UC Berkeley

UC Berkeley Previously Published Works

Title

How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy

Permalink

<https://escholarship.org/uc/item/69z8f7qf>

Authors

Srinivasan, Mahesh
Rabagliati, Hugh

Publication Date

2014-12-01

Peer reviewed

Available online at www.sciencedirect.com

ScienceDirect

Lingua xxx (2015) xxx–xxx

Lingua

www.elsevier.com/locate/lingua

How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy

Mahesh Srinivasan^{a,1,*}, Hugh Rabagliati^{b,c,1,**}

^a University of California, Berkeley, United States

^b Harvard University, United States

^c University of Edinburgh, United States

Received 27 September 2013; received in revised form 5 December 2014; accepted 6 December 2014

Abstract

Words often have multiple distinct but related senses, a phenomenon called polysemy. For instance, in English, words like *chicken* and *lamb* can label animals and their meats while words like *glass* and *tin* can label materials and artifacts derived from those materials. In this paper, we ask why words have some senses but not others, and thus what constrains the structure of polysemy. Previous work has pointed to two different sources of constraints. First, polysemy could reflect conceptual structure: word senses could be derived based on how ideas are associated in the mind. Second, polysemy could reflect a set of arbitrary, language-specific conventions: word senses could be difficult to derive and might have to be memorized and stored. We used a large-scale cross-linguistic survey to elucidate the relative contributions of concepts and conventions to the structure of polysemy. We explored whether 27 distinct patterns of polysemy found in English are also present in 14 other languages. Consistent with the idea that polysemy is constrained by conceptual structure, we found that almost all surveyed patterns of polysemy (e.g., animal for meat, material for artifact) were present across languages. However, consistent with the idea that polysemy reflects language-specific conventions, we also found variation across languages in how patterns are instantiated in specific senses (e.g., the word for glass material is used to label different glass artifacts across languages). We argue that these results are best explained by a “conventions-constrained-by-concepts” model, in which the different senses of words are learned conventions, but conceptual structure makes some types of relations between senses easier to grasp than others, such that the same patterns of polysemy evolve across languages. This opens a new view of lexical structure, in which polysemy is a linguistic adaptation that makes it easier for children to learn word meanings and build a lexicon.

© 2014 Elsevier B.V. All rights reserved.

Keywords: Polysemy; Pragmatics; Lexical semantics; Cross-linguistic variation; Conceptual development; Theory theory

1. Introduction

All natural languages include a repertoire of words to express a large set of basic ideas, from concrete concepts of animals, objects, and materials, to more abstract notions like events and beliefs. Interestingly, however, rather than labeling each idea with a unique word, languages systematically group sets of related ideas – or *senses* – under a single

* Corresponding author at: Department of Psychology, University of California, Berkeley, Berkeley, CA 94720-1650, United States.

** Corresponding author at: Department of Psychology, University of Edinburgh, Edinburgh EH1 3NS, UK. Tel.: +44 1316513454.

E-mail addresses: srinivasan@berkeley.edu (M. Srinivasan), hugh.rabagliati@ed.ac.uk (H. Rabagliati).

¹ The authors contributed equally to the project.

word, a phenomenon called polysemy (Breál, 1897). Polysemy is important not only because it is ubiquitous, but also because it provides a source of linguistic creativity: to express new ideas, we needn't invent new words, but can instead extend existing words beyond their original meanings. In English, such creativity has yielded systematic patterns of senses: for instance, the same words are often used to label an animal or its meat (e.g., *chicken*, *lamb*, etc.), or a material and an artifact derived from that material (e.g., *glass*, *tin*, etc.).

Here, we explore what representations and processes might account for the structure of polysemy, i.e., for how word meanings are extended, and thus why senses are grouped together in particular ways. Previous work has suggested two potential sources of constraints on polysemy. One line of work has focused on the role of conceptual structure, and has suggested that the ways in which senses are grouped together reflect the relations we perceive between different ideas, given the situational context (e.g., Fauconnier, 1985; Nunberg, 1979, 1995; Papafragou, 1996; Wilson, 2003). Thus, *chicken* may have animal and meat senses because we find the relation between the animal and its meat particularly noteworthy or salient. A second line of work has focused on the role of conventions, and has suggested that because many word senses are related in seemingly arbitrary and opaque ways, they must each be learned and stored within the mental lexicon (e.g., Klein and Murphy, 2001; Lehrer, 1990; Murphy, 1997, 2007; Pinker, 2007). For example, the fact that the English word *glass* labels a glass drinking vessel – rather than a glass window or mirror – may be a relatively arbitrary fact that we have to learn.

The relative contributions of concepts and conventions to the structure of polysemy have important implications for the relationship between conceptual and lexical structure, but there is currently little consensus as to what those contributions are. The present paper aims to provide critical data to remedy the situation, by documenting cross-linguistic regularity and variation in polysemy. Broadly speaking, if polysemy is tightly constrained by conceptual structure, it should manifest quite similarly across languages, but if polysemy corresponds to arbitrary lexicalized conventions, it should be quite variable across languages. Based on our data, we will argue that polysemy is best explained by a model that incorporates both concepts and conventions, and in particular by a model in which the senses of polysemous words are learned conventions that are shaped by the cognitive biases of learners.

Below, we set out and discuss possible constraints on polysemy, and explain how they might influence variation in polysemy across languages. Then, we review findings from previous cross-linguistic studies of polysemy, and present our own large-scale cross-linguistic study.

1.1. Constraints on polysemy

Two important features characterize the structure of polysemy, at least in English. First, linguists have identified a number of systematic patterns of polysemy, wherein multiple words have sets of senses that are related in similar ways (see, e.g., Copestake and Briscoe, 1995; Lakoff and Johnson, 1980; Ostler and Atkins, 1992; Pustejovsky, 1995). Table 1 presents examples of some of these patterns. As can be seen, they often include sets of senses that cross different semantic categories, alternately labeling people, animals, objects, substances, actions, and more. Some of these patterns invoke metaphorical relations, such as when body part names are used to label parts of objects (e.g., “the chair's *arm* is broken”). These patterns sometimes also include sets of senses that cross lexical categories, as when words are used as nouns to label objects and substances, and used as verbs to label actions involving those objects (e.g., “He *buttered* the bread”, “She *shoveled* the snow”).²

The second important feature of polysemy is that these patterns vary in how freely they permit generalizations. In a number of cases, patterns can be easily extended to create new senses, with minimal supporting linguistic and extra-linguistic context. The animal for meat pattern provides a good example of this. We can easily extend this pattern to label the meat of animals that aren't typically thought of as edible. Thus, it sounds natural (though culinarily odd) to say “he ate some *seagull*.” However, not every pattern is similarly generative. While *glass* and *tin* both describe materials and artifacts, it sounds distinctly odd to say “He bought a *plastic*,” even though we know that plastic is a material out of which many artifacts are made. This use of *plastic* would seem to require significantly more contextual support to be felicitous (much like contextual innovations such as “ham sandwich”, see footnote 2). These two types of patterns – generative and non-generative patterns – are typically referred to in the literature as *regular* and *irregular* polysemy, respectively (see e.g., Apresjan, 1974; Ostler and Atkins, 1992).

² Importantly, these examples of polysemy require minimal linguistic context to be felicitous, and do not depend heavily on the extra-linguistic context or prior discourse for their meanings. As such, these examples are typically distinguished from *contextual innovations*, such as the creative use of “ham sandwich” to label a restaurant patron who ordered a ham sandwich (Nunberg, 1979). Contextual innovations have provided evidence that we can stretch word meanings quite dramatically, by reasoning pragmatically and drawing on the linguistic and extra-linguistic context. However, it is a matter of debate as to whether the mechanisms underlying our interpretation of contextual innovations like “ham sandwich” also support the relatively more context-independent examples of polysemy provided in Table 1, which will be our focus here.

Table 1
Examples of polysemy in English.

Patterns and their Senses	Examples
Animal for Meat (<i>chicken, turkey, fish</i> , etc.)	The <i>chicken</i> walked on the grass/ The <i>chicken</i> was well-salted
Material for Artifact (<i>glass, tin, iron</i> , etc.)	There is broken <i>glass</i> on the floor/ She drank milk from the <i>glass</i>
Object for Representational Content (<i>book, magazine, DVD</i> , etc.)	The <i>book</i> is very light to carry/ The <i>book</i> is very interesting
Container for Contents (<i>pot, bowl, box</i> , etc.)	She washed the <i>pot</i> after dinner/ She stirred the <i>pot</i> with a spoon
Body Part for Object Part (<i>leg, arm, back</i> , etc.)	He broke his <i>leg</i> last year/ That chair has a broken <i>leg</i>
Artist for Product (<i>Picasso, Camus, Mozart</i> , etc.)	<i>Picasso</i> was born in 1881/ That museum has a <i>Picasso</i>
Place for Institution (<i>White House, Wall Street, City Hall</i> , etc.)	The <i>White House</i> is being renovated/ The <i>White House</i> should make a decision
Place for Event (<i>Vietnam, Waterloo, Woodstock</i> , etc.)	<i>Vietnam</i> shares a border with China/ He championed civil rights during <i>Vietnam</i>
Substance for Placing Substance at Goal (<i>butter, salt, water</i> , etc.)	He bought some <i>butter</i> from the store/ He is going to <i>butter</i> the bread
Instrument for Action Involving Instrument (<i>shovel, hammer, rake</i> , etc.)	She has a red <i>shovel</i> / She is going to <i>shovel</i> the snow

How can we account for why some sets of senses follow patterns, and why some patterns are generative? As noted before, some work has focused on the role of conceptual structure, and has suggested that the senses of polysemous words can be derived according to relations we find noteworthy or salient (e.g., Fauconnier, 1985; Nunberg, 1979, 1995; Papafragou, 1996; Wilson, 2003; Wilson and Carston, 2007), while other work has focused on the role of linguistic conventions, and has suggested that because senses are arbitrary and opaque, they must be memorized and stored within the lexicon (e.g., Klein and Murphy, 2001; Lehrer, 1990; Murphy, 1997, 2007; Pinker, 2007). Below, we describe each of these ideas, and discuss whether either of them can account for the structure of polysemy on their own.

If conceptual structure provides tight constraints on polysemy, speakers may not need to store all of the individual senses of polysemous words in memory, but could instead derive these senses on-line, from a single represented meaning. Ruhl (1989) provides perhaps the most extreme example of this idea, arguing that most polysemous words actually only have a single core meaning that captures the essence of the concept, and that can be adjusted to suit the surrounding context. For example, a core meaning of *glass* could denote a material that can be used to form solid objects, and context would then be used to fill in the details, such as whether the word is being used to label the material itself (as in “He bought a sheet of *glass*”), or instead an object composed of that material (as in “He poured water into the *glass*”).

In order to make contextual adjustments to core meanings, listeners and speakers could reason pragmatically, based on their general knowledge of the world, the linguistic context and discourse, and their knowledge of the intentions of interlocutors. The structure of polysemy could therefore be, in part, a function of the structure and content of concepts. For example, Nunberg (1995) proposes that senses can be derived based on a principle of noteworthiness: when there is a noteworthy conceptual relationship between a core meaning and another possible sense, that sense becomes plausible.³ For example, one noteworthy aspect of a glass drinking vessel might be its material composition, allowing the word *glass* to label both the material and artifact within the appropriate contexts (e.g., “broken *glass*” versus “drinking *glass*”). Critically, then, language users may not need to learn the different senses of polysemous words like *glass*: instead, different senses could be derived using world knowledge and conceptual relations like noteworthiness.

Conceptual structure could account not only for how senses are derived from core meanings, but also for the fact that sets of senses are organized into patterns, and that some of these patterns are generative. For example, under Nunberg's theory, a pattern arises when the same noteworthy relationship is seen over and over again in the world. Thus, just as the material composition of glasses may be noteworthy and explain the different senses of *glass*, the material composition of tins, sponges, and irons may also be noteworthy and explain the material and artifact senses corresponding to *tin*, *sponge*, and *iron*. Generativity could also follow naturally from conceptual structure because a new sense of a word could be derived when it stands in a noteworthy relationship to the core meaning of the word. Thus, when a relationship is

³ To explain how senses are derived, other, similar theories appeal to factors such as relevance (Papafragou, 1996; Falkum, 2011; Wilson, 2003; Wilson and Carston, 2007), idealized cognitive models (Fauconnier, 1985; Lakoff, 1987), and cue validity (Nunberg, 1979).

noteworthy in a novel context – e.g., that some meat being served is from a seagull – a novel sense can be coined (e.g., “That *seagull* could use some salt”).

However, one problem for the idea that conceptual structure constrains polysemy is that, because there are many noteworthy relationships between concepts, conceptual structure may fail to account for the relatively constrained nature of polysemy. For example, while conceptual factors like noteworthiness may help explain the senses we do use, they have trouble explaining why we do *not* use many other senses. For instance, in English, *glass* only labels one kind of artifact (i.e., drinking vessels), even though there are many other artifacts that are also noteworthy for being composed of glass material (e.g., mirrors, windows, etc.). To take another example, although English permits the use of names for animals to label their meats, it does not allow animal names to label products associated with those animals, outside of their fur or meat. Thus, although it seems noteworthy that eggs are laid by chickens, *chicken* cannot label an egg. These facts suggest that a theory that appeals only to conceptual structure may be too unconstrained to explain how polysemy is actually realized. Indeed, consistent with this, Rabagliati et al. (2011) tested whether a number of conceptual metrics could predict the acceptability of different possible senses, and found that these metrics were poor predictors of acceptability judgments.

Citing the apparent arbitrariness of polysemy, other work has proposed that some word senses cannot be derived via conceptual structure, and are instead *conventions* that members of a linguistic community must learn, one-by-one (Lehrer, 1990; Murphy, 1997, 2007; Pinker, 2007). By this account, senses are initially coined by an individual speaker and learned by individual listeners. When they are useful for members of the linguistic community, they become more frequent and more widely used. Conceptual structure may therefore play only a limited role in constraining polysemy, and place weak constraints on how new senses are coined: speakers and listeners have to grasp the relationship between the new and old sense, but otherwise, these senses could be related in any number of ways. Thus, there may be no principled reason as to why, for example, *glass* labels drinking vessels as opposed to windows, or why *chicken* labels chicken meat but not an egg: these are merely facts about language that speakers must master.

The idea that senses are learned conventions provides a natural explanation for why we only use words with specific senses and not others. In particular, if a sense has not been coined and learned, then it should not be used (like the use of *glass* in English to label a window). In addition, the establishment of certain senses in a language might “pre-empt” the emergence of other possible senses. For instance, because speakers have already learned to use *glass* to label one artifact in English (a drinking vessel), they might be wary of using it to label another artifact (i.e., the drinking vessel sense may pre-empt emergence of the window sense). Critically, this mechanism of pre-emption is not consistent with theories in which all senses are derived on-line via conceptual structure, because pre-emption depends on some senses being conventional, and stored within the lexicon.

Further evidence for conventionalized senses comes from the phenomenon of semantic drift: Once a new sense for a word has been memorized and has entered the language, its relationship to other senses does not need to remain transparent. This means that senses can remain in a language long after the initial communicative motivation that created them has died away. An example of this is the fossilized use of *iron* to describe a tool for pressing clothes. When this sense of *iron* was first coined, clothes were mainly pressed using large pieces of flattened iron (i.e., flatirons). However, that technology is now obsolete, and *iron* can now be used to label pressing machines that do not contain any iron. Similarly, it is possible for *glasses* to be made of plastic, to *land* on water, and to *shelve* books on a windowsill (see Clark and Clark, 1979; Kiparsky, 1997).

However, while conventions do a good job of explaining arbitrary properties of polysemy, they do not naturally account for the two features of polysemy that we laid out earlier. In particular, the presence of patterns, in which sets of senses are related in similar ways, does not naturally follow from a theory in which there are only weak conceptual constraints on the senses that can be formed. Further, the presence of generativity is unexpected if senses have to be individually stored in memory and cannot be derived.

In response to these points, Murphy (2007) has speculated that patterns emerge when new senses of a word are coined via analogy to existing senses of other words. For example, the inspiration for referring to seagull meat using its animal name could come from comparisons to the existing animal and meat senses of *chicken* or *fish*. Additionally, if a language has many words whose senses are related in similar ways, speakers may form generative rules by analogical comparison, allowing language users to produce and understand novel senses of words following the same patterns (e.g., such that new animal names can label meat; see also Copestake and Briscoe, 1995; Strigin, 1998). However, this proposal has not been directly tested, and thus remains speculative.

To sum up, there are reasons to think that both concepts and conventions could play important roles in explaining the structure of polysemy, but that neither of these sources of constraints are sufficient on their own. In particular, while conceptual structure provides a natural explanation for why sets of senses form systematic patterns and can be extended to create new senses, it has trouble explaining other, seemingly arbitrary properties of polysemy. Conversely, while conventions make sense of the arbitrary aspects of polysemy, they do not provide a natural account for why senses form generative patterns.

The above discussion points toward a middle-ground: both conceptual structure *and* conventions may be involved to some degree in explaining polysemy. In particular, senses of polysemous words may indeed be learned as conventions – rather than derived online using conceptual structure – which would explain arbitrary properties of polysemy, such as the pre-emption of possible but unattested senses, and semantic drift. But the process by which these senses are coined and learned could itself be shaped by conceptual structure, accounting for why senses form specific, generative patterns. For example, on one possible version of this model, conceptual structure could have its effect when senses are learned.⁴ In particular, if concepts bias learners to find some sets of senses to be easier to acquire than others, then this would limit which sets of senses enter the language, resulting in the formation of patterns of polysemy that loosely correspond to aspects of conceptual structure.

This particular combination of concepts and conventions – which we will describe in detail later – provides a compelling account of the structure of polysemy. However, at present, there is little direct evidence that supports it, relative to models that invoke only concepts or only conventions. As we describe below, cross-linguistic regularity and variation in polysemy could provide evidence to adjudicate between the different models.

1.2. Cross-linguistic predictions of different models of polysemy

In this section we describe the distinct predictions that different models make about cross-linguistic variation in polysemy. In particular, we contrast a model in which concepts alone explain the structure of polysemy (“concepts-only model”), with one in which conventions alone do this work (“conventions-only model”), and with one that relies on both conventions and concepts (“conventions-constrained-by-concepts model”).

If the senses of polysemous words are directly derived from conceptual structure and communicative context, then we would expect a great deal of uniformity in how polysemy is expressed across different languages. In particular, because speakers of different languages are likely to share broadly similar conceptual repertoires and world knowledge, they should generally find the same conceptual relationships to be salient or noteworthy and so, in similar communicative situations, they should derive senses from core meanings in similar ways. This would predict that the same patterns of polysemy (e.g., animal for meat, material for product, etc.) should be present in different languages, because these patterns conform to types of conceptual relations between senses that should be similarly noteworthy across linguistic communities. Further, it would also predict that within a particular pattern, similar sets of senses should exist across languages, because these senses all stand in equally noteworthy relationships to their core meanings. For example, the relations between glass material and a drinking vessel, and between tin material and a cookie tin, should be similarly noteworthy for speakers of different languages, such that across languages, words for glass material and tin material should also have the corresponding artifact senses. Of course, cultural and technological differences between speakers of different communities may cause some variation. For instance, cultures that do not use glass vessels to drink from would not have that particular sense of *glass*. But, by and large, if conceptual structure accounts for the structure of polysemy, there should be little variability with respect to both senses and patterns across languages.

According to a concepts-only model of polysemy, there should also be little variability with respect to generativity across languages—i.e., whether new senses for patterns (e.g., “tasty *seagull*”) can be coined. Specifically, a new sense for a word should be able to be coined whenever it stands in a noteworthy relationship to an existing sense of the word. This predicts that the same patterns should be generative across languages: Because speakers of different languages are likely to have similar concepts and world knowledge, they should find the same relations between new and old senses to be noteworthy, and thus derive novel senses in the same ways. However, the predictions of this model with respect to *which* patterns will be generative across languages are unclear. On one hand, it is possible that all patterns that are present in a language will also be generative, because if a pattern is present in a language, this would mean that speakers find the relation among attested senses of the pattern to be noteworthy (e.g., between a chicken animal and chicken meat), such that they should also find the same relation between an old and new sense to be noteworthy (e.g., between a seagull and seagull meat). On the other hand, whether a relation between a new and old sense is perceived as sufficiently noteworthy may also depend on speakers’ knowledge of the novel concept: e.g., if speakers know little about seagull meat, its relation to seagulls may not be noteworthy enough to warrant coining a new sense.

In contrast to the concepts-only model, a model that invokes only conventions to explain polysemy predicts that language- and culture-specific variation in senses and patterns should be the norm, not the exception. First, because conceptual structure plays only a limited role in how senses are coined and learned, there should be significant cross-linguistic variation in senses: languages should develop senses in unique ways, to address their own communicative needs. For instance, in some languages, it may have been communicatively beneficial to create a convention by which

⁴ We will return to alternative versions of this model in the discussion, but their predictions do not strongly differ from the account focused on here.

glass can label a drinking vessel, but this might not have been true of other languages. Further variation could be caused by the earlier-discussed phenomenon of pre-emption, in which the emergence of one sense prevents additional senses from arising (e.g., a drinking vessel sense could preempt the use of *glass* to label windows). Note that the phenomenon of pre-emption would be difficult to explain for a concepts-only model, because on this model senses are not stored and can be derived whenever they stand in noteworthy relations to core meanings. Given that there are therefore no stored, conventionalized senses, the acceptability of one sense cannot pre-empt another similar sense from emerging.⁵

If conventions alone account for the structure of polysemy, then languages should either fail to exhibit patterns, or should vary greatly in the patterns of polysemy they include. For example, if Murphy (2007) is correct, and patterns arise as speakers draw analogies from existing sets of senses to coin new ones, then the first sets of senses coined – which could vary dramatically by language – should determine which additional sets of senses are coined, and thus which patterns emerge. If a pattern emerges, speakers could also form a generative rule for the pattern, as long as there are a sufficient number of similar sets of senses following that pattern to form a basis for coining the rule. Thus in contrast to the concepts-only model, a model that appeals only to conventions predicts a great degree of variation with respect to both senses and patterns across languages.

Finally, within a model in which senses are conventions that are shaped by conceptual structure, the same patterns of polysemy should be present across languages, although the sets of senses that instantiate these patterns may vary. For example, according to one version of this model, the sets of senses that emerge in a language will be those that are more learnable, given the cognitive biases of learners (we will discuss the nature of these cognitive biases in section 4.3). Because these cognitive biases are likely to be shared by members of all linguistic communities, different languages should develop similar patterns, containing easily learnable sets of senses that are adapted to these cognitive biases. For example, learners may find the relation between materials and artifacts composed of those materials easy to conceptualize, allowing them to learn artifact senses for multiple material words that form a pattern. Thus, by this account, patterns should only be absent from a language when they have not had time to evolve, or when other forms in the language pre-empt them. For example, in some languages, morphemes or compounds may express equivalent meanings (e.g., the German morpheme *fleisch*, which can be added to animal names to denote the meat derived from animals), and the use of these devices may obviate the need to coin and learn additional senses.

However, although cognitive biases may constrain learners to acquire sets of senses of certain types (e.g., such that material words can also label artifacts, animal words can also label meat, etc.), in many cases these biases may leave open the specific sets of senses that will emerge in a language. For example, because there are many artifacts that are made of glass (e.g., drinking vessels, windows, mirrors, etc.), and because the relations between each of these artifacts and glass material are similarly easy to conceptualize, learners could in principle acquire senses that apply to any of these artifacts. Consequently, languages could vary in which of these artifact senses first enter the language as conventions, and these conventions could then pre-empt other possible artifact senses from emerging. The result would be that, while all languages would include sets of senses that follow a material for artifact pattern, the specific sets of senses instantiating those patterns might vary, e.g., such that one language might use a word for glass material to label a drinking vessel, while another language might use the word to label a window.

However, not all patterns will be as loosely constrained as the material-artifact pattern. For instance, patterns like animal for meat appear to be more tightly constrained: Knowing that a *chicken* is a type of animal, we can be confident in predicting exactly what *chicken* refers to when it refers to meat (i.e., chicken meat). Critically, the degree to which a pattern's sets of senses are constrained could have important implications for whether the same sets of senses appear across languages or not. In particular, sets of senses that follow tightly constraining patterns like animal for meat are likely to be present across languages (e.g., the word for a chicken animal should typically also label chicken meat in other languages), while sets of senses that follow loosely constraining patterns should be more variable across languages (e.g., the word for glass material could variously label a drinking vessel, window, etc., across different languages).

Importantly, the degree to which a pattern is tightly constraining may determine not only cross-linguistic variability in the pattern's sets of senses, but also whether the pattern is generative. In particular, patterns that constrain their senses more tightly, like animal for meat, should result in sets of senses that can be more easily aligned when they are compared. This will make it easier for learners to abstract a common underlying relation, allowing them to discover a higher-order generalization. For example, by realizing that the sets of senses corresponding to *chicken*, *lamb*, *fish*, are each related in the same way, learners could make the higher-order generalization that all names for animals can describe meats. This would make the pattern generative, and allow speakers to coin new senses for words, such as "tasty *seagull*". However,

⁵ Proponents of a concepts-only model could argue that when a sense has been derived from a core meaning on many occasions, this process of derivation may become routinized, such that the sense is stored and can then pre-empt other senses from being derived. Because this version of a concepts-only model is very similar to the conventions-constrained-by-concepts model (i.e., it includes stored, conventionalized routines), we do not consider it further here.

Table 2

The predictions of different models of polysemy about cross-linguistic variation in patterns, senses, and generativity.

Model of polysemy	Cross-linguistic variability in patterns	Cross-linguistic variability in senses	Cross-linguistic generativity of patterns
<i>Concepts only</i>	Low variability; The same patterns should be present in all languages	Low variability; Similar senses should be present in all languages assuming that communicative contexts are similar	Low variability; Patterns may either be generative or non-generative across languages
<i>Conventions only</i>	High variability; Patterns develop when senses are coined by analogy, therefore different languages have different patterns	High variability; Senses should vary considerably across languages	High variability; If patterns exist, they are language-specific and so should generate novel senses in different ways
<i>Conventions constrained by concepts</i>	Low variability; The same patterns should be present in all languages	Moderate variability; More loosely constraining patterns should have more variable senses across languages	Low variability; More tightly-constraining patterns should be more generative across languages

it should be harder to abstract a higher-order generalization when a pattern is loosely constrained, like the material for artifact pattern, as the relations between sets of senses may be more difficult to align (e.g., a glass, a tin, and an iron are all made of different materials and are used for different purposes). If this line of reasoning is correct, then patterns that tightly constrain their senses across languages should also be generative, allowing speakers to easily coin novel, analogous senses.

In sum, the concepts-only, conventions-only, and conventions-constrained-by-concepts models of polysemy make distinct predictions about cross-linguistic variability in polysemy, and these predictions are summarized Table 2. In the next section, we review the existing literature on how polysemy varies across languages, and show how it fails to distinguish between the different models of polysemy we have discussed.

1.3. Existing cross-linguistic data on polysemy

To evaluate the cross-linguistic predictions of the different models of polysemy described in the previous section, we would need to assess a large number of patterns of polysemy across a large set of languages, and test for the presence of many sets of senses per pattern. However, to our knowledge, such a survey has not yet been conducted. As we review below, studies that have explored a large set of languages have often focused on just one pattern of polysemy, while studies that have assessed a large set of patterns have often focused on a limited set of languages, or have probed only a small set of senses for each pattern.

For example, Boyeldieu (2008) provides a cross-linguistic analysis of the use of the word *animal* to mean *meat*. Greenberg (1983) was the first to note that many of the Niger-Congo languages of west and southern Africa collapse the meanings *animal* and *meat* into a single word. The same conflation is also found in two Tibeto-Burman languages (Matisoff, 1978, reported in Boyeldieu, 2008) and in Warlpiri, where the same word is used for meat and edible animals. Boyeldieu argues for a conventionalized, culture-driven explanation for why this example of polysemy emerged. But because this claim is based on only the polysemy of a single word and a limited number of languages, it provides a weak foundation for drawing general conclusions about the nature of polysemy across languages.

Studies conducted on a broader scale have found evidence for both cross-linguistic regularities in patterns of polysemy, as well as variation in the particular sets of senses that exemplify those patterns, consistent with the predictions of the conventions-constrained-by-concepts model. Perhaps the best known of these studies is Viberg's (1984) survey of perception verbs, which was conducted across 52 languages. Based on his data, Viberg proposed a hierarchy for characterizing how, across languages, verbs for one sensory modality can be extended to describe another. For example, verbs that originally described seeing were often extended to describe hearing, but not vice versa, and in turn, verbs that meant hearing were sometimes further extended to describe smelling, but again, not the reverse. This hierarchy can be thought of as a skeleton that permits languages to extend verb senses in certain ways, but not others. However, extensions are not required, allowing for cross-linguistic variation. These data are most consistent with the conventions-constrained-by-concepts account, but it is unclear if these results generalize to other patterns of polysemy.

Away from verbs of perception, Peters (2003) writes about an unpublished study by Seto that explored cross-linguistic variation in the container-contents pattern of polysemy, which is evidenced by the use of the English word *kettle* to label a container (*cast-iron kettle*), as well as the contents of that container (*boiling kettle*). Peters reports that Seto found the pattern to be present across a wide set of languages, including Korean, Mongolian, Javanese, Italian and English. Because we do not know whether this pattern was always instantiated in sets of senses in the same way across

languages, these findings are consistent with both the concepts-only model, which predicts that the same patterns and senses should be present across languages, and the conventions-constrained-by-concepts account, which predicts that the same patterns, but not necessarily the same senses, should appear across languages.

Although the studies of Viberg (1984) and Seto explored a large set of languages and assessed multiple sets of senses, they each focused on only a single pattern of polysemy, and so their findings may not generalize to other patterns of polysemy. To resolve this issue, some researchers have conducted broader surveys, assessing multiple patterns of polysemy across languages. For instance, Kamei and Wakao (1992) found differences in how English, Mandarin Chinese and Japanese speakers rated the acceptability of 25 different sentences that exhibited several different patterns of polysemy, including container for contents and producer for product. While they do not give a detailed report of their data, their summary of their findings suggests interesting evidence of cross-linguistic variability. For example, Mandarin speakers judged that the use of a word for a producer to describe their product was unacceptable, exemplified by the sentence *He read Mao* (which is acceptable in English).

The findings of Kamei and Wakao (1992) fit with a conventions-only model, in which different languages develop different patterns of polysemy, but they are difficult to explain for the other models. In particular, the concepts-only model predicts that the use of *Mao* to label his writings should be possible in Mandarin, because the same senses are predicted to be present across all languages. Similarly, the conventions-constrained-by-concepts model predicts that this use of *Mao* should be possible in Mandarin, because the producer-product pattern tightly constrains its senses. However, it is not clear how much weight should be placed on Kamei and Wakao's findings, because they only tested a small set of senses in total, which makes it risky to generalize beyond that set. For instance, subjects' judgments could have been affected by cultural norms involving the specific tested senses, e.g., Chinese individuals may not want to refer to Mao in careless ways. Judgments about specific senses could also have been exceptions to larger patterns. For example, a survey that asked English speakers to rate "The man ate pig" or "The man ate cow" might erroneously conclude that English does not use the same labels for animals and the meat derived from them. But in fact, this pattern is common English. Exceptions like *cow* and *pig* can be explained through the presence of synonymous terms, like *beef* and *pork*, which pre-empt the regular pattern, as discussed above.

The above considerations suggest that it is critical not only to assess a large number of patterns of polysemy across languages, but also to probe a large set of senses for each pattern. Peters (2003) attempted to do this by comparing sets of senses in English, Dutch and Spanish, using a large dataset: the cross-linguistic thesaurus known as EuroWordNet (Vossen, 1998). Peters reported that a number of patterns of polysemy (e.g., plant for food) are shared across languages, though not every set of senses found in English is also attested in Dutch or Spanish. These findings are consistent with the predictions of the conventions-constrained-by-concepts model, which predicts that, while patterns of polysemy should be present across languages, the sets of senses that instantiate those patterns may be more variable. However, Peters' findings are limited, in that his data were confined to three languages. Additionally, it is possible that Peters' findings resulted from the structure of EuroWordNet. Specifically, the English thesaurus had by far the broadest coverage, raising the possibility that many senses do exist in other languages, but were not listed in the less extensive thesauri of those other languages.

Finally, a recent paper by Zhu and Malt (2014) argues that there are cognitive constraints on the senses that different languages develop. The authors explored whether the senses corresponding to 36 different words in English (e.g., *head*) are also present in translation-equivalents of the core senses of these words in Mandarin Chinese (e.g., in the Mandarin word for the body part sense of *head*). Interestingly, roughly half of the English senses were also attested in Mandarin, and those senses that were more semantically-related to core meanings (e.g., the leader or decision-maker sense of *head*) were also more likely to be shared across the two languages. These results appear to provide support for the conventions-constrained-by-concepts model, because they suggest both that polysemy has a cognitive basis (because some senses are shared across languages) and that different languages develop their own conventionalized senses (because some senses are not shared across languages). However, from these data, it is difficult to draw robust inferences about cross-linguistic variability in patterns and senses because these authors did not test multiple sets of senses for each pattern, and focused on only two languages.

To review, although previous cross-linguistic studies have provided intriguing data, they do not yield strong conclusions about the nature of polysemy across languages. As such, it is difficult to evaluate the possible roles of concepts and conventions in constraining the structure of polysemy. Thus, to test the predictions of these different models, we conducted a new assessment of how polysemy varies across languages, which differed from previous surveys in several critical ways.

First, in contrast to previous studies, which focused on testing a limited number of patterns of polysemy across a large set of languages or vice versa, our study assessed both a large set of languages (15 including English) and a large set of polysemy patterns (27 patterns found in English). Second, to provide a rigorous test of whether each language included a particular pattern of polysemy, we tested a large number of sets of senses for each pattern (from three to seventeen), including not only examples of attested senses in English (e.g., *chicken* and *lamb* for the animal for meat pattern), but also examples of exceptions in English which might not constitute exceptions in other languages (e.g., *cow* and *pig*). We also

asked participants to report other senses that fit the patterns that came to mind, and supplemented all of these data by consulting dictionaries. Third, we probed the generativity of each pattern, by asking participants to judge novel senses. This allowed us to explore a critical prediction of the conventions-constrained-by-concepts model, that there should be a link between cross-linguistic variation in a pattern's senses and the generativity of that pattern, because both factors are determined by how tightly the pattern constrains its senses.

2. Methods

2.1. Participants

Because our goal was to probe whether senses and patterns in English are also attested in other languages, we selected participants who were native speakers of languages other than English, but who could also read, write, and speak English to a reasonable degree of fluency. Further, because we wanted respondents to comment on whether, in lieu of polysemy, their language used morphological devices like derivation and compounding, we sought out respondents who had some background training in linguistics and would thus be familiar with these concepts.

We were able to recruit 36 participants to complete the 4-hour-long survey. Four were native English speakers, who validated our judgments about polysemy in English. The remaining 32 were speakers of 17 different languages (for details about these participants, see Table 3, other participants were excluded for failing to complete the survey). For most languages, we collected responses from multiple participants. In general, we recruited additional participants for a language if our existing participants lacked a background in linguistics. Thus, although we only had 1 participant each for French, Hungarian, Italian, and Turkish, we were confident in their judgments because they had received training in linguistics. However, for each of three other languages (Arabic, Sindhi and Marathi) we were only able to get responses from a single speaker who lacked training in linguistics; since we were not confident in the robustness of these data, we do not report them here. This left us with 29 speakers of languages other than English, providing data on 14 languages. All participants were drawn from the communities around Harvard University and the University of California, San Diego, and included students as well as full-time researchers.

2.2. Materials

We selected 27 patterns of polysemy in English that have previously been identified in the literature and for which we could generate multiple examples (see, e.g., Baker, 1968; Clark and Clark, 1979; Copestake and Briscoe, 1995; Jackendoff, 1997; Klein and Murphy, 2002; Krifka, 2001; Lakoff and Johnson, 1980; Ostler and Atkins, 1992; Pustejovsky, 1995). Drawing on this literature and our own intuitions, we identified and tested multiple sets of English senses that followed each pattern, and in some cases, also included examples of exceptions to these patterns. For example, for the animal-meat pattern, we probed not only sets of senses like *chicken* and *lamb* that are attested in English, but also

Table 3
Background information about our participants and references to consulted dictionaries.

Native language	Number of participants	Participants with linguistics background	Average age began learning English	Average number of years speaking fluent English	Dictionary analysis
English	4	1	0	19	NA
Cantonese	3	0	4	10	dict.youdao.com
Farsi	3	1	2	19	NA
French	1	1	14	4	Harrap's (2001)
Hindi	3	1	5	15	http://www.hinkhoj.com/
Hungarian	1	1	11	7	Orszargh et al. (1998)
Indonesian	2	0	8	8	Stevens and Schmidgall-Tellings (2004)
Italian	1	1	5	16	Reynolds (1981)
Japanese	2	1	6	14	http://www.csse.monash.edu.au/~jwb/cgi-bin/wwwjdic.cgi?1C
Korean	2	0	5	16	Martin et al. (1967)
Mandarin	3	2	7	8	dict.youdao.com
Russian	2	2	5	16	http://en.bab.la/dictionary/english-russian/
Spanish	3	0	6	13	Galimberti Jarman et al. (2008), Velázquez de la Cadena et al. (2003)
Turkish	1	1	11	15	Iz et al. (1992)
Vietnamese	2	1	7	7	United States Joint Publications Research Service (1966)

exceptions like *cow* and *pig*. For each pattern, we drew on previous accounts, as well as our own intuitions, to decide which was the base sense (e.g., the animal meaning of *chicken*) and which was the extended sense (e.g., the meat meaning of *chicken*).

Table 4 lists the patterns and sets of senses assessed in our survey, along with the shorthand notation for each pattern, which we use in our data figures.⁶ We used a broad definition of polysemy when selecting these patterns, which resulted in the inclusion of patterns that are not always classified in the literature as examples of polysemy. For example, noun-verb alternations are often discussed as examples of morphological conversion, rather than as polysemy. However, there are reasons to think that polysemy and morphological conversion can be analyzed similarly (see Pylkkänen et al., 2006), and thus should not be separated. For example, some patterns of polysemy are paralleled by morphological alternations in other languages and even within the same language (e.g., “I drank a *bottle* (of wine)” versus “I drank a *bottleful* (of wine)”); see Copestake and Briscoe, 1995 for discussion). So as to not presume what should and should not count as polysemy prior to empirical testing, we decided to include these more controversial cases in our survey.

2.3. Procedure

The survey was hosted online using the LimeSurvey package. Participants completed it at their own pace, and in a place of their choosing. Upon completion, they received a gift voucher.

The survey was broken into 27 sections, based on the 27 English patterns that we evaluated. Each section began with a description of the pattern under investigation, instructing participants about the kinds of words they would be making judgments about. For example, before beginning the animal for meat section, participants read: “In this section, we will ask you a series of questions about words that, in English, can be both the name for an animal, and the name for the meat from that animal. In addition, we will ask you about words that can be the name of an animal, but not the name for the meat of that animal.”

After reading the instructions for a section (which described the pattern), participants answered a series of questions designed to investigate (1) whether translation-equivalents of attested English senses existed in their language (we also tested some English exceptions, e.g., *cow* does not label beef), and (2) whether they were willing to coin novel senses that follow the pattern. For each pattern, we first asked detailed questions about three or four particular sets of word senses that follow the pattern in English (e.g., the animal/meat senses of *chicken*, *lamb*), and elicited participants’ judgments about whether these sets of senses also existed in their language. For a number of additional senses, including English exceptions, we asked participants less-detailed questions. We used these judgments to test predictions about cross-linguistic variability in both the presence of patterns, and the presence of specific sets of senses. Finally, for each pattern, participants completed a question that probed their willingness to coin new senses following the pattern. We describe each type of question more fully below:

- 1) Detailed judgments about attested English words. These items tested whether translation-equivalents to sets of attested English senses existed in the tested language, using several detailed questions. An example from the word *chicken* is provided in Fig. 1.

As can be seen, for each tested word, participants were first given an explanation of the base sense of the English word (e.g., the animal sense of *chicken*), as well as an example of a sentence in which this sense of the word was used (e.g., “The *chicken* drank some water.”). Participants were then asked to provide a translation of the critical word, as it was used in the sentence.

Following this, participants were given an explanation of the extended English sense (e.g., the meat sense of *chicken*), as well as a sentence in which the English word was used in the extended sense (e.g., “The *chicken* is tasty”).⁷ They were then asked whether the translation-equivalent of the base sense they had identified could be similarly extended, and provided a *naturalness rating* for this extended sense on a 1 (Not Natural) to 5 (Perfectly Natural) scale.

Next, if participants deemed this extended sense “natural”, they were asked to translate an English sentence using the extended sense (“The *chicken* is tasty”) into one box, and if they deemed it unnatural, they were asked to provide an alternative translation of the sentence, in a second box (see Fig. 1). We refer to participants’ decision of which box to use as a *binary acceptability judgment*. When reading the initial instructions to the survey, participants were also told to note down, in this box, if the translation of the second sense involved adding a morpheme to the original word, or the creation of a compound.

⁶ Place for event polysemy was added after the survey was completed by the Italian and Spanish speakers, and so it is excluded from some of our analyses, as indicated below.

⁷ Note that, because of our focus on relatively context-independent examples of polysemy (as opposed to contextual uses like “The *ham sandwich* is ready for his check”), the extended senses were presented to participants with minimally supporting linguistic contexts.

Table 4

The patterns, attested English senses, and exceptions assessed in the survey.

Pattern used (base sense listed first)	Short name	Attested English senses (starred words were used on the list task)	Example sentences (used in the survey)	Exceptions to pattern
Animal for fur derived from animal	AnFur	Mink, chinchilla, rabbit, beaver, raccoon*, alpaca*, crocodile*	The <i>mink</i> drank some water/ She likes to wear <i>mink</i>	Sheep, cow, goose, elephant, oyster
Animal (or object) for personality property	AniPro	Chicken, sheep, pig, snake, star*, rat*, doll*	The <i>chicken</i> drank some water/ He is a <i>chicken</i>	
Animal for meat derived from animal	AnMe	Chicken, lamb, fish, shrimp, salmon*, rabbit*, lobster*	The <i>chicken</i> drank some water/ The <i>chicken</i> is tasty	Cow, pig, deer, calf, sheep
Artifact for activity involving artifact	ArtAct	Shower, bath, sauna, baseball,	The <i>shower</i> was leaking/ The <i>shower</i> was relaxing	Classroom, racket, toilet
Body part for object part	BdyObj	Arm, leg, hand, face, back*, head*, foot*, shoulder*, lip*, heart*, eye*, tongue*, wing*	John's <i>arm</i> was tired/ The <i>arm</i> was discolored [referring to arm of a chair]	Hair, calf, wrist, bones
Building for people in the building	BldPers	Church, factory, school, airplane,	The <i>church</i> was built 20 years ago/ The <i>church</i> sang a song	
Predicate for predicate with verbal complement (complement coercion)	CmpCoer	Begin, start, finish, try	John <i>began</i> reading the book/ John <i>began</i> the book [see Pustejovsky, 1995]	
Container for contents	ConCon	Bottle, can, pot, pan, bowl*, plate*, box*, bucket*	The <i>bottle</i> is made of steel/ He drank half of the <i>bottle</i>	
Word for question involving word (concealed question)	ConcQ	Price, weight, speed	The <i>price</i> of the coffee was low/ John asked the <i>price</i> of the coffee [what the price of the coffee was, see Baker, 1968]	Picture, car, shoe
Figure for Ground	FigGrd	Window, door, gate, goal	The <i>window</i> is broken/ The cat walked through the <i>window</i> [contrasting the window pane with the window frame]	
Object for substance constituting that object	Grinding	Apple, chair, fly	The <i>apple</i> was tasty/ There is <i>apple</i> all over the table	
Instrument for action performed by instrument	InsAct	Hammer, brush, shovel, tape, lock*, bicycle*, comb*, saw*	The <i>hammer</i> is heavy/ She <i>hammered</i> the nail into the wall	Car, broom, over, razor, scissors, spade, jug
Instance of an entity for kind of entity it is	InsKnd	Tennis, soccer, cat, dog, class*, dinner*, chair*, table*	<i>Tennis</i> was invented in England/ <i>Tennis</i> was fun today [contrasting the type of thing something is, to a token of that thing, see Carlson, 1977]	
Location for placing object in location	Location	Bench, land, floor, ground, box*, bottle*, jail*	The <i>bench</i> was made of pine/ The coach <i>benched</i> the player [sent the player to sit on the bench]	Garage, oven, hive
Object/substance for placing object/substance at goal	LocGoal	Water, paint, salt, butter, frame*, dress*, oil*	The <i>water</i> is cold/ He <i>watered</i> the plant.	Blanket, shirt, ring, letter
Object/substance for taking object/substance from source	LocSrc	Milk, dust, weed, peel, pit*, skin*, juice*	The <i>milk</i> tastes good/ He <i>milked</i> the cow	Lint, fleas
Material for artifact	MatArt	Tin, iron, china, glass, linen*, rubber*, nickel*, fur*	Watch out for the broken <i>glass</i> / He filled the <i>glass</i> with water	Copper, aluminum, silver, clay, cement, wool, yarn, cotton
Object for color	ObjCol	Orange, violet, peach, rose, gold*, amber*, lavender*, turquoise*	She ate an <i>orange</i> / She has an <i>orange</i> t-shirt [see Casson, 1994]	Yellow, green, pink, black, brown, purple, white, blue, scarlet
Occupation for a role played in action	OccRol	Boss, nurse, guard, tutor,	My <i>boss</i> is nice/ He <i>bossed</i> me around	Chef, lawyer, priest
Place for an event	PIEv	Vietnam, Korea, Waterloo, Iraq	It is raining in <i>Vietnam</i> / John was shot during <i>Vietnam</i>	
Place for an institution	PIIns	White House, Washington, Hollywood, Pentagon, Wall Street*, Silicon Valley*, Supreme Court*	The <i>White House</i> is being repainted/ The <i>White House</i> made an announcement	

Table 4 (Continued)

Pattern used (base sense listed first)	Short name	Attested English senses (starred words were used on the list task)	Example sentences (used in the survey)	Exceptions to pattern
Plant for food or material	PlntFd	Corn, broccoli, coffee, cotton, lettuce*, eggs*, oak*, pine*	The large field of <i>corn</i> / The <i>corn</i> is delicious	Grape, orange, apple, olive, chickpea
Substance for portioning of that substance	Portion	Water, beer, jam	She drank some <i>water</i> / She bought three <i>waters</i>	
Publishing institution for product created by institution	Prd	Newspaper, magazine, encyclopedia, Wall Street Journal*, New York Times*, People*	The <i>newspaper</i> is badly printed/ The <i>newspaper</i> fired three employees	Book, car, toy, hat
Artist for product created by artist	PrPr	Writer, artist, composer, Shakespeare, Dickens*, Mozart*, Picasso*	The <i>writer</i> drank a lot of wine/ The <i>writer</i> is hard to understand [the writer's works are hard to understand]	
Container for representational contents	RepCont	Book, CD, DVD, TV*, magazine*, newspaper*	The heavy, leather-bound <i>book</i> / The <i>book</i> is funny.	Hardcover, cassette, sheet
Object for something that is visually or functionally-related	VisFun	Beam, belt, column, stick, bug*, leaf*	Most of the weight in the structure rests on the <i>beam</i> / There was a <i>beam</i> of light	

Animal-Meat-Chicken

***In English, a "chicken" is a type of bird, and one that typically lives on a farm. For instance, you can say "The chicken drank some water." How would you translate the word "chicken" as it is used in this sentence?**

***In English, it is also possible to use the same word "chicken" that refers to the animal, to also refer to the meat that is derived from the flesh of the animal. For instance, you can say "The chicken is tasty" when you mean that the meat derived from the chicken was tasty to eat. In your language, how natural is it to use the word for the animal that you provided above to also refer to the meat, as it is used in English in "The chicken is tasty"? (Rate this between 1 (Not Natural) and 5 (Perfectly Natural))**

☐ 1
 ☐ 2
 ☐ 3
 ☐ 4
 ☐ 5

If it is natural, please translate the above phrase into your language:

If this use is non-natural, please write down how you would instead translate this second sense.

If there is another way of using "chicken" that immediately comes to mind, please describe it below.

Fig. 1. An example of the judgments elicited for the word *chicken*, in the animal for meat pattern.

Finally, participants were asked to list any other additional senses of the translation-equivalent they had provided for the base sense of the target word (e.g., “If there is another way of using “chicken” that immediately comes to mind, please describe it below”; see Fig. 1). This was to probe for the existence of other possible senses within the same pattern (e.g., the use of *glass* to label a mirror, as opposed to a drinking vessel), as well as other, possibly language-specific patterns of polysemy.

- 2) *Brief judgments about additional attested senses.* For some patterns, we also included an additional list of attested English words whose senses follow the pattern (see Table 4 for an indication of which patterns included these questions). Participants were asked if translation-equivalents of these words had analogous senses in their language. In the interests of time, participants were not asked to give naturalness ratings. Instead, we simply asked participants to provide a translation for the base sense of each example, and then asked them to judge whether their provided translation also had an extended sense in line with the pattern in question or not.
- 3) *Brief judgments about exceptions.* These items tested whether words that are exceptions to patterns in English (e.g., *cow* cannot be used to label beef) are also exceptions in other languages. For some patterns, we could not identify any exceptions to patterns, but for others, we provided participants with a list of exceptions, and asked them questions about these exceptions (see Table 4). Participants provided translations for each exception and indicated whether these translations could be extended (e.g., whether their translation-equivalent of *cow* could label beef), just as they did for the judgments about additional examples. If a translation-equivalent did not have an extended sense, participants were asked to provide the distinct word in their language that corresponded to this extended sense.
- 4) *Free recall of attested senses and exceptions.* With these items, we asked participants to provide a list of any additional words in their languages that either did or did not have senses that followed the pattern in question.
- 5) *Generalization judgments.* Finally, we tested whether participants were willing to coin new senses following the pattern in question. Participants read about a newly coined word that corresponded to the base sense of the pattern, and then rated whether this word could be felicitously extended (on a 1–5 scale). For example, for the animal for meat pattern, participants read: “Imagine that a new animal was discovered called a “dax”. Imagine that a person was eating the meat derived from this animal and found it to be tasty. How acceptable would it be to say, in your language, that “The dax is tasty?” We included generalization questions for all patterns, except for the complement coercion pattern, for which we could not easily construct a candidate novel example.

2.4. Dictionary analysis

We complemented participants’ responses by exploring whether additional sets of senses, for each pattern, were documented in bilingual dictionaries (see Table 3 for references). To do so, for each language we took one participants’ translations of base senses (e.g., of the animal sense of *chicken*), and noted whether additional senses for those words were listed in dictionaries, and whether those senses fit the target patterns.⁸

3. Results

Here, we confine our analyses to our questions of interest (our full dataset and analysis scripts are available upon request). Specifically, we present our findings as they bear on (1) whether the same patterns are present across languages, (2) whether the sets of senses that instantiate these patterns are the same across languages, and (3) whether the same patterns are generative across languages and how this relates to cross-linguistic variability in senses. All analyses were conducted using the R software package (R Development Core Team, 2014)

3.1. Cross-linguistic variability in patterns

As described before (see Table 2), according to the concepts-only and conventions-constrained-by-concepts models, there should be little cross-linguistic variability in patterns of polysemy: the same patterns should be present across languages, such that English patterns should be present in other languages. However, according to the conventions-only model, there should be high cross-linguistic variability in patterns of polysemy: to the extent that patterns emerge at all, they should vary across languages, such that English patterns should not be present in other languages.

To distinguish between these predictions, we evaluated whether each of the target patterns found in English were also present in each of the tested languages. We considered a pattern to be present in a language if we found evidence that at

⁸ We were unable to perform a dictionary analysis for Farsi: Our Farsi respondents all transliterated into the Roman alphabet, and attempts to transliterate them back into the Persian alphabet (using <http://www.behnevis.com/en/>) did not produce words that could be found in Farsi dictionaries.

Fig. 2. Evidence for the presence of patterns across languages. A white box indicates evidence for a pattern in a language, and a gray box indicates no evidence for the pattern. The figure is ordered such that languages with evidence for more patterns are toward the top, and patterns that are attested across more languages are toward the right. The place for event pattern is excluded from this analysis (see footnote 5). Table 4 contains a legend for the pattern names.

least one set of senses in the language followed the pattern. More specifically, we judged a pattern to be present if one of the following was true:

1. If participants judged an extended sense following the pattern as acceptable in one of the judgment tasks (i.e., in the binary acceptability judgment task, or in the judgment tasks about additional English attested senses and exceptions).⁹
2. If participants listed a word from their language whose senses followed the pattern in the free recall task.
3. If the dictionary analysis uncovered attested sets of senses in the language that fit the pattern.

Fig. 2 depicts the results for each of the patterns, across each of the languages. As can be seen, the data suggest that patterns of polysemy that are present in English are also generally present in other languages. In particular, only 23 pattern/language combinations were unattested out of a total of 390 possible combinations, a rate of 6%.¹⁰ This rate of absent pattern/language combinations is much lower than would be expected by a conventions-only model of polysemy, which would predict that if patterns exist at all, they should be variable across languages. Thus, our data provide strong evidence against the idea that the structure of polysemy boils down to a set of learned conventions: instead, conceptual structure may also play a role in constraining polysemy.

Although our criteria considered a pattern to be “present” in a language when one set of senses following that pattern could be identified, most patterns were evidenced by more than one set of senses. Indeed, as indicated in Table 5, native speakers often indicated that there were multiple polysemous words whose senses followed each pattern in their language. Together, these data are consistent both with the concepts-only model (which predicts that the same patterns and similar senses should be present across languages) and the conventions-constrained-by concepts model (which predicts that patterns should be present across languages, though the senses that instantiate those patterns may vary).

⁹ Naturalness ratings were not used in this analysis because they were redundant with participants’ binary acceptability judgments.

¹⁰ We also analyzed how often the 390 pattern/language combinations were only met by one of our three criteria for considering a pattern “present”. One combination (visual/functional metaphors in Mandarin) was considered present only via a dictionary analysis. 18 combinations were considered present due only to data from the free recall task. Finally, 75 combinations were considered present due only to data from our judgment tasks (24 of these combinations involved the complement coercion and concealed question patterns).

Table 5
Average number of extended senses judged or listed by native speakers as acceptable, for each pattern and language (data from the dictionary analysis and for place for event are not included).

	LocSrc	Location	OccRol	LocGoal	Portion	Prd	AnFur	Grinding	InsAct	VisFun	AniPro	AnMe	ArtAct	BdyObj	BldPers	CmpCoer	ConCon	ConcQ	FigGrd	InsKnd	MatArt	ObjCol	PlIns	PlntFd	PrPr	RepCont
Spanish	3.5	2.7	3.3	3.5	4.0	5.1	3.7	1.3	3.7	7.0	8.3	9.0	4.0	7.0	3.0	3.7	10.7	4.0	5.0	6.0	5.7	7.0	4.0	6.2	4.7	7.3
Farsi	5.0	3.5	6.7	7.7	5.0	5.7	4.7	3.0	7.0	3.3	7.8	9.3	7.7	7.3	3.3	3.3	10.0	3.0	7.0	9.7	4.0	6.3	8.7	10.7	6.7	9.3
Turkish	1.0	3.0	0.0	6.0	3.0	3.7	5.0	3.0	12.0	5.0	9.0	9.0	4.0	9.0	4.0	4.0	9.0	3.0	5.0	6.0	4.0	4.0	6.0	5.0	7.0	7.0
Russian	0.0	1.0	6.0	7.5	1.0	6.3	4.0	2.0	3.5	9.0	9.0	7.5	2.0	12.0	4.5	1.5	10.0	6.0	5.5	8.0	3.0	6.5	6.5	9.0	9.0	5.0
Japanese	0.0	0.5	1.5	2.5	1.5	3.7	6.0	1.5	1.5	2.0	0.5	6.5	3.0	4.5	1.0	2.0	3.5	4.0	3.5	7.0	3.5	5.0	2.5	6.5	7.0	7.5
Indonesian	0.5	1.0	2.0	3.0	0.5	0.0	2.0	1.0	5.0	2.0	3.5	5.0	2.0	2.5	2.5	1.0	5.0	4.0	4.5	3.5	7.5	2.0	5.0	5.5	4.5	10.0
Hindi	0.3	1.0	2.0	2.7	0.0	3.0	8.3	2.0	3.8	2.3	6.8	10.0	0.3	2.3	2.0	1.0	5.3	1.3	4.0	4.3	2.3	3.3	6.0	5.0	1.7	6.5
French	1.0	4.0	0.0	3.0	2.0	5.7	9.0	3.0	13.0	7.0	12.0	12.0	6.0	14.0	5.0	4.0	9.0	6.0	4.0	9.0	14.0	4.0	8.0	12.0	4.0	10.0
Cantonese	0.3	0.0	1.3	0.8	1.8	3.3	2.0	3.0	12.3	4.3	7.5	11.3	3.3	5.5	3.3	2.0	10.0	2.8	4.0	6.0	3.8	5.8	6.3	10.0	4.5	8.4
Korean	0.0	0.0	2.0	2.0	1.5	1.7	2.0	1.0	5.0	0.5	2.0	5.5	3.5	3.0	1.0	2.5	5.0	3.0	4.0	4.0	1.0	5.0	3.5	7.5	5.0	7.0
Italian	4.0	1.0	0.0	4.0	2.0	2.3	7.0	0.0	4.0	1.0	5.0	6.0	2.0	10.0	1.0	3.0	9.0	3.0	5.0	4.0	3.0	4.0	3.0	4.0	6.0	6.0
Hungarian	0.0	3.0	4.0	0.0	4.0	4.7	3.0	2.0	0.0	4.0	5.0	9.0	1.0	9.0	3.0	3.0	4.0	2.0	4.0	9.0	6.0	2.0	6.0	8.0	7.0	9.0
Vietnamese	0.0	0.0	2.0	1.5	0.0	3.7	0.0	2.0	2.0	3.5	5.5	7.5	4.0	8.0	3.0	2.0	10.0	6.0	4.5	8.5	2.5	7.5	4.0	8.0	2.5	9.4
Mandarin	0.0	0.0	1.0	0.0	2.0	0.0	1.0	2.5	5.0	0.0	4.0	5.0	4.5	1.5	2.0	1.5	1.0	2.5	3.0	6.0	0.5	2.0	6.0	5.5	3.0	3.9

But what can we make of the fact that some English patterns appeared to be present across more languages than others? For example, patterns like plant for food (“The *corn* is delicious”) and container for representational contents (“The *book* is funny”), were present across all languages (Fig. 2), and participants were able to report many examples of senses that follow those patterns (Table 5). In contrast, other patterns, like locatum verbs describing sources (“John *milked* the cow”) and location verbs describing where objects are placed (“She *boxed* the books”) were more absent across languages, with fewer reported sets of senses. One question raised by these data is whether it is valid to say that some patterns were more absent across languages than others, or instead whether this observed variability in patterns was an artifact of limitations in our methods. For example, it is possible that we did not find evidence for some pattern-language combinations not because patterns do not exist in some languages, but instead because our participants forgot the appropriate translations or because we could only assess a limited number of possible senses per pattern.¹¹

We reasoned that if the observed absent-pattern language combinations were due to random variation given our methods, and do not reflect that some patterns are more likely to be absent than others, then the absences we observed (i.e., the gray squares in Fig. 2) should be randomly-distributed across both languages and patterns, as opposed to concentrated around specific patterns. To explore if this was the case, we used a resampling approach to estimate how likely it would be, given the observed number of absent patterns in each language, that a single pattern would be absent across multiple languages. Thus, we took the data from each language and, within that language, randomly shuffled the pattern labels (e.g., “material for artifact”) so that each label was associated with a randomly chosen data point.¹² Consequently, there was a 1 out of 26 chance that the “material for artifact” label would be associated with the data for the “material for artifact” pattern, as opposed to the data corresponding to the other 25 patterns. We did this for each language, and then recorded the largest number of languages for which a particular pattern was absent (e.g., in our original dataset this number would be 6). We repeated this process 10,000 times, and then compared the resulting empirical distribution of “most absent” languages for patterns to our actual dataset.

The most absent pattern in our dataset was the set of locatum verbs describing sources (e.g., “*weed* the garden”), which was absent in 6 languages (see Fig. 2, Table 5). In our simulations, it was very rare for a pattern to be absent in 6 or more languages: this occurred only 1 out of 10,000 times. The probability of a pattern being absent in 6 or more languages of our sample is therefore approximately .0001, which meets standard criteria for statistical reliability, even when corrected for multiple comparisons ($\alpha = .002$), suggesting that the locatum source pattern is less attested than we might expect by chance. The next most absent pattern, was the set of verbs describing the location something goes to (e.g., “*box* the books”) – this pattern was absent in 4 languages. Our simulations suggest that this result was also quite improbable ($p = 0.024$), although this was not significant after correcting for multiple comparisons. But after this, the probability of a pattern being absent in fewer than 4 languages by chance, which was true of several other patterns in our dataset, was above standard thresholds for statistical significance.¹³

The analysis described above suggests that the absent pattern-language combinations in our data were not randomly distributed. For example, an unexpectedly high number of languages appeared to lack locatum source verbs. Although this could be explained by limitations in our methods, we think this is unlikely. Instead, we believe that there is a good reason for the high number of languages that did not include locatum source verbs, which we set out in the discussion section.

Our analyses do not speak to whether those patterns that were absent in fewer than four languages in our data are indeed absent in those languages. On one hand, as noted above, it remains possible that we did not find evidence for those patterns because we failed to ask participants about actual, attested senses. On the other hand, it is also possible that these patterns were indeed absent in the languages indicated by our data. For instance, some of these patterns may not have been present in some languages due to restrictions those languages place on syntactic flexibility. This would help explain why most of the more absent patterns involved senses that cross lexical categories, appearing as nouns or verbs (e.g., in the case of locatum and location verbs), or as count or mass nouns (e.g., in the case of grinding and portioning).¹⁴ Future research, probing a more exhaustive set of possible senses, will be necessary to determine whether these patterns are indeed absent in languages, as suggested by our data.

¹¹ The latter seems unlikely, because there was no correlation between the number of examples assessed per pattern and the proportion of languages in which that pattern was absent in ($r(24) = .02$, *ns*).

¹² We excluded the place for event pattern from this analysis, as we did not have data from Spanish and Italian for this pattern.

¹³ To ensure that this result was not driven by data from the dictionary analysis (which potentially may not reflect speakers’ intuitions), we excluded these data and repeated the analysis. The results were the same.

¹⁴ However, a grammatical restriction on noun-verb flexibility is unlikely to fully explain the observed variation. In our data set, for example, we found evidence that all probed languages exhibited at least one of the five forms of noun-verb conversion we probed (i.e., locatum source verbs, locatum goal verbs, location verbs, instrument verbs, or occupation verbs; see Table 2). Thus, it does not seem to be the case that in these languages, there are grammatical restrictions that wholly prevent noun-verb conversion, although languages may vary in the productivity of these conversion processes.

To sum up, although our data provide evidence of some cross-linguistic variation with respect to patterns of polysemy, most patterns that are present in English were also generally present across languages. Indeed, all patterns were present in multiple languages, and most were present in nearly all probed languages. While these data are difficult to explain for a conventions-only model, they are consistent with a concepts-only model and a conventions-constrained-by-concepts model. The next section provides evidence that distinguishes between these latter two models.

3.2. Cross-linguistic variability in senses

Although both the concepts-only model and the conventions-constrained-by-concepts model posit that the same patterns should recur across languages, they make different predictions about whether the sets of senses that instantiate those patterns should vary (see Table 2). The concepts-only model predicts low cross-linguistic variability in senses: Within each pattern, the same sets of senses should be derivable across languages when communicative contexts are kept similar (as was the case in our survey). In contrast, the conventions-constrained-by-concepts model predicts that cross-linguistic variability in senses should depend on the pattern, such that for more tightly constrained patterns (e.g., animal for meat) the same sets of senses should be present across languages, but for more loosely constrained patterns (e.g., material for artifact), languages could develop different sets of senses. In this section, we focus on whether senses are indeed variable across languages, and whether this depends on the pattern in question.

To examine whether the senses observed in English are also present in analogous words in other languages, we looked at subjects' judgments of whether translation-equivalents of the base senses of English words (e.g., of *chicken* [animal]) also had the same extended senses (e.g., referring to chicken meat). We analyzed items that all subjects received (and thus excluded free recall responses) using two measures: (1) naturalness ratings (i.e., 1–5 ratings of whether an extended sense was natural) and (2) binary acceptability judgments (i.e., whether extended senses were judged acceptable or not). For each language and each probed sense, we computed average naturalness ratings and binary acceptability judgments across respondents for that language. Fig. 3 depicts the mean naturalness ratings across languages, separated by pattern.

Fig. 3. The mean naturalness ratings of senses following each pattern, by language. Average rating is color-coded (white = most natural, black = least natural). White squares with crosses indicate missing data. The heatmap is ordered such that the languages with the highest average naturalness of senses (across patterns) are at the top, and patterns with the highest average naturalness of senses (across languages) are on the right.

To get a broad sense of the degree to which these extended English senses are also attested in other languages, we first explored whether the average naturalness and binary acceptability judgments for attested English senses by our native English-speaking participants were significantly higher than the average judgments of these possible senses in the other 14 languages (i.e., by speakers of those languages). As a look at Fig. 3 would suggest, our English raters typically gave higher naturalness ratings and also judged senses as acceptable more often than did our informants from other languages. For the naturalness ratings, all 14 of the pair-wise comparisons (e.g., English ratings versus Farsi ratings) reached significance (Bonferroni corrected p value of 0.0036, median $t = .5.6$, range: 3.2–8.3; median $p = .000002$), while 12 of 14 comparisons were significant for the binary acceptability judgments (Median $t = 4.7$, range: 1.9–7.6; median $p = .00008$). These analyses suggest that in other languages, translation-equivalents of some English senses are less acceptable, and perhaps absent.

This analysis suggests that there is significant cross-linguistic variation in the sets of senses that instantiate patterns. This is inconsistent with the concepts-only model, which predicts uniformity in senses across languages (Table 2). But does this variation in senses depend on the pattern in question? Recall that the conventions-constrained-by-concepts model predicts that some patterns – i.e., those that more tightly constrain their senses – should have more similar sets of senses across languages than other patterns have.

Consistent with the idea that cross-linguistic variation in senses may depend on pattern, an inspection of Fig. 3 suggests that for some patterns, naturalness ratings of senses across the non-English languages were quite similar to the ratings of English senses, but that this was not as true of other patterns. For example, the plant for food pattern had similarly high naturalness ratings for its senses across languages.¹⁵ The ratings for senses following the animal for meat pattern were also quite similarly high across languages, although there was also variation, perhaps because a number of languages can label meat by compounding animal names with the word for “meat” (e.g., Farsi, Chinese, Japanese, Korean, Indonesian).¹⁶ In contrast, however, a number of patterns, like material for artifact, were instantiated in very different senses across languages, as evidenced by relatively low naturalness ratings. For example, material words for iron, tin, glass and rubber often do not name the same artifacts in English as they do in other languages. However, as noted in the previous section, this pattern was still present across languages: it was just instantiated in different senses (e.g., one respondent noted that the Russian word for *rubber* names a car tire).

From looking at Fig. 3, it seems plausible that the patterns for which English-like senses are acceptable in other languages are also patterns that place tighter constraints on their senses, consistent with the conventions-constrained-by-concepts model. For example, knowing that a plant name labels food, or that an animal name labels meat, it is relatively easy to determine what specific food or meat the name refers to, which may explain why the sets of senses that instantiate these patterns in English are also natural in other languages. In contrast, knowing that a material word labels an artifact is not of much help in determining what artifact the word refers to, which could potentially explain why the particular artifact senses of these words found in English are not acceptable in other languages.

To formally test whether cross-linguistic variability in naturalness ratings of senses is dependent on pattern, we used a resampling analysis to ask whether the distribution observed in Fig. 3 could be a product of random variability, rather than dependent on pattern. In particular, we asked whether the differences between patterns in which English-like senses were natural across languages (e.g., animal for meat) and patterns in which English-like senses were less natural across languages (e.g., material for artifact) were larger than would be expected if we had computed our averages over random groups of senses, as opposed to grouping senses by their pattern. That is to say, are there real differences in the ways in which patterns are instantiated in senses across languages?

To conduct this analysis, we first measured the difference between patterns that have more English-like senses across languages, and patterns that have less English-like senses across languages. First, we averaged the mean naturalness ratings of senses for each pattern across languages excluding English, which served as a measure of how natural English-like senses are across languages, for each pattern. Then, we ranked these patterns by their average naturalness ratings, and regressed the rank order of each pattern against its average naturalness rating. The resulting regression slope served as our measure of difference between patterns in how natural their corresponding English-like senses are in other languages: the steeper the slope, the larger the difference between patterns.

Critically, if patterns play no role, then the steepness of this regression slope should be no different when senses are not linked to their patterns. We therefore compared our observed slope ($\beta = 0.12$) against a distribution of simulated

¹⁵ Still, this was not without exceptions: e.g., while most languages use the same word for the cotton plant and material, Hungarian differentiates the two (*gyapot* describes the plant; *pamut* the material).

¹⁶ However, respondents often felt that use of the animal name on its own was still acceptable: e.g., a Korean respondent wrote that “In Korean, one can say the chicken is tasty. However, it is more proper to say the chicken meat is tasty”).

regression slopes that did not invoke patterns. Specifically, within each language, we first shuffled pattern labels for the probed senses to create random groups of senses, e.g., such that “material for artifact” could come to be associated with an animal for meat sense like *chicken*. Then, we calculated regression slopes based on the mean naturalness ratings for these random groups of senses. We repeated this process 10,000 times, and then compared the regression slope derived from our data – which was based on grouping senses by patterns – to the empirical distribution of slopes – which used random groups of senses. Strikingly, the regression slope derived from our data was steeper than each of the 10,000 slopes derived from the simulations, indicating that cross-linguistic differences in English sense naturalness between patterns were indeed larger than would be expected by chance.¹⁷

Together, these findings suggest that although patterns of polysemy are generally present across languages (as shown in the previous section), there is large cross-linguistic variation in how these patterns are instantiated in senses, with some patterns having more variable and less English-like senses across languages than others. These findings are consistent with the conventions-constrained-by-concepts model, which proposes that while patterns should be present across languages, cross-linguistic variation in sense acceptability across patterns should also exist, due to variation in how tightly patterns constrain their senses. The next section explores whether patterns that have more similar, English-like senses across languages, are also more generative.

3.3. Generativity of patterns across languages

The conventions-constrained-by-concepts model predicts that, when a pattern tightly constrains its senses, it should have very similar senses across languages. The previous section provided some evidence for this idea, by showing that, across languages, some patterns have more English-like senses than others. But showing that these differences between patterns stem from the degree to which patterns constrain their senses is more difficult. For example, obtaining a metric of how tightly the material for artifact pattern constrains its senses would require us to count the number of possible referents of material words – e.g., the number of things made of glass, or sponge, or iron, etc., across cultures.

However, the role of patterns in constraining their senses can be tested indirectly, because it is predicted to affect not only how similar a pattern's senses will be across languages, but also that pattern's generativity. Specifically, when a pattern tightly constrains its senses, language users may also find it easier to make a generalization about the senses that instantiate the pattern (e.g., realizing that any word for an animal can label its meat), allowing them to coin novel senses that follow that pattern. Our final analysis tested whether this prediction of the conventions-constrained-by-concepts model holds true: Are patterns with more similar senses across languages also more likely to be generative?

We measured each pattern's generativity using the survey's generalization questions, which asked participants to rate the acceptability of novel senses following patterns on a 1–5 scale. Fig. 4 depicts the average acceptability ratings of these novel senses, for each language and pattern. As can be seen, although there was considerable cross-linguistic variation in ratings of the novel senses *across* different patterns, there was also consistency in the ratings of novel senses *within* patterns. To confirm this, we correlated ratings for novel English senses (by our English-speaking subjects) – where all patterns are known to exist – with ratings of these senses in each of the other 14 languages, and found reliable correlations in each case (mean $r(24) = .58$, range: .42–.79; mean $p = .008$, range: .000003–.04).¹⁸ To follow this up, we looked at the intercorrelations of the novel sense ratings amongst all languages. The result was strikingly similar, with a mean correlation coefficient of 0.57 (range 0.16–0.87). This suggests that there may be universal intuitions with respect to how easy it is to coin new senses for patterns.

Next, we conducted an analysis to directly address whether the generativity of a pattern is related to the degree to which it is instantiated in similar senses across languages. First, we created a *generativity score* for each pattern, by averaging together generalization ratings from all of the non-English languages for each pattern.¹⁹ Thus, higher generativity scores indicate that participants, across languages, judged the novel senses as more acceptable. Second, to create a measure of the degree to which patterns have similar senses across languages, we averaged the naturalness ratings of the different senses from each pattern, across all of the non-English languages. Higher average naturalness ratings thus indicate more similar – or more precisely, more English-like – senses across languages.

Fig. 5 plots the relationship between generativity scores and average naturalness ratings of English-like senses across the non-English languages. Strikingly, the relationship between these two variables was very strong, $r(24) = .91$, $p < .001$.

¹⁷ We found the same result when we performed this analysis on respondents' binary acceptability judgments.

¹⁸ Note that we excluded the Place for Event pattern from this analysis, as it was not assessed in Italian and Spanish, as well as the complement coercion pattern, for which we did not test generalization.

¹⁹ English ratings were excluded from this analysis because participants' naturalness ratings were at ceiling.

Fig. 4. The mean acceptability ratings of novel senses for each pattern, by language. Higher scores – which are progressively whiter – indicate that the novel sense was rated as being more acceptable, while lower scores – which are progressively blacker – indicate that this sense was rated less acceptable. White squares with crosses indicate missing data.

Thus, patterns that have more similar senses across languages are also more generative.²⁰ This relationship is difficult to explain for the conventions-only and concepts-only models. In particular, because these models predict that all patterns should be instantiated either in cross-linguistically similar senses (the concepts-only model) or variable senses (the conventions-only model), they cannot explain why cross-linguistic similarity in a pattern's senses should co-vary with the generativity of that pattern. However, this relationship is consistent with the conventions-constrained-by-concepts model, which predicts that the generativity of a pattern and cross-linguistic variability in its senses should be related, because they are each determined by a common factor: how tightly the pattern constrains its senses. Specifically, when a pattern constrains its senses more tightly, languages will develop similar conventions regarding which senses can follow the pattern. Further, when a pattern constrains its senses tightly, the relation between each set of senses that follows the pattern will be easier to abstract and generalize to novel senses.

4. Discussion

The present study explored the roles of concepts and conventions in the structure of the lexicon, by documenting cross-linguistic regularity and variation in polysemy. We reasoned that if polysemy is a direct reflection of conceptual structure, the same patterns and senses should be present across languages, but if polysemy corresponds to arbitrary lexicalized conventions, patterns and senses should be highly variable across languages. Our findings suggest that the structure of polysemy cannot be explained by either concepts or conventions on their own. Specifically, across 15 languages and 26 patterns of polysemy, we found very few instances where a language showed no evidence of having a

²⁰ Readers may be concerned that this analysis risks circularity. For example, average naturalness ratings of English-like senses might be lower when a pattern is rarely used across languages, and so generativity scores would also be correspondingly lower. To assuage this concern, we repeated the analysis by comparing the average naturalness ratings to *English* generativity scores (where all the patterns are known to exist). Again, we observed a robust correlation ($r(24) = .65, p < .01$), consistent with the proposal that the generativity of a pattern and cross-linguistic variability in its senses are linked.

Fig. 5. The relationship between a pattern's generativity score across languages and the naturalness ratings across languages of senses that follow the pattern in English.

particular pattern of polysemy (like the use of material words to label artifacts), which provides evidence against a conventions-only model. However, contrary to a concepts-only model, we found that many patterns are instantiated by different sets of senses across languages (e.g., *glass* labeling a drinking vessel, a mirror, etc.).²¹

Instead, our data are best explained by a model that we call “conventions-constrained-by-concepts”. In this model, the senses of polysemous words are learned conventions, explaining why we found that languages develop different sets of senses. However, conceptual biases make it easier for members of all linguistic communities to learn some sets of senses, compared to others, explaining why we found some commonalities across languages. In particular, because learners’ cognitive biases make it easier for them to learn sets of senses that are related in particular ways, these sets of senses will be similar to one another and form patterns, explaining why we found that the same patterns of polysemy (e.g., animal for meat, material for artifact) are largely present across languages. Finally, this model also predicts – as we found – that some patterns will have more variable senses across languages than others, and that this will be linked to how generative those patterns are. In sum, the model provides a reasonable account of our data as well as a satisfying integration of previous theoretical claims. Much of the rest of the discussion is devoted to describing this theory in more detail, relating it to previous work in linguistic and cognitive development, and laying out its predictions for future work. But first, we discuss possible limitations of our data.

²¹ Our editors note that pragmatic theories, which we classify together with conceptual theories, also include an important notion of background communicative context: The sense of a word can only be determined by attending to that background context as part of the process of determining a speaker’s intended meaning. They argue that these communicative contexts might explain some of the apparent arbitrariness of polysemy, and that this may also explain differences in polysemy that we observed across languages (see the paper by Falkum in this issue). We agree that variability in communicative context could cause differences in how word senses are used across languages. However, we doubt that the everyday communicative needs of, e.g., English or Russian speakers are so distinct. Moreover, the responses in our study were provided by informants who were bilingual, and who were completing the survey in English. If communicative context drove their responses, we would expect them to answer in a similar manner to English speakers. Still, the degree to which communicative contexts differ across language groups remains an empirical question that can be answered by future research.

4.1. Limitations of our data

To our knowledge, our study provides the most comprehensive cross-linguistic survey of polysemy undertaken to date. We went beyond previous work by focusing on a wide variety of patterns of polysemy across a wide variety of languages. We assessed multiple sets of senses of each pattern of polysemy, and (for almost all languages) used multiple informants to provide ratings. Our data therefore provide the clearest picture yet of how polysemy does and does not vary across different languages; we believe that this picture only makes sense under the conventions-constrained-by-concepts model.

That said, there were limits to our study design that could be improved upon in future work. For example, our finding that patterns of polysemy were generally present across languages could have been influenced by Anglo-centricity: we focused on whether obtained translation-equivalents of English words in other languages exhibit English-like polysemy. This raises the possibility that polysemy in other languages has developed due to contact with English, resulting in broad similarities. We find this explanation unlikely for two reasons. First, many of the languages we tested are historically unrelated to English (e.g., Japanese, Mandarin, Indonesian, Turkish), casting doubt on the idea that contact with English created the parallels we observed. Second, although our data revealed that the same patterns are present across languages, it also revealed considerable cross-linguistic variation in the sets of senses that follow these patterns. For instance, while English uses *glass* to describe a material and a drinking vessel made from the material, this is not true across languages: Spanish uses *glass* to label a car window while Russian uses *glass* to label a mirror (meanwhile, Turkish, Cantonese, Farsi, Japanese and Korean demand that speakers use a whole phrase, the equivalent of “glass cup”). If sets of senses across languages have developed due to contact with English, we would not expect this variability, suggesting that the presence of the same patterns across languages is not due to the influence of English. Still, in order to test our suggestion that patterns are generally conserved across languages, future cross-linguistic work should explore whether there are patterns found in other languages that are not found in English, as our method could not probe for these patterns.

Future work should also confront a number of logistical concerns we faced, to further corroborate our data. First, we were only able to recruit a limited number of informants per language (2.5 on average). While we believe that our informants' linguistic judgments are robust and clear-cut (particularly those of our informants who were trained in linguistics), replication via additional informants would be useful. Second, an anonymous reviewer suggested that our use of bilingual English-speaking participants may have biased the results, as these speakers may suffer lexical interference. This shortcoming could be overcome by translating the survey into different languages and testing monolingual informants.

Finally, although we tested a large number of patterns of polysemy in a large number of languages, our survey was not exhaustive, meaning that our conclusions can only be tentative until additional languages and patterns are included. In addition, time constraints meant that, on average, we assessed only eight examples of senses per pattern (see Table 4). This could have skewed our results in two ways: We may have fortuitously tested senses that happen to be cross-linguistically common (making polysemy seem more similar across languages than it actually is), or we may have probed senses that happen to be cross-linguistically rare (making polysemy seem more different across languages than it actually is).

The above discussion suggests that a more extensive survey could be helpful to solidify our conclusions about the regularity of patterns and senses across languages. Such an effort would be important, as robust conclusions about universal versus language-specific patterns of polysemy could provide insights into how conceptual structure constrains polysemy. Of course, there are reasons for why a pattern could be unattested in a language that do not invoke conceptual structure. For example, as described before, some patterns may be pre-empted by the presence of morphological rules or compounding, and grammatical restrictions on the productivity of morphological conversion could prevent some patterns from emerging in languages in the first place (e.g., some languages may not permit noun-verb conversions).

Still, the finding that a pattern is relatively unattested across many languages may also raise questions about whether the relation it encodes is less conceptually salient than those encoded by other patterns. One possible example of this is the most absent pattern in our data, the set of locatum verbs describing the transfer of substances from sources (“John *milked* the cow”), which was absent in six languages. As noted in the results section, this level of absence was more than would have been expected by chance, suggesting that there may be good reasons for its absence. And indeed, recent work has suggested that children learning English have a surprising level of difficulty learning verbs like this, because they initially assume that they label the transfer of substance to goals rather than from sources (similar to verbs like *salt* or *butter*; Srinivasan and Barner, 2013a). This results in striking errors, wherein children initially assume that “*milking* the cow” involves putting milk onto a cow! This bias likely stems from a more fundamental goal bias in language and thought, which has been documented extensively in previous work (e.g., Lakusta and Landau, 2005; Lakusta et al., 2007; Papafraou, 2010; Regier, 1997). This example suggests that cross-linguistic data from polysemy can provide evidence and perhaps new hypotheses about the nature of conceptual structure.

4.2. The “conventions-constrained-by-concepts” model of polysemy

To properly understand the relationship between polysemy and conceptual structure, it is necessary to have a model of the structure of polysemy, and understand why words have the senses that they do. We believe that the current best candidate is the conventions-constrained-by-concepts model, which we now describe in more detail. By this model, senses of polysemous words are learned as conventions, explaining why different languages may develop different senses. But the set of conventions that are easy to learn is constrained by conceptual structure, accounting for why sets of senses may form the same patterns across languages. This model also predicts – as we found – that some patterns will have more variable senses across languages than others, and that this will be linked to how generative the pattern is. Taken together, our results suggest that the structure of polysemy is mutually constrained by both concepts and conventions.

Thus far, we have focused on just one version of the conventions-constrained-by-concepts model, in which conceptual structure constrains how senses are first learned. By this account, the sets of senses that emerge in a language will be those that are more learnable, given the conceptual biases of children. Because these cognitive biases are likely to be shared by children of all linguistic communities, different languages will develop similar patterns, containing easily learnable sets of senses that are adapted to these cognitive biases. Patterns that tightly constrain possible senses, like animal for meat, will result in sets of senses that can be more easily aligned, setting the stage for higher-order generalizations that permit generative use of the pattern. This type of explanation, in which children’s learning biases lead to characteristic structural features of language, has precedent in empirical work showing that children impose grammatical structure onto language in the face of impoverished input (e.g., Bickerton, 1984; Goldin-Meadow and Mylander, 1990; Senghas and Coppola, 2001), that children’s expectation that words will contrast may explain the lack of true synonyms in languages (e.g., Clark, 1988), as well as in theoretical and experimental studies of how learning biases, iterated across generations of learners, can influence the structure of the lexicon (Silvey et al., 2014; Smith, 2004).

However, it is also possible that conceptual structure plays its constraining role not when children are learning new senses, but when these senses are coined by adult language users. Thus, because adult speakers and listeners across different linguistic communities may find certain conceptual relations salient – like the relation between an animal and its meat or a material and its artifact – they may be more likely to coin sets of senses that fit these relations, explaining why the same patterns are present across languages. Children, by this view, would make no contribution to this process, but would simply memorize widely used senses as conventions. This account could explain not only why the same patterns are present across languages, but also why some patterns have variable senses across languages. In particular, for loosely constraining patterns, there could be many possible sets of senses that could be coined (e.g., *glass* to label a window, drinking vessel, mirror, etc.), such that the first senses coined could differ between languages. But once a set of senses has been coined and memorized (e.g., *glass* to label a drinking vessel), adult speakers would avoid coining other possible senses (such as *glass* to label a window), to limit ambiguity.

One possibility is that both versions of the model make some contribution: concepts could constrain conventions when children are first learning senses and also when adult speakers are coining these senses. However, we think that the developmental effects on the structure of the lexicon are likely to be stronger for two reasons. First, the conceptual biases of adult speakers are likely to be culturally influenced – or at least more likely to be influenced than those of children – which is hard to reconcile with our finding that patterns were nearly universal across languages. Second, and perhaps more importantly, there is an emerging body of evidence suggesting that children have strong expectations about how sets of senses should be grouped under the same label, and that these expectations may have their roots in universal and early-developing cognitive biases.

This recent evidence motivates a new view of polysemy, in which the structure of the lexicon is, in part, a consequence of how children approach the problem of learning mappings between words and concepts. In particular, children may sidestep the task of learning mappings one-by-one, and instead expect that words will label multiple concepts in systematic and constrained ways. We discuss this theory and the evidence that motivates it in the next section.

4.3. Polysemy as a tool for building a lexicon

By adulthood, humans can use words to express an extraordinary range of basic concepts. But if each of these basic concepts had to be represented by a separate word, the resulting lexicon would be vast, and building it would be slow and difficult. This is because mappings between word forms and concepts are arbitrary (Saussure, 1959), such that each mapping has to be learned on its own. We propose that polysemy arises as a way of reducing the arbitrariness of this mapping process, and speeding up learning. In particular, after learning one sense of a polysemous word, children could be well-equipped to guess its other senses.

The proposal that polysemy aids word learning might initially seem surprising: shouldn’t it be confusing to learn a language whose words conflate different concepts? On the contrary, we suggest that, because the senses of polysemous words are related, learning one sense of a word could provide a clue for learning its other senses. Imagine, for example, a

parent who would like to warn their child about a shard of glass on the floor. The child has not yet learned the name of this material, but has learned that *glass* can label a kind of drinking vessel. If the parent tells the child “there is *glass* on the floor”, the child may reason that the referent of *glass* may be related in some way to the drinking vessel – e.g., perhaps they share the same material – and may constrain her hypotheses on this basis. However, without such polysemy, the parent would have to use a different word for the material, e.g., “there is *dax* on the floor”, which could potentially refer to anything on the floor. This suggests that learning multiple senses should be much easier than learning multiple words.

Moreover, polysemy may also allow children to *spontaneously* infer new senses, rather than just learn through observation. In particular, if children make higher-order generalizations about tightly constraining patterns of polysemy, (e.g., by realizing that animal names also label meats after learning the senses of *chicken*, *lamb* and *fish*), then they would be in a position to infer new meanings without direct, ostensive evidence. For example, having learned that *seagull* labels an animal, children could spontaneously infer that it can also label the derived meat, even in absence of ostensive evidence that it can.²² Such generalizations would greatly simplify the process of learning word-concept mappings, because children would only need to learn one label to express two concepts.

Consistent with these ideas, there is now good evidence that, from early in development, children are able to learn multiple senses for words, and even expect words to have specific sets of senses. By at least age four, children are sensitive to the relations between the senses of polysemous words (e.g., the use of *book* to label an object or its abstract content), and distinguish them from unrelated homophones (e.g., the use of *bat* to label an animal or baseball equipment; Srinivasan and Snedeker, 2011, 2014). They also show a sophisticated ability to determine the correct meaning for an ambiguous word in context (Rabagliati et al., 2013). Moreover, even before age four, children appear to expect words to be used in innovative ways. For example, children creatively use words for space to describe time (“Mommy, can I have some reading *behind* dinner”; Bowerman, 1983), words for instruments to describe actions involving those instruments (“Don’t *broom* my mess”; Clark, 1982), and words for abstract content to describe objects (e.g., agreeing that a *movie* can be round; Rabagliati et al., 2010). Importantly, these innovations are related to attested senses of polysemous words (e.g., the use of *broom* is similar to attested uses of *hammer* and *shovel*), suggesting that children have formed higher-order generalizations that allow them to coin novel senses. Consistent with this, recent evidence indicates that four- and five-year-olds spontaneously expect new words to have multiple senses, according to existing patterns (e.g., to label tools and functional uses of those tools, as words like *hammer* and *shovel* do; Srinivasan and Barner, 2013b).

Together, the findings described above confirm that children can learn and even infer new word senses. But what kinds of conceptual biases might constrain children’s inferences about senses, and how would this affect the structure of polysemy? For children to guess the new senses of a polysemous word, they first have to understand how the concepts labeled by that word are both related and distinct. For instance, to learn the different senses of *glass*, children have to understand that a drinking glass is made from glass material. Further, to use the senses of *glass* correctly, children have to understand how they differ from one another. For example, when *glass* is used in its material sense, it can label any entity that is composed of glass, irrespective of its particular form or function, but when *glass* is used in its artifact sense, it labels an entity with a specific form, that was created with a particular function in mind (Bloom, 1996; Malt, 2010; Malt and Johnson, 1992). Thus, to learn the different senses of polysemous words like *glass*, children need to be able to flexibly conceptualize entities in multiple ways, differentially focusing on properties such as form, function, material composition, and origin.

Flexible conceptualization is a key component of the successful “Theory Theory” of cognitive development, which argues that, from early in life, children *do* have the ability to construe the world in different ways, based on a set of intuitive framework theories (e.g., Carey, 1985; Gelman and Wellman, 1991; Gopnik and Meltzoff, 1997; Keil, 1989, 1994; Leslie, 1994; Murphy and Medin, 1985; Prasada, 2000). These framework theories are argued to help children understand how an entity’s properties might be related to one another, and how this depends on the domain an entity belongs to – e.g., whether it is an artifact, substance, natural kind, and so on. Thus, guided by an intuitive theory of artifacts, children could understand that an object, like a drinking glass, has the function it does due to its process of creation (e.g., it may hold liquid because it was designed to), and could understand its form by invoking its function (e.g., it may be hollow so that

²² This example is interesting because it is likely that, to make the generalization that words can label animals and their meats, children will first need to understand (either implicitly or explicitly) that meat *comes from* animals. Indeed, prior to this, it is possible that children treat the different senses of these words as homophones (see Srinivasan and Snedeker, 2014 for discussion of this issue). Importantly, however, the kinds of lexical or conceptual structures that constrain children’s expectations about how word senses are related and that allow them to capture generalizations may still be present from early in acquisition. By this account, acquiring relevant world knowledge may be necessary for changing how children represent a specific pattern of flexibility, and such changes may take place for different patterns of flexibility at different times. For example, children may learn early on that the object and intellectual content senses of words like *book* and *magazine* are related, allowing them to abstract a higher-order generalization to capture this pattern. But this might happen only later in life for the animal for meat pattern because children – especially those that grow up in urban contexts – may initially fail to recognize the relation between the animal and its meat (indeed, it may be the lexical overlap that initially alerts them to this relation).

it can hold liquid). An intuitive theory of natural substances, in contrast, may specify different relations between these explanatory factors: e.g., to understand why a piece of glass has the properties it does, it may not make sense to appeal to its intended function.

Young children appear to have little difficulty focusing on different explanatory factors such as function, process of creation, form, and material composition (see Gelman, 2003 for a review). For example, Gelman and Bloom (2000) showed that three-year-olds can construe the same object in terms of its intended function as an artifact, or instead as a material, depending on the information they are provided with. Specifically, children will label a sharp piece of plastic as a “knife” if told that it was intentionally shaped that way, but will label it as “plastic” if told that it was created accidentally. Similar flexibility has also been observed in other domains, and may arise even before children acquire relevant aspects of language. For example, even before children acquire a syntactic mass-count distinction, they are able to alternately focus on the forms of objects, and on their material compositions. When a substance is non-solid, children treat material composition as more central to its kind membership than shape, and vice versa for solid substances (Soja et al., 1991). Prior to learning language, children are also able to alternately construe the same physical entity as an inanimate object or intentional agent. For instance, infants as young as 12 months will follow the “gaze” of a faceless object, if it interacts with them in a contingent, “animate” way (Johnson et al., 1998).

Children’s ability to adopt different explanatory construals raises the possibility that this ability could also guide their early expectations about the different senses of words. This proposal—that explanatory frameworks underlie polysemy—has not been directly tested. But it is interesting to note that, independently, a similar account of polysemy has been proposed in the lexical semantics literature. In particular, several scholars have argued that lexical items include explanatory schemes that give rise to and constrain a variety of linguistic phenomena, including polysemy (Keil, 1994; Moravcsik, 1981, 1990; Pustejovsky, 1995; Prasada, 2000). For example, the systematic alternation of words like *glass* between materials and products made from those materials, and of words like *chicken* between animals and meat, could stem from an ability to view an entity both as an object with a particular form (and intended function, in the case of a glass), and as something composed of a particular material (Moravcsik, 1981).

This overlap between theories of children’s concepts and theories of the lexical semantic representation of polysemy points to a compelling explanation of our findings. In particular, if children find it easier to learn sets of senses that correspond to the application of different explanatory schemas, then as different languages change and evolve over generations of learners, those languages should each develop patterns of easily learnable sets of senses that correspond to children’s flexible construals. This would in turn facilitate children’s acquisition of the lexicon, because by learning one sense of a word, children would be well-equipped to guess its other senses, especially for tightly-constraining patterns for which children can make higher-order generalizations. Thus, in our view, children actively shape the structure of polysemy and, in return, polysemy makes it easier for children to build a lexicon. This model is illustrated in Fig. 6.

4.4. Empirical consequences of our proposal

We believe that our proposal, that children’s cognitive biases shape polysemy, has several advantages over a proposal in which polysemy is constrained only when adult speakers coin new senses. First, as reviewed above, there is a wealth of evidence from the study of cognitive development to suggest that children have cognitive biases that could constrain how they learn new word senses. In contrast, studies that have measured conceptual structure in adults have found that it is a poor predictor of polysemy (Rabagliati et al., 2011). Second, an appeal to the cognitive biases of young children may provide a stronger source of constraints on possible universals, than an appeal to adult conceptual structure. In particular, compared to young children, adults growing up in different cultural communities may be more likely to have different cognitive biases, which would predict more variability in polysemy than we found in the present data. Third, and most centrally, our proposal provides a reason for *why* polysemy exists in natural language: Polysemy may help children build a lexicon because learning one sense of a word could provide information about its other possible senses.

However, a precise test of our theory will require further study. In particular, our cross-linguistic data cannot itself adjudicate between the developmental and non-developmental versions of the conventions-constrained-by-concepts model. One way to assess the developmental proposal is to explore whether children’s earliest hypotheses about word senses are guided by the same flexible, non-linguistic conceptual structures that allow them to construe the world in terms of form, material, function, and origin. For instance, we could assess whether non-linguistic cues that bias functional or form-based construals of an entity also bias word learning.

Another way to assess our theory is to focus on how children form higher-order generalizations about sets of senses, and see how this affects their acquisition and representation of polysemous words. As described before, higher-order generalizations may be formed as young children learn multiple sets of senses following a pattern (e.g., the animal/meat senses of *chicken*, *lamb*, etc.), and abstract their underlying relation (e.g., that words for animals can label their meats). This proposal raises a number of testable questions. First, how quickly are children able to make these generalizations, and thus, to what extent can these generalizations constrain children’s early hypotheses about polysemy? On one hand,

Fig. 6. How children’s flexible, explanatory construals might constrain the structure of polysemy over the course of language development and change.

it is possible that the process of forming such generalizations is a slow one, such that these generalizations do not play an important role in the acquisition of polysemous senses. However, it is also possible that such generalizations are quickly abstracted, and facilitate learning from early in life. Second, how might higher-order generalizations affect how children represent polysemous words? One possibility, for example, is that children initially represent sets of senses using separate lexical items, but are able to derive senses from one another on-line, once they have formed higher-order generalizations. This proposal is similar to previous accounts of how children acquire regular past-tense inflections for verbs (e.g., Marcus et al., 1992; O’Donnell et al., 2011; Pinker, 1991).

Moving beyond these targeted predictions, our proposal raises questions about how polysemy relates to other linguistic systems. Perhaps the clearest relationship is between polysemy and morphology. Our study used a working definition of polysemy in which the same word form is used to label different meanings, but the functions served by polysemy bear a clear resemblance to those served by morphological paradigms (see e.g., Pykkänen et al., 2006). Indeed, in some cases, polysemous alternations are mirrored by parallel morphological rules, e.g., just as one can say “He drank a bottle of whiskey”, using bottle as a measure of the contents of a bottle, one can also say “He drank a bottleful of whiskey”, using the suffix –ful to indicate a measure (Copestake and Briscoe, 1995). Thus, future research should explore how morphological phenomena are related to polysemy across languages. Semitic languages could provide a particularly interesting case study, because they include polysemy-like morphological paradigms, through which roots can be adjusted to encode different meanings, and which children master early in life (see, e.g., Berman, 1999, for the case of Hebrew).

Finally, our proposal raises an interesting question about the relationship between children’s conceptual knowledge of the world and their knowledge of polysemy. Above, we have assumed that children’s world knowledge – e.g., that a drinking glass is made of glass – supports their expectations about polysemy. However, it is also possible that learning

word senses teaches children about the world – e.g., that because a drinking vessel and transparent material share the same label, they must be somehow related. Indeed, there are anecdotes that support this idea: for example, parents have recounted that their children became vegetarians upon learning that it was no accident that the same word, *chicken*, labels an animal and its meat (Foer, 2010; Gelman, 2003)! Our theory does not rule out the possibility that polysemy teaches children about the world. However, we suspect that, at least early in life, it is likely that children's knowledge about the world will precede their knowledge of how words relate to one another, given that young children have difficulty explicitly reflecting upon how words are used (e.g., Bialystok, 1986; Gombert, 1992). Still, it remains possible that, in some cases, learning to use words flexibly might invite children to form relations between concepts. If this is the case, polysemy may open a window onto how language shapes cognition, including whether speakers of different languages may come to think differently about the world, by virtue of learning language-specific forms of polysemy.

Acknowledgments

We would like to thank all of the respondents who endured our exceptionally long linguistic survey, and in particular Eva Wittenberg for her considerable help during pilot testing. Gregory L. Murphy and our anonymous reviewers provided exceptionally detailed comments that greatly improved the manuscript. Jesse Snedeker not only provided helpful discussion, but also generously supported the work via funding from Harvard University. Hugh Rabagliati was supported in part by a grant from Harvard University's Mind, Brain, Behavior Initiative.

References

- Apresjan, J.D., 1974. Regular polysemy. *Linguistics* 142, 5–32.
- Baker, C.L., 1968. *Indirect Questions in English* (Doctoral dissertation). University of Illinois.
- Berman, R., 1999. Children's innovative verbs versus nouns: structured elicitations and spontaneous coinages. In: Menn, L., Bernstein Ratner, N. (Eds.), *Methods for Studying Language Production*. Erlbaum, Mahwah, NJ, pp. 69–93.
- Bialystok, E., 1986. Children's concept of word. *J. Psycholinguist. Res.* 15 (1), 13–32.
- Bickerton, D., 1984. The language bioprogram hypothesis. *Behav. Brain Sci.* 7, 173–188.
- Bloom, P., 1996. Intention, history, and artifact concepts. *Cognition* 60, 1–29.
- Bowerman, M., 1983. Hidden meanings: the role of covert conceptual structures in children's development of language. In: Rogers, D.R., Sloboda, J.A. (Eds.), *The Acquisition of Symbolic Skills*. Plenum, New York, pp. 445–470.
- Boyeldieu, P., 2008. From semantic change to polysemy. The cases of meat/animal and drink. In: Vanhove, M. (Ed.), *From Polysemy to Semantic Change. Towards a Typology of Lexical Semantic Associations*. John Benjamins, Amsterdam/Philadelphia, pp. 303–315.
- Breál, M., 1897. *Essai de sémantique: Science des significations*. Hachette, Paris.
- Carey, S., 1985. *Conceptual Change in Childhood*. MIT Press, Cambridge.
- Carlson, G., 1977. *Reference to Kinds in English* (Doctoral dissertation). University of Massachusetts.
- Casson, R.W., 1994. Russett, rose, and raspberry: the development of English secondary color terms. *J. Linguist. Anthropol.* 4 (1), 5–22.
- Clark, E.V., 1982. The Young Word Maker: a case study of innovation in the child's lexicon. In: Wanner, E., Gleitman, L. (Eds.), *Language Acquisition*. University Press, Cambridge, pp. 390–428.
- Clark, E.V., 1988. On the logic of contrast. *J. Child Lang.* 15, 317–335.
- Clark, E.V., Clark, H.H., 1979. When nouns surface as verbs. *Language* 55 (4), 767–811.
- Copestake, A., Briscoe, T., 1995. Semi-productive polysemy and sense extension. *J. Semant.* 12, 15–67.
- Falkum, I.L., 2011. *The Semantics and Pragmatics of Polysemy: A Relevance-theoretic Account* (Ph.D. thesis). University College London, London.
- Fauconnier, G., 1985. *Mental Spaces*. MIT Press, Cambridge, MA.
- Foer, J.S., 2010. *Eating Animals*. Little, Brown, and Company, New York.
- Galimberti Jarman, B., Russell, R., Rollin, N., & Carvajal, C. S. (Eds.). (2008). *The Oxford Spanish Dictionary: Spanish-English, English-Spanish*. (4th/edited by Nicholas Rollin, Carol Styles Carvajal ed.). New York, NY: Oxford University Press.
- Gelman, S.A., 2003. *The Essential Child: Origins of Essentialism in Everyday Thought*. Oxford University Press, New York.
- Gelman, S.A., Bloom, P., 2000. Young children are sensitive to how an object was created when deciding what to name it. *Cognition* 76, 91–103.
- Gelman, S.A., Wellman, H.M., 1991. Insides and essences: early understanding of the non-obvious. *Cognition* 38, 213–244.
- Goldin-Meadow, S., Mylander, C., 1990. Beyond the input given: the child's role in the acquisition of language. *Language* 66 (2), 323–355.
- Gombert, J.E., 1992. *Metalinguistic Development*. University of Chicago Press, Chicago, IL.
- Gopnik, A., Meltzoff, A., 1997. *Words, Thoughts, and Theories*. MIT Press, Cambridge, MA.
- Greenberg, J.H., 1983. Some areal characteristics of African languages. In: Dihoff, I.R. (Ed.), *Current Approaches to African Linguistics*, vol. 1. Foris, Dordrecht, pp. 3–21.
- Harrap's., 2001. *Harrap's Unabridged French-English/English-French Dictionary*. London, UK.
- Iz, F., Alderson, A.D., Hony, H.C., 1992. *The Oxford Turkish Dictionary*. Oxford University Press, Oxford, UK.
- Jackendoff, R., 1997. *The Architecture of the Language Faculty*. MIT Press, Cambridge, MA.
- Johnson, S.C., Slaughter, V., Carey, S., 1998. Whose gaze will infants follow? Features that elicit gaze-following in 12-month-olds. *Dev. Sci.* 1, 233–238.
- Kamei, S.-I., Wakao, T., 1992. Metonymy: Reassessment, Survey of Acceptability, and Its Treatment in a Machine Translation System. Paper presented at the Proceedings of the 30th Annual Meeting of the Association for Computational Linguistics.

- Keil, F.C., 1989. *Concepts, Kinds, and Cognitive Development*. MIT Press, Cambridge, MA.
- Keil, F.C., 1994. Explanation based constraints on the acquisition of word meaning. *Lingua* 92, 169–196.
- Kiparsky, P., 1997. Remarks on denominal verbs. In: Alsina, A., Bresnan, J., Sells, P. (Eds.), *Complex Predicates*. CSLI Publications, Palo Alto, CA, pp. 473–499.
- Klein, D.E., Murphy, G.L., 2001. The representation of polysemous words. *J. Mem. Lang.* 45, 259–282.
- Klein, D.E., Murphy, G.L., 2002. Paper has been my ruin: conceptual relations of polysemous senses. *J. Mem. Lang.* 47, 548–570.
- Krifka, M., 2001. Lecture Notes on Polysemy. Humboldt-Universität zu Berlin Retrieved from: http://amor.cms.hu-berlin.de/~h2816i3x/Lehre/2001_HS_LexikalischeSemantik/LexSemantik5.pdf.
- Lakoff, G., 1987. *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. The Chicago University Press, Chicago.
- Lakoff, G., Johnson, M., 1980. *Metaphors We Live By*. University of Chicago Press, Chicago.
- Lakusta, L., Landau, B., 2005. Starting at the end: the importance of goals in spatial language. *Cognition* 96, 1–33.
- Lakusta, L., Wagner, L., O’Hearn, K., Landau, B., 2007. Conceptual foundations of spatial language: evidence for a goal bias in infants. *Lang. Learn. Dev.* 3, 179–197.
- Lehrer, A., 1990. Polysemy, conventionality, and the structure of the lexicon. *Cognit. Linguist.* 1, 207–246.
- Leslie, A.M., 1994. ToMM, ToBY, and agency: core architecture and domain specificity. In: Hirschfeld, L.A., Gelman, S.A. (Eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture*. Cambridge University Press, New York.
- Malt, B.C., 2010. Naming artifacts: patterns and processes. In: Ross, B. (Ed.), *The Psychology of Learning and Motivation: Advances in Research and Theory*. Elsevier, Amsterdam, pp. 1–38.
- Malt, B.C., Johnson, E.C., 1992. Do artifact concepts have cores? *J. Mem. Lang.* 31, 195–217.
- Marcus, G.F., Pinker, S., Ullman, M., Hollander, M., Rosen, T.J., Xu, F., 1992. Overregularization in language acquisition. *Monogr. Soc. Res. Child Dev.* 57 (4, Serial No. 228).
- Martin, S.E., Lee, Y.-H., Chang, S.-U. (Eds.), 1967. *A Korean-English Dictionary*. Yale University Press, New Haven, CT.
- Moravcsik, J.M.E., 1981. How do words get their meanings? *J. Philos.* 78, 5–24.
- Moravcsik, J.M.E., 1990. *Thought and Language*. Routledge.
- Murphy, G.L., 1997. Polysemy and the creation of novel word meanings. In: Ward, T.B., Smith, S., Vaid, J.J. (Eds.), *Creative Thought: An Investigation of Conceptual Structures and Processes*. American Psychological Association, Washington, DC, pp. 235–265.
- Murphy, G.L., 2007. Parsimony and the psychological representation of polysemous words. In: Rakova, M., Petho, G., Rákosi, C. (Eds.), *The Cognitive Basis of Polysemy*. Peter Lang Verlag, Frankfurt am Main, Germany.
- Murphy, G.L., Medin, D.L., 1985. The role of theories in conceptual coherence. *Psychol. Rev.* 92, 289–316.
- Nunberg, G., 1979. The non-uniqueness of semantic solutions: polysemy. *Linguist. Philos.* 3, 143–184.
- Nunberg, G., 1995. Transfers of meaning. *J. Semant.* 12, 109–132.
- O’Donnell, T.J., Snedeker, J., Goodman, N.D., Tenenbaum, J.B., 2011. Productivity and reuse in language. In: *Proceedings of the Thirty-third Annual Conference of Cognitive Science Society*.
- Orszagh, L., Futasz, D., Kövecses, Z., 1998. *Magyar Angol Nagyszotar/Hungarian-English Dictionary*. Akadémiai Kiadó, Budapest, Hungary.
- Ostler, N., Atkins, B., 1992. Predictable meaning shift: some linguistic properties of lexical implication rules. *Lex. Semant. Knowl. Represent.* 627, 87–100.
- Papafragou, A., 1996. On metonymy. *Lingua* 99 (4), 169–195.
- Papafragou, A., 2010. Source-goal asymmetries in motion representation: implications for language production and comprehension. *Cognit. Sci.* 34, 1064–1092.
- Peters, W., 2003. Metonymy as a cross-lingual phenomenon. In: *Proceedings of the ACL 2003 Workshop on Lexicon and Figurative Language*, Association for Computational Linguistics, pp. 1–9.
- Pinker, S., 1991. Rules of language. *Science* 253, 530–535.
- Pinker, S., 2007. *The Stuff of Thought: Language as a Window into Human Nature*. Viking, New York.
- Prasada, S., 2000. Acquiring generic knowledge. *Trends Cogn. Sci.* 4 (2), 66–72.
- Pustejovsky, J., 1995. *The Generative Lexicon*. MIT Press, Cambridge.
- Pylkkänen, L., Llinas, R., Murphy, G.L., 2006. The representation of polysemy: MEG evidence. *J. Cogn. Neurosci.* 18, 1–13.
- R Development Core Team, 2014. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria URL: <http://www.R-project.org>.
- Rabagliati, H.A., Marcus, G.E., Pylkkänen, L., 2010. Shifting senses in semantic development. *Cognition* 117 (1), 17–37.
- Rabagliati, H., Marcus, G.F., Pylkkänen, L., 2011. Rules, radical pragmatics and restrictions on regular polysemy. *J. Semant.* 28, 485–512.
- Rabagliati, H., Pylkkänen, L., Marcus, G.F., 2013. Top-down influence in young children’s linguistic ambiguity resolution. *Dev. Psychol.* 49, 1076–1089.
- Regier, T., 1997. Constraints on the learning of spatial terms: a computational investigation. In: Goldstone, R., Schyns, P., Medin, D. (Eds.), *Psychology of Learning and Motivation*, Vol. 36: Mechanisms of Perceptual Learning. Academic Press, San Diego, pp. 171–271.
- Reynolds, B. (Ed.), 1981. *The Cambridge Italian Dictionary*. Cambridge University Press, Cambridge, UK.
- Ruhl, C., 1989. *On Monosemy: A Study in Linguistic Semantics*. SUNY Press, Albany.
- Saussure, F.de, 1959. *Course in General Linguistics*. McGraw-Hill.
- Senghas, A., Coppola, M., 2001. Children creating language: how Nicaraguan Sign Language acquired a spatial grammar. *Psychol. Sci.* 12 (4), 323–328.
- Silvey, C., Kirby, S., Smith, K., 2014. Word meanings evolve to selectively preserve distinctions on salient dimensions. *Cognit. Sci.* <http://dx.doi.org/10.1111/cogs.12150>
- Smith, K., 2004. The evolution of vocabulary. *J. Theor. Biol.* 228, 127–142.
- Soja, N.N., Carey, S., Spelke, E.S., 1991. Ontological categories guide young children’s inductions on word meaning: object terms and substance terms. *Cognition* 38, 179–211.
- Srinivasan, M., Barner, D., 2013a. The Amelia Bedelia effect: world knowledge and the goal bias in language acquisition. *Cognition* 128, 431–450.

- Srinivasan, M., Barner, D., 2013b. *Daxing with a dax: 4- and 5-year-olds Expect Words to Label Instruments and the Actions they Support*. Poster Presented at the Biennial Meeting of the Society for Research in Child Development, Seattle, WA.
- 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. *Cognit. Psychol.* 62, 245–272.
- Srinivasan, M., Snedeker, J., 2014. Polysemy and the taxonomic constraint: children's representations of words with taxonomically-different meanings. *Lang. Learn. Dev.* 10, 97–128.
- Stevens, A.M., Schmidgall-Tellings, A. (Eds.), 2004. *A Comprehensive Indonesian English Dictionary*. Ohio University Press, Athens, OH.
- Strigin, A., 1998. *Lexical rules as hypotheses generators*. *J. Semant.* 15 (2), 163–190.
- United States Joint Publications Research Service, 1966. *Vietnamese English Dictionary*. National Technical Information Service, U.S. Department of Commerce, Washington, DC.
- Velázquez de la Cadena, M., Gray, E., Iribas, J.L., Navarro Hinojosa, I., Blanco-González, M., Nelson, R.J. (Eds.), 2003. *Velázquez Spanish and English Dictionary*. Velázquez Press, El Monte, CA.
- Viberg, A., 1984. *The verbs of perception: a typological study*. *Linguistics* 21 (1), 123–162.
- Vossen, P., 1998. *EuroWordNet: A Multilingual Database with Lexical Semantic Networks*. Kluwer Academic.
- Wilson, D., 2003. *Relevance and lexical pragmatics*. *Ital. J. Linguist.* 15, 273–292.
- Wilson, D., Carston, R., 2007. *A unitary approach to lexical pragmatics: relevance, inference and ad hoc concepts*. In: Burton-Roberts, N. (Ed.), *Pragmatics*. Palgrave Macmillan, London, pp. 230–259.
- Zhu, H., Malt, B.C., 2014. *Cross-linguistic evidence for cognitive foundations of polysemy*. In: *Proceedings of the 36th Annual Conference of the Cognitive Science Society*. Cognitive Science Society, Austin, TX, pp. 934–939.