THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Word senses

Citation for published version:

Rabagliati, H & Srinivasan, M 2019, Word senses. in N Katsos & C Cummins (eds), *Oxford Handbook of Experimental Semantics and Pragmatics*. Oxford University Press.
<https://doi.org/10.1093/oxfordhb/9780198791768.013.20>

Digital Object Identifier (DOI):

[10.1093/oxfordhb/9780198791768.013.20](https://doi.org/10.1093/oxfordhb/9780198791768.013.20)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Oxford Handbook of Experimental Semantics and Pragmatics

Publisher Rights Statement:

This is a draft of a chapter/article that has been accepted for publication by Oxford University Press in The Oxford Handbook of Experimental Semantics and Pragmatics edited by Chris Cummins and Napoleon Katsos, published in 19/3/2019

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Word Senses

Hugh Rabagliati

School of Philosophy, Psychology and Language Sciences, University of Edinburgh

hugh.rabagliati@ed.ac.uk

Mahesh Srinivasan

Department of Psychology, University of California, Berkeley

srinivasan@berkeley.edu

A fundamental goal for semantics is to understand why words have the meanings that they do. Words can be thought of as arbitrary conventions that are shared by a community, and that meet the communicative needs of that community. However, these conventions may also reflect the ways in which speakers conceptualize the world around them. Concepts and conventions provide a framework for understanding variation in the lexical inventories of languages (Malt & Majid, 2013; Regier, Kemp, & Kay, 2015), for instance in how languages divide up the space of colours (Berlin & Kay, 1969; Kay & Regier, 2006; Roberson, Davidoff, Davies, & Shapiro, 2005), or the component parts of the human body (Andersen, 1978; Majid, 2010; Wierzbicka, 2007).

But words don't just name single things. *Blue*, for example, names both a colour and a sad state of mind, while *eye* names both a part of the body and a part of a needle. The malleability of a word's meaning can be quite striking. Consider the word *chicken*. We typically think of chickens as animals that are frequently found on farms, but a *rubber chicken* is a toy, *poached chicken* is a type of food, and the derogatory label for a cowardly person is also *chicken*. Almost all common words are associated with a diversity of different-but-related meanings – which researchers call “senses” – and this holds for both open class words like nouns (e.g., *chicken*) and verbs (e.g., the different ways in which we *grasp* objects and ideas), as well as for closed-class words like prepositions (e.g., Mary can be *over* the moon, *over* her ex, or discussing matters *over* a meal).

Table 1. Examples of lexical flexibility in English.

Patterns and their Senses	Examples
Material for Artifact (<i>glass, tin, iron, etc.</i>)	There is shattered <i>glass</i> on the street/ She poured beer into the <i>glass</i>
Animal for Meat (<i>chicken, turkey, fish, etc.</i>)	The <i>chicken</i> strutted through the coop/ The <i>chicken</i> was overcooked
Object for Representational Content (<i>book, magazine, DVD, etc.</i>)	The <i>book</i> is water damaged/ The <i>book</i> is very funny
Container for Contents (<i>pot, bowl, box, etc.</i>)	She cleaned the <i>pot</i> after dinner / She stirred the <i>pot</i> with a spoon
Body Part for Object Part (<i>leg, arm, back, etc.</i>)	He broke his <i>leg</i> last year / That table has a broken <i>leg</i>
Artist for Product (<i>Picasso, Camus, Mozart, etc.</i>)	<i>Picasso</i> was born in 1881 / That collection has a <i>Picasso</i>

Place for Institution (<i>White House, Wall Street, City Hall, etc.</i>)	The <i>White House</i> is painted white / The <i>White House</i> released a statement
Place for Event (<i>Vietnam, Waterloo, Woodstock, etc.</i>)	<i>Vietnam</i> shares a border with China / He fought during <i>Vietnam</i>
Substance for Placing Substance at Goal (<i>butter, salt, water, etc.</i>)	He found some <i>butter</i> in the fridge/ He is going to <i>butter</i> the toast
Instrument for Action Involving Instrument (<i>shovel, hammer, rake, etc.</i>)	She has a green <i>shovel</i> / She is going to <i>shovel</i> the snow

Table 1 illustrates a number of ways in which the meanings of English words can change in context. As can be seen, the names of materials are used to describe artifacts, the names of artists are used to describe their works, and the names of body parts are used to describe parts of inanimate objects.¹ These are examples of phenomena that have been given a variety of different names by researchers, including polysemy, systematic polysemy, metonymy, logical metonymy, coercion, denominalization, class extension, and more. Here, we treat these different phenomena as variants on a single theme, as examples of a broader phenomenon that we call “lexical flexibility”. Rather than list all of the ways that a word can be flexible, our focus in this chapter is to explain the core organizational characteristics that, we believe, may underlie all of these different phenomena.

To do this, we pose, and attempt to answer, three questions. First, is lexical flexibility a structured and predictable phenomenon, and how does this structure vary across languages? Second, how are flexible word senses represented in the mental lexicon? Finally, we speculate on a somewhat deeper question: why do languages have flexible words in the first place?

1. The structure of lexical flexibility

Lexical flexibility, at least in English, has two key characteristics. First, word senses appear to fall into patterns, such that words from similar semantic fields exhibit similar flexibility (Table 1, Apresjan, 1974; Lehrer, 1990; Nunberg, 1979; Srinivasan & Rabagliati, 2015). For example, *chicken, lamb, duck, and salmon* are all names for animals, and they can also all be used to name the meat that comes from those animals. Likewise, words for substances can also be used to label objects made from those substances (*drinking glass, biscuit tin, clothes iron* etc.). Second, these patterns vary in how freely they permit generalization (Apresjan, 1974; Copestake & Briscoe, 1995; Srinivasan & Rabagliati, 2015). In particular, for some patterns, word meanings can be extended quite freely. Almost any name for an animal can be extended to describe its meat; This holds for animals that are not typically eaten (*John ate stewed wildebeest*) as well as for entirely novel examples of animals (e.g., if a new species of animal, the *wug* is discovered, it would be natural to say *John ate wug for dinner*).

¹ Interestingly, in many of these examples, word meanings vary alongside syntactic category: Nouns and verbs often share a common root (*hammer, water*) and nouns often have different senses when used in mass (*some glass*) or count (*a glass*) syntax.

This pattern allows generalization in much the same way as a productive morphological alternation does (e.g., the use of the *-ly* morpheme in *sad/sadly*), and like these alternations, it also permits for exceptions (e.g., *cow/beef*)². However, for other patterns, generalization is more constrained. While, as noted above, English contains many instances where a name for a substance has been extended to describe an object, this pattern of extension is not productive. For example, imagine that a new type of carbon is invented, called *wug*; it may not be immediately natural to use *wug* to describe a specific, newly-created artifact made from that material.

In this section, we aim to explain these two features of lexical flexibility: i.e., why word senses cluster in patterns, and what determines whether a particular pattern is productive. In particular, our goal is to understand how these two features of lexical flexibility can be explained by conceptual organization, and the linguistic conventions that facilitate communication.

1.1 Explaining patterns of flexibility

“Conceptual” explanations of patterns of flexibility have taken a variety of forms, but perhaps the most prominent argument is that patterns of flexibility reflect deep-set connections between different types of concepts. One particular example of this approach is Conceptual Metaphor Theory (Lakoff & Johnson, 1980a, 1980b), which argues that concepts stand in asymmetric, metaphorical relations. The argument from Conceptual Metaphor Theory, in brief, is that certain types of concepts are grounded in concrete, sensorimotor experience (e.g., our perception of colour allows us to represent the concept of blue) and “image schemas”, and these concrete concepts are used to make sense, metaphorically, of more abstract concepts. For instance, we might ground our understanding of the otherwise ineffable concept of sadness in the concrete foundation of the colour blue. Conceptual Metaphor Theory proposes that this tendency to represent abstract concepts through concrete concepts is reflected in the lexicon: e.g., Because the colour blue grounds our representation of sadness, the color label *blue* has been extended to describe sadness as well.

As a theory of concepts, Conceptual Metaphor Theory is somewhat unique: More standard theories of concepts do not assume that abstract concepts are metaphorically structured or constituted in terms of sensorimotor primitives. Instead, other theories, like Prototype Theory (e.g., Hampton, 1995; Rosch & Mervis, 1975), Exemplar theory (Medin & Schaffer, 1978), and Theory theory (Carey, 1985; Gopnik & Meltzoff, 1997; Murphy & Medin, 1985) assume that concepts are assembled by combining more basic features in particular ways: e.g., The concept of chicken might be built by combining features such as “has feathers”, “found on farms”, “lays eggs”, etc. This means that, unlike Conceptual Metaphor theory, these approaches do not naturally predict the existence of lexical flexibility: e.g., Because there is no “grounding” of one concept in terms of another (e.g., sadness in color), there is no reason to predict that two concepts will share the same label (e.g., the two senses of

² Indeed, in some languages, the animal-meat pattern is marked by a process of compounding (e.g., German and Japanese compound the animal name with their language’s word for meat).

blue). Nevertheless, while these theories may not provide a first-principles explanation for the existence of lexical flexibility, they do generate predictions about how flexibility should be structured, and why patterns might exist. In particular, word senses could be extended based on overlap in important conceptual features: For example, each of the different senses of *chicken* could share features with the core sense that refers to the animal. Under this approach, patterns of lexical flexibility emerge because similar pairs of concepts are related in similar ways; that is to say, chicken meat is to a chicken animal, as turkey meat is to a turkey animal (see, for examples, Nunberg, 1979; Papafragou, 1996; Wilson, 2003).

Regardless of their different approaches toward concepts, the different theories described above each predict that, so long as members of different cultures share similar concepts, lexical flexibility should also be similar across different languages. Because patterns of lexical flexibility reflect important properties of conceptual structure, the same patterns should appear across different languages. As we discuss below, this prediction is not made so strongly under conventionalization accounts.

The most radical form of a conventionalization account would state that the different senses that a word can take are simply arbitrary conventions that speakers must learn. But this account is a non-starter, since word senses are clearly not arbitrary but are instead related to one another in systematic patterns. However, one possibility is that these systematic patterns arise via analogy. For example, Murphy (2007) speculates that the animal and meat senses of *chicken* first entered English as conventions, and that their presence provided a basis for coining analogous senses (e.g., rather than invent a new word for the meat sense of *turkey*, adults coined a new sense by analogy to *chicken*).

If patterns of lexical flexibility result from analogies about conventionalized senses, then we might expect that different patterns would arise in different languages, depending on which senses get coined first in each language. For example, the animal for meat pattern may have arisen in English because the senses of *chicken* became a part of the language, but different events could have unfolded in other languages. In sum, the conventionalization account predicts that different patterns of flexibility should be found in different languages, while the conceptual accounts reviewed above predict that patterns of flexibility should be similar across languages.

To test these predictions, Srinivasan and Rabagliati (2015) investigated whether twenty-seven patterns of flexibility found in English (such as the patterns in Table 1) were also found in fourteen other languages (Cantonese, Farsi, French, Hindi, Hungarian, Indonesian, Italian, Japanese, Korean, Mandarin, Russian, Spanish, Turkish and Vietnamese). Bilingual participants rated whether translation-equivalents of English flexible words showed the same flexibility in their native languages. For example, participants were asked to translate the wordform that, in their language, referred to the animal chicken, and then were asked to rate whether that same wordform could be used to refer to chicken meat. These judgments were then verified and augmented with dictionary analyses. We found evidence that sixteen of the assessed patterns were exemplified by at least one word in all fourteen languages. Of the remaining eleven patterns, statistical analyses suggested that only one was absent more frequently than might be expected due to chance and sampling error. These findings provide some evidence for a conceptual view of lexical flexibility, by

suggesting that patterns of flexibility recur across languages – even across quite typologically distinct languages.

To our knowledge, Srinivasan and Rabagliati (2015) is the only large-scale survey of patterns of flexibility, but other typological work has also provided evidence that conceptual structure helps explain lexical flexibility. For example, Viberg (1983) surveyed flexible uses of verbs of perception across 52 languages, and discovered a hierarchy in how verbs from one modality (sight, hearing) are extended to refer to other modalities (taste, touch), more often than the reverse. As an example, the Italian verb *sentire* can refer to acts of hearing (*sentire il concerto* [to hear the concert]) and acts of feeling (*sentire dolore* [feel pain]). This sensory hierarchy in lexical flexibility has motivated claims about how different senses dominate one another, outside of language. Additional, smaller-scale typological work is reviewed in Srinivasan and Rabagliati (2015).

The robust cross-linguistic presence of patterns of flexibility, such as animal-for-meat or material-for-artifact, is consistent with the conceptual approach. However, there is also evidence for cross-linguistic variability in the specific words that instantiate these patterns (Srinivasan and Rabagliati, 2015). Some patterns of flexibility, like the material-for-object pattern, are instantiated in quite different ways across languages: The Russian word for *glass*, for example, is not used to describe cups, but is instead used to describe car windows. Such variability is also often present in metaphorical patterns, such as in how names for body parts like *head* or *arm* are extended to label other kinds of objects. This cross-linguistic variability suggests that lexical flexibility only reflects a bird's-eye view of conceptual structure. That is to say, conceptual structure may provide constraints on which patterns of flexibility are present in languages, but does not provide tight controls on the manner in which those patterns are lexicalized. Thus, although the meanings of flexible words are subject to cognitive constraints, they are also conventions that reflect a society's communicative needs.

However, not all patterns of lexical flexibility are instantiated across languages in variable ways. For example, many languages use names for animals to refer to their meat, and use names for media (*video*, *book*, *newspaper*, etc.) to refer to their contents (*the hour-long video*). This suggests either that some conventions are so useful that they are required by all cultures, or, perhaps more plausibly, that some patterns place tighter constraints on which senses can plausibly be coined. Consider, for instance, the material-for-object pattern. Knowing that a material word could, in principle, refer to any object made out of that material, does not help in predicting exactly which object will have that label (for example, in English, *glass* cannot be used to label vases, jars, etc.), and this is reflected in cross-linguistic variation. By contrast, knowing that an animal's name should be associated with its meat provides a tight constraint on how that name should be used, as there is typically only one type of meat associated with each type of animal.

1.2 Explaining generalization

The second key feature of lexical flexibility is its productivity: For some patterns of flexibility, it is much easier to coin novel senses than for other patterns. For instance, patterns like animal-for-meat can easily be extended (*John ate stewed wildebeest*), but other patterns,

like material-for-object, are more difficult to extend (*John bought a carbon*). Srinivasan and Rabagliati (2015) argued that such differences in generalizability are determined by the same factors that determine whether a pattern will show variation in its cross-linguistic instantiation. That is, the same factors that explain why the possible senses of *glass* vary across languages also explain why it is hard to coin novel senses for the material-object pattern.

Cross-linguistic comparisons provide evidence for this. Srinivasan and Rabagliati's (2015) survey uncovered a robust correlation between whether a pattern of flexibility permits generalization, and whether that pattern shows variability across languages: The difference between the animal-for-meat and material-for-object patterns we have discussed – in terms of both generalization and cross-linguistic variability – holds up for a large range of patterns. Thus, it is possible to predict a Russian speaker's willingness to coin a novel sense for a pattern from knowing whether that pattern shows a high or low degree of cross-linguistic variability.

One way to explain these data is to assume that patterns like animal-for-meat are semantically transparent, such that it is relatively easy for speakers to see how different exemplars of the pattern are equivalent. This transparency allows them to form generalizations about those words – e.g., that words that can label animals can also label their meat – and use these generalizations to create novel senses. In contrast, it might be more difficult for speakers to form generalizations about words that follow less regular patterns, like material-for-object. As discussed in the next section, these generalizations could be represented as lexical rules and allow speakers to infer new senses without receiving direct evidence for their existence. In this way, they would act similarly to morphological rules, e.g., -y suffixation: e.g., Speakers could infer that *wildebeest* (animal) can become *wildebeest* (meat) in the same way that they can infer that *wildebeest* (animal) becomes *wildebeesty* (trait). The distinction between patterns of flexibility that do and do not permit generalizations could thus be similar to the distinction between morphological affixes that are productive and affixes that are non-productive.

However, an alternative interpretation is that the ability to coin new word senses does not depend on having formed generalizations about patterns, but instead on whether speakers can perceive a semantic relationship between the original word sense and the new sense. This is similar to the idea encountered earlier in the chapter, i.e., that the different senses of a word follow from conceptual organization. For example, Nunberg (1979, 2004) argues that new word senses can be coined whenever the new sense stands in a “noteworthy” relationship to an existing sense (for related proposals see Falkum, 2015; Papafragou, 1996; Wilson, 2003)). Thus, the editor of a handbook might say *The words senses chapter has asked for a deadline extension*, in order to refer to the authors of the chapter on word senses.

Although Nunberg's proposal is interesting, the evidence to date suggests that it is not a correct account of pattern-based generalization (though it may explain more idiosyncratic word uses). In particular, Rabagliati, Marcus and Pytkänen (2011) found that acceptability judgments of word senses were better explained by patterns of flexibility than by metrics of conceptual structure (e.g., similarity, relatedness, etc.). In their study, participants rated

whether the meaning of one word, such as *DVD*, could be changed to refer to the meaning of another word, such as *movie*, and vice versa. Independent sets of participants then rated whether DVDs were related to movies, whether movies were related to DVDs, whether DVDs were similar to movies, whether movies were similar to DVDs, and so forth. Interestingly, these metrics were poor predictors of the acceptability of different word senses, contra the approach suggested by Nunberg. Instead, participants' judgments were best explained by their adherence to specific patterns: e.g., names for media (*DVDs*, *books*, *newspapers*, etc) were easily able to stand for their contents (*movies*, *stories*, *information*) and not vice versa, no matter how related or unrelated participants judged the two things (media and contents) to be.

Together, the typological and experimental data presented in this section indicate that patterns of flexibility have a large influence on the lexicon: They are predictive of variability in whether word senses are constrained across languages, and variability in how readily novel senses can be coined. But how might patterns be represented in the lexicon, and are more generative and less generative patterns represented in similar ways?

2. Lexical representation of word senses

Historically, psychological theories of lexical representation have posited two types of words: 1) A large set of unambiguous words, defined as each having only one meaning, and 2) A smaller set of homophones, defined as each having multiple, unrelated meanings. Under these theories, the lexicon is assumed to consist of a set of form-meaning mappings, with most forms associated with a single meaning, and a small set of forms associated with multiple meanings (and, perhaps, some meanings associated with multiple synonymous forms). However, it has been clear for some time that such an account of the lexicon – often referred to as a sense enumeration lexicon – is a gross over-simplification.

A sense enumeration lexicon ignores not only the fact that most words are flexible and have multiple related meanings, but also that speakers can coin and understand new word senses in ways that generalize from existing patterns (see prior section). Thus, just as speakers know that *chicken*, *duck*, and *lamb* can describe both an animal and its meat, they can understand that “a plate of wildebeest” refers to wildebeest meat, rather than the animal itself. This suggests that the lexicon contains multiple, systematic form-meaning mappings for many words, and also contains mechanisms for generating novel senses. What types of representational structures and processes might explain these properties of the lexicon?

One approach has been to augment standard theories of the lexicon with generative “lexical rules”, akin to morphological or syntactic rules, which generate novel word senses when they are applied to existing words. First, the lexicon is assumed to contain a large store of “irregular” ambiguous words: Mappings between forms and meanings that cannot be easily predicted by a rule, and thus must be learned. Second, and alongside these stored pairings, are a relatively small set of lexical rules that transform word meanings in a principled fashion (Borer, 2005; Copestake & Briscoe, 1995; Pustejovsky, 1995). Such rules generate the “regular” senses of words (e.g., *chicken*, *lamb*, etc.), and so reduce the need to store redundant information in the lexicon. They also provide a mechanism that allows new word meanings to be coined and comprehended (e.g., *plate of wildebeest*) with minimal difficulty.

How might the patterns of flexibility discussed before – e.g., animal-for-meat and material-for-object – be represented within this view of the lexicon? One possibility is that all patterns are supported by lexical rules, such that the patterns themselves are specifically represented in the lexicon. However, this theory faces difficulty, because it assumes that generative forms of flexibility (e.g., animal-for-meat) will be represented in the same fashion as non-generative forms (e.g., material-for-object). As we discuss below, this assumption appears to be incorrect. Instead, we argue that only generative patterns of flexibility are represented by lexical rules, and that the senses that follow non-generative patterns must be stored.

Under the approach we favour, conceptual biases can be thought of as prior beliefs about how words are likely to be used. These beliefs guide the creation of an optimally structured lexicon: One that both stores word senses that are hard to predict, and employs lexical rules to represent and generate predictable, regular senses, obviating the need to store these senses. Under this account, lexical rules reflect conceptual relations but are still distinct from them, in the same way that morphological rules are distinct from conceptual relations, even though morphological rules reflect conceptually-important distinctions.

Critically, this approach to lexical structure assumes an important representational difference between generative “regular” flexibility, where a word’s senses are generated by lexical rules, and non-generative “irregular” flexibility, where a word’s senses are stored. Considerable experimental evidence now supports this distinction, although some caveats remain. This distinction makes two key experimental predictions. First, because the senses of “irregular” flexible words are separately stored, in much the same way as the meanings of homophones (e.g., “baseball bat” and “animal bat”), irregular senses and homophonous meanings should rely on very similar mechanisms and processes for access and use. By contrast, because the senses of “regular” flexible words are generated through rules, the representations and mechanisms that are used for processing these senses should be very different from those used for processing homophones and irregular flexible words.

Evidence for this distinction between “regular” and “irregular” flexible words has come from a variety of different paradigms. Some of the earliest work examined how such words are processed in context, based on a long-standing result in the psycholinguistics literature: When homonymous words are read in unconstraining contexts, readers tend to default to the most frequent interpretation of that word (e.g., the “toy” meaning of the word *ball* is more likely than the “party” meaning). This effect has been demonstrated as follows. If the context after an ambiguous word suggests its less-frequent meaning (e.g., the “party” meaning of *ball*), this causes a slowdown in participants’ reading speeds (a disambiguation effect, which can be measured via eye-tracking). Similarly, if prior context before an ambiguous word suggests its less frequent meaning, then participants will be slower to read that word, compared to when the context suggests its more frequent meaning (a so-called subordinate bias effect).

Interestingly, studies employing these paradigms have suggested that some flexible words may be represented and processed similarly to homophones such as *ball* (Foraker & Murphy, 2012), while other flexible words may be represented and processed differently

than homophones (Frazier & Rayner, 1990; Frisson, 2009; Frisson & Pickering, 1999). For example, Frazier and Rayner (1990) found that the disambiguation effect was strongly reduced for flexible words like *newspaper* or *book*, compared to homonymous words such as *ball* or *bat*. Similarly, Frisson and Pickering (1999) found that the subordinate bias effect was reduced for words that follow the place name for event pattern, such as the use of *Vietnam* to refer to the war in Vietnam. In contrast, Foraker and Murphy (2012) found large disambiguation effects for other flexible words, suggesting that their senses may be represented and processed similarly to distinct and unrelated homophones. Critically, a major difference between these studies was the flexible words that were used as stimuli: Frazier and Rayner, and Frisson and Pickering, used productive patterns of lexical flexibility, while Foraker and Murphy used a variety of words drawn from multiple patterns, many of which were not productive. Thus, taken together, the findings suggest that the senses of regular flexible words and the senses of irregular flexible words are represented in different ways, with irregular flexible words represented similarly to homophones.

Priming studies provide further evidence for this distinction. If word senses are stored separately, then the use of one sense should not prime access to the other sense, but if one sense is generated from the other sense, then they should prime each other. Using the same set of flexible words as Foraker and Murphy (2012), Klein and Murphy (2001) found that participants were faster to process simple phrases when they were primed by phrases using the same sense, and slower when the sense was switched; For instance, *polluted atmosphere* primes *thin atmosphere*, but inhibits the processing of *tense atmosphere*. Importantly, Klein & Murphy found similarly-sized effects for flexible words and homophones, consistent with the idea that these (mostly irregular) flexible words are stored as separate and unrelated words. Complementary work by Klepousniotou, Titone and Romero (2008), however, has suggested that for other flexible words, different senses do prime one another: They found that the sense-switching effect was greatly reduced for flexible words whose senses had a strong semantic relationship, and these words were almost uniformly “regular” (e.g., *young/juicy chicken*).

However, these priming studies should be interpreted with caution: It is unclear whether they indicate a true representational distinction between types of words, or instead differences in semantic associations between phrases. For example, while the word *polluted* in *polluted atmosphere* has a strong association with the gaseous sense of *atmosphere*, it has very little association with the social sense of *atmosphere*. By contrast, the word *young* in “young chicken” has a strong association with not only the animal sense of *chicken*, but is also reasonably associated with the meat sense of *chicken*. Thus, for words like *chicken*, the cost of switching senses may be reduced by associative priming. This potential influence of semantic associations makes it difficult to say whether or not previous priming effects indicate representational distinctions amongst types of lexical flexibility (Klepousniotou, Pike, Steinhauer, & Gracco, 2012; MacGregor, Bouwsema, & Klepousniotou, 2015).

Thankfully, evidence for a regular-irregular distinction holds under other paradigms. Rabagliati and Snedeker (2013) used a simple referential communication task, modeled after Ferreira, Slevc and Rogers (2005), who had discovered that when participants are asked to produce descriptions of target pictures surrounded by distractors, they provide informative descriptions when targets and distractors depict the same type of thing (e.g., if

both depicted chicken animals, the target might be called large chicken), but often provide uninformative descriptions when targets and distractors depict different things that share a homophonous name (e.g., if an animal bat and baseball bat were depicted, the target might only be called bat). Ferreira and colleagues interpreted these findings to suggest that participants do not notice ambiguous “bats” because homophonous meanings are represented separately. After replicating these findings, Rabagliati and Snedeker showed that participants also differentiate between regular and irregular flexible words. For regular flexible words like *chicken*, participants provided informative descriptions when both senses (i.e., a chicken animal and meat) were depicted as target and distractor, and to the same degree that they did when the target and distractor depicted the same type of thing (e.g., two different chickens). By contrast, for irregular flexible words like *glass*, participants often provided uninformative descriptions, describing a glass cup as *glass* when a piece of glass acted as a distractor, similar to how they behaved for homophones. Importantly, participants’ ability to produce informative descriptions was better explained by the distinction between regular and irregular flexibility than by independent ratings of whether the senses of these words were similar. These data therefore suggest that while irregular senses are separately stored, regular senses are generated through a shared, core meaning.

More direct evidence for the existence of generative mechanisms comes from developmental work by Srinivasan and Snedeker (2011; 2014). They used a novel paradigm: Children were taught words from a “muppet language” that labeled one sense of a flexible word, or one meaning of a homophone pair, and were assessed on whether they would be willing to extend those novel words to label the other sense or meaning. For instance, Elmo might label the physical object sense of the regular flexible word *book* as *blicket*, or label the animal meaning of the homophone *bat* as *wug*. Interestingly, 4- and 5-year-olds in this study were willing to extend the labels of flexible words, but not the labels of homophones. For example, they accepted that *long blicket* could describe a long story, but they would not accept that *black blicket* might label a black baseball bat (Srinivasan & Snedeker, 2011). More impressively, in other studies, children appeared to extend the senses of flexible words in a spontaneous and automatic fashion (Srinivasan & Snedeker, 2014; Srinivasan, Al-Mughairy, Foushee & Barner, 2017). One study has focused on children’s knowledge of instrument-activity flexibility. For example, after watching a video in which a character was shoveling some sand, and hearing the novel verb *dax* to describe it (e.g., She is *daxing* the sand), children spontaneously assumed that *dax* could also label the shovel itself (i.e., in line with flexible instrument-activity words like *shovel*, *hammer*, *bicycle*, etc.). In particular, after being taught the verbal, activity sense of *dax*, children were given a forced-choice task, where they had to decide whether a second word, *kiv*, referred to the shovel, or to the sand. Children reliably chose the sand, suggesting that 1) they spontaneously assumed that *dax* referred to the shovel, and 2) they had used the mutual exclusivity principle (Markman & Wachtel, 1988) to infer that *kiv* must therefore refer to the sand. Further studies indicated that children’s spontaneous generalization between senses isn’t limited to familiar words like *shovel*, but instead extends to entirely novel words that label novel activities and instruments (Srinivasan, Al-Mughairy, Foushee & Barner, 2017).

In sum, a variety of experimental data supports an important representational distinction between “regular” flexible words, whose senses are generated through lexical rules, and

“irregular” flexible words, whose senses are separately stored, just like homophonous meanings. This is consistent with a view of the lexicon in which word meanings are efficiently represented through a combination of storage and generative mechanisms.

The storage plus generativity account does contrast with some other accounts of how flexible word meanings are represented. For example, Pykkänen, Llinas and Murphy (2006) suggest that, while the senses of many flexible words may be stored, they are stored in a fundamentally different way from how the meanings of homophones are stored. In particular, senses are stored inside a single “root” form of a word, while homophonous meanings are stored in separate roots. As evidence for this, they show that the neurophysiological response to reading a word differs depending on whether it has multiple senses or multiple meanings. However, their paradigm used both irregular and regular flexible words as stimuli, and so it remains possible that this effect was only driven by the regular flexible words.

A quite different account comes from Rodd, Gaskell and Marslen-Wilson (2004), who describe a connectionist model of lexical representations in which the forms of words are linked to their meanings via a connectionist network. The model is interesting because it does not draw a distinction between storage and productive generalization, in analogous fashion to connectionist models of inflectional morphology (McClelland & Patterson, 2002). Instead, it learns links between distributed representations of a word’s phonetic form, and distributed representations of its meanings. These distributed meaning representations cause the network to behave quite differently when learning about flexible words and when learning about homophones. For flexible words, whose multiple senses overlap in terms of features, the model is easily able to learn a mapping from a word’s form to its meanings. For homophones, whose meanings do not share features, it is harder to activate a stable semantic representation, because its two meanings have very distinct patterns of activation, and they will compete more strongly.

Rodd and colleagues’ model has been successfully used to explain a surprising finding about how quickly people access the meanings of flexible words: lexical decision times are differentially affected by the number of flexible senses a word has, as opposed to its number of homophonous meanings (Rodd, Gaskell, & Marslen-Wilson, 2002). In particular, while decision times are reduced for words with more senses, they are increased for words with more homophonous meanings. However, this experiment did not distinguish between regular and irregular flexible words, and so it is possible that this effect was driven by the presence of regular flexible senses.

It is currently unclear whether the model of Rodd and colleagues can account for the various other findings about flexible word representations discussed so far. For instance, the model does not have an obvious mechanism for generalization; indeed, generalization outside of a training space is often problematic for connectionist models (Marcus, 2001). Additionally, when training their model, Rodd and colleagues did not use a realistic approximation of the similarity structure of word senses. Instead, they assumed that all of a word’s senses are random deviations from a single core meaning. As we discuss below, word senses tend to be related to one another, rather than similar to each other, and thus often do not share many surface features. It is unclear if this would affect the ease with

which the network maps between form and meaning. Nevertheless, Rodd and colleagues' model is important, and suggests how clear benefits can be gained from explicit computational modeling of theories. Other theories, such as the storage plus generativity model that we have discussed so far, would likely accrue additional benefits from such a well-specified implementation (see O'Donnell, 2015, for a possible framework in which such models could be implemented).

3. Why do words have distinct senses?

Thus far, we have characterized the types of word senses that exist across languages, and how these senses may be represented in the lexicon. Left untouched, however, is a deeper problem: why does natural language contain lexical flexibility in the first place? After all, lexical ambiguity is typically considered problematic for other types of languages, such as programming languages. Could lexical flexibility be a useful feature of natural language? We have already touched on one potential reason for the existence of lexical flexibility, i.e., that it may reflect conceptual organization. An alternative approach, however, takes a different tack, and emphasizes that lexical flexibility could play an important functional role in making language easier to process and to learn.

Functional accounts of language assume that linguistic characteristics, such as grammatical or lexical structure, are adaptations that are sculpted by two core pressures, communication and learning (Givón, 1995; Kirby, 1999; Regier et al., 2015). Pressures from communication arise from how a language is used in a community. Languages face pressures to be expressive – allowing many different meanings to be transmitted – and also face pressures to be efficient: in an optimal language, it should be easy to express any meaning (e.g., expressions should be concise), and hard to misinterpret any expression (due, for instance, to a noisy environment). Meanwhile, pressures from learning derive from how languages are transmitted to new members of a community: Languages that are easier to learn are more likely to be accurately recovered by learners, and then will be passed down to new learners in turn.

We suggest that lexical flexibility is likely to be an adaptation to both the pressures of communication, and the pressures of learning (see also Regier et al., 2015). As we detail below, either pressure alone seems unlikely to fully explain lexical flexibility, but the prevalence of lexical flexibility can be explained by the joint action of these two pressures. Moreover, the more precise characteristics of word sense relations – why words have the senses that they do – may be explained by how these pressures act to create a communication system that is suitably adapted to help express the types of concepts that adults and, especially, language-learning children are able to represent.

3.1 Functional pressures and lexical flexibility

Communicative pressures can be understood in terms of costs. For every word we speak or understand, we pay a small cost. This cost can be explicit, such as the metabolic energy required to articulate a syllable or to retrieve a word from memory. Or, it can constitute a more implicit, opportunity cost: By spending our time producing or comprehending words,

we take our efforts away from other, potentially more important tasks. The precise costs involved in using a word depend upon some of its features. For example, the time and energy that are required to articulate a word depend upon its length, while the time and energy required to retrieve a word from memory depend upon its frequency of use. Given that we spend so much of our lives speaking and listening, a rational response to these costs would be to develop a lexicon that minimizes them.

The communicative lexicon hypothesis (Piantadosi, Tily, & Gibson, 2012) argues that lexical ambiguity is one manifestation of that rational response. In particular, if it is the case that speakers and listeners prefer words that are short and frequent, but also need to use those words to communicate a variety of ideas, then one solution is to increase the number of meanings that each short and frequent word can be used to stand for. Such ambiguity might be assumed to cause confusion, another cost, but as long as the context around each word is sufficient to disambiguate which meaning is intended, then such a system will be communicatively efficient. Consistent with this, Piantadosi and colleagues found that, across languages, words that are shorter, that are more frequent, and that tend to occur in predictable contexts (which will make them easier to retrieve from memory) also tend to be associated with more meanings.

The communicative lexicon hypothesis is important and interesting, and provides a plausible reason for the prevalence of lexical ambiguity. However, it falls short of providing a full explanation of why lexical flexibility in particular is so commonplace in natural language, because it incorrectly predicts that flexible words should be less prevalent than homophony. In particular, the communicative lexicon approach predicts that it is only beneficial for words to have multiple meanings when those meanings can be easily resolved from context. Because homophones like *bat* have very distinct meanings, their different uses tend to occur in distinctive and informative contexts; as such, this type of ambiguity should be commonplace, because ambiguity resolution will be easy. Meanwhile the related senses of flexible words, like *chicken*, are likely to occur in more related contexts, precisely because they are related, and the more subtle differences between these contexts will make it easier to become confused about which sense is being used. However, and contrary to this prediction of the communicative lexicon hypothesis, flexible words are far more prevalent than homophones. For example, Rodd et al (2004) report that in the English Wordsmyth dictionary (Parks, Ray, & Bland, 1998), 7.4% of words can be classified as homonyms, while 84% of words can be classified as flexible.

To understand why lexical flexibility is more prevalent than homonymy, it may be useful to consider additional pressures on the evolution of the lexicon that go beyond the pressure to transmit information quickly. For example, language learners need to rapidly acquire their language's lexicon, and interlocutors often need to stretch their existing lexicon to express new ideas. These pressures may have shaped the lexicon over time, leading multiple senses to be packed into words in ways that minimize the costs associated with learning a lexicon and expressing novel ideas. Such costs could be minimized if the multiple senses of a word are highly associable and easy for learners and interlocutors to compute, as is the case for flexible words.

For example, when having a conversation, speakers may find it easier to express a new idea

by re-using an existing word form that has a related meaning (which may already be activated by the context), than by creating an entirely new word with that meaning, and listeners should find it easier to compute a new intended meaning if they can use a highly associated existing word sense to constrain their interpretation. Similarly, children may find it relatively easy to learn a flexible lexicon, because learning a label for one concept can give the child a cue as to how other related concepts might be labeled. By contrast, when the lexicon is unambiguous or only contains homophones, learning the label for one concept provides no hint as to how other concepts should be labeled. As an example, consider a child who has learned that the word *glass* is used to describe a drinking vessel. If that child wanders into the kitchen after a glass vase has shattered on the floor, then they may be able to guess what is being referred to if they are told “don’t go near the glass!” (referring to the material). By contrast, if the two senses of glass used different labels (e.g., “don’t go near the dax!”), then children would need to rely on context alone in order to determine what was being referred to.

Consistent with the idea that the lexicon has been shaped by the demands of expressing and learning new meanings, recent findings suggest that words have developed new senses over history in predictable ways (Xu, Malt & Srinivasan, 2017; Ramiro, Malt, Srinivasan & Xu, 2017). For example, drawing on an historical database of the English lexicon that dates back a millennium, Ramiro, Malt, Srinivasan & Xu (2017) developed computational models to predict how new senses of a word have “attached” to existing senses of a word over history, on the basis of semantic relations between senses. Strikingly, these models were able to recapitulate the order in which word senses have emerged over history better than chance. Notably, the most successful model at explaining the historical data employed an algorithm known as *chaining* (Lakoff, 1987), which ensures that novel senses develop from existing ones that are closest in meaning and which may thus minimize the collective costs associated with expressing and learning new word meanings.

Lexical flexibility may therefore allow interlocutors and learners to leverage their existing knowledge of how speakers use words, to improve their odds of understanding additional ways in which speakers use words. This ability will be useful during childhood, when the learner is first establishing their lexicon, and also later on, in learning words for newly invented concepts. For example, flexible neologisms, like using *google* to refer to searching the internet, may well have spread through the language because their meaning was easy to understand. Similarly, when a speaker does not know the correct label for a concept (e.g., a child trying to express the term for glass material, or an adult trying to label “searching the internet”), then they can use flexibility to guess what label they should use, and have a reasonable chance of either producing the correct form (if that concept does indeed have a label) or being understood (if that concept does not have a label, see Zhu, 2013; Zhu & Malt, 2014). In this way, flexibility allows learners to easily understand and produce new senses.

But lexical flexibility also provides a second, perhaps more powerful learning mechanism. Not only does flexibility make it easier to guess how people have used words, it also allows the learner to anticipate how people might use words flexibly in future. In particular, learners who have realized that some word senses participate in generative patterns can use those patterns to infer additional senses without ever being provided with direct evidence for those senses. For example, knowing that *dax* labels an activity the learner can

infer that *dax* also labels the instrument that enables that activity. Indeed, the studies described before by Srinivasan & Snedeker (2013) and Srinivasan, Al-Mughairy, Foushee & Barner (2017) provide evidence that young children can spontaneously infer new senses, allowing them to bypass observational learning.

The examples above show how the relative frequencies of flexibility and homophony are better justified by an account in which functional pressures come from both communication and learning, rather than just communication alone. In particular, functional pressures that come from communication should result in homophonous lexicons, which we rarely find, but the addition of pressures from learning should result in flexible lexicons, which we do find. In fact, this account even predicts the types of word senses that we should find in a well-adapted language. In particular, those word senses should be related enough to facilitate learning, but not so similar that they result in confusion and misinterpretation during communication.

The examples of flexibility seen in this chapter appear to fit that description. The senses of *glass* for example, are clearly related to one another (there is an obvious link between the material sense of glass and the drinking vessel sense), but this does not mean that they are similar, or drawn from the same category of things. A drinking glass has a number of characteristic properties: it holds liquid, it can be held in the hand, it can be easily balanced on a surface, and so on. The material sense of *glass* has none of these properties. Instead, it is a brittle, transparent substance, that can be used to make a variety of distinct artifacts (such as drinking glasses). Indeed, many flexible words appear to have senses that are related but not similar. Klein and Murphy (2001, 2002) found that the senses of flexible words were often given surprisingly low similarity ratings, and that, even for the most similar pairs of senses, subjects were unwilling to classify the two senses as describing the same type of thing. Similarly, Rabagliati et al (2011) found that relatedness, rather than similarity, was a better predictor of whether two plausible word senses existed.³

Behind this idea is the implicit assumption that related-yet-not-similar word senses are easy to learn, and in particular are easy for children to learn. After all, the majority of language learners are children. For this theory to be correct, it must be the case that children have no difficulty in learning word senses that are only related through a flexible conceptualization (e.g., conceptualizing *glass* as both a substance and an artifact) or that are related through an abstract link (e.g., a metaphorical relation, such as using *chicken* to describe a

³ We should note that similarity-based flexibility may in fact occur in human languages, but that this type of ambiguity tends to be classified as vagueness. Potential examples of borderline flexible/vague words can be found in semantic domains for which detailed cross-linguistic data has been collected, such as kinship terms. For example, in English the word *Aunt* is used to label the sisters of your parents, but other languages make finer grained distinctions: e.g., Korean distinguishes maternal aunts (*imo*) and paternal aunts (*gomo*). This suggests that the English *Aunt* may in some sense be polysemous, even if it is not classified as such by English lexicographers. Unfortunately, we do not currently have a good method or test for distinguishing when a word is polysemous from when a word is vague, apart from the existence of more fine-grained distinctions in other languages. As such, it is hard to evaluate this possibility's explanatory potential.

personality). This assumption requires investigation, because it demands that young children can flexibly reason about different concepts, e.g., to understand not only how the different senses of words like *glass* are related, but also to understand how these word senses are distinct.

3.2 Mechanisms and constraints for learning senses

Interestingly, there is considerable evidence that children are able to flexibly reason about, and conceptualize, objects, individuals, and events in the world around them, even before they master flexible words. For example, Gelman and Bloom (2000) demonstrated that 3-year-old children can construe the same object as an intentionally created artifact, or as a piece of material, depending on what background information they are given about that object. For instance, they will label a sharp piece of plastic as a “knife” if told that the plastic was intentionally-shaped, but they will label it as “plastic” if told that its shape arose accidentally. Similarly, two-year-old children, who have not acquired a syntactic mass/count distinction, can flexibly switch their attention between shape and material composition: When dealing with solid objects, they reason that shape is a more important predictor of kind membership than material composition, and vice versa when dealing with substances (Soja, Carey, & Spelke, 1991). Finally, children are able to reason about the same physical object as either an artifact or as an intentional agent, long before they have mastered intentional language, or even multi-word syntax. For instance, 12-month-old infants will follow the “gaze” of a faceless object if that object interacts with them by contingently reacting to their own movements and vocalisations (Johnson, Slaughter, & Carey, 1998). These data, and many other reports, support the idea that children make sense of the world through a small set of “framework theories” that allow them to explain observed properties of the world in different ways (Carey, 1985; Gopnik & Meltzoff, 1997; Keil, 1994). For example, these framework theories would allow children to understand chickens in terms of their form (e.g., they have wings, they have beaks), in terms of their purpose (e.g., they are bred for eating), or in terms of their intentional behavior (e.g., they are cowardly).

We suggest that these early-developing cognitive capacities allow children to learn flexible words. Indeed, it may well be the case that patterns of flexibility (like animal-for-meat) are concrete manifestations of how children apply different framework theories to the world, in order to make sense of coherent superordinate classes. For instance, animal-for-meat may reflect children’s ability to construe animals as both intentional agents and substances. An important corollary of this is that such cognitive biases will not only allow children to learn certain senses, but will also make it harder for children to learn other senses. Indeed, if children were not constrained in which kinds of word senses they could learn, one might expect that different languages would develop different kinds of lexical flexibility, since all forms of lexical flexibility could be learned and transmitted across generations. But the existence of cross-linguistic regularities in flexibility – documented by Srinivasan & Rabagliati (2015) – suggests that this cannot be the case.

Indeed, there is evidence for an important interaction between children’s cognitive biases, their ability to learn novel word senses, and the types of lexical flexibility that are found across languages. In particular, children’s conceptual biases may lead them to be better at

learning certain types of word senses compared to others, leading these more learnable senses to be more frequent across different languages. The best evidence for this idea comes from work on denominal verbs. Previous work has identified a number of distinct classes of denominal verbs (Clark & Clark, 1979), with these classes partially determined by the relationship between the meaning of the noun and the thematic role structure of the derived verb. For example, denominal verbs like *water* or *butter* are referred to as “goal” verbs, because their direct objects (*water the garden*, *butter the bread*) indicate a goal for the substance denoted by the noun (water is placed in the garden, butter is placed on the bread). By contrast, denominal verbs like *weed* or *milk* are referred to as “source” verbs, because their direct objects (*weed the garden*, *milk the cow*) indicate the source of the substance denoted by the noun (weeds are taken from the garden, milk is taken from the cow).

Interestingly, children find certain classes of denominal verbs easier to learn than other classes, and this is consistent not only with children’s cognitive biases, but also with the typological patterning of these verb classes. In particular, English-learning children find it more difficult to learn source verbs like *milk* or *weed*, than to learn goal verbs like *water* or *butter*. Srinivasan & Barner (2013) found that four- and five-year-olds often assume that source verbs denote the transfer of a substance to a goal – e.g., such that “milk the cow” entails putting milk onto a cow – and do so even when they understand that such meanings are implausible (e.g., because milk comes from a cow). Moreover, children show this goal bias not simply in their misinterpretations of familiar verbs, but also in their learning of novel verbs.

This bias toward learning goal verbs may reflect a broader and early-developing cognitive bias to focus on goals. It is well-established that a cognitive goal bias colours children’s event perception: When asked to describe events involving an agent moving from a source to a goal, both children and adults are more likely to mention the goal than the source (Lakusta & Landau, 2005, 2012). Even prelinguistic infants exhibit a goal bias, as they are more accurate in detecting changes to events that involve goals, as opposed to sources (Lakusta, Wagner, O’Hearn, & Landau, 2007). Although it has not been proven that a nonlinguistic, cognitive goal bias causes children to have a goal bias in learning denominal verbs, the coordination of these two biases seems unlikely to be a coincidence.

Moreover, and consistent with the broader idea that easy-to-learn senses are more likely to be present cross-linguistically, source verbs appear to be less typologically robust than goal verbs. In Srinivasan and Rabagliati’s (2015) cross-linguistic survey of flexibility (described before), source verbs were the least attested across different languages, and were the only pattern of flexibility that was absent more often than expected due to sampling error. As such, the distinction between source and goal verbs – in how they are learned by children, and in their relative presence across languages – provides intriguing evidence for the idea that children’s learning biases shape lexical flexibility.

Future work should explore whether children have learning biases that make sense of other features of cross-linguistic variation in lexical flexibility. For example, as discussed earlier, there are clear patterns in whether verbs relating to one sensory modality can be extended

to describe another modality or not (Viberg, 1983). This generates predictions about asymmetries in learnability that could be tested in the future.

4. Discussion

Why do words have the senses that they do? In this chapter, we have argued that the senses of flexible words can be understood from at least three perspectives. They can be understood as reflections of conceptual biases and culture-specific conventions, as representations in the mental lexicon, and as mechanisms that allow languages to be more learnable, and more efficient for communication. While these three perspectives are in many ways independent, they also interlink in important ways.

For example, we have argued that word senses are culture-specific conventions, constrained by conceptual biases, and that these biases have their roots in children's earliest cognitive development. We began this chapter by discussing the existence of patterns of flexibility in English, of which there are many. Across many languages we found that, while word senses may vary, which indicates that senses are cultural conventions, nevertheless these same patterns appear to recur across languages. We have suggested that these patterns may reflect the flexible ways in which even children are able to reason about the world. For example, children's ability to flexibly reason about form, function, or intention could allow them to construe individuals, objects or events in multiple different ways, and thus to understand links between word senses that are related but not similar. Under this proposal, therefore, the same intuitive theories that explain how children acquire concepts might explain how children learn about the meanings of flexible words, and thus influence which patterns of flexibility tend to recur across languages. We suggest that patterns of lexical flexibility represent the outcome of applying these different intuitive framework theories to the world, in order to make sense of coherent superordinate classes. For example, the animal-for-meat pattern may reflect children's ability to construe the animals as both agents and substances (Johnson et al., 1998), while the material-for-artifact pattern may reflect children's ability to reason about both an object's material, and its intended form (Soja et al., 1991).

As another example of how different perspectives toward lexical flexibility complement one another, we have argued that flexible words are represented as a combination of stored senses and generative lexical rules. These lexical rules are important, as they allow children to enlarge their expressive vocabulary, by spontaneously guessing how novel concepts might be labeled (e.g., on learning that an animal is called a *dax*, children can guess that its meat is also called *dax*). This example shows how a consideration of the potential functions of lexical flexibility (in this case, helping children to master the conventionalized lexicon of their language) can provide reasons for preferring certain models of how words are mentally represented over others. In particular, the utility (and existence) of spontaneous generalization in young children suggests that their lexicons must possess mechanisms for generating new senses, such as lexical rules.

These examples suggest that a broad perspective is necessary to fully understand lexical flexibility and word senses: A perspective under which theories are examined from both structuralist and functionalist perspectives, and where data is gathered from a variety of

sources: typological investigation, psycholinguistic experiments, and developmental studies. We believe that this type of broad, interdisciplinary approach will be fruitful in elucidating other facets of semantics and pragmatics as well.

References

- Andersen, E. S. (1978). Lexical universals of body-part terminology. In J. H. Greenberg, C. A. Ferguson, & E. A. Moravcsik (Eds.), *Universals of human language* (Vol. 3, pp. 335-368). Stanford, CA: Stanford University Press.
- Apresjan, J. (1974). Regular polysemy. *Linguistics*, 142, 5-32.
- Berlin, B., & Kay, P. (1969). *Basic color terms: Their universality and evolution*. Berkeley, CA: Univ of California Press.
- Borer, H. (2005). *Structuring sense*: Oxford University Press.
- Carey, S. (1985). *Conceptual change in childhood*. Cambridge, MA: MIT Press.
- Clark, E. V., & Clark, H. H. (1979). When Nouns Surface as Verbs. *Language*, 55(4), 767-811.
- Copestake, A., & Briscoe, E. J. (1995). Semi-productive polysemy and sense extension. *Journal of Semantics*, 12(1), 15-67.
- Falkum, I. L. (2015). The how and why of polysemy: A pragmatic account. *Lingua*, 157, 83-99.
- Ferreira, V. S., Slevc, L. R., & Rogers, E. S. (2005). How do speakers avoid ambiguous linguistic expressions? *Cognition*, 96(3), 263-284.
- Foraker, S., & Murphy, G. L. (2012). Polysemy in sentence comprehension: Effects of meaning dominance. *Journal of Memory and Language*, 67(4), 407-425.
- Frazier, L., & Rayner, K. (1990). Taking on semantic commitments: Processing multiple meanings vs. multiple senses. *Journal of Memory and Language*, 29(2), 181-200.
- Frisson, S. (2009). Semantic underspecification in language processing. *Language and Linguistics Compass*, 3(1), 111-127.
- Frisson, S., & Pickering, M. J. (1999). The processing of metonymy: Evidence from eye movements. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 25(6), 1366-1383.
- Gelman, S. A., & Bloom, P. (2000). Young children are sensitive to how an object was created when deciding what to name it. *Cognition*, 76(2), 91-103.
- Givón, T. (1995). *Functionalism and grammar*: John Benjamins Publishing.
- Gopnik, A., & Meltzoff, A. N. (1997). *Words, thoughts, and theories* (Vol. 1). Cambridge, MA: MIT Press.
- Hampton, J. A. (1995). Testing the prototype theory of concepts. *Journal of Memory and Language*, 34(5), 686-708.
- Johnson, S., Slaughter, V., & Carey, S. (1998). Whose gaze will infants follow? The elicitation of gaze-following in 12-month-olds. *Developmental Science*, 1(2), 233-238.
- Kay, P., & Regier, T. (2006). Language, thought and color: recent developments. *Trends Cogn Sci*, 10(2), 51-54.
- Keil, F. C. (1994). Explanation, association, and the acquisition of word meaning. *Lingua*, 92, 169-196.
- Kirby, S. (1999). *Function, selection, and innateness: The emergence of language universals*: OUP Oxford.

- Klein, D. E., & Murphy, G. L. (2001). The representation of polysemous words. *Journal of Memory and Language*, 45(2), 259-282.
- Klein, D. E., & Murphy, G. L. (2002). Paper has been my ruin: conceptual relations of polysemous senses. *Journal of Memory and Language*, 47(4), 548-570. doi:Pii S0749-596x(02)00020-7
- Klepousniotou, E., Pike, G. B., Steinhauer, K., & Gracco, V. (2012). Not all ambiguous words are created equal: An EEG investigation of homonymy and polysemy. *Brain and language*, 123(1), 11-21.
- Klepousniotou, E., Titone, D., & Romero, C. (2008). Making sense of word senses: The comprehension of polysemy depends on sense overlap. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(6), 1534-1543.
- Lakoff, G., & Johnson, M. (1980a). The metaphorical structure of the human conceptual system. *Cognitive Science*, 4(2), 195-208.
- Lakoff, G., & Johnson, M. (1980b). *Metaphors we live by*. Chicago, IL: University of Chicago Press.
- Lakoff, G., (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago: University of Chicago Press.
- Lakusta, L., & Landau, B. (2005). Starting at the end: The importance of goals in spatial language. *Cognition*, 96(1), 1-33.
- Lakusta, L., & Landau, B. (2012). Language and memory for motion events: Origins of the asymmetry between source and goal paths. *Cognitive Science*, 36(3), 517-544.
- Lakusta, L., Wagner, L., O'Hearn, K., & Landau, B. (2007). Conceptual foundations of spatial language: Evidence for a goal bias in infants. *Language Learning and Development*, 3(3), 179-197.
- Lehrer, A. (1990). Polysemy, conventionality, and the structure of the lexicon. *Cognitive Linguistics*, 1, 207 - 246.
- MacGregor, L. J., Bouwsema, J., & Klepousniotou, E. (2015). Sustained meaning activation for polysemous but not homonymous words: Evidence from EEG. *Neuropsychologia*, 68, 126-138.
- Majid, A. (2010). Words for parts of the body. In B. C. Malt & P. Wolff (Eds.), *Words and the mind: How words capture human experience* (pp. 58-71). Oxford: OUP.
- Malt, B. C., & Majid, A. (2013). How thought is mapped into words. *Wiley Interdisciplinary Reviews: Cognitive Science*, 4(6), 583-597.
- Marcus, G. F. (2001). *The algebraic mind: Integrating connectionism and cognitive science*. Cambridge, MA: MIT press.
- Markman, E. M., & Wachtel, G. F. (1988). Children's use of mutual exclusivity to constrain the meanings of words. *Cognitive psychology*, 20(2), 121-157.
- McClelland, J. L., & Patterson, K. (2002). Rules or connections in past-tense inflections: What does the evidence rule out? *Trends Cogn Sci*, 6(11), 465-472.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological review*, 85(3), 207-238.
- Murphy, G. L. (2007). Parsimony and the psychological representation of polysemous words. In M. Rakova, G. Petho, & C. Rákosi (Eds.), *The cognitive basis of polysemy*. Frankfurt am main, Germany: Peter Lang Verlag.
- Murphy, G. L., & Medin, D. L. (1985). The Role of Theories in Conceptual Coherence. *Psychological review*, 92(3), 289-316.

- Nunberg, G. (1979). The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy*, 3(2), 143-184.
- Nunberg, G. (2004). The pragmatics of deferred interpretation. In L. Horn & G. Ward (Eds.), *Handbook of Pragmatics* (pp. 344 - 364). Oxford, England: Blackwell.
- O'Donnell, T. J. (2015). *Productivity and reuse in language: A theory of linguistic computation and storage*: MIT Press.
- Papafragou, A. (1996). On metonymy. *Lingua*, 99(4), 169-195.
- Piantadosi, S. T., Tily, H., & Gibson, E. (2012). The communicative function of ambiguity in language. *Cognition*, 122(3), 280-291.
- Pustejovsky, J. (1995). *The Generative Lexicon*. Cambridge, MA: MIT Press.
- Pylkkänen, L., Llinas, R., & Murphy, G. L. (2006). The representation of polysemy: MEG evidence. *J Cogn Neurosci*, 18(1), 97-109.
- Rabagliati, H., Marcus, G. F., & Pylkkänen, L. (2011). Rules, Radical Pragmatics and Restrictions on Regular Polysemy. *Journal of Semantics*, 28(4), 485-512.
- Rabagliati, H., & Snedeker, J. (2013). The Truth About Chickens and Bats Ambiguity Avoidance Distinguishes Types of Polysemy. *Psychological Science*, 24(7), 1354-1360.
- Regier, T., Kemp, C., & Kay, P. (2015). Word Meanings across Languages Support Efficient Communication. In B. MacWhinney (Ed.), *The handbook of language emergence* (Vol. 87, pp. 237).
- Roberson, D., Davidoff, J., Davies, I. R., & Shapiro, L. R. (2005). Color categories: Evidence for the cultural relativity hypothesis. *Cognitive psychology*, 50(4), 378-411.
- Rodd, J. M., Gaskell, G., & Marslen-Wilson, W. (2002). Making sense of semantic ambiguity: Semantic competition in lexical access. *Journal of Memory and Language*, 46(2), 245-266.
- Rodd, J. M., Gaskell, M. G., & Marslen-Wilson, W. D. (2004). Modelling the effects of semantic ambiguity in word recognition. *Cognitive Science*, 28(1), 89-104.
- Rosch, E., & Mervis, C. (1975). Family resemblance: Studies in the internal structure of categories. *Cognitive psychology*, 7, 573-605.
- Soja, N. N., Carey, S., & Spelke, E. S. (1991). Ontological Categories Guide Young Children's Inductions of Word Meaning - Object Terms and Substance Terms. *Cognition*, 38(2), 179-211.
- Srinivasan, M., Al-Mughairy, S., Foushee, R., & Barner, D. (2017). Learning language from within: Children use semantic generalizations to infer word meanings. *Cognition*, 159, 11-24.
- Srinivasan, M., & Barner, D. (2013). The Amelia Bedelia effect: World knowledge and the goal bias in language acquisition. *Cognition*, 128(3), 431-450.
- Srinivasan, M., & Rabagliati, H. (2015). How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy. *Lingua*, 157, 124-152.
- Srinivasan, M., & Snedeker, J. (2011). Judging a book by its cover and its contents: The representation of polysemous and homophonous meanings in four-year-old children. *Cognitive psychology*, 62(4), 245.
- Srinivasan, M., & Snedeker, J. (2014). Polysemy and the Taxonomic Constraint: Children's Representation of Words that Label Multiple Kinds. *Language Learning and Development*, 10(2), 97-128.
- Xu, Y., Malt, B., & Srinivasan, M. (2017). Evolution of word meanings through metaphorical mapping: Systematicity over the past millenium. *Cognitive Psychology*, 96, 41-53.

- Ramiro, C., Malt, B., Srinivasan, M. & Xu, Y., (2017). Mental algorithms in the historical emergence of word meanings. In *Proceedings of the 39th Annual Conference of the Cognitive Science Society*, London, UK.
- Viberg, Å. (1983). The verbs of perception: a typological study. *Linguistics*, 21(1), 123-162.
- Wierzbicka, A. (2007). Bodies and their parts: An NSM approach to semantic typology. *Language Sciences*, 29(1), 14-65.
- Wilson, D. (2003). Relevance and lexical pragmatics. *Italian Journal of Linguistics*, 15, 273-292.
- Zhu, H. (2013). *A Cognitive and Cross-Linguistic Approach to Polysemy*. (Ph.D), Lehigh University, Bethlehem, PA.
- Zhu, H., & Malt, B. C. (2014). *Cross-linguistic evidence for cognitive foundations of polysemy*. Paper presented at the Proceedings of the 36th Annual Conference of the Cognitive Science Society. Cognitive Science Society, Austin, TX.