dan*$_{\text{PL}}$ over *-nem*$_{\text{DIM}}$ to expresss DIM.PL in Experiment I and disfavors it in Experiments II-III, Dan consistently favors using *some* PL suffix to express DIM.PL. That suffix, however, is *-dan*$_{\text{PL}}$ in Experiment I and *-sil*$_{\text{PL}}$ in Experiments II-III. These results therefore suggest that participants exposed to the Dan language are consistently more likely to construe DIM.PL as PL compared to participants exposed to the Nem Language.

The possible accounts of this effect of meaning frequency are constrained by the fact that the preference to express the frequent semantic feature holds only in production. There is no preference for expressing the frequent semantic feature of DIM.PL in the forced form choice task. The semantic feature frequency effect is therefore mediated by form accessibility. It is therefore not the case that the learners of Dan find it more important to express PL than the learners of Nem do. Rather, the greater salience of PL in Dan results in a stronger activation of the associated forms, *-dan*$_{\text{PL}}$ and *-sil*$_{\text{PL}}$. When the

activation differences are leveled (in the forced choice task), construction frequency no longer affects what aspects of the DIM.PL meaning are profiled.

More work is needed to determine whether construction frequency actually affects how DIM.PL is perceived when form choice is not at issue (cf., Levinson, 1996; Papafragou, Hilbert & Trueswell, 2008; Papafragou, Massey & Gleitman, 2002; Pederson et al., 1998; Slobin, 2004). The present results are consistent with either position on this question. Form accessibility benefits from the accessibility of the associated semantics as forms compete for production, and accessibility of the semantics of a construction at that moment is influenced by construction frequency. Whether this would also be the case at other times, e.g., before utterance formulation commences, is an interesting question for future research.

7.9.2 Compositionality

The training data in Experiments IIII are consistent with both languages obligatorily expressing both DIM and PL features: BIG.SG is always unsuffixed, while DIM.SG is always suffixed with -*nem*$_{\text{DIM}}$ or -*shoon*; and BIG.PL is always suffixed with -*dan*$_{\text{PL}}$ or -*sil*$_{\text{PL}}$. One might therefore argue that participants learning either language ought to learn that both features should be obligatorily expressed. However, participants do not act as if they had learned such a generalization. If they did, they would express both features of DIM.PL using a combination of a diminutive suffix and a plural suffix, producing forms like -*dannem*, -*silshoon*, etc. Such forms account for only between 1–2% of productions. Instead, participants extend one of the DIM or PL suffixes to the DIM.PL meaning.

I can only speculate as to the reasons for the absence of compositional productions. It could be that participants learn this from the language because the suffixes occur in complementary distribution in the input and never co-occur. It could also be that this reluctance to combine suffixes is due to the relative poverty of English morphology. Learners may also make comparisons between the suffixed and unsuffixed forms to learn what contrasts need to be expressed (e.g., Kapatsinski, 2013). Finally, participants may need to be nudged not to collapse semantic distinctions, possibly by feedback from the interlocutor (Kirby, Cornish & K. Smith, 2008). The necessary conditions for the emergence of compositionality remains an important direction for future work.

7.9.3 Prior biases and learning rate

Across Experiments I–III, participants are more likely to express the DIM.PL meaning using a PL suffix than using a DIM suffix. This bias in favor of expressing PL may be the result of first language experience, since English is much more likely to express PL than DIM. It would be an interesting question to ask whether speakers of languages in which diminutive marking is more common, such as Spanish or Russian, would also show a preference for expressing DIM.PL with plural markers. This bias too disappears in the forced form choice task, indicating that its influence is mediated by form accessibility.

In Experiment IV, there *is* a bias in the forced form choice task, and it is in the opposite direction from the production bias in Experiments I–III: in Experiment IV, participants tend to choose diminutive suffixes to express DIM.PL. This was not an expected result. A possible post hoc explanation is that the change in the language between experiments made DIM more salient by explicitly pairing it with a form, *-nem*,

that expresses the feature whether the referent is plural or singular. This may draw attention to the fact that DIM.PL is DIM. When DIM.PL is marked as DIM, this is relatively unexpected, compared to marking it as PL. According to error-driven learning theory, surprise leads to faster learning (Rescorla & Wagner, 1972). In particular, Tassoni (1995) has argued that learning is faster when it involves unexpected cues. If this holds for form–meaning mappings, then learning rate may be faster on DIM.PL ~ -*nem*$_{\text{DIM}}$ training trials compared to DIM.PL ~ -*dan*$_{\text{PL}}$ training trials. As a result, the DIM suffix would be associated with DIM.PL more strongly than the PL suffix after Experiment IV training.

7.9.4 Robustness to leveling accessibility differences

The fact that the bias to treat DIM.PL as DIM in Experiment IV is robust to leveling form accessibility differences, whereas the bias to treat DIM.PL as PL in Experiments I–III is not, suggests that the two biases do not have the same source. This is reflected in attributing the former to differences in meaning–form connection strengths at the end of training, and the latter to greater prior accessibility of PL compared to DIM. If the PL ~ -*dan*$_{\text{PL}}$ connection is of the same strength as the DIM ~ -*nem*$_{\text{DIM}}$ connection, -*dan*$_{\text{PL}}$ will receive more activation from PL than -*nem*$_{\text{DIM}}$ will from DIM. However, when accessibility differences are leveled by forced choice, the difference between the form activations disappears.

The differences between forms that do *not* disappear in forced form choice are the ones attributable to the experimental contingencies: forms are chosen in response to meanings with which they were paired more than in response to the meaning that shares one feature with those meanings (DIM.PL). They are chosen in response to DIM.PL more

than in response to the wrong meaning. In addition, as just mentioned, DIM forms are chosen more than PL forms. All of these differences can be plausibly captured by the connections between meanings and forms, rather than by form or meaning activation levels. Intuitively, it seems that differences in activations resulting from learning experimental contingencies would be more robust than differences resulting from differences in activation of cues. However, it is important to note that this does not directly follow from the models of learning I have explored in this dissertation: when choosing among forms, the learner has no access to what the differences in activation levels between these forms are based on. If this generalization is upheld in future experiments, the contrast between open-set and closed-set choice tasks may prove quite informative about the nature of learning.

7.9.5 Niche-seeking and push chains

Entrenchment of a frequent form in a specific meaning (DIM.SG or BIG.PL) pushes the synonymous form out of the shared meaning. As just discussed, in Experiments I-III, the semantic feature borne by the frequent form and its synonym is preferentially expressed in referring to DIM.PL during production. In Experiments II-III, the synonym of a frequent form becomes the dominant expression of DIM.PL in production, after being assigned to it during the comprehension test.

An important question regarding this result is whether it constitutes a push chain in semantic space or merely semantic narrowing (Blank, 1999). There are few clear instances of push chains, except for the domain of future reference (Torres Cacoullos & Walker, 2009) while narrowing is relatively ubiquitous and uncontroversial. The difference has to do with whether the synonym of a frequent form would have had the

single-feature meaning DIM or PL if its competitor were infrequent. If it would, then assignment of the form to DIM.PL is merely narrowing: the entrenching frequent form takes over part of the space previously occupied by its synonym, leaving the synonym part of its erstwhile niche. If it would not, then the frequent form pushes its synonym into an area of semantic space it would not otherwise occupy.

It is difficult to argue for a push on the basis of DIM.PL because competitors of infrequent forms do tend to be used to express both DIM.PL and the meaning with which they are paired in training, and to be mapped onto both meanings in comprehension (though they are mapped onto the original meanings more often). A possible piece of evidence for a push, however, comes from the fact that competitors of frequent forms are mapped onto the 'wrong' meaning in comprehension more than competitors of infrequent forms are. In Experiments II and III, the form–meaning mappings settled on in comprehension are also used to express the wrong meaning in production more often than competitors of infrequent forms are. Because competitors of infrequent forms are almost never mapped onto the wrong meaning or used to express it, extension of the same forms to these meanings when they have a frequent competitor suggests that a frequent competitor pushes its synonym out of their comfort area. The lack of a competitor frequency effect on accuracy in the forced form choice task suggests that extension of competitors of frequent forms to 'wrong' meanings in comprehension is not just random noise, and may reflect activation of the meaning from the form.

How could extension to the wrong meaning happen? Note that the wrong meaning shares a feature with the DIM.PL meaning, which all forms can be used to express. A frequent form could then push its competitor from, say, DIM.SG into DIM.PL. The

competitor could then be reinterpreted as meaning PL, and extended to be used with both DIM.PL and BIG.PL. The results of the present experiments provide some tentative support for this line of developments but much more research would be needed to fully substantiate it, either in experimental contexts or in real-life language change.

7.10 Future directions

As noted earlier, the effect of forced choice on leveling accessibility differences between alternatives opens up avenues for future work on the mechanisms of learning. At the same time, it also raises important questions about whether differences in accessibility that have different causes are equally prone to leveling. In particular, participants in the comprehension task appear to be quite successful at matching meaning probabilities conditioned on the presented form. In contrast, participants in the form choice task are not sensitive to form probabilities conditioned on the meaning. This result occurred despite the fact that participants were presented with meanings before forms in training, and therefore received input conducive to conditioning form choice on meaning rather than vice versa. Given that the comprehension task was also a forced choice task, it is not clear why participants were more successful matching meaning probabilities compared to form probabilities. It is possible that the fact that the comprehension task involved four choices rather than two may be important. Alternatively, form choice may inherently impose a greater working memory load because the first alternative form needs to be held in memory until the second form is perceived before the choice can be made. In contrast, the pictures representing alternative meanings are available simultaneously and remain available throughout the choice process. It has been suggested that matching probabilities requires remembering the responses one has made so far (Azab, Ruskin, &

Kidd, 2016). If this hypothesis is true, working memory load may interfere with probability matching in the form choice task.

Whereas this dissertation has focused on form–meaning associations and form accessibility as determinants of form choice, form choice is also influenced by other forms that have already been selected for production. Future work should explore how semantic and syntagmatic cues interact in determining the selection of a form.

While extension, narrowing and entrenchment are well-documented in natural language, niche-seeking remains somewhat mysterious. Many examples of niche-seeking I considered could ultimately be accounted for as a simple narrowing, as a result of a stronger competitor taking over part of a form's semantic niche. It is relatively difficult to find clear examples of a form actually pushing another form into an area of semantic space it did not occupy prior to the other form's encroachment.

Another interesting direction involves generalizing the present work to continuous semantic dimensions. Koranda, Zettersten, and MacDonald (2018) have recently replicated frequency-driven semantic extension in a continuous space. However, the effects of entrenchment, pre-emption, and niche-seeking would be interesting to explore in continuous domains.

More work is needed to elucidate the interaction between accessibility-driven extension and entrenchment in comprehension within a communicative setting. Given that speakers tend to extend frequent forms to new uses, but listeners may consider such extensions to be over-extensions because they are highly confident that the form is not used in that way, what determines whether the innovation will catch on? Social factors

likely play a major role in determining whether the extension is perceived to be acceptable, and whether it is, in the end, accepted.

For accessibility-driven choices to drive language change, comprehenders should be adopting the extended form. Iterated learning studies are a good way to test whether extensions would be taken up and transmitted across generations of learners (Kirby et al., 2008). Another important step that can be taken towards understanding the nature of extension in language is collecting the over-extension errors of adult native speakers. These are more frequent than the literatures on over-extension and language acquisition make us believe. In general, in both linguistics and psychology, there is a tendency to idealize the adult native speaker's perception and production. But if adults continue learning, native speaker's production and comprehension cannot be ideal (see Dabrowska, 2012). Clark and Clark (1979) provide many examples of adult (over-)extensions, referred to as *innovations*. However, there has not been any systematic study of the properties of these innovations. For example, is it really the case that all innovations that are accepted have no pre-emptor constructions (Goldberg, 2019)? At the end, the true test may only be diachronic.

7.11 Conclusion

Why are frequent forms extended to new meanings? I believe the answer to lie in the effect of frequency on lexical access in production. Under the pressures of real-time language use, the speaker is likely to reach for the tool that is easiest for her to access, a frequent form, as long as that tool is known to be at least somewhat suitable for the purpose at hand. When accessibility differences between frequent and rare forms are leveled, frequency no longer has an effect on form choice.

A frequent form is a stronger competitor than an infrequent form. It can therefore successfully pre-empt another form from mapping onto the same area of semantic space. However, frequent exposure to a form in a particular context causes the listener to believe that the form is restricted to that context. Entrenchment of frequent forms provides the semantic competitors of these forms a niche to survive in. As a frequent form retracts from a part of the space it shares with another form, the abandoned niche is left for the competing form to occupy. Conversely, when the semantic range of a frequent form covers the entire range of another form, that form's survival is in question.

I have argued that a dissociation between the effects of frequency in production and comprehension is predicted both by Hebbian associative learning models and Bayesian inference. The work has also documented this dissociation in human behavior: the form that is believed to be restricted to the familiar context may be the one most likely to be used in new contexts in production. However, under some circumstances, it has also been shown that entrenchment can constrain frequency-driven extension. Language change is the result of a tug of war between production pressures, which favor extension of frequent forms, and the stabilizing force of entrenchment, which may help prevent extension from becoming over-extension in a developing system.

REFERENCES CITED

Aaron, J. E. (2006). *Variation and change in Spanish future temporal expression*. Ph.D. Dissertation, University of New Mexico.

Aguado-Orea, J., & Pine, J. M. (2015). Comparing different models of the development of verb inflection in early child Spanish. *PLoS ONE, 10*, e0119613.

Aitchison, J. (1991). *Language change: Progress or decay?* New York: Cambridge University Press.

Akhtar, N. (1999). Acquiring basic word order: Evidence for data-driven learning of syntactic structure. *Journal of Child Language*, *26*, 339–356.

Albright, A., & Hayes, B. (2003). Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition*, *90*, 119–161.

Schoneberger, T. (2010). Three myths from the language acquisition literature. *The Analysis of Verbal Behavior*, *26*, 107–131.

Givón, T. (1979). *On understanding grammar*. San Diego: Academic Press.

Ambridge, B. (2013). How do children restrict their linguistic generalizations?: an (un-)grammaticality judgment study. *Cognitive Science, 37*, 508–543.

Ambridge, B., & Lieven, E. V. (2011). *Child language acquisition: Contrasting theoretical approaches*. Cambridge, UK: Cambridge University Press.

Ambridge, B., Bidgood, A., Twomey, E., Pine, J. M., Rowland, C. F., & Freudenthal, D. (2015). Preemption versus Entrenchment: Towards a construction-general solution to the problem of the retreat from verb argument structure overgeneralization. *PLoS ONE, 10*, e0123723.

Ambridge, B., Pine, J. M., Rowland, C. F., & Young, C. R. (2008). The effect of verb semantic class and verb frequency (entrenchment) on children's and adults' graded judgments of argument-structure overgeneralisation errors, *Cognition, 106*, 87–129.

Ambridge, B., Pine, J. M., Rowland, C. F., & Chang, F. (2012). The roles of verb semantics, entrenchment and morphophonology in the retreat from dative argument structure overgeneralization errors. *Language, 88*, 45–81.

Anglin, J. M. (1977). *Word, object, and conceptual development*. New York: WW Norton.

Arcediano, F., Escobar, M., & Miller, R. R. (2003). Temporal integration and temporal backward associations in human and nonhuman subjects. *Learning & Behavior*, *31*, 242–256.

Aronoff, M. (1976). *Word formation in generative grammar*. Cambridge, MA: MIT Press.

Aronoff, M.(2016). Competition and the lexicon. In *Livelli di Analisi e fenomeni di interfaccia. Atti del XLVII congresso internazionale della società di linguistica Italiana* (pp. 39–52). Roma: Bulzoni Editore.

Lindsay, M., & Aronoff, M. (2013). Natural selection in self-organizing morphological systems. *Morphology in Toulouse: Selected Proceedings of Décembrettes*, *7*, 133–153.

Asch, S. E., & Ebenholtz, S. M. (1962). The principle of associative symmetry. *Proceedings of the American Philosophical Society*, *106*, 135–163.

Azab, H., Ruskin, D., & Kidd, C. (2016). Adults' guesses on probabilistic tasks reveal incremental representativeness biases. *Proceedings of the Annual Conference of the Cognitive Science Society*, *38*, 2831–2836.

Baayen, H. (1992). Quantitative aspects of morphological productivity. In G. E. Booij et al. (Eds.), *Yearbook of Morphology 1991* (pp. 109–149). Dordrecht: Springer.

Baayen, R. H., Milin, P., Đurđević, D. F., Hendrix, P., & Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological Review, 118*, 438–481.

Baerman, M., Brown, D., & Corbett, G. G. (2005). *The syntax-morphology interface: A study of syncretism*. Cambridge University Press.

Baker, C. L. (1979). Syntactic theory and the projection principle. *Linguistic Inquiry*, *10*, 533–581.

Barðdal, J. (2008). *Productivity: Evidence from case and argument structure in Icelandic*. Amsterdam: John Benjamins.

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language, 68*, 255–278.

Bates, D., Maechler, M., Bolker, B. M. & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software, 67*, 1–48.

Bates, E., & Goodman, J. C. (1997). On the inseparability of grammar and the lexicon: Evidence from acquisition, aphasia and real-time processing. *Language and Cognitive Processes*, *12*, 507–584.

Bates, E., & MacWhinney, B. (1989). Functionalism and the competition model. *The crosslinguistic study of sentence processing*, *3*, 73–112.

Berg, T. (1998). *Linguistic structure and change: An explanation from language processing*. Oxford: Clarendon Press.

Jay, T. (1992). *Cursing in America: A Psycholinguistic Study of Dirty Language in the Courts, in the Movies, in the Schoolyards, and on the Streets*. Philadelphia/Amsterdam: John Benjamins.

Blank, A. (1999). Why do new meanings occur? A cognitive typology of the motivations for lexical semantic change. In A. Blank & P. Koch (Eds.), *Historical semantics and cognition*, (pp. 61–90). Berlin: Mouton de Gruyter.

Bloom, L. (1973). *One word at a time: The use of single-word utterances before syntax*. The Hague: Mouton.

Blything, R. P., Ambridge, B., & Lieven, E. V. (2014). Children use statistics and semantics in the retreat from overgeneralization. *PloS ONE*, *9*, e110009.

Boersma, P. & Weenink, D. (2016). Praat: doing phonetics by computer [Computer program]. Version 6.0.17.

Booij, G. (2010). Construction morphology. *Language and Linguistics Compass, 4*, 543–555.

Boudelaa, S., & Gaskell, M. G. (2002). A re-examination of the default system for Arabic plurals. *Language and Cognitive Processes*, *17*, 321–343.

Bowerman, M. (1982). Starting to talk worse: Clues to language acquisition from children's late speech errors. In Stavy, R. (Ed.), *U shaped behavioral growth* (pp. 101–145). San Diego: Academic Press.

Bowerman, M. (1988). The "no negative evidence" problem: how do children avoid constructing an overly general grammar? In: Hawkins, J. A. (Ed.), *Explaining language universals*, (pp. 73–101). Oxford: Blackwell.

Boyd, J. K., & Goldberg, A. E. (2011). Learning what not to say: The role of statistical preemption and categorization in a-adjective production. *Language*, *87*, 55–83.

Braine, M. D. S. (1971). On two types of models of the internalisation of grammars. In D. I. Slobin (Ed.), *The ontogenesis of grammar: A theoretical symposium*. New York: Academic Press.

Braine, M. D. S., & Brooks, P. (1995). Verb argument structure and the problem of avoiding an overgeneral grammar. In M. Tomasello & W. Merriman (Eds.), *Beyond names for things: Young children's acquisition of verbs* (pp. 353–376). Hillsdale, NJ: Erlbaum.

Brooks, P. J., & Tomasello, M. (1999). How children constrain their argument structure constructions. *Language, 75*, 720–738.

Brooks, P. J., Tomasello, M., Dodson, K., and Lewis, L. (1999). Young children's overgeneralizations with fixed transitivity verbs, *Child Development*, *70*, 1325–1337.

Brooks, P. J., & Zizak, O. (2002). Does preemption help children learn verb transitivity? *Journal of Child Language, 29*, 759–781.

Browman, C. P., & Goldstein, L. (1992). Articulatory phonology: An overview. *Phonetica*, *49*, 155–180.

Brown, R., & Hanlon, C. (1970). Derivational complexity and order of acquisition in child speech. In Hayes, J. R. (Ed.), *Cognition and the development of language* (pp. 11–53). New York: Wiley.

Burke, M. D., Kester Locantore, J., Austin, A. A., Chae, B. (2004). *Cherry pit* primes *Brad Pitt*: Homophone priming effects on young and older adults' production of proper names. *Psychological Sciences, 15*, 164–170.

Bush, R. R., & Mosteller, F. (1951). A mathematical model for simple learning. *Psychological Review*, *58*, 313–323.

Buz, E., Tanenhaus, M. K., & Jaeger, T. F. (2016). Dynamically adapted context-specific hyper-articulation: Feedback from interlocutors affects speakers' subsequent pronunciations. *Journal of Memory and Language*, *89*, 68–86.

Bybee, J. (1985). *Morphology: A study of the relation between meaning and form*. Amsterdam: John Benjamins.

Bybee, J. (1995). Regular morphology and the lexicon. *Language and Cognitive Processes*, *10*, 425–455.

Bybee, J. (2001). Main clauses are innovative, subordinate clauses are conservative: Consequences for the nature of constructions. In J. Bybee & M. Noonan (Eds.), *Complex sentences in grammar and discourse: Essays in honor of Sandra A. Thompson*, (pp. 1–18). Amsterdam: John Benjamins.

Bybee, J. (2003). Cognitive processes in grammaticalization. In M. Tomasello (Ed.), *The new psychology of language* (Vol. 2, pp. 145–167). Mahwah, N.J.: Lawrence Erlbaum.

Bybee, J. (2006). From usage to grammar: The mind's response to repetition. *Language*, *82*, 711–733.

Bybee, J. (2010). *Language, Usage and Cognition*. Cambridge, UK: Cambridge University Press.

Bybee, J. L., & Brewer, M. A. (1980). Explanation in morphophonemics: changes in Provençal and Spanish preterite forms. *Lingua*, *52*, 201–242.

Bybee, J. L., Perkins, R. D., & Pagliuca, W. (1994). *The evolution of grammar: Tense, aspect, and modality in the languages of the world*. Chicago: University of Chicago Press.

Bybee, J. L., & Slobin, D. I. (1982). Rules and schemas in the development and use of the English past tense. *Language*, *58*, 265–289.

Bybee, J., & Eddington, D. (2006). A usage-based approach to Spanish verbs of 'becoming'. *Language*, *82*, 323–355.

Bybee, J., & Pagliuca, W. (1987). The evolution of future meaning. In *Papers from the 7th international conference on historical linguistics* (pp. 109–122). Amsterdam: John Benjamins.

Caballero, G., & Inkelas, S. (2013). Word construction: Tracing an optimal path through the lexicon. *Morphology*, *23*, 103–143.

Campbell, L. (2000). What's wrong with grammaticalization? *Language Sciences*, *23*, 113–161.

Cazden, C. B. (1972). *Child language and education*. New York: Holt, Rinehart & Winston.

Chouinard, M. M., & Clark, E. V. (2003). Adult reformulations of child errors as negative evidence. *Journal of Child Language*, *30*, 637–669.

Christiansen, M. H., & Chater, N. (2016). *Creating language: Integrating evolution, acquisition, and processing*. Cambridge, MA: MIT Press.

Clahsen, H. (1999). Lexical entries and rules of language: A multidisciplinary study of German inflection. *Behavioral and Brain Sciences*, *22*, 991–1013.

Clark, E. V. (1973). What's in a word? On the child's acquisition of semantics in his first language. In T. E. Moore (Ed.), *Cognitive development and the acquisition of language* (pp. 65–110). New York: Academic.

Clark, E. (1987). The principle of contrast: a constraint on language acquisition. In B. MacWhinney (Ed.), *Mechanisms of language acquisition* (pp. 1–33). Hillsdale, NJ: Lawrence Erlbaum Associates

Clark, E. V., & Clark, H. H. (1979). When Nouns Surface as Verbs. *Language, 55*, 767–811.

Clopper, C. G., Pisoni, D. B., & Tierney, A. T. (2006). Effects of open-set and closed-set task demands on spoken word recognition. *Journal of the American Academy of Audiology, 17*, 331–349.

Coetzee, A. W., & Pater, J. (2011). The place of variation in phonological theory. In J. Goldsmith, J. Riggle and A. C. L. Yu (Eds.), *The handbook of phonological theory, 2nd Edition* (pp. 401–434). Oxford: Wiley.

Croft, W., & Cruse, D. A. (2004). *Cognitive linguistics*. Cambridge, UK: Cambridge University Press.

Dąbrowska, E. (2012). Different speakers, different grammars: Individual differences in native language attainment. *Linguistic Approaches to Bilingualism*, *2*, 219–253.

Dąbrowska, E., & Szczerbiński, M. (2006). Polish children's productivity with case marking: The role of regularity, type frequency, and phonological diversity. *Journal of Child Language, 33*, 559–597.

Dahan, D., & Tanenhaus, M. K. (2005). Looking at the rope when looking for the snake: Conceptually mediated eye movements during spoken-word recognition. *Psychonomic Bulletin & Review*, *12*, 453–459.

Dale, R., & Lupyan, G. (2012). Understanding the origins of morphological diversity: The linguistic niche hypothesis. *Advances in Complex Systems*, *15*(03n04), 1150017.

Danks, D. (2003). Equilibria of the Rescorla–Wagner model. *Journal of Mathematical Psychology*, *47*, 109–121.

Davies, M. (2008). *The corpus of contemporary American English*. Brigham Young University.

Dell, G. S. (1985). Positive feedback in hierarchical connectionist models: Applications to language production. *Cognitive Science*, *9*, 3–23.

Dell, G. S. (1986). A spreading-activation theory of retrieval in sentence production. *Psychological Review, 93*, 283–321.

Dell, G. S., & Reich, P. A. (1981). Stages in sentence production: An analysis of speech error data. *Journal of Verbal Learning and Verbal Behavior*, *20*, 611–629.

Dewey, J. (1894). The psychology of infant language. *Psychological Review*, *1*, 63–66.

Diessel, H. (2007). Frequency effects in language acquisition, language use, and diachronic change. *New Ideas in Psychology*, *25*, 108–127.

Dromi, E. (1987). *Early lexical development*. Cambridge, UK: Cambridge University Press.

Ellis, N. C., & Ferreira-Junior, F. (2009). Constructions and their acquisition: Islands and the distinctiveness of their occupancy. *Annual Review of Cognitive Linguistics*, *7*, 188–221.

Engelmann, F., Granlund, S., Kolak, J., Szreder, M., Ambridge, B., Pine, J., Theakston, A. & Lieven, E. (2019). How the input shapes the acquisition of verb morphology: Elicited production and computational modelling in two highly inflected languages. *Cognitive Psychology*, *110*, 30–69.

Ervin, S. M., & Miller, W. R. (1963). Language development, In H. W. Stevenson (Ed.), *Child psychology, the 62nd Yearbook of the National Society for the Study of Education Part 1*. Chicago, IL: University of Chicago Press.

Ferdinand, V., Thompson, B., Kirby, S., & Smith, K. (2013). Regularization behavior in a non-linguistic domain. *Proceedings of the Annual Conference of the Cognitive Science Society*, *36*, 436–441.

Ferreira, V. S., & Griffin, Z. M. (2003). Phonological influences on lexical (mis)selection. *Psychological Science*, *14*, 86–90.

Fremgen, A., & Fay, D. (1980). Overextensions in production and comprehension: A methodological clarification. *Journal of Child Language*, *7*, 205–211.

Freudenthal, D., Pine, J. M., Aguado-Orea, J., & Gobet, F. (2007). Modeling the developmental patterning of finiteness marking in English, Dutch, German, and Spanish using MOSAIC. *Cognitive Science*, *31*, 311–341.

Gershkoff-Stowe, L. (2001). The course of children's naming errors in early word learning. *Journal of Cognition and Development, 2*, 131–155.

Gershkoff-Stowe, L. (2002). Object naming, vocabulary growth, and the development of word retrieval abilities. *Journal of Memory and Language,46*, 665–87.

Gershkoff-Stowe, L., & Smith, L. B. (1997). A curvilinear trend in naming errors as a function of early vocabulary growth. *Cognitive Psychology*, *34*, 37–71.

Gershkoff-Stowe, L., Connell, B., & Smith, L. (2006). Priming overgeneralizations in two- and four-year-old children. *Journal of Child Language*, *33*, 461–486.

Givón, T. (1979). *On understanding grammar*. Orlando, FL: Academic Press.

Gold, E. M. (1967). Language identification in the limit. *Information and Control*, *10*, 447–474.

Goldberg, A. E. (2006). *Constructions at work*. Oxford: Oxford University Press.

Goldberg, A. E. (2011). Corpus evidence of the viability of statistical preemption. *Cognitive Linguistics*, *22*, 131–153.

Goldberg, A. E. (2019). *Explain Me This: Creativity, competition, and the partial productivity of constructions*. Princeton, NJ: Princeton University Press.

Goldberg, A. E., Casenhiser, D. M., & Sethuraman, N. (2004). Learning argument structure generalizations. *Cognitive Linguistics*, *15*, 289–316.

Gollan, T. H., Slattery, T. J., Goldenberg, D., Van Assche, E., Duyck, W., & Rayner, K. (2011). Frequency drives lexical access in reading but not in speaking: the frequency-lag hypothesis. *Journal of Experimental Psychology: General*, *140*, 186–209.

Goodman, N. D., Tenenbaum, J. B., Feldman, J., & Griffiths, T. L. (2008). A rational analysis of rule-based concept learning. *Cognitive Science*, *32*, 108–154.

Griffin, Z. M., & Bock, K. (1998). Constraint, word frequency, and the relationship between lexical processing levels in spoken word production. *Journal of Memory and Language, 38*, 313–338.

Hare, M., & Elman, J. L. (1995). Learning and morphological change. *Cognition*, *56*, 61–98.

Hare, M., Elman, J. L., & Daugherty, K. G. (1995). Default generalisation in connectionist networks. *Language and Cognitive Processes*, *10*, 601–630.

Harley, B., & King, M. L. (1989). Verb lexis in the written compositions of young L2 learners. *Studies in Second Language Acquisition*, *11*, 415–439.

Harmon, Z., & Kapatsinski, V. (2017). Putting old tools to novel uses: The role of form accessibility in semantic extension. *Cognitive Psychology, 98*, 22–44.

Harmon, Z., & Kapatsinski, V. (2016). Fuse to be used: A weak cue's guide to attracting attention. *Proceedings of the Annual Conference of the Cognitive Science Society, 38*, 520–525.

Haspelmath, M. (1999). Why is grammaticalization irreversible? *Linguistics*, *37*, 1043–1068.

Hay, J. (2001). Lexical frequency in morphology: Is everything relative?. *Linguistics*, *39*, 1041–1070.

Hebb, D. O. (1949). *The organization of behavior: A neuropsychological approach*. Malden, MA: Wiley.

Heine, B. (2011). Grammaticalization in African languages. In H. Narrog & B. Heine (Eds.), *The Oxford handbook of grammaticalization* (pp. 696–707). New York: Oxford University Press.

Heine, B., & Kuteva, T. (2002). *World lexicon of grammaticalization*. Cambridge, UK: Cambridge University Press.

Hopper, P. J., & Thompson, S. A. (1980). Transitivity in grammar and discourse. *Language*, *56*, 251–299.

Hopper, P. J., & Traugott, E. C. (2003). *Grammaticalization*. Cambridge, UK: Cambridge University Press.

Hudson Kam, C. L., & Newport, E. L. (2005). Regularizing unpredictable variation: The roles of adult and child learners in language formation and change. *Language Learning and Development*, *1*, 151–195.

Hudson, J., & Nelson, K. (1984). Play with language: Overextensions as analogies. *Journal of Child Language*, *11*, 337–346.

Huttenlocher, J. (1974). The origins of language comprehension. In R. L. Solso (Ed.), *Theories in cognitive psychology* (pp. 331–368). Potomac, MD: Erlbaum.

Jurafsky, D. (2003). Probabilistic modeling in psycholinguistics: Linguistic comprehension and production. In Bod, R., Hay, J., & Jannedy, S. (Eds.), *Probabilistic linguistics* (pp. 39–96). Cambridge, MA: The MIT Press.

Kahana, M. J. (2002). Associative symmetry and memory theory. *Memory & Cognition*, *30*, 823–840.

Kahneman, D. (2011). *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

Kapatsinski, V. (2009). Adversative conjunction choice in Russian (*no*, *da*, *odnako*): Semantic and syntactic influences on lexical selection. *Language Variation and Change, 21*, 157–173.

Kapatsinski, V. (2010). Velar palatalization in Russian and artificial grammar: Constraints on models of morphophonology. *Laboratory Phonology*, *1*, 361–393.

Kapatsinski, V. (2013). Conspiring to mean: Experimental and computational evidence for a usage-based harmonic approach to morphophonology. *Language*, *89*, 110–148.

Kapatsinski, V. (2018). On the intolerance of the Tolerance Principle. *Linguistic Approaches to Bilingualism*, *8*, 738–742.

Kapatsinski, V. (2018a). *Changing minds changing tools: From learning theory to language acquisition to language change*. Cambridge, MA: MIT Press.

Kapatsinski, V. (2018b). Learning morphological constructions. In G. E. Booij (Ed.), *The Construction of Words* (pp. 547–581). Cham: Springer.

Kapatsinski, V. (2018c). Words versus rules (Storage versus on-line production/processing) in morphology. In M. Aronoff (Ed.), *Oxford Research Encyclopedia of Linguistics*. Oxford: Oxford University Press.

Kapatsinski, V. M. (2005). Characteristics of a rule-based default are dissociable: Evidence against the Dual Mechanism Model. In *Formal approaches to Slavic linguistics* (Vol. 13, pp. 136–146).

Kapatsinski, V., & Harmon, Z. (2017). A Hebbian account of entrenchment and (over)-extension in language learning. *Proceedings of the Annual Conference of the Cognitive Science Society, 39*, 2366–2371.

Kapatsinski, V., & Janda, R. (2011). It's around here: Residential history and the meaning of 'Midwest'. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 33). 2983–2988.

Kapatsinski, V., Easterday, S., & Bybee, J. (2019). Vowel reduction: A usage-based perspective. *Rivista di Linguistica*.

Kay, D. A., & Anglin, J. M. (1982). Overextension and underextension in the child's expressive and receptive speech. *Journal of Child Language*, *9*, 83–98.

Keuleers, E., Stevens, M., Mandera, P., & Brysbaert, M. (2015). Word knowledge in the crowd: Measuring vocabulary size and word prevalence in a massive online experiment. *The Quarterly Journal of Experimental Psychology*, *68*, 1665–1692.

Kiparsky, P. (1973). Elsewhere in phonology. In S. Anderson, & P. Kiparsky (Eds.), *A Festschrift for Morris Halle* (pp. 93–106). New York: Holt, Rhinehart, and Winston.

Kiparsky, P. (1982). Word-formation and the lexicon. In F. Ingemann (Ed.), *Proceedings of the 1982 Mid-America Linguistics Conference* (pp. 3–29). Lawrence: University of Kansas.

Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences*, *105*, 10681–10686.

Kirov, C., & Cotterell, R. (2018). Recurrent neural networks in linguistic theory: Revisiting Pinker and Prince (1988) and the past tense debate. *Transactions of the Association for Computational Linguistics*, *6*, 651–665.

Kleinschmidt, D. F., & Jaeger, T. F. (2016). Re-examining selective adaptation: Fatiguing feature detectors, or distributional learning?. *Psychonomic Bulletin & Review*, *23*, 678–691.

Koepcke, K. M. (1988) Schemes in German plural formation. *Lingua*, *74*, 303–335.

Koranda, M. Zettersten, M. & MacDonald, M. (2018). Word frequency can affect what you choose to say. *Proceedings of the Cognitive Science Society*, *40*, 631–636.

Kuczaj, S. A. (1977). The acquisition of regular and irregular past tense forms. *Journal of Verbal Learning and Verbal Behavior*, *16*, 589–600.

Kuczaj, S. A. (1981). More on children's initial failures to relate specific acquisitions. *Journal of Child Language*, *8*, 485–487.

Kuczaj, S. A. (1982). Young children's overextensions of object words in comprehension and/or production: Support for a prototype theory of early object word meaning. *First Language*, *3*, 93–105.

Kurtz, K. J., Levering, K. R., Stanton, R. D., Romero, J., & Morris, S. N. (2013). Human learning of elemental category structures: Revising the classic result of Shepard, Hovland, and Jenkins (1961). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*, 552–572.

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. (2016). lmerTest: Tests in Linear Mixed Effects Models. R package version 2.0–33. https://CRAN.R-project.org/package=lmerTest.

Labov, W. (1966). *The social stratification of English in New York City*. Cambridge, UK: Cambridge University Press.

Labov, W. (1969). Contraction, deletion, and inherent variability of the English copula. *Language*, *45*, 715–762.

Labov, W. (1973). The boundaries of words and their meanings. *New ways of analyzing variation in English*.

Lakoff, G. (1987). *Women, fire and dangerous things: What categories reveal about the mind*. Chicago, IL: University of Chicago Press.

Lambert, W. E., & Jakobovits, L. A. (1960). Verbal satiation and changes in the intensity of meaning. *Journal of Experimental Psychology*, *60*, 376–383.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*, 211–240.

Langacker, R. W. (1987). *Foundations of cognitive grammar: Theoretical prerequisites* (Vol. 1). Stanford University Press.

Langacker, R. W. (1990). Subjectification. *Cognitive Linguistics*, *1*, 5–38.

Leonard, L. B., Caselli, M. C., & Devescovi, A. (2002). Italian children's use of verb and noun morphology during the preschool years. *First Language, 22*, 287–304.

Levelt, W. J. M. (1989). *Speaking: From intention to articulation.* Cambridge, MA: MIT Press.

Levinson, S. (1996). Frames of reference and Molyneux's question: Crosslinguistic evidence. In: Bloom, P., Peterson, M., Nadel, L., & Garrett, M. (Eds.) *Language and space* (pp. 109–170). Cambridge, MA: MIT Press.

Lewis, M. M. (1957). *How children learn to speak*. London: George G. Harrap.

Lightfoot, D. (1999). *The development of language: Acquisition, change, and evolution*. Wiley-Blackwell.

Lobben, M. (1991). Pluralization of Hausa nouns, viewed from psycholinguistic experiments and child language data. (M.A. Thesis, University of Oslo).

Luce, P. A., & Pisoni, D. B. (1998). Recognizing spoken words: The neighborhood activation model. *Ear and Hearing*, *19*, 1–36.

Luce, R. D. (1963). A threshold theory for simple detection experiments. *Psychological Review*, *70*, 61–79.

MacDonald, M. C. (2013). How language production shapes language form and comprehension. *Frontiers in Psychology, 4,* 226.

MacKay, D. G. (1992). Awareness and error detection. *Consciousness & Cognition*, *1*, 199–225.

MacWhinney, B. (1978). The acquisition of morphophonology. *Monographs of the Society for Research in Child Development*, 1–123.

MacWhinney, B. (1987). The competition model. In B. MacWhinney (Ed.), *Mechanisms of language acquisition* (pp. 249–308). Hillsdale, NJ: Lawrence Erlbaum.

MacWhinney, B. (1993). The (il)logical problem of language acquisition. *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum.

MacWhinney, B. (2004). A multiple process solution to the logical problem of language acquisition. *Journal of child language*, *31*, 883–914.

MacWhinney, B., & Snow, C. (1985). The child language data exchange system. *Journal of Child Language*, *12*, 271–295.

Madlener, K. (2016). Input Optimization. Effects of Type and Token Frequency Manipulations in Instructed Second Language Learning. In H. Behrens & S. Pfaender (Eds.). *Experience counts: frequency effects in language* (pp. 133–173). Berlin: Walter De Gruyter.

Marchman, V. A. (1997). Children's productivity in the English past tense: The role of frequency, phonology, and neighborhood structure. *Cognitive Science, 21*, 283–304.

Marcus, G. F. (1993). Negative evidence in language acquisition. *Cognition*, *46*, 53–85.

Marcus, G. F., Pinker, S., Ullman, M., Hollander, M., Rosen, T. J., Xu, F., & Clahsen, H. (1992). Overregularization in language acquisition. *Monographs of the Society for Research in Child Development*, i–178.

Markman, E. M., & Wachtel, G. F. (1988). Children's use of mutual exclusivity to constrain the meanings of words. *Cognitive Psychology*, *20*, 121–157.

Martin, A. T. (2007). *The evolving lexicon*. Ph.D. Dissertation, UCLA.

Martinet, A. (1952). Function, structure and sound change. *Word*, *8*, 1–32.

Matthews, D. E., Lieven, E., Theakston, A., & Tomasello, M. (2005). The role of frequency in the acquisition of English word order. *Cognitive Development*, *20*, 121–136.

Matthews, D. E., & Theakston, A. L. (2006). Errors of omission in English-speaking children's production of plurals and the past tense: The effects of frequency, phonology, and competition. *Cognitive Science, 30*, 1027–1052.

Matzel, L. D., Held, F. P., & Miller, R. R. (1988). Information and expression of simultaneous and backward associations: Implications for contiguity theory. *Learning and Motivation*, *19*, 317–344.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, *88*, 375–407.

McDonough, K., & Nekrasova-Becker, T. (2014). Comparing the effect of skewed and balanced input on English as a foreign language learners' comprehension of the double-object dative construction. *Applied Psycholinguistics*, *35*, 419–442.

McDonough, K., & Trofimovich, P. (2013). Learning a novel pattern through balanced and skewed input. *Bilingualism: Language and Cognition*, *16*, 654–662.

McKenzie, C. R., & Mikkelsen, L. A. (2007). A Bayesian view of covariation assessment. *Cognitive Psychology*, *54*, 33–61.

McMurray, B., Horst, J. S., & Samuelson, L. K. (2012). Word learning emerges from the interaction of online referent selection and slow associative learning. *Psychological Review*, *119*, 831–877.

McNeill, D. (1966). Developmental psycholinguistics. In F. Smith, & G. Miller, (Eds.), *The Genesis of Language* (pp. 15–84). Cambridge: MIT Press.

Medin, D. L., & Edelson, S. M. (1988). Problem structure and the use of base-rate information from experience. *Journal of Experimental Psychology: General*, *117*, 68–85.

Meillet, A. (1912). L'évolution des formes grammaticales. *Scientia*, *6*(12), 384.

Merriman, W. E., & Bowman, L. L. (1989). The mutual exclusivity bias in children's word learning. *Monographs of the Society for Research in Child Development*, i–129.

Mervis, C. B. (1987). Child-basic object categories and early lexical development. In U. Neisser (Ed.), *Emory symposia in cognition, 1. Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 201–233). New York, NY: Cambridge University Press.

Mervis, C. B., & Bertrand, J. (1994). Acquisition of the novel name–nameless category (N3C) principle. *Child Development*, *65*, 1646–1662.

Mervis, C. B., & Canada, K. (1983). On the existence of competence errors in early comprehension: a reply to Fremgen & Fay and Chapman & Thomson. *Journal of Child Language*, *10*, 431–440.

Miller, R. R., & Matzel, L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. *Psychology of Learning and Motivation*, *22*, 51–92.

Miller, W., & Ervin, S. (1964). The development of grammar in child language. *Monographs of the Society for Research in Child Development*, 9–34.

Mirman, D. (2011). Effects of near and distant semantic neighbors on word production. *Cognitive, Affective, & Behavioral Neuroscience*, *11*, 32–43.

Moscoso del Prado Martin, F., & Brendel, C. (2016). Case and cause in Icelandic: Reconstructing causal networks of cascaded language changes. In *Proceedings of the ACL*.

Mowrey, R., & Pagliuca, W. (1995). The reductive character of articulatory evolution. *Rivista di Linguistica*, *7*, 37–124.

Müller, A., Storto, L., & Coutinho-Silva, T. (2006). Number and the mass/count distinction in Karitiana. In S. Armoskaite & J. J. Thompson (Eds.). *Proceedings of WSCLA XI: The Eleventh Workshop on Structure and Constituency in Languages of the Americas* (pp. 122–135).

Naigles, L. G., & Gelman, S. A. (1995). Overextensions in comprehension and production revisited: Preferential-looking in a study of dog, cat, and cow. *Journal of Child Language*, *22*, 19–46.

Nakamura, D. (2012). Input skewedness, consistency, and order of frequent verbs in frequency-driven second language construction learning: A replication and extension of Casenhiser and Goldberg (2005) to adult second language acquisition. *International Review of Applied Linguistics in Language Teaching*, *50*, 1–37.

Nelson, K., Rescorla, L., Gruendel, J., & Benedict, H. (1978). Early lexicons: What do they mean? *Child Development*, *49*, 960–968.

Norris, D. (2006). The Bayesian Reader: Explaining word recognition as an optimal Bayesian decision process. *Psychological Review*, *113*, 327–357.

Norris, D., & McQueen, J. M. (2008). Shortlist B: a Bayesian model of continuous speech recognition. *Psychological Review*, *115*, 357–395.

O'Donnell, T. J. (2015). *Productivity and reuse in language: A theory of linguistic computation and storage*. Cambridge, MA: MIT Press.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.

Oldfield, R., & Wingfield, A. (1965). Response latencies in naming objects. *The Quarterly Journal of Experimental Psychology, 17*, 273–81.

Olejarczuk, P. (2018). *Phonotactic generalizations and the metrical parse*. Ph.D. Dissertation, University of Oregon.

Oppenheim, G. M., Dell, G. S., & Schwartz, M. F. (2010). The dark side of incremental learning: A model of cumulative semantic interference during lexical access in speech production. *Cognition*, *114*, 227–252.

Osgood, C. E. & Sebeok, T. A., (1954). *Psycholinguistics: A survey of theory and research problems*. Bloomington, IN: Indiana University Press.

Pavlov, I. P. (1927). *Conditioned reflexes. An Investigation of the physiological activity of the cerebral cortex*. Mineola, NY: Dover.

Pederson, E., Danziger, E., Wilkins, D., Levinson, S., Kita, S., & Senft, G. (1998). Semantic typology and spatial conceptualization. *Language*, *74*, 557–589.

Pereira, A. F., Smith, L. B., & Yu, C. (2014). A bottom-up view of toddler word learning. *Psychonomic Bulletin & Review*, *21*, 178–185.

Perek, F., & Goldberg, A. E. (2015). Generalizing beyond the input: The functions of the constructions matter. *Journal of Memory and Language*, *84*, 108–127.

Perfors, A. (2016). Adult regularization of inconsistent input depends on pragmatic factors. *Language Learning and Development*, *12*, 138–155.

Perfors, A., Ransom, K., & Navarro, D. (2014). People ignore token frequency when deciding how widely to generalize. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *36*, 2759–2764.

Pertsova, K. (2011). Grounding systematic syncretism in learning. *Linguistic Inquiry*, *42*, 225–266.

Piantadosi, S.T., Tily, H., & Gibson, E. (2012). The communicative function of ambiguity in language. *Cognition, 122*, 280–291.

Pinker, S. (1989). *Learnability and cognition: The acquisition of verb-argument structure*. Cambridge, MA: Harvard University Press.

Pinker, S. (1991). Rules of language. *Science*, *253*(5019), 530–535.

Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition*, *28*, 73–193.

Plaster, K., & Polinsky, M. (2010). Features in categorization, or a new look at an old problem. In A. Kibort & G. G. Corbett (Eds.), *Features: Perspectives on a key notion in linguistics* (pp. 109–142). Oxford: Oxford University Press.

Plaut, D. C., & Booth, J. R. (2000). Individual and developmental differences in semantic priming: empirical and computational support for a single-mechanism account of lexical processing. *Psychological Review*, *107*, 786–823.

Poplack, S., & Tagliamonte, S. (1999). The grammaticization of going to in (African American) English. *Language Variation and Change*, *11*, 315–342.

Prasada, S., & Pinker, S. (1993). Generalisation of regular and irregular morphological patterns. *Language and Cognitive Processes*, *8*, 1–56.

R Core Team (2015). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL http://www.R-project.org/.

Radford, A. (1996). Towards a structure-building model of acquisition. *Language Acquisition and Language Disorders*, *14*, 43–90.

Rainer, F. (1988). Towards a theory of blocking: the case of Italian and German quality nouns. *Yearbook of Morphology*, *1*, 155–185.

Ramscar, M., Dye, M., & Klein, J. (2013). Children value informativity over logic in word learning. *Psychological Science*, *24*, 1017–1023.

Ramscar, M., Hendrix, P., Shaoul, C., Milin, P., & Baayen, H. (2014). The myth of cognitive decline: Non-linear dynamics of lifelong learning. *Topics in Cognitive Science*, *6*, 5–42.

Ramscar, M., Yarlett, D., Dye, M., Denny, K., & Thorpe, K. (2010). The effects of feature-label-order and their implications for symbolic learning. *Cognitive Science*, *34*, 909–957.

Räsänen, S. H. M., Ambridge, B., & Pine, J. M. (2016). An elicited-production study of inflectional verb morphology in child Finnish. *Cognitive Science, 40*, 1704–1738.

Reder, L. M., Anderson, J. R., & Bjork, R. A. (1974). Asemantic interpretation of encoding specificity. *Journal of Experimental Psychology*, *102*, 648–656.

Regier, T., & Gahl, S. (2004). Learning the unlearnable: The role of missing evidence. *Cognition*, *93*, 147–155.

Rescorla, L. A. (1980). Overextension in early language development. *Journal of Child Language*, *7*, 321–335.

Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*, *66*, 1–5.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.

Rizzi, L. (1993). Some notes on linguistic theory and language development: The case of root infinitives. *Language Acquisition*, *3*, 371–393.

Robenalt, C., & Goldberg, A. E. (2015). Judgment evidence for statistical preemption: It is relatively better to vanish than to disappear a rabbit, but a lifeguard can equally well backstroke or swim children to shore. *Cognitive Linguistics*, *26*, 467–503.

Robenalt, C., & Goldberg, A. E. (2016). Nonnative speakers do not take competing alternative expressions into account the way native speakers do. *Language Learning, 66*, 60–93.

Rosch, E. (1978). Principles of categorization. In E. Margolis & S. Laurence (Eds.), *Concepts: core readings* (pp. 189–206). Cambridge, MA: MIT Press.

Rubino, R. B., & Pine, J. M. (1998). Subject–verb agreement in brazilian portuguese: What low error rates hide. *Journal of Child Language, 25*, 35–59.

Rumelhart, D. E., & McClelland, J. L. (1986). On learning the past tenses of English verbs. In J. L. McClelland, D. E. Rumelhart, & PDP Research Group (Eds.). *Parallel distributed processing. Explorations in the microstructure of cognition* 2 (pp. 216–271). Cambridge, MA: MIT Press.

Schnorr, J. A., & Atkinson, R. C. (1970). Study position and item differences in the short- and long-term retention of paired associates learned by imagery. *Journal of Verbal Learning and Verbal Behavior*, *9*, 614–622.

Schriefers, H., Meyer, A. S., & Levelt, W. J. (1990). Exploring the time course of lexical access in language production: Picture-word interference studies. *Journal of Memory and Language*, *29*, 86–102.

Schuchardt, H. (1885). *Ueber die Lautgesetze: Gegen die Junggrammatiker*. Berlin: Oppenheim. [Translated by Vennemann, T., Schuchardt, H. E. M., Wilbur, T. H., & Schuchardt, H. E. M. *Schuchardt, the Neogrammarians, and the transformational theory of phonological change: Four essays* (pp. 39–72). Frankfurt: Athenäum.]

Schumacher, R. A., & Pierrehumbert, J. B. (2017). Prior expectations in linguistic learning: a stochastic model of individual differences. *Proceedings of the Annual Conference of the Cognitive Science Society*. *39*, 1059.

Schwab, J. F., Lew-Williams, C., & Goldberg, A. E. (2018). When regularization gets it wrong: children over-simplify language input only in production. *Journal of Child Language*, *45*, 1054–1072.

Skousen, R. (1989). *Analogical modeling of language*. Dordrecht: Springer.

Slobin, D. I. (1982). Universal and particular in acquisition. In L. R. Gleitman & E. Wanner (Eds.), *Language acquisition: The state of the art* (pp. 128–170). Cambridge, UK: Cambridge University Press.

Slobin, D. I. (1987). Thinking for speaking. *Berkeley Linguistics Society*, *13*, 435–445.

Slobin, D. I. (1994). Talking perfectly. In Pagliuca, W. (Ed.), *Discourse origins of the present perfect* (pp. 119–133). Amsterdam: John Benjamins.

Slobin, D. I. (2002). Language evolution, acquisition and diachrony: Probing the parallels. In Givon, T., & Malle, B. F. (Eds.), *The evolution of language out of pre-language* (pp.375–392). Amsterdam: John Benjamins.

Slobin, D. I. (2004) The many ways to search for a frog: Linguistic typology and the expression of motion events. In Strömqvist, S., Verhoeven, L. (Eds.). *Relating events in narrative: Typological and contextual perspectives* (pp. 219–257). Mahwah, NJ: Erlbaum.

Smith, L., & Klein, R. (1990). Evidence for semantic satiation: Repeating a category slows subsequent semantic processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 852–861.

Sommers, M. S., Kirk, K. I., & Pisoni, D. B. (1997). Some considerations in evaluating spoken word recognition by normal-hearing, noise-masked normal-hearing, and cochlear implant listeners. I: The effects of response format. *Ear and Hearing, 18*, 89–99.

Stefanowitsch, A. (2008). Negative entrenchment: A usage-based approach to negative evidence. *Cognitive Linguistics, 19*, 513–531.

Stefanowitsch, A., & Gries, S. T. (2003). Collostructions: Investigating the interaction of words and constructions. *International Journal of Corpus Linguistics*, *8*, 209–243.

Suttle, L., & Goldberg, A. E. (2011). The partial productivity of constructions as induction. *Linguistics*, *49*, 1237–1269.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning.* Cambridge, MA: MIT Press.

Tassoni, C. J. (1995). The least mean squares network with information coding: A model of cue learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 193–204.

Theakston, A. L. (2004). The role of entrenchment in children's and adults' performance limitations on grammaticality judgment tasks, *Cognitive Development, 19*, 15–34.

Theakston, A. L., Lieven, E. V., & Tomasello, M. (2003). The role of the input in the acquisition of third person singular verbs in English. *Journal of Speech, Language, and Hearing Research, 46*, 863–877.

Theakston, A. L., & Rowland, C. F. (2009). The acquisition of auxiliary syntax: A longitudinal elicitation study. Part 1: Auxiliary be. *Journal of Speech, Language, and Hearing Research, 52*, 1449–1470.

Thomson, J. R., & Chapman, R. S. (1977). Who is 'Daddy' revisited: the status of two-year-olds' over-extended words in use and comprehension. *Journal of Child Language, 4*, 359–375.

Tiersma, P. M. (1982). Local and general markedness. *Language*, *58*, 832–849.

Tomasello, M. (2003). *Constructing a language. A usage-based approach.* Cambridge, MA: Harvard University Press.

Torres Cacoullos, R., & Walker, J. A. (2009). The present of the English future: Grammatical variation and collocations in discourse. *Language*, *85*, 321–354.

Traugott, E. C. (1988). Pragmatic strengthening and grammaticalization. *Berkeley Linguistics Society*, *14*, 406–416.

Traugott, E. C., & Dasher, R. B. (2001). *Regularity in semantic change*. Cambridge, UK: Cambridge University Press.

Traugott, E. C., & Trousdale, G. (2013). *Constructionalization and constructional changes*. Oxford: Oxford University Press.

Treffers-Daller, J., & Calude, A. (2015). The role of statistical learning in the acquisition of motion event construal in a second language. *International Journal of Bilingual Education and Bilingualism*, *18*, 602–623.

van Hamme, L. J., & Wasserman, E. A. (1994). Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements. *Learning and Motivation*, *25*, 127–151.

van Rij, J., Wieling, M., Baayen, R., & van Rijn, H. (2016). itsadug: Interpreting time series and autocorrelated data using GAMMs. R package version 2.2.

Vihman, M. M. (1980). Sound change and child language. In E. C. Traugott, R. Labrum, & S. Shepherd (Eds.), *Papers from the 4th International Conference on Historical Linguistics* (pp. 303–320). Amsterdam: John Benjamins.

Wagenmakers, E. J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review, 14*, 779–804.

Wasserman, E. A., Dorner, W. W., & Kao, S. F. (1990). Contributions of specific cell information to judgments to interevent contingency. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*, 509–521.

Wexler, K. (1998). Very early parameter setting and the unique checking constraint: A new explanation of the optional infinitive stage. *Lingua*, *106*, 23–79.

Wickham, H. (2009). *ggplot2: Elegant graphics for data analysis*. New York: Springer.

Wonnacott, E., Newport, E. L., Tanenhaus, M. K. (2008). Acquiring and processing verb argument structure: Distributional learning in a miniature language. *Cognitive Psychology, 56*, 165–209.

Wood, S. N. (2011) Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society B*, *73*, 3–36.

Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review, 114*, 245–272.

Yang, C. (2016). *The price of linguistic productivity: How children learn to break the rules of language*. Cambridge, MA: MIT Press.

Yu, C., & Smith, L. B. (2012). Modeling cross-situational word–referent learning: Prior questions. *Psychological Review*, *119*, 21–39.

Zeldes, A. (2012). *Productivity in argument selection: From morphology to syntax*. Berlin: Walter de Gruyter.

Zipf, G. K. (1949). *Human behavior and the principle of least effort: An introduction to human ecology*. Cambridge, MA: Addison-Wesley Press.

Zwicky, A. M. (1970). Auxiliary reduction in English. *Linguistic Inquiry*, *1*, 323–336.